



Brocade[®] Fabric OS[®] Administration Guide, 9.1.x

Administration Guide
January 11, 2024

Table of Contents

Introduction.....	27
Supported Hardware and Software.....	27
Contacting Technical Support for Your Brocade® Product.....	28
Document Feedback.....	28
Understanding Fibre Channel Services.....	29
Management Server.....	29
Application Server.....	30
Platform Services.....	30
Platform Services and Virtual Fabrics.....	30
Enabling Platform Services.....	31
Disabling Platform Services.....	31
Management Server Database.....	31
Displaying the Management Server ACL.....	31
Adding a Member to the ACL.....	32
Deleting a Member from the ACL.....	33
Viewing the Contents of the Management Server Database.....	34
Clearing the Management Server Database.....	34
Supporting the New Management Server ACL Interface.....	35
Configuring the Management Server ACL.....	35
Topology Discovery.....	37
Displaying Topology Discovery Status.....	37
Enabling Topology Discovery.....	37
Disabling Topology Discovery.....	37
Device Login.....	38
Principal Switch.....	38
E_Port Login Process.....	39
Fabric Login Process.....	39
Port Login Process.....	39
RSCNs.....	39
Duplicate Port World Wide Name.....	40
Device Recovery.....	40
High Availability of the Daemon Process.....	40
Performing Basic Configuration Tasks.....	42
Command Line Interface.....	42
Console Sessions Using the Serial Port.....	42
Connecting to Fabric OS through the Serial Port.....	42

Management Session.....	43
Rules for Management Connections.....	43
Connecting to Fabric OS Using the Management Interface.....	44
Configuring the Idle Session Timeout.....	44
Getting Help on a Command.....	45
Viewing the History of Command Line Entries.....	45
CLI History.....	46
clihistory --show.....	47
clihistory --showuser username.....	47
clihistory --showall.....	47
clihistory --help.....	47
Considerations When Viewing CLI Logs.....	48
Using fosexec to Run Commands on the Remote Switch or Domain.....	48
Modification of the Default Password.....	50
Default Account Password.....	51
Changing the Default Account Password of the Switch.....	51
Switch Ethernet Interface.....	52
Brocade Directors.....	52
Brocade Switches.....	52
Virtual Fabrics and the Ethernet Interface.....	52
Management Ethernet Port Bonding.....	53
Supported Brocade Directors.....	53
Setting Up the Second Ethernet Port on a CPX7 or CPX6 or CPX Blade.....	54
Displaying the Network Interface Settings.....	54
Resetting the Error Counter for the Management Network Interface.....	55
Setting the TCP Timeout Level.....	55
Management Port Tracker.....	55
Configuring Global LLDP Parameters for the Management Port.....	57
Management Port Ingress Rate Limiting.....	59
Static Ethernet Addresses.....	60
Setting a Static Address for the Ethernet Network Interface.....	60
Setting a Static Address for the Chassis Management IP Interface.....	61
DHCP Activation.....	61
Configuring a Static IPv6 Gateway Address.....	66
IPv6 Autoconfiguration.....	66
Setting IPv6 Autoconfiguration.....	67
Setting the Ethernet Interface Mode and Speed.....	67
Date and Time Settings.....	68
Setting the Date and Time.....	68
Time Zone Settings.....	69

Setting the Time Zone.....	70
Setting the Time Zone Interactively.....	70
Network Time Protocol.....	71
Synchronizing the Local Time with an External Source.....	72
Configuring Symmetric Authentication Keys.....	73
Distributing the NTP Configuration to Access Gateways.....	75
Domain IDs.....	75
Domain ID Issues.....	76
Displaying the Domain IDs.....	76
Setting the Domain ID.....	77
Registered Organization Name.....	77
Setting and Displaying the Registered Organization Name.....	77
Switch Name.....	78
Customizing the Switch Name.....	78
Switch Type.....	78
Chassis Name.....	79
Customizing the Chassis Name.....	79
Fabric Name.....	79
Configuring the Fabric Name.....	80
High Availability Considerations for Fabric Names.....	80
Upgrade and Downgrade Considerations for Fabric Names.....	80
Switch Activation and Deactivation.....	80
Disabling a Switch.....	81
Enabling a Switch.....	81
Disabling a Chassis.....	81
Enabling a Chassis.....	82
Device Shutdown.....	82
Powering Off a Brocade Switch.....	82
Powering Off a Brocade Director.....	82
Basic Connections.....	83
Device Connection.....	83
Switch Connection.....	83
DNS.....	84
Performing Advanced Configuration Tasks.....	85
Port Identifiers and PID Binding Overview.....	85
Unified Addressing Mode.....	85
Core PID Addressing Mode.....	86
Fixed Addressing Mode.....	87
10-Bit Addressing (Mode 0).....	87
256-Area Addressing (Mode 1 and Mode 2).....	87

Zero-Based Addressing (Mode 1).....	88
Port-Based Addressing (Mode 2).....	88
Port Address Binding.....	88
WWN-Based PID Assignment.....	88
Virtual Fabrics Considerations for WWN-Based PID Assignment.....	89
NPIV.....	89
Enabling Automatic PID Assignment.....	89
Assigning a Static PID.....	90
Clearing a PID Binding.....	90
Showing PID Assignments.....	90
Ports.....	90
Port Types.....	90
Director Port Blades.....	91
Setting Port Names.....	92
Port Identification by Slot and Port Number.....	92
Port Identification by Index.....	92
Dynamic Port Name.....	92
Dynamic Port Name Format.....	95
Enabling a Port.....	97
Disabling a Port.....	97
Port Decommissioning.....	98
Port Decommissioning without Disabling.....	98
Setting Port Speeds.....	99
Setting All Ports on a Switch to the Same Speed.....	99
Setting Port Speed for a Port Octet.....	100
Setting the Maximum Autonegotiated Port Speed.....	102
Decoding fdmishow Command Output.....	102
Blade Terminology and Compatibility.....	104
CP Blades.....	105
Core Blades.....	105
Port and Application Blade Compatibility.....	105
Disabling Switches.....	106
Configuring Fibre Channel Link Aggregation.....	106
Power Management.....	108
Powering Off a Port Blade or Core Blade.....	108
Powering On a Port Blade or Core Blade.....	108
Persistently Powering Off a Port Blade or Core Blade.....	109
Equipment Status.....	109
Checking Switch Operation.....	110
Verifying High-Availability Features in Directors.....	110

Verifying Fabric Connectivity.....	111
Verifying Device Connectivity.....	111
Audit Log Configuration.....	112
Verifying the Host Syslog Prior to Configuring the Audit Log.....	113
Configuring an Audit Log for Specific Event Classes, Severity Levels, or Message ID Filters.....	113
Configuring Remote Syslog Servers.....	114
Using Secure Syslog CA Root Certificates.....	115
Host Name Support for the syslogAdmin Command.....	115
Configuring Quiet Time for RASLog and Audit Log Messages.....	116
Duplicate PWWN Handling during Device Login.....	119
Setting the Behavior for Handling Duplicate PWWNs.....	120
Forward Error Correction.....	121
FEC Limitations.....	121
Disabling FEC.....	121
Enabling or Disabling FEC for Long-Distance Ports.....	122
Displaying and Clearing the FEC Counters.....	122
Fabric-Assigned PWWN.....	124
User-Assigned and Auto-Assigned FA-PWWN Behavior.....	125
Configuring an FA-PWWN for an HBA Connected to an Access Gateway.....	126
Configuring an FA-PWWN for an HBA Connected to an Edge Switch.....	127
Supported Switches and Configurations for FA-PWWN.....	128
Configuration Upload and Download Considerations for FA-PWWN.....	128
Security Considerations for FA-PWWN.....	128
Restrictions of FA-PWWN.....	128
Access Gateway N_Port Failover with FA-PWWN.....	129
N_Port ID Virtualization.....	130
Unified Addressing Mode.....	130
Fixed Addressing Mode.....	131
10-Bit Addressing Mode.....	131
Configuring NPIV.....	131
Enabling and Disabling NPIV.....	132
Base Device Logout.....	133
Difference in Device Logout Behaviors.....	133
Enabling Base Device Logout.....	134
Use Cases and Dependencies.....	134
Viewing the Base Device Logout Setting.....	135
Viewing NPIV Port Configuration Information.....	136
Viewing Virtual PID Login Information.....	138
Maintaining the Switch Configuration File.....	139

Configuration File Format.....	139
Chassis Section.....	140
Switch Section.....	140
Configuration File Backup.....	140
Special Characters in Server Credentials.....	141
Uploading a Configuration File in Interactive Mode.....	141
Configuration File Restoration.....	142
Restrictions When Downloading a Configuration File.....	142
Configuration Download without Disabling a Switch.....	144
Restoring a Configuration.....	144
Configurations across a Fabric.....	145
Downloading a Configuration File from One Switch to Another Switch of the Same Switch Type.....	145
Security Considerations.....	145
Configuration Management for Virtual Fabrics.....	146
Uploading a Configuration File from a Switch with Virtual Fabrics Enabled.....	146
Restoring a Logical Switch Configuration Using configDownload.....	146
Configuration Upload and Download Restrictions When Virtual Fabrics Is Enabled.....	147
Brocade Configuration Form.....	148
Managing User Accounts.....	149
Role Management.....	149
Role-Based Access Control.....	149
Default Fabric OS Roles.....	150
Role Permissions.....	150
Managing User-Defined Roles.....	151
Creating a User-Defined Role.....	151
Assigning a User-Defined Role to a User.....	152
Maximum Session Login Limits.....	153
Local Database User Accounts.....	153
Default Accounts.....	153
Displaying Account Information.....	154
Managing a User-Defined Account.....	154
Local Account Passwords.....	155
Changing the Password for the Current Login Account.....	155
Changing the Password for a Different Account.....	156
Password Policies.....	156
Password Strength Policy.....	156
Password History Policy.....	157
Password Expiration Policy.....	158
Expiring Passwords Immediately.....	159
Account Lockout Policy.....	159

Enabling the Admin Lockout Policy.....	160
Unlocking an Account.....	160
Disabling the Admin Lockout Policy.....	160
Denial of Service Implications.....	160
Configuring the Password Hash Type.....	160
Local User Account Database Distribution.....	161
Distributing the Local User Database.....	162
Accepting Distributed User Databases on the Local Switch.....	162
Rejecting Distributed User Databases on the Local Switch.....	162
Remote User.....	162
Overview.....	163
Client/Server Model.....	163
Authentication Server Data.....	163
Switch Configuration.....	163
Remote User Configuration on the Switch.....	164
Adding an AAA Server to the Switch Configuration.....	164
Deleting an Authentication Server from the Configuration.....	165
Changing an Authentication Server Configuration.....	165
Changing the Order in Which Authentication Servers Are Contacted for Service.....	165
Displaying the Current Authentication Configuration.....	165
Setting the Switch Authentication Mode.....	166
Remote Server Configuration.....	167
RADIUS Server Data Configuration.....	167
LDAP Server Data Configuration.....	173
TACACS+ Server Data Configuration.....	178
Configuring Protocols.....	180
Secure Copy.....	181
Setting Up SCP for Configuration Uploads and Downloads.....	181
Secure Shell Protocol.....	182
SSH Public Key Authentication.....	182
Allowed-User.....	183
Configuring Incoming SSH Authentication.....	183
Configuring Outgoing SSH Authentication.....	184
Deleting Public Keys for Inbound Connections on the Switch.....	184
Deleting Private Keys for Outbound Connections on the Switch.....	185
Generating and Installing Hostkeys on a Switch.....	185
Configuring the Rekey Interval on a Switch.....	185
Managing Known Hosts on a Switch.....	186
SecCryptoCfg Templates.....	186
SecCryptoCfg Template Management.....	190

Default SecCryptoCfg Templates.....	191
Certificate Revocation Check Enforcement.....	197
Client Authentication for TLS Client: RADIUS/LDAP/SYSLOG.....	198
Configuring the Ciphers, KEX, and MAC Algorithms.....	198
Secure Sockets Layer Protocol.....	199
Browser Support.....	199
SSL Configuration Overview.....	199
Certificate Authorities.....	200
Managing the Security Certificates Using the SecCertMgmt Command.....	200
Generating and Storing a Certificate Signing Request.....	203
Obtaining CA-Signed Certificates.....	205
Installing a Switch Certificate.....	205
Generating Self-Signed Certificates.....	206
Creating a Complete Chain of CA Certificates.....	207
Configuring the Browser for HTTPS.....	208
Checking and Installing Root Certificates on Mozilla Firefox.....	208
Simple Network Management Protocol.....	208
SNMP Manager.....	208
SNMP Agent.....	208
Management Information Base.....	209
Basic SNMP Operation.....	209
Configuring SNMP Using the CLI.....	210
Configuring the SNMP Security Level.....	210
Supported Protocol Configurations for SNMPv3 Users.....	210
Configuring SNMPv3 Users and Traps.....	211
Encrypting SNMPv3 User Passwords.....	219
Changing Encrypted SNMPv3 User Passwords to Unencrypted Form.....	220
Managing SNMPv1 Communities and SNMPv3 Users.....	220
Telnet Protocol.....	224
Blocking Telnet.....	224
Port Configuration.....	226
Unblocking Telnet.....	226
Listener Applications.....	227
Ports and Applications Used by Switches.....	227
Configuring Security Policies.....	228
How ACL Policies Are Stored.....	228
Virtual Fabric Considerations for ACL Policies.....	228
Policy Members.....	228
ACL Policy Management.....	229
Displaying ACL Policies.....	229

Saving Changes without Activating the Policies.....	229
Activating ACL Policy Changes.....	229
Deleting an ACL Policy.....	230
Adding a Member to an Existing ACL Policy.....	230
Removing a Member from an ACL Policy.....	231
Abandoning Unsaved ACL Policy Changes.....	231
FCS Policies.....	231
FCS Policy Restrictions.....	232
Ensuring That the Fabric Domains Share the Same Policy.....	232
Creating an FCS Policy.....	233
Modifying the Order of FCS Switches.....	233
FCS Policy Distribution.....	234
Device Connection Control Policies.....	234
Virtual Fabrics Considerations for Device Connection Control Policies.....	235
DCC Policy Restrictions.....	235
Creating a Device Connection Control Policy.....	235
Deleting a DCC Policy.....	236
DCC Policy Behavior with Fabric-Assigned PWWNs.....	237
SCC Policies.....	238
Virtual Fabrics Considerations for SCC Policies.....	240
Creating an SCC Policy.....	240
Authentication Policy for Fabric Elements.....	240
Virtual Fabrics Considerations for Authentication Policies.....	242
E_Port Authentication.....	242
Virtual Fabrics Considerations for E_Port Authentication.....	242
Configuring E_Port Authentication.....	242
Re-Authenticating E_Ports.....	243
Device Authentication Policy.....	244
Virtual Fabrics Considerations for the Authentication Policy.....	244
Configuring Device Authentication.....	244
AUTH Policy Restrictions.....	244
Authentication Protocols.....	245
Viewing the Current Authentication Parameter Settings for a Switch.....	245
Setting the Authentication Protocol.....	245
Secret Key Pairs for DH-CHAP.....	246
Characteristics of a Secret Key Pair.....	246
Viewing the List of Secret Key Pairs in the Current Switch Database.....	246
Setting a Secret Key Pair.....	247
FCAP Configuration Overview.....	248
Generating the Key and CSR for FCAP.....	248

Exporting the CSR for FCAP.....	249
Importing the CA for FCAP.....	249
Importing the FCAP Switch Certificate.....	249
Starting FCAP Authentication.....	249
Fabric-Wide Distribution of the Authorization Policy.....	249
IP Filter Policy.....	250
Virtual Fabrics Considerations for IP Filter Policies.....	250
Creating an IP Filter Policy.....	250
Cloning an IP Filter Policy.....	250
Displaying an IP Filter Policy.....	250
Saving an IP Filter Policy.....	251
Activating an IP Filter Policy.....	251
Deleting an IP Filter Policy.....	251
IP Filter Policy Rules.....	251
Source Address.....	252
Destination Port.....	252
Protocol.....	253
Action.....	253
Implicit Filter Rules.....	253
Default Policy Rules.....	254
Brocade Custom Policy Rules.....	254
IP Filter Policy Enforcement.....	255
Adding a Rule to an IP Filter Policy.....	255
Deleting a Rule from an IP Filter Policy.....	255
Aborting an IP Filter Transaction.....	255
Distributing an IP Filter Policy.....	256
Policy Database Distribution.....	256
Database Distribution Settings.....	257
Displaying the Database Distribution Settings.....	257
Enabling Local Switch Protection.....	258
Disabling Local Switch Protection.....	258
ACL Policy Distribution to Other Switches.....	258
Distributing the Local ACL Policies.....	258
Fabric-Wide Enforcement.....	258
Displaying the Fabric-Wide Consistency Policy.....	259
Setting the Fabric-Wide Consistency Policy.....	259
Notes on Joining a Switch to the Fabric.....	260
Matching Fabric-Wide Consistency Policies.....	260
Nonmatching Fabric-Wide Consistency Policies.....	261
In-Flight Encryption and Compression.....	263

Gen 7 and Gen 6 In-Flight Encryption and Compression Features	264
Supported Use Cases for the Brocade G630 Switch	266
Additional In-Flight Encryption and Compression Restrictions	273
Availability Considerations for Encryption and Compression	274
Virtual Fabrics Considerations for Encryption and Compression	274
How In-Flight Encryption and Compression Are Enabled	274
Authentication and Key Generation for Encryption and Compression	275
In-Flight Compression on Long-Distance Ports	275
Compression Ratios for Compression-Enabled Ports	276
Configuring In-Flight Encryption and Compression on an EX_Port	276
Configuring In-Flight Encryption and Compression on an E_Port	277
Viewing the Encryption and Compression Configuration	278
Configuring and Enabling Authentication for In-Flight Encryption	278
Enabling In-Flight Encryption	281
Enabling In-Flight Compression	282
Disabling In-Flight Encryption	283
Disabling In-Flight Compression	283
Encryption and Trunking Support on FC32-48 Blades	284
Buffer-to-Buffer Credits and Credit Recovery	286
Buffer-to-Buffer Flow Control	286
Optimal Buffer Credit Allocation	287
Considerations for Calculating Buffer Credits	287
Fibre Channel Gigabit Values Reference Definition	287
Buffer Credit Allocation Based on Full-Size Frames	288
Fibre Channel Data Frames	288
Allocating Buffer Credits Based on Full-Sized Frames	288
Calculating the Number of Buffers Required Based on Full-Size Frames	289
Allocating Buffer Credits Based on Average-Size Frames	290
Configuring Buffers for a Single Port Directly	291
Configuring Buffers Using Frame Size	291
Calculating the Number of Buffers Required Given the Distance, Speed, and Frame Size	291
Allocating Buffer Credits for F_Ports	293
Monitoring Buffers in a Port Group	293
Buffer Credits per Switch or Blade Model	295
Maximum Configurable Distances for Extended Fabrics	297
Configuring Credits for a Single VC	299
Increasing Credits for Normal Distance E_Ports	300
Buffer Credit Recovery	300
Buffer Credit Recovery over an E_Port	300
Buffer Credit Recovery over an F_Port	300

Buffer Credit Recovery over an EX_Port.....	301
Enabling and Disabling Buffer Credit Recovery.....	301
Credit Loss Detection.....	302
Enabling Backend Credit Loss Detection and Recovery for Link Reset Thresholds.....	302
Performing a Link Reset for a Loss of Sync from Peer Ports.....	302
Backend Credit Loss Detection and Recovery for Link Faults.....	303
Optimizing Fabric Behavior.....	305
Traffic Optimizer.....	305
Profiles.....	306
Profile Compatibility.....	307
Configuring Profiles.....	308
Speed and Protocol-Based Classification.....	313
Fabric Notifications.....	314
QoS.....	317
QoS on E_Ports.....	318
QoS over FC Routers.....	318
vTap and QoS High Priority Zoning Compatibility Mode.....	319
Virtual Fabrics Considerations for QoS Zone-Based Traffic Prioritization.....	320
Traffic Prioritization Based on QoS Zones.....	321
Setting QoS Zone-Based Traffic Prioritization.....	323
Setting QoS Zone-Based Traffic Prioritization over FC Routers.....	324
Disabling QoS Zone-Based Traffic Prioritization.....	324
Supported Configurations for QoS Zone-Based Traffic Prioritization.....	325
Limitations and Restrictions for QoS Zone-Based Traffic Prioritization.....	325
CS_CTL-Based Frame Prioritization.....	325
Supported Configurations for CS_CTL-Based Frame Prioritization.....	326
Enabling CS_CTL-Based Frame Prioritization on Ports.....	326
Disabling CS_CTL-Based Frame Prioritization on Ports.....	327
Using CS_CTL Auto Mode at the Chassis Level.....	327
Considerations for Using CS_CTL-Based Frame Prioritization.....	328
Managing Trunking Connections.....	329
Types of Trunking.....	329
Controllerless Trunking.....	329
License Requirements for Trunking.....	330
Port Groups for Trunking.....	330
Supported Platforms for Trunking.....	330
Supported Configurations for Trunking.....	330
ISLs between G630 Switches and Other Devices.....	331
QSFP-to-QSFP ISLs.....	331

QSFP-to-SFP ISLs.....	333
High-Availability Support for Trunking.....	335
Requirements for Trunk Groups.....	335
Recommendations for Trunk Groups.....	336
Configuring Trunk Groups.....	336
Enabling Trunking.....	337
Disabling Trunking.....	337
Displaying Trunking Information.....	338
ISL Trunking over Long-Distance Fabrics.....	339
EX_Port Trunking.....	340
Controllerless EX_Port Trunking.....	340
Supported Configurations and Platforms for EX_Port Trunking.....	340
Backward Compatibility Support.....	341
Configuring EX_Port Trunking.....	341
Displaying EX_Port Trunking Information.....	341
F_Port Trunking.....	341
F_Port Trunking for an Access Gateway.....	342
Requirements for F_Port Trunking on an Access Gateway.....	343
Configuring F_Port Trunking for an Access Gateway.....	343
F_Port Trunking for Brocade Adapters.....	344
Configuring F_Port Trunking for a Brocade Adapter.....	344
F_Port Trunking Considerations.....	344
F_Port Trunking in Virtual Fabrics.....	346
Displaying F_Port Trunking Information.....	346
Disabling F_Port Trunking.....	347
Enabling the DCC Policy on a Trunk Area.....	347
Inter-Chassis Links.....	348
License Requirements for ICLs.....	348
ICLs between Brocade X6 Directors.....	348
ICL Trunking between Brocade X6 Directors.....	351
Using QSFPs That Support 2 km on ICL Ports.....	352
ICLs between X7 Directors.....	352
ICL Trunking between Brocade X7 Directors.....	356
ICLs between an X7 Director and an X6 Director.....	359
ICL Trunking between an X7 Director and an X6 Director.....	360
ICLs between a Gen 7 Director and Gen 5 Directors.....	362
Virtual Fabrics Considerations for ICLs.....	362
Supported Topologies for ICL Connections.....	362
Mesh Topology.....	363
Core-Edge Topology.....	365

Managing Long-Distance Fabrics.....	367
Extended Fabrics Device Limitations.....	367
Long-Distance Link Modes.....	367
Configuring an Extended ISL.....	368
Enabling Long Distance When Connecting to TDM Devices.....	370
Forward Error Correction on Long-Distance Links.....	370
Enabling FEC on a Long-Distance Link.....	370
Disabling FEC on a Long-Distance Link.....	371
Configuring Long-Distance Support for G730 Switch.....	371
Configuration of Long-Distance Support On a G730 Switch.....	371
Configuring NVMe.....	378
NVMe over Fibre Channel.....	379
NVMe Name Server Support.....	380
NVMe Devices and Zoning.....	381
Administering Flow Vision.....	382
Flow Vision Features.....	382
Flow Vision Licensing.....	383
Firmware Upgrade and Downgrade Considerations for Flow Vision.....	383
Firmware Upgrade Considerations.....	383
Firmware Downgrade Considerations.....	383
Configuration File Upload and Download Considerations.....	383
Flow Vision Configuration Setup.....	384
High Availability and Flow Vision.....	384
Flow Vision Flows.....	384
Flow Definitions.....	385
Flow Definition Parameters and Rules.....	385
User-Defined Flows.....	386
Static versus Learned Flows.....	386
Predefined Flows.....	386
Resetting Flow Statistics.....	387
Clearing All Flow Statistics.....	387
Clearing the Statistics for Specific Flow Features.....	387
Number of Flows Supported.....	387
Flow Vision Integration with MAPS.....	388
System Event Handling.....	389
Flow Vision Limitations.....	390
Identifying All Zoned Devices in a Flow.....	391
Examples of show Command Output for Zoned Devices.....	391
Limitations on Viewing Zoned Devices.....	393

Flow Monitor.....	393
Automated Flow Monitoring.....	394
Support for Multiple Virtual Fabrics.....	396
Flow Metrics.....	397
Flow-Level Errors and Exception Metrics.....	400
Flow Granularity.....	401
Viewing Flows.....	402
Flow Monitoring with IO Insight Capability.....	416
I/O Statistics Counters.....	416
Supported Flow Types for SCSI I/O Latency Statistics.....	417
SCSI Exception Metrics.....	417
SCSI I/O Operation Classification.....	418
Example of IT Flow Metrics.....	420
Supported NVMe I/O Latency Statistics.....	421
Flow Monitoring with VM Insight Capability.....	424
VM Insight Requirements.....	425
Flow Definition Metrics for VM Insight.....	425
Flow Monitoring with VMID+.....	425
Working with Flow Monitor.....	427
Streaming Monitored Flows.....	455
Integration with MAPS.....	456
Integration with SANnav.....	456
Flow Monitor Limitations.....	456
Flow Generator.....	457
Flow Generator Setup.....	458
Predefined Flow Generator Flows.....	458
Notes on Predefined Flows.....	460
Creating Flow Generator Flows.....	460
Parameter Usage Exceptions.....	461
Creating an Inactive Flow in Flow Generator.....	461
Activating Flow Generator Flows.....	461
Automatic Activation of a Flow Generator Flow.....	462
Learning in Flow Generator Flows.....	462
Flow Generator Learning Flows on Egress Ports.....	462
Viewing Flow Generator Flows.....	462
Displaying the Status of a Single Flow Generator Flow.....	463
Viewing the Output of a Learned Flow Generator Flow.....	463
Notes on Displaying the Status of a Flow Generator Flow.....	463
Repeating Flow Generator Output.....	464
Sorting Flow Generator Output.....	465

Deactivating Flow Generator Flows.....	466
Customizing Flow Generator Flows.....	466
Frame Payload Size.....	466
Frame Payload Pattern.....	467
Flow Generator Examples.....	467
Creating a Flow from a Specific Source ID to a Specific Destination ID.....	467
Commands Related to Flow Generator.....	468
SIM Port Attributes and Configuration.....	470
SIM Port Criteria.....	470
Identifying SIM Ports.....	471
Sending Traffic Using a Fabric-Assigned WWN.....	472
Flow Generator and High Availability.....	472
Flow Generator and MAPS.....	472
Flow Generator Limitations and Considerations.....	472
Flow Mirror.....	473
Creating Flow Mirror Flows.....	474
Flow Mirror Limitations and Restrictions.....	475
Local Flow Mirroring.....	477
Creating an Inactive Flow in Flow Mirror.....	480
Remote Flow Mirroring.....	480
CFM and LFM Flow Support for Trunk Groups.....	481
Activating Flow Mirror Flows.....	482
Automatic Activation of a Flow Mirror Flow.....	482
Viewing Flow Mirror Flows.....	482
Summary Information View of a Flow Mirror Flow.....	482
Verbose Information View of a Flow Mirror Flow.....	483
Viewing a Flow Mirror Flow in Time Blocks.....	484
Learning in Flow Mirror Flows.....	486
Repeating Flow Mirror Output.....	486
Sorting Flow Mirror Output.....	488
Deactivating Flow Mirror Flows.....	490
Customizing Flow Mirror CFM Flow Frame Retention.....	490
Mirroring Traffic Flowing to Remote Fabrics.....	491
Using Flow Mirror with Access Gateway.....	492
Troubleshooting Using Flow Mirror.....	492
Tracking Latency between a Host and All Connected Targets.....	493
Diagnosing Excessive SCSI Reserve and Release Activity.....	494
Diagnosing a Slow-Draining F_Port.....	494
Flow Mirror and High Availability.....	495
Understanding the Fabric Flow Dashboard.....	495

Viewing Fabric Information with the Fabric Flow Dashboard.....	497
Limitations and Restrictions on Fabric Flow Dashboard Flows.....	500
Fabric Flow Dashboard Use Cases.....	500
Managing Virtual Fabrics.....	501
Logical Switch Overview.....	501
Default Logical Switch.....	501
Logical Switches and Fabric IDs.....	503
Port Assignment in Logical Switches.....	504
Logical Switches and Connected Devices.....	506
Management Model for Logical Switches.....	507
Logical Fabric Overview.....	507
Logical Fabric and ISLs.....	508
Base Switch and Extended ISLs.....	509
Base Fabric.....	512
Logical Ports.....	512
Logical Fabric Formation.....	513
Account Management and Virtual Fabrics.....	513
Setting Up IP Addresses for a Logical Switch.....	513
Logical Switch Login Context.....	514
Supported Platforms for Virtual Fabrics.....	515
Supported Port Configurations in Fixed-Port Switches.....	516
Supported Port Configurations in Brocade Directors.....	516
Virtual Fabrics Interaction with Other Fabric OS Features.....	516
Limitations and Restrictions of Virtual Fabrics.....	517
Restrictions on Using XISLs.....	518
Restrictions on Moving Ports.....	518
Enabling Virtual Fabrics Mode.....	518
Disabling Virtual Fabrics Mode.....	519
Configuring Logical Switches to Use Basic Configuration Values.....	520
Creating a Logical Switch or Base Switch.....	520
Executing a Command in a Different Logical Switch Context.....	522
Deleting a Logical Switch.....	523
Adding and Moving Ports on a Logical Switch.....	523
Displaying the Logical Switch Configuration.....	524
Changing the Fabric ID of a Logical Switch.....	525
Changing a Logical Switch to a Base Switch.....	526
Configuring a Logical Switch for XISL Use.....	527
Changing the Context to a Different Logical Fabric.....	527
Administering Advanced Zoning.....	529

Approaches to Zoning.....	530
Zone Types.....	531
Zone Objects.....	531
Naming Zone Objects.....	532
Zoning Schemes.....	532
Zone Configurations.....	533
Zoning Enforcement.....	533
Identifying the Enforced Zone Type.....	534
Considerations for Zoning Architecture.....	534
Best Practices for Zoning.....	535
Zone Aliases.....	535
Creating an Alias.....	536
Adding Members to an Alias.....	536
Removing Members from an Alias.....	537
Deleting an Alias.....	537
Viewing an Alias in the Defined Configuration.....	538
Zone Creation and Maintenance.....	538
Displaying Existing Zones.....	538
Creating a Zone.....	539
Adding Devices (Members) to a Zone.....	540
Removing Devices (Members) from a Zone.....	542
Replacing Zone Members.....	544
Notes and Restrictions.....	544
Deleting a Zone.....	545
Viewing a Zone.....	546
Viewing Zone Configuration Names without Case Distinction.....	548
Examining Changes in the Zone Database.....	548
Validating a Zone.....	549
Inconsistencies between the Defined and Effective Configurations.....	551
Default Zoning Mode.....	552
Setting the Default Zoning Mode.....	552
Viewing the Current Default Zone Access Mode.....	553
Zone Database Size.....	553
Zone Configurations.....	555
Creating a Zone Configuration.....	555
Adding Zones to a Zone Configuration.....	556
Removing Members from a Zone Configuration.....	556
Enabling a Zone Configuration.....	557
Disabling a Zone Configuration.....	557
Validating Zone Members in the Zone Configuration.....	558

Deleting a Zone Configuration.....	560
Deleting Invalid Zones.....	560
Abandoning Zone Configuration Changes.....	560
Zone Fabric Locking.....	561
Viewing All Zone Configuration Information.....	564
Viewing Selected Zone Configuration Information.....	565
Viewing the Zone Aliases in the Zone Configuration.....	565
Viewing the Configuration in the Effective Zone Database.....	567
Clearing All Zone Configurations.....	567
Zone Object Group Member Sorting.....	568
Zone Object Maintenance.....	569
Copying a Zone Object.....	569
Deleting a Zone Object.....	570
Renaming a Zone Object.....	571
Zone Configuration Management.....	571
Security and Zoning.....	572
Zone Merging.....	572
Fabric Segmentation and Zoning.....	573
Zone Merging Scenarios.....	574
Concurrent Zone Transactions.....	576
Viewing Zone Database Transactions.....	577
Peer Zoning.....	578
Peer Zoning Compared to Other Zoning Types.....	579
Advantages of Peer Zoning.....	581
Peer Zone Connectivity Rules.....	581
Firmware Upgrade and Downgrade Considerations for Peer Zoning.....	581
LSAN and QoS Peer Zoning Considerations.....	581
Peer Zone Configuration.....	582
Creating Peer Zones.....	582
Adding Devices to a Peer Zone.....	582
Replacing a Device in a Peer Zone.....	582
Removing Devices from a Peer Zone.....	583
Deleting Peer Zones.....	583
Viewing Peer Zones.....	583
Alias Support for Peer Zoning.....	585
Displaying the Property Members of Peer Zones.....	588
Target-Driven Zoning.....	588
Limitations and Considerations for Target-Driven Zoning.....	589
Target-Driven Zoning Configuration.....	589
Enabling Target-Driven Zoning.....	590

Disabling Target-Driven Zoning.....	590
Viewing Target-Driven Zoning.....	591
Alias Support for Target-Driven Zoning.....	593
Supported Commands for Peer Zones and Target-Driven Peer Zones.....	594
Boot LUN Zoning.....	595
Setting Up Boot LUN Zoning.....	596
Configuring the Server and HBA for a BLUN Zone.....	596
Configuring the Brocade Switch for a BLUN Zone.....	597
Routing Traffic.....	599
Paths and Route Selection.....	599
FSPF.....	599
Fibre Channel NAT.....	601
Inter-Switch Links.....	601
Buffer Credits.....	603
Virtual Channels.....	603
Gateway Links.....	604
Configuring a Link through a Gateway.....	605
Routing Policies.....	606
Considerations for Routing Policies.....	606
Displaying the Current Routing Policy.....	606
Port-Based Routing.....	606
Exchange-Based Routing.....	607
Device-Based Routing.....	607
Dynamic Path Selection.....	607
Displaying a Dynamic Path Selection Group for Reachable Domains.....	607
Route Selection.....	608
Dynamic Load Sharing.....	608
Setting DLS.....	609
Frame Order Delivery.....	609
Forcing In-Order Frame Delivery across Topology Changes.....	610
Restoring Out-of-Order Frame Delivery across Topology Changes.....	610
Enabling Frame Viewer.....	610
Using Frame Viewer to Understand Why Frames Are Dropped.....	611
Displaying Discarded Frames by a Backend Port in Frame Viewer.....	612
Using Frame Viewer for Fabric Notifications.....	612
Lossless Dynamic Load Sharing on Ports.....	612
Lossless Core.....	613
ICL Limitations.....	613
Configuring Lossless Dynamic Load Sharing.....	614
Lossless Dynamic Load Sharing in Virtual Fabrics.....	614

How DLS Affects Other Logical Switches in the Fabric.....	614
Two-Hop Lossless DLS Route Update.....	615
Configuring Two-Hop Lossless DLS.....	615
Using FC-FC Routing to Connect Fabrics.....	617
License Requirements for FC-FC Routing.....	617
Supported Platforms for FC-FC Routing.....	617
Configuring EX_Ports on QSFP Ports of a Brocade G630 Switch.....	618
Supported Configurations for FC-FC Routing.....	619
Fibre Channel Routing Concepts.....	620
Proxy Devices.....	623
FC-FC Routing Topologies.....	625
Phantom Domains.....	625
Identifying and Deleting Stale Xlate Domains.....	628
Domain ID Range for Front and Translate Phantom Domains.....	628
FC Router Authentication.....	629
Setting Up FC-FC Routing.....	629
Verifying the Setup for FC-FC Routing.....	629
Backbone Fabric IDs.....	631
Assigning Backbone Fabric IDs.....	631
Assigning Alias Names to Fabric IDs.....	632
FCIP Tunnel Configuration.....	632
Inter-Fabric Link Configuration.....	633
Configuring an IFL for Both Edge and Backbone Connections.....	633
Configuring EX_Ports on an ICL.....	636
FC Router Port Cost Configuration.....	639
Port Cost Considerations.....	639
Setting the Router Port Cost for an EX_Port.....	640
Shortest IFL Cost Configuration.....	640
Configuring the Shortest IFL Cost.....	642
EX_Port Trunking Configuration.....	644
LSAN Zone Configuration.....	645
Zone Definition and Naming.....	645
LSAN Zones and Fabric-to-Fabric Communications.....	645
Controlling Device Communication with the LSAN.....	645
Configuring Interconnectivity between Backbone and Edge Fabrics.....	649
Setting the Maximum LSAN Count.....	649
HA and Downgrade Considerations for LSAN Zones.....	650
LSAN Zone Policies Using LSAN Tagging.....	651
Enforce Tag.....	651
Speed Tag.....	653

Rules for LSAN Tagging.....	654
Configuring an Enforce LSAN Tag.....	655
Configuring a Speed LSAN Tag.....	655
Removing an LSAN Tag.....	655
Displaying the LSAN Tag Configuration.....	656
LSAN Zone Binding.....	656
LSAN Zone Binding Considerations.....	658
How LSAN Zone Binding Works.....	658
FC Router Matrix Definition.....	658
LSAN Fabric Matrix Definition.....	659
Setting Up LSAN Zone Binding.....	659
Viewing the LSAN Zone Binding Matrices.....	660
Location-Embedded LSAN Zones.....	660
Creating Location-Embedded LSAN Zones.....	660
Migrating LSAN Zones to Location-Embedded LSAN Zones in the Edge Fabric.....	662
Limitations of Location-Embedded LSAN Zones.....	663
Peer LSAN Zone Support.....	663
Proxy PID Configuration.....	663
Fabric Parameter Considerations.....	663
Inter-Fabric Broadcast Frames.....	664
Displaying the Current Broadcast Configuration.....	664
Enabling Broadcast Frame Forwarding.....	664
Disabling Broadcast Frame Forwarding.....	664
Resource Monitoring.....	664
FC-FC Routing and Virtual Fabrics.....	666
Logical Switch Configuration for FC Routing.....	667
Backbone-to-Edge Routing with Virtual Fabrics.....	668
Upgrade and Downgrade Considerations for FC-FC Routing.....	669
How Replacing Port Blades Affects EX_Port Configuration.....	669
Displaying the Range of Output Ports Connected to xlate Domains.....	669
FCR Migration Mode.....	670
Administering FICON Fabrics.....	672
FICON Requirements.....	672
FICON Best Practices.....	672
FICON and FEC-via-TTS.....	674
Considerations for a FICON Logical Switch.....	674
Preparing a Switch for FICON.....	675
Creating a FICON Logical Switch.....	675
Configuring FICON.....	677
Converting an Existing Logical Switch to a FICON Logical Switch.....	683

Enabling the Insistent Domain ID.....	685
Creating and Activating the SCC Policy.....	685
Enabling the Fabric-Wide Consistency Policy.....	685
Enabling High-Integrity Fabric Mode.....	685
FCR and FICON Cascading.....	686
FICON and FICON CUP.....	686
Blocking and Unblocking Ports.....	687
Clearing the FICON Management Database.....	688
Automating CS_CTL Mapping.....	688
Latency Guideline.....	690
FICON CUP Overview.....	690
FICON CUP Restrictions.....	691
CUP Configuration Recommendations.....	691
Zoning and PDCM Considerations.....	692
Port and Switch Naming Standards for FMS Mode.....	692
FICON CUP Fabric OS Commands.....	692
Configuring FICON CUP.....	693
Determining Physical Port Assignment.....	694
FMS Mode and FICON CUP.....	694
Fabric OS Command Limitations and Considerations.....	694
Displaying FMS Mode.....	695
Enabling FMS Mode.....	695
Disabling FMS Mode.....	695
Upgrade Considerations.....	695
Mode Register Bit Settings.....	696
Considerations for Setting the Mode Register Bits.....	696
Setting the Mode Register Bits.....	697
Setting the MIHPTO Value.....	697
Persistently Enabling and Disabling Ports for CUP.....	697
Administering FICON Extension Services.....	698
Platforms Supporting FICON Extension over IP.....	698
FICON Emulation Overview.....	698
IBM z/OS Global Mirror Emulation.....	698
Tape Emulation.....	700
Teradata Emulation.....	701
FCIP Configuration Requirements for FICON Extension.....	702
Configuration Requirements for Switches and Directors.....	702
High-Integrity Fabric Requirements for Cascaded Configurations.....	702
FICON Emulation Requirement for a Determinate Path.....	702
Configuring FICON Emulation.....	703

Configuration Examples.....	704
Displaying FICON Emulation Configuration Values.....	704
Modifying FICON Emulation.....	704
Displaying FICON Emulation Performance Statistics.....	705
FICON Emulation Monitoring.....	705
Options for Displaying Statistics.....	706
Platforms Supporting FICON.....	707
Currently Supported Platforms.....	707
Supported Brocade Blades.....	708
Port Indexing.....	709
Brocade Switch and Director Sensors.....	716
Brocade Blade Temperature Sensors.....	716
System Temperature Monitoring.....	717
Appendix A: Setting Up the AAA Server Configuration.....	718
Setting Up a RADIUS Server.....	718
Configuring RADIUS Server Support with Linux.....	718
Adding the Brocade Attributes to the Server.....	718
Creating the User.....	719
Enabling Clients.....	719
Configuring RADIUS Server Support with Windows 2012.....	720
LDAP Configuration and Microsoft Active Directory.....	723
Configuring Microsoft Active Directory LDAP Service.....	724
Adding a Brocade VSA in the Windows AD Schema.....	725
Adding a Virtual Fabric List.....	728
LDAP Configuration and OpenLDAP.....	729
Supported LDAP Options.....	729
OpenLDAP Server Configuration Overview.....	729
Enabling Group Membership.....	730
Adding Entries to the Directory.....	731
Assigning a User to a Group.....	731
Modifying an Entry.....	732
TACACS+ Service.....	732
TACACS+ Configuration Overview.....	732
Configuring the TACACS+ Server on Linux.....	732
Configuring a Windows TACACS+ Server.....	733
Appendix B: Configuring Factory Reset.....	734
Supported Switches and Directors.....	734
FOS Configurations That Support Fabric Reset.....	734
Resetting a Device to Factory Default Settings.....	734

Appendix C: Supported Cipher Suites in Fabric OS	738
Appendix D: FICON	744
Basic Switch Configuration	744
Address Binding Examples	746
Sequential Address Binding	746
Binding Ports Sequentially (Mode 1)	747
Unbinding Multiple Ports	748
Configuration Information Record	750
EBCDIC Page	750
Appendix E: Hexadecimal Conversion	752
Example Conversion of the Hexadecimal Triplet 0x610600	752
Decimal-to-Hexadecimal Conversion Table	753
Revision History	755
Documentation Legal Notice	756

Introduction

This document describes the administration of Brocade® Fibre Channel Services (FCS) features and the configuration tasks of Brocade storage area networking (SAN) products that use the Fabric OS® command line interface (CLI).

This FOS administration guide describes general Fabric OS configuration tasks, security, and user and switch management. Refer to the following guides for information on more specific functions and management.

- *Brocade Fabric OS Software Licensing User Guide*
- *Brocade Fabric OS Software Upgrade User Guide*
- *Brocade Fabric OS Command Reference Manual*
- *Brocade Fabric OS Message Reference Manual*
- *Brocade Fabric OS Extension User Guide*
- *Brocade Fabric OS MAPS User Guide*
- *Brocade Fabric OS Web Tools User Guide*
- *Brocade Fabric OS Access Gateway User Guide*
- *Brocade Fabric OS MIB Reference Manual*
- *Brocade Fabric OS REST API Reference Manual*
- *Brocade Fabric OS FCoE User Guide*

Supported Hardware and Software

The following hardware platforms are supported by Brocade Fabric OS 9.1.x.

Brocade Gen 7 (64G) Fixed-Port Switches

- Brocade G720 Switch
- Brocade G730 Switch

Brocade Gen 7 (64G) Directors

- Brocade X7-4 Director
- Brocade X7-8 Director

Brocade Gen 6 (32G) Fixed-Port Switches

- Brocade G610 Switch
- Brocade G620 Switch
- Brocade G630 Switch
- Brocade 7810 Extension Switch
- Brocade G648 Blade Server SAN I/O Module
- Brocade MXG610 Blade Server SAN I/O Module

Brocade Gen 6 (32G) Directors

- Brocade X6-4 Director
- Brocade X6-8 Director

Contacting Technical Support for Your Brocade® Product

If you purchased Brocade® product support from a Broadcom® OEM or solution provider, contact your OEM or solution provider for all your product support needs.

- OEM and solution providers are trained and certified by Broadcom to support Brocade products.
- Broadcom provides backline support for issues that cannot be resolved by the OEM or solution provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information on this option, contact Broadcom or your OEM.
- For questions regarding service levels and response times, contact your OEM or solution provider.

If you purchased Brocade product support directly from Broadcom, use one of the following methods to contact the Technical Assistance Center 24x7. For product support information and the latest information on contacting the Technical Assistance Center, go to www.broadcom.com/support/fibre-channel-networking/contact-brocade-support.

Online	Telephone
For nonurgent issues, the preferred method is to log on to the Support portal at support.broadcom.com . (You must initially register to gain access to the Support portal.) Once registered, log on and then select Hardware > Brocade Storage Networking . You can now navigate to any of the Brocade tools.	For Severity 1 (critical) issues, call Brocade Fibre Channel Networking Global Support at one of the phone numbers listed at www.broadcom.com/support/fibre-channel-networking/contact-brocade-support .


Document Feedback

Quality is our first concern. We have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission or if you think that a topic needs further development, we want to hear from you. Send your feedback to documentation.pdl@broadcom.com. Provide the publication title; topic heading; publication number and page number (for PDF documents); URL (for HTML documents); and as much detail as possible.

Understanding Fibre Channel Services

This section describes about Fibre Channel that defines the service function resides at well-known addresses. A *well-known address* is a reserved 3-byte address for each service. Services are provided to either nodes or management applications in the fabric.

Figure 1: Well-Known Addresses



- Fabric log in – The fabric login server assigns a fabric address to a fabric node, which allows it to communicate with services on the switch or other nodes in the fabric. The fabric address is a 24-bit address (0x000000) containing three 3-byte nodes. Reading from left to right, the first node (0x 000000) represents the domain ID, the second node (0x 000000) represents the port area number of the port where the node is attached, and the third node (0x 000000) represents the arbitrated loop physical address (AL_PA), if applicable.
- Directory server – The directory server or name server registers fabric and public nodes and conducts queries to discover other devices in the fabric.
- Fabric controller – The fabric controller provides State Change Notifications (SCNs) to registered nodes when a change in the fabric topology occurs.
- Time server – The time server sends the time to the member switches in the fabric from either the principal switch or, if configured, the primary fabric configuration server switch. See [Configuring Security Policies](#) for additional information on FCS policies.
- Management server – The management server provides a single point for managing the fabric. The Management server is the only service that users can configure. See [Management server](#) for more details.
- Alias server – The alias server keeps a group of nodes that are registered as one name to handle multicast groups.
- Broadcast server – The broadcast server is optional. When frames are transmitted to this address, they are broadcast to all operational N_Ports and NL_Ports.

When registration and query frames are sent to a well-known address, a different protocol service, Fibre Channel Common Transport (FC-CT), is used. This protocol provides a simple, consistent format and behavior when a service provider is accessed for registration and query purposes.

Management Server

The Brocade Fabric OS management server (MS) allows a SAN management application to retrieve information and administer interconnected switches, servers, and storage devices. The management server assists in the autodiscovery of switch-based fabrics and their associated topologies.

A client of the management server can find basic information about the switches in the fabric and use this information to construct topology relationships. The management server also allows you to obtain certain switch attributes and, in some cases, modify them. For example, logical names that identify switches can be registered with the management server.

The management server provides several advantages for managing a Fibre Channel fabric:

- The management server is accessed by an external Fibre Channel node at the well-known address FFFFFAh, so an application can access information about the entire fabric management with minimal knowledge of the existing configuration.
- It is replicated on every Brocade switch within a fabric.
- It provides an unzoned view of the overall fabric configuration. This fabric topology view exposes the internal configuration of a fabric for management purposes; it contains interconnect information about switches and devices connected to the fabric. Under normal circumstances, a device (typically an FCP initiator) queries the name server for storage devices within its member zones. Because this limited view is not always sufficient, the management server provides the application with a list of the entire name server database.

NOTE

Management server command queries are not supported for ETH port configuration.

Application Server

The application server is a key component of the VM Insight flow monitoring feature for Flow Vision.

As a component of the management server, the application server stores entity IDs for VM instances registered by HBAs and allocates application IDs to HBAs for identifying VM traffic in the fabric. With traffic for individual VM flows identified, monitors can be configured to provide flow statistics to help isolate VM performance issues.

The application server is a distributed database, similar to the name server and management server, and it is propagated to all switches in the fabric that support the application server. The application server is supported only by Gen 6 platforms running Fabric OS 8.1.0 or later and by Gen 7 platforms.

You can use the `appserver --show` command to display contents of the application server database, including N_Port IDs, application IDs, port IDs, and domain information associated with all registered VM entity IDs. Refer to the *Brocade Fabric OS Command Reference Manual* for details.

Platform Services

By default, all management services except platform services are enabled; the MS platform service and topology discovery are disabled.

You can activate and deactivate the platform services throughout the fabric. Activating the platform services attempts to activate the MS platform service for each switch in the fabric. The change takes effect immediately and is committed to the configuration database of each affected switch. MS activation is persistent across power cycles and reboots.

Platform Services and Virtual Fabrics

Each logical switch has a separate platform database. All platform registrations done to a logical switch are valid only in that particular logical switch's virtual fabric.

Activating the platform services on a switch activates the platform services on all logical switches in a virtual fabric. Similarly, deactivating the platform services deactivates the service on all logical switches in a virtual fabric. The `msplatshow` command displays all platforms registered in a virtual fabric.

Enabling Platform Services

When the FCS policy is enabled, the `msplmgtactivate` command can be issued only from the primary FCS switch.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `mscapabilityshow` command to verify that all switches in the fabric support the MS platform service; otherwise, the next step fails.
3. Enter the `msplmgtactivate` command, as in the following example.

```
switch:admin> msplmgtactivate
Request to activate MS Platform Service in progress.....
*Completed activating MS Platform Service in the fabric!
```

Disabling Platform Services

Use the following procedure to disable platform services.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `msplmgtdeactivate` command.
3. Enter `y` to confirm the deactivation, as in the following example.

```
switch:admin> msplmgtdeactivate

MS Platform Service is currently enabled.
This will erase MS Platform Service configuration
information as well as database in the entire fabric.
Would you like to continue this operation? (yes, y, no, n): [no] y
2
Request to deactivate MS Platform Service in progress.....
*Completed deactivating MS Platform Service in the fabric!
```

Management Server Database

You can control access to the management server database.

An access control list (ACL) of WWN addresses determines which systems have access to the management server database. The ACL typically contains those WWNs of host systems that are running management applications.

If the list is empty (the default), the management server is accessible to all systems connected in-band to the fabric. For more access security, you can specify WWNs in the ACL so that access to the management server is restricted to only those WWNs listed.

NOTE

The management server is logical-switch-capable. All management server features are supported within a logical switch.

Displaying the Management Server ACL

Use the following procedure to display the management server ACL.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `msconfigure` command.

The command becomes interactive.

3. At the `select` prompt, enter 1 to display the access list.

A list of WWNs that have access to the management server is displayed.

The following is an example of an empty access list.

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 1
MS Access list is empty.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
```

Adding a Member to the ACL

Use the following procedure to add a member to the ACL.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `msconfigure` command.
The command becomes interactive.
3. At the `select` prompt, enter 2 to add a member based on its port/node WWN.
4. At the `Port/Node WWN` prompt, enter the WWN of the host to be added to the ACL.
5. At the `select` prompt, enter 1 to display the access list. You can verify that the WWN has been added to the ACL.
6. After verifying that the WWN is added correctly, enter 0 at the prompt to end the session.
7. At the `Update the FLASH?` prompt, enter `y`.
8. Press `Enter` to update the nonvolatile memory and end the session.

The following is an example of adding a member to the management server ACL.

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 2
Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully added to the MS ACL.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (14): {
  20:00:00:20:37:65:ce:aa
```

```

20:00:00:20:37:65:ce:bb
20:00:00:20:37:65:ce:ff
20:00:00:20:37:65:ce:11
20:00:00:20:37:65:ce:22
20:00:00:20:37:65:ce:33
20:00:00:20:37:65:ce:44
10:00:00:60:69:04:11:24
10:00:00:60:69:04:11:23
21:00:00:e0:8b:04:70:3b
10:00:00:60:69:04:11:33
20:00:00:20:37:65:ce:55
20:00:00:20:37:65:ce:66
00:00:00:00:00:00:00:00
}
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH? (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.

```

Deleting a Member from the ACL

When you delete a member from the ACL, the member will not be able to access the management server.

NOTE

If you delete the last member of the ACL, leaving the ACL list empty, the management server will be accessible to all systems connected in-band to the fabric.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `msconfigure` command.
The command becomes interactive.
3. At the `select` prompt, enter 3 to delete a member based on its port/node WWN.
4. At the `Port/Node WWN` prompt, enter the WWN of the member to be deleted from the ACL.
5. At the `select` prompt, enter 1 to display the access list so you can verify that the WWN you entered was deleted from the ACL.
6. After verifying that the WWN was deleted correctly, enter 0 at the `select` prompt to end the session.
7. Press `Enter` to update the nonvolatile memory and end the session.

The following is an example of deleting a member from the management server ACL.

```

switch:admin> msconfigure

0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 3

Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 10:00:00:00:c9:29:b3:84

```

```
*WWN is successfully deleted from the MS ACL.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [3] 1
```

```
MS Access list is empty
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
```

Viewing the Contents of the Management Server Database

Use the following procedure to view the contents of the management server database.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `msplatshow` command.

The following example shows the contents of the management server platform database.

```
switch:admin> msplatshow

-----
Platform Name: [9] "first obj"
Platform Type: 5 : GATEWAY
Number of Associated M.A.: 1
[35] "http://java.sun.com/products/plugin"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:71
-----

Platform Name: [10] "second obj"
Platform Type: 7 : HOST_BUS_ADAPTER
Number of Associated M.A.: 1
Associated Management Addresses:
[30] "http://java.sun.com/products/1"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:75
```

Clearing the Management Server Database

Use the following procedure to clear the management server database.

1. Connect to the switch and log in using an account that is assigned to the admin role.
2. Enter the `msplcleardb` command.
3. Enter `y` to confirm the deletion.

The management server platform database is cleared.

Supporting the New Management Server ACL Interface

From Fabric OS 9.1.0 and above, a new CLI command `portcfgmsacl` is introduced to configure the Management Services ACL (MS ACL). Access to the Fabric Configuration Server (FCS) and Enhanced Fabric Configuration Server (EFCS) depends upon both the `msconfigure` and `portcfgmsacl` interfaces.

NOTE

Read-write access of the MS ACL is defined per port (and applies to all devices behind the port) and is defined per management service.

The new Management Server ACL interface `portcfgmsacl` supports the following management servers:

- Application Server
- Enhanced Fabric Configuration Server
- Fabric Configuration Server
- Fabric Device Management Interface (FDMI)
- Fabric Zone Server
- Unzoned Name Server

Configuring the Management Server ACL

A new CLI command `portcfgmsacl` is used to provide access to different management servers. The following table highlights the service and short arguments for each of the configured management servers.

Table 1: Management Server Arguments

Service Argument	Short Argument	Service Name
<code>application_server</code>	<code>appsrvr</code>	Application Server
<code>enhanced_fabric_configuration_server</code>	<code>efcs</code>	Enhanced Fabric Configuration Server
<code>fabric_configuration_server</code>	<code>fcs</code>	Fabric Configuration Server
<code>fabric_device_management_interface</code>	<code>fdmi</code>	Fabric Device Management Interface
<code>fabric_zone_server</code>	<code>zone</code>	Fabric Zone Server
<code>unzoned_name_server</code>	<code>unzn</code>	Unzoned Name Server

The following table displays access subargument options that correspond to the read, read-write, or no access of the given management service.

Table 2: Accessing the Management Server

Service Argument	Short Argument
<code>no_access</code>	<code>no</code>
<code>read_only</code>	<code>ro</code>
<code>read_write</code>	<code>rw</code>

NOTE

The default configuration for Unzoned Name Server is read-only mode. The default configuration for the remaining services is read-write access mode.

The following examples show how to configure the Management Server ACL using the `portcfgmsacl` command.

```

switch:admin> portcfgmsacl --configure application_server -access no_access 10
switch:admin> portcfgmsacl --configure appsrvr -access no 10
switch:admin> portcfgmsacl --configure fcs -access rw 9

```

The following examples display the Management Server ACL access in different ports using the `portcfgshow` command.

```

switch:admin> portcfgshow 0
Area Number:          0
FC Speed Level:       AUTO (HW)
AL_PA Offset 13:      OF
Trunk Port            ON
....
....
....
10G/16G FEC:         ON
16G FEC via TTS:      OFF
App Header:          OFF
App Srvr Access:      RO
EFCS Access:         RW
FCS Access:          RW
FDMI Access:         RO
Zone Access:         RO
Unzoned NS Access:   NO

```

```
switch:admin> portcfgshow
```

Ports of Slot 0	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Speed	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN
AL_PA Offset 13	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Trunk Port	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
Long Distance	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Persistent Disable	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
LOS TOV mode	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
NPIV capability	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
NPIV PP Limit	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126
NPIV FLOGI Logout	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
QOS Port	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE
EX Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Mirror Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Credit Recovery	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
D-Port over DWDM	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Compression	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Encryption	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
10G/16G FEC	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
16G FEC via TTS	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Fault Delay	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
App Header	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
App Srvr Access	RO	RW	RW	RW	NO	RO	RW	RW	RW	RW	RW	RW	RO	RW	RW	RW

EFCS Access	RO	RW	RW	RW	NO	RO	RW	RW	RW	RW	RW	RW	RO	RW	RW	RW
FCS Access	RO	RW	RW	RW	NO	RO	RW	RW	RW	RW	RW	RW	RO	RW	RW	RW
FDMI Access	RO	RW	RW	RW	NO	RO	RW	RW	RW	RW	RW	RW	RO	RW	RW	RW
Zone Access	RO	RW	RW	RW	NO	RO	RW	RW	RW	RW	RW	RW	RO	RW	RW	RW
Unzoned NS Access	RO	RO	RO	RO	NO	RO	RO	RO	RO	RO	RO	RO	RO	RO	RO	RO

Topology Discovery

The topology discovery feature can be displayed, enabled, and disabled; it is disabled by default.

Displaying Topology Discovery Status

Use the following procedure to display the status of the topology discovery.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `mstdreadconfig` command.

```
switch:admin> mstdreadconfig
```

```
*MS Topology Discovery is Enabled.
```

Enabling Topology Discovery

Use the following procedure to enable topology discovery.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter one of the following commands based on how you want to enable discovery:
 - For the local switch, enter the `mstdenable` command.
 - For the entire fabric, enter the `mstdenable all` command.

The following is an example of enabling discovery.

```
switch:admin> mstdenable
```

```
Request to enable MS Topology Discovery Service in progress....
```

```
*MS Topology Discovery enabled locally.
```

```
switch:admin> mstdenable all
```

```
Request to enable MS Topology Discovery Service in progress....
```

```
*MS Topology Discovery enabled locally.
```

```
*MS Topology Discovery Enable Operation Complete!!
```

Disabling Topology Discovery

To disable topology discovery, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter one of the following commands based on how you want to disable discovery.
 - For the local switch, enter the `mstddisable` command.
 - For the entire fabric, enter the `mstddisable all` command.

A warning displays stating that all NID entries might be cleared.

3. Enter `y` to disable topology discovery.

NOTE

Topology discovery is disabled by default.

ATTENTION

Disabling discovery of the management server topology might erase all node ID entries.

Example of Disabling Discovery

The following is an example of disabling topology discovery.

```
switch:admin> mstddisable
This may erase all NID entries. Are you sure? (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.

switch:admin> mstddisable all
This may erase all NID entries. Are you sure? (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.
*MS Topology Discovery Disable Operation Complete!!
```

Device Login

A device can be storage, a host, or a switch. When new devices are introduced into the fabric, they must be powered on and, if they are a host or storage device, connected to a switch. Switch-to-switch logins (using the E_Port) are handled differently than storage and host logins. E_Ports exchange different frames than the following ones with the fabric controller to access the fabric. After storage and host devices are powered on and connected, the following logins occur:

1. FLOGI – The Fabric Login command establishes a 24-bit address for the device logging in and establishes buffer-to-buffer credits and the class of service supported.
2. PLOGI – The Port Login command logs the device into the name server to register its information and query for devices that share its zone. During the PLOGI process, information is exchanged between the new device and the fabric. Some of the following types of information exchanges occur:
 - SCR – State Change Registration registers the device for State Change Notifications. If a change in the fabric occurs, such as a zoning change or a change in the state of a device to which this device has access, the device receives a Registered State Change Notification (RSCN).
 - Registration – A device exchanges registration information with the name server.
 - Query – Devices query the name server for information about the device it can access.

Principal Switch

In a fabric with one or more switches interconnected by an inter-switch link (ISL) or inter-chassis links (ICLs), a principal switch is automatically elected. The principal switch provides the following capabilities:

- Maintains time for the entire fabric. Subordinate switches synchronize their time with the principal switch. Changes to the clock server value on the principal switch are propagated to all switches in the fabric.
- Manages domain ID assignment within the fabric. If a switch requests a domain ID that has been used before, the principal switch grants the same domain ID unless it is in use by another switch.

E_Port Login Process

An E_Port does not use a FLOGI to log in to another switch. Instead, the new switch exchanges frames with the neighboring switch to establish the E_Port of the new switch and that it has information to exchange. If everything is acceptable to the neighboring switch, it replies to the new switch with an SW_ACC (accept) frame. The initializing frame is an Exchange Link Parameters (ELP) frame that allows an exchange of parameters between two ports, such as flow control, buffer-to-buffer credits, RA_TOV, and ED_TOV. This is not a negotiation. If one or the other port's link parameters do not match, a link does not occur. When an SW_ACC frame is received from the neighboring switch, the new switch sends an Exchange Switch Capabilities (ESC) frame. The two switches exchange routing protocols and agree on a common routing protocol. An SW_ACC frame is received from the neighboring switch, and the new switch sends an Exchange Fabric Parameters (EFP) frame to the neighboring switch, requesting the principal switch priority and the domain ID list. Buffer-to-buffer credits for the device and switch ports are exchanged in the SW_ACC command sent to the device in response to the FLOGI.

Fabric Login Process

A device performs a fabric login (FLOGI) to determine if a fabric is present. If a fabric is detected, the device exchanges service parameters with the fabric controller. A successful FLOGI sends back the 24-bit address for the device in the fabric. The device must issue and successfully complete a FLOGI command before communicating with other devices in the fabric.

Because the device does not know its 24-bit address until after the FLOGI, the source ID (SID) in the frame header of the FLOGI request is zeros (0x000000).

Port Login Process

The steps in the port initialization process occur as the result of a protocol that discovers the type of device connected, establishes the port type, and negotiates the port speed. See [Port Types](#) for a discussion of available port types.

The Fibre Channel Protocol (FCP) auto discovery process enables private storage devices that accept the process login (PRLI) to communicate in a fabric.

If device probing is enabled, the embedded port performs a PLOGI and attempts a PRLI into the device to retrieve information to enter into the Name Server (NS). This enables the private devices that do not explicitly register and enters in the NS and receive full fabric access.

A fabric-capable device registers its information with the NS during a FLOGI. These devices typically register information with the NS before querying for a device list. The embedded port still performs a PLOGI and attempts a PRLI with these devices.

If a port decides to end the current session, it initiates a logout. A logout concludes the session and terminates any work in progress associated with that session.

To display the contents of the NS, use the `nsshow` or `nsallshow` commands. For more information, refer to *Brocade Fabric OS Command Reference Manual*.

RSCNs

A Registered State Change Notification (RSCN) is a notification frame that is sent to devices that are zoned together and are registered to receive a State Change Notification (SCN). The RSCN is responsible for notifying all devices of fabric changes. The following general list of actions can cause an RSCN to be sent through your fabric:

- A new device has been added to the fabric.
- An existing device has been removed from the fabric.
- A zone has changed.
- A switch name or an IP address has changed.
- Nodes have left or joined the fabric, such as zoning, powering on or shutting down a device, or zoning changes.

NOTE

Fabric reconfigurations with no domain change do not cause an RSCN.

Duplicate Port World Wide Name

According to Fibre Channel standards, the Port World Wide Name (PWWN) of a device cannot overlap with that of another device; thus having duplicate PWWNs within the same fabric is an illegal configuration.

If a PWWN conflict occurs with two devices attached to the same domain, Fabric OS handles device login in such a way that only one device may be logged in to the fabric at a time. For more information, see [Duplicate PWWN Handling during Device Login](#).

If a PWWN conflict occurs and two duplicate devices are attached to the fabric through different domains, the devices are removed from the NS database and a RASlog is generated.

Device Recovery

To recover devices that have been removed from the Name Server database due to duplicate PWWNs, you must log in the devices to the fabric again. This is true for any device—for example, a device on an F_Port, NPIV devices, or devices attached to a switch in Access Gateway (AG) mode.

High Availability of the Daemon Process

Starting noncritical daemons is automatic; you cannot configure the startup process. The following sequence of events occurs when a noncritical daemon fails:

1. A RASlog and an AUDIT event message are logged.
2. The daemon is automatically started again.
3. If the restart is successful, another message is sent to RASlog and AUDIT reports the successful restart status.
4. If the restart fails, another message is sent to RASlog and no further attempts are made to restart the daemon. Rebooting the switch restarts the noncritical daemon. You can schedule the downtime and reboot the switch at your convenience.

The following table lists the daemons that are considered noncritical and that are automatically restarted on failure.

Table 3: Automatic Restart Daemons

Daemon	Description
arrd	Asynchronous Response Router daemon, which is used to send management data to hosts when the switch is accessed through APIs (FA API or SMI-S).
cald	Common Access Layer daemon, which is used by manageability applications.
npd	Flow Vision daemon.
raslogd	Reliability, Availability, and Supportability daemon, which logs error detection, reporting, handling, and presentation of data into a format readable by you and management tools.
snmpd	Simple Network Management Protocol daemon.

Daemon	Description
traced	Trace daemon, which provides the trace entry date and time translation to the trace device at startup and when the date/time is changed by a command. Maintains the trace dump trigger parameters in a trace device. Performs the trace background dump, the trace automatic FTP, and the FTP "aliveness check" if auto-FTP is enabled.
trafd	Traffic daemon, which implements bottleneck detection.
webd	Webserver daemon, which is used for Web Tools (includes httpd as well).
weblinkerd	Weblinker daemon, which provides an HTTP interface to manageability applications for switch management and fabric discovery.

Performing Basic Configuration Tasks

This section describes how to configure your Brocade SAN using the Fabric OS command line interface (CLI). Before you can configure a storage area network (SAN), you must power up the chassis platform or switch and blades and then set the IP addresses of those devices.

Although this section focuses on configuring a SAN using the CLI, you can also use the following methods to configure a SAN:

- Web Tools – For Web Tools procedures, refer to the *Brocade Fabric OS Web Tools User Guide*.
- Brocade SANnav Management Portal – For additional information, refer to the *Brocade SANnav Management Portal User Guide*.
- REST API – For additional information, refer to the *Brocade Fabric OS REST API Reference Manual*.

Due to the differences between fixed-port and variable-port devices, procedures sometimes differ among Brocade models. As new Brocade models are introduced, new features sometimes apply only to those models.

When procedures or parts of procedures apply to some models but not others, this guide identifies the specifics for each model. For example, various procedures that apply only to variable-port devices are found in [Performing Advanced Configuration Tasks](#).

The hardware reference manuals for Brocade products describe how to power up devices and set their IP addresses. After the IP address is set, you can use the CLI procedures that are contained in this guide. For additional information, refer to *Brocade Fabric OS Command Reference Manual*.

Command Line Interface

Fabric OS uses Role-Based Access Control (RBAC) to control access to all Fabric OS operations. Each feature is associated with an RBAC role, and you need to know which role is allowed to run a command, make modifications to the switch, or view the output of the command. To determine which RBAC role you need to run a command, review [Role-Based Access Control](#).

Note the following about the command display in this guide:

- The entire command line (both commands and options) is case-sensitive. Selected command names and options also support Java-style capitalization. Java-style capitalization means that while `bannerShow` and `bannershow` will both work, and `BANNERSHOW` and `BannerShow` will not work. Refer to the *Brocade Fabric OS Command Reference Manual* for explicit instructions on supported capitalization for each command.

Console Sessions Using the Serial Port

Be aware of the following behaviors for serial connections:

- Some procedures require that you connect through the serial port; for example, setting the IP address or setting the boot PROM password.
- Brocade X7 and X6 Director families: You can connect to CP0 or CP1 using either of the two serial ports. One console will be the active CP, and the other will be the standby CP. The standby CP has reduced command line capabilities.

Connecting to Fabric OS through the Serial Port

Use the following procedure to connect to the Fabric OS using the serial port.

NOTE

If you connect the serial console port to a terminal server, ensure that flow control is disabled.

1. Connect the serial cable to the serial port on the switch and to an RS-232 serial port on the workstation.

If the serial port on the workstation is an RJ-45 port, instead of RS-232, remove the adapter on the end of the serial cable and insert the exposed RJ-45 connector into the RJ-45 serial port on the workstation.

2. Open a terminal emulator application (such as HyperTerminal in a Windows environment or TERM, TIP, or Kermit in a UNIX environment), and configure the application as follows:

- In a UNIX environment, enter the following string at the prompt:

```
tip /dev/ttyb -9600
```

You may use `ttya` at the prompt, if `tttyb` is already in use.

```
tip /dev/ttya -9600
```

- In a Windows environment, enter the following parameters:

Table 4: Parameters to Configure a Windows Terminal Emulator Application

Parameter	Value
Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

Management Session

You can connect to the Fabric OS through an SSH or telnet connection or by using a console session on the serial port. The switch must also be physically connected to the network. If the switch network interface is not configured or the switch has been disconnected from the network, use a console session on the serial port as described in [Console Sessions Using the Serial Port](#).

During the initial configuration login to the device, you must log in using the admin account and change all the default passwords.

NOTE

To automatically configure the network interface on a DHCP-enabled switch, plug the switch into the network and power it on. The DHCP client automatically gets the IP and gateway addresses from the DHCP server. The DHCP server must be on the same subnet as the switch. See [DHCP Activation](#) for more information.

Rules for Management Connections

The following rules are followed to make SSH or telnet connections to your switch:

- Never change the IP address of the switch while SSH or telnet sessions are active; if you do, the existing SSH or telnet sessions does not respond and your next attempt to log in will fail. To recover, gain access to the switch by one of the following methods:

- You can use Web Tools to perform a fast boot (POST tests are skipped). When the switch comes up, the SSH or telnet quota is cleared. For instructions on performing a fast boot with Web Tools, refer to the *Brocade Fabric OS Web Tools User Guide*.
- If you have the required privileges, you can connect through the serial port, log in as admin, and use the `killTelnet` command to identify and stop the telnet, SSH, or un-responsive sessions without disrupting the fabric.
- For accounts with an admin role, Fabric OS limits the number of simultaneous SSH or telnet sessions per switch to four like other users. For more details on session limits, see [Managing User Accounts](#).

Connecting to Fabric OS Using the Management Interface

Use the following procedure to connect to the Fabric OS using the management interface.

1. Connect through a serial port to the switch that is appropriate for your fabric:
 - If Virtual Fabrics is enabled, log in using an admin account assigned the chassis-role permission.
 - If Virtual Fabrics is not enabled, log in using an account assigned to the admin role.
2. Verify that the switch's network interface is configured and that it is connected to the IP network through the RJ-45 Ethernet port.

Switches in the fabric that are not connected through the Ethernet port can be managed through switches that are using IP over Fibre Channel. The embedded port must have an assigned IP address.
3. Log off the switch's serial port.
4. From a management station, open a SSH or telnet connection using the IP address of the switch to which you want to connect.

The login prompt is displayed when the SSH or telnet connection finds the switch in the network.

5. Enter the account ID at the login prompt.
6. Enter the password.

If you do not change the default system password, you are prompted to change it upon login. For more information on system passwords, see [Default Account Password](#).

7. Verify that the login was successful.

The prompt displays the switch name and user ID to which you are connected.

```
login: admin
password: xxxxxxxx
```

Configuring the Idle Session Timeout

The existing `timeout` command for management sessions initiates the timer only when there is no CLI command that is being executed. This prevents the sessions from timing out when the CLI is waiting for user input for a long time. The enhanced session timeout feature helps to monitor the idle session and times out the session (in minutes) using the new `--session` option added to the `timeout` command, even if the CLI is in execution and waiting for your input.

NOTE

The enhanced timeout feature is applicable only for SSH sessions.

NOTE

If the session timeout is configured for long-running commands such as `firmwaredownload` and `supportsave`, it is recommended to configure the session timeout value as 60 mins.

The following example configures the SSH session timeout feature in minutes:


```
sw0:admin> timeout --session 15
SSH daemon will be restarted and all SSH based session(s) will be terminated
Do you want to proceed (yes, y, no, n) [no]? yes
Session Idle Timeout Changed to 15 minutes.

Broadcast message from admin@sw0 (ttyS0) (Tue Nov 26 11:18:16 2019):
All SSH accounts will be logged out
```

The following example shows how to disable the SSH timeout feature:

```
sw0:admin> timeout --session 0
SSH daemon will be restarted and all SSH based session(s) will be terminated
Do you want to proceed (yes, y, no, n) [no]? yes
Session Timeout is disabled.
```

Getting Help on a Command

You can display a list of all command help topics for a given login level. For example, if you log in as a user and enter the `help` command, a list of all user-level commands that can be executed is displayed. The same rule applies to the admin, securityAdmin, and switchAdmin roles.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `help` command with no specific command. All commands are displayed.

The optional `|more` argument displays the commands one page at a time.

For command-specific information, enter the `help` command and specify the command for which you need specific information.

```
help
  command
|more
```

The commands in the following table provide help files for the indicated specific topics.

Table 5: Help Topic Contents

Topic Name	Help Contents Description
diagHelp	Diagnostic help information
ficonHelp	FICON help information
mapsHelp	MAPS help information
routeHelp	Routing help information
zoneHelp	Zoning help information

Viewing the History of Command Line Entries

The CLI command history log file saves the last 1680 commands from all users on a first-in, first-out (FIFO) basis, and this log persists reboot and firmware downloads. The `clihistory` command is also supported for the standby control processor (CP).

The log records the following information whenever a command is entered in the switch CLI:

- Timestamp
- User name
- IP address of the telnet session
- Options
- Arguments

Use the following procedure to view the CLI command log.

1. Connect to the switch and log in.
2. Enter the `clihistory` command with the desired argument. Entering no specific argument displays only the command-line history of the currently logged-in user.

CLI History

Entering `clihistory` with no keywords displays the command line history for the currently logged-in user only. This is true for the admin user account as well.

The following example shows the output of the `clihistory` command for the admin login.

```
switch:admin> clihistory

CLI history
Date & Time                Message
Thu Feb 27 10:14:41 2020    admin, 10.1.1.1, switchshow | grep switchType
Thu Feb 27 10:15:09 2020    admin, 10.1.1.1, version
Thu Feb 27 10:15:35 2020    admin, 10.1.1.1, ipaddrshow
Thu Feb 27 10:15:59 2020    admin, 10.1.1.1, ag --modeshow
```

The following example shows the output of the `clihistory` command with fabric IDs (FIDs) in virtual fabric mode.

```
switch:admin> clihistory

CLI history
Date & Time                Message
Fri Feb 19 09:41:08 2020    admin, FID 128, console, clihistory --clear
Fri Feb 19 09:41:41 2020    admin, FID 128, console, lscfg --create 10
Fri Feb 19 09:42:10 2020    admin, FID 128, console, lscfg --create 120
Fri Feb 19 09:42:54 2020    admin, FID 128, console, lscfg --create 30
```

The following example shows the output of the `clihistory` command without the FIDs in non-virtual-fabric mode. The output shows that the `firmwareshow` and `errdump` commands were executed in non-virtual-fabric mode.

```
switch:admin> clihistory

Fri Feb 19 09:43:53 2020    admin, FID 10, console, tracedump
Fri Feb 19 09:43:59 2020    admin, FID 10, console, coreshow
Fri Feb 19 09:44:21 2020    admin, , console, firmwareshow
Fri Feb 19 09:44:25 2020    admin, , console, errdump
```

clihistory --show

Using the `--show` keyword displays the same results as entering `clihistory` alone.

clihistory --showuser username

Use the `clihistory --showuser username` command to display the command-line history for the named user. This argument is available only to accounts with Admin and Securityadmin RBAC role permissions.

The following example shows the output of the `clihistory --showuser` for the `admin42` user account.

```
switch:admin> clihistory --showuser admin42

CLI history
Date & Time          Message
Thu Sep 27 10:14:41 2019  admin42, 10.1.1.1, switchshow | grep switchType
Thu Sep 27 10:14:48 2019  admin42, 10.1.1.1, version
Thu Sep 27 10:15:00 2019  admin42, 10.1.1.1, clihistory --show
```

clihistory --showall

Use the `clihistory --showall` command to display the command-line history for all users. This argument is available only to accounts with Admin and Securityadmin RBAC role permissions.

The following example displays the command history for all users.

```
switch:admin> clihistory --showall | more
CLI history

Date & Time          Message
Mon Oct 5 10:44:05 2015  admin42, FID 128, 172.18.1.1, setcontext 20
Mon Oct 5 10:44:17 2015  admin01, FID 20, 172.18.1.1, firmwareshow
Mon Oct 5 10:44:32 2015  admin01, FID 20, 172.18.1.1, islshow
Mon Oct 5 10:44:20 2015  admin01, FID 20, 172.18.1.1, version
Mon Oct 5 10:44:36 2015  admin01, FID 20, 172.18.1.1, switchshow
Mon Oct 5 10:44:40 2015  admin01, FID 20, 172.18.1.1, trunkshow
```

clihistory --help

Use the `--help` argument to display a list of the available command arguments.

The following example shows the output of the `clihistory --help` command.

```
switch:admin> clihistory --help
clihistory usage:
clihistory:
    Displays the CLI History of the current user
clihistory --show:
    Displays the CLI History of the current user
clihistory --showuser <username>:
    Displays the CLI History of the given user
clihistory --showall:
    Displays the CLI History of all users
```

```

clihistory --enable -shellbuffer:
    CLI history shell CLI buffer is enabled.
clihistory --disable -shellbuffer:
    CLI history shell CLI buffer is disabled.
clihistory --display -shellbuffer:
    Displays CLI history shell CLI buffer status.
clihistory --help:
    Displays the usage

```

Considerations When Viewing CLI Logs

You should be aware of the following considerations when viewing CLI logs:

- The CLI command log is collected as part of any `supportSave` operation. The command log record of such an operation is the equivalent of running `cliHistory --showall`.
- For commands that require a password (for example, the `firmwaredownload`, `configupload`, `configdownload`, and `supportSave` commands), only the command (and no arguments) is stored (refer to the following example for an illustration).

```

switch:admin> firmwaredownload -s -p scp 10.1.1.1,MyUser,/dist,MyPassword
Server IP: 10.1.1.1, Protocol IPv4
Checking system settings for firmwaredownload...
Failed to access scp://fvt:*****@10.1.1.1//dist/release.plist

```

```

switch:admin> clihistory
Date & Time          Message
Wed Nov 23 03:39:37 2016    admin, console, firmwaredownload

```

Using foexec to Run Commands on the Remote Switch or Domain

The `foexec` command allows you to run Fabric OS commands on remote switches or domains across the fabric. Both the local and remote switches must be configured to send and receive the remote command execution. You do not need to log in to the remote switch locally. The outputs of the commands are displayed in the local switch. The commands that you run using `foexec` are also captured in the CLI history and audit logs of the remote switch.

The `foexec` feature is configured by using the `configure` command. The `foexec` feature is disabled by default and it is enabled for both sending and receiving switches or domains. The configuration is checked when sending any `foexec` request to a remote switch and also when receiving such a request from a remote switch. You can execute this command either on a specific switch or domain or on all the switches or domains in the fabric. For a specific domain or switch, you must provide the domain ID (DID), which must be from 1 through 240.

The `configure` command is used to turn the remote `foexec` on or off on a particular switch. The applicable element is shown in **bold**.

```
switch:admin> configure
```

```

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

```

```
Configure...
```

```

Fabric parameters (yes, y, no, n): [no] y

WWN Based persistent PID (yes, y, no, n): [no]
Location ID: (0..4) [0]
High Integrity Fabric Mode (yes, y, no, n): [no]
Edge Hold Time(Low(80ms), Medium(220ms), High(500ms), UserDefined(80-500ms): (80..500)
[220]
Remote fosexec feature: (on, off): [off] on
D-Port Parameters (yes, y, no, n): [no] 2014/02/10-12:39:49, [CONF-1043], 1815, FID 128, INFO,
sw_85, Fabric Configuration Parameter Remote fosexec feature changed to enabled
Zoning Operation parameters (yes, y, no, n): [no]
Commands executed via remote fosexec is captured in clihistory and audit logs of remote switch

```

The fosexec syntax is as follows:

```

fosexec --domain DID -cmd "[args]"
fosexec --domain all -cmd "[args]"

```

When nested quotes or strings must be provided, instead of using nested quotes, precede the string with an escape character or backward slash with quotes, as shown in the following example:

```

switch:admin>fosexec --domain all -cmd "CLIname --set \"SWAT Setup\""

```

From Fabric OS 8.2.0 or later, the remote domain execution of fosexec feature allows permission for the user to execute the portcfgshow command as shown in the following example:

```

switch:admin>fosexec --ag all -cmd "portcfgshow 1"

```

The fosexec command has the following limitations:

- Remote execution is allowed based on RBAC checks or permissions for the remote domain role. The RBAC permission denied message is displayed for unsupported commands.
- fosexec can be executed only by users with the fabricadmin/admin roles and RBAC permissions.
- fosexec does not work from the remote domain for switch HA failover or reboot.
- The command syntax must match the command syntax that is supported by the remote switch.
- Commands with parameters such as a slot/port specific to a fixed-port switch or director are not supported with the all option.
- fosexec does not support interactive commands that require inputs in the middle of execution.
- Commands that take longer time to execute are not supported. The timeout period is 15 seconds.
- fosexec command returns a maximum of 64 KB of data from the remote switch. Any data beyond this limit is truncated.
- When you run supportSave through fosexec, wait for the execution to complete. To know whether the execution has completed, refer to the RASLOG [SS-1000] on all domains or the specific domains where fosexec is being run.
- fosexec is not supported in Access Gateway (AG) mode.
- The fosexec command is not supported with inline SSH.

NOTE

Running the fosexec command does not log you out from the current session in your local switch.

Example 1

This example shows the output of running fosexec on all domains.

```

admin> fosexec --domain all -cmd "islshow"

```

```

Domain 13
=====
 1: 0-> 8 10:00:50:eb:1a:b7:11:24 (incompatible)
 2: 13-> 13 10:00:50:eb:1a:b7:16:a4 124 SW6505_114124      sp: 16.000G bw: 64.000G TRUNK QOS
CR_RECOV FE
 3: 16-> 88 10:00:50:eb:1a:b8:66:96 166 SW6520_114166      sp: 16.000G bw:128.000G TRUNK QOS
CR_RECOV FE

Domain 17
=====
 1: 16-> 16 10:00:50:eb:1a:b5:a4:01 120 SW6505_114120      sp: 16.000G bw:128.000G TRUNK QOS
CR_RECOV FEC
 2: 25-> 13 10:00:50:eb:1a:b5:27:f8 121 SW6505_114121      sp:  8.000G bw:  8.000G TRUNK QOS
CR_RECOV
 3: 26-> 14 10:00:50:eb:1a:b5:27:f8 121 SW6505_114121      sp: 16.000G bw: 32.000G TRUNK QOS
CR_RECOV FEC

Domain 36
=====
 1:1366->457 10:00:00:05:33:0d:52:01 54 EGEN6_8__114054_LS1 sp:----- bw:852.000G
 2:1372->464 10:00:00:05:33:10:ca:01 72 EDCX16G_130072_LS1 sp:----- bw:852.000G

Domain 39
=====
 1:1556->1610 10:00:00:27:f8:f3:3b:41 13 Core-X6-8_130013_LS1      sp:----- bw:512.000G
 2:1570->1406 10:00:88:94:71:ba:cd:e2 239 Core-X7-8_130239_LS1      sp:----- bw:848.000G
 3:1590->1581 10:00:88:94:71:6e:af:6b 101 E-X7-8_Upgraded_130101_LS1 sp:----- bw:768.000

```

Example 2

This example shows the output of running `fosexec` only on domain 13.

```

admin> fosexec --domain 13 -cmd "islshow"
 1: 0-> 8 10:00:50:eb:1a:b7:11:24 (incompatible)
 2: 13-> 13 10:00:50:eb:1a:b7:16:a4 124 SW6505_114124      sp: 16.000G bw: 64.000G TRUNK QOS
CR_RECOV FE
 3: 16-> 88 10:00:50:eb:1a:b8:66:96 166 SW6520_114166      sp: 16.000G bw:128.000G TRUNK QOS
CR_RECOV FE

```

Modification of the Default Password

The switch automatically prompts you to change the default account password after logging in for the first time. When you initially log in as admin, or user, you must immediately change the password for the default accounts. For a maintenance user, you must enable the user first and log in as maintenance to change the password. If you cancel the password change using `Ctrl+C`, the system logs out, and once you log in, you are again prompted to change the password.

NOTE

When you log in as admin and the default password is not changed from the manufacturer default, you are prompted to change the default passwords only for the admin and user accounts.

Use the admin account when you log in to the switch for the first time and to perform basic configuration tasks. The user account is primarily used for basic system monitoring. For more information on the default accounts, see [Default Accounts](#).

NOTE

If an account password is at default, if it has expired or if the password hash has changed, and a command is entered as a single line using an SSH login, then the command is blocked and the following message is displayed: Your account password is at default, expired, or password hash has changed. Kindly log into the switch and change the password to a non-default value before proceeding.

NOTE

When you log in as the admin or the user account in the standby CP, and your password is at default on the active CP, the following message is displayed: Password for admin/user is at default. Kindly login to Active CP and change the password to a non-default value before proceeding. When the default password is changed for the admin or the user account on the active CP, all corresponding sessions on the standby CP are logged out.

Default Account Password

The prompt to change the default account password is a string that begins with the message `Please change passwords for switch default accounts now`. User-defined passwords are 8 through 40 characters. They are case-sensitive and are not displayed when entered on the command line.

NOTE

If you are changing the password string for the admin, user, and maintenance account to `password`, the system sends the warning message `You have chosen the default password and it must be changed at the next switch login. Please acknowledge the temporary use of the password or use Control+C to exit`. The system detects that the password is set to the manufacturer's default password. Even though you are allowed to set the default password at the next login, it prompts you to change the password immediately.

NOTE

Maintenance account is only to be used when directed by your service provider.

Record the passwords exactly as entered and store them in a secure place because recovering passwords requires significant effort and fabric downtime.

For more information, see [Password Policies](#).

Changing the Default Account Password of the Switch

Use the following procedure to change the default account password of the switch.

1. Connect to the switch and log in using the default administrative account.
2. At each of the `Enter new password` prompts, enter the new password.

The following example shows the output for changing the default password.

```
Fabric OS (sw0)
sw0 login: admin
Password:
Please change passwords for switch default accounts now.
for user - admin
Changing password for admin
Enter old password:
Enter new password:
Re-type new password:
passwd: all authentication tokens updated successfully
Please change passwords for switch default accounts now.
for user - user
```

```
Changing password for user
Enter new password:
Re-type new password:
passwd: all authentication tokens updated successfully
Saving passwords to stable storage.
Passwords saved to stable storage successfully
```

Switch Ethernet Interface

The Ethernet (network) interface provides management access, including direct access to the Fabric OS CLI, and it allows other tools, such as Web Tools, to interact with the switch. You can use either Dynamic Host Configuration Protocol (DHCP) or static IP addresses for the Ethernet network interface configuration.

Brocade Directors

On Brocade X7 and X6 Directors, you must set IP addresses for the following components:

- Both control processors (CP0 and CP1)
- Chassis management

Brocade Switches

On Brocade switches, you must set the Ethernet and chassis management IP interfaces.

Setting the chassis management IP address eliminates the need to know which CP is active and automatically connects the requestor to the currently active CP.

You can continue to use a static Ethernet addressing system or allow the DHCP client to automatically acquire Ethernet addresses. Configure the Ethernet interface IP address, subnet mask, and gateway addresses in one of the following manners:

- Using static Ethernet addresses (see [Static Ethernet Addresses](#))
- Activating DHCP (see [DHCP Activation](#))

NOTE

When you change the Ethernet interface settings, open connections such as SSH or telnet may be dropped. Reconnect using the new Ethernet IP address information or change the Ethernet settings using a console session through the serial port to maintain your session during the change. You must connect through the serial port to set the Ethernet IP address if the Ethernet network interface is not configured already. For details, see [Connecting to Fabric OS through the Serial Port](#).

Virtual Fabrics and the Ethernet Interface

On the Brocade X7 Director and Brocade X6 Director, the single-chassis IP address and subnet mask are assigned to the management Ethernet ports on the front panels of the CPs. These addresses allow access to the chassis—more specifically, the active CP of the chassis—and not individual logical switches. The IP addresses can also be assigned to each CP individually. This allows for direct communication with a CP, including the standby CP. Each CP has two management Ethernet ports on its front panel. These two physical ports are bonded together to create a single, logical Ethernet port, and it is the logical Ethernet port to which IP addresses are assigned.

IPv4 addresses assigned to individual virtual fabrics are assigned to IP over Fibre Channel (IPFC) network interfaces. In Virtual Fabrics environments, a single chassis can be assigned to multiple fabrics, each of which is logically distinct and separate from one another. Each IPFC point of connection to a given chassis needs a separate IPv4 address and prefix to be accessible to a management host. For more information, see [Managing Virtual Fabrics](#).

NOTE

The 10G port located on each CPX, CPX6, and CPX7 control processor blade is currently unsupported and must be left disconnected.

Management Ethernet Port Bonding

The two external Ethernet ports of a CP (CPX, CPX6, or CPX7) blade can be bound together as a single logical network interface. This configuration uses an active-standby failover model to provide automatic failover support for the primary Ethernet port on the blade. If the primary Ethernet port fails (due to something other than power loss), the second Ethernet port immediately takes over to ensure that link layer communication is retained.

One of the physical Ethernet ports is selected as the active interface. The second interface is set as the standby interface. All traffic is transmitted over the active interface. No traffic is transmitted over the standby interface, unless the active interface is determined to be no longer connected; at which point, the second interface is made active.

When active, all Fabric OS kernel modules and applications on the CP blade use the logical network interface named “bond0” instead of “eth0”.

NOTE

On reboot, physical port eth0 is always made active if it is connected.

The CP blade contains multiple Ethernet devices (including eth0 and eth3), which map to the two Ethernet ports on the front of the CP blade. Other Ethernet devices on the blade are reserved for use by the operating system.

The CP blade enables eth0 by default. If an error is encountered on eth0, it is treated the same as for any other port, unless the error causes the eth0 port to go down. If eth0 goes down, the eth3 interface becomes active and will remain active until eth0 comes back. Perform one of the following actions to restore eth0 as the active interface:

- Unplug the network cable, wait 5 seconds, and then plug it back in.
- Perform a High Availability (HA) failover routine.
- Power down the switch and then power it back up again.

ATTENTION

Performing an HA failover and powering down the switch will cause a disruptive delay in content delivery.

Supported Brocade Directors

Management Ethernet port bonding is available on the following CP blades:

- CPX7 or CPX blade installed on a Brocade X7 Director
- CPX6 or CPX blade installed on a Brocade X6 Director

Setting Up the Second Ethernet Port on a CPX7 or CPX6 or CPX Blade

The port speed and duplex mode between the Ethernet ports should always match. Both ports should be set at a fixed speed or set to autonegotiate.

1. Make sure that the speed and link operating mode settings are the same for both eth3 and eth0. See [Setting the Ethernet Interface Mode and Speed](#) for instructions on setting port modes, and [Setting Port Speeds](#) for instructions on setting port speeds.
2. Physically connect the second Ethernet port to the same network as the primary Ethernet port.

Displaying the Network Interface Settings

NOTE

If an IP address has not been assigned to the network interface (Ethernet), you must connect to the Fabric OS CLI using a console session on the serial port. For more information, see [Console Sessions Using the Serial Port](#). Otherwise, connect using SSH or telnet.

1. Connect to the switch, and log in using an account assigned to the admin role.
2. Enter `ipAddrShow`.

The following example shows the output of `ipAddrShow` for a Brocade X6 Director.

```
admin> ipaddrshow

CHASSIS
Ethernet IP Address: 10.1.1.1
Ethernet Subnet mask: 255.255.240.0

CP0
Ethernet IP Address: 10.1.1.2
Ethernet Subnet mask: 255.255.240.0
Host Name: cp0
Gateway IP Address: 10.155.1.1

CP1
Ethernet IP Address: 10.1.1.3
Ethernet Subnet mask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.155.1.1
DHCP: Off
IPv6 Autoconfiguration Enabled: No
Local IPv6 Addresses:
cp 0 link local fe80::c6f5:7cff:fe01:ab09/64
cp 1 link local fe80::c6f5:7cff:fe01:acd1/64
IPv6 Gateways:
DHCPv6: Off
```

The display of the Ethernet IP address, subnet mask, and gateway address shows that the network interface is configured. Verify that the information on your switch is correct. If DHCP is enabled, the network interface information was acquired from the DHCP server.

NOTE

You can use either IPv4 or IPv6 with a classless inter-domain routing (CIDR) block notation (also known as a *network prefix length*) to set up your IP addresses.

Resetting the Error Counter for the Management Network Interface

You can reset the management network interface error counter on the Brocade Gen 7 and Gen 6 platforms. To reset the error counter, complete the following steps.

1. Enter `ethif --reseterror <interface_ID [eth0|eth3]>`.

This command clears the counter and displays the completion message.

```
switch:admin> ethif --reseterror eth0
Statistics cleared for eth0
```

2. To confirm that the error counter has been cleared, enter `ethif --show <interface_ID [eth0|eth3]>`.

This command displays the counter for that interface. The errors counter will be zero (0).

```
switch:admin> ethif --show eth0
eth0 interface:
Link mode: negotiated 1000baseT-FD, link ok
MAC Address: C4:F5:7C:01:AC:D1

eth0      Link encap:Ethernet HWaddr C4:F5:7C:01:AC:D1
          UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1
          RX packets:8188743 errors:0 dropped:0 overruns:0 frame:0
          TX packets:11421394 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:704473789 (671.8 Mb) TX bytes:13583614219 (12954.3 Mb)
          Memory:fe5e6000-fe5e6fff
```

Setting the TCP Timeout Level

Based on network requirements, you can configure or set the timeout value for TCP half-open connections using the `ethif --set tcptimeoutlevel` command.

The new `ethif --set tcptimeoutlevel [1|2|3|4|5]` command allows you to configure the timeout level range from 1 to 5, which modifies the effective timeouts. This configuration is applicable for both IPv4 and IPv6 connections. For directors, the timeout level configured will be synced to the standby Ethernet interfaces as well. The approximate timeout value for each level will be 3, 7, 15, 31, and 63 seconds, respectively.

```
switch:admin> ethif --set -tcptimeoutlevel 3
TCP timeout level changed to 3 (~15 secs)
```

To view the current value set for the TCP timeout level, use the `ethif --show tcptimeoutlevel` command.

```
switch:admin> ethif --show -tcptimeoutlevel
3 (~15 secs)
```

Management Port Tracker

From FOS 9.1.0, management port tracker is used to periodically scan the network stack to get the network connection details of the switch.

The management port tracker records the following network connection details:

- The local IP address/port in use on the switch
- The remote IP address/port of the connection
- The protocol that is used
- The name of the program using the connection
- The time at which the connection is first seen or detected by the tool
- The last updated time of the connection state (State Update Time) or closed time
- Both IPv4 and IPv6 connections

The information that is recorded by the management port tracker is used to identify the network services running on the switch. The port tracker also maintains the history for a maximum of 2000 network connections.

Enabling or Disabling Management Port Tracker

You can enable or disable the management port tracker using the `ethif --enable mgmtporttrack` or `ethif --disable mgmtporttrack` commands. By default, the feature is enabled and connections are scanned periodically every 30 seconds.

NOTE

The network connections that open and close between two poll intervals cannot be tracked by the port tracker.

```
switch:admin> ethif --enable mgmtporttrack
```

Configuring the Management Port Tracker Interval

You can set the interval for querying the network stack to monitor connections using `ethif --set mgmtporttrack -interval <30s-300s>`. By default, the scan interval is set to 30 seconds. You can change the port scan interval up to 5 minutes (300s).

```
switch:admin> ethif --set mgmtporttrack -interval <30s-300s>
```

Displaying Network Connection Information

You can use the `ethif --show mgmtporttrack <filter_options>` command to display the different network connections.

Usage:

```
ethif --show mgmtporttrack -all - Displays most recent X number of connections
ethif --show mgmtporttrack -port <port> - Displays most recent X number of connections to local port <port>
ethif --show mgmtporttrack -proto tcp|udp - Displays most recent X number of tcp/udp connections
ethif --show mgmtporttrack -family ipv4|ipv6 - Displays most recent X number of v4/v6 connections
ethif --show mgmtporttrack -state active|inactive - Displays X number of current/past connections
ethif --show mgmtporttrack -config - Display the current configuration state
```

```
switch:admin> ethif --show mgmtporttrack -all
```

Start Time PID/Program	State Update Time	Local Address	Remote Address	State	Protocol
07/23/21-04:40:11 2568/snmpd	-----	0.0.0.0:161	0.0.0.0:0	CLOSE	UDP
07/23/21-04:40:11 2549/webd	-----	0.0.0.0:52357	0.0.0.0:0	CLOSE	UDP

07/23/21-04:40:11 1702/inetd	-----	0.0.0.0:69	0.0.0.0:0	CLOSE	UDP
07/29/21-10:35:19 ----	-----	10.1.1.29:22	10.2.67.71:12345	ESTABLISHED	TCP
07/29/21-10:08:49 ----	-----	10.1.1.29:22	10.2.67.71:67890	ESTABLISHED	TCP
07/29/21-09:56:19 ----	-----	10.1.1.29:22	10.2.45.58:13579	ESTABLISHED	TCP
07/23/21-04:40:41 1702/inetd	-----	0.0.0.0:23	0.0.0.0:0	LISTEN	TCP
07/23/21-04:40:41 2395/sshd	-----	0.0.0.0:22	0.0.0.0:0	LISTEN	TCP

Configuring Global LLDP Parameters for the Management Port

From Fabric OS 9.1.0, the Link Layer Discovery Protocol (LLDP) is extended to the physical management ports. LLDP discovers the Ethernet neighbor at the link level, and it advertises the switch information with the peer device. The default transmit interval is 30 seconds.

NOTE

LLDP is enabled by default when a management port is up and running, and the global parameters are applied to all management ports.

To configure LLDP global parameters, perform the following steps:

1. If LLDP is not enabled, enable it on the switch using the `lldp --enable` command.

```
switch:admin> lldp --enable
```

2. Configure LLDP global parameters using the `lldp --config` command.

```
switch:admin> lldp --config -sysname system_19
switch:admin> lldp --config -sysdesc X6_default_vf
switch:admin> lldp --config -mx 5
switch:admin> lldp --config -txintvl 40
```

3. Enable the required type length values (TLVs) globally using the `lldp --enable -tlv` command.

```
switch:admin> lldp --enable -tlv sys-desc
```

NOTE

The `sys-name` and `port-desc` TLVs are globally enabled by default.

The following TLVs are supported:

- `dot3` – IEEE 802.3 Organizationally Specific
- `mgmt-addr` – Management Address
- `port-desc` – Port Description
For example, for the TLV `port-desc`, the string format is "Switch Model Name: Port Description". For example, Brocade_X6_4: ACTIVE MGMT 1G eth0.
- `sys-cap` – System Capabilities
- `sys-desc` – System Description
For example, for the TLV `sys-desc`, the default system description string in the TLV is advertised for the switch. For example, Brocade_G720_Fabric_A.
The maximum size of the description string is 255 characters. Also, the string that arises from the `sys-desc` TLV on a neighboring switch is accepted as is.
- `sys-name` – System Name

4. Use the following command to display the global LLDP statistics for the local switch.

```
switch:admin> lldp --show
LLDP Global Information
  system-name: G720_29
  system-description: Brocade_G720_Fabric_A
  State: Enabled
  Mode: Receive/Transmit
  Advertise Transmitted: 30 seconds
  Hold time for advertise: 120 seconds
  Tx Delay Timer: 1 seconds
  Transmit TLVs: Chassis ID      Port ID
                  TTL           Port Description
                  System Name
```

5. Use the following command to display the LLDP statistics based on the Ethernet port.

```
switch:admin> lldp --show -port eth0
LLDP information for eth0
  State: Enabled
  Mode: Receive/Transmit
  Advertise Transmitted: 30 seconds
  Hold time for advertise: 120 seconds
  Tx Delay Timer: 1 seconds
  Transmit TLVs: Chassis ID      Port ID
                  TTL           Port Description
                  System Name
```

6. Use the following command to display the LLDP neighbors and statistics.

```
switch:admin> lldp --show -nbr
Local Intf    Dead Interval Remaining Life Remote Intf    Chassis ID    Tx    Rx    System
Name
eth0          120          104          Gi1/0/11      0005.1e8f.fd04 2134 2140 RS1.nw-
am.broadcom.net
3/41          120          110          TenGigabitEthernet 1/0/19 8c60.4ff7.edc6 2155 2130 sw0
3/42          120          95           port0         8c60.4ff7.edc7 2154 2039 (null)
```

7. Use the following command to display all the LLDP statistics.

```
switch:admin> lldp --show -stats
LLDP Interface statistics for eth0
Frames transmitted: 2134
Frames Aged out: 0
Frames Discarded: 0
Frames with Error: 0
Frames Received: 2140
TLVs discarded: 0
TLVs unrecognized: 0
LLDP Interface statistics for 3/41
Frames transmitted: 2155
Frames Aged out: 0
Frames Discarded: 0
Frames with Error: 0
```

```

Frames Received:      2130
TLVs discarded:       0
TLVs unrecognized:    0
LLDP Interface statistics for 3/42
Frames transmitted:   2154
Frames Aged out:      0
Frames Discarded:     0
Frames with Error:    0
Frames Received:      2093
TLVs discarded:       0
TLVs unrecognized:    0
LLDP Interface statistics for 3/43
Frames transmitted:   0
Frames Aged out:      0
Frames Discarded:     0
Frames with Error:    0
Frames Received:      0
TLVs discarded:       0
TLVs unrecognized:    0

```

NOTE

On Gen 7 or Gen 6 directors, the backup port agent that is connected on each CP is not visible in the neighbor switch LLDP output. On your switch LLDP output, the neighbor port is still visible.

Management Port Ingress Rate Limiting

A Denial of Service (DoS) attack is one of the primitive type of mechanisms where the network and computing resources are exhausted by the attacker and genuine applications are denied their share of services. The following resources are targeted:

- Overutilizing the network bandwidth
- Overutilizing the CPU's bandwidth
- Overutilizing some specific resource by flooding with only one type of packet to an Ethernet node

The management port ingress rate limiting mechanism is enforced on the eth0 and eth3 interfaces of Gen 6 and Gen 7 directors to prevent the CPU from stalling or being overwhelmed during a DoS attack.

Use the `configure` command to configure the management port ingress rate limiting feature. By default, this feature is disabled.

Configuring Management Port Ingress Rate Limiting

Use `Management Port Ingress Rate Limiting`, listed under the `System services` parameters of the `configure` command, to enable or disable the feature.

NOTE

This is applicable or available only when executed from the default FID.

```
switch:admin> configure
```

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

```
Configure...
```

```
Fabric parameters (yes, y, no, n): [no]
D-Port Parameters (yes, y, no, n): [no]
RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24)[1]
System services (yes, y, no, n): [no]yes

Management Port Ingress Rate Limiting (on, off): [off]on
..... <output truncated>
```

Static Ethernet Addresses

Use static Ethernet network interface addresses in environments where DHCP service is not available. To use static addresses for the Ethernet interface, you must first disable DHCP. You can enter static Ethernet information and disable DHCP at the same time. For more information, see [DHCP Activation](#).

You can choose not to use DHCP or to specify an IP address for your switch Ethernet interface, by entering `none` or `0.0.0.0` in the Ethernet IP address field.

On an application blade, configure the two external Ethernet interfaces to two different subnets. If two subnets are not present, configure one of the interfaces and leave the other interface unconfigured. Otherwise, the following message displays and blade status may go into a faulty state after a reboot.

```
Neighbor table overflow.
print: 54 messages suppressed
```

Setting a Static Address for the Ethernet Network Interface

Use the following procedure to set an Ethernet network interface static address.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Perform the appropriate action based on whether you have a switch or director:
 - If you are setting the IP address for a switch, enter the `ipaddrset` command.
 - If you are setting the IP address for a director, enter the `ipaddrset -chassis` command.
3. Enter the network information in dotted-decimal notation for the Ethernet IPv4 address or in semicolon-separated notation for IPv6.
4. Enter the value for Ethernet Subnet Mask.
5. Enter the Gateway IP Address.
6. Enter `off` to disable DHCP.

The following is an example of setting an IPv4 address for a switch.

```
switch:admin> ipaddrset
DHCP [OFF]:
Ethernet IP Address [10.1.18.5]:10.1.18.9
Ethernet Subnet mask [255.255.240.0]:
Gateway IP Address [10.1.18.1]:
IP address is being changed...
(User gets disconnected from mgmt. interface and will need to re-login with new IP address)
```

The following is an example of setting an IPv4 address for a chassis.


```
switch:admin> ipaddrset -chassis
DHCP [OFF]:
Ethernet IP Address [10.1.18.7]:10.1.15.5
Ethernet Subnet mask [255.255.240.0]:
IP address is being changed...
(User gets disconnected from mgmt. interface and will need to re-login with new IP address)
```

The following is an example of setting an IPv6 address on a switch:

```
switch:admin> ipaddrset -ipv6 --add 1080::8:800:200C:417A/64

IP address is being changed...Done.
```

For more information, see [Managing Virtual Fabrics](#).

Setting a Static Address for the Chassis Management IP Interface

Use the following procedure to set a chassis management IP interface static address.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `ipAddrSet -chassis` command.
3. Enter the network information either in dotted-decimal notation for an Ethernet IPv4 address or in semicolon-separated notation for an IPv6 address.
4. Enter the value for `Ethernet Subnetmask`.

```
switch:admin> ipaddrset -chassis
DHCP [Off]:
Ethernet IP Address [10.1.18.14]:
Ethernet Subnetmask [255.255.240.0]:
Committing configuration...Done.
```

DHCP Activation

The Dynamic Host Configuration Protocol (DHCP) is not enabled on Brocade switches by default.

The Fabric OS DHCP client conforms to the latest IETF Draft Standard RFCs for IPv4, IPv6, and DHCP. You can use the following parameters to activate DHCP:

- External Ethernet port IP addresses and subnet masks
- Default gateway IP address

Notes on DHCP

Keep the following items in mind when working with DHCP:

- The DHCP client can obtain stateful IPv6 addresses.
- When DHCP is enabled, the switch name is registered to the Domain Name System (DNS) server automatically, and the switch is assigned an IP address from the DHCP server. Dynamic DNS (DDNS) is a method of automatically updating a name server in the DNS, often in real time, with the active DDNS configuration of its configured host names, addresses, or other information.

Enabling DHCP for IPv4 in Fixed-Port Devices

When you connect a DHCP-enabled switch to the network and power on the switch, the switch automatically obtains the Ethernet IP address, Ethernet subnet mask, and default gateway address from the DHCP server.

NOTE

The DHCP client can only connect to a DHCP server on the same subnet as the switch. Do not enable DHCP if the DHCP server is not on the same subnet as the switch.

Enabling DHCP after the Ethernet information has been configured releases the current Ethernet network interface settings. These include the Ethernet IP address, Ethernet subnet mask, and gateway IP address. The Fibre Channel IP address and subnet mask are static and are not affected by DHCP; for instructions on setting the FC IP address, see [Static Ethernet Addresses](#).

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `ipAddrSet` command.

NOTE

Alternatively, you can enable DHCP for IPv4 by entering `ipaddrset -ipv4 -add -dhcp ON` as a single command. If you do so, you do not need to follow the subsequent steps.

3. Enter `On` to enable DHCP.
4. Confirm that the change has been made using the `ipAddrShow` command.

The following example enables DHCP for IPv4.

```
switch:admin> ipaddrset
DHCP [Off]:On
IP address is being changed...
Done.
```

The following example enables DHCP for IPv4 using a single command.

```
switch:admin> ipaddrset -ipv4 -add -dhcp ON
switch:admin> ipaddrshow

SWITCH
Ethernet IP Address: 10.1.18.219
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.1.18.1
DHCP: On
```

Disabling DHCP for IPv4 in Fixed-Port Devices

When you disable DHCP, the previously configured static Ethernet IP address, subnet mask, and default gateway address are restored as the active configuration.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `ipAddrSet` command.

NOTE

Alternatively, you can disable DHCP for IPv4 by entering `ipaddrset -ipv4 -add -dhcp OFF` as a single command. If you do so, you do not need to perform the subsequent steps.

3. Enter `Off` to disable DHCP.
4. Confirm that the change has been made using the `ipAddrShow` command.

The following example disables DHCP for IPv4.

```
switch:admin> ipaddrset
DHCP [On]:Off
IP address is being changed...
Done.
```

The following example disables DHCP for IPv4 using a single command.

```
switch:admin> ipaddrset -ipv4 -add -dhcp OFF

switch:admin> ipaddrshow

SWITCH
Ethernet IP Address: 10.1.18.219
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.1.18.1
DHCP: Off
```

Enabling DHCP for IPv4 in Directors

To enable DHCP for IPv4 on Brocade X7 Directors and Brocade X6 Directors, perform the following steps.

1. Enter the `ipAddrSet -chassis` command.

```
switch:admin> ipaddrset -chassis -ipv4 -add -dhcp on
Info : This feature got enabled on all Management Interfaces of this chassis
Done.
```

2. To display the configured IP address, enter the `ipAddrShow` command.

```
switch:admin> ipaddrshow

CHASSIS
Ethernet IP Address: 10.1.18.231
Ethernet Subnetmask: 255.255.240.0

CP0
Ethernet IP Address: 10.1.18.228
Ethernet Subnetmask: 255.255.240.0
Host Name: cp0
Gateway IP Address: 10.1.18.1

CP1
Ethernet IP Address: 10.1.18.229
Ethernet Subnetmask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.1.18.1
```

3. To enable DHCP for IPv4 on directors, perform the following steps.
 - In interactive mode:

```
switch:admin> ipaddrset -chassis
DHCP [On]:On
IP address is being changed...
Done
```

- In a single command:

```
switch:admin> ipaddrset -chassis -ipv4 -add -dhcp on
```

Configuring DHCPv6 for IPv6

You can configure the IPv6 address from the DHCPv6 server with the following procedure.

1. Enter the `ipAddrSet -ipv6 -dhcpv6` command.
2. To display the configured IPv6 address, enter the `ipAddrShow` command.
 - For fixed-port switches, the following output is displayed:

```
switch:admin> ipaddrshow
SWITCH
Ethernet IP Address: 10.1.18.33
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.1.18.1
DHCP: Off
IPv6 Autoconfiguration Enabled: No
Local IPv6 Addresses:
dhcpv6 2001:100:0:f603::f9ff/64 preferred
link local 2001::205:33ff:fe3a:77f9/64
IPv6 Gateways:
2001::db81b:edff:fe0b:9000
DHCPv6: On
```

- For directors, the following output is displayed:

```
switch:admin> ipaddrshow
CHASSIS
Local IPv6 Addresses:
dhcpv6 2001:100:0:f603::f912/64 preferred
link local 2001::205:33ff:fe3a:77f9/64

CP0
Local IPv6 Addresses:
dhcpv6 2001:100:0:f603::f9fb/64 preferred
link local 2001::205:33ff:fe3a:77f8/64
IPv6 Gateways:
2001::21b:edff:fe0b:9000
DHCPv6: On

CP1
Local IPv6 Addresses:
dhcpv6 2001:100:0:f603::f9fa/64 preferred
link local 2001::205:33ff:fe3a:77f7/64
IPv6 Gateways:
2001::21b:edff:fe0b:9000
DHCPv6: On
```

3. To disable DHCPv6 for IPv6, run the following command.

```
ipaddrset -ipv6 -nodhcpv6
```

Configuring Stateless DHCPv6

The configuration parameters, such as the DNS recursive name servers, the DNS search list, and the SIP servers, must be configured using stateless DHCPv6.

1. Enter the `ipAddrSet -ipv6 -auto` command.
2. To display the configured IP address details, enter the `ipAddrShow` command.
 - For fixed-port switches, the following output is displayed:

```
switch:admin> ipaddrshow
SWITCH
Ethernet IP Address: 10.1.18.33
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.1.18.1
DHCP: Off
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
dhcpv6 2001:100:0:f603::f9ff/64 preferred
stateless 2001:db8:0:f602:227:f8ff:fe0:1800/64 preferred
link local 2001::db8:33ff:fe3a:77f9/64
IPv6 Gateways:
2001::db8:edff:fe0b:9000
DHCPv6: On
```

- For directors, the following output is displayed:

```
switch:admin> ipaddrshow
CHASSIS
Ethernet IP Address: 10.1.18.217
Ethernet Subnetmask: 255.255.240.0

CP0
Ethernet IP Address: 10.1.18.240
Ethernet Subnetmask: 255.255.240.0
Host Name: cp0
Gateway IP Address: 10.1.18.1

CP1
Ethernet IP Address: 10.1.18.241
Ethernet Subnetmask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.1.18.1
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
chassis 0 dhcpv6 2001:100:0:f603::f1a2/64 preferred
cp 0 link local 2001::db8:1eff:febf:7069/64
chassis 0 stateless 2001:100:0:f603:205:1eff:feb7:3c00/64 preferred
cp 0 static 2001:100:0:f603::f1a3/64 preferred
cp 0 stateless 2001:100:0:f603:205:1eff:febf:7069/64 preferred
cp 0 dhcpv6 2001:100:0:f603::f773/64 preferred
cp 1 link local 2001::db8:1eff:febf:7068/64
```

```

cp 1 stateless 2001:100:0:f603:205:1eff:febf:7068/64 preferred
cp 1 dhcpv6 2001:100:0:f603::f49f/64 preferred
IPv6 Gateways:
cp 0 2001::db8:edff:fe0b:3c00
cp 0 2001::db8:edff:fe0b:9000
DHCPv6: On

```

3. To display the configuration parameters of the DNS name servers in `supportSave`, run the `dnsconfig --show` command.

```
switch:admin> dnsconfig --show
```

Domain Name Server Configuration Information

```

Domain Name           = englab.brocade.com.
Name Server IP Address = 2001:db8:4:e200::10
Name Server IP Address = 2001:db8:4:e200::11

```

Configuring a Static IPv6 Gateway Address

You can configure a static IPv6 gateway address using the `ipaddrset -ipv6 --add` command. When the static IPv6 gateway is configured, auto-assignment of the gateway address is not blocked; however, the static address remains active.

Consider the following when configuring the static IPv6 gateway address:

- Both the CPs in the chassis must always have the same gateway IP address.
- When you delete the static IPv6 gateway address, a chassis reboot is required for directors and a reboot is required for fixed-port switches for the autoconfigured IPv6 gateway address to become active.
- Only one IPv6 static gateway can be configured at a time. Configuring another overwrites the existing one.
- On platforms with blade processors, using the option `-cp cp_number` for IPv6 static gateway configuration updates both CPs, no matter what number is used for `cp_number`.

1. Configure the static IPv6 gateway address using the following syntax:

```
ipaddrset -ipv6 --add -gwyip gateway_ip [-cp(x)]
```

2. Delete the static IPv6 gateway address using the following syntax:

```
ipaddrset -ipv6 --del -gwyip gateway_ip [-cp(x)]
```

IPv6 Autoconfiguration

IPv6 can assign multiple IP addresses to each network interface. Each interface is configured with a link local address in almost all cases, but this address is accessible only from other hosts on the same network. To provide for wider accessibility, interfaces are typically configured with at least one additional global scope IPv6 address. IPv6 autoconfiguration allows more IPv6 addresses, the number of which is dependent on the number of routers serving the local network and the number of prefixes they advertise.

There are two methods of autoconfiguration for IPv6 addresses: stateless autoconfiguration and stateful autoconfiguration. *Stateless autoconfiguration* allows an IPv6 host to obtain a unique address using the IEEE 802 MAC address. *Stateful autoconfiguration* uses a DHCPv6 server, which keeps a record of the IP address and other configuration information for the host. Whether a host engages in autoconfiguration and which method it uses are dictated by the routers serving the local network, not by a configuration of the host. There can be multiple routers serving the network, each potentially advertising multiple network prefixes. Thus, the host is not in full control of the number of IPv6

addresses that it configures, much less the values of those addresses, and the number and values of addresses can change as routers are added to or removed from the network.

When IPv6 autoconfiguration is enabled, the platform engages in stateless IPv6 autoconfiguration. When IPv6 autoconfiguration is disabled, the platform relinquishes usage of any autoconfigured IPv6 addresses that it may have acquired while it was enabled. This same enable or disable state also enables or disables the usage of a link local address for each managed entity, although a link local address continues to be generated for each nonchassis-based platform and for each CP of a chassis-based platform because those link local addresses are required for router discovery. The enabled or disabled state of autoconfiguration is independent of whether any static IPv6 addresses have been configured.

Setting IPv6 Autoconfiguration

Use the following procedure to enable IPv6 autoconfiguration:

1. Connect to the switch and log in using an account with admin permissions.
2. Take the appropriate action based on whether you want to enable or disable IPv6 autoconfiguration:
 - Enter the `ipAddrSet -ipv6 -auto` command to enable IPv6 autoconfiguration for all managed entities on the target platform.
 - Enter the `ipAddrSet -ipv6 -noauto` command to disable IPv6 autoconfiguration for all managed entities on the target platform.

NOTE

A manually configured DNS configuration has precedence over a DNS configuration broadcast by the DHCP or DHCPv6 server. So any DNS broadcast to a switch with a manual DNS configuration is ignored. Manual DNS configuration can overwrite an existing broadcast DNS configuration.

Setting the Ethernet Interface Mode and Speed

Network interfaces can be set to use one of the link operating modes: full duplex or autonegotiate.

Changing the link operating mode is not supported for all network interfaces or for all Ethernet network interfaces. On the CP blade in a Brocade director, the supported interfaces are eth0 and eth3. On all other platforms, only eth0 is supported.

For dual-CP systems, the `ethIf` command affects only the CP to which you are currently logged in. Therefore, to set the link operating mode on the active CP, you must issue the `ethIf` command on the active CP; and to set the mode on the standby CP, you must issue the `ethIf` command on the standby CP. During failover, the mode is retained separately for each CP because the physical links may be set to operate in different modes.

ATTENTION

Forcing the link to an operating mode not supported by the network equipment to which it is attached may result in an inability to communicate with the system through its Ethernet interface. It is recommended that the `ethIf` command be used only from the serial console port. When used through an interface other than the serial console port, the command displays a warning message and prompts for verification before continuing. This warning is not displayed and you are not prompted when the command is used through the serial console port.

Use the following procedure to set the mode of a port.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `ethIf` command.

Example of Setting the Port Mode to Full AutoNegotiate

The following example sets the mode for eth3 to autonegotiate and permits full duplex modes to be selected at 10, 100, and 1000Mb/s.

```
switch:admin> ethif --set eth3 -an on -speed 1000 -duplex full
an:on
speed:1000
cap:full
MII_CMD:-A
ADVERTISE:Advertise
DEFMODE:yes
auto:1
MII_MODE:1000baseTx-FD,
Committing configuration...done.
```

Example of Setting the Port Mode to 100Mb/s Full-Duplex Operation

The following example forces the link for the eth0 interface from autonegotiation to 100Mb/s full-duplex operation.

```
switch:admin> ethif --set eth0 -an on -speed 100 -duplex full
an:on
speed:100
cap:full
MII_CMD:-A
ADVERTISE:Advertise
DEFMODE:yes
auto:1
MII_MODE:100baseTx-FD,
Committing configuration...done.
```

Date and Time Settings

Switches maintain the current date and time inside a battery-backed real-time clock (RTC) circuit that receives the date and time from the fabric's principal switch. Date and time are used for logging events. Switch operation does not depend on the date and time; a switch with an incorrect date and time value functions properly. However, because the date and time are used for logging, error detection, and troubleshooting, you must set them correctly.

Authorization access to set or change the date and time for a switch is role-based. For an understanding of role-based access, see [Role-Based Access Control](#).

Setting the Date and Time

Use the following procedure to set the device date and time.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `date` command, using the following syntax:

```
date [MMDDhhmm[YY]]
```

The values represent the following:

- MM is the month; valid values are 01 through 12.
- DD is the date; valid values are 01 through 31.
- hh is the hour; valid values are 00 through 23.
- mm is minutes; valid values are 00 through 59.
- YY is the year; valid values are 00 through 37 and 70 through 99 (year values from 70 through 99 are interpreted as 1970 through 1999; year values from 00 through 37 are interpreted as 2000 through 2037).

The following example shows how to set a date.

```
switch:admin> date
```

```
Fri Feb 9 17:01:48 UTC 2019
```

```
switch:admin> date "0204101020"
```

```
Mon Feb 4 10:10:00 UTC 2020
```

A system-generated message `Date change within 64 seconds from previous change is not allowed` appears, if you try to change the date within 64 seconds of a previous change.

NOTE

From Fabric OS 9.1.0, the system is monitored and protected from disproportionate time changes; consequently the date cannot be changed beyond plus/minus 7 days using the `date` command. This function is called Restrictive Time Adjustment (RTA).

Time Zone Settings

You can set the time zone for a switch by name. You can specify the setting using country and city or time zone parameters. Switch operation does not depend on a date and time setting. However, having an accurate time setting is needed for accurate logging and audit tracking.

If the time zone is not set with new options, the switch retains the offset time zone settings. The `tstimezone` command includes an option to revert to the prior time zone format.

When you set the time zone for a switch, you can perform the following tasks:

- Display all of the time zones supported in the firmware.
- Set the time zone based on a country and city combination or based on a time zone ID, such as Pacific Standard Time (PST).

The time zone setting has the following characteristics:

- Users can view the time zone settings. However, only those with administrative permissions can set the time zones.
- The setting automatically adjusts for Daylight Savings Time.
- Changing the time zone on a switch updates the local time zone setup and is reflected in local time calculations.
- By default, all switches are set to Greenwich Mean Time (GMT) (0,0). If all switches in a fabric are in one time zone, it is possible for you to keep the time zone setup at the default setting.
- System services that have already started reflect the time zone changes after the next reboot.
- Time zone settings persist across failover for high availability.
- Setting the time zone on any dual domain backbone has the following characteristics:
 - Updating the time zone on any switch updates the entire chassis.
 - The time zone of the entire chassis is the time zone of switch 0.

Setting the Time Zone

The following procedure describes how to set the time zone for a switch. You must perform the procedure on *all* switches for which the time zone must be set. However, you need to set the time zone only once on each switch because the value is written to nonvolatile memory.

1. Connect to the switch and log in using an account assigned to the admin role and with the chassis-role permission.
2. Use the `tsTimeZone` command to set the time zone for the device.
 - Enter `tsTimeZone` with no parameters to display the current time zone setting.
 - Enter the `--interactive` option to list all of the time zones supported by the firmware.
 - Enter `timeZone_fmt` to set the time zone by country or city or by time zone ID, such as Pacific Standard Time (PST).

The following example displays and changes the time zone to US/Central.

```
switch:admin> tstimezone

Time Zone : US/Pacific
switch:admin> tstimezone US/Central
System Time Zone change will take effect at next reboot. On Chassis, reboot both CPs simultaneously.

switch:admin> tstimezone

Time Zone : US/Central
```

Setting the Time Zone Interactively

Use the following procedure to set the current time zone to the African time zone using interactive mode.

1. Connect to the switch and log in using an account assigned to the admin role and with the chassis-role permission.
2. Enter the `tsTimeZone --interactive` command.

You are prompted to select a general location.

```
Please identify a location so that time zone rules can be set correctly.
Please select a continent or ocean.
1) Africa
2) Americas
3) Antarctica
4) Arctic Ocean
5) Asia
6) Atlantic Ocean
7) Australia
8) Europe
9) Indian Ocean
10) Pacific Ocean
Enter number or control-D to quit ?1
```

3. Enter the appropriate number or press `Ctrl-D` to quit.

```
Please select a country.
```

```

1) Algeria          18) Gabon          35) Rwanda
2) Angola           19) Gambia         36) Sao Tome & Principe
3) Benin            20) Ghana          37) Senegal
4) Botswana         21) Guinea         38) Sierra Leone
5) Burkina Faso     22) Guinea-Bissau  39) Somalia
6) Burundi          23) Kenya        40) South Africa
                    41) South Sudan
7) Cameroon         24) Lesotho        42) Spain
8) Central African Rep. 25) Liberia       43) Sudan
9) Chad             26) Libya          44) Swaziland
10) Congo (Dem. Rep.) 27) Malawi         45) Tanzania
11) Congo (Rep.)     28) Mali           46) Togo
12) Cote d'Ivoire    29) Mauritania     47) Tunisia
13) Djibouti         30) Morocco        48) Uganda
14) Egypt            31) Mozambique     49) Western Sahara
15) Equatorial Guinea 32) Namibia        50) Zambia
16) Eritrea          33) Niger           51) Zimbabwe
17) Ethiopia         34) Nigeria

Enter number or control-D to quit ?1
The following information has been given:

```

Algeria

```

Therefore TZ='Africa/Algiers' will be used.
Local time is now:           Wed Mar 25 15:01:38 CET 2020.
Universal Time is now:      Wed Mar 25 14:01:38 UTC 2020.
Is the above information OK?
1) Yes
2) No

Enter number or control-D to quit ?1
System Time Zone change will take effect at next reboot. On Chassis, reboot both CPs simultaneously.

```

4. Enter the appropriate number at the prompt to specify the time zone region. Enter 1 to confirm or 2 to restart the time zone selection again. At any time, use **Ctrl-D** to quit.

Network Time Protocol

From Fabric OS 9.1.0, the time server module runs `ntpd` daemon replacing `ntpdate` to query the NTP server for the time to update the system clock on the principal or primary switch in the fabric.

The `ntpd` daemon allows the continuous adjustment of time on the switch that sets and maintains the system time. The NTP service runs only on the principal or primary switch when one or more external servers are configured as the switch clock source. Legacy mode is a new feature that is used on the `ntpd` daemon to allow the user to continue to use the NTP servers as the legacy behavior.

When legacy mode is enabled, only the active server is used to synchronize the time. If the active server fails, then the next server in the server list becomes the active server and so on. When legacy mode is disabled, the `ntpd` daemon simultaneously uses all the servers in the configured server list to synchronize the time.

You can synchronize the local time of the principal or primary FCS switch to a maximum of eight external NTP servers.

All switches in the fabric maintain the current NTP clock server value in nonvolatile memory. By default, this value is the local clock (LOCL) of the principal or primary FCS switch. Changes to the clock server value on the principal or primary FCS switch are propagated to all switches in the fabric.

When a new switch enters the fabric, the principal or primary FCS switch sends the new switch the time and the addresses of all existing clock servers. When a switch enters the fabric, it stores the list and the active servers.

NOTE

In a fabric, if the FCS policy is enabled and all the switches are not defined in the FCS policy, it is possible that the fabric principal switch is not the same as the primary switch and not in the FCS policy. In this scenario, `ntpd` would run in both the primary switch and the principal switch of the fabric.

NOTE

In non-legacy mode, the active NTP server is chosen by NTP based on its selection algorithm. If no server is chosen yet, the active NTP server is displayed as NONE. The non-primary or non-principal switch always displays the active NTP server as NONE in non-legacy mode.

If virtual fabrics are enabled, the switch behavior is as follows:

- All switches in a given chassis must be configured for the same set of NTP servers. This ensures that the time does not go out of sync in the chassis. Do not configure LOCL in the NTP server list.
- The logical switches in chassis get their clock information from the default logical switch, and not from the principal or primary FCS switch.

Synchronizing the Local Time with an External Source

The `tsClockServer` command accepts multiple NTP server addresses in IPv4, IPv6, or DNS name formats. When multiple NTP server addresses are specified and legacy mode is enabled, `tsClockServer` sets the first synchronized server as the active NTP server. The remaining addresses are stored as backup servers that can take over if the active NTP server fails. When legacy mode is disabled, all servers in the clock server list are used as NTP servers. The principal or primary FCS switch continuously polls the NTP server at an interval from 64 seconds through 1024 seconds to synchronize and adjust its time with the NTP server. As an NTP client, a FOS switch operates in nonburst mode, and a single packet is sent on every poll. This causes initial synchronization with the NTP server to take between 6 and 8 minutes.

The time server daemon synchronizes fabric time by periodically sending updates of the principal or primary FCS local switch time to every switch in the fabric. The time server daemon runs on all switches in the fabric, but only the principal switch (when an FCS policy is not enabled) or the primary FCS switch (when an FCS policy is enabled) broadcasts the time updates across the fabric.

The following example sets the NTP server to a specified address.

```
switch:admin> tsclockserver 192.0.12.1
Updating Clock Server configuration...done.
Updated with the NTP servers
```

The following example displays the NTP server and its parameters.

```
switch:admin> tsclockserver

Active NTP Server          198.51.100.12
Configured NTP Server List 198.51.100.12
Configured NTP Key Index List NONE
Configured NTP Authspec Mode SYMMETRIC AUTHENTICATION ENABLED
Configured NTP Legacy Mode LEGACY MODE ENABLED
```

The following example configures multiple NTP servers.

```
switch:admin> tsclockserver "198.51.100.12; 198.51.100.13"
Updating Clock Server configuration...done.
Updated with the NTP servers
```

The following example shows how to enable legacy mode. By default, the legacy mode is enabled.

```
switch:admin> tsclockserver --legacy enable
```

The following example displays the NTP server authentication state.

```
switch:admin> tsclockserver --show

Active NTP Server           LOCL
Configured NTP Server List  LOCL
Configured NTP Key Index List  NONE
Configured NTP Authspec Mode  NTP AUTHENTICATION DISABLED
Configured NTP Legacy Mode    LEGACY MODE ENABLED
```

NOTE

From Fabric OS 9.1.0, the system is monitored and protected from disproportionate time changes; consequently the date cannot be changed beyond plus/minus 7 days using the `tsclockserver` command. This function is called Restrictive Time Adjustment (RTA).

Configuring Symmetric Authentication Keys

Symmetric authentication plays a major role in identifying the secure NTP servers from unauthorized servers. Symmetric authentication is based on a secret key that is shared by both client and server. This symmetric key is used for authentication. The keys must use only hexadecimal characters (0 1 2 3 4 5 6 7 8 9 A B C D E F a b c d e f).

The following table lists the symmetric keys that are used for symmetric authentication.

Table 6: Symmetric Keys and Their Minimum and Maximum Lengths

Symmetric Keys	Minimum Length	Maximum Length
MD5	20	64
HMAC-SHA1	40	64
HMAC-SHA256	64	64
CMAC-AES-128	32	64

The existing `tsClockServer` command is enhanced to enable or disable symmetric authentication. The enhanced `tsClockServer` command is used to support symmetric key index, authspec mode, and legacy mode configurations. By default, symmetric authentication is disabled. You can use `tsClockServer --authspec symmetric` to enable symmetric authentication or `tsClockServer --authspec noauth` to disable symmetric authentication. Symmetric authentication is enabled or disabled whether or not legacy mode is enabled.

The following example shows how to add the symmetric key:

```
switch:admin> tsclockserver --addkey -index 16 -type HMAC-SHA1 -key ****abcd****efgh****1234****ijkl****5678

Adding authentication key ...
```

The following example shows how to delete the symmetric key:

```
switch:admin> tsclockserver --delkey -index 16
```

```
Deleting authentication key ...
```

The following example displays the configured symmetric key in hidden format:

```
switch:admin> tsclockserver --showkeys
NTP Authentication Keys:
100 HMAC-SHA1 ****abcd****efgh****1234****ijkl****5678
```

The following example displays the configured symmetric key in plain text format. Use the `-s` option to view the key in plain text output. The following command is restricted to the admin user only.

```
switch:admin> tsclockserver --showkeys -s
NTP Authentication Keys:
100 HMAC-SHA1 ****abcd****efgh****1234****ijkl****5678
```

The following example shows how to configure the clock server with symmetric authentication.

```
switch:admin> tsclockserver --set "10.11.2.80" -index 100
Updating Clock Server configuration...
```

```
switch:admin> tsclockserver
Active NTP Server           10.1.1.1
Configured NTP Server List  10.1.2.2
Configured NTP Key Index List 100
Configured NTP Authspec Mode NTP AUTHENTICATION DISABLED
Configured NTP Legacy Mode   LEGACY MODE ENABLED
```

The following example shows how to modify the symmetric authentication key for the configured server.

```
switch:admin> tsclockserver --change "10.1.1.1" -index 110
Updating Clock Server key index...
```

The following example shows how to enable symmetric authentication.

```
switch:admin> tsclockserver --authspec symmetric
Updating Clock Server authspec...
```

```
switch:admin> tsclockserver
Active NTP Server           10.1.1.2
Configured NTP Server List  10.1.1.2
Configured NTP Key Index List 110
Configured NTP Authspec Mode SYMMETRIC AUTHENTICATION ENABLED
Configured NTP Legacy Mode   LEGACY MODE ENABLED
```

The following example shows how to disable authentication.

```
switch:admin> tsclockserver --authspec noauth
Updating Clock Server authspec...
```

```
switch:admin> tsclockserver
Active NTP Server           10.1.1.2
Configured NTP Server List  10.1.1.2
Configured NTP Key Index List 110
Configured NTP Authspec Mode NTP AUTHENTICATION DISABLED
Configured NTP Legacy Mode   LEGACY MODE ENABLED
```

Distributing the NTP Configuration to Access Gateways

Any switch can distribute the NTP server configuration to core Access Gateway (AG) devices connected to the fabric. The core AG devices distribute the NTP configuration to other cascaded or edge AG devices. However, AG devices are not capable of distributing the NTP configuration to other switches in the fabric.

In switch mode, the principal or primary FCS switch synchronizes its time with the external NTP server every 64 seconds and sends time updates to other switches in the fabric. The time updates are not sent in-band to AG devices. Each AG device separately synchronizes with the external NTP server and updates the time. If the AG device is connected to more than one fabric, the latest clock server request received is used.

Consider the following points when distributing the NTP server configuration:

- The `tsClockServer` command distributes the NTP server configuration to all switches within the fabric and to AG devices connected to the same fabric. The `tsClockServer` configuration distribution from the fabric switches to the core AGs is limited to only the NTP clock servers. The symmetric key details, authspec mode, and legacy mode configuration are excluded from the fabric switch distribution to the connected AGs.

NOTE

However, you can configure NTP symmetric authentication in individual AG switches running Fabric OS 9.1.0.

- Already distributed NTP server configurations will persist on the AG device after downgrading the firmware from Fabric OS 8.0.1 or later to an earlier version on supported platforms. However, an AG device cannot distribute any configuration to edge AG devices.

Domain IDs

Although domain IDs are assigned dynamically when a switch is enabled, you can change them manually so that you can control the ID number or can resolve a domain ID conflict when you merge fabrics.

If a switch has a domain ID when it is enabled, and that domain ID conflicts with another switch in the fabric, the conflict is automatically resolved if the other switch domain ID is not persistently set. The process can take several seconds, during which time traffic is delayed. If both switches have their domain IDs persistently set, one of them must have its domain ID changed to a domain ID not used within the fabric.

NOTE

The default domain ID for Brocade switches is 1. When additional switches are enabled, which assigns the default ID, it is recommended that you change the default ID to an insistent domain ID. This will avoid future conflicts and unexpected changes.

The xlate domain ID (XD) or front domain ID (FD) for an FC router is assigned in the following ranges, if the insistent domain ID is not configured.

- The FD range for an FC router is 160 through 199. The default front domain ID requested by an FC router is 160.
- The XD range for an FC router is 200 through 239. The default xlate domain ID requested by an FC router is 200.

NOTE

If there is no free domain ID available in the specified range for XD or FD, the principal switch assigns the lowest available domain ID starting from 1.

The principal switch overrides the specified domain ID range if the following conditions are true:

- The insistent domain ID is configured for FD or XD.
- The requested domain ID is not in the specified range.
- The requested domain ID is available for domain ID assignment.

When a switch other than the FC router FD or XD requests a domain ID within the new specified range for the FC router FD or XD, the principal switch assigns the domain ID, if it is available.

NOTE

It is recommended that the insistent domain ID be configured when deploying features that rely on domain IDs, such as [D, I] zoning.

Domain ID Issues

Keep the following restrictions in mind when working with domain IDs:

- Do not use domain ID 0. Using this domain ID can cause the switch to reboot continuously.
- Avoid changing the domain ID on the FCS switch in secure mode.
- To minimize downtime, change the domain IDs on the other switches in the fabric.

Displaying the Domain IDs

Use the following procedure to display device domain IDs.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `fabricShow` command.

The following is an example of output of fabric information, including the domain ID (D_ID).

The principal switch is determined by the arrow (>) next to the name of the switch.

```
switch:admin> fabricshow
```

Switch ID	Worldwide Name	Enet IP Addr	FC IP Addr	Name
2: fffc02	10:00:00:60:69:e0:01:46	10.3.18.2	0.0.0.0	"ras001"
3: fffc03	10:00:00:60:69:e0:01:47	10.3.18.3	0.0.0.0	"ras002"
5: fffc05	10:00:00:05:1e:34:01:bd	10.3.18.5	0.0.0.0	"ras005"
	fec0:60:69bc:63:205:1eff:fe34:1bd			
6: fffc06	10:00:00:05:1e:34:02:3e	10.3.18.6	0.0.0.0	>"ras006"
7: fffc07	10:00:00:05:1e:34:02:0c	10.3.18.7	0.0.0.0	"ras007" (output truncated)

The Fabric has 26 switches
Fabric Name: Non_VF

The following table displays the `fabricShow` fields.

Table 7: fabricShow Fields

Fields	Description
Switch ID	The switch domain_ID and embedded port D_ID. The numbers are broken down as shown in the following example: 64: fffc40 64 is the switch domain_ID fffc40 is the hexadecimal format of the embedded port D_ID.
Worldwide Name	The switch WWN.
Enet IP Addr	The switch Ethernet IP address for IPv4-configured and IPv6-configured switches. For IPv6 switches, only the static IP address displays.
FC IP Addr	The switch Fibre Channel IP address.
Name	The switch symbolic or user-created name in quotes.

Setting the Domain ID

Use the following procedure to set the domain ID:

1. Connect to the switch and log in on an account assigned to the admin role.
2. Enter the `switchDisable` command to disable the switch.
3. Enter the `configure` command.
4. Enter `y` after the `Fabric parameters` prompt.

```
Fabric parameters (yes, y, no, n): [no] y
```

5. Enter a unique domain ID at the `Domain` prompt. Use a domain ID value from 1 to 239 for normal operating mode (FCSW-compatible).

```
Domain: (1..239) [1] 3
```

6. Respond to the remaining prompts, or press **Ctrl-D** to accept the other settings and exit.
7. Enter the `switchEnable` command to re-enable the switch.

Registered Organization Name

A registered organization name (RON) is a 32-byte string stored in a World Wide Name (WWN) card.

The following considerations apply to the registered organization name:

- The RON can begin with a printable ASCII letter, number, or symbol. When upgrading to Fabric OS 8.2.0, if the switch is already configured with an IP address, you must change the IP address to permit RON configuration.
- The maximum length for the organization name is 28 characters.
- Once the organization name is set, it cannot be field-modified.
- The encrypted content is collected as part of supportSave, and cannot be modified.
- The `ron` command has the following operands:

```
--set org_name : Sets the registered organization name with the specified name.
--show : Displays the registered organization name configured in the system.
--help : Displays the command information and usage.
```

Setting and Displaying the Registered Organization Name

This section sets and displays the RON, which is stored in a WWN card.

1. Connect to the switch, and log in using an account assigned to the admin role.
2. Enter `ron --set org_name` to set the registered organization name. Once set, the organization name cannot be modified.

```
switch:admin> ron --set "XXXX_Organization"
Registered Organization Name will be set to: XXXX_Organization
Once changes are committed, it cannot be modified.
Are you sure you want to commit these changes? (Y/N)?y
Registered Organization Name is set successfully.
switch:admin> ron --set "YYYY_Organization"
Registered Organization Name is already set.
```

3. Enter the `ron --show` command to display the registered organization name.

```
switch:admin> ron --show
Registered Organization Name: XXXX_Organization
Registration complete on: Sep 2017
```

4. Record the registered organization name for future reference.

Switch Name

A switch can be identified by an IP address, a domain ID, a World Wide Name (WWN), or a customized switch name.

The following considerations apply to switch names:

- Switch names can be from 1 to 30 characters long.
- Case is recorded, but is not enough to make a switch name unique. (“MySwitch24” is considered the same as “myswitch24”).
- All switch names must begin with a letter or number and can contain letters, numbers, hyphens, periods, or underscore characters.
- A switch name that begins with a numeric character must also contain at least one of the following:
 - An alphabetic (A-Z, a-z) character
 - An underscore (_)
 - A dash (-)
 - A period (.)
- Switch names can be identical, but Brocade recommends keeping them unique.
- Changing the switch name causes a domain address format RSCN to be issued and might be disruptive to the fabric.

NOTE

When FICON Management Server (FMS) mode is enabled, the switch name can be a maximum of 24 characters long. All other name restrictions remain in effect.

Customizing the Switch Name

This section describes how to customize the switch name.

1. Connect to the switch, and log in using an account assigned to the admin role.
2. Enter the `switchName newname` command using a new name for the switch.

The prompt does not change to the new switch name until after you log in again.

```
switch:admin> switchname myswitch
Committing configuration...
Done.
Switch name has been changed. Please re-login into the switch for the change to be applied.
switch:FID128:# admin>
```

3. Record the new switch name for future reference.

Switch Type

The `switchType` attribute is a unique device model identifier that is displayed when you issue the `switchshow` command. When you are gathering information for your switch support provider, you may be asked for the Brocade product name. Use the information in the following table to convert the `switchType` identifier to a Brocade product name.

Table 8: Mapping Switch Type to Brocade Product Name

Switch Type	Brocade Product Name	Description
165	Brocade X6-4	Chassis with a bandwidth of 10Tb/s. It supports 256 device ports with a 32G data rate plus 16 4xGen 6 ICLs.
166	Brocade X6-8	Chassis with a bandwidth of 20Tb/s. It supports 512 device ports with a 32G data rate plus 32 4xGen 6 ICLs.
162 or 183	Brocade G620	Fixed-port switch with configurable scaling from 24 to 64 ports. It supports 32G, 16G, and 10G media types.
170	Brocade G610	Fixed-port switch with configurable scaling from 8 to 24 ports. It supports 32G, 16G, and 10G media types.
173 or 184	Brocade G630	Fixed-port switch with configurable scaling from 48 ports to 128 ports. It supports 32G, 16G, and 10G media types.
178	Brocade 7810 Extension Switch	Extension switch with configurable scaling of 12 ports and supports 32G, 16G, 8G, and 4G media types. It supports 1GbE and 10GbE Ethernet interfaces.
179	Brocade X7-4 Director	Chassis with a bandwidth of 15.5Tb/s. It supports 192 device ports with a 64G data rate plus 16 4xGen 7 ICLs.
180	Brocade X7-8 Director	Chassis with a bandwidth of 31Tb/s. It supports 384 device ports with a 64G data rate plus 32 4xGen 7 ICLs.
181	Brocade G720 Switch	Fixed-port switch with configurable scaling from 24 to 64 ports. It supports 64G, 32G, 16G, and 10G media types.
189	Brocade G730 Switch	Fixed-port switch with configurable scaling from 48 ports to 128 ports. It supports 64G, 32G, 16G, and 10G media types.

Chassis Name

Brocade recommends that you customize the chassis name for each platform. Some system logs identify devices by platform names; if you assign meaningful platform names, logs are more useful. All chassis names support 31 characters. Chassis names must begin with an alphabetic character and can include alphabetic and numeric characters, and the underscore (_).

Customizing the Chassis Name

Use the following procedure to customize the chassis name:

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `chassisName` command.

```
ecp:admin> chassisname newname
```

3. Record the new chassis name for future reference.

Fabric Name

You can assign a alphanumeric name to identify and manage a logical fabric that formerly could only be identified by a fabric ID. The fabric name does not replace the fabric ID or its usage. The fabric continues to have a fabric ID, in addition to the assigned alphanumeric fabric name.

The following considerations apply to fabric naming:

- Each name must be unique for each logical switch within a chassis; duplicate fabric names are not allowed.
- A fabric name can be from 1 through 128 alphanumeric characters (a–z, 0–9).
- All switches in a logical fabric must be running Fabric OS v7.2.0 or later. Switches running earlier versions of the firmware can coexist in the fabric, but do not show the fabric name details.
- You must have admin permissions to configure the fabric name.

Configuring the Fabric Name

To set and display the fabric name, use the `fabricname` command.

```
switch:admin> fabricname --set myfabric1
```

Using the `fabricname --set` command without a fabric name takes the existing fabric name and synchronizes it across the entire fabric. An error message displays if no name is configured.

To set a fabric name that includes spaces, enclose the fabric name in double quotes.

```
switch:admin> fabricname --set "my new fabric"
Fabric Name is Set to my new fabric.
```

To set a fabric name that includes special metacharacters, enclose the fabric name in single quotes.

```
switch:admin> fabricname --set 'red fabric $$'
```

To clear the fabric name, use the `fabricname --clear` command.

High Availability Considerations for Fabric Names

Fabric names that are locally configured or obtained from a remote switch are saved in the configuration database and then synchronized to the standby CP on dual-CP-based systems.

Upgrade and Downgrade Considerations for Fabric Names

Fabric names are lost during a firmware downgrade. No default fabric name is provided. If a fabric name is needed, it must be configured after the upgrade.

Switch Activation and Deactivation

By default, the switch is enabled after power is applied and diagnostics and switch initialization routines have finished. You can disable and re-enable the switch as necessary.

When you enable or disable a switch, the affected ports depend on whether Virtual Fabrics is enabled. The following table describes which ports are affected for each type of enable or disable operation.

Table 9: Ports Affected When You Enable or Disable a Switch in VF or Non-VF Mode

Operation	Virtual Fabrics Enabled	Virtual Fabrics Not Enabled
Enable switch	Enables all ports on the logical switch	Enables all ports on the physical chassis
Enable chassis	Enables all ports on the physical chassis	Not allowed
Disable switch	Disables all ports on the logical switch	Disables all ports on the physical chassis

Operation	Virtual Fabrics Enabled	Virtual Fabrics Not Enabled
Disable chassis	Disables all ports on the physical chassis	Not allowed

Disabling a Switch

You must disable a switch before making configuration changes or before running offline diagnostic tests.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `switchDisable` command.

```
switch:admin> switchdisable
```

All Fibre Channel ports on the switch are taken offline. If the switch is part of a fabric, the fabric is reconfigured.

If Virtual Fabrics is enabled, only the ports allocated to the logical switch are disabled. To disable all of the ports, you must disable the entire chassis.

Enabling a Switch

The switch is enabled by default after it is powered on and switch initialization routines have finished. You must re-enable the switch after making configuration changes or running offline diagnostics.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `switchenable` command.

```
switch:admin> switchenable
```

All Fibre Channel ports that passed the power-on self-test (POST) are enabled. If the switch has inter-switch links (ISLs) to a fabric, it joins the fabric.

If Virtual Fabrics are enabled, only the ports allocated to the logical switch are enabled. To enable all of the ports, you must enable the entire chassis.

Disabling a Chassis

Disabling a chassis disables all Fibre Channel ports on all logical switches in the chassis. You must disable a chassis before making chassis-wide configuration changes or before running offline diagnostic tests.

1. Connect to any logical switch in the chassis, and log in using an account assigned to the admin role.
2. Enter the `chassisdisable` command.

```
switch:FID128:admin> chassisdisable
```

```
This command can cause disruption to multiple logical switches.
```

```
Are you sure you want to disable all chassis ports now? (yes, y, no, n): [no]y
```

```
switch:FID128:admin>
```

All Fibre Channel ports on all logical switches are taken offline. If the logical switches are in fabrics, the fabrics are reconfigured.

NOTE

After a `chassisDisable`, if you want to do an `haFailover`, you should wait at least 30 seconds.

Enabling a Chassis

Enabling a chassis enables all Fibre Channel ports on all logical switches in the chassis. The chassis is enabled by default after it is powered on and switch initialization routines have finished. You must re-enable the chassis after making fabric-wide configuration changes or running offline diagnostics.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `chassisEnable` command.

```
switch:FID128:admin> chassisenable
```

For all logical switches in the chassis, all Fibre Channel ports are enabled. If any of the logical switches has inter-switch links (ISLs) to a fabric, it joins the fabric.

Device Shutdown

To avoid corrupting your file system, you must perform graceful shutdowns of Brocade switches and directors by using the following instructions.

Powering Off a Brocade Switch

Use the following procedure to gracefully shut down a Brocade switch.

1. Connect to the switch and log in using an account that is assigned to the admin role.
2. Enter the `sysShutdown` command.
3. Enter `y` at the prompt.

```
switch:admin> sysshutdown
```

```
This command will shutdown the operating systems on your switch.  
You are required to power-cycle the switch in order to restore operation.  
Are you sure you want to shutdown the switch [y/n]?y
```

4. Wait until the following message displays:

```
The system is going down for system halt NOW !!  
INIT: Switching to runlevel: 0  
INIT: Sending processes the TERM signal  
Unmounting all filesystems.  
The system is halted  
flushing ide devices: hda  
Power down.
```

5. Power off the switch.

Powering Off a Brocade Director

Use the following procedure to power off Brocade X7 and X6 Directors:

1. From the active CP in a dual-CP platform, enter the `sysShutdown` command.

NOTE

When the `sysShutdown` command is issued on the active CP, the active CP, the standby CP, and any application blades are all shut down.

2. Enter `y` at the prompt.
3. Wait until the following message displays:

```
FID128:admin> sysshutdown

This command will shutdown the operating systems on your switch.
You are required to power-cycle the switch in order to restore operation.
Are you sure you want to shutdown the switch [y/n]?y
HA is disabled
Stopping blade 10
Shutting down the blade....
Stopping blade 12
Shutting down the blade....
...
...
The system is going down for system halt NOW !!
```

4. Power off the director.

Basic Connections

Before connecting a switch to a fabric that contains switches running different firmware versions, you must first set the same port identification (PID) format on all switches. The presence of different PID formats in a fabric causes fabric segmentation.

- For information on PID formats and related procedures, see [Performing Advanced Configuration Tasks](#).
- For information on configuring the routing of connections, see [Routing Traffic](#).
- For information on configuring extended inter-switch connections, see [Managing Long-Distance Fabrics](#).

Device Connection

To minimize port logins, power off all devices before connecting them to the switch. When powering the devices back on, wait for each device to complete the fabric login before powering on the next one.

For devices that cannot be powered off, first use the `portDisable` command to disable the port on the switch, connect the device, and then use the `portEnable` command to enable the port.

Switch Connection

Refer to the hardware reference manual of your specific switch for ISL connection and cable management information. The standard or default ISL mode is L0. ISL mode L0 is a static mode with the following maximum ISL distances:

- 10 km at 1G
- 5 km at 2G
- 2.5 km at 4G
- 1 km at 8G
- 1 km at 10G
- 625m at 16G
- 150m at 32G

For more information on extended ISL modes, which enable long-distance inter-switch links, see [Managing Long-Distance Fabrics](#).

DNS

You can define a domain name service (DNS) domain on a Brocade switch and thereby recognize all hosts within that domain. After specifying a domain name, the Brocade switch automatically appends the appropriate domain to the host and forwards it to the domain name server. You can configure DNS parameters such as the domain name and the name server IP address for primary and secondary name servers.

The `switchname` command that is used to display or configure the host name conforms to RFC 1123. White space, symbols other than letters, numbers, hyphens, periods, and underscore characters are not allowed for a host name. A host name can contain either a letter or number. The host name that begins with a numeric character must at least have an alphabetic character, an underscore, or a dash or a period.

Performing Advanced Configuration Tasks

This section describes about Port Identifiers (PIDs), different types of addressing modes, different port types, port management, FC Link Aggregation, blade types, and power management.

Port Identifiers and PID Binding Overview

Port identifiers (PIDs, also called *fabric addresses*) are used by the routing and zoning services in Fibre Channel fabrics to identify ports in the network.

All devices in a fabric must use the same PID format. When you add new equipment to the SAN, you might need to change the PID format on legacy equipment.

Many scenarios cause a device to receive a new PID: for example, unplugging the device from one port and plugging it into a different port as part of fabric maintenance; or changing the domain ID of a switch, which might be necessary when merging fabrics; or changing compatibility mode settings.

Some device drivers use the PID to map logical disk drives to physical Fibre Channel counterparts. Most drivers can either change PID mappings dynamically, also called *dynamic PID binding*, or use the WWN of the Fibre Channel disk for mapping, also called *WWN binding*.

Some older device drivers behave as if a PID uniquely identifies a device; they use *static PID binding*. These device drivers should be updated, if possible, to use WWN binding or dynamic PID binding instead, because static PID binding creates problems in many routine maintenance scenarios. Fortunately, very few device drivers still behave this way. Many current device drivers enable you to select static PID binding as well as WWN binding. You should select static PID binding only if there is a compelling reason and only after you have evaluated the effect of doing so.

Unified Addressing Mode

Unified addressing mode (UAM) provides a single addressing mode for an 8-bit area, 10-bit area, or shared area, and it provides other addressing modes for different combinations of platform and logical switch type. UAM simplifies these addressing mode combinations and features into a single unified addressing mode. UAM is introduced in FOS 9.0.0 and is the default address allocation mode for next generation platforms such as the Brocade X7-8 Director, X7-4 Director, and G720 Switch.

The following table describes how UAM works for FOS features:

Table 10: Difference between UAM and pre-UAM Mode

Feature	UAM Mode	Pre-UAM Mode
F-Port	A port is assigned only a base area.	A port is assigned only a base area.
F_Port with NPIV	For F_Ports with NPIV less than 63 devices, port is assigned only a base area. If the 64th NPIV device is received in F_Port, then a new additional area is assigned to the existing base area. Every time a port exhausts an area after 64 additional devices login, another additional area is assigned. The maximum number of NPIVs supported is 255 devices.	The 10-bit addressing mode utilizes the 8-bit area ID and the borrowed upper 2 bits from the AL_PA portion of the PID. Areas 0x00 through 0x8F use only 8 bits for the port address and support upto 255 NPIV devices. Areas 0x90 through 0xFF use an additional 2 bits from the AL_PA for the port address. Therefore, these ports support only 64 NPIV devices per port.

Feature	UAM Mode	Pre-UAM Mode
F_Port Trunking	Same as F_Port and F_Port with NPIV. In FOS 9.0, UAM allocates additional areas to both F_Port trunking and Emulex HBA trunking.	F_Port trunking uses an 8-bit area in VF mode enabled logical switch. Default switch in VF or non-VF mode uses existing port address.
EX_Port	EX_Port in a base switch is assigned a 10-bit area.	EX_Port in a base switch uses an existing port address as per the current switch addressing mode.
EX_Port on ICL Ports	When an ICL port is configured as an EX_Port, a new area is assigned to the ICL port from the pool. This is not applicable for ICL ports that are not configured as an EX_Port. From the FOS 9.0 release, when an ICL port is configured as an EX_Port, the new areas are assigned in both VF mode and non-VF mode.	When an ICL is configured as an EX_Port, the switch is assigned an 8-bit area in VF. EX_Port over ICL is not supported in non-VF mode.
EX_Port Trunking	In UAM mode, an EX_Port trunk behaves in a controllerless manner both in the base switch and in non-VF mode.	In pre-UAM mode, an EX_Port trunk behaves in a controllerless manner only in VF mode.
FCoE	In UAM mode, FCoE address assignment is simple. Whenever Ethernet flexport requires an area, it will be assigned an area from the pool of unused areas.	In pre-UAM mode, whenever Ethernet flexport requires an area, it will be assigned an area from the pool of unused areas. For default switches in VF mode or non-VF mode, the switches are converted to a 10-bit area mode as soon as FCoE enodes are configured.
Port Address Binding	In platforms that are enabled in UAM mode, the port address binding feature is supported in both VF and non-VF mode.	In pre-UAM mode, the port address binding feature is supported only in logical switches in VF enabled mode.
Persistent PID	In platforms that are enabled in UAM mode, the persistent PID feature is supported in both VF and non-VF mode.	In pre-UAM mode, the persistent PID feature is supported only in logical switches in VF enabled mode.

Core PID Addressing Mode

The core PID is the default PID format for Brocade platforms. It uses the entire 24-bit address space of the domain ID, area ID, and AL_PA to determine an object's address within the fabric.

The core PID is a 24-bit address built from the following three 8-bit fields:

- Domain ID, written in hex and the numeric range is from 01 through ee (1 through 239)
- Area ID, written in hex and the numeric range is from 01 through ff (1 through 255)
- AL_PA

For example, if a device is assigned an address of 0f1e00, the following would apply:

- 0f is the domain ID.
- 1e is the area ID.
- 00 is the assigned AL_PA.

From this information, you can determine which switch the device resides on from the domain ID, which port the device is attached to from the area ID, and if this device is part of a loop from the AL_PA number.

Fixed Addressing Mode

With fixed addressing mode, each port has a fixed address assigned by the system based on the port number.

Fixed addressing mode is the default addressing mode used in all platforms that do not have Virtual Fabrics enabled. When Virtual Fabrics is enabled on a director, fixed addressing mode is used only on the default logical switch.

NOTE

Fixed addressing mode is supported only on the default switch of Gen 6 platforms.

10-Bit Addressing (Mode 0)

The 10-bit addressing mode is the default mode for all logical switches. This addressing scheme is flexible to support a large number of F_Ports.

NOTE

The default switch in Brocade directors uses the fixed addressing mode.

The 10-bit addressing mode utilizes the 8-bit area ID and the borrowed upper 2 bits from the AL_PA portion of the PID. Areas 0x00 through 0x8F use only 8 bits for the port address and support up to 255 NPIV devices. The supported 10-bit addressing mode configuration in X6 directors is as follows:

- For a logical switch in an X6-4, it supports up to 144 ports that can each support 256 devices. Areas running from 0x90 through 0xFF use an additional 2 bits from the AL_PA for the port address. Therefore, these ports support only 64 NPIV devices per port.
- For a logical switch in an X6-8, it supports up to 112 ports that can each support 256 devices. Areas running from 0x70 through 0xFF use an additional 2 bits from the AL_PA for the port address.

The 10-bit addressing mode provides the following features:

- A PID is dynamically allocated only when the port is first moved to a logical switch, and, thereafter, it is persistently maintained.
- PIDs are assigned in each logical switch starting with 0xFFC0, and can go to 0x8000 in the case of 64-port blades.
- Shared area limitations are removed on 48-port and 64-port blades.
- Any port on a 48-port or 64-port blade can support up to 255 NPIV devices (in fixed addressing mode, only 128 NPIV devices are supported in non-VF mode and 64 NPIV devices are supported in VF mode on a 48-port blade).
- Any port on a 48-port or 64-port blade can support hard port zoning.
- The port index is not guaranteed to be equal to the port area ID.

NOTE

The 10-bit addressing mode is supported only on logical switches of Gen 6 platforms.

256-Area Addressing (Mode 1 and Mode 2)

The 256-area addressing mode is available only in a logical switch on a Brocade director. In this mode, only 256 ports are supported, and each port receives a unique 8-bit area address. This mode can be used in FICON environments, which have strict requirements for 8-bit area FC addresses.

There are two types of area assignment modes with 256-area addressing: zero-based and port-based.

ATTENTION

On default logical switches with an address mode other than mode 1, any 48-port and 64-port blades are disabled if FICON Management Server (FMS) is enabled.

NOTE

The 256-area addressing mode is supported only on the logical switch of Gen 6 platforms.

See the [Administering FICON Fabrics](#) chapter for additional details.

Zero-Based Addressing (Mode 1)

With zero-based addressing, unique area assignments begin at zero regardless of where the port is physically located, which allows FICON users to make use of high port count blades with port indexes greater than 256.

Zero-based addressing assigns areas when the ports are added to the logical switch, beginning at area 0x00. When a port is assigned to a logical switch, the next free PID starting from 0x00 is assigned. This mode allows FICON customers to make use of the upper ports of a 48-port or 64-port blade.

Zero-based mode is supported on the default switch.

NOTE

In Gen 7 platforms, this mode is not configurable and is available only when creating a FICON logical switch.

Port-Based Addressing (Mode 2)

With port-based addressing, unique area assignments are based on the port index.

Port-based addressing mode does not allow FICON users to make use of ports with an index greater than 256 (high ports of a high port count blade), but this mode is compatible with domain-index zoning.

Port-based addressing has the following restrictions:

- Port-based addressing mode is not supported on the default switch.
- Port-based addressing mode is supported only on ports 0–31 of the blade in the X6-8 chassis.

NOTE

This mode is not supported in Gen 7 platforms.

Port Address Binding

In an FC switch, a user can bind any 8-bit address in both VF and non-VF mode. In the chassis, a user can bind a 10-bit address in both the base switch (0x80 through 0xFFC0) and the logical switch (0x90 through 0xFFC0). This feature is not supported in non-VF mode. In UAM mode, a user can bind a 10-bit address in both VF and non-VF mode across any address range.

WWN-Based PID Assignment

WWN-based PID assignment is disabled by default. When the feature is enabled, bindings are created dynamically; as new devices log in, they automatically enter the WWN-based PID database. The bindings exist until you explicitly unbind the mappings through the CLI or change to a different addressing mode. If there are any existing devices when you enable the feature, you must manually enter the WWN-based PID assignments through the CLI.

This feature also allows you to configure a PID persistently using a device WWN. When the device logs in to the switch, the PID is bound to the device WWN. If the device is moved to another port in the same switch, or a new blade is hot-plugged, the device receives the same PID (area) at its next login.

After WWN-based PID assignment is enabled, you must manually enter the WWN-based PID assignments through the CLI for any existing devices.

NOTE

When the WWN-based PID assignment feature is enabled, area assignment is dynamic and has no predetermined order or precedence. This means that on the first fabric login of a PID, the PID will be assigned the next available WWN area and it will remain mapped to that WWN after that point (including subsequent logins). Static WWN-area bindings are preserved, but are not used for dynamic assignments.

PID assignments are supported for a maximum of 4096 devices; this includes both point-to-point and N-Port Identifier Virtualization (NPIV) devices. The number of point-to-point devices supported depends on the areas available. For example, 448 areas are available on directors and 256 areas are available on switches. When the number of entries in the WWN-based PID database reaches 4096 areas used up, the oldest unused entry is purged from the database to free up the reserved area for the new FLOGI.

Virtual Fabrics Considerations for WWN-Based PID Assignment

WWN-based PID assignment is disabled by default, and it is supported in the default switch on Brocade X6 Directors, but it does not require the 8-bit dynamic area mode be enabled. Enabling Dynamic Area Mode requires the switch to be disabled. This feature is not supported on the SX6 blades. The total number of ports in the default switch must be 256 or less.

When the WWN-based PID assignment feature is enabled and a new blade is plugged into the chassis, the ports for which the area is not available are disabled.

NPIV

If any N_Port ID Virtualization (NPIV) devices have static PIDs configured and the acquired area is not the same as the one being requested, the FDISC coming from that device is rejected and the error is noted in the RASlog.

If the NPIV device has Dynamic Persistent PID set, the same AL_PA value in the PID is used. This guarantees that NPIV devices get the same PID across reboots and that AL_PAs assigned for the device do not depend on the order in which the devices come up. For more information on NPIV, see [NPIV](#).

Enabling Automatic PID Assignment

NOTE

To activate WWN-based PID assignment, you do not need to disable the switch.

Use the following procedure to enable automatic PID assignment.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `configure` command.
3. At the `Fabric parameters` prompt, type `y`.
4. At the `WWN Based persistent PID` prompt, type `y`.
5. Press `Enter` to bypass the remaining prompts without changing them.

Example of Activating PID Assignments

```
switch: admin> configure

Configure...
Fabric parameters (yes, y, no, n): [no] y
WWN Based persistent PID (yes, y, no, n): [no] y

System services (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
cfgload attributes (yes, y, no, n): [no]
webtools attributes (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no]
```

```
system attributes (yes, y, no, n): [no]
```

Assigning a Static PID

Use the following procedure to assign a static PID.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `wnnaddress --bind` command to assign a 16-bit PID to a given WWN.

Clearing a PID Binding

Use the following procedure to clear a PID binding.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `wnnaddress --unbind` command to clear the PID binding for the specified WWN.

Showing PID Assignments

Use the following procedure to display PID assignments.

1. Connect to the switch and log in using an account with admin permissions.
2. Based on what you want to display, enter the appropriate command.
 - To display the assigned WWN-PID bindings, enter the following:

```
wnnaddress --show
```

- To display the PID assigned to a specific device WWN, enter the following:

```
wnnaddress --findPID[WWN]
```

Ports

Ports provide either a physical or virtual network connection point for a device. Brocade devices support a wide variety of ports.

Port Types

The following is a list of port types that might be part of a Brocade device:

- **D_Port:** A diagnostic port lets you isolate the link to diagnose link-level faults. This port runs only specific diagnostics tests and does not carry any fabric traffic.
- **E_Port:** An expansion port that is assigned to ISL links to expand a fabric by connecting it to other switches. Two connected E_Ports form an inter-switch link (ISL). When E_Ports are used to connect switches, those switches merge into a single fabric without an isolation demarcation point. ISLs are non-routed links.
- **EX_Port:** A type of E_Port that connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, an EX_Port appears as a normal E_Port. It follows applicable Fibre Channel standards as do other

E_Ports. However, the router terminates EX_Ports rather than allowing different fabrics to merge as would happen on a switch with regular E_Ports. An EX_Port cannot be connected to another EX_Port.

- F_Port: A fabric port that is assigned to fabric-capable devices, such as SAN storage devices.
- G_Port: A generic port that acts as a transition port for non-loop fabric-capable devices.
- L_Port or FL_Port: A loop or fabric loop port that connects loop devices. L_Ports are associated with private loop devices and FL_Ports are associated with public loop devices.
- U_Port: A universal Fibre Channel port. This is the base Fibre Channel port type, and all unidentified or uninitiated ports are listed as U_Ports.
- VE_Port: A virtual E_Port that is a gigabit Ethernet switch port configured for an FCIP tunnel.
- VF_Port: An FCoE enabled ETH port supported on the FC32-64 port blade only.

Director Port Blades

The procedure differs between directors and fixed-port switches mainly because the X6 Director contains interchangeable port blades. For example, fixed-port models identify ports using the port number, while directors identify ports by *slot/port* notation.

NOTE

For detailed information on Brocade X7 and Brocade X6 Directors, refer to the respective hardware reference manuals.

The different blades that can be inserted into a chassis are described as follows:

- Control processor (CP) blades contain communication ports for system management, and they are used for low-level, platform-wide tasks.
- Core blades are used for intra-chassis switching as well as interconnecting two directors.
- Port blades are used for host, storage, and interswitch connections.
- Application (AP) blades are used for Extension (FCIP and IP) support.

NOTE

On each port blade, a particular port must be represented by both slot number and port number.

The Brocade X7-8 has 12 slots that contain control processor, core, port, and AP blades:

- Slot numbers 1 and 2 contain CPs.
- Slot numbers 7 and 8 contain core blades.
- Slot numbers 3 through 6 and 9 through 12 contain port and AP blades.

The Brocade X7-4 has 8 slots that contain control processor, core, port, and AP blades:

- Slot numbers 1 and 2 contain CPs.
- Slot numbers 5 and 6 contain core blades.
- Slot numbers 3 and 4, and 7 and 8 contain port and AP blades.

The Brocade X6-8 has 12 slots that contain control processor, core, port, and AP blades:

- Slot numbers 1 and 2 contain CPs.
- Slot numbers 7 and 8 contain core blades.
- Slot numbers 3 through 6 and 9 through 12 contain port and AP blades.

The Brocade X6-4 has 8 slots that contain control processor, core, port, and AP blades:

- Slot numbers 1 and 2 contain CPs.
- Slot numbers 5 and 6 contain core blades.
- Slot numbers 3 and 4, and 7 and 8 contain port and AP blades.

When you have port blades with different port counts in the same directors, (for example, 32-port blades and 48-port blades, or 48-port blades and 64-port blades), the area IDs no longer match the port numbers. [Blade terminology and compatibility](#) lists the port numbering schemes for the blades.

Setting Port Names

Perform the following steps to specify a port name. For directors, specify the slot number where the blade is installed.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `portname` command.

```
switch:admin> portName 1/0 -n trunk1
```

Port Identification by Slot and Port Number

The port number is assigned to an external port to give it a unique identifier in a switch.

To select a specific port in a director, you must identify both the slot number and the port number using the format `slot number/port number`. No spaces are allowed between the slot number, the slash (/), and the port number.

The following example enables port 4 on a blade in slot 2.

```
switch:admin> portEnable 2/4
```

Port Identification by Index

With the introduction of 48-port and 64-port blades, indexing was introduced. The port indexing method is the preferred method to identify the ports.

A number of fabric-wide databases supported by Fabric OS (including ZoneDB, and the ACL DCC) allow a port to be designated by the use of a "D,P" (*domain,port*) notation. While the "P" component appears to be the port number, for up to 255 ports it is actually the *area* assigned to that port.

Dynamic Port Name

When the Dynamic Port Name feature is enabled, the port name can be dynamically populated with a few default fields.

This feature simplifies the switch configuration and quickly associates ports with the switch or device. This feature dynamically populates the port name with various fields, such as switch name, port type, port index, and alias name.

The Dynamic Port Name feature is disabled by default and can be enabled using the `configure` command.

The FDMI Host Name feature and the Dynamic Port Name feature are mutually exclusive. If you try to enable the Dynamic Port Name feature along with the FDMI Host Name feature, the `configure` command fails and exits with an error message. You can configure the Dynamic Port Name feature when the switch is online.

The Dynamic Port Name feature is not supported on a switch if AG mode is enabled. The Dynamic Port Name feature is not supported in FICON Management Server (FMS) mode.

The Monitoring and Alerting Policy Suite (MAPS) rules and alerts use the dynamically populated port name to report the thresholds.

The port name field for E_Ports and F_Ports displays the following information:

- E_Port: Switch name, port type, and port index as part of the port name
- F_Port: Switch name, port type, port index, and alias name as part of the port name

When displaying the port name of an F_Port assigned with multiple aliases, the following order of preferences picks the alias name:

1. Alias with one member in the WWN format
2. Alias with one member in the D,I format
3. Alias with multiple members in the WWN format
4. Alias with multiple members in the D,I format

Example to Configure the Dynamic Portname Feature

The `configure` command configures the Dynamic Port Name. The option to configure the Dynamic Port Name appears under Fabric parameters.

```
switch:admin> configure
```

```
Configure...
```

```
Fabric parameters (yes, y, no, n): [no] yes
  WWN Based persistent PID (yes, y, no, n): [no]
  Location ID: (0..4) [0]
  Dynamic Portname (on, off): [off] on
  Edge Hold Time (Low(80ms), Medium(220ms), High(500ms), UserDefined(80-500ms): (80..500) [220]
  Remote Fosexec feature: (on, off): [on]
  High Integrity Fabric Mode (yes, y, no, n): [no]
D-Port Parameters (yes, y, no, n): [no]
RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24) [0]
System services (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
cfgload attributes (yes, y, no, n): [no]
webtools attributes (yes, y, no, n): [no]
```

This example shows that the FDMI Host Name and Dynamic Port Name features are mutually exclusive. If you try to enable the Dynamic Port Name feature along with the FDMI Host Name feature, the `configure` command fails and exits with an error message.

```
switch:admin> configure
```

```
Configure...
```

```
Fabric parameters (yes, y, no, n): [no] yes

  Domain: (1..239) [1]
  WWN Based persistent PID (yes, y, no, n): [no]
  F-Port Device Update Mode: (on, off): [on]
  Allow XISL Use (yes, y, no, n): [yes]
  R_A_TOV: (4000..120000) [10000]
  E_D_TOV: (1000..5000) [2000]
  WAN_TOV: (0..30000) [0]      MAX_HOPS: (7..19) [7]
  Data field size: (256..2112) [2112]
  Sequence Level Switching: (0..1) [0]
  Disable Device Probing: (0..1) [0]
  Suppress Class F Traffic: (0..1) [0]
  Per-frame Route Priority: (0..1) [0]
  Long Distance Fabric: (0..1) [0]
  BB credit: (1..27) [16]
```

```

Insistent Domain ID Mode (yes, y, no, n): [no]
Disable Default Portname (yes, y, no, n): [no]
Display FDMI Host Name (yes, y, no, n): [yes]
Dynamic Portname (on, off): [off] on

```

Error: "Dynamic Portname" and "Display FDMI Host Name" are mutually exclusive.
Please set "Display FDMI Host Name" to "no" to enable "Dynamic Portname".

Example to Show the portname Output

The `portname` command displays the switch name, port type, port index, and alias name in the output. The value *none* indicates an error in the alias name and *null* indicates that the alias name is not present.

```

switch:admin> portname
port 0: EDGE1_sw76.E_PORT.0
port 1: EDGE1_sw76..1
port 2: EDGE1_sw76..2
port 3: EDGE1_sw76..3
port 4: EDGE1_sw76..4
port 5: EDGE1_sw76..5
port 6: EDGE1_sw76.F_PORT.6.emlx
port 7: EDGE1_sw76.F_PORT.7.
port 8: EDGE1_sw76.F_PORT.8.

```

Example to Show the portshow Output

The `portshow` command displays the switch name, port type, port index, and alias name in the output.

```

switch:admin> portshow 28
portIndex: 28
Portname: sw78.E_PORT.28
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x1000090b      PRESENT ACTIVE E_PORT T_PORT T_MASTER G_PORT U_PORT LOGICAL_ONLINE LOGIN
LocalSwcFlags: 0x0
portType: 18.0
POD Port: Port is licensed
portState: 1      Online
Protocol: FC
portPhys: 6      In_Sync      portScn: 16      E_Port      Trunk master port Flow control mode 4
port generation number: 0
state transition count: 1
[output truncated]

```

Example to Show the switchshow Output

The `switchshow -portname` command displays the switch name, port type, port index, and alias name in the output.

```

switch:admin> switchshow -portname

switchName:      sw0
switchType:      181.0
switchState:      Online
switchMode:      Native
switchRole:      Principal

```

```

switchDomain: 151
switchId: fffc01
switchWwn: 10:00:00:05:1e:82:3c:2a
zoning: ON (fos6k)
switchBeacon: OFF
FC Router: OFF
Fabric Name: Scale
Allow XISL Use: OFF
LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]

```

Index	Slot	Port	PortWWN	Name
0	7	0	20:00:00:05:1e:82:3c:2a	slot7 port0
1	7	1	20:01:00:05:1e:82:3c:2a	slot7 port1
2	7	2	20:02:00:05:1e:82:3c:2a	slot7 port2
3	7	3	20:03:00:05:1e:82:3c:2a	slot7 port3
4	7	4	20:04:00:05:1e:82:3c:2a	slot7 port4
5	7	5	20:05:00:05:1e:82:3c:2a	slot7 port5
6	7	6	20:06:00:05:1e:82:3c:2a	slot7 port6
7	7	7	20:07:00:05:1e:82:3c:2a	slot7 port7
8	7	8	20:08:00:05:1e:82:3c:2a	slot7 port8
9	7	9	20:09:00:05:1e:82:3c:2a	slot7 port9
10	7	10	20:0a:00:05:1e:82:3c:2a	slot7 port10
11	7	11	20:0b:00:05:1e:82:3c:2a	slot7 port11
12	7	12	20:0c:00:05:1e:82:3c:2a	slot7 port12
13	7	13	20:0d:00:05:1e:82:3c:2a	slot7 port13
14	7	14	20:0e:00:05:1e:82:3c:2a	slot7 port14

[output truncated]

Dynamic Port Name Format

You can configure the port name format for the dynamically populated port names. You can add or delete any fields from the supported list. At least one field should be given as input to set the dynamic port name format. Use the `portname -d` command with the option to configure the dynamic port name format. The dynamic port name format is a switch-wide configuration, and this format is used for all the ports in the switch. The dynamic port name format is stored in the configuration.

The supported port name fields are switch name, port type, port index, slot number/port number, F_Port alias, FDMI hostname, and remote switch name. A control key represents the supported port name fields. The period (.), hyphen (-), and underscore (_) are the supported field separators. Only these control keys and field separators are allowed as inputs to set the dynamic port name format.

NOTE

You can define the dynamic port name format even if the Dynamic Port Name feature is not enabled. However, the customized dynamic port name format is displayed only if the dynamic port name feature is enabled.

NOTE

You cannot downgrade to an earlier version than Fabric OS 8.0.1 when the dynamic port name format is set to a non-default dynamic port name format using the dynamic port name format feature.

The following table lists the supported port name fields and their corresponding control keys.

Table 11: Control Keys Associated with Supported Port Name Fields

Control Key (Case-sensitive)	Port Name Field
S	Switch name
T	Port type
I	Port index
C	Slot number/port number
A	F-Port alias
F	FDMI hostname
R	Remote switch name

Consider the following conditions while configuring the dynamic port name format.

- The default format string is "S.T.I.A" and the default field separator is a period (.). "S.T.I.A" represents <switchname>.<Port type>.<Port index>.<F_Port alias name>
- The maximum number of characters can be 128. Exceeding characters are truncated.
- Only the control keys and supported separators are allowed in the format string, and the control keys are case sensitive.
- The dynamic port name format is not supported in Access Gateway (AG) and FICON Management Server (FMS) modes.
- The format string should not start or end with field separators. For example, ".S.T.I.A."
- All field separators must be of same type. Using different types of field separators is not allowed in a format string. For example, "S.T_I-A".
- Consecutive field separators are not allowed. For example, "S..T..I.A".
- Consecutive control keys are not allowed. For example, "ST.A".
- Repeated control keys are not allowed in the format string. For example, "S.T.S.I.T".
- The "C" and "I" control keys are mutually exclusive.
- At least one control key should be available to set the dynamic port name format.
- The fields without values will be left empty. For example, "my_switch..246" for format string "S.T.I.A", while the port is offline.

The following is an example:

```
switch:admin> portName -d "S.T.I.R.A"
switch:admin> portName -d
"S.T.I.R.A"
```

```
switch:admin> portName
port 0: sw0..0..
port 1: sw0..1..
port 2: sw0..2..
port 3: sw0.E_PORT.3.pluto_218.
port 4: sw0..4..
port 5: sw0..5..
// (output truncated) //
port 14: sw0..14..
port 15: sw0..15..
port 16: sw0.E_PORT.16.sw-46.
port 17: sw0.F_PORT.17..
```

```
port 18: sw0..18..
port 19: sw0.F_PORT.19..
port 20: sw0..20..
// (output truncated) //
```

Enabling a Port

By default, all licensed ports are enabled. You can disable and re-enable them as needed.

Ports that you activate with the Ports-on-Demand license must be enabled explicitly, as described in the *Brocade Fabric OS Software Licensing User Guide*.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the appropriate command based on the current state of the port and whether it is necessary to specify a slot number:
 - To enable a port that is disabled, enter the following command:

```
portEnable [ slot / ] port
```

NOTE

To enable an FCoE VF_Port, use the `portenable -i` command.

- To enable a port that is persistently disabled, enter the following command:

```
portCfgPersistentEnable [ slot / ] port
```

In FMS mode, you cannot use the `portcfgpersistentenable` command, so you must use the `portenable` command instead.

If you change port configurations during a switch failover, the ports might become disabled. To bring the ports online, re-issue the `portEnable` command after the failover is complete.

NOTE

You can also use the `portcfgpersistent` command while enabling or disabling a port.

NOTE

Do not use the `portcfgdefault` command on the Blade Server SAN I/O Modules.

Disabling a Port



CAUTION

If you disable the last E_Port or ISL connecting the switch to the fabric, the fabric reconfigures, the switch segments from the fabric, and all traffic flowing between the switch and the fabric is lost.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the appropriate command based on the current state of the port and whether it is necessary to specify a slot number:
 - To disable a port that is enabled, enter the `portdisable` command.
 - To disable an FCoE enabled VF_Port, enter the `portdisable -i` command.
 - To disable a port that is persistently enabled, enter the `portcfgpersistentdisable` command.

In FMS mode, you cannot use the `portcfgpersistentdisable` command, so you must use the `portdisable` command instead.

- To set a persistently disabled port to normal disabled without enabling the port first, enter the `portcfgpersistent --set -persistentenable` command.

In this case, the port becomes normal disabled and not persistently disabled.

In FMS mode, you cannot use the `portcfgpersistent` command.

The following example disables port 3 and persistently disables port 4, both in slot 2.

```
switch:admin> portdisable 2/3
switch:admin> portcfgpersistentdisable 2/4
```

The following example changes port 4 to be normal disabled instead of persistently disabled.

```
switch:admin> portcfgpersistence --set -persistentenable 2/4
```

Port Decommissioning

Port decommissioning provides an automated mechanism to remove an E_Port or E_Port trunk port from use.

The port decommissioning feature identifies the target port and communicates the intention to decommission the port to those systems within the fabric affected by the action. Each affected system can agree or disagree with the action, and these responses are automatically collected before a port is decommissioned.

NOTE

All members of a trunk group must have an equal link cost value for any of the members to be decommissioned. If any member of a trunk group does not have an equal cost, requests to decommission a trunk member will fail and an error reminding the caller of this requirement is produced.

The following restrictions apply to port decommissioning:

- Port decommissioning is not supported on ports with DWDM, CWDM, or TDM.
- Port decommissioning requires that the lossless feature is enabled on both the local switch and the remote switch.

Use the `portdecom` command to begin the decommission process.

Port Decommissioning without Disabling

When an error threshold is reached on a port, instead of fencing the port by decommissioning and disabling the link, the port is decommissioned without disabling and marked as *Impaired*. An impaired port is removed from all the routes from both sides of the link and prevented from being used in further route decisions, but allowed to remain online and be used only if it is the last link available to a neighboring domain. The port stops carrying any data traffic unless no other shortest path routes are available. An impaired port is automatically removed from trunks. Port impairs is supported on both trunked and non-trunked E_Ports. You can clear the impaired state of the port and using the `portimpair --clear` command.

NOTE

A port can be impaired only if the `mapsconfig --decomcfg impair mode` is set.

1. To clear the impaired state of a port, run the following command.

```
switch:admin> portimpair --clear 0
The local port and peer port impaired state has been cleared.
portImpair: To rejoin a trunk group, the port must be toggled.
```

If the port was part of a trunk earlier, the port joins the trunk again.

2. The `switchshow` output displays the impaired state of a port as follows:

```
3  3  210100  id  N16  Online  FC  E-Port  10:00:c4:f5:7c:99:8c:97  (Impaired)
```

3. The `portshow` and `islshow` commands mark the impaired port as `IMPAIRED`.

For example:

```
portFlags: 0x10004103      PRESENT ACTIVE E_PORT G_PORT U_PORT LOGIN LED IMPAIRED
```

For example:

```
1: 3-> 3 10:00:c4:f5:7c:99:8c:97 3 w86_f1 sp: 16.000G bw: 16.000G QOS IMPAIRED
```

Setting Port Speeds

Use the following procedure to set port speeds.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `portcfgspeed` command.

The following example sets the speed for port 3 on slot 2 to 16G:

```
switch:admin> portcfgSpeed 2/3 16
```

The following example sets the speed for port 3 on slot 2 to auto-negotiate:

```
switch:admin> portcfgSpeed 2/3 0
```

NOTE

The supported port speed depends on the SFP used. The `sfpshow` command displays the SFP information. The `sfpshow <PortNumber> -link` command displays the peer port SFP information for a link between an N_Port on the device and an FX_Port or F_Port on the switch, but it does not display the peer port SFP information in case of ISL and ICL links.

Setting All Ports on a Switch to the Same Speed

Use the following procedure to set all ports on a switch to the same speed.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `switchcfgspeed` command.

If the specified speed is not supported on a particular port or blade, that port or blade is skipped, and the speed is not changed.

The following example sets the speed for all ports on the switch to 32G:

```
switch:admin> switchcfgspeed 32

Committing configuration...done.
```

The following example sets the speed for all ports on the switch to auto-negotiate:

```
switch:admin> switchcfgspeed 0

Committing configuration...done.
```

Setting Port Speed for a Port Octet

The `portcfgoctetspeedcombo` command sets the octet mode for all ports in an octet. This command is not supported on the Brocade 7810, Brocade G610, and Brocade Gen 7 platforms.

Fabric OS supports the following port octet speed combinations. All ports in the same octet can support only the port speeds listed in the applicable column of the following table for that combination.

Table 12: Port Octet Combinations and Modes

Port Octet in Combination	Gen 6 Octet Modes
1 (default)	32G 16G 8G 4G
2	10G 8G 4G
3	Not supported

NOTE

A port octet can be set to a single octet mode. The ports in the octet can run on speeds supported by its octet combination. This applies to both auto-speed-negotiation and fixed speeds.

To change the octet speed combination you have to ensure the following:

- All the fixed-speed ports must be configured at a speed that is supported for the port octet combinations.
- All the online ports in auto-negotiation mode must have a negotiated speed that is supported for the port octet combinations. If the port is running at a speed not supported in the specified combination, you must disable the port or change the speed to a supported fixed speed to change the combo.
- All the ports that are in Media Invalid state should be disabled before the octet combo is changed.

If the conditions are not met, the combination change operations will fail with an appropriate error message. When you receive the error message, you can configure the ports to a supported speed in the specific combo or disable the online ports.

Example 1: If you have the octet of ports 0 to 7 set to combination 1, port 0 can be set to ASN up to 32G, port 1 to ASN up to 16G, port 2 to fixed 8G, port 3 to fixed 4G, as so on, as long as only the port speeds of 4G, 8G, 16G, or 32G are configured.

Example 2: If you have the octet of ports 8 to 15 set to combination 2, port 8 can be set to fixed 10G, port 9 to ASN up to 8G, port 10 to ASN up to 4G, and so on, as long as only the port speeds of 4G, 8G, or 10G are configured.

The octet can be specified by any port within the octet. To change the first octet, for example, any port from 0 through 7 can be used as the port argument value.

This command is nondisruptive. If any of the ports do not meet the port octet conditions, the operation fails, and you must change the port speed or disable the ports and retry the operation.

NOTE

Be aware that in a Virtual Fabrics environment, this command configures the speed of a port octet chassis-wide and not only on the logical switch.

This feature is supported on Gen 6 (32G) switches and ports, except for the core blades.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `portcfgoctetspeedcombo` command, specifying the octet and mode.

This example configures the ports in the first octet for combination 3 (supporting autonegotiated or fixed-port speeds of 16G and 10G):

```
switch:admin> portcfgoctetspeedcombo 1 3
```


If the combinations configurations cannot be switched without disrupting the ports, the following message is displayed indicating which ports must be disabled and which speed configurations must be changed. Octet combination information is stored in the configuration file. All the configuration operations are supported with octet speed combinations.

```
switch:admin> portcfgoctetspeedcombo 8/1 2
The following ports must be disabled or speed configuration needs to be change.
Please retry the operation after taking appropriate action.
```

Speeds supported for octet combo 2 - [10G,8G,4G,AN]

```
Port Index      Speed
=====
192             AN
193             N16+
194             N16+
195             N16+
196             N16+
197             N32+
198             AN
199             AN
```

```
* Port speed configuration or SFP must be changed
+ Port must be disabled
Setting octet speed combo failed
```

3. Enter `portcfgspeed`.

This command checks if the desired speed is allowed based on the combination configured for that octet. If the speed is invalid, a warning message is displayed with the suggested combination to achieve the desired speed and the command is aborted. You must run `portcfgoctetspeedcombo` to set the suggested combination before re-running `portcfgspeed` with your desired speed setting. For example, if the octet of port 3 is in combination 1, and you specified 10, the command is rejected with the suggested combination of 2. You must change the octet speed combination to 2 before retrying 10.

The supported speed for each combination is shown in the following table for Gen 6 devices.

Table 13: Supported Speed Arguments and Command Strings for Octet Combination

Octet Combination	Supported <speed> Argument Setting	portCfgShow Combination String	portCfgShow Speed String
1 (32G 16G 8G 4G)	0	1	AN
	4		4G
	8		8G
	16		16G
	32		32G
2 (10G 8G 4G)	0	2	AN
	4		4G
	8		8G

Octet Combination	Supported <speed> Argument Setting	portCfgShow Combination String	portCfgShow Speed String
	10		10

The final speed achieved also depends on the SFP used.

Example 1: If the octet of ports 0 to 7 is set to combination 1, port 0 is set to 0, and 8G SFP is inserted, it can negotiate only up to 8G; if 10G SFP is inserted, port will not go online, because combination 1 does not support 10G.

Example 2: If the octet of ports 8 to 15 is set to combination 2, port 8 is set to 10, but 16G SFP is inserted, the port will not go online as the configured speed (10G), because that port is not supported by the SFP (16G|8G|4G).

Setting the Maximum Autonegotiated Port Speed

If you do not know exactly at what speed a new device will connect, but you want to ensure that nothing connects faster than a certain speed, then you can configure the maximum autonegotiated speed.

For example, if a port is configured with a maximum autonegotiated speed of 16G, and the SFP is 32G, then the attempted speed negotiations are 16G and 8G.

NOTE

A 32G optical transceiver in a Gen 6 device can only autonegotiate to 32G, 16G, or 8G, and a 16G optical transceiver in a Gen 6 device can only autonegotiate to 16G, 8G, or 4G; this capability is determined by the speed of the optical transceiver at the other end of the link.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `portcfgspeed -m` command.

The following example sets the maximum autonegotiated speed to 16G for port 3 on slot 2.

```
switch:admin>portCfgSpeed 2/3 0 -m 16
```

Decoding fdmishow Command Output

In Fabric OS, the `fdmishow` command automatically decodes certain hexadecimal port attribute data to ASCII format. However, if you still need to view the data in hexadecimal format, you can use the `fdmishow -hexoutput` command.

You use this command to display Fabric-Device Management Interface (FDMI) information for all HBAs and ports.

Detailed FDMI information is displayed for local HBAs and ports. The devices from certain vendors may register the following extended Vendor-Specific attributes and the output may vary depending on the device-registered values.

- **0xF047** – End-to-End Version
- **0xF100** – Service Category
- **0xF101** – GUID
- **0xF102** – Version
- **0xF103** – Product Name
- **0xF104** – Port Info
- **0xF105** – QOS Support
- **0xF106** – Security Support
- **0xF107** – Connected Ports
- **0xF108** – Storage Array Family
- **0xF109** – Storage Array Name
- **0xF10A** – Storage Array System Model
- **0xF10B** – Storage Array OS
- **0xF10C** – Storage Array Number of Nodes
- **0xF10D** – Storage Array Nodes

The following example highlights the difference in the `fdmishow` command output:

```
swl:admin> fdmishow
Local HBA database contains (port 16):
10:00:00:10:9b:7d:b4:df
Ports: 1
10:00:00:10:9b:7d:b4:df
Port attributes:
FC4 Types: 0x01 FCP FC-CT FC-NVMe
Supported Speed: 8 16 32 Gb/s
Port Speed: 32 Gb/s
Max Frame Size: 2048 bytes
Host Name: SJS0946W050919
Node Name: 20:00:00:10:9b:7d:b4:df
Port Name: 10:00:00:10:9b:7d:b4:df
Port Type: N_Port (0x1)
Port Symb Name: Emulex PPN-10:00:00:10:9B:7D:B4:DF
Class of Service: 3
Fabric Name: 20:10:d8:1f:cc:05:3f:ba
FC4 Active Type: 0x01 FCP FC-CT
Port State: 0x2
Discovered Ports: 0x7
Port Identifier: 0xd11000
HBA attributes:
Node Name: 20:00:00:10:9b:7d:b4:df
Manufacturer: Emulex Corporation
Serial Number: FP91779946
Model: LPe35002-M2
Model Description: Emulex LPe35002-M2 2-Port 32Gb Fibre Channel Adapter
Hardware Version: 0000000D
Driver Version: 12.4.243.4
Option ROM Version: 12.2.212.14
Firmware Version: 12.2.212.14
OS Name and Version: Windows 2019
Max CT Payload Length: 524288 words
Symbolic Name: Emulex LPe35002-M2 FV12.2.212.14 DV12.4.243.4 HN:SJS0946W050919 OS:Windows 2019
Number of Ports: 1
Fabric Name: 10:00:d8:1f:cc:05:3f:ba
Bios Version: 12.2.212.14
Vendor Identifier: Emulex

Local Port database contains (port 16):
10:00:00:10:9b:7d:b4:df

Remote HBA database contains:
.
.
20:13:8c:7c:ff:07:1b:00
20:04:00:11:0d:68:9e:02
10:00:04:90:fa:99:89:3d
```

20:10:00:05:1e:74:7d:71

Blade Terminology and Compatibility

Before configuring a chassis, familiarize yourself with the platform CP blade and port blade nomenclature, as well as the port blade compatibilities.

Table 14: Core and CP Blade Terminology and Platform Support

Blade	Blade ID (slotshow)	Supported on X7 Family	Supported on X6 Family	Supported on Upgraded X6 Family	Definition
CPX7	220	Yes	No	No	A Brocade Gen 7 director family control processor blade. The control processor hosts the FOS that manages all hardware within the Gen 7 director.
CPX ^a	220	Yes	X6 directors with switch type 165.5 or 166.5	No	A Brocade director family control processor blade used on X7 directors and on X6 directors with switch ID 165.5 or 166.5. The control processor hosts the FOS that manages all hardware within the director.
CR64-8	215	Yes X7-8 only	No	Yes X6-8 only	A core blade that has 16×4 QSFPs per blade. It can be connected to another CR64-8 or a CR64-4 core blade.
CR64-4	214	Yes X7-4 only	No	Yes X6-4 only	A core blade that has 8×4 QSFPs per blade. It can be connected to another CR64-4 or a CR64-8 core blade.
CPX6+	175	No	No	Yes, Upgraded with a Gen 7 FPGA	A Brocade Gen 6 director control processor blade originally integrated into a Gen 6 director chassis as a CPX6 that has been field upgraded to support Gen 7 functionality.
CPX6	175	No	Yes	No	A Brocade Gen 6 director family control processor blade. The control processor hosts the FOS that manages all hardware within the Gen 6 director.
CR32-8	177	No	Yes X6-8 only	No	A core blade that has 16×4 QSFPs per blade. It can be connected to another CR32-8 or a CR32-4 core blade.
CR32-4	176	No	Yes X6-4 only	No	A core blade that has 8×4 QSFPs per blade. It can be connected to another CR32-4 or a CR32-8 core blade.

a. In the Fabric OS 9.1.0b release, the CPX blade is displayed as CPX7 for X6 directors with switch type 165.5 or 166.5.

Table 15: Port Blade Terminology, Numbering, and Platform Support

Blade	Blade ID (slotshow)	Supported on X7 Family	Supported on X6 Family	Definition
FC64-48	216	Yes	No	A Brocade Gen 7 director FC port blade. A 48-port front-end 64G FC SFP+ with edge switch interconnect to the core blades. It supports 8/10/16/32/64G port speeds.
FC32-X7-48	218	Yes	No	A Brocade Gen 7 director FC port blade. A 48-port front-end 32G FC SFP+ with edge switch interconnect to the core blades. It supports 4/8/10/16/32G port speeds. The 4G port speed supports F_Port and E_Port connectivity.
FC32-48	178	Yes	Yes	A 48-port, 32G port blade that supports 4/8/10/16/32G port speeds.
FC32-64	204	Yes	Yes	A 16 QSFP-port, 4×32G port blade that supports 4×10G, 4×16G, and 4×32G port speeds. The ports can also work in ETH (FCoE) mode to support 16×40G, 64×25G, and 65×10G port speeds.
SX6	186	Yes	Yes	Extension blade with 32G Fibre Channel, FCIP, and 10GbE technology.

CP Blades

The control processor (CP) blade provides redundancy and acts as the main controller on the Brocade director.

The CP blades in the Brocade X7 and X6 Directors are hot-swappable within their director family. The Brocade X7-8 and X7-4 Directors support the CPX7 blades that are hot-swappable and interchangeable. The Brocade X6-8 and X6-4 Directors support the CPX6 blades that are hot-swappable and interchangeable. The Brocade X6-8 Director with switch type 166.5 and the Brocade X6-4 Director with switch type 165.5 support CPX blades that are hot-swappable and interchangeable.

Each CP (primary and secondary partition) should maintain the same firmware version.

Core Blades

Core blades provide intra-chassis switching and inter-chassis link (ICL) connectivity between X7 and X6 platforms.

- The Brocade X7-8 Director supports two CR64-8 core blades.
- The Brocade X7-4 Director supports two CR64-4 core blades.
- The Brocade X6-8 Director supports two CR32-8 core blades.
- The Brocade X6-4 Director supports two CR32-4 core blades.

The core blades for each platform are not interchangeable or hot-swappable with the core blades for any other platform. If you try to interchange the blades, they may be physically incompatible or become faulty.

Port and Application Blade Compatibility

The maximum number of blades supported is eight on an 8-slot chassis and four on a 4-slot chassis. However, the maximum number of a specific blade type is limited.

The following table lists the maximum supported limits of SX6 blade for a specific platform. Software functions are not supported across application blades.

NOTE

The hardware limit is enforced by software.

Table 16: Port and Application Blade Compatibility within Brocade Director Families

Intelligent blade	X7-8	X7-4	X6-8	X6-4
SX6	4	4	4	4

Disabling Switches

Switches are enabled by default. In some cases, you may need to disable a switch to perform diagnostic testing. This action ensures that diagnostic activity does not interfere with normal fabric traffic.

Use the following procedure to disable a switch.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter `switchcfgpersistentdisable --setdisablestate`.

This procedure sets the switch to the disabled state without disabling it. On reset, the switch is in disabled state, and must be enabled.

Configuring Fibre Channel Link Aggregation

Fabric OS switches allow Fibre Channel end devices to log in to the switches through different mechanisms such as using direct connections, using Access Gateway (AG) as an intermediate, and also using third-party devices such as Cisco Unified Computing System (UCS) servers and Fabric Interconnect (FI) devices. Starting with Fabric OS 9.1.0, the Fibre Channel Link Aggregation (FC-LAG) feature provides FC-LAG capability between Brocade switches and Cisco UCS-FIs that are configured in end-host mode. This enhancement on Brocade switches supports three different connectivity options for UCS-FIs.

The FC-LAG feature provides the advantage of automatic load distribution based on the SID, DID, and OXID across all physical links in the FC-LAG, and it provides the ability to add and remove physical links to the FC-LAG without application disruption. When the FC-LAG feature is enabled, a link aggregation protocol runs between the remote FI device and the Brocade switch and configures all the underlying physical links into a single aggregated link. The bandwidth is adjusted based on the available physical links, which improves load balancing and reduces congestion events.

When a FOS switch port is configured to be a member of a FC-LAG, the link aggregation protocol is used to communicate with its peer to establish and manage the FC-LAG.

FC-LAG is supported on the following platforms:

- X7-8 Director
- X7-4 Director
- G730 Switch
- G720 Switch
- X6-8 Director
- X6-4 Director
- G630 Switch

NOTE

The FC-LAG feature is supported across multiple blades and a FC-LAG is formed using multiple ports across different blades.

Creating the FC-LAG Configuration

You can create the FC-LAG configuration using the `fclag --create <fclag_name>` command.

```
fclag --create port123
```

Deleting the FC-LAG Configuration

You can delete the FC-LAG configuration using the `fclag --delete <fclag_name>` command.

```
fclag --delete port123
```

Enabling the FC-LAG Configuration

You can enable the FC-LAG using the `fclag --enable <fclag_name>` command. When an FC-LAG is created, it is enabled by default.

```
fclag --enable port123
```

Disabling the FC-LAG Configuration

You can disable the FC-LAG using the `fclag --disable <fclag_name>` command. When an FC-LAG is created, it is enabled by default.

```
fclag --disable port123
```

Adding a Member Port to the FC-LAG Configuration

When an FC-LAG is created, you can add the physical ports to the FC-LAG using the `fclag --add_mbr <fclag_name> -port <slot/port|port-range>` command.

When a port is added to the FC-LAG, the port is disabled. You must use the `portenable` command to enable the port.

```
fclag --add_mbr port123 -port 8/6
```

```
fclag --add_mbr port123 -port 8/7-10
```

Removing a Member Port from the FC-LAG Configuration

Ports can be removed from the FC-LAG using the `fclag --remove_mbr <fclag_name> -port <slot/port|port-range>` command.

When a port is removed from the FC-LAG, traffic is adjusted evenly to the remaining ports in the FC-LAG.

```
fclag --remove_mbr port123 -port 8/6
```

```
fclag --remove_mbr port123 -port 8/7-10
```

Displaying the FC-LAG Configuration

You can verify the existing FC-LAG details using the `fclag --show [<fclag_name>]` command.

```
fclag --show port123
```

Power Management

To manage power and ensure that more critical components are the least affected by power changes, you can specify the order in which the components are powered off by using the `powerofflistset` command.

The power monitor compares the available power with the power that is required to determine if there will be enough power to operate. If it is predicted that there is less power available than required, the power-off list is processed until there is enough power for operation. By default, the processing begins with slot 3 and proceeds to the last slot in the chassis. As power becomes available, slots are powered up in the reverse order. When a chassis is initially powered-up, when the `slotpoweron` command is used, or when a blade is inserted, the available power is compared to the required power before power is applied to the blade.

NOTE

Some FRUs in the chassis may use significant power, yet they cannot be powered off through software.

The `powerofflistshow` command displays the power-off order.

NOTE

Before issuing the `slotpoweron` or `slotpoweroff` command, verify whether there are any member ports that are configured with the `Short` timeout using the `portchannel --show -detail` command. You must change the timeout to `Long` on all the member ports. You can revert the timeout setting to `Short` once the blade is up and running after issuing the `slotpoweron` command.

Powering Off a Port Blade or Core Blade

Blades cannot be powered off when POST or AP initialization is in progress.

The CP blades cannot be powered off from the CLI. You must manually power off the CP blades by lowering the slider or removing power from the chassis. If the CP is not up and running, then physical removal or powering off the chassis is required.

ATTENTION

Powering off the last operational core blade disables the chassis.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `slotpoweroff` command with the slot number of the port blade or core blade you want to power off.

```
ecp:admin> slotpoweroff 3
```

```
Slot 3 is being powered off
```

```
ecp:admin> slotpoweroff 5
```

```
You have chosen a core blade to power off,  
if this is the last operational core blade, then  
the chassis will disable itself.
```

```
Would you like to proceed? (yes, y, no, n):[no] y
```

Powering On a Port Blade or Core Blade

All blades are powered on by default when the switch chassis is powered on.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `slotPowerOn` command with the slot number of the port blade or core blade you want to power on.

```
ecp:admin> slotPowerOn 3
```

```
Powering on slot 3
```


Persistently Powering Off a Port Blade or Core Blade

The `slotpoweroff` command powers down a blade and keeps the blade powered off even after an HA failover. However, when the chassis is rebooted, the blade is powered on. Then, you can persistently power down a blade so that the slot remains powered down even after reboot, using the `slotcfgpersistence --poweroff` command. The persistent power off works even in an empty slot.

When a blade is inserted into a persistently powered-off slot, the slot powers on. However, when the chassis is rebooted, the slot powers off again.

1. Use the following command to persistently power off a slot:

```
switch:admin> slotcfgpersistence --poweroff 4
```

2. Use the following command to display the persistent power flag for all the slots:

```
switch:admin> slotcfgpersistence
```

3. Use the following command to display the Persistent flag for the slot.

```
switch:admin> slotshow -m
```

Slot	Blade Type	ID	Model Name	Status
1	CP BLADE	175	CPX6	ENABLED
2	CP BLADE	175	CPX6	ENABLED
3	SW BLADE	178	FC32-48	INSERTED, NOT POWERED ON(Persistent)
4	SW BLADE	216	FC64-48	ENABLED
5	SW BLADE	178	FC32-48	ENABLED
6	SW BLADE	178	FC32-48	ENABLED
7	CORE BLADE	215	CR64-8	ENABLED
8	CORE BLADE	215	CR64-8	ENABLED
9	SW BLADE	204	FC32-64	ENABLED
10	SW BLADE	204	FC32-64	ENABLED
11	SW BLADE	216	FC64-48	ENABLED
12	SW BLADE	216	FC64-48	ENABLED

4. Use the following command to persistently power on a slot:

```
switch:admin> slotcfgpersistence --poweron 4
```

NOTE

After you issue the `configdownload` command, the system must be rebooted for `slotcfgpersistence` settings to take effect. You can use the `slotpoweron` command to temporarily power on a blade that is persistently disabled. After the system is powered up or rebooted, the blade will be powered off.

NOTE

If a blade is removed and re-inserted or a new blade is re-inserted in the same slot, the blade will be online. The `slotcfgpersistence` flag setting will remain unchanged and will be applied after reboot.

Equipment Status

You can check the status of switch operation, High Availability features, and fabric connectivity.

Checking Switch Operation

Use the following procedure to check switch operation.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `switchShow` command. This command displays a switch summary and a port summary.
3. Check that the switch and ports are online.

Verifying High-Availability Features in Directors

High-Availability (HA) features provide maximum reliability and nondisruptive management of key hardware and software modules.

Use the following procedure to verify HA features for a Brocade X7 Director or Brocade X6 Director.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `chassisshow` command to verify the model of the X7 or X6 and obtain a listing of all the FRUs.
3. Enter the `chassisshow --sort` command to display all the FRUs in a sorted order.
4. Enter the `chassisshow --help` command to display the command usage.
5. Enter the `hashow` command to verify HA is enabled, the heartbeat is up, and that the HA state is synchronized between the active and standby CP blades.
6. Enter the `haredundancy --show` command to display the CP redundancy settings and switch uptime.

```
X6-8:FID165:admin> haredundancy --show
=== HA Redundancy Statistics ===
    HA State synchronized
    Current Active Session:
    Active Slot = CP0 (Local), Expected Recovered
    Standby Slot = CP1 (Remote)
    Start Time: 17:55:33 UTC Fri Jan 03 2014

    Previous Active Session:
    Active Slot = CP1, Expected Recovered
    Standby Slot = CP0
    Start Time: 17:49:46 UTC Fri Jan 03 2014
    End Time: 17:54:10 UTC Fri Jan 03 2014

    System Uptime: 17:42:11 UTC Fri Jan 03 2014
```

```
Core-X7-8_239:FID128:admin> haredundancy --show
=== HA Redundancy Statistics ===
    HA State synchronized
    Current Active Session:
    Active Slot = CP0 (Local), Expected Recovered
    Standby Slot = CP1 (Remote)
    Start Time: 17:22:21 PDT Wed Mar 18 2020

    Previous Active Session:
    Active Slot = CP1, Expected Recovered
    Standby Slot = CP0
```

Start Time: 17:15:58 PDT Wed Mar 18 2020

End Time: 17:20:26 PDT Wed Mar 18 2020

System Uptime: 17:20:00 PDT Mon Mar 16 2020

7. Enter the `fanShow` command to display the status and speed of each fan in the system. Refer to the hardware reference manual of your system to determine the appropriate values.
8. Enter the `psShow` command to display the status of the switch power supplies. Refer to the hardware reference manual of your system to determine the appropriate values.
9. Enter the `slotShow -m` command to display the inventory and the status of each slot in the system.

Example of the Slot Information Displayed for a X7-8 Director Chassis

```
switch:admin> slotshow -m
```

Slot	Blade Type	ID	Model Name	Status
1	CP BLADE	220	CPX7	ENABLED
2	CP BLADE	220	CPX7	ENABLED
3	SW BLADE	216	FC64-48	ENABLED
4	SW BLADE	216	FC64-48	ENABLED
5	SW BLADE	216	FC64-48	ENABLED
6	SW BLADE	216	FC64-48	ENABLED
7	CORE BLADE	215	CR64-8	ENABLED
8	CORE BLADE	215	CR64-8	ENABLED
9	AP BLADE	186	SX6	ENABLED
10	SW BLADE	204	FC32-64	ENABLED
11	SW BLADE	204	FC32-64	ENABLED
12	AP BLADE	186	SX6	ENABLED

Verifying Fabric Connectivity

Use the following procedure to verify fabric connectivity.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `fabricShow` command. This command displays a summary of all the switches in the fabric.

The output of the `fabricShow` command is described in [Domain IDs](#).

Verifying Device Connectivity

Use the following procedure to verify device connectivity.

1. Connect to the switch and log in using an account with admin permissions.
2. Optional: Enter the `switchShow` command to verify that devices, hosts, and storage are connected.
3. Optional: Enter the `nsShow` command to verify that devices, hosts, and storage have successfully registered with the name server.
4. Enter the `nsAllShow` command to display the 24-bit Fibre Channel addresses of all devices in the fabric.

```
switch:admin> nsallshow
```

```
{
  010e00 012fe8 012fef 030500 030b04 030b08 030b17 030b18
```

```

030b1e 030b1f 040000 050000 050200 050700 050800 050de8
050def 051700 061c00 071a00 073c00 090d00 0a0200 0a07ca
0a07cb 0a07cc 0a07cd 0a07ce 0a07d1 0a07d2 0a07d3 0a07d4
0a07d5 0a07d6 0a07d9 0a07da 0a07dc 0a07e0 0a07e1 0a0f01
0a0f02 0a0f0f 0a0f10 0a0f1b 0a0f1d 0b2700 0b2e00 0b2fe8
0b2fef 0f0000 0f0226 0f0233 0f02e4 0f02e8 0f02ef 210e00
211700 211fe8 211fef 2c0000 2c0300 611000 6114e8 6114ef
611600 620800 621026 621036 6210e4 6210e8 6210ef 621400
621500 621700 621a00
75 Nx_Ports in the Fabric }

```

The number of devices listed should reflect the number of devices that are connected.

Audit Log Configuration

When managing SANs, you may want to audit certain classes of events to ensure that you can view and generate an audit log for what is happening on a switch, particularly for security-related event changes. These events include login failures, zone configuration changes, firmware downloads, and other configuration changes; in other words, critical changes that have a serious effect on the operation and security of the switch.

Important information related to event classes is also tracked and made available. For example, you can track changes from an external source by the user name, IP address, or type of management interface used to access the switch.

Auditable events are generated by the switch and streamed to an external host through a configured system message log daemon (syslog). You specify a filter on the output to select the event classes that are sent through the system message log. The filtered events are streamed chronologically and sent to the system message log on an external host in the specified audit message format. This ensures that they can be easily distinguished from other system message log events that occur in the network. Then, at some regular interval of your choosing, you can review the audit events to look for unexpected changes.

Before you configure audit event logging, familiarize yourself with the following audit event log behaviors and limitations:

- By default, *all event classes* are configured for audit; to create an audit event log for specific events, you must explicitly set a filter with the `class` operand and then enable it.
- Audited events are generated specific to a switch and have no negative impact on performance.
- The last 8192 messages are persistently saved in the audit log, but all the audit events are sent to the system message log, assuming there are no bottlenecks that is forwarded to your syslog server.
- The audit log depends on the system message log facility and IP network to send messages from the switch to a remote host. Because the audit event log configuration has no control over these facilities, audit events can be lost if the system message log and IP network facilities fail.
- If too many events are generated by the switch, the system message log becomes a bottleneck and audit events are dropped by the Fabric OS.
- If the user name, IP address, or user interface is not transported, "*None*" is used instead for each of the respective fields.
- For High Availability, the audit event logs exist independently on both active and standby CPs. The configuration changes that occur on the active CP are propagated to the standby CP and take effect.
- Audit log configuration is also updated through a configuration download.

Before configuring an audit log, you must select the event classes you want audited.

NOTE

Only the active CP can generate audit messages because event classes being audited occur only on the active CP. Audit messages cannot originate from other blades in a Director.

Switch names are logged for switch components and Director names for Director components. For example, a Director name may be FWDL or RAS, and a switch component name may be zone, name server, or SNMP.

Refer to the *Brocade Fabric OS Message Reference Manual* for details on event classes and message formats.

NOTE

If an audit message is logged from the CLI, any environment variables will be initialized with proper values for login, interface, IP, and other session information. Refer to the *Brocade Fabric OS Message Reference Manual* for more information.

Verifying the Host Syslog Prior to Configuring the Audit Log

Audit logging assumes that your syslog is operational and running. Before configuring an audit log, you must perform the following steps to ensure that the host syslog is operational.

1. Set up an external host machine with a system message log daemon running to receive the audit events that will be generated.
2. On the switch where the audit configuration is enabled, enter the `syslogAdmin` command to add the IP address of the host machine so that it can receive the audit events.

You can use IPv4, IPv6, or DNS names for the `syslogAdmin` command.

3. Ensure that the network is configured with a network connection between the switch and the remote host.
4. Check the host syslog configuration. If all error levels are not configured, you may not see some of the audit messages.

Configuring an Audit Log for Specific Event Classes, Severity Levels, or Message ID Filters

1. Connect to the switch from which you want to generate an audit log and log in using an account with admin permissions.
2. Enter the `auditcfg --class` command, which defines the specific event classes to be filtered.

NOTE

By default, audit log is enabled for all event classes. Valid values are: 1-ZONE, 2-SECURITY, 3-CONFIGURATION, 4-FIRMWARE, 5-FABRIC, 7-LS (Logical Switch), 8-CLI, and 9-MAPS. The filter is specified by its numeric value. To add more than one filter, the numeric values must be separated by commas and spaces are not permitted. This operand is required.

```
switch:admin> auditcfg --class 2,4
```

```
Audit filter is configured.
```

3. Enter the `auditcfg --severity` command, which defines the severity levels for audit messages.
4. Enter the `auditcfg --enable -all` command, which enables audit event logging based on the classes configured in Step 2.

```
switch:admin> auditcfg --enable -all
```

```
Audit filter is enabled.
```

To set filters for specific message IDs, enter `auditcfg --enable [-log message id]` command.

To disable an audit event configuration, enter the `auditcfg --disable -all` command.

5. Enter the `auditcfg --show -filter` command to view the filter configuration and confirm that the correct event classes are being audited and that the correct filter state appears (enabled or disabled).

```
switch:admin> auditcfg --show -filter
```

```
Audit filter is enabled.
2-SECURITY
4-FIRMWARE
9-MAPS
Severity level: INFO
```

6. Enter the `auditDump -s` command to confirm that the audit messages are being generated.

Example of the Syslog (System Message Log) Output for the Audit Logging

```
Oct 10 08:52:06 10.1.18.7 raslogd: AUDIT, 2008/10/10-08:20:19 (GMT), [SEC-3020], INFO, SECURITY,
admin/admin/10.1.18.13/telnet/CLI, ad_0/ras007/FID 128, , Event: login, Status: success, Info:
Successful login attempt via REMOTE, IP Addr: 10.1.18.13.
Oct 10 08:52:23 10.1.18.7 raslogd: 2008/10/10-08:20:36, [CONF-1001], 13, WWN
10:00:00:05:1e:34:02:0c | FID 128, INFO, ras007, configUpload completed successfully. All config
parameters are uploaded.
Oct 10 09:00:04 10.1.18.7 raslogd: AUDIT, 2008/10/10-08:28:16 (GMT), [SEC-3021], INFO, SECURITY,
admin/NONE/10.1.18.13/None/CLI, None/ras007/FID 128, , Event: login, Status: failed, Info: Failed
login attempt via REMOTE, IP Addr: 10.1.18.13.
```

Configuring Remote Syslog Servers

Fabric OS supports configuring a switch to forward all error log entries to a remote syslog server, to set the syslog facility to a specified log file, to remove a syslog server, and to display the list of configured syslog servers. Brocade switches use the syslog daemon, a process available on most UNIX systems that reads and forwards system messages to the appropriate log files or users, depending on the system configuration. Up to six servers are supported.

By default, the switch uses the UDP protocol to send the error log messages to the syslog server. The default UDP port is 514. A user-defined port number can be configured to the UDP protocol. Use the `-secure` option to configure the switch to send the error log messages securely using the Transport Layer Security (TLS) protocol. TLS is an encryption protocol over the TCP/IP network protocol that can be used only with the TCP-based destinations: `tcp()` and `tcp6()`. While enabling secure syslog mode, you can specify a port that is configured to receive the log messages from the switch. The default TLS port is 6514.

NOTE

If the syslog TLS cipher string is modified by applying a new template, you must remove and add back the syslog server or the crypto string is not written to the syslog conf file.

Refer to the following examples to configure syslog server hosts.

To configure an IPv4 secure syslog server to which error log messages are sent:

```
switch:admin> syslogadmin --set -ip 172.26.26.173 -secure -port 2000
```

To configure an IPv6 nonsecure syslog server:

```
switch:admin> syslogadmin --set -ip 2001:db8:69bc:92:218:8bff:fe40:15c4
```

To set the syslog facility to LOG_LOCAL2:

```
switch:admin> syslogadmin --set -facility 2
switch:admin> syslogadmin --show -facility
Syslog facility: LOG_LOCAL2
```

To display all syslog IP addresses configured on a switch:

```
switch:admin> syslogadmin --show -ip
syslog.1 172.26.26.173
```

```
syslog.2 2001:db8:69bc:92:218:8bff:fe40:15c4
```

To remove the IP address fec0:60:69bc:92:218:8bff:fe40:15c4 from the list of servers to which error log messages are sent:

```
switch:admin> syslogadmin --remove -ip 2001:db8:69bc:92:218:8bff:fe40:15c4
```

To configure unsecured UDP port while configuring syslog server:

```
switch:admin> syslogadmin --set -ip 10.1.18.40 -port 512
Changing the unsecured port
Syslog IP address 10.1.18.40 and port 512 added
```

Using Secure Syslog CA Root Certificates

You can import and export syslog root CA certificates. You can have only one certificate installed.

- Use the following command to import syslog CA certificates.
seccertmgmt import -ca -server syslog
- Use the following command to delete syslog CA certificates.
seccertmgmt delete -ca -server syslog
- Use the following command to export syslog CA certificates.
seccertmgmt export -ca -server syslog
- Use the following command to display the syslog CA certificates.
seccertmgmt show -ca -server syslog

NOTE

You can also generate third-party signed certificates to be installed as client certificates.

Host Name Support for the syslogAdmin Command

The `syslogAdmin` command accepts either an IP address or a host name as input for the syslog server. The system does not validate the host name. You must ensure that you enter the correct host name.

- The following example shows the command using a host name and then using an IP address as input:

```
switch:admin> syslogadmin --set -ip win2k8-58-113
Syslog host name win2k8-58-113 added
switch:admin>
switch:admin> syslogadmin --set -ip 10.1.18.113
Syslog IP address 10.1.18.113 added
```

NOTE

The host name in the signed certificates must match the switch host name or IP address.

- To add a secure syslog server, use the following command:
switch:admin> syslogadmin --set -ip <ipaddress> -secure -port <portnumber>
- The following example shows the output of the syslog server information:
switch:admin> syslogadmin --show -ip
syslog.1 10.1.18.113
syslog.2 win2k8-58-113
- The following example shows the command used to remove the syslog server host name:
switch:admin> syslogadmin --remove -ip win2k8-58-113
Syslog host name win2k8-58-113 removed

Configuring Quiet Time for RASLog and Audit Log Messages

RASLog messages and audit log messages can be suppressed and sent only to Compact Flash (CF) storage (but not to other destinations such as a syslog server, SNMP, Telnet, or a console) by entering `rasadmin --quiet` on the active CP. The quiet time settings are preserved after both HA failover and switch reboot, but are overwritten when the command is run again for the same message type (RASLog messages, audit log messages, or both). When the quiet time is configured, a log message is generated that indicates that quiet time has been configured to enable or disable LOG or AUDIT messages. After that, LOG or AUDIT messages will be disabled or enabled depending on the configuration settings.

To configure quiet time for RASLog and audit log messages, complete the following steps.

1. Enter `rasadmin --quiet` followed by the desired quiet time options.

Use the following syntax options to specify the quiet time settings:

- `-enable log_type`. Valid values for `log_type` are: 1 (audit messages), 2 (RASLog messages), and 3 (both audit and RASLog messages).
- `-stime start_time`. The format for `start_time` is: HH:MM, based on a 24-hour clock. This is when the quiet time should start.
- `-etime end_time`. The format for `end_time` is: HH:MM, based on a 24-hour clock. This is when the quiet time should end.
- `dow Days_of_week`. Valid values for `Days_of_week` are: 1 (Monday) through 7 (Sunday). Indicate multiple days by separating the numbers using a comma. For example, Monday, Thursday, and Saturday would be 1,4,6.

2. Enter `rasadmin --quiet -show` to display the quiet time settings.

```
device:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	ON	06:00	12:00	MON, TUE
RASLOG	OFF	NA	NA	NA

NOTE

If a message is part of both RASLog messages and audit log messages, and the quiet time is set only for the audit log messages, the RASLog message is sent to the console, the syslog server, the local CF, and SNMP. If a message is part of both message types, and the quiet time is set only for the RASLog messages, the audit log message is sent to the local CF and the syslog server.

The quiet time configuration keys are visible in the `configshow` command output.

```
device:admin> configshow
(output truncated)
raslog.quiettime.enable
raslog.quiettime.starttime
raslog.quiettime.endtime
raslog.quiettime.dow
.
.
.
auditlog.quiettime.enable
auditlog.quiettime.starttime
auditlog.quiettime.endtime
auditlog.quiettime.dow
(output truncated)
```


NOTE

The configuration keys are set to their default values when you upgrade to Fabric OS 8.1.0 or later. The configuration keys are removed when you downgrade to an earlier version of Fabric OS.

Examples

The following example displays the configured quiet time.

```
device:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	ON	06:00	12:00	MON, TUE
RASLOG	ON	07:00	12:00	EVERYDAY

The following example immediately suppresses the audit log messages. Audit log messages will not be displayed until the setting is changed.

```
device:admin> rasadmin --quiet -enable 1
2016/11/28-11:19:46, [LOG-1012], 2070, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is enabled
for Audit LOG, for ever.
```

The following example enables quiet time for LOG and AUDIT messages without an end date.

```
device:admin> rasadmin --quiet -enable 3
2016/11/28-11:19:46, [LOG-1012], 2072, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is enabled
for ALL, for ever.
```

```
switch:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	ON	NA	NA	FOREVER
RASLOG	ON	NA	NA	FOREVER

The following example disables quiet time for LOG and AUDIT messages.

```
device:admin> rasadmin --quiet -disable 3
2016/11/28-11:17:46, [LOG-1013], 2071, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is disabled
for ALL.
```

```
switch:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	OFF	NA	NA	NA
RASLOG	OFF	NA	NA	NA

The following example automatically enables quiet time for LOG and AUDIT messages every day at 7 a.m. (07:00), and disables them at noon (12:00 pm).

```
device:admin> rasadmin --quiet -enable 3 -stime 07:00 -etime 12:00
2016/11/28-11:24:44, [LOG-1012], 2073, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is enabled
for ALL, with stime and etime.
```

```
switch:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	ON	07:00	12:00	EVERYDAY
RASLOG	ON	07:00	12:00	EVERYDAY

The following example automatically enables quiet time for LOG and AUDIT messages every day at 10 p.m. (22:00), and disables them the next day at 2 a.m. (02:00).

```
device:admin> rasadmin --quiet -enable 3 -stime 22:00 -etime 02:00
```

2016/11/28-11:29:44, [LOG-1012], 2075, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is enabled for ALL, with stime and etime.

```
switch:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	ON	22:00	02:00	EVERYDAY
RASLOG	ON	22:00	02:00	EVERYDAY

The following example automatically enables quiet time for both LOG and AUDIT messages on Mondays and Thursdays at 7 a.m. (07:00), and disables them at Noon (12:00 pm).

```
device:admin> rasadmin --quiet -enable 3 -stime 07:00 -etime 12:00 -dow 1,4
```

2016/11/28-11:31:27, [LOG-1012], 2076, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is enabled for ALL, with stime, etime and dow.

```
switch:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	ON	07:00	12:00	Mon Thu
RASLOG	ON	07:00	12:00	Mon Thu

The following example automatically enables quiet time for both LOG and AUDIT messages on Tuesdays and Fridays at 10 p.m. (22:00), and disables them the next day at 2 a.m. (02:00).

```
device:admin> rasadmin --quiet -enable 3 -stime 22:00 -etime 02:00 -dow 2,5
```

2016/11/28-11:33:04, [LOG-1012], 2077, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is enabled for ALL, with stime, etime and dow.

```
switch:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	ON	22:00	02:00	Tue Fri
RASLOG	ON	22:00	02:00	Tue Fri

The following example enables quiet time for AUDIT messages forever and also automatically enables quiet time for LOG messages on Mondays and Wednesdays at 10 p.m. (22:00), and disables them the next day at 2 a.m. (02:00).

```
device:admin> rasadmin --quiet -enable 1
```

2016/11/28-11:35:20, [LOG-1012], 2079, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is enabled for Audit LOG, for ever.

```
device:admin> rasadmin --quiet -enable 2 -stime 22:00 -etime 02:00 -dow 1,3
```

2016/11/28-11:35:02, [LOG-1012], 2078, SLOT 1 CHASSIS, INFO, Sherman_58, Ras Quiet Time is enabled for RAS LOG, with stime, etime and dow.

```
switch:admin> rasadmin --quiet -show
```

Type	QuietTime	StartTime	EndTime	DayOfWeek
AUDIT	ON	NA	NA	FOREVER
RASLOG	ON	22:00	02:00	Mon Wed

Duplicate PWWN Handling during Device Login

If a device attempts to log in with the same port WWN (PWWN) as another device on the same switch, you can configure whether the new login or the existing login takes precedence.

Note that if a device attempts to log in with the same PWWN as another device on a different switch, both devices are logged out. This section describes what happens with devices logging in on the same switch.

You can configure how duplicate PWWNs are handled by selecting an option in the Enforce FLOGI/FDISC login prompt of the `configure` command:

- Setting 0: First login takes precedence over second login (default behavior).
- Setting 1: Second login overrides first login.
- Setting 2: The port type determines whether the first or second login takes precedence.

Setting 0, First Login Precedence

When setting 0 is selected, the first login takes precedence over the second. This is the default behavior. The following table describes the behavior when setting 0 is selected.

Table 17: Duplicate PWWN Behavior: First Login Takes Precedence over Second Login

Input Port	First Port Login is NPIV Port	First Port Login is F_Port
FLOGI received	The new login is rejected and the new port is persistently disabled.	The new login is rejected and the new port is persistently disabled.
FDISC received	The new FDISC is rejected.	The new FDISC is rejected.

Setting 1, Second Login Precedence

When setting 1 is selected, the second login takes precedence over the first. The following table describes the behavior when setting 1 is selected.

Table 18: Duplicate PWWN Behavior: Second Login Overrides First Login

Input Port	First Port Login is F_Port	First Port Login is NPIV Port
FLOGI received	New login forces an explicit logout of original login on the previous F_Port. The previous F_Port is persistently disabled.	New login forces an explicit logout of original FDISC on the previous NPIV port. If Base Device Logout is enabled on the NPIV port, only the base device is logged out and the remaining NPIV devices stay logged in.
FDISC received	New FDISC forces an explicit logout of original login on the previous F_Port. The previous F_Port is persistently disabled.	New FDISC forces an explicit logout of original FDISC on the previous NPIV port. If Base Device Logout is enabled on the NPIV port, only the base device is logged out and the remaining NPIV devices stay logged in.

Setting 2, Mixed Precedence

When setting 2 is selected, the precedence depends on the port type of the first login:

- If the previous port is an F_Port, the first login takes precedence.
- If the previous port is an NPIV port, the second login overrides the first login.

Table 19: Duplicate PWWN Behavior: Port Type Determines which Login Takes Precedence

Input Port	First Port Login is NPIV Port	First Port Login is F_Port
FLOGI received	New login forces an explicit logout of original FDISC on the previous NPIV port. If Base Device Logout is enabled on the NPIV port, only the base device is logged out and the remaining NPIV devices stay logged in.	New login is rejected and the new port is persistently disabled.
FDISC received	New FDISC forces an explicit logout of original FDISC on the previous NPIV port. If Base Device Logout is enabled on the NPIV port, only the base device is logged out and the remaining NPIV devices stay logged in.	New FDISC is rejected.

Setting the Behavior for Handling Duplicate PWWNs

You can configure how duplicate port WWNs (PWWNs) are handled if a device attempts to log in with the same PWWN as another device on the switch.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `switchDisable` command to disable the switch.
3. Enter the `configure` command.
4. Enter **y** after the F_Port login parameters prompt.

```
F-Port login parameters (yes, y, no, n): [no] y
```

5. Enter one of the following options at the Enforce FLOGI/FDISC login prompt to select the behavior for handling duplicate PWWNs.
 - Enter **0** to have the first login take precedence over the second login (default).
 - Enter **1** to have the second login override the first login.
 - Enter **2** to have the port type determine the behavior.

If a duplicate login is received on an F_Port, the duplicate login is rejected and the old login is preserved; if a duplicate login is received on an NPIV port, the newer login is accepted.

```
Enforce FLOGI/FDISC login: (0..2) [0] 1
```

6. Respond to the remaining prompts, or press **Ctrl + D** to accept the other settings and exit.
7. Enter the `switchEnable` command to re-enable the switch.

With any of these settings, detection of duplicate PWWNs results in a RASLog. Ports that are restricted become persistently disabled, marked with the reason "Duplicate Port WWN detected".

Forward Error Correction

Forward error correction (FEC) provides a data transmission error control method by including redundant data (error-correcting code) to ensure error-free transmission on a specified port or port range.

The supported data transmissions by FEC are:

- When 10/16G FEC is enabled, it can correct one burst of up to 11-bit errors in every 2112-bit transmission, whether the error is in a frame or a primitive.
- When 32G FEC is enabled, it can correct up to 7 symbols in every 5280-bit transmission. A symbol consists of 10 bits, so there are 528 symbols in every 5280-bit transmission.
- When 64G FEC is enabled, it can correct up to 15 symbols in every 5440-bit transmission. A symbol consists of 10 bits, so there are 544 symbols in every 5440-bit transmission.

Because FEC is optional at 10G and 16G speeds, the Transmitter Training Signal (TTS) was extended to include a means to negotiate FEC capabilities. FEC is negotiated and activated when both sides of the link have FEC enabled. The FEC active indicator in the Fabric OS indicates whether or not FEC was successfully negotiated or not. FEC uses unused bits within the signaling protocol to generate an error correcting code (ECC) and correct bits as needed.

The following considerations apply to FEC:

- FEC is always active on ports operating at 64G and 32G speed based on the standard, regardless of the FEC settings on the switch.
- FEC is supported on E_Ports on 16G-capable switches at 10G/16G speeds.
- FEC is supported on the N_Ports and F_Ports of an Access Gateway using RDY, Normal (R_RDY), or Virtual Channel (VC_RDY) flow control modes.
- FEC is supported on F_Ports if the device attached to it supports FEC.
- FEC is enabled by default.
- FEC enables automatically when negotiation with a switch detects FEC capability.
- FEC persists after driver reloads and system reboots.
- FEC functions with features, such as QoS, trunking, and BB_Credit recovery.

FEC Limitations

The following limitations apply to Forward Error Correction (FEC):

- FEC is configurable on 64G and 32G-capable switches (Brocade G720, G610, G620, G630, 7810 switches and the Brocade X7 and X6 Director family). 64G is supported only from Fabric OS 9.0.0 or later whereas 32G is supported from Fabric OS 8.0.0 and later. FEC is mandatory for ports operating at 32G or higher speed and hence enabled by default.
- For switch-to-adapter connections, FEC is supported on QLogic BR-1860 and BR-1867 with driver v3.2.x on the Fabric Adapter ports operating in HBA mode and Emulex Lancer G6 and Prism based HBAs connected to 16G and 32G Brocade switches.
- FEC is supported only on link speeds of 10G, 16G, 32G, and 64G regardless of whether the platform is FEC-capable.
- To connect between a switch and an HBA at 16G or 32G or 64G, both sides must be in the same mode (port speed, and FEC on or off) for them to communicate at that rate. If only one port has FEC enabled, neither port will be able to see the other. If the ports are in dynamic mode, then they may connect, but not at 16G or 32G or 64G.
- Some Dense Wavelength Division Multiplexing (DWDM) in transparent mode can support FEC. For details, check with the DWDM vendor

Disabling FEC

Use the following procedure to disable FEC.

ATTENTION

Disabling FEC is disruptive to traffic and can be disabled at 10G, 16G, and 32G. However, FEC is always active (even if disabled) at 32G.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `portCfgFec --disable` command, specifying the port or range of ports on which FEC is to be enabled.

```
portcfgfec --disable -FEC slot/port
```

3. Enter the `portCfgFec --show` command to display the current FEC configuration.

```
portcfgfec --show slot/port
```

Disabling FEC on a Single Port

```
switch:admin> portcfgfec --disable -FEC 1
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
```

```
switch:admin> portcfgfec --show 1
```

```
Port: 1
FEC Capable: YES
10G/16G FEC Configured: OFF
16G FEC via TTS Configured: OFF
FEC State: Inactive
```

Enabling or Disabling FEC for Long-Distance Ports

To enable or disable FEC for long-distance ports, use `portCfgLongDistance` with the `-fecEnable` or `-fecDisable` parameter as required.

```
switch:admin> portcfglongdistance 12/6 LS 1 -distance 100 -fecenable
```

See [Managing Long-Distance Fabrics](#) for more details on working with long-distance ports.

Displaying and Clearing the FEC Counters

Both correctable and uncorrectable FEC counters help you to identify link degradation and take actions.

1. Enter the `portstatsshow` and `portstats64show` commands to display the counters.

```
switch:admin> portstatsshow 16 | grep fec
fec_cor_detected      1111      Count of blocks that were corrected by FEC
fec_uncor_detected    0        Count of blocks that were left uncorrected by FEC
switch:admin> portstats64show 16
...
stat64_fec_cor  0          top_int : FEC corrected errors detected
                1111      bottom_int : FEC corrected errors detected
stat64_fec_uncor 0          top_int : FEC uncorrected errors detected
                0          bottom_int : FEC uncorrected errors detected
```

NOTE

Gen 7 and Gen 6 platforms display `fec_uncor_detected` and `fec_corrected_rate`.

2. Enter the `portstatsclear` command to clear the counters.

Fabric-Assigned PWWN

This section describes about Fabric-Assigned PWWN (FA-PWWN) that simplifies the server deployment in a Fibre Channel SAN (FC SAN) environment by using a virtual port World-Wide Name (PWWN) instead of a physical PWWN to configure zoning and LUN mapping and masking.

Server deployment typically requires multiple administrative teams such as server and storage teams. These teams perform configuration tasks such as zone creation in the fabric, LUN mapping, and masking on the storage device. The configuration tasks must be completed before the server is deployed. Before you can configure WWN zones and LUN masks, you must determine the physical PWWN of the server. The administrative teams cannot start their configuration tasks until the physical server arrives (and its physical PWWN is known). Because the configuration tasks are sequential and interdependent across various administrative teams, it may take several days before the server is deployed in an FC SAN.

You can use a FA-PWWN to configure services such as zoning and LUN masking before you have a physical server. An FA-PWWN is a "virtual" port WWN that can be used instead of the physical PWWN. When the server is later attached to the SAN, the FA-PWWN is assigned to the server.

For example, you can use the FA-PWWN feature to perform the following tasks:


- Replace one server with another server, or replace failed HBAs or adapters within a server, without having to change any zoning or LUN mapping and masking configurations.
- Easily move servers across ports or Access Gateways by way of reassigning the FA-PWWN to another port.
- Use the FA-PWWN to represent a server in boot LUN zone configurations so that any physical server that is mapped to the FA-PWWN can boot from that LUN, thus simplifying boot over SAN configuration.

NOTE

The server must use a supported HBA or adapter and must be enabled to use the FA-PWWN feature. Refer to the release notes for the HBA or adapter versions that support this feature.

You can configure an FA-PWWN for the following topologies:

- An FA-PWWN for an HBA device that is connected to an Access Gateway switch
- An FA-PWWN for an HBA device that is connected directly to an edge switch

Figure 2: Fabric-assigned port World-Wide Name provisioning scenarios

User-Assigned and Auto-Assigned FA-PWWN Behavior

Each switch port and Access Gateway port can have up to two FA-PWWNs, one assigned automatically and one assigned by the user. FA-PWWNs must be unique, and only one FA-PWWN can be active at any given time.

The automatically assigned FA-PWWN is created by default if you enable the feature without explicitly providing a virtual PWWN.

The user-assigned FA-PWWN takes precedence over the automatically assigned FA-PWWN. This means the switch will bind the user-assigned FA-PWWN to the port if both a user-assigned and an automatically assigned FA-PWWN are available. If you want to select the automatically assigned FA-PWWN over the user-assigned FA-PWWN, you must delete the user-assigned FA-PWWN from the port to which it has been assigned.

The switch ensures that automatically assigned FA-PWWNs are unique in a fabric. However, it is your responsibility to ensure that user-assigned FA-PWWNs are also unique throughout the fabric.

ATTENTION

You must ensure that the same user-assigned FA-PWWN is not used in multiple chassis. There is no fabric-wide database, so adding the same FA-PWWNs in multiple chassis causes duplicate PWWNs.

Configuring an FA-PWWN for an HBA Connected to an Access Gateway

To configure an FA-PWWN, assign the FA-PWWN on the Access Gateway switch. The FA-PWWN feature is enabled by default on the HBA.

1. Log in to the edge switch to which the Access Gateway is directly connected.
2. Assign the FA-PWWN.

- If you are manually assigning a WWN, enter the following command:

```
fapwwn --assign -ag AG_WWN -port AG_port -v Virtual_PWWN
```

- If you want the WWN to be automatically assigned, enter the following command:

```
fapwwn --assign -ag AG_WWN -port AG_port
```

3. Display the FA-PWWN.

```
fapwwn --show -ag all
```

You should see output similar to the following sample. The FA-PWWNs are in the Virtual Port WWN column. (In this example, long lines of output are shown split across two lines, for better readability.)

```
-----
AG Port          Port      Device Port WWN  \
-----
10:00:00:05:1e:65:8a:d5/16  --      --:--:--:--:--:--:-- \
10:00:00:05:1e:d7:3d:dc/8   20      20:08:00:05:1e:d7:2b:74 \
10:00:00:05:1e:d7:3d:dc/9   20      20:09:00:05:1e:d7:2b:73 \
10:00:00:05:1e:d7:3d:dc/16  --      --:--:--:--:--:--:-- \
-----
Virtual Port WWN  PID      Enable MapType
-----
52:00:10:00:00:0f:50:30  --      Yes   AG/Auto
11:22:33:44:55:66:77:88  11403   Yes   AG/User
52:00:10:00:00:0f:50:33  11404   Yes   AG/Auto
52:00:10:00:00:0f:50:38  --      Yes   AG/Auto
```

4. Display the FA-PWWN on the HBA.

The following steps are to be executed on the server and not the switch.

- a) Log in to the server as admin.
- b) Enter the following command to display the FA-PWWN.

```
# bcu port --faa_query port_id
```

```
FAA state: Enabled
PWWN: 11:22:33:44:55:66:77:88
PWWN source: Fabric
```

The HBA retains the FA-PWWN until rebooted. This means you cannot unplug and plug the cable into a different port on the Access Gateway. You must reboot the HBA before moving the HBA to a different port.

If you move an HBA to a different port on a switch running Fabric OS 8.1.0 or later, the HBA will disable its port.

Configuring an FA-PWWN for an HBA Connected to an Edge Switch

To configure an FA-PWWN, assign the FA-PWWN on the edge switch. The FA-PWWN feature is enabled by default on the HBA.

1. Log in to the edge switch to which the device is connected.
2. Assign the FA-PWWN.
 - If you are manually assigning a WWN, enter the following command:

```
fapwwn --assign -port [ slot/] port -v Virtual_PWWN
```

- If you want the WWN to be automatically assigned, enter the following command:

```
fapwwn --assign -port [ slot/] port
```

3. Display the FA-PWWN.

```
fapwwn --show -port all
```

You should see output similar to the following sample. The FA-PWWNs are in the VPWWN column.

```
-----
Port      PPWWN                                VPWWN                                PID Enable MapType
-----
0  --:--:--:--:--:--:--:--:--:--:--  52:00:10:00:00:0f:50:30 10101 Yes  Port/Auto
1  --:--:--:--:--:--:--:--:--:--:--  11:22:33:44:33:22:11:22 --   Yes  Port/User
                                     52:00:10:00:00:0f:50:44
10 --:--:--:--:--:--:--:--:--:--:--  52:00:10:00:00:0f:50:45 --   Yes  Port/Auto
-----
```

4. Display the FA-PWWN on the HBA.

The following steps are to be executed on the server and not the switch.

- a) Log in to the server as admin.
- b) Enter the following command to display the FA-PWWN.

```
# bcu port --faa_query port_id
```

```
FAA state: Enabled
PWWN: 11:22:33:44:33:22:11:22
PWWN source: Fabric
```

The HBA retains the FA-PWWN until it is rebooted. This means you cannot unplug and plug the cable into a different port on the switch. You must reboot the HBA before moving the HBA to a different port.

If you move an HBA to a different port on a switch running Fabric OS v7.0.0 or later, the HBA will disable its port. The port remains disabled even if you then move the HBA to a port on a switch running a version of Fabric OS earlier than 7.0.0.

Supported Switches and Configurations for FA-PWWN

The FA-PWWN feature is supported only on switches running Fabric OS 8.1.x or later and only on Brocade HBAs and adapters. The HBA can be connected to an edge switch or to an Access Gateway switch.

The FA-PWWN feature is supported on the following platforms:

- Gen 7 (64G) platforms running Fabric OS 9.1.x or later
 - Brocade G730
- Gen 7 (64G) platforms running Fabric OS 9.0.x or later
 - Brocade G720
 - Brocade X7 Directors
- Gen 6 (32G) platforms running Fabric OS 8.1.x or later
 - Brocade G610
 - Brocade G620
 - Brocade G630
 - Brocade X6 Directors

Configuration Upload and Download Considerations for FA-PWWN

The configuration upload and download utilities can be used to import and export the FA-PWWN configuration.

ATTENTION

You must delete all FA-PWWNs from the switch with the configuration being replaced before you upload or download a modified configuration. This is to ensure no duplicate FA-PWWNs are in the fabric.

Security Considerations for FA-PWWN

If security is a concern, ensure that only authorized users can configure FA-PWWNs. Device authentication and DCC policies provide additional security between the switch and the server.

The FA-PWWN feature can be enabled only by authorized administrators. Thus, existing user-level authentication and authorization mechanisms should be used to ensure only authorized users can configure this feature.

If you are concerned about security for FA-PWWNs, you should configure device authentication. You can use authentication at the device level to ensure security between the switch and the server. See [Device Authentication Policy](#) for information about configuring device authentication.

You can also use the Device Connection Control (DCC) policy to ensure that only an authorized physical server can connect to a specific switch port.

NOTE

When creating the DCC policy, use the physical device WWN and not the FA-PWWN.

If you use DCC, a policy check is done on the physical PWWN on the servers. In the case of an HBA, the FA-PWWN is assigned to the HBA only after the DCC check is successful. See [DCC Policy Behavior with Fabric-Assigned PWWNs](#) for additional information.

Restrictions of FA-PWWN

The FA-PWWN feature is not supported with some Fibre Channel fabric features.

FA-PWWN is not supported for the following:

- FCoE devices
- Cascaded Access Gateway topologies
- FICON/FMS mode
- With F_Port trunking on directly attached Brocade HBAs or adapters

NOTE

FA-PWWN is supported with F_Port trunking on the supported Access Gateway platforms.

Access Gateway N_Port Failover with FA-PWWN

If an Access Gateway is connected to multiple switches, you should configure the same FA-PWWNs on both switches to avoid having to reboot the host in case of failover.

If the same FA-PWWNs are not configured on the switches, and if an FA-PWWN F_Port on an Access Gateway fails over to an N_Port that is connected to a different switch, the FA-PWWN assigned to the Access Gateway F_Port following the failover will be different than it was before the failover occurred. This situation may require the host to reboot to bring it back online. Even after the reboot, the host may potentially go into a different zone because the FA-PWWN is different.

N_Port ID Virtualization

This section describes the N_Port ID Virtualization (NPIV) that enables a single Fibre Channel protocol port to appear as multiple, distinct ports, providing separate port identification within the fabric for each operating system image behind the port (as if each operating system image had its own unique physical port).

NPIV assigns a different virtual port ID to each Fibre Channel protocol device. NPIV is designed to enable you to allocate virtual addresses without affecting your existing hardware implementation. The virtual port has the same properties as an N_Port, and it registers with all the services of the fabric. This chapter does not describe the Access Gateway feature. For more information on the Access Gateway feature, refer to the *Brocade Fabric OS Access Gateway User Guide*.

Each NPIV device has a unique device PID, Port WWN, and Node WWN, and behaves the same as all other physical devices in the fabric. In other words, multiple virtual devices that are emulated by NPIV appear no different from regular devices that are connected to a non-NPIV port.

The same zoning rules apply to NPIV devices as non-NPIV devices. Zones can be defined by *domain,port* notation, by WWN-based zoning, or both. However, to perform zoning to the granularity of the virtual N_Port IDs, you must use WWN-based zoning.

If you are using *domain, port* zoning for an NPIV port, and all the virtual PIDs associated with the port are included in the zone, a port login (PLOGI) to a non-existent virtual PID is not blocked by the switch. It is delivered to the device attached to the NPIV port. In cases where the device is not capable of handling such unexpected PLOGIs, you must use WWN-based zoning.

The following example shows the number of NPIV devices in the output of the `switchshow` command. The number of NPIV devices is equal to the sum of the base port plus the number of NPIV public devices. The base port is the N_Port listed in the `switchshow` output. Based on the formula, index 010000 shows only 1 NPIV device and index 010300 shows a total of 222 NPIV devices (one N_Port FLOGI device and 221 NPIV devices).

```
switch:admin> switchshow
switchName:      switch180
switchType:      180.0
switchState:     Online
switchMode:      Access Gateway Mode
switchWwn:       10:00:00:05:1e:41:49:3d
switchBeacon:    OFF
Index Slot Port Address Media Speed State Proto
=====
0      3      0              010000 id    N32   Online FC F-Port  20:0c:00:05:1e:05:de:e4 0xa06601

1      3      1              010100 id    N32   Online FC F-Port  1 N Port + 4 NPIV public
2      3      2              010200 id    N32   Online FC F-Port  1 N Port + 119 NPIV public
3      3      3              010300 id    N32   Online FC F-Port  1 N Port + 221 NPIV public
```

NOTE

When an Access Gateway is connected to the switch, the Access Gateway is counted as the base device or base login and it is not included in the NPIV device count.

Unified Addressing Mode

Unified addressing mode (UAM) is the default addressing mode in all Gen 7 platforms irrespective of whether Virtual Fabrics is enabled or not.

The number of NPIV devices supported in both VF and non-VF mode is 255 devices.

NOTE

For more information on UAM, see the [Unified Addressing Mode](#) section in Performing Advanced Configuration Tasks chapter.

Fixed Addressing Mode

Fixed addressing mode is the default addressing mode used in all Gen 6 platforms that do not have Virtual Fabrics enabled.

The number of NPIV devices supported on shared area ports (48-port blades) is reduced to 64 from 128 when Virtual Fabrics mode is enabled.

10-Bit Addressing Mode

10-bit addressing mode is the default mode for all logical switches created in the Brocade X6 Director. The number of NPIV devices supported on a port is 64.

In an X6-8 Director, the switch can dynamically assign up to 112, 8-bit areas. In X6-4 Director, the switch can dynamically assign up to 144, 8-bit areas, after that 10-bit areas are assigned. So the 8-bit areas can have up to 255 devices, and the 10-bit areas can have up to 63 devices.

The following table shows the number of NPIVs devices supported on the Brocade X6 Directors.

Table 20: Number of Supported NPIV devices Per Port

Platform	VF Disabled	VF-Enabled Default Switch	VF-Enabled Logical Switch	VF-Enabled Base Switch
X6-4	Yes, 255	Yes, 255	Yes, 255*	No NPIV support
X6-8	Yes, 127	Yes, 63	Yes, 255**	No NPIV support

NOTE

- All numbers indicate the virtual device limit.
- *The first 112 ports present in the logical switch using 10-bit addressing allows up to 255 logical devices. The physical ports above 112 are limited to login of 63 devices.
- **If port is assigned a 10-bit area, the number of NPIV devices supported is 63.

Configuring NPIV

The NPIV feature is enabled by default. You can set the number of virtual N_Port IDs per port to a value from 1 through 255 per port. The default setting is 126.

The `portCfgNPIVPort` command specifies the maximum number of virtual N_port IDs per port on a switch. It can also enable or disable NPIV. When NPIV is enabled on the port, you can specify the number of logins per port. If the NPIV feature has been disabled, the NPIV port configuration does not work.

The addressing mode can limit the maximum number of NPIV logins to 127 or 63 depending on the mode. The `portCfgNPIVPort` command can set the maximum number of NPIV logins limit to anything from 1 through 255, regardless of the addressing mode. Whichever of these two (addressing mode or the value configured through `portCfgNPIVPort`) is lower will be the maximum number that can be logged in.

**CAUTION**

The `portDisable` command disables the port and stops all traffic flowing to and from the port. Use this command during a scheduled maintenance.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portDisable` command.
3. Enter the `portCfgNPiVPort --setloginlimit` command with the port number and the number of logins per port.
4. Enter the `portEnable` command to enable the port.

```
switch:admin> portcfgnpivport --setloginlimit 4/12 255
NPIV Limit Set to 255 for Port 108
switch:admin> portcfgshow 4/12
Area Number:                108
Octet Speed Combo:          1 (32G|16G|8G|4G)
FC Speed Level:              AUTO (HW)
Trunk Port                   ON
Long Distance                OFF
VC Link Init                 OFF
Disabled E_Port              OFF
Locked E_Port                OFF
ISL R_RDY Mode               OFF
RSCN Suppressed              OFF
Persistent Disable           OFF
LOS TOV mode                 0 (OFF)
NPIV capability              ON
QOS Port                     AE
Port Auto Disable:           OFF
EX Port                      OFF
Mirror Port                  OFF
SIM Port                     OFF
Credit Recovery              ON
F_Port Buffers               OFF
E_Port Credits               OFF
Fault Delay:                 0 (R_A_TOV)
NPIV PP Limit:               255
NPIV FLOGI Logout:           OFF
CSCTL mode:                  OFF
TDZ mode:                    OFF
Clean Address Bit:           OFF
D-Port mode:                 OFF
D-Port over DWDM:            OFF
Compression:                 OFF
10G/16G FEC:                 ON
16G FEC via TTS:             OFF
```

Enabling and Disabling NPIV

NPIV is enabled for every port.

NOTE

NPIV is a requirement for FCoE.

1. Connect to the switch and log in using an account assigned to the admin role.
2. To enable or disable NPIV on a port, enter the `portCfgNPIVPort` command with either the `--enable` or `--disable` option.

The following example shows NPIV being enabled on port 10 of a Brocade G630:

```
switch:admin> portCfgNPIVPort --enable 10
```

NOTE

If the NPIV feature is disabled, the port is toggled if NPIV devices are logged in from that F_Port (a true NPIV port). Otherwise, the firmware considers that port as an F_Port even though the NPIV feature was enabled.

Base Device Logout

Base device logout is a Fibre Channel - Lin Service 2 (FC-LS2) standard-based feature in which the Fabric OS firmware allows NPIV devices to remain logged in after the base device logs out. This feature is available on all the switches and Access Gateways starting with Fabric OS 7.3.0. Any device attached to an NPIV port is able to log in and log out without affecting the login status of the other devices on the port. This section explains the changes in F-port behavior, use case scenarios including possible changes in existing Fabric OS show commands, new CLIs and configuration commands associated with this feature. This section explains also the possible risks, compatibility (FW upgrade/downgrade) and HA aspects of Fabric OS related to this feature.

From their inception, Brocade switches have treated a base device logout on a given port as a logout of all devices on that port. For NPIV devices, this means that each NPIV device is logged out when the base device is logged out, which is referred to as the legacy functionality. When the final FC-LS2 standard was released, the functionality it defined had changed from the initial understanding. Instead of the base device logout causing all NPIV devices to log out, the base device logout affects only the base device; the NPIV devices stay logged in.

Difference in Device Logout Behaviors

A base device is a device on an F_Port that has the base PID. The base device logs in with a FLOGI. An NPIV device is a device on an F_Port that has an NPIV PID. An NPIV device logs in with an FDISC.

The following table summarizes the difference in the base and NPIV device logout behaviors with and without the base device logout feature enabled.

Table 21: Base and NPIV Logout Behaviors

Scenario	Without Base Device Logout Enabled	With Base Device Logout Enabled
No device is logged in.	Requires a device with a base PID to log in to the F_Port before any NPIV device can log in.	Requires a device with a base PID to log in to the F_Port before any NPIV device can log in.
Base device is logged in.	- Allows NPIV devices to log in and log out. - Permanent Port Name (PPN) is assigned with PWWN of the base device for all the devices logged-in on that port (for both base and NPIV devices).	- Allows NPIV devices to log in and log out. - Permanent Port Name (PPN) is assigned according to FC-GS-6 (5.2.3.13) standard.

Scenario	Without Base Device Logout Enabled	With Base Device Logout Enabled
- Base device is logged in. - At least one NPIV device is logged in.	- If the base device is logged out, then all the NPIV devices are also logged out. - Removes routing information when the base device logs out.	- Allows base device to log out while NPIV devices are still logged in. Indicates that the base device is logged out in the <code>switchShow</code> command output. - Does not remove the routing information.
- Base device is logged out. - At least one NPIV device is logged in.	N/A	- Allows base device to log back in as a base or NPIV device. - Allows NPIV devices to log in and log out. - Does not remove the routing information.
- Base device is logged out. - All NPIV devices are also logged out.	Requires a device with a base PID to log in to the F_Port before any NPIV device can log in.	- Removes all the routing information. - Requires a device with a base PID to log in to the F_Port before any NPIV device can log in.

Enabling Base Device Logout

Both the active and the standby switches/AGs should be upgraded to Fabric OS 7.3.0 or later.

The purpose of this feature is to make it possible for all devices; including base device and NPIV devices on a NPIV port to log out and log in without disrupting the remaining logged on devices. By default, the base device logout option is disabled in all the ports.

1. Enable NPIV on the required ports.

Ports that do not have NPIV capability cannot have the base device logout option enabled.

2. Disable the ports for which you want to enable base device logout, as this is a disruptive operation.
3. Enable base device logout using the `portCfgFLOGILogout` command.

The following example enables base device logout on all the ports in the logical switch.

```
portcfgflogillogout --enable -all
```

The following example enables base device logout on only specific ports in the logical switch.

```
portcfgflogillogout --disable 7/1-5
```

Use Cases and Dependencies

The following use cases and dependencies apply to the base device logout feature:

- When the base device logs out, the remaining NPIV devices on that port remain logged in.
- The base device can log in again later on the same port as the base device with a FLOGI. However, all the existing NPIV devices will be dropped and only the new FLOGI information will be retained.
- The base device can log in again later on the same port with the same PWWN, but as an NPIV device. At this point, it is no longer referred to as the base device, and it becomes a logged-in NPIV device.
- NPIV devices can log in or log out on the same port as long as at least one device (either a base device or an NPIV device) remains logged in on the port.
- The logged-out base device can log in on a different port as either the base device (with a FLOGI) or as an NPIV device (with an FDISC). If the base device logs in on a different port as a base device, it no longer is the base device for the remaining NPIV devices on the original port.
- The logged-out base device can log in to a port on a different switch as either the base device (with a FLOGI) or as an NPIV device (with an FDISC). If the base device logs in to a port on a different switch as a base device, it is no longer the base device for the remaining NPIV devices on the original port.
- When all devices (base device and NPIV devices) have logged out of a port, a base device must be the next device to log in on that port with a FLOGI.
- Fabric OS provides the ability to update NS entries when these base devices perform logout. You can turn on this update mode to ensure NS and login database consistency using the "F-port Device Update Mode" option under the "Fabric parameters" section of the configure CLI. This option can only be changed when a switch is disabled. After enabling this option, all other changes for this feature are completely transparent except minor changes in `fcProbeShow` output. In other words, the firmware treats all the base devices (AL_PAs 0x00, 0x40, 0x80, 0xC0) in the same way.
- The base device logout feature is backward compatible when you do not enable base device logout for the port.
- Default switch behavior is the legacy functionality (base device logout is disabled).
- All checks for duplicate WWN remain valid.
- The base device logout feature can coexist with other switches that do not have base device logout enabled, even if they are in the same fabric.
- Any port-related Fabric OS features that depend on the base device cannot work when the base device logs out. One example of this is device probing.
- Trunking is not affected by base device logout, because only Brocade switches utilize F_Port trunking, and the related Access Gateway N_Ports do not keep NPIV devices logged in when the base device logs out.
- QoS, FEC, and credit recovery features are not affected by base device logout.
- Port swap is not affected by the base device logout feature.
- Connectivity to NPIV devices is not affected when the base device logs out.
- Base device logout is not supported on ICL or VE/GE (E/Ex/LG) ports.
- Base device logout must be disabled on all the ports before downgrading to a lower Fabric OS version.

Viewing the Base Device Logout Setting

The following `portcfgShow` command output shows ON if the base device logout option is enabled, and a "." if it is disabled:

Ports of Slot 1	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Speed	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN
AL_PA Offset 13	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Trunk Port	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
Long Distance	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
VC Link Init	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Locked L_Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Locked G_Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Disabled E_Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..

Locked E_Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
ISL R_RDY Mode	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
RSCN Suppressed	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Persistent Disable	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
LOS TOV enable	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
NPIV capability	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
NPIV PP Limit	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126
NPIV FLOGI Logout	ON	ON	ON	..	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
QOS E_Port	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE
EX Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Mirror Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Rate Limit	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Credit Recovery	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
Fport Buffers	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Port Auto Disable	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
CSCTL mode	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
D-Port mode	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Fault Delay	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

The `switchShow` command displays the number of base devices as "1" and the number of NPIV devices. If the base device is logged out, the output displays only the number of NPIV devices. When base device logout is enabled and the base device has logged out, the `portShow` command output displays a FLOGI_LOGO with other port flags, and the PWWNs of the NPIV devices. When base device is logged out, the `agShow`, `fdmiShow`, and `nsShow` commands do not list the base device information.

Viewing NPIV Port Configuration Information

Fabric OS allows you to see N_Port ID Virtualization (NPIV) port configuration information using the `portcfgshow` and `switchshow` commands.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter `portcfgshow` to view the switch ports information.

The following example shows whether a port is configured for NPIV:

```
switch:admin> portcfgshow
```

Ports of Slot 0	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Speed	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN
Trunk Port	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
Long Distance	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
VC Link Init	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Locked L_Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Locked G_Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Disabled E_Port	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
ISL R_RDY Mode	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
RSCN Suppressed	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
Persistent Disable..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..
NPIV capability	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON

3. Enter `switchshow` to view NPIV information for a given port.

If a port is an F_Port, and you enter `switchshow`, the port WWN of the N_Port is returned. For an NPIV F_Port, there are multiple N_Ports, each with a different port WWN. The `switchshow` command output indicates whether or not a port is an NPIV F_Port, and identifies the number of virtual N_Ports behind it.

The following example displays the typical output from the `switchshow` command.

```
switch:admin> switchshow

switchName:    switch_G720
switchType:    181
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   29
switchId:       fffclD
switchWwn:      10:00:00:05:1e:82:3c:2a
zoning:         OFF
switchBeacon:   OFF
FC Router:      OFF
LS Attributes:  128
Index Port Address Media Speed State      Proto
=====
  0   0   1d0000   id    N64   Online          F-Port  1 Nport + 1 NPIV devices.
  1   1   1d0100   id    N64   No_Light
  2   2   1d0200   id    N64   Online          F-Port  20:0e:00:05:1e:0a:16:59
  3   3   1d0300   id    N64   No_Light
  4   4   1d0400   id    N64   No_Light(output truncated)
```

4. Use the `portshow port_number` command to view the NPIV attributes and all the N_Port (physical and virtual) port WWNs that are listed under “portWwn” for connected devices.

The following example displays the typical output for this command.

```
switch:admin> portshow 2

portName: 02
portHealth: HEALTHY
Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x24b03    PRESENT ACTIVE F_PORT G_PORT NPIV LOGICAL_ONLINE LOGIN NOELP LED ACCEPT
portType: 10.0
portState: 1    Online
portPhys: 6    In_Sync
portScn: 32    F_Port
port generation number: 148
portId: 630200
portIfId: 43020005
portWwn: 20:02:00:05:1e:35:37:40
portWwn of device(s) connected:
  c0:50:76:ff:fb:00:16:fc
  c0:50:76:ff:fb:00:16:f8
  ...(output truncated)
```

```

...
c0:50:76:ff:fb:00:16:80
50:05:07:64:01:a0:73:b8
Distance:   normal
portSpeed:  N2Gbps
Interrupts:      0          Link_failure: 16          Frjt:      0
Unknown:         0          Loss_of_sync: 422        Fbsy:      0
Lli:            294803      Loss_of_sig: 808
Proc_rqrd:       0          Protocol_err: 0
Timed_out:       0          Invalid_word: 0
Rx_flushed:      0          Invalid_crc:  0
Tx_unavail:      0          Delim_err:    0
Free_buffer:     0          Address_err: 1458
Overrun:         0          Lr_in:        15
Suspended:       0          Lr_out:       17
Parity_err:      0          Ols_in:       16
2_parity_err:    0          Ols_out:      15
CMI_bus_err:     0

```

Viewing Virtual PID Login Information

You can use the `portloginshow` command to display the login information for the virtual PIDs of a port. The following example is sample output from the `portloginshow` command.

```

switch:admin> portloginshow 2
Type  PID      World Wide Name      credit df_sz cos
=====
fe  630240  c0:50:76:ff:fb:00:16:fc  101  2048  c  scr=3
fe  63023f  c0:50:76:ff:fb:00:16:f8  101  2048  c  scr=3
fe  63023e  c0:50:76:ff:fb:00:17:ec  101  2048  c  scr=3
...
(output truncated)
...
ff  630202  c0:50:76:ff:fb:00:17:70  192  2048  c  d_id=FFFFFFC
ff  630201  c0:50:76:ff:fb:00:16:80  192  2048  c  d_id=FFFFFFC

```

To view the details of the devices that were last logged in to a particular port, you can use the `history` option in the `portloginshow` command. The device details such as type, PID and WWN along with the logged out timestamp in UTC format will be displayed.

```

switch:admin> portloginshow 1/0 -history
Type  PID      World Wide Name      logout time
=====
fd  550002  30:0c:02:05:1e:61:23:8f  09/17/2014 13:56:01
fd  550001  30:0c:01:05:1e:61:23:8f  09/17/2014 13:56:02
fe  550000  30:0c:00:05:1e:61:23:8f  09/17/2014 13:56:02
fd  550001  30:0c:01:05:1e:61:23:8f  09/18/2014 05:49:37

```

Maintaining the Switch Configuration File

Maintain the consistent configuration settings on all switches in the same fabric, because inconsistent parameters, such as inconsistent PID formats, can cause a fabric segmentation.

As part of standard configuration maintenance procedures, you must back up all important configuration data for every switch on a host computer server.

To view configuration settings for a switch in a Brocade fabric:

- Issue the `configShow -all` command.
To display configuration settings, connect to the switch, log in as admin, and enter the `configShow -all` command. The configuration settings vary depending on a switch type and configuration. This command does not show as much configuration information as the text file created from the `configUpload` command.
- Issue the `configUpload -all` command to upload an ASCII text file from the switch or switch module.
You can open the text file with any text editor to view the configuration information of the switch.



CAUTION

You must not edit the uploaded file. If an edited file is downloaded, it results in system errors.

If your user account has chassis account permissions, you can use any of the following options when uploading or downloading a configuration file:

Table 22: Configuration Management Options

fid	To upload the specified FID configuration.
all	To upload all the system configuration file including the chassis and switch sections for all logical switches. Use this parameter when obtaining a complete capture of the switch configuration in a switch that has a Virtual Fabrics mode disabled.
chassis	To upload only the chassis section of the system configuration file.
-switch	To upload the switch configuration only, if the Virtual Fabrics mode is disabled.
cra	Challenge Response Authentication (CRA) is to be used with Secure Copy Protocol (SCP). You can run this command without arguments for interactive mode. Only a password is supported as a challenge.
map	Uploads the port-to-area addressing mode configuration files. This command should be used in the FICON environment before replacing both the CP blades.
vf	Upload the virtual fabric data.

Configuration File Format

The configuration file is divided into three areas: the header, the chassis section, and one or more logical-switch sections.

Chassis Section

There is only one chassis section within a configuration. It defines configuration data for chassis components that affect the entire system, not just one individual logical switch. The chassis section is included in non-Virtual Fabric modes only if you use the `configUpload -all` command.

The chassis section specifies characteristics for the following software components:

- FC Routing–Fibre Channel Routing
- Chassis configuration–Chassis configuration
- FCOE_CH_CONF–FCoE chassis configuration
- UDROLE_CONF–User-defined role configuration
- LicensesDB–License Database (newer-format license storage)
- DMM_WWN–Data migration manager World Wide Name configuration
- License Database - (older-format license storage)
- AGWWN_MAPPING_CONF–Access Gateway WWN mapping configuration
- LicensesLservc–Sentinel License configuration
- GE blade mode–GigE Mode configuration
- FRAME_LOG–Frame log configuration (enable/disable)
- DMM_TB–Data migration manager configuration
- MOTD–Message of the day

Switch Section

At least one switch section for the default switch or a switch has Virtual Fabrics mode disabled, and additional sections correspond to each additionally defined logical switch instance on a switch with Virtual Fabrics mode enabled. This data is switch-specific and affects only that logical switch behavior.

The switch section of the configuration file contains information for all of the following:

- Boot parameters
- Configuration
- Bottleneck configuration
- Flow Vision configuration
- FCoE software configuration
- Zoning
- Defined security policies
- Active security policies
- iSCSI
- CryptoDev
- FICU saved files
- VS_SW_CONF
- MAPS configuration
- Banner

Configuration File Backup

You should keep a backup configuration file. You should keep individual backup files for all switches in the fabric and avoid copying configurations from one switch to another. The `configUpload` command, by default, only uploads the switch context configuration for the logical switch context in which the command is executed.

In non-Virtual Fabric mode, you must use the `configUpload -all` command to include both the switch and the chassis information. In Virtual Fabric mode, the `all` command can be selected to upload all logical switches and the chassis configuration. Only administrators with chassis permissions are allowed to upload other FIDs or the chassis configuration.

The following information is *not saved* in a backup:

- `dnsConfig` command information
- User accounts
- Passwords

Before you upload a configuration file, verify that you can reach the FTP server from the switch. Using a Telnet connection, save a backup copy of the configuration file from a logical switch to a host computer.

Secure File Transfer Protocol (SFTP) is now an option when uploading a configuration file. SFTP is analogous to Secure Copy Protocol (SCP). SFTP can be used for the `configupload`, `configdownload`, `supportsave`, and `auto FFDC/trace upload (supportftp)` commands.

NOTE

If you download a configuration file that was uploaded from a switch whose default FID (128) was changed, you should download the VF configuration file first before you download the normal configuration file to ensure that all the logical switches are the same in both switches. Follow these steps:

1. Run the `configupload -vf` command from the old setup.
2. Run the `configupload -all` command from the old setup.
3. Run the `configDownload -vf` command in the new setup.
4. Run the `configDownload -all` command in the new setup.

Special Characters in Server Credentials

The FTP, SCP, or SFTP server credentials may include special characters (also referred to as meta-characters) that are members of a special character set that are interpreted differently when used in command-line mode. The characters shown in following table have an alternate meaning or are designated to perform a special instruction.

Table 23: Special Characters Used in FTP, SCP, or SFTP Server Credentials

Symbols	Description
@ or \	Special characters used in a user name to separate the user name and domain.

Uploading a Configuration File in Interactive Mode

1. Verify that the FTP, SFTP, or SCP service is running on the host computer.
2. Connect to the switch and log in using an account with admin permissions.
3. Enter the `configUpload` command. The command becomes interactive and you are prompted for the required information.
4. Store a soft copy of the switch configuration information in a safe place for future reference.

```
switch:admin> configupload
Protocol (scp, ftp, sftp, local) [ftp]: sftp
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: UserFoo
Path/Filename [<home dir>/config.txt]: switchConfig.txt
Section (all|chassis|FID# [all]): chassis
```

```
username@10.1.2.3's password:
Password: <hidden>
configUpload complete
```

Configuration File Restoration

When you restore a configuration file, you overwrite the existing configuration with a previously saved backup configuration file.



CAUTION

Make sure that the configuration file you are downloading is compatible with your switch type. Downloading a configuration file from a different switch type or from a different firmware could cause your switch to fail.



CAUTION

If you have Virtual Fabrics enabled, you must follow the procedure in [Configuration Management for Virtual Fabrics](#) to restore the logical switches.



CAUTION

User-defined roles must be deleted prior to uploading the configuration. Any user-defined roles created after the last generated `configupload` command must be reapplied.

NOTE

- For FOS switches that are not enabled with Virtual Fabrics, configuration downloads are allowed for the same generation or major version across FOS 9.0.x or 9.1.x.
- When executing the `configdownload` command with a file that contains unsorted zone membership, any unsorted members are automatically sorted in the system. The automatic sorting can lead to zone merge conflicts when the switch is re-enabled after the configuration download is completed.

If a `configdownload` command is issued on a non-FC router, any FC router parameters can be viewed in the downloaded data.

You can use the `configdownload` command to restore the switch configuration.

```
configdownload -vf config_VF file
configdownload -all config_All file
```

Restrictions When Downloading a Configuration File

Gen 7 platforms, including Gen 7 platforms that were upgraded from Gen 6 platforms, do not support the configuration files from Gen 6 or older platforms. For Brocade G620 and G630 platforms, the configuration file uploaded from a switch can be downloaded onto another switch of the same switch type (162, 173, 183, and 184).

Supported Configuration Files

Configuration files are only supported between the same switch type. For more information on switch types that are mapped to Brocade product names, see [Switch Type](#).

The following table lists restrictions for some of the options of the `configDownload` command.

NOTE

If you modify the configuration attributes of a user in the key file, the configuration download is blocked.

NOTE

If the admin domain commands are blocked, you cannot download any admin domain configuration from the configuration file.

Table 24: Restrictions for Using the Options of the configDownload Command

Option	Restrictions
-chassis	The number of switches defined in the downloaded configuration file must match exactly the number of switches currently defined on the switch.
-fid <i>FID</i>	The FID must be defined in both the downloaded configuration file and the current system. You should disable a switch before downloading a configuration file. If you plan to download a configuration file while the switch is enabled, see Configuration Download without Disabling a Switch .
-fid <i>FID</i> -sfid <i>FID</i>	The FID must be defined on the switch, and the source FID must be defined in the downloaded configuration file.
-all	The number of switches or FIDs defined in the downloaded configuration file must match exactly the number of switches or FIDs currently defined on the switch. The switches must be disabled first. If they are not, the <code>configDownload</code> command will download the configuration for as many switches as possible until a non-disabled switch is found. If a non-disabled switch is found, the downloading process stops. Before running the <code>configDownload</code> command, verify if any switches must be disabled. If you are performing a configuration download due to a configuration error, run the <code>configDefault</code> command before running the <code>configDownload</code> command. See Configuration Download without Disabling a Switch for more information on non-disruptive configuration downloads. If Virtual Fabrics mode is enabled, the <code>chassisDisable</code> and <code>chassisEnable</code> commands disable all logical switches on the affected switch. This process bypasses the need to disable and enable each switch individually after the configuration download has completed. Non-Virtual Fabric configuration files downloaded to a Virtual Fabric system have a configuration applied only to the default switch. If multiple logical switches are created in a Virtual Fabric-enabled system, there may be problems if there are ports that belong to the default switch in a Virtual Fabric-disabled system, but are now assigned to logical switches in a Virtual Fabric-enabled system. Only configurations related to ports within the default switch are applied.

If you must set up your switch again, run the commands listed in table below and save the output in a file format. Store the files in a safe place for emergency reference.

Table 25: CLI Commands to Display or Modify Switch Configuration Information

Command	Displays
<code>configShow</code>	System configuration parameters, settings, and license information.
<code>fcLunQuery</code>	LUN IDs and LUNs for all accessible targets.
<code>fcrRouterPortCost</code>	FC Router route information.
<code>fcrXlateConfig</code>	Translate (xlate) domain's domain ID for both EX_Port-attached fabric and backbone fabric.
<code>fosConfig</code>	Fabric OS features.
<code>ipAddrShow</code>	IP address.
<code>isnscCfg</code>	Configuration state of the iSNS client operation.
<code>license --show</code>	License keys installed with more detail than the license information from the <code>configShow</code> command.
<code>portCfgEXPort</code>	EX_Port configuration parameters.

**CAUTION**

Though the switch itself has advanced error checking, the configdownload feature within Fabric OS was not designed for users to edit, and it is limited in its ability. Edited files can become corrupted, and this corruption can lead to switch failures.

NOTE

If you download a configuration file with either the `portCfgEportCredits` or `portCfgFportBuffers` entries manually removed, the switch will retain the existing configuration even after the configuration has downloaded.

Configuration Download without Disabling a Switch

You can download configuration files to a switch while the switch is enabled; that is, you do not need to disable the switch for changes in the SNMP, MAPS, AAA, Authutil, IPFilter, or ACL parameters. Configure the DH-CHAP or FCAP certificates for Authutil and the DNS for AAA using the FQDN on the switch before downloading. However, if there is any changed parameter that does not belong to SNMP, MAPS, AAA, Authutil, IPFilter, or ACL, you must disable the switch. When you use the `configDownload` command, you are prompted to disable the switch *only when necessary*.

ATTENTION

The configuration download process can restore only logical switches that already exist and that use the same FIDs. It cannot clone or repair the current switch because the `configDownload` command cannot create logical switches if they do not exist.

Restoring a Configuration

**CAUTION**

Using the SFID parameter erases all configuration information on the logical switch. Use the SFID parameter only when the logical switch has no configuration information that you want to save.

1. Verify that the FTP/SFTP/SCP service is running on the server where the backup configuration file is located.
2. Connect to the switch and log in using an account with admin permissions, and if necessary, with chassis permissions.
3. If any changed parameters are in the configuration file that do not belong to SNMP, Fabric Watch, or ACL, disable the switch by entering the `switchDisable` command.
4. Enter the `configDownload` command.

The command becomes interactive, and you are prompted for the required information.

5. At the `Do you want to continue [y/n]` prompt, enter **y**.

Wait for the configuration to be restored.

6. If you disabled the switch, enter the `switchEnable` command when the process is finished.

NOTE

Always perform a reboot after you download a configuration file. On dual-CP platforms, you must reboot both CPs simultaneously.

```
switch:admin> configdownload
Protocol (scp, ftp, sftp, local) [ftp]:sftp
Do you want to continue with CRA (Y/N) [N]: Y
SCP/SFTP Server Port Number [22]: 22
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: User1
Path/File name [<home dir>/config.txt]: a.txt
Section (all|chassis|FID# [all]): all
*** CAUTION ***

This command is used to download a backed-up configuration
for a specific switch. If using a file from a different
switch, this file's configuration settings will override
```

any current switch settings. Downloading a configuration file, which was uploaded from a different type of switch, may cause this switch to fail.

A switch reboot is required for the changes to take effect. Please reboot both the CPs in the chassis simultaneously after configurations are successfully downloaded.

Please make sure all the switches are disabled by using "chassisdisable" command. Downloading configuration to an online switch may result in some configuration not being downloaded to that switch.

Please reboot the system for the configuration to be effective

configDownload operation may take several minutes to complete for large files.

Do you want to continue [y/n]:y

Password: <hidden>

configDownload complete.

Example of a Non-Interactive Download of All Configurations (Chassis and Switches)

```
configdownload -a -ftp 10.1.2.3,UserFoo,/pub/configurations/config.txt,password
```

Configurations across a Fabric

To save time when configuring fabric parameters and software features, you can save a configuration file from one switch and download it to other switches of the same model type.

Do not download a configuration file from one switch to another switch that is a different switch type or runs a different firmware version, because it can cause the switch to fail. If you need to reset affected switches, issue the `configDefault` command after download is completed but before the switch is enabled. If a switch is enabled with a duplicate domain ID, the switch becomes segmented.

NOTE

The `configDefault` command is not recommended on the Blade Server SAN I/O Modules. For more information, refer to the *Brocade Fabric OS Command Reference Manual*.

Downloading a Configuration File from One Switch to Another Switch of the Same Switch Type

1. Configure one switch.
2. Use the `configUpload` command to save the configuration information. See [Configuration File Backup](#) for more information.
3. Run `configDefault` on each of the target switches, and then use the `configDownload` command to download the configuration file to each of the target switches. See [Configuration File Restoration](#) for more information.

Security Considerations

Security parameters and the switch identity cannot be changed by the `configDownload` command. Parameters, such as the switch name and IP address (lines in the configuration file that begin with "boot"), are ignored. Security parameters (lines in the configuration file that begin with "sec"), such as secure mode setting and version stamp, are ignored. For more detailed information on security, see [Configuring Protocols](#).

Configuration Management for Virtual Fabrics

You can use the `configUpload -vf` or `configDownload -vf` command to restore configurations to a logical switch. The `-vf` option only restores the virtual fabrics configuration information on to a switch of the same model and same release.

The virtual fabrics configuration on the switch defines all of the logical switches allowed and configured for a particular platform.

Uploading a Configuration File from a Switch with Virtual Fabrics Enabled

The `configUpload` command with the `-vf` option specifies that configuration upload will upload the Virtual Fabrics configuration instead of the non-Virtual Fabrics configuration information.

You must specify a file name with the `configUpload -vf` command. Do not use `config.txt` for a file name because this name can be confused with a normal uploaded configuration file.

Example of configUpload on a Switch with Virtual Fabrics

```
FID128:admin> configupload -vf
Protocol (scp, ftp, sftp, local) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: anonymous
Path/Filename [<home dir>/config.txt]: 720_vf.txt
configUpload complete: VF config parameters are uploaded
2014/07/20-09:13:40, [LOG-1000], 225, SLOT 7 | CHASSIS, INFO, Brocade720, Previous message repeated 7 time(s)
2014/07/20-10:27:14, [CONF-1001], 226, SLOT 7 | FID 128, INFO, 720_80, configUpload completed successfully for
VF config parameters.
```

Example of configUpload on a Logical Switch Configuration

```
Sprint5100:FID128:admin> configupload
Protocol (scp, ftp, sftp, local) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: UserFoo
Path/Filename [<home dir>/config.txt]: 5100.txt
Potentially remote file may get overwritten
Section (all|chassis|FID# [all]):
Password: <hidden>
configUpload complete: All selected config parameters are uploaded
```

Restoring a Logical Switch Configuration Using configDownload

The `configDownload -vf` command specifies that the Virtual Fabrics configuration download file is downloaded instead of the regular configuration. After the Virtual Fabrics configuration file is downloaded, the switch is automatically rebooted.

On dual-CP platforms, if CPs are incompatible (HA not in sync), the Virtual Fabrics configuration file is not propagated to the standby CP. If CPs are compatible, the active CP attempts to remain active after the reboot, and the new Virtual Fabrics configuration file is then propagated to the standby CP.

**CAUTION**

You must issue the `configDownload` command on the switch after restoring the Virtual Fabrics configuration to fully restore your switch or chassis configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `configDownload -vf` command.
3. Respond to the prompts.

Wait for the configuration file to download on to the switch. You may need to reconnect to the switch.

4. Enter the `configDownload -all` command.
5. Respond to the prompts.
6. Verify the logical inter-switch links (LISL) ports are set up correctly.

Wait for the configuration file to download to the switch.

Example of a Non-Interactive Download from a Switch with FID = 8 and SFID = 10

```
configdownload -fid 8 -sfid 10 -ftp 10.1.2.3
,UserFoo
,config.txt,password
```

Example of configDownload on a Switch

```
FID128:admin> configdownload -vf
Protocol (scp, ftp, sftp, local) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: UserFoo
Path/Filename [<home dir>/config.txt]: 720_FID89.txt
```

*** CAUTION ***

This command is used to download the VF configuration to the switch. Afterwards, the switch will be automatically rebooted and the new VF settings will be used. You will then need to run `configdownload` again to install the configuration(s) for any logical switch(s) that are setup in the new VF configuration.

Do you want to continue [y/n]: y

(output truncated)

Configuration Upload and Download Restrictions When Virtual Fabrics Is Enabled

NOTE

A reboot is required after downloading a configuration for all parameters to take effect.

The following restrictions apply when using the `configupload` or `configdownload` commands when Virtual Fabrics mode is enabled:

- The `-vf` option is incompatible with the `-fid`, `-sfid`, or `-all` options. Any attempt to combine it with any of the other three will cause the configuration upload or download operation to fail.
- You are not allowed to modify the Virtual Fabrics configuration file after it has been uploaded. Only minimal verification is done by the `configdownload` command to ensure it is compatible, much like the normal downloaded configuration file.
- After the `configdownload -vf` command completes and reboots your switch, you must then download the matching regular configuration using the `configdownload -all` command. This ensures proper behavior of the system and logical switches.

All of the attributes of the Virtual Fabrics configuration file will be downloaded to the system and take effect. This includes, but is not limited to, logical switch definitions, whether Virtual Fabrics is enabled or disabled, and the F_Port trunking ports, except the LISL ports. The LISL ports on the system are not affected by the Virtual Fabrics configuration file download.

You can restore Virtual Fabrics configurations only to a switch of the same model and same release. For example, a Virtual Fabrics configuration file for Fabric OS 8.2.x cannot be used on a Fabric OS 8.1.x switch and vice versa.

Brocade Configuration Form

Complete the second column in the following table, and save it as a hard-copy reference for your configuration information.

Table 26: Brocade Configuration and Connection Form

Basic Settings	Values that You Chose
IP address	
Gateway address	
Chassis configuration option	
Management Connections	Values that You Chose
Serial cable tag	
Ethernet cable tag	
Configuration Information	Values that You Chose
Domain ID	
Switch name	
Ethernet IP address	
Ethernet subnet mask	
Total number of local devices (<code>nsShow</code>)	
Total number of devices in fabric (<code>nsAllShow</code>)	
Total number of switches in the fabric (<code>fabricShow</code>)	

Managing User Accounts

In addition to the admin, maintenance, and user accounts, Fabric OS allows you to create up to 252 additional user accounts. These accounts expand your ability to track account access and audit administrative activities.

NOTE

Access to the root account is removed in 9.1.x and the account cannot be activated. References to the root may persist in FOS and should be ignored.

Each user account is associated with the following parameters:

- Permissions – Associate roles with each user account to determine the functional access levels.
- Virtual Fabric list – Specifies the Virtual Fabric that a user account is allowed to log in.
- Home Virtual Fabric – Specifies the Virtual Fabric that the user is logged in to, if available. The home Virtual Fabric must be a member of the user's Virtual Fabric list. If the home Virtual Fabric ID is not available, the user is logged in to the default switch if that is a member of the user's Virtual Fabric list. If the default switch is also not available, the valid lowest Virtual Fabric ID from the user's VF list is used.
- LF Permission List – Determines functional access levels within the bounds of the user's Virtual Fabrics.
- Chassis role – Similar to switch-level roles, but applies to a different subset of commands.

For more information about Virtual Fabrics, see [Managing Virtual Fabrics](#).

Fabric OS provides four options for authenticating users: remote RADIUS service, remote LDAP service, remote TACACS+ service, and the local-switch user database. All options allow users to be managed centrally by the following methods:

- Remote RADIUS service – Users are managed in a remote RADIUS server. All switches in the fabric can be configured to authenticate against the centralized remote database.
- Remote LDAP service – Users are managed in a remote LDAP server. All switches in the fabric can be configured to authenticate against the centralized remote database. The remote LDAP server can run Microsoft Active Directory or OpenLDAP.
- Remote TACACS+ service – Users are managed in a remote TACACS+ server. All switches in the fabric can be configured to authenticate against the centralized remote database.
- Local user database – Users are managed by the local user database. The local user database is manually synchronized by the `distribute` command to push a copy of the switch's local user database to all other switches in the fabric. If users with the user-defined roles exist on the sending switch or on any receiving switch, the `distribute` command is blocked.

Role Management

This section describes role-based access control, default rules, role permissions, and managing user-defined roles.

Role-Based Access Control

Role-Based Access Control (RBAC) specifies the permissions that a user account has on the basis of the role the account has been assigned. For each role, a set of predefined permissions determine the jobs and tasks that can be performed on a fabric and its associated fabric elements. Fabric OS uses RBAC to determine which command a user is allowed to access.

When you log in to a switch, your user account is associated with a predefined role or a user-defined role. The role that your account is associated with determines the level of access you have on that switch and in the fabric. The chassis role can also be associated with user-defined roles; it has permissions for RBAC classes of commands that are configured when user-defined roles are created. The chassis role is similar to a switch-level role, except that it affects a different subset of commands. You can use the `userConfig` command to add this permission to a user account.

Default Fabric OS Roles

The following table outlines the Fabric OS predefined (default) roles.

Table 27: Default Fabric OS Roles

Role Name	Duties	Description
Admin	All administration	All administrative commands
BasicSwitchAdmin	Restricted switch administration	Mostly monitoring with limited switch (local) commands
FabricAdmin	Fabric and switch administration	All switch and fabric commands, excluding user management commands
Operator	General switch administration	Routine switch-maintenance commands
SecurityAdmin	Security administration	All switch security and user management functions
SwitchAdmin	Local switch administration	Most switch (local) commands, excluding security, user management, and zoning commands
User	Monitoring only	Non-administrative use such as monitoring system activity
ZoneAdmin	Zone administration	Zone management commands only
Maintenance	Service support and debugging	Additional maintenance commands

NOTE

The keywords `smmsp`, `utmp`, `users`, `udrole`, and `nobody` must not be used to create roles.

Role Permissions

The following table describes the types of permissions that are assigned to roles.

Table 28: Permission Types

Abbreviation	Definition	Description
O	Observe	The user can run commands by using options that display information only, such as running <code>userconfig --show -a</code> to show all users on a switch.
M	Modify	The user can run commands by using options that create, change, and delete objects on the system, such as running the <code>userconfig --change</code> command with the <code>-r</code> option to change a user's role.
OM	Observe and Modify	The user can run commands by using both observe and modify options; if a role has modify permissions, it almost always has observe permissions.
N	None	The user is not allowed to run commands in a given category.

To view the permission type for categories of commands, use the `classconfig` command.

- Enter the `classconfig --show -classlist` command to list all command categories.
- Enter the `classconfig --showroles` command with the command category of interest as the argument to list all the roles.

This command shows the permissions that apply to all commands in a specific category.

```
switch:admin> classconfig --showroles authentication
Roles that have access to the RBAC Class 'authentication' are:
  Role name      Permission
  -----
  Admin          OM
  Root           OM
  Security Admin OM
  Maintenance    OM
```

NOTE

Access to the root account is removed from Fabric OS 9.1.0 release and the account cannot be activated. References to the root may persist in FOS and should be ignored.

You can also use the `classconfig --showcli` command to show the permissions that apply to a specific command.

Managing User-Defined Roles

Fabric OS provides an extensive toolset for managing user-defined roles:

- The `roleconfig` command defines new roles, deleting created roles, or viewing information about user-defined roles.
- The `classconfig` command displays RBAC information about each category or class of commands, and includes an option to show all roles associated with a given RBAC command category.
- The `userconfig` command assigns a user-defined role to a user account.



CAUTION

You can add a maximum of only 64 user-defined roles for Fabric OS and other management applications, such as SANnav, Brocade Network Advisor, and Web Tools.

NOTE

The command passed as arguments in `classconfig --showcli` must use only lowercase.

Creating a User-Defined Role

You can define a role as long as it has a unique name that is not the same as any of the Fabric OS default roles, any other user-defined role, or any existing user account name.

The following conditions also apply:

- A role name is case-insensitive and contains only letters.
- The role name should have a minimum of four letters and can be up to 16 letters long.
- The maximum number of user-defined roles that are allowed across all platforms is 64.

The `roleconfig` command defines unique roles. You must have chassis-level access and permissions to execute this command. The following example creates a user-defined role called `mysecurityrole`. The RBAC class `Security` is added to the role, and the `Observe` permission is assigned:

```
switch:admin> roleconfig --add mysecurityrole -class security -perm O
```

Role added successfully

The assigned permissions can be no higher than the admin role permission assigned to the class. The admin role permission for the Security class is Observe/Modify. Therefore, the Observe permission is valid.

The `roleconfig --show` command is available to view the permissions assigned to a user-defined role. You can also use the `classconfig --showroles` command to see that the role was indeed added with Observe permission for the security commands.

```
switch:admin> classconfig --showroles security
Roles that have access to RBAC Class 'security' are:
```

Role Name	Permissions
-----	-----
User	O
Admin	OM
Root	OM
SwitchAdmin	O
FabricAdmin	OM
BasicSwitchAdmin	O
SecurityAdmin	OM
Maintenance	OM

To delete a user-defined role, use the `roleconfig --delete` command.

NOTE

Note: Access to the root account is removed from Fabric OS 9.1.0 release and the account cannot be activated. References to the root may persist in FOS and should be ignored.

Assigning a User-Defined Role to a User

The `userconfig` command allows you to assign a user-defined role to a user.

To assign a user-defined role to a user, complete the following steps.

1. Connect to the device and log in using an account with admin permissions.
2. You have multiple options for assigning a user-defined role to a user:
 - To create a new user account and assign a role: `userconfig --add user_account -r role_name .`
 - To change a user-defined role or add a new one to an existing user account: `userconfig --change user_account -r role_name .`
 - To create a new user account and assign a chassis role: `userconfig --add user_account -c chassis_role_name .`
 - To add a chassis role to an account: `userconfig --change user_account -c chassis_role_name .`

The commands can be combined; the following example assigns the `mysecurityrole` role to the existing `anewuser` account and also adds the admin chassis role.

```
switch:admin> userconfig --change anewuser -r mysecurityrole -c admin
```

For more information on `userconfig` command options, refer to the *Brocade Fabric OS Command Reference Manual*.

Maximum Session Login Limits

The management channel is the communication established between the management workstation and the switch. The following table shows the number of simultaneous login sessions allowed for each user when authenticated locally. The roles are displayed in alphabetic order, which does not reflect their importance.

Table 29: Maximum Number of Simultaneous Sessions

User Role	Maximum Sessions
Admin	4
BasicSwitchAdmin	4
FabricAdmin	4
Operator	4
SecurityAdmin	4
SwitchAdmin	4
User	4
ZoneAdmin	4
Maintenance	4

NOTE

The total number of sessions on a switch cannot exceed 32.

NOTE

If an account has reached maximum sessions allowed and a command is entered as a single line using a SSH login, then the command will be blocked and the following message will be displayed: `Max remote sessions for <user> has been reached.`

Local Database User Accounts

User `add`, `change`, and `delete` operations are subject to the *subset* rule: an admin with LList 1 to 10 cannot perform operations on an admin, user, or any role with LList 11 to 128. The user account being changed must have an LList that is a subset of the account that is making the change.

In addition to the default administrative and user accounts, Fabric OS supports up to 252 user-defined accounts in each switch (domain). These accounts expand your ability to track account access and audit administrative activities.

Default Accounts

The following table lists the predefined accounts offered by Fabric OS that are available in the local-switch user database. The password for all default accounts should be changed during the initial installation and configuration of each switch.

Table 30: Default Local User Accounts

Account Name	Role	Logical Fabric	Default Status	Description
admin	Admin	LF1 to 128 home: 128	Enabled	Most commands have <i>observe/modify</i> permission.
user	User	LF: 128 home: 128	Enabled	Most commands have <i>observe-only</i> permission.

Account Name	Role	Logical Fabric	Default Status	Description
maintenance	Maintenance	LF1 to 128 home: 128	Disabled	This is mainly reserved for service support and debugging.

NOTE

When using Virtual Fabrics, administrators can act on other accounts only if that account has a Logical Fabric list that is a subset of the administrator.

NOTE

A remote or user defined account can manage a default account if they share the same switch specific role, chassis role, and have access to the same VF list, which are the attributes that can be modified.

```
-e [account enable/disable]
-at[access-time]
-x [password expiration]
-u [unlock]
```

For all users, an asterisk (*) is displayed next to the user name if the password is still a default password when you run `userconfig -show` command.

Displaying Account Information

1. Connect to the switch and log in using an account with admin permissions, or an account associated with a user-defined role with permissions for the UserManagement class of commands.
2. Enter the appropriate `show` operands for the account information you want to display:

- `userconfig --show -a` to show all account information for a switch.

NOTE

A * is displayed next to the user names that still have the default password set.

- `userconfig --show username` to show account information for the specified account.
- `userconfig --showlf -l logicalFabric_ID` for each LF in an LF_ID_list, which displays a list of users that include that LF in their LF permissions.

Managing a User-Defined Account

This section describes user-defined account management.

Creating an Account

1. Connect to the switch and log in using an account with admin permissions or an account associated with a user-defined role with permissions for the UserManagement class of commands.
2. Enter the `userconfig --add` command.

```
> userconfig --add user1 -l 1-128 -h 128 -r admin -c admin
```

This example creates a user account for a user with the following properties:

- Access to virtual fabrics 1 through 128
- Default home logical switch set to 128
- Admin role permissions
- Admin chassis role permissions

3. In response to the prompt, enter a password for the account.

The password is not displayed when you enter it interactively.

Deleting an Account

This procedure can be performed on local user accounts.

1. Connect to the switch and log in using an account with admin permissions or an account associated with a user-defined role with permissions for the UserManagement class of commands.
2. Enter the `userconfig --delete` command.
You cannot delete the default accounts. An account cannot delete itself. All active CLI sessions for the deleted account are logged out.
3. At the prompt for confirmation, enter `y`.

Changing Account Parameters or LF Permissions

This procedure can be performed on local user accounts.

1. Connect to the switch and log in using an account with admin permissions or an account associated with a user-defined role with permissions for the UserManagement class of commands.
2. Enter the `userconfig --change` command to change account parameters, or enter the `userconfig --addlf` or `userconfig --deletelf` command to add or delete LF permissions.
For more information on `userconfig` subparameters, refer to *Brocade Fabric OS Command Reference Manual*.

Local Account Passwords

The following rules apply to changing passwords:

- When changing your own password, you must provide the current password.
- To change the password for another account requires admin permissions or an account associated with a user-defined role with Modify permissions for the UserManagement RBAC class of commands.
- The user account being changed must have an VList that is a subset of the account that is making the change.
- A new password must have at least one character different from the previous password. However, with the default password policy, this rule can be changed.
- You cannot change passwords using SNMP.

Changing the Password for the Current Login Account

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `passwd` command.
3. Enter the requested information at the prompts.

The following example shows the output for changing password using the `passwd` command.

```
sw0 login: admin

Password:####
sw0:FID128:admin> passwd
Changing password for admin
Enter old password: ####
Enter new password: ####
Re-type new password: ####
```

```
passwd: all authentication tokens updated successfully
Saving password to stable storage.
Password saved to stable storage successfully.(output truncated)
```

Changing the Password for a Different Account

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `passwd` command specifying the name of the account for which the password is being changed.
For example, `passwd username`.

NOTE

The super user can also update tokens for named users using the `-old` and `-new` password options.

3. Enter the requested information at the prompts.

Password Policies

The password policies described in this section apply to the local-switch user database only. Configured password policies (and all user account attribute and password state information) are synchronized across CPs and remain unchanged after an HA failover. Password policies can also be manually distributed across the fabric. For more information, see [Local User Account Database Distribution](#).

NOTE

To recover any password, contact your switch provider or Brocade Customer Support.

Password Strength Policy

The password strength policy is enforced across all user accounts, and enforces a set of format rules to which new passwords must adhere. The password strength policy is enforced only when a new password is defined. The total of the other password strength policy parameters (lowercase, uppercase, digits, and punctuation) must be less than or equal to the value of the `MinLength` parameter.

Use the following attributes to the `passwdCfg` command to set the password strength policy:

- **Lowercase**
Specifies the minimum number of lowercase alphabetic characters that must appear in the password. The default value is zero. The maximum value must be less than or equal to the `MinLength` value.
- **Uppercase**
Specifies the minimum number of uppercase alphabetic characters that must appear in the password. The default value is zero. The maximum value must be less than or equal to the `MinLength` value.
- **Digits**
Specifies the minimum number of numeric digits that must appear in the password. The default value is zero. The maximum value must be less than or equal to the `MinLength` value.
- **Punctuation**
Specifies the minimum number of punctuation characters that must appear in the password. All printable, non-alphanumeric punctuation characters except the colon (:) are allowed. The default value is zero. The maximum value must be less than or equal to the `MinLength` value.
- **MinLength**
Specifies the minimum length of the password. The minimum can be from 8 through 40 characters. New passwords must be between the minimum length specified and 40 characters. The default value is 8. The maximum value must be greater than or equal to the `MinLength` value.
- **CharSet**

Specifies the minimum length of the character set. Character set includes alphabets and special characters. The default value is 0. For example, if the `-charset` option is set to 3, the password should contain at least 3 characters (can be 2 alphabets and 1 special character, or 3 alphabets or any other combination). However, there can be more than three characters and include digits as part of the password. But the password cannot have less than three characters.

- **AllowUser**

Specifies whether to allow the user name to be included in the password either in the forward or reverse order of characters. The default value is set to "yes" for "-allowusername" option. If the option is set to "no", the user name both in forward and reverse direction is not allowed to be part of the password. The user names are validated as case sensitive.

- **Repeat**

Specifies the length of repeated character sequences that will be disallowed. For example, if the "repeat" value is set to 3, a "passAAAword" password is disallowed because it contains the repeated sequence "AAA". A password of "passAAword" would be allowed because no repeated character sequence exceeds two characters. The range of allowed values is from 1 through 40. The default value is 1. When set to 1, character repetition check is not enforced.

- **Sequence**

Specifies the length of sequential character sequences that will be disallowed. A sequential character sequence is defined as a character sequence in which the ASCII value of each contiguous character differs by one. The ASCII value for the characters in the sequence must all be increasing or decreasing. For example, if the "sequence" value is set to 3, a password "passABCword" is disallowed because it contains the sequence "ABC". A password of "passABword" would be allowed because it contains no sequential character sequence exceeding two characters. The range of allowed values is from 1 through 40. The default value is 1. When set to 1, sequential characters are not enforced.

- **Reverse**

Activates or deactivates the validation check to determine whether the password is an exact reverse string of the user name. This option is disabled by default.

Example of a Password Strength Policy

The following example shows a password strength policy that requires passwords to contain at least 3 uppercase characters, 4 lowercase characters, and 2 numeric digits; the minimum length of the password is 9 characters. The password cannot be an exact reverse string of the username.

```
switch:admin> passwdcfg --set -uppercase 3 -lowercase 4 -digits 2 -minlength 9 -reverse 1
```

Password History Policy

The password history policy prevents reuse of a recently used password. The password history policy is enforced across all user accounts when users are setting their own password or when an administrator sets the password for another administrator.

The password history policy is enhanced to prevent users from configuring a new password that is similar to the current password. A password similarity check is applied using the `-minDiff` option to measure the difference between the new and current password. This check ensures that the new password is not a simple replacement of some characters in the current password.

Use the following attributes of the `passwdCfg` command to set the password history policy:

- **History**

This option specifies the number of past password values that are disallowed when setting a new password. The allowable password history values range from 0 through 24. If the value is set to 0, the new password cannot be set to the current password, but it can be set to the most recent password. The default value is 1, which means that the

current password and one previous password cannot be reused. The value 2 means that the current password and the two previous passwords cannot be used (and so on, up to 24 passwords).

This option does not verify that a new password meets a minimal standard of difference from prior passwords; rather, it only determines whether a newly specified password is identical to one of the specified number (1 to 24) of previously used passwords.

- **MinDiff**

This option defines whether the new password is acceptable by measuring how similar it is to the user's current password. The password policy determines the minimum number of changes required (where a change may be inserting, deleting, or replacing a character) to transform from one string into another. This policy is used to prevent users from making only minor changes to their current password when setting a new password.

The configuration range is between 0 to 40 and must always be less than the configured `-minlength` value. The default value is 0, and if it is set to 0, the password similarity check is off.

The following are the requirements for using the MinDiff parameter in the password history policy:

- The MinDiff configuration value denotes the *minimum number of differences* expected between the current password and the new password.
- The password similarity check passes if the calculated difference between the current password and the new password is greater than or equal to the configured MinDiff value.
- The password similarity check fails if the calculated difference between the current password and the new password is less than the configured MinDiff value.
- The MinDiff is not supported during the distribution of the password policy.
- The MinDiff policy is enforced only when the user is logged in and changes the user's own password. The administrator changing a password of a user is exempted from this policy.

The following table outlines the different combinations between the current password and the new password and their results.

Table 31: Different MinDiff Combinations between the Current Password and New Password

Current Password	New Password	Configured MinDiff Value	Measured Value of Character Changes	Result	How It Is Calculated
justification	jurisdiction	4	5	Pass	Replace <i>s</i> with <i>r</i> , replace <i>t</i> with <i>i</i> , replace <i>i</i> with <i>s</i> , replace <i>f</i> with <i>d</i> and delete <i>a</i> in the middle.
required	acquired	4	2	Fail	Replace <i>r</i> with <i>a</i> in the front and replace <i>e</i> with <i>c</i> .
password	pass123word	3	3	Pass	Insert <i>123</i> in the middle.
manifesto	maniFESTo	5	4	Fail	Replace <i>f</i> with <i>F</i> , replace <i>e</i> with <i>E</i> , replace <i>s</i> with <i>S</i> , and replace <i>t</i> with <i>T</i> .

Example of a Password History Policy

The following example shows a password policy that requires passwords to contain at least three uppercase characters, and four lowercase characters; the minimum length of the password is 9 characters. The measurement value difference between the current and the new password must be five characters.

```
switch:admin> passwdcfg --set -uppercase 3 -lowercase 4 -minlength 9 -mindiff 5
```

Password Expiration Policy

The password expiration policy forces the expiration of a password after a configurable period of time. The expiration policy can be enforced across all user accounts or on specified users only. When the user logs in, a warning message appears before the password is expired. When a password expires, the user must change the password to complete the

authentication process and open a user session. You can specify the number of days prior to password expiration during which warnings will commence. Password expiration does not disable or lock out the account.

Use the following attributes to the `passwdCfg` command to set the password expiration policy:

- **MinPasswordAge**
Specifies the minimum number of days that must elapse before a user can change a password. MinPasswordAge values range from 0 through 999. The default value is zero. Setting this parameter to a nonzero value discourages users from rapidly changing a password to circumvent the password history setting to select a recently used password. The MinPasswordAge policy is not enforced when an administrator changes the password for another user.
- **MaxPasswordAge**
Specifies the maximum number of days that can elapse before a password must be changed, and it is also known as the password expiration period. MaxPasswordAge values range from 0 through 999. The default value is zero. Setting this parameter to zero disables password expiration.
- **Warning**
Specifies the number of days prior to password expiration that a warning about password expiration is displayed. Warning values range from 0 through 999. The default value is 0 days.

NOTE

When MaxPasswordAge is set to a nonzero value, MinPasswordAge and Warning must be set to a value that is less than or equal to MaxPasswordAge.

Example of Password Expiration Policies

The following example configures a password expiration policy for the user1 user account. This user must change the password within 90 days of setting the current password and no sooner than 10 days after setting the current password. The user will start to receive warning messages 3 days before the 90-day limit, if the password is not already changed.

```
switch:admin> passwdcfg --setuser user1 -minpasswordage 10 -maxpasswordage 90 -warning 3
```

The following example configures a password expiration policy for all users.

```
switch:admin> passwdcfg --set -minpasswordage 5 -maxpasswordage 30 -warning 5
```

Expiring Passwords Immediately

You can expire the password immediately for a specific user account or all user accounts using the `passwdcfg` command. The users will be forced to configure a new password at the next login. Expiring the passwords immediately does not change the MaxPasswordAge, MinPasswordAge, and Warning configurations.

1. To expire the password immediately for all users, use the following command.

```
Switch#admin> passwdcfg --set -expire
```

2. To expire the password immediately for a specific user, use the following command.

```
Switch#admin> passwdcfg --setuser user1 -expire
```

Account Lockout Policy

The account lockout policy disables a user account when that user exceeds a specified number of failed login attempts, and it is enforced across all the user accounts. You can configure this policy to keep the account locked until explicit administrative action is taken to unlock it, or the locked account can be automatically unlocked after a specified period. Administrators can unlock a locked account at any time.

A failed login attempt counter is maintained for each user on each switch instance. The counters for all user accounts are reset to zero when the account lockout policy is enabled. The counter for an individual account is reset to zero when the account is unlocked after a lockout duration period expires, or when the account user logs in successfully.

The admin account can also have the lockout policy enabled on it. The admin account lockout policy is disabled by default and uses the same lockout threshold and duration as the other accounts. It can be automatically unlocked after the lockout duration passes or when it is manually unlocked by either a user account that has a securityAdmin or other admin permissions.

Virtual Fabrics considerations: The home logical fabric context validates user enforcement for the account lockout policy.

Note that the account-locked state is distinct from the account-disabled state.

Use the following attributes to set the account lockout policy:

- **LockoutThreshold**
Specifies the number of times a user can attempt to log in using an incorrect password before the account is locked. The number of failed login attempts is counted from the last successful login. LockoutThreshold values range from 0 through 999, and the default value is 0. Setting the value to 0 disables the lockout mechanism.
- **LockoutDuration**
Specifies the time, in minutes, after which a previously locked account is automatically unlocked. LockoutDuration values range from 0 through 99999, and the default value is 30. Setting the value to 0 disables lockout duration, and requires a user to seek administrative action to unlock the account. The lockout duration begins with the first login attempt after the LockoutThreshold has been reached. Subsequent failed login attempts do not extend the lockout period.

Enabling the Admin Lockout Policy

1. Log in to the switch using an account that has admin or securityAdmin permissions.
2. Enter the `passwdCfg --enableadminlockout` command.

Unlocking an Account

1. Log in to the switch using an account that has admin or securityAdmin permissions.
2. Enter the `userConfig --change account_name -u` command, specifying the `-u` option to unlock the account.

Disabling the Admin Lockout Policy

1. Log in to the switch using an account that has admin or securityAdmin permissions.
2. Enter the `passwdCfg --disableadminlockout` command.

Denial of Service Implications

The account lockout mechanism creates a denial of service condition when a user repeatedly attempts to log in to an account by using an incorrect password. Selected privileged accounts are exempted from the account lockout policy to prevent users from being locked out from a denial of service attack. However, these privileged accounts may then become the target of password-guessing attacks. Examine audit logs to monitor if such attacks are attempted.

Configuring the Password Hash Type

With Fabric OS 8.0.1 and later, SHA256 and SHA512 hash types are also supported. Whenever you change the hash type, you will be prompted to change the password in the next login. You can also use the `manual` option that does

not trigger any password change enforcement and expects you to manually execute `passwd` command to change the password after the hash type configuration.

After a new password hash configuration has been set, by default, you are forced to change the password in the next login. When the `-manual` option is used, you can change the password using the `passwd` command.

To change the password hash type, complete the following steps:

1. Log in to the device using the username and password.
2. Enter the `passwdcfg --hash {md5|sha256|sha512} [-manual]` command to configure the hash type for all user accounts.
3. Do one of the following:
 - If you do not use the `-manual` option, log out the user session and at the time of re-login, enter the new password and confirm the same.
 - If you use the `-manual` option, explicitly change the passwords for the account using `passwd` command.
4. Use the following commands to display the password hash type setting of the switch, current password hash type of specified user or current password hash type of all users configured in the device.
 - Use the `passwdcfg --showhash` command to display the configured hash type on the switch.
 - Use the `passwdcfg --showhash <username>` command to display the current password hash for the username.
 - Use the `passwdcfg --showhash -all` to display the current password hash for all the users on the switch.

Sample output:

```
switch(mode)# passwdcfg --hash sha256 -manual
switch(mode)# passwdcfg -showhash -all
<Command output>
```

Local User Account Database Distribution

Fabric OS allows you to distribute the user database and passwords to other switches in the fabric. When the switch accepts a distributed user database, it replaces the local user database with the user database it receives.

By default, switches accept the user databases and passwords distributed from other switches. The "Locked" status of a user account is not distributed as part of local user database distribution.

When the user database is distributed, it may be rejected by a switch for one of the following reasons:

- One of the target switches does not support local account database distribution.
- One of the target switch's user databases is protected.
- One of the remote switches has logical switches defined.
- Either the local switch or one of the remote switches has user accounts associated with user-defined roles.

Distributing the Local User Database

When the local user database is distributed, all user-defined accounts residing in the receiving switches are logged out of any active sessions.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `distribute -p PWD -d <domain id>` command.

NOTE

- If virtual fabrics mode is enabled and logical switches are defined other than the default logical switch, distributing the password database to switches is not supported.
- The `distribute` command is not supported in VF mode with more than one logical switch. The database distribution is blocked if the initiating switch has VF enabled and the receiving switch has more than one logical switch enabled.
- When a password database update is received, the hash type is not reset if the received hash configuration is weaker than the existing hash configuration. For example, if the switch is currently using SHA512, and the configuration request is for SHA256 using the `distribute` command, then the hash type is retained as SHA512.
- If the receiving switch is running pre-9.0 firmware, and if `admin` account is disabled, then the distribution is blocked.

Accepting Distributed User Databases on the Local Switch

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `fddCfg --localaccept PWD` command.

Rejecting Distributed User Databases on the Local Switch

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `fddCfg --localreject PWD` command.

Remote User

Fabric OS supports remote users through one of the following external Authentication, Authorization, and Accounting (AAA) services:

- Remote Authentication Dial-In User Service (RADIUS)
- Lightweight Directory Access Protocol (LDAP).
- Terminal Access Controller Access-Control System Plus (TACACS+)

The following table shows the latest supported OS version for the AAA server:

Table 32: Supported Software Versions of the AAA Server

AAA Server	Windows Version	Linux Version	RSA 6.x
RADIUS	Windows 2012	FreeRADIUS	steelbelt
LDAP	AD 2012: Schema Version 69 AD 2016: Schema Version 87	OpenLDAP	N/A
TACACS+	N/A	tacplus	N/A

Overview

A switch can be configured to try one of the supported remote authentication services (RADIUS, LDAP, or TACACS+) and local switch authentication. The switch can also be configured to use only a remote authentication service or only local switch authentication.

Client/Server Model

When configured to use one of the supported remote authentication services, the switch acts as a Network Access Server (NAS) and RADIUS, LDAP, or TACACS+ client. The switch sends all authentication, authorization, and accounting (AAA) service requests to the authentication server. The authentication server receives the request, validates the request, and sends its response back to the switch.

The supported management access channels that integrate with RADIUS, LDAP, and TACACS+ include serial port, Telnet, SSH, Web Tools, and REST API interface. All these access channels require the switch IP address or name to connect except serial port connection. RADIUS, LDAP, and TACACS+ servers accept both IPv4 and IPv6 address formats. For accessing both the active and standby CP blades, and for the purpose of HA failover, both CP IP addresses of a director should be included in the authentication server configuration.

NOTE

For systems, such as a Brocade director, the switch IP addresses are aliases of the physical Ethernet interfaces on the CP blades. If both IPv4 and IPv6 addresses are provisioned for the switch, both the addresses need to be specified while adding client IP address on the AAA server

Authentication Server Data

When configured for remote authentication, a switch becomes a RADIUS, LDAP, or TACACS+ client. In any of these configurations, authentication records are stored in the authentication host server database. Login and logout account name, assigned permissions, and time-accounting records are also stored on the authentication server for each user.

Switch Configuration

The authentication services uses the switch's local database as the remote authentication services are disabled by default. To enable remote authentication, you should access the CLI through an SSH connection so that the shared secret is protected. Multiple login sessions can configure simultaneously, and the last session to apply a change leaves its configuration in effect. After a configuration is applied, it persists after a reboot or an HA failover. To enable the secure LDAP service, you can also install a certificate from the Microsoft Active Directory server or the OpenLDAP server. By default, the LDAP service does not require certificates.

The configuration applies to all switches. On a Director, the configuration replicates itself on a standby CP blade if one is present. It is saved in a configuration upload and applied in a configuration download.

You should configure at least two authentication servers, so that if one fails, the other will assume service. Up to five servers are supported. You can set the configuration with any one of the supported authentication services and local authentication enabled, so that if the authentication servers do not respond because of a power failure or network problems, the switch uses local authentication.

From Fabric OS 8.1.0, RADIUS server authentication is supported by installing a CA certificate for the RADIUS server on the switch.

Consider the effects of the use of a remote authentication service on other Fabric OS features. For example, when a remote authentication service is enabled, all account passwords must be managed on the authentication server. The Fabric OS mechanisms for changing switch passwords remain functional; however, such changes affect only the involved switches locally. They do not propagate to the authentication server, nor do they affect any account on the authentication server. Authentication servers also support notifying users of expiring passwords.

When RADIUS, LDAP, or TACACS+ is set up for a fabric that contains a mix of switches with and without RADIUS, an LDAP, and a TACACS+ support, the way a switch authenticates users depends on whether a RADIUS, an LDAP, or a TACACS+ server is set up for that switch. For a switch with remote authentication support and configuration, authentication bypasses the local password database. For a switch without remote authentication support or configuration, authentication uses the switch's local account names and passwords.

Remote User Configuration on the Switch

At least one RADIUS, LDAP, or TACACS+ server must be configured before you can enable the AAA service. You can configure the AAA service even if it is disabled on the switch. You can configure up to five AAA services. You must be logged in as admin or switchAdmin to configure the remote authentication service.

NOTE

On X7 and X6 Directors, the switch sends its RADIUS, LDAP, or TACACS+ request using the IP address of the active CP. If both IPv4 and IPv6 addresses are provisioned for the switch, both the addresses need to be specified while adding the client IP address on the AAA server.

NOTE

LDAP authentication in BASIC and STRICT crypto modes of operation requires a valid DNS server configuration to support hostname validation. DNS is required to be configured for forward as well as reverse lookups. In BASIC mode, authentication will not fail even if a hostname mismatch error is found. In STRICT mode, authentication will fail if a hostname mismatch error occurs. DNS configuration on the switch is not required if LDAP is configured using the server's IP address. If this configuration is used, the server's certificate must be configured with the server's IP address in the Common Name (CN) field. If configured for FQDN, it is mandatory to have the FQDN in the server's certificate in the Common Name (CN) and/or Subject Alternative Name (SAN) field.

NOTE

The existing hostname validation is performed first irrespective of the LDAP server configured using the IP address or FQDN. If the hostname validation fails and the server is configured using the IP address, the validation falls back to the IP address. IP address validation validates the configured server address (IP) against the received certificate's `ipAddress` field under the Subject Alternative Name (SAN). If one of the checks passes, server validation succeeds.

RADIUS, LDAP, or TACACS+ configuration is chassis-based configuration data. On platforms that contain multiple switch instances, the configuration applies to all instances. The configuration is persistent across reboots and firmware downloads. On a chassis-based system, the command must replicate the configuration to the standby CP.

Multiple login sessions can invoke the `aaaConfig` command simultaneously. The last session that applies the change is the one whose configuration is in effect. This configuration is persistent after an HA failover.

The AAA servers are contacted in the order they are listed, starting from the top of the list and moving to the bottom.

Adding an AAA Server to the Switch Configuration

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aaaConfig --add` command.

At least one AAA server must be configured before you can enable the RADIUS, LDAP, or TACACS+ service.

If no RADIUS, LDAP, or TACACS+ configuration exists, turning on the authentication mode triggers an error message. When the command succeeds, the event log indicates that the configuration is enabled or disabled.

NOTE

When the RADIUS FQDN is configured with PEAP-MSCHAPv2 authentication, RADIUS client communicates first with the IPv4 address of the FQDN(if exists), and if the RADIUS server does not respond, IPv6(if exists)

is attempted. This will help to provide a seamless migration for FQDN PEAP-MsCHAPv2, as PEAP was only supported on IPv4 until 9.0 release.

Deleting an Authentication Server from the Configuration

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aaaConfig --remove` command.

When the command succeeds, the event log indicates that the server is removed.

Changing an Authentication Server Configuration

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aaaConfig --change` command.

Changing the Order in Which Authentication Servers Are Contacted for Service

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aaaConfig --move` command.

When the command succeeds, the event log indicates that a server configuration is changed.

Displaying the Current Authentication Configuration

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aaaConfig --show` command.

If a configuration exists, its parameters are displayed. If the RADIUS, LDAP, or TACACS+ services are not configured, only the parameter heading line is displayed. The following table shows the parameters.

Table 33: Current Authentication Server Configuration Parameters

Parameter	Description
Position	The order in which servers are contacted to provide service.
Server	The server names or IPv4 or IPv6 addresses.
Port	The server ports.
Secret	The shared secrets.
Timeouts	The length of time servers needs to respond before the next server is contacted.
Domain	The domain name of the LDAP server.
Authentication	The type of authentication being used on servers.
Encryption Level	The encryption algorithm of the stored sharedsecret value.
TLS Mode	The TLS mode for connection establishment with the LDAP server.

Setting the Switch Authentication Mode

This section describes the steps to set the authentication mode.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aaaConfig --authspec` command.

The following table outlines the `aaaConfig` command options used to set the authentication mode.

Table 34: Authentication Configuration Options

aaaConfig Options	Description
<code>--authspec "local"</code>	Default setting. This option authenticates management connections against the local database only. If the password does not match or the user is not defined or not authorized, the login fails.
<code>--authspec "radius"</code>	This option authenticates management connections against any RADIUS databases only. If the credentials do not match, the login fails. If the RADIUS service is not available, the authentication will not fall back to the AAA server if configured.
<code>--authspec "radius;local"</code>	This option authenticates management connections to the first reachable RADIUS database. If RADIUS fails for any reason, <code>--authspec "radius;local"</code> option authenticates against the local user database.
<code>--authspec "radius;local" -backup</code>	This option authenticates management connections against the first available RADIUS database. If RADIUS fails because the service is not available, it then authenticates against the local user database. The <code>-backup</code> option directs the service to try the secondary authentication database only if the primary authentication database is not available. For example, if user1 has entered a wrong password on the primary authentication database and the secondary authentication database has the correct password for user1, authentication does not query the secondary database until the primary database is reachable. Only when the primary database is not reachable, the secondary database is queried.
<code>--authspec "ldap"</code>	This option authenticates management connections against any LDAP databases only. If the LDAP service is not available or the credentials do not match, the login fails.
<code>--authspec "ldap;local"</code>	This option authenticates management connections against any LDAP databases first. If LDAP fails for any reason, it then authenticates against the local user database.
<code>--authspec "ldap;local" -backup</code>	This option authenticates management connections against any LDAP databases first. If LDAP fails for any reason, it then authenticates against the local user database. The <code>--backup</code> option tries the secondary authentication database only if the primary authentication database is not available.
<code>--authspec "tacacs+"</code>	This option authenticates management connections against any TACACS+ databases only. If the TACACS+ service is not available or the credentials do not match, the login fails.

aaaConfig Options	Description
<code>--authspec "tacacs+;local"</code>	This option authenticates management connections against any TACACS+ databases first. If TACACS+ fails for any reason, it then authenticates against the local user database.
<code>--authspec "tacacs+;local" -backup</code>	This option authenticates management connections against any TACACS+ databases first. If TACACS+ fails for any reason, it then authenticates against the local user database. The <code>--backup</code> tries the secondary authentication database only if the primary authentication database is not available.
<code>--authspec -nologout</code>	This option prevents users from being logged out when you change authentication. The default behavior is to log out users when you change authentication.
<code>--authspec -logpriauth [yes/no]</code>	This option suppresses log messages for authentication failure by the primary AAA service if authentication is to be done using the secondary AAA service, which is the local switch database. The default value is set to "Yes" and always displays login failure messages from the primary AAA service. The <code>--authspec -logpriauth [yes/no]</code> option is allowed for LDAP, RADIUS, and TACACS+.

NOTE

The `aaaconfig --authspec` command to authenticate only remote users may take more time to complete depending on number of local users configured on the switch.

Remote Server Configuration

RADIUS, LDAP, and TACACS+ servers allow you to set up user accounts by their true network-wide identities rather than by the account names created on a Fabric OS switch. With each account name, assign the appropriate switch access permissions. For LDAP servers, you can use the `ldapCfg --maprole` and `--mapattr` commands to map LDAP server permissions.

NOTE

Non-VF switches are required to have the role mapped to an LDAP group in Fabric OS 9.0.0 or later. If not, authentication will fail regardless of how AD is configured.

Users must enter their assigned RADIUS, LDAP, or TACACS+ account name and password when logging in to a switch that has been configured with remote authentication. After the remote authentication (RADIUS, LDAP, or TACACS+) server authenticates a user, it responds with the assigned switch role, chassis role, and LF permission in a *Brocade Vendor-Specific Attribute* (VSA). In case of RADIUS and TACACS+, if the response does not have the VSA permissions then the user role is assigned with the default LF permission or the authentication is rejected. In case of LDAP, if the response does not have the VSA permissions and if there is no role mapping performed on the switch, then the authentication is rejected.

NOTE

The role of the remote user is not applicable in non-VF mode.

RADIUS Server Data Configuration

This section describes FOS user configuration on the RADIUS server. It also describes RADIUS configuration with the virtual fabrics and the password expiry date.

The table below describes the syntax used for assigning VSA-based account switch roles on a RADIUS server.

Table 35: Syntax for VSA-based Account Roles on the RADIUS Server

Item	Value	Description
Type	26	1 octet
Length	7 or higher	1 octet, calculated by the server
Vendor ID	1588	4 octet, Brocade SMI Private Enterprise Code
Vendor type	1	1 octet, Brocade-Auth-Role; valid attributes for the Brocade-Auth-Role are: Admin BasicSwitchAdmin FabricAdmin Operator SecurityAdmin SwitchAdmin User ZoneAdmin User-defined role
	2	<i>Optional:</i> Specifies the Virtual Fabric member list. For more information on Virtual Fabrics, see RADIUS Configuration with Virtual Fabrics . Brocade-AVPairs1
	3	Brocade-AVPairs2
	4	Brocade-AVPairs3
	5	Brocade-AVPairs4
	6	Brocade-Passwd-ExpiryDate. This only accepts the string value of 6 in MM/DDYYYY in UTC.
	7	Brocade-Passwd-WarnPeriod. This only accepts the integer value of 6 in days.
Vendor length	2 or higher	1 octet, calculated by server, including vendor-type and vendor-length
Attribute-specific data	ASCII string	Multiple octet, maximum 253, indicating the name of the assigned role and other supported attribute values.

RADIUS Configuration with Virtual Fabrics

When configuring users with Virtual Fabrics, you must also include the Virtual Fabric member list. This section describes the way that you configure attribute types for this configuration.

The values for these attribute types use the syntax `key=val[;key=val]`, where `key` is a text description of attributes, `val` is the attribute value for the given key, the `equal sign (=)` is the separator between key and value, and the `semicolon (;)` is an optional separator for multiple key-value pairs.

Multiple key-value pairs can appear for one Vendor-Type code. Key-value pairs with the same key name may be concatenated across multiple Vendor-Type codes. You can use any combination of the Vendor-Type codes to specify key-value pairs. Note that a switch always parses these attributes from *Vendor-Type code 2* to *Vendor-Type code 4*.

Only the following keys are accepted; all other keys are ignored.

- **HomeLF** is the designated home Virtual Fabric for the account. The valid values are from 1 through 128 and chassis context. The first valid HomeLF key-value pair is accepted by the switch; additional HomeLF key-value pairs are ignored.
- **LFRoleList** is a comma-separated list of Virtual Fabric ID numbers of which this account is a member. Valid numbers range from 1 through 128. A dash between two numbers specifies a range. Multiple Virtual Fabric list key-value pairs within the same or across different Vendor-Type codes are concatenated. Multiple occurrences of the same Virtual Fabric ID number are ignored.
- **ChassisRole** is the account access permission at the chassis level. The chassis role allows the user to execute chassis-related commands in a Virtual Fabrics-enabled environment. Valid chassis roles include the default roles and any of the user-defined roles.

RADIUS authentication requires that the account have valid permissions through the attribute type tuple of HomeLF, LFRoleList or Brocade-Auth-Role. When both (LFRolelist, HomeLF) and Brocade-Auth-Role is provided, the Brocade-Auth-Role will be ignored. When only Brocade-Auth-Role is used, its permission is assigned to default VF. If both are unspecified, the account can log in with default VF as its member list and home Virtual Fabric with user permission. When an invalid values or out of range values are specified for LFRolelist, and HomeLF, the user will be denied login.

In the following example, on a Linux FreeRADIUS Server, the user has the "zoneadmin" role permissions for VFs and "switchadmin" chassis role, with VFlist 2, 4, 5, 6, 7, 8, 10, 11, 12, 13, 15 17, 19, 22, 23, 24, 25, 29, 31 and HomeLF 1.

```
user300 Auth-Type := Local, User-Password == "password"
Brocade-Auth-Role = "zoneadmin",
Brocade-AVPairs1 = "HomeLF=1;LFRoleList=securityadmin:2,4-8,10"
Brocade-AVPairs2 = "LFRoleList=admin:11-13, 15, 17, 19;user:22-25,29,31"
Brocade-AVPairs3 = "ChassisRole=switchadmin"
```

RADIUS Configuration for Password Expiry Date

Fabric OS allows you to configure a password expiration date for each user account and to configure a warning period for notifying the user that the account password is about to expire.

To configure these value, set the following attributes.

- **Brocade-Passwd-ExpiryDate** sets the password expiration date in UTC and mm/dd/yyyy format.
- **Brocade-Passwd-WarnPeriod** sets the warning period in days as an integer value.

You specify either both attributes or none. If you specify a single attribute or if there is any syntax error in the attributes, the password expiration warning will not be issued. If your RADIUS server maintains its own password expiration attributes, you must set the exact date twice to use this feature, once on the RADIUS server and once on the VSA. If the dates do not match, the RADIUS server authentication fails.

The following example sets the password expiration date for an account. It also specifies that a warning be sent to the user 30 days before the password is due to expire.

```
Auth-Type = System
Brocade-Auth-Role = "admin",
Brocade-AVPairs1 = "HomeLF=70",
Brocade-AVPairs2 = "LFRoleList=admin:2,4-8,70,80,128",
Brocade-AVPairs3 = "ChassisRole=switchadmin",
Brocade-Passwd-ExpiryDate = "11/10/2008",
Brocade-Passwd-WarnPeriod = "30"
```


Configuring Fabric OS Users on the RADIUS Server

All existing Fabric OS mechanisms for managing local-switch user accounts and passwords remain functional when the switch is configured to use RADIUS. Changes made to the local switch database do not propagate to the RADIUS server, nor do the changes affect any account on the RADIUS server.

Windows 2012 IAS

To configure a Windows 2012 Internet Authentication Service (IAS) server to use VSA to pass the admin role to the switch in the dial-in profile, the configuration specifies the Vendor code (1588), Vendor-assigned attribute number (1), and attribute value (admin), as shown in the following figure.

Figure 3: Windows 2012 VSA Configuration



Adding the Brocade Attributes Using the Linux FreeRADIUS Server

1. Create and save the `$PREFIX/etc/raddb/dictionary.brocade` file with the following information:

```
# dictionary.brocade
#
VENDOR Brocade 1588
#
# attributes #
ATTRIBUTE   Brocade-Auth-Role           1    string Brocade
ATTRIBUTE   Brocade-AVPairs1             2    string Brocade
ATTRIBUTE   Brocade-AVPairs2            3    string Brocade
ATTRIBUTE   Brocade-AVPairs3            4    string Brocade
ATTRIBUTE   Brocade-AVPairs4            5    string Brocade
ATTRIBUTE   Brocade-Passwd-ExpiryDate    6    string Brocade
```

```
ATTRIBUTE      Brocade-Passwd-WarnPeriod      7      string Brocade
```

This information defines the Brocade vendor ID as 1588, Brocade attribute 1 as Brocade-Auth-Role, Brocade attribute 6 as Brocade-Passwd-ExpiryDate, and Brocade attribute 7 as Brocade-Passwd-WarnPeriod.

2. Open the `$PREFIX/etc/raddb/dictionary` file in a text editor and add the line:

```
#INCLUDE dictionary.brocade
```

As a result, the `dictionary.brocade` file is located in the RADIUS configuration directory and loaded for use by the RADIUS server.

3. Define the permissions for the user in a configuration file. For example, you must add the following statement to the configuration file to grant the user admin permissions.

```
swladmin      Auth-Type := Local, User-Password == "myPassword" Brocade-Auth-Role = "admin",
Brocade-AVPairs1 = "HomeLF=70",
Brocade-AVPairs2 =      "LFRoleList=admin:2,4-8,70,80,128;ChassisRole=admin", Brocade-Passwd-ExpiryDate =
      "11/10/2011",
Brocade-Passwd-WarnPeriod = "30"
```

RSA RADIUS Server

Traditional password-based authentication methods are based on *one-factor* authentication, where you confirm your identity using a memorized password. Two-factor authentication increases the security by using a second factor to corroborate identification. The first factor is either a PIN or password and the second factor is the RSA SecurID token.

RSA SecurID with an RSA RADIUS server is used for user authentication. The Brocade switch does not communicate directly with the RSA Authentication Manager, so the RSA RADIUS server is used in conjunction with the switch to facilitate communication.

For more information on how to install and configure the RSA Authentication Manager and the RSA RADIUS server, visit www.rsa.com.

Setting Up the RSA RADIUS Server

This section describes how to configure the RSA RADIUS server.

1. Create user records in the **RSA Authentication Manager**.
2. Configure the RSA Authentication Manager by adding an agent host.
3. Configure the RSA RADIUS server.

Setting up the RSA RADIUS server involves adding RADIUS clients, users, and vendor-specific attributes to the RSA RADIUS server.

- a) Add the following data to the `vendor.ini` file:

```
vendor-product = Brocade
dictionary = brocade
ignore-ports = no
port-number-usage = per-port-type
help-id = 2000
```

b) Create a `brocade.dct` file that must be added into the `dictiona.dcm` file located in the following path:

`C:\Program Files\RSA Security\RSA Authentication Manager\radius\Service`

Example of a `brocade.dct` file:

```
#####
# brocade.dct -- Brocade Dictionary
#
# (See readme.dct for more details on the format of this file)
#####
#
# Use the Radius specification attributes in lieu of the Brocade one:
#
@radius.dct
MACRO Brocade-VSA(t,s) 26 [vid=1588 type1=%t% len1=+2 data=%s%]
ATTRIBUTE Brocade-Auth-Role Brocade-VSA(1,string) r
ATTRIBUTE Brocade-AVPairs1 Brocade-VSA(2,string) r
ATTRIBUTE Brocade-AVPairs2 Brocade-VSA(3,string) r
ATTRIBUTE Brocade-AVPairs3 Brocade-VSA(4,string) r
ATTRIBUTE Brocade-AVPairs4 Brocade-VSA(5,string) r
ATTRIBUTE Brocade-Passwd-ExpiryDate Brocade-VSA(6,string) r
ATTRIBUTE Brocade-Passwd-WarnPeriod Brocade-VSA(7,integer) r
#####
# brocade.dct -- Brocade Dictionary
```

Example of the `dictiona.dcm` file:

```
#####
# dictiona.dcm
#####
# Generic Radius
@radius.dct
#
# Specific Implementations (vendor specific)
#
@3comsw.dct
@aat.dct
@acc.dct
@accessbd.dct
@agere.dct
@agns.dct
@airespace.dct
@alcatel.dct
@altiga.dct
@annex.dct
@aptis.dct
@ascend.dct
@ascndvsa.dct
@axc.dct
@bandwagn.dct
@brocade.dct <-----
```


Example of a `brocade.dct` file shows what the `dictionary.dcm` file should look like and Example of the `dictionary.dcm` file shows what needs to be modified in the `dictionary.dcm` file.

NOTE

The dictionary files for the RSA RADIUS server must remain in the installation directory. Do not move the files to other locations on your computer.

- c) Add Brocade-VSA macro and define the attributes as follows:
 - `vid (Vendor-ID): 1588`
 - `type1 (Vendor-Type): 1`
 - `len1 (Vendor-Length): >=2`
- d) When selecting items from the **Add Return List Attribute**, select the Brocade VSA attribute to configure (such as **Brocade-Auth-Role**, **Brocade-AVPairs1**, **Brocade-AVPairs2**, **Brocade-AVPairs3**, and **Brocade-AVPairs4**) and type the corresponding string value. See [RADIUS Server Data Configuration](#) for details on Brocade VSA attributes.
- e) Add the Brocade profile.
- f) In **RSA Authentication Manager**, edit the user records that are authenticated using RSA SecurID.

LDAP Server Data Configuration

This section describes the data configuration on the LDAP server. It also describes the LDAP configuration with the virtual fabrics.

The following table describes the keys used for assigning VSA-based account switch roles on the LDAP server.

Table 36: Keys for VSA-based Account Roles on the LDAP Server

Include	Key	Description
ATTRIBUTE	AdminDescription	Windows only attribute for user objects information.
	memberOf	Windows only attribute for user-specific membership information. When a user is authenticated, the role of the user is obtained from the <code>memberOf</code> attribute, which determines group membership.
	brcdAdVfData	2,3,4,5 string. Brocade vendor-specific user attribute to provide virtual fabric member list.
	UPN	Windows only attribute for the user principal name.
	CN	Windows only attribute for the common name.
	sAMAccountName	Windows only attribute for the user, which is unique across the domain. If the authentication based on both UPN and DN fails, then <code>sAMAccountName</code> will be used for authentication.

LDAP Configuration with Virtual Fabrics

If you are using a virtual fabric, the `brcdAdvfData` attribute can be used to map a role to a virtual fabric. Enter the value of the logical fabrics where the user has access. The following three value fields can be specified, separated by a semicolons (;):

- The `HomeLF` field specifies the user's home logical fabric.
- The `LFRoleList` field specifies the additional logical fabric, where the user has access and the user's access permissions for those logical fabrics. Logical fabric numbers are separated by commas (,) and the range is separated by an hyphen (-). The roles are separated by a semicolon (;). When UDR is listed in `LFRoleList`, the user must not delete the specified UDR in FOS.
- The `ChassisRole` field designates the permissions that apply to the `ChassisRole` subset of commands.

In the following example, the logical switch that would be logged in to by default is 10. If 10 is not available, the lowest FID available is chosen. The user is given permission to enter logical switches 1 through 10 in an user, and 11 through 128 in an admin role. Additionally, it is also given the chassis role permission of admin.

```
brcdAdvfData: HomeLF=10;LFRoleList=user:1-10;admin:11-128;ChassisRole=admin
```

Configuring Fabric OS Users on the LDAP Server

All existing Fabric OS mechanisms for managing local-switch user accounts and passwords remain functional when the switch is configured to use LDAP. Changes made to the local switch database do not propagate to the LDAP server, nor do the changes affect any account on the LDAP server.

To add the `brcdAdvfData` attribute in the schema, see [Adding a Brocade VSA in the Windows AD Schema](#).

Configuring Fabric OS Users Using Windows 2012 Server

Perform the following steps to configure Fabric OS users on the LDAP server using Windows 2012 server.

1. From the Windows start menu, select **Programs > Administrative Tools > ADSI.msc**.

ADSI is a Microsoft Windows Resource Utility. This utility must be installed to proceed with the rest of the setup. For Windows 2003, this utility comes with Service Pack 1 or you can download it from the Microsoft website.

2. Navigate to **CN=Users**.
3. Select **Properties**. Click the **Attribute Editor** tab.
4. Double-click the **adminDescription** attribute.

NOTE

The attribute can be added to user objects only.

The **String Attribute Editor** dialog is displayed.

5. Enter the values of the logical fabrics separated by a semicolon into the **Value** field.
Example for adding Virtual Fabrics:

```
HomeLF=10;LFRoleList=admin:128,10;ChassisRole=admin
```

In this example, the default logical switch is 10. If 10 is not available, the lowest fabric ID (FID) available will be chosen. You would have permission to enter logical switch 128 and 10 in an admin role, and you must have the chassis role permission of admin.

NOTE

AD users can have multiple AD group memberships (with one group as primary), out of which only a non-primary group should be used for switch role mapping on the switch. LDAP user authentication

has the following guidelines: The LDAP role is not assigned by default in non-VF mode, and it will not implicitly assign the memberOf value to the switch role without the LDAP role mapping. (For example, if the memberOf value is “securityadmin”, then it will not implicitly assign the “securityadmin” switch role; it will still mandate LDAP role mapping from the LDAP role to the switch role whether or not the name has a matching switch role.) During firmware upgrade, if LDAP authentication is configured as part of the authentication spec on the switch, then a warning message will be displayed to the user: `WARNING: LDAP authentication is enabled on this switch.`

If Brocade VSA is not configured for users on the LDAP server, LDAP authentication uses role mapping in the target firmware version; and if role mapping is also unavailable, then the role of the local user account with the same name as the LDAP user is used for authorization. Ensure that the Brocade VSA configuration exists on the server, or map the AD group to a switch role with the `ldapcfg --maprole` command, or add a local user account with the same name as the LDAP user name to avoid a login failure for LDAP users after an upgrade.

Adding the Brocade Attributes Using the Linux OpenLDAP Server

Use the `brcdAdVfData` attribute to map a role to a Virtual Fabric. To perform this operation, you must modify the schema to include the definition of the `brcdAdVfData` attribute and the definition of a user class that can use this attribute. You can then add this attribute to user entries in the LDAP directory.

When a user is authenticated, the role of the user is obtained from the memberOf attribute, which determines group membership. This feature is supported in OpenLDAP through the memberOf overlay. You must use this overlay on the OpenLDAP server to assign membership information.

Perform the following steps to add the Brocade attribute using the OpenLDAP server:

1. In a schema file, assign the `brcdAdVfData` attribute to a user class.

The following sample schema file defines a new objectClass named `user` with optional attributes `brcdAdVfData` and `description`.

```
#New attr brcdAdVfData
attributetype ( 1.3.6.1.4.1.1588.6.2
    NAME ( 'brcdAdVfData' )
    DESC 'Brocade specific data for LDAP authentication'
    EQUALITY caseIgnoreIA5Match
    SUBSTR caseIgnoreIA5SubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26{256} )
objectclass ( 1.3.6.1.4.1.1588.6.1 NAME 'user'
    DESC 'Brocade switch specific person'
    SUP top AUXILIARY
    MAY ( brcdAdVfData $ description ) )
```

2. Include the schema file in the `slapd.conf` file.

The following example `slapd.conf` line assumes that `local.schema` contains the attribute definition.

```
include /usr/local/etc/openldap/schema/local.schema
```

3. Include the `brcdAdVfData` attribute in a user entry in the LDAP directory.

The following fragment from a file named `test4.ldif` provides an entry for a user with Virtual Fabric access roles.

```
# Organizational Role for Users
dn: cn=Users,dc=mybrocade,dc=com
```

```

objectClass: organizationalRole
cn: Users
description: User
# User entries
dn: cn=Sachin,cn=Users,dc=mybrocade,dc=com
objectClass: user
objectClass: person
objectClass: uidObject
cn: Sachin
sn: Mishra
description: First user
brcdAdVfData: HomeLF=30;LFRoleList=admin:1-128;ChassisRole=admin
userPassword: pass
uid: mishras@mybrocade.com

```

The following command adds the user to the LDAP directory.

```
[User_CentOS users]# ldapadd -D cn=Sachin,dc=mybrocade,dc=com -x -w secret -f test4.ldif
```

Configuring Role-Based LDAP User Attributes

Fabric OS-specific attributes, such as chassis role, home LF, and LF list for authorization in LDAP, are configured on the server while configuring the user. These attributes must be configured for every user on the server. However, you can also configure these attributes on the switch per role level using the `ldapcfg --mapattr` command for the role that is already mapped using the `ldapcfg --maprole` command.

NOTE

From Fabric OS 9.0.0, mapping the LDAP role to the switch role using the `ldapcfg` command is mandatory for non-VF switches.

If the Fabric OS firmware is unable to assign a role to the user using `brcdAdVfData` and mapping of the `memberOf` attribute to the switch role, then if any local user with the same name is created in the switch, that user's role is assigned to the user.

Follow these considerations when configuring role-based LDAP user attributes:

- If these attributes exist on both the server and the switch, the attribute values from the server are used.
- If these attributes do not exist on either the server or the switch, the login fails.

The configuration of role-based LDAP user attributes is performed in the following priority:

1. Configuration on the server.
2. Mapping of the `memberOf` attribute to the switch role.
 - a. Mapping of the Nested group attribute to the switch role. This is applicable only to the Microsoft Windows Active Directory. If there are *Direct* groups available for the user, but none of them are mapped to a switch role, then authentication is done based on the Nested group.

NOTE

If User A is a member of Group A and Group A is a member of Group B, then Group A is a direct group and Group B is an indirect or Nested group.

3. Local user with the same name created in the switch. The following are the limitations for adding the local user name:

- Only a user name that is compatible with the Fabric OS user name policy is supported. For example, a user name that starts with `_` is not supported.
- If there are multiple users with the same sAMAccountName in the list tree, then only one user can be configured in the switch.
- Since `@` is not supported for the user name, the full UPN cannot be added as the user.
- If the sAMAccountName and login name are different for a user, then the administrator must provide the same role to both user names.
- Local authorization is not supported for a remote user where the user name is the same as the default accounts in the Fabric OS firmware.

NOTE

From Fabric OS 9.0.1, if the user exists in the LDAP server with no Brocade VSA data and a valid LDAP group of the user with role mapping on the switch, the remote user can log in to the switch with LDAP credentials, if there is a local user with the same name on the switch. The role assigned to the login user is the role of the local user with the same name. In pre-FOS 9.0.1 versions, the login fails for the user.

Perform the following steps to configure the role-based LDAP user attributes:

1. Map the LDAP role to the switch role.

NOTE

You can map a maximum of 16 LDAP roles using the command `ldapcfg --maprole`.

```
switch:admin> ldapcfg --maprole <LDAP rolename> <switch rolename>
```

2. Add attributes to an existing mapping of LDAP role.

```
switch:admin> ldapcfg --mapattr <LDAP rolename> {[ -l <LF_ID list> ] | [ -h <LF_ID> ] | [ -c <chassis role> ] }
```

The `-l` option maps a role or VF pair. For example:

```
switch:admin> ldapcfg --mapattr <LDAP rolename> -l "user=1-10;admin=11-128" -h 25 -c user
```

The `-h` option maps a home VF. The `-c` option maps a chassis role.

NOTE

If any of the following characters is carried forward from earlier versions of LDAP role, then the mapping is ignored in the Fabric OS v9.0.0 and later.

- `/` – forward slash
- `[ ]` opening and closing square brackets
- `:` colon
- `;` - semicolon
- `|` - pipe
- `=` - equal to
- `,` - comma
- `+` - plus sign
- `*` - asterisk
- `?` - question mark
- `>` - greater than
- `<` - less than

The `ldapcfg --mapattr` configuration is ignored in non-VF mode.

NOTE

If the user configures any duplicated role in the VF list, the first mapping from the list is used during the LDAP authentication and authorization.

TACACS+ Server Data Configuration

The TACACS+ server is configured in the `tac_plus.cfg` file. Open the file by using the editor of your choice and customize the file as needed.

You must add users into this file and provide some attributes specific to the Brocade implementation. The below table lists and defines attributes specific to Brocade.

Table 37: Brocade Custom TACACS+ Attributes

Attribute	Purpose
<code>brcd-role</code>	Role assigned to the user account
<code>brcd-AV-Pair1</code>	The Virtual Fabric member list, and chassis role
<code>brcd-AV-Pair2</code>	The Virtual Fabric member list, and chassis role
<code>brcd-passwd-expiryDate</code>	The date on which the password expires
<code>brcd-passwd-warnPeriod</code>	The time before expiration for the user to receive a warning message

Configuring Virtual Fabric Lists

If your network uses Virtual Fabrics, you should create Virtual Fabric lists for each user to identify the Virtual Fabrics to which the account has access.

Assign the following key-value pairs to the `brcd-AV-Pair1` and, optionally, `brcd-AV-Pair2` attributes to grant the account access to the Virtual Fabrics:

- `HomeLF` is the designated home Virtual Fabric for the account. The valid values are from 1 through 128 and chassis context. The first valid `HomeLF` key-value pair is accepted by the switch. Additional `HomeLF` key-value pairs are ignored.
- `LFRoleList` is a comma-separated list of Virtual Fabric ID numbers to which this account is a member, and specifies the role the account has on those Virtual Fabrics. Valid numbers range from 1 through 128. A - between two numbers specifies a range.

The following example sets the home Virtual Fabric for the `userVF` account to 30 and allows the account admin role access to Virtual Fabrics 1, 3, and 4 and securityAdmin access to Virtual Fabrics 5 and 6.

```
user = userVF {
    pap = clear "password"
    service = shell {
        set brcd-role = zoneAdmin
        set brcd-AV-Pair1 = "homeLF=30;LFRoleList=admin:1,3,4;securityAdmin:5,6"
        set brcd-AV-Pair2 = "chassisRole=admin"
    }
}
```

Configuring the Password Expiration Date

Fabric OS allows you to configure a password expiration date for each user account and to configure a warning period for notifying the user that the account password is about to expire. To configure these values, set the following attributes:

- `brcd-passwd-expiryDate` sets the password expiration date in `mm/dd/yyyy` format.
- `brcd-passwd-warnPeriod` sets the warning period as a number of days.

The following example sets the password expiration date for the `fosuser5` account. It also specifies that a warning be sent to the user 30 days before the password is due to expire.

```
user = fosuser5 {
    pap = clear "password"
    chap = clear "password"
    password = clear "password"
    service = shell {
        set brcd-role = securityAdmin
        set brcd-passwd-expiryDate = 03/21/2014;
        set brcd-passwd-warnPeriod = 30;
    }
}
```

Configuring Fabric OS Users on the TACACS+ Sever

When adding a user to the `tac_plus.cfg` file, you should at least provide the `brcd-role` attribute. The value assigned to this attribute should match a role defined for the switch. When a login is authenticated, the role specified by the `brcd-role` attribute represents the permissions granted to the account. If no role is specified, or if the specified role does not exist on the switch, the account is granted user role permissions only.

See [Role-Based Access Control](#) for details about roles.

The following fragment from a `tac_plus.cfg` file adds a user named `fosuser1` and assigns the `securityAdmin` role to the account.

```
user = fosuser1 {
    chap = cleartext "my$chap$pswrd"
    pap = cleartext "pap-password"
    service = shell {
        brcd-role = securityAdmin;
    }
}
```

Configuring TACACS+ Server Data in RAW Format

When adding a user to applications that support passing *raw* data to FOS, remove the spaces and quotes from what would normally be added to `tac_plus.cfg`. Traditional TACACS+ on linux will parse the data and send the correct *raw* format to FOS. When an application is used, it may or may not parse the data in the same manner. If the provided format in the previous example fails authorization, verify if the application is using direct raw format and adjust the configuration as shown below:

Verify the profile type is "Shell"
Example of RAW view for profile attributes

```
brcd-role=admin
brcd-AV-Pair1=homeLF=128;LFRoleList=admin:1-128
brcd-AV-Pair2=chassisRole=admin
```

Configuring Protocols

Security protocols provide endpoint authentication and communications privacy using cryptography.

Endpoint authentication can be either single authentication or mutual authentication. When a user is authenticated to the switch, which is a single authentication. When a Transport Layer Security (TLS) client and a server can authenticate each other, which is a mutual authentication. The user is clear with which switch they are communicating.

The following table describes the secure protocols that are supported by Fabric OS.

Table 38: Secure Protocol Support

Protocol	Description
CHAP	Challenge Handshake Authentication Protocol (CHAP) uses shared secrets to authenticate switches.
HTTPS	HTTPS is a Uniform Resource Identifier scheme that is used to indicate a secure HTTP connection. Web Tools supports the use of Hypertext Transfer Protocol over SSL (HTTPS).
LDAP	Lightweight Directory Access Protocol (LDAP) with Transport Layer Security (TLS) uses a certificate authority (CA). With the import of signed certificates, you can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL)/TLS technology with LDAP.
RADIUS	RADIUS uses TLS when configured with PEAP-MSCHAPv2. The CA root certificate for RADIUS can be imported using the <code>seccertmgmt</code> command.
SCP	Secure Copy (SCP) is a means of securely transferring computer files between a local and a remote host or between two remote hosts, using the Secure Shell (SSH) protocol. Configuration upload and download support the use of SCP.
Secure Syslog	Secure syslog requires importing syslog CA certificates using the <code>seccertmgmt</code> command.
SFTP	Secure File Transfer Protocol (SFTP) is a network protocol for securely transferring files on a network. Configuration upload and download support the use of SFTP.
SNMP	Simple Network Management Protocol (SNMP) is used in network management systems to monitor network-attached devices for conditions that warrants an administrative attention. Supports SNMPv1 and v3.
SSH	Secure Shell (SSH) is a network protocol that allows data to be exchanged over a secure channel between two network devices. Encryption provides confidentiality and integrity of data. SSH uses the public-key cryptography to authenticate the remote network device and allow the remote network device to authenticate the user, if necessary.
SSL	Fabric OS uses Secure Socket Layer (SSL) to support LDAP and RADIUS peap-mschapv2, syslog, and HTTPS. A certificate must be generated and installed on each switch to enable the HTTPS server communication. Supports 128-bit encryption by default. Also supports TLS 1.0, TLS 1.1, and TLS 1.2. TLS 1.2 is recommended for a stronger encryption. HTTPS supports TLS 1.3.

Fabric OS supports the following features:

- Challenge Response Authentication (CRA) as the client is supported in SCP, SFTP, and SSH sessions.
- Authentication of RADIUS, LDAP, and Syslog client and validation of chain of certificates if a connection to an external server is established over a secure TLS connection. For more information, see [Managing the Security Certificates Using the SecCertMgmt Command](#).
- TLS protocol version configuration and cipher configuration using the `seccryptocfg` command template.

Secure Protocol Deployment Requirements

The following table describes additional software or certificates that you must obtain to deploy secure protocols.

Table 39: Items Needed to Deploy Secure Protocols

Protocol	Host Side	Switch Side
SSHv2	Secure shell client	None
HTTPS	No requirement on the host side except a browser that supports HTTPS	Switch IP certificate for SSL or fully qualified host name (FQHN) of switch.
SCP	SSH daemon, SCP server	None
SFTP	SFTP server	None
SNMPv3	None	None

Security Protocol Scenarios

The following table describes the security protocols that are designed with the four main use cases.

Table 40: Main Security Scenarios

Fabric	Management Interfaces	Comments
Nonsecure	Nonsecure	No special setup is required to use Telnet or HTTP.
Nonsecure	Secure	Secure protocols may be used. An SSL switch certificate must be installed if HTTPS is used.
Secure	Secure	Switches running earlier versions of Fabric OS can be part of the secure fabric, but they do not support the secure management. Secure management protocols must be configured for each participating switch. Nonsecure protocols may be disabled on nonparticipating switches. If SSL is used, then certificates must be installed. For more information on installing certificates, see Installing a Switch Certificate .
Secure	Nonsecure	You must use SSH because Telnet is not allowed with some features.

Secure Copy

The Secure Copy protocol (SCP) runs on port 22. It encrypts data during transfer, thereby avoiding packet sniffers that attempt to extract useful information during data transfer. SCP relies on SSH to provide authentication and security.

Setting Up SCP for Configuration Uploads and Downloads

Use the following procedure to configure SCP for configuration uploads and downloads.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter one of the following commands:
 - If Virtual Fabrics is enabled, enter the `configurechassis` command.
 - If Virtual Fabrics is not enabled, enter the `configure` command.
3. Enter `y` at the `cfgload attributes` prompt.
4. Enter `y` at the `Enforce secure configUpload/Download` prompt.

Example of setting up SCP for `configUpload/Download`:

```
switch:admin# configurechassis
Configure...
```

```
cfgload attributes (yes, y, no, n): [no] y

Enforce secure config Upload/Download (yes, y, no, n): [no] y
Add Suffix to the uploaded file name (yes, y, no, n): [no]
Do you want to enable auto firmwaresync (yes, y, no, n): [yes]

ssl attributes (yes, y, no, n): [no]
webtools attributes (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no]
system attributes (yes, y, no, n): [no]
fos attributes (yes, y, no, n): [no]
```

Secure Shell Protocol

To ensure security, Fabric OS supports Secure Shell (SSH) encrypted sessions. SSH encrypts all messages, including the client transmission of the password during login.

The SSH package contains a daemon (sshd), which runs on the switch. The daemon supports a wide variety of encryption algorithms, such as Digital Encryption Standard (DES) and Advanced Encryption Standard (AES).

NOTE

To maintain a secure network, you should avoid using Telnet or any other unprotected application when you are working on the switch.

Commands that require a secure login channel must originate from an SSH session. Fabric OS supports OpenSSH version 7.9p1 and OpenSSL version 1.1.1d with the disabled heartbeat feature.

If you set up a message of the day (MOTD), the MOTD is displayed either before or after the login prompt, depending on the SSH client implementation. Fabric OS does not control when the message is displayed.

NOTE

Interactive commands are not supported when you use the inline ssh login.

SSH Public Key Authentication

SSH login retry attempts are restricted to three. The clients with a single public-private key pair are allowed with one public key authentication attempt and two password retry attempts. The clients without public-private key pairs are allowed three password retry attempts.

NOTE

If you have three private-public key pairs under a home directory on an SSH client, all three retry attempts are used by the public key authentication. If the public key authentication fails on the switch, there will not be any successive password retries, and you cannot log in to the switch.

OpenSSH public key authentication provides password-less logins, known as SSH authentication that uses public and private key pairs for incoming and outgoing authentication.

This feature allows only one *allowed-user* to be configured to set up outgoing OpenSSH public key authentication. Any admin user can set up incoming OpenSSH public key authentication.

Using OpenSSH Rivest-Shamir-Adleman (RSA), Digital Signature Algorithm (DSA), and Elliptic Curve Digital Signature Algorithm (ECDSA), the authentication protocols are based on a pair of specially generated cryptographic keys called the private key and the public key. The advantage of using these key-based authentication systems is that, in many cases, it is possible to establish secure connections without having to depend on passwords for security.

Incoming authentication is used when the remote host must authenticate to the switch. Outgoing authentication is used when the switch must authenticate to a server or remote host, such as when running the `configupload` or `configdownload` commands or performing firmware download. Both password and public key authentications can coexist on the switch.

Allowed-User

For outgoing authentication, the default admin user must set up the allowed-user with admin permissions. By default, the admin is the configured allowed-user. The allowed-user must have admin permissions to generate a key pair for an outgoing connection, to export the public key, and to delete the public and private keys.

NOTE

No special privileges are required for the incoming allowed-user.

Configuring Incoming SSH Authentication

1. Log in to your remote host.
2. Generate a key pair for host-to-switch (incoming) authentication by verifying that SSH v2 is installed and working (refer to your host's documentation as necessary) by entering the following command:

```
ssh-keygen -t rsa
```

The following example generates an RSA key pair.

```
anyuser@mymachine: ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/users/anyuser/.ssh/id_rsa
):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /users/anyuser/.ssh/id_rsa.
Your public key has been saved in /users/anyuser/.ssh/id_rsa.pub.
The key fingerprint is:
32:9f:ae:b6:7f:7e:56:e4:b5:7a:21:f0:95:42:5c:d1 anyuser@mymachine
```

3. Import the public key to the switch by logging in to the switch as any user with the admin role and entering the `sshUtil importpubkey` command to import the key.

The following example adds the public key to the switch.

```
switch:anyuser> sshutil importpubkey
Enter user name for whom key is imported: aswitchuser
Enter IP address:192.168.38.244
Enter remote directory:~auser/.ssh
Enter public key name(must have .pub suffix):id_rsa.pub
Enter login name:auser
Password:
Public key is imported successfully.
```

4. Test the setup by logging in to the switch from a remote device.

NOTE

The permissions of the remote file will be reflected in Fabric OS.

Configuring Outgoing SSH Authentication

After the allowed-user is configured, the remaining setup steps must be completed by the allowed-user.

Use the following procedure to configure outgoing SSH authentication.

1. Log in to the switch as the default admin.
2. Change the allowed-user's permissions to admin, if applicable.

```
switch:admin> userconfig --change username -r admin
```

The `username` variable is the name of the user who can set up SSH public key authentication, and who can export, and delete keys.

3. Set up the allowed-user by typing the following command:

```
switch:admin> sshutil allowuser username
```

The `username` variable is the name of the user who has admin privileges.

4. Generate a key pair for switch-to-host (outgoing) authentication by logging in to the switch as the allowed-user and entering the `sshUtil genkey [-rsa|-dsa|-ecdsa]` command.

Example of generating a key pair on the switch:

```
switch:alloweduser> sshutil genkey -rsa
Key pair generated successfully.
```

5. Export the public key to the host by logging in to the switch as the allowed-user and entering the `sshUtil exportpubkey [-rsa|-dsa|-ecdsa]` command to export the key.

Example of exporting a public key from the switch:

```
switch:alloweduser> sshutil exportpubkey -rsa
Enter IP address:10.1.1.244
Enter remote directory:~auser/.ssh
Enter login name:auser
Password:
public key out_going.pub is exported successfully.
```

6. Append the public key to a remote host by logging in to the remote host, locating the directory where authorized keys are stored, and appending the public key to the file.

You may need to refer to the host's documentation to locate where the authorized keys are stored.

7. Test the setup by using a command that uses SCP and authentication, such as `firmwareDownload` or `configUpload`.

Deleting Public Keys for Inbound Connections on the Switch

Use the following procedure to delete public keys from the switch.

1. Connect to the switch and log in using an account with admin permissions.
2. Use the `sshutil delpubkeys -all` command to delete public keys.

Example of deleting the public keys on the switch:

```
switch:admin> sshutil delpubkeys -all
WARNING: It deletes all the ssh public keys for all users. Do you want to proceed(yes, y, no, n)[no]?yes
ssh public keys associated to all users are deleted.
```

Deleting Private Keys for Outbound Connections on the Switch

Use the following procedure to delete the public-private key pair from the switch.

1. Log in to the switch as the allowed-user.
2. Use the `sshutil delprivkey [-rsa|-dsa|-ecdsa]` command to delete the private key.

Example of deleting a private key from the switch:

```
switch:alloweduser> sshutil delprivkey -rsa
private key is deleted successfully.
```

Generating and Installing Hostkeys on a Switch

You can generate, install, display, and delete SSH hostkeys on a switch.

1. Enter the `sshutil genhostkey [ -rsa | -dsa | -ecdsa ]` command to generate and install the SSH hostkey.
2. Enter the `sshutil showhostkey` command to display current SSH hostkeys installed on the switch.
3. Enter the `sshutil delhostkey [ -rsa | -dsa | -ecdsa ]` command to delete the selected SSH hostkeys on the switch.

Configuring the Rekey Interval on a Switch

You can configure, display, and disable the SSH rekey interval on a switch.

1. Enter the `sshutil rekeyinterval seconds` command to configure the SSH rekey interval from 900 to 3600 seconds. In the following example, the SSH rekey interval is set to 1200.

```
switch:admin> sshutil rekeyinterval 1200
```

2. Enter the `sshutil showrekey` command to display the SSH rekey interval.

```
switch:admin> sshutil showrekey
```

3. To disable the SSH rekey interval configuration, set the SSH rekey interval to 0.

```
switch:admin> sshutil rekeyinterval 0
```

Managing Known Hosts on a Switch

You can add, display, and delete remote host keys of algorithm type MD5, SHA1, SHA256, SHA384, and SHA512.

1. Log in to the switch as the allowed user.
2. Enter the `sshutil addknownhost <IP_Address/FQDN> -fp <fingerprint> -f` command to add the host key of the specified remote host.
3. Enter the `sshutil showknownhost -all` command to display known hosts installed on the switch.
4. Enter the `sshutil delknownhost -all` command to delete the known hosts on the switch.

You need to generate hostkeys on the remote server in the following file formats only:

- `ssh_host_rsa_key`
- `ssh_host_dsa_key`
- `ssh_host_ecdsa_key`

for any keysize. There is no support for ed25519 keys.

SecCryptoCfg Templates

SecCryptoCfg templates permit you to predefine standardized sets of values for switch cipher configurations. These templates can be easily loaded onto a switch to meet specific security requirements.

Fabric OS supports templates for TLS, SSH, and FIPS configurations. Templates consist of key-value pairs for configuring values for application-specific cipher and protocol version configurations, certificate validation mode, and FIPS-specific configurations.

A template can be specific for the requirements of a certification or the template can be based on the definition of the security configurations for various security levels. For example, a high-security configuration template can enforce high-security strengths that are not FIPS approved. You cannot overwrite the default templates, but you can upload the templates, edit them, and then download them with a different name. You can create a new template similar to default templates, download it, and apply it. Fabric OS supports a maximum of eight templates, including the default templates.

NOTE

The maximum supported size for importing the SecCryptoCfg template is 20 KB.

The following default templates are provided:

- Default configuration (default_generic)
- Secured configuration (default_strong)
- FIPS configuration (default_fips)
- CC configuration (default_cc)

See [Default secCryptoCfg Templates](#) for an example of each template.

Template Structure

The following table lists the groups within each template and the available options. See [Default secCryptoCfg Templates](#) for illustrations.

Table 41: SecCryptoCfg Template Groups and Options

Group	Option	Function
Version	Number	The version number of the template.
SSH	Kex, Mac, Enc	SSH key exchange, MAC, and encryption modes.

Group	Option	Function
AAA	RAD_Ciphers, LDAP_Ciphers, and RSA_Ciphers RAD_Protocol, LDAP_Protocol, and RSA_Protocol	TLS encryption ciphers and the protocol version for RADIUS, LDAP, and RSA SecurID.
LOG	Syslog_Ciphers, Syslog_Protocol	TLS ciphers and the protocol versions for syslog.
HTTPS	Ciphers, Protocol, Ciphers_tlsv1.3	TLS ciphers and the protocol versions for HTTPS.
Compliance	BootUpSelfTests, Zeroize, FIPSInside, SelfTests, CryptoIntegrityTest, and FIPSExtn	This option is present only in the FIPS template. The valid options for these FIPS functions are <code>on</code> and <code>off</code> for each function.
X509v3	Validation	The X509v3 certificate validation permits Fabric OS to enable or disable certificate validation during the certificate import and session establishment. User-defined templates may have either the <code>Basic</code> or <code>Strict</code> validation mode.

NOTE

In case of HTTPS, the `Any` configuration for protocol includes TLSv1.3.

The following points should be considered when working with these templates:

- If an option is mentioned multiple times, the last value is considered.
- Lines starting with `/*` are considered to be comments and must be terminated with `*/`.
- Configuration upload and download do not include the template files. To import, export, apply, and delete templates, use the template commands described in [SecCryptoCfg Template Management](#).
- To explicitly configure or deconfigure cipher suites, the input must be provided using the SSL cipher suite name. Using the TLS cipher suite name does not work. For example, to configure the TLS cipher suite `TLS_DH_RSA_WITH_AES_256_GCM_SHA384`, use the SSL cipher suite name `DH-RSA-AES256-GCM-SHA384`. Using the TLS cipher suite name `TLS_DH_RSA_WITH_AES_256_GCM_SHA384` does not work.
- The TLS cipher strings that can yield at least one cipher suite are considered valid input. Invalid ciphers in the input are ignored.
- Cipher and protocol configurations for RSA are not part of any default templates. But cipher and protocol configurations for RSA can be added or updated in imported templates. Below is an example of how the configuration can look like:
 - `RSA_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!AESCCM8:!AESCCM:!ARIAGCM:!CAMELLIA:!CHACHA20:!SEED:!RC4`
 - `RSA_Protocol:Any`

Certificate Terminology

Term	Description
Root CA certificate	A public key certificate that identifies the root certificate authority (CA). Root certificates are self-signed.
Intermediate certificate	A subordinate certificate issued by a trusted root CA, specifically to issue identity certificates.
Identity certificate	A client or server certificate used to identify the client or server and used to authenticate the client to the server or the server to the client.
Certificate chain	A chain of certificates that contain the authority relationship for the certificate authority.
Online Certificate Status Protocol (OCSP)	The protocol used to verify the revocation of the certificate.

X509v3 Validation Modes

The following table identifies the differences between Basic mode and Strict mode for the X509v3 certificate validation.

Table 42: Difference in X509v3 Validation Modes

Serial No.	Basic Mode	Strict Mode
1.	OCSP validation for identity certificates is not performed.	OCSP validation for identity certificates is performed.
2.	Audit logs for establishing and terminating TLS session are not printed.	An audit log for establishing and terminating TLS session is printed for each TLS session established. In a switch, this applies both as client and as server.
3.	In SSH sessions with a switch, the 1-GB data transfer limit is not enforced for rekeying. Incoming blocks exceeding incoming data limit (or) (2 * outgoing blocks) exceeding outgoing data limit.	In SSH sessions with a switch, rekeying is enforced just before the 1-GB data transfer limit is crossed. Incoming blocks exceeds incoming data limit (or) (2 * outgoing blocks) exceeds outgoing data limit (or) Incoming and outgoing blocks exceeds outgoing data limit (or) Incoming and outgoing blocks exceeds incoming data limit.
4.	For RADIUS, LDAP, or Syslog secure authentication using a certificate: - A host name mismatch error is ignored. - Host name validation is not performed. - Over TLS using a certificate, if a host name mismatch error occurs, no audit log entry is generated.	For RADIUS, LDAP, or Syslog secure authentication using a certificate: - A host name mismatch error causes authentication failure. - Host name validation is enforced and a host name mismatch error causes authentication failure (host name resolution failure also causes authentication failure). - Over TLS using a certificate, if there is a host name mismatch error, an audit log entry is generated.

X509v3 Certificate Validation

Both CA certificates and identity certificates must be validated for common criteria compliance. The following table provides more details. To be compliant with the CC mode, you must configure the validation mode value to be 'strict'.

Table 43: X509v3 Certificate Validation

X509 Fields	Valid NDcPP 1.0 Values	Validation Behavior During Certificate Import	Validation Behavior During Session Establishment
Not Before: Not After:	Certificate that is validated for time and date during import and is accepted only if the time and date are within the allowed range.	Certificates that are expired or invalid are not allowed to be imported. Import fails.	Session establishment fails if CA certificate or identity certificate that is used has expired or is not yet valid. If there was Web Tools, the session can continue if an exception is added to the browser.
CN=	Must be an FQN device name or IP. Wildcards are allowed only for one level of sub-domain and not allowed for the main domain.	Validates identity certificates against the IP or host name of the host during import. If the common name does not match the IP or host name, or the common name has a wildcard in any field other than the first level of sub-domain, then the import fails. The validation does not apply to commoncert.	Validates identity certificates against the IP or host name of the host. This field applies to all identity certificates that are used for session establishment including commoncert. Wildcards are allowed only for the first level of sub-domain.
Public-Key	CC-compliant configuration values are documented. However, no restriction is imposed on using non compliant values.	No validation is performed or audit log generated. Import succeeds.	Audit logs are generated for certificates that have a <code>keysize</code> less than 2048. The session is established successfully.
Key Usage: ^a	Must have <code>Certificate Sign</code> if there is CA certificates and <code>Digital Signature</code> in case of identity certificates. To support RSA-based ciphers for TLS sessions, <code>Key encipherment</code> must be set.	Validated for CA certificates and identity certificates (for <code>Certificate Sign</code> and <code>Digital Signature</code> , respectively, and if validation fails, import fails.	Validated for CA certificates and identity certificates (for <code>Certificate Sign</code> and <code>Digital Signature</code> , respectively), and if validation fails, session establishment fails.
X509v3 Extended Key Usage: ^b	You must rightly indicate whether it is for use as <code>server certificate</code> or <code>client certificate</code> . If it is incorrect, connection is not allowed.	Validated for correct EKU (<code>server</code> for identity certificates for servers and <code>client</code> for identity certificates of clients) and if incorrect, import fails.	Validated for correct EKU (<code>server</code> for identity certificates for servers and <code>client</code> for identity certificates of clients) and, if incorrect, session establishment fails.

a. Key Usage is mandatory in basic mode as well.

b. It is not mandatory in basic mode. But if present must rightly indicate whether it is for use as server certificate or client certificate.

X509 Fields	Valid NDcPP 1.0 Values	Validation Behavior During Certificate Import	Validation Behavior During Session Establishment
Authority Information Access:	You must have valid OCSP server respond affirmatively. If this field is unavailable, the OCSP check is not performed.	Checked during the import of HTTPS and syslog certificates. If present, then a revocation check is performed. If the revocation check fails, then the import fails. If the OCSP URL is not found in certificate, then check is not performed."	Checked for OCSP URL in certificate during the session establishment. If present, then a revocation check is performed and connection is established only if revocation check passes. If no OCSP URL found in certificate, then connection is established successfully.
Signature Algorithm:	CC compliant configuration values are documented. However no restriction is imposed on using non-compliant values.	No validations are performed at import.	No validations are performed during session establishment.
Subject Alternative Name	Not a mandatory attribute. If present, the values that are stored in SAN takes priority over the CN in Subject attribute.	Validated during import for all certificates. Validates identity certificates against the host name of the host during import. If SAN DNS does not match with hostname or has a wildcard in any field other than the first level of subdomain, then the check fails. The validation applies to HTTPS and syslog only.	Validated during sessions for all certificates. Validates identity certificates against the hostname of the host. Applies to all identity certificates that are used for session establishment. Wildcards are allowed only for the first level of sub-domain. If SAN DNS does not match with host name or has a wildcard in any field other than the first level of subdomain, then the check fails. Not applicable to FCAP or commoncert.
Basic Constraints Attribute ^c	Attribute must be present and must set the CA field.	Validated for CA certificates at import. Import fails if the attribute does not exist or the attribute is not set to TRUE.	Validated for all CA certificates that are used during the session establishment and validation fails if the attribute does not exist or attribute is not set to TRUE.

SecCryptoCfg Template Management

Use the following commands to manage the template files.

- To apply the configurations in a template file, use the `seccryptocfg --apply template_file` command.
- To import a template file, use the `seccryptocfg --import` command.

```
switch:admin> seccryptocfg -import <template filename> [-server <ipaddr> -name <user> -proto <scp|sftp|ftp>
-file <remote filename>]
```

- To export a template file, use the `seccryptocfg --export template_file` command.

c. BasicConstraints must be present and set to TRUE in basic validation mode as well.

```
seccryptocfg --export <template filename> -server <ipaddr> -name <user> -proto <scp|sftp|ftp> -file <remote
file name>
```

- To verify the running configuration against a required configuration specified in the template file, use the `seccryptocfg --verify template_file` command.
- To display the existing ciphers, protocol, validation mode, Kex, and Mac, enter `seccryptocfg --show`.
- To display the configuration contents in a template file, use the `seccryptocfg --show template_file` command.
- To delete a template file, use the `seccryptocfg --delete template_file` command.
- To list all the template files, enter `seccryptocfg --lstemplates`.

Default SecCryptoCfg Templates

The following default `seccryptocfg` templates are available and supported. See [Appendix C: Supported Cipher Suites in Fabric OS](#) for more information on supported TLS ciphers and SSH Macs, ciphers, and KEX.

Default_generic

This template is an unhardened default template suitable for initial setup.

```

/*****
* Brocade - Generic Template for Security Crypto Configuration
*
* Desc:
*
* Default values for security crypto configurations for operation in
* generic mode
*
*****/

[Ver] 0.2

/*
* Group : SSH
* Rules : Comma Separated
* Example : aes128-ctr,aes192-ctr -> Note, no space before and after comma.
* Valid options: Kex, Mac, Enc
*/

[SSH]
Enc:aes128-ctr,aes192-ctr,aes256-ctr,aes128-cbc,3des-cbc,aes192-cbc,aes256-cbc
Kex:ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521,diffie-hellman-group-exchange-sha256,diffie-
hellman-group-exchange-sha1,diffie-hellman-group14-sha1,diffie-hellman-group1-sha1
Mac:hmac-sha1,hmac-sha2-256,hmac-sha2-512

/*
* Group : AAA
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: RAD_Ciphers, LDAP_Ciphers
*/

[AAA]
RAD_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!AESCCM8:!AESCCM:!ARIAGCM:!CAMELLIA:!CHACHA20:!
SEED:!RC4

```

```

LDAP_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!AESCCM8:!AESCCM:!ARIAGCM:!CAMELLIA:!CHACHA20:!
SEED:!RC4
RAD_Protocol:Any
LDAP_Protocol:Any

/*
* Group : LOG
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[LOG]
Syslog_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!AESCCM8:!AESCCM:!ARIAGCM:!CAMELLIA:!
CHACHA20:!SEED:!RC4
Syslog_Protocol:Any

/*
* Group : HTTPS
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[HTTPS]
Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!AESCCM8:!AESCCM:!ARIAGCM:!CAMELLIA:!CHACHA20:!
SEED:!RC4
Protocol:Any
Ciphers_tlsv1.3:TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256

/*
* Group : X509v3
* Rules : Textual X509v3 validation options
* Example: Validation:Strict
* Valid options: Ciphers
*/
[X509v3]
Validation:Basic

```

Default_strong

Strong configuration consists of TLS 1.2 or higher and SHA above 256.

```

/*****
* Brocade - Template for High Security Crypto Configuration
*
* Desc:
*
* Default values for security crypto configurations for high security
*
*****/

[Ver] 0.2

/*
* Group : SSH

```

```

* Rules : Comma Separated
* Example : aes128-ctr,aes192-ctr -> Note, no space before and after comma.
* Valid options: Kex, Mac, Enc
*/

[SSH]
Enc:aes128-ctr,aes192-ctr,aes256-ctr
Kex:ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521,diffie-hellman-group-exchange-sha256,diffie-
hellman-group14-sha256,diffie-hellman-group16-sha512,diffie-hellman-group18-sha512,curve25519-sha256
Mac:hmac-sha2-256,hmac-sha2-512

/*
* Group : AAA
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: RAD_Ciphers, LDAP_Ciphers
*/

[AAA]
RAD_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM
LDAP_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM
RAD_Protocol:TLSv1.2
LDAP_Protocol:TLSv1.2

/*
* Group : LOG
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/

[LOG]
Syslog_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM
Syslog_Protocol:TLSv1.2

/*
* Group : HTTPS
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/

[HTTPS]
Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!AESCCM
Protocol:TLSv1.3
Ciphers_tlsv1.3:TLS_AES_256_GCM_SHA384:TLS_AES_128_GCM_SHA256:TLS_AES_128_CCM_8_SHA256:TLS_AES_128_CCM_SHA256

/*
* Group : X509v3
* Rules : Textual X509v3 validation options
* Example: Validation:Strict
* Valid options: Ciphers
*/

```

```
[X509v3]
Validation:Basic
```

Default_fips

The FIPS template has the ciphers that are certified for Brocade products. It consists of key-value pairs for configuring FIPS, such as zeroize, enable, bootprom, and selftests.

```

/*****
 * Brocade - FIPS Template for Security Crypto Configuration
 *
 * Desc:
 *
 * Default values for security crypto configurations for FIPS compliance
 *
 *****/

[Ver] 0.3

/*
 * Group : SSH
 * Rules : Comma Separated
 * Example : aes128-ctr,aes192-ctr -> Note, no space before and after comma.
 * Valid options: Kex, Mac, Enc
 */
[SSH]
Enc:aes128-ctr,aes192-ctr,aes256-ctr,aes128-cbc,aes192-cbc,aes256-cbc
Kex:ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521,diffie-hellman-group-exchange-sha256
Mac:hmac-sha1,hmac-sha2-256,hmac-sha2-512

/*
 * Group : AAA
 * Rules : Textual openssl cipherlist (colon,comma and space separated)
 * Example: ALL:-MD5:!PSK
 * Valid options: RAD_Ciphers, LDAP_Ciphers
 */
[AAA]
RAD_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM:!AESGCM
LDAP_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM:!AESGCM
RAD_Protocol:TLSv1.2
LDAP_Protocol:TLSv1.2

/*
 * Group : LOG
 * Rules : Textual openssl cipherlist (colon,comma and space separated)
 * Example: ALL:-MD5:!PSK
 * Valid options: Ciphers
 */
[LOG]
```

```

Syslog_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM:!AESGCM
Syslog_Protocol:TLSv1.2

/*
* Group : HTTPS
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[HTTPS]
Ciphers:ECDSA:ECDH:RSA:AES:!3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!TLSv1:!AESCCM:!
AESGCM:!SEED:!RC4
Protocol:TLSv1.2
Ciphers_tls1.3:TLS_AES_256_GCM_SHA384:TLS_AES_128_GCM_SHA256:TLS_AES_128_CCM_8_SHA256:TLS_AES_128_CCM_SHA256

/*
* Group : X509v3
* Rules : Textual X509v3 validation options
* Example: Validation:Strict
* Valid options: Ciphers
*/
[X509v3]
Validation:Basic

/*
* Group : Compliance
* Rules : "yes" for configuration enabled and "no" for disabled. For SelfTests, all or default or required
self test name can be provided. Also, note that CryptoVersion is a read-only attribute.
* Example: "yes" OR "no". For SelfTests, following tests are supported if 'all' is provided, or required self
test name can be provided from these supported tests as well:
* DH, DRBG, AES, AES_GCM, AES_CCM, RSA_ENCRYPT, RSA_SIGN, DES, DSA, SHA1, SHA256, SHA384, SHA512, HMAC_SHA1,
HMAC_SHA224, HMAC_SHA256, HMAC_SHA384, HMAC_SHA512, SHA3, ECDSA, ECDH
* Also, following SelfTests are supported if 'default' is provided: AES, AES_GCM, AES_CCM, RSA_SIGN, SHA512,
HMAC_SHA512, ECDSA
* Valid options: SelfTests, Zeroize, Inside, FIPSExtn
*/
[Compliance]
BootUpSelfTests:yes
Zeroize:yes
FIPSInside:yes
SelfTests:default
CryptoIntegrityTest:yes
FIPSExtn:yes

```

Default_cc

This template is required for NDcPP configurations.

NOTE

The ECDSA certificates are not compatible with the default_cc HTTPS cipher string.

```

/*****

```

```

* Brocade - Common Criteria (CC) Template for Security Crypto Configuration
*
* Desc:
*
* Default values for security crypto configurations for CC compliance
*
*****/

[Ver] 0.3

/*
* Group : SSH
* Rules : Comma Separated
* Example : aes128-ctr,aes192-ctr -> Note, no space before and after comma.
* Valid options: Kex, Mac, Enc
*/
[SSH]
Enc:aes128-cbc,aes256-cbc
Kex:diffie-hellman-group14-sha1,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521
Mac:hmac-sha1,hmac-sha2-256,hmac-sha2-512

/*
* Group : AAA
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: RAD_Ciphers, LDAP_Ciphers
*/
[AAA]
RAD_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM
LDAP_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM
RAD_Protocol:TLSv1.2
LDAP_Protocol:TLSv1.2

/*
* Group : LOG
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[LOG]
Syslog_Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!
AESCCM
Syslog_Protocol:TLSv1.2

/*
* Group : HTTPS
* Rules : Textual openssl cipherlist (colon,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/

```



```
[HTTPS]
Ciphers:ECDSA:ECDH:RSA:AES:3DES:!RSAPSK:!DHEPSK:!PSK:!DSS:!ARIAGCM:!CAMELLIA:!CHACHA20:!SSLv3:!TLSv1:!AESCCM
Protocol:Any
Ciphers_tlsv1.3:TLS_AES_256_GCM_SHA384:TLS_AES_128_GCM_SHA256:TLS_AES_128_CCM_8_SHA256:TLS_AES_128_CCM_SHA256

/*
* Group : X509v3
* Rules : Textual X509v3 validation options
* Example: Validation:Strict
* Valid options: Strict/Basic
*/
[X509v3]
Validation:Strict

/*
* Group : Compliance
* Rules : "yes" for configuration enabled and "no" for disabled. For SelfTests, all or default or required
self test name can be provided. Also, note that CryptoVersion is a read-only attribute.
* Example: "yes" OR "no". For SelfTests, following tests are supported if 'all' is provided, or required self
test name can be provided from these supported tests as well:
* DH, DRBG, AES, AES_GCM, AES_CCM, RSA, RSA2048_SHA256, DES, DSA, SHA1, SHA256, SHA384, SHA512, HMAC_SHA1,
HMAC_SHA224, HMAC_SHA256, HMAC_SHA384, HMAC_SHA512, SHA3, ECDSA, ECDH
* Also, following SelfTests are supported if 'default' is provided: AES, AES_GCM, AES_CCM, RSA, SHA512,
HMAC_SHA512, ECDSA
* Valid options: BootUpSelfTests, SelfTests, Zeroize, FIPSInside, CryptoIntegrityTest, FIPSExtn
*/
[Compliance]
BootUpSelfTests:yes
Zeroize:no
FIPSInside:no
SelfTests:default
CryptoIntegrityTest:yes
FIPSExtn:no
```

Certificate Revocation Check Enforcement

Certificate revocation check enforcement is achieved using the Online Certificate Status Protocol (OCSP).

If a certificate contains a OCSP Uniform Resource Identifier (URI), a revocation check is performed during certificate import and TLS session establishment when validation is set to `strict`.

- All TLS-based applications on the switch that support certificate validation of the peer, verify the peer certificate with the OCSP responder for revocation, if the OCSP URI is already present as part of the peer certificate.
- If the peer certificate does not have OCSP URI, the certificate is not verified for revocation and connection status depends on other pre-existing validations of the peer certificate alone.
- The OCSP revocation check is implicit based on the peer certificate having OCSP URI in it. There are no changes to any user interfaces or CLIs.
- The certificate that signs the OCSP response for a peer is the same as the CA certificate that is imported for that peer. No separate OCSP certificate is on the switch for verifying the OCSP response.
- The switch certificates that are sent by the switch to its peer for validation do not have OCSP URI and will not support revocation check by the peer.

The TLS session is established only if the OCSP response is good. The following table describes the responses and the outcome.

OCSP Responder Status	TLS Session Status
Reachable; Response = Good	Pass
Reachable; Response = Signature verification fails	Fail
Reachable; Response = Revoked	Fail
Reachable; Response = Unknown	Fail
Reachable; No Response	Fail
Not reachable	Fail

Client Authentication for TLS Client: RADIUS/LDAP/SYSLOG

The following notes are applicable to all TLS clients. Only SYSLOG is NDcPP compliant.

- The `secCertMgmt` command provides options to support the generation and import of certificates for RADIUS, LDAP, and SYSLOG.
- With this feature, when RADIUS/LDAP/SYSLOG communication over TLS is initiated, if the RADIUS/LDAP/SYSLOG client certificate has been imported on the switch and the RADIUS/LDAP/SYSLOG server has been configured to verify the peer certificate, the TLS connection is successfully established only if the client certificate validation passes.
- If the client certificate for RADIUS/LDAP/SYSLOG has not been imported on the switch, the connection proceeds without client certificate validation (provided that the server does not mandate client certificate validation).
- Use the following CLI commands to import the client CA certificate:
 - a. Import the client CA certificate.

```
secCertMgmt import -ca -client <application> -protocol <protocol>
                  -ipaddr <remote IP address> -certname <certificate name>
                  -login <username> -password <password>
```

where <application> can be radius or ldap or syslog and <protocol> can be scp or ftp .

NOTE

As part of importing the client CA certificate, you must import the complete CA chain of certificates that is required to validate the client identity.

- b. Import the client certificate.

```
secCertMgmt import -cert <application> -protocol <protocol>
                  -ipaddr <remote IP address> -certname <certificate name>
                  -login <username> -password <password>
```

where <application> can be radius or ldap or syslog and <protocol> can be scp or ftp .

Configuring the Ciphers, KEX, and MAC Algorithms

You can also manually configure (without using the templates) the SSH ciphers, key exchange (KEX), message authentication code (MAC) algorithms, and HTTPS ciphers dictated by your security policies.

1. To configure the ciphers and KEX and MAC algorithm for SSH, use the `secCryptoCfg` command.

```
secCryptoCfg --replace -type SSH [-cipher cipher string|-kex value|-mac value] -force
```

The following example configures the ciphers, and KEX and MAC algorithms.

```
seccryptocfg --replace -type SSH -cipher 3des-cbc,aes128-cbc,aes192-cbc -kex diffie-hellman-group-exchange-sha1 -mac hmac-sha2-256
```

To enforce the default algorithm using SSH, use the following command.

```
seccryptocfg --default -type SSH -force
```

2. To configure the cipher for HTTPS, use the `seccryptocfg` command.

```
seccryptocfg --replace -type https -cipher cipher string
```

To enforce the default algorithm using HTTPS, use the following command.

```
seccryptocfg --default -type https -force
```

3. To display the configured algorithm, use the following command.

```
secCryptoCfg --show
```

For crypto configuration, you should use `seccryptocfg` CLI command.

Secure Sockets Layer Protocol

The Secure Sockets Layer (SSL) protocol provides secure access to a fabric through web-based management tools, such as Web Tools. SSL support is a standard Fabric OS feature.

Switches configured for SSL grant access to management tools through Hypertext Transfer Protocol over SSL links (which begin with `https://`) instead of standard links (which begin with `http://`).

SSL uses public key infrastructure (PKI) encryption to protect data transferred over SSL connections. PKI is based on digital certificates obtained from an Internet certificate authority (CA) that acts as the trusted key agent.

Certificates are based on the switch IP address or fully qualified domain name (FQDN), depending on the issuing CA. If you change a switch IP address or an FQDN after activating an associated certificate, you may have to obtain and install a new certificate. Check with the CA to verify this possibility, and plan these types of changes accordingly.

Browser Support

Fabric OS supports the following web browsers for SSL connections:

- Mozilla Firefox v49 or later
- Google Chrome v53 or later

NOTE

Review the release notes for the latest information and to verify if your platform and browser are supported.

SSL Configuration Overview

You configure SSL access for a switch by obtaining, installing, and activating digital certificates. Certificates are required on all switches that are to be accessed through SSL.

Also, for HTTPS (HTTP over TLS), you might need to add a CA certificate to your web browser.

Configuring for SSL involves these main steps, which are described in detail in the next sections.

1. Choose a CA.
2. Generate the public and private key and certificate-signing request (CSR) by using the `seccertmgmt generate -csr https` command.
3. Store the CSR on a file server by using the `seccertmgmt export` command.
4. Obtain the certificates from the CA.

You can request a certificate from a CA through a web browser. After you request a certificate, the CA either sends certificate files by email (public) or gives access to them on a remote host (private).

5. On each switch, install both the CA and switch certificates. When the certificate is loaded on the switch, the HTTPS starts automatically.
6. If necessary, install the root certificate to the browser on the management workstation.

Certificate Authorities

To ease maintenance and allow secure out-of-band communication between switches, consider using one certificate authority (CA) to sign all management certificates for a fabric. If you use different CAs, management services operate correctly, but the Web Tools **Fabric Events** button is unable to retrieve events for the entire fabric.

Each CA (for example, Verisign or GeoTrust) has slightly different requirements; for example, some generate certificates based on IP address, while others require an FQDN, and most require a 1024-bit public/private key pair while some may accept a 2048-bit key. Consider your fabric configuration, check the CA websites for requirements, and gather all the information that the CA requires.

Managing the Security Certificates Using the SecCertMgmt Command

You can manage all operations, such as generating a CSR or certificate, importing a certificate, exporting a CSR or certificate, deleting a CSR or certificate, and displaying a CSR or certificate, for any application using the `seccertmgmt` command.

The following options are available:

```
switch:admin> seccertmgmt
      generate - Generate keypairs and CSR or Certificate
      delete   - Delete CSR or Certificates and associated Keypair
      import   - Import Certificate
      export   - Export CSR or Certificate
      show     - List or Display Certificates
      help     - Shows the usage
```

Generating the Certificates

The following options generate the certificates.

```
switch:admin> seccertmgmt generate
      -csr <fcap|commoncert|https|radius|ldap|syslog|extn>
      -cert <https|extn>
      [-type <rsa|dsa|ecdsa>]
      [-keysize <1024|2048|4096|8192|P384>]
      [-hash <sha1|sha256|sha384|sha512>]
      [-years <x>]
      -keypair_tag <keypair_name>
```

[-f] - no warning

NOTE

- The `-keypair_tag` option works only with the `-extn` type.
- The `-type ecdsa` is supported only with the `-extn` and `-https` types.
- The keysize 4096 and 8192 are not supported for `fcap`.

Importing the Certificates

The following options can be used to import the certificates.

```
switch:admin> seccertmgmt import
-cert <fcap|commoncert|https|radius|ldap|syslog|extn>
-ca -client|-server <fcap|commoncert|https|radius|ldap|syslog|extn|asc>
-keypair_tag <keypair_name>
-protocol <scp|ftp>
-ipaddr <IP address>
-remotedir <remote directory>
-certname <certificate name>
-cacert <preimported_local_ca_cert>
-login <login name>
-password <password>
```

NOTE

- The `-keypair_tag` option works only with the `-extn` type.
- The `-keypair_tag` option must not be present for the `-ca extn` options.
- The `<certificate_name>` must not contain more than one period (dot) before the extension of the file name.
- The `https` allows you to import an externally generated keypair and its certificate. In this case, both the keypair and the certificate must be embedded in the same input file.

NOTE

From Fabric OS 9.1.0, the `seccertmgmt` command is enhanced to import the externally generated private key and its certificate. This enhancement is to enable wildcard certificate support across FOS switches. It is restricted to HTTPS only.

Exporting the Certificates

The following options export the certificates.

```
switch:admin> seccertmgmt export
-cert <fcap|commoncert|https|radius|ldap|syslog|extn|mgmtip>
-ca -client|-server <fcap|commoncert|https|radius|ldap|syslog|extn>
-csr <fcap|commoncert|https|radius|ldap|syslog|extn>
-keypair_tag <keypair_name>
-protocol <scp|ftp>
-ipaddr <IP address>
-remotedir <remote_directory>
-certname <certificate_name>
-login <login_name>
-password <password>
```

NOTE

- The `-keypair_tag` option works only with the `-extn` type.
- The `-keypair_tag` option must not be present for the `-ca` `extn` options.

Deleting the Certificates

The following options delete the certificates.

```
switch:admin> seccertmgmt delete
```

Usage:

```
seccertmgmt delete
-cert <fcap|commoncert|https|radius|ldap|syslog|extn (<certificate name> | -keypair_tag
<keypair_tag>)|all>
-ca -client|-server <fcap|commoncert|https|radius|ldap|syslog|extn <certificate name> |asc|all>
-csr <fcap|commoncert|https|radius|ldap|syslog|extn -keypair_tag <keypair_tag>>
-all <default|fcap|commoncert|https|radius|ldap|syslog|extn>
[-f] - no warning
```

NOTE

- The `-keypair_tag` option works only with the `-extn` type.
- The `-keypair_tag` option should not be present for the `-ca` `extn` options.
- The `<certificate_name>` must not contain more than one period (dot) before the extension of the file name.

Displaying the Certificates

The following options display the certificates:

```
switch:admin> seccertmgmt show
-cert <fcap|commoncert|https|radius|ldap|syslog|extn <certificate_name>>
-ca -client|-server <fcap|commoncert|https|radius|ldap|syslog|extn
<certificate_name>|asc|kafka>
-csr <fcap|commoncert|https|radius|ldap|syslog|extn <csr_name>>
-keypair_tag <keypair_name>
-hexdump - Dump the hexadecimal output
-verbose - Dump the verbose output
-all
```

NOTE

- The `-keypair_tag` option works only with the `-extn` type.
- The `<certificate_name>` must not contain more than one period (dot) before the extension of the file name.

Example:

```
switch:admin> seccertmgmt show -all
```

```
ssh private key:
```

```
Exists
```

```
ssh public keys available for users:
```

```
user
```

Certificate Files:

Protocol	Client CA	Server CA	SW	CSR	PVT Key	Passphrase
FCAP	Empty	NA	Empty	Empty	Empty	Empty
RADIUS	Empty	Empty	Empty	Exist	Exist	NA
LDAP	Empty	Exist	Empty	Empty	Empty	NA
SYSLOG	Empty	Empty	Empty	Empty	Empty	NA
HTTPS	NA	Empty	Exist	Empty	Exist	NA
KAKFA	NA	Empty	NA	NA	NA	NA
ASC	NA	Empty	NA	NA	NA	NA
EXTN	NA	Empty	Empty	Empty	Empty	NA

The `-keypair_tag` option and the certificate names with the `-cacert`, `-certname`, and `-csr` options support the following characters.

Table 44: Supported Characters on Keypair Tag and Certificate Names

Name of the Character	Character
lower-case alphabets	a–z
upper-case alphabets	A–Z
numerals	0–9
tilde	~
percent sign	%
at sign	@
plus sign	+
underscore	_
hyphen	-
colon	:

Generating and Storing a Certificate Signing Request

This section describes the method to generate and store a certificate signing request (CSR) with HTTPS as an example.

1. Connect to the switch, and log in using an account with admin permissions.
2. Enter `seccertmgmt generate -csr https`.
3. Enter the requested information.

The following example generates a CSR.

NOTE

You can input a maximum of 16 DNS names and 9 IP addresses.

```
switch:admin> seccertmgmt generate -csr https
Generating a CSR will automatically do the following:
Delete all existing CSRs.
```

Warning:

Key-pair generation is CPU intensive and can cause high CPU usage
Private IPs and hostnames should not be part of SCN and/or SAN per CA/Browser forum.

```
Continue (yes, y, no, n): [no] y
Country Name (2 letter code, eg, US):US
State or Province Name (full name, eg, California):California
Locality Name (eg, city name):San Jose
Organization Name (eg, company name):*****
Organizational Unit Name (eg, department name):***
Common Name (Fully qualified Domain Name, or IP address):**.**.**.**
Email Address:abc@xyz.com
Subject Alternative Name, DNS (Fully Qualified Domain Name, or IP address):*****.com
Subject Alternative Name, DNS (Fully Qualified Domain Name, or IP address):<enter>
Subject Alternative Name, IPAddress (IP v4 or v6 address): 10.1.1.3
Subject Alternative Name, IPAddress (IP v4 or v6 address): <enter>
```

Generating CSR, file name is: 10.1.1.3.csr

Your CA might require specific codes for Country, State or Province, Locality name, Organization name, Organizational Unit name, common name, email address, and subject alternative names. Make sure that your spelling is correct and matches the CA requirements. If the CA requires that the Common Name be specified as an FQDN, make sure that the FQDN is set on the domain name switch or director. The IP address or FQDN is the switch where the certificate is installed.

You can also specify the hash type using the `-hash <sha1|sha256|sha384|sha512>` option while generating the CSR. For example:

```
switch:admin> seccertmgmt generate -csr https -type rsa -keysize 2048 -hash sha256 -years 1 -f
```

NOTE

- FCAP CSR generation supports only a keysize of 1024 and 2048 and a hash algorithm of sha1 and sha256.
- Generating a new CSR will not remove the installed certificate and key to maintain the HTTPS state. The key and the certificate will be replaced when the newly signed certificate is imported.

4. Enter `seccertmgmt export` to store the CSR.

5. Enter the requested information. You can use either File Transfer Protocol (FTP) or Secure Copy Protocol (SCP).

FTP or SCP is used based on the setup in your machine. Enter the IP address and remote directory name of the file server to export the CSR. Enter your account name and password of the server.

The following example exports a CSR.

```
switch:admin> seccertmgmt export -csr https
Select protocol [ftp or scp]: ftp
Enter IP address: 192.1.2.3
Enter remote directory: path_to_remote_directory
Enter Login Name: your account
Enter Password: your password
Success: exported CSR.
```


Obtaining CA-Signed Certificates

When you have generated a CSR, you will need to follow the instructions on the website of the certificate issuing authority that you want to use; and then obtain the certificate.

Fabric OS supports the following types of files from the certificate authority (CA):

- `.cer` (binary/text)
- `.crt` (binary/text)
- `.pem` (text)
- `.psk` (text)

Typically, the CA provides the certificate files listed in the following table.

Table 45: SSL Certificate Files

Certificate File	Description
<code>name .pem</code>	The identity certificate of the host (the switch certificate when the switch is the entity being authenticated).
<code>name Root .pem</code>	The root CA certificate. To be installed on the host that will perform identity certificate validation. If being installed on the switch, this file must be concatenated with <code>CA .pem</code> .
<code>name CA .pem</code>	The intermediate CA certificate. It must be installed on the host that performs the certificate validation.

Use the following procedure to obtain a security certificate:

1. Generate and store the CSR as described in [Generating and Storing a Certificate Signing Request](#).
2. Open a web browser window on the management workstation, and go to the CA website. Follow the instructions to request a certificate. Locate the area in the request form into which you are to paste the CSR.
3. Connect to the switch, and log in using an account with admin permissions.
4. Enter the `seccertmgmt show -csr <application>` command. The contents of the CSR are displayed.
5. Locate the section that begins with `BEGIN CERTIFICATE REQUEST` and ends with `END CERTIFICATE REQUEST`.
6. Copy and paste this section (including the `BEGIN` and `END` lines) into the area provided in the request form; then, follow the instructions to complete and send the request.

It may take several days to receive the certificates. If the certificates arrive by email, save them to an FTP server. If the CA provides access to the certificates on an FTP server, make note of the path name and make sure that you have a login name and password on the server.

Installing a Switch Certificate

Before you import a switch certificate, be aware of the following:

- Certificate Authorities may provide their certificates in different encodings and different extensions. Be sure to save the certificate with the applicable file extension before you import the certificate to the switch.
For example, certificates that contain lines similar to the following are usually `.pem` encoded:

"----BEGIN REQUEST----" and "----END REQUEST----" (and may include the strings "x509" or "certificate")

- For Certificate Authorities that request information regarding the type of web server, Fabric OS uses the Apache web server running on Linux.
- If you try to import certificates of different sizes for a given switch, the import fails. If this occurs, remove the previous certificate and then import the new certificate.

Use the following procedure to install a security certificate on a switch.

NOTE

You must perform this procedure on each switch.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `seccertmgmt import -ca -server <application>` or `seccertmgmt import -ca -client <application>` command to import the CA certificate for server or client respectively.
For information on CA chain of certificates, see [Creating a Complete Chain of CA Certificates](#).
3. Select a protocol, enter the IP address of the host on which the CA certificate is saved, and enter your login name and password.
4. Enter the `seccertmgmt import -cert <application>` command to import the switch certificate.
5. Select a protocol, enter the IP address of the host on which the switch certificate is saved, and enter your login name and password.
 - a. In case of HTTPS certificate installation, to optimize the HTTPS session handshake, enable the HTTP keep alive mode using the `mgmtapp --enable keepalive` CLI command.

Example of Installing a CA and Switch Certificate in Non-Interactive Mode

```
switch:admin> seccertmgmt import -ca -server https -protocol SCP -ipaddr 10.1.18.17 -remotedir /
share/certs -login fvt -certname PADL12R2SHA256CA.pem
Password:
Success: imported https server CA certificate [PADL12R2SHA256CA.pem].
```

```
switch:admin> seccertmgmt import -cert https -protocol SCP -ipaddr 10.1.18.17 -remotedir /share/
certs -login fvt -certname PADL12R2SHA256.pem
Password:
Success: imported https certificate [PADL12R2SHA256.pem].
Certificate file in configuration has been updated.
Secure http has been enabled.
```

Generating Self-Signed Certificates

You can use the `seccertmgmt generate` command to generate a new key pair, sign the generated certificate with the private key, install the self-signed identity certificate within the switch for HTTPS and extn, and enable the HTTPS and extn service. Also, using Monitoring and Alerting Policy Suite (MAPS), you can configure the time period for notification when certificates expire.

NOTE

Generating self-signed certificates for Fibre Channel Authentication Protocol (FCAP) is not supported.

1. Log in to the device.
2. Run the `seccertmgmt generate -cert <application>` command.

```
switch:admin> seccertmgmt generate -cert https
Generating a new certificate will automatically do the following
1. Delete existing switch certificate(s).
2. Disable secure protocol HTTPS

Warning: Certificate generation is CPU intensive and can cause high CPU usage

Continue (yes, y, no, n): [no] y
Generating .....Generated self-signed https certificate successfully.
```

From Fabric OS 9.1.0, the switches are shipped with self-signed SSL certificate with a placeholder IP address as the Common Name and the Subject Alternative Name. This is to enable HTTP over TLS (HTTPS) by default. The factory-shipped self-signed certificate may be replaced later with the another self-signed certificate that has been generated. The newly generated self-signed certificate has the Common Name based on the configured switch IP address and has the Subject Alternative Name field containing all configured switch IP addresses.

Client applications may enforce restrictions during validation of a self-signed server certificate for HTTPS. In such cases, to continue using HTTPS with a self-signed certificate, SSL certificate verification may have to be turned off or an exception may have to be added on the client application. You can generate a third-party signed SSL certificate (with the configured IP address or FQDN of the switch as the Common Name/Subject Alternative Name) and use it for HTTPS for enhanced security.

Creating a Complete Chain of CA Certificates

The Fabric OS switch sends the identity certificate whenever there is a request for a client/server certificate for the SSL client/server authentication. To comply with this requirement, the complete chain of certificate authority (CA) certificates that sign the identity certificate must be installed on the remote machine in the corresponding server/client application. You are required to import the whole CA certificate chain to the Fabric OS switch prior to importing the client/server certificates, to validate the client/server certificates. The CA certificate chain can be imported using the `seccertmgmt` command with the `-ca` option. The Fabric OS switch also validates the chain of CA certificates during the import operation.

Import the complete hierarchy of CA certificates by concatenating the certificates. For example, follow this signing hierarchy:

1. The public key certificate is `root.pem`, which identifies root CA. Root certificates are self-signed.
2. The first-level Intermediate certificate is `Level1-ca.pem`.

NOTE

An intermediate certificate is a subordinate certificate issued by the trusted root CA, specifically to issue identity certificates.

3. The second-level Intermediate certificate is `Level2-ca.pem`.
4. The third-level Intermediate `Level3-ca.pem` is the third-level certificate.
5. An Identity certificate is a client or server certificate that identifies the client or server. Use this certificate to authenticate the client to the server or the server to the client.

In the preceding chain of certificates, if you use the intermediate certificate `Level3-ca.pem` to sign the switch or server certificate, you can create a chain certificate on a Linux host machine. Use the following command: `cat Level3-ca.pem Level2-ca.pem Level1-ca.pem root.pem > chain-ca.pem`

Import the `chain-ca.pem` certificate as the CA certificate for the switch.

Configuring the Browser for HTTPS

The root certificate may already be installed on your browser; if not, you must install it. To see whether it is already installed, check the certificate store on your browser.

The next procedures are guides for installing root certificates on Chrome and Firefox browsers. For more detailed instructions, refer to the documentation that came with the certificate.

Checking and Installing Root Certificates on Mozilla Firefox

Use the following procedure to check and install a root security certificate on a switch using Firefox:

1. Select **Tools** > **Options**.
2. Click **Advanced**.
3. Click the **Encryption** tab.
4. Click **View Certificates** > **Authorities** and scroll the list to see if the root certificate is listed. For example, its name may have the form *name*Root.crt. Take the appropriate following action based on whether you find the certificate:
 - If the certificate is listed, you do not need to install it. You can skip the rest of this procedure.
 - If the certificate is not listed, click **Import**.
5. Browse to the certificate location and select the certificate. For example, select *name*Root.crt.
6. Click **Open** and follow the instructions to import the certificate.

Simple Network Management Protocol

Simple Network Management Protocol (SNMP) is a set of protocols for managing complex networks. SNMP protocols are application layer protocols. Using SNMP, devices within a network send messages, called protocol data units (PDUs), to different parts of a network. Network management using SNMP requires three components:

- SNMP manager
- SNMP agent
- Management information base (MIB)

SNMP Manager

The SNMP manager can communicate to the devices within a network using the SNMP protocol. Typically, SNMP managers are network management systems (NMS) that manage networks by monitoring the network parameters and, optionally, setting parameters in managed devices. Normally, the SNMP manager sends read requests to the devices that host the SNMP agent, to which the SNMP agent responds with the requested data. In some cases, the managed devices can initiate the communication, and send data to the SNMP manager using asynchronous events called traps.

SNMP Agent

The SNMP agent is software that resides in the managed devices in the network, and collects data from these devices. Each device hosts an SNMP agent. The SNMP agent stores the data, and sends it when requested by an SNMP manager. In addition, the agent can asynchronously alert the SNMP manager about events, by using special PDUs called traps.

Management Information Base


SNMP agents in the managed devices store the data about these devices in a database called a management information base (MIB). The MIB is a hierarchical database that is structured on the standard specified in the RFC 2578 (Structure of Management Information Version 2 [SMIV2]).

The MIB is a database of objects that a network management system uses to manage and monitor devices on the network. The MIB can be retrieved by a network management system that uses SNMP. The MIB structure determines the scope of management access allowed by a device. By using SNMP, a manager application can issue read or write operations within the scope of the MIB.

Basic SNMP Operation


Every Brocade device carries an agent and MIB. The agent accesses information about a device and makes it available to an SNMP network management station.

Figure 4: SNMP Structure



When active, the management station can get information or set information when it queries an agent. SNMP commands, such as `get`, `set`, and `getnext` are sent from the management station, and the agent replies after the value is obtained or modified. Agents use variables, to report such data as the number of bytes and packets in and out of the device, or the number of broadcast messages sent and received. These variables are also known as *managed objects*. All managed objects are contained in the MIB.

Figure 5: SNMP Query



The management station can also receive *traps*, which are unsolicited messages from the switch agent if an unusual event occurs.

Figure 6: SNMP Trap



The agent can receive queries from one or more management stations and can send traps to up to six management stations.

Configuring SNMP Using the CLI

For more information on configuring SNMP, refer to the *Brocade Fabric OS Command Reference Manual*.

Configuring the SNMP Security Level

The following example sets the SNMP security level to 1 (authentication only). This setting allows all SNMPv1 users to perform GET and SET operations on MIBs, but creates an exception for SNMPv3 users who do not have authentication and privacy privileges (noAuthnoPriv).

```
switch:admin> snmpconfig --set secLevel -snmpget 1 -snmpset 3
```

Select SNMP GET security Levels are 0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 = No Access.

Select SNMP SET security Levels are 0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 = No Access.

NOTE

The default options are `No security` for GET and `No access` for SET.

The following table shows the security level options.

Table 46: Security Level Options

Security Level	Protocol	Query Behavior	Traps
No security [0] (noAuthnoPriv)	SNMPv1 SNMPv3	Allowed. Allowed.	Sent. Sent.
Authentication only [1] (authNoPriv)	SNMPv1 SNMPv3	Allowed. All SNMPv3 users allowed except noAuthNoPriv users.	Sent. Sent for all SNMPv3 users except noAuthNoPriv users.
Authentication and Privacy [2] (authPriv)	SNMPv1 SNMPv3	Not allowed. Only SNMPv3 users with authPriv privilege are allowed.	Not sent. Sent only for authPriv users.
No Access [3]	SNMPv1 SNMPv3	Not allowed. Not allowed.	Not sent. Not sent.

Supported Protocol Configurations for SNMPv3 Users

The following table shows the authentication and privacy protocols supported for configuring SNMPv3 users.

Table 47: Supported Protocol Options

Protocol	Options
Auth protocols	MD5 SHA noAuth
Priv protocols	DES noPriv AES128 AES256

Configuring SNMPv3 Users and Traps

The following examples list how to configure SNMPv3 users and traps.

NOTE

If you use SNMPGET to poll the switches, for example, HiTrack, the SNMPv3 user must be configured based on the logical switch that is being monitored. If SNMP monitoring is for the default switch, you do not need the SNMP user to be configured as a user on the switch. Whereas, when SNMP monitoring is VF specific, you must create a user with permission for the corresponding VF and add the created user into the SNMP database as an SNMPv3 user. After the SNMPv3 user is successfully added, you can use the same user name in the management applications for SNMP queries to the corresponding logical switch.

1. Create a user on the switch in a non-VF context using the `userconfig` command with the required role.

```
switch:admin> userconfig --add fa_adm -r fabricadmin
Setting initial password for fa_adm
Enter new password:*****
Re-type new password:*****
Account fa_adm has been successfully added.
```

Create a user on the switch in a VF context using the `userconfig` command with the required role.

```
switch:admin> userconfig --add sa_user -r switchadmin -l 1-128 -hl -c admin
Setting initial password for sa_user
Enter new password:*****
Re-type new password:*****
Account sa_user has been successfully added.
```

2. Enter `snmpconfig --add snmpv3` to create the SNMPv3 user.

```
sw0:admin> snmpconfig --add snmpv3 -index 6 -user test6 -groupname rw -auth_proto 1 -auth_passwd
myauthpassword -priv_proto 1 -priv_passwd myprivpassword

Committing configuration.....done.
```

3. Enter `snmpconfig --show snmpv3` to list the created SNMPv3 user.

```
sw0:admin> snmpconfig --show snmpv3

SNMP Informs = 1 (ON)

SNMPV3 user password encrypted = 1 (ON)

SNMPv3 USM configuration:
User 1:
User 2:
User 3:
User 4:
User 5:
User 6 (ro): test6
    Auth Protocol: noAuth
    Priv Protocol: noPriv
    Engine ID: 00:00:00:00:00:00:00:00
User 7:
```

```

User 8:
User 9:
User 10:
User 11:
User 12 (rw): testv3
    Auth Protocol: MD5
    Priv Protocol: DES
    Engine ID: 00:00:00:00:00:00:00:00:00

SNMPv3 Trap/Informs configuration:
Trap Entry 1:      10.1.18.10
    Trap Port: 162
    Trap User: testv3
    Trap recipient Severity level: 0
    Notify Type: TRAP(1)
Trap Entry 2:      No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 3:      No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 4:      No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 5:      No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 6:      No trap recipient configured yet
    Notify Type: TRAP(1)

```

4. Enter `snmpconfig --set snmpv3` to modify the SNMPv3 user.

```
sw0:admin> snmpconfig --set snmpv3 -index 6 -user test6 -groupname rw
```

```
SNMP Informs Enabled (true, t, false, f): [true]
```

```
SNMPv3 user configuration(snmp user not configured in FOS user database will have physical AD and admin
role as the default):
```

```

User (rw): [password]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [1]
New Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (1..4) [3]
New Priv Passwd:
Engine ID: [80:00:05:23:01:ac:1a:1a:ac]

```

```

SNMPv3 trap/inform recipient configuration:
Trap Recipient's IP address : [172.26.26.172]

```

```

Notify Type [TRAP(1)/INFORM(2)]: (1..2) [1] 2
UserIndex: (1..6) [1] 1
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [172.26.26.172]

```



```
Committing configuration.....done.
```

The following is an example of the SNMPv3 user trap recipients configured with DNS names and IPv4 address.

```
snmpconfig --set snmpv3 -index <index> -default
sw0:admin> snmpconfig --set snmpv3 -trap_index 1 -user_index 12 -host 10.1.18.10 -port 162 -severity 0 -
notify_type trap
Committing configuration.....done.
```

```
sw0:admin> snmpconfig --show snmpv3
```

```
SNMP Informs = 1 (ON)
```

```
SNMPV3 user password encrypted = 1 (ON)
```

```
SNMPv3 USM configuration:
```

```
User 1:
User 2:
User 3:
User 4:
User 5:
User 6 (ro): test6
    Auth Protocol: noAuth
    Priv Protocol: noPriv
    Engine ID: 80:00:06:34:03:ac:1a:1a:00
User 7:
User 8:
User 9:
User 10:
User 11:
User 12 (rw): testv3
    Auth Protocol: MD5
    Priv Protocol: AES128
    Engine ID: 80:00:05:23:01:ac:1a:1a:ac
```

```
SNMPv3 Trap/Informs configuration:
```

```
Trap Entry 1:    No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 2:    No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 3:    No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 4:    No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 5:    No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 6:    No trap recipient configured yet
```

Notify Type: TRAP(1)

To display the test traps associated with the real MIBs supported in Fabric OS:

```
sw0:admin> snmptraps --show
```

#	Mib Name	Supported Traps
---	-----	-----
001	SW-MIB	sw-fc-port-scn
		sw-ip-v6-change-trap
		sw-pmgr-event-trap
		sw-event-trap
		sw-fabric-reconfig-trap
		sw-fabric-segment-trap
		sw-state-change-trap
		sw-zone-config-change-trap
		sw-port-move-trap
		sw-brcd-generic-trap
		sw-device-status-trap
002	FICON-MIB	link-rnid-device-registration
		link-rnid-device-deregistration
		link-lirr-listener-added
		link-lirr-listener-removed
		link-rlir-failure-incident
003	FA-MIB	conn-unit-status-change
		conn-unit-port-status-change
		conn-unit-event-trap
004	MIB-2	cold-restart-trap
		warm-restart-trap
005	IF-MIB	if-link-up-trap
		if-link-down-trap
006	RFC1157	snmp-authetication-trap
007	HA-MIB	fru-status-change-trap
		fru-history-trap
		cp-status-change-trap
008	MAPS-MIB	maps-trap
		maps-quiet-time-trap
009	T11-FC-ZONE-SERVER-MIB	t11ZsRequestRejectNotify
		t11ZsMergeSuccessNotify
		t11ZsMergeFailureNotify
		t11ZsDefZoneChangeNotify
		t11ZsActivateNotify

NOTE

The T11-FC-ZONE-SERVER-MIB and its traps are for internal use only.

To send all traps to the configured recipients:

```
switch:admin> snmptraps --send
Number of traps sent : 30
```

To send the sw-fc-port-scn trap to the configured recipients:

```
switch:admin> snmptraps --send -trap_name sw-fc-port-scn
Number of traps sent : 1
```

To unblock port traps on all the ports or on a specific port:

```
switch:admin> snmptraps --unblock -port ALL
switch:admin> snmptraps --unblock -port 1/10
```

To block port traps on slot 1 and port 10:

```
Switch:admin> snmptraps --block -port 1/10
```

NOTE

Blocking of all ports using the ALL option is not allowed.

The following example displays the accessControl configuration.

```
switch:admin> snmpconfig --set accessControl
SNMP access list configuration:
Access host subnet area in dot notation: [0.0.0.0] 192.168.0.0
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0] 10.1.18.2
Read/Write? (true, t, false, f): [true] f
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0] 10.1.18.3
Read/Write? (true, t, false, f): [true] f
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration...done.
```

To remove all the SNMP users and traps from the switch, execute `snmpconfig --default snmpv3/snmpv1`.

```
switch:admin> snmpconfig --set mibCapability -mib_name SW-MIB -bitmask 0x10
Operation succeeded
```

```
switch:admin> snmpconfig --show mibCapability
[...]
SW-MIB: NO
swFault: NO
swSensorScn: NO
swFCPortScn: NO
swEventTrap: YES
    DesiredSeverity:None
swIPv6ChangeTrap: NO
swPmgrEventTrap: NO
swFabricReconfigTrap: NO
swFabricSegmentTrap: NO
swExtTrap: NO
swStateChangeTrap: NO
swPortMoveTrap: NO
swBrcdGenericTrap: NO
swDeviceStatusTrap: NO
```

```
swZoneConfigChangeTrap: NO(output truncated)
```

To enable the SW-MIB only without changing the current trap configuration:

```
switch:admin> snmpconfig --enable mibCapability -mib_name SW-MIB
Operation succeeded
```

```
switch:admin> snmpconfig --show mibCapability
[...]
SW-MIB: YES
swFault: NO
swSensorScn: NO
swFCPortScn: NO
swEventTrap: YES
    DesiredSeverity:None
swTrackChangesTrap: NO
swIPv6ChangeTrap: NO
swPmgrEventTrap: NO
swFabricReconfigTrap: NO
swFabricSegmentTrap: NO
swExtTrap: NO
swStateChangeTrap: NO
swPortMoveTrap: NO
swBrcdGenericTrap: NO
swDeviceStatusTrap: NO
swZoneConfigChangeTrap: NO(output truncated)
```

To re-enable all traps under the SW-MIB category after they were disabled:

```
switch:admin> snmpconfig --set mibCapability -mib_name SW-MIB -bitmask 0xFFFF
Operation succeeded
```

```
switch:admin> snmpconfig --show mibCapability
[...]
SW-MIB: YES
swFault: YES
swSensorScn: YES
swFCPortScn: YES
swEventTrap: YES
    DesiredSeverity:None
swTrackChangesTrap: YES
swIPv6ChangeTrap: YES
swPmgrEventTrap: YES
swFabricReconfigTrap: YES
swFabricSegmentTrap: YES
swExtTrap: YES
swStateChangeTrap: YES
swPortMoveTrap: YES
swBrcdGenericTrap: YES
swDeviceStatusTrap: YES
swZoneConfigChangeTrap: YES(output truncated)
```

To display the configuration for all MIBs and associated traps:

```
sw0:admin> snmpconfig --show mibcapability
FE-MIB: YES
SW-MIB: YES
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
FCIP-MIB: YES
IF-MIB: YES
BROCADE-MAPS-MIB: YES
T11-FC-ZONE-SERVER-MIB: NO
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
        DesiredSeverity:None
    swIPv6ChangeTrap: YES
    swPmgrEventTrap: YES
    swFabricReconfigTrap: YES
    swFabricSegmentTrap: YES
    swExtTrap: NO
    swStateChangeTrap: NO
    swPortMoveTrap: NO
    swBrcdGenericTrap: YES
    swDeviceStatusTrap: YES
    swZoneConfigChangeTrap: NO
FA-TRAP: YES
    connUnitStatusChange: YES
    connUnitEventTrap: YES
    connUnitPortStatusChange: YES
FICON-TRAP: YES
    linkRNIDDeviceRegistration: YES
    linkRNIDDeviceDeRegistration: YES
    linkLIRRLListenerAdded: YES
    linkLIRRLListenerRemoved: YES
    linkRLIRFailureIncident: YES
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: YES
IF-TRAP: YES
    linkDown: YES
    linkUp: YES
MAPS-TRAP: YES
    mapsTrapAM: YES
    mapsQuietTimeExpirationTrap: YES
T11-FC-ZONE-SERVER-TRAP: NO
    t11ZsRequestRejectNotify: NO
    t11ZsMergeFailureNotify: NO
    t11ZsMergeSuccessNotify: NO
    t11ZsDefZoneChangeNotify: NO
    t11ZsActivateNotify: NO
```

To set the system group:

```
switch:admin> snmpconfig --set systemgroup -syscontact "Sys contact" -authTrapEnabled false -sysdescr
"Filed Support" -syslocation "Sys Location"
Committing configuration.....done.
switch:admin> snmpconfig --show systemgroup
sysDescr = Filed Support
sysLocation = Sys Location
sysContact = Sys contact
authTrapEnabled = 0 (OFF)
```

5. Set the security level.

```
switch:admin> snmpconfig --set secLevel -snmpget 2 -snmpset 2
switch:admin> snmpconfig --show secLevel
GET security level = 2, SET level = 2
SNMP GET Security Level: Authentication and Privacy
SNMP SET Security Level: Authentication and Privacy
```

If you select the SNMP GET security level to "3" or "No Access", you will receive the following warning message shown in the example.

```
switch:admin> snmpconfig --set secLevel
Select SNMP GET Security Level
(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 = No Access): (0..3) [0] 3
Warning: Modifying the security access level to No Access for GET operation might impact the SNMP GET/SNMP
Walk query triggered by the applications monitoring the system.
```

To set the security level to the default:

```
switch:admin> snmpconfig --default seclevel
GET security level = 0, SET level = 3
SNMP GET Security Level: No security
SNMP SET Security Level: No Access
SNMP GET Security Level will be set to 'No Security'
SNMP SET Security Level will be set to 'No Access'
Do you want to continue? (yes, y, no, n): [no] y

switch:admin> snmpconfig --show seclevel
GET security level = 0, SET level = 3
SNMP GET Security Level: No security
SNMP SET Security Level: No Access
```

6. Set the audit interval.

```
switch:admin> snmpconfig --set auditinterval -interval 31
Committing configuration.....done.
switch:admin> snmpconfig --show auditInterval
SNMP Audit Interval (in min): 31
```

Set the default audit interval.

```
switch:admin> snmpconfig --default auditInterval
*****

This command will reset the agent's audit log interval configuration back to factory default
*****

SNMP Audit Interval (in min): 31
```

```
*****
Are you sure? (yes, y, no, n): [no] y
switch:admin> snmpconfig --show auditInterval
SNMP Audit Interval (in min): 60
```

NOTE

The Success and Failure count Audit messages are only for SNMPv3 logins and not for SNMPv1 logins.

7. In the manager (SNMP browser), create a user snmpadmin1 with the authentication protocol as noAuth, the privacy protocol as noPriv, set the password, and set the trap port as 162. (The same values are set as in the switch SNMPv3 configuration.)

Notes:

- SNMPv3 supports the AES-128, AES-256, and DES protocols.
- For resolving the AES-256 protocol in the User-based Security Model (USM) MIB walk, the eso Consortium MIB must be loaded.

Encrypting SNMPv3 User Passwords

The SNMPv3 user authentication and privacy passwords stored in the configuration database can be stored in either an encrypted or unencrypted form.

If password encryption is enabled, the SNMPv3 user authentication and privacy passwords are encrypted and stored in the configuration database. The encrypted password cannot be decrypted. If an encrypted SNMPv3 user password is disabled, the authentication and privacy passwords are reset to their default values. The SNMPv3 user password encryption option is only available in the CLI.

**CAUTION**

If the SNMPv3 user authentication and privacy passwords are encrypted, do not modify the configuration key attributes value for that user in the configuration file. Doing so will cause the configuration download to fail.

To encrypt SNMPv3 user passwords, complete the following steps.

1. Connect to the device and log in using an account with admin permissions.
2. Enable the SNMPv3 encryption. The command that you will use depends on the mode in which you are operating.
 - If you are in interactive mode, enter `snmpconfig --set snmpv3`, and then enter `t` or `true`, followed by `Yes`.

```
device:admin> snmpconfig --set snmpv3
SNMPv3 user password encryption Enabled (true, t, false, f): [false]
Warning "The encrypted password cannot be decrypted. Do you want to continue (Yes / No)?"
```

```
sw0:admin> snmpconfig --set snmpv3

SNMP Informs Enabled (true, t, false, f): [false]

SNMPV3 Password Encryption Enabled (true, t, false, f): [false]
```

```
SNMPv3 user configuration(snmp user not configured in FOS user database will have default VF context and
admin role as the default):
```

- If you are in non interactive mode, enter `snmpconfig --set snmpv3 -enable passwd_encryption`.
- ```
device:admin> snmpconfig --set snmpv3 -enable passwd_encryption
```

When you enable the encryption option during the SNMPv3 user configuration, the password encryption configuration key for each configured user is automatically enabled internally by the software. This key identifies whether the passwords should be encrypted for each user.

## Changing Encrypted SNMPv3 User Passwords to Unencrypted Form

When encrypted, you cannot remove the encryption from the SNMPv3 user authentication and privacy passwords stored in the configuration database. You can only replace them with new unencrypted passwords.

To change an account's passwords to unencrypted form, you must first disable the encrypted password, and then replace it. When an encrypted SNMPv3 password is disabled, the authentication and privacy passwords are reset to their default values. You then specify new values for the passwords that will not be encrypted.

### NOTE

The SNMPv3 user password decryption option is available only in the CLI.

To change encrypted SNMPv3 user passwords to unencrypted passwords, complete the following steps.

1. Connect to the device and log in using an account with admin permissions.
2. Disable the SNMPv3 encryption. The command that you will use depends on the mode in which you are operating.
  - If you are in interactive mode, enter `snmpconfig --set snmpv3`, and then enter `f` or `false`, followed by `yes`.  

```
device:admin> snmpconfig --set snmpv3
SNMPv3 user password encryption Enabled (true, t, false, f): [false]
Warning "The encrypted password cannot be decrypted. Do you want to continue (Yes / No)?"
```
  - If you are in non interactive mode, enter `snmpconfig --set snmpv3 -disable passwd_encryption`.  

```
device:admin> snmpconfig --set snmpv3 -disable passwd_encryption
```
3. If you want to change from the default passwords, enter the new unencrypted passwords. See [Local Account Passwords](#) for instructions on this task.

## Managing SNMPv1 Communities and SNMPv3 Users

Fabric OS v8.2.x or earlier releases supported six default SNMPv1 communities and SNMPv3 users. You could use the `--set` command to modify the default SNMPv1 communities and SNMPv3 users.

Beginning with Fabric OS v9.0.0, you can manage the SNMPv1 communities and SNMPv3 users with the `--add` or `--delete` options. A new suboption `trap_index` is added for SNMPv3 user. You must use the `trap_index` only, while adding an SNMPv3 user to the host. Fabric OS v9.0.0 supports 12 SNMPv3 users. The default accounts will be removed after performing the `netinstall` or by issuing the `configremoveall` command.

Beginning with Fabric OS v9.0.0, each group name can be configured with either read-only or read-write permission. Resetting the group name of all SNMPv1 and v3 to read-only is not supported in Fabric OS v9.0.0.

To access the switch using SNMP, you must configure the SNMPv1 communities or SNMPv3 users. From Fabric OS v9.0.0, you can add or remove the SNMPv1 communities and SNMPv3 users using the `--add` or `--delete` options.



## Adding SNMPv1 Communities or SNMPv3 Users

To add an SNMPv1 or SNMPv3 account, complete the following steps:

1. Execute the `snmpconfig --add snmpv1` or `snmpconfig --add snmpv3` command with a group name to create an SNMPv1 community or SNMPv3 user.
  - `snmpconfig --add snmpv1 -index <index> -community <community name> -groupname <rw/ro> -host <ipv4/ipv6/dns> -port <udp port> -severity <severity level>`
  - `snmpconfig --add snmpv3 -index <index> -user <user_name> -groupname <rw/ro> -auth_proto <auth_protocol> -auth_passwd <auth_password> -priv_proto <priv_protocol> -priv_passwd <priv_password> -engine_id <engine_id>`
2. Set the read-only or read-write group name for each SNMPv1 or v3 user using the `--set` command.
  - `snmpconfig --set snmpv1 -index <index> -groupname <rw/ro>`
  - `snmpconfig --set snmpv1 -index <index> -community <community name> -groupname <rw/ro> -host <ipv4/ipv6/dns> -port <udp port> -severity <severity level>`
  - `snmpconfig --set snmpv1 -index <index> -default`
  - `snmpconfig --set snmpv3 -index <index> -groupname <rw/ro>`
  - `snmpconfig --set snmpv3 -enable informs`
  - `snmpconfig --set snmpv3 -disable informs`
  - `snmpconfig --set snmpv3 -enable passwd_encryption`
  - `snmpconfig --set snmpv3 -disable passwd_encryption`
  - `snmpconfig --set snmpv3 -index <index> -user <user_name> -groupname <rw/ro> -auth_proto <auth_protocol> -auth_passwd <auth_password> -priv_proto <priv_protocol> -priv_passwd <priv_password> -engine_id <engine_id>`
  - `snmpconfig --set snmpv3 -trap_index <trap_index> -user_index <index> -host <ipv4/ipv6/dns> -port <udp_port> -severity <severity_level> -notify_type <trap/informs>`
  - `snmpconfig --set snmpv3 -index <index> -default`
  - `snmpconfig --set snmpv3 -index <index> -user <user_name> -groupname <rw/ro> -auth_proto <auth_protocol> -auth_passwd <auth_password> -priv_proto <priv_protocol> -priv_passwd <priv_password> -engine_id <engine_id>`
  - `snmpconfig --set snmpv3 -trap_index <trap_index> -user_index <index> -host <ipv4/ipv6/dns> -port <udp port> -severity <severity level> -notify_type <trap/informs>`

The following sample output shows the addition of an SNMPv1 community:

```
sw0:FID128:admin> snmpconfig --show snmpv1
Community 1: test1 (rw)
 No trap recipient configured yet
Community 2: test2 (rw)
 No trap recipient configured yet
Community 3: test (rw)
 No trap recipient configured yet
Community 4: test (ro)
 No trap recipient configured yet
Community 5:
 No trap recipient configured yet
Community 6:
 No trap recipient configured yet
SNMPv1:Enabled
sw0:FID128:admin> snmpconfig --add snmpv1 -index 5 -community private -groupname rw
Committing configuration.....done.
```

```
sw0:FID128:admin> snmpconfig --show snmpv1
Community 1: test1 (rw)
 No trap recipient configured yet
Community 2: test2 (rw)
 No trap recipient configured yet
Community 3: test (rw)
 No trap recipient configured yet
Community 4: test (ro)
 No trap recipient configured yet
Community 5: private (rw)
 No trap recipient configured yet
Community 6:
 No trap recipient configured yet
SNMPv1:Enabled
```

The following sample output shows the addition of a read-write group name to an SNMPv3 user with an index of 9:

```
sw0:FID128:admin> snmpconfig --add snmpv3 -index 9 -user mytest1 -groupname rw
Committing configuration.....done.
sw0:FID128:admin> snmpconfig --show snmpv3
SNMP Informs = 0 (OFF)
SNMPV3 user password encrypted = 0 (OFF)
SNMPv3 USM configuration:
User 1:
User 2 (ro): test1
Auth Protocol: noAuth
Priv Protocol: noPriv
User 3 (rw): test2
Auth Protocol: noAuth
Priv Protocol: noPriv
User 4 (ro): test3
Auth Protocol: noAuth
Priv Protocol: noPriv
User 5 (ro): test4
Auth Protocol: noAuth
Priv Protocol: noPriv
User 6 (ro): test5
Auth Protocol: noAuth
Priv Protocol: noPriv
User 7 (ro): test6
Auth Protocol: noAuth
Priv Protocol: noPriv
User 8 (ro): test7
Auth Protocol: noAuth
Priv Protocol: noPriv
User 9 (rw): mytest1
Auth Protocol: noAuth
Priv Protocol: noPriv
User 10 (ro): test9
Auth Protocol: noAuth
Priv Protocol: noPriv
User 11 (ro): test10
```

```

 Auth Protocol: noAuth
 Priv Protocol: noPriv
User 12 (ro): test11
 Auth Protocol: noAuth
 Priv Protocol: noPriv
SNMPv3 Trap/Informs configuration:
Trap Entry 1: No trap recipient configured yet
 Notify Type: TRAP(1)
Trap Entry 2: No trap recipient configured yet
 Notify Type: TRAP(1)
Trap Entry 3: No trap recipient configured yet
 Notify Type: TRAP(1)
Trap Entry 4: No trap recipient configured yet
 Notify Type: TRAP(1)
Trap Entry 5: No trap recipient configured yet
 Notify Type: TRAP(1)
Trap Entry 6: No trap recipient configured yet
 Notify Type: TRAP(1)
sw0:FID128:admin>

```

## Deleting SNMPv1 Communities or SNMPv3 Users

To delete an SNMPv1 community or an SNMPv3 account, complete the following steps.

1. Connect to a device and log in using an account with admin permissions.
2. Remove an SNMPv1 community or SNMPv3 user using the `snmpconfig --delete` command. Use the index value to delete an SNMPv1 community or SNMPv3 user corresponding to the index.
  - `snmpconfig --delete snmpv1 -index index`
  - `snmpconfig --delete snmpv3 -index index`

The following sample output shows the deletion of an SNMPv1 community:

```

sw0:FID128:admin> snmpconfig --delete snmpv1 -index 1
Committing configuration.....done.
sw0:FID128:admin> snmpconfig --show snmpv1
Community 1:
 No trap recipient configured yet
Community 2: test2 (rw)
 No trap recipient configured yet
Community 3: test (rw)
 No trap recipient configured yet
Community 4: test (ro)
 No trap recipient configured yet
Community 5: private (rw)
 No trap recipient configured yet
Community 6:
 No trap recipient configured yet
SNMPv1:Enabled

```

## Firmware Considerations

A firmware upgrade from Fabric OS v8.2.x to v9.0.x is nondisruptive. A successful firmware upgrade from Fabric OS v8.2.x to v9.0.x will carry forward all the default and nondefault SNMP accounts.

When you perform a firmware downgrade, you must verify the group name and proceed. You must retain the default group name as “rw” for the first three entries and “ro” for the next three entries.

If you perform a firmware downgrade with more than six SNMPv3 entries, the download will be blocked and an following error message displays:

```
SNMPv3 has more than 6 users. Please keep only 6 user entries and delete the remaining users for successful
downgrade. Please use snmpconfig --delete snmpv3 -index <7>
to delete the users.
```

You can delete the index by changing the index value using the `snmpconfig --delete snmpv1 -index` command. The following sample output shows the deletion of an index.

```
sw0:FID128:admin> snmpconfig --delete snmpv1 -index 1
Committing configuration.....done.
sw0:FID128:admin> snmpconfig --show snmpv1
Community 1:
 No trap recipient configured yet
Community 2: test2 (rw)
 No trap recipient configured yet
Community 3: test (rw)
 No trap recipient configured yet
Community 4: test (ro)
 No trap recipient configured yet
Community 5: private (rw)
 No trap recipient configured yet
Community 6:
 No trap recipient configured yet
SNMPv1:Enabled
```

The following snippet is a sample output message that displays when you perform a firmware downgrade with a nondefault group name:

```
Checking system settings for firmwaredownload...
SNMPV1 has non default groupname. That is the first 3 entries should have the group name as (rw)
and the next 3 entries should have (ro). Please check and update default groupname for all the entries.
The following item(s) need to be addressed before downloading the specified firmware:
SNMPv1 has non default groupname. That is the first 3 entries should have the group name as (rw)
and the next 3 entries should have (ro). Please check and update default groupname for all the entries.
Firmwaredownload failed
```

## Telnet Protocol

Telnet is enabled by default. To prevent passing clear text passwords over the network when connecting to the switch, you can block the Telnet protocol using an IP filter policy. For more information on IP filter policies, see [IP Filter policy](#).

### ATTENTION

Before blocking Telnet, make sure that you have an alternate method of establishing a connection with the switch.

## Blocking Telnet

If you create a new policy using commands with just one rule, all the missing rules have an implicit deny and you lose all IP access to the switch, including Telnet, SSH, and management ports.

Use the following procedure to block Telnet access.

1. Connect to the switch and log in using an account with admin permissions.
2. Clone the default policy by entering the `ipfilter --clone` command.

```
switch:admin> ipfilter --clone BlockTelnet -from default_ipv4
```

3. Save the new policy by entering the `ipfilter --save` command.

```
switch:admin> ipfilter --save BlockTelnet
```

4. Verify that the new policy exists by entering the `ipfilter --show` command.

```
switch:admin> ipfilter --show
```

5. Add a rule to the policy, by entering the `ipfilter --addrule` command.

```
switch:admin> ipfilter --addrule BlockTelnet -rule 1 -sip any -dp 23 -proto tcp -act deny
```

### ATTENTION

The rule number assigned must precede the default rule number for this protocol. For example, in the defined policy, the Telnet rule number is 2. Therefore, to effectively block Telnet, the rule number to assign must be 1. If you choose not to use 1, you must delete the Telnet rule number 2 after adding this rule. See [Deleting a rule from an IP Filter policy](#) for more information on deleting IP filter rules.

6. Save the new IP filter policy by entering the `ipfilter --save` command.
7. Verify that the new policy is correct by entering the `ipfilter --show` command.
8. Activate the new IP filter policy by entering the `ipfilter --activate` command.

```
switch:admin> ipfilter --activate BlockTelnet
```

9. Verify that the new policy is active (the `default_ipv4` policy should be displayed as defined ).

```
switch:admin> ipfilter --show
```

```
Name: default_ipv4, Type: ipv4, State: defined
```

| Rule | Source IP | Protocol | Dest Port  | Action |
|------|-----------|----------|------------|--------|
| 1    | any       | tcp      | 22         | permit |
| 2    | any       | tcp      | 23         | permit |
| 3    | any       | tcp      | 80         | permit |
| 4    | any       | tcp      | 443        | permit |
| 5    | any       | udp      | 161        | permit |
| 6    | any       | udp      | 123        | permit |
| 7    | any       | tcp      | 600 - 1023 | permit |
| 8    | any       | udp      | 600 - 1023 | permit |

```
Name: default_ipv6, Type: ipv6, State: active
```

| Rule | Source IP | Protocol | Dest Port | Action |
|------|-----------|----------|-----------|--------|
|------|-----------|----------|-----------|--------|

|   |     |     |            |        |
|---|-----|-----|------------|--------|
| 1 | any | tcp | 22         | permit |
| 2 | any | tcp | 23         | permit |
| 3 | any | tcp | 80         | permit |
| 4 | any | tcp | 443        | permit |
| 5 | any | udp | 161        | permit |
| 6 | any | udp | 123        | permit |
| 7 | any | tcp | 600 - 1023 | permit |
| 8 | any | udp | 600 - 1023 | permit |

Name: BlockTelnet, Type: ipv4, State: active (modified)

| Rule | Source IP | Protocol | Dest Port  | Action |
|------|-----------|----------|------------|--------|
| 1    | any       | tcp      | 23         | deny   |
| 2    | any       | tcp      | 22         | permit |
| 3    | any       | tcp      | 80         | permit |
| 4    | any       | tcp      | 443        | permit |
| 5    | any       | udp      | 161        | permit |
| 6    | any       | udp      | 123        | permit |
| 7    | any       | tcp      | 600 - 1023 | permit |
| 8    | any       | udp      | 600 - 1023 | permit |

## Port Configuration

The following table provides information on ports that the switch uses. When configuring the switch for various policies, take into consideration firewalls and other devices that may sit between switches in the fabric and your network or between the managers and the switch.

**Table 48: Port Information**

| Port  | Type | Common Use | Comment                                                                                                              |
|-------|------|------------|----------------------------------------------------------------------------------------------------------------------|
| 22    | TCP  | SSH, SCP   | N/A                                                                                                                  |
| 23    | TCP  | Telnet     | Use the <code>ipfilter</code> command to block the port.                                                             |
| 80    | TCP  | HTTP       | Use the <code>ipfilter</code> command to block the port.                                                             |
| 123   | UDP  | NTP        | N/A                                                                                                                  |
| 161   | UDP  | SNMP       | Disable the SNMP service on the remote host if you do not use it, or filter incoming UDP packets going to this port. |
| 443   | TCP  | HTTPS      | Use the <code>ipfilter</code> command to block the port.                                                             |
| 52357 | UDP  | EZ Server  | Use the <code>ipfilter</code> command to block the port.                                                             |

## Unblocking Telnet

Use the following procedure to unblock Telnet access.

1. Connect to the switch through a serial port or SSH and log in as admin.
2. Enter the `ipfilter --delete` command to delete the IP filter policy that blocks telnet or use `ipfilter --delrule` to delete the particular IP filter rule.

See [Deleting a rule from an IP Filter policy](#) for more information on deleting IP filter rules.

3. To permanently delete the policy, enter the `ipfilter --save` command.

### ATTENTION

If you deleted the rule to deny Telnet, you must add a rule to permit Telnet.

## Listener Applications

Brocade switches block Linux subsystem listener applications that are not used to implement supported features and capabilities.

The following table lists the listener applications that Brocade switches either block or do not start. Note that RPC ports are blocked.

**Table 49: Blocked Listener Applications**

| Listener Application | Brocade X7 and X6 Director Families | Brocade Switches |
|----------------------|-------------------------------------|------------------|
| chargen              | Disabled                            | Disabled         |
| daytime              | Disabled                            | Disabled         |
| discard              | Disabled                            | Disabled         |
| echo                 | Disabled                            | Disabled         |
| ftp                  | Disabled                            | Disabled         |
| rexec                | Block with packet filter            | Disabled         |
| rlogin               | Block with packet filter            | Disabled         |
| rsh                  | Block with packet filter            | Disabled         |
| rstats               | Disabled                            | Disabled         |
| rusers               | Disabled                            | Disabled         |
| time                 | Block with packet filter            | Disabled         |

## Ports and Applications Used by Switches

The following table lists the defaults for accessing hosts, devices, switches, and zones.

**Table 50: Default Accessing for Host, Device, Switch and Zone**

| Application   | Default Access                                                                                                                                                                                                                                                                                                     |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hosts         | <ul style="list-style-type: none"> <li>Any host can access the fabric by SNMP.</li> <li>Any host can Telnet to any switch in the fabric.</li> <li>Any host can establish an HTTP connection to any switch in the fabric.</li> <li>Any host can establish an API connection to any switch in the fabric.</li> </ul> |
| Devices       | <ul style="list-style-type: none"> <li>All devices can access the management server.</li> <li>Any device can connect to any FC port in the fabric.</li> </ul>                                                                                                                                                      |
| Switch access | <ul style="list-style-type: none"> <li>Any switch can join the fabric.</li> <li>All switches in the fabric can be accessed through a serial port.</li> </ul>                                                                                                                                                       |
| Zoning        | No zoning is enabled.                                                                                                                                                                                                                                                                                              |

## Configuring Security Policies

---

This section describes the supported Access Control List (ACL) policy which is identified by a specific name, and only one policy of each type can exist, except for DCC policies.

Policy names are case-sensitive and must be entered in all uppercase. Fabric OS provides the following policies:

- **Fabric configuration server policy:** Restricts which switches can change the configuration of the fabric.
- **Device connection control (DCC) policy:** Restricts which Fibre Channel device ports can connect to which Fibre Channel switch ports.
- **Switch connection control (SCC) policy:** Restricts which switches can join with a switch.

### How ACL Policies Are Stored

The policies are stored in a local database. The database contains the ACL policy types of FCS, DCC, SCC, and IPFilter. The number of policies that may be defined is limited by the size of the database. FCS, SCC, and DCC policies are all stored in the same database.

The limit for the security policy database size is set to 1 MB. The policies are grouped by state and type. A policy can be in either of the following states:

- Active, which means that the policy is being enforced by the switch.
- Defined, which means that the policy has been set up but is not enforced.

#### NOTE

The maximum members allowed for SCC policy is 9999. If you add the members beyond 9999 to SCC policy, you will receive an error message "Maximum limit for policy members reached". This limitation still applies irrespective of the security policy database size limit is set to 1 MB.

Policies with the same state are grouped together in a `Policy Set`. Each switch has the following two sets:

- `Active policy set`, which contains ACL policies being enforced by the switch.
- `Defined policy set`, which contains a copy of all ACL policies on the switch.

When a policy is activated, the defined policy either replaces the policy with the same name in the active set or becomes a new active policy. If a policy appears in the defined set but not in the active set, the policy was saved but has not been activated. If a policy with the same name appears in both the defined and active sets but they have different values, then the policy has been modified but the changes have not been activated.

### Virtual Fabric Considerations for ACL Policies

ACL policies, such as DCC, SCC, and FCS, can be configured on each logical switch. The limit for the security policy database size is set to 1 MB per logical switch.

### Policy Members

The FCS, DCC, and SCC policy members are specified by device port WWN, switch WWN, domain IDs, or switch names, depending on the policy. The valid methods for specifying policy members are listed below.



**Table 51: Valid Methods for Specifying Policy Members**

| Policy Name    | Device Port WWN<br>or Fabric Port WWN | Switch WWN | Domain ID | Switch Name |
|----------------|---------------------------------------|------------|-----------|-------------|
| FCS_POLICY     | No                                    | Yes        | Yes       | Yes         |
| DCC_POLICY_ddd | Yes                                   | Yes        | Yes       | Yes         |
| SCC_POLICY     | No                                    | Yes        | Yes       | Yes         |

## ACL Policy Management

All policy modifications are temporarily stored in volatile memory until those changes are saved or activated. You can create multiple sessions to the switch from one or more hosts. You should make changes from one switch only to prevent multiple transactions from occurring. Each logical switch has its own ACL.

The FCS, SCC, and DCC policies in Secure Fabric OS are not interchangeable with Fabric OS FCS, SCC, and DCC policies. You should upload and save a copy of the Fabric OS configuration after creating policies. For more information on configuration uploads, see [Maintaining the Switch Configuration File](#).

### NOTE

All changes, including the creation of new policies, are saved and activated on the local switch only—unless the switch is in a fabric that has a strict or tolerant fabric-wide consistency policy for the ACL policy type for SCC or DCC. See [Policy database distribution](#) for more information on the database settings and the fabric-wide consistency policy.

## Displaying ACL Policies

You can view the active and defined policy sets at any time. In addition, within a defined policy set, any policies created in the same login session will be listed, but these policies are automatically deleted if you log out without saving them.

To view the active and defined policy sets, complete the following steps.

1. Connect to the switch and log in using either an account with admin permissions, or an account with “O” permission for the Security RBAC class of commands.
2. Enter `secPolicyShow`.

The following example shows the command and a typical response.

```
switch:admin> secPolicyShow
```

---

ACTIVE POLICY SET

---

DEFINED POLICY SET

## Saving Changes without Activating the Policies

1. Connect to the switch and log in using either an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Enter the `secPolicySave` command.

## Activating ACL Policy Changes

You can implement changes to the ACL policies using the `secPolicyActivate` command. This action saves the changes to the active policy set and activates all policy changes since the last time the command was issued. You cannot activate policies on an individual basis; all changes to the entire policy set are activated by the command. Until a

`secPolicySave` or `secPolicyActivate` command is issued, all policy changes are in volatile memory only and are lost upon rebooting.

1. Connect to the switch and log in using either an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Type the `secPolicyActivate` command.

Example of activating policy changes.

```
switch:admin> secpolicyactivate
About to overwrite the current Active data.
ARE YOU SURE (yes, y, no, n): [no] y
```

## Deleting an ACL Policy

1. Connect to the switch and log in using either an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Enter the `secPolicyDelete` command, specifying the name of the ACL policy to delete.

```
secpolicydelete " policy_name"
```

3. Save and activate the policy deletion by entering the `secPolicyActivate` command.

Example of deleting an ACL policy

```
switch:admin> secpolicydelete "DCC_POLICY_010"
About to delete policy Finance_Policy.
Are you sure (yes, y, no, n):[no] y
Finance_Policy has been deleted.
```

## Adding a Member to an Existing ACL Policy

As soon as a policy has been activated, the aspect of the fabric managed by that policy is enforced.

1. Connect to the switch and log in using either an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Enter the `secPolicyAdd` command.
3. To implement the change immediately, enter the `secPolicyActivate` command.

For example, to add a member to the SCC\_POLICY using the switch WWN:

```
switch:admin> secpolicyadd "SCC_POLICY", "12:24:45:10:0a:67:00:40"
Member(s) have been added to SCC_POLICY.
```

Example of adding members to the DCC policy.

To add two devices to the DCC policy, and to attach domain 3 ports 1 and 3 (WWNs of devices are 11:22:33:44:55:66:77:aa and 11:22:33:44:55:66:77:bb):

```
switch:admin> secpolicyadd "DCC_POLICY_abc",
"11:22:33:44:55:66:77:aa;11:22:33:44:55:66:77:bb;3(1,3)"
```

## Removing a Member from an ACL Policy

As soon as a policy has been activated, the aspect of the fabric managed by that policy is enforced.

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Enter the `secPolicyRemove` command.
3. To implement the change immediately, enter the `secPolicyActivate` command.

For example, to remove a member that has a WWN of 12:24:45:10:0a:67:00:40 from the SCC\_POLICY:

```
switch:admin> secpolicyremove "SCC_POLICY", "12:24:45:10:0a:67:00:40"
Member(s) have been removed from SCC_POLICY.
```

## Abandoning Unsaved ACL Policy Changes

You can abandon all ACL policy changes that have not yet been saved.

1. Connect to the switch and log in using either an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Enter the `secPolicyAbort` command.

Example of aborting unsaved changes.

```
switch:admin> secpolicyabort
Unsaved data has been aborted.
```

All changes since the last time the `secPolicySave` or `secPolicyActivate` commands were entered are aborted.

## FCS Policies

The fabric configuration server policy in base Fabric OS may be performed on a local switch basis and may be performed on any switch in the fabric.

The FCS policy is not present by default, but it must be created. When the FCS policy is created, the WWN of the local switch is automatically included in the FCS list. Additional switches can be included in the FCS list. The first switch in the list becomes the primary FCS switch.

Switches in the fabric are designated as either a primary FCS, backup FCS, or non-FCS switch. Only the primary FCS switch is allowed to modify and distribute the database within the fabric. Automatic distribution is supported and you can either configure the switches in your fabric to accept the FCS policy or manually distribute the FCS policy. Changes made to the FCS policy are saved to permanent memory only after the changes have been saved or activated; they can be aborted later if you have set your fabric to distribute the changes manually.

**Table 52: FCS Policy States**

| Policy State                        | Characteristics                                                                                                                                                                                  |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No active policy                    | Any switch can perform fabric-wide configuration changes.                                                                                                                                        |
| Active policy with one entry        | A primary FCS switch is designated (local switch), but there are no backup FCS switches. If the primary FCS switch becomes unavailable for any reason, the fabric is left without an FCS switch. |
| Active policy with multiple entries | A primary FCS switch and one or more backup FCS switches are designated. If the primary FCS switch becomes unavailable, the next switch in the list becomes the primary FCS switch.              |

## FCS Policy Restrictions

The backup FCS switches usually cannot modify the policy. However, if the primary FCS switch in the policy list is not reachable, a backup FCS switch is allowed to modify the policy.

When an FCS policy is configured and distributed across the fabric, only the primary FCS switch can perform certain operations. Operations that affect fabric-wide configuration are allowed only from the primary FCS switch. Backup and non-FCS switches cannot perform security, zoning, and AD operations that affect the fabric configuration. The following error message is returned if a backup or non-FCS switch tries to perform these operations:

Can only execute this command on the Primary FCS switch.

Operations that do not affect the fabric configuration, such as `show` or `local switch` commands, are allowed on backup and non-FCS switches.

FCS enforcement applies only for user-initiated fabric-wide operations. Internal fabric data propagation because of a fabric merge is not blocked. Consequently, a new switch that joins the FCS-enabled fabric could still propagate the AD and zone database.

The table below shows the commands for primary FCS enforcement switch operations.

**Table 53: FCS Switch Operations**

| Allowed on FCS Switches                                                                                          | Allowed on All Switches                                                                                                                           |
|------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>secPolicyAdd</code> (Allowed on all switches for SCC and DCC policies as long as it is not fabric-wide)    | <code>secPolicyShow</code>                                                                                                                        |
| <code>secPolicyCreate</code> (Allowed on all switches for SCC and DCC policies as long as it is not fabric-wide) | <code>fddCfg --localaccept</code> or <code>fddCfg --localreject</code>                                                                            |
| <code>secPolicyDelete</code> (Allowed on all switches for SCC and DCC policies as long as its not fabric-wide)   | <code>userconfig</code> , <code>Passwd</code> , <code>Passwdcfg</code> (Fabric-wide distribution is not allowed from a backup or non-FCS switch.) |
| <code>secPolicyRemove</code> (Allowed on all switches for SCC and DCC policies as long as its not fabric-wide)   | <code>secPolicyActivate</code>                                                                                                                    |
| <code>fddCfg -- fabwideset</code>                                                                                | <code>secPolicySave</code>                                                                                                                        |
| Any fabric-wide commands                                                                                         | <code>secPolicyAbort</code>                                                                                                                       |
| All zoning commands except the <code>show</code> commands                                                        | SNMP commands                                                                                                                                     |
| All AD commands                                                                                                  | <code>configupload</code>                                                                                                                         |
| —                                                                                                                | Any local-switch commands                                                                                                                         |

## Ensuring That the Fabric Domains Share the Same Policy

To create new FCS policies or manage your current FCS policies, you must follow certain steps to ensure that the domains throughout your fabric have the same policy.

The local switch WWN cannot be deleted from the FCS policy.

1. Create the FCS policy using the `secPolicyCreate` command.
2. Activate the policy using the `secPolicyActivate` command.

If the command is not entered, the changes are lost when the session is logged out.

3. Distribute the policy using the `distribute -p` command.

```
distribute -p policy_list -d switch_list
```

You can specify an asterisk (\*) for the *switch\_list* to send the policy to all switches.

4. Set the FCS policy as a fabric-wide policy using the `fddcfg --fabwideset fcs` command.

This setting prevents the admin from creating a second FCS controller in the fabric.

## Creating an FCS Policy

1. Connect to the switch and log in either using an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Enter the `secPolicyCreate "FCS_POLICY" "2;4"` command.

Example of creating an FCS policy

The following example creates an FCS policy that allows a switch with domain ID 2 to become a primary FCS and domain ID 4 to become a backup FCS.

```
switch:admin> secpolicycreate "FCS_POLICY", "2;4"
FCS_POLICY has been created
```

3. To save or activate the new policy, enter either the `secPolicySave` or the `secPolicyActivate` command. After the policy is activated, you can distribute the policy.

### NOTE

FCS policy must be consistent across the fabric. If the policy is inconsistent in the fabric, you will not be able to perform any fabric-wide configurations from the primary FCS.

## Modifying the Order of FCS Switches

1. Log in to the Primary FCS switch using an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Type `secPolicyShow "Defined", "FCS_POLICY"`.

This command displays the WWNs of the current Primary FCS switch and backup FCS switches.

3. Enter the `secPolicyFCSMove` command.

```
secpolicyfcsmove from to
```

Specify the existing and new positions of the switch in the list using the *from* and *to* options. Alternatively, you can specify the positions interactively at the prompts.

Example of moving an FCS policy.

The following example moves a backup FCS switch from position 2 to position 3 in the FCS list, using interactive mode:

```
primaryfcs:admin> secpolicyfcsmove
Pos Primary WWN DId swName.
=====
1 Yes 10:00:00:60:69:10:02:18 1 switch5.
2 No 10:00:00:60:69:00:00:5a 2 switch60.
3 No 10:00:00:60:69:00:00:13 3 switch73.
Please enter position you'd like to move from : (1..3) [1] 2
Please enter position you'd like to move to : (1..3) [1] 3
```

---

DEFINED POLICY SET

| FCS_POLICY |         |                         |     |           |
|------------|---------|-------------------------|-----|-----------|
| Pos        | Primary | WWN                     | Did | swName    |
| 1          | Yes     | 10:00:00:60:69:10:02:18 | 1   | switch5.  |
| 2          | No      | 10:00:00:60:69:00:00:13 | 3   | switch73. |
| 3          | No      | 10:00:00:60:69:00:00:5a | 2   | switch60. |

4. Enter the `secPolicyActivate` command to activate and save the new order.

## FCS Policy Distribution

The FCS policy can be automatically distributed using the `fddCfg --fabwideset` command or it can be manually distributed to the switches using the `distribute -p` command. Each switch that receives the FCS policy must be configured to receive the policy. To configure the switch to accept distribution of the FCS policy, refer to [Database Distribution Settings](#).

Database distributions may be initiated from only the Primary FCS switch. FCS policy configuration and management is performed using the command line or a manageability interface.

Only the Primary FCS switch is allowed to distribute the database. The FCS policy can be manually distributed across the fabric using the `distribute -p` command.

FCS enforcement for the `distribute` command is handled differently for FCS and other databases in an FCS fabric:

- For an FCS database, the enforcement allows any switch to initiate the distribution, which supports FCS policy creation specifying a remote switch as Primary.
- For other database distributions, only the Primary FCS switch can initiate the distribution.

The FCS policy distribution can be distributed from a switch in the FCS list. However, if none of the FCS switches in the existing FCS list are reachable, receiving switches accept distribution from any switch in the fabric. To learn more about how to distribute policies, refer to [ACL Policy Distribution to Other Switches](#).

Local switch configuration parameters are needed to control whether a switch accepts or rejects distributions of FCS policy and whether the switch can initiate distribution of an FCS policy. A configuration parameter controls whether the distribution of the policy is accepted or rejected on the local switch. Setting the configuration parameter to accept indicates distribution of the policy will be accepted and distribution may be initiated using the `distribute -p` command. Setting the configuration parameter to reject indicates the policy distribution is rejected, and the switch may not distribute the policy.

The default value for the distribution configuration parameter is *accept*, which means the switch accepts all database distributions and is able to initiate a distribute operation for all databases.

**Table 54: Distribution Policy States**

| Fabric OS            | State                                                                                                                                             |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Configured to accept | Target switch accepts distribution and a fabric state change occurs.                                                                              |
| Configured to reject | Target switch explicitly rejects the distribution, and the operation fails. The entire transaction is aborted, and no fabric state change occurs. |

## Device Connection Control Policies

Multiple device connection control (DCC) policies restrict the device ports that can connect to the switch ports. The devices can be initiators, targets, or intermediate devices, such as SCSI routers and loop hubs. By default, all device ports

can connect to all switch ports; no DCC policies exist until they are created. For information regarding DCC policies and F\_Port trunking, refer to the *Brocade Fabric OS Access Gateway User Guide*.

Each device port can be bound to one or more switch ports; the same device ports and switch ports may be listed in multiple DCC policies. After a switch port is specified in a DCC policy, it permits connections only from designated device ports. Device ports that are not specified in a DCC policy can connect only to switch ports that are not specified in any DCC policies.

When a DCC violation occurs, the related port is automatically disabled and must be re-enabled using the `portEnable` command.

The following table displays the possible DCC policy states:

**Table 55: DCC Policy States**

| Policy state           | Characteristics                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No policy              | Any device can connect to any switch port in the fabric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Policy with no entries | Any device can connect to any switch port in the fabric. An empty policy is the same as no policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Policy with entries    | <p>If a device WWN or Fabric port WWN is specified in a DCC policy, then it only allows the access to the switch if connected by a switch port listed in the same policy.</p> <p>If a switch port is specified in a DCC policy, it only permits connections from devices that are listed in the policy.</p> <p>Devices with WWNs that are not specified in a DCC policy are allowed to connect to the switch at any switch ports that are not specified in a DCC policy.</p> <p>Switch ports and device WWNs may exist in multiple DCC policies.</p> <p>Proxy devices are always granted full access and can connect to any switch port in the fabric.</p> |

## Virtual Fabrics Considerations for Device Connection Control Policies

The DCC policies that have entries for the ports that are being moved from one logical switch to another will be considered *stale* and will not be enforced. You can choose to keep *stale* policies in the current logical switch or delete the *stale* policies after the port movements. Use the `secPolicyDelete` command to delete stale DCC policies.

## DCC Policy Restrictions

The following restrictions apply when using DCC policies:

- Some older private-loop host bus adapters (HBAs) do not respond to port login from the switch and are not enforced by the DCC policy. This restriction does not create a security problem because these HBAs cannot contact any device outside of their immediate loop.
- DCC policies cannot manage or restrict iSCSI connections, that is, an FC Initiator connection from an iSCSI gateway.
- You cannot manage proxy devices with DCC policies. Proxy devices are always granted full access, even if the DCC policy has an entry that restricts or limits access of a proxy device.

## Creating a Device Connection Control Policy

DCC policies must follow the naming convention `DCC_POLICY_ nnn`, where *nnn* represents a unique string. The maximum length is 30 characters, including the prefix `DCC_POLICY_`.

Device ports must be specified by port WWN. Switch ports can be identified by the switch WWN, domain ID, or switch name followed by the port or area number. To specify an allowed connection, enter the device port WWN, a semicolon, and the switch port identification.

The following methods of specifying an allowed connection are possible:

- *deviceportWWN;switchWWN* (port or area number)
  - *deviceportWWN;domainID* (port or area number)
  - *deviceportWWN;switchname* (port or area number)
1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
  2. Enter `secpolicycreate "DCC_policyname"` to create a new policy.
  3. To save or activate the new policy, enter the appropriate command:
    - To save the policy, enter `secpolicysave`.
    - To save and activate the policy, enter `secpolicyactivate`.

If neither of these commands is entered, your changes are lost when the session is logged out.

### Examples of Creating DCC Policies

To create the DCC policy "DCC\_POLICY\_server" that includes device 11:22:33:44:55:66:77:aa and port 1 and port 3 of switch domain 1:

```
switch:admin> secpolicycreate "DCC_POLICY_server", "11:22:33:44:55:66:77:aa;1(1,3)"
DCC_POLICY_server
has been created
```

To create the DCC policy "DCC\_POLICY\_storage" that includes device port WWN 22:33:44:55:66:77:11:bb, all ports of switch domain 2, and all currently connected devices of switch domain 2:

```
switch:admin> secpolicycreate "DCC_POLICY_storage", "22:33:44:55:66:77:11:bb;2[*]"
DCC_POLICY_storage
has been created
```

To create the DCC policy "DCC\_POLICY\_abc" that includes device 33:44:55:66:77:11:22:cc and ports 1 through 6 and port 9 of switch domain 3:

```
switch:admin> secpolicycreate "DCC_POLICY_abc", "33:44:55:66:77:11:22:cc;3(1-6,9)"
DCC_POLICY_abc
has been created
```

To create the DCC policy "DCC\_POLICY\_example" that includes devices 44:55:66:77:22:33:44:dd and 33:44:55:66:77:11:22:cc, ports 1 through 4 of switch domain 4, and all devices currently connected to ports 1 through 4 of switch domain 4:

```
switch:admin> secpolicycreate "DCC_POLICY_example",
"44:55:66:77:22:33:44:dd;33:44:55:66:77:11:22:cc;4[1-4]"
DCC_POLICY_example
has been created
```

## Deleting a DCC Policy

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Enter the `secPolicyDelete` command.

Example of deleting stale DCC policies.

```
switch:admin> secpolicydelete ALL_STALE_DCC_POLICY
```



```
About to clear all STALE DCC policies
ARE YOU SURE (yes, y, no, n): [no] y
```

## DCC Policy Behavior with Fabric-Assigned PWWNs

A DCC policy check is always performed for the physical port WWN of a device when the HBA has established that the device is attempting a normal FLOGI and has both a fabric-assigned port WWN (FA-PWWN) and a physical port WWN.

DCC policies created with FA-PWWNs will result in the disabling of FA-PWWN assigned ports on subsequent FLOGI. Create policies with the physical PWWN

DCC policies created with the lockdown feature result in DCC policies with FA-PWWNs. Avoid using the lockdown feature in fabrics that are using FA-PWWNs.

A DCC policy created with a device WWN for a specific port allows the device to log in only on the same port. The same device will not be allowed to log in on a different port. For devices that log in across an Access Gateway (AG), the policy should be created with all the NPIV ports, so even if failover occurs, the device will be allowed to log in on a different NPIV port.

The table below lists the behavior of the DCC policy with FA-PWWNs in the fabric when the DCC policy is created using lockdown support.

**Table 56: DCC Policy Behavior with FA-PWWN when Created using Lockdown Support**

| Configuration                                                                                                                                                                          | WWN Seen on DCC Policy List | Behavior when DCC Policy Activates     | Behavior on portDisable and portEnable                  |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|----------------------------------------|---------------------------------------------------------|
| <ul style="list-style-type: none"> <li>FA-PWWN has logged into the switch.</li> <li>DCC policy creation with lockdown (uses FA-PWWN).</li> <li>DCC policy activation.</li> </ul>       | FA-PWWN                     | Traffic is not disrupted. <sup>a</sup> | Ports are disabled for security violation. <sup>b</sup> |
| <ul style="list-style-type: none"> <li>DCC policy creation with lockdown (uses physical PWWN).</li> <li>FA-PWWN has logged into the switch.</li> <li>DCC policy activation.</li> </ul> | Physical PWWN               | Traffic is not disrupted.              | Ports come up without security issues.                  |
| <ul style="list-style-type: none"> <li>DCC policy creation with lockdown (uses physical PWWN).</li> <li>DCC policy activation.</li> <li>FA-PWWN has logged into the switch.</li> </ul> | Physical PWWN               | Traffic is not disrupted.              | Ports come up without security issues.                  |

The table below shows the behavior of a DCC policy created manually with the physical PWWN of a device. The configurations shown in this table are the recommended configurations when an FA-PWWN is logged into the switch.

a. Indicates a security concern, because devices that are logged in with FA-PWWNs will not be disabled after activation of DCC policies that are created with FA-PWWNs, which is done to avoid disturbing any existing management.

b. Any disruption in the port will disable the port for a security violation. As the traffic is already disrupted for this port, you must enforce the DCC policy for a physical device WWN; otherwise, the device will not be allowed to log in again.

**Table 57: DCC Policy Behavior when Created Manually with PWWN**

| Configuration                                                                                                                                                                             | WWN Seen on DCC Policy List | Behavior when DCC Policy Activates | Behavior on portDisable and portEnable |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|------------------------------------|----------------------------------------|
| <ul style="list-style-type: none"> <li>FA-PWWN has logged into the switch.</li> <li>DCC policy creation manually with physical PWWN of device.</li> <li>DCC policy activation.</li> </ul> | PWWN                        | Traffic is not disrupted.          | Ports come up without security issues. |
| <ul style="list-style-type: none"> <li>DCC policy creation manually with physical PWWN.</li> <li>FA-PWWN has logged into the switch.</li> <li>DCC policy activation.</li> </ul>           | PWWN                        | Traffic is not disrupted.          | Ports come up without security issues. |
| <ul style="list-style-type: none"> <li>DCC policy creation manually with physical PWWN.</li> <li>DCC policy activation.</li> <li>FA-PWWN has logged into the switch.</li> </ul>           | Physical PWWN               | Traffic is not disrupted.          | Ports come up without security issues. |

**NOTE**

FCoE devices are excluded from DCC control.

## SCC Policies

The switch connection control (SCC) policy restricts which switches can access each other. Only one SCC policy (named SCC\_POLICY) can be created per switch and each switch might have a different SCC policy. SCC\_POLICY does not exist until it is created and it accepts members listed as WWNs, domain IDs, or switch names. Switches are checked against this policy whenever an E\_Port-to-E\_Port connection is made.

When connecting a Fibre Channel router to a fabric or switch that has an active SCC policy, the front domain switch WWN of the Fibre Channel router must be included in the SCC policy. You can obtain the front domain switch WWN from the RASLog SEC-1187 on the edge fabric switch, which is generated when the IFL link to the Fibre Channel router is disabled upon connecting to the edge fabric.

If you are configuring the fabric for the first time, do the following to obtain the domain WWN:

1. Delete the SCC policy from the edge fabric.
2. Save and distribute the SCC policy state across all the switches in the fabric.
3. Configure an IFL link between the edge fabric and the new backbone switch.
4. Verify that the IFL link is online. (The front domain switch WWN is available only if the link is online.)
5. Do one of the following:
  - Use the `switchshow | grep fcr` command to obtain the front domain switch WWN from the edge fabric switch that has an IFL from the backbone switch connected to itself.

```
switch:admin> switchshow | grep fcr
12 12 010c00 id N8 Online FC E-Port 50:00:53:36:5a:60:4e:05 "fcr_fd_160" (Trunk
master)
```

- Use the `fabricshow` command.

```
switch:admin> fabricshow
```

| Switch ID   | Worldwide Name                     | Enet IP Addr | FC IP Addr     | Name             |
|-------------|------------------------------------|--------------|----------------|------------------|
| 1: fffc01   | 10:00:00:05:33:e5:58:2b            | 10.1.18.14   | 0.0.0.0        | "6520_14"        |
| 74: fffc4a  | 10:00:00:05:33:13:9d:da            | 10.1.18.25   | 220.220.220.10 | "6510_25"        |
| 98: fffc62  | 10:00:00:05:1e:b7:10:00            | 10.1.18.30   | 0.0.0.0        | >"Pluto_30"      |
|             | 2620:100:4:f400:205:1eff:feb7:1000 |              |                |                  |
| 160: fffca0 | 50:00:53:36:5a:60:4e:05            | 0.0.0.0      | 0.0.0.0        | "fcr_fd_160"     |
| 161: fffca1 | 50:00:51:e4:4d:20:0e:05            | 0.0.0.0      | 0.0.0.0        | "fcr_fd_161"     |
| 200: fffcc8 | 50:00:51:e4:4d:32:0f:1a            | 0.0.0.0      | 0.0.0.0        | "fcr_xd_200_12"  |
| 201: fffcc9 | 50:00:53:36:5a:64:4f:01            | 0.0.0.0      | 0.0.0.0        | "fcr_xd_201_105" |
| 202: fffcca | 50:00:51:e4:4d:32:0f:44            | 0.0.0.0      | 0.0.0.0        | "fcr_xd_202_114" |
| 203: fffccb | 50:00:51:e4:4d:32:0f:43            | 0.0.0.0      | 0.0.0.0        | "fcr_xd_203_9"   |

The Fabric has 9 switches  
Fabric Name: PROWLER\_EDGE\_5

- Use the `portcfgexport` command on the port that connects to the edge fabric switch.

```
switch:admin> switchshow | grep EX-P
```

| Port | Port ID | Port Name | Port Type | Port Status | Port Mode | Port Fabric | Port Fabric ID | Port Fabric Name                                                   |
|------|---------|-----------|-----------|-------------|-----------|-------------|----------------|--------------------------------------------------------------------|
| 4    | 4       | 1a0500    | id        | N8          | Online    | FC          | EX-Port        | 10:00:00:05:33:e5:58:2b "6520_14" (fabric id = 5 ) (Trunk master)  |
| 5    | 5       | 1a0800    | id        | N8          | Online    | FC          | EX-Port        | (Trunk port, master is Port 4 )                                    |
| 6    | 6       | 1a0700    | id        | N8          | Online    | FC          | EX-Port        | (Trunk port, master is Port 4 )                                    |
| 7    | 7       | 1a0600    | id        | N8          | Online    | FC          | EX-Port        | (Trunk port, master is Port 4 )                                    |
| 12   | 12      | 1a0b00    | id        | N8          | Online    | FC          | EX-Port        | 10:00:00:05:1e:b7:10:00 "Pluto_30" (fabric id = 5 ) (Trunk master) |

```
switch:admin> portcfgexport 4
```

| Port                         | 4 | info                    |
|------------------------------|---|-------------------------|
| Admin:                       |   | enabled                 |
| State:                       |   | OK                      |
| Pid format:                  |   | core(N)                 |
| Operate mode:                |   | Brocade Native          |
| Edge Fabric ID:              |   | 5                       |
| Front Domain ID:             |   | 160                     |
| Front WWN:                   |   | 50:00:53:36:5a:60:4e:05 |
| Principal Switch:            |   | 98                      |
| Principal WWN:               |   | 10:00:00:05:1e:b7:10:00 |
| Fabric Parameters:           |   | Auto Negotiate          |
| R_A_TOV:                     |   | 10000 (N)               |
| E_D_TOV:                     |   | 2000 (N)                |
| Authentication Type:         |   | None                    |
| DH Group:                    |   | N/A                     |
| Hash Algorithm:              |   | N/A                     |
| Encryption:                  |   | OFF                     |
| Compression:                 |   | OFF                     |
| Forward Error Correction:    |   | OFF                     |
| Edge fabric's primary wwn:   |   | N/A                     |
| Edge fabric's version stamp: |   | N/A                     |

## 6. Do either of the following:

- Manually add the front domain switch WWN to the SCC policy.
- Use the `secpolicycreate SCC_POLICY "*"` command to automatically add all switches in the fabric (including the front and xlate domains).

7. Activate and distribute the SCC policy.

SCC policy states are shown in the table below.

**Table 58: SCC Policy States**

| Policy State                      | SCC Policy Enforcement                                                  |
|-----------------------------------|-------------------------------------------------------------------------|
| No active policy                  | All switches can connect to the switch with the specified policy.       |
| Active policy that has no members | All neighboring switches are segmented.                                 |
| Active policy that has members    | The neighboring switches not specified in the SCC policy are segmented. |

## Virtual Fabrics Considerations for SCC Policies

In a logical fabric environment, the SCC policy enforcement is not done on the logical ISL. For a logical ISL-based switch, the SCC policy enforcement is considered as the reference, and the logical ISL is formed if the SCC enforcement passes on the extended ISL. The following changes occur:

- A logical switch supports an SCC policy. You can configure and distribute an SCC policy on a logical switch.
- SCC enforcement is performed on a ISL, based on the SCC policy present on the logical switch.

For more information on Virtual Fabrics, see [Managing Virtual Fabrics](#).

## Creating an SCC Policy

1. Connect to the switch and log in either using an account with admin permissions, or an account with OM permissions for the Security RBAC class of commands.
2. Enter the `secPolicyCreate "SCC_POLICY"` command.
3. Save or activate the new policy by entering either the `ssecPolicySave` or the `secPolicyActivate` command.

If neither of these commands is entered, the changes are lost when the session is logged out.

Example of creating an SCC policy.

For example, to create an SCC policy that allows switches that have domain IDs 2 and 4 to join the fabric:

```
switch:admin> secpolicycreate "SCC_POLICY", "2;4"
SCC_POLICY has been created
switch:admin> secpolicysave
```

## Authentication Policy for Fabric Elements

By default, Fabric OS use Diffie Hellman - Challenge Handshake Authentication Protocol) (DH-CHAP) or Fibre Channel Authentication Protocol (FCAP) for authentication.

These protocols use shared secrets and digital certificates, based on switch WWN and public key infrastructure (PKI) technology, to authenticate switches. Authentication automatically defaults to FCAP if both switches are configured to accept FCAP protocol in authentication, unless ports are configured for in-flight encryption, in which case authentication

defaults to DH-CHAP if both switches are configured to accept the DH-CHAP protocol in authentication. To use FCAP on both switches, PKI certificates must be installed.

The DH-CHAP and FCAP authentication protocols used by Brocade switches are FC-SP2 standard compliant.

#### NOTE


The Fabric Authentication feature is available in base Fabric OS. No license is required.

When you configure DH-CHAP authentication, you also must define a *pair of shared secrets* known to both switches as a *secret key pair*. The figure below illustrates how the secrets are configured. A *secret key pair* consists of a local secret and a peer secret. The local secret uniquely identifies the local switch. The peer secret uniquely identifies the entity to which the local switch authenticates. Every switch can share a *secret key pair* with any other switch or host in a fabric.

To use DH-CHAP authentication, a *secret key pair* has to be configured on both switches. For more information on setting up secret key pairs, see [Setting a Secret Key Pair](#).

When configured, the *secret key pair* is used for authentication. During a state change, authentication occurs for the switch or port. The state change can be due to a switch reboot, a switch or port disable and enable, or the activation of a policy.

**Figure 7: DH-CHAP Authentication**



If you use DH-CHAP authentication, a *secret key pair* must be installed only in connected fabric elements. However, as connections are changed, new *secret key pairs* must be installed between newly connected elements. Alternatively, a *secret key pair* for all possible connections may be initially installed, enabling links to be arbitrarily changed while still maintaining a valid *secret key pair* for any new connection.

The switch authentication (AUTH) policy initiates DH-CHAP/FCAP authentication on all E\_Ports. This policy is persistent across reboots, which means authentication will be initiated automatically on ports or switches brought online if the policy is set to activate authentication. The AUTH policy is distributed by command; automatic distribution of the AUTH policy is not supported.

The default configuration directs the switch to attempt FCAP authentication first, DH-CHAP second. The switch may be configured to negotiate FCAP, DH-CHAP, or both.

The DH group is used in the DH-CHAP protocol only. The FCAP protocol exchanges the DH group information, but does not use it.

## Virtual Fabrics Considerations for Authentication Policies

If Virtual Fabrics is enabled, all AUTH module parameters, such as shared secrets and shared switch and device policies, are logical switch-wide. Therefore, you must configure shared secrets and policies separately on each logical switch, and the shared secrets and policies must be set on each switch prior to authentication. During logical switch creation, authentication takes the default values for policies and other parameters. FCAP certificates are installed on a chassis but are configured on each logical switch.

## E\_Port Authentication

The authentication (AUTH) policy allows you to configure DH-CHAP authentication on Fabric OS switches. By default, the policy is set to PASSIVE, and you can change the policy. All changes to the AUTH policy take effect during the next authentication request. This includes starting authentication on all E\_Ports on the local switch if the policy is changed to ON or ACTIVE, and clearing the authentication if the policy is changed to OFF. The authentication configurations will be effective only on subsequent E\_Port and F\_Port initialization.

### ATTENTION

A *secret key pair* must be installed prior to changing the policy. For more information on setting up secret key pairs, refer to [Setting a Secret Key Pair](#).

If you must disable authentication on a port that has in-flight encryption or compression configured, you must first disable in-flight encryption or compression on the port, and then disable authentication.

## Virtual Fabrics Considerations for E\_Port Authentication

The switch authentication policy applies to all E\_Ports in a logical switch, which includes ISLs and extended ISLs. Authentication of extended ISLs between two base switches is considered peer-chassis authentication. Authentication between two physical entities is required, so the extended ISL that connects the two chassis must be authenticated. The corresponding extended ISL for a logical ISL authenticates the peer-chassis; therefore, the logical ISL authentication is not required. Because the logical ISLs do not carry actual traffic, they do not need to be authenticated. Authentication on re-individualization is also blocked on logical ISLs. The following error message is printed on the console when you execute the `authUtil --authinit` command on logical-ISLs:

```
Failed to initiate authentication. Authentication is not supported on logical ports
<port#>.
```

For more information on Virtual Fabrics, refer to [Managing Virtual Fabrics](#).

## Configuring E\_Port Authentication

1. Connect to the switch and log in either using an account with admin permissions, or an account with OM permissions for the Authentication RBAC class of commands.
2. Enter the `authUtil` command to set the switch policy mode.

### Example of Configuring E\_Port Authentication

The following example shows how to enable Virtual Fabrics and configure the E\_Ports to perform authentication using the AUTH policies `authUtil` command.

```
switch:admin> fosconfig -enable vf
WARNING: This is a disruptive operation that requires a reboot to take effect.
All EX ports will be disabled upon reboot.
Would you like to continue [Y/N] y
switch:admin> authutil --authinit 2,3,4
```

**CAUTION**

If data input has not been completed and a failover occurs, the command is terminated without completion and your entire input is lost. If data input has completed, the enter key pressed, and a failover occurs, data may not be replicated to the other CP depending on the timing of the failover. Log in to the other CP after the failover is complete, and verify the data was saved. If data was not saved, run the command again.

**Example of Setting the Policy to Active Mode**

```
switch:admin> authutil --policy -sw active
Warning: Activating the authentication policy requires
either DH-CHAP secrets or PKI certificates depending
on the protocol selected. Otherwise, ISLs will be
segmented during next E-port bring-up.
ARE YOU SURE (yes, y, no, n): [no] y
Auth Policy is set to ACTIVE
```

**NOTE**

This authentication-policy change will not affect online EX\_Ports.

**Re-Authenticating E\_Ports**

Use the `authUtil --authinit` command to re initiate the authentication on selected ports. It provides flexibility to initiate authentication for specified E\_Ports, a set of E\_Ports, or all E\_Ports on the switch. This command does not work on loop, NPIV, and FICON devices, or on ports configured for in-flight encryption. The `authUtil` command can re initiate authentication only if the device was previously authenticated. If the authentication fails because shared secrets do not match, the port is disabled.

This command works independently of the authentication policy, which means you can initiate the authentication even if the switch is in PASSIVE mode. This command restarts authentication after changing the DH-CHAP group, hash type, or shared secret between a pair of switches.

**ATTENTION**

This command may bring down E\_Ports if the DH-CHAP shared secrets are not installed correctly.

1. Log in to the switch either using an account with admin permissions, or an account with OM permissions for the Authentication RBAC class of commands.
2. Enter the `authUtil --authinit` command.

**Example for Specific Ports on the Switch**

```
switch:admin> authutil --authinit 2,3,4
```

**Example for All E\_Ports on the Switch**

```
switch:admin> authutil --authinit allE
```

**Example for Directors Using the Slot/Port Format**

```
switch:admin> authutil --authinit 1/1, 1/2
```

## Device Authentication Policy

Device authentication policy can also be categorized as an F\_Port, a node port, or an HBA authentication policy. Fabric-wide distribution of the device authentication policy is not supported because the device authentication requires manual interaction in setting the HBA shared secrets and switch shared secrets, and most of the HBAs do not support the defined DH groups for use in the DH-CHAP protocol.

### NOTE

Authentication is supported from Brocade fabric switches in native mode to Access Gateway switches and from Access Gateway switches to HBAs. For more information, refer to the *Brocade Fabric OS Access Gateway User Guide*.

By default, the device policy is in the OFF state, which means the switch clears the security bit in the FLOGI (fabric login). The `authUtil` command provides an option to change the device policy mode to select PASSIVE policy, which means the switch responds to authentication from any device and does not initiate authentication to devices.

When the policy is set to ON, the switch expects an FLOGI with the Fibre Channel Security Protocols (FC-SP) bit set. If not, the switch rejects the FLOGI with reason `LS_LOGICAL_ERROR (0x03)`, explanation `Authentication Required (0x48)`, and disables the port. Regardless of the policy, the F\_Port is disabled if the DH-CHAP protocol fails to authenticate.

If the HBA sets the FC-SP bit during FLOGI and the switch sends an FLOGI accept with the FC-SP bit set, the switch expects the HBA to start the AUTH\_NEGOTIATE. From this point on, until the AUTH\_NEGOTIATE is completed, all ELS and CT frames, except the AUTH\_NEGOTIATE ELS frame, are blocked by the switch. During this time, the Fibre Channel driver rejects all other ELS frames. The F\_Port does not form until the AUTH\_NEGOTIATE is completed. It is the HBA's responsibility to send an Authentication Negotiation ELS frame after receiving the FLOGI accept frame with the FC-SP bit set.

## Virtual Fabrics Considerations for the Authentication Policy

Because the device authentication policy has switch and logical switch-based parameters, each logical switch is set when Virtual Fabrics is enabled. Authentication is enforced based on each logical switch's policy settings.

## Configuring Device Authentication

1. Connect to the switch and log in using an account with admin permissions or an account with OM permissions for the Authentication RBAC class of commands.
2. Enter the `authUtil` command to set the device policy mode.

### Example of Setting the Device Policy to Passive Mode

```
switch:admin> authutil --policy -dev passive
Warning: Activating the authentication policy requires
DH-CHAP secrets on both switch and device. Otherwise,
the F-port will be disabled during next F-port
bring-up.
ARE YOU SURE (yes, y, no, n): [no] y
Device authentication is set to PASSIVE
```

## AUTH Policy Restrictions

All fabric element authentication configurations are performed on a local switch basis.

Device authentication policy supports devices that are connected to the switch in a point-to-point manner and it is visible to the entire fabric. The following devices are not supported:



- Public loop devices
- Single private devices
- Private loop devices
- Mixed public and private devices in loop
- NPIV devices
- FICON channels
- Config upload and download are not supported for the following AUTH attributes: auth type, hash type, group type.

**NOTE**

For information about how to use authentication with Access Gateway, refer to the *Brocade Fabric OS Access Gateway User Guide*.

## Authentication Protocols

Use the `authUtil` command to perform the following tasks:

- Display the current authentication parameters.
- Select the authentication protocol used between switches.
- Select the Diffie-Hellman (DH) group for a switch.

Run the `authUtil` command on the switch you want to view or change. Use the following options to specify which DH group you want to use:

- **00** – DH Null option
- **01** – 1024 bit key
- **02** – 1280 bit key
- **03** – 1536 bit key
- **04** – 2048 bit key

## Viewing the Current Authentication Parameter Settings for a Switch

1. Log in to the switch using either an account with admin permissions, or an account with the O permission for the Authentication RBAC class of commands.
2. Enter the `authUtil -- show` command.

Example of output from the `authUtil -- show` command.

```

AUTH TYPE HASH TYPE GROUP TYPE

fcap,dhchap sha1,md5 0, 1, 2, 3, 4
Switch Authentication Policy: PASSIVE
Device Authentication Policy: OFF

```

## Setting the Authentication Protocol

1. Log in to the switch using an account with admin permissions, or an account with OM permissions for the Authentication RBAC class of commands.
2. Enter the `authUtil --set -a` command specifying `fcap`, `dhchap`, or `all`.

Example of setting the DH-CHAP authentication protocol.

```
switch:admin> authutil --set -a dhchap
```

Authentication is set to dhchap.

When using DH-CHAP, make sure that you configure the switches at both ends of a link.

#### NOTE

When you set the authentication protocol to FCAP, ensure that the certificates are present at both ends.

#### NOTE

If you set the authentication protocol to DH-CHAP or FCAP, have not configured shared secrets or certificates, and authentication is checked (for example, you enable the switch), switch authentication will fail. You must also use the `-g` option to set the DH group value to group 4 or all groups.

```
switch# authutil --set -g *
```

or

```
switch# authutil --set -g 4
```

## Secret Key Pairs for DH-CHAP

When you configure the switches at both ends of a link to use DH-CHAP for authentication, you must also define a secret key pair—one for each end of the link. Use the `secAuthSecret` command to perform the following tasks:

- View the WWN of switches with a secret key pair.
- Set the secret key pair for switches.
- Remove the secret key pair for one or more switches.

#### NOTE

The DH-CHAP secrets are stored using AES-256 encryption.

## Characteristics of a Secret Key Pair

- The secret key pair must be set up locally on every switch. The secret key pair is not distributed fabric-wide.
- If a secret key pair is not set up for a link, authentication fails. The `Authentication Failed (reason code 05h)` error will be reported and logged.
- The minimum length of a shared secret is 8 characters and the maximum length is 40 characters. If the E\_Port is to carry in-flight encrypted traffic, use a shared secret of at least 32 characters. The value that you specify cannot have spaces. All other characters are allowed.

#### NOTE

When setting a secret key pair, note that you are entering the shared secrets in plain text. Use a secure channel (for example, SSH or the serial console) to connect to the switch on which you are setting the secrets.

## Viewing the List of Secret Key Pairs in the Current Switch Database

1. Log in to the switch either using an account with admin permissions, or an account with the O permission for the Authentication RBAC class of commands.
2. Enter the `secAuthSecret --show` command.

The output displays the WWN, domain ID, and name (if known) of the switches with defined shared secrets.

| WWN                     | DId | Name    |
|-------------------------|-----|---------|
| 10:00:00:60:69:80:07:52 |     | Unknown |
| 10:00:00:60:69:80:07:5c | 1   | switchA |

## Note about Access Gateway Switches

Because Domain ID and name are not supported for Access Gateway, `secAuthSecret --show` output for Access Gateway appears as follows.

| WWN                     | DId   | Name    |
|-------------------------|-------|---------|
| -----                   | ----- | -----   |
| 10:00:8C:7C:FF:03:9E:00 | -1    | Unknown |
| 10:00:8C:7C:FF:03:9E:01 | -1    | Unknown |
| 10:00:8C:7C:FF:0D:AF:01 | -1    | Unknown |

When setting and removing the secret for a switch or device on Access Gateway, only the WWN can be used.

## Setting a Secret Key Pair

1. Log in to the switch either using an account with admin permissions, or an account with OM permissions for the Authentication RBAC class of commands.
2. Enter the `secAuthSecret --set` command.

The command enters interactive mode. The command returns a description of itself and needed input; then it loops through a sequence of switch specification, peer secret entry, and local secret entry.

### NOTE

The value that you specify should not contain spaces. All the other characters are allowed.

### NOTE

From Fabric OS 9.0, the secrets must not be the same. To avoid link disruption after upgrading from pre-9.0.0 to 9.0.0, secrets must be swapped for HBA or AG devices.

To exit the loop, press **Enter** for the switch name; then type **y**.

Example of setting a secret key pair

```
switchA:admin> secauthsecret --set
This command is used to set up secret keys for the DH-CHAP authentication.
The minimum length of a secret key is 12 characters and maximum 40
characters. Setting up secret keys does not initiate DH-CHAP
authentication. If switch is configured to do DH-CHAP, it is performed
whenever a port or a switch is enabled.
Warning: Please use a secure channel for setting secrets. Using
an insecure channel is not safe and may compromise secrets.
Following inputs should be specified for each entry.
1. WWN for which secret is being set up.
2. Peer secret: The secret of the peer that authenticates to peer.
3. Local secret: The local secret that authenticates peer.
Press Enter to start setting up shared secrets > <cr>
Enter WWN, Domain, or switch name (Leave blank when done): 10:20:30:40:50:60:70:80
Enter peer secret: <hidden>
Re-enter peer secret: <hidden>
Enter local secret: <hidden>
Re-enter local secret: <hidden>
Enter WWN, Domain, or switch name (Leave blank when done): 10:20:30:40:50:60:70:81
Enter peer secret: <hidden>
Re-enter peer secret: <hidden>
```

```

Enter local secret: <hidden>
Re-enter local secret: <hidden>
Enter WWN, Domain, or switch name (Leave blank when done): <cr>
Are you done? (yes, y, no, n): [no] y
Saving data to key store... Done.

```

3. Disable and enable the ports on a peer switch using the `portDisable` and `portEnable` commands.

## FCAP Configuration Overview

Beginning with Fabric OS 7.0.0, you must configure the switch to use third-party certificates for authentication with the peer switch.

To perform authentication with FCAP protocol with certificates issued from third party, the user must perform following steps:

1. Choose a certificate authority (CA).
2. Generate a public key, a private key, a passphrase, and a CSR on each switch.
3. Store the CSR from each switch on a file server.
4. Obtain the certificates from the CA.

You can request a certificate from a CA through a Web browser. After you request a certificate, the CA either sends certificate files by email (public) or gives access to them on a remote host (private). Typically, the CA provides the certificate files listed in the following table.

### ATTENTION

Only the `.pem` file type is supported for FCAP authentication.

**Table 59: Certificate Files Provided by Certificate Authority**

| Certificate File         | Description                                                                                                                                               |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| name <code>CA.pem</code> | The CA root certificate. It must be installed on the remote and local switch to verify the validity of the switch certificate or switch validation fails. |
| name <code>.pem</code>   | The switch certificates: switch certificate.                                                                                                              |

5. On each switch, install the CA certificate before installing the switch certificate.
6. After the CA certificate is installed, install the switch certificate on each switch.
7. Update the switch database for peer switches to use third-party certificates.
8. Use the newly installed certificates by starting the authentication process.

## Generating the Key and CSR for FCAP

The public/private key and CSR has to be generated for the local and remote switches that will participate in the authentication. In FCAP, one command generates the public/private key the CSR, and the passphrase.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having OM permissions for the PKI RBAC class of commands.
2. Enter the `seccertmgmt generate -csr fcap` command on the local switch.

3. Repeat Step 2 on the remote switch.

## Exporting the CSR for FCAP

You will need to export the CSR file created in [Generating the Key and CSR for FCAP](#) section and send it to a certificate authority (CA). The CA will in turn provide two files as outlined in [FCAP Configuration Overview](#).

1. Log in to the switch either using an account with admin permissions or an account associated with the chassis role and having OM permissions for the PKI RBAC class of commands.
2. Enter the `seccertmgmt export -csr fcap` command.

## Importing the CA for FCAP

When you receive the files back from the Certificate Authority with the CA root certificate, you must install or import them onto the local and remote switches.

1. Log in to the switch either using an account with admin permissions or an account associated with the chassis role and having OM permissions for the PKI RBAC class of commands.
2. Enter the `seccertmgmt import -ca -client fcap` command, and verify that the CA certificates are consistent on both local and remote switches.

## Importing the FCAP Switch Certificate

### ATTENTION

The CA certificates must be installed prior to installing the FCAP switch certificate by using `seccertmgmt import -ca -client fcap` command.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having OM permissions for the PKI RBAC class of commands.
2. Enter the `seccertmgmt import -cert fcap` command.

## Starting FCAP Authentication

1. Log in to the switch either using an account with admin permissions, or an account with OM permissions for the Authentication RBAC class of commands.
2. Enter the `authUtil --authinit` command to start the authentication using the newly imported certificates (This command is not supported in Access Gateway mode.).
3. Enter the `authUtil --policy -sw` command with either the `active` or `on` option.

```
authutil --policy -sw active
```

This makes the changes permanent and forces the switch to request authentication. (For Access Gateway mode, the defaults for `sw policy` and `dev policy` are `off`, and there is no `passive` option for `sw policy`.)

### NOTE

This authentication-policy change does not affect online EX\_Ports.

## Fabric-Wide Distribution of the Authorization Policy

The AUTH policy can be manually distributed to the fabric using a command; there is no support for automatic distribution. To distribute the AUTH policy, refer to [Distributing the local ACL policies](#) for instructions.

Local Switch configuration parameters are needed to control whether a switch accepts or rejects distributions of the AUTH policy using the `distribute` command and whether the switch can initiate distribution of the policy. To set the local switch configuration parameter, refer to [Policy\\_database\\_distribution](#).

**NOTE**

This capability is not supported for Access Gateway mode.

## IP Filter Policy

The IP Filter policy is a set of rules applied to the IP management interfaces as a packet filtering firewall. The firewall permits or denies the traffic going through the IP management interfaces according to the policy rules.

Fabric OS supports multiple IP Filter policies to be defined at the same time. Each IP Filter policy is identified by a name and has an associated type. Two IP Filter policy types, IPv4 and IPv6, exist to provide separate packet filtering for IPv4 and IPv6. You cannot specify an IPv6 address in the IPv4 filter, nor can you specify an IPv4 address in the IPv6 filter. There can be up to six different IP Filter policies defined for both types. Only one IP Filter policy for each IP type can be activated on the affected management IP interfaces.

Audit messages are generated for any changes to the IP Filter policies.

The rules in the IP Filter policy are examined one at a time until the end of the list of rules. For performance reasons, specify the most commonly used rules at the top.

On a chassis system, changes to persistent IP Filter policies are automatically synchronized to the standby CP when the changes are saved persistently on the active CP. The standby CP will enforce the filter policies to its management interface after policies are synchronized with the active CP.

## Virtual Fabrics Considerations for IP Filter Policies

Each logical switch cannot have its own different IP Filter policies. IP Filter policies are treated as a chassis-wide configuration and are common for all the logical switches in the chassis.

## Creating an IP Filter Policy

You can create an IP Filter policy specifying any name and using type IPv4 or IPv6. The policy created is stored in a temporary buffer, and is lost if the current command session logs out. The policy name is a unique string composed of a maximum of 20 alphabetic, numeric, and underscore characters. The names `default_ipv4` and `default_ipv6` are reserved for default IP filter policies. The policy name is case-insensitive and always stored as lowercase. The policy type identifies the policy as an IPv4 or IPv6 filter. There can be a maximum of six IP Filter policies.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having OM permissions for the `ipfilter` RBAC class of commands.
2. Enter in the `ipfilter` command.

## Cloning an IP Filter Policy

You can create an IP Filter policy as an exact copy of an existing policy. The policy created is stored in a temporary buffer and has the same type and rules as the existing defined or active policy.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having OM permissions for the `ipfilter` RBAC class of commands.
2. Enter the `ipfilter --clone` command.

## Displaying an IP Filter Policy

You can display the IP Filter policy content for the specified policy name, or all IP Filter policies if a policy name is not specified.

For each IP Filter policy, the policy name, type, persistent state and policy rules are displayed. The policy rules are listed by the rule number in ascending order. There is no pagination stop for multiple screens of information. Pipe the output to the `|more` command to stop this.

If a temporary buffer exists for an IP Filter policy, the `--show` subcommand displays the content in the temporary buffer, with the persistent state set to `no`.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having the O permission for the `ipfilter` RBAC class of commands.
2. Enter the `ipfilter --show` command.

## Saving an IP Filter Policy

You can save one or all IP Filter policies persistently in the defined configuration.

Only the CLI session that owns the updated temporary buffer may run this command. Modification to an active policy cannot be saved without being applied. Hence, the `--save` subcommand is blocked for the active policies. Use `--activate` instead.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having the OM permissions for the `ipfilter` RBAC class of commands.
2. Enter the `ipfilter --save` command.

## Activating an IP Filter Policy

IP Filter policies are not enforced until they are activated. Only one IP Filter policy per IPv4 and IPv6 type can be active. If there is a temporary buffer for the policy, the policy is saved to the defined configuration and activated at the same time. If there is no temporary buffer for the policy, the policy existing in the defined configuration becomes active. The activated policy continues to remain in the defined configuration. The policy to be activated replaces the existing active policy of the same type. Activating the default IP Filter policies returns the IP management interface to its default state. An IP Filter policy without any rule cannot be activated. This subcommand prompts for a user confirmation before proceeding.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having OM permissions for the `ipfilter` RBAC class of commands.
2. Enter the `ipfilter --activate` command.

## Deleting an IP Filter Policy

You can delete a specified IP Filter policy. Deleting an IP Filter policy removes it from the temporary buffer. To permanently delete the policy from the persistent database, run `ipfilter --save` command. An active IP Filter policy cannot be deleted.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having the OM permissions for the `ipfilter` RBAC class of commands.
2. Enter the `ipfilter --delete` command.
3. To permanently delete the policy, enter the `ipfilter --save` command.

## IP Filter Policy Rules

An IP Filter policy consists of a set of rules. Each rule has an index number identifying the rule. There can be a maximum of 256 rules within an IP Filter policy.

Each rule contains the following elements:

- Source address: A source IP address or a group prefix.
- Destination port: The destination port number or name, such as Telnet, SSH, HTTP, or HTTPS.
- Protocol: The protocol type. Supported types are TCP or UDP.
- Action: The filtering action taken by this rule, either Permit or Deny.

A destination IP can also be specified

## Source Address

For an IPv4 filter policy, the source address has to be a 32-bit IPv4 address in dot decimal notation. The group prefix has to be a classless inter-domain routing (CIDR) block prefix representation. For example, 208.130.32.0/24 represents a 24-bit IPv4 prefix starting from the most significant bit. The special prefix 0.0.0.0/0 matches any IPv4 address. In addition, the `any` keyword is supported to represent any IPv4 address.

For an IPv6 filter policy, the source address must be a 128-bit IPv6 address, in a format acceptable in RFC 3513. The group prefix must be a CIDR block prefix representation. For example, 12AB:0:0:CD30::/64 represents a 64-bit IPv6 prefix starting from the most significant bit. In addition, the `any` keyword is supported to represent any IPv6 address.

## Destination Port

For the destination port, a single port number or a port number range can be specified. According to IANA ( <http://www.iana.org> ), ports 0 to 1023 are well-known port numbers, ports 1024 to 49151 are registered port numbers, and ports 49152 to 65535 are dynamic or private port numbers. Well-known and registered ports are normally used by servers to accept connections, while dynamic port numbers are used by clients.

For an IP Filter policy rule, you can select port numbers in the port number range between 1 and 65535, inclusive. This means that you have the ability to control how to expose the management services hosted on a switch, but not the ability to affect the management traffic that is initiated from a switch. A valid port number range is represented by a dash, for example 7-30. Alternatively, service names can also be used instead of port number. Table below lists the supported service names and their corresponding port numbers.

**Table 60: Destination Port Details**

| Service Name | Port Number |
|--------------|-------------|
| echo         | 7           |
| discard      | N/A         |
| systat       | 11          |
| daytime      | 13          |
| netstat      | 15          |
| chargen      | 19          |
| ftp data     | 20          |
| ftp          | 21          |
| fsp          | 21          |
| ssh          | 22          |
| telnet       | 23          |
| smtp         | 25          |
| time         | 27          |
| name         | 42          |
| whois        | 43          |



| Service Name | Port Number |
|--------------|-------------|
| domain       | 53          |
| bootps       | 67          |
| bootpc       | 68          |
| tftp         | 69          |
| http         | 80          |
| kerberos     | 88          |
| hostnames    | 101         |
| sftp         | 115         |
| ntp          | 123         |
| snmp         | 161         |
| snmp trap    | 162         |
| https        | 443         |
| ssmtp        | 465         |
| exec         | 512         |
| login        | 513         |
| shell        | 514         |
| uucp         | 540         |
| biff         | 512         |
| who          | 513         |
| syslog       | 514         |
| route        | 520         |
| timed        | 525         |

## Protocol

TCP and UDP protocols are valid protocol selections. Fabric OS v6.2.0 and later do not support the configuration to filter other protocols. Implicitly, ICMP type 0 and type 8 packets are always allowed to support an ICMP echo request and reply on commands, such as `ping` and `traceroute`.

## Action

For the action, only `permit` and `deny` are valid.

## Implicit Filter Rules

For every IP Filter policy, the two rules listed in below table are always assumed to be appended implicitly to the end of the policy. This ensures that TCP and UDP traffic to dynamic port ranges is allowed, so that management IP traffic initiated from a switch, such as syslog, radius and ftp, is not affected.

**Table 61: Implicit IP Filter Rules**

| Source Address | Destination Port | Protocol | Action |
|----------------|------------------|----------|--------|
| Any            | 1024-65535       | TCP      | Permit |
| Any            | 1024-65535       | UDP      | Permit |

## Default Policy Rules

Switches have a default IP filter policy for IPv4 and IPv6 and it cannot be deleted or changed. When an alternative IP filter policy is activated, the default IP filter policy becomes deactivated. The following table lists the rules of the default IP filter policy.

**Table 62: Default IP Policy Rules**

| Rule Number | Source Address | Destination Port | Protocol | Action |
|-------------|----------------|------------------|----------|--------|
| 1           | Any            | 22               | TCP      | Permit |
| 2           | Any            | 23               | TCP      | Permit |
| 6           | Any            | 80               | TCP      | Permit |
| 3           | Any            | 443              | TCP      | Permit |
| 4           | Any            | 161              | UDP      | Permit |
| 10          | Any            | 123              | UDP      | Permit |
| 11          | Any            | 600-1023         | TCP      | Permit |
| 12          | Any            | 600-1023         | UDP      | Permit |

## Brocade Custom Policy Rules

Brocade will add new default policies in future releases based on the rules mentioned in the following table. Customer can configure the rules using `ipfilter` command.

**Table 63: Brocade Custom IP Policy Rules**

| Rule Number | Source Address | Destination Port | Protocol | Action |
|-------------|----------------|------------------|----------|--------|
| 1           | Any            | 22               | TCP      | Permit |
| 2           | Any            | 23               | TCP      | Permit |
| 3           | Any            | 80               | TCP      | Permit |
| 4           | Any            | 443              | TCP      | Permit |
| 5           | Any            | 161              | TCP      | Permit |
| 6           | Any            | 123              | TCP      | Permit |
| 7           | Any            | 32769-61000      | TCP      | Permit |
| 8           | Any            | 1024-32768       | TCP      | Deny   |
| 9           | Any            | 61001-65535      | TCP      | Deny   |

## IP Filter Policy Enforcement

An active IP Filter policy is a filter applied to the IP packets through the management interface. IPv4 management traffic passes through the active IPv4 filter policy, and IPv6 management traffic passes through the active IPv6 filter policy. The IP Filter policy applies to the incoming (ingress) management traffic only. When a packet arrives, it is compared against each rule, starting from the first rule. If a match is found for the source address, destination port, and protocol, the corresponding action for this rule is taken, and the subsequent rules in this policy are ignored. If there is no match, it is compared to the next rule in the policy. This process continues until the incoming packet is compared to all rules in the active policy.

If none of the rules in the policy matches the incoming packet, the two implicit rules are matched to the incoming packet. If the rules still do not match the packet, the default action, which is to deny, is taken.

When the IPv4 or IPv6 address for the management interface of a switch is changed through the `ipAddrSet` command or manageability tools, the active IP Filter policies automatically become enforced on the management IP interface with the changed IP address.

### NOTE

If a switch is part of a LAN behind a Network Address Translation (NAT) server, depending on the NAT server configuration, the source address in an IP Filter rule may be the NAT server address.

## Adding a Rule to an IP Filter Policy

You can create the maximum number of 256 rules for an IP Filter policy. The change to the specified IP Filter policy is not saved to the persistent configuration until a `save` or `activate` subcommand is run.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having the OM permissions for the `ipfilter` RBAC class of commands.
2. Enter the `ipfilter --addrule` command.

## Deleting a Rule from an IP Filter Policy

Deleting a rule in the specified IP Filter policy causes the rules following the deleted rule to shift up in rule order. The change to the specified IP Filter policy is not saved to persistent configuration until a `save` or `activate` subcommand is run.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having the OM permissions for the `ipfilter` RBAC class of commands.
2. Enter the `ipfilter --delrule` command.

## Aborting an IP Filter Transaction

A transaction is associated with a command line or manageability session. It is opened implicitly when the `--create`, `--addrule`, `--delrule`, `--clone`, and `--delete` subcommands are run. The `--transabort`, `--save`, or `--activate` subcommands explicitly end the transaction owned by the current command line or manageability session. If a transaction is not ended, other command line or manageability sessions are blocked on the subcommands that would open a new transaction.

1. Log in to the switch either using an account with admin permissions, or an account associated with the chassis role and having the OM permissions for the `ipfilter` RBAC class of commands.
2. Enter the `ipfilter --transabort` command.

### NOTE

The `ipfilter --transabort` command does not end an open transaction of the previous session after the recursive login in the console. This is due to the characteristics of the console. You must reboot the switch to end the open transaction of the previous session.

## Distributing an IP Filter Policy

The IP Filter policy is manually distributed using a command. The distribution includes both active and defined IP Filter policies. All policies are combined as a single entity to be distributed and cannot be selectively distributed. However, you can choose the time at which to implement the policy for optimization purposes. If a distribution includes an active IP Filter policy, the receiving switches activate the same IP Filter policy automatically. When a switch receives IP Filter policies, all uncommitted changes left in its local transaction buffer are lost, and the transaction is aborted.

The IP Filter policy can be manually distributed to the fabric using a command; automatic distribution is not supported. To distribute the IP Filter policy, see [Distributing the Local ACL Policies](#) for instructions.

You can accept or deny IP Filter policy distribution using the `fddCfg --localaccept` or `fddCfg --localreject` commands. See [Policy Database Distribution](#) for more information about distributing the IP Filter policy.

### NOTE

To distribute the IP Filter policy in a logical fabric using Virtual Fabrics, use the `chassisDistribute` command.

### NOTE

IP filter policy distribution from the pre-FOS 9.0 is blocked, if the distributed IP filter policies contains packet forwarding rules.

## Policy Database Distribution

Fabric OS lets you manage and enforce the ACL policy database on either a per-switch or fabric-wide basis. The local switch distribution setting and the fabric-wide consistency policy affect the switch ACL policy database and related distribution behavior.

The ACL policy database is managed as follows:

- **Switch database distribution setting** – Controls whether the switch accepts or rejects databases distributed from other switches in the fabric. The `distribute` command sends the database from one switch to another, overwriting the target switch database with the distributed one. To send or receive a database the setting must be `accept`. For configuration instructions, refer to [Database Distribution Settings](#).

**Virtual Fabric considerations:** FCS, DCC, SCC, and AUTH databases can be distributed using the `-distribute` command, but the PWD is blocked from distribution if non-default VFs exist. The IPFILTER databases can be distributed using the FID in a VF environment.

- **Manually distribute an ACL policy database** – Use the `distribute` command to push the local database of the specified policy type to target switches. Refer to [ACL Policy Distribution to Other Switches](#).
- **Fabric-wide consistency policy** – Use this policy to ensure that switches in the fabric enforce the same policies. Set a strict or tolerant fabric-wide consistency policy for each ACL policy type to automatically distribute that database when a policy change is activated. If a fabric-wide consistency policy is not set, the policies are managed on a per-switch basis. For configuration instructions, refer to [Fabric-Wide Enforcement](#).

**Virtual Fabric considerations:** Fabric-wide consistency policies are configured on a per-logical switch basis and are applied to the fabrics connected to the logical switches. Automatic policy distribution behavior for DCC, SCC, and FCS is the same and are configured on a per-logical switch basis.

The following table explains how the local database distribution settings and the fabric-wide consistency policy affect the local database when the switch is the target of a `distribute` command.

**Table 64: Interaction between Fabric-Wide Consistency Policy and Distribution Settings**

| Distribution Setting | Fabric-Wide Consistency Policy                                                                                                                                                                                                                                                 |                                                                                                                                                                                              |                                                                                                                                                                         |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Absent (default)     | Tolerant                                                                                                                                                                                                                                                                       | Strict                                                                                                                                                                                       | Invalid configuration.                                                                                                                                                  |
| Reject               | Database is protected, it cannot be overwritten.<br>May not match other databases in the fabric.                                                                                                                                                                               | Invalid configuration.<br>An error is returned indicating that the distribution setting must be Accept before you can set the fabric-wide consistency policy.                                | An error is returned indicating that the distribution setting must be accept before you can set the fabric-wide consistency policy.                                     |
| Accept (default)     | Database is not protected, the database can be overwritten.<br>If the switch initiating a <code>distribute</code> command has a strict or tolerant fabric-wide consistency policy, the fabric-wide policy is also overwritten.<br>May not match other databases in the fabric. | Database is not protected.<br>If the fabric-wide consistency is set as "strict" for a particular policy, the manual distribution is blocked.<br>May not match other databases in the fabric. | Database is not protected.<br>Automatically distributes activated changes to all switches in the fabric.<br>Active database is the same for all switches in the fabric. |

## Database Distribution Settings

The distribution settings control whether a switch accepts or rejects distributions of databases from other switches and whether the switch may initiate a distribution. Configure the distribution setting to reject when maintaining the database on a per-switch basis.

The table below lists the databases supported in FOS switches.

**Table 65: Supported Policy Databases**

| Database Type                  | Database Identifier (ID) |
|--------------------------------|--------------------------|
| Authentication policy database | AUTH                     |
| DCC policy database            | DCC                      |
| FCS policy database            | FCS                      |
| IP Filter policy database      | IPFILTER                 |
| Password database              | PWD                      |
| SCC policy database            | SCC                      |

Use the `chassisDistribute` command to distribute IP Filter policies. To distribute other security policies, use the `distribute` command. The `distribute` command distributes the database also to Access Gateways in non-VF mode.

## Displaying the Database Distribution Settings

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the FabricDistribution RBAC class of commands.
2. Enter the `fddCfg --showall` command.

The following sample output shows the database distribution settings:

```
switch:admin> fddcfg --showall
```

```

Local Switch Configuration for all Databases:-
 DATABASE - Accept/Reject

 SCC - accept
 DCC - accept
 PWD - accept
 FCS - accept
 AUTH - accept
 IPFILTER - accept
Fabric Wide Consistency Policy:- ""

```

## Enabling Local Switch Protection

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the FabricDistribution RBAC class of commands.
2. Enter the `fddCfg --localreject` command.

## Disabling Local Switch Protection

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the FabricDistribution RBAC class of commands.
2. Enter the `fddCfg --localaccept` command.

## ACL Policy Distribution to Other Switches

This section explains how to manually distribute local ACL policy databases. The `distribute` command has the following dependencies:

- All target switches must accept the database distribution (refer to [Database Distribution Settings](#)).
- The fabric must have a tolerant or no (absent) fabric-wide consistency policy (refer to [Fabric-Wide Enforcement](#)).

If the fabric-wide consistency policy for a database is strict, the database cannot be manually distributed. When you set a strict fabric-wide consistency policy for a database, the distribution mechanism is automatically invoked whenever the database changes.

- The local distribution setting must be accepted. To initiate the `distribute` command, set the local distribution to `Accept`.

## Distributing the Local ACL Policies

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the FabricDistribution RBAC class of commands.
2. Enter the `distribute -p` command.

## Fabric-Wide Enforcement

The fabric-wide consistency policy enforcement setting determines the distribution behavior when changes to a policy are activated. Using the tolerant or strict fabric-wide consistency policy ensures that changes to local ACL policy databases are automatically distributed to other switches in the fabric.

### NOTE

To completely remove all fabric-wide policy enforcement from a fabric, enter the `fddCfg --fabwideset ""` command.

When you set the fabric-wide consistency policy using the `fddCfg` command with the `--fabwideset database_id` option, both the fabric-wide consistency policy and specified database are distributed to the fabric. The active policies of the specified databases overwrite the corresponding active and defined policies on the target switches.

Policy changes that are saved but not activated are stored locally until a policy database change is activated. Activating a policy automatically distributes the Active policy set for that policy type (SCC, DCC, FCS, or any combination of the three) to the other switches in the fabric.

The following table describes the fabric-wide consistency settings.

**Table 66: Fabric-Wide Consistency Policy Settings**

| Setting  | Value                 | When a Policy is Activated                                                                                                            |
|----------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Absent   | null                  | Database is not automatically distributed to other switches in the fabric.                                                            |
| Tolerant | <i>database_id</i>    | All updated and new policies of the type specified (SCC, DCC, FCS, or any combination) are distributed to all switches in the fabric. |
| Strict   | <i>database_id</i> :S | All updated and new policies of the type specified (SCC, DCC, FCS, or any combination) are distributed to all switches in the fabric. |

## Displaying the Fabric-Wide Consistency Policy

1. Connect to the switch and log in using an account with admin permissions, or an account with O permission for the FabricDistribution RBAC class of commands.
2. Enter the `fddCfg --showall` command.

The following sample output shows policies for a fabric where no consistency policy is defined.

```
switch:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
 DATABASE - Accept/Reject

 SCC - accept
 DCC - accept
 PWD - accept
 FCS - accept
 AUTH - accept
 IPFILTER - accept
Fabric Wide Consistency Policy:- ""
```

## Setting the Fabric-Wide Consistency Policy

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the FabricDistribution RBAC class of commands.
2. Enter the `fddCfg --fabwideset` command.

The following example shows how to set a strict SCC and tolerant DCC fabric-wide consistency policy.

```

switch:admin> fddcfg --fabwideset "SCC:S;DCC"
switch:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
 DATABASE - Accept/Reject

 SCC - accept
 DCC - accept
 PWD - accept
 FCS - accept
 AUTH - accept
 IPFILTER - accept
Fabric Wide Consistency Policy:- "SCC:S;DCC"

```

## Notes on Joining a Switch to the Fabric

When a switch is joined to a fabric with a tolerant SCC, DCC, or FCS fabric-wide consistency policy, the joining switch must have a matching tolerant SCC, DCC, or FCS fabric-wide consistency policy. If the tolerant SCC, DCC, or FCS fabric-wide consistency policies do not match, the switch can still join the fabric, but an error message notifies you about the mismatch. If the tolerant SCC, DCC, and FCS fabric-wide consistency policies match, the corresponding SCC, DCC, and FCS ACL policies are compared.

The enforcement of fabric-wide consistency policy involves comparison of the Active policy set. If the ACL policies match, the switch joins the fabric successfully. If the ACL policies are absent either on the switch or on the fabric, the switch joins the fabric successfully, and the ACL policies are copied automatically from where they are present to where they are absent. The Active policy set where it is present overwrites the Active and Defined policy set where it is absent. If the ACL policies do not match, the switch cannot join the fabric and the neighboring E\_Ports are disabled.

Use the `fddCfg --fabwideset` command on either the switch or the fabric to set a matching strict SCC, DCC, or FCS fabric-wide consistency policy. Use ACL policy commands to delete the conflicting ACL policy from one side to resolve the ACL policy conflict. If neither the fabric nor the joining switch is configured with a fabric-wide consistency policy, no ACL merge checks are required. Under both conflicting conditions, `secPolicyActivate` is blocked in the merged fabric. Use the `distribute` command to explicitly resolve conflicting ACL policies.

The descriptions also apply to joining two fabrics. In this context, the joining switch becomes a joining fabric.

## Matching Fabric-Wide Consistency Policies

This section describes the interaction between the databases with active SCC and DCC policies and combinations of fabric-wide consistency policy settings when fabrics are merged.

For example: Fabric A with SCC:S;DCC (strict SCC and tolerant DCC) joins Fabric B with SCC:S;DCC (strict SCC and tolerant DCC), the fabrics can merge as long as the SCC policies match, including the order SCC:S;DCC and if both are set to strict.

The following table describes the impact of merging fabrics with the same fabric-wide consistency policy that have SCC, DCC, or both policies.

**Table 67: Merging Fabrics with Matching Fabric-Wide Consistency Policies**

| Fabric-Wide Consistency Policy | Fabric A ACL Policies | Fabric B ACL Policies | Merge Results | Database Copied         |
|--------------------------------|-----------------------|-----------------------|---------------|-------------------------|
| None                           | None                  | None                  | Succeeds      | No ACL policies copied. |



| Fabric-Wide Consistency Policy | Fabric A ACL Policies      | Fabric B ACL Policies      | Merge Results | Database Copied                                                                                     |
|--------------------------------|----------------------------|----------------------------|---------------|-----------------------------------------------------------------------------------------------------|
|                                | None                       | SCC/DCC                    | Succeeds      | No ACL policies copied.                                                                             |
| Tolerant                       | None                       | None                       | Succeeds      | No ACL policies copied.                                                                             |
|                                | None                       | SCC/DCC                    | Succeeds      | ACL policies are copied from B to A.                                                                |
|                                | SCC/DCC                    | SCC/DCC                    | Succeeds      | If A and B policies do not match, a warning displays and policy commands are disabled. <sup>a</sup> |
| Strict                         | None                       | None                       | Succeeds      | No ACL policies copied.                                                                             |
|                                | None                       | SCC/DCC                    | Succeeds      | ACL policies are copied from B to A.                                                                |
|                                | Matching SCC/DCC           | Matching SCC/DCC           | Succeeds      | No ACL policies copied.                                                                             |
|                                | Different SCC/DCC policies | Different SCC/DCC policies | Fails         | Ports are disabled.                                                                                 |

## Nonmatching Fabric-Wide Consistency Policies

You may encounter one of the following two scenarios described in Table 66 and Table 67, where you are merging a fabric with a strict policy to a fabric with an absent, tolerant, or nonmatching strict policy and the merge fails and the ports are disabled.

The table below shows merges that are not supported.

**Table 68: Examples of Strict Fabric Merges**

| Fabric-Wide Consistency Policy Setting |             |           | Expected Behavior                         |
|----------------------------------------|-------------|-----------|-------------------------------------------|
| —                                      | Fabric A    | Fabric B  | —                                         |
| Strict/Tolerant                        | SCC:S;DCC:S | SCC;DCC:S | Ports that connect switches are disabled. |
|                                        | SCC;DCC:S   | SCC:S;DCC |                                           |
|                                        | SCC:S;DCC   | SCC:S     |                                           |
| Strict/Absent                          | SCC:S;DCC:S | —         |                                           |
|                                        | SCC:S       | —         |                                           |
|                                        | DCC:S       | —         |                                           |
| Strict/Strict                          | SCC:S       | DCC:S     |                                           |

The table below lists the matrix of merging fabrics with tolerant and absent policies.

**Table 69: Fabric Merges with Tolerant and Absent Combinations**

| Fabric-Wide Consistency Policy Setting |          | Expected Behavior |   |
|----------------------------------------|----------|-------------------|---|
| Tolerant/Absent                        | Fabric A | Fabric B          | — |

a. To resolve the policy conflict, manually distribute the database you want to use to the switch with the mismatched database. Until the conflict is resolved, commands, such as `fddCfg --fabwideset` and `secPolicyActivate` are blocked.

| Fabric-Wide Consistency Policy Setting |         | Expected Behavior |                                                                                       |
|----------------------------------------|---------|-------------------|---------------------------------------------------------------------------------------|
|                                        | SCC;DCC | —                 | Error message logged                                                                  |
|                                        | DCC     | —                 | Run the <code>fddcfg --fabwideset &lt;policyID&gt;</code>                             |
|                                        | SCC;DCC | SCC               | from any switch with the desired configuration to fix the conflict.                   |
|                                        | DCC     | SCC               | The <code>secpolicyactivate</code> command is blocked until the conflict is resolved. |

## In-Flight Encryption and Compression

This section describes In-flight encryption, which provides security for frames while they are in the flight between two switches. In-flight compression provides better bandwidth use on the ISLs/IFLs, especially over long distance.

The in-flight encryption and compression features allow frames to be encrypted or compressed at the egress point of an ISLs/IFLs between two Brocade switches, and then to be decrypted or decompressed at the ingress point of the ISLs/IFLs. Frames are never left in an encrypted or compressed state when delivered to an end device.


These features use port-based encryption and compression. You can enable the in-flight encryption and compression features for both E\_Ports and EX\_Ports on a per-port basis. By default, these features are disabled for all ports on a switch.

### NOTE

No additional license is required to enable an in-flight encryption or compression.

Encryption and compression can be enabled at the same time, or you can enable either encryption or compression selectively. The following figure shows 32G links connecting three Brocade switches. One link is configured with encryption and compression, one with just encryption, and one with just compression.

**Figure 8: Encryption and Compression on 32G ISLs/IFLs**



## Gen 7 and Gen 6 In-Flight Encryption and Compression Features

**Table 70: Gen 7 (64G) In-Flight Encryption and Compression Features (Supported from Fabric OS 9.0.0)**

| Feature                                                                                                     | G720 (Switch Type 181)                                                                                                | FC64-48 (Blade ID 216) |
|-------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|------------------------|
| Number of front-end port groups (ASICs)                                                                     | 1 port group                                                                                                          | 2 port groups          |
| Number of ports per ASIC/port group that can be enabled with encryption or compression at any speed         | 4 ports                                                                                                               |                        |
| Total number of ports per switch/blade that can be enabled with encryption or compression (row 1 x row 2)   | 4 ports                                                                                                               | 8 ports                |
| Port types that support encryption or compression                                                           | E_Ports, EX_Ports, XISL ports in VF mode, inter-switch link/inter-fabric link (ISL/IFL) between Gen 7 and Gen 6 ports |                        |
| Port types that do not support encryption or compression but can exist along the I/O path                   | VE_Ports, GE_Ports, F_Ports, F_Port trunks, ICLs, D_Ports, AE_Ports, FCoE/Ethernet                                    |                        |
| Maximum total bandwidth of the ports online per port group with only encryption and compression enabled     | 128Gb/s per port group                                                                                                |                        |
| All the ports in a port group must be connected to the same domain or can be connected to different domains | Different domains                                                                                                     |                        |
| Other features supported on the in-flight encryption or compression enabled ports                           | Trunking, QoS, and long distance                                                                                      |                        |
| Earliest Fabric OS version that supports encryption                                                         | 9.0.0/9.1.0                                                                                                           | 9.0.0                  |
| Earliest Fabric OS version that supports compression                                                        | 9.0.0/9.1.0                                                                                                           | 9.0.0                  |
| Ports that must be disabled when in-flight encryption is enabled on any supported ports                     | None                                                                                                                  | None                   |
| Supported flow control modes                                                                                | R_RDY, VC_RDY, and EXT_VC_RDY                                                                                         |                        |
| Ports that do not support in-flight encryption or compression                                               | None                                                                                                                  | None                   |
| Number of buffers required for a non-QoS encryption-enabled online E_Port                                   | 1400                                                                                                                  | 1400                   |
| Number of buffers required for a QoS encryption-enabled online E_Port                                       | 1400                                                                                                                  | 1400                   |

### NOTE

- The Brocade G730 switch, any switch in Access Gateway mode, and the CR64-8 and CR64-4 core routing blades do not support in-flight encryption or compression.
- Encrypted trunk groups on the FC64-48 blades are supported in Fabric OS 9.0.0 and later releases.

**Table 71: Gen 6 (32G) In-Flight Encryption and Compression Features (Supported from Fabric OS 9.0.0)**

| Feature                                                                                                   | G630 (Switch Type 184) | G620 (Switch Type 183) | FC32-X7-48 (Blade ID 218) |
|-----------------------------------------------------------------------------------------------------------|------------------------|------------------------|---------------------------|
| Number of front-end port groups (ASICs)                                                                   | 4 port groups          | 1 port group           | 2 port groups             |
| Number of ports per ASIC/port group that can be enabled with encryption or compression at any speed       | 4 ports                |                        |                           |
| Total number of ports per switch/blade that can be enabled with encryption or compression (row 1 x row 2) | 16 ports               | 4 ports                | 8 ports                   |

| Feature                                                                                                     | G630 (Switch Type 184)                                                                                                | G620 (Switch Type 183) | FC32-X7-48 (Blade ID 218) |
|-------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|------------------------|---------------------------|
| Port types that support encryption or compression                                                           | E_Ports, EX_Ports, XISL ports in VF mode, inter-switch link/inter-fabric link (ISL/IFL) between Gen 7 and Gen 6 ports |                        |                           |
| Port types that do not support encryption or compression but can exist along the I/O path                   | VE_Ports, GE_Ports, F_Ports, F_Port trunks, ICLs, D_Ports, AE_Ports, FCoE/Ethernet                                    |                        |                           |
| Maximum total bandwidth of the ports online per port group with only encryption and compression enabled     | 64Gb/s per port group                                                                                                 |                        |                           |
| All the ports in a port group must be connected to the same domain or can be connected to different domains | Different domains                                                                                                     |                        |                           |
| Other features supported on the in-flight encryption or compression enabled ports                           | Trunking, QoS, and long distance                                                                                      |                        |                           |
| Earliest Fabric OS version that supports encryption                                                         | 9.0.0                                                                                                                 | 9.0.0                  | 9.0.0                     |
| Earliest Fabric OS version that supports compression                                                        | 9.0.0                                                                                                                 | 9.0.0                  | 9.0.0                     |
| Ports that must be disabled when in-flight encryption is enabled on any supported ports                     | None                                                                                                                  | None                   | None                      |
| Supported flow control modes                                                                                | R_RDY, VC_RDY, and EXT_VC_RDY                                                                                         |                        |                           |
| Ports that do not support in-flight encryption or compression                                               | None                                                                                                                  | None                   | None                      |
| Number of buffers required for a non-QoS encryption-enabled online E_Port                                   | 1400                                                                                                                  | 1400                   | 1400                      |
| Number of buffers required for a QoS encryption-enabled online E_Port                                       | 1400                                                                                                                  | 1400                   | 1400                      |

**NOTE**

Encrypted trunk groups on the FC32-X7-48 is supported in Fabric OS 9.0.0 and later releases.

**Table 72: Gen 6 (32G) In-Flight Encryption and Compression Features (Supported from Fabric OS 8.0.x)**

| Feature                                                                                                     | G630 (Switch Type 173)                                                                                                | G620 (Switch Type 162)   | FC32-48SX6 – Compression (Blade ID 178) | Only (Blade ID 186) |
|-------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|--------------------------|-----------------------------------------|---------------------|
| Number of front-end port groups (ASICs)                                                                     | 3 port groups                                                                                                         | 1 port group             | 2 port groups                           | 1 port group        |
| Number of ports per ASIC/port group that can be enabled with encryption or compression at any speed         | 4 ports                                                                                                               |                          |                                         |                     |
| Total number of ports per switch/blade that can be enabled with encryption or compression (row 1 x row 2)   | 12 ports                                                                                                              | 4 ports                  | 8 ports                                 | 4 ports             |
| Port types that support encryption or compression                                                           | E_Ports, EX_Ports, XISL ports in VF mode, inter-switch link/inter-fabric link (ISL/IFL) between Gen 7 and Gen 6 ports |                          |                                         |                     |
| Port types that do not support encryption or compression but can exist along the I/O path                   | VE_Ports, GE_Ports, F_Ports, F_Port trunks, ICLs, D_Ports, AE_Ports, FCoE/Ethernet                                    |                          |                                         |                     |
| Maximum total bandwidth of the ports online per port group with only encryption and compression enabled     | 64Gb/s per port group                                                                                                 |                          |                                         |                     |
| All the ports in a port group must be connected to the same domain or can be connected to different domains | Same domain                                                                                                           | Can be different domains |                                         |                     |

| Feature                                                                                 | G630 (Switch Type 173)           | G620 (Switch Type 162)                                            | FC32-48SX6 – Compression Only (Blade ID 178) | FC32-48SX6 – Compression Only (Blade ID 186) |
|-----------------------------------------------------------------------------------------|----------------------------------|-------------------------------------------------------------------|----------------------------------------------|----------------------------------------------|
| Other features supported on the in-flight encryption or compression enabled ports       | Trunking, QoS, and long distance | QoS and long distance                                             | Trunking, QoS, and long distance             | Trunking, QoS, and long distance             |
| Earliest Fabric OS version that supports encryption                                     | 8.2.0a                           | 8.2.0                                                             | 8.1.0                                        | Not supported                                |
| Earliest Fabric OS version that supports compression                                    | 8.2.0a                           | 8.0.1                                                             | 8.1.0                                        | 8.2.0                                        |
| Ports that must be disabled when in-flight encryption is enabled on any supported ports | None                             | SFP+ ports 44 to 47<br>(Also, 40 to 43 cannot be in octet mode 2) | None                                         | None                                         |
| Supported flow control modes                                                            | R_RDY, VC_RDY, and EXT_VC_RDY    |                                                                   |                                              |                                              |
| Ports that do not support in-flight encryption or compression                           | All QSFP ports                   | SFP ports 44 to 47                                                | None                                         | None                                         |
| Number of buffers required for a non-QoS encryption-enabled online E_Port               | 448                              | 175                                                               | 175                                          | N/A                                          |
| Number of buffers required for a QoS encryption-enabled online E_Port                   | 448                              | 210                                                               | 210                                          | N/A                                          |

**NOTE**

- The Brocade G610 switch, any switch in Access Gateway mode, the FC32-64 port blade, the CR32-8 and CR32-4 core routing blades, and the Brocade 7810 extension switch do not support in-flight encryption or compression.
- Encrypted trunk groups on the FC32-48 blade are supported in Fabric OS 8.2.1a and later releases.
- Trunking is not supported on ports that are enabled with in-flight encryption or compression for the switch type 162.


**Supported Use Cases for the Brocade G630 Switch**

The following are the Brocade G630 Switch encryption/compression SFP port groups as you see them in the front panel:

**Table 73: Brocade G630 Encryption/Compression SFP Port Groups**

|                                                                              |                                                                  |
|------------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>Encryption/compression port group 1</b><br>SFP ports 0 to 15 and 48 to 63 | <b>Encryption/compression port group 2</b><br>SFP ports 16 to 47 |
|                                                                              | <b>Encryption/compression port group 3</b><br>SFP ports 64 to 95 |


In this section, the following format is used to represent the Brocade G630 Switch with any four ports enabled with encryption, compression, or both in each port group.

**Figure 9: Encryption/Compression Port Groups****Between Two Data Centers**

Only one domain per port group is supported. Use all the ports in a port group before going to the next port group to plan for future data center connections.

**NOTE**


In all the illustrations in this section, the Brocade G630 Switch is considered to be in data center 1.

**Figure 10: Between Two Data Centers**

### **Among Three Data Centers**

Only one domain per port group is supported. Use all the ports in a port group before going to the next port group to plan for future data center connections.


**Figure 11: Among Three Data Centers**



### **Among Four Data Centers**

Only one domain per port group is supported.





**Figure 12: Among Four Data Centers****With FCR Support on the G630**

Only one domain per port group is supported. Only one domain per all the same FID EX\_Ports is allowed in a port group.

**NOTE**


No host or device can be connected in a logical SAN (LSAN) zone on the same port group with EX\_Ports in non-VF mode. However, this configuration is supported in VF mode.

**Figure 13: Supported Configuration with Fibre Channel Routing (FCR)****Figure 14: Unsupported Configuration with FCR**

### **With Virtual Fabrics Support on the G630**

Only one domain per logical switch, default switch, or base switch is supported. You cannot have a logical switch and the base switch in the same port group.


**Figure 15: With Virtual Fabrics Support on G630**




- BS – Base Switch
- DS – Default Switch
- LS1 – Logical Switch 1
- LS2 – Logical Switch 2

### **With Virtual Fabrics and FCR Support on the G630**

Only one domain per logical switch, default switch, or base switch is allowed. Only one domain per FID/edge fabric is allowed in a port group. You cannot have a logical switch and the base switch in the same port group.


**Figure 16: Supported Configuration with Virtual Fabrics and FCR**

- BS – Base Switch
- DS – Default Switch
- LS1 – Logical Switch 1
- LS2 – Logical Switch 2

**Figure 17: Unsupported Configuration with Virtual Fabrics and FCR**

- BS – Base Switch
- LS1 – Logical Switch 1
- LS2 – Logical Switch 2

**Figure 18: Rectified Configuration with Virtual Fabrics and FCR**



- BS – Base Switch
- LS1 – Logical Switch 1
- LS2 – Logical Switch 2

## Additional In-Flight Encryption and Compression Restrictions

Consider the following additional restrictions:

- In a configuration with two switches with multiple ISLs connecting them, an encrypted and non-encrypted pair of links between two chassis is not allowed to work concurrently. Also, the transition from non-encrypted to encrypted ISLs between the two switches is disruptive.
- Port mirroring through any encryption-enabled or compression-enabled port is not supported.
- You cannot enable both trunking and encryption on G620 ports. Attempting to do so will segment the link with an error directing you to disable trunking or encryption on the port before you can proceed. During a non-disruptive upgrade from a previous Fabric OS version to 8.2.1, trunking will be disabled on G620 ports that are enabled for encryption.
- Although you can enable trunking and encryption on FC32-48 blade ports when using Fabric OS 8.2.1a and later, you must consider the additional restrictions.
- FCP data frames and non-FCP data frames except extended link service (ELS) and basic link service (BLS) frames are supported.
  - FCP data frames are of Type = 0x8. For encryption, R\_CTL = 0x1 and R\_CTL = 0x4 are supported. For compression, only R\_CTL = 0x1 is supported.
  - Non-FCP data frames are of Type != 0x8. Non-FCP frames with ELS/BLS (R\_CTL == 0x2 || R\_CTL == 0x8) are not supported.

## Availability Considerations for Encryption and Compression

To provide redundancy in the event of encryption or compression port failures, connect each ISL or trunk group to different ASICs on the peer switch. See the following examples.

- If the two ports configured for encryption or compression within the same ASIC on a switch or port blade are not configured for trunking, connect each ISL to a different ASIC on the peer switch. Configure the two ports on the other ASIC similarly. If the ports are configured for trunking, connect each trunk group to different ASICs on the peer switch.
- Gen 6 switches (except Brocade G630) do not support trunking with encryption enabled ports. The `portCfgDefault` command is also blocked on encryption enabled ports. You must disable the encryption configuration on these ports before running the `portCfgDefault` command. However, compression enabled ports support trunking.

### NOTE

If any port with encryption enabled encounters rare error conditions that require error recovery to be performed on the encryption engine within that ASIC, all encryption or compression enabled ports on that ASIC go offline.

## Virtual Fabrics Considerations for Encryption and Compression

The E\_Ports and EX\_Ports in the user-created logical switch, base switch, or default switch, and the EX\_Ports on base switches can support encryption and compression, with some exceptions.

You can configure encryption on extended ISL (XISL) ports, but not on logical ISL (LISL) ports. However, frames from the LISL ports are implicitly encrypted or compressed as they pass through encryption- or compression-enabled XISL ports.

You cannot move a port from one logical switch to another logical switch if in-flight encryption or compression is enabled on the port. You must disable the encryption and compression configurations before moving the port, and then enable encryption and compression after the port has moved.

## How In-Flight Encryption and Compression Are Enabled

Encryption and compression capabilities and configurations from each end of the ISL are exchanged during E\_Port or EX\_Port initialization. Capabilities and configurations must match; otherwise, port segmentation or disablement occurs.

If the port was configured for compression, the compression feature is enabled.

If the port was configured for encryption, authentication is performed and the keys needed for encryption are generated. The encryption feature is enabled if authentication is successful. If authentication fails, the ports are segmented.

### ATTENTION

The port is disabled if the authentication fails and when there is any mismatch in the port configuration at the other end of the IFL.

The most common reasons for E\_Port or EX\_Port segmentation include the following situations:

- Port authentication fails. One of the following error messages is displayed:  
`Authentication Rejected`  
`Authentication Failure`
- Encryption or compression configurations do not match at both ends. For example, if at one end there is a switch that does not support encryption or compression, the port will be disabled. One of the following error messages is displayed:  
`Compression configuration mismatch`  
`Encryption configuration mismatch`
- An encryption or compression configuration is enabled but resources are not available, or there are other failures preventing encryption or compression from being enabled. The following error message is displayed.

Encryption/Compression enabled but resource unavailable

- The number of available ports has reached the bandwidth limitation.

#### NOTE

If trunking is enabled, be aware that the ports creating the bandwidth limitation will form a trunk group, while the rest of the ports will be segmented.

You can also decommission any port that has in-flight encryption and compression enabled. See [Port Decommissioning](#) for details on decommissioning ports.

## Authentication and Key Generation for Encryption and Compression

The following points apply to authentication and key generation on the supported devices:

- Authentication and key generation only apply to ports that are configured for encryption. They do not apply to ports that are only configured for compression.
- Authentication using FCAP or DH-CHAP is supported between E\_Ports and F\_Ports connecting Brocade switches and Access Gateway devices. However, EX\_Ports only support authentication using DH-CHAP. This means that if an inter-fabric link (IFL) must be authenticated, the authentication type on edge fabric should be set to either DH-CHAP or ALL. Configuration done in FCR using the `authUtil` command is applicable only for Backbone fabric (E\_Ports in FCR, not EX\_Ports). Authentication parameters for FCR cannot be modified.
  - Default configuration:
    - Authentication Type – DH-CHAP
    - Hash Type – SHA256, SHA1, md5
    - Group Type – 0, 1, 2, 3, 4
    - Switch Authentication Policy – Passive
- The in-flight encryption protocol supports the Advanced Encryption Standard – Galois Counter Mode (AES-GCM) authenticated encryption block cipher mode. A key, Initial Vector (IV), segment number, and salt are required to encrypt the data before it is transmitted, and to decode the data after it is received on the other end of the link.
- In-flight encryption feature uses DH-CHAP (Diffie Hellman – Challenge Handshake Authentication Protocol) or Fibre Channel Authentication Protocol (FCAP) for authentication and key generation for secure frame transaction.
- In the in-flight encryption process, a session key is generated during authentication phase, and it is used in internet key exchange (IKE) protocol session to generate and exchange encryption/decryption keys for packet encryption/decryption between two devices.
- For in-flight encryption using DH-CHAP, DH-CHAP must be configured along with DH group 4 and pre-shared secret keys on the devices on both ends of the ISL as a pre-requisite. Authentication secrets greater than 32 characters are recommended for stronger encryption keys. When the link is authenticated, the keys are generated and exchanged.
- For in-flight encryption using FCAP, FCAP must be configured along with DH group 4 and certificates (CA and switch) at both ends of ISL as a prerequisite.
- The encryption keys never expire. While the port remains online, the keys generated for the port remain the same. When a port is disabled, segmented, or taken offline, a new set of keys is generated when the port is enabled again.
- All members of a trunk group use the same set of keys as the controller port. Agent ports do not exchange keys. If the controller port goes offline causing an E\_Port or EX\_Port change, the trunk continues to use the same set of keys.

## In-Flight Compression on Long-Distance Ports

When configuring in-flight compression on long-distance ports, configure the long-distance ports with double the number of buffers.

Configure the port to use the long-distance LS mode, and specify the number of buffers to allocate to the port. You can see what the average compression ratio and the average frame size values are and adjust the allocated credit accordingly using the `portEncCompShow` and `portBufferShow` commands. You can then use the `portBufferCalc` command to estimate the assigned credit value to optimize performance.



## Compression Ratios for Compression-Enabled Ports

An average compression ratio of 2:1 is provided. The compression ratio value is recalculated every 5 seconds, and is the ratio between the accumulated original length and the compressed length of the data over the previous 5 seconds.

When a port is configured for compression, entering `portStatsShow` displays the port's compression ratio. The value shown by `portStatsShow` is a 5 second average. Your results depend on the pattern of the payload data.

The ASIC compression block can compress data only if there is at least 3 bytes of data.

The `portBufferShow` command shows the average frame size for both received (Rx) and transmitted (Tx) frames. The Rx values are after compression and the Tx values are before compression.

Because encryption adds more payload to the port in addition to compression, the compression ratio calculation is significantly affected on ports configured for both encryption and compression because the compressed length then also includes the encryption header. This overhead affects the ratio calculation. To obtain accurate compression ratio data, enable ports for compression only.

## Configuring In-Flight Encryption and Compression on an EX\_Port

When you configure in-flight encryption and compression across an IFL, first configure the EX\_Port and then configure the E\_Port. The encryption and compression settings must match at either end of the IFL.

The following steps summarize how to enable in-flight encryption or compression on an EX\_Port. Perform these steps on the switch or the Director.

1. Determine which ports are available for encryption or compression.

See [Viewing the Encryption and Compression Configuration](#) for instructions.

2. Enter `fcredgeshow` to learn the WWN of the edge switch.

You will need this WWN when you set up the secret key.

```
switch:admin> fcredgeshow
```

| FID | EX-port | E-port | Neighbor Switch (PWWN, SWWN)                    | Flags |
|-----|---------|--------|-------------------------------------------------|-------|
| 20  | 1       | 1      | 20:01:00:05:33:13:70:3e 10:00:00:05:33:13:70:3e |       |

3. Enter `portcfgexport` to learn the WWN of the front phantom domain.

You will need this WWN when you set up the secret key on the E\_Port on the other end of the IFL.

```
FCR:admin> portcfgexport 1
Port 1 info
Admin: enabled
State: OK
Pid format: core(N)
Operate mode: Brocade Native
Edge Fabric ID: 20
Front Domain ID: 160
Front WWN: 50:00:53:31:37:43:ee:14
Principal Switch: 8(output truncated)
```

You can also specify a port range per blade in chassis-based FCR. However, the range should not include a SIM\_Port.

For example, the following sample sets the range of EX\_Ports as FCR ports:

```
switch:admin> portcfgexport 1/1-2 -a 1
```



```
2014/09/15-07:09:12, [FCR-1071], 11763, SLOT 6 | FID 128, INFO, DCX, Port 1/1 is changed from non FCR port
to FCR port.
2014/09/15-07:09:13, [FCR-1071], 11764, SLOT 6 | FID 128, INFO, DCX, Port 1/2 is changed from non FCR port
to FCR port.
```

The following sample sets the range of EX\_Ports as non-FCR ports:

```
switch:admin> portcfgexport 1/1-2 -a 2
2014/09/15-07:09:17, [FCR-1072], 11765, SLOT 6 | FID 128, INFO, DCX, Port 1/1 is changed from FCR port to
non FCR port.
2014/09/15-07:09:18, [FCR-1072], 11766, SLOT 6 | FID 128, INFO, DCX, Port 1/2 is changed from FCR port to
non FCR port.
```

Following successful port initialization, the configured features are enabled and active. Use the `fcledgedshow` command to check that the EX\_Port has come online with encryption or compression enabled.

4. If you are enabling encryption on the port, configure port level authentication for the port.

Omit this step if you want to enable only compression on the port.

See [Configuring and Enabling Authentication for In-Flight Encryption](#) for instructions.

5. Enable encryption on the port.

See [Enabling In-Flight Encryption](#) for instructions.

6. Enable compression on the port.

See [Enabling In-Flight Compression](#) for instructions.

7. Configure encryption and compression on the E\_Port at the other end of the IFL.

## Configuring In-Flight Encryption and Compression on an E\_Port

To enable and configure in-flight encryption or compression on an E\_Port, perform the following steps to configure the E\_Port.

1. Determine which ports are available for encryption or compression.

See [Viewing the Encryption and Compression Configuration](#) for instructions.

2. If you are enabling encryption on the port, configure port level authentication for the port.

Omit this step if you want to enable only compression on the port.

See [Configuring and Enabling Authentication for In-Flight Encryption](#) for instructions.

3. Enable encryption on the port.

See [Enabling In-Flight Encryption](#) for instructions.

4. Enable compression on the port.

See [Enabling In-Flight Compression](#) for instructions.

Following successful port initialization, the configured features are enabled and active. Use the `islshow` command to check that the E\_Port has come online with encryption or compression enabled. Alternatively, you can use the `portenccompshow` command to see which ports are active.

If port initialization is not successful, you can check for port segmentation errors with the `switchshow` command. This command will tell you if the segmentation was due to mismatched encryption or compression configurations on the ports at either end of the ISL, if port-level authentication failed, or if a required resource was not available.

## Viewing the Encryption and Compression Configuration

Before enabling ports for in-flight encryption or compression, you should determine which ports are available. Enabling encryption or compression fails if you try to exceed the number of allowable ports available for encryption or compression on the ASIC.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `portEncCompShow` command.

The following example shows the output for two ASICs.

ASIC 1 (below the line of dashes) already has compression configured and active on port index numbers 348 and 349. Depending on switch or blade where ports are configured, ASIC 1 may have no more ports available or three ports available for encryption or compression. See [Gen 7 and Gen 6 In-Flight Encryption and Compression Features](#) to determine the ports available per ASIC.

ASIC 0 (above the dashed line) has no ports configured for either encryption or compression. Depending on switch or blade where ports are configured, ASIC 0 may have two or four ports available for encryption or compression. See [Gen 7 and Gen 6 In-Flight Encryption and Compression Features](#) to determine the ports available per ASIC.

```
switch:admin> portenccompshow
```

| Index | Encryption<br>Configured | Active | Compression<br>Configured | Active | Config<br>Speed |                       |
|-------|--------------------------|--------|---------------------------|--------|-----------------|-----------------------|
| ----  | -----                    | -----  | -----                     | -----  | -----           |                       |
| 17    | No                       |        | No                        | No     | No              | 4G                    |
| 18    | No                       |        | No                        | No     | No              | 4G                    |
| 19    | No                       |        | No                        | No     | No              | 4G (output truncated) |
| 149   | No                       |        | No                        | No     | No              | 4G                    |
| 150   | No                       |        | No                        | No     | No              | 4G                    |
| 151   | No                       |        | No                        | No     | No              | 4G                    |
| ----- |                          |        |                           |        |                 |                       |
| 88    | No                       |        | No                        | No     | No              | 4G                    |
| 89    | No                       |        | No                        | No     | No              | 4G                    |
| 90    | No                       |        | No                        | No     | No              | 4G (output truncated) |
| 348   | No                       |        | No                        | Yes    | Yes             | 4G                    |
| 349   | No                       |        | No                        | Yes    | Yes             | 4G                    |
| 350   | No                       |        | No                        | No     | No              | 4G                    |
| 351   | No                       |        | No                        | No     | No              | 4G                    |

The output displays the user port number. For bladed switches, use the `switchShow` command to determine the slot number of a specific user port.

## Configuring and Enabling Authentication for In-Flight Encryption

Authentication and a secret key must be configured and established before configuring in-flight encryption.

To enable authentication between an FC router and an edge fabric switch, you must first bring all EX\_Ports online without using authentication. After this, the front WWN of any online EX\_Port connected to the same switch can be used to configure the secret keys in the edge fabric switch.

You must obtain the WWN of the peer switch to configure the secret key. If you are configuring an EX\_Port on an FC router, you can use the `fcredgeshow` command to obtain the WWN of the switch at the other end of the IFL.

**NOTE**

Only DH-CHAP authentication is supported for in-flight encryption of EX\_Ports.

1. Log in to the switch either using an account with admin permissions or an account with OM permissions for the Authentication RBAC class of commands.

**ATTENTION**

When setting a secret key pair, you are entering the shared secrets in plain text. Use a secure channel, such as SSH or the serial console, to connect to the switch on which you are setting the secrets.

2. Configure DH-CHAP or FCAP for authentication using the `authutil --set` command with the `-a` option.

```
switch:admin> authutil --set -a dhchap
```

```
Authentication is set to dhchap.
```

You can specify any one of the following options:

- `dhchap`
- `fcap`
- `all`

The `dhchap` option sets authentication protocol to DH-CHAP. The `fcap` option sets authentication protocol to FCAP. Although, `all` enables both FCAP and DH-CHAP, the active protocol defaults to FCAP for all ports configured for in-flight encryption.

If `dhchap` is specified, all switches in the fabric must enable DH-CHAP and establish pre-shared secrets. If `fcap` is specified, all switches in the fabric must enable FCAP and use certificates (CA and switch) installed on them. If the protocol is set to `all`, you must establish pre-shared secrets or certificates based on the encryption method selected (DH-CHAP or FCAP).

3. Set the DH group to group 4 using the `authutil --set` command with the `-g` option.

```
switch:admin> authutil --set -g "4"
```

```
DH Group was set to 4.
```

You can specify either "4" or "\*" . The "4" option explicitly enables DH group 4. Although "\*" enables all DH groups (0 through 4), the DH group defaults to group 4 for all ports configured for in-flight encryption.

4. Configure pre-shared keys or certificates based on the encryption method selected (DH-CHAP or FCAP):
  - If DH-CHAP is the configured authentication protocol, use the `secauthsecret --set` command to establish pre-shared secret key at each end of the ISL. It is recommended to use a 32-bit secret for an ISL carrying encrypted or compressed traffic.

```
switch:admin> secauthsecret --set
```

When prompted, enter the WWN for the remote switch and secret strings for the local switch and the remote switch.

- If FCAP is the configured authentication protocol, use the `seccertmgmt` command to generate the public or private key, the CSR, and the passphrase and then import certificates (CA and switch) at both the ends of ISL.

```
switch:admin> seccertmgmt
```

For more information, see [Generating the Key and CSR for FCAP](#).

5. Activate the configured authentication using the `authutil --policy` command to set the switch policy mode to Active or On.

```
switch:admin> authutil --policy -sw active
```

If you are configuring authentication on an EX\_Port, there is no need to set the authentication policy to Active or On. EX\_Ports can operate on any switch authentication policy.

6. Verify the authentication configuration using the `authutil --show` command.

The following example sets up authentication in preparation for in-flight encryption. Specifically, it configures DH-CHAP for authentication, sets the DH group to group 4, sets up a pre-shared secret key, and activates authentication.

```
switch:admin> authutil --set -a dhchap
```

```
Authentication is set to dhchap.
```

```
switch:admin> authutil --set -g "4"
```

```
DH Group was set to 4.
```

```
switch:admin> secauthsecret --set
```

This command is used to set up secret keys for the DH-CHAP authentication. The minimum length of a secret key is 12 characters and maximum 40 characters. Setting up secret keys does not initiate DH-CHAP authentication. If switch is configured to do DH-CHAP, it is performed whenever a port or a switch is enabled.

Warning: Please use a secure channel for setting secrets. Using an insecure channel is not safe and may compromise secrets. Following inputs should be specified for each entry.

1. WWN for which secret is being set up.
2. Peer secret: The secret of the peer that authenticates to peer.
3. Local secret: The local secret that authenticates peer.

```
Press enter to start setting up secrets >
```

```
Enter peer WWN, Domain, or switch name (Leave blank when done): 10:00:00:05:1e:e5:cb:00
```

```
Enter peer secret:
```

```
Re-enter peer secret:
```

```
Enter local secret:
```

```
Re-enter local secret:
```

```
Enter peer WWN, Domain, or switch name (Leave blank when done):
```

```
Are you done? (yes, y, no, n): [no] y
```

```
Saving data to key store... Done.
```

```
switch:admin> secauthsecret --show
```

```
WWN DIId Name

10:00:00:05:1e:e5:cb:00 150 dcx_150
switch:admin> authutil --policy -sw active
```

Warning: Activating the authentication policy requires either DH-CHAP secrets or PKI certificates depending on the protocol selected. Otherwise, ISLs will be segmented during next E-port bring-up.

```
ARE YOU SURE (yes, y, no, n): [no] y
```

```
Auth Policy is set to ACTIVE
switch:admin> authutil --show
```

```

AUTH TYPE HASH TYPE GROUP TYPE

dhchap md5 4
Switch Authentication Policy: ACTIVE
Device Authentication Policy: OFF
```

For additional information about configuring DH-CHAP and FCAP authentication protocols, see [Authentication Policy for Fabric Elements](#).

## Enabling In-Flight Encryption

Enable in-flight encryption to provide security for frames while they are in flight between two switches. Frames are encrypted at the egress point of an ISL and then decrypted at the ingress point.

Enabling encryption is an offline event. Ports must be disabled first, and then re-enabled after.

Before performing this procedure, it is recommended that you check for port availability. Enabling encryption fails if you try to exceed the number of allowable ports available for encryption or compression on the ASIC. See [Viewing the Encryption and Compression Configuration](#) for details.

You must also authenticate the port as described in [Configuring and Enabling Authentication for In-Flight Encryption](#).

1. Connect to the switch and log in using an account with secure admin permissions, or an account with OM permissions for the EncryptionConfiguration RBAC class of commands.
2. Enter the `portdisable` command to disable the port on which you want to configure encryption.

### NOTE

Also, on Brocade G620, disable the ports from 44 to 47.

3. Enter the `portcfgencrypt --enable` command.
4. Enter the `portenable` command to enable the port.

After manually enabling the port, the new configuration becomes active.

5. Verify the configuration using the `portcfgshow` command.

The following example enables the encryption on port 15 of an FC32-48 blade in slot 4 of the X6-8 director:

```

switch:admin> portdisable 4/15
switch:admin> portcfgencrypt --enable 4/15
switch:admin> portenable 4/15

switch:admin> portcfgshow 4/15

Area Number: 111
FC Speed Level: AUTO(HW)
Trunk Port ON
Octet Speed Combo: 3 (16G,10G) (output truncated)
D-Port mode: OFF
D-Port over DWDM ..
Compression: OFF
```

```
Encryption: ON
10G/16G FEC: ON
16G FEC via TTS: OFF
```

## Enabling In-Flight Compression

Enable in-flight compression to provide better bandwidth use on the ISLs, especially over long distance. Frames are compressed at the egress point of an ISL and then decompressed at the ingress point.

Enabling compression is an offline event. Ports must be disabled first, and then re-enabled after.

Before performing this procedure, check for port availability. Enabling compression fails if you try to exceed the number of allowable ports available for encryption or compression on the ASIC. See [Viewing the Encryption and Compression Configuration](#) for details.

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the SwitchPortConfiguration RBAC class of commands.
2. Enter the `portdisable` command to disable the port on which you want to configure compression.
3. Enter the `portcfgcompress --enable` command to enable compression.
4. Enter the `portenable` command to enable the port.

After enabling the port, the new configuration becomes active.

5. Verify the configuration using the `portcfgshow` command.

The following example enables the compression on port 0 of the G630 switch:

```
switch:admin> portdisable 0
switch:admin> portcfgcompress --enable 0
switch:admin> portenable 0

switch:admin> portcfgshow 0

Area Number: 10
FC Speed Level: AUTO(HW)
Trunk Port ON (output truncated)
D-Port mode: OFF
D-Port over DWDM ..
Compression: ON
Encryption: OFF
10G/16G FEC: ON
16G FEC via TTS: OFF
```

## Disabling In-Flight Encryption

Disabling encryption is an offline event. Ports must be disabled first, and then re-enabled after.

1. Connect to the switch and log in using an account with secure admin permissions, or an account with OM permissions for the EncryptionConfiguration RBAC class of commands.
2. Disable the port using the `portdisable` command.
3. Enter the `portcfgencrypt --disable` command.
4. Enter the `portenable` command to enable the port.

After manually enabling the port, the new configuration becomes active.

5. Verify the configuration using the `portcfgshow` command.

The following example disables the encryption on port 15 of an FC32-48 blade in slot 4 of the X6-8 director:

```
switch:admin> portdisable 4/15
switch:admin> portcfgencrypt --disable 4/15
switch:admin> portenable 4/15

switch:admin> portcfgshow 4/15

Area Number: 111
FC Speed Level: AUTO(HW)
Trunk Port ON
Octet Speed Combo: 3 (16G,10G) (output truncated)
D-Port mode: OFF
D-Port over DWDM ..
Compression: OFF
Encryption: OFF
10G/16G FEC: ON
16G FEC via TTS: OFF
```

## Disabling In-Flight Compression

Disabling compression is an offline event. Ports must be disabled first, and then re-enabled after.

### NOTE

Set the DHCPAP protocol using the `authutil --set` command before downgrade.

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the SwitchPortConfiguration RBAC class of commands.
2. Disable the port using the `portdisable` command.
3. Disable compression on the port using the `portcfgcompress --disable` command.
4. Enable the port using the `portenable` command.
5. Verify the configuration using the `portcfgshow` command.

The following example disables the compression on port 0 of the G630 switch:

```
switch:admin> portdisable 0
switch:admin> portcfgcompress --disable 0
switch:admin> portenable 0
```

```
switch:admin> portcfgshow 0

Area Number: 10
FC Speed Level: AUTO(HW)
Trunk Port ON (output truncated)
D-Port mode: OFF
D-Port over DWDM ..
Compression: OFF
Encryption: OFF
10G/16G FEC: ON
16G FEC via TTS: OFF
```

## Encryption and Trunking Support on FC32-48 Blades

Fibre Channel ports that are configured for encryption can be configured in a trunk group on the FC32-48 blade starting with Fabric OS v8.2.1a and later. This is supported on EX\_Ports and E\_Ports.

The following requirements must be met per ASIC.

- Only one encrypted trunk is allowed per ASIC. If a second encrypted trunk is configured within the same ASIC, the second trunk is disabled. If you want to reenable a port within the second encrypted trunk, you must first disable trunking on the port (`portcfgtrunkport 0`), then enable the port (`portenable`).
- A maximum of four ports can be configured within an encrypted trunk to the same destination domain.

For general trunking and trunk group requirements, see [Managing Trunking Connections](#). For requirements for configuring trunking on FC32-48 blade encrypted ports, see the following "Recommendations and Considerations" section.

### NOTE

The `portenccompshow` command displays ports by index number that are within a specific ASIC/group and identifies those ports that are configured and active for compression and encryption. See [Viewing the Encryption and Compression Configuration](#) for more information.

### Trunking Behavior for Encrypted Ports

The following are trunking behaviors that you should note for encrypted ports. These behaviors also exist when both encryption and compression are enabled on a port being enabled for trunking.

- Only the first ports to form a trunk on a logical switch can form a trunk. Ports on a different logical switch will not form a trunk. For example, encrypted ports 0 and 1 are in logical switch A and encrypted ports 2 and 3 are in logical switch B. If ports 0 and 1 form a trunk before ports 2 and 3, ports 2 and 3 is disabled when attempting to form a trunk.
- If trunking and encryption are enabled on FC32-48 blade ports connecting to a local switch, but the remote switch has trunking disabled, no trunking license, or ISL R\_RDY Mode is enabled, the connection will not come online. The port is disabled with an appropriate error.
- When you configure encryption on a port, trunking will be enabled by default. The port attempts to form a trunk when connected to a neighboring switch.
- Non-trunking encrypted ports and trunking encrypted ports can exist on the same ASIC/encryption port group.
- A cold boot does not disable trunking on the blade if both trunking and encryption are enabled.
- Ports that are enabled for compression have always been supported for trunking in previous Fabric OS versions. With Fabric OS 8.2.1a and later, ports that are configured for both compression and encryption or encryption only can be enabled for trunking.



## Recommendations and Considerations

Following are important recommendations and considerations when configuring and using encrypted trunks on the FC32-48 blade.

- If Fabric OS 8.2.1 or later is installed, do not enable trunking on encrypted ports (function is not blocked) to prevent accidental use of the encryption and trunking features.
- If an encrypted trunk group is already present and an encrypted port tries to form another trunk on the same ASIC, the encrypted port will be disabled. To re-enable encryption on the port, disable encryption, disable trunking, and then re-enable encryption.
- Only one port type (E\_Port or EX\_Port) is supported on an ASIC, hence only one port type is available for encryption and trunking. When an E\_Port and EX\_Port configured for encryption and trunking try to come up at the same time, the second port that comes up in order of being enabled will be hard faulted with RASLog message C4-1047. You must first disable encryption on the port to bring it back online. To re-enable the port for encryption, you must first disable trunking on the port.
- When ISL RDY and trunking are configured on FC32-48 blade port, an attempt to configure encryption using `portcfgEncrypt` will be blocked. Encryption is not supported when both ISL R\_RDY mode and trunking are enabled on FC32-48 ports. Either ISL R\_RDY mode (`portCfgISLMode`) or trunking (`portCfgTrunkPort`) must first be disabled before you enable encryption (`portCfgEncrypt`). To enable ISL R\_RDY mode on an encrypted trunk, you must first disable encryption, disable trunking, and then re-enable encryption and ISL R\_RDY mode. Note that ports enabled for ISL R\_RDY mode do not support trunking anyway.
- Encryption is not supported if trunking is configured without a trunking license on FC32-48 ports. If trunking is enabled on ports without a license, disable trunking, and then enable encryption.
- When encryption is configured on an FC32-48 port, an attempt to enable or disable trunking on the port using `portcfgtrunkport` or `switchcfgtrunkport` will be blocked. To enable encryption and trunking on these ports, you must disable encryption, enable trunking, then re-enable encryption.
- Before creating an ISL between the FC32-48 blade and the G620 switch, it is recommend that you disable trunking on encrypted E\_Ports for the ISL since encryption and trunking are not supported simultaneously on G620 ports.

## Firmware Considerations

Following are considerations for firmware support of encrypted trunks on FC32-48 blades:

- To support encryption and trunking on FC32-48 blade ports, both ends of a trunk link must connect to a device running Fabric OS v8.2.1a or later. If one end runs an earlier Fabric OS than v8.2.1a, only one trunk with single port will come online. In this case, disable trunking on encrypted ports for backward compatibility.
- Downgrade to a Fabric OS release prior to v8.2.1a will be prevented if encrypted ports are enabled for trunking. You must disable encryption or trunking on these ports before downgrade.
- An encrypted port on a FC32-48 blade will have trunking disabled during a non-disruptive firmware upgrade from a previous Fabric OS version to v8.2.1a.
- Disable encryption on FC32-48 port blades when migrating Fabric OS from v8.1.x or v8.2.0 to v8.2.1 or later or when downgrading to 8.1.x or 8.2.0 from 8.2.1 or later.

## Buffer-to-Buffer Credits and Credit Recovery

This section describes Buffer-to-buffer credit management, which affects performance over distances; therefore, allocating a sufficient number of buffer credits for the long-distance traffic is essential to performance.

To prevent a target device (either host or storage) from being overwhelmed with frames, the Fibre Channel architecture provides flow-control mechanisms that are based on a system of credits. Each of these credits represents the ability of the device to accept additional frames. If a recipient issues no credits to the sender, no frames can be sent. Pacing the transport of subsequent frames that is based on the credit system helps prevent the loss of frames and reduces the frequency of entire Fibre Channel sequences needing to be retransmitted across the link.

If the number of buffer credits available for use within each port group is limited, configuring buffer credits for extended links may affect the performance of the other ports in the group that is used for core-to-edge connections. You must balance the number of long-distance ISL connections and core-to-edge ISL connections within a switch.

### NOTE

Configuring long-distance ISLs between core and edge switches is possible, but it is not a recommended practice.

All switch ports provide protection against the buffer depletion through buffer limiting. A buffer-limited port reserves a minimum of eight buffer credits, allowing the port to continue to operate rather than being disabled because of a lack of buffers.

Buffer-limited operations are supported for the static mode (LS) and dynamic mode (LD) extended ISL modes only. For LD, distance in kilometers is the smaller of the distance that is measured during the port initialization versus the *desired\_distance* value. For LS, distance in kilometers is always the *desired\_distance* value.

## Buffer-to-Buffer Flow Control

Buffer-to-buffer (BB) credit flow control is implemented to limit the amount of data that a port may send, and it is based on the number and size of the frames sent from that port. Buffer credits represent finite physical-port memory. Within a fabric, each port may have a different number of buffer credits. Within a connection, each side may have a different number of buffer credits.

Buffer-to-buffer flow control is flow control between adjacent ports in the I/O path, for example, transmission control over individual network links. A separate, independent pool of credits manages buffer-to-buffer flow control. A sending port uses its available credit supply and waits to have the credits replenished by the port on the opposite end of the link. These buffer credits are used by Class 2 and Class 3 services and rely on the Fibre Channel Receiver-Ready (R\_RDY) control word to be sent by the receiving link port to the sender. The rate of frame transmission is regulated by the receiving port, and it is based on the availability of buffers to hold received frames.

If Virtual Channel technology is in use, the VC\_RDY or EXT\_VC control word is used instead of the R\_RDY control word to manage buffer credits. For Virtual Channels, the buffer credits are managed for each Virtual Channel, and not for the entire physical link.

The Virtual Channels used in VC\_RDY flow-control mode range from VC0 through VC7.

- **Gen 6 or Gen 7 link to Gen 6 link:** When QoS is enabled, EXT\_VC\_RDY flow-control mode allocates VC0 through VC14. VC8 through VC14 are allocated specifically for QoS VCs.
- **Gen 7 link to Gen 7 link:** When QoS is enabled, EXT\_VC\_RDY flow-control mode allocates VC0 through VC27. VC8 through VC14 are allocated specifically for QoS VCs.

Upon arriving at a receiver, a frame goes through several steps. It is received, deserialized, and decoded, and it is stored in a receive buffer where it is processed by the receiving port. If another frame arrives while the receiver is processing the first frame, a second receive buffer is needed to hold this new frame. Unless the receiver is capable of processing frames

as fast as the transmitter is capable of sending them, it is possible for all of the receive buffers to fill up with received frames. At this point, if the transmitter should send another frame, the receiver will not have a receive buffer available, and the frame is lost. Buffer-to-buffer flow control provides consistent and reliable frame delivery of information from sender to receiver.

## Optimal Buffer Credit Allocation

The optimal number of buffer credits is determined by the distance (frame delivery time), the processing time at the receiving port, the link signaling rate, and the size of the frames being transmitted. As the link speed increases, the frame delivery time is reduced and the number of buffer credits must be increased to obtain full link utilization, even in a short-distance environment.

For each frame that is transferred, the hardware at the other end must acknowledge that the frame has been received before a successful transmission occurs. This flow requires enough capacity in the hardware to allow continuous transmission of frames on the link, while waiting for the acknowledgment to be sent by the receiver at the other end.

As the distance between switches and the link speed increases, additional buffer credits are required for the ports used for long-distance connections. Distance levels define how buffer credits are allocated and managed for extended ISLs. Buffer credits are managed from a common pool available to a group of ports on a switch. The buffer credit can be changed for specific applications or operating environments, but it must be in agreement among all switches to allow formation of the fabric.

Smaller frame sizes need more buffer credits. Two commands are available to help you determine whether you need to allocate more buffer credits to handle the average frame size. The `portBufferShow` command calculates the average frame size. The `portBufferCalc` command uses the average frame size with the speed and link distance to determine the number of buffer credits needed.

## Considerations for Calculating Buffer Credits

The following are the considerations that you need to follow for calculating the number of ports that can be configured for long distance on all Gen 6 or later capable switch modules:

- Each port is part of a port group that includes a pool of buffer credits that can be used. This port group is not the same as the port groups used for ISL trunking.
- Each user port reserves eight buffer credits when online or offline.
- Any remaining buffers can be reserved by any port in the port group.
- When QoS is enabled and the port is online, additional buffers are allocated to that port. See [Calculating the Number of Buffers Required Based on Full-Size Frames](#) and [Configuring Buffers for a Single Port Directly](#) for more information.

## Fibre Channel Gigabit Values Reference Definition

The following table shows the Fibre Channel gigabit values that you can use to calculate buffer requirements.

**Table 74: Fibre Channel Gigabit Values**

| Gigabit Value | Line Rate |
|---------------|-----------|
| 1G            | 1.0625    |
| 2G            | 2.125     |
| 4G            | 4.25      |
| 8G            | 8.5       |
| 10G           | 10.625    |
| 16G           | 14.025    |

| Gigabit Value | Line Rate |
|---------------|-----------|
| 32G           | 28.05     |
| 64G           | 57.8      |

## Buffer Credit Allocation Based on Full-Size Frames

Assuming that the frame is a full-size frame, one buffer credit allows a device to send one payload up to 2112 bytes (2148 with headers). Assuming that each payload is 2,112, you need one credit per 1 km of link length at 2G (smaller payloads require additional buffer credits to maintain link utilization). See [Allocating Buffer Credits Based on Average-Size Frames](#) for additional information.

## Fibre Channel Data Frames

The final frame size must be a multiple of 4 bytes. If the data (payload) needs to be segmented, it will be padded with 1 to 3 "fill-bytes" to achieve an overall 4-byte frame alignment. The standard frame header size is 24 bytes. If applications require extensive control information, up to 64 additional bytes (for a total of an 88-byte header) can be included. Because the total frame size cannot exceed the maximum of 2,148 bytes, the additional header bytes will subtract from the data segment by as much as 64 bytes (per frame). This is why the maximum data (payload) size is 2112 (because  $[2112 - 64] = 2,048$ , which is 2 kb of data). The final frame, after it is constructed, is passed through the 8-byte-to-10-byte conversion process.

The following table describes fibre channel data frames.

**Table 75: Fibre Channel Data Frames**

| Fibre Channel frame fields | Field size              | Final frame size        |
|----------------------------|-------------------------|-------------------------|
| Start of frame             | 4 bytes                 | 32 bits                 |
| Standard frame header      | 24 bytes                | 192 bits                |
| Data (payload)             | 0 to 2112 bytes         | 0 to 16,896 bits        |
| CRC                        | 4 bytes                 | 32 bits                 |
| End of frame               | 4 bytes                 | 32 bits                 |
| Total (number bits/frame)  | <b>36 to 2148 bytes</b> | <b>288 to 7184 bits</b> |

## Allocating Buffer Credits Based on Full-Sized Frames

You can allocate buffer credits based on distance using the `portCfgLongDistance` command. The long-distance link modes allow you to select the dynamic mode (LD) or the static mode (LS) to calculate the buffer credits.

For LD, the estimated distance in kilometers is the smaller of the distance measured during port initialization versus the `desired_distance` parameter, which is required when a port is configured as an LD or an LS mode link. A best practice is to use LS instead of LD. The assumption that Fibre Channel payloads are consistently 2112 bytes is not realistic in practice. To gain the proper number of buffer credits with the LS mode, enough buffer credits must be available in the pool, because Fabric OS checks before accepting a value.

### NOTE

The `desired_distance` parameter of the `portCfgLongDistance` command's is the upper limit of the link distance and calculates buffer availability for other ports in the same port group. When the measured distance exceeds the value of `desired_distance`, this value allocates the buffers. In this case, the port operates in degraded mode instead of being disabled as a result of insufficient buffer availability. In LS mode, the actual link distance is not measured; instead, the `desired_distance` value allocates the buffers required for the port.

See the tables listed in [Buffer Credits per Switch or Blade Model](#) to get the total ports in a switch or blade, the number of user ports in a port group, and the unreserved buffer credits available per port group.

## Calculating the Number of Buffers Required Based on Full-Size Frames

Use the following procedure to calculate the number of buffers required for a long-distance connection:

1. Determine the desired distance in kilometers of the switch-to-switch connection.
2. Determine the speed that you will use for the long-distance connection.
3. Use one of the following formulas to calculate the reserved buffers for distance:

- If QoS is enabled:

$$(\text{Reserved\_Buffer\_for\_Distance\_Y}) = (\times * \text{LinkSpeed} / 2) + 6 + 14$$

- If QoS is not enabled:

$$(\text{Reserved\_Buffer\_for\_Distance\_Y}) = (\times * \text{LinkSpeed} / 2) + 6$$

The formulas use the following parameters:

$\times$  = The distance determined in Step 1 (in km).

*LinkSpeed* = The speed of the link determined in Step 2.

6 = The number of buffer credits reserved for fabric services, multicast, and broadcast traffic. This number is static.

14 = The number of buffer credits reserved for QoS. This number is static.

Using 50 km as the desired distance of the switch-to-switch connection and 2G as the speed of the long-distance connection, insert the numbers into the appropriate formula. The formula should read as follows:

$(50 \text{ km} * 2\text{G} / 2) + 6 = 56$  buffers, which is the number of buffers reserved for distance.

The following examples use different speeds, all based on a distance of 50 km. The distances and speeds are variables that can change depending on how your network is set up.

- If you have a distance of 50 km at 1G,  $(50 \text{ km} * 1\text{G} / 2) + 6 = 31$  buffers.
- If you have a distance of 50 km at 2G,  $(50 \text{ km} * 2\text{G} / 2) + 6 = 56$  buffers.
- If you have a distance of 50 km at 4G,  $(50 \text{ km} * 4\text{G} / 2) + 6 = 106$  buffers.
- If you have a distance of 50 km at 8G,  $(50 \text{ km} * 8\text{G} / 2) + 6 = 206$  buffers.
- If you have a distance of 50 km at 10G,  $(50 \text{ km} * 10\text{G} / 2) + 6 = 256$  buffers.
- If you have a distance of 50 km at 16G,  $(50 \text{ km} * 16\text{G} / 2) + 6 = 406$  buffers.
- If you have a distance of 50 km at 32G,  $(50 \text{ km} * 32\text{G} / 2) + 6 = 806$  buffers.

### Example

Consider the Brocade 300, which has a single 24-port port group and a total of 676 buffer credits for that port group. The formulas use the following parameters:

24 = The number of user ports in a port group

8 = The number of reserved credits for each user port

676 = The number of buffer credits available in the port group

The maximum remaining number of buffer credits for the port group, after each port reserves its 8 buffer credits, is obtained from the following formula:

$$676 - (24 * 8) = 484 \text{ unreserved buffer credits}$$

492 buffers to a single port (484 + 8 [8 for the reserved buffers already allocated to that user port]), you can calculate the maximum single-port extended distance supported:

$$\text{Maximum\_Distance} \times (\text{in km}) = (\text{BufferCredits} + 6) * 2 / \text{LinkSpeed}$$

$$498 \text{ km} = (492 + 6 \text{ buffers for Fabric Services}) * 2 / 2\text{G}$$

If you have a distance of 50 km at 8G, then  $484 / (206 - 8) = 2$  ports.

The following values are used in the example:

- 484 – The total number of unreserved buffer credits
- 206 – Buffer credits needed for 50 km at 8G
- 8 – The number of reserved buffer credits already allocated to that port

The resulting number is rounded down to the next whole number because fractions of a port are not allowed.

If you have a distance of 50 km at 1G, then  $484 / (31 - 8) = 21$  ports.

## Allocating Buffer Credits Based on Average-Size Frames

In cases where the frame size is average, for example, 1024 bytes, you must allocate twice the buffer credits or configure twice the distance in the long-distance LS configuration mode. Refer to [Fibre Channel Gigabit Values Reference Definition](#) for an approximation of the calculated number of buffer credits.

1. Use the following formula to calculate the value for the *desired\_distance* parameter needed for Fabric OS to determine the number of buffer credits to allocate:

$$\text{desired\_distance} = \text{roundup} [(\text{real\_estimated\_distance} * 2112) / \text{average\_payload\_size}]$$

The *average\_payload\_size* in this equation uses 1024 bytes.

If the real estimated distance is 100 km, the *desired\_distance* is 207.

$$\text{desired\_distance} = \text{roundup} [(100 * 2112) / 1024] = 207$$

When configuring the LS mode with the `portCfgLongDistance` command, enter a *desired\_distance* value of 207 for an actual 100-km link connected to an 8G E\_Port, which causes Fabric OS to allocate the correct number of buffer credits.

2. Determine the speed you will use for the long-distance connection. This example uses 8G.
3. Look up the *data\_rate* value for the speed of the connection. Refer to [Fibre Channel Gigabit Values Reference Definition](#) to determine the *data\_rate* value.

For example, the *data\_rate* is 8.5 for a speed of 8G.

4. Use the following formula to calculate the number of buffer credits to allocate:

$$\text{buffer\_credits} = [\text{desired\_distance} * (\text{data\_rate} / 2.125)]$$

With the values for *desired\_distance* and *data\_rate* from Step 1 and Step 3, the value for buffer credits is calculated as follows:

$$\text{buffer\_credits} = [207 * (8.5 / 2.125)] = 828$$

### NOTE

This *buffer credits* formula does not work with LD mode because LD mode checks the distance and limits the estimated distance to the real value of 100 km. LS mode allows for the necessary *desired\_distance* value based on the data size entered, regardless of the distance.

If buffer credit recovery is enabled, Fabric OS supports a *BB\_SC\_N* range of 1 to 8; therefore, it is impossible for the *desired\_distance* value to be more than the number of buffer credits available in the pool as determined by

the previous calculations. The distance for buffer credit recovery is well within the range of all possible connections. An estimated distance of 32,768 is considerably higher than the available buffer credits and only lower values of *desired\_distance* are permitted by Fabric OS.

## Configuring Buffers for a Single Port Directly

To configure the number of buffers directly, use the `-buffers` option of the `portCfgLongDistance` command. Fabric OS uses this value to calculate the total number of buffers according to the following formula:

$$\text{Total Buffers} = \text{Configured Buffers} + \text{QOS\_VC\_Credits} + \text{Non-data\_VC\_Credits}$$

Seven Virtual Channels (VCs) are required for each QoS port. Each VC requires two buffers. Thus, the total number of QoS buffers required for a port is 14 (7\*2). An additional six VCs are required for nondata transmission (for example, control traffic). As a consequence, for a QoS port, 20 buffers are added. For a non-QoS port, six buffers are added.

For example, if the configured number of buffers is 100, then the total number of buffers allocated for a QoS port is 120, as shown in the following example.

$$\text{Total\_Buffers} = 100 + 14 + 6 = 120$$

If the configured number of buffers is 100, the total number of buffers allocated for a non-QoS port is 106, as shown in the following example.

$$\text{Total\_Buffers} = 100 + 6 = 106$$

### NOTE

You cannot use the `-buffers` option with the `-distance` option or the `-frameSize` option.

```
switch:admin> portcfglongdistance 2/35 LS 1 -buffers 400
Reserved Buffers = 420
```

## Configuring Buffers Using Frame Size

You can configure the number of buffers by using the `-frameSize` option of the `portCfgLongDistance` command along with the `-distance` option. Fabric OS calculates the number of buffers from the `-frameSize` option value according to the following formula:

$$\text{buffers\_required} = (2048/\text{framesize}) * \text{data\_vc\_credits}$$

If you enter the average frame size of 1024, Fabric OS allocates almost twice as many buffers as for the maximum frame size of 2048.

The `-frameSize` option value is persistent across reboots and HA failover.

For example,

```
switch:admin> portcfglongdistance 2/35 LS 1 -distance 100 -framesize 1024
```

## Calculating the Number of Buffers Required Given the Distance, Speed, and Frame Size

If you know the distance, speed, and frame size for a given port, you can use the `portBufferCalc` command to calculate the number of buffers required. If you omit the distance, speed, or frame size, the command uses the currently configured values for the port. Given the buffer requirement and port speed, you can use the same distance and frame size values when using the `portCfgLongDistance` command.



To determine the number of buffers required, complete the following steps:

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portBufferCalc` command and provide values for the distance, port speed, and frame size.

The following example calculates the number of buffers required for an 8G port on a 100-km link with an average frame size of 512 bytes.

```
switch:admin> portbuffercalc 9/4 -distance 100 -speed 8 -framesize 512
1606 buffers required for 100km at 8G and framesize of 512 bytes
```

#### NOTE

The `portBufferCalc` and `portCfgEportCredits` are applicable to the 2-km ICL QSFP ports as well.

The following examples show calculation of buffers for 2-km ICL ports at 16G and 32G:

```
switch_6510:admin> portbuffercalc 1 -distance 2 -speed 16 -framesize 2048
22 buffers required for 2km at 16G and framesize of 2048bytes
```

```
switch_6510:admin> portbuffercalc 1 -distance 2 -speed 16 -framesize 1024
38 buffers required for 2km at 16G and framesize of 1024bytes
```

```
switch_6510:admin> portbuffercalc 1 -distance 2 -speed 16 -framesize 512
70 buffers required for 2km at 16G and framesize of 512bytes
```

```
switch_6510:admin> portbuffercalc 1 -distance 1 -speed 16 -framesize 2048
14 buffers required for 1km at 16G and framesize of 2048bytes
```

```
switch_6510:admin> portbuffercalc 1 -distance 1 -speed 16 -framesize 1024
22 buffers required for 1km at 16G and framesize of 1024bytes
```

```
switch_6510:admin> portbuffercalc 1 -distance 1 -speed 16 -framesize 512
38 buffers required for 1km at 16G and framesize of 512bytes
```

```
switch_G620:admin> portbuffercalc 1 -distance 2 -speed 32 -framesize 2048
38 buffers required for 2km at 32G and framesize of 2048bytes
```

```
switch_G620:admin> portbuffercalc 1 -distance 2 -speed 32 -framesize 1024
70 buffers required for 2km at 32G and framesize of 1024bytes
```

```
switch_G620:admin> portbuffercalc 1 -distance 2 -speed 32 -framesize 512
134 buffers required for 2km at 32G and framesize of 512bytes
```

```
switch_G620:admin> portbuffercalc 1 -distance 1 -speed 32 -framesize 2048
22 buffers required for 1km at 32G and framesize of 2048bytes
```

```
switch_G620:admin> portbuffercalc 1 -distance 1 -speed 32 -framesize 1024
38 buffers required for 1km at 32G and framesize of 1024bytes
```

```
switch_G620:admin> portbuffercalc 1 -distance 1 -speed 32 -framesize 512
70 buffers required for 1km at 32G and framesize of 512bytes
```



## Allocating Buffer Credits for F\_Ports

The default configured F\_Port buffer credit is fixed at 28 buffers for all Gen 7 devices and 20 buffers for all Gen 6 devices, except the Brocade G610, which is fixed at eight buffers. Use the `portcfgfportbuffers` command to configure a given port with the specified number of buffers.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portcfgfportbuffers` command.

```
switch:admin> portcfgfportbuffers --enable 2/44 12
```

Notice that in the sample commands provided in the following procedure, 12 buffers are configured for an F\_Port.

To disable the port buffer configuration and return to the default buffer allocation, use the `--disable` option.

```
switch:admin> portcfgfportbuffers --disable 2/44
```

### NOTE

The configured number of buffers for a given port is stored in the configuration database and is persistent across reboots. The F\_Port buffer feature does not support EX\_Port, Long-Distance, L\_Port, FastWrite, QoS, and Trunk Area enabled ports. F\_Port buffers are mutually exclusive to E\_Port credits

## Monitoring Buffers in a Port Group

Use the `portBufferShow` command to monitor the remaining buffers on a long-distance link and to monitor the average frame size and average buffer usage for a given port.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portBufferShow` command.

The average frame size in bytes is shown in parentheses with the average buffer usage for packet transmission and reception.

```
switch:admin> portbuffershow
```

| User Port | Port Type | Lx Mode | Max/Resv Buffers | Avg Buffer Usage & FrameSize |         | Buffer Needed |         | Link Distance | Remaining Buffers |
|-----------|-----------|---------|------------------|------------------------------|---------|---------------|---------|---------------|-------------------|
|           |           |         |                  | Tx                           | Rx      | Usage         | Buffers |               |                   |
| 0         | F         | -       | 91               | - ( - )                      | - ( - ) | 91            | -       | -             |                   |
| 1         | F         | -       | 91               | - ( - )                      | - ( - ) | 91            | -       | -             |                   |
| 2         |           | -       | 28               | - ( - )                      | - ( - ) | 0             | -       | -             |                   |
| 3         | E         | -       | 28               | - ( - )                      | - ( - ) | 188           | 188     | <2km          |                   |
| 4         | E         | -       | 28               | - ( - )                      | - ( - ) | 94            | 94      | <2km          |                   |
| 5         | E         | -       | 28               | - ( - )                      | - ( - ) | 160           | 160     | <2km          |                   |
| 6         | E         | -       | 28               | - ( - )                      | - ( - ) | 70            | 70      | <2km          |                   |
| 7         |           | -       | 91               | - ( - )                      | - ( - ) | 91            | -       | -             |                   |
| 8         | F         | -       | 91               | - ( - )                      | - ( - ) | 91            | -       | -             |                   |
| 9         |           | -       | 28               | - ( - )                      | - ( - ) | 0             | -       | -             |                   |
| 10        |           | -       | 94               | - ( - )                      | - ( - ) | 0             | -       | -             |                   |
| 11        |           | -       | 28               | - ( - )                      | - ( - ) | 0             | -       | -             |                   |
| 12        |           | -       | 28               | - ( - )                      | - ( - ) | 0             | -       | -             |                   |
| 13        |           | -       | 28               | - ( - )                      | - ( - ) | 0             | -       | -             |                   |
| 14        |           | -       | 28               | - ( - )                      | - ( - ) | 0             | -       | -             |                   |
| 15        |           | -       | 28               | - ( - )                      | - ( - ) | 0             | -       | -             |                   |

|    |   |    |      |         |         |      |      |       |
|----|---|----|------|---------|---------|------|------|-------|
| 16 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 17 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 18 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 19 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 28 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 21 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 22 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 23 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 24 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 25 | D | LS | 8006 | - ( - ) | - ( - ) | 8006 | 8006 | 500km |
| 26 | D |    | 28   | - ( - ) | - ( - ) | 188  | 188  | <2km  |
| 27 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 28 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 29 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 30 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 31 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 32 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 33 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 34 | E | -  | 28   | - ( - ) | - ( - ) | 94   | 94   | <2km  |
| 35 | E | -  | 28   | - ( - ) | - ( - ) | 94   | 94   | <2km  |
| 36 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 37 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 38 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 39 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 40 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 41 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 42 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 43 | F | -  | 91   | - ( - ) | - ( - ) | 91   | -    | -     |
| 44 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 45 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 46 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 47 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 48 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 49 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 50 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 51 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 52 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 53 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 54 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 55 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 56 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 57 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 58 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 59 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 60 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 61 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 62 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |
| 63 |   | -  | 28   | - ( - ) | - ( - ) | 0    | -    | -     |

11887

## Buffer Credits per Switch or Blade Model

The following tables show the total FC ports in a switch or blade, the number of user ports in a port group, and the unreserved buffer credits available per port group.

### Formula for Gen 7 64G Devices

Remaining buffers = Total Buffers - [(Total FE Ports \* 28) + (Total BE/BI Ports \* Buffers for BE/BI Port) + Internal buffer usages (approximately 286).

**Table 76: Total FC Ports, Ports per Port Group, and Unreserved Buffer Credits per Port Group**

| Switch/<br>Blade Model | Total FC<br>Ports (per<br>Switch/Blade) | User Port<br>Group Size | Unreserved Buffer<br>Credits per Port<br>Group (with Minimum<br>Buffer Allocation) | Unreserved Buffer<br>Credits per Port<br>Group (with QoS) | Unreserved Buffer<br>Credits per Port<br>Group (without QoS) |
|------------------------|-----------------------------------------|-------------------------|------------------------------------------------------------------------------------|-----------------------------------------------------------|--------------------------------------------------------------|
| G730                   | 128                                     | 16                      | 658                                                                                | 50                                                        | 50                                                           |
| G720                   | 64                                      | 64                      | 22338                                                                              | 6658                                                      | 16578                                                        |
| FC32-X7-48             | 48                                      | 24                      | 12582                                                                              | 8550                                                      | 11382                                                        |
| FC64-48                | 48                                      | 24                      | 16466                                                                              | 10586                                                     | 14306                                                        |
| CR64-8                 | 64                                      | 16                      | 5650                                                                               | 4766                                                      | 4766                                                         |
| CR64-4                 | 32                                      | 16                      | 5650                                                                               | 4766                                                      | 4766                                                         |

If you enable ECB on a G720 on any port of that port group, the numbers for that port group are changed as per the following table.

**Table 77: Total FC Ports, Ports per Port Group, and Unreserved Buffer Credits per Port Group**

| Switch/<br>Blade Model | Total FC<br>Ports (per<br>Switch/Blade) | User Port<br>Group Size | Unreserved Buffer<br>Credits per Port<br>Group (with Minimum<br>Buffer Allocation) | Unreserved Buffer<br>Credits per Port<br>Group (with QoS) | Unreserved Buffer<br>Credits per Port<br>Group (without QoS) |
|------------------------|-----------------------------------------|-------------------------|------------------------------------------------------------------------------------|-----------------------------------------------------------|--------------------------------------------------------------|
| G720                   | 64                                      | 64                      | 20938                                                                              | 5258                                                      | 15178                                                        |
| FC32-X7-48             | 48                                      | 24                      | 11182                                                                              | 7150                                                      | 9982                                                         |
| FC64-48                | 48                                      | 24                      | 15066                                                                              | 9186                                                      | 12906                                                        |

### Formula for Gen 6 32G Devices

Remaining buffers = Total Buffers - [(Total FE Ports \* 20) + (Total BE/BI Ports \* Buffers for BE/BI Port) + Internal buffer usages (approximately 152).

**Table 78: Total FC Ports, Ports per Port Group, and Unreserved Buffer Credits per Port Group**

| Switch/<br>Blade Model | Total FC<br>Ports (per<br>Switch/Blade) | User Port<br>Group Size | Unreserved Buffer<br>Credits per Port<br>Group (with Minimum<br>Buffer Allocation) | Unreserved Buffer<br>Credits per Port<br>Group (with QoS) | Unreserved Buffer<br>Credits per Port<br>Group (without QoS) |
|------------------------|-----------------------------------------|-------------------------|------------------------------------------------------------------------------------|-----------------------------------------------------------|--------------------------------------------------------------|
| G610                   | 24                                      | 24                      | 1728                                                                               | 960                                                       | 1296                                                         |
| G620                   | 64                                      | 64                      | 13928                                                                              | 8488                                                      | 10728                                                        |
| G630                   | 128                                     | 32                      | 6328                                                                               | 3608                                                      | 4728                                                         |

| Switch/<br>Blade Model | Total FC<br>Ports (per<br>Switch/Blade) | User Port<br>Group Size | Unreserved Buffer<br>Credits per Port<br>Group (with Minimum<br>Buffer Allocation) | Unreserved Buffer<br>Credits per Port<br>Group (with QoS) | Unreserved Buffer<br>Credits per Port<br>Group (without QoS) |
|------------------------|-----------------------------------------|-------------------------|------------------------------------------------------------------------------------|-----------------------------------------------------------|--------------------------------------------------------------|
| FC32-48                | 48                                      | 24                      | 8152                                                                               | 6112                                                      | 6952                                                         |
| FC32-64                | 64                                      | 32                      | 7848                                                                               | 5128                                                      | 6248                                                         |
| CR32-8                 | 64                                      | 16                      | 4376                                                                               | 3016                                                      | 3576                                                         |
| CR32-4                 | 32                                      | 16                      |                                                                                    |                                                           |                                                              |
| SX6                    | 16                                      | 16                      | 5432                                                                               | 4072                                                      | 4632                                                         |

If you enable ECB on a G630 on any port of that port group, the numbers for that port group are changed as per the following table.

**Table 79: Total FC Ports, Ports per Port Group, and Unreserved Buffer Credits per Port Group**

| Switch/<br>Blade Model | Total FC<br>Ports (per<br>Switch/<br>Blade) | User Port<br>Group Size | Unreserved Buffer<br>Credits per Port<br>Group (with Minimum<br>Buffer Allocation) | Unreserved<br>Buffer Credits<br>per Port Group<br>(with QoS) | Unreserved<br>Buffer Credits<br>per Port Group<br>(without QoS) | Unreserved<br>Buffer Credits<br>per Port Group<br>(with ECB Port) |
|------------------------|---------------------------------------------|-------------------------|------------------------------------------------------------------------------------|--------------------------------------------------------------|-----------------------------------------------------------------|-------------------------------------------------------------------|
| 7810                   | 12                                          | 12                      | 1376                                                                               | 1064                                                         | 1160                                                            | N/A                                                               |
| G630                   | 128                                         | 32                      | 6328                                                                               | 3608                                                         | 4728                                                            | 1816                                                              |
| FC32-64                | 64                                          | 32                      | 7848                                                                               | 5128                                                         | 6248                                                            | N/A                                                               |

For the FC32-xx port blade and CR32-xx core blades, the first number in the "Unreserved buffer credits per port group" column designates the number of unreserved buffers per port group without buffer optimized mode; the second number designates the unreserved buffers with buffer optimized mode that is enabled on the slot. Use the `bufOpMode` command to display or change the buffer optimized mode. Only `bufOpMode --setall` and `bufOpMode --resetall` commands are allowed in Gen 6 devices. Downgrade is also blocked in `bufOpMode`.

#### Considerations and limitations:

- On the FC32-48 port blade and the G620 switch, when a port is configured with in-flight encryption, an additional 105 buffer credits are reserved for the encryption purpose. For example, the required buffer credits for an in-flight encryption-enabled, non-QoS, online E\_Port is 175 (70 + 105) buffer credits and for an in-flight encryption-enabled, QoS-enabled, online E\_Port is 210 (105 + 105) buffer credits.
- You must use the `portcfgencrypt` command with the `-force` option for the ports 40 through 47 on the G620 switch, and it has the following limitations:
  - You cannot have both Octet Combo Mode 2 and in-flight encryption that is enabled at the same time.
  - You cannot have both SIM port (using the `flow -control -enable` command) and in-flight encryption that is enabled at the same time.
  - Before you enable in-flight encryption, you must save the configuration, because when encryption is enabled, the ports are set to the default configuration, which wipes out any current configuration.
- If in-flight encryption is enabled on any port in the G620 switch, ports 40 through 47 cannot be configured to Octet Combo Mode 2.

## Minimum Buffers Allocated per Port Type

**Table 80: Minimum Buffers Allocated per Port Type**

| Port Type                                                                     | Minimum Buffers Allocated per Port |
|-------------------------------------------------------------------------------|------------------------------------|
| 64G Gen 7- F_Port                                                             | 28                                 |
| 64G Gen 7- E_Port (non QoS) or F_Port trunk without QoS                       | 118                                |
| 64G Gen 7- E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS         | 273                                |
| 32G Gen 7- F_Port                                                             | 20                                 |
| 32G Gen 7- E_Port (non QoS) or F_Port trunk without QoS                       | 70                                 |
| 32G Gen 7- E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS         | 188                                |
| 16/10/8/4G Gen 7 - F_Port                                                     | 8                                  |
| 16/10/8/4G Gen 7- E_Port (non QoS), or F_Port trunk without QoS               | 28                                 |
| 16/10/8/4G Gen 7 - E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS | 94                                 |
| 32G Gen 6 - F_Port                                                            | 20                                 |
| 32G Gen 6 - E_Port (non QoS) or F_Port trunk without QoS                      | 70                                 |
| 32G Gen 6 - E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS        | 105                                |
| 16/10/8/4G Gen 6 - F_Port                                                     | 8                                  |
| 16/10/8/4G Gen 6 - E_Port (non QoS), or F_Port trunk without QoS              | 28                                 |
| 16/10/8/4G Gen 6 - E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS | 34                                 |

## Minimum Buffers Allocated per Port Type for G730 Platform

**Table 81: Minimum Buffers Allocated per Port Type**

| Port Type                                                                  | Minimum Buffers Allocated per Port |
|----------------------------------------------------------------------------|------------------------------------|
| 64G Gen 7- F_Port                                                          | 20                                 |
| 64G Gen 7- E_Port (non-QoS) or F_Port trunk without QoS                    | 20                                 |
| 64G Gen 7- E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS      | 58                                 |
| 32G Gen 7- F_Port                                                          | 20                                 |
| 32G Gen 7- E_Port (non-QoS) or F_Port trunk without QoS                    | 20                                 |
| 32G Gen 7- E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS      | 58                                 |
| 16/10/8 Gen 7 - F_Port                                                     | 8                                  |
| 16/10/8 Gen 7- E_Port (non-QoS), or F_Port trunk without QoS               | 8                                  |
| 16/10/8 Gen 7 - E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS | 58                                 |

## Maximum Configurable Distances for Extended Fabrics

The following table shows the maximum supported extended distances (in kilometers) that can be configured for one port on a specific switch or blade at different speeds.

**Formula for Gen 7 64G devices**

The remaining buffers = (Total FE Ports \* 28) + (Total BE/BI Ports \* Buffers for BE/BI Port) + Internal buffer usages (approximately 304)

Long distance Calculation: Buffers = distance \* speed \* 0.5

Distance = (Buffers \* 2) / speed

**Table 82: Configurable Distances for Gen 7 Extended Fabrics**

| Switch/<br>Blade Model | Maximum Distances (km) that can be Configured (Assuming a 2112-byte Frame Size) |                                                                   |                                                                   |                                                                   |                                                                  |                                                                |
|------------------------|---------------------------------------------------------------------------------|-------------------------------------------------------------------|-------------------------------------------------------------------|-------------------------------------------------------------------|------------------------------------------------------------------|----------------------------------------------------------------|
|                        | 4G                                                                              | 8G                                                                | 10G                                                               | 16G                                                               | 32G                                                              | 64G                                                            |
| G730                   | N/A                                                                             | 160                                                               | 125                                                               | 80                                                                | 40                                                               | 20                                                             |
| G720                   | N/A                                                                             | 5584 (with min. buffers)<br>1644 (with QoS)<br>4144 (without QoS) | 4467 (with min. buffers)<br>1331 (with QoS)<br>3315 (without QoS) | 2792 (with min. buffers)<br>832 (with QoS)<br>2072 (without QoS)  | 1396 (with min. buffers)<br>416 (with QoS)<br>1036 (without QoS) | 698 (with min. buffers)<br>208 (with QoS)<br>518 (without QoS) |
| FC64-48                | N/A                                                                             | 4116 (with min. buffers)<br>2646 (with QoS)<br>3576 (without QoS) | 293 (with min. buffers)<br>2117 (with QoS)<br>2861 (without QoS)  | 2058 (with min. buffers)<br>1323 (with QoS)<br>1788 (without QoS) | 1029 (with min. buffers)<br>661 (with QoS)<br>894 (without QoS)  | 514 (with min. buffers)<br>330 (with QoS)<br>447 (without QoS) |
| FC32-X7-48             | 6291 (with min. buffers)<br>4275 (with QoS)<br>5691 (without QoS)               | 3145 (with min. buffers)<br>2137 (with QoS)<br>2845 (without QoS) | 2516 (with min. buffers)<br>1710 (with QoS)<br>2276 (without QoS) | 1572 (with min. buffers)<br>1068 (with QoS)<br>1422 (without QoS) | 786 (with min. buffers)<br>534 (with QoS)<br>711 (without QoS)   | N/A                                                            |

**NOTE**

The distances in the previous table assume that QoS is enabled. If QoS is disabled, the maximum supported distances are higher, because QoS requires an additional 20 buffer credits per active port.

*Estimated maximum equally distributed distance = 1-port maximum distance / Number of ports*

**Formula for Gen 6 32G devices**

The remaining buffers = (Total FE Ports \* 20) + (Total BE/BI Ports \* Buffers for BE/BI Port) + Internal buffer usages (approximately 152)

Long distance Calculation: Buffers = distance \* speed \* 0.5

Distance = (Buffers \* 2) / speed

**Table 83: Configurable Distances for Gen 6 Extended Fabrics**

| Switch/Blade Model | Maximum Distances (km) that can be Configured (Assuming a 2112-byte Frame Size)     |                                                                                    |                                                                                   |                                                                                  |                                                                                  |
|--------------------|-------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
|                    | 4G                                                                                  | 8G                                                                                 | 10G                                                                               | 16G                                                                              | 32G                                                                              |
| 7810               | *                                                                                   | *                                                                                  | NA                                                                                | *                                                                                | *                                                                                |
| G610               | 861 (with min. buffers)<br>549 (with QoS)<br>645 (without QoS)                      | 430 (with min. buffers)<br>274 (with QoS)<br>322 (without QoS)                     | N/A                                                                               | 215 (with min. buffers)<br>137 (with QoS)<br>161 (without QoS)                   | 107 (with min. buffers)<br>59 (with QoS)<br>80 (without QoS)                     |
| G620               | 6068                                                                                | 3034                                                                               | 2427                                                                              | 1517                                                                             | 758                                                                              |
| G630               | 3164 (with min. buffers)<br>1804 (with QoS)<br>2364 (without QoS)<br>908 (with ECB) | 1582 (with min. buffers)<br>902 (with QoS)<br>1182 (without QoS)<br>454 (with ECB) | 1265 (with min. buffers)<br>721 (with QoS)<br>945 (without QoS)<br>363 (with ECB) | 791 (with min. buffers)<br>451 (with QoS)<br>591 (without QoS)<br>227 (with ECB) | 395 (with min. buffers)<br>225 (with QoS)<br>295 (without QoS)<br>113 (with ECB) |
| FC32-48            | 4073                                                                                | 2036                                                                               | 1629                                                                              | 1018                                                                             | 509                                                                              |
| FC32-64            | 3924 (with min. buffers)<br>2564 (with QoS)<br>3124 (without QoS)                   | 1962 (with min. buffers)<br>1282 (with QoS)<br>1562 (without QoS)                  | 1569 (with min. buffers)<br>1025 (with QoS)<br>1249 (without QoS)                 | 981 (with min. buffers)<br>641 (with QoS)<br>781 (without QoS)                   | 490 (with min. buffers)<br>320 (with QoS)<br>390 (without QoS)                   |
| SX6                | 2713                                                                                | 1356                                                                               | 1085                                                                              | 678                                                                              | 339                                                                              |

**NOTE**

The distances in the previous table assume that QoS is enabled. If QoS is disabled, the maximum supported distances are higher, because QoS requires an additional 20 buffer credits per active port.

*Estimated maximum equally distributed distance = 1-port maximum distance / Number of ports.* For example, for three ports running at 32G on a Brocade G620 switch, the maximum equally distributed distance is calculated as  $758/3 = 252$  km.

## Configuring Credits for a Single VC

You can alter the default credit allocation for a normal distance E\_Port or EX\_Port so that a specific number of credits is allocated to a port. When you allocate a specific number of credits to an E\_Port or EX\_Port, the number of credits specified override the default credit allocation. When this feature is disabled, the default credit model is restored. Only a normal distance E\_Port and EX\_Port can utilize the new credit model, and the allocated credits are reserved only for that port.

When this feature is enabled, the E\_Port credit configuration is persistent across system reboots and High Availability (HA) failover.

This feature is supported on E\_Ports and EX\_Ports. It does not support ports configured as F\_Ports, Mirror Ports, L\_Ports, and Longdistance Ports. If E\_Port credits are configured on ports, you cannot move the ports from one logical switch to another. You can also allocate buffers for 2-km ICL ports.

For Gen 7 devices (64G-capable blades), the minimum credit allocation is 5 and the maximum is 40. For ICL ports of Gen 7 devices, the valid range is from 5 through 16. For Gen 6 devices, the credit allocation is allowed in the range between 5 to 160 for both ICL ports and normal ISL and EX\_Port links. In Gen 6 devices, if the specified credit allocation is less than the default allocation, the default allocation scheme is enforced (in Gen 6 devices, except for the Brocade G610, 32G links are reserved 20 credits per VC). The configured credits will be allocated for each of the medium VCs for the non-QoS ports. For QoS ports, after sharing, both the medium and high VCs will have the configured credits allocated.

## Increasing Credits for Normal Distance E\_Ports

Use the following steps to allocate credits to an E\_Port.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portCfgEPortCredits --enable 12/6 14` command to allocate credits to an E\_Port. In the following example, 14 credits are allocated to each of the medium VCs for non-QoS ports, and to both the medium and high VCs for QoS ports.

```
switch:admin> portcfgportcredits --enable 12/6 14
Success
```

3. Enter the `portCfgEPortCredits --show 12/6` command to verify that the credits have been allocated to the E\_Port. In the following example, it is verified that 14 credits per VC have been allocated to the E\_Port.

```
switch:admin> portcfgportcredits --show 12/6
E-Port Credit Configured : 14
Success.
```

## Buffer Credit Recovery

Buffer credit recovery allows links to recover after buffer credits are lost when the buffer credit recovery logic is enabled. The buffer credit recovery feature also maintains performance. If a credit is lost, a recovery attempt is initiated. During link reset, the frame and credit loss counters are reset without performance degradation.

Credit recovery is supported on E\_Ports, F\_Ports, and EX\_Ports. Buffer credit recovery is enabled automatically across any long-distance connection for which the E\_Port, F\_Port, or EX\_Port buffer credit recovery mechanism is supported.

For 64G (Gen 7) and 32G (Gen 6) FC switches and blades, use the `portCfgCreditRecovery` command to disable or enable buffer credit recovery on a port.

### NOTE

The credit recovery mode is supported when a link is in ISL\_RDY mode and VC Init is off (only on VC0). If VC init is off, and ISL RDY mode is set, VC0 supports credit recovery.

## Buffer Credit Recovery over an E\_Port

To support buffer credit recovery, E\_Ports must be connected between devices that support 64G, 32G, 16G or between devices that support 8G. If a device that supports 64G or 32G or 16G is connected to a device that supports only 8G, buffer credit recovery is disabled, even if both devices are running 8G.

A limited amount of credit recovery that occurs between 64G, 32G, 16G platforms and 8G platforms. For buffer credit recovery to occur between 64G devices and 32G devices, or 16G devices and 8G devices, the long-distance mode must be enabled for the link.

The buffer credit recovery feature for E\_Ports is enabled for the following flow-control modes:

- Normal (R\_RDY)
- Virtual Channel (VC\_RDY)
- Extended VC (EXT\_VC\_RDY)

## Buffer Credit Recovery over an F\_Port

Buffer credit recovery for F\_Ports is supported for F\_Port-to-N\_Port links between a Brocade switch and Access Gateway, between a Brocade switch and an adapter, and between an Access Gateway and an adapter. For an F\_Port on a Brocade switch connected to an Access Gateway, the following conditions must be met:



- Both devices must run Fabric OS v8.1 or later.
- Fabric OS must support buffer credit recovery at either end of the link.
- If both devices support 16G, the flow-control mode can be either Normal mode (R\_RDY) or VC mode (VC\_RDY); otherwise, the flow-control mode must be R\_RDY.

For an F\_Port on a Brocade switch or Access Gateway connected to an adapter, the following conditions must be met:

- The Brocade switch or Access Gateway must run Fabric OS v8.1 or later.
- Fabric OS must support buffer credit recovery at both ends of the link.
- The adapter must be running QLogic or Emulex HBAs. For more information, refer to *QLogic or Emulex HBA documentation*.
- The flow-control mode must be R\_RDY.

The feature is enabled automatically during a link reset if the conditions are met. If the conditions for buffer credit recovery are not met, the link will come up, but buffer credit recovery will not be enabled.

## Buffer Credit Recovery over an EX\_Port

Buffer credit recovery is supported on a Fibre Channel router (FCR) EX\_Port that connects over an inter-fabric link (IFL) to an edge fabric E\_Port when the following conditions are met:

- The FCR and the switch at the other end of the IFL must both run Fabric OS v8.1 or later.
- The FCR and the switch at either end of the IFL must be Gen 7 (64G) or Gen 6 (32G) platforms.
- Either end of the IFL must support buffer credit recovery.
- Long-distance mode can be enabled or disabled on devices that support 8G or above.
- Virtual Channel flow control (VC\_RDY) mode or Extended VC flow control (EXT\_VC\_RDY) mode must be in use. Buffer credit recovery is not supported for EX\_Ports when normal (R\_RDY) flow control mode is in use.

The feature is enabled automatically during a link reset if the conditions are met. If the capabilities at either end of the EX\_Port-to-E\_Port link are not matched, the link will come up with the buffer credit recovery feature enabled.

## Enabling and Disabling Buffer Credit Recovery

### NOTE

By default, buffer credit recovery is not enabled on any of the platforms.

1. To disable buffer credit recovery on a port, perform the following steps.
  - a) Connect to the switch and log in using an account assigned to the admin role.
  - b) Enter the `portCfgCreditRecovery` command and include the `-disable` option.

The following example disables buffer credit recovery on port 1/20.

```
switch:admin> portcfgcreditrecovery 1/20 -disable
```

2. To enable buffer credit recovery on a port for which it is disabled, perform the following steps.
  - a) Connect to the switch and log in using an account assigned to the admin role.
  - b) Enter the `portCfgCreditRecovery` command and include the `enable` option.

The following example enables buffer credit recovery on port 1/20.

```
switch:admin > portcfgcreditrecovery 1/20 -enable
```

## Credit Loss Detection

Fabric OS supports the credit loss detection for backend ports and core blades, although the support is slightly different on each device.

### Enabling Backend Credit Loss Detection and Recovery for Link Reset Thresholds

Credit loss detection and recovery is enabled and disabled using the `creditrecovmode` command. This is only supported on the backend ports of 64G and 32G-capable Fibre Channel blades in the chassis of Brocade X7 and X6 Directors respectively.

In Fabric OS, `creditrecovmode --fault` makes this fault option applicable to all backend link failures. Which blade fault is registered due to a backend link failure is decided by the blade fault option. If credit recovery mode is off and there are any backend link failures, Fabric OS will fault the core blade; only one core blade exists, Fabric OS will fault the edge blade. Refer to [Back-end credit loss detection and recovery for link faults](#) for additional information about using `creditrecovmode --fault`.

To enable backend credit loss detection and recovery, perform the following steps.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter `creditrecovmode --cfg` to enable credit recovery of backend ports.

In the following example, backend port credit loss recovery is enabled with the link reset only option.

```
switch:admin> creditrecovmode --cfg onLrOnly
```

3. Enter `creditrecovmode --show` to display information about the backend port credit recovery configuration. In the following example, backend port credit loss recovery is enabled with the link reset only option.

```
switch:admin> creditrecovmode --show
Internal port credit recovery is Enabled with LrOnly
C2 FE Complete Credit Loss Detection is Enabled
```

### Performing a Link Reset for a Loss of Sync from Peer Ports

You can perform a link reset on all the 8G, 16G, and 32G backend ports if there is a loss of sync from peer ports.

The link reset action is decided based on the `--lrthreshold`, `--crdLoss`, and `--losync` settings as per the following matrix.

**Table 84: Link Reset Actions Based on Settings**

| lrthreshold | crdloss | losync | Link Reset Action                                |
|-------------|---------|--------|--------------------------------------------------|
| On          | Off     | On     | Link resets with threshold logic to fault blade. |
| Off         | Off     | On     | Link resets with no threshold logic.             |
| On          | On      | Off    | Link resets with threshold logic to fault blade. |
| Off         | On      | Off    | Link resets with no threshold logic.             |
| On          | On      | On     | Link resets with threshold logic to fault blade. |
| Off         | On      | On     | Link resets with no threshold logic.             |
| Off/On      | Off     | Off    | No link reset.                                   |

For example, to activate both credit loss and loss of sync with link reset and No threshold, use the following command:

```
switch:admin> creditrecovmode --cfg onLrOnly
```

For example, to activate both credit loss and loss of sync with link reset and threshold (10), use the following commands:

```
switch:admin> creditrecovmode --cfg onLrThresh --lrthreshold 2
switch:admin> creditrecovmode -show
Internal port credit recovery is Enabled with LrThresh
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

For example, to enable link reset for credit loss alone, use the following commands.

```
switch:admin> creditrecovmode --be_crdloss on
switch:admin> creditrecovmode --show
Internal port credit recovery is Enabled with LrThresh
Back end port Loss of Sync's Link Reset is Disabled
LR threshold (currently activated): 2
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

For example, to disable link reset for credit loss alone, use the following commands.

```
switch:admin> creditrecovmode --be_crdloss off
switch:admin> creditrecovmode --show
Internal port credit recovery is Disabled
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

For example, to enable link reset for loss of sync alone, use the following commands.

```
switch:admin> creditrecovmode --be_losync on
switch:admin> creditrecovmode --show
Internal port credit recovery is Enabled with LrThresh
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

For example, to disable link reset for loss of sync alone, use the following commands.

```
switch:admin> creditrecovmode --be_losync off
switch:admin> creditrecovmode --show
Internal port credit recovery is Disabled
Back end port Loss of Sync's Link Reset is Disabled
LR threshold (not currently activated): 2
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

## Backend Credit Loss Detection and Recovery for Link Faults

The following table identifies the blade that is faulted based on certain factors.

**Table 85: Blade Fault Scenarios**

| Fault Option  | Backend Link Failure Location | Edge ASIC Fault Count | Core ASIC Fault Count | Faulted Blade When Director Has Single-Active Core | Faulted Blade When Director Has Dual-Active Cores |
|---------------|-------------------------------|-----------------------|-----------------------|----------------------------------------------------|---------------------------------------------------|
| edgeblade     | Edge Blade                    | N/A                   | N/A                   | Edge Blade                                         | Edge Blade                                        |
| edgeblade     | Core Blade                    | N/A                   | N/A                   | Edge Blade                                         | Edge Blade                                        |
| coreblade     | Edge Blade                    | N/A                   | N/A                   | Edge Blade                                         | Core Blade                                        |
| coreblade     | Core Blade                    | N/A                   | N/A                   | Edge Blade                                         | Core Blade                                        |
| edgecoreblade | Edge Blade                    | N/A                   | <2                    | Edge Blade                                         | Edge Blade                                        |
| edgecoreblade | Edge Blade                    | N/A                   | >2                    | Edge Blade                                         | Core Blade                                        |
| edgecoreblade | Core Blade                    | N/A                   | <2                    | Edge Blade                                         | Edge Blade                                        |
| edgecoreblade | Core Blade                    | N/A                   | >2                    | Edge Blade                                         | Core Blade                                        |

### Examples of Credit Recovery Mode Fault Commands

The following examples illustrate the use of the `creditrecovmode --fault` command in various settings.

```
switch:admin> creditrecovmode --fault edgeblade
switch:admin> creditrecovmode -show
Internal port credit recovery is Disabled
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : EDGEBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

```
switch:admin> creditrecovmode --fault edgecoreblade
switch:admin> creditrecovmode --show
Internal port credit recovery is Disabled
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : EDGECOREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

```
switch:admin> creditrecovmode --fault coreblade
switch:admin> creditrecovmode -show
Internal port credit recovery is Disabled
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

## Optimizing Fabric Behavior

---

This section describes optimizing fabric behaviors using Traffic Optimizer (TO), Fabric Notifications, Quality of Service (QoS), and CS\_CTL-based Frame Prioritization features.

When congestion occurs in storage networks, the performance impact can be significant, and the symptoms are difficult to isolate and mitigate. Brocade switches delivers the various range of new features that help to identify and alleviate congestion in Fibre Channel storage area networks (SANs) and to improve the overall performance.

The two new features that are introduced in Fabric OS 9.0.x are Traffic Optimizer, and Fabric Notifications. The QoS feature was introduced in Fabric 8.x or later.

- **Traffic Optimizer**

Traffic Optimizer provides automated traffic flow policies that are based on predefined performance groups. For more information, see [Traffic Optimizer](#).

- **Fabric Notifications**

Fabric Notifications provides congestion alerts that are initiated from the fabric to end devices when congestion conditions are detected. Fabric Notifications also notifies on other impacts such as link integrity and frame loss. For more information, see [Fabric Notifications](#).

- **Quality of Service (QoS)**

QoS allows you to categorize the traffic flow between a host and a target as having a high, medium, or low priority. QoS does not require a license. See [QoS](#) for more information about this feature.

## Traffic Optimizer

Traffic Optimizer is a new feature, introduced in Fabric OS 9.0.x on Brocade Gen 7 platforms that optimizes SAN traffic by segregating traffic flows with different characteristics. Traffic Optimizer automatically groups traffic flows with similar predefined attributes, such as flow destination speed and priority, as a dedicated performance group. Each performance group uses a dedicated resource path within the fabric. By organizing flows by performance group and preventing mixed traffic, Traffic Optimizer prevents slower traffic from obstructing the higher-speed or higher-priority traffic on that path.

From Fabric OS 9.1.0 and above, the Traffic Optimizer feature is enhanced to provide the following advanced features to efficiently optimize SAN traffic:

1. Protocol-based (SCSI/NVMe) traffic flow classification
2. Oversubscribed device isolation (OS)

### Performance Groups

Traffic flows with similar predefined attributes, such as flow destination speed, protocol, and priority, are grouped logically as a performance group (PG). With Gen 7 platforms, Traffic Optimizer provides a way to optimize the shared resource allocation (for example, VCs, buffers) for each performance group, and it automatically maps traffic flows to predefined performance groups.

The predefined performance groups provide support for speed and protocol-based classification, oversubscribed device isolation, credit-stalled device isolation using Slow Drain Device Quarantine (SDDQ), compatibility with the Adaptive Networking (QoS Zone and CSCTL) feature, and compatibility with Gen 6 platforms. Traffic Optimizer is enabled by default on all Gen 7 platforms and Gen 7 blades.

## Profiles

In Fabric OS 9.0.x, a predefined set of PGs were used to provide the optimal performance for most SAN customers. However, with a diverse customer base, it is hard to define a common set of PGs that suits every customer.

From Fabric OS 9.1.x, the following predefined profiles are supported to optimize different customer environments:

- Sys\_TrafOpt\_Version1
- Sys\_TrafOpt\_Version2

### Sys\_TrafOpt\_Version1

This system profile contains PGs that are created based on the device speed. This profile version is optimal for fabrics that need isolation across multiple generation speed devices.

The Sys\_TrafOpt\_Version1 profile supports the following PGs:

- Speed-based traffic flow classification
- Credit-stalled device isolation using Slow Drain Device Quarantine (SDDQ)
- Compatibility with Gen 6 platforms that are connected to the same fabric
- Compatibility with Adaptive Networking (QoS Zone and CSCTL) features

**Table 86: Supported Performance Groups in the Sys\_TrafOpt\_Version1 Profile**

| Flow Classification | Performance Group Name | Speed Range (G)  | Assigned Priority |
|---------------------|------------------------|------------------|-------------------|
| Speed Based         | PG_4G_8G_16G           | 0 < speed <= 16  | Medium            |
|                     | PG_SCSI_32G            | 16 < speed <= 32 | Medium            |
|                     | PG_SCSI_64G            | 32 < speed <= 64 | Medium            |
| SDDQ                | PG_SDDQ_QOS_LOW_1      | 0 < speed <= 64  | Low               |
|                     | PG_SDDQ_QOS_LOW_2      | 0 < speed <= 64  | Low               |
| QoS Zones           | PG_QOS_HIGH_1          | 0 < speed <= 64  | High              |
|                     | PG_QOS_HIGH_2          | 0 < speed <= 64  | High              |
|                     | PG_QOS_HIGH_3          | 0 < speed <= 64  | High              |
|                     | PG_QOS_HIGH_4          | 0 < speed <= 64  | High              |
|                     | PG_QOS_HIGH_5          | 0 < speed <= 64  | High              |

### Sys\_TrafOpt\_Version2

This system profile contains PGs that are created based on the device speed, protocol, and congestion state. This profile version is optimal for fabrics that need isolation across multiple generation speeds and protocols and also protection from congestion events.

The Sys\_TrafOpt\_Version2 profile supports the following PGs:

- Speed-based traffic flow classification
- Protocol-based traffic flow classification
- Credit-stalled device isolation using Slow Drain Device Quarantine (SDDQ)
- Oversubscribed device isolation (OS)
- Compatibility with Gen 6 platforms that are connected to the same fabric
- Compatibility with Adaptive Networking (QoS Zone and CSCTL) features

**Table 87: Supported Performance Groups in the Sys\_TrafOpt\_Version2 Profile**

| Flow Classification    | Performance Group Name      | Protocol  | Speed Range (G)  | Assigned Priority |
|------------------------|-----------------------------|-----------|------------------|-------------------|
| Speed & Protocol Based | PG_4G_8G_16G                | SCSI/NVMe | 0 < speed <= 16  | Medium            |
|                        | PG_SCSI_32G                 | SCSI      | 16 < speed <= 32 | Medium            |
|                        | PG_SCSI_64G                 | SCSI      | 32 < speed <= 64 | Medium            |
|                        | PG_NVMe_32G                 | NVMe      | 32               | Medium            |
|                        | PG_NVMe_64G                 | NVMe      | 32 < speed <= 64 | Medium            |
| Oversubscription       | PG_OVER_SUBSCRIPTION_4G_16G | N/A       | 0 < speed <= 16  | N/A               |
|                        | PG_OVER_SUBSCRIPTION_32G    | N/A       | 16 < speed <= 32 | N/A               |
|                        | PG_OVER_SUBSCRIPTION_64G    | N/A       | 32 < speed <= 64 | N/A               |
| SDDQ                   | PG_SDDQ_QOS_LOW_1           | N/A       | 0 < speed <= 64  | Low               |
|                        | PG_SDDQ_QOS_LOW_2           | N/A       | 0 < speed <= 64  | Low               |
| QoS Zones              | PG_QOS_HIGH_1               | N/A       | 0 < speed <= 64  | High              |
|                        | PG_QOS_HIGH_2               | N/A       | 0 < speed <= 64  | High              |
|                        | PG_QOS_HIGH_3               | N/A       | 0 < speed <= 64  | High              |
|                        | PG_QOS_HIGH_4               | N/A       | 0 < speed <= 64  | High              |
|                        | PG_QOS_HIGH_5               | N/A       | 0 < speed <= 64  | High              |

## Profile Compatibility

When two switches with different TO profiles form the fabric, the PG groups are appropriately remapped based on the PG attributes. When two switches with the same TO profile form the fabric, the PG groups are merged, so the superset of PG groups is available in the fabric.

As part of the fabric formation, the profile information from either end is exchanged to the other side and the profile merge operation is triggered. If the profiles are merged, once the fabric is formed, the new PGs are created on both the ends of the fabric.

**Table 88: Profile Compatibility between Different Switches or Firmware**

| Switch 1 (Fabric OS Version/Profile Version) | Switch 2 (Fabric OS Version/Profile Version) | From Switch 1                                                                                              | From Switch 2                           |
|----------------------------------------------|----------------------------------------------|------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| FOS 9.1.x / Sys_TrafOpt_Version2             | FOS 9.0.x                                    | FOS 9.0.x PGs are supported. New FOS 9.1.x PGs are remapped to FOS 9.0.x.                                  | FOS 9.0.x PGs are supported.            |
| FOS 9.1.x / Sys_TrafOpt_Version2             | FOS 9.1.x / Sys_TrafOpt_Version1             | Sys_TrafOpt_Version1 PGs are supported. New Sys_TrafOpt_Version2 PGs are remapped to Sys_TrafOpt_Version1. | Sys_TrafOpt_Version1 PGs are supported. |
| FOS 9.1.x / Sys_TrafOpt_Version1             | FOS 9.1.x / Sys_TrafOpt_Version1             | Sys_TrafOpt_Version1 PGs are supported.                                                                    | Sys_TrafOpt_Version1 PGs are supported. |

| Switch 1 (Fabric OS Version/Profile Version) | Switch 2 (Fabric OS Version/Profile Version) | From Switch 1                                                                         | From Switch 2                                        |
|----------------------------------------------|----------------------------------------------|---------------------------------------------------------------------------------------|------------------------------------------------------|
| FOS 9.1.x / Sys_TrafOpt_Version2             | FOS 9.1.x / Sys_TrafOpt_Version2             | Sys_TrafOpt_Version2 PGs are supported.                                               | Sys_TrafOpt_Version2 PGs are supported.              |
| FOS 9.1.x / Sys_TrafOpt_Version2             | Pre-FOS 9.0.x / Pre Gen 7 Link               | 1. VC remapping to PG_4G_8G_16G.<br>2. Same priority (low and high QoS is supported). | 1. Legacy flow support.<br>2. QoS VCs are supported. |

## Configuring Profiles

The new `trafopt` command is introduced to configure which policy to activate, to manage the activation schedule, or to abort the scheduled activation.

### Activating a Profile

You can use the `trafopt --activate <profile_name>` command to activate any one of the predefined profiles.

For example:

```
switch:admin> trafopt --activate Sys_TrafOpt_Version1
```

### Displaying an Active Profile

You can use the `trafopt --show` command to display a summary of the active profile.

For example:

```
switch:admin> trafopt --show
```

All Performance Group statistics of active profile

Active Profile: Sys\_TrafOpt\_Version2

| Performance  |      | Flow   |              | IO Error |         | MAX IO      |           | Throughput |        | Aggregate |      |
|--------------|------|--------|--------------|----------|---------|-------------|-----------|------------|--------|-----------|------|
| Groups       |      | Active | High Latency | Count    | Latency | (bytes/sec) | (IOs/sec) |            |        |           |      |
|              |      | count  | violation    |          |         |             |           |            |        |           |      |
|              |      | 5min   | 24hr         | 5min     | 24hr    | 5min        | 24hr      | 5min       | 24hr   | 5min      | 24hr |
| PG_4G_8G_16G | 1593 | 0      | 0            | 0        | 0       | 1.061s      | 1.855s    | 22.06G     | 2.899G | 44.20K    |      |
| 5.807K       |      |        |              |          |         |             |           |            |        |           |      |
| PG_SCSI_32G  | 1449 | 0      | 0            | 0        | 0       | 76.46m      | 246.5m    | 24.61G     | 3.063G | 52.09K    |      |
| 8.154K       |      |        |              |          |         |             |           |            |        |           |      |
| PG_SCSI_64G  | 0    | 0      | 0            | 0        | 0       | 0           | 0         | 0          | 0      | 0         | 0    |



|                         |     |   |   |   |   |        |        |        |        |        |     |
|-------------------------|-----|---|---|---|---|--------|--------|--------|--------|--------|-----|
| PG_NVME_32G             | 1   | 0 | 0 | 0 | 0 | 5.657m | 3.208m | 3.263G | 1.922G | 6.527K |     |
| 3.845K                  |     |   |   |   |   |        |        |        |        |        |     |
| PG_NVME_64G             | 0   | 0 | 0 | 0 | 0 | 0      | 0      | 0      | 0      | 0      | 0   |
|                         |     |   |   |   |   |        |        |        |        |        |     |
| PG_OVERSUBSCRIPTION_32G | 0   | 0 | 0 | 0 | 0 | 0      | 0      | 0      | 0      | 0      | 0   |
|                         |     |   |   |   |   |        |        |        |        |        |     |
| PG_SDDQ_QOS_LOW_1       | 225 | 0 | 0 | 0 | 0 | 86.39m | 11.99m | 2.205G | 322.3M | 4.410K | 644 |
|                         |     |   |   |   |   |        |        |        |        |        |     |
| PG_SDDQ_QOS_LOW_2       | 207 | 0 | 0 | 0 | 0 | 77.88m | 12.88m | 3.740G | 1.515G | 7.481K |     |
| 3.030K                  |     |   |   |   |   |        |        |        |        |        |     |
| PG_QOS_HIGH_1           | 49  | 0 | 0 | 0 | 0 | 0      | 559.2m | 0      | 3      | 0      | 0   |
|                         |     |   |   |   |   |        |        |        |        |        |     |
| PG_QOS_HIGH_2           | 48  | 0 | 0 | 0 | 0 | 0      | 2u     | 0      | 14     | 0      | 0   |
|                         |     |   |   |   |   |        |        |        |        |        |     |
| PG_QOS_HIGH_3           | 45  | 0 | 0 | 0 | 0 | 0      | 2u     | 0      | 1      | 0      | 0   |
|                         |     |   |   |   |   |        |        |        |        |        |     |
| PG_QOS_HIGH_4           | 45  | 0 | 0 | 0 | 0 | 0      | 1u     | 0      | 3      | 0      | 0   |
|                         |     |   |   |   |   |        |        |        |        |        |     |
| PG_QOS_HIGH_5           | 44  | 0 | 0 | 0 | 0 | 0      | 559.2m | 0      | 22     | 0      | 0   |
|                         |     |   |   |   |   |        |        |        |        |        |     |

-----

Available Traffic Optimizer profiles to activate

Number of Profiles: 2

| Name                  | Description                                                                                                           | Use-cases                                                                                                                   |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Sys_TrafOpt_Version1  | This System profile contains Performance Groups that are created based on device speed                                | Optimal for fabrics which needs isolation across multiple generation speed devices                                          |
| Sys_TrafOpt_Version2* | This System profile contains Performance Groups that are created based on device speed, protocol and congestion state | Optimal for fabrics which need isolation across multiple generations, protocols and also protection from congestion events. |

-----

Active profile is marked with \*.

-----

No scheduled Profile to Activate

-----

### **Scheduling an Upgrade to Another Profile**

You can schedule an upgrade from the Sys\_TrafOpt\_Version1 profile to the Sys\_TrafOpt\_Version2 profile using the `trafopt --schedule <profile_name>` command.

For example,

```
switch:admin> trafopt --schedule Sys_TrafOpt_Version2 -delay 10
Profile migration scheduled successfully.
```

```
switch:admin> trafopt --show
```

All Performance Group statistics of active profile

Active Profile: Sys\_Trafopt\_Version1

| Performance |  | Flow   |  |              |  | Aggregate |  |             |  |       |  |
|-------------|--|--------|--|--------------|--|-----------|--|-------------|--|-------|--|
| Aggregate   |  |        |  | IO Error     |  | MAX IO    |  | Throughput  |  | IOPS  |  |
|             |  |        |  | Count        |  | Latency   |  | (bytes/sec) |  | (IOs/ |  |
| Groups      |  | Active |  | High Latency |  |           |  |             |  |       |  |
| sec)        |  | count  |  | violation    |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |
|             |  |        |  |              |  |           |  |             |  |       |  |

```
|
| ed based on device speed, proto | lation speeds, protocols and als |
|
| col and congestion state | o protection from congestion ev |
|
|
| ents. |
```

-----

Active profile is marked with \*.

-----

Profile scheduled to Activate: Sys\_TrafOpt\_Version2  
 Activate Delay: 10 days 0 hours 0 minutes

-----

## NOTE

If you are upgrading from Fabric OS 9.0.x to Fabric OS 9.1.0, the Sys\_TrafOpt\_Version1 profile remains active for 30 days. At the end of the 30th day, it automatically upgrades to the Sys\_TrafOpt\_Version2 profile. After the Sys\_TrafOpt\_Version2 profile becomes active, a firmware downgrade from Fabric OS 9.1.0 to Fabric OS 9.0.x is disruptive. Downgrading while the Sys\_TrafOpt\_Version1 is active is non-disruptive.

## Displaying the PG Groups in a Profile

You can use the `trafopt --show -profile <profile_name>` command to display the performance groups associated with a profile.

For example,

```
switch:admin> trafopt --show -profile Sys_TrafOpt_Version2
Performance Groups associated with profile
```

```
Profile Name: Sys_TrafOpt_Version2
```

```
=====
| Performance Groups Category | Description |
=====
| PG_4G_8G_16G | Flows destined to 4G, 8G and 16G devices |

| PG_SCSI_32G | Flows destined to SCSI 32G devices |

| PG_SCSI_64G | Flows destined to SCSI 64G devices |

| PG_NVME_32G | Flows destined to NVME 32G devices |

| PG_NVME_64G | Flows destined to NVME 64G devices |

PG_OVER_SUBSCRIPTION_4G_16G	Flows destined to oversubscribed devices
PG_OVER_SUBSCRIPTION_32G	
PG_OVER_SUBSCRIPTION_64G	

| PG_SDDQ_QOS_LOW_1 | User Configured QoS-Low and Credit Stall Flows |
| PG_SDDQ_QOS_LOW_2 | |

PG_QOS_HIGH_1	User configured QoS-High flows
PG_QOS_HIGH_2	
PG_QOS_HIGH_3	
PG_QOS_HIGH_4	
```

```
PG_QOS_HIGH_5
```

## Displaying Active Flows Associated with the Performance Group

You can use the `trafopt --show [-flow <-pg <performance_group_name>>]` command to display the active flows that are associated with the specified performance group.

For example,

```
switch:admin> trafopt --show -flow -pg PG_SCSI_64G
```

```
=====
Name : sys_flow_monitor
Performance Group : PG_SCSI_32G
Active Flow : 5

```

```
Timebase : 5 Minutes

```

```
I/O Brief Metrics
```

|            |             | RD          |             |      |      |       | Max Time   | Max Time   | Avg IOs/sec | Avg |
|------------|-------------|-------------|-------------|------|------|-------|------------|------------|-------------|-----|
| Bytes/sec  | Max Count   | Total Count | Total Count |      |      |       |            |            |             |     |
| SID        | DID         | WR          | Device      | Flow | Port | Proto | ECT        | FRT        | IOPS        |     |
| BPS        | Pending IOs | Exceptions  | Violations  |      |      |       |            |            |             |     |
|            | Type        | Hierarchy   | Index       |      |      |       | 5min / All | 5min / All | 5min / All  |     |
| 5min / All | 5min / All  | 5min / All  | 5min / All  |      |      |       |            |            |             |     |

|                    |       |       |         |               |               |           |
|--------------------|-------|-------|---------|---------------|---------------|-----------|
| 6d5c00  6e0700  RD | I     | IT->L | 0  SCSI | 4.512m/30.19m | 1.880m/20.74m | 319 / 317 |
| 1.157T/432.1G      | 3 / 5 | /     | /       |               |               |           |
| 6d5c00  6e0700  WR | I     | IT->L | 0  SCSI | 2.574m/30.65m | 918u/20.98m   | 349 / 347 |
| 1.157T/495.5G      | 4 / 5 | /     | /       |               |               |           |
| 6d5c00  6e0700  RD | I     | IT->L | 0  SCSI | 4.512m/30.19m | 1.880m/20.74m | 319 / 317 |
| 1.157T/432.1G      | 3 / 5 | /     | /       |               |               |           |
| 6d5c00  6e0700  WR | I     | IT->L | 0  SCSI | 2.574m/30.65m | 918u/20.98m   | 349 / 347 |
| 1.157T/495.5G      | 4 / 5 | /     | /       |               |               |           |
| 6d5c00  6e0700  RD | I     | IT->L | 0  SCSI | 4.512m/30.19m | 1.880m/20.74m | 319 / 317 |
| 1.157T/432.1G      | 3 / 5 | /     | /       |               |               |           |
| 6d5c00  6e0700  WR | I     | IT->L | 0  SCSI | 2.574m/30.65m | 918u/20.98m   | 349 / 347 |
| 1.157T/495.5G      | 4 / 5 | /     | /       |               |               |           |
| 6d5c00  6e0700  RD | I     | IT->L | 0  SCSI | 4.512m/30.19m | 1.880m/20.74m | 319 / 317 |
| 1.157T/432.1G      | 3 / 5 | /     | /       |               |               |           |
| 6d5c00  6e0700  WR | I     | IT->L | 0  SCSI | 2.574m/30.65m | 918u/20.98m   | 349 / 347 |
| 1.157T/495.5G      | 4 / 5 | /     | /       |               |               |           |

## Displaying IT Flows

You can use the `trafopt --show [-flow <-srcdev">"]` command to display the IT flows that are associated with the specified source FCID or specified destination FCID.

For example,

```
switch:admin> trafopt --show -flow -srcdev "*"
All the flows and their associated Performance Group

Performance SrcDev DstDev VC num
Group

PG_4G_8G_16G 0xef4d00 0x460000 2
PG_4G_8G_16G 0xef4200 0x460100 3
PG_4G_8G_16G 0xef2900 0x93df00 5
PG_4G_8G_16G 0xef4200 0x93be80 4
PG_4G_8G_16G 0xef4d00 0xab4a40 4

```

## Speed and Protocol-Based Classification

From Fabric 9.1.0, the `Sys_Trafopt_Version2` profile is introduced to classify flows that are based on the destination speed and protocol (for 32G and 64G speeds) and place the flows into the respective performance group. This profile isolates different flow speeds and protocols into different VCs allocated for each performance group to eliminate performance interference of lower-speed devices to higher-speed devices and SCSI to NVMe when they connect to the same fabric.

As long as the source device and target device both register as NVMe capable in the fabric, all the flows between these devices are categorized as NVMe. In FOS 9.1.0, two PGs are available for NVMe isolation (32G and 64G). If a device registers as NVMe capable but the supported speed is less than 32G, flows destined to such device are segregated into default speed PG.

NVMe classification is enforced only for flows that starts or ends with devices that perform the login process after upgrading to Fabric OS 9.1.0 firmware and updating with the FOS 9.1.0 default profile. For example, if the flow has already started while the switches are running FOS 9.0.x or below, to upgrade to the FOS 9.1.0 default profile (`Sys_Trafopt_Version2`), the flows continue to be in the speed-based categorization.

### Speed-Based Classification

Flows are classified into one of three performance groups based on the destination speed. Automatic speed classification isolates traffic flows into different speed-based performance groups and eliminates speed mismatches that cause performance interference when lower-speed devices collide with higher-speed devices that are connected in the same fabric.

#### NOTE

Speed-based classification is applicable only to the `Sys_Trafopt_Version1` profile.

### Oversubscription Flow Isolation

Similar to SDDQ classification, TO uses MAPS to identify devices that are oversubscribed. TO collects complete data for performance, health, and utilization of PGs from Flow Vision. When MAPS reports an OS event, TO initiates a fabric-wide action that creates the OS PG and moves flows to the OS PG. TO can create up to three speed-based OS PGs (0-16G OS-PG, 32G OS-PG, and 64G OS-PG) based on resources.

## **SDDQ Isolation**

A *slow-drain device*, sometimes also referred to as a *credit-stalled device*, is a misbehaving device that promptly stops returning R\_RDY signals (buffer credits) to the switch, causing the switch to stop sending frames to the device. A credit-stalled device that stops returning credits to the switch for hundreds of milliseconds or more causes frame drops and link resets. This typically happens when the end device is not working as per its specification. An impaired flow that is associated with credit-stalled devices is automatically mapped to a low-priority performance group: PG\_SDDQ\_QOS\_LOW\_1 or PG\_SDDQ\_QOS\_LOW\_2.

## **Compatibility with QoS Zones and CS\_CTL**

Gen 7 platforms maintain compatibility if the QoS feature is enabled. If QoS Low or QoS High priority is defined, traffic flows are assigned to one of the QoS performance groups. If CS\_CTL is enabled on a port, traffic flows on that port retain the legacy CS\_CTL-based classification and is not classified into any performance group.

## **Compatibility with Gen 6 Platforms**

Gen 6 platforms in Fabric OS 9.0.x or later continue the legacy behavior and maintain compatibility with Gen 7 platforms in the same fabric. If a Gen 7 switch or blade is connected to a port on a Gen 6 platform and the traffic flow source port is on the Gen 7 switch, the traffic flow is speed-based classified within the Gen 7 switch until reaching the port on the Gen 6 platform, and then the traffic flow is handled with legacy behavior starting from the Gen 6 port.

## **Traffic Optimizer and QoS Mode**

Both Traffic Optimizer and QoS mode are automatically enabled by default on all Gen 7 platforms. Traffic Optimizer's advanced classification (for example, mapping a traffic flow to a predefined performance group) requires QoS mode to be enabled on all ports of the supported platforms. It is not recommended to disable QoS mode on any ports. If QoS mode is disabled on a port, Traffic Optimizer's advanced classification will not work on the port, and the port's traffic classification will fall back to legacy PID to VC mapping.

## **FCR Considerations for Traffic Optimizer**

Traffic Optimizer is supported only on an FCR edge-to-edge fabric topology; it is not supported on a backbone-to-edge fabric topology. For Traffic Optimizer to run over FCR, all switches and FC routers in both edge fabrics and backbone fabrics must be running Fabric OS 9.0.x or later.

Starting from Fabric OS 9.0.x, Traffic Optimizer over FC routers is supported for both a Virtual Fabric enabled configuration and a Virtual Fabric disabled configuration in the backbone fabric. For more information on the operating modes of backbone fabrics, see [Using FC-FC Routing to Connect Fabrics](#). If the traffic flow in the metaSAN encounters a switch platform that does not support TO, then the traffic flow is mapped and retains the legacy behavior.

# **Fabric Notifications**

Fabric Notifications is a facility for proactively notifying devices in a fabric of performance-impacting behaviors that affect the normal flow of I/O. Fabric Performance Impact (FPI) behaviors such as signal integrity degradation, frame loss, or congestion are detected by either the fabric or the participating devices. Signaling, through Fibre Channel primitives, and notifications through Extended Link Services (ELS) are then used to communicate the type and severity of the FPI to the affected members of the fabric.

Fabric Notifications include the following components:

- **Registered Capabilities**  
A set of ELS commands exchanged between two ports arbitrates the range of Fabric Notifications capabilities that are supported by each port.
- **Congestion Detection Signaling**

Hardware-generated primitive signals, supported by Gen 7 platforms only, are sent toward the point of congestion and indicate that the queue of frames to be transmitted has reached a level where the fabric performance may be impacted.


- **FPI Notifications (FPINs)**

FPINs are software that is generated using ELS commands. If a device has registered to utilize them, administrators can be notified of events that might affect fabric performance.

These components allow members of the fabric to indicate their ability to interpret and react to signals and notifications. Participating devices can respond by actions to mitigate and reduce the impacts of congestion, frame loss, and signal integrity degradation.

Fabric OS is the primary generator of signals and notifications due to its awareness of I/O flows, physical layer impacts, and device responses. The relationship between event detection, signaling, and notification operations is depicted in the following figure.

**Figure 19: Signal or Notification Operations**



When an event is detected (1), the fabric initiates the appropriate response. In the case where frame transmission is stalled, signal events are generated (2) toward the point of congestion (that is a credit-stalled or oversubscribed device). In the case of a transmission failure or sustained congestion event, notification events are generated (3) and distributed to peer devices (4). The fabric determines the distribution of the notification event from the zone member information of the triggering N\_Port. If there is a transmission failure or software-based congestion signaling, notification events are sent to the attached N\_Port (5) to engage the device in mitigating activities.

### **FPIN Description**

FPIN ELS descriptors are used to categorize the type of impact that has been detected. A device may register for each descriptor individually. The following descriptors are supported with Fabric OS 9.0.x:

- **Congestion Descriptor**

Sent to a device that is on a congested port detected by Fabric Performance Impact (FPI) monitoring. The cause of the congestion is conveyed via an event type (for example, credit-stall or oversubscription).

- **Peer Congestion Descriptor**

Sent to devices that are in the same zone with a device that is on a congested port.

- **Link Integrity Descriptor**

Indicates that CRC or Invalid-Transmission-Word (ITW) errors have been detected at the affected N\_Port. The notification is sent to devices where the error has been detected and is also sent to peer devices in the same zone.

- **Delivery Descriptor**

Sent to a device when a SCSI-FCP command frame that originated from the device has been dropped within the fabric.

The fabric keeps track of which FPIN descriptors have been sent or received (to and from a device, respectively). These counters can be viewed using the `fabricnotifications` command. Associated counters are cleared when the device logs out of the fabric.

### **Registering for Signals and Notifications**

The signaling and notification capabilities of a device must be registered with the fabric. The registered capabilities of a device can be viewed using the `fabricnotifications` command. For more information, refer to the *Brocade Fabric OS Command Reference Manual*.

#### **NOTE**

It may be necessary to enable the Fabric Notifications at the device level. Consult the device (HBA, storage) manufacturer for support and instructions to enable these features.

### **Performance Impact Detection Mechanism**

In many cases, performance-impacting behavior is detected by the fabric (F\_Port) through hardware or software. Hardware-level detection and signaling may be enabled or disabled using the `portcfgcongestionssignal` command, which is enabled by default. For more information, refer to *Brocade Fabric OS Command Reference Manual*. The Monitoring and Alert Policy Suite (MAPS) and Frame Viewer application may be configured to enable the FPIN action, which enables Fabric Notifications in response to events. For more information, refer to the *Brocade Fabric OS MAPS User Guide* and [Using Frame Viewer for Fabric Notifications](#).

In addition, devices (N\_Ports) may also detect impacts and notify the fabric so that it will, in turn, forward notifications to affected members of the fabric.

- **Response to Signals from a Device**

The fabric tracks signals from a device (N\_Port) and can be viewed using the `portShow` command (see the value for the `Cong_Prim_in` field).

- **Response to Notifications from a Device**

Notifications transmitted from an N\_Port to an F\_Port indicate a device-detected event and are propagated to the peer N\_Ports of the detecting N\_Port by the fabric controller. The notifications include the event information to be relayed to the peer N\_Ports. N\_Port peers are those devices that are zoned-in with the detecting N\_Port.

- **Response to Signals from the Fabric**

Signals transmitted from an F\_Port to an N\_Port indicate that the accumulation of frames pending transmission to the device is impacting fabric resources. The signal indicates to the N\_Port that a backlog of frames exists and the pace of frame consumption is not keeping up with the arrival rate. The device may report the condition to the end user and can also take actions to mitigate the frame backlog (for example, reducing the rate of I/O).

- **Response to Notifications from the Fabric**

Notifications are transmitted from an F\_Port to the peer N\_Ports when an event is detected at the attached N\_Port. The notifications include the event type and the associated data (that is, affected port names, event type, frame header, and so on). The notification indicates to the receiving N\_Port that a condition affecting the flow of I/O to the affected N\_Port exists.

A device may record the fact that a notification has been received and may also attempt to mitigate the condition. For example, link integrity events are used to accelerate I/O failure processing by marking the path as defective and employing alternative I/O paths. Similarly, congestion events may trigger a device to take action to reduce the overall rate of I/O.

In any case, the objective is to reduce the response time associated with I/O failures, eliminate application retry delays, shift I/O to unimpaired paths, or alleviate frame backlogs (congestion). These activities vastly improve the overall resiliency and performance of the fabric and create a more stable environment for applications.



## Fabric Notifications Support on Access Gateway

From Fabric OS 9.1.0, Fabric notifications has been extended to Access Gateway (AG) switches. AG switches gain the ability to drive FPINs for FPIs detected locally on the AG through native mode switches running Fabric OS 9.1.0 and above. The AG mode switches support hardware-level congestion detection signaling for devices that are attached directly to the AG, irrespective of the native mode switch. For more information, refer to the *Brocade Fabric OS Access Gateway User Guide*.

## FCR Considerations

Fabric notifications are not exchanged across routed fabrics.

## QoS

Quality of Service (QoS) allows you to categorize the traffic flow between a host and a target as having a high, or low priority.

Fabric OS supports two types of prioritization:

- **Class-Specific Control (CS\_CTL)-based frame prioritization**  
Each frame between a host and a target is assigned a specific priority, depending on the value of the CS\_CTL field in the frame header.
- **QoS zone-based traffic prioritization**  
All traffic between a host and a target is assigned a specific priority, depending on the name you define for the QoS zone.

CS\_CTL-based frame prioritization and QoS zone-based traffic prioritization are mutually exclusive. If you enable CS\_CTL-based frame prioritization on F\_Ports, QoS zone-based traffic prioritization cannot be used between any devices connected to the F\_Ports.

CS\_CTL-based frame prioritization takes precedence over QoS zone-based traffic prioritization. If you enable CS\_CTL-based frame prioritization on F\_Ports that are defined in a QoS zone, CS\_CTL-based frame prioritization takes precedence over the QoS zones.

The following table compares CS-CTL-based frame prioritization and QoS zone-based traffic prioritization. See [CS\\_CTL-Based Frame Prioritization](#) and QoS zone-based traffic prioritization for detailed information about each type of prioritization scheme.

**Table 89: Comparison between CS\_CTL-Based and QoS Zone-Based Prioritization**

| CS_CTL-Based Frame Prioritization                                                                                                | QoS Zone-Based Traffic Prioritization                                                                                                                                                            |
|----------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Must be manually enabled.                                                                                                        | Automatically enabled.                                                                                                                                                                           |
| No zones are required.                                                                                                           | Requires you to create QoS zones.                                                                                                                                                                |
| Enabled on F_Ports.                                                                                                              | Enabled on E_Ports.                                                                                                                                                                              |
| Takes precedence over QoS zone-based traffic prioritization.                                                                     | Is overridden by CS_CTL-based frame prioritization.                                                                                                                                              |
| Priority is defined by the CS-CTL field in the frame header.                                                                     | Priority is defined by the name of the QoS zone.                                                                                                                                                 |
| Prioritization is on a frame basis.                                                                                              | Prioritization is on a flow basis.                                                                                                                                                               |
| Setup steps:<br>1. Enable CS_CTL mode on F_Ports.<br>2. Ensure that the CS_CTL-mode-enabled host and storage are zoned together. | Setup steps:<br>1. Create QoS zones with host/target members.<br>2. Add the QoS zones to the zone configuration.<br>3. Save and then enable the zone configuration.<br>4. Enable QoS on E_Ports. |


## QoS on E\_Ports

In addition to configuring the hosts and targets in a zone, you must also enable QoS on individual E\_Ports that might carry traffic between the host and target pairs. Path selection between the "host,target" pairs is governed by Fabric Shortest Path First (FSPF) rules and is not affected by QoS priorities. For example, in the following figure, QoS should be enabled on the encircled E\_Ports.

### NOTE

By default, QoS is enabled on 8G or higher ports. QoS is disabled by default on all long-distance ports.

**Figure 20: QoS with E\_Ports Enabled**



You must enable QoS on the E\_Ports on both ISLs between domain 3 and domain 4, as either path might be selected to carry traffic.

You do *not* need to enable QoS on the E\_Ports on the ISLs between domain 1 and domain 2 and between domain 2 and domain 3, because these are not the shortest paths between the hosts and the targets. However, if the ISL between domain 1 and domain 3 is broken, the path through domain 2 would be used.

To guarantee traffic priority, enable QoS on all possible E\_Ports.

If QoS is not enabled on an E\_Port, the traffic prioritization stops at that point. For example, in the figure, if you disabled QoS on E\_Ports "3,12" and "3,13", the traffic from H1 and H2 to S3 would be low priority from the hosts to domain 3, but would change to the default (medium) priority from domain 3 to the target S3.

## QoS over FC Routers

QoS over FC routers uses QoS traffic prioritization between devices in edge fabrics over an FC router. See [Using FC-FC Routing to Connect Fabrics](#) for information about FC routers, phantom switches, and the FC-FC routing service.

To establish QoS over FC routers, perform the following tasks:

1. Define QoS zones in each edge fabric.
2. Define logical SAN (LSAN) zones in each edge fabric.
3. Enable QoS on the E\_Ports in each edge fabric.
4. Enable QoS on the EX\_Ports in the backbone fabric.

See [Setting QoS Zone-Based Traffic Prioritization over FC Routers](#) for detailed instructions.

The following are requirements for establishing QoS over FC routers:

- QoS over FC routers is supported in Brocade native mode only. It is not supported in interopmode 2 or interopmode 3.
- QoS over FC routers is supported for the following configurations:
  - Edge-to-edge fabric configuration: Supported on all platforms.
  - Backbone-to-edge fabric configuration: Supported on both Gen 7 and Gen 6 platforms. For all other platforms, you cannot prioritize the flow between a device in an edge fabric and a device in the backbone fabric.
- QoS over FC routers is supported only if Virtual Fabrics is disabled in the backbone fabric. QoS over FC routers cannot be enabled if Virtual Fabrics is also enabled in the backbone fabric.
- The port WWN of the host or target and the port WWN of the proxy device must be in both an LSan zone and a QoS zone.
- QoS over FC routers is supported on EX\_Ports.
- QoS zones must use WWN notation only; D, I notation is not supported for QoS over FC Routers.

## vTap and QoS High Priority Zoning Compatibility Mode

The vTap analytics feature uses VC 14 for remote flow mirroring, and it requires that this VC be reserved exclusively for mirrored traffic in order to ensure accurate results. When vTap is enabled, QoS high priority zoning ("QoSH/QoSH5") must operate in a manner that avoids assigning flows to VC 14. You can control this behavior by a configuration setting at each relevant switch in the fabric. If you require vTap and QoS high priority zoning, you must enable vTap/QoSH compatibility on the switches having QoS high-priority devices, the switches, or both where vTap must be configured. You can use the `configureChassis` command to enable vTap and QoS High Priority Zone Compatibility mode.

### NOTE

vTap and QoS compatibility mode are not supported in Gen 7 platforms.

```
sw:admin> configurechassis

Configure...

cfgload attributes (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no]
system attributes (yes, y, no, n): [no]
fos attributes (yes, y, no, n): [no] y

Reboot needed to effect new CSCTL Mode
CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [0]
Chassis SDDQ Limit: (0..32) [10]
vTap and QOS High Priority Zone Compatibility Mode (on, off): [off]
```

If vTap/QoSH compatibility is enabled and the vTap features (that is, remote flow mirroring) are also active on the local switch, vTap/QoSH compatibility cannot be disabled. vTap features must be deactivated on all logical switches in the chassis prior to disabling vTap/QoSH compatibility.

By default, vTap and QoS High Priority Zone Compatibility mode are disabled. When they are disabled, the switch assigns QoS high-priority zoning VCs in the usual manner.

### NOTE

The chassis must be disabled to enable the new compatibility mode, but only when the effective zoning configuration has QOSH (default VC assignment) and/or QOSH5 zones, or both.

You can issue the `configureChassis` command from any logical switch partition/context, but it affects a chassis-wide scope regardless of switch context.

## QoS/QoS5 Zones

QoS/QoS5 zones created from non-vTap/QoS-compatible switches can be pushed to CM-mode-enabled switches as part of a zone update. QoS/QoS5 zones added to the effective zoning configuration using the `configDownload` command are permitted when vTap/QoS compatibility is enabled. However, creation of new QoS5 zones using the `zoneCreate` command is not allowed when vTap/QoS compatibility is enabled. This enforcement is applied only on and for the local switch. For example, if a QoS5 zone create, rename, or update command is run on a switch that is *not* configured for vTap/QoS compatibility, the zone update is *accepted* on those switches where vTap/QoS compatibility is enabled.

### NOTE

The presence or absence of QoS5 zones does not have any effect on the fabric build, regardless of the vTap and QoS compatibility settings.

The `zoneCreate` command does not allow creation of QoS5 zones when vTap and QoS compatibility is enabled; the following message is displayed.

```
error: Operation is not allowed for QOSH5 zone "< zone name>" when vTap and QOS High Priority Zoning
Compatibility Mode is enabled
```

The `zoneRename` command does not allow the renaming of a given zone such that it becomes a QoS5 zone when vTap/QoS compatibility is enabled; the following message is displayed.

```
error: Operation is not allowed for QOSH5 zone "<zone name>" when vTap and QOS High Priority Zoning
Compatibility Mode is enabled.
```

Activation of vTap remote flow mirroring is prevented unless vTap and QoS compatibility is enabled.

```
sw:admin>flow --activate sys_analytics_vtap
Enable vTap and QoS High Priority Zone Compatibility mode to activate vTap flow. Please use the
configureChassis command to enable this compatibility mode.
```


To avoid having to disable the chassis, perform the following steps:

1. Temporarily enable a zoning configuration that does not have any QoS high-priority zones.
2. Enable vTap/QoS compatibility using the `configureChassis` command on all switches in the fabric.
3. Re-enable a zoning configuration that has the desired QoS high-priority zones, with the recommendation to avoid QoS5 zones with the new configuration.

## Virtual Fabrics Considerations for QoS Zone-Based Traffic Prioritization

You can prioritize flows between devices in a logical fabric. The priority is retained for traffic going across ISLs and through the base fabric XISLs.

For example, the following figure shows a logical fabric that includes H1 and S1. To set the traffic between H1 and S1 to high priority, create a QoS zone in the logical fabric with H1 and S1 as members. Then enable QoS on all E\_Ports shown circled in the figure, including all E\_Ports in the XISLs (ports 10, 11, 12, 13, 14, 15, 16, and 17).

**Figure 21: Traffic Prioritization in a Logical Fabric**

## Traffic Prioritization Based on QoS Zones

Quality of service (QoS) zones are user-defined zones that allow you to manage the traffic priority between specified host–target pairs. You assign these pairs a high or low QoS-level priority by configuring a QoS zone for that level and then identifying those pairs as members of the appropriate zone. A host–target pair can belong to only one QoS zone. By default, traffic in non-QoS zones is assigned medium priority.

QoS zones can contain WWN members (WWNN or WWPNN) or domain, index (D,I) members. If you use D,I notation in your QoS zones, see [Limitations and Restrictions for QoS Zone-Based Traffic Prioritization](#) for some considerations.

A QoS zone has a special name format to differentiate it from a regular zone, and this name format determines the priority of the traffic flow. The format of the QoS zone name is shown in the following table; *id* is a flow identifier that designates a specific VC for the traffic flow, and *xxxxxx* is the user-defined portion of the name. The switch automatically sets the priority for the “host,target” pairs specified in the zones according to the priority level (H, M, or L) in the zone name. The *id* is optional; if it is not specified, the VCs are allocated by means of a round-robin scheme.

**Table 90: Name Formats for Priorities of QoS Zones**

| Priority | Name Format   |
|----------|---------------|
| High     | QOSHid_xxxxxx |
| Low      | QOSLid_xxxxxx |

As examples, “QOSH3\_HighPriorityTraffic” and “QOSL1\_LowPriorityZone” are both valid QoS zone names.

Each priority level is allocated to different VCs. High-priority flows receive more fabric resources than medium-priority flows, which receive more resources than low-priority flows. For example, you could assign online transaction processing

(OLTP) to a high-priority zone and backup traffic to a low-priority zone. The flow *id* allows you to have control over the VC assignment and balancing the flows throughout the fabric.


The *id* range is as follows:

- 1 to 5 for high-priority traffic, which corresponds to VCs 10 to 14.
- 1 to 2 for low-priority traffic, which corresponds to VCs 8 and 9.

The following figure shows a fabric with two hosts (H1 and H2) and three targets (S1, S2, and S3). The traffic prioritization is as follows:

- Traffic between H1 and S1 has high priority.
- Traffic between H1 and S3 and between H2 and S3 has low priority.
- All other traffic has medium priority, which is the default.

**Figure 22: QoS traffic prioritization**



For this fabric, you could set up the following QoS zones:

- QOSH\_Zone1 Members: H1 and S1
- QOSL\_Zone3 Members: H1, H2, and S3

### Notes on QoS Zoning

Keep the following items in mind when working with QoS zoning:

- For new switches, QoS mode is automatically enabled on the E\_Ports, except for long-distance E\_Ports. For long-distance E\_Ports, you must manually enable QoS mode.
- If a QoS zone name prefix is specified in an LSAN zone (a zone beginning with the prefix "LSAN\_"), the QoS tag is ignored. Only the first prefix in a zone name is recognized. For example, a zone with the name "LSAN\_QOSH\_zone1" is recognized as an LSAN zone and not a QoS zone. See [QoS over FC Routers](#) for additional considerations when using QoS to prioritize traffic between device pairs in different edge fabrics.
- If there is a single low-priority flow to a destination ID (DID) and there are several medium-priority flows to that same DID, the medium-priority flows might have less bandwidth. This is because they have to share the medium-priority fabric resources, whereas the low-priority flow would have a separate set of fabric resources for its exclusive use.

## Setting QoS Zone-Based Traffic Prioritization

This section describes the steps to set the QoS zone-based traffic prioritization.

1. Connect to the switch, and log in using an account with admin permissions.
2. Enter the `zonecreate` command to create zones for high-priority and low-priority traffic.

- For high-priority traffic, use the following syntax:

```
zonecreate "QOSHid_zonename", "member[; member...]"
```

- For low-priority traffic, use the following syntax:

```
zonecreate "QOSLid_zonename", "member[; member...]"
```

The `id` range is from 1 through 5 for high-priority traffic, which corresponds to VCs 10 through 14. For low-priority traffic, the `id` range is from 1 through 2, which corresponds to VCs 8 and 9. The `id` is optional; if it is not specified, the virtual channels are allocated through a round-robin scheme.

3. Enter the `cfgadd` command to add the QoS zone to the zone configuration:

```
cfgadd "cfgname", "QOSzonename"
```

4. Enter the `cfgsave` command to save the change to the defined configuration.
5. Enter the `cfgenable` command for the appropriate zone configuration to make the change effective.

```
cfgenable "cfgname"
```

6. Enter the `portcfgqos` command to enable QoS on the E\_Ports:

```
portcfgqos --enable [slot/] port
```

The `portcfgqos` command does not affect QoS prioritization. It only enables or disables the link to pass the QoS priority traffic.

### NOTE

QoS is enabled by default on all ports (except long-distance ports). If you use the `portcfgqos` command to enable QoS on a specific port, the port is toggled to apply this configuration, even though the port already has QoS enabled. The port is toggled because the user configuration changed, even though the actual configuration of the port did not change. If you use the `portcfgqos` command later to enable QoS on the same port again, the port is *not* toggled because the configuration did not change.

```
sw0:admin> zonecreate "QOSH1_zone", "10:00:00:00:10:00:00:00; 10:00:00:00:20:00:00:00"
sw0:admin> zonecreate "QOSL2_zone", "10:00:00:00:30:00:00:00; 10:00:00:00:40:00:00:00"
sw0:admin> cfgadd "cfg1", "QOSH1_zone"
sw0:admin> cfgadd "cfg1", "QOSL2_zone"
sw0:admin> cfgshow
Defined configuration:
cfg: cfg1 QOSH1_zone; QOSL2_zone
zone: QOSH1_zone
 10:00:00:00:10:00:00:00; 10:00:00:00:20:00:00:00
zone: QOSL2_zone
 10:00:00:00:30:00:00:00; 10:00:00:00:40:00:00:00
Effective configuration:
No Effective configuration: (No Access)
```

```

sw0:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no] y
Updating flash ...

sw0:admin> cfgenable cfg1
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'cfg1' configuration (yes, y, no, n): [no] y
zone config "cfg1" is in effect
Updating flash ...

sw0:admin> portcfgqos --enable 3

```

## Setting QoS Zone-Based Traffic Prioritization over FC Routers

This section describe the steps to set the QoS zone-based traffic prioritization over FC routers.

1. Connect to the switch in the edge fabric, and log in using an account with admin permissions.
2. Create QoS zones in the edge fabric.

The QoS zones must have WWN members only, not D,I members.

3. Create LSAN zones in the edge fabric.

See [Controlling Device Communication with the LSAN](#) for instructions.

4. Enter the `portCfgQos --enable` command to enable QoS on the E\_Ports.

```
portcfgqos --enable [slot/] port
```

5. Repeat Step 1 through Step 3 to create QoS zones and LSAN zones on the other edge fabric.
6. Connect to the FC router in the backbone fabric, and log in using an account with admin permissions.
7. Enter the `portCfgQos --enable` command to enable QoS on the EX\_Ports.

```
portcfgqos --enable [slot/] port
```

## Disabling QoS Zone-Based Traffic Prioritization

This section describe the steps to disable QoS zone-based traffic prioritization.

1. Connect to the switch, and log in using an account with admin permissions.
2. Enter the `cfgRemove` command to remove the QoS zones from the current zone configuration.

```
cfgremove "configname" "qos_zonename"
```

3. Enter the `cfgEnable` command for the appropriate zone configuration to make the change effective.

```
cfgenable "configname"
```



4. Enter the `portCfgQos --disable` command to disable QoS on the E\_Ports.

```
portcfgqos --disable [slot/] port
```

## Supported Configurations for QoS Zone-Based Traffic Prioritization

The following configuration rules apply to QoS zone-based traffic prioritization:

- By default, all devices are assigned medium priority. To assign high or low priority, hosts and targets must be connected to the Brocade 32G or 64G switches or port blades.
- QoS is enabled by default on 8G or later generation switches on E\_Ports and F\_Ports that connect to AG.

## Limitations and Restrictions for QoS Zone-Based Traffic Prioritization

- Enabling and disabling QoS is potentially disruptive to the I/O on the affected port.
- If a host and target are included in two or more QoS zones with different priorities, the following priorities take precedence:
  - High-priority QoS zones and non-QoS zones = High-priority
  - High-priority QoS zones and Low-priority QoS zones = Low-priority
  - Low-priority QoS zones and non-QoS zones = Low-priority
  - High-priority QoS zones, Low-priority QoS zones, and non-QoS zones = Low-priority

For example, if an effective zone configuration has QOSH\_z1 (H,T) and QOSL\_z2 (H,T), the traffic flow between H and T will be of low QoS priority.
- If QOSH\_z1 (H,T) overlaps with a D,I (domain,index) zone at the H port, the traffic flow between H and T is dropped to medium priority, and the H port is marked as a session-based zoning port. This behavior applies only to Gen 6 platforms.
- Traffic prioritization is enforced on the egress ports only, not on the ingress ports.
- Traffic prioritization is not supported on mirrored ports.
- Traffic prioritization is not supported over LSAN zones beyond L2. The traffic is always medium priority in the ingress edge fabric, the backbone fabric, and the egress edge fabric.
- QoS zones that use D,I notation are not supported for QoS over FCR.
- QoS zones that use D,I notation should not be used for loop or NPIV ports.
- If QoS is enabled, an additional 16 buffer credits are allocated per port for 8G ports in Extended Mode (LE). See [Managing Long-Distance Fabrics](#) for information about buffer credit allocation in extended fabrics.
- If some ports in a trunk group have QoS enabled and some ports have QoS disabled, then two different trunks are formed, one with QoS enabled and one with QoS disabled.
- QoS mode cannot be explicitly enabled on simulation ports (SIM ports). CS\_CTL mode can be enabled on SIM ports.

## CS\_CTL-Based Frame Prioritization

CS\_CTL-based frame prioritization allows you to prioritize the frames between a host and a target as having high, medium, or low priority, depending on the value of the CS\_CTL field in the FC frame header.

The CS\_CTL field in the FC header assigns a priority to a frame. This field can be populated by selected end devices (storage and host) and then honored by the switch, which assigns the frame, based on the value in the CS\_CTL field, to allocate appropriate resources throughout the fabric. This method of establishing QoS is an alternative to the switch-controlled assignment, which uses zone-based QoS.

**ATTENTION**

Check with your host and storage manufacturers to determine whether they support Fibre Channel CS\_CTL prioritization on their devices.

High-, medium-, and low-priority frames are allocated to different sets of fabric resources. High-priority frames are assigned more fabric resources than medium-priority frames, which in turn are assigned more fabric resources than low-priority frames. The resources are allocated according to the CS\_CTL value, as shown in the following table. The values are enabled by default to ensure backward compatibility.

**Table 91: Mapping of CS\_CTL Values to QoS Priority for Frame Prioritization in CS\_CTL Default Mode**

| CS_CTL Value | Priority |
|--------------|----------|
| 1 to 8       | Low      |
| 9 to 16      | Medium   |
| 17 to 24     | High     |

Alternatively, the user can apply CS\_CTL auto mode. CS\_CTL auto mode uses only three CS\_CTL values, as illustrated in the following table.

**Table 92: Mapping of CS\_CTL Values to QoS Priority for Frame Prioritization in CS\_CTL Auto Mode**

| CS_CTL Value | Priority |
|--------------|----------|
| 1            | Low      |
| 2            | Medium   |
| 3            | High     |

**NOTE**

The values in the tables represent chassis-level configurations. For configuration details, see [Using CS\\_CTL Auto Mode at the Chassis Level](#), and [Considerations for Using CS\\_CTL-based Frame Prioritization](#).

## Supported Configurations for CS\_CTL-Based Frame Prioritization

CS\_CTL-based frame prioritization is supported on all 8G, 16G, 32G, and 64G platforms.

## Enabling CS\_CTL-Based Frame Prioritization on Ports

Make sure that the CS\_CTL-mode-enabled source and destination are zoned together.

When you enable CS\_CTL-based frame prioritization, you must enable it on both the source port and the destination port. This makes the frames that are returned from the destination port for a given exchange always have the same CS\_CTL prioritization as the frames that originate from the source port.

1. Connect to the switch, and log in to an account that has admin permissions.
2. Enable CS\_CTL mode by entering the following command.

```
switch:admin> portcfgqos --enable 60 csctl_mode
Enabling CSCTL mode flows casuses QoS zone flows to lose priority on such ports.
Enabling CSCTL mode takes precedence over quarantine of traffic destined to zoned slow drain device.

Do you want to proceed?(y/n):y
```

3. Enter `y` at the prompt to override QoS zone-based traffic prioritization.

## Disabling CS\_CTL-Based Frame Prioritization on Ports

When you disable CS\_CTL-based frame prioritization, QoS zone-based traffic prioritization is restored if it had been previously enabled.

1. Connect to the switch, and log in to an account that has admin permissions.
2. Disable CS\_CTL mode using the `portCfgQos` command with the `--disable` option:

```
portcfgqos --disable [slot/] port csctl_mode
```

Alternatively, you can use the `--default` option to disable CS\_CTL mode and set QoS to auto-enable (AE):

```
portcfgqos --default [slot/] port
```

## Using CS\_CTL Auto Mode at the Chassis Level

You can use the CS\_CTL QoS mode options for the `configureChassis` command to change the chassis-wide default mode, as in the following example.

```
switch:admin> configurechassis
Configure...
cfgload attributes (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no]
system attributes (yes, y, no, n): [no]
fos attributes (yes, y, no, n): [no] y
 CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [0] 1
```

Set CSCTL QoS Mode to 1 to enable auto mode, establishing the settings shown in [Table](#). Set CSCTL QoS Mode to 0 to disable auto mode and revert to the default settings shown in [Table](#).

This is a chassis-level configuration. It does not provide options to enable CS\_CTL QoS on the ports.

### ATTENTION

After changing the CS\_CTL QoS mode in a chassis, you must run the `slotPowerOff/On` commands for all the edge blades; whereas, in a fixed-port switch, you must reboot the switch. This is required for the new CS\_CTL QoS mode to become effective, because this mode change affects the persistent storage in the switch/chassis. To determine the current mode, use the following command:

```
switch:admin> configshow -all | grep csctl
Default mode - fos.csctlMode:0
Auto mode - fos.csctlMode:1
```

## Considerations for Using CS\_CTL-Based Frame Prioritization

To use CS\_CTL for QoS on a given port for a given flow, perform the following steps.

1. Determine whether to use the default mode or the auto mode. Not choosing results in the default mode.
2. In either case, ensure that the switch port connected to the initiator host and the switch port connected to the target host have CS\_CTL-mode enabled. See [Enabling CS\\_CTL-Based Frame Prioritization on Ports](#) for more information.
3. For a chassis-based system, use the `slot poweroff / slot poweron` command pair for all edge blades in the chassis. For a fixed-port switch, reboot the switch.

### ATTENTION

Changing the CS\_CTL QoS mode using the `configureChassis` command is a disruptive operation. This command updates the chassis configuration and stores the information that you provide in the chassis-level persistent database. When the ASIC is initialized, it retrieves this information to initialize its CSCTL-VC table. You must power-cycle all edge blades in the chassis for the ASIC to register the new information. For fixed-port switches, you must reboot the switch. To determine the current mode, use the following command:

```
switch:admin> configshow -all | grep csctl
Default mode - fos.csctlMode:0
Auto mode - fos.csctlMode:1
```

## Managing Trunking Connections

---

This section describes how to manage trunking connections. Trunking optimizes the use of bandwidth by allowing a group of links to merge into a single logical link, called a trunk group.

Traffic is distributed dynamically and in order over this trunk group, achieving greater performance with fewer links. Within the trunk group, multiple physical ports appear as a single port, thus simplifying management. Trunking also improves system reliability by maintaining in-order delivery of data and avoiding I/O retries, if one of the link within the trunk group fails.

Trunking is frame-based instead of exchange-based. A frame is much smaller than an exchange, which means that frame-based trunks are more granular and better balanced than exchange-based trunks. Frame-based trunks provide the maximum utilization of links.

The Trunking license is required for any type of trunking, and must be installed on each switch that participates in trunking. For details on obtaining and installing licensed features, refer to the *Brocade Fabric OS Software Licensing User Guide*.

### Types of Trunking

Trunking can be between two switches, between a switch and an Access Gateway module, or between a switch and a Brocade adapter. The types of trunking are as follows:

- ISL trunking, or E\_Port trunking, is configured on an inter-switch link (ISL) between two Fabric OS switches and is applicable only to E\_Ports.
- ICL trunking is configured on an inter-chassis link (ICL) between two Brocade X7 Directors, X6 Directors, or both, and it is applicable only to ports on the core blades.  
See [Inter-Chassis Links](#) for detailed information about ICL trunking.
- EX\_Port trunking is configured on an inter-fabric link (IFL) between an FC router (EX\_Port) and an edge fabric (E\_Port). The trunk ports are EX\_Ports connected to E\_Ports.  
See [EX\\_Port Trunking](#) for additional information about EX\_Port trunking.
- F\_Port trunking is configured on a link between a switch and either an Access Gateway module or a Brocade adapter. The trunk ports are F\_Ports (on the switch) connected to N\_Ports (on the Access Gateway or adapter).  
See [F\\_Port Trunking](#) for additional information about F\_Port trunking.
- N\_Port trunking is configured on a link between a switch and either an Access Gateway module or a Brocade adapter. It is similar to F\_Port trunking. The trunk ports are N\_Ports (on the Access Gateway or adapter) connected to F\_Ports (on the switch).  
See [Configuring F\\_Port Trunking for a Brocade Adapter](#) for additional information.

#### NOTE

This chapter uses the term *F\_Port trunking* to refer to a trunk between the F\_Ports on a switch and the N\_Ports on either an Access Gateway module or a Brocade adapter.

### Controllerless Trunking

Controllerless trunking means that if the controller port goes offline, one of the agent ports automatically becomes the new controller port, thus avoiding traffic disruption. The new controller port uses the old controller port area and the old controller port is assigned a new, unused area. In this way, the port identifier (PID) of the trunk does not change if the controller port goes offline.

If trunking is Non-Controllerless, and if the controller port goes offline, traffic disruption can occur because the agent ports in the trunk group goes offline to select the new controller port and then come back online.

**Table 93: Supported Controllerless Trunking for Gen 6 and Gen 7 Platforms**

| Ports            | Gen 6 Switches with VF Disabled | Gen 6 Switches with VF Enabled | Gen 6 Directors with VF Disabled | Gen 6 Directors with VF Enabled | Gen 7 Platforms Irrespective of VF Enabled or Disabled |
|------------------|---------------------------------|--------------------------------|----------------------------------|---------------------------------|--------------------------------------------------------|
| E_Port Trunking  | Controllerless                  | Controllerless                 | Controllerless                   | Controllerless                  | Controllerless                                         |
| EX_Port Trunking | Controllerless                  | Controllerless                 | Controllerful                    | Controllerless                  | Controllerless                                         |
| F_Port Trunking  | Controllerless                  | Controllerless                 | Controllerless                   | Controllerless                  | Controllerless                                         |
| ICL Trunking     | Controllerless                  | Controllerless                 | Controllerless                   | Controllerless                  | Controllerless                                         |
| ISL Trunking     | Controllerless                  | Controllerless                 | Controllerless                   | Controllerless                  | Controllerless                                         |

## License Requirements for Trunking

Trunking of non-ICL ports (E\_Ports, EX\_Ports, and F\_Ports) requires the Trunking license. This license must be installed on each switch that participates in trunking.

Trunking of ICL ports (E\_Ports and EX\_Ports) does not require a Trunking license.

### ATTENTION


After you add the Trunking license, to enable trunking functionality, you must disable and then re-enable each port to be used in trunking, or disable and reenable the switch.

Refer to the *Brocade Fabric OS Software Licensing User Guide* for information about activating licenses.

## Port Groups for Trunking

For trunk groups to form, several conditions must be met. One of the conditions is that all of the ports in a trunk group must belong to the same port group. A *port group* is a group of eight ports, based on the user port number, such as 0 to 7, 8 to 15, 16 to 23, and up to the number of ports on the switch. The maximum number of port groups is platform-specific.

Ports in a port group are usually contiguous, but they may not be. Refer to the hardware reference manual for your switch for information about which ports can be used in the same port group for trunking.

**Figure 23: Port Group Configuration for the Brocade G620**

## Supported Platforms for Trunking

Trunking is supported on the FC ports of all Brocade platforms and blades. EX\_Port trunking is supported only on those platforms that support EX\_Ports. See [Supported platforms for FC-FC routing](#) for more information.

## Supported Configurations for Trunking

- Links within a trunk can be 4, 8, 10, 16, 32 or Gen 7 ICL depending on the Brocade platform.

### NOTE

4G QSFP ports is supported only in Gen 6 platforms.

### NOTE

Gen 7 ICL QSFP ports can be used to create trunks between a pair of X7 Directors for ICL. For more information on ICL trunking between Gen 7 Directors, see [ICL Trunking Between X7 Directors](#)

**NOTE**

32G QSFP ports can be used to create trunks between a pair of Brocade G620 for ISL, and between a pair of Brocade X6 Directors for ICL. For more information on ICL trunking between Gen 6 Directors, see [ICL trunking between Brocade X6 Directors](#).

- The maximum number of ports per trunk is generally eight, but the number of trunks per switch depends on the Brocade platform.
- You can have up to eight ports in one trunk group to create high-performance ISL trunks between switches, providing up to 256G (based on a 32G port speed).
- An E\_Port or EX\_Port trunk can be up to eight ports wide. All the ports must be adjacent to each other, in the clearly marked groups on the front of the switch.

Trunks operate best when the cable length of each trunked link is roughly equal to the length of the others in the trunk. For optimal performance, use no more than 30 meters difference. Trunks are compatible with both short-wavelength (SWL) and long-wavelength (LWL) fiber-optic cables and transceivers.

Trunking is performed according to the Quality of Service (QoS) configuration on the controller and the agent ports. That is, in a given trunk group, if some ports with QoS are enabled and some with QoS are disabled, they form two different trunks, one with QoS enabled and the other with QoS disabled.

## ISLs between G630 Switches and Other Devices

ISLs can be created between G630 switches and other devices to form a core-edge topology, departmental SAN, native FC extension, and integrated routing topologies.

The following are the allowed trunk port groups on the G630 switch:

**Table 94: G630 QSFP and SFP Trunking Port Groups**

| Trunk Groups       | Eight Port Trunks | Allowed Member Ports      |
|--------------------|-------------------|---------------------------|
| SFP trunk group 1  | Trunk 1           | 0 to 7                    |
|                    | Trunk 2           | 8 to 15                   |
|                    | Trunk 3           | 48 to 55                  |
|                    | Trunk 4           | 56 to 63                  |
| SFP trunk group 2  | Trunk 5           | 16 to 23                  |
|                    | Trunk 6           | 24 to 31                  |
|                    | Trunk 7           | 32 to 39                  |
|                    | Trunk 8           | 40 to 47                  |
| SFP trunk group 3  | Trunk 9           | 64 to 71                  |
|                    | Trunk 10          | 72 to 79                  |
|                    | Trunk 11          | 80 to 87                  |
|                    | Trunk 12          | 88 to 95                  |
| QSFP trunk group 1 | Trunk 13          | 96 to 99 and 100 to 103   |
|                    | Trunk 14          | 104 to 107 and 108 to 111 |
| QSFP trunk group 2 | Trunk 15          | 112 to 115 and 116 to 119 |
|                    | Trunk 16          | 120 to 123 and 124 to 127 |


## QSFP-to-QSFP ISLs

QSFP-to-QSFP ISL connections are supported between G630 and the following devices:

- Another G630 switch.
- With a G620 switch.
- With an FC32-64 port blade on the X6 chassis.

The QSFP ports in a Brocade G630 switch can form ISLs between G630 switches, freeing up SFP+ ports for device connectivity.

- A three-switch full mesh configuration provides 288 (3 x 96) device ports.
- Eight links in 2 x QSFP ports can be trunked to get up to 256Gb/s bandwidth on ISLs.




### **Guidelines for QSFP-to-QSFP ISL Connections**

- You must select the QSFP ports from the two separate QSFP port groups. For example, refer to the following illustration for the trunks and allowed QSFP ports.




- Connection 1 between switch 1 and switch 2 using ports 96 to 99 (Group 1) and ports 112 to 115 (Group2)
- Connection 2 between switch 1 and switch 3 using ports 104 to 107 (Group 1) and ports 120 to 123 (Group 2)
- Connection 3 between switch 2 and switch 3 using ports 104 to 107 (Group 1) and ports 124 to 127 (Group 2)
- QSFP-to-QSFP connection from a G630 switch to QSFP ports on the X6 Director core routing blades are not supported.



## QSFP-to-SFP ISLs


QSFP ports in Brocade G630 can form dense ISLs between G630 switches and Director or any other device SFP+ ports simplifying cabling.

- For all QSFP-to-SFP connections, you must use only the Brocade qualified breakout cables.
- Up to eight links in 2 x QSFP ports can be trunked to get up to 256Gb/s bandwidth on ISLs.
- When the G630 switch is used as an edge switch, the ISL connections to the Directors should be made observing the trunk boundaries.
- You should buy QSFP POD licenses for the Brocade G630 switch and use them for connections to storage ports (F\_Ports) for a fully populated G630 switch.
- You should use the SFP+ ports in the G630 switch to connect to host ports and Brocade Analytics Monitoring Platform.
- Connections between QSFP ports on the G630 switch and ICL ports on Director core blades are not supported.



### **Guidelines for QSFP-to-SFP+ ISL Connections**

- You must select the QSFP ports from the two separate QSFP port groups. For example, refer to the following illustration for the trunks and allowed QSFP ports.
  - Connection 1 between switch 1 and Director 1 using ports 100 to 103 (Group 1) and ports 112 to 115 (Group 2)
  - Connection 2 between switch 1 and switch 2 using ports 104 to 105 (Group 1) and ports 120 to 121 (Group 2)
  - Connection 3 between switch 1 and switch 3 using ports 106 to 107 (Group 1) and ports 122 to 123 (Group 2)
  - Connection 4 from switch 1 to storage devices using any of the free ports.
- At a minimum, there should be at least two ISL breakout connections per domain.

**NOTE**

While forming a four-port trunk between G630 QSFP ports and G620 or FC32-48 blade, the breakout ports must be in the same trunk group in the destination device as well.

**NOTE**

You should use the SFP ports for connecting hosts and storage devices but use the QSFP ports only for storage connections.

## High-Availability Support for Trunking

Trunking is a High-Availability (HA) supported feature. The HA protocol for trunking is as follows:

- If trunking is disabled prior to the HA failover, it remains disabled after the HA failover.
- If trunking is enabled prior to the HA failover, it remains enabled after the HA failover.

## Requirements for Trunk Groups

The following requirements apply to all types of trunking:

- All of the ports in a trunk group must belong to the same port group.
- All of the ports in a trunk group must meet the following conditions:

- They must be running at the same speed.
- They must be configured for the same distance.
- They must have the same QoS and FEC state.
- They must have the same encryption and compression state.
- Trunk groups must be between Brocade switches (or Brocade adapters in the case of F\_Port trunking).
- A direct connection must exist between participating switches.
- Trunking cannot be performed if ports are in ISL R\_RDY mode. (You can disable this mode by using the `portCfgIslMode` command.)
- Trunking is supported only on FC ports. Virtual FC ports (VE\_Ports) do not support trunking.

## Recommendations for Trunk Groups

To identify the most useful trunk groups, consider the following recommendations along with the standard guidelines for SAN design:

- Evaluate the traffic patterns within the fabric.
- Place trunking-capable switches adjacent to each other.  
This maximizes the number of trunk groups that can form. If you are using a core and edge topology, place trunking-capable switches at the core of the fabric and any switches that are not trunking-capable at the edge of the fabric.
- When connecting two switches with two or more ISLs, ensure that all trunking requirements are met to allow a trunk group to form.
- Determine the optimal number of trunk groups between each set of linked switches, depending on traffic patterns and port availability.  
The goal is to avoid traffic congestion without unnecessarily using ports that could attach other switches or devices.
- Each physical ISL uses two ports that could otherwise attach node devices or other switches.
- Trunk groups can be used to resolve ISL oversubscription if the total capability of the trunk group is not exceeded.
- Consider how the addition of a new path will affect existing traffic patterns:
  - A trunk group has the same link cost as the controller ISL of the group, regardless of the number of ISLs in the group, which allows agent ISLs to be added or removed without causing data to be rerouted, because the link cost remains constant.
  - The addition of a path that is shorter than existing paths causes traffic to be rerouted through that path.
  - The addition of a path that is longer than existing paths may not be useful, because the traffic will choose the shorter paths first.
- Plan for future bandwidth addition to accommodate increased traffic.  
For trunk groups over which traffic is likely to increase as business requirements grow, consider leaving one or two ports in the group available for the future nondisruptive addition of bandwidth.
- Consider creating redundant trunk groups where additional ports are available or paths are particularly critical.  
Redundant trunk groups help to protect against oversubscription of trunk groups, multiple ISL failures in the same group, and the rare occurrence of an ASIC failure.
- To provide the highest level of reliability, deploy trunk groups in redundant fabrics to help ensure that ISL failures do not disrupt business operations.

## Configuring Trunk Groups

After you install the Trunking license, you must reinitialize the ports that are to be used in trunk groups so that they recognize that trunking is enabled. This procedure must be performed only once, and is required for all types of trunking.

To reinitialize the ports, you can either disable and then reenable the switch, or disable and then reenable the affected ports.

1. Connect to the switch and log in using an account that is assigned to the admin role.
2. Enter the `islShow -slotport` command to determine which slot/port are used for ISLs.

```
switch:admin> islshow -slotport
1:0/116-> 3/48 10:00:00:27:f8:f2:76:f8 57 Core-X7-8_130057 sp: 32.000G bw:256.000G TRUNK QOS CR_RECOV FEC
2:0/120-> 3/228 10:00:00:27:f8:f3:3b:40 13 Core-X6-8_130013 sp: 16.000G bw: 64.000G TRUNK QOS CR_RECOV FEC
3:0/124-> 3/192 10:00:88:94:71:ba:cd:e1 239 Core-X7-8_130239 sp: 32.000G bw: 32.000G QOS CR_RECOV FEC
4:0/125-> 3/193 10:00:88:94:71:ba:cd:e1 239 Core-X7-8_130239 sp: 32.000G bw: 96.000G TRUNK QOS CR_RECOV
FEC
```

3. Enter the `portDisable` command for each port to be used in a trunk group.

Alternatively, you can enter the `switchDisable` command to disable all ports on the switch.

4. Enter the `portEnable` command for each port that you disabled in Step 3, or enter the `switchEnable` command to enable all the ports on the switch.

#### NOTE

F\_Port trunking requires additional steps to configure the Trunk Area (TA). Refer to [Configuring F\\_Port Trunking for an Access Gateway](#) or [Configuring F\\_Port Trunking for a Brocade Adapter](#) for information.

## Enabling Trunking

You can enable trunking for a single port or for an entire switch. Because trunking is automatically enabled when you install the Trunking license, you need to use this procedure only if trunking has been subsequently disabled on a port or switch. Enabling trunking disables and reenables the affected ports. As a result, traffic through these ports may be temporarily disrupted.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portCfgTrunkPort` command to enable trunking on a port. Enter the `switchCfgTrunk` command to enable trunking on all ports on the switch.

```
portcfgtrunkport[slot/] port mode
switchcfgtrunk mode
```

Mode 1 enables trunking.

In the following example, trunking is being enabled on slot 1, port 3.

```
switch:admin> portcfgtrunkport 1/3 1
```

## Disabling Trunking

You can disable trunking for a single port or for an entire switch. Disabling trunking disables and reenables the affected ports. As a result, traffic through these ports may be temporarily disrupted.

Trunking on ICLs is always enabled and cannot be disabled.

Disabling trunking fails if a Trunk Area (TA) is enabled on the port. Use the `portTrunkArea` command to remove the TA before disabling trunking.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portCfgTrunkPort` command to disable trunking on a port. Enter the `switchCfgTrunk` command to disable trunking on all ports on the switch.

```
portcfgtrunkport[slot/] port mode
switchcfgtrunk mode
```

Mode 0 disables trunking.

```
switch:admin> switchcfgtrunk 0
```

## Displaying Trunking Information

Use the `trunkshow` command to view the following:

- All trunks and members of a trunk.
- The controller port connection for the trunk group.
- Correct formation of trunks.
- Trunking information of a switch that is part of an FC router backbone fabric interlinking several edge fabrics.
- Bandwidth and throughput of all trunk groups in a switch.

Use the `portperfshow` command to monitor problem areas where there are congested paths or dropped links to determine whether you must adjust the fabric design by adding, removing, or reconfiguring ISLs and trunking groups.

To view detailed information about F\_Port trunking, see [Displaying F\\_Port Trunking Information](#).

Use the following procedure to view trunking information:

1. Connect to the switch, and log in using an account that is assigned to the admin role.
2. Enter the `trunkshow` command.

The following example shows trunking groups 1 and 2; ports 25 and 35 are controller ports.

```
switch:admin> trunkshow
1: 25->694 10:00:00:05:33:e7:2d:60 4 deskew 1 E-Port MASTER
 26->695 10:00:00:05:33:e7:2d:60 4 deskew 1
2: 35->704 10:00:00:05:33:e7:2d:60 4 deskew 0 E-Port MASTER
```

### NOTE

- Deskew is the difference between the time for traffic to travel over each ISL and the time through the shortest ISL in the group plus the minimum deskew value. The value is expressed in nanoseconds divided by 10.
- For Brocade Gen 7 platforms, the minimum deskew value is from 1 through 10 and is set based on the link speed.
- For Brocade Gen 6 platforms, the minimum deskew value is from 4 through 14 and is set based on the link speed.

The following example show F\_Port trunking information along with the bandwidth, capacity, and utilization for all trunk groups in a G730 switch.

```
switch:admin> trunkshow -perf

1: 14-> 1 10:01:00:10:9b:8f:2d:9f deskew 3 F-Port MASTER
 13-> 0 10:00:00:10:9b:8f:2d:9f deskew 9
Tx: Bandwidth 64.00Gbps, Capacity 54.40Gbps, Utilization 29.00Gbps (53.30%)
Rx: Bandwidth 64.00Gbps, Capacity 54.40Gbps, Utilization 28.99Gbps (53.29%)
Tx+Rx: Bandwidth 128.00Gbps, Capacity 108.80Gbps, Utilization 57.99Gbps (53.30%)
```

The following example show E\_Port trunking information along with the bandwidth, capacity, and utilization for all trunk groups in an X7-8 Director.

```
switch:admin> trunkshow -perf

1:105-> 27 10:00:d8:1f:cc:03:a5:43 10 deskew 0 E-Port MASTER
Tx: Bandwidth 32.00Gbps, Capacity 27.20Gbps, Utilization 17.33Gbps (63.72%)
Rx: Bandwidth 32.00Gbps, Capacity 27.20Gbps, Utilization 14.09Gbps (51.80%)
Tx+Rx: Bandwidth 64.00Gbps, Capacity 54.40Gbps, Utilization 31.42Gbps (57.76%)
```

From Fabric OS 9.1.0, the `trunkshow` command is added with the new option `-slotport`, which displays both local and neighbor details in slot/port format. The following example shows trunking groups with the new slot/port format.

```
switch:admin> trunkshow -slotport

1: 0/31 -> 0/75 10:00:00:05:33:e7:2d:60 22 deskew 0 E-Port MASTER

2: 0/45 -> 0/17 10:00:00:05:33:e7:2d:60 22 deskew 1 E-Port MASTER
 0/44 -> 0/16 10:00:00:05:33:e7:2d:60 22 deskew 1
```

#### NOTE

During non-disruptive firmware migration, the ports which are already online do not show the slot/port details. To get this details, the port has to go through E\_Port initialization. For more information, see [E\\_Port Login Process](#).

## ISL Trunking over Long-Distance Fabrics

In long-distance fabrics, if a port speed is set to autonegotiate, the maximum speed is assumed for reserving buffers for the port. If the port is running at 8G or 4G, this speed wastes buffers. For long-distance ports, specify the port speed instead of setting it to autonegotiate.

In addition to the criteria listed in [Supported Configurations for Trunking](#), observe the following criteria for trunking over long-distance fabrics:

- Extended Fabrics and Trunking licenses are required on all participating switches.
- When configuring long distance, you must configure the `portCfgLongDistance --vc_translation_link_init` parameter to be the same on all ports in a long-distance fabric.

For additional information on configuring long distance, see [Configuring an Extended ISL](#).

The below table summarizes the support for trunking over long-distance fabrics for the X6 and X7 Directors and supported blades.

**Table 95: Trunking over Long Distance for the Brocade Directors and Blades**

| Long-distance Mode | Distance | Number of 4G Ports              | Number of 8G Ports     |
|--------------------|----------|---------------------------------|------------------------|
| LE                 | 10 km    | 48 (six 8-port trunks)          | 48 (six 8-port trunks) |
| L0                 | Normal   | See note                        | 48 (six 8-port trunks) |
| LD                 | 200 km   | 4 (one 2-port trunk per switch) | 0                      |
| LD                 | 250 km   | 4 (one 2-port trunk per switch) | 0                      |
| LD                 | 500 km   | 0                               | 0                      |
| LS                 | Static   | See note                        |                        |

**NOTE**

The L0 mode supports up to 5 km at 2G, up to 2 km at 4G, and up to 1 km at 8/10G, 625 m at 16G, and 150 m at 32G. The distance for the LS mode is static. You can specify any distance greater than 10 km.

The distance that is supported depends on the available buffers, the number of back-end ports, and the number of ports that are offline.

## EX\_Port Trunking

You can configure EX\_Ports to use trunking just as you do regular E\_Ports. EX\_Port trunking support is designed to provide the best utilization and balance of frames transmitted on each link between the FC router and the edge fabric. You should trunk all ports connected to the same edge fabrics.

The FC router front domain has a higher node WWN, derived from the FC router, than that of the edge fabric. Therefore, the FC router front domain initiates the trunking protocol on the EX\_Port.

After initiation, the first port from the trunk group that comes online is designated as the controller port. The other ports that come online on the trunk group are considered to be the agent ports. Adding or removing a agent port does not cause frame drop. However, removing a agent port causes the loss of frames in transit.

If router port cost is used with EX\_Port trunking, the controller port and agent ports share the router port cost of the controller port.

The restrictions for EX\_Port frame trunking are the same as for E\_Ports. All the ports must be adjacent to each other, in the clearly marked groups on the front of the switch.

Refer to [Using FC-FC Routing to Connect Fabrics](#) for more information about EX\_Ports and the FC router.

## Controllerless EX\_Port Trunking

If controllerless EX\_Port trunking is not in effect and the controller port goes offline, the entire EX\_Port-based trunk reforms and is taken offline for a short period of time. If there are no other links to the edge fabric from the backbone, the controller port going offline may cause a traffic disruption in the backbone fabric. For more information, see [Controllerless Trunking](#).

## Supported Configurations and Platforms for EX\_Port Trunking

EX\_Port trunking is a Fibre Channel Routing (FCR) software feature and requires that you have a Trunking license installed on the FC router and on the edge fabric connected to the other side of the trunked EX\_Ports. The Trunking license is not required for EX\_Ports on an ICL.

EX\_Port trunking is supported only with Brocade edge fabrics.

You can use EX\_Port frame trunking in the following configurations and cases:



- For ports with speeds of 4G up to maximum speed and trunking over long distance.
- In the edge fabric, when the FC router is connected to a switch that supports eight ports from the trunkable group.
- When the FC router is connected to an edge fabric through a mix of trunked and nontrunked EX\_Ports; all will share the same front domain.
- In edge-to-edge, backbone-to-edge, and dual-backbone configurations.

Controllerless EX\_Port trunking has additional configuration requirements. See [Controllerless EX\\_Port Trunking](#) for these additional requirements.

#### NOTE

QoS and EX\_Port trunking can coexist. However, if some ports in the trunk group have QoS enabled and some have QoS disabled, two trunk groups will form: one with QoS enabled and one with QoS disabled. If the EX\_Ports are coming from a base switch, then QoS will not work on the EX\_Ports.

## Backward Compatibility Support

For backward compatibility, an FC router that supports EX\_Port trunking can continue to interoperate with older FC routers and all previously supported Brocade switches in the backbone fabric or Brocade edge fabric.

## Configuring EX\_Port Trunking

With EX\_Port trunking, you use the same CLI commands as you do for E\_Port trunking. See [Configuring Trunk Groups](#) for instructions.

## Displaying EX\_Port Trunking Information

1. Log in as an admin and connect to the switch.
2. Enter the `switchShow` command to display trunking information for the EX\_Ports.

The following is an example of a controller EX\_Port and an agent EX\_Port displayed in `switchShow` output.

```
switch:admin> switchshow
Index Slot Port Address Media Speed State
=====
16 2 0 ee1000 id N4 No_Light
17 2 1 ee1100 id N4 Online EX_Port (Trunk port, master is Slot 2 Port 2)
18 2 2 ee1200 id N4 Online EX_Port 10:00:00:05:1e:35:bb:32
 "MtOlympus_82" (fabric id = 2) (Trunk master)
19 2 3 ee1300 id N4 No_Light
20 2 4 ee1400 id N4 Online EX_Port (Trunk port, master is Slot 2 Port 7)
21 2 5 ee1500 id N4 Online EX_Port (Trunk port, master is Slot 2 Port 7)
22 2 6 ee1600 id N4 Online EX_Port (Trunk port, master is Slot 2 Port 7)
23 2 7 ee1700 id N4 Online EX_Port 10:00:00:60:69:80:1d:bc
 "MtOlympus_72" (fabric id = 2) (Trunk master)
```

## F\_Port Trunking

You can configure F\_Port trunking in the following scenarios:

- Between F\_Ports on a Fabric OS switch and N\_Ports on an Access Gateway module
- Between F\_Ports on a Fabric OS switch and N\_Ports on adapters that are compatible with the Brocade Trunking technology.

For F\_Port trunking, you must create a Trunk Area (TA) within the trunk group. When you assign an area within a trunk group, that group is enabled for F\_Port trunking. The TA that you assign must be within the eight-port trunk group beginning with port 0 (zero). After you assign a TA to a port, the port immediately acquires the TA as the area of its PID. Likewise, after you remove a TA from a port, the port immediately acquires the default area as its PID. F\_Port trunking prevents reassignments of the Port ID (also referred to as the Address Identifier) when F\_Ports go offline, and it increases the F\_Port bandwidth.

#### NOTE

Ensure that the Trunking license is enabled on the connecting Access Gateway module before creating a TA on the edge switch. You must not create a TA if the connecting Access Gateway module has no trunking capability (that is, no Trunking license or the Trunking license is disabled) or else the corresponding N\_ports on the Access Gateway module is disabled.

Refer to the *Brocade Fabric OS Access Gateway User Guide* for information about configuring the corresponding N\_Port trunking on the Access Gateway.

## F\_Port Trunking for an Access Gateway

You can configure trunking between the F\_Ports on an edge switch and the N\_Ports on an Access Gateway (AG) module.

#### NOTE

You cannot configure F\_Port trunking on the F\_Ports of an Access Gateway module.

F\_Port trunking keeps F\_Ports from becoming disabled when they are mapped to an N\_Port on a switch in AG mode. With F\_Port trunking, any link within a trunk can go offline or become disabled, but the trunk remains fully functional and there are no reconfiguration requirements.

The following figures show a switch in AG mode with and without F\_Port controllerless trunking.

**Figure 24: Switch in Access Gateway Mode with F\_Port Controllerless Trunking**





**Figure 25: Switch in Access Gateway Mode without F\_Port Controllerless Trunking****NOTE**

You do not need to map the host to the controller port manually because the Access Gateway performs a cold failover to the controller port.

**Requirements for F\_Port Trunking on an Access Gateway**

In addition to the requirements listed in [Requirements for Trunking Groups](#), refer to the *Brocade Fabric OS Access Gateway User Guide* for additional requirements that are specific to F\_Port trunking on an Access Gateway.

**Configuring F\_Port Trunking for an Access Gateway**

Access Gateway trunking configuration is mostly on the edge switch. On the Access Gateway module, you only need to ensure that the Trunking license is applied and enabled.

Use the following procedure on the edge switch connected to the Access Gateway module to configure F\_Port trunking.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portCfgShow` command to ensure that the ports have trunking enabled. If trunking is not enabled, enter the `portCfgTrunkPort port 1` command.
3. Enter the `portDisable` command for each port to be included in the trunk area.
4. Enter the `portTrunkArea --enable` command to enable the trunk area.

For example, the following command creates a TA for ports 36 to 39 with index number 37.

```
switch:admin> porttrunkarea --enable 36-39 -index 37
Trunk index 37 enabled for ports 36, 37, 38 and 39.
```

When you assign a trunk area on a port, trunking is automatically enabled on the F\_Ports. The `portTrunkArea` command does not unassign a TA if its previously assigned Area\_ID is the same Address Identifier (Area\_ID) of the TA unless all the ports in the trunk group are specified to be unassigned.

5. Enter the `portEnable` command to reenable the ports in the TA.

## F\_Port Trunking for Brocade Adapters

You can configure trunking between the F\_Ports on an edge switch and the Brocade adapters.

In addition to the requirements listed in [Requirements for Trunk Groups](#), note the following requirements, which are specific to F\_Port trunking for Brocade adapters:

- The edge switch must be running in Native mode. You cannot configure trunking between the Brocade adapters and the F\_Ports of an Access Gateway module.
- You can configure only two F\_Ports in one trunk group.

### Configuring F\_Port Trunking for a Brocade Adapter

F\_Port trunking for Brocade adapters requires configuration on the FC switch as well as on the Brocade HBAs. This section describes the configuration steps you perform on the switch.

1. On the switch side, perform the following steps:
  - a) Configure both ports for trunking by using the `portCfgTrunkPort` command.
 

```
switch:admin> portcfgtrunkport 3/40 1
switch:admin> portcfgtrunkport 3/41 1
```
  - b) Disable the ports to be used for trunking by using the `portDisable` command.
 

```
switch:admin> portdisable 3/40
switch:admin> portdisable 3/41
```
  - c) Enable the trunk on the ports by using the `portTrunkArea` command.
 

```
switch:admin> porttrunkarea --enable 3/40-41 -index 296
Trunk index 296 enabled for ports 3/40 and 3/41.
```
2. On the host side, enable the trunking.
3. On the switch side, enable the ports by using the `portEnable` command.
 

```
switch:admin> portenable 3/40
switch:admin> portenable 3/41
```

## F\_Port Trunking Considerations

The following table describes the F\_Port controllerless trunking considerations.

**Table 96: F\_Port Controllerless Trunking Considerations**

| Category        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Area assignment | You statically assign the area within the trunk group on the edge switch. That group is the F_Port trunk. The static trunk area you assign must fall within the ASIC's trunk group of the switch or blade starting from port 0, and must be one of the port's default areas of the trunk group.                                                                                                                                                                                 |
| Authentication  | Authentication occurs only on the F_Port trunk controller port and only once per the entire trunk. This behavior is the same as E_Port trunk controller authentication. Because only one port in the trunk does FLOGI to the switch, and authentication follows FLOGI on that port, only that port displays the authentication details when you issue the <code>portShow</code> command.<br>Switches in Access Gateway (AG) mode support authentication in F_Ports and N_Ports. |

| Category           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| configdownload     | If you issue the <code>configDownload</code> command for a port configuration that is not compatible with F_Port trunking, and the port is Trunk Area-enabled, then the port will be persistently disabled. F_Port trunks will never be restored through <code>configDownload</code> .<br>Long distance, port mirroring, non-CORE_PID, and FastWrite are not compatible with F_Port trunking.                                                                                                                                                                                                                                                                                 |
| domain,index (D,I) | Creating a Trunk Area may remove the Index (I) from the switch to be grouped to the Trunk Area. All ports in a Trunk Area share the same index. This means that a <code>domain,index</code> (D,I), which refers to an index that might have been removed, will no longer be part of the switch.<br>Be sure to include zoning and DCC when creating a Trunk Area.<br>You can remove the port from the Trunk Area to place the index back in effect. D,I behaves as normal, but you may see the effects of grouping ports into a single index.<br>Also, D,I continues to work for Trunk Area groups. The index can be used in D,I if it was the index for the Trunk Area group. |
| DCC Policy         | DCC policy enforcement for the F_Port trunk is based on the Trunk Area; the FDISC requests to a trunk port are accepted only if the WWN of the attached device is part of the DCC policy against the Trunk Area. The PWWN of the FLOGI sent from the Access Gateway will be dynamic for the F_Port trunk controller. Because you do not know ahead of time what PWWN the Access Gateway will use, the PWWN of the FLOGI will not go through DCC policy check on an F_Port trunk controller. However, the PWWN of the FDISC continues to go through DCC policy check.                                                                                                          |
| Default Area       | Port X is a port that has its Default Area the same as its Trunk Area. The only time you can remove port X from the trunk group is when the entire trunk group has the Trunk Area disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Downgrade          | You can have trunking on, but you must disable the trunk ports before performing a firmware downgrade. Removing a Trunk Area on ports running traffic is disruptive because you must disable the port to disable the Trunk Area on the port.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| FastWrite          | When you assign a Trunk Area to a trunk group, the trunk group cannot have FastWrite enabled on those ports. If a port is FastWrite-enabled, the port cannot be assigned a Trunk Area.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| FICON              | FICON is not supported on F_Port trunk ports. However, FICON can still run on ports that are not F_Port trunked within the same switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Long Distance      | Long distance is not allowed on F_Port trunks, which means that a Trunk Area is not allowed on long-distance ports. You cannot enable long distance on ports that have a Trunk Area assigned to them.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Management Server  | Registered Node ID (RNID), Link Incident Record Registration (LIRR), and Query Security Attribute (QSA) extended link services (ELSSs) are not supported on F_Port trunks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| NPIV               | N_Port ID virtualization (NPIV) is supported on the F_Port controller trunk.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| PID format         | F_Port trunking is supported only in the CORE PID format.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Port mirroring     | Port mirroring is not supported on Trunk Area ports or on the PID of an F_Port trunk port.<br>Port mirroring is not supported on the Brocade Encryption Switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Port Swap          | When you assign a Trunk Area to a trunk group, the Trunk Area cannot be port swapped. If a port is swapped, you cannot assign a Trunk Area to that port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Port Types         | Only F_Port trunk ports are allowed on a Trunk Area port. All other port types are persistently disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| PWWN               | The entire Trunk Area trunk group shares the same Port WWN within the trunk group. The initial 5 is the default value for an F_Port trunk port in a virtual fabric and non-virtual fabric. The valid trunk port format is 25:xx:nn:nn:nn:nn:nn:nn                                                                                                                                                                                                                                                                                                                                                                                                                             |
| QoS                | QoS is supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Routing            | Routing will route against the F_Port trunk controller. Bandwidth information will be modified accordingly as the F_Port trunk forms.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Trunk Controller   | No more than one trunk controller is allowed in a trunk group. The second trunk controller will be persistently disabled with the reason "Area has been acquired".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## F\_Port Trunking in Virtual Fabrics

F\_Port trunking functionality performs the same in Virtual Fabrics as it does in non-Virtual Fabrics platforms. Fabric OS uses a 10-bit addressing model.

On the Brocade X6, F\_Port trunk ports dynamically receive an 8-bit area address that remains persistent. After F\_Port trunking configurations are removed from a port in a logical switch, that port returns to the default 10-bit area addressing model, which supports up to 1024 F\_Ports in a logical switch.

On the Brocade X7, F\_Ports and F-Port trunking ports dynamically receive an 10-bit area address that remains persistent in UAM mode.

The following are the F\_Port trunking considerations for Virtual Fabrics:

- If a port is enabled for F\_Port trunking, you must disable the configuration before you can move a port from the logical switch.
- If the user-bound area for a port is configured by means of the `portAddress` command, the port cannot be configured as an F\_Port trunk port. You must explicitly remove the user-bound area before enabling F\_Port trunking.
- F\_Port trunks are not allowed on the base switch because you cannot have F\_Ports on the base switch.
- If F\_Port trunking is enabled on some ports in the default switch, and you disable Virtual Fabrics, all of the F\_Port trunking information is lost.
- All of the ports in an F\_Port trunk must belong to a single trunk group of ports on the platform and must also belong to the same logical switch.

Refer to [Managing Virtual Fabrics](#) for detailed information about Virtual Fabrics.

## Displaying F\_Port Trunking Information

Use the following commands on the edge switch to verify the F\_Port trunking configuration.

- Enter `switchShow` to display the switch and port information.
- Enter `portTrunkArea --show enabled` to display the Trunk Area-enabled port configuration.

```
switch:admin> porttrunkarea --show enabled
```

| Port | Type   | State  | Master | TI | DI |
|------|--------|--------|--------|----|----|
| 36   | F-port | Master | 36     | 37 | 36 |
| 37   | F-port | Slave  | 36     | 37 | 37 |
| 38   | F-port | Slave  | 36     | 37 | 38 |
| 39   | F-port | Slave  | 36     | 37 | 39 |

- Enter `portTrunkArea --show trunk` to display the trunking information.

```
switch:admin> porttrunkarea --show trunk
```

```
Trunk Index 37: 39->0 sp: 8.000G bw: 16.000G skew 15 MASTER
Tx: Bandwidth 16.00Gbps, Throughput 1.63Gbps (11.84%)
Rx: Bandwidth 16.00Gbps, Throughput 1.62Gbps (11.76%)
Tx+Rx: Bandwidth 32.00Gbps, Throughput 3.24Gbps (11.80%)
38->1 sp: 8.000G bw: 8.000G skew 15
Tx: Bandwidth 16.00Gbps, Throughput 1.63Gbps (11.84%)
Rx: Bandwidth 16.00Gbps, Throughput 1.62Gbps (11.76%)
Tx+Rx: Bandwidth 32.00Gbps, Throughput 3.24Gbps (11.80%)
37->1 sp: 8.000G bw: 8.000G skew 15
Tx: Bandwidth 16.00Gbps, Throughput 1.63Gbps (11.84%)
Rx: Bandwidth 16.00Gbps, Throughput 1.62Gbps (11.76%)
Tx+Rx: Bandwidth 32.00Gbps, Throughput 3.24Gbps (11.80%)
36->1 sp: 8.000G bw: 8.000G skew 15
```

```
Tx: Bandwidth 16.00Gbps, Throughput 1.63Gbps (11.84%)
Rx: Bandwidth 16.00Gbps, Throughput 1.62Gbps (11.76%)
Tx+Rx: Bandwidth 32.00Gbps, Throughput 3.24Gbps (11.80%)
```

## Disabling F\_Port Trunking

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portDisable` command to disable the ports that are to be removed from the Trunk Area.
3. Enter the `portTrunkArea --disable` command to remove ports from the trunk area.

This command does not unassign a TA if its previously assigned Area\_ID is the same address identifier (Area\_ID) of the TA unless all the ports in the trunk group are specified to be unassigned.

```
switch:admin> portdisable 0-2
switch:admin> porttrunkarea --disable 0-2
Trunk index 2 disabled for ports 0, 1, and 2.
```

## Enabling the DCC Policy on a Trunk Area

After you assign a Trunk Area, the `portTrunkArea` command checks whether there are any active DCC policies on the port with the index TA, and then issues a warning to add all the device WWNs to the existing DCC policy with index as TA.

All DCC policies that refer to an index that no longer exists will not be in effect.

1. Add the WWN of all the devices to the DCC policy against the TA.
2. Enter the `secPolicyActivate` command to activate the DCC policy.

For security to enforce the DCC policy on the trunk ports, you must enable the TA *before* issuing the `secPolicyActivate` command.

3. Turn on the trunk ports.

Turn on the trunk ports *after* issuing the `secPolicyActivate` command to prevent the ports from becoming disabled in case there is a DCC security policy violation.

You can configure authentication on all Brocade trunking configurations. For more information on authentication, refer to [Configuring Security Policies](#).

## Inter-Chassis Links

---

An inter-chassis link (ICL) is a licensed feature that is used to interconnect the following chassis platforms. ICL ports in the core blades interconnect the chassis platforms, potentially increasing the number of usable ports in the chassis platforms.

- Two Brocade X6 Directors
- Two Brocade X7 Directors
- A Brocade X6 Director with a Brocade X7 Director

**NOTE**

To connect an X6 Director to an X7 Director, you must fix the port speed on the X7 Director to 32G and must use a 32G QSFP.

- A Brocade Gen 5 Director with a Brocade X7 Director

**NOTE**

To connect a Gen 5 Director to an X7 Director, you must fix the port speed on the X7 Director to 16G and must use a 32G QSFP.

When two chassis platforms are interconnected by ICLs, each chassis requires a unique domain and it is managed as a separate switch.

The ICL ports appear as regular ports, with some restrictions. All port parameters that are associated with the ICL ports are static, and all `portCfg` commands are blocked from changing any of the ICL port parameters, except for EX\_Port and FEC port parameters. The only management that is associated with ICL ports and cables is monitoring the status of the LEDs on the ICL ports and any maintenance if the Attention LED is blinking yellow.

The ICL ports are managed as E\_Ports. You can also configure EX\_Ports on the ICLs. See [Using FC-FC Routing to Connect Fabrics](#) for more information.

When you connect two chassis platforms, the following features are supported:

- Trunking
- Buffer-to-buffer credit
- QoS

**NOTE**

A Brocade Trunking license is not required for trunking on ICL connections.

Refer to the specific hardware reference manuals for additional information about LED status and ICL connections.

## License Requirements for ICLs

Brocade ICL links operate between the core blades of the Brocade X7 and Brocade X6 family. The ICL Port On Demand feature is used on X6-4, X6-8, X7-4 and X7-8.

After the addition or removal of a license, the license enforcement is performed on the ICL ports only after being disabled and enabled. An ICL license must be installed on the enterprise platforms at both ends of the ICL connection.

For more information on how license enforcement occurs, refer to the *Brocade Fabric OS Software Licensing User Guide*.

## ICLs between Brocade X6 Directors

You can connect up to nine directors in a full mesh topology and 12 directors in a core-edge topology using inter-chassis links (ICLs) on core routing blades that are installed in X6 Directors.

Each ICL connects the core blades of two Brocade X6 Director chassis and provides up to 128G of throughput within a single cable.



You can have up to 32 QSFP ports in an X6-8 chassis or 16 QSFP ports in an X6-4 chassis, with up to 4Tb/s of an ICL bandwidth and support for up to 100 meters over OM4 cable.

The 100-meter ICL is supported beginning in Fabric OS 8.0.1, when using 100-meter-capable QSFPs over OM4 cable.

The Brocade X6-8 has four port groups on the CR32-8 core blade. The Brocade X6-4 has two port groups on the CR32-4 core blade. Each port group has four QSFP connectors, and each QSFP connector maps to four user ports.

The following table shows the mappings from the numbered QSFP ports on the face of the core blade to the individual FC port numbers as shown by the `slotShow` command.


**Table 97: External Port to slotshow Port Mapping for Core Blades**

| External Port Number | slotShow FC Port Numbers | External Port Number | slotShow FC Port Numbers |
|----------------------|--------------------------|----------------------|--------------------------|
| 0                    | 0 to 3                   | 8                    | 32 to 35                 |
| 1                    | 4 to 7                   | 9                    | 36 to 39                 |
| 2                    | 8 to 11                  | 10                   | 40 to 43                 |
| 3                    | 12 to 15                 | 11                   | 44 to 47                 |
| 4                    | 16 to 19                 | 12                   | 48 to 51                 |
| 5                    | 20 to 23                 | 13                   | 52 to 55                 |
| 6                    | 24 to 27                 | 14                   | 56 to 59                 |
| 7                    | 28 to 31                 | 15                   | 60 to 63                 |

Following are ICL configuration guidelines for trunking bandwidth and high availability:

- A minimum of four ICL ports (two on each core blade) must be connected between each chassis pair.
- The dual connections on each core blade must reside within the same ICL trunk boundary on the core blades.
- If more than four ICL connections are required between a pair of X6 chassis, add ICL connections in pairs, one on each core blade with each additional pair residing within the same trunk boundary.

For high-availability, you need at least two ICLs from each core blade. The following illustration shows the connection between an X6-8 chassis and an X6-4 chassis with full redundancy using four ICL connections:

**Figure 26: Minimum Configuration for 128G ICLs**

1. X6-8 Chassis with two CR32-8 blades
2. X6-4 Chassis with two CR32-4 blades

- The maximum number of ICLs between two X6-4 chassis or between an X6-8 and an X6-4 is 16.
- The maximum number of ICLs between two X6-8 chassis is 32.

Because the FSPF routing logic uses only the first 16 paths to come online, only 16 ICLs are utilized. With Virtual Fabrics, however, you can define two logical switches on the chassis and have 16 ICLs in each.

**NOTE**

You should have a maximum of eight ICLs connected to the same neighboring domain, with a maximum of four ICLs from each core blade.

- The ICLs can connect to both the core blades in the neighboring chassis. The QSFP ICLs need not be cross-connected.

**NOTE**

QSFP ICLs and ISLs in the same logical switch and connected to the same neighboring switch are not supported. This topology restriction is with ICLs and any ISLs that are E\_Ports or VE\_Ports. If Virtual Fabrics is

enabled, you can have ICLs and ISLs between a pair of chassis if the ICLs are in a different logical switch than the ISLs.

## ICL Trunking between Brocade X6 Directors

Trunking optimizes the use of ICL bandwidth by allowing a group of links to merge into a single logical link, called a trunk. Traffic is distributed dynamically and in order over this trunk, achieving greater performance with fewer links. Within the trunk, multiple physical ports appear as a single route, thus simplifying management. Trunking also improves system reliability by maintaining in-order delivery of data and avoiding I/O retries if one link within the trunk fails.

Each QSFP provides four 32G Fibre Channel ports. Because each port within a QSFP terminates on a different ASIC within each core blade, an ICL trunk cannot be formed using the individual FC ports within the same QSFP. A trunk must be formed from individual FC ports in different QSFPs. These QSFPs must reside in the same trunk group. To form an ICL trunk between two devices, a minimum of two QSFPs within a port trunk group on a core blade installed in one device must be connected to a pair of QSFPs within a trunk group on a core blade in another device using ICLs.

ICL trunks form automatically but additional licenses may be required for enabling all ICL ports or for larger ICL configurations. For more information about ICL licensing options, refer to the *Brocade Fabric OS Software Licensing User Guide*.


### NOTE

Each QSFP cable has four ports, each terminating on a different ASIC. These ports cannot form a trunk with each other, but they can form trunks only with corresponding ports on another QSFP.

Each core blade in a Brocade X6-8 Director contains 16 ICL trunk groups. Each core blade in a Brocade X6-4 Director contains eight ICL trunk groups. Each ICL trunk group contains four user ports, one from each QSFP. To establish ICL trunking between platforms in the Brocade X6 Director family, the QSFP cables must be in the same trunk group. Up to four 4x32G QSFP ports in a trunk group can form a 512G trunk between two CR32-4 or CR32-8 blades.

The following figure marks the 4-port ICL trunking boundaries or port groups for Brocade X6 core routing blades (CR32-8 and CR32-4 blades).

**Figure 27: ICL Trunking for Brocade X6 Core Routing Blades**



Ports belonging to the same trunking groups are indicated with the same color border under the ports on the blade faceplate. These colors are also applied to the port map labels on each blade faceplate to indicate ports belonging to

the same trunking groups. Although a connection between two QSFPs on one core routing blade and two QSFPs on another blade in a different device is required to form an ICL trunk, you can create trunks consisting of two to four QSFP connections between these blades.

#### NOTE

Fabric OS does not permit a fabric configuration having E\_Port-to-E\_Port ISLs on port blades and at the same time E\_Port-to-E\_Port connections through ICLs on core blades between a pair of domains on two directors.

The following table summarizes the possible trunking connections between Brocade X6 Directors.

**Table 98: Trunking Connections between Brocade X6 Directors**

| Connecting Blades | 2- to 4-Port ICL Trunks                                                |
|-------------------|------------------------------------------------------------------------|
| CR32-8 and CR32-8 | Supported with QSFPs located within the same trunk group on each blade |
| CR32-8 and CR32-4 | Supported with QSFPs located within the same trunk group on each blade |
| CR32-4 and CR32-4 | Supported with QSFPs located within the same trunk group on each blade |

Refer to the specific Brocade X6 hardware reference manuals for information about port numbering, port trunk groups, and connecting the ICL cables.

## Using QSFPs That Support 2 km on ICL Ports

ICL ports support a 2-km distance. To support this distance, you must use specific QSFPs and allocate a greater number of buffer credits per port.

Consider the following points if you are attempting to support a 2-km links on ICL ports:

- The new credit model does not affect the HA configuration.
- On X6-8 core blades, only a maximum of eight ports can be configured for a 2-km link with all frame sizes at line rate. The remaining eight ports can be configured for a 2-km link with full size frames at line rate.
  - For each port that requires a 2-km line rate for all frame sizes, assign 96 buffers using the `portCfgEportCredits` command.
  - For each port that requires a 2-km line rate for full frame sizes, assign 37 buffers using the `portCfgEportCredits` command.
- On X6-4 core blades, only a maximum of eight ports can be configured for a 2-km link with all frame sizes at line rate. The remaining eight ports can be configured for a 2-km link with full size frames at line rate.
  - For each port that requires a 2-km line rate for all frame sizes, assign 48 buffers using the `portCfgEportCredits` command.
    - For each port that requires a 2-km line rate for full frame sizes, assign 18 buffers using the `portCfgEportCredits` command.
- Due to the 2-km QSFP module limitation, the link failure counter is not reliable during module or cable removal or insertion.

## ICLs between X7 Directors


Each ICL connects the core blades of two Brocade X7 director chassis and provides up to 256G of throughput within a single cable.

The following figures show acceptable cabling configurations for the inter-chassis link (ICL) feature between QSFP ports in core routing blades. The recommended topology is the parallel type where at least four QSFP cables connect between both core routing blades on any two directors, which provides ICL trunking between directors, and ensuring redundancy. Parallel connections between core blades are recommended.

The maximum numbers of directors that you can connect using QSFP ICLs:

- Up to nine Gen 7 directors can be connected in a full-mesh topology.
- Up to nine Gen 7 directors are recommended in a full-mesh topology if connecting a mixture of directors.
- Up to 12 Gen 7 directors can be connected in a core-to-edge topology.

**Figure 28: ICL Cable Connections for Gen 7 director (Sample Configuration)**




1. Chassis 1 – X7-8
2. Chassis 2 – X7-4

To connect multiple Gen 7 directors using ICLs, a minimum of four ICL ports (two on each core routing blade) must be connected between each chassis. The dual connections on each core blade must reside within the same trunk group. If more than four ICL connections are required between a pair of directors, additional ICL connections should be added in pairs (one on each core blade). ISLs and ICLs can co-exist between the same pair of directors if the ISLs and ICLs are on different logical switches.

A Gen 7 director can be connected in a core-edge configuration. The following figure shows two core and four edge devices. Although X7-8 directors are shown in the figure, the devices can be either X7-4 or X7-8 directors. The cabling

scheme should follow the parallel example that is shown in the previous figure. Each line in the example actually represents four cables running between the devices.

**Figure 29: Brocade X7 director Core-Edge ICL Topology**



You can have up to 32 QSFP ports in an X7-8 chassis or 16 QSFP ports in an X7-4 chassis, with up to 4Tb/s ICL bandwidth and support for up to 100 meters over OM4 cable.

The Brocade X7-8 chassis has four port groups on the CR64-8 core blade. The Brocade X7-4 chassis has two port groups on the CR64-4 core blade. Each port group has four QSFP connectors, and each QSFP connector maps to four user ports.

The following table shows the mappings from the numbered QSFP ports on the face of the core blade to the individual FC port numbers as shown by the `slotShow` command.

**Table 99: External Port to slotshow Port Mapping for Core Blades**


| External Port Number | slotShow FC Port Numbers | External Port Number | slotShow FC Port Numbers |
|----------------------|--------------------------|----------------------|--------------------------|
| 0                    | 0 to 3                   | 8                    | 32 to 35                 |
| 1                    | 4 to 7                   | 9                    | 36 to 39                 |
| 2                    | 8 to 11                  | 10                   | 40 to 43                 |
| 3                    | 12 to 15                 | 11                   | 44 to 47                 |
| 4                    | 16 to 19                 | 12                   | 48 to 51                 |
| 5                    | 20 to 23                 | 13                   | 52 to 55                 |
| 6                    | 24 to 27                 | 14                   | 56 to 59                 |
| 7                    | 28 to 31                 | 15                   | 60 to 63                 |

Following are ICL configuration guidelines for trunking bandwidth and high availability:

- A minimum of four ICL ports (two on each core blade) must be connected between each chassis pair.
- The dual connections on each core blade must reside within the same ICL trunk boundary on the core blades.
- If more than four ICL connections are required between a pair of X7 chassis, add ICL connections in pairs, one on each core blade with each additional pair residing within the same trunk boundary.

For high availability, you need at least two ICLs from each core blade. The following illustration shows one X7-8 chassis and one X7-4 connected with full redundancy using four ICL connections:

**Figure 30: Minimum Configuration for Gen 7 ICLs**



1. X7-8 Chassis with two CR64-8 Blades
2. X7-4 Chassis with two CR64-4 Blades

- The maximum number of ICLs between two X7-4 chassis or between an X7-8 and an X7-4 is 16.
- The maximum number of ICLs between two X7-8 chassis is 32.

Because the FSPF routing logic uses only the first 16 paths to come online, only 16 ICLs are utilized. With Virtual Fabrics, however, you can define two logical switches on the chassis and can have 16 ICLs in each.

**NOTE**

You should have a maximum of eight ICLs connected to the same neighboring domain, with a maximum of four ICLs from each core blade.

**NOTE**

The ICLs can connect to both the core blades in the neighboring chassis. The QSFP ICLs need not be cross-connected.

**NOTE**

QSFP ICLs and ISLs in the same logical switch and connected to the same neighboring switch are not supported. This topology restriction is with ICLs and any ISLs that are E\_Ports or VE\_Ports. If Virtual Fabrics is enabled, you can have ICLs and ISLs between a pair of chassis if the ICLs are in a different logical switch than the ISLs.

## ICL Trunking between Brocade X7 Directors

Each core blade in the Brocade X7-8 Director contains Gen 7 ICL trunk groups. Each core blade in a Brocade X7-4 Director contains Gen 7 ICL trunk groups. Each ICL trunk group contains four user ports, one from each QSFP. To establish ICL trunking between platforms in the Brocade X7 Director family, the QSFP cables must be in the same trunk group. Up to four Gen 7 ICL QSFP ports in a trunk group can form a Gen 7 ICL trunk between two CR64-4 or CR64-8 blades.

The following figures mark the 8-port and 4-port ICL trunking boundaries or port groups for Brocade X7 core routing blades.



Figure 31: ICL Trunking between Two CR64-8 Core Routing Blades





**Figure 32: ICL Trunking between CR64-4 and CR64-8 Core Routing Blades**

Ports that belong to the same trunking groups are indicated with the same color border under the ports on the blade faceplate. These colors are also applied to the port map labels on each blade faceplate to indicate ports that belong to the same trunking groups. Although a connection between two QSFPs on one core routing blade and two QSFPs on another blade in a different device are required to form an ICL trunk, you can create trunks that consist of two to four QSFP connections between these blades.

**NOTE**

Fabric OS does not permit a fabric configuration that has E\_Port-to-E\_Port ISLs on port blades and at the same time E\_Port-to-E\_Port connections through ICLs on core blades between a pair of domains on two directors.

The following table summarizes the possible trunking connections between Brocade X6 Directors.

**Table 100: Trunking Connections between Brocade X7 Directors**

| Connecting Blades | 4- to 8-Port ICL Trunks                                                |
|-------------------|------------------------------------------------------------------------|
| CR64-8 and CR64-8 | Supported with QSFPs located within the same trunk group on each blade |
| CR64-8 and CR64-4 | Supported with QSFPs located within the same trunk group on each blade |
| CR64-4 and CR64-4 | Supported with QSFPs located within the same trunk group on each blade |

Refer to the *Brocade X7-8 Director Hardware Installation Guide* or *Brocade X7-4 Director Hardware Installation Guide* for information about port numbering, port trunk groups, and connecting the ICL cables.

## ICLs between an X7 Director and an X6 Director


Each ICL connects the core blade of a Brocade X7 Director chassis with a Brocade X6 Director chassis and provides up to 128G of throughput within a single cable. You must set the port speed to 32G and must use the 32G QSFP to interconnect on the X6 Director.

You can have up to 32 QSFP ports in a Brocade X7-8/Brocade X6-8 chassis, and 16 QSFP ports in a Brocade X7-4/Brocade X6-4 chassis, with up to 4Tb/s of the ICL bandwidth and support for up to 100 meters over OM4 cable.

The 100-meter ICL is supported beginning in Fabric OS 8.0.1, when using 100-meter-capable QSFPs over OM4 cable.

The Brocade X7-8/Brocade X6-8 has four port groups on the CR64-8/CR32-8 core blade. The Brocade X7-4/Brocade X6-4 has two port groups on the CR64-4/CR32-4 core blade. Each port group has four QSFP connectors, and each QSFP connector maps to four user ports. Refer to the hardware reference manuals for details about the port groups. Notice the port numbers that form a port group in Gen 7 versus Gen 6 CR blades, as represented in the following illustration. Every four ports that have the same color representation form a port group.

**Figure 33: Port Groups in Gen 7 CR Blades versus Gen 6 CR Blades**



The following are ICL configuration guidelines for trunking bandwidth and high availability:

- ICLs must be installed in groups of two. Each pair of ICLs must be in the same port group.
- The recommended minimum number of ICLs between a Brocade X7 Director and a Brocade X6 Director is four. Add ICLs in increments of two, one on each core blade.
- For high availability, you should have at least two ICLs from each core blade.
- The maximum number of ICLs between two X7-4/X6-4 chassis or between a Brocade X7-8/X6-8 chassis and a Brocade X7-4/Brocade X6-4 chassis is 16.
- The maximum number of ICLs between two X7-8/X6-8 chassis is 32.

Because the FSPF routing logic uses only the first 16 paths to come online, only 16 ICLs are utilized. With Virtual Fabrics, however, you can define two logical switches on the chassis and can have 16 ICLs in each.

**NOTE**

You should have a maximum of eight ICLs connected to the same neighboring domain, with a maximum of four ICLs from each core blade.

**NOTE**

QSFP ICLs and ISLs in the same logical switch and connected to the same neighboring switch are not supported. This is a topology restriction with 16G ICLs and any ISLs that are E\_Ports or VE\_Ports. If Virtual Fabrics is enabled, you can have ICLs and ISLs between a pair of X7/X6 Director chassis if the ICLs are in a different logical switch than the ISLs.

## ICL Trunking between an X7 Director and an X6 Director

ICL trunks form automatically, but additional licenses might be required for enabling all ICL ports or for larger ICL configurations. For more information about ICL licensing options, refer to the *Brocade Fabric OS Software Licensing User Guide*.

**NOTE**

Each QSFP cable has four ports, each terminating on a different ASIC. These ports cannot form a trunk with each other but can form trunks with corresponding ports on another QSFP.

Each core blade in the X7-8 and X6-8 contains 16 ICL trunk groups. Each core blade in the Brocade X7-4 and X6-4 contains 8 ICL trunk groups. Each ICL trunk group contains 4 user ports, one from each QSFP.

To establish ICL trunking between platforms in the Brocade X7 Director family and Brocade X6 Director family, the QSFP cables must be in the same trunk group. Each ICL that connects the core blades between the Gen 7 and Gen 6 platforms provides up to 128G of throughput within a single cable.

**NOTE**



When you connect X7 Director ICL ports with X6 Director ICL ports, you must use either the Gen 7 4x32G QSFPs or the Gen 6 4x32G QSFPs on the Gen 7 ICL ports, and you must set the port speed to 32G on the X7 Director.

The following tables summarize the possible trunking connections between an X7 Director and an X6 Director:

**Table 101: Trunking Connections between an X7 Director and an X6 Director**

| Connecting Blades | 4-Port to 8-Port ICL Trunks                                            |
|-------------------|------------------------------------------------------------------------|
| CR64-8 and CR32-8 | Supported with QSFPs located within the same trunk group on each blade |
| CR64-4 and CR32-8 | Supported with QSFPs located within the same trunk group on each blade |
| CR64-8 and CR32-4 | Supported with QSFPs located within the same trunk group on each blade |
| CR64-4 and CR32-4 | Supported with QSFPs located within the same trunk group on each blade |

The following figures are an example of 8-port and 4-port ICL trunks between an X7 Director and an X6 Director.

**Figure 34: 8-Port ICL Trunking between an X7 Director and an X6 Director****Figure 35: 4-Port ICL Trunking between an X7 Director and an X6 Director**

Refer to the specific hardware reference manuals for information about port numbering, port trunk groups, and connecting the ICL cables.

## ICLs between a Gen 7 Director and Gen 5 Directors

Each ICL connects the core blade of a Brocade DCX 8510 chassis with an X7 Director chassis and provides up to 64G of throughput within a single cable. You must set the port speed to 16G and use the 16G QSFP to interconnect on the X7 Director.

The 100-meter ICL is supported beginning in Fabric OS 8.0.1, when using 100-meter-capable QSFPs over OM4 cable only.

The following are ICL configuration guidelines for trunking bandwidth and high availability:

- ICLs must be installed in groups of two. Each pair of ICLs must be in the same port group.
- The recommended minimum number of ICLs between a DCX 8510 chassis and an X7 Director chassis is four. Add ICLs in increments of two, one on each core blade.
  - For high availability, you should have at least two ICLs from each core blade.
- The maximum number of ICLs between two DCX 8510-4/X7-4 chassis or between a Brocade DCX 8510-8/Brocade X6-8 chassis and a Brocade DCX 8510-4/Brocade X6-4 chassis is 16.
- The maximum number of ICLs between two DCX 8510-8/X7-8 chassis is 32.

### NOTE

QSFP ICLs and ISLs in the same logical switch and connected to the same neighboring switch are not supported. This is a topology restriction with 16G ICLs and any ISLs that are E\_Ports or VE\_Ports. If Virtual Fabrics is enabled, you can have ICLs and ISLs between a pair of DCX 8510 Director/X7 Director chassis if the ICLs are in a different logical switch than the ISLs.

## Virtual Fabrics Considerations for ICLs

For Virtual Fabrics configured on supported Gen 7 and Gen 6 platforms, the ICL ports can be split across the logical switch, base switch, and default switch. The triangular topology requirement must be met for each fabric individually. In addition, the individual QSFP ports of Gen 6 platforms are split into the logical switch, base switch, or default switch.

When deploying ICLs between X7-4 and X7-8 directors using Gen 7 QSFPs, the individual ports on the QSFP cannot be split across different logical switches and must all reside in the same logical switch. This limitation does not apply when connecting Gen 7 directors with Gen 6 QSFPs for backward compatibility with Gen 6 directors.

The following restrictions apply when the Allow XISL Use switch attribute is enabled:

- When one (or more) ports on a QSFP is in a base switch, another port on that QSFP cannot be moved to a logical switch.
- When one (or more) ports on a QSFP is in a logical switch, another port on that QSFP cannot be moved to a base switch.

In addition, be aware of the following when configuring ICLs and logical switches:

- When a logical switch uses one ICL port to reach Domain1 and another ICL port to reach Domain2, both of those ICL ports must be from the same ASIC. This restriction allows local switching for the frames that are transmitted between these two domains.
- Make sure that any single domain is reachable through both core blades for redundancy. Therefore, a minimum of one ICL port from each core blade must be moved to the logical switch.

## Supported Topologies for ICL Connections

You can connect the Brocade X7 Directors, X6 Directors, or both, in a mesh topology or a core-edge topology. Following are the maximum numbers of directors that you can connect using ICLs:

- Up to nine Gen 7 or Gen 6 directors can be connected in a full-mesh topology.
- Up to nine Gen 7 or Gen 6 directors are recommended in a full-mesh topology if connecting a mixture of directors.
- Up to 12 Gen 7 or Gen 6 directors can be connected in a core-to-edge topology.


Refer to the *Brocade SAN Scalability Guidelines* for details about maximum topology configurations.

## Mesh Topology

You can connect the Brocade directors in a mesh topology, in which every chassis is connected to every other chassis. The mesh topology can consist of all homogeneous chassis or a mix of Gen 7 and Gen 6 chassis that have the same slot count.

A simple form of the mesh topology is the triangular topology (shown in the following illustration). The triangular topology is supported by three Brocade director chassis. The chassis for each topology must all be from the same family.


- Brocade X7 Director family (Brocade X7-8 or Brocade X7-4)
- Brocade X6 Director family (Brocade X6-8 or Brocade X6-4)

**Figure 36: Three Brocade X6-8 Chassis in a Full-Mesh, Triangular ICL Topology**

During an ICL break in the triangular topology, the chassis that has the connections of the other two is the main chassis. Any error messages relating to a break in the topology appear in the RASLog of the main chassis.

Another form is the full nine-mesh topology that is shown in the following figure. This topology is supported by a mix of Brocade X7-4 and Brocade X7-8 Directors, a mix of Brocade X6-4 and Brocade X6-8 Directors, or a mix of Gen 7 and Gen 6 Directors that have the same slot count.




**Figure 37: Full Nine-Mesh Topology**

## Core-Edge Topology

You can connect Brocade X7 Directors or Brocade X6 Directors in a core-edge topology. For example, the following figure shows six chassis that are connected in a core-edge topology (four edges and two cores). Each chassis can be either a Brocade X7-4, a Brocade X7-8, a Brocade X6-4, or a Brocade X6-8. You can have up to eight edges using Brocade X7-8 or Brocade X6-8 cores or up to four edges using Brocade X7-4 or Brocade X6-4 cores.

Each line in the following figure represents four QSFP cables.

**Figure 38: Illustration of an ICL Core-Edge Topology Using Brocade Chassis Platforms**

## Managing Long-Distance Fabrics

---

This section discusses on how to manage long-distance SAN fabrics. The most effective configuration for implementing long-distance SAN fabrics is to deploy Fibre Channel switches at each location in the SAN.

Each switch handles local interconnectivity and multiplexes traffic across a long-distance dark fiber or wave-length division multiplexing (WDM) links. The Brocade Extended Fabrics software enables SAN management over long distances.

Brocade Extended Fabrics is an optional licensed feature for the Brocade SAN deployment over distances beyond 10 km. A Brocade Extended Fabrics license is required before you can implement long-distance dynamic (LD) and long-distance static (LS) distance levels. The LD and LS settings are necessary to achieve maximum performance results over inter-switch links (ISLs/IFLs) that are greater than 10 km.

For details about obtaining and installing licensed features, refer to the *Brocade Fabric OS Software Licensing User Guide*.

The Extended Fabrics feature enables the following functionality:

- Fabric interconnectivity over Fibre Channel at longer distances  
ISLs/IFLs can use long-distance dark fiber connections to transfer data. Wavelength-division multiplexing, such as dense wavelength-division multiplexing (DWDM), coarse wavelength-division multiplexing (CWDM), and time-division multiplexing (TDM), can increase the capacity of the links. As the Fibre Channel speeds increase, the maximum distance decreases for each switch.  
The Extended Fabrics feature extends the distance the ISLs/IFLs can reach over an extended fiber. This extension is accomplished by providing enough buffer credits on each side of the link to compensate for latency introduced by the extended distance.
- Simplified management over distance  
Each device attached to the SAN appears as a local device, which simplifies the deployment and administration.
- Optimized switch buffering  
When Extended Fabrics is installed on gateway switches (with E\_Port connectivity from one switch to another), the ISLs/IFLs (E\_Ports) are configured with a large pool of buffer credits. The enhanced switch buffers help ensure that data transfer can occur at the near-full bandwidth to use the connection over the extended links efficiently. This efficiency ensures the highest possible performance on ISLs/IFLs.

## Extended Fabrics Device Limitations

Do not use the FC8-64 and FC16-64 port blades for long distance because of their limited buffers. These blades do not support long-wavelength (LWL) fiber optics and only support limited distance. However, you can use the `portCfgLongDistance` command to reserve frame buffers for the ports intended to be used in long-distance mode through DWDM.

A limited number of reserved buffers are used for long distance for each blade. If some ports are configured in long-distance mode and have buffers reserved for them, insufficient buffers may remain for the other ports. In this case, some of the remaining ports may come up in degraded mode.

## Long-Distance Link Modes

Use the `portCfgLongDistance` command to support long-distance links and to allocate sufficient numbers of full-size frame buffers on a specific port. Changes made by this command are persistent across switch reboots and power cycles.

The `portCfgLongDistance` command supports the following long-distance link modes:

- Normal Mode (L0) — L0 is the normal (default) mode for an E\_Port. It configures the E\_Port as a standard (not long-distance) ISL. A total of 20 full-size frame buffers are reserved for data traffic, regardless of the E\_Port's operating speed. The maximum supported link distance is up to 5-km at 2G, up to 2-km at 4G, and up to 1-km at 8G, 10G, and 16G.
- Extended Mode (LE) — LE configures the distance for an E\_Port when that distance is greater than 5-km and up to 10-km. LE does not require an Extended Fabrics license. The baseline for the buffer credit calculation is one buffer credit per km at 2G. This allocation yields the following values for 10 -km:
  - 10 buffer credits per port at 2G
  - 20 buffer credits per port at 4G
  - 40 buffer credits per port at 8G
  - 50 buffer credits per port at 10G
  - 80 buffer credits per port at 16G
  - 160 buffer credits per port at 32G
- Dynamic Mode (LD) — LD calculates buffer credits based on the distance measured during port initialization. Brocade switches use a proprietary algorithm to estimate distance across an ISL. The estimated distance determines the buffer credits required in LD (dynamic) extended link mode based on a maximum Fibre Channel payload size of 2,112 bytes. You can place an upper limit on the calculation by providing a *desired\_distance* value. Fabric OS confines user entries to no larger than what it has estimated the distance to be. When the measured distance is more than the specified desired distance, the desired distance (the smaller value) is used in the calculation.
- Static Mode (LS) — LS calculates a static number of buffer credits based only on a user-defined *desired\_distance* value. LS mode also assumes that all FC payloads are 2,112 bytes. Specify LS mode to configure a static long-distance link with a fixed buffer allocation greater than 10-km.

## Configuring an Extended ISL

Before configuring an extended ISL, ensure that the following conditions are met:

- The ports on both ends of the ISL are operating at the same port speed and can be configured for the same *distance\_level* without compromising local switch performance.

### NOTE

A long-distance link also can be configured to be part of a trunk group. Two or more long-distance links in a port group form a trunk group when they are configured for the same speed and distance, and their link distances are nearly equal. For information on trunking concepts and configurations, refer to [Managing Trunking Connections](#).

- Only qualified Brocade SFP transceivers are used. Only Brocade-branded or certain Brocade-qualified SFP transceivers are supported.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `switchdisable` command.
3. Enter the `configure` command to set the switch fabric-wide configurations.

The following table shows the fabric-wide settings of fabric-wide configuration.

**Table 102: Fabric-Wide Settings**

| Field   | Type   | Default | Range                 |
|---------|--------|---------|-----------------------|
| Domain  | Number | 1       | Varies                |
| R_A_TOV | Number | 10000   | E_D_TOV * 2 to 120000 |
| E_D_TOV | Number | 2000    | 1000 to R_A_TOV/2     |
| WAN_TOV | Number | 0       | 0 to R_A_TOV/4        |

| Field    | Type   | Default | Range   |
|----------|--------|---------|---------|
| MAX_HOPS | Number | 7       | 7 to 19 |

4. Enter the `portcfglongdistance` command.

```
portcfglongdistance[slot/] port [distance_level] [vc_translation_link_init
[-distance desired_distance]
```

5. Both the local and remote extended ISL ports must be configured to the same `distance_level`. When the connection is initiated, the fabric will reconfigure.

The following example configures slot 1, port 2 to support a 100-km link in LS mode and to use the extended link initialization sequence.

```
switch:admin> portcfglongdistance 1/2 LS 1 -distance 100

Reserved Buffers = 406
Warning: port may be reserving more credits depending on port speed.
switch:admin> portshow 1/2
portName:
portHealth: OFFLINE
Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x1 PRESENT U_PORT
portType: 17.0
portState: 2 Offline
Protocol: FC
portPhys: 2 No_Module
portScn: 0
port generation number: 0
portId: 010200
portIfId: 4312003b
portWwn: 20:02:00:05:1e:94:0f:00
portWwn of device(s) connected:
Distance: static (desired = 100 Km)
portSpeed: N8Gbps
LE domain: 0
FC Fastwrite: OFF
Interrupts: 0 Link_failure: 0 Frjt: 0
Unknown: 0 Loss_of_sync: 0 Fbsy: 0
Lli: 0 Loss_of_sig: 3
Proc_rqrd: 5 Protocol_err: 0
Timed_out: 0 Invalid_word: 0
Rx_flushed: 0 Invalid_crc: 0
Tx_unavail: 0 Delim_err: 0
Free_buffer: 0 Address_err: 0
Overrun: 0 Lr_in: 0
Suspended: 0 Lr_out: 0
Parity_err: 0 Ols_in: 0
2_parity_err: 0 Ols_out: 0
CMI_bus_err: 0
```

## Enabling Long Distance When Connecting to TDM Devices

Use this procedure when connecting to time-division multiplexing (TDM) devices and your Brocade switch has QoS and buffer credit recovery enabled.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Disable QoS.

```
switch:admin> portcfgqos --disable [slot/] port
```

If you do not disable QoS, after the second or third link reset (LR), ARB fill words display.

3. Disable buffer credit recovery. Buffer credit recovery is not compatible with the IDLE mode. If you do not disable buffer credit recovery, it continues to perform a link reset.

```
switch:admin> portcfgcreditrecovery --disable [slot/] port
```

4. Configure the port to support long-distance links.

```
switch:admin> portcfglongdistance [slot/] port,LS,0,-distance 100
```

## Forward Error Correction on Long-Distance Links

Forward error correction (FEC) on user ports is supported for LD and LS long-distance modes. Use the `portCfgLongDistance` command with the `-fecEnable` or `-fecDisable` options to enable or disable FEC, respectively, on a user port. Alternatively, you can use the `portCfgFec` command with the `--enable` or `--disable` option as you would for any regular port.

For additional details about FEC, refer to [Forward Error Correction](#).

### Enabling FEC on a Long-Distance Link

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portCfgLongDistance` command and include the `-fecEnable` option, or issue the `portCfgFec` command with the `--enable` option.
3. Enter the `portCfgFec --show` command to verify the configuration.

```
switch:admin> portcfglongdistance 1/20 LS 1 -distance 122 -fecenable
FEC has been enabled.
Reserved Buffers = 982
Warning: port (132) may be reserving more credits depending on port speed.
switch:admin> portcfgfec --show 1/20
Port: 20
FEC Capable: YES
10G/16G FEC Configured: ON
16G FEC via TTS Configured: OFF
FEC State: Inactive
```

## Disabling FEC on a Long-Distance Link

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `portCfgLongDistance` command and include the `-fecDisable` option, or issue the `portCfgFec` command with the `--disable` option.
3. Enter the `portCfgFec --show` command to verify the configuration.

```
switch:admin> portcfglongdistance 1/20 LS 1 -buffers 500 -fecdisable
FEC has been disabled.
Reserved Buffers = 982
Warning: port (132) may be reserving more credits depending on port speed.
switch:admin> portcfgfec --show 1/20
Port: 20
FEC Capable: YES
10G/16G FEC Configured: OFF
16G FEC via TTS Configured: OFF
FEC State: Inactive
```

## Configuring Long-Distance Support for G730 Switch

From Fabric OS 9.1.0, when a user configures a long-distance port on a G730 switch using `portcfglongdistance` command, all of the following requirements apply to the same port group.

When a long-distance is configured on a G730 switch, the following port types are not supported in the same port group:

- E\_Ports and EX\_Ports
- F\_Port trunking with AG
- F\_Port QoS with AG

When a long-distance is configured on a G730 switch, the following port types are supported in the same port group:

- F\_Ports
- F\_Port trunk with ECD

When a long-distance is configured on a G730 switch, the following conditions must be met before the configuration:

- The port must be changed from auto-negotiate mode to fixed speed. You can set the port to fixed speed using the `portcfgspeed` command.
- If any unsupported port types exist in the port group, then `portcfglongdistance` command prompts with the list of ports and their logical switches.
- If any unsupported port types exist in the port group, you must remove the configuration or disable the port types using the `portdisable` command.
- If the buffers are not available on the port group, you are recommended to change buffers on the other non long-distance ports (including offline ports) using the `portcfgfportbuffers` command.

## Configuration of Long-Distance Support On a G730 Switch

To configure long-distance support on a G730 switch, perform the following steps:

1. Verify the port speed using the `switchshow` command.

```
switch:admin> switchshow
```

```

switchName: sw0

switchType: 189.3

switchState: Online

switchMode: Native

switchRole: Principal

switchDomain: 3

switchId: fffc03

switchWwn: 2001:db8:c4:f5:7c:02:2c:18

zoning: ON (cfg_vk)

switchBeacon: OFF

FC Router: OFF

Fabric Name: \1234

HIF Mode: OFF

Allow XISL Use: OFF

LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]

```

| Index | Port | Address | Media | Speed | State     | Proto |                                            |
|-------|------|---------|-------|-------|-----------|-------|--------------------------------------------|
| 0     | 0    | 030000  | id    | N32   | Online    | FC    | E-Port 2001:db8:c4:f5:7c:02:21:18 "sw0"    |
| 1     | 1    | 030100  | id    | N32   | Online    | FC    | E-Port 2001:db8:c4:f5:7c:02:21:18 "sw0"    |
| 2     | 2    | 030200  | --    | N64   | No_Module | FC    |                                            |
| 3     | 3    | 030300  | --    | N64   | No_Module | FC    |                                            |
| 4     | 4    | 030400  | --    | N64   | No_Module | FC    |                                            |
| 5     | 5    | 030500  | --    | N64   | No_Module | FC    |                                            |
| 6     | 6    | 030600  | --    | N64   | No_Module | FC    |                                            |
| 7     | 7    | 030700  | --    | N64   | No_Module | FC    |                                            |
| 8     | 8    | 030800  | id    | N16   | Online    | FC    | F-Port 2001:db8:0f:00:27:f8:8f:64:25 (AoQ) |
| 9     | 9    | 030900  | id    | N16   | Online    | FC    | F-Port 1 N Port + 1 NPIV public (AoQ)      |
| 10    | 10   | 030a00  | id    | N16   | Online    | FC    | F-Port 1 N Port + 1 NPIV public (AoQ)      |
| 11    | 11   | 030b00  | id    | N64   | No_Light  | FC    |                                            |
| 12    | 12   | 030c00  | --    | N64   | No_Module | FC    |                                            |
| 13    | 13   | 030d00  | --    | N64   | No_Module | FC    |                                            |
| 15    | 15   | 030f00  | --    | N64   | No_Module | FC    |                                            |



2. Configure the long-distance in port 0 using the `portcfglongdistance` command. You receive the warning messages and you must meet the conditions specified in the warning messages.

```
switch:admin> portcfglongdistance 0 LS 1 -distance 25
Operation Failed:
Please set the port to fixed speed via portcfgspeed.
Please disable the following port(s) [10, 9, 8, 1] configured in the Logical Switch (128)

switch:admin> portcfgspeed 0 32

switch:admin> portcfglongdistance 0 LS 1 -distance 25

Operation Failed:
Please disable the following port(s) [1, 8, 9, 10] configured in the Logical Switch (128)

switch:admin> portdisable 1
switch:admin> portdisable 8-10

switch:admin> portcfglongdistance 0 LS 1 -distance 25

Warning: one or more E_Ports may come up in buffer limited mode.
```

3. Verify the long-distance configuration in port 0 using the `portbuffershow` command.

```
switch:admin> portbuffershow 0
```

| User | Port | Lx   | Max/Resv | Avg Buffer Usage & FrameSize | Buffer Needed | Link  | Remaining |
|------|------|------|----------|------------------------------|---------------|-------|-----------|
| Port | Type | Mode | Buffers  | Tx                           | Rx            | Usage | Buffers   |
| 0    | E    | LS   | 396      | 6( 136)                      | 17( 356)      | 396   | 396       |
| 1    | -    | -    | 20       | - ( - )                      | - ( - )       | 0     | -         |
| 2    | -    | -    | 20       | - ( - )                      | - ( - )       | 0     | -         |
| 3    | -    | -    | 20       | - ( - )                      | - ( - )       | 0     | -         |
| 4    | -    | -    | 20       | - ( - )                      | - ( - )       | 0     | -         |
| 5    | -    | -    | 20       | - ( - )                      | - ( - )       | 0     | -         |
| 6    | -    | -    | 20       | - ( - )                      | - ( - )       | 0     | -         |
| 7    | -    | -    | 20       | - ( - )                      | - ( - )       | 0     | -         |

|    |   |    |         |         |   |   |   |     |
|----|---|----|---------|---------|---|---|---|-----|
| 8  | - | 20 | - ( - ) | - ( - ) | 0 | - | - |     |
| 9  | - | 20 | - ( - ) | - ( - ) | 0 | - | - |     |
| 10 | - | 20 | - ( - ) | - ( - ) | 0 | - | - |     |
| 11 | - | 20 | - ( - ) | - ( - ) | 0 | - | - |     |
| 12 | - | 20 | - ( - ) | - ( - ) | 0 | - | - |     |
| 13 | - | 20 | - ( - ) | - ( - ) | 0 | - | - |     |
| 14 | - | 20 | - ( - ) | - ( - ) | 0 | - | - |     |
| 15 | - | 20 | - ( - ) | - ( - ) | 0 | - | - | 282 |

#### 4. Configure the long-distance in port 1 using the `portcfglongdistance` command.

```
switch:admin> portcfglongdistance 1 LS 1 -distance 25
```

Operation Failed: Configure average buffers (13) in the following ports

[ 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 15] configured in the Logical Switch (128)

[ 14 ] configured in the Logical Switch (10)

A warning message appears to configure average buffers for the remaining ports. You can change buffers on the other non long-distance ports (including offline ports) using the `portcfgfportbuffers` command.

#### 5. Configure the average buffers in all the ports using the `portcfgfportbuffers` command.

```
switch:admin> portcfgfportbuffers --enable 2 13
```

Success.

```
switch:admin> portcfgfportbuffers --enable 3 13
```

Success.

#### 6. Verify the reserved buffers for all the ports using the `portbuffersshow` command.

```
switch:admin> portbuffersshow 0
```

| User | Port | Lx   | Max/Resv | Avg Buffer Usage & FrameSize |         | Buffer Needed |         | Link     | Remaining |
|------|------|------|----------|------------------------------|---------|---------------|---------|----------|-----------|
| Port | Type | Mode | Buffers  | Tx                           | Rx      | Usage         | Buffers | Distance | Buffers   |
| ---- | ---- | ---- | ----     | -----                        | -----   | -----         | -----   | -----    | -----     |
| 0    | E    | LS   | 396      | - ( - )                      | - ( - ) | 396           | 396     | 25km     |           |
| 1    |      | -    | 20       | - ( - )                      | - ( - ) | 0             | -       | -        |           |

|    |   |    |         |         |   |   |   |     |
|----|---|----|---------|---------|---|---|---|-----|
| 2  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 3  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 4  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 5  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 6  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 7  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 8  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 9  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 10 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 11 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 12 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 13 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 14 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |     |
| 15 | - | 13 | - ( - ) | - ( - ) | 0 | - | - | 380 |

## 7. Configure the long-distance in port 1 using the `portcfglongdistance` command.

```
switch:admin> portcfglongdistance 1 LS 1 -distance 25
```

Warning: one or more E\_Ports may come up in buffer limited mode.

## 8. Verify the long-distance ports using the `portbuffershow` command. You can also use `portbuffershow -ls <user port number>` command to list all the user ports that are part of the same port group.

```
switch:admin> portbuffershow 1
```

| User | Port | Lx   | Max/Resv | Avg Buffer Usage & FrameSize |         | Buffer Needed |         | Link     | Remaining |
|------|------|------|----------|------------------------------|---------|---------------|---------|----------|-----------|
| Port | Type | Mode | Buffers  | Tx                           | Rx      | Usage         | Buffers | Distance | Buffers   |
| ---- | ---- | ---- | ----     | -----                        | -----   | -----         | -----   | -----    | -----     |
| 0    | E    | LS   | 396      | - ( - )                      | - ( - ) | 396           | 396     | 25km     |           |
| 1    | E    | LS   | 396      | - ( - )                      | - ( - ) | 396           | 396     | 25km     |           |
| 2    | -    |      | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |           |

|    |   |    |         |         |   |   |   |   |
|----|---|----|---------|---------|---|---|---|---|
| 3  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 4  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 5  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 6  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 7  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 8  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 9  | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 10 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 11 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 12 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 13 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 14 | - | 13 | - ( - ) | - ( - ) | 0 | - | - |   |
| 15 | - | 13 | - ( - ) | - ( - ) | 0 | - | - | 4 |

```
switch:admin> portbuffershow -ls 1
```

| User      | LS   | Port | Lx   | Max/Resv | Avg Buffer Usage & FrameSize |         | Buffer Needed |         | Link     |
|-----------|------|------|------|----------|------------------------------|---------|---------------|---------|----------|
| Remaining |      |      |      |          |                              |         |               |         |          |
| Port      | FID  | Type | Mode | Buffers  | Tx                           | Rx      | Usage         | Buffers | Distance |
| Buffers   |      |      |      |          |                              |         |               |         |          |
| ----      | ---- | ---- | ---- | -----    | -----                        | -----   | -----         | -----   | -----    |
| 0         | 11   |      | LS   | 396      | - ( - )                      | - ( - ) | 0             | -       | -        |
| 1         | 11   |      | LS   | 396      | - ( - )                      | - ( - ) | 0             | -       | -        |
| 2         | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 3         | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 4         | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 5         | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 6         | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 7         | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 8         | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 9         | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 10        | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 11        | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 12        | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 13        | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 14        | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 15        | 11   |      | -    | 13       | - ( - )                      | - ( - ) | 0             | -       | -        |
| 4         |      |      |      |          |                              |         |               |         |          |

When a long distance port is already configured in a particular port group and then an unsupported port type ports are brought in the same port group then those unsupported port types are disabled and `switchshow` shows these ports as disabled along with the following reason.

```
8 8 030800 id N64 No_Sync FC Disabled (Long distance ports configured on this chip)
9 9 030900 id N64 No_Sync FC Disabled (Long distance ports configured on this chip)
10 10 030a00 id N64 No_Sync FC Disabled (Long distance ports configured on this chip)
```

The following RASLog SWCH-1007 appears in the above condition:

```
2021/08/31-04:52:30:538578 (GMT), [SWCH-1007], 2/2, FID 128, WARNING, sw0, Switch port 8 disabled due to
"Long distance ports configured on this chip"., OID:0x42000000, SPOID:0x42000000, switch_obj.c, line: 6018,
comp:SWITCH_0, ltime:2021/08/31-04:52:30:537422
```

**NOTE**

For long distance configuration using 64G speeds beyond 10KM distance, you are recommended to use the G720 or X7 platforms.

## Configuring NVMe

---

This section discusses in-depth details on the non-volatile memory express (NVMe). Flash and Solid-State Devices (SSDs) are a type of non-volatile memory (NVM). NVMe is a logical device interface specification for accessing NVM storage media that are attached with a PCI Express (PCIe) bus, which removes SCSI from the I/O stack.

By its design, NVMe allows the host hardware and software to utilize the parallelism found in SSDs. As a result, NVMe reduces I/O overhead and brings performance improvements to the logical device interfaces. These improvements include multiple command queues and reduced latency. By comparison, SCSI interfaces have a single processing queue and are designed to support the hard disk drive (rotating) media which has several challenges that do not exist in SSDs.


The high performance and reduced cost of SSDs mean that they are increasingly incorporated in data center strategies. SSD media can deliver data with 100  $\mu$ s to 200  $\mu$ s of latency; however, the protocol overhead that is associated with SCSI can significantly reduce the effective speed of SSDs, particularly for latency-sensitive applications that require high I/O, and low delay. Next-generation SSD technologies have access latencies of 10  $\mu$ s or less, making high I/O and low delay even more important. By using NVMe, the protocol overhead of SCSI is bypassed.

The NVMe interface is designed with following key attributes:

- Support for up to 64K I/O queues with minimal command overhead
- Each I/O queue supports 64K I/O operations
- Each I/O queue is designed for simultaneous multi-threaded processing
- NVMe protocol enables hardware automated queues
- NVMe commands and structures are transferred end-to-end
- The NVMe protocol can be transported across multiple network fabric types

From FOS 9.1.0, Cisco UCS NVMe host over FC or FC-LAG uplink to a FOS switch is supported.

NVMe over Fibre Channel runs on existing Gen 7 and Gen 6 fabrics and can coexist with a SCSI-based traffic in both cases. It is seamlessly integrated with Gen 6 fabrics IO Insight features on the Brocade G630 Switch and Brocade FC32-64 Port Blade.

**Figure 39: NVMe Fabric Support**

The Fibre Channel standard is a link level network transport that supports multiple protocols. Fibre Channel only needs new upper layer protocols (ULPs) to support NVMe over FC. The NVMe protocol is defined by the FC standards committee. Support for new ULPs needed by NVMe needs no changes to FC switching logic, which provides backward compatibility with existing FC switches. However, to support NVMe ULPs in the FC network, new OS, and hypervisor device drivers are required along with new Gen 6 HBAs.

Additional information and background on NVMe is available at <http://www.nvmexpress.org/>.

## NVMe over Fibre Channel


In NVMe over FC, the FC acts as the transport protocol for the NVMe protocol. NVMe frames are encapsulated in FC and transmitted over the SAN. The following table shows the mappings between the NVMe protocol and the FC protocol. NVMe has a simple set of thirteen commands.

| NVMe Entity                  | FC Mapping                                                          |
|------------------------------|---------------------------------------------------------------------|
| Namespace / NSID             | LUN                                                                 |
| I/Os                         | Exchanges                                                           |
| Command                      | Command routing-control (RCTL) = 6                                  |
| Status                       | Status RCTL = 7                                                     |
| Extended/intermediate status | E_Status RCTL = 7                                                   |
| NVMe subsystem reset         | Unsupported                                                         |
| No Recovery, only retries    | Sequence retransmission request (SRR) / read exchange concise (REC) |
| Association creation         | Extended link service (ELS)                                         |
| IO Q creation                | ELS                                                                 |

| NVMe Entity | FC Mapping |
|-------------|------------|
| Controller  | Target     |
| Namespace   | LUN        |
| Rd          | Read       |
| Wr          | Write      |

The following figure shows namespace and namespace IDs in an NVMe subsystem. In an NVMe subsystem, the controllers represents the equivalent of FCP\_SCSI targets, and the namespace represents the LUNs. A namespace can be exported through multiple controllers, which would be the multipath I/O (MPIO) topology. In this case, the Namespace-B is exported through both Controller-1 and Controller-2. The Namespace-B NSID is different for both Controller-1 and Controller-2.

**Figure 40: NVMe Namespace**



## NVMe Name Server Support

When an NVMe device logs in on the fabric, a generic Name Server entry is created. This also updates the additional information about the device registers and Name Server entries. The updated information includes but is not limited to the following:

- FC4 type: NVMe device (0x28)
- Device type: None, host, target, or both host and target
- Discovery services: Absent or present (specific to NVMe devices only)
- Port symbolic name
- Node symbolic name



**NOTE**

This information is displayed in the `nsshow` command as shown in the following example.

```
switch:admin> nsshow
{
 Type Pid COS PortName NodeName SCR
N 010800; 2,3;30:03:50:eb:1a:bb:6c:fd;10:00:50:eb:1a:bb:6c:fd; 0x0000000f
 SCR: Fabric-Detected Nx-Port-Detected Fabric-Name Peer-Zone
FC4s: FCP FC-NVMe Application-Services
 Fabric Port Name: 20:08:c4:f5:7c:00:a3:20
 Permanent Port Name: 30:03:50:eb:1a:bb:6c:fd
 Device type: Physical Initiator
 Port Index: 8
 Share Area: No
 Redirect: No
 Partial: No
 LSAN: No
 Slow Drain Device: No
 Device link speed: 16G
 Connected through AG: No
 Real device behind AG: No
 FCoE: No
 FC4 Features [FCP]: Initiator
 FC4 Features [FC-NVMe]: Initiator Discovery-Service

 The Local Name Server has 1 entries }
switch:admin>
```

## NVMe Devices and Zoning

In a large fabric, NVMe devices can coexist with other types of devices. Device World Wide Name (WWN), Target Driven Zoning (TDZ), and [D, I] zoning are all supported, and they are no different at the switch-level from the existing FCP and FICON device support. You can create zones for the NVMe devices to isolate them for the following reasons:

- To keep NVMe devices separate from other devices
- To form multiple groups of NVMe devices
- To form select groups of NVMe devices and FCP devices

The administrator can create zones by using command-line commands, SANnav, Web Tools, or Brocade Network Advisor. Alternatively, the NVMe device itself can use TDZ mode to create a peer zone. Refer to [Target-Driven Zoning](#) for more information.

# Administering Flow Vision

---

Flow Vision is a Brocade Fibre Channel SAN network diagnostic tool that is supported on most Fabric OS and SANnav platforms.

Flow Vision provides a comprehensive view of fabric traffic flows, providing in-depth insight and the ability to analyze performance, bottlenecks, bandwidth utilization, and other fabric behaviors. Flow Vision also provides a test flow generation capability that you can use to pretest a SAN infrastructure for robustness. This test flow generation capability is also useful for testing the internal connections on a switch before deploying the switch into a production environment.

Automatic monitoring provides the SAN administrator with deep visibility at each F\_Port. Brocade's IO Insight capability adds built-in device input/output (I/O) latency and performance instrumentation, providing proactive, nonintrusive, real-time monitoring and alerting of the storage I/O health and performance. This additional visibility delivers deep insights into problems and concerns that affect the ability to maintain service levels.

Brocade's VM Insight capability for Fabric OS 9.1.x allows discovering and monitoring flows between virtual machines (VMs) with greater granularity. This granularity allows you to gather Fibre Channel, SCSI I/O performance, NVMe I/O performance, and I/O latency statistics about individual traffic flows for VM instances that originate from a single N\_Port ID (PID) to destination targets. Using Flow Vision, you can enforce the monitoring of virtualized applications that are running behind a remote host. Data about learned VM flows can help isolate VMs that are being impacted by performance issues, isolate the source of performance issues, and help plan for a VM migration for load balancing or other purposes.

## Flow Vision Features

Flow Vision has three main components: Flow Monitor, Flow Generator, and Flow Mirror.

### Flow Monitor

Flow Monitor provides performance metrics for flows and the gathering of frame statistics and I/O statistics for fabric application flows, including the ability to automatically learn (discover) flows through a specified port. See [Flow Monitor](#) for a complete description and sample use cases.

With Fabric OS 9.1.x, Flow Monitor provides automated and dynamic monitoring on Gen 6 and Gen 7 switches and directors.

### Flow Generator

Flow Generator simulates and generates test-load traffic in specific flows, which allows you to validate hardware components and connectivity and to verify performance. See [Flow Generator](#) for a complete description and sample use cases.

#### **NOTE**

SANnav does not manage Flow Generator.

### Flow Mirror

Flow Mirror can nondisruptively create copies of application flow frames which can be captured for the in-depth analysis of their contents. See [Flow Mirror](#) for a complete description and sample use cases.

#### **NOTE**

SANnav does not manage Flow Mirror.

## Flow Vision Licensing

The Fabric Vision® (FV) license is required to use Flow Vision features and functionality.

Refer to the *Brocade Fabric OS Software Licensing User Guide* for more specific information about licensing and how to check the status or obtain the needed license keys.

## Firmware Upgrade and Downgrade Considerations for Flow Vision

Flow Vision has specific considerations that apply to upgrading or downgrading the firmware on a switch with Flow Vision configured. Both user-defined and system-defined flows support the configuration upload and download procedures.

With Fabric OS 9.1.x, Flow Vision supports only those devices that are listed in [Supported Hardware and Software](#).

Flow Vision functionality is not affected if a switch that is running Fabric OS 9.1.x is connected to a switch that is running Fabric OS 9.0.x or later.

## Firmware Upgrade Considerations

The following considerations apply when you upgrade to Fabric OS 9.1.x or later versions:

- The logical partition information and the state of the automated flow sys\_flow\_monitor across VFs are stored and replayed.
- Port pending I/O values are hardware-driven on all supported platforms.
- Violation monitoring is automatically enabled by default.

## Firmware Downgrade Considerations

The following considerations apply when you downgrade to Fabric OS 9.0.x or earlier versions:

- Flow is replayed on the default switch or the first created logical partition in Fabric OS 9.0.x.
- I/O Exceptions statistics are not supported.
- Only aggregated weighted average statistics are provided for the Pending I/O statistics in pre-Fabric OS 9.1.x releases.
- Real-time violation is not supported and software violations are provided only for port statistics on downgrade to pre-Fabric OS 9.1.x releases.
- Firmware downgrade is blocked if the VMID+ configuration is enabled on any port of the switch or chassis.
- User-defined flows are not supported in pre-Fabric OS 9.0.x versions. Ensure that the user-defined flows are deleted or deactivated before downgrading.

## Configuration File Upload and Download Considerations

When uploading from Fabric OS 9.0.x or downloading to Fabric OS 9.1.x, or, conversely, when uploading from Fabric OS 9.1.x and downloading to Fabric OS 9.0.x, the following considerations apply:

- Partition information and the state of the automated flow sys\_flow\_monitor across VFs are stored and replayed when downloaded on Fabric OS 9.1.x.
- IO Exceptions for SCSI READ and SCSI WRITE commands are supported. NVMe Exceptions are not supported.
- Supported predefined flows and user-defined flows are replayed.
- Violation monitoring is automatically supported when downloaded on Fabric OS 9.1.x.

## Flow Vision Configuration Setup

When a switch goes offline or comes online, Flow Vision reads the configuration files and then deletes, creates, and activates flows.

After a switch goes offline, any active flows at the time it went offline are reactivated when it comes back online. Statistics are reset, and new statistics are generated once the source and destination devices that are defined in the flow are online. This includes predefined flows.

Use the following commands to upload, download, and delete configurations:

- To download a Flow Vision configuration to the switch, use the `configDownload` command.
- To save the Flow Vision configuration, use the `configUpload` command.
- To delete all flows and simulation ports (SIM ports) from a switch, use the `flow --delete all` or `flow --control -simport '*' -disable` or `configDefault` command.

### NOTE

Statistical data for flows is not saved in the configuration database.

## High Availability and Flow Vision

High Availability (HA) preserves the Flow Vision configuration settings only through an HA failover, an HA reboot, or a power cycle and reboot. If a failover occurs, feature-related data (for example, statistics) is not saved, and user-defined or system-defined (predefined) flows are not restored, provided the failover occurs to a version of Fabric OS that does not support these flows.

If a standby command processor (CP) with the down-level code comes online and any flows (active or non-active) are configured, then the HA is out of sync. If a standby CP with the down-level code comes online and no flows (active or non-active) are configured, then HA is in sync, but flow creation fails.

See the following sections for information on how each Flow Vision feature is treated under HA:


- [Flow Generator and High Availability](#)
- [Flow Mirror and High Availability](#)

## Flow Vision Flows

A flow is a set of Fibre Channel (FC) frames or packets that share similar traits, such as an ingress port or egress port identifier or any other data that can be used to differentiate one set of related frames or packets from another. These parameters are specified as part of the `flow` command, and include:

- **Port Parameters:** Also called the Point of Interest, or the location from where the behavior you want to examine is traversing. These consist of an ingress port (ingrport) or an egress port (egrport). Only one can be specified when defining a flow.
- **Frame Parameters:** These include the following: Source Device identification (SID) or Destination Device identification (DID). At least one frame parameter must be present to define a flow.
- **VM-Flow Parameters:** These include the VM entity ID (srceid), destination device ID, and the ingress port ID. Flows are monitored from the HBA N\_Port, through the fabric, to the target.
- **Direction:** A direction is implicitly defined from an ingress port to an egress port, or a source device (srcdev) to a destination device (dstdev). For example, `srcdev=x, dstdev=y` indicates traffic flowing from x to y unless the flow is bi-directional.

The following figure illustrates how the frame and port parameters apply to a flow.

**Figure 41: Frame and Port Parameters**

## Flow Definitions

To define a flow and configure Flow Vision (Mirror and Generator, specifically), you must provide a unique flow name and must specify the flow parameters. These parameters identify the sets of related frames that comprise the flow. You can explicitly define them or allow Flow Vision to learn them through observation.

### NOTE

These flow definitions are stored on the switch on which the flow is created, and are not distributed across the fabric. This means that each switch (logical or physical) knows only its unique flows and does not know what flows exist on other switches.

When viewing a flow, you can specify any combination of the three features (fabinfo, mirror, generator) in the `flow` command.

Once you choose to use an option, the next argument must be correct, irrespective of the order of the operands. For the command `flow --activate sys_gen_all_simports -feature generator`, as `sys_gen_all_simports` is the name of the flow, it must immediately follow `--activate`.

## Flow Definition Parameters and Rules

The rules listed in the following table identify the parameters that can be used to define a flow.

**Table 103: Flow Definition Rules**

| Parameters | Field Names                                                                                                        | Rules                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port       | <ul style="list-style-type: none"> <li>ingrport</li> <li>egrport</li> </ul>                                        | <ul style="list-style-type: none"> <li>One field only must be specified.</li> <li>Values must be explicit.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Frame      | <ul style="list-style-type: none"> <li>srcdev</li> <li>dstdev</li> <li>LUN/<br/>NSID</li> <li>Entity ID</li> </ul> | <ul style="list-style-type: none"> <li>At least one field must be specified.</li> <li>Values for srcdev and dstdev can be explicit or "*" . ("*" indicates learned flows, and must include the quotation marks.)</li> <li>Values for the logical unit number (LUN) or Namespace ID equivalent to LUN in NVMe monitoring. LUN and NSID cannot coexist in a flow with frame parameters such as Source Device Identification (SID) and Destination Device Identification (DID).</li> <li>Values for an identifier associated with an entity (VM). For example, the VM Instance UUID in the case of VMware.</li> </ul> |

The following restrictions apply to the flow definition parameters and rules:

- Available frame monitoring flow classifiers include: `ingrport`, `egrport`, `dstdev`, `srcdev`, `LUN/NSID`, and `Entity ID`.
- NVMe flow is supported only on the Brocade Gen 6 (G630 Switch, FC32-64, and FC32-X7-48 blades) and Gen 7 platforms.
- SCSI and NVMe Flow Monitor and Flow Mirror features are supported only on F\_Ports.
- Latency and performance metrics are not available on the AG mode.
- Latency statistics are not supported on the Fibre Channel Router (FCR).

## User-Defined Flows

User-defined flows are flows created by users using the `create` command to monitor flows. From Fabric OS 9.0.x and later, Flow Monitor no longer supports user-defined flows. However, user-defined flow mirror and flow generator flows are supported.

## Static versus Learned Flows

Flow Vision creates a learned flow when you replace the source device, the destination device, or both devices with an asterisk (\*). This allows you to discover which flows are active on a port without having to explicitly identify all the devices.

The following considerations apply when constructing learning flows:

- Learning is dictated by the parameters that are replaced with an asterisk. For parameters for which learning is not requested, cumulative data is presented.
- When you enter an asterisk as part of the command to indicate a learned flow, you must enclose it in single or double quotes, like this: `"*"`.
- Learning source device (`srcdev`) or destination device (`dstdev`) values are supported on Gen 6 and Gen 7 (Mirror and Generator only) Fiber Channel ports.
- Each Flow Vision feature uses learning as follows:
  - a. Flow Monitor can learn all the source device and destination device pairs passing through the ingress or egress port defined in a flow. See [Learning in Flow Monitor Flows](#) for the additional information.
  - b. Flow Generator can generate traffic to or from every source or destination device that shares the zone with the ingress or egress port defined in a flow. See [Learning in Flow Generator Flows](#) for the additional information.
  - c. Flow Mirror can capture all the source device and destination device pairs passing through the ingress or egress port defined in a flow. See [Learning in Flow Mirror Flows](#) for the additional information.
- Because hardware resources for a learned flow are limited, the number of actual subflows for a flow definition might exceed the permitted number of subflows. Flow Monitor refreshes the subflows when that number reaches or exceeds the limit. See [Learning in Flow Monitor Flows](#) for more information.

## Predefined Flows

Fabric OS 9.1.x supports automated and dynamic monitoring where no user configuration is required. In previous releases, flow monitoring was user-defined.

To accomplish automated monitoring, Flow Vision employs the predefined flow `sys_flow_monitor` that non-intrusively learns flows on the host F\_Ports (including F\_Port trunk and target F\_Ports in supported platforms) of a chassis. This flow monitors a single F\_Port, logical group, or initiator/target and provides frame and I/O metrics on all the application flows on the monitored switch.

The following guidelines apply:

- On activation, the `sys_flow_monitor` flow automatically discovers application flows based on the direction and frame parameters:
  - Direction parameters: `ingress` or `egress`
  - Frame parameters: `source device` and `destination device`

- I/O latency and performance metrics are provided in a single group for all I/O sizes.
- The new flow supports monitoring on all host and target F\_Ports, which is always enforced on the primary port on all platforms.
- Support for the `sys_mon_all_fport` flow has been removed.
- `sys_flow_monitor` supports Show, Activate, Deactivate, Modify, Reset, Control, and VM Enable/Disable CLI configuration parameters. Create and Delete are not supported.

## Resetting Flow Statistics

Flow Vision allows you to clear (reset) the flow statistics record for each feature individually or as a group.

### NOTE

Clearing the Flow Mirror statistics for a flow also clears the mirrored frames.

## Clearing All Flow Statistics

To clear all the statistics for a flow, perform the following steps:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `flow --reset all` command.

You are not asked to confirm this action.

The following example clears all the statistics for the flow named Flow4.

```
switch:admin> flow --reset Flow4 -feature all
```

## Clearing the Statistics for Specific Flow Features

To clear the statistics for specified features of a flow, perform the following steps:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `flow -reset flow_name -feature feature_list` command. Replace `feature_list` with either an individual feature or a comma-separated list of features (for example, "generator,monitor" or "mir,mon").

You are not prompted to confirm this action.

The following example clears only the Flow Generator statistics for the flow named Flow4.

```
switch:admin> flow --reset Flow4 -feature sys_flow_generator
```

## Number of Flows Supported

Flow Vision has both overall limits on the number of flows supported on a switch and separate limits for Flow Generator and Flow Mirror. The following table shows the limits for only the user-defined flows.

**Table 104: Total Flow Count Limits**

| Platform | User-Defined Static and Learning Flows (Combined) | Learned Flows |
|----------|---------------------------------------------------|---------------|
| Director | 512                                               | 2048          |
| Switch   | 128                                               | 512           |

**Table 105: Per-Port Flow Count Limits**

| Platform | User-Defined Static and Learning Flows (Combined) | Learned Flows |
|----------|---------------------------------------------------|---------------|
| Director | 64                                                | 512           |
| Switch   | 64                                                | 128           |

The limits for the user-defined static flows and learning flows include both active and inactive flows. The combined limit applies to all features on a per-port basis. For example, two different learning flows on a given port can each have 64 learned flows. Beyond this, there are limits for each individual feature, as described in the following table. Refer to the individual features for other feature-specific restrictions.

**NOTE**

For predefined flows, the learned flow limits are reduced by 1, because one subflow is used internally for managing the flows.

**Table 106: Feature-Specific Flow Count Restrictions in Flow Vision**

| Flow Types     | Description                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Flow Monitor   | Up to 64 active flows per port, including static flows and learning flows.                                                                                                                                                                                                                                                                                                                                                            |
| Flow Generator | Up to 39 active flows per port for ingress ports and 64 active flows per port for egress ports.<br>The Brocade G630 switch and the Brocade FC32-64 blades support four flows from an ingress port including static flows, learned flows, and <code>sys_gen_all_simports</code> . The <code>sys_gen_all_simports</code> flow sends traffic only to four VCs at a time, but over time, it iterates through all the VCs, four at a time. |
| Flow Mirror    | Multiple instances of the predefined flow <code>sys_analytics_vtap</code> can be active across the Brocade Brocade G630, G620, and G610 switches and Brocade X6-4 and X6-8 directors. However, only one Flow Mirror flow can be active per virtual fabric (VF). No other fixed-port switches support multiple instances of the <code>sys_analytics_vtap</code> flow.                                                                  |

**NOTE**

To ensure that no identical flows exist, each flow is verified when activated; duplicate flows cannot be activated simultaneously. To get around this restriction and create a new "duplicate" flow, you must use the `-noactivate` keyword as part of the `flow --create` command. Refer to the "Creating an inactive flow" section of each feature for details.

## Flow Vision Integration with MAPS

Statistics generated using Flow Vision can be monitored with the Monitoring and Alerting Policy Suite (MAPS) threshold service.

Predefined flow support from flow vision is integrated with the collection configuration (through SANnav).

Refer to the individual features for information on how that feature interacts with MAPS, and the Brocade Fabric OS MAPS User Guide for more details on MAPS in general.



## System Event Handling

Flow Vision handles the following system events.

Specific to `sys_flow_monitor`:

- **SLOT ON:** Start monitoring flows on all the F\_Ports when a slot is plugged in. There is a delay of three polling cycles (~30 seconds) before you start to see the data on the CLI output. The number of flows also depends on the number of resources / total chassis resources count available.
- **SLOT OFF:** Stops monitoring flows on the F\_Ports when a slot is removed.
- **SWITCH ENABLE:** Start monitoring flows on all the F\_Ports on the switch. There is a delay of three polling cycles (~30 seconds) before you start to see the data on the CLI output. The number of flows that are learned for the switch also depends on the number of resources / total chassis resources count available.
- **SWITCH DISABLE:** Stops monitoring the flows on the F\_Ports on that switch. All learned flows related to the switch are deleted.
- **F\_PORT ONLINE:** Start monitoring flows on the F\_Port. There is a delay of three polling cycles (~30 seconds) before you start to see the data on CLI output provided resources are available.
- **F\_PORT OFFLINE:** No action. Flows are aged out in the subsequent aging cycle.
- **PORT ONLINE - PORT TYPE CHANGE:** Stop monitoring if the port type changes to anything other than F\_Port.
- **VM Enable -** Enables the VM mode that allows visibility to the virtualized applications running behind the VM instances. Once you enable the VM mode, individual traffic flows from source to destination target for multiple virtual machine (VM) instances that originate from a single N\_Port ID (PID) are monitored through automated flow monitoring.
- **VM Disable -** Disables the VM mode that allows visibility to the virtualized applications running behind the VM instances.

Specific to Mirror and Generator:

- **FPORT\_CHANGE:** When an F\_Port is changed, provided the resources are available, any flows specifying that port are installed in the ASIC and made active. If the resources are unavailable (for example, a flow currently uses that port), the flow stays deactivated.

An F\_Port trunk has the same trunk index for both the primary and secondary ports (as displayed in the output for `switchshow`, below). To create the same flow definition on the primary and secondary ports when `portIdMode` is set to "index", the domain/Index can be obtained with the `porttrunkarea -show all` command (as shown below). If you create a flow using the trunk index, then depending upon whether the trunk index maps to the primary or the secondary port the flow may or may not be installed. If the secondary port becomes the primary, it is installed.

```
switch:admin> switchshow
switchName: X7-8_113
switchType: 180.0
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 113
switchId: fffc71
switchWwn: 10:00:88:94:71:df:c0:cd
zoning: ON (cfg_Chew_Awing)
switchBeacon: OFF
FC Router: OFF
Fabric Name: StatsMON
HIF Mode: OFF
Allow XISL Use: OFF
LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]
Index Slot Port Address Media Speed State Proto
=====
```

```

 8 4 3 010400 id N8 Online FC F-Port (Trunk master)
 9 4 4 010400 id N8 Online FC F-PORT (Trunk port, master is Port 4)
 96 4 0 716000 -- 16G Online FC SIM-Port 20:60:88:94:71:df:c0:cd
 101 4 5 716500 -- 16G Online FC SIM-Port 20:65:88:94:71:df:c0:cd
 102 4 6 716600 -- 16G Online FC SIM-Port 20:66:88:94:71:df:c0:cd
 103 4 7 716700 -- 16G Online FC SIM-Port 20:67:88:94:71:df:c0:cd
 104 4 8 716800 -- 16G Online FC SIM-Port 20:68:88:94:71:df:c0:cd

```

----- Output truncated -----

```

1804 -1 1804 7106c0 -- -- Offline FCoE
1805 -1 1805 7106c0 -- -- Offline FCoE
1806 -1 1806 7106c0 -- -- Offline FCoE
1807 -1 1807 7106c0 -- -- Offline FCoE

```

```

switch:admin> porttrunkarea --show all
Slot Port Type State Master TI DI

 4 0 -- -- -- 0
 4 1 -- -- -- 1
 4 2 -- -- -- 2
 4 3 -- -- -- 3
 4 4 F-port Master 4 4 4
 4 5 F-port Slave 4 4 5

```

## Flow Vision Limitations

The following limitations apply to Flow Monitor:

- User-defined flow monitors are not supported; only the system-defined `sys_flow_monitor` flow is supported.
- Support for the `sys_flow_monitor` flow is limited to host F\_Ports on Gen 6 devices. The `sys_flow_monitor` flow is supported on host and target F\_Ports on Gen 7 devices.
- When the VM mode is enabled using `sys_flow_monitor` on Brocade G620 Switch and Brocade FC32-48 Port Blade, the following metrics are not available:
  - Throughput (Bytes per sec)
  - Total Bytes
  - Average IO Size
- Brocade Gen 6 systems support only IT flows. Brocade Gen 7 systems support IT and ITL flows.
- If the domain or PID is changed:
  - Learned flows of the previous PID remains idle until it attains the configured aging limit.
  - A newly associated PID is automatically learned.
  - Deactivate and activate the flow to clear all learned flows.
- Enable a persistent PID on the switch and AG to avoid a PID change.
- I/O-size metrics buckets are not supported.
- Device and port ID mode are not supported; user port and PID are the only parameters that can be modified.
- SIM port monitoring is not supported.
- Domain change during a reboot must be handled in SANnav.
- Port-swap functionality is not supported. Therefore, you cannot create flows on ports that are swapped, and you cannot swap ports on which flows are currently defined.

## Identifying All Zoned Devices in a Flow

Flow Vision allows you to identify all the flows (actual and possible) passing through a specified F\_Port, subject to zoning restrictions.

This feature is valid for all Flow Vision flows. If a flow is already defined, you can enter the `flow --show` command and completely define in this command the parameters for the flow definition and include the `-allzoned` parameter. This allows you to see all the zoned devices matching the flow definition without creating an actual flow. The advantage of these techniques is that it does not interfere with any existing flows. One use case for this feature is to assist you in determining a host-to-target ratio or device-to-ISL ratio and keeping such a ratio to a user-determined value for the optimal performance based on standard best practice recommendations for SANs.

The following table lists the combinations of flow definition values (Source ID, Destination ID, and flow terminus) for flows that let you identify devices in a zone.

**Table 107: Supported Flow Definition Values for Zoned Device Identification**

| Source ID       | Destination ID  | Monitorable Port          |
|-----------------|-----------------|---------------------------|
| *               | WWN or PID      | Ingress or Egress (fixed) |
| WWN or PID      | *               | Ingress or Egress (fixed) |
| *               | *               | Ingress or Egress (fixed) |
| * or WWN or PID | Not defined     | Ingress or Egress (fixed) |
| Not defined     | * or WWN or PID | Ingress or Egress (fixed) |
| WWN or PID      | WWN or PID      | Ingress or Egress (fixed) |

## Examples of show Command Output for Zoned Devices

The following example displays all devices that are zoned with the device at port 11 (an F\_Port) as the destination port.

```
switch:admin> flow --show -allzoned -srcdev "*" -dstdev 0x190b00 -egr 11
Active Flows:
```

```

| SrcDev | DstDev |

| 0x21400 0x21e00 0x21f00 0x31000 0x31100 | 0x190b00 |

```

The following example displays all devices that are zoned with the device at port 11 (an F\_Port) as the source port.

```
switch:admin> flow --show -allzoned -srcdev "*" -dstdev "*" -ingrport 11
Active Flows:
```

```

| SrcDev | DstDev |

| 0x190b00 | 0x21400 0x21e00 0x21f00 0x31000 0x31100 |

```

To display all the zoned devices for a flow, enter the `flow --show -allzoned <port_options> <frame_options>` command. Do not specify other flow parameters. For more details, see [Identifying all Zoned Devices in a Flow](#). The following example displays all devices that are zoned with the device at port 11 (an F\_Port) as the destination port.

```
switch:admin> flow --show -allzoned -srcdev "*" -dstdev 0x190b00 -egr 11
```

```
Active Flows:

| SrcDev | DstDev |

| 0x21400 0x21e00 0x21f00 0x31000 0x31100 | 0x190b00 |

```

The following example shows how to create the flow "AllHostMirror" and then to display all the zoned devices in the flow with the `flow --show flow_name -allzoned` command.

```
switch:admin> flow --create AllHostMirror -fe mon -srcdev "*" -dstdev 0x190b00 -egrport 11
switch:admin> flow --show AllHostMirror -allzoned
Active Flows:
```

```

| SrcDev | DstDev |

| 0x21400 0x21e00 0x21f00 0x31000 0x31100 | 0x190b00 |

```

To display only the information for one feature in a flow that contains multiple features, enter the `flow --show <flow_name> -feature <feature_name>` command. For root and static flows, this command shows both the Source ID–Destination ID pair and the cumulative frame count for the ingress or egress port that is specified in the flow definition.

- The following example shows the status of the Flow Generator flow named "fgflow12":

```
switch:admin> flow --show fgflow12 -feature generator
Name : fgflow12 Features: gen(Active) noConfig: Off
Definition: IngrPort(4),DstDev(*)
Flow Generator (Activated):

| SrcDev | DstDev |

| 0x204000 | 0x040700 |

| 0x204000 | 0x040800 |

| 0x204000 | 0x050900 |

| 0x204000 | 0x051000 |

Number of frames generated from IngrPort : 82.21M
```

- The following example displays the summary information that is recorded for the Flow Mirror flow named "fmshow". The LUN column is a learned field in this example.

```
switch:admin> flow --show fmshow -feature mirror
=====
Name : fmshow Features: mir(Activated) noConfig: Off
Definition: IngrPort(2),SrcDev(*)
Flow Mirror (Activated):

| SID(*) | DID(*) | OXID | RXID | SOF | EOF | Frame_type | LUN(*) | Dir | Time-Stamp |
|-----|
| 040200 | 040f00 | 0ca0 | ffff | SOFi3 | EOFt | SCSIRead | 0000 | Rx | Jul 12 06:29:13:639|
```

```

| 040200 | 040f00 | 0bea | 0bel | SOFn3 | EOFn | Data | ---- | Rx | Jul 12
06:29:13:639|
| 040200 | 040f00 | 0f42 | ffff | SOFi3 | EOFt | SCSIWrite| 0000 | Rx | Jul 12
06:29:13:639|
| 040200 | 040f00 | 0dc0 | ffff | SOFi3 | EOFt | SCSIRead | 0000 | Rx | Jul 12
06:29:13:639|
| 040200 | 040f00 | 0592 | ffff | SOFi3 | EOFt | SCSIRead | 0000 | Rx | Jul 12
06:29:13:639|
.....
| 000000 | fffffe | 801e | ffff | SOFi3 | EOFt | FLOGI | ---- | Rx | Jul 12
06:29:45:292|
| 040200 | fffffc | 8024 | ffff | SOFi3 | EOFt | PLOGI | ---- | Rx | Jul 12
06:29:49:411|
| 040200 | fffffd | 8028 | ffff | SOFi3 | EOFt | ELSframe | ---- | Rx | Jul 12
06:29:49:411|
| 040200 | fffc04 | 88b9 | ffff | SOFi3 | EOFt | 01 | ---- | Rx | Jul 12
06:29:51:614|
| 040200 | fffc04 | 802c | ffff | SOFi3 | EOFt | PRLI | ---- | Rx | Jul 12
06:29:51:614
No of Mirrored Frames : 528, No of RX Mirrored Frames : 528, No of TX Mirrored Frames : 0

=====

```

Specifying `fabinfo` as the feature for the `flow --show` command displays a summary topology of all the paths taken by the flows in the flow definition. This includes switch level information and MAPS rule violations for each of the switches the flow passes through. See [Viewing Fabric Information with the Fabric Flow Dashboard](#) for additional information on this feature. Refer to the individual features to see feature-specific output, and to the *Brocade Fabric OS Command Reference Manual* for a complete description of the `flow` command options.

## Limitations on Viewing Zoned Devices

The following limitations apply to viewing zoned devices through Flow Vision flows:

- This feature is supported only on F\_Ports.
- Only primary trunk ports are supported. Secondary ports are not supported.
- The maximum number of flows that can be discovered is the maximum number of subflows supported for a flow. This count is not affected by the maximum number of subflows that is supported across the chassis.
- This feature is not available on switches running in Access Gateway mode.

## Flow Monitor

To provide more effective fabric design and capacity planning, storage administrators require the ability to monitor a complete fabric using predefined flows. Flow Monitor enables you to do that, providing Fibre Channel-level statistics, general Fibre Channel statistics, and I/O statistics (Performance and Latency metrics for the SCSI and NVMe protocols).

With Fabric OS 9.1.x, flow monitoring supports dynamic and automated "real-time" monitoring of all traffic passing through host and target (Gen 7) F\_Ports using any hardware-supported flow parameters. Flow monitoring also lets you define your monitoring flows using combinations of ingress and egress ports, source and destination devices, and frame types to create a flow definition for a specific use case.

The following considerations also apply:

- This new capability requires only moderate use of switch resources and can adapt to the high-scale requirements of the newer hardware.
- The automated Flow Monitor employs a use-case-driven CLI. This means that the interface is hierarchical for ease of use.
- Automated monitoring supports multiple protocols.

## Automated Flow Monitoring

Flow Monitor running Fabric OS 9.1.x supports automated monitoring, which enables you to learn SCSI and NVMe traffic and obtain Read/Write Frame Count and Read/Write Data statistics. A cohesive view of metrics provides you all necessary information in one output.

Fabric OS 9.1.x supports the predefined `sys_flow_monitor` flow to provide automated nonintrusive monitoring. When you upgrade to Fabric OS 9.1.x, the `sys_flow_monitor` flow is activated by default.

Once activated, the flow enables automatic learning for Gen 6 and Gen 7 switches and blades running Fabric OS 9.1.x, and it automatically clears the associated flow statistics and force-deletes all user-defined flows. The predefined flow automatically discovers application flows based on the direction and port parameters. With the direction either ingress or egress, this flow gathers the following details:

- DID
- SID
- Protocol
- LUN/NSID
- Entity ID

Brocade Gen 6 X6-4 and X6-8 directors and Brocade Gen 6 switches support only ingress monitoring for IO Insight and VM Insight. Brocade Gen 7 X7-4 and X7-8 directors and Brocade Gen 6 directors upgraded to Gen 7 support ingress and egress monitoring if they have Brocade Gen 7 blades. Brocade Gen 7 switches also support the ingress and egress monitoring.

The following table lists support for automatic monitoring on the Brocade Gen 6 and Gen 7 platforms.

**Table 108: Automatic Monitoring/Platform Support Matrix**

|                     | Gen 6 Platforms                                                                                                                                                                                                                                                                                                                  | Gen 7 Platforms                                                                                                                         |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Supported Platforms | <ul style="list-style-type: none"> <li>• Brocade G610 Switch</li> <li>• Brocade G620 Switch</li> <li>• Brocade G630 Switch</li> <li>• Brocade 7810 Switch</li> <li>• Brocade FC32-48 Port Blade</li> <li>• Brocade FC32-64 Port Blade</li> <li>• Brocade SX6 Extension Blade</li> <li>• Brocade FC32-X7-48 Port Blade</li> </ul> | <ul style="list-style-type: none"> <li>• Brocade G720 Switch</li> <li>• Brocade G730 Switch</li> <li>• Brocade FC64-48 Blade</li> </ul> |
| ASIC                | <ul style="list-style-type: none"> <li>• Brocade FC32-48 Port Blade</li> <li>• Brocade FC32-64 Port Blade</li> <li>• Brocade G630 Switch</li> </ul>                                                                                                                                                                              | <ul style="list-style-type: none"> <li>• Brocade G720 64-Port Switch</li> <li>• ASIC used in the Gen 7 platforms</li> </ul>             |

The following table lists the details that are gathered by the flow type for all monitored flows between host and target.

**Table 109: Flow Definition per Flow Type**

| Flow Type     | Definition                                                                                        |
|---------------|---------------------------------------------------------------------------------------------------|
| IT            | Destination device PID, direction of flow, protocol, and source device PID                        |
| I (Initiator) | Source or target device PID (statistics aggregated at the I level)                                |
| T (Target)    | Target device PID (statistics aggregated at the T level)                                          |
| ITL/ITN       | Destination device PID, direction of flow, protocol, source device PID, and LUN/NSID              |
| Entity ID     | VM entity ID (srceid), destination device PID, and ingress port ID                                |
| Port          | Name of the flow and port number. For example, <code>flow --show sys_flow_monitor -port 33</code> |

The following table lists the supported flow types for all monitored flows for Gen 6 and Gen 7 switches, chassis systems, and supported blades.

**Table 110: Flow Types for Monitored Flows**

| Flow Type  | Brocade G610 Switch | Brocade G620 Switch and Brocade FC32-48 and SX6 Blades | Brocade G630 Switch and Brocade FC32-64 Blades | Brocade FC32 -X7-48 Blades | Brocade G720 Switch and Brocade FC64-48 Blades | Brocade G730 Switch |
|------------|---------------------|--------------------------------------------------------|------------------------------------------------|----------------------------|------------------------------------------------|---------------------|
| IT         | Y                   | Y                                                      | Y                                              | Y                          | Y                                              | Y                   |
| ITL (SCSI) | N                   | N                                                      | N                                              | N                          | Y                                              | Y                   |
| ITN (NVMe) | N                   | N                                                      | N                                              | N                          | Y                                              | Y                   |
| VIT        | Y                   | Y                                                      | Y                                              | Y                          | N                                              | N                   |
| VITL/VITN  | N                   | N                                                      | N                                              | N                          | Y                                              | Y                   |
| Initiator  | Y                   | Y                                                      | Y                                              | Y                          | Y                                              | Y                   |
| Target     | N                   | N                                                      | N                                              | Y                          | Y                                              | Y                   |
| Port       | Y                   | Y                                                      | Y                                              | Y                          | Y                                              | Y                   |

The following table lists flow limits based on flow type.

**Table 111: Flow Limits per Flow Type**

| Flows          | Brocade Gen 6 Switches | Brocade Gen 6 Chassis Upgraded to Gen 7 | Brocade Gen 7 Directors | Brocade Gen 7 G720 and G730 Switches | Flow Granularity                                                                                                                       |
|----------------|------------------------|-----------------------------------------|-------------------------|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| Total Flows    | 2048                   | 2048                                    | 73,728 (72k)            | 7168                                 | <ul style="list-style-type: none"> <li>IT - 5 minutes</li> <li>ITL/ITN - 6 hours</li> <li>Real-time monitoring - 10 seconds</li> </ul> |
| IT             | 2048                   | 2048                                    | 8192 (8k)               | 1024                                 | 5 minutes                                                                                                                              |
| ITL/ITN        | N/A                    | N/A                                     | 64k                     | 6k                                   | 6 hours                                                                                                                                |
| Real Time      | 25                     | 25                                      | 25                      | 25                                   | 10 seconds                                                                                                                             |
| ITL/ITN per IT | N/A                    | N/A                                     | FCFS*                   | FCFS*                                | N/A                                                                                                                                    |

| Flows     | Brocade Gen 6 Switches | Brocade Gen 6 Chassis Upgraded to Gen 7 | Brocade Gen 7 Directors | Brocade Gen 7 G720 and G730 Switches | Flow Granularity |
|-----------|------------------------|-----------------------------------------|-------------------------|--------------------------------------|------------------|
| VIT       | FCFS*                  | FCFS*                                   | N/A                     | N/A                                  | N/A              |
| VITL/VITN | N/A                    | N/A                                     | FCFS*                   | FCFS*                                | N/A              |

\* FCFS – first-come, first-served

Flow Vision cannot monitor a zoned device with more than 124 devices. If a device is zoned with more than 124 devices, there will be no disruptions to the data path. However, Flow Vision cannot monitor the device.

VITL/VITN flows in Gen 7 directors share the chassis-level or ASIC-level maximum flow limits with ITL/ITN flows. Flows beyond these limits are dropped. VIT flows in Gen 6 directors share the chassis-level or ASIC-level maximum flow limits with IT flows. The flows beyond these limits are dropped.

The following table lists `sys_flow_monitor` support based on switches or port blades.

**Table 112: `sys_flow_monitor` Support by Switch/Port Blade**

| Statistic Type | Protocol | Brocade G620 Switch and Brocade FC32-48 and SX6 Blades | Brocade G630, G720, and G730 Switches and Brocade FC32-64, FC 64-48, and FC32-X7-48 Blades | Brocade G610 Switch and 7810 Extension Switch |
|----------------|----------|--------------------------------------------------------|--------------------------------------------------------------------------------------------|-----------------------------------------------|
| FC             | —        | Y                                                      | Y                                                                                          | Y                                             |
| I/O            | SCSI     | Y                                                      | Y                                                                                          | N                                             |
| I/O            | NVMe     | N                                                      | Y                                                                                          | N                                             |

Automated flow monitoring supports a mix of blades on the supported platforms.

Automated flow monitoring performs threshold-based monitoring and alerting based on flows through integration with the Monitoring and Alerting Policy Suite (MAPS) service. Newly learned flows are discovered, and no parallel user configuration is required.

For detailed information on integration with MAPS, refer to the *Brocade Fabric OS MAPS User Guide, 9.1.x*.

## Support for Multiple Virtual Fabrics

Flow Vision in Fabric OS 9.1.x supports automatic monitoring of all F\_Ports on multiple virtual fabrics (VFs). These VFs are physical switches that can be partitioned into independently managed logical switches, each with its own data, control, and management paths. The automatic monitoring flow `sys_flow_monitor` is enabled by default on all VFs. Operations such as activate, deactivate, modify, and reset are done at the VF level. These operations in one VF are independent from the operations of other VFs. All these operations and any modifications to these operations can be done in VFs during runtime without any disruption. For example, if you have a FID 128 with 10 F\_Ports and another FID 10 with 2 F\_Ports, you can choose to monitor 6 F\_Ports in FID 128 and monitor a single F\_Port in FID 10. You can move one F\_Port from one logical partition to the other logical partition, create a VF, or delete the VF while the flows continue to be monitored.

Support for monitoring multiple VFs is available in Brocade Gen 6 and Gen 7 platforms. On Gen 6 platforms, disruption may occur during deactivating, resetting, or moving monitored ports across VFs if an ASIC is shared between multiple FIDs. Gen 6 platforms support operations such as activate, deactivate, modify, and reset during runtime without any disruption.

The following table describes the support of reset, deactivate, and moving logical port operations for multiple VFs on various Brocade platforms.



**Table 113: Support for Multiple VFs Support on Brocade Platforms**

| Operation                                        | Brocade Gen 7 Switches                                                                                                                                                                                   | Brocade Gen 6 Switches                                                                                                                                                              | Brocade Gen 7 Directors with FC32-X7-48 and FC64-48 Blades                                                                                                                                               | Brocade Gen 6 Directors                                                                                                                                                             | Brocade Directors with Combination of Gen 7 Blades and Condor-4 ASIC                                                                                                                                                                                                                                |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Resetting/<br>Deactivating a single VF           | <ul style="list-style-type: none"> <li>All data is reset for the monitored F_Ports in the current VF.</li> <li>No disruption occurs in the current sample for monitored F_Ports in other VFs.</li> </ul> | <ul style="list-style-type: none"> <li>All data is reset for the monitored F_Ports in the current VF.</li> <li>A few samples may be lost for F_Ports in the shared ASIC.</li> </ul> | <ul style="list-style-type: none"> <li>All data is reset for the monitored F_Ports in the current VF.</li> <li>No disruption occurs in the current sample for monitored F_Ports in other VFs.</li> </ul> | <ul style="list-style-type: none"> <li>All data is reset for the monitored F_Ports in the current VF.</li> <li>A few samples may be lost for F_Ports in the shared ASIC.</li> </ul> | <ul style="list-style-type: none"> <li>All data is reset for the monitored F_Ports in the current VF.</li> <li>No disruption occurs in the current sample for monitored F_Ports in the shared Condor-5 ASIC.</li> <li>A few samples may be lost for F_Ports in the shared Condor-4 ASIC.</li> </ul> |
| Moving a logical port to a new logical partition | <ul style="list-style-type: none"> <li>The old flow remains in the source logical partition.</li> <li>The new logical partition rediscovers the flow.</li> </ul>                                         | <ul style="list-style-type: none"> <li>The old flow remains in the source logical partition.</li> <li>The new logical partition rediscovers the flow.</li> </ul>                    | <ul style="list-style-type: none"> <li>The old flow remains in the source logical partition.</li> <li>The new logical partition rediscovers the flow.</li> </ul>                                         | <ul style="list-style-type: none"> <li>The old flow remains in the source logical partition.</li> <li>The new logical partition rediscovers the flow.</li> </ul>                    | <ul style="list-style-type: none"> <li>The old flow remains in the source logical partition.</li> <li>The new logical partition rediscovers the flow.</li> </ul>                                                                                                                                    |

## Flow Metrics

Automated flow monitoring in Fabric OS 9.1.x provides FC frame and I/O metrics on all application flows in the monitored switch that is managed by SANnav and Fabric Analytics.

The following table describes the supported FC key metrics on Gen 6 and Gen 7 platforms that run Fabric OS 9.1.x.

**Table 114: Supported FC Key Metrics**

| Supported FC Metric | Support on Brocade Gen 6 Platforms                                                                                                                                                          | Support on Brocade Gen 7 Platforms                                                                                                                                                                                |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Frame Count         | <ul style="list-style-type: none"> <li>Only ingress (Rx) flow monitoring is supported.</li> <li>Only host ports are monitored.</li> <li>Access Gateway supports only FC metrics.</li> </ul> | <ul style="list-style-type: none"> <li>Both ingress and egress (Rx/Tx) flow monitoring is supported.</li> <li>Host and target F_Ports are monitored.</li> <li>Access Gateway supports only FC metrics.</li> </ul> |
| Word Count          |                                                                                                                                                                                             |                                                                                                                                                                                                                   |
| Frame Rate          |                                                                                                                                                                                             |                                                                                                                                                                                                                   |
| Frame Size          |                                                                                                                                                                                             |                                                                                                                                                                                                                   |
| Throughput          |                                                                                                                                                                                             |                                                                                                                                                                                                                   |

The following table describes the supported I/O metrics on Gen 6 and Gen 7 platforms that run Fabric OS 9.1.x.

**Table 115: Supported I/O Metrics**


| Supported I/O Metric                          | Support on Brocade Gen 6 Platforms                                                                                                                     | Support on Brocade Gen 7 Platforms                                                                                                                             |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| First Response Time Accumulated (Avg and Max) | <ul style="list-style-type: none"> <li>SCSI and NVMe Read (RD) and Write (WR) commands are supported.</li> <li>Only IT flows are monitored.</li> </ul> | <ul style="list-style-type: none"> <li>SCSI and NVMe Read (RD) and Write (WR) commands are supported.</li> <li>Both IT and ITL flows are monitored.</li> </ul> |
| First Response Time Max                       |                                                                                                                                                        |                                                                                                                                                                |
| Exchange Completion Time Accumulated          |                                                                                                                                                        |                                                                                                                                                                |
| Exchange Completion Time Max                  |                                                                                                                                                        |                                                                                                                                                                |
| I/O Count                                     |                                                                                                                                                        |                                                                                                                                                                |
| IO Size                                       |                                                                                                                                                        |                                                                                                                                                                |
| Pending I/O                                   |                                                                                                                                                        |                                                                                                                                                                |
| Throughput                                    |                                                                                                                                                        |                                                                                                                                                                |
| Other Command Errors and Exceptions           |                                                                                                                                                        |                                                                                                                                                                |
| Errors and Exceptions                         |                                                                                                                                                        |                                                                                                                                                                |
| Real-time I/O Violations                      |                                                                                                                                                        |                                                                                                                                                                |

Latency metrics, which include Exchange Completion Time, First Response Time, Pending I/O, and Real-time I/O are at the core of metric support in Fabric OS 9.1.x.

For Read I/O, the First Response Time is taken when the first data frame is transmitted by the target after receiving the command frame. For Write I/O, the First Response Time is taken when the target sends the Transfer (Xfer) Ready frame after receiving the command frame. Completion Time is taken when the target transmits the status frame to complete an exchange. Refer to the following figure for an illustration.


Read latency metrics include the maximum and average time values for monitored First Response Time and Command Completion Latency Time for Read command frames.

- First Response Read Latency is the time difference between when a Read command frame (FCP command) is issued by the host and when the corresponding First Data Response frame is returned by the target.
- Read Completion Latency is the time difference between when a Read command frame is issued by the host and when the corresponding Status frame is returned by the target.

**Figure 42: Read Latency**

Write latency metrics include the maximum and average time values for the monitored First Response Time and Command Completion Latency Time for Write command frames. Refer to the following image to understand what follows.

- First Response Write Latency is the time difference between when the Write command frame is issued by the host and when the corresponding First Transfer Ready frame is returned by the target.
- Write Command Completion latency is the time difference between when a Write command frame is issued by the host and when the corresponding Status frame is returned by the target.

**Figure 43: Write Latency****Pending I/Os**

- Pending I/Os are the average number of I/Os yet to be completed when a command frame is received at the target port.
- Total Pending I/Os reflect the actual IT queue depth. Pending I/Os are learned at the IT flow level for 10 seconds.

**Flow-Level Errors and Exception Metrics**

Fabric OS 9.1.x provides visibility into the I/O Exceptions Count and the Other (non Read/Write) commands in I/O exchanges between an initiator and a target. I/O and latency metrics for replication traffic flows for the READ and WRITE commands are captured and reported separately. Other (non Read/Write) commands for SCSI are reported separately depending on platforms. The counts displayed under the Other (non Read/Write) commands category helps you to discover the SCSI devices, troubleshoot the issues that are associated with the SCSI commands, and faster error recovery.

The following Other (non Read/Write) commands counts are supported for SCSI flows:

- Total
- Reserve
- Release
- Inquiry
- Request Sense
- Test Unit Ready
- Read Capacity

Flow Vision in Fabric OS 9.1.x supports the I/O exceptions measure type for flow investigation. The following table describes the supported measures for SCSI flows.

**Table 116: I/O Exceptions Counts Supported for SCSI**

| Flows  | Supported I/O Exceptions Count                                                                                                                                                                                                                                                            |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SCSI   | <ul style="list-style-type: none"> <li>• Busy</li> <li>• Check Conditions</li> <li>• Reserve Conflict</li> <li>• Task Abort</li> <li>• Task Set Full</li> <li>• Timeout</li> </ul> <p>The SCSI exceptions that are not listed in the above list are presented in the OTHERS category.</p> |
| OTHERS | <ul style="list-style-type: none"> <li>• Auto Count Align (ACA)</li> <li>• Conditions Met</li> </ul>                                                                                                                                                                                      |

The frames of SCSI Other (non Read/Write) commands are captured and mirrored by ASIC to CPU. Flow Vision monitors the mirrored frames from CPU and combines with the automated `sys_flow_monitor` flow. The counts of the mirrored frames are presented under the Other (non Read/Write) commands category.

The following table describes the support of flow statistics across the platforms that run Fabric OS 9.1.x.

**Table 117: Flow Statistics Support Matrix**

| Statistics                  | Brocade G610 Switch and Brocade 7810 Extension Switch | Brocade G620 Switch and Brocade FC32-48 and SX6 Blades | Brocade G630 Switch and Brocade FC32-64 Blades | Brocade FC32-X7-48 Blades | Brocade G720 Switch and Brocade FC64-48 Blades | Brocade G730 Switch |
|-----------------------------|-------------------------------------------------------|--------------------------------------------------------|------------------------------------------------|---------------------------|------------------------------------------------|---------------------|
| FC                          | Y                                                     | Y                                                      | Y                                              | Y                         | Y                                              | Y                   |
| SCSI                        | N                                                     | Y                                                      | Y                                              | Y                         | Y                                              | Y                   |
| NVMe                        | N                                                     | N                                                      | Y                                              | Y                         | Y                                              | Y                   |
| IO RD/WR Stats              | N                                                     | Y                                                      | Y                                              | Y                         | Y                                              | Y                   |
| Other Command Counts (SCSI) | N                                                     | N                                                      | N                                              | Y                         | Y                                              | Y                   |
| I/O Exceptions              | N                                                     | N                                                      | N                                              | Y                         | Y                                              | Y                   |
| Port Pending I/O            | N                                                     | Y                                                      | Y                                              | Y                         | Y                                              | Y                   |

The flow statistics are platform-dependent. On Gen 7 platforms, the Port Pending I/O statistics are hardware-driven.

## Flow Granularity

In Fabric OS 9.1.x, the flow granularity is flexible. An IT flow can reside at 5-minute granularity, and an ITL flow can reside at 6-hour granularity by default. 10-second granularity is supported only when the IT or ITL flow is moved to real-time investigation mode through SANnav. 5-minute granularity supports maximum of 72k IT flows.

## Viewing Flows

Flow Vision allows you to view the configuration parameters for each flow on a switch.

**Table 118: Valid show Commands for sys\_flow\_monitor**

| Command                                                                                                                           | Action                                                                                             |
|-----------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| <code>flow --show</code>                                                                                                          | Show a list of all existing flows.                                                                 |
| <code>flow --show sys_flow_monitor</code>                                                                                         | Show details on the top violations for monitored flows.                                            |
| <code>flow --show sys_flow_monitor -stats</code>                                                                                  | Show details on the top 10 ECT for monitored flows.                                                |
| <code>flow --show sys_flow_monitor -port</code>                                                                                   | Show details of the ingress or egress port on which <code>sys_flow_monitor</code> is applied.      |
| <code>flow --show sys_flow_monitor -resource -fid</code>                                                                          | Show the maximum number of IT and ITL flows learned by <code>sys_flow_monitor</code> on a chassis. |
| <code>flow --show sys_flow_monitor -srcdev "*" -dstdev "*"</code>                                                                 | Show brief metrics of flow I/O statistics.                                                         |
| <code>flow --show sys_flow_monitor -srcdev SID -destdev DID</code>                                                                | Show per-flow I/O statistics.                                                                      |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -lun "*"</code>                                    | Show brief metrics for the ITL flows in a specific IT flow.                                        |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -nsid "*"</code>                                   | Show brief metrics for the ITN flows in a specific IT flow.                                        |
| <code>flow --show sys_flow_monitor --srcdev &lt;devId&gt;</code>                                                                  | Show a summary of metrics for a host flow.                                                         |
| <code>flow --show sys_flow_monitor --dstdev &lt;devId&gt;</code>                                                                  | Show a summary of metrics for a target flow.                                                       |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -lun &lt;LUN ID&gt;</code>                         | Show detailed metrics for a specific ITL flow.                                                     |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -nsid &lt;NS ID&gt;</code>                         | Show detailed metrics for a specific ITN flow.                                                     |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -srceid "*"</code>                                 | Show brief metrics for VMIT flows in a specific IT flow.                                           |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -lun &lt;LUN ID&gt; -srceid "*"</code>             | Show brief metrics for VMITL flows in a specific ITL flow.                                         |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -nsid &lt;NS ID&gt; -srceid "*"</code>             | Show brief metrics for VMITN flows in a specific ITN flow.                                         |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -srceid &lt;SRCE ID&gt;</code>                     | Show detailed metrics for a specific VMIT flow.                                                    |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -lun &lt;LUN ID&gt; -srceid &lt;SRCE ID&gt;</code> | Show detailed metrics for a specific VMITL flow.                                                   |
| <code>flow --show sys_flow_monitor -srcdev &lt;devId&gt; -dstdev &lt;devId&gt; -nsid &lt;NS ID&gt; -srceid &lt;SRCE ID&gt;</code> | Show detailed metrics for a specific VMITN flow.                                                   |

| Command                                                                                                                                                                                                                                                                                                                   | Action                                                                                                                                                                                                           |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• flow --show sys_flow_monitor -top frt</li> <li>• flow --show sys_flow_monitor -bottom frt</li> <li>• flow --show sys_flow_monitor -top ect</li> <li>• flow --show sys_flow_monitor -top pending_io -n &lt;count&gt;</li> <li>• flow --show sys_flow_monitor -top iops</li> </ul> | Show the top and bottom 10 flows or a configurable number of flows of type Exchange Completion Time (ECT), First Response Time (FRT), Bytes Per Second (BPS), I/O Operations per Second (IOPS), and Pending I/O. |

To display all existing flows, enter the `flow --show` command. Flow names with a "sys\_" prefix are predefined flows, and a "+" in the Feature column indicates that the named feature is active for the flow.

```
switch:admin> flow --show
```

| Flow Name            | Feature   | SrcDev | DstDev | IngrPt | EgrPt | BiDir | LUN |
|----------------------|-----------|--------|--------|--------|-------|-------|-----|
| FrameType            | SFID DFID | SEID   | MirPt  |        |       |       |     |
| sys_gen_all_simports | gen       | *      | *      | *      | *     | no    | -   |
| -                    | -         | -      | -      |        |       |       |     |
| sys_flow_monitor     | mon+      | *      | *      | *      | *     | no    | *   |
| -                    | -         | -      | -      |        |       |       |     |
| flow_generator1      | gen       | 632c00 | 632800 | 44     | -     | no    | -   |
| -                    | -         | -      | -      |        |       |       |     |
| flow_generator2      | gen       | 631a00 | 637900 | 26     | -     | no    | -   |
| -                    | -         | -      | -      |        |       |       |     |
| flow_mirror1         | mir       | *      | *      | 33     | -     | yes   | -   |
| abts                 | -         | -      | -      |        |       |       |     |
| flow_mirror1         | mir       | *      | *      | 33     | -     | yes   | -   |
| sing                 | -         | -      | -      |        |       |       |     |

+ Denotes feature is currently activated for the flow  
The flow name with prefix sys\_ denotes predefined flow

To display the definition for a specific flow, enter the `flow --show <flow_name>` command. Flow data, the flow definition for the specified flow, and feature-specific information are displayed.

To show details on the `sys_flow_monitor` flow, enter the `flow --show sys_flow_monitor` command. The output shows the `sys_flow_monitor` flow details with the top 10 threshold violations (if there are any violations observed).

The following example shows a typical output of this command with the top 10 threshold violations.

```
switch:admin> flow --show sys_flow_monitor
```

```
=====
Name : sys_flow_monitor
VM Mode : No
Monitored port(s): (*)
Host Port Count : 1
Target Port Count: 0
SCSI : 1
NVMe : 0
No of IT : 1
=====
```

```
| TOP 10 Threshold Violations |
```

```
=====
| | | RD| | Total Count | Violations | Violations | Violations | Violations |
Total Count |
```

| Port       | PID    | WR | Device | Proto | Violations | ECT        | FRT        | Pending IOs | Exceptions |
|------------|--------|----|--------|-------|------------|------------|------------|-------------|------------|
| Errors     |        |    |        |       |            |            |            |             |            |
| Index      |        |    | Type   |       | 5min / All | 5min / All | 5min / All | 5min / All  | 5min / All |
| 5min / All |        |    |        |       |            |            |            |             |            |
| 33         | 632100 | RD | I      | SCSI  | /2.316K    | /1.152K    | /1.164K    | /           | /          |
| /          | /      |    |        |       |            |            |            |             |            |
| 33         | 632100 | WR | I      | SCSI  | /2.490K    | /1.361K    | /1.129K    | /           | /          |
| /          | /      |    |        |       |            |            |            |             |            |

To view the top 10 brief IT flow metrics, enter the `flow --show sys_flow_monitor -stats` command.

The following example shows the top 10 IT flow metrics. In the following example, VM-I indicates the port in which applications are registered.

```
switch:admin> flow --show sys_flow_monitor -stats
```

```
Name : sys_flow_monitor
```

```
VM Mode : Yes
```

```
Monitored port(s): (*)
```

```
Host Port Count : 6
```

```
Target Port Count: 5
```

```
SCSI : 7
```

```
NVMe : 4
```

```
No of IT : 11
```

```

```

| SID           | DID         | WR         | Device     | Flow      | Port | Proto | ECT           | FRT           | IOPS          |
|---------------|-------------|------------|------------|-----------|------|-------|---------------|---------------|---------------|
| BPS           | Pending IOs | Exceptions | Violations |           |      |       |               |               |               |
| 5min / All    | 5min / All  | 5min / All | 5min / All |           |      |       |               |               |               |
| 71af40        | 718d40      | WR         | VM-I       | IT->L->VM | 431  | SCSI  | 124.4m/144.3m | 70.97m/91.88m | 935 / 923     |
| 233.9M/230.7M | 64 / 64     | /          | /          | /         | /    | /     |               |               |               |
| 71af40        | 718d40      | RD         | VM-I       | IT->L->VM | 431  | SCSI  | 86.90m/92.53m | 86.79m/91.82m | 935 / 923     |
| 233.9M/230.7M | 48 / 48     | /          | /          | /         | /    | /     |               |               |               |
| 290f00        | 718d40      | WR         | T          | IT->L->VM | 397  | SCSI  | 86.84m/93.65m | 86.75m/93.40m | 24.70K/24.32K |
| 211.8M/208.8M | 9 / 48      | /          | /          | /         | /    | /     |               |               |               |
| 290f00        | 718d40      | RD         | T          | IT->L->VM | 397  | SCSI  | 86.71m/93.63m | 86.70m/93.62m | 24.69K/24.32K |
| 211.8M/208.8M | 5 / 22      | /          | /          | /         | /    | /     |               |               |               |
| 719b40        | 9f0200      | RD         | I          | IT->L     | 411  | SCSI  | 585u/10.63m   | 159u/10.61m   | 12.19K/12.10K |
| 362.7M/360.2M | 9 / 10      | /          | /          | /         | /    | /     |               |               |               |
| 290700        | 718d00      | RD         | T          | IT        | 141  | NVMe  | 264u/ 631u    | 131u/ 631u    | 7.877K/7.841K |
| 238.1M/237.1M | 7 / 7       | /          | /          | /         | /    | /     |               |               |               |
| 290600        | 718d00      | RD         | T          | IT        | 141  | NVMe  | 228u/1.927m   | 125u/1.927m   | 7.496K/7.520K |
| 230.4M/231.1M | 7 / 7       | /          | /          | /         | /    | /     |               |               |               |
| 719b40        | 9f0200      | WR         | I          | IT->L     | 411  | SCSI  | 215u/10.69m   | 143u/10.62m   | 11.43K/11.35K |
| 357.3M/354.9M | 4 / 6       | /          | /          | /         | /    | /     |               |               |               |



|               |        |     |   |   |    |     |      |             |             |               |
|---------------|--------|-----|---|---|----|-----|------|-------------|-------------|---------------|
| 290600        | 718d00 | WR  | T |   | IT | 141 | NVMe | 186u/36.17m | 123u/36.11m | 7.298K/7.321K |
| 228.0M/228.8M |        | 3 / | 3 | / |    | /   |      |             |             |               |
| 290700        | 718d00 | WR  | T |   | IT | 141 | NVMe | 184u/36.18m | 123u/36.14m | 7.462K/7.433K |
| 233.2M/232.3M |        | 3 / | 3 | / |    | /   |      |             |             |               |

To view the maximum number of IT and ITL flows learned by `sys_flow_monitor` on a chassis, enter the `flow --sys flow monitor -resource` command.

## NOTE

The ITL flows are not displayed on a Gen 6 chassis.

The following example shows the values for IT, ITL, and both IT and ITL flows (total limits): currently used flows, reserved flows, and maximum flows on a Gen 7 chassis.

```
swtich:admin> flow --show sys flow monitor -resource
```

| Flows     | Current Configuration | Chassis Maximum |
|-----------|-----------------------|-----------------|
|           | Used                  |                 |
| =====     |                       |                 |
| IT count  | 1                     | 8192            |
| ITL count | 4                     | 16384           |
| Total     | 5                     | 24576           |

\*ITL limits are derived based on IT reservation limit  
Timebase:

| Time Bases | Used | Max   |
|------------|------|-------|
| 10 sec     | 0    | 25    |
| 5 min      | 1    | 8192  |
| 6 hr       | 4    | 16384 |

Slot Resource:

| Slot    | Total     | Initiator |      |       | Target |      |       | IT      |      |
|---------|-----------|-----------|------|-------|--------|------|-------|---------|------|
| ITL     |           |           |      |       |        |      |       |         |      |
| No      | F-Port(s) | SCSi      | NVMe | Total | SCSi   | NVMe | Total | Current | Max  |
| Current | Max       |           |      |       |        |      |       |         |      |
| =====   |           |           |      |       |        |      |       |         |      |
| 3       | 1         | 1/        | 0/   | 1     | 0/     | 0/   | 0     | 1/      | 4096 |
| 4/      | 8192      |           |      |       |        |      |       |         |      |
| -----   |           |           |      |       |        |      |       |         |      |
| 4       | 0         | 0/        | 0/   | 0     | 0/     | 0/   | 0     | 0/      | 4096 |
| 0/      | 8192      |           |      |       |        |      |       |         |      |

\* These values are system defined and not configurable

The following example shows the number of IT flows on a Gen 6 chassis:

```
swtich:admin> flow --show sys flow monitor -stats
```

Name : sys\_flow\_monitor

VM Mode : Yes

Monitored port(s): (\*)

Host Port Count : 7

SCSI : 2

NVMe : 2

No of IT : 4

-----  
TOP 10 ECT |

| Bytes/sec     |        | Max Count |              | RD   | Device    | Flow  | Port | Proto | ECT           | Max Time      | Max Time   | Avg I/Os/sec  | Avg |
|---------------|--------|-----------|--------------|------|-----------|-------|------|-------|---------------|---------------|------------|---------------|-----|
| SID           | DID    | WR        | Pending I/Os | Type | Hierarchy | Index |      |       |               | 5min / All    | 5min / All | 5min / All    |     |
| 290f00        | 718d40 | RD        | 12 / 14      | VM-I | IT->VM    | 15    | SCSI |       | 85.73m/145.4m | 85.72m/145.4m |            | 24.37K/24.44K |     |
| 811.2G/1.103T |        |           |              |      |           |       |      |       |               |               |            |               |     |
| 290f00        | 718d40 | WR        | 17 / 20      | VM-I | IT->VM    | 15    | SCSI |       | 85.69m/145.3m | 85.64m/145.1m |            | 24.36K/24.44K |     |
| 811.2G/1.103T |        |           |              |      |           |       |      |       |               |               |            |               |     |
| 290700        | 718d00 | RD        | 6 / 7        | I    | IT        | 7     | NVMe |       | 284u/1000m    | 152u/1000m    |            | 7.844K/7.768K |     |
| 1.187T/594.5G |        |           |              |      |           |       |      |       |               |               |            |               |     |
| 295200        | 9f0600 | WR        | 2 / 2        | I    | IT        | 82    | SCSI |       | 282u/17.78m   | 241u/17.74m   |            | 6.764K/6.673K |     |
| 409.8G/550.5G |        |           |              |      |           |       |      |       |               |               |            |               |     |
| 290600        | 718d00 | RD        | 6 / 9        | I    | IT        | 6     | NVMe |       | 248u/36.20m   | 151u/36.17m   |            | 7.502K/7.453K |     |
| 1.187T/578.9G |        |           |              |      |           |       |      |       |               |               |            |               |     |
| 290600        | 718d00 | WR        | 2 / 12       | I    | IT        | 6     | NVMe |       | 209u/36.30m   | 146u/36.25m   |            | 7.304K/7.320K |     |
| 1.187T/577.6G |        |           |              |      |           |       |      |       |               |               |            |               |     |
| 290700        | 718d00 | WR        | 2 / 3        | I    | IT        | 7     | NVMe |       | 206u/36.31m   | 147u/36.25m   |            | 7.433K/7.431K |     |
| 1.187T/587.1G |        |           |              |      |           |       |      |       |               |               |            |               |     |
| 295200        | 9f0600 | RD        | 3 / 3        | I    | IT        | 82    | SCSI |       | 135u/17.59m   | 121u/17.57m   |            | 7.140K/7.046K |     |
| 573.6G/549.6G |        |           |              |      |           |       |      |       |               |               |            |               |     |

To display the resource consumption for each fabric ID (FID), enter the flow `--show sys_flow_monitor -resource -fid` command.

The following example shows flow `--show sys_flow_monitor -resource -fid` command output. The output is separated into flows by the FID being monitored and lists them by counts of IT and ITL flows. The `-fid` operand is optional and is available only with the `-resource` option.

```
switch:admin> flow --show sys_flow_monitor -resource -fid
```

Resource consumption by FID:

| FID | Timebase | IT Count | ITL Count | Total |
|-----|----------|----------|-----------|-------|
| 128 | 10 sec   | 0        | 0         | 0     |
|     | 5 min    | 0        | 0         | 0     |
|     | 6 hr     | 0        | 0         | 0     |

|       |       |  |        |  |    |  |    |  |    |  |
|-------|-------|--|--------|--|----|--|----|--|----|--|
|       | 44    |  | 10 sec |  | 0  |  | 0  |  | 0  |  |
|       |       |  | 5 min  |  | 1  |  | 0  |  | 1  |  |
|       |       |  | 6 hr   |  | 0  |  | 1  |  | 1  |  |
| ----- |       |  |        |  |    |  |    |  |    |  |
|       | 88    |  | 10 sec |  | 0  |  | 0  |  | 0  |  |
|       |       |  | 5 min  |  | 14 |  | 0  |  | 14 |  |
|       |       |  | 6 hr   |  | 0  |  | 14 |  | 14 |  |
| ----- |       |  |        |  |    |  |    |  |    |  |
|       | 1     |  | 10 sec |  | 0  |  | 0  |  | 0  |  |
|       |       |  | 5 min  |  | 0  |  | 0  |  | 0  |  |
|       |       |  | 6 hr   |  | 0  |  | 0  |  | 0  |  |
| ----- |       |  |        |  |    |  |    |  |    |  |
|       | Total |  | ALL    |  | 15 |  | 15 |  | 30 |  |
| ----- |       |  |        |  |    |  |    |  |    |  |

To view the statistics of the ingress or egress port on which `sys_flow_monitor` is applied, enter the `flow --show sys_flow_monitor -port <port>` command. The selected port can be a wildcard port, a single port index, or a MAPS logical group port.

The following example shows port-specific details for an initiator port on a Gen 7 chassis.

```
switch:maintenance-service> flow --show sys_flow_monitor -port 59
```

```
=====
Name : sys_flow_monitor
Definition : Port(59)
Port Speed : SrcDev(16G)
Active Flow : 1
Timebase : 5Min
Protocol : SCSI
=====
```

```
Monitor Time : | Thu Jan 01 00:00:00 GMT 1970 |
RD Elapsed Time : | 0s |
WR Elapsed Time : | 0s |
```

```

FC Statistics |
```

| Frame Count |      |       | Frames Per Second |      |       | Byte Count |      |       | Throughput (BPS) |      |       | Frame |
|-------------|------|-------|-------------------|------|-------|------------|------|-------|------------------|------|-------|-------|
| Size(Bytes) |      |       |                   |      |       |            |      |       |                  |      |       |       |
| Tx /        | Rx / | Total | Tx /              | Rx / | Total | Tx /       | Rx / | Total | Tx /             | Rx / | Total | Tx /  |
| Rx          |      |       |                   |      |       |            |      |       |                  |      |       |       |
| 0 /         | 0 /  | 0     | 0 /               | 0 /  | 0     | 0 /        | 0 /  | 0     | 0 /              | 0 /  | 0     | 0 /   |
|             |      |       |                   |      |       |            |      |       |                  |      |       | 0     |

```

I/O Latency |
```

| Metric             |  | Max     |        | Avg     |        |
|--------------------|--|---------|--------|---------|--------|
|                    |  | Current | All    | Current | All    |
| RD Completion Time |  | /       | 2.959s | /       | 436.8m |

|                      |  |   |        |   |        |
|----------------------|--|---|--------|---|--------|
| WR Completion Time   |  | / |        | / |        |
| RD 1st Response Time |  | / | 2.097s | / | 465.7m |
| WR 1st Response Time |  | / |        | / |        |
| RD Pending IOs       |  | / | 528    | / | 528    |
| WR Pending IOs       |  | / |        | / |        |

## I/O Performance |

| Metric      |           | Avg IO Size |   | Avg (IOPS)    |   | Avg (BPS)     |   | Total IO  |  |
|-------------|-----------|-------------|---|---------------|---|---------------|---|-----------|--|
| Total Bytes |           | Current     |   | Current / All |   | Current / All |   | Current / |  |
| All         | Current / | All         |   |               |   |               |   |           |  |
| RD IO       |           |             | / | 178           | / | 1.357M        | / |           |  |
| 1.780K      | /         | 13.58M      |   |               |   |               |   |           |  |
| WR IO       |           |             | / |               | / |               | / |           |  |
|             | /         |             |   |               |   |               |   |           |  |

## I/O Violation Count |

| RD           | Total      | Violations | Violations | Violations  |
|--------------|------------|------------|------------|-------------|
| WR  Timebase | Violations | ECT        | FRT        | Pending IOs |
| RD  Current  |            | /          | /          | /           |
| All          | 1.529K     | 720 /      | 809 /      | /           |
| WR  Current  |            | /          | /          | /           |
| All          |            | /          | /          | /           |

## I/O Exceptions Count |

| RD          | SCSI       |         |      |          |           |       |          |  |  |  |
|-------------|------------|---------|------|----------|-----------|-------|----------|--|--|--|
|             |            |         |      |          |           |       |          |  |  |  |
| Timebase    | Total      |         |      |          |           |       |          |  |  |  |
| WR          | Exceptions | TIMEOUT | BUSY | RESERVE  | CHECK     | TASK  | TASK SET |  |  |  |
| OTHERS      |            |         |      | CONFLICT | CONDITION | ABORT | FULL     |  |  |  |
| RD  Current |            |         |      |          |           |       |          |  |  |  |
| All         |            |         |      |          |           |       |          |  |  |  |

```

WR	Current						
	All						

```

#### SCSI Other Command Counts |

```

| Timebase | Total | INQUIRY | RESERVE | RELEASE | REQ SENSE | TUR |
| READ CAP | | | | | | |

Current						
All						

=====

```

The following example shows port-specific details for a target port on a Gen 7 chassis.

```
switch:admin> flow --show sys_flow_monitor -port 397
```

```

=====
Name : sys_flow_monitor
Definition : Port(397)
Port Speed : DstDev(32G)
Active Flow : 2
Elapsed Time : | 4m:43s |

```

#### FC Statistics |

```

| Frame Count | Frames Per Second | Byte Count | Throughput (BPS) |Frame
| Size(Bytes) |
| Tx / Rx / Total | Tx / Rx / Total | Tx / Rx / Total | Tx / Rx / Total | Tx /
Rx
6.29M/ 1.79M/ 8.09M	623.68k/177.54k/801.23k	1.49G/ 1.25G/ 2.75G	151.73M/127.78M/279.51M	256/ 756

```

#### I/O Latency |

```

| Metric | Max | Avg |
| Current / All | Current / All |

|RD Completion Time | 315u/ 55.03m| 57u/ 56u|

|WR Completion Time | 363u/ 55.08m| 99u/ 97u|

|RD 1st Response Time | 299u/ 55.02m| 53u/ 51u|

|WR 1st Response Time | 284u/ 55.06m| 51u/ 50u|

```

```

|RD Pending IOs | 4 / 4 | 1 / 1 |

```

```
|WR Pending IOs | 4 / 4 | 1 / 1 |

```

```

I/O Performance
```

```

| Metric |Avg IO Size | Avg (IOPS) | Avg (BPS) | Total IO
| | Total Bytes | | | | | |
| | | | | | | | |
| All | Current / | All | Current / | All | Current / | All | Current /

```

```

|RD IO | 3.106K| 38.80K/ 9.435K| 117.7M/ 28.66M| 10.98M/
707.3M | 32.54G/ 2.048T|

```

```

|WR IO | 3.106K| 38.80K/ 9.435K| 117.7M/ 28.66M| 10.98M/
707.3M | 32.54G/ 2.048T|

```

```

I/O Violation Count
```

```

| RD| | Total | Violations | Violations | Violations |
| WR| Timebase | Violations | ECT | FRT | Pending IOs |

| RD| Current | | | | |
| | All | | | | |

| WR| Current | | | | |
| | All | | | | |

```

```

I/O Exceptions Count
```

```

| RD| | | | | | | | |
| | Timebase | Total | | | | | | |
|-----|
| WR| | Exceptions | TIMEOUT | BUSY | RESERVE | CHECK | NAME SPACE | RESERVE |
| OTHERS | | | | | | | | |
| | | | | | | | | |
|-----|
| RD| Current | | | | | | | |
| | All | | | | | | | |

```

[illegible]

In this example, both source and destination are specified for IT flows.

```
switch:admin> flow --show sys flow monitor -srcdev 632100 -dstdev 718640
```

```
Name : sys_flow_monitor
Definition : SrcDev(632100)DstDev(718640)
Port Speed : SrcDev(32G), DstDev(16G)
Active Flow : 1
Timebase : 5Min
Protocol : SCSI
Port : 390
```

```
Monitor Time : | Mon Aug 16 04:23:26 MDT 2021 |
RD Elapsed Time : | 3m:33s |
WR Elapsed Time : | 3m:33s |
```

FC Statistics |

| Frame Count  |         |       | Frames Per Second |        |         | Byte Count |         |         | Throughput (BPS) |        |        | Frame |     |
|--------------|---------|-------|-------------------|--------|---------|------------|---------|---------|------------------|--------|--------|-------|-----|
| Size (Bytes) |         |       |                   |        |         |            |         |         |                  |        |        |       |     |
| Tx           | Rx      | Total | Tx                | Rx     | Total   | Tx         | Rx      | Total   | Tx               | Rx     | Total  | Tx    | Rx  |
| -----        |         |       |                   |        |         |            |         |         |                  |        |        |       |     |
| 482.14k      | 602.70k | 1.08M | 47.73k            | 59.67k | 107.40k | 498.42M    | 498.44M | 996.87M | 49.34M           | 49.35M | 98.69M | 1084  | 868 |

I/O Latency |

| Metric               | Max     |        | Avg     |     |
|----------------------|---------|--------|---------|-----|
|                      | Current | All    | Current | All |
| RD Completion Time   | 1.677m/ | 9.114m | 35u/    | 34u |
| WR Completion Time   | 1.700m/ | 8.388s | 43u/    | 43u |
| RD 1st Response Time | 1.631m/ | 9.065m | 19u/    | 19u |
| WR 1st Response Time | 1.642m/ | 9.031m | 16u/    | 16u |
| RD Pending IOs       | 7 /     | 9      | 1 /     | 1   |

```

|WR Pending IOs | 7 / 9 | 1 / 1 |

=====
```

```

I/O Performance |
```

```

| Metric |Avg IO Size | Avg(IOPS) | Avg(BPS) | Total IO
| | Total Bytes | | | | | |
| | | Current | Current / All | Current / All | Current /
All | Current / All |

|RD IO | 4.000K| 11.83K/ 9.745K| 46.21M/ 38.06M| 2.519M/
36.74M| 9.612G/ 140.1G|

|WR IO | 4.000K| 11.83K/ 9.745K| 46.21M/ 38.06M| 2.519M/
36.74M| 9.612G/ 140.1G|

```

```

I/O Violation Count |
```

```

| Cmd | Timebase | Total | Violations | Violations |
Violations | | | | |
| | | Violations | ECT | FRT | Pending
IOs
RD	Current		/ 2.000m	/ 1.750m	/
	ALL	10	5 / 2.000m	5 / 1.750m	/

| WR | Current | | / 2.000m | / 1.750m | /
| | | | | |
| | ALL | 23 | 12 / 2.000m | 11 / 1.750m | /
```

```

I/O Exceptions Count |
```

```

| RD| | | | SCSI | NVME |
| |
| | Timebase | Total
|-----|
| WR| | Exceptions | TIMEOUT | BUSY | RESERVE | CHECK | NAME SPACE | RESERVE |
OTHERS |
| | | | | | CONFLICT | CONDITION | NOT READY | CONFLICT |
RD	Current							
```



```

| | All | | | | | | | |
|

| WR| Current | | | | | | | |
|
| | All | | | | | | | |
|

```

In this example, both source and destination are specified for IT flows on a Gen 6 switch.

```
switch:admin> flow --show sys_flow_monitor -srcdev 295200 -dstdev 9f0600
```

```

=====
Name : sys_flow_monitor
Definition : SrcDev(295200) DstDev(9f0600)
Port Speed : SrcDev(32G), DstDev(16G)
Active Flow : 1
Timebase : 5Min
Protocol : SCSI
Port : 82

```

```

Monitor Time : | Mon Aug 16 04:25:02 MDT 2021 |
RD Elapsed Time : | 20s |
WR Elapsed Time : | 20s |

```

```

FC Statistics |
```

```

| Rx Frames Count | Rx Frames per Second | Rx Bytes Count | Rx Throughput (BPS) | Avg Rx Frm Sz(Bytes) |

| 1.22M | 121.03k | 2.10G | 213.67M | 1852 |

```

```

I/O Latency |
```

```

| Metric | Max | Avg |
| Current / All | Current / All |

|RD Completion Time | 136u/ 17.59m| 65u/ 65u|

|WR Completion Time | 168u/ 17.78m| 79u/ 80u|

|RD 1st Response Time | 134u/ 17.57m| 34u/ 139u|

|WR 1st Response Time | 100u/ 17.74m| 32u/ 115u|

|RD Pending IOs | 3 / 3 | 1 / 1 |

|WR Pending IOs | 2 / 2 | 1 / 1 |

=====

```

| I/O Performance |           |             |  |               |        |               |        |               |  |
|-----------------|-----------|-------------|--|---------------|--------|---------------|--------|---------------|--|
| Metric          |           | Avg IO Size |  | Avg (IOPS)    |        | Avg (BPS)     |        | Total IO      |  |
| Total Bytes     |           | Current     |  | Current / All |        | Current / All |        | Current / All |  |
| All             | Current / | All         |  |               |        |               |        |               |  |
| -----           |           |             |  |               |        |               |        |               |  |
| RD IO           |           | 116.7M      |  | 7.186K/       | 7.046K | 819.4G/       | 550.1G | 143.7K/       |  |
| 1.974G          | 16.00T/   | 146.9P      |  |               |        |               |        |               |  |
| -----           |           |             |  |               |        |               |        |               |  |
| WR IO           |           | 61.68M      |  | 6.803K/       | 6.673K | 409.8G/       | 550.2G | 136.0K/       |  |
| 1.869G          | 8.003T/   | 147.0P      |  |               |        |               |        |               |  |
| -----           |           |             |  |               |        |               |        |               |  |

To display the top 10 First Response Time (FRT) flows for a selected metric in ascending order, enter the flow `--show sys_flow_monitor -top frt` command. You can use this command with ECT, BPS, IOPS, and pending IOs metrics.

The following example shows the top 10 FRT metrics.

```
switch:admin> flow --show sys_flow_monitor -top frt
```

```
Name : sys_flow_monitor
```

```
VM Mode : Yes
```

```
Monitored port(s): (*)
```

```
Host Port Count : 6
```

```
Target Port Count: 2
```

```
SCSI : 48
```

```
NVMe : 2
```

```
No of IT : 50
```

```

```

```
TOP 10 ECT |
```

| Bytes/sec     |             | RD          |            | Max Time    |            | Max Time    |               | Avg IOs/sec   |               | Avg         |  |
|---------------|-------------|-------------|------------|-------------|------------|-------------|---------------|---------------|---------------|-------------|--|
| Max Count     |             | Total Count |            | Total Count |            | Total Count |               | Total Count   |               | Total Count |  |
| SID           | DID         | WR          | Device     | Flow        | Port       | Proto       | ECT           | FRT           | IOPS          |             |  |
| BPS           | Pending IOs | Exceptions  | Violations |             |            |             |               |               |               |             |  |
| Type          | Hierarchy   | Index       |            | 5min / All  | 5min / All | 5min / All  |               |               |               |             |  |
| 5min / All    | 5min / All  | 5min / All  | 5min / All |             |            |             |               |               |               |             |  |
| d71a00        | 6d1201      | RD          | T          | IT->N       | 18         | NVMe        | 12.21m/12.21m | 12.21m/12.21m | 8.462K/8.491K |             |  |
| 4.131M/4.145M | 5 /         | 5           | /          | /           | /          | /           |               |               |               |             |  |
| d71a00        | 6d1200      | RD          | T          | IT->N       | 18         | NVMe        | 12.04m/12.04m | 12.04m/12.04m | 8.472K/8.490K |             |  |
| 4.136M/4.145M | 6 /         | 6           | /          | /           | /          | /           |               |               |               |             |  |
| d71a00        | 6d1201      | WR          | T          | IT->N       | 18         | NVMe        | 11.70m/11.70m | 11.69m/11.69m | 8.466K/8.490K |             |  |
| 4.133M/4.145M | 5 /         | 5           | /          | /           | /          | /           |               |               |               |             |  |
| d71a00        | 6d1200      | WR          | T          | IT->N       | 18         | NVMe        | 11.49m/11.49m | 11.48m/11.48m | 8.456K/8.491K |             |  |
| 4.128M/4.145M | 6 /         | 6           | /          | /           | /          | /           |               |               |               |             |  |
| 6d2600        | 830700      | RD          | I          | IT->L       | 38         | SCSI        | 991u/ 991u    | 990u/ 990u    | 2.947K/2.999K |             |  |
| 11.51M/11.71M | 2 /         | 2           | /          | /           | /          | /           |               |               |               |             |  |

```

| 6d0100| 830600| RD| I | IT->L| 1| SCSI| 990u/37.01m| 989u/37.01m| 1.963K/1.955K|
7.670M/7.638M| 2 / 2 | / | / |
| 6d0300| 830400| RD| I | IT->L| 3| SCSI| 990u/ 991u| 988u/ 989u| 5.723K/5.753K|
22.35M/22.47M| 2 / 2 | / | / |
| 6d2100| 830600| RD| I | IT->L| 33| SCSI| 988u/19.47m| 987u/19.47m| 1.970K/2.001K|
7.696M/7.819M| 2 / 2 | / | / |

```

The following example shows the bottom 10 FRT metrics.

```
switch:admin> flow --show sys_flow_monitor -bottom ect
```

```
=====
Name : sys_flow_monitor

```

```
VM Mode : No
```

```
Monitored port(s): (*)
```

```
Host Port Count : 1
```

```
Target Port Count: 0
```

```
SCSI : 1
```

```
NVMe : 0
```

```
No of IT : 1
```

```

```

```
TOP 10 ECT |
```

```

| | RD| | | | | | | | |
Bytes/sec| Max Count | Total Count | Total Count | | | | | |
| SID | DID | WR| Device | Flow | Port |Proto| ECT | FRT | IOPS |
BPS | Pending IOs | Exceptions | Violations |
| | | Type | Hierarchy | Index| | 5min / All | 5min / All | 5min / All |
5min / All | 5min / All | 5min / All | 5min / All |

```

```

| 632100| 718640| RD| I | IT->L| 33| SCSI| 931u/8.388s| 886u/2.097s| 12.16K/10.69K|
47.51M/41.79M| 7 / 7 | / 4 | / |
| 632100| 718640| WR| I | IT->L| 33| SCSI| 947u/8.388s| 895u/2.097s| 12.16K/10.69K|
47.51M/41.79M| 7 / 8 | / | / |

```

The following example shows the top pending IO metrics.

```
switch:admin> flow --show sys_flow_monitor -top pending_io -n 2
```

```
=====
Name : sys_flow_monitor

```

```
VM Mode : No
```

```
Monitored port(s): (*)
```

```
Host Port Count : 1
```

```
Target Port Count: 0
```

```
SCSI : 1
```

```
NVMe : 0
```

```
No of IT : 1
```

```

```

```
TOP 10 ECT |
```

| Bytes/sec     |              | Max Count  | Total Count | Total Count | Max Time   |            | Max Time    | Avg I/Os/sec | Avg           |
|---------------|--------------|------------|-------------|-------------|------------|------------|-------------|--------------|---------------|
| SID           | DID          | WR         | Device      | Flow        | Port       | Proto      | ECT         | FRT          | IOPS          |
| BPS           | Pending I/Os | Exceptions | Violations  | Type        | Hierarchy  | Index      | 5min / All  | 5min / All   | 5min / All    |
| 5min / All    | 5min / All   | 5min / All | 5min / All  | 5min / All  | 5min / All | 5min / All | 5min / All  | 5min / All   | 5min / All    |
| 632100        | 718640       | RD         | I           | IT->L       | 33         | SCSI       | 112u/8.388s | 91u/2.097s   | 10.89K/10.68K |
| 42.56M/41.74M | 7 /          | 7          | /           | 4           | /          | /          |             |              |               |
| 632100        | 718640       | WR         | I           | IT->L       | 33         | SCSI       | 130u/8.388s | 104u/2.097s  | 10.89K/10.68K |
| 42.56M/41.74M | 7 /          | 8          | /           | /           | /          | /          |             |              |               |

## Flow Monitoring with IO Insight Capability

Flow Monitor supports SCSI I/O command latency monitoring. SCSI I/O latency is the time that it takes for an I/O operation to complete between an initiator and a target.

Flow Performance Monitor allows the flexibility to select from predefined flows, static flows, or learning flows for monitoring purposes. It provides both Fibre Channel-level statistics and SCSI I/O operations per second (IOPS) performance counters. Fabric OS software enables flow latency monitoring at the SCSI I/O exchange level using a new set of counters. Flow Monitor supports the I/O latency metrics only on Gen 6-based and Gen 7-based F\_Ports.

The following table describes the platform support for Flow Vision and IO Insight in Fabric OS 9.1.x.

**Table 119: Platform Support Matrix for Flow Vision and IO Insight**

| Brocade Platforms     | Brocade Blades                                     | Flow Vision | IO Insight    |
|-----------------------|----------------------------------------------------|-------------|---------------|
| Brocade G610 Switch   | N/A                                                | Supported   | Not Supported |
| Brocade G620 Switch   | N/A                                                | Supported   | Supported     |
| Brocade G630 Switch   | N/A                                                | Supported   | Supported     |
| Brocade 7810          | N/A                                                | Supported   | Not Supported |
| Brocade X6-4 and X6-8 | FC32-48<br>FC32-64<br>SX6                          | Supported   | Supported     |
| Brocade G720          | N/A                                                | Supported   | Supported     |
| Brocade G730          | N/A                                                | Supported   | Supported     |
| Brocade X7-4 and X7-8 | FC32-48<br>FC32-64<br>FC32-X7-48<br>FC64-48<br>SX6 | Supported   | Supported     |

## I/O Statistics Counters

Flow Monitor provides the following counters for monitoring I/O statistics:

- **First Data Time:** The time difference of either of the situations:

- When the read command frame is transmitted and the first data frame is received.
- When the write command frame is transmitted and the first Xfer Ready frame is received.
- **Status Time:** The time difference of either of the situations:
  - When the read command frame is transmitted and the status frame is received.
  - When the write command frame is transmitted and the status frame is received.
- **Pending I/Os (Queue Depth):** The average number of outstanding I/O operations to be completed.

## Supported Flow Types for SCSI I/O Latency Statistics

The following flow types are supported by SCSI I/O latency statistics:

- Initiator-Target (IT) flow: IO Insight metrics are provided for SCSI READ and WRITE commands with SID and DID parameters.
  - An IT flow monitors traffic between any source device and a Gen 6 or Gen 7 destination device.
  - The source or destination device can be in an unspecified condition (DC).
  - A fully qualified IT flow uses fixed `srcdev` and `dstdev` options.
  - A fully qualified IT flow shows complete IOPS metrics, but partially defined IT flows do not show SCSI statistics counters.
- Initiator-Target-LUN (ITL) flow: IO Insight metrics are provided for SCSI READ and WRITE commands with SID, DID, and LUN parameters.
  - An ITL flow monitors traffic between a source device and an individual LUN on a destination device.
  - The source or destination device can be in an unspecified condition.
  - Flow learning is supported in Fabric OS 9.1.x using the `flow --show sys_flow_monitor -srcdev <devId> -dstdev <devId> -lun "*" and flow --show sys_flow_monitor -srcdev <devId> -dstdev <devId> -nsid "*" commands.`
  - ITL flows do not have Fibre Channel statistics.
  - A fully qualified ITL flow uses fixed `srcdev`, `dstdev`, and LUN options.
  - A fully qualified ITL flow shows complete IOPS metrics, but partially defined ITL flows do not show SCSI statistics counters.

## SCSI Exception Metrics

Brocade Gen 7 platforms capture frames of other commands that are non-READ and non-WRITE commands with non-zero status counts known as *exceptions*. For each learned flow, either SCSI or NVMe, the ASIC provides five programmable opcode counters and a sixth counter with the aggregated count of nonprogrammable opcode counters. These aggregated exception metrics are monitored at the level of flow (IT/ITL/ITN/VMITL/VMITN), port, and device.

The following table describes how the opcode counters are shared optimally for SCSI and NVMe flows.

**Table 120: SCSI and NVMe Opcode Counters**

| Flows       | Programmable Opcode Counters                                                                              |
|-------------|-----------------------------------------------------------------------------------------------------------|
| SCSI        | Default – 5 SCSI opcodes + 1 aggregated count                                                             |
| SCSI + NVMe | 3 SCSI opcodes + 2 NVMe opcodes + 1 aggregated count (even if a single NVMe flow is learned on the ASIC). |
| NVMe only   | 5 NVMe opcodes + 1 aggregated count                                                                       |

The following table describes the support for exception metrics across the flow, port, device, and SCSI/NVMe protocols. These exception metrics are supported only on Gen 7 platforms.

**Table 121: Support for Exception Counts**


| Total Exception Count | Flow | Port | Device | Protocol |      |
|-----------------------|------|------|--------|----------|------|
|                       |      |      |        | SCSI     | NVMe |
| BUSY                  | Y    | Y    | Y      | Y        | N    |
| RESERVE CONFLICT      | Y    | Y    | Y      | Y        | N    |
| ABORTS                | Y    | Y    | Y      | Y        | Y    |
| CHECK CONDITION       | Y    | Y    | Y      | Y        | N    |
| TASK ABORTED          | Y    | Y    | Y      | Y        | N    |
| NAME SPACE NOT READY  | Y    | Y    | Y      | N        | Y    |
| LBA OUT OF RANGE      | Y    | Y    | Y      | N        | Y    |
| CAPACITY EXCEEDED     | Y    | Y    | Y      | N        | Y    |
| FORMAT IN PROGRESS    | Y    | Y    | Y      | N        | Y    |
| TIMEOUT               | Y    | Y    | Y      | Y        | Y    |

## SCSI I/O Operation Classification

I/O operations are classified based on type and size, resulting in various combinations for which I/O latency metrics are provided. SCSI I/O supports the SCSI READ and SCSI WRITE operation types.

### SCSI Read I/O


The following figure shows the frame exchange for a SCSI Read I/O sequence.

**Figure 44: SCSI Read I/O Latency**

- **Read First Data Time:** The time difference between when the read command frame is transmitted and when the first data frame is returned.
- **Read Status Time:** The time difference between when the read command frame is transmitted and when the status frame is returned.
- **Pending I/Os (Queue Depth):** The average number of outstanding read operations to be completed.

#### SCSI Write I/O

The following figure shows the frame exchange for a SCSI Write I/O sequence.

**Figure 45: SCSI Write I/O Latency**

- **Write First Data Time:** The time difference between when the write command frame is transmitted and when the Xfer Ready frame is returned.
- **Write Status Time:** The time difference between when the write command frame is transmitted and when the status frame is returned.
- **Pending I/Os (Queue Depth):** The average number of outstanding write operations to be completed.

## Example of IT Flow Metrics

The following example shows the Fibre Channel, IOPS, and I/O performance and latency metrics in the output of the `flow --show sys_flow_monitor` command for IT flows for which the ingress port (IngPort), the source device (SrcDev), and the target or destination device (DstDev) are fixed.

I/O latency metrics are recorded in units of milliseconds (m) and microseconds ( $\mu$ ).

### NOTE

If you install the monitor on the ingress port of a Gen 6-based fixed-port switch, the I/O performance metrics and the I/O latency metrics are not displayed. However, the Fibre Channel statistics and the SCSI statistics are displayed.

```

switch:maintenance-service> flow --show sys_flow_monitor
=====
Name : sys_flow_monitor
VM Mode : Yes
Monitored port(s): (*)

```



Host Port Count : 3  
Target Port Count: 3  
SCSI : 1  
NVMe : 0  
No of IT : 1


-----  
| TOP 10 Threshold Violations |

|             |        |      |        |       |            |             |            |             |            |            |
|-------------|--------|------|--------|-------|------------|-------------|------------|-------------|------------|------------|
|             |        | RD   |        |       |            | Total Count | Violations | Violations  | Violations | Violations |
| Total Count |        |      |        |       |            |             |            |             |            |            |
| Port        | PID    | WR   | Device | Proto | Violations | ECT         | FRT        | Pending IOs | Exceptions |            |
| Errors      |        |      |        |       |            |             |            |             |            |            |
| Index       |        | Type |        |       | 5min / All | 5min / All  | 5min / All | 5min / All  | 5min / All |            |
| 5min / All  |        |      |        |       |            |             |            |             |            |            |
| -----       |        |      |        |       |            |             |            |             |            |            |
| 59          | 023b00 | RD   | I      | SCSI  | /1.529K    | / 720       | / 809      | /           | /          |            |
| /           |        |      |        |       |            |             |            |             |            |            |


Supported NVMe I/O Latency Statistics

NVMe I/O latency is the time that is taken by an I/O operation to complete between an initiator and NVMe-based storage. The I/O latency statistics support the following types metrics:

- **NVMe Read I/O Latency**



- **Read First Data Time:** The time difference between when the read command frame is transmitted and when the first data frame is returned.
- **Read Status Time:** The time difference between when the read command frame is transmitted and when the status frame is returned.
- **Pending I/Os (Queue Depth):** The average number of outstanding read operations to be completed.
- **NVMe Write I/O Latency**




- **Write First Data Time**: The time difference between when the write command frame is transmitted and when the Xfer Ready frame is returned.
- **Write Status Time**: The time difference between when the write command frame is transmitted and when the status frame is returned.
- **Pending I/Os (Queue Depth)**: The average number of outstanding write operations to be completed.

## Flow Monitoring with VM Insight Capability

The VM Insight feature allows you to monitor individual traffic flows from the source to the destination target for multiple virtual machine (VM) instances that originate from a single N\_Port ID (PID). The following figure illustrates these flows.

**Figure 46: Monitoring VM Flows**



Statistics that are gathered by monitoring the application-to-target performance path in the fabric using predefined flows allow you to perform the following tasks:

- Isolate VMs that are being impacted by slow-draining devices, problematic cables, faulty SFPs, or other factors.
- Determine whether VM performance issues are originating in the host, fabric, or storage array.
- Gather detailed statistics on target devices to plan the allocation of VMs per node or to determine when VMs must be migrated for load balancing based on input/output operations per second (IOPS) or historical latency measurements.

VM Insight identifies VM traffic flows by tagging frames in each flow with an application header containing a unique *application ID*. This ID is based on the entity ID assigned to the VM instance by the hypervisor, such as the UUID assigned by VMware, and the PID assigned by the fabric to the HBA initiator N\_Port.

### NOTE

When defining a flow, you use the entity ID for the VM instance and not the application ID.

The application server is a key component of VM Insight. Based on a new FC standard (FC-GS), application server allows entities (VMs) to be registered with the fabric. An application server stores entity IDs for VM instances that are registered by HBAs, and it allocates application IDs to HBAs for identifying VM traffic in the fabric. The application server is a distributed database, similar to the name server and management server, and database is propagated to all switches in the fabric that support the application server.

You can use the `appserver` command to display contents of the application server database, including N\_Port IDs, application IDs, PIDs, and domain information associated with registered entity IDs. Refer to the *Brocade Fabric OS Command Reference Manual* for details.

The following steps summarize the process used by the fabric to identify VM traffic frames with the application ID:

1. The HBA (initiator) logs on to the fabric, and the fabric assigns a PID to the N\_Port.
2. The HBA fetches the entity IDs for each VM instance from the VM traffic flows and registers these entity IDs with the application server.
3. The application server returns a unique application ID for each VM instance to the HBA.
4. The HBA tags each traffic flow for a specific VM with an application header that contains the application ID.

Created VM flows can also be imported to MAPS, which can monitor statistics and generate alerts for defined out-of-bound values. For details, refer to the *Brocade Fabric OS MAPS User Guide*.

## VM Insight Requirements

The following are requirements for VM Insight:

- The initiator HBA must support VM Insight by communicating with the application server to register VM entity IDs and tagging VM traffic with an application header that contains unique application IDs. Contact your Brocade representative for the latest information on HBA products and firmware with VM Insight support.
- The target must support VM Insight by consuming a frame that is tagged with the application header containing the application ID and sending responses using this ID as the destination. Contact your Brocade representative for the latest information on storage devices and firmware with VM Insight support.
- VMware ESXi 7 supports NVMe over FC. Fabric OS application server and Flow Vision support VM Insight for FC-NVMe devices using the `sys_flow_monitor` flow for FCP device VM Insight support. Contact your adapter vendor to determine support for your specific adapter.
- Hypervisor support must be through VMware vSphere Hypervisor 5.5, 6.x, and later. Contact your Brocade representative for the latest support details.

## Flow Definition Metrics for VM Insight

In general, the same type of statistics are displayed for VM flows as for other Flow Monitor flows, but statistics are tracked by entity ID instead of port ID. The following metrics are reported from VM Insight flows.

Fibre Channel Statistics:

- **Frame:** Frame count, rate, and average frame size.
- **Throughput:** Frame count and throughput (bytes per sec).

SCSI I/O statistics: Brief or detailed I/O statistics are reported from VM Insight flows, depending on the platform (fixed-port switch or chassis) and the configured flow monitor:

- **Detailed Statistics:** Classified by the type of operation (READ/WRITE) and the size of the data block.
- **Brief Statistics:** Output for learning flows defined with an asterisk ( \* ) for a target device. Only maximum and average I/O data are reported for the last 6 seconds and the lifetime of the flows. Statistics for various sizes of data blocks are not reported.
- **I/O Performance Statistics:** I/O count and I/O operations per second (IOPS). This includes average and maximum data. Latency is measured from 6 seconds to the lifetime of the flow.
- **I/O Latency Statistics:** Queue depth. Command status latency and command-first response latency. This includes average and maximum data. Latency is measured from 6 seconds to the lifetime of the flow.

## Flow Monitoring with VMID+

From Fabric OS 9.1.x, Flow Vision and VM Insight on Gen 7 platforms support the VMID+ feature. VMID+ does not require any storage array support for VM Insight to capture VM flow metrics, whereas VM Insight on Gen 6 platforms requires storage array support for VM Insight.

VM Insight allows flows between a virtual machine (VM) and target to be tracked and identified with high granularity. VM instances can include an optional header in FC frames that contains a virtual machine identification (VMID) known as an *application header*. The application header identifies the specific VM instance that is initiating the I/O and subsequent IOs that are destined to the target. However, VMs tag the application header only if the target supports it. The application header support is advertised during the device discovery phase (PLOGI exchange). VMID+ helps to emulate a VM-capable target to the host even if the target is not VM-capable. Hereafter, VMs tag the application header in all the I/Os destined to the non-VM target. However, the application header is removed in the egress of the target F\_Ports, and the I/O exchanges behave identically to the non-VMID tagged exchanges. The endpoint switches can track the VM statistics through the VM Insight feature.

You can view the VM Insight metrics only on Gen 7 endpoints when the VMID+ is enabled on the target F\_Port. You can enable VMID+ for a maximum of 4 ports per logical switch, 16 ports for a non-VF switch, and 64 ports for non-VF chassis-based platforms.

For more information on CLI commands related to the application header (`portcfgappheader`), refer to *Brocade Fabric OS Command Reference Manual, 9.1.x*.

The following table describes the support for Flow Vision and VM Insight on various platforms that run Fabric OS 9.1.x.

**Table 122: Platform Support Matrix for Flow Vision and VM Insight**

| Brocade Platforms     | Blades                                      | Flow Vision | VM Insight                                                   |
|-----------------------|---------------------------------------------|-------------|--------------------------------------------------------------|
| Brocade G610 Switch   | N/A                                         | Supported   | Not Supported                                                |
| Brocade G620 Switch   | N/A                                         | Supported   | Supported (VMID)                                             |
| Brocade G630 Switch   | N/A                                         | Supported   | Supported (VMID)                                             |
| Brocade 7810          | N/A                                         | Supported   | Not Supported                                                |
| Brocade X6-4 and X6-8 | FC32-48<br>FC32-64                          | Supported   | Supported (VMID)                                             |
| Brocade G720          | N/A                                         | Supported   | Supported (VMID and VMID+)                                   |
| Brocade G730          | N/A                                         | Supported   | Supported (VMID and VMID+)                                   |
| Brocade X7-4 and X7-8 | FC32-48<br>FC32-64<br>FC32-X7-48<br>FC64-48 | Supported   | Supported<br>(VMID+ is supported only on the FC64-48 blades) |

### Limitations and Restrictions of VMID+

- The VMID+ configuration should not be enabled on ports that are connected with application ID-capable devices.
- The VMID+ is not supported on any ports that are connected to an Access Gateway. When a port connected to an AG is attempted to be configured with a VMID+, a RASLog message is generated, and the VMID+ configuration succeeds. If an AG is connected to a switch port that is already configured with the VMID+, then the VMID+ does not work, and no RASLog messages are generated.
- Enabling a VMID+ is a non-disruptive operation. However, the port should be toggled to allow for the application ID configuration to take effect. Disabling the application header is a disruptive process, and the associated port is implicitly disabled during the disable operation.
- The VMID+ feature is mutually exclusive with extended ISL (XISL) usage on logical switches. Therefore, you should disable the VMID+ to allow the XISL-usage feature to work, and you should turn off the XISL usage feature for the

VMID+ to work. For more information on configuring XISL usage, refer to the "Configuring a Logical Switch for XISL Use" section of the *Brocade Fabric OS Administration Guide, 9.1.x*.

- The VMID+ feature is not affected if an F\_Port is converted to other port types after enabling the application header. For example, a non-F\_Port or D\_Port.
- The VMID+ feature is not supported with Fibre Channel Router (FCR).
- If the VMID+ is enabled on a port in VF mode, then moving the port out of the logical switch is not allowed. Therefore, disable the VMID+ before moving the port out of the logical switch.
- Configuring an EX\_Port and F\_Port with the application header on the same chassis is not supported in VF and non-VF mode. However, the configuration is not blocked.
- The VMID+ feature is not supported on FICON logical switch ports.
- The configuration should be same on all trunk ports on an F\_Port trunk to form a new trunk. Ensure that the same configuration is enabled on all ports of an existing trunk as well.
- Enabling the VMID+ configuration on F\_Ports connected to encryption-supported third-party devices is not supported.

## Working with Flow Monitor

This section describes the various CLI commands that are associated with the `sys_flow_monitor` automated flow, and it includes examples.

### Working with Flows

You can activate the `sys_flow_monitor` flow, you can deactivate it, and you can display its output. You can also reset flow logs and monitor flow resources.

**Table 123: Valid Commands for `sys_flow_monitor`**

| Action              | Command                                                                                                             |
|---------------------|---------------------------------------------------------------------------------------------------------------------|
| Activate the flow   | <code>flow --activate sys_flow_monitor</code>                                                                       |
| Deactivate the flow | <code>flow --deactivate sys_flow_monitor</code>                                                                     |
| Modify the flow     | <code>flow --modify sys_flow_monitor -port &lt;&gt;</code>                                                          |
| Reset the flow      | <code>flow --reset sys_flow_monitor</code>                                                                          |
| Resource            | <code>flow --show sys_flow_monitor -resource</code>                                                                 |
| Show                | <code>flow -show sys_flow_monitor</code>                                                                            |
| Control             | <code>flow --control sys_flow_monitor</code>                                                                        |
| VM mode             | <code>flow --modify sys_flow_monitor -vmenable</code> and<br><code>flow --modify sys_flow_monitor -vmdisable</code> |

The following considerations apply to this table:

- Because `sys_flow_monitor` is a predefined flow, it does not support the `create` and `delete` options for the `flow` command.
- A `sys_flow_monitor` flow is automatically activated on upgrade, and its setting is retained after a reboot or failover if no other user-defined or predefined flows for monitor and mirror are active.
- Upon activation, a predefined flow automatically discovers application flows based on the direction and port parameters. Deactivation is a logical switch operation; there is no deactivate support at the chassis level.

### Activating `sys_flow_monitor` on Host F\_Ports and on a Logical Group

All Host F\_Ports

```
switch:admin> flow --deactivate sys_flow_monitor
```

```

Monitor feature(s) have been deactivated.
switch:admin> flow --modify sys_flow_monitor -port all_host_ports
switch:admin> flow --activate sys_flow_monitor
Monitor feature(s) have been activated.

```

### Single F\_Port

```

switch:admin> flow --deactivate sys_flow_monitor
Monitor feature(s) have been deactivated.
switch:admin> flow --modify sys_flow_monitor -port 1
switch:admin> flow --activate sys_flow_monitor
Monitor feature(s) have been activated

```

### MAPS Logical Group

```

switch:admin> flow --deactivate sys_flow_monitor
Monitor feature(s) have been deactivated.
switch:admin> flow --modify sys_flow_monitor -port "*"
switch:admin> flow --activate sys_flow_monitor
Monitor feature(s) have been activated.
switch:admin> flow --modify sys_flow_monitor -port "*"
Operation not supported on Active flow.

```

### Resetting a Predefined Flow

Resetting is a logical switch operation; there is no reset support at the chassis level. This resets all flows in hardware; new flows are learned.

#### NOTE

Learning the same flow between resets depends on resource availability.

```
switch:admin> flow --reset sys_flow_monitor
```

### Monitoring Flow Resources

The following example shows how to display the total number of IT flows learned by the flow `sys_flow_monitor`.

```

switch:admin> flow --show sys_flow_monitor -resource
=====
Flows | Current Configuration | Chassis Maximum |
 | Used |
=====
IT count | 1 | 8192 |
ITL count | 4 | 16384 |
Total | 5 | 24576 |
=====
*ITL limits are derived based on IT reservation limit
Timebase:

=====
Time Bases | Used | Max |
=====
10 sec | 0 | 25 |
5 min | 1 | 8192 |
6 hr | 4 | 16384 |
=====

```

Slot Resource:



| =====                                                  |  |           |  |                     |  |  |                     |  |  |               |  |  |
|--------------------------------------------------------|--|-----------|--|---------------------|--|--|---------------------|--|--|---------------|--|--|
| Slot                                                   |  | Total     |  | Initiator           |  |  | Target              |  |  | IT            |  |  |
| ITL                                                    |  |           |  |                     |  |  |                     |  |  |               |  |  |
| No                                                     |  | F-Port(s) |  | SCSI / NVMe / Total |  |  | SCSI / NVMe / Total |  |  | Current / Max |  |  |
| Current /                                              |  | Max       |  |                     |  |  |                     |  |  |               |  |  |
| =====                                                  |  |           |  |                     |  |  |                     |  |  |               |  |  |
| 3                                                      |  | 1         |  | 1 / 0 / 1           |  |  | 0 / 0 / 0           |  |  | 1 / 4096      |  |  |
| 4 /                                                    |  | 8192      |  |                     |  |  |                     |  |  |               |  |  |
| -----                                                  |  |           |  |                     |  |  |                     |  |  |               |  |  |
| 4                                                      |  | 0         |  | 0 / 0 / 0           |  |  | 0 / 0 / 0           |  |  | 0 / 4096      |  |  |
| 0 /                                                    |  | 8192      |  |                     |  |  |                     |  |  |               |  |  |
| -----                                                  |  |           |  |                     |  |  |                     |  |  |               |  |  |
| * These values are system defined and not configurable |  |           |  |                     |  |  |                     |  |  |               |  |  |

Enabling or Disabling VM Monitoring

You can modify the sys\_flow\_monitor to enable or disable the monitoring of applications that are running behind the VM instances.

To enable the VM monitoring, enter the flow --modify sys\_flow\_monitor -vmenable command. The following example shows the output when a VM instance is enabled.


```
switch:admin> flow --modify sys_flow_monitor -vmenable
VM learning Enabled for the flow sys_flow_monitor.
```

To disable the VM monitoring, enter the flow --modify sys\_flow\_monitor -vmdisable command. The following example shows the output when a VM instance is disabled.

```
switch:admin> flow --modify sys_flow_monitor -vmdisable
VM learning Disabled for the flow sys_flow_monitor.
```

CLI Workflow

In the following figure, you can see that the detailed and list views in Fabric OS 9.1.x are provided only for IT flows and initiator ports.

**Figure 47: CLI Workflow**

## View Options

The following are the view options of the `flow --show` command.

### NOTE

IT and ITL flows are monitored on Gen 7 chassis systems. Gen 6 chassis systems are limited to only IT flows.

### IT Detail View

- IT flows can be monitored at 5-min (default) granularity.
- Each monitored flow has two samples (Previous and Current) that are aligned with the system clock. After the 10-sec, 5-min, or 6-hr timer expires, a snapshot of the statistics is taken. This snapshot is associated with the Previous sample, whereas the Current sample starts over. The snapshot is used for streaming, and the statistics that are associated with the Current sample are used in the CLI. At any given point, the time duration of the Current statistics is always less than that of the Monitored time (the time when the last statistics were updated).
- Monitored and Elapsed time (the amount of time that has passed) are used to indicate the details on the Flow Monitor duration. When flows move across granularities, the CLI reflects that the Elapsed time might be more or less than the Monitored time.
- A consequence of flexible granularity is that a flow can move between two timebases in consecutive CLI requests.

The following example shows the detailed view of the ITL flow with the complete details of the traffic between a source device and an individual LUN on a destination device.

```
switch:admin> flow --show sys_flow_monitor -srcdev 632100 -dstdev 718640
```

```
=====
Name : sys_flow_monitor
Definition : SrcDev(632100) DstDev(718640)
Port Speed : SrcDev(32G), DstDev(16G)
Active Flow : 1
Timebase : 5Min
Protocol : SCSI
Port : 390
=====
```

```

Monitor Time : | Mon Aug 16 04:23:26 MDT 2021 |
RD Elapsed Time : | 3m:33s |
WR Elapsed Time : | 3m:33s |

```

```

FC Statistics |

```

```

| Frame Count | Frames Per Second | Byte Count | Throughput (BPS) | Frame
| Size(Bytes) |
| Tx / Rx / Total | Tx / Rx / Total | Tx / Rx / Total | Tx / Rx / Total | Tx /
Rx
482.14k/602.70k/ 1.08M	47.73k/ 59.67k/107.40k	498.42M/498.44M/996.87M	49.34M/ 49.35M/ 98.69M	1084/ 868

```

```

I/O Latency |

```

```

| Metric | Max | Avg |
| Current / All | Current / All |

|RD Completion Time | 1.677m/ 9.114m| 35u/ 34u|

|WR Completion Time | 1.700m/ 8.388s| 43u/ 43u|

|RD 1st Response Time | 1.631m/ 9.065m| 19u/ 19u|

|WR 1st Response Time | 1.642m/ 9.031m| 16u/ 16u|

|RD Pending IOs | 7 / 9 | 1 / 1 |

|WR Pending IOs | 7 / 9 | 1 / 1 |

=====

```

```

I/O Performance |

```

```

| Metric | Avg IO Size | Avg (IOPS) | Avg (BPS) | Total IO
| Total Bytes |
| Current | Current / All | Current / All | Current /
| All | Current / All |

|RD IO | 4.000K| 11.83K/ 9.745K| 46.21M/ 38.06M| 2.519M/
| 36.74M| 9.612G/ 140.1G|

|WR IO | 4.000K| 11.83K/ 9.745K| 46.21M/ 38.06M| 2.519M/
| 36.74M| 9.612G/ 140.1G|

```

```

I/O Violation Count |

```

| Cmd                  | Timebase   | Total      | Violations  | Violations  |           |
|----------------------|------------|------------|-------------|-------------|-----------|
| Violations           |            | Violations | ECT         | FRT         | Pending   |
| IOs                  |            |            |             |             |           |
| RD                   | Current    |            | / 2.000m    | / 1.750m    | /         |
|                      | ALL        | 10         | 5 / 2.000m  | 5 / 1.750m  | /         |
| WR                   | Current    |            | / 2.000m    | / 1.750m    | /         |
|                      | ALL        | 23         | 12 / 2.000m | 11 / 1.750m | /         |
| -----                |            |            |             |             |           |
| I/O Exceptions Count |            |            |             |             |           |
| RD                   |            |            | SCSI        | NVME        |           |
|                      | Timebase   | Total      |             |             |           |
| WR                   | Exceptions | TIMEOUT    | BUSY        | RESERVE     | CHECK     |
| OTHERS               |            |            | CONFLICT    | CONDITION   | NOT READY |
|                      |            |            |             |             |           |
| RD                   | Current    |            |             |             |           |
|                      | All        |            |             |             |           |
| WR                   | Current    |            |             |             |           |
|                      | All        |            |             |             |           |

## IT List View

- IT flows are monitored at 5-min granularity.
- Information is presented as two tables for grouping frames across granularity.

The following example shows the list view of the IT flow that monitors traffic between a source device and a destination device.

```
switch:admin> flow --show sys_flow_monitor -srcdev '*' -dstdev '*'
=====
Name : sys_flow_monitor
Definition : SrcDev(*)DstDev(*)

Timebase : 5 Minutes
```

| -----             |             |             |             |            |          |             |  |             |          |               |     |
|-------------------|-------------|-------------|-------------|------------|----------|-------------|--|-------------|----------|---------------|-----|
| I/O Brief Metrics |             |             |             |            |          |             |  |             |          |               |     |
| -----             |             |             |             |            |          |             |  |             |          |               |     |
|                   |             | RD          |             |            |          |             |  | Max Time    | Max Time | Avg IOs/sec   | Avg |
| Bytes/sec         | Max Count   | Total Count | Total Count |            |          |             |  |             |          |               |     |
| SID               | DID         | WR  Device  | Flow        | Port       | Proto    | ECT         |  | FRT         |          | IOPS          |     |
| BPS               | Pending IOs | Exceptions  | Violations  |            |          |             |  |             |          |               |     |
|                   |             | Type        | Hierarchy   | Index      |          | 5min / All  |  | 5min / All  |          | 5min / All    |     |
| 5min / All        | 5min / All  | 5min / All  | 5min / All  | 5min / All |          |             |  |             |          |               |     |
| -----             |             |             |             |            |          |             |  |             |          |               |     |
| 632100            | 718640      | RD          | I           | IT->L      | 33  SCSI | 257u/8.388s |  | 244u/2.097s |          | 26.87K/10.77K |     |
| 1.573T/68.90G     | 7 /         | 7           | /           | 4          | /2.754K  |             |  |             |          |               |     |
| 632100            | 718640      | WR          | I           | IT->L      | 33  SCSI | 194u/8.388s |  | 116u/2.097s |          | 26.87K/10.77K |     |
| 1.573T/68.82G     | 7 /         | 8           | /           |            | /2.983K  |             |  |             |          |               |     |
| =====             |             |             |             |            |          |             |  |             |          |               |     |

## ITL/ITN Detail View

- ITL/ITN flows are monitored at 6-hr (default) granularity.
- Monitored and Elapsed time (the amount of time that has passed) indicate the details on the Flow Monitor duration. When flows move across granularities, the CLI reflects that the Elapsed time might be more or less than the Monitored time.
- Each monitored ITL/ITN (with LUN/NSID) flow displays the following performance metrics for a specific host or storage device to diagnose I/O operational issues:
  - Total I/Os
  - First Response Time maximum and average (max/average)
  - I/O latency (Exchange Completion Time or ECT) max/average
  - Outstanding I/Os max/average

The following example shows the detailed view of the ITL flow with the complete details of the traffic between a source device and an individual LUN on a destination device.

```
switch:admin> flow --show sys_flow_monitor -srcdev 632100 -dstdev 718640 -lun 2
```

```
=====
Name : sys_flow_monitor
Definition : SrcDev(632100)DstDev(718640),LUN(2)
Port Speed : SrcDev(32G), DstDev(16G)
Active Flow : 1
Timebase : 6Hrs
Protocol : SCSI
Port : 390
=====
```

```
Monitor Time : | Mon Aug 16 04:26:09 MDT 2021 |
RD Elapsed Time : | 1h:5m |
WR Elapsed Time : | 1h:5m |
=====
```

```
I/O Latency |
```

```
=====
| Metric | Max | Avg |
| Current / All | Current / All |
=====
|RD Completion Time | 8.915m/ 8.915m| 34u/ 34u|
=====
```

|                      |  |         |        |      |     |
|----------------------|--|---------|--------|------|-----|
| WR Completion Time   |  | 9.087m/ | 9.087m | 43u/ | 43u |
| -----                |  |         |        |      |     |
| RD 1st Response Time |  | 8.873m/ | 8.873m | 19u/ | 19u |
| -----                |  |         |        |      |     |
| WR 1st Response Time |  | 9.031m/ | 9.031m | 15u/ | 15u |
| -----                |  |         |        |      |     |
| RD Pending IOs       |  | 4 /     | 4      | 1 /  | 1   |
| -----                |  |         |        |      |     |
| WR Pending IOs       |  | 4 /     | 4      | 1 /  | 1   |
| -----                |  |         |        |      |     |

-----  
I/O Performance |

| Metric      |           | Avg IO Size |        | Avg (IOPS) |        | Avg (BPS) |        | Total IO |  |
|-------------|-----------|-------------|--------|------------|--------|-----------|--------|----------|--|
| Total Bytes |           |             |        |            |        |           |        |          |  |
| Current     |           | Current     |        | All        |        | Current   |        | All      |  |
| All         | Current / | All         |        |            |        |           |        |          |  |
| -----       |           |             |        |            |        |           |        |          |  |
| RD IO       |           |             | 4.000K | 2.430K/    | 2.430K | 9.492M/   | 9.492M | 9.558M/  |  |
| 9.558M      | 36.46G/   | 36.46G      |        |            |        |           |        |          |  |
| -----       |           |             |        |            |        |           |        |          |  |
| WR IO       |           |             | 4.000K | 2.430K/    | 2.430K | 9.492M/   | 9.492M | 9.558M/  |  |
| 9.558M      | 36.46G/   | 36.46G      |        |            |        |           |        |          |  |
| -----       |           |             |        |            |        |           |        |          |  |

-----  
I/O Violation Count |

| Cmd        |  | Timebase |         | Total |   | Violations |        | Violations |        |
|------------|--|----------|---------|-------|---|------------|--------|------------|--------|
| Violations |  |          |         |       |   |            |        |            |        |
| Current    |  | Current  |         | All   |   | Current    |        | All        |        |
| IOs        |  |          |         |       |   |            |        |            |        |
| -----      |  |          |         |       |   |            |        |            |        |
| RD         |  |          | Current |       | 2 | 1 /        | 2.000m | 1 /        | 1.750m |
|            |  |          |         |       |   |            |        |            |        |
|            |  |          | ALL     |       | 2 | 1 /        | 2.000m | 1 /        | 1.750m |
|            |  |          |         |       |   |            |        |            |        |
| -----      |  |          |         |       |   |            |        |            |        |
| WR         |  |          | Current |       | 6 | 3 /        | 2.000m | 3 /        | 1.750m |
|            |  |          |         |       |   |            |        |            |        |
|            |  |          | ALL     |       | 6 | 3 /        | 2.000m | 3 /        | 1.750m |
|            |  |          |         |       |   |            |        |            |        |
| -----      |  |          |         |       |   |            |        |            |        |

-----  
I/O Exceptions Count |

|        |          |            |         |      |         |       |      |          |
|--------|----------|------------|---------|------|---------|-------|------|----------|
| RD     |          | SCSI       |         |      |         |       |      |          |
|        |          |            |         |      |         |       |      |          |
|        | Timebase | Total      |         |      |         |       |      |          |
| -----  |          |            |         |      |         |       |      |          |
| WR     |          | Exceptions | TIMEOUT | BUSY | RESERVE | CHECK | TASK | TASK SET |
| OTHERS |          |            |         |      |         |       |      |          |

|  |    |         |  |  |  |  |          |           |       |  |      |  |
|--|----|---------|--|--|--|--|----------|-----------|-------|--|------|--|
|  |    |         |  |  |  |  | CONFLICT | CONDITION | ABORT |  | FULL |  |
|  |    |         |  |  |  |  |          |           |       |  |      |  |
|  | RD | Current |  |  |  |  |          |           |       |  |      |  |
|  |    |         |  |  |  |  |          |           |       |  |      |  |
|  |    | All     |  |  |  |  |          |           |       |  |      |  |
|  |    |         |  |  |  |  |          |           |       |  |      |  |
|  | WR | Current |  |  |  |  |          |           |       |  |      |  |
|  |    |         |  |  |  |  |          |           |       |  |      |  |
|  |    | All     |  |  |  |  |          |           |       |  |      |  |
|  |    |         |  |  |  |  |          |           |       |  |      |  |

## ITL/ITN List View

- ITL/ITN flows are monitored in real time or 6-hour granularity.
- Information is presented as three tables for grouping frames across granularity.

The following example shows the list view of the ITL flow that monitors traffic between a source device and an individual LUN on a destination device.

```
switch:admin> flow --show sys_flow_monitor -srcdev 632100 -dstdev 718640 -lun '*'
```

```
=====
Name : sys_flow_monitor
Definition : SrcDev(632100)DstDev(718640),LUN(*)
Protocol : SCSI
Port Speed : SrcDev(32G), DstDev(16G)
LUN count : 10
ECT(0) FL(0)
```

```

Timebase : 6 Hours
RD Elapsed Time : | 60s |
WR Elapsed Time : | 60s |

```

```
| I/O Brief Metrics |
```

|  | RD         |             |      | Max Time      | Max Time      | Avg IOs/sec   | Avg Bytes/sec | Max Count   | Total  |
|--|------------|-------------|------|---------------|---------------|---------------|---------------|-------------|--------|
|  | Count      | Total Count |      |               |               |               |               |             |        |
|  | WR         | LUN         | Port | ECT           | FRT           | IOPS          | Data Rate     | Pending IOs |        |
|  | Exceptions | Violations  |      |               |               |               |               |             |        |
|  |            |             |      | 6Hrs / All    | 6Hrs / All    | 6Hrs / All    | 6Hrs / All    | 6Hrs / All  | 6Hrs / |
|  | All        | 6Hrs / All  |      |               |               |               |               |             |        |
|  | RD         | 2           | 390  | 8.915m/8.915m | 8.873m/8.873m | 2.435K/2.435K | 9.511M/9.511M | 4 / 4       | /      |
|  |            | 2 / 2       |      |               |               |               |               |             |        |
|  | WR         | 2           | 390  | 9.087m/9.087m | 9.031m/9.031m | 2.435K/2.435K | 9.511M/9.511M | 4 / 4       | /      |
|  |            | 6 / 6       |      |               |               |               |               |             |        |
|  | RD         | 0           | 390  | 8.771m/8.771m | 8.734m/8.734m | 2.376K/2.376K | 9.282M/9.282M | 4 / 4       | /      |
|  |            | 2 / 2       |      |               |               |               |               |             |        |
|  | WR         | 0           | 390  | 9.034m/9.034m | 8.980m/8.980m | 2.376K/2.376K | 9.282M/9.282M | 3 / 3       | /      |
|  |            | 6 / 6       |      |               |               |               |               |             |        |
|  | RD         | 3           | 390  | 9.097m/9.097m | 9.048m/9.048m | 2.513K/2.513K | 9.816M/9.816M | 4 / 4       | /      |
|  |            | 2 / 2       |      |               |               |               |               |             |        |

|    |       |     |               |               |               |               |       |   |
|----|-------|-----|---------------|---------------|---------------|---------------|-------|---|
| WR | 3     | 390 | 8.388s/8.388s | 8.863m/8.863m | 2.513K/2.513K | 9.816M/9.816M | 4 / 4 | / |
|    | 7 / 7 |     |               |               |               |               |       |   |
| RD | 1     | 390 | 9.114m/9.114m | 9.065m/9.065m | 2.297K/2.297K | 8.975M/8.975M | 4 / 4 | / |
|    | 4 / 4 |     |               |               |               |               |       |   |
| WR | 1     | 390 | 8.995m/8.995m | 8.937m/8.937m | 2.297K/2.297K | 8.975M/8.975M | 4 / 4 | / |
|    | 4 / 4 |     |               |               |               |               |       |   |

## VMITL/VMITN View

- VMITL/VMITN flows are monitored at 6-hour granularity.
- Monitored and Elapsed time (the amount of time that has passed) are used to indicate the details on Flow Monitor duration. When flows move across granularities, the CLI reflects that the Elapsed time might be more or less than the Monitored time.
- Each monitored VMITL/VMITN flow displays the details of the VM instances (with LUN/NSID) on total I/Os, first response time maximum and average (max/average), I/O latency (Exchange Completion Time, or ECT) max/average, and outstanding I/Os max/average performance metrics for a specific host or storage device to diagnose I/O operational issues.

### NOTE

VMITL/VMITN views are available only on Gen 7 chassis systems.

The following example shows the output of the VMITL list view with SID, DID, and LUN parameters.

```
switch:admin> flow --show sys_flow_monitor -srcdev 71af40 -dstdev 718d40 -lun 80 -srceid *
```

---

```
Name : sys_flow_monitor
Definition : SrcDev(71af40)DstDev(718d40),LUN(80),SEID(*)
Protocol : SCSI
Port Speed : SrcDev(32G), DstDev(32G)
```

---

```
Timebase : 6 Hours
```

---

```
| I/O Brief Metrics |
```

---

| Avg I/Os/sec | Avg Bytes/sec | Max Count    |  | RD | No. |     | Max Time | Max Time   |            |
|--------------|---------------|--------------|--|----|-----|-----|----------|------------|------------|
| SEID         |               |              |  | WR | LUN | of  | Port     | ECT        | FRT        |
| IOPS         | Data Rate     | Pending I/Os |  |    |     | VMs |          | 6Hrs / All | 6Hrs / All |
| 6Hrs / All   | 6Hrs / All    | 6Hrs / All   |  |    |     |     |          |            |            |

---

```
| 50 0f 43 27 23 b9 a5 a7-44 a8 10 e9 ec da ca 13| RD| 80| 1| 431| 111.0m/111.0m| 56.46m/56.46m|
309 / 309 | 77.40M/77.40M| 32 / 32 |
| 50 0f 43 27 23 b9 a5 a7-44 a8 10 e9 ec da ca 13| WR| 80| 1| 431| 104.9m/104.9m| 35.79m/35.79m|
309 / 309 | 77.40M/77.40M| 32 / 32 |
```

---

The following example shows the output of the VMITL detail view.

```
switch:admin> flow --show sys_flow_monitor -srcdev 71af40 -dstdev 718d40 -lun 80 -srceid '50 0f 43 27 23 b9 a5
a7-44 a8 10 e9 ec da ca 13'
```

---

```
Name : sys_flow_monitor
```



```

Definition : SrcDev(71af40)DstDev(718d40),LUN(80),SEID(50 0f 43 27 23 b9 a5 a7-44 a8 10 e9 ec da ca 13)
Port Speed : SrcDev(32G), DstDev(32G)
Active Flow : 1
Timebase : 6Hrs
Protocol : SCSI
Port : 431

```

```

Monitor Time : | Mon Aug 16 04:33:34 MDT 2021 |
RD Elapsed Time : | 1h:12m |
WR Elapsed Time : | 1h:12m |

```

#### I/O Latency |

| Metric               | Max     |        | Avg     |        |
|----------------------|---------|--------|---------|--------|
|                      | Current | All    | Current | All    |
| RD Completion Time   | 111.0m/ | 111.0m | 1.245m/ | 1.245m |
| WR Completion Time   | 104.9m/ | 104.9m | 35.78m/ | 35.78m |
| RD 1st Response Time | 56.46m/ | 56.46m | 903u/   | 903u   |
| WR 1st Response Time | 35.79m/ | 35.79m | 210u/   | 210u   |
| RD Pending IOs       | 32 /    | 32     | 5 /     | 5      |
| WR Pending IOs       | 32 /    | 32     | 11 /    | 11     |

#### I/O Performance |

| Metric | Total Bytes | Avg IO Size |     | Avg (IOPS) |        | Avg (BPS) |     | Total IO |     |
|--------|-------------|-------------|-----|------------|--------|-----------|-----|----------|-----|
|        |             | Current     | All | Current    | All    | Current   | All | Current  | All |
|        |             |             |     |            |        |           |     |          |     |
| RD IO  | 256.0K      | 309 /       | 309 | 77.41M/    | 77.41M | 1.356M/   |     |          |     |
| 1.356M | 331.0G/     | 331.0G      |     |            |        |           |     |          |     |
| WR IO  | 256.0K      | 309 /       | 309 | 77.41M/    | 77.41M | 1.355M/   |     |          |     |
| 1.355M | 331.0G/     | 331.0G      |     |            |        |           |     |          |     |

#### I/O Violation Count |

| Cmd | Timebase | Total  | Violations |        | Violations |        |
|-----|----------|--------|------------|--------|------------|--------|
|     |          |        | ECT        | FRT    | Pending    |        |
| RD  | Current  | 514.0K | 275.0K /   | 2.000m | 239.0K /   | 1.750m |

|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|----------------------|----|--|------------|--|---------|--|----------|--|----------|--|-----------|--|------------|--|----------|
|                      |    |  | ALL        |  | 514.0K  |  | 275.0K / |  | 2.000m   |  | 239.0K /  |  | 1.750m     |  | /        |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      | WR |  | Current    |  | 32.94K  |  | 32.93K / |  | 2.000m   |  | 5 /       |  | 1.750m     |  | /        |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  | ALL        |  | 32.94K  |  | 32.93K / |  | 2.000m   |  | 5 /       |  | 1.750m     |  | /        |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
| -----                |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
| I/O Exceptions Count |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      | RD |  |            |  |         |  |          |  | SCSI     |  |           |  | NVME       |  |          |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  | Timebase   |  | Total   |  |          |  |          |  |           |  |            |  |          |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      | WR |  | Exceptions |  | TIMEOUT |  | BUSY     |  | RESERVE  |  | CHECK     |  | NAME SPACE |  | RESERVE  |
|                      |    |  | OTHERS     |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  |            |  |         |  |          |  | CONFLICT |  | CONDITION |  | NOT READY  |  | CONFLICT |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      | RD |  | Current    |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  | All        |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      | WR |  | Current    |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  | All        |  |         |  |          |  |          |  |           |  |            |  |          |
|                      |    |  |            |  |         |  |          |  |          |  |           |  |            |  |          |

The following example shows the output of the VMITN list view with SID, DID, and NSID parameters.

```
switch:admin> flow --show sys_flow_monitor -srcd 290400 -dstd 718d00 -nsid '*'
```

```
=====
Name : sys_flow_monitor
Definition : SrcDev(290400), DstDev(718d00), NSID(*)
Protocol : NVMe
Port Speed : SrcDev(32G), DstDev(32G)
NSID count : 1
ECT(0) FL(0)
```

```
Timebase : 6 Hours
```

```
| I/O Brief Metrics |
```

|  |       |  |             |  |             |  |            |  |            |  |                |  |               |  |           |
|--|-------|--|-------------|--|-------------|--|------------|--|------------|--|----------------|--|---------------|--|-----------|
|  | RD    |  |             |  | Max Time    |  | Max Time   |  | Max Time   |  | Avg IOs/sec    |  | Avg Bytes/sec |  | Max       |
|  | Count |  | Total Count |  | Total Count |  |            |  |            |  |                |  |               |  |           |
|  | WR    |  | NSID        |  | Port        |  | ECT        |  | FRT        |  | Fabric Latency |  | IOPS          |  | Data Rate |
|  |       |  | Exceptions  |  |             |  | Violations |  |            |  |                |  |               |  | Pending   |
|  |       |  |             |  | 6Hrs / All  |  | 6Hrs / All |  | 6Hrs / All |  | 6Hrs / All     |  | 6Hrs / All    |  | 6Hrs /    |
|  | All   |  | 6Hrs / All  |  | 6Hrs / All  |  |            |  |            |  |                |  |               |  |           |

```

| RD| 1| 141| 958u/ 958u| 898u/ 898u| / | 6.775K/6.775K| 208.8M/208.8M| 44 /
 44 | / | / |
| WR| 1| 141| 503u/ 503u| 129u/ 129u| / | 6.621K/6.621K| 206.9M/206.9M| 19 /
 19 | / | / |
=====
| RD| * | * | 958u/ 958u| 898u/ 898u| / | 6.775K/6.775K| 208.8M/208.8M| 44 /
 44 | / | / |
| WR| * | * | 503u/ 503u| 129u/ 129u| / | 6.621K/6.621K| 206.9M/206.9M| 19 /
 19 | / | / |
=====
=====

```

The following example shows the output of the VMITN detail view.

```
switch:admin> flow --show sys_flow_monitor -srcd 290400 -dstd 718d00 -nsid 1
```

```

=====
Name : sys_flow_monitor
Definition : SrcDev(290400), DstDev(718d00), NSID(1)
Port Speed : SrcDev(32G), DstDev(32G)
Active Flow : 1
Timebase : 6Hrs
Protocol : NVMe

```

```
Port: (141)
```

```

Monitor Time : | Wed Nov 25 03:05:49 MST 2020 |
Elapsed Time : | 30s |

```

```
I/O Latency |
```

```

=====
Metric	Max	Avg			
	Current / All	Current / All			
=====					
RD Completion Time	958u/ 958u	34u/ 34u			
=====					
WR Completion Time	716u/ 716u	47u/ 47u			
=====					
RD 1st Response Time	898u/ 898u	19u/ 19u			
=====					
WR 1st Response Time	134u/ 134u	18u/ 18u			
=====					
RD Pending IOs	44 / 44	1 / 1			
=====					
WR Pending IOs	24 / 24	1 / 1			
=====

```

```
I/O Performance |
```

```

=====
Metric	Avg IO Size	Avg (IOPS)	Avg (BPS)	Total IO	
	Current	Current / All	Current / All	Current / All	
All | Current / All |
=====

```

```

|RD IO | 31.56K| 7.326K/ 7.326K| 225.8M/ 225.8M| 219.8K/
219.8K| 6.615G/ 6.615G|

|WR IO | 32.00K| 7.158K/ 7.158K| 223.6M/ 223.6M| 214.7K/
214.7K| 6.553G/ 6.553G|

```

## Host View

The host view provides an aggregated view of all flows originating from the initiator. The following example from a Gen 7 platform shows the output of the aggregated flow from the initiator.

```
switch:admin> flow --show sys_flow_monitor -srcdev 71af40
```

```

Name : sys_flow_monitor
Definition : SrcDev(71af40)
Port Speed : SrcDev(32G)
Target Count : 1
Active Flow : 1
Protocol : SCSI
Timebase: 5Min

```

```

Monitor Time : | Mon Aug 16 04:34:46 MDT 2021 |
RD Elapsed Time : | 4m:54s |
WR Elapsed Time : | 4m:54s |

```

### FC Statistics |

```

| Frame Count | Frames Per Second | Byte Count | Throughput (BPS) |Frame
Size(Bytes) |
| Tx / Rx / Total | Tx / Rx / Total | Tx / Rx / Total | Tx / Rx / Total | Tx /
Rx
1.49M/ 10.33M/ 11.82M	146.83k/ 1.01M/ 1.16M	2.37G/ 2.80G/ 5.18G	239.71M/282.84M/522.56M	1712/ 292

```

### I/O Latency |

```

| Metric | Max | Avg |
| Current / All | Current / All |

|RD Completion Time | 90.10m/ 111.4m| 1.483m/ 1.510m|

|WR Completion Time | 130.2m/ 144.3m| 36.19m/ 36.33m|

|RD 1st Response Time | 89.44m/ 110.9m| 1.142m/ 1.168m|

|WR 1st Response Time | 87.37m/ 91.88m| 363u/ 357u|

```

|                      |          |             |         |            |           |            |            |            |         |
|----------------------|----------|-------------|---------|------------|-----------|------------|------------|------------|---------|
| -----                |          |             |         |            |           |            |            |            |         |
| I/O Performance      |          |             |         |            |           |            |            |            |         |
| -----                |          |             |         |            |           |            |            |            |         |
| Metric               |          | Avg IO Size |         | Avg(IOPS)  |           | Avg(BPS)   |            | Total IO   |         |
| Total Bytes          |          |             |         |            |           |            |            |            |         |
|                      |          | Current     |         | Current    | /         | All        |            | Current    | /       |
| All                  |          | Current     | /       | All        |           |            |            |            |         |
| -----                |          |             |         |            |           |            |            |            |         |
| RD IO                |          | 256.0K      |         | 923        | /         | 923        |            | 230.9M/    | 230.7M  |
| 4.108M               | 66.51G/  | 1003G       |         |            |           |            |            |            | 272.4K/ |
| -----                |          |             |         |            |           |            |            |            |         |
| WR IO                |          | 256.0K      |         | 923        | /         | 923        |            | 230.9M/    | 230.7M  |
| 4.108M               | 66.51G/  | 1003G       |         |            |           |            |            |            | 272.4K/ |
| -----                |          |             |         |            |           |            |            |            |         |
| -----                |          |             |         |            |           |            |            |            |         |
| I/O Violation Count  |          |             |         |            |           |            |            |            |         |
| -----                |          |             |         |            |           |            |            |            |         |
| Cmd                  |          | Timebase    |         | Total      |           | Violations |            | Violations |         |
| Violations           |          |             |         |            |           |            |            |            |         |
|                      |          |             |         | Violations |           | ECT        |            | FRT        |         |
| IOs                  |          |             |         |            |           |            |            | Pending    |         |
| -----                |          |             |         |            |           |            |            |            |         |
| RD                   |          | Current     |         | 104.3K     | 55.61K /  | 2.000m     | 48.75K /   | 1.750m     | /       |
|                      |          |             |         |            |           |            |            |            |         |
|                      |          | ALL         |         | 1.634M     | 868.5K /  | 2.000m     | 765.9K /   | 1.750m     | /       |
| -----                |          |             |         |            |           |            |            |            |         |
| WR                   |          | Current     |         | 8.946K     | 7.250K /  | 2.000m     | 1.696K /   | 1.750m     | /       |
|                      |          |             |         |            |           |            |            |            |         |
|                      |          | ALL         |         | 202.8K     | 178.4K /  | 2.000m     | 24.39K /   | 1.750m     | /       |
| -----                |          |             |         |            |           |            |            |            |         |
| I/O Exceptions Count |          |             |         |            |           |            |            |            |         |
| -----                |          |             |         |            |           |            |            |            |         |
| RD                   |          |             |         | SCSI       |           |            | NVME       |            |         |
|                      |          |             |         |            |           |            |            |            |         |
|                      | Timebase | Total       |         |            |           |            |            |            |         |
| -----                |          |             |         |            |           |            |            |            |         |
| WR                   |          | Exceptions  | TIMEOUT | BUSY       | RESERVE   | CHECK      | NAME SPACE | RESERVE    |         |
| OTHERS               |          |             |         |            |           |            |            |            |         |
|                      |          |             |         | CONFLICT   | CONDITION | NOT READY  | CONFLICT   |            |         |
| -----                |          |             |         |            |           |            |            |            |         |
| RD                   | Current  |             |         |            |           |            |            |            |         |
|                      |          |             |         |            |           |            |            |            |         |
|                      | All      |             |         |            |           |            |            |            |         |
| -----                |          |             |         |            |           |            |            |            |         |
| WR                   | Current  |             |         |            |           |            |            |            |         |
|                      |          |             |         |            |           |            |            |            |         |

```
| | All | | | | | | | | | |
|
```

The following example from a Gen 6 platform shows the output of the aggregated flow from the initiator.

```
switch:admin> flow --show sys_flow_monitor -srcdev 290f00
```

```
=====
Name : sys_flow_monitor
Definition : SrcDev(290f00)
Port Speed : SrcDev(32G)
Active Flow : 1
Timebase : 5Min
Protocol : SCSI
=====
```

```
Monitor Time : | Mon Aug 16 04:35:45 MDT 2021 |
RD Elapsed Time : | 60s |
WR Elapsed Time : | 60s |
```

```

FC Statistics |
```

```
=====
| Rx Frames Count | Rx Frames per Second | Rx Bytes Count | Rx Throughput (BPS) | Avg Rx Frm Sz (Bytes) |
=====
| 9.93M | 983.80k | 2.55G | 259.52M | 276 |
=====
```

```

I/O Latency |
```

```
=====
Metric	Max	Avg			
Current / All	Current / All				
=====					
RD Completion Time	93.35m/	145.4m	124u/	137u	
=====					
WR Completion Time	93.37m/	145.3m	375u/	415u	
=====					
RD 1st Response Time	93.34m/	145.4m	13.41s/	5.658s	
=====					
WR 1st Response Time	93.15m/	145.1m	13.41s/	5.690s	
=====
```

```

I/O Performance |
```

```
=====
Metric	Avg IO Size	Avg (IOPS)	Avg (BPS)	Total IO	
Total Bytes		Current / All	Current / All	Current / All	
All	Current / All				
=====					
RD IO	51.18M	24.34K/ 24.44K	82.85G/ 1.103T	1.475M/	
8.541G| 72.01T/ 376.6P|
=====
```

|        |         |        |         |        |         |        |         |
|--------|---------|--------|---------|--------|---------|--------|---------|
| WR IO  |         | 51.18M | 24.34K/ | 24.44K | 82.85G/ | 1.103T | 1.475M/ |
| 8.541G | 72.01T/ | 376.6P |         |        |         |        |         |

## Target View

The target view provides an aggregated view of all flows originating from the target.

### NOTE

If both the initiator and the target are in the same ASIC in a switch, Flow Monitor counts it as an initiator flow. It is not added to the target aggregation count.

The following example shows the output of the aggregated flow from the target.

```
switch:admin> flow --show sys_flow_monitor -dstdev 718640
```

```
=====
Name : sys_flow_monitor
Definition : DstDev(718640)
Port Speed : DstDev(16G)
Host Count : 1
LUN Count : 11
Active Flow : 1
Protocol : SCSI
Timebase: 5Min
```

```

Monitor Time : | Mon Aug 16 04:36:48 MDT 2021 |
RD Elapsed Time : | 2m:1s |
WR Elapsed Time : | 2m:1s |
```

### FC Statistics |

| Frame Count             | Frames Per Second      | Byte Count              | Throughput (BPS)       | Frame Size (Bytes) |
|-------------------------|------------------------|-------------------------|------------------------|--------------------|
| Tx / Rx / Total         | Tx / Rx / Total        | Tx / Rx / Total         | Tx / Rx / Total        | Tx / Rx / Total    |
| 379.24k/474.06k/853.30k | 37.54k/ 46.93k/ 84.48k | 392.05M/392.06M/784.11M | 38.81M/ 38.81M/ 77.63M | 1084/ 868          |

### I/O Latency |

| Metric               | Max           | Avg           |
|----------------------|---------------|---------------|
|                      | Current / All | Current / All |
| RD Completion Time   | 88u/ 9.114m   | 33u/ 34u      |
| WR Completion Time   | 104u/ 8.388s  | 42u/ 43u      |
| RD 1st Response Time | 71u/ 9.065m   | 18u/ 19u      |
| WR 1st Response Time | 66u/ 9.031m   | 15u/ 15u      |

```

|RD Pending IOs | 7 / 9 | 3 / 1 |

|WR Pending IOs | 7 / 9 | 3 / 1 |

=====

```

```

I/O Performance

```

| Metric      |           | Avg IO Size | Avg (IOPS) |        | Avg (BPS) |        | Total IO |     |
|-------------|-----------|-------------|------------|--------|-----------|--------|----------|-----|
| Total Bytes |           |             |            |        |           |        |          |     |
|             |           | Current     | Current    | All    | Current   | All    | Current  | All |
| All         | Current / | All         |            |        |           |        |          |     |
| RD IO       |           | 4.000K      | 9.181K/    | 9.434K | 35.86M/   | 36.85M | 1.116M/  |     |
| 43.14M      | 4.259G/   | 164.5G      |            |        |           |        |          |     |
| WR IO       |           | 4.000K      | 9.181K/    | 9.434K | 35.86M/   | 36.85M | 1.116M/  |     |
| 43.14M      | 4.259G/   | 164.5G      |            |        |           |        |          |     |

```

I/O Violation Count

```

| Cmd        | Timebase | Total      | Violations  |             | Violations |  |         |
|------------|----------|------------|-------------|-------------|------------|--|---------|
| Violations |          | Violations | ECT         |             | FRT        |  | Pending |
| IOs        |          |            |             |             |            |  |         |
| RD         | Current  |            | / 2.000m    | / 1.750m    | /          |  |         |
|            |          |            |             |             |            |  |         |
|            | ALL      | 10         | 5 / 2.000m  | 5 / 1.750m  | /          |  |         |
|            |          |            |             |             |            |  |         |
| WR         | Current  |            | / 2.000m    | / 1.750m    | /          |  |         |
|            |          |            |             |             |            |  |         |
|            | ALL      | 23         | 12 / 2.000m | 11 / 1.750m | /          |  |         |
|            |          |            |             |             |            |  |         |

```

I/O Exceptions Count

```

| RD     |            |         |      | SCSI     |           |            | NVME     |  |  |
|--------|------------|---------|------|----------|-----------|------------|----------|--|--|
|        | Timebase   | Total   |      |          |           |            |          |  |  |
| WR     | Exceptions | TIMEOUT | BUSY | RESERVE  | CHECK     | NAME SPACE | RESERVE  |  |  |
| OTHERS |            |         |      | CONFLICT | CONDITION | NOT READY  | CONFLICT |  |  |





```

| | | RD| | | Total Count | Violations | Violations | Violations |
Violations | Total Count |
| Port | PID | WR| Device |Proto| Violations | ECT | FRT | Pending IOs |
Exceptions | Errors |
| Index| | Type | | 5min / All | 5min / All | 5min / All | 5min / All |
5min / All | 5min / All |

```

```

| | 59| 023b00| RD| | I | SCSI| | /1.529K| | / 720 | | / 809 | | / |
/ | | / |

```

## Top Flows

To display the top 10 flows or a configurable number of flows of type Exchange Completion Time (ECT), First Response Time (FRT), Bytes Per Second (BPS), I/O Operations Per Second (IOPS), and Pending I/O, use the `flow --show sys_flow_monitor` command. The ECT Current Max metric is used for top ranking. The following example shows the top 10 number of flows of BPS type.

```
switch:admin> flow --show sys_flow_monitor -top bps
```

```
=====
Name : sys_flow_monitor

```

```

VM Mode : Yes
Monitored port(s): (*)
Host Port Count : 6
Target Port Count: 2
SCSI : 48
NVMe : 2
No of IT : 50

```

```
TOP 10 ECT |
```

```

| | | RD| | | | | | Max Time | Max Time | Avg IOs/sec | Avg
Bytes/sec| Max Count | Total Count | Total Count |
| SID | DID | WR| Device | Flow | Port |Proto| ECT | FRT | IOPS |
BPS | Pending IOs | Exceptions | Violations |
| | | Type | Hierarchy | Index| | 5min / All | 5min / All | 5min / All |
5min / All | 5min / All | 5min / All | 5min / All |

| 7a1b00| 6d1300| WR| T | IT->L| 19| SCSI| 139u/36.15m| 125u/36.14m| 6.833K/6.865K|
26.69M/26.81M| 2 / 2 | / | / |
| 015200| 6d1300| WR| T | IT->L| 19| SCSI| 158u/36.10m| 146u/36.09m| 6.824K/6.866K|
26.65M/26.82M| 2 / 2 | / | / |
| 015200| 6d1300| RD| T | IT->L| 19| SCSI| 159u/ 189u| 157u/ 188u| 6.823K/6.861K|
26.65M/26.80M| 2 / 2 | / | / |
| 7a1b00| 6d1300| RD| T | IT->L| 19| SCSI| 123u/ 190u| 122u/ 189u| 6.815K/6.865K|
26.62M/26.81M| 2 / 2 | / | / |

```

```

| 6d0200| 7a2200| WR| I | IT->L| 2| SCSI| 175u/36.15m| 162u/36.13m| 6.623K/6.647K|
25.87M/25.96M| 2 / 2 | / | / |
| 6d0200| 7a2200| RD| I | IT->L| 2| SCSI| 143u/ 180u| 142u/ 179u| 6.620K/6.649K|
25.86M/25.97M| 2 / 2 | / | / |
| 6d0100| d7ce00| RD| I | IT->L| 1| SCSI| 123u/36.12m| 122u/36.12m| 6.399K/6.424K|
24.99M/25.09M| 2 / 2 | / | / |
| 6d0100| d7ce00| WR| I | IT->L| 1| SCSI| 150u/ 191u| 136u/ 178u| 6.397K/6.426K|
24.98M/25.10M| 2 / 2 | / | / |
=====

```

The following example shows the top 10 number of First Response Time (FRT) flows.

```
switch:admin> flow --show sys_flow_monitor -top frt
```

```
=====
Name : sys_flow_monitor

```

```

VM Mode : Yes
Monitored port(s): (*)
Host Port Count : 6
Target Port Count: 2
SCSI : 48
NVMe : 2
No of IT : 50

```

```
TOP 10 ECT |
```

```

| | RD| | | | | | | | | |
Bytes/sec| Max Count | Total Count | Total Count | | | | | | |
| SID | DID | WR| Device | Flow | Port |Proto| ECT | FRT | IOPS |
BPS | Pending IOs | Exceptions | Violations |
| | | | Type | Hierarchy | Index| | 5min / All | 5min / All | 5min / All |
5min / All | 5min / All | 5min / All | 5min / All |

| d71a00| 6d1201| RD| T | IT->N| 18| NVMe| 12.21m/12.21m| 12.21m/12.21m| 8.462K/8.491K|
4.131M/4.145M| 5 / 5 | / | / |
| d71a00| 6d1200| RD| T | IT->N| 18| NVMe| 12.04m/12.04m| 12.04m/12.04m| 8.472K/8.490K|
4.136M/4.145M| 6 / 6 | / | / |
| d71a00| 6d1201| WR| T | IT->N| 18| NVMe| 11.70m/11.70m| 11.69m/11.69m| 8.466K/8.490K|
4.133M/4.145M| 5 / 5 | / | / |
| d71a00| 6d1200| WR| T | IT->N| 18| NVMe| 11.49m/11.49m| 11.48m/11.48m| 8.456K/8.491K|
4.128M/4.145M| 6 / 6 | / | / |
| 6d2600| 830700| RD| I | IT->L| 38| SCSI| 991u/ 991u| 990u/ 990u| 2.947K/2.999K|
11.51M/11.71M| 2 / 2 | / | / |
| 6d0100| 830600| RD| I | IT->L| 1| SCSI| 990u/37.01m| 989u/37.01m| 1.963K/1.955K|
7.670M/7.638M| 2 / 2 | / | / |
| 6d0300| 830400| RD| I | IT->L| 3| SCSI| 990u/ 991u| 988u/ 989u| 5.723K/5.753K|
22.35M/22.47M| 2 / 2 | / | / |
| 6d2100| 830600| RD| I | IT->L| 33| SCSI| 988u/19.47m| 987u/19.47m| 1.970K/2.001K|
7.696M/7.819M| 2 / 2 | / | / |

=====

```

The following example shows the top 10 number of IOPS flows using the flow `--show sys_flow_monitor -top iops` command.

```
switch:admin> flow --show sys_flow_monitor -top iops
```

```
=====
```

```
Name : sys_flow_monitor
```

```

```

```
VM Mode : Yes
```

```
Monitored port(s): (*)
```

```
Host Port Count : 6
```

```
Target Port Count: 2
```

```
SCSI : 48
```

```
NVMe : 2
```

```
No of IT : 50
```

```

```

```
TOP 10 ECT |
```

```

```

| Bytes/sec          | Max Count    | Total Count | Total Count | Max Time      | Max Time      | Avg I/Os/sec  | Avg        |
|--------------------|--------------|-------------|-------------|---------------|---------------|---------------|------------|
| SID                | DID          | WR          | Device      | Flow          | Port          | Proto         | ECT        |
| BPS                | Pending I/Os | Exceptions  | Violations  | Type          | Hierarchy     | Index         | IOPS       |
| 5min / All         | 5min / All   | 5min / All  | 5min / All  | 5min / All    | 5min / All    | 5min / All    | 5min / All |
| d71a00  6d1201  RD | T            | IT->N       | 18  NVMe    | 12.21m/12.21m | 12.21m/12.21m | 8.463K/8.486K |            |
| 4.131M/4.143M      | 6 / 6        | /           | /           |               |               |               |            |
| d71a00  6d1200  WR | T            | IT->N       | 18  NVMe    | 11.49m/11.49m | 11.48m/11.48m | 8.461K/8.487K |            |
| 4.130M/4.143M      | 6 / 6        | /           | /           |               |               |               |            |
| d71a00  6d1200  RD | T            | IT->N       | 18  NVMe    | 12.04m/12.04m | 12.04m/12.04m | 8.460K/8.484K |            |
| 4.130M/4.142M      | 6 / 6        | /           | /           |               |               |               |            |
| d71a00  6d1201  WR | T            | IT->N       | 18  NVMe    | 11.70m/11.70m | 11.69m/11.69m | 8.457K/8.485K |            |
| 4.128M/4.142M      | 5 / 5        | /           | /           |               |               |               |            |
| 7a1b00  6d1300  WR | T            | IT->L       | 19  SCSI    | 161u/36.15m   | 146u/36.14m   | 6.837K/6.863K |            |
| 26.70M/26.80M      | 2 / 2        | /           | /           |               |               |               |            |
| 015200  6d1300  WR | T            | IT->L       | 19  SCSI    | 158u/36.10m   | 146u/36.09m   | 6.832K/6.864K |            |
| 26.68M/26.81M      | 2 / 2        | /           | /           |               |               |               |            |
| 015200  6d1300  RD | T            | IT->L       | 19  SCSI    | 163u/ 189u    | 162u/ 188u    | 6.829K/6.858K |            |
| 26.67M/26.79M      | 2 / 2        | /           | /           |               |               |               |            |
| 7a1b00  6d1300  RD | T            | IT->L       | 19  SCSI    | 158u/ 190u    | 156u/ 189u    | 6.821K/6.862K |            |
| 26.64M/26.80M      | 2 / 2        | /           | /           |               |               |               |            |

```

```

```
=====
```

The following example shows the top 10 number of IOPS flows using the flow `--show sys_flow_monitor -top pending_io` command.

```
switch:admin> flow --show sys_flow_monitor -top pending_io
```

```
=====
```

```
Name : sys_flow_monitor
```

```

```

```
VM Mode : No
```

```
Monitored port(s): (*)
```

```
Host Port Count : 1
```

```
Target Port Count: 0
```

```

SCSI : 1
NVMe : 0
No of IT : 1

TOP 10 ECT
		RD								
Bytes/sec	Max Count	Total Count	Total Count						
SID	DID	WR	Device	Flow	Port	Proto				
BPS	Pending IOs	Exceptions	Violations						
		Type	Hierarchy	Index					
5min / All	5min / All	5min / All	5min / All						

632100	718640	RD	I	IT->L	33	SCSI	77u/7.524m	66u/7.495m	4.814K/6.256K
18.80M/24.43M	5 / 5	/	/						
632100	718640	WR	I	IT->L	33	SCSI	92u/8.388s	66u/7.559m	4.814K/6.256K
18.80M/24.43M	5 / 5	/	/						

=====

```

You can use the following set of commands to view the number of top flows (n) that are defined manually:

- `flow --show sys_flow_monitor -top ect -n 25`
- `flow --show sys_flow_monitor -top frt -n 5`
- `flow --show sys_flow_monitor -top bps -n 20`
- `flow --show sys_flow_monitor -top iops -n 15`
- `flow --show sys_flow_monitor -top pending_io -n 1`

### Bottom Flows

You can also use the `flow --show sys_flow_monitor` command to display the bottom 10 flows or a configurable number of flows of type Exchange Completion Time (ECT), First Response Time (FRT), Bytes Per Second (BPS), I/O Operations Per Second (IOPS), and Pending I/O.

The ECT Current Max metric is used for bottom ranking.

For this set of examples, the bottom 10 flows (the default) are displayed.

- `flow --show sys_flow_monitor -top ect`
- `flow --show sys_flow_monitor -top frt`
- `flow --show sys_flow_monitor -top bps`
- `flow --show sys_flow_monitor -top iops`
- `flow --show sys_flow_monitor -top pending_io`

The `flow --control sys_flow_monitor -n <n>` command helps to set `n`, which is the number of entries to be displayed in the brief output of `sys_flow_monitor`.

For this set of examples, the number of bottom flows (n) is set manually.

- `flow --show sys_flow_monitor -top ect -n 25`
- `flow --show sys_flow_monitor -top frt -n 5`
- `flow --show sys_flow_monitor -top bps -n 20`
- `flow --show sys_flow_monitor -top iops -n 15`
- `flow --show sys_flow_monitor -top pending_io -n 1`

The following example shows the bottom 10 Exchange Completion Time (ECT) flows.

```
switch:admin> flow --show sys_flow_monitor -bottom ect
```

```
=====
```

```
Name : sys_flow_monitor
```

```

```

```
VM Mode : No
```

```
Monitored port(s): (*)
```

```
Host Port Count : 1
```

```
Target Port Count: 0
```

```
SCSI : 1
```

```
NVMe : 0
```

```
No of IT : 1
```

```

```

```
TOP 10 ECT |
```

```

```

|            |             |             | RD          |            |      |       |            | Max Time   | Max Time   | Avg IOs/sec | Avg |
|------------|-------------|-------------|-------------|------------|------|-------|------------|------------|------------|-------------|-----|
| Bytes/sec  | Max Count   | Total Count | Total Count |            |      |       |            |            |            |             |     |
| SID        | DID         | WR          | Device      | Flow       | Port | Proto | ECT        | FRT        | IOPS       |             |     |
| BPS        | Pending IOs | Exceptions  | Violations  |            |      |       |            |            |            |             |     |
|            |             | Type        | Hierarchy   | Index      |      |       | 5min / All | 5min / All | 5min / All |             |     |
| 5min / All | 5min / All  | 5min / All  | 5min / All  | 5min / All |      |       |            |            |            |             |     |

|               |        |    |   |       |    |      |             |            |               |  |  |
|---------------|--------|----|---|-------|----|------|-------------|------------|---------------|--|--|
| 632100        | 718640 | RD | I | IT->L | 33 | SCSI | 85u/7.524m  | 71u/7.495m | 4.975K/6.143K |  |  |
| 19.43M/23.99M | 5 / 5  | /  | / |       |    |      |             |            |               |  |  |
| 632100        | 718640 | WR | I | IT->L | 33 | SCSI | 183u/8.388s | 74u/7.559m | 4.975K/6.143K |  |  |
| 19.43M/23.99M | 5 / 5  | /  | / |       |    |      |             |            |               |  |  |

## Switch-Level Brief Display

In this example of I/O metrics, where no specified source or destination is learned, brief statistics are displayed.

```
switch:admin> flow --show sys_flow_monitor -srcdev "*" -dstdev ""
```

```
=====
```

```
Name : sys_flow_monitor
```

```
Definition : SrcDev(*)DstDev(*)
```

```

```

```
Timebase : 5 Minutes
```

```

```

```
| I/O Brief Metrics |
```

```

```

|            |             |             | RD          |            |      |       |            | Max Time   | Max Time   | Avg IOs/sec | Avg |
|------------|-------------|-------------|-------------|------------|------|-------|------------|------------|------------|-------------|-----|
| Bytes/sec  | Max Count   | Total Count | Total Count |            |      |       |            |            |            |             |     |
| SID        | DID         | WR          | Device      | Flow       | Port | Proto | ECT        | FRT        | IOPS       |             |     |
| BPS        | Pending IOs | Exceptions  | Violations  |            |      |       |            |            |            |             |     |
|            |             | Type        | Hierarchy   | Index      |      |       | 5min / All | 5min / All | 5min / All |             |     |
| 5min / All | 5min / All  | 5min / All  | 5min / All  | 5min / All |      |       |            |            |            |             |     |

|               |        |    |   |    |     |      |             |             |               |  |  |
|---------------|--------|----|---|----|-----|------|-------------|-------------|---------------|--|--|
| 290600        | 718d00 | RD | T | IT | 141 | NVMe | 224u/1.927m | 128u/1.927m | 7.558K/7.521K |  |  |
| 232.3M/231.1M | 7 / 7  | /  | / |    |     | 990  |             |             |               |  |  |
| 290600        | 718d00 | WR | T | IT | 141 | NVMe | 185u/36.17m | 127u/36.11m | 7.355K/7.322K |  |  |
| 229.8M/228.8M | 3 / 3  | /  | / |    |     | 89   |             |             |               |  |  |
| 290700        | 718d00 | RD | T | IT | 141 | NVMe | 260u/ 672u  | 129u/ 672u  | 7.864K/7.841K |  |  |
| 237.8M/237.1M | 7 / 7  | /  | / |    |     | 69   |             |             |               |  |  |

```

| 290700| 718d00| WR| T | IT| 141| NVMe| 181u/36.18m| 122u/36.14m| 7.456K/7.433K|
233.0M/232.3M| 3 / 3 | / | / 4 |
| 632100| 718640| RD| T | IT->L| 390| SCSI| 710u/9.114m| 668u/9.065m| 10.80K/9.650K|
42.21M/37.69M| 7 / 9 | / | / 10 |
| 632100| 718640| WR| T | IT->L| 390| SCSI| 722u/8.388s| 668u/9.031m| 10.80K/9.650K|
42.21M/37.69M| 7 / 9 | / | / 23 |
| 719b40| 9f0200| RD| I | IT->L| 411| SCSI| 584u/10.63m| 168u/10.61m| 12.11K/12.09K|
360.7M/360.1M| 9 / 10 | / | /1.102K|
| 719b40| 9f0200| WR| I | IT->L| 411| SCSI| 209u/10.69m| 159u/10.62m| 11.37K/11.35K|
355.4M/354.8M| 5 / 6 | / | / 733 |
| 9f0a00| 718740| RD| T | IT->N| 391| NVMe| 123u/ 197u| 122u/ 196u| 21.54K/21.29K|
42.07M/41.59M| 5 / 7 | / | / |
| 9f0a00| 718740| WR| T | IT->N| 391| NVMe| 182u/ 226u| 120u/ 205u| 21.54K/19.48K|
42.07M/38.04M| 3 / 5 | / | / |
| 9f0b00| 718740| RD| T | IT->N| 391| NVMe| 124u/ 204u| 123u/ 204u| 21.28K/21.44K|
41.56M/41.89M| 4 / 7 | / | / |
| 9f0b00| 718740| WR| T | IT->N| 391| NVMe| 170u/ 223u| 121u/ 207u| 21.28K/19.61K|
41.56M/38.31M| 4 / 7 | / | / 7 |
| 71af40| 718d40| RD| VM-I | IT->L->VM| 431| SCSI| 80.10m/111.4m| 79.81m/110.9m| 926 / 923 |
231.5M/230.8M| 48 / 48 | / | 33.43K/1.344M|
| 71af40| 718d40| WR| VM-I | IT->L->VM| 431| SCSI| 108.8m/144.3m| 78.42m/91.88m| 926 / 923 |
231.5M/230.8M| 64 / 64 | / | 5.962K/173.5K|
| 290f00| 718d40| RD| T | IT->L->VM| 397| SCSI| 91.32m/93.63m| 91.32m/93.62m| 24.45K/24.33K|
210.1M/208.8M| 5 / 22 | / | 174 /8.551K|
| 290f00| 718d40| WR| T | IT->L->VM| 397| SCSI| 91.39m/93.65m| 91.23m/93.40m| 24.45K/24.33K|
210.1M/208.8M| 9 / 48 | / | 306 /14.21K|

```

## Detailed Show - IT View

In this example of a detailed IT view, both source and destination are specified for IT flows.

```
flow --show sys_flow_monitor -srcdev 632100 -dstdev 718640
```

```

Name : sys_flow_monitor
Definition : SrcDev(632100)DstDev(718640)
Port Speed : SrcDev(32G), DstDev(16G)
Active Flow : 1
Timebase : 5Min
Protocol : SCSI
Port : 390

```

```

Monitor Time : | Mon Aug 16 04:23:26 MDT 2021 |
RD Elapsed Time : | 3m:33s |
WR Elapsed Time : | 3m:33s |

```

```
FC Statistics |
```

```

| Frame Count | Frames Per Second | Byte Count | Throughput (BPS) | Frame
Size (Bytes) |

```

|                                                                                                                |
|----------------------------------------------------------------------------------------------------------------|
| Tx / Rx / Total   Tx / Rx / Total   Tx / Rx / Total   Tx / Rx / Total   Tx / Rx /                              |
| Rx                                                                                                             |
| -----                                                                                                          |
| 482.14k/602.70k/ 1.08M   47.73k/ 59.67k/107.40k   498.42M/498.44M/996.87M   49.34M/ 49.35M/ 98.69M   1084/ 868 |
|                                                                                                                |
| -----                                                                                                          |

-----  
I/O Latency |

|                                                  |
|--------------------------------------------------|
| Metric   Max   Avg                               |
| Current / All   Current / All                    |
| -----                                            |
| RD Completion Time   1.677m/ 9.114m   35u/ 34u   |
| -----                                            |
| WR Completion Time   1.700m/ 8.388s   43u/ 43u   |
| -----                                            |
| RD 1st Response Time   1.631m/ 9.065m   19u/ 19u |
| -----                                            |
| WR 1st Response Time   1.642m/ 9.031m   16u/ 16u |
| -----                                            |
| RD Pending IOs   7 / 9   1 / 1                   |
| -----                                            |
| WR Pending IOs   7 / 9   1 / 1                   |
| -----                                            |
| =====                                            |

-----  
I/O Performance |

|                                                                         |
|-------------------------------------------------------------------------|
| Metric   Avg IO Size   Avg (IOPS)   Avg (BPS)   Total IO                |
| Total Bytes   Current / All   Current / All   Current / All   Current / |
| All   Current / All                                                     |
| -----                                                                   |
| RD IO   4.000K   11.83K/ 9.745K   46.21M/ 38.06M   2.519M/              |
| 36.74M   9.612G/ 140.1G                                                 |
| -----                                                                   |
| WR IO   4.000K   11.83K/ 9.745K   46.21M/ 38.06M   2.519M/              |
| 36.74M   9.612G/ 140.1G                                                 |
| -----                                                                   |

-----  
I/O Violation Count |

|                                                  |
|--------------------------------------------------|
| Cmd   Timebase   Total   Violations   Violations |
| Violations                                       |
| Violations   ECT   FRT   Pending                 |
| IOs                                              |
| -----                                            |
| RD   Current   / 2.000m   / 1.750m   /           |
|                                                  |



|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|----------------------|-------|--|------------|--|---------|--|------|--------|----------|------|-----------|--|----------------------|
|                      |       |  | ALL        |  | 10      |  | 5 /  | 2.000m |          | 5 /  | 1.750m    |  | /                    |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      | WR    |  | Current    |  |         |  | /    | 2.000m |          | /    | 1.750m    |  | /                    |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  | ALL        |  | 23      |  | 12 / | 2.000m |          | 11 / | 1.750m    |  | /                    |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
| -----                |       |  |            |  |         |  |      |        |          |      |           |  |                      |
| I/O Exceptions Count |       |  |            |  |         |  |      |        |          |      |           |  |                      |
| -----                |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      | RD    |  |            |  |         |  |      | SCSI   |          |      | NVME      |  |                      |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  | Timebase   |  | Total   |  |      |        |          |      |           |  |                      |
|                      | ----- |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      | WR    |  | Exceptions |  | TIMEOUT |  | BUSY |        | RESERVE  |      | CHECK     |  | NAME SPACE   RESERVE |
|                      |       |  | OTHERS     |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  |            |  |         |  |      |        | CONFLICT |      | CONDITION |  | NOT READY   CONFLICT |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
| -----                |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      | RD    |  | Current    |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  | All        |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
| -----                |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      | WR    |  | Current    |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  | All        |  |         |  |      |        |          |      |           |  |                      |
|                      |       |  |            |  |         |  |      |        |          |      |           |  |                      |
| -----                |       |  |            |  |         |  |      |        |          |      |           |  |                      |

## Detailed Show - Initiator Flow

In this example of a detailed view, only the initiator is specified.

```
switch:admin> flow --show sys_flow_monitor -srcdev 71af40
```

```
=====
Name : sys_flow_monitor
Definition : SrcDev(71af40)
Port Speed : SrcDev(32G)
Target Count : 1
Active Flow : 1
Protocol : SCSI
Timebase: 5Min
```

```

Monitor Time : | Mon Aug 16 04:34:46 MDT 2021 |
RD Elapsed Time : | 4m:54s |
WR Elapsed Time : | 4m:54s |
```

```

FC Statistics |
```

```

| Frame Count | Frames Per Second | Byte Count | Throughput (BPS) | Frame
| Size(Bytes) |
```

| Tx / Rx / Total       | Tx / Rx / Total       | Tx / Rx / Total     | Tx / Rx / Total         | Tx / Rx / Total |
|-----------------------|-----------------------|---------------------|-------------------------|-----------------|
| 1.49M/ 10.33M/ 11.82M | 146.83k/ 1.01M/ 1.16M | 2.37G/ 2.80G/ 5.18G | 239.71M/282.84M/522.56M | 1712/ 292       |

## I/O Latency |

| Metric               | Max            | Avg            |
|----------------------|----------------|----------------|
| Current / All        | Current / All  | Current / All  |
| RD Completion Time   | 90.10m/ 111.4m | 1.483m/ 1.510m |
| WR Completion Time   | 130.2m/ 144.3m | 36.19m/ 36.33m |
| RD 1st Response Time | 89.44m/ 110.9m | 1.142m/ 1.168m |
| WR 1st Response Time | 87.37m/ 91.88m | 363u/ 357u     |

## I/O Performance |

| Metric                | Avg IO Size | Avg (IOPS)    | Avg (BPS)      | Total IO      |
|-----------------------|-------------|---------------|----------------|---------------|
| Total Bytes           | Current     | Current / All | Current / All  | Current / All |
| RD IO                 | 256.0K      | 923 / 923     | 230.9M/ 230.7M | 272.4K/       |
| 4.108M  66.51G/ 1003G |             |               |                |               |
| WR IO                 | 256.0K      | 923 / 923     | 230.9M/ 230.7M | 272.4K/       |
| 4.108M  66.51G/ 1003G |             |               |                |               |

## I/O Violation Count |

| Cmd        | Timebase | Total      | Violations      | Violations      |
|------------|----------|------------|-----------------|-----------------|
| Violations |          | Violations | ECT             | FRT             |
| IOs        |          |            |                 | Pending         |
| RD         | Current  | 104.3K     | 55.61K / 2.000m | 48.75K / 1.750m |
|            | ALL      | 1.634M     | 868.5K / 2.000m | 765.9K / 1.750m |
| WR         | Current  | 8.946K     | 7.250K / 2.000m | 1.696K / 1.750m |

|                      |    |         |          |            |        |         |          |        |  |          |        |            |   |
|----------------------|----|---------|----------|------------|--------|---------|----------|--------|--|----------|--------|------------|---|
|                      |    |         | ALL      |            | 202.8K |         | 178.4K / | 2.000m |  | 24.39K / | 1.750m |            | / |
|                      |    |         |          |            |        |         |          |        |  |          |        |            |   |
| -----                |    |         |          |            |        |         |          |        |  |          |        |            |   |
| I/O Exceptions Count |    |         |          |            |        |         |          |        |  |          |        |            |   |
| -----                |    |         |          |            |        |         |          |        |  |          |        |            |   |
|                      | RD |         |          |            |        |         |          | SCSI   |  |          |        | NVME       |   |
|                      |    |         |          |            |        |         |          |        |  |          |        |            |   |
|                      |    |         | Timebase |            |        |         | Total    |        |  |          |        |            |   |
| -----                |    |         |          |            |        |         |          |        |  |          |        |            |   |
|                      | WR |         |          | Exceptions |        | TIMEOUT |          | BUSY   |  | RESERVE  |        | CHECK      |   |
|                      |    |         |          |            |        |         |          |        |  |          |        | NAME SPACE |   |
|                      |    |         |          | OTHERS     |        |         |          |        |  |          |        | RESERVE    |   |
|                      |    |         |          |            |        |         |          |        |  | CONFLICT |        | CONDITION  |   |
|                      |    |         |          |            |        |         |          |        |  |          |        | NOT READY  |   |
|                      |    |         |          |            |        |         |          |        |  |          |        | CONFLICT   |   |
| -----                |    |         |          |            |        |         |          |        |  |          |        |            |   |
|                      | RD | Current |          |            |        |         |          |        |  |          |        |            |   |
|                      |    |         |          |            |        |         |          |        |  |          |        |            |   |
|                      |    |         | All      |            |        |         |          |        |  |          |        |            |   |
| -----                |    |         |          |            |        |         |          |        |  |          |        |            |   |
|                      | WR | Current |          |            |        |         |          |        |  |          |        |            |   |
|                      |    |         |          |            |        |         |          |        |  |          |        |            |   |
|                      |    |         | All      |            |        |         |          |        |  |          |        |            |   |
| -----                |    |         |          |            |        |         |          |        |  |          |        |            |   |

Unsupported Uses of the flow --show Command

The following examples are applicable for Gen 6 platforms only.

```
switch:admin> flow --show sys_flow_monitor -dstdev 160800
Target port monitoring not supported on this blade.
switch:admin> flow --show sys_flow_monitor -srcdev 160100 -dstdev 160800 -lun ""
LUN or NSID monitoring not supported on this blade
```

Streaming Monitored Flows

Streaming of flow definitions and statistics, device-level statistics, and violation statistics are supported for all monitored flows.

Telemetry streaming is applicable only to SANnav. The following table describes this support in more detail.

Table 124: Telemetry Support

| Streaming Type                                               | Details                                                                                                                                                                                                 |
|--------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Default support                                              | <ul style="list-style-type: none"><li>10-sec flows are guaranteed to be streamed at the next second.</li><li>5-min flows at the next one minute.</li><li>6-hr flows at the next 30 minutes.</li></ul>   |
| Violation statistics streamed only for flows with violations | <ul style="list-style-type: none"><li>Streamed at 5 min and 6 hour flows.</li><li>Streamed at 10 sec for flows that are monitored in real time.</li></ul> Violation metrics are not streamed to SANnav. |

| Streaming Type                                                                               | Details                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Flows that are streamed at system clock time irrespective of the elapsed statistics duration | <ul style="list-style-type: none"> <li>Global timer configured for 10 sec, 5 min, and 6 hr.</li> <li>Data for streaming is a snapshot per flow at the timer expiration.</li> <li>Snapshot data is moved for streaming, and the current stats counters are reset.</li> <li>The CLI uses the current counters.</li> </ul> |
| Newly learned flows at 5 min or 6 hr                                                         | The statistics duration of the first stressed data can always be less than the monitored time.                                                                                                                                                                                                                          |

## Integration with MAPS

Integration with MAPS allows automatically detected flows to be monitored using the predefined rules. Hardware real-time violations are flagged at the I/O (ECT and FRT) and FC level (frame latency).

### NOTE

In Fabric OS 9.1.x, user-defined flow configurations are *not* supported.

Refer to the *Brocade Fabric OS MAPS User Guide, 9.1.x*, for details on MAPS.

## Integration with SANnav

Automated monitoring is supported by SANnav.

Refer to the *Brocade SANnav Management Portal User Guide, 2.2.x*, for more details on SANnav.

## Flow Monitor Limitations

The following limitations and considerations are applicable when you use Flow Monitor:

- Automatic flow monitoring does not support the following features:
  - Create and Delete options
  - E/EX\_Port monitoring
  - Frame type
  - Monitoring of target ports (only supported on Gen 7)
  - Metrics based on the I/O data size
  - User-defined flows
  - Visibility into other commands
  - XISL monitoring
- For Flow Monitor learned flows, if the domain or PID is changed in a remote switch, consider the following:

- The learned flows for the previously defined PID pair remain idle in the flows, they are not actively monitored, and statistics are not generated.
- However, the learned flows for the newly defined PID pair are updated in the flow, they are monitored, and statistics are generated.
- To remove outdated entries, use the `flow --reset` command.
- Flow mirroring is activated automatically. To deactivate it, you can deactivate vTAP with the command `flow --deactivate <flow_name>`.
- In Fabric OS 9.1.x, `sys_flow_monitor` is the *only* monitor that can be activated. It is mutually exclusive with other predefined flows.
- VM Insight does not support SANnav.
- An entity should be registered in the application server to monitor and report statistics. The entity can be registered at multiple N\_Ports within a fabric.
- Flow latency is supported only on F\_Ports. However, it is not supported on F\_Ports on a backbone if the flow is an inter-fabric flow.
- I/O latency statistics are not supported in the Access Gateway platform.

## Flow Generator

Flow Generator is a test traffic generator that allows you to pre-test a SAN infrastructure (including internal connections) for robustness before deploying it.

### NOTE

With Fabric OS 9.1.x, Flow Generator supports Gen 6 and Gen 7.



### CAUTION

Do not use Flow Generator in an active production environment. Traffic created by Flow Generator can saturate the links, which might have a negative impact on any other traffic sharing links.

Flow Generator lets you perform the following:

- Configure a 64G Gen 7 Fibre Channel-capable port as a simulated device that can transmit frames at a full 64G line rate.
- Emulate a 64G SAN without actually having any 64G hosts or targets or SAN-testers.
- Pre-test the entire SAN fabric at the full line rate, including optics and cables on ISLs as well as internal connections within a switch.

Flow Generator uses simulation mode (SIM) ports. SIM ports behave like normal F\_Ports, but are used only for testing. By using SIM ports, Flow Generator traffic is terminated at the destination port and does not leave the fabric.

Flow Generator can generate standard frames or can create custom frames with sizes and patterns that you specify. A sample use case would be to create a traffic flow from a Source ID (SID) to a Destination ID (DID) to validate routing and throughput. [Creating a Flow from a Specific Source ID to a Specific Destination ID](#) provides an example of the command and the results for this use case.

Flow Generator supports predefined flows to generate traffic between all configured SIM ports.

After you activate a Flow Generator flow, the flow stays active until you deactivate the flow. The flow stays active even though the SID and DID SIM-ports are offline. When SID and DID SIM-ports are online, traffic starts.

Flow Generator also supports generating simulated traffic to a remote fabric. To generate traffic, you specify the imported proxy ID or the real PWWN of the device. Flow Generator support is similar to Flow Monitor support, except Flow Generator only supports flows from edge to edge but not from backbone to edge or from edge to backbone.


## Flow Generator Setup

Flow Generator offers several flow control options that you can configure, including the ability to specify both the frame size and the frame payload pattern. Header parameters and other control parameters can also be added as part of the definition. The OXID value for frames cannot be user-specified.

Flow Generator generates and receives traffic only from simulated ingress and egress ports (SIM ports), which emulate device entries in the Name Server database, so that they are treated as real devices and can be used to evaluate various switch and fabric operations such as QoS and Traffic Isolation. For more information on working with SIM ports, see [SIM Port Attributes and Configuration](#).

Flow Generator flows are defined using a combination of the source device (srcdev), destination device (dstdev), ingress port (ingrport), and egress port (egrport) parameters. All of these must be SIM ports. The source device is the origination point of the test traffic. The destination device is the destination of the test traffic; for Flow Generator flows, it may be remote from the switch. The port that transmits the simulation traffic must be a 32G-capable or 64G-capable Fibre Channel port. The port that receives the simulated traffic can be an 8, 16, 32 or 64G-capable Fibre Channel port. The following figure illustrates this concept.

**Figure 48: A Flow Generator flow structure**



## Predefined Flow Generator Flows

Predefined flows are system-defined flows that support the most common testing scenarios. These flows are automatically created and defined based on your network structure as part of the upgrade to Fabric OS 9.1.x or later. These flows use the port ID and device ID modes and other control parameters configured for Flow Generator.

Fabric OS 9.1.x supports one predefined flow named `sys_gen_all_simports`, which generates traffic from each SIM port on a switch to all the other SIM ports configured on that switch. This allows you to easily stress-test the hardware, ASIC path, backplanes, front-end, switch, and fabric operations. The stress-testing operation is handled in iterations. At each iteration, all SIM ports generate traffic to a set of destination ports, determined by their order in the physical ASIC interleave, so that all the ASICs are tested at the beginning of the test. In this flow, every SIM port sends traffic to four destinations using nine Virtual Channels (VCs) for each destination, but does not generate traffic to the originating port. This uses 36 VCs; the remaining three VCs are not used by the flow.

You can configure this flow to run on all the ports on the switch or on a specific slot, port range, or list of ports.

- To run the test on all the ports on the local switch, set all the ports on the switch to SIM ports.
- To run the test on one slot only, set only the ports of the slot to SIM ports.
- To run the test on any other set of ports, set these specific ports to SIM ports.

The frame payload size and pattern control parameters apply to the `sys_gen_all_simports`. The control parameters can be the default, user-defined default, or specific to the `sys_gen_all_simports` flow. See [Customizing Flow Generator Flows](#) for more details.

To start this flow, you must first enable SIM ports (specifying the ports you want used), and then you must activate it. The following example enables SIM ports 10 through 20 and then activates the flow.

```
switch:admin> flow --control -simport 10-20 -enable
switch:admin> flow --activate sys_gen_all_simports -feature generator
```

**NOTE**

Once started, the `sys_gen_all_simports` flow runs until you deactivate it.

To stop the flow, enter the flow `--deactivate sys_gen_all_simports -feature generator` command. To reset the flow metrics, enter the flow `--reset sys_gen_all_simports -feature generator` command. This resets the ingress and egress counts for all the tested SIM ports.

The following restrictions apply specifically to this stress-test flow:

- You must configure the SIM ports you are going to test before activating the flow.
- There is a limit of 256 flows (64 flows at four flows per port) that can be generated per iteration.
- This flow bypasses all zoning, as SIM ports do not need to be zoned for the flows to be generated.
- Activating the predefined flow is mutually exclusive with all other active Flow Generator flows. If there are any static flows or sub-flows activated on the switch, you cannot activate the predefined flow. If the predefined flow is activated, no other Flow Generator flow can be activated.

Issuing the flow `--show sys_gen_all_simports` command displays the following:

- The SID-DID pairs tested
- The number of complete runs where all the SID-DID pairs have been tested
- The completion percentage of the current run
- The total number of frames generated from all the ingress ports

The following example shows the typical results of the flow `--show` command when the `sys_gen_all_simports` flow is active:

```
switch:admin> flow --show
```

| Flow Name            | Feature   | SrcDev | DstDev | IngrPt | EgrPt | BiDir | LUN |
|----------------------|-----------|--------|--------|--------|-------|-------|-----|
| FrameType            | SFID DFID | SEID   | MirPt  |        |       |       |     |
| sys_gen_all_simports | gen       | *      | *      | *      | *     | no    | -   |
| -                    | -         | -      | -      | -      | -     | -     | -   |
| sys_flow_monitor     | mon+      | *      | *      | *      | *     | no    | *   |
| -                    | -         | -      | -      | -      | -     | -     | -   |
| flow_generator1      | gen       | 632c00 | 632800 | 44     | -     | no    | -   |
| -                    | -         | -      | -      | -      | -     | -     | -   |
| flow_generator2      | gen       | 631a00 | 637900 | 26     | -     | no    | -   |
| -                    | -         | -      | -      | -      | -     | -     | -   |
| flow_mirror1         | mir       | *      | *      | 33     | -     | yes   | -   |
| abts                 | -         | -      | -      | -      | -     | -     | -   |
| flow_mirror2         | mir       | *      | *      | 33     | -     | yes   | -   |
| sinq                 | -         | -      | -      | -      | -     | -     | -   |

+ Denotes feature is currently activated for the flow

The flow name with prefix `sys_` denotes predefined flow

The following example shows the results of the flow `--show sys_gen_all_simports` command.

```
switch:admin> flow --show sys_gen_all_simports
```

```
=====
Name : sys_gen_all_simports Features: gen(Activated) noConfig: Off
Definition : IngrPort(*),EgrPort(*),SrcDev(*), DstDev(*)
```

```
Flow Generator (Activated):
```

```
Number of complete runs: 0
```

```

Percent complete of the current run: 67

| SrcDev | DstDev |

| 0x2dec80 | 0x2dee40 |

| 0x2dee40 | 0x2dec80 |

Number of frames generated from IngrPort : 64.57M
=====

```

## Notes on Predefined Flows

The following considerations apply when working with predefined flows:

- If the `sys_gen_all_simports` flow was active before a reboot, it will be reactivated after the reboot.
- If the port, SIM device, slot, or switch goes offline, the subflows on the offline ports will be stopped. Traffic will continue to run on the online ports.
- If all SIM ports in the test go offline, the flow will stay active but traffic will not start, and entering the `flow --show` command returns the message `no sim devices`. In this case, the flow traffic resumes when any SIM ports in the flow come online.
- To prevent dropped frames when offline SIM ports come back on line, those ports included in the test flow are only added back to the test flow during the pause between iterations.
- Activation, deactivation, reset, show, and control operations are supported for predefined flows, but creation and deletion operations are not supported.
- In those cases where the device ID (DID) is equal to the source (SID), that DID will be skipped and the next DID (from the next iteration) is selected to run. This way, every SID will send traffic to four destinations at a time for every iteration.

See [Flow Generator and MAPS](#) for information on how the `sys_gen_all_simports` flow behaves with MAPS, and see [Flow Generator and High Availability](#) for information on how the `sys_gen_all_simports` flow behaves in HA situations.

## Creating Flow Generator Flows

To create a Flow Generator flow, enter the `flow --create <flow_name> -feature generator <parameters>` command.

The following table lists the parameters available to Flow Generator.

**Table 125: Flow Parameter Combinations Supported by Flow Generator**

| Parameter                          | Field Name            | Description                                                                                                                                                                                                                                                                                                              |
|------------------------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port                               | -ingrport<br>-egrport | <ul style="list-style-type: none"> <li>• One field only must be specified.</li> <li>• Values must be explicit.</li> <li>• Must be a SIM port local to the switch for a flow to generate traffic.</li> </ul>                                                                                                              |
| Frame                              | -srcdev<br>-dstdev    | <ul style="list-style-type: none"> <li>• At least one field must be specified.</li> <li>• Values for srcdev and dstdev can be explicit or "*" ("*" indicates learned flows).</li> <li>• Must be a SIM port PID for a flow to generate traffic.</li> <li>• The parameters LUN and frametype are not supported.</li> </ul> |
| <b>Optional Keyword Parameters</b> |                       |                                                                                                                                                                                                                                                                                                                          |
|                                    | -bidir                | Not supported directly. To emulate this function, you must create two flows (one in each direction).                                                                                                                                                                                                                     |



| Parameter | Field Name  | Description                                                                        |
|-----------|-------------|------------------------------------------------------------------------------------|
|           | -noactivate | Adding this keyword creates the flow without activating it.                        |
|           | -noconfig   | Adding this keyword creates the flow without saving the flow to the configuration. |


See [Customizing Flow Generator Flows](#) for information on changing the flow frame size and contents.

## Parameter Usage Exceptions

When you create a flow, it is automatically activated unless you use the `-noactivate` keyword for the `flow --create` command. Refer to [Creating an inactive flow in Flow Generator](#) for an example.

The following illustrated example creates a Flow Generator flow named “Flow11” and generates traffic using the ingress SIM port 1/1 from device 040100 (Domain 4) to device 050200 (Domain 5).

**Figure 49: A Flow Generator Flow Between Two Switches**



```
switch123:admin> flow --create Flow11 -feature generator -ingrport 1/1 -srcdev 040100 -dstdev 050200
```

## Restrictions

The following restrictions apply to the `-feature generator` parameter usage:

- If the `srcdev` and `ingrport` parameters are both defined, they must both be local to the switch and refer to the same source.
- If the `dstdev` and `egrport` parameters are both defined for the flow, they must both be local to the switch and represent the same destination.

## Creating an Inactive Flow in Flow Generator

To create an inactive Flow Generator flow, enter the `flow --create flow_name -feature generator flow_parameters -noactivate` command.

See [Activating Flow Generator Flows](#) for information on activating a Flow Generator flow.

The following example creates an inactive Flow Generator flow named “superflow238” from device 020a00 to device 01c000 ingressing through SIM\_Port 110.

```
switch:admin> flow --create superflow238 -feature generator -ingrport 110 -srcdev 0x020a00 -dstdev 0x01c000 -noactivate
```

## Activating Flow Generator Flows

When a flow is activated, traffic is generated by the ingress port or source device for that flow and any subflows associated with it as soon as all SIM ports and devices defined in the flow are online.

Activating a flow does not automatically clear the flow statistics for that flow; the existing statistic counters resume counting using the resumed flow data. If you are activating a learned flow, the subflows considered for the flow statistics

are refreshed based on the first 39 Source ID-Destination ID SIM device pairs in the zone database that are registered in the Name Server. A flow can be activated that includes SIM ports that are either offline or that have not yet been created. If such a flow is activated, an alert message is displayed, noting that the activated flow is not enforced.

To activate an inactive Flow Generator flow, perform the following steps:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `flow --activate flow_name -feature generator` command.

The following example activates the Flow Generator flow named "Flow1."

```
switch:admin> flow --activate Flow1 -feature generator
```

## Automatic Activation of a Flow Generator Flow

Flow Generator automatically activates a generated flow under the following conditions:

- On flow creation, unless the flow is being created using the `-noactivate` keyword as part of the `flow` command. See [Creating an Inactive Flow in Flow Generator](#) for information on this procedure.
- On slot power-on, if the specified port belongs to the slot being powered on and the flow was active when the slot was powered off.
- On a High Availability (HA) failover, or an HA reboot, if the flow was active when the HA event occurred.

## Learning in Flow Generator Flows

To apply learning to a Flow Generator flow, use an asterisk inside of single quotation marks `'*'` to specify the parameter to be learned.

When Flow Generator activates learned flows, it queries the Name Server database to identify source and destination devices that are zoned together. These pairings are not automatically changed if either member of the pair changes zones. If either member of the pair changes zones, you must deactivate the flow and then reactivate it to use the new zone values. Flow Generator allocates the flows per source ID to zoned destination IDs based on the maximum permitted number of flows. Once this limit has been reached, the rest of the destination IDs are not tested. For learned flows, zone enforcement is applied to both the source and destination SIM-Ports.

In the following example, the only flows that are activated are the ones where the destination devices share a zone with the source device (0x021800) and use ingress port 24. To view the data generated by this flow, see [Viewing the Output of a Learned Flow Generator Flow](#).

```
switch:admin> flow --create fgflow12 -feature generator -ingrport 24 -srcdev 0x021800 -dstdev "*"
```

## Flow Generator Learning Flows on Egress Ports

Creating an active Flow Generator flow on an egress port starts a Flow Generator flow from the srcdev port to the designated egrport. This is the opposite of what happens when creating a Flow Generator flow on an ingress port. If the value for the srcdev is `"*"`, a Flow Generator flow will be initiated to the egress port from every SIM port on the switch that shares a zone with the egress port. This may result in a large number of flows from a large number of ports going to the egress port; the exact number depends on how many SIM ports share the zone with the egress port and the maximum permitted number of flows for the switch.

## Viewing Flow Generator Flows

To display Flow Generator flows, enter the `flow --show flow_name -feature generator` command. For root and static flows, this command shows the Source ID-Destination ID pairs and the cumulative frame count on the ingress or egress port specified in the flow definition.

## Displaying the Status of a Single Flow Generator Flow

The following example displays the status of the Flow Generator flow named "f2."

```
switch:admin> flow --show f2 -feature generator
=====
Name : f2 Features: gen(Active) noConfig: Off
Definition: IngrPort(1/9),SrcDev(0x010900),DstDev(0x01c100)
Flow Generator (Activated):

| SrcDev | DstDev |

| 0x010900 | 0x01c100 |

Number of frames generated from IngrPort : 595.41M
Note: More than 1 flow active on this port.
=====
```

## Viewing the Output of a Learned Flow Generator Flow

To view the output of a learned Flow Generator flow, enter the `flow --show flow_name -feature generator` command.

When you view the output of a learned flow, the Name line displays the flow name and flow features with their respective states. The Definition line displays the port of interest and the device. In the table under Flow Generator, each row is an individual subflow, and the column shows the individual device IDs. The last line displays the number of frames (in units of 1000 (K), 1,000,000 (M), or 1,000,000,000 (G)) that have passed through that port measured from the time the port became active.

The following example shows the output of the Flow Generator flow named "fgflow12":

```
switch:admin>flow --show fgflow12 -feature generator
Name : fgflow12 Features: gen(Active) noConfig: Off
Definition: IngrPort(4),DstDev(*)
Flow Generator (Activated):

| SrcDev | DstDev |

| 0x204000 | 0x040700 |

| 0x204000 | 0x040800 |

| 0x204000 | 0x050900 |

| 0x204000 | 0x051000 |

Number of frames generated from IngrPort : 82.21M
```

## Notes on Displaying the Status of a Flow Generator Flow

- Flow Generator will append a note at the bottom of the output if there is more than one static flow or a learned flow active on the port.

## Repeating Flow Generator Output

You can configure the Flow Generator feature to repeat the flow output. This allows you to view sample frames or output over time to look for differences as the situation evolves. You can specify that a flow output repeats from one to ten times.

To specify the number of times a flow output should be repeated, perform the following steps:

1. Connect to the switch and log in using an account with admin permissions.
2. Use the `flow --show flow_name -feature feature_name -count num` command.

The `num` value can range from 1 through 10. The default value is 1.

The following example creates a Flow Generator flow named “simflow\_1”, and then repeats the output three times.

```
switch:admin> flow --create simflow_1 -feature generator
 -srcdev 07f000 -dstdev 371400 -ingrport 12/16

switch:admin> flow --show simflow_1 -feature generator -count 3
=====
Name : simflow_1 Features: gen(Activated),mon(Activated) noConfig: Off
Definition: IngrPort(12/16),SrcDev(0x07f000),DstDev(0x371400)
Flow Generator (Activated):

| SrcDev | DstDev |

| 0x07f000 | 0x371400 |

Number of frames generated from IngrPort : 2.57G
=====
=====
Name : simflow_1 Features: gen(Activated),mon(Activated) noConfig: Off
Definition: IngrPort(12/16),SrcDev(0x07f000),DstDev(0x371400)
Flow Generator (Activated):

| SrcDev | DstDev |

| 0x07f000 | 0x371400 |

Number of frames generated from IngrPort : 2.57G
=====
=====
Name : simflow_1 Features: gen(Activated),mon(Activated) noConfig: Off
Definition: IngrPort(12/16),SrcDev(0x07f000),DstDev(0x371400)
Flow Generator (Activated):

| SrcDev | DstDev |

| 0x07f000 | 0x371400 |

Number of frames generated from IngrPort : 2.58G
=====
```

## Sorting Flow Generator Output

In Flow Vision, frames can be sorted whether or not subflows are present. Sorting the output allows you to highlight a selected aspect of the flow data.

To sort the flow output, perform the following steps:

1. Connect to the switch and log in using an account with admin permissions.
2. Use the `flow --show flow_name -feature feature_name -sortby column column_num` command. The column `column_name` value is the number of the output column on which the data is to be sorted. There is no space between “column” and the column number.

### NOTE

The `-sortby` parameter can only be applied when there is only one feature (monitor, mirror, or generator) specified in the `flow --show flow_name` command.

The following example creates a Flow Generator flow, and then shows the output sorted by column 2.

```
switch:admin> flow --create fgSort -fea gen -ing 10 -srcdev "*" -dstdev "*"
```

### Output before sorting:

```
switch:admin> flow --show fgSort
=====
Name : gen1 Features: gen(Activated) noConfig: Off
Definition: IngrPort(10),SrcDev(*),DstDev(*)

Flow Generator (Activated):

| SrcDev | DstDev |

| 0x040a00 | 0x041c00 |

| 0x040a00 | 0x040b00 |

| 0x040a00 | 0x041d00 |

| 0x040a00 | 0x012500 |

Number of frames generated from IngrPort : 2.33G
```

### Output after sorting:

```
switch:admin> flow --show fgSort -feature generator -sortby column2
=====
Name : gen1 Features: gen(Activated) noConfig: Off
Definition: IngrPort(10),SrcDev(*),DstDev(*)

Flow Generator (Activated):

| SrcDev | DstDev |

| 0x040a00 | 0x012500 |

| 0x040a00 | 0x040b00 |

| 0x040a00 | 0x041c00 |
```

```

| 0x040a00 | 0x041d00 |

Number of frames generated from IngrPort : 2.42G
=====

```

## Deactivating Flow Generator Flows

You can deactivate Flow Generator flows without deleting them. This allows you to have a “library” of flows that you can activate as needed without recreating them.

### NOTE

When a flow is deactivated, traffic stops for that flow and any subflows associated with it. When a flow is deactivated, the definition remains, but Flow Generator does not populate it with traffic.

To manually deactivate a Flow Generator flow, perform the following steps:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the flow `--deactivate flow_name -feature generator` command.

The following example deactivates the Flow Generator flow named “Flow1.”

```
switch:admin> flow --deactivate Flow1 -feature generator
```

## Customizing Flow Generator Flows

Flow Generator allows you to specify the frame payload size and pattern to be used for the Flow Generator flows by using the `flow --control` command.

### Frame Payload Size

Flow Generator allows you to define the frame payload size in bytes. The frame payload size value must be either 0 (which produces frames of random size using a multiple of 4 between 64 and 2048 bytes) or a multiple of 4 in the range from 64 through 2048 (64, 68, 72, 120, 140, 320, 512, and so on). The default payload size value is 2048.

There are two ways to change the frame payload size: you can change the default payload size or the payload size for a single flow.

To see the current generic payload size, enter the `flow --show -ctrlcfg` command; to see the payload size for a specific flow, enter the `flow --show flow_name -ctrlcfg` command.

### NOTE

There are limitations for Flow Generator learning flows and the `sys_gen_all_simports` flow; see [Flow Generator Limitations and Considerations](#) for specific details.

### Changing the Default Frame Payload Size

Assigning a frame payload size without defining a flow creates a default that applies to all flows created afterward. The following example sets the payload size default for all new flows to 512 bytes.

```
switch:admin> flow --control -feature generator -size 512
```

### Changing the Frame Payload Size for a Single Flow

Assigning a frame payload size explicitly to a flow affects only that flow and overrides the existing payload size for that flow. Changing a flow’s payload size can only be done to an inactive flow. The following example changes the payload size for the flow “F1024” to 1024 bytes.

```
switch:admin> flow --control F1024 -feature generator -size 1024
```

## Frame Payload Pattern

Flow Generator allows you to define the pattern to be used as the frame payload. The frame payload pattern must be an alphanumeric ASCII string between 1 and 32 characters in length. The default frame payload pattern value is 0, which produces a random pattern of alphanumeric ASCII characters with a variable string length between 1 and 32 characters.

There are two ways to change a frame payload pattern; you can change the default payload pattern, or you can change the payload pattern for a single flow.

To see the current generic payload pattern, enter the `flow --show -ctrlcfg` command; to see the payload pattern for a specific flow, enter the `flow --show flow_name -ctrlcfg` command.

### Changing the Default Frame Payload Pattern

Assigning a frame payload pattern without defining a flow creates a default that applies to all subsequently-created flows. The following example sets the default payload pattern for all new flows to "TestFlow."

```
switch:admin> flow --control -feature generator -pattern "TestFlow"
```

### Changing the Frame Payload Pattern for a Specific Flow

Explicitly assigning a frame payload pattern to a flow overrides the existing frame payload pattern for that flow, and affects only that flow. Changing a payload pattern can only be done to an inactive flow. The following example sets the default payload pattern for the flow F2 to "a5a5a5."

```
switch:admin> flow --control F2 -feature generator -pattern "a5a5a5"
```

## Flow Generator Examples


The following examples describe how to work with Flow Generator flows.

### Creating a Flow from a Specific Source ID to a Specific Destination ID

To create a flow between a specific source ID (SID) and a specific destination ID (DID), perform the following steps:

1. Create two SIM ports.
2. Create an active flow from the SID to the DID.

**Figure 50: A Flow from a Specific Source ID to a Specific Destination ID**



The following example creates the flow as shown in the figure above. SIM port 1/1 is the source port, and SIM port 1/2 is the destination device. The `flow --show flowCase1 -feature generator` command displays the SID frame count and then the DID frame count.

The second `flow --control` command must be run on the remote switch.

```

switch4:admin> flow --control -simport 1/1 -enable

switch5:admin> flow --control -simport 1/2 -enable <--On the remote switch

switch4:admin> flow --create flowcase1 -feature generator
 -ingrport 1/1 -srcdev 0x040100 -dstdev 0x050200

switch4:admin> flow --show flowCase1 -feature generator
Name: flowcase1 Features: gen(Activated) noConfig: Off
Definition: IngrPort(1/1),SrcDev(040100),DstDev(050200)
Flow Generator (Activated):

| SrcDev | DstDev |

| 0x040100 | 0x050200 |

Number of frames generated from ingrport: 19.46M

switch4:admin> flow --show -ctrlcfg
SimPort Information
-----|-----|-----|-----|-----|-----|
Slot	Port	PID	PWWN	SID Frame Count	DID Frame Count
1	2	050200	20:02:00:05:1e:e2:8e:00	0K	19480K
-----|-----|-----|-----|-----|-----|

Addressing mode information
Port Addressing Mode: index
Device Addressing Mode: PID
Flow Generator Information
Size: 2048
Pattern: Random (Default)
Flow mirror Information
enable_wrap

```

## Commands Related to Flow Generator

The following commands affect or are affected by Flow Generator. Refer to the *Brocade Fabric OS Command Reference Manual* for more detailed information.

### **portcfgshow**

Entering the `portcfgshow` command for a specified port shows the configuration of the SIM port. (Called out for illustration.)

```

switch:admin> portcfgshow 8/0
Area Number: 8
Octet Speed Combo: 1(32G|16G|8G|4G)
Speed Level: AUTO(HW)
AL_PA Offset 13: OFF
Trunk Port ON
Long Distance OFF
VC Link Init OFF
Locked L_Port OFF

```



|                    |                                  |
|--------------------|----------------------------------|
| Locked G_Port      | OFF                              |
| ISL R_RDY Mode     | OFF                              |
| RSCN Suppressed    | OFF                              |
| Persistent Disable | OFF                              |
| LOS TOV mode       | 0 (OFF)                          |
| NPIV capability    | ON                               |
| QOS Port           | AE                               |
| Port Auto Disable: | OFF                              |
| Mirror Port        | OFF                              |
| SIM Port           | ON <b>&lt;-- SIM Port status</b> |
| Credit Recovery    | ON                               |
| F_Port Buffers     | OFF                              |
| Fault Delay:       | 0 (R_A_TOV)                      |
| NPIV PP Limit:     | 126                              |
| NPIV FLOGI Logout: | OFF                              |
| CSCTL mode:        | OFF                              |
| TDZ mode:          | OFF                              |
| D-Port mode:       | OFF                              |
| D-Port over DWDM:  | OFF                              |
| 10G/16G FEC:       | ON                               |
| 16G FEC via TTS:   | OFF                              |

### **portperfshow**

When you run the `portperfshow` command, an asterisk (\*) represents traffic running on either an ingress port or an egress port from a Flow Generator flow. For the `portperfshow` command, the Transmit Throughput value represents the egress port, and the Receive Throughput value represents the ingress port.

### **portstatsclear**

Entering the `portstatsclear` command for a port clears the Flow Generator frame count for all flows sharing this port. This is equivalent to the `flow --reset` command.

### **portstatsshow**

When you run the `portstatsshow` command, the Transmit Word Count value represents the egress port, and the Receive Word Count value represents the ingress port. The frame size used for `portstatsshow` is 2048 bytes, regardless of what value has been set for the flow.

### **slotstatsclear**

Entering the `slotstatsclear` command for a slot clears the Flow Generator frame count for all flows sharing the ports on that slot.

### **switchshow**

Entering the `switchshow` command generates output showing which ports are set as simulation mode ports (SIM ports) and displays the WWN for each emulated device.

```
switch:admin> switchshow | grep SIM
192 8 0 046000 id 16G Online FC SIM-Port 20:c0:00:05:1e:99:61:00
193 8 1 046100 id 16G Online FC SIM-Port 20:c1:00:05:1e:99:61:00
194 8 2 046200 id 16G Online FC SIM-Port 20:c2:00:05:1e:99:61:00
195 8 3 046300 id 16G Online FC SIM-Port 20:c3:00:05:1e:99:61:00
```

```
196 8 4 046400 id 16G Online FC SIM-Port 20:c4:00:05:1e:99:61:00
```

## SIM Port Attributes and Configuration

Flow Vision requires that the source device and destination device ports to be in simulation mode (SIM port mode) before activating the predefined or static test flows and checks for this before activating the test flows.

Prior to creating and activating flows, you must enter `flow --control` on the local switch to set the source device (srcdev) and the destination device (dstdev) ports as SIM ports. This ensures that test flows are not unintentionally transmitted to real devices. After the source device and destination device ports are configured to be SIM ports, you can create and activate the flow.

The following restrictions will affect your use of SIM ports:

- Zoning is bypassed on SIM ports for both predefined and static flows. In these cases, traffic will reach its destination regardless of the zoning configuration.
- Zones are used to gather the Source ID-Destination ID pairs for learning flows.

## SIM Port Criteria

Flow Generator simulation (SIM) ports must meet the following criteria to be valid:

- SIM ports are supported on ASICs that support either 8, 16, 32, or 64G-capable Fibre Channel ports. Source devices or ingress ports can only be on 16G-capable, 32G-capable or 64G-capable Fibre Channel ports. Destination devices or egress ports can be on either 8, 16, 32, or 64G-capable Fibre Channel ports.
- SIM ports cannot be in the base switch or Access Gateway.
- SIM ports, E\_Port, and EX\_Port cannot be configured on a port that is online and connected to a real device.

### NOTE

If a port is connected to a real device, you can disable the port, configure the SIM port, and then re-enable the port. The port will be a SIM port; the real device will not join the fabric.

- Existing SIM ports are added to Device Connection Control (DCC) policies when created with a wildcard (\*) but are not adhered to. These SIM port entries must be deleted if a new WWN is connected.
- SIM ports cannot be configured as any of the following port types; these restrictions also apply at the time a SIM port is enabled:
  - Any port running Encryption or Compression
  - Any F\_Port connected to a real device (unless the port is disabled)
  - Any port configured with F\_Port buffers
  - D\_Port (Diagnostic Port)
  - F\_Port trunked
  - Fastwrite port
  - FCoE port
  - ICL port
  - L\_Port
  - M\_Port (Mirror Port)
- If a port is configured with CSCTL\_mode and/or QoS enabled, you can configure it as a SIM port.
- If a port has an Ingress Rate Limit set, you can configure it as a SIM port.
- If a port is configured as a SIM port, you can set an Ingress Rate Limit.
- The following features of a SIM port are persistent across a reboot:
  - Each SIM port is assigned a PID and is displayed in a `switchshow` command.
  - Each SIM port's Port Worldwide Name by default is the switch Port Worldwide Name (PWVN), unless a user-defined Virtual Port Worldwide Name is assigned to it.
  - Each SIM port registers itself into the Name Server database.

## Identifying SIM Ports

To distinguish SIM ports from other devices, the Name Server commands show Port Properties: SIM Port, as called out in the following example:

```
switch:admin> nscamshow
nscam show for remote switches:
Switch entry for 105
state rev owner cap_available
known v730 0xffffc6e 1
Device list: count 12
Type Pid COS PortName NodeName
N 691000; 3; 10:00:00:00:00:0f:00:00;10:00:00:00:00:0f:00:00;
Fabric Port Name: 20:10:00:05:1e:57:dc:b3
Permanent Port Name: 10:00:00:00:00:0f:00:00
Port Index: 16
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No ← Port Properties value is not shown for non-SIM ports
N 691400; 2,3; 20:14:00:05:1e:57:dc:b3;20:14:00:05:1e:57:dc:b3;
Fabric Port Name: 20:14:00:05:1e:57:dc:b3
Permanent Port Name: 20:14:00:05:1e:57:dc:b3
Port Index: 20
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
Port Properties: SIM Port ← Port Properties value shows "SIM Port"
(output truncated)
```

```
switch:admin> nsshow
Type Pid COS PortName NodeName TTL(sec)
N 691000; 3; 10:00:00:00:00:0f:00:00;10:00:00:00:00:0f:00:00; na
Fabric Port Name: 20:10:00:05:1e:57:dc:b3
Permanent Port Name: 10:00:00:00:00:0f:00:00
Port Index: 16
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
LSAN: No
Device link speed: 8GConnected through AG: YesReal device behind AG: Yes ← Port Properties value is
not shown for non-SIM ports
N 691400; 2,3; 20:14:00:05:1e:57:dc:b3;20:14:00:05:1e:57:dc:b3; na
Fabric Port Name: 20:14:00:05:1e:57:dc:b3
Permanent Port Name: 20:14:00:05:1e:57:dc:b3
Port Index: 20
Share Area: No
Device Shared in Other AD: No
```

```

Redirect: No
Partial: No
LSAN: No
Device link speed: 8GConnected through AG: YesReal device behind AG: YesPort Properties: SIM Port <-
Port Properties value shows "SIM Port"
(output truncated)

```

## Sending Traffic Using a Fabric-Assigned WWN

If you want to use a Fabric-Assigned WWN (FA-WWN), you need to set the FA-WWN on the SIM port using the Dynamic Fabric Provisioning command, `fapwwn - assign`. For details, refer to the "Dynamic Fabric Provisioning" section of the *Brocade Fabric OS Administration Guide*.



### CAUTION

If the `fapwwn` command is used to assign a user-defined Port WWN to a SIM port, it is the responsibility of the person making the assignment to not assign a Port WWN that duplicates one already in the fabric. If there is a duplicated WWN, both entries will be removed from the Name Server database. This has a high probability of disrupting traffic.

## Flow Generator and High Availability

On a High Availability (HA) failover, an HA reboot, or a power-cycle and reboot, both local and remote flows remain active. When SIM ports come back online from an HA failover or HA reboot event, local flows are re-created and reactivated, and local traffic is restarted. Inactive flows are re-created but not activated by an HA failover, HA reboot, or power-cycle and reboot.

### ATTENTION

A HA failover, an HA reboot, or a power-cycle and reboot event may cause different subflows to be re-created, as the flow order depends on the zone database. If the number of zoned SID-DID pairs is greater than the maximum permitted number of flows, then any pairs that appear in the zone database after that limit is reached will be ignored.

If the `sys_gen_all_simports` flow was active prior to an HA failover, the flow will be replayed after the HA failover with all SIM ports that are configured on the switch.

See [High Availability and Flow Vision](#) and [Notes on Predefined Flows](#) for more information.

## Flow Generator and MAPS

The Brocade Monitoring and Alerting Policy Suite (MAPS) can be used to monitor SIM port traffic thresholds while Flow Generator flows are running.

MAPS treats SIM ports as F\_Ports, so MAPS can issue warnings on these ports if threshold values are triggered. If you do not want to see MAPS warnings for SIM ports, you must disable MAPS monitoring for those ports.

Flow Generator traffic will also impact E\_Ports; this may cause MAPS warnings for E\_Port throughput levels. Refer to the *Brocade Fabric OS MAPS User Guide* for more information about working with MAPS.

You can use MAPS to monitor traffic in the `sys_gen_all_simports` flow using the same default and custom policies as used for F\_Ports. However, there will not be a check in MAPS for maximum throughput for SIM ports, because Flow Generator always runs at maximum throughput.

## Flow Generator Limitations and Considerations

The following limitations apply specifically to Flow Generator:

- Flow Generator traffic competes with any other traffic on a fabric. So, E\_Ports and FCIP links can become congested when using Flow Generator, leading to the throughput degradation. FCIP links are more prone to congestion than E\_Ports. If used on a live production system, this can cause traffic issues.
- `sys_flow_monitor` is enabled in all FIDs by default. Disable `sys_flow_monitor` in all FIDs and activate generator to enable a generator.
- The maximum number of Flow Generator flows allowed per ingress port varies by platform type from 4 to 39.
- User-defined flows for supported cases like a flow generator can be created and activated only on online ports.
- Activating more than five static flows on a fixed-port switch or five static flows per blade in a chassis switch might result in a DMA buffer full error message. Before upgrading, make sure that you will not exceed these limits after upgrading.
- Flow Generator flows can only be mirrored at the ingress port; they cannot be mirrored at the egress port.
- Flow Generator is not supported on Access Gateways.
- If a SIM-port configuration is deleted while the port is online, Flow Vision automatically stops all Flow Generator data flowing to those ports, but the flows themselves are not deactivated.
- Zoning is not enforced. Sources and destinations can be in different zones except for the learning flow.
- Flow Generator gathers source and destination pairs from the zoning database for learning flows only when the flow is activated. Subsequent changes to this database will not be registered until a flow is reactivated.
- Flow Generator flows support a frame size range from 64 bytes to 128 bytes. You can set this value from 64 to 2048 using either BNA or the CLI. Use the command `flow --control flow_name -ctrlcfg` to display the active frame size. The frame size is restricted internally when the flow is activated.

For Flow Generator learning flows, if the domain or PID is changed in a remote switch, the following points should be considered:

- Flow Generator indicates that the device is offline.
- The flow is not enforced.

### **Considerations When Using Flow Generator Across FCR**

The following considerations apply when you use Flow Generator across Fibre Channel Routing (FCR):

- In backbone fabrics, only FCR-enabled domains that are running Fabric OS 9.1.x or a later version can sync SIM-port and FC4 feature parameters with each other.
- For Fabric OS 9.1.x, this capability is supported only for edge-to-edge topologies; it is not supported in backbone-to-edge topologies. Also, All the switches from edge to edge must be running Fabric OS 9.1.x or later.

## **Flow Mirror**

As storage networks get larger and more complicated, non-intrusive diagnostic tools are becoming increasingly important to help identify problems without disturbing the operating fabric. Flow mirroring is a diagnostic feature within Flow Vision that addresses this need.

With Flow Mirror, you can do the following:

- Non-disruptively create copies of application flows that can optionally be captured for in-depth analysis.
- Define a traffic pattern and create a real-time copy of this traffic, allowing you to analyze a live system without disturbing existing connections. You can also use this feature to view traffic passing through a port.
- Use a MAPS logical port group for either an ingress or egress port. Note that logical port groups are not supported on CPU flow mirroring (CFM) or logical flow mirroring (LFM).
- Use a predefined flow with the `flow --show` command to mirror SCSI command, first response, status, and ABTS frames from all the Gen 6 F\_Ports on a switch. For example: `flow --show sys_analytics_vtap`.

Flow Mirror duplicates the specified frames in a user-defined flow and sends them to a sink. This sink could be either:


- The local switch control central processor unit (CPU); this form is called “CPU flow mirroring” (CFM), and has a limit of 256 frames per second.
- An analyzer/packet sniffer connected through a port in the metaSAN. The bandwidth limit for a flow using this type of link is the bandwidth of the mirror destination port. This form is called “Local flow mirroring” (LFM), and mirrors the flow to a port on the same physical switch. This requires that a loopback SFP be plugged in at the other end of the analyzer or on the port configured as a mirror port, which must be in the same domain.

**NOTE**

All forms of mirroring are mutually exclusive. Only one form of mirroring (LFM, CFM, or RFM) can be active at a time.

Flow Mirror flows can be in an active or inactive state. If the mirror flow is “active”, mirroring starts immediately; if the flow is “inactive”, the flow must be activated (by using the `flow --activate` command) for mirroring to start. Mirrored flows can be unidirectional or bidirectional. A sample use case would be to mirror the traffic flow from a slow-draining F\_Port to see what is causing this condition. The following figure provides a diagram of a flow that mirrors to the CPU the traffic passing into the switch through the ingress port. Flow Mirror can similarly mirror the egress port, but only one port (ingrport or egrport) can be mirrored per flow. To mirror from one port in both flow directions (left to right and right to left in the figure), the `-bidir` keyword must be used in the flow definition.

**Figure 51: Illustration of CPU Flow Mirroring**

**NOTE**

Flow Mirror does not support any NVMe specific keywords such as `nvme`, `nsid`, `nvmerread`, or `nvmerwrite`.

## Creating Flow Mirror Flows

To create a Flow Mirror flow, use the `flow --create flow_name -feature mirror parameters` command.

When you create a flow, it is automatically activated unless you use the `-noactivate` keyword as part of the `flow --create` command. Refer to [Creating an inactive flow in Flow Mirror](#) for an example.

The following table shows the supported Flow Mirror flow parameter combinations.

**Table 126: Flow Parameter Combinations Supported by Flow Mirror**

| Parameters                         | Field Names                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port                               | ingrport<br>egrport                  | <ul style="list-style-type: none"> <li>One field only must be specified.</li> <li>Values must be explicit.</li> <li>Can only be an F_Port local to the switch, a Gen 6 (32G) F_Port or Gen 6 F_Port trunk on the local domain.</li> <li>For CFM and LFM, the values for ingrport and egrport can only be a single F_Port or F_Port trunk.</li> <li>For the predefined flow, sys_analytics_vtap, the values for ingrport and egrport can be a single F_Port, an F_Port trunk, an asterisk in quotation marks ( "*" ), or a MAPS logical port group. Mirroring is active only on the F_Ports, even though the port is specified as "*" or as a MAPS logical group.</li> </ul> |
| Frame                              | srcdev<br>dstdev<br>LUN<br>frametype | <ul style="list-style-type: none"> <li>Only one field can be specified.</li> <li>Values for srcdev and dstdev can be explicit or "*" ( "*" indicates learned flows).</li> <li>Values for LUN and frametype must be explicit.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Optional keyword parameters</b> |                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|                                    | -bidir                               | Adding this keyword makes the application mirror traffic in both directions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|                                    | -noactivate                          | Adding this keyword creates the flow without activating it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|                                    | -noconfig                            | Adding this keyword creates the flow without saving the flow to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

The following example creates a Flow Mirror flow named "flowmirror14" that mirrors traffic flowing from device "080e00" to device "080f00" ingressing through the port 14 on the switch on which this command was run.

```
switch1234:admin> flow --create flowmirror14 -feature mirror -ingrport 14 -srcdev 080e00 -dstdev 080f00
Mirror feature(s) have been activated.
```

## Flow Mirror Limitations and Restrictions

The following limitations and restrictions apply specifically to Flow Mirror flows and flow mirroring:

- Flow Mirror does not support NVMe frames.
- `sys_flow_monitor` is enabled in all FIDs by default. To enable a mirror, disable `sys_flow_monitor` in all the FIDs and activate mirror.
- Flow Mirror is not supported on the Brocade G630 ports 96 to 127 in the VF mode. Flow Mirror is supported on these ports in non-VF mode.
- ASL connectivity is not supported on the Brocade G630 ports 96 to 127. If ASL is enabled on these ports, the ASL ports are segmented.
- Mirror Port is not supported on the Brocade G630 ports 96 to 127. If these ports are enabled as Mirror Ports, then the ports are disabled.
- If a CFM flow is defined on a F\_Port trunk group, then the CFM flow must be deactivated before making the entire trunk offline or before running the `switchdisable` or `chassisdisable` command.
- Flow Mirror is supported only on Gen 6 and Gen 7 Fibre Channel platforms. Refer to [Supported Hardware and Software](#) for a list of these platforms.
- Activation of the predefined Flow Mirror flow, `sys_analytics_vtap`, fails if the chosen `AE_port` is not a Gen 6 port.
- The `sys_analytics_vtap` flow can be enabled simultaneously on multiple logical fabrics for Gen 6 slotted chassis and on the Brocade G630 platform.
- Before activating the RFM flow on Gen 6 platforms, turn on the chassis-level configuration feature vTAP and QoS High Priority Zone Compatibility mode. The configuration feature is required in order for RFM and QoS High Priority Zones to co-exist. Use the following configuration:

```
switch:admin> configurechassis
```

```
Configure...
```

```
cfgload attributes (yes, y, no, n): [no]
```

```
Custom attributes (yes, y, no, n): [no]
```

```
system attributes (yes, y, no, n): [no]
```

```
fos attributes (yes, y, no, n): [no] y
```

```
Reboot needed to effect new CSCTL Mode
```

```
CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [0]
```

```
Chassis SDDQ Limit: (0..32) [10]
```

```
ChassisDisble followed by ChassisEnable needed to effect new Mode
```

```
vTap and QOS High Priority Zone Compatibility Mode (on, off): [on]
```

```
switch:admin>
```

- The mirror port is specified using `-mirrorport` in the `flow` command. Not specifying this option indicates that the flow is a CFM flow. For an LFM flow, the `-mirrorport` is the index of a local port in the chassis. An analyzer can be connected on this port. For the `sys_analytics_vtap` flow, the `-mirrorport` refers to the logical Fabric AF\_Port of the AMP that is connected to the logical switch or fixed-port switch.
- Of the three types of Flow Mirror flows—CFM, LFM, and RFM (predefined)—only one type of flow can be active, because these types are exclusive of each other.
- The predefined Flow Mirror flow is not supported in the base switch. You must deactivate the predefined Flow Mirror flow before you convert a logical switch to a base switch.
- A mirror port cannot be either the ingress or egress port specified in an LFM flow.
- The flow ingress or egress ports must refer to ports on the current switch or a MAPS logical group.
- You can use the following options of the `flow --control` command only for CPU flow mirroring:



- `flow --control enable_wrap`
- `flow --control disable_wrap`
- For an LFM flow, if the combined rate of bidirectional traffic exceeds the limit of the mirror port traffic limit, then the Flow Mirror flow is not permitted. (For example, the combined rate of a bidirectional flow from a 16G port is 32G, so it cannot be mirrored to a 16G mirror port.)
- Mirroring of exported and imported devices is supported only on edge fabrics, not on backbone or base fabrics.
- Flow Mirror is not supported on SIM ports that are specified as ingress or egress ports in the flow definition.
- When using Flow Mirror in Access Gateway mode, you cannot configure internal ports on SAN I/O modules as mirror ports.
- Flow Mirror cannot mirror:
  - Frames belonging to device-switch communication (for example, a FLOGI or PLOGI action).
  - Link Primitives, discarded frames, frames from a remote Control Unit Port (CUP), Link Control Frames, or frames containing domain controller addresses used as source IDs.
- For flows mirrored to a CPU, only the first 256 frames of a second are mirrored. If a greater number of frames that match the flow definition within a second are identified, those later frames are not mirrored.  
For example, on a fixed-port switch, if 500 frames meet the flow definition in the first second of the Flow Mirror operation, only the first 256 frames are mirrored. In the next second, frame mirroring will begin with the 501st frame that matches the flow definition. The intervening frames will not be mirrored, even though they match the flow definition.
- Although the system attempts to handle misconfigurations of zoning, do not place mirror ports in zones with any other port or device.

### **CFM Frame Rates and Frame Capacity**

The following table shows the maximum frame rate and mirroring capacity for each platform type.

**Table 127: Flow Mirror CFM Frame Rates and Frame Capacity**

| Platform Type         | Maximum Rate (frames per second) | Maximum Capacity (frames) |
|-----------------------|----------------------------------|---------------------------|
| Fixed-port switch     | 256                              | 1280                      |
| Chassis-based systems | 256                              | 5120                      |


## **Local Flow Mirroring**

Local flow mirroring (LFM) allows you to mirror a flow to a port in the same domain that the flow has been defined in. This mirrored data can then be analyzed through an external analyzer/frame sniffer connected to the port.

To mirror traffic to a local port, the port must be configured as a mirror port before it is used in the flow definition. A loopback SFP should be plugged in at the other end of the analyzer or on the port configured as a mirror port. Here is the pattern for a command creating a flow definition to mirror a local flow:

- **First:** Use the `portcfg mirrorport port_num --enable` command to configure the mirror port.
- **Next:** Use the `flow --create flow_name -feature mirror -srcDev SrcDevID -dstDev DestDevID -ingrport portID -bidir -mirrorport port_num` command to create the flow that is mirrored to that port.

A typical configuration for local flow mirroring is shown below.

**Figure 52: Sample Local Flow Mirroring Configuration**

The following example configures port 193 as a mirror port, then creates a flow for all traffic flowing from device 02ca40 to device 080200 ingressing through port 214 that sends the mirrored frames to the mirror port, and finally displays a list of the flows followed by the results of the flow.

```
switch:admin> portcfg mirrorport 193 --enable
```

```
flow --create fctest -feature mirror -srcdev 02ca40 -dstdev 080200 -ingrport 214 -mirrorport 193
Mirror feature(s) have been activated.
```

```
switch:admin> flow --show
```

| Flow Name            | Feature   | SrcDev         | DstDev | IngrPt | EgrPt | BiDir | LUN |
|----------------------|-----------|----------------|--------|--------|-------|-------|-----|
| NSID                 | FrameType | SFID DFID SEID | MirPt  |        |       |       |     |
| sys_gen_all_simports | gen       | *              | *      | *      | *     | no    | -   |
| sys_analytics_vtap   | mir       | *              | *      | *      | *     | yes   | -   |
| sys_flow_monitor     | mon       | *              | *      | *      | *     | no    | -   |
| fctest               | mir+      | 02ca40         | 080200 | 214    | 193   | no    | -   |

+ Denotes feature is currently activated for the flow  
The flow name with prefix sys\_ denotes predefined flow

```
switch:admin> flow --show fctest
```

```
Name : fctest Features: mir(Activated) noConfig: Off
Definition : IngrPort(214),SrcDev(0x02ca40), DstDev(0x080200),MirPort(193)
```

```
Flow Mirror (Activated):
```

```
No of Mirrored Frames : 16431161, No of RX Mirrored Frames : 16431161, No of TX Mirrored Frames : 0
```

After a port is configured as a mirror port, the `switchshow` command output looks similar to the following example.

```
switch:admin> switchshow|grep Port
```

| Index | Slot | Port | Address | Media | Speed | State     | Proto          |                                |
|-------|------|------|---------|-------|-------|-----------|----------------|--------------------------------|
| 40    | 3    | 40   | 02f5c0  | id    | N16   | Online    | FC F-Port      | 10:00:00:90:fa:94:28:0d        |
| 41    | 3    | 41   | 02f580  | id    | N16   | Online    | FC F-Port      | 1 N Port + 2 NPIV public       |
| 44    | 3    | 44   | 02f4c0  | --    | N32   | No_Module | FC D-Port      |                                |
| 128   | 7    | 0    | 02dfc0  | id    | N8    | Online    | FC F-Port      | 10:00:00:05:33:48:c7:9b        |
| 146   | 7    | 18   | 02db40  | id    | N8    | Online    | FC F-Port      | 20:07:00:11:0d:7a:01:00        |
| 147   | 7    | 19   | 02db00  | id    | N8    | Online    | FC F-Port      | 20:06:00:11:0d:7a:00:00        |
| 159   | 7    | 31   | 02d800  | id    | N8    | Online    | FC F-Port      | 1 N Port + 3 NPIV public       |
| 164   | 7    | 36   | 020000  | id    | N16   | Online    | FC F-Port      | 1 N Port + 4 NPIV public (AoQ) |
| 166   | 7    | 38   | 02d640  | id    | N8    | Online    | FC F-Port      | 1 N Port + 3 NPIV public       |
| 193   | 8    | 1    | 02cf80  | id    | 16G   | Online    | FC Mirror Port |                                |
| 194   | 8    | 2    | 02cf40  | id    | N32   | Online    | FC F-Port      | 10:00:00:90:fa:94:36:c5        |
| 203   | 8    | 11   | 02cd00  | id    | N8    | Online    | FC F-Port      | 20:01:00:11:0d:6a:01:00        |
| 204   | 8    | 12   | 02ccc0  | id    | N8    | Online    | FC F-Port      | 20:00:00:11:0d:6a:00:00        |
| 214   | 8    | 22   | 02ca40  | id    | 32G   | Online    | FC F-Port      | 10:00:00:90:fa:94:27:9b        |

After a port is configured as a mirror port, the `portcfgshow` command output looks similar to the following example.

```
switch:admin> portcfgshow 8/1
Area Number: 207
Octet Speed Combo: 1(32G|16G|8G|4G)
FC Speed Level: 16G
Trunk Port ON
Long Distance OFF
VC Link Init OFF
Disabled E_Port OFF
Locked E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable OFF
LOS TOV mode 0(OFF)
DWDM losync mode OFF
NPIV capability ON
QOS Port AE
Port Auto Disable: OFF
EX Port OFF
Mirror Port ON
SIM Port OFF
Credit Recovery ON
F_Port Buffers 32
E_Port Credits OFF
Fault Delay: 0(R_A_TOV)
NPIV PP Limit: 126
NPIV FLOGI Logout: OFF
CSCTL mode: OFF
TDZ mode: OFF
Clean Address Bit: OFF
D-Port mode: OFF
D-Port over DWDM: OFF
Compression: OFF
Encryption: OFF
10G/16G FEC: ON
16G FEC via TTS: OFF
App Srvr Access: RW
EFCS Access: RW
```

|                    |    |
|--------------------|----|
| FCS Access:        | RW |
| FDMI Access:       | RW |
| Zone Access:       | RW |
| Unzoned NS Access: | RO |

## Creating an Inactive Flow in Flow Mirror

To create an inactive Flow Mirror flow, enter the `flow --create flow_name -feature mirror flow_parameters -noactive` command.

The following example creates an inactive Flow Mirror flow named “mirroredflow120” from device 020a00 to device 01c000 ingressing through port 120.

```
switch:admin> flow --create mirroredflow120 -noactivate -feature mirror -ingrport 120 -srcdev 0x020a00 -dstdev 0x01c000
```

See [Activating Flow Mirror Flows](#) for information on activating a Flow Mirror flow.

## Remote Flow Mirroring

Remote Flow Mirroring (RFM) permits you to route a flow from a switch to a Brocade Analytics Monitoring Platform so that the flow frame data can be monitored and analyzed.

### NOTE

This feature is only supported when utilizing the Brocade Analytics Monitoring Platform, and the mirrored frames can be sent only to the analytics platform.

To correctly route the mirrored frames, Remote Flow Mirroring must be activated on the source switch F\_Port, and the destination port for the mirrored flow must be a Brocade Analytics Monitoring Platform analytics port (AF\_Port). The RFM flow can monitor either one F\_Port, all F\_Ports, or a specific group of F\_Ports defined using a MAPS logical group. Flow Vision will automatically learn the RFM mirror port based on the connected analytics platform. The mirrored traffic uses VC14.

If RFM is already active in the logical switch, then the RFM filter flow is updated to include the new F\_Port. The VTAP frames belonging to the new F\_Port being mirrored can be viewed as soon as it is online.

Remote flow monitoring also allows for simultaneous mirroring of VTAP frames (SCSI CMD, STS, 1st XFER\_RDY, and 1st RD\_RSP) that belong to multiple logical fabrics (VFs). Simultaneous mirroring is only supported on Brocade G630 switch and Gen 6 chassis systems.

### NOTE

vtap is supported only on Gen 6 platforms running Fabric OS 9.0.x: the Brocade X6-4 and X6-8 directors, the Brocade G610, G620, and G630 switches.

The following restrictions and limitations apply to Remote Flow Mirroring (RFM).

- RFM is not supported on F\_Ports 96–127 of the Brocade G630 in VF mode.
- RFM is supported on all the F\_Ports, including 96–127 of the Brocade G630 in non-VF mode.
- ASLs and mirror ports are not supported on ports 96 to 127 of the Brocade G630.
- While enabling RFM on an F\_Port trunk controller or agent ingress port, the following conditions apply:
  - If the port is an F\_Port trunk controller and the entire F\_Port trunk goes down and comes up with a new controller port, the mirroring is stopped.
  - While activating the flow, if the port is an F\_Port trunk agent and if that port becomes the controller, then the `sys_analytics_vtap` flow should be deactivated and reactivated for the mirroring to resume.
  - During the activation of the `sys_analytics_vtap` flow, if the port is an F\_Port trunk controller and the port becomes an agent or goes offline, consider the following:
    - The mirroring does not stop if one of the following applies:

A member of the F\_Port trunk is online and traffic is flowing through them.

The same port comes up as controller without the entire trunk going offline.

- The mirroring stops if one of the following applies:
  - None of the members in that trunk group are online.

The entire trunk goes offline and then comes online with a new controller port. If so, the `sys_analytics_vtap` flow must be deactivated and reactivated for the mirroring to resume.
- RFM can only mirror frames to the Brocade Analytics Monitoring Platform.
- RFM is not supported in Access Gateway mode.
- RFM is supported on Gen 6 SX6 FCIP blade.
- No user-defined flow creation is supported for RFM. You must use the existing predefined `sys_analytics_vtap` flow.
- Before activating an RFM flow, you must turn on the chassis-level configuration feature "vTAP and QoS High Priority Zone Compatibility Mode." With this feature activated, RFM/vTAP and QoS High Priority can co-exist.
- Because `sys_analytics_vtap` is a flow that is pre-defined in the system, it cannot be deleted nor can a new one be created. The permitted operations on this flow are "show," "modify," "deactivate," "activate," and "reset."
- When set up as a learning flow, the predefined `sys_analytics_vtap` flow uses the first mirror port it learns about as the remote mirror port.
- Remote Flow mirroring does not function if the remote domain or AF\_Port is not accessible.
- You can only modify the ingress port and mirror port of the predefined `sys_analytics_vtap` flow. The ingress port for a RFM can be "\*" , a single F\_Port, or a MAPS logical port group.
- RFM traffic always uses only VC14 for reaching a Brocade Analytics Monitoring Platform.
- Flow Mirror supports using only Gen 6 ports as AE\_Ports to connect a switch and a Brocade Analytics Monitoring Platform.
- RFM multi-VF feature is not supported on fixed-port switches, except for the Brocade G630 switch.
- RFM multi-VF feature requires AE\_Port DISLs in each logical switch for which a predefined `sys_analytics_vtap` flow is supposed to be activated. This feature is not supported with LISLs.
- RFM uses only a single AE\_Port DISL per fabric; that is, DPS is not supported for mirrored VTAP frames. When multiple AE\_Ports are present in the logical switch, only one AE\_Port is picked up to transport mirrored frames. However, an AE\_Port trunk is supported.
- Activation of features for the predefined `sys_analytics_vtap` flow cannot be enforced if the AF\_port is either unreachable or offline.
- The following limitations apply when using in-flight encryption or compression with RFM:
  - When in-flight encryption and compression are not enabled on any of the ports on an ASIC, then RFM functions at full capacity on that ASIC.
  - When in-flight encryption or compression is enabled, and RFM exists on the same ASIC, and the number of frames being mirrored exceeds 250K IOPS, then the RFM is unenforced on the F\_Ports of the ASIC.

#### NOTE

Certain versions of Fabric OS support connections to the Brocade Analytics Monitoring Platform. For more information about the Brocade Analytics Monitoring Platform, including which versions of Fabric OS can interact with it, refer to the Brocade Analytics Monitoring Platform documentation or MyBrocade.com.

## CFM and LFM Flow Support for Trunk Groups

CFM and LFM flows support mirroring of frames on trunk ports subject to the following conditions:

- You must create the same flow on all trunk member ports.
- If you create a flow on a agent port without using the `-noactivate` keyword in the flow command, the flow is then automatically activated when the agent port becomes the controller port.
- After a switch initialization or a cold or warm recovery, existing flows are re-created on both controller and agent ports, but only those flows associated with the controller port are activated.

To create CFM or LFM flow mirror on a trunk group, perform the following steps:

1. Identify the trunk group members using the `switchshow` command.
2. Create individual CFM or LFM flows for each member of the trunk group using the `flow --create` command with same flow definition.
3. Run the `flow --show` command with the `flow_name` of the controller port to view the mirror data for controller port.
4. If you run the `flow --show` command with the `flow_name` of the agent ports, then no mirror data is displayed.

## Activating Flow Mirror Flows

To activate an inactive Flow Mirror flow, perform the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter `flow --activate flow_name -feature mirror`.

The following example activates the Flow Mirror flow named "Flow1."

```
switch:admin> flow --activate Flow1 -feature mirror
```

### NOTE

Activating a Flow Mirror flow automatically clears all the flow statistics for that flow.

## Automatic Activation of a Flow Mirror Flow

Flow Mirror automatically activates a mirroring flow under the following conditions:

- On flow creation, unless the flow is created using the `-noactivate` keyword as part of the `flow` command.
- On slot power-on, if the port specified is part of the slot being powered on and the flow was active when the slot was powered off.
- On a High Availability (HA) failover, HA reboot, or a power cycle, if the flow was active when the HA event occurred.

### NOTE

Flow Mirror will not automatically re-activate a flow if the port types are other than Gen 6 (32G) F\_Ports or F\_Port trunks.

## Viewing Flow Mirror Flows

The following sections describe the different ways of viewing frames mirrored by Flow Mirror.

### Summary Information View of a Flow Mirror Flow

To display the summary view of a Flow Mirror flow, enter `flow --show flow_name -feature mirror`. In the summary information view, the first output line lists the flow name and the flow features; the second line lists the source and destination devices and ports and the flow's directionality; the third line identifies the active features. The following example presents a table with a line for each mirrored frame listing the destination ID, the OXID (originator exchange identifier) of the flow, the RXID (responder exchange identifier) of the flow, the start of frame and end of frame values, the frame type, the LUN, the direction of the frame, and the timestamp of the frame. A learned field column (indicated by an asterisk (\*) adjacent to the column name) is added to the output only if the flow definition does not contain the keyword associated with the field for this column. The last data line displays the total number of frames sent to the CPU, the total

number of received frames, and the total number of transmitted frames. The following example displays the summary information recorded for the Flow Mirror flow named "fmshow." The **LUN** column is a learned field in this example.

```
switch:admin> flow --show fmshow -feature mirror
=====
Name : fmshow Features: mir(Activated) noConfig: Off
Definition: IngrPort(2),SrcDev(*)
Flow Mirror (Activated):

| SID(*) | DID(*) | OXID | RXID | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp |

040200	040f00	0ca0	ffff	SOFi3	EOFt	SCSIRead	0000	Rx	Jul 12 06:29:13:639
040200	040f00	0bea	0bel	SOFn3	EOFn	Data	----	Rx	Jul 12 06:29:13:639
040200	040f00	0f42	ffff	SOFi3	EOFt	SCSIWrite	0000	Rx	Jul 12 06:29:13:639
040200	040f00	0dc0	ffff	SOFi3	EOFt	SCSIRead	0000	Rx	Jul 12 06:29:13:639
040200	040f00	0592	ffff	SOFi3	EOFt	SCSIRead	0000	Rx	Jul 12 06:29:13:639
.....									
000000	fffffe	801e	ffff	SOFi3	EOFt	FLOGI	----	Rx	Jul 12 06:29:45:292
040200	fffffc	8024	ffff	SOFi3	EOFt	PLOGI	----	Rx	Jul 12 06:29:49:411
040200	fffffd	8028	ffff	SOFi3	EOFt	ELSframe	----	Rx	Jul 12 06:29:49:411
040200	fffc04	88b9	ffff	SOFi3	EOFt	01	----	Rx	Jul 12 06:29:51:614
040200	fffc04	802c	ffff	SOFi3	EOFt	PRLI	----	Rx	Jul 12 06:29:51:614

No of Mirrored Frames : 528, No of RX Mirrored Frames : 528, No of TX Mirrored Frames : 0
=====
```

## Verbose Information View of a Flow Mirror Flow

To display all the information recorded for a flow, enter `flow --show flow_name -feature mirror -verbose`. In the verbose information view, the first output line lists the flow name and the flow features; the second line lists the source and destination devices and ports; the third line identifies the active features. The following lines list for each frame the time stamp of the frame, the direction of the frame, the start of frame and end of frame values, the frame type, and the first 64 bytes (16 words) of the frame. The last data line displays the number of frames sent to the switch Control Processor Unit (CPU), the number of received frames, and the number of transmitted frames. If any learned field (indicated by an asterisk (\*) adjacent to the column name) is part of the flow definition, then that field is displayed in the show output (there will be a column in the output representing this field). The following example displays all the information recorded for the Flow Mirror flow named "fmshow." The **LUN** column is a learned field in this example, and the **Frame Contents** column has been trimmed.

```
switch:admin> flow --show fmshow -feature mirror -verbose
=====
Name : fmshow Features: mir(Activated) noConfig: Off
Definition: IngrPort(2),SrcDev(*)
Flow Mirror (Activated):

| Time-Stamp | Dir | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Frame Contents |

Jul 12 06:29:13:637	Rx	SOFn3	EOFn	Data	----	01040f00 00040200 08000008 ...
Jul 12 06:29:51:614	Rx	SOFi3	EOFt	PRLI	----	22fffc04 00040200 01290000 ...
Jul 12 06:29:51:622	Rx	SOFi3	EOFt	01	----	23fffc04 00040200 01990000 ...
Jul 12 06:29:51:625	Rx	SOFi3	EOFt	01	----	23fffc04 00040200 01990000 ...
Jul 12 06:30:10:951	Rx	SOFi3	EOFt	Abort	----	81ffffffc 00040200 00090000 ...
Jul 12 06:30:12:970	Rx	SOFi3	EOFt	20	----	02ffffffc 00040200 20290000 ...
```

(output truncated)

```

No of Mirrored Frames : 530, No of RX Mirrored Frames : 530, No of TX Mirrored Frames : 0

=====
```

## Viewing a Flow Mirror Flow in Time Blocks

To display all the information recorded for a Flow Mirror flow blocked out using a specific time interval, enter `flow --show flow_name -feature mirror -t num`.

The `-t` parameter applies only to Flow Mirror flows. The `num` value is the number of seconds between samples, and can be 6, 7, 8, 9, or 10. The default value is 7. Using this parameter updates the output on the console at the specified time interval until you press **Ctrl+C**. In time interval output, only frames that were mirrored in the time window between `t` and `t+t` (for example, at 6, 12, 18, 24... seconds) are displayed. The following examples show the frame rates for the Flow Mirror flow named "fmshow" at 6-second intervals. The timestamp is for the frame, not the time block. At the bottom of each interval block, a data line displays the following:

- The total number of mirrored frames sent to the CPU
- The total number of received mirrored frames
- The total number of frames mirrored during the previous interval

**Example of Mirroring the Flow to the CPU (CFM)** This example shows the result of mirroring the flow to the CPU (CFM).

```
switch:admin> flow --show fmtime_cfm -feature mirror -time 6
```

```
=====
Name : fmtime_cfm Features: mir(Activated) noConfig: Off
Definition: EgrPort(15),SrcDev(0x010e00),DstDev(0x010f00),BiDir
```

Flow Mirror (Activated):

```

| OXID | RXID | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp |

| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 09:45:34:100 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 09:45:34:100 | (output
truncated)
0001	ffff	SOFn3	EOFn	Data	----	Tx	Jun 05 09:45:38:110
0044	ffff	SOFn3	EOFn	Data	----	Rx	Jun 05 09:45:38:110
0001	ffff	SOFn3	EOFn	Data	----	Tx	Jun 05 09:45:38:110

```

```
No of Mirrored Frames : 1280, No of RX Mirrored Frames : 640, No of TX Mirrored Frames : 640

=====
```

```
=====
Name : fmtime_cfm Features: mir(Activated) noConfig: Off
Definition: EgrPort(15),SrcDev(0x010e00),DstDev(0x010f00),BiDir
```

Flow Mirror (Activated):

```

| OXID | RXID | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp |

| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 09:45:41:100 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 09:45:41:100 | (output
truncated)
| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 09:45:45:109 |

```



```

| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 09:45:45:109 |

No of Mirrored Frames : 1280, No of RX Mirrored Frames : 640, No of TX Mirrored Frames : 640

=====
Name : fmtime_cfm Features: mir(Activated) noConfig: Off
Definition: EgrPort(15),SrcDev(0x010e00),DstDev(0x010f00),BiDir

Flow Mirror (Activated):

| OXID | RXID | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp |

| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 09:45:47:100 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 09:45:47:100 |
| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 09:45:47:100 | (output
truncated)
| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 09:45:51:109 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 09:45:51:109 |

No of Mirrored Frames : 1280, No of RX Mirrored Frames : 640, No of TX Mirrored Frames : 640

===== (output truncated)

```

**Example of Mirroring the Flow to a Local Port (LFM)** This example shows the result of mirroring the flow to a local port (LFM).

```

switch:admin> flow --show fmtime_lfm -feature mirror -time 6
=====
Name : fmtime_lfm Features: mir(Activated) noConfig: Off
Definition: EgrPort(15),SrcDev(0x010e00),DstDev(0x010f00),MirPort(16)

Flow Mirror (Activated):

No of Mirrored Frames : 481442, No of RX Mirrored Frames : 0, No of TX Mirrored Frames : 481442

=====
Name : fmtime_lfm Features: mir(Activated) noConfig: Off
Definition: EgrPort(15),SrcDev(0x010e00),DstDev(0x010f00),MirPort(16)

Flow Mirror (Activated):

No of Mirrored Frames : 716244, No of RX Mirrored Frames : 0, No of TX Mirrored Frames : 716244

=====
Name : fmtime_lfm Features: mir(Activated) noConfig: Off
Definition: EgrPort(15),SrcDev(0x010e00),DstDev(0x010f00),MirPort(16)

Flow Mirror (Activated):

No of Mirrored Frames : 951046, No of RX Mirrored Frames : 0, No of TX Mirrored Frames : 951046

```

```
===== (output truncated)
```

## Learning in Flow Mirror Flows

Flow Mirror supports learning for both source and destination devices. To specify learning in a Flow Mirror flow, identify the parameter to be learned using an asterisk inside quotation marks ("\*") for the device identifier.

For a flow using learning, if the frame type is specified in the flow definition, the value for `-frametype` must be a fixed value for the flow to work. The following example creates a Flow Mirror flow using the learning capability to mirror traffic from any device to any device that is ingressing through port 1/20.

```
switch:admin> flow --create fmir_learn1-20 -feature mirror -ingrport 1/20 -srcdev "*" -dstdev "*"
```

## Repeating Flow Mirror Output

You can configure the Flow Mirror feature to repeat its flow output. This allows you to view sample frames or output over time to look for differences as the situation evolves. You can specify that a flow output repeats from one to ten times.

To specify the number of times, a flow output should be repeated, perform the following steps:

1. Connect to the switch and log in using an account with admin permissions.
2. Use the `flow --show flow_name -feature feature_name -count num` command.

The `num` value can range from 1 through 10. The default value is 1.

The following example creates a bidirectional Flow Mirror flow named "count\_cfm" that is mirrored to the CPU and repeats the output three times.

```
switch:admin> flow --create count_cfm -feature mir -ingrport 10
 -srcdev 010403 -dstdev 020504 -bidir
```

```
switch:admin> flow --show count_cfm -fea mir -count 3
```

```
=====
Name : count_cfm Features: mir(Activated) noConfig: Off
Definition: IngrPort(14),SrcDev(0x010e00),DstDev(0x010f00),BiDir
```

```
Flow Mirror (Activated):
```

```

| OXID | RXID | SOF | EOFn | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp |

| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 07:54:27:100 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 07:54:27:100 |
```

```
.
.
.
```

```
| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 07:54:31:109 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 07:54:31:109 |

```

```
No of Mirrored Frames : 1280, No of RX Mirrored Frames : 640, No of TX Mirrored Frames : 640

```

```
=====
Name : count_cfm Features: mir(Activated) noConfig: Off
Definition: IngrPort(14),SrcDev(0x010e00),DstDev(0x010f00),BiDir
```

```
Flow Mirror (Activated):
```

```

| OXID | RXID | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp |

| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 07:54:34:100 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 07:54:34:100 |
.
.
.
| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 07:54:38:109 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 07:54:38:109 |

No of Mirrored Frames : 1280, No of RX Mirrored Frames : 640, No of TX Mirrored Frames : 640

=====
Name : count_cfm Features: mir(Activated) noConfig: Off
Definition: IngrPort(14),SrcDev(0x010e00),DstDev(0x010f00),BiDir

Flow Mirror (Activated):

| OXID | RXID | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp |

| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 07:54:40:100 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 07:54:40:100 |
.
.
.
| 0001 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 07:54:44:109 |
| 0044 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 07:54:44:109 |

No of Mirrored Frames : 1280, No of RX Mirrored Frames : 640, No of TX Mirrored Frames : 640

=====

```

The following example creates a Flow Mirror flow named “count16\_lfm” that is mirrored to local port 16, and it repeats the output five times.

```

switch:admin> flow --create count16_lfm -feature mir -ingrport 14
 -srcdev 010403 -dstdev 020504 -mirrorport 16

```

```

switch:admin> flow --show count16_lfm -feature mir -count 5

```

```

=====
Name : count16_lfm Features: mir(Activated) noConfig: Off
Definition: IngrPort(14),SrcDev(0x010e00),DstDev(0x010f00),MirPort(16)

Flow Mirror (Activated):

No of Mirrored Frames : 1032316, No of RX Mirrored Frames : 1032316, No of TX Mirrored Frames : 0

=====

Name : count16_lfm Features: mir(Activated) noConfig: Off
Definition: IngrPort(14),SrcDev(0x010e00),DstDev(0x010f00),MirPort(16)

```

```

Flow Mirror (Activated):

No of Mirrored Frames : 1267119, No of RX Mirrored Frames : 1267119, No of TX Mirrored Frames : 0

=====
Name : count16_lfm Features: mir(Activated) noConfig: Off
Definition: IngrPort (14), SrcDev (0x010e00), DstDev (0x010f00), MirPort (16)

Flow Mirror (Activated):

No of Mirrored Frames : 1501921, No of RX Mirrored Frames : 1501921, No of TX Mirrored Frames : 0

=====
Name : count16_lfm Features: mir(Activated) noConfig: Off
Definition: IngrPort (14), SrcDev (0x010e00), DstDev (0x010f00), MirPort (16)

Flow Mirror (Activated):

No of Mirrored Frames : 1736723, No of RX Mirrored Frames : 1736723, No of TX Mirrored Frames : 0

=====
Name : count16_lfm Features: mir(Activated) noConfig: Off
Definition: IngrPort (14), SrcDev (0x010e00), DstDev (0x010f00), MirPort (16)

Flow Mirror (Activated):

No of Mirrored Frames : 1971525, No of RX Mirrored Frames : 1971525, No of TX Mirrored Frames : 0

=====

```

## Sorting Flow Mirror Output

In Flow Vision, frames can be sorted whether or not subflows are present. Sorting the output allows you to highlight a selected aspect of the flow data.

To sort the flow output, perform the following steps:

1. Connect to the switch and log in using an account with admin permissions.
2. Use the `flow --show flow_name -feature feature_name -sortby columncolumn_num` command. The `columncolumn_name` value is the number of the output column on which the data is to be sorted. There is no space between "column" and the column number.

### NOTE

The `-sortby` parameter can only be applied when there is only one feature (monitor, mirror, or generator) specified in the `flow --show flow_name` command.

The following example creates a Flow Mirror flow "sortMirror," and then shows the output sorted by column 3, the **OXID**.

```
switch:admin> flow --create sortMirror -feature mirror -egrport 15 -srcdev "*" -bidir
```

```
switch:admin> flow --show sortMirror -feature mirror -sortby column3
```

```
=====
```

Name : sortMirror Features: mir(Activated) noConfig: Off  
 Definition: EgrPort(15),SrcDev(\*),BiDir

Flow Mirror (Activated):

```

| SID(*) | DID(*) | OXID | RXID | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp
|

| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:100
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:101
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:102
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:103
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:104
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:105
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:106
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:107
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:108
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:109
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:50:110
|
| 010e00 | 010f00 | 0001 | ffff | SOFn3 | EOFn | Data | ---- | Tx | Jun 05 10:16:51:104
|
| 010f00 | 010e00 | 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 10:16:51:111
|
| 010f00 | 010e00 | 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 10:16:52:104
|
| 010f00 | 010e00 | 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 10:16:52:104
|
| 010f00 | 010e00 | 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 10:16:53:106
|
| 010f00 | 010e00 | 0044 | ffff | SOFn3 | EOFn | Data | ---- | Rx | Jun 05 10:16:54:109
|
.
.
.

```

No of Mirrored Frames : 1280, No of RX Mirrored Frames : 640, No of TX Mirrored Frames : 640

## Deactivating Flow Mirror Flows

Flow Mirror flows can be deactivated without deleting them. This allows you to have a “library” of flows that you can activate as needed without having to create them repeatedly.

To manually deactivate a Flow Mirror flow, enter `flow --deactivate flow_name -feature mirror`. The following example deactivates the Flow Mirror flow named “Flow1”:

```
switch:admin> flow --deactivate Flow1 -feature mirror
```

Flow Vision automatically deactivates a user-defined Flow Mirror flow if any of the following changes to a port defined as part of the mirrored flow occur:

- The port type of an ingress or egress port defined in the flow has been changed to other than an F\_Port or F\_Port trunk. You must correct the port type error and then manually reactivate the flow.
- The slot power has been powered off for ingress or egress ports. Reactivation occurs automatically when the power is restored.

## Customizing Flow Mirror CFM Flow Frame Retention

You can change how frames mirrored to a CPU are retained in the Flow Mirror buffer when it is full.

To have the Flow Mirror buffer overwrite existing frames in the buffer on a first-in-first-out basis when full (replacing the oldest frames with newer ones), enter the `flow --control -feature mirror -enable_wrap` command. To have the Flow Mirror buffer discard any additional mirrored frames once the buffer is full, enter the `flow --control -feature mirror -disable_wrap` command. The `-enable_wrap` and `-disable_wrap` keywords affect only Flow Mirror flows, but they apply to all Flow Mirror flows, so you cannot specify a flow name. By default, `-enable_wrap` is active.

### ATTENTION

All Flow Mirror flows must be inactive to use `-enable_wrap` or `-disable_wrap` as part of a `flow --control` command. If any Flow Mirror flow is active when you run the command, it fails, and an error message is displayed in the interface.

To see the current buffer setting, enter the `flow --show ctrlcfg` command (the buffer status appears in the last line of the following example and is called out).

```
switch:admin> flow --show -ctrlcfg
```

```
SimPort Information
```

| Slot | Port | PID    | PWWN                    | SID Frame Count | DID Frame Count |
|------|------|--------|-------------------------|-----------------|-----------------|
| 1    | 2    | 050200 | 20:02:00:05:1e:e2:8e:00 | 0K              | 19480K          |

```
Addressing mode information
```

```
Port Addressing Mode : index
```

```
Device Addressing Mode: PID
```

```
Flow Generator Information
```

```
Size: 2048
```

```
Pattern: Random (Default)
```

```
Flow mirror Information
```


```
enable_wrap <- current buffer setting
```

## Mirroring Traffic Flowing to Remote Fabrics

Flow Vision allows you to mirror traffic that is flowing to a remote fabric.

To mirror traffic that is flowing to a remote fabric, in the flow definition, you must specify the imported port ID or real PWWN of the device in the other fabric. The following illustration shows how such a mirroring situation might be constructed. In this illustration, if the flow definition is made on Device A (a switch in Fabric 1), then you must use the imported port ID for B' to mirror to a port on device B in Fabric 2. If the flow definition is given in WWN mode, then the real PWWN of Device B can be used instead.

Figure 53: Mirroring Traffic Flowing to a Remote Fabric



The following example shows the proxy identifications and output of such a flow.

```
switchP1:admin> fcrproxydevshow
Proxy WWN Proxy Device Physical State
Created
in Fabric PID Exists PID
in Fabric in Fabric

1 20:00:00:11:0d:07:00:00 02f001 2 010300 Imported
Total devices displayed: 1

switchP2:admin> fcrproxydevshow
Proxy WWN Proxy Device Physical State
Created
in Fabric PID Exists PID
in Fabric in Fabric

2 10:00:8c:7c:ff:25:9b:00 02f001 1 010000 Imported
Total devices displayed: 1

switchP2:admin> flow --create fm_edge1_edge2 -srcdev 010000 -dstdev 02f001
 -feature mirror -ingrport 3/22 -mirrorport 3/43
Mirror feature(s) have been activated.

switch2:admin> flow --show

```

| Flow Name            | Feature | SrcDev |  | DstDev | IngrPt EgrPt BiDir LUN FrameType | SFID DFID MirPt |      |
|----------------------|---------|--------|--|--------|----------------------------------|-----------------|------|
| sys_gen_all_simports | gen     | *      |  | *      | * * no - -                       | no - -          |      |
| fm_edge1_edge2       | mir+    | 010000 |  | 02f001 | 3/22 - no - -                    | no - -          | 3/43 |

+ Denotes feature is currently activated for the flow

```
switchP2:admin> flow --show fm_edge1_edge2 -feature mirror
=====
Name : fm_edge1_edge2 Features: mir(Activated) noConfig: Off
Definition: IngrPort(3/22),SrcDev(0x010000),DstDev(0x02f001),MirPort(3/43)

Flow Mirror (Activated):
=====
No of Mirrored Frames : 1903144, No of RX Mirrored Frames : 1903144, No of TX Mirrored Frames : 0
=====
```

## Using Flow Mirror with Access Gateway

With Fabric OS 9.1.x and later, Access Gateway supports Flow Mirror for Brocade Gen 6 and Gen 7 switches. Flow Mirror is supported for N\_Ports and F\_Ports as ingress or egress ports in the flow definition.

Enable a mirror port on the Access Gateway, and then configure a flow to the mirrored port from an Access Gateway port exactly as you would for any switch. See [Local Flow Mirroring](#) for general steps.

The following are limitations and considerations for configuring Flow Mirror on Access Gateways:

- Local flow mirroring (LFM) and CPU flow mirroring (CFM) flows are supported in AG mode. Remote flow mirroring (RFM) is not supported.
- A port should be in the disabled state before configuring it as a mirror port. The mirror port configuration will be blocked if the port is enabled.
- You do not need to clear the N\_Port to F\_Port mappings when a port is configured as a mirror port.
- The mirror port can be specified by its port or index number because both are identical in AG mode. Display ports and their index numbers using the `switchShow` command.
- You cannot configure internal ports on SAN I/O modules as mirror ports.
- Enabling and disabling AG mode is blocked if flows are configured on the switch.
- During a high-availability failover, all flows will be deleted and then reinitiated after failover. If any Flow Mirror flow is active, then mirroring resumes after the failover.
- If an online port is part of the flow definition, Flow Mirror will be enforced.
- If an offline port is part of the flow definition, Flow Mirror is activated, but not enforced.

For more information on Flow Mirror related to Access Gateway, refer to the *Brocade Fabric OS Access Gateway Administration Guide*.

## Troubleshooting Using Flow Mirror


The following use cases describe how to use Flow Mirror to troubleshoot typical fabric performance problems.



Tracking Latency between a Host and All Connected Targets

To smooth out application performance and apply load balancing, you can track the latency of SCSI initiator-target pairs. The following examples capture all the SCSI commands and their status frames initiated by device H1 ingressing through port F1, as illustrated by the following figure. This will capture all the SCSI commands and status frames that can be viewed using Flow Vision. By viewing the output of the Flow Mirror flow using the `flow --show` command and viewing the mirrored frame data for a CPU mirrored flow, you can deduce the latency between a SCSI command and its respective status frame for a CPU-mirrored flow at the host port.

Figure 54: SCSI Command Tracking for Latency



Tracking Latency Using a CPU-Mirrored Flow

To track the latency using a CPU-mirrored flow, you create the flow and then use the `flow --show` command to view the captured results. In this case, the latency is an approximate latency (best effort latency), because the latency is an approximate latency block completion time. You can use the timestamps on the appropriate frames and deduce the latency from those values. The following example shows a typical command for this purpose; be aware that the output is dependent on the fabric topology.

```
switch:admin> flow --show cfm -verbose
=====
Name : cfm Features: mir(Activated) noConfig: Off
Definition : IngrPort(30),SrcDev(*)

Flow Mirror (Activated):

| Time-Stamp | Dir | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Frame Contents
|

No of Mirrored Frames : 0, No of RX Mirrored Frames : 0, No of TX Mirrored Frames : 0

=====
```

```
switch:admin> flow --show cfm -fea mir -sortby column1
```

```

| SID(*) | DID(*) | OXID | RXID | SOF | EOF | Frame_type | LUN(*)/NSID(*) | Dir | Time-Stamp
|

```

```

No of Mirrored Frames : 0, No of RX Mirrored Frames : 0, No of TX Mirrored Frames : 0

=====
```

### Tracking Latency Using a Local Mirrored Flow

To track the latency using a locally mirrored flow, you must first specify the mirror port, and then create the flow using that mirror port. You can then deduce the latency by viewing the mirrored frames. In this case, the latency is close to the exact latency because the latency I/O Block completion time is close to the actual I/O Block time. The following example shows a typical command for this purpose; the output can be viewed on an analyzer.

```
switch:admin> portcfg mirrorport 9 --enable
```

```
switch:admin> flow --create LocalMirrorflow -feature mirror -ingrport F1 -srcdev H1
 -frametype scsicmdsts -mirrorport 9
```

### Diagnosing Excessive SCSI Reserve and Release Activity

If there are excessive SCSI reserve and release activity in a virtualized environment, you can use Flow Mirror to identify the affected LUNs.

The following example creates a flow to mirror all the SCSI release frames from multiple servers to LUNs on the target on port 1/20. You can then analyze the mirrored frames to determine the impacted LUNs.

```
switch:admin> flow --create flow_scsi -feature mirror -egrport 1/20
 -srcdev "*" -dstdev "*" -frametype scsiresrel
```

```

| SID(*) | DID(*) | OXID | RXID | SOF | EOF | Frame_type | LUN(*) | Dir | Time-Stamp
|


b28600	a2bd00	0f27	ffff	SOFi3	EOFt	SCSI3_Rel	0003	Tx	Jul 16 17:21:47:253
b28000	a2bd00	09de	ffff	SOFi3	EOFt	SCSI3_Rel	0002	Tx	Jul 16 17:21:47:253
b2c680	a2bd00	0afe	ffff	SOFi3	EOFt	SCSI3_Rel	0001	Tx	Jul 16 17:21:47:253
b28600	a2bd00	0f51	ffff	SOFi3	EOFt	SCSI3_Rel	0005	Tx	Jul 16 17:21:47:253
b28000	a2bd00	09f0	ffff	SOFi3	EOFt	SCSI3_Res	0002	Tx	Jul 16 17:21:47:253
b28600	a2bd00	0f1f	ffff	SOFi3	EOFt	SCSI3_Rel	0004	Tx	Jul 16 17:21:47:253
```

```
(output truncated)
```

### Diagnosing a Slow-Draining F\_Port

A slow-draining F\_Port can be challenging to diagnose. Bottleneck detection allows you to identify the slow-draining F\_Port, and Flow Mirroring helps you identify the affected LUN on that F\_Port.

The following example creates a flow to mirror traffic passing in both directions from device 0x020200 to F\_Port 15 on device 0x040F00, and then displays the output. (The **Frame Contents** column has been trimmed in this example.) The following image provides an illustration of what is happening in the example.

**Figure 55: A Flow Mirror Revealing a Slow Drain**

## Flow Mirror and High Availability


On a High Availability (HA) failover, HA reboot, or power-cycle, Flow Mirror will stop mirroring frames until the system recovers; at which point it will resume mirroring. This could be as early as when the `hashow` command indicates that both control processor units (CPUs) are in sync, but it could occur after HA sync, in which case the output of the `switchshow` command would then indicate the correct switch port status. All flow statistics are cleared and reset after a failover recovery. See [High Availability and Flow Vision](#) for more information.

## Understanding the Fabric Flow Dashboard

The following illustration displays a sample fabric that is used as a foundation for the Fabric Flow Dashboard explanation. In this discussion, the `flow` command is executed on Switch 2. In the switch identification format Switch #(X), Switch # is the switch name, and X is the Domain ID. If SAN administrators see a performance issue affecting Host 1, they can then

check the Fabric Flow Dashboard to determine all the targets connected to Host 1. Once the targets are known, then each of them can be checked to determine the problematic target.

**Figure 56: Sample Fabric**



For the figure above, the Fabric Flow Dashboard identifies the flows as follows; notice that there are multiple potential routes for each flow.

- Flow 1: Target 1 to Host 1 (Red flow)
- Flow 2: Target 2 to Host 1 (Green flow)
- Flow 3: Target 3 to Host 1 (Blue flow)

The topology information for the identified flows uses the following format:

```
[SrcDev: DeviceID] <-> (Domain1/inPort, Domain1/outPort) <-> (D2/inP, D2/outP) <-> (D3/inP, D3/outP) ... <->
[DstDev: DeviceID]
```

For Flow 1 in the example above, the paths are:

- [SrcDev:Target 1] <-> (SW1/F1, SW1/E5) <-> (SW2/E7, SW2/F17) <-> [DstDev:Host 1]

For Flow 2 and Flow 3, the paths are:

- srcDev(Target 2, Target 3), dstDev(Host 1)
- [srcDev: \*] <-> (SW4/\*, SW4/E2) <-> (SW1/E1, SW1/E5) <-> (SW2/E7, SW2/17#) <-> [dstDev: Host 1]
- [srcDev: \*] <-> (SW4/\*, SW4/E3) <-> (SW1/E3, SW1/E1#) <-> (SW2/E4, SW2/17#) <-> [dstDev: Host 1]
- [srcDev: \*] <-> (SW4/\*, SW4/E4) <-> (SW1/E4, SW1/E1#) <-> (SW2/E4, SW2/17#) <-> [dstDev: Host 1]

The # symbol indicates that one or more MAPS rule violations are present in a port.

The following extract from a Fabric Flow Dashboard provides a typical example of a situation where multiple hosts/targets are connected to the source and destination switches.

```
srcDev (0x170500), dstDev(0x031100)
[srcDev: 23/1] <-> (23/1#, 23/5) <-> (3/7, 3/17) <-> [dstDev: 3/0x031100]

srcDev (0x190a00), dstDev(0x031100, 0x031200, 0x031300, 0x031400, 0x031500)
srcDev (0x190b00), dstDev(0x031100, 0x031200, 0x031300, 0x031400, 0x031500)
srcDev (0x190c00), dstDev(0x031100, 0x031200, 0x031300, 0x031400, 0x031500)
srcDev (0x190d00), dstDev(0x031100, 0x031200, 0x031300, 0x031400, 0x031500)
srcDev (0x190e00), dstDev(0x031100, 0x031200, 0x031300, 0x031400, 0x031500)
[srcDev: 25/*] <-> (25/*, 25/2) <-> (23/2#, 23/5) <-> (3/7, 3/*) <-> [dstDev: 3/*]
[srcDev: 25/*] <-> (25/*, 25/3) <-> (21/3, 21/1#) <-> (3/4, 3/*) <-> [dstDev: 3/*]
```

```
[srcDev: 25/*] <-> (25/*, 25/4) <-> (21/4, 21/1#) <-> (3/4, 3/*) <-> [dstDev: 3/*]
```

### NOTE

If the flow definition contains an E\_Port as an ingress or egress port, the path will contain information starting from the E\_Port until reaching the destination. For the figure above, if the `flow -show -feature fabinfo` command is run on Switch 1 with the aim to identify all paths from Target 2 to Host 1 using egress port 5, the paths are identified as follows:

```
[srcDev:Target2] <-> (SW1/*, SW1/E5) <-> (SW2/E7, SW2/F17) <-> [dstDev:Host1]
```

## Viewing Fabric Information with the Fabric Flow Dashboard

The Flow Vision Fabric Flow Dashboard provides SAN administrators with the ability to have a consolidated view of a flow across a fabric, allowing them to quickly isolate network issues.

You can view the Fabric Flow Dashboard by specifying `fabinfo` as the feature for the `flow --show` command, for example:

```
flow --show -feature fabinfo -srcdev "*" -egrport 17
```

The output of this command has two parts.

- **Summary Topology:** The first part presents a summary topology of all the paths taken by the flows in the flow definition.
- **Switch-specific Information and MAPS Rule Violations:** The second part presents information and MAPS rule violations for each of the switches the flow passes through. While reporting the MAPS data, preference is given to the data for ports that are in the flow path. If there are no violations on any ports in the flow path, only the category and violation count is displayed without additional details. Additional information about MAPS history data and about fenced, decommissioned, and quarantined ports is displayed only for switches running Fabric OS 9.1.x.

The following example shows the results of running the command described above. The column headings in the example have been edited slightly to allow them to display clearly.

```
switch:admin> flow --show -feature fabinfo -srcdev "*" -egrport 17
Flow Dashboard Information:
=====
Topology Data:

srcDev(0x170500), dstDev(0x31100)
srcDev(0x170500) <-> (23/1#, 23/5) <-> (3/7, 3/17) <-> dstDev(0x31100)

srcDev(0x190a00), dstDev(0x31100)
srcDev(0x190a00) <-> (25/10, 25/2) <-> (23/2#, 23/5) <-> (3/7, 3/17#) <-> dstDev(0x31100)
srcDev(0x190a00) <-> (25/10, 25/3) <-> (21/3, 21/1#) <-> (3/4, 3/17#) <-> dstDev(0x31100)
srcDev(0x190a00) <-> (25/10, 25/4) <-> (21/4, 21/1#) <-> (3/4, 3/17#) <-> dstDev(0x31100)

"Indicates there are MAPS violation on these ports"

MAPS Violations on source and destination ports:
Category(RuleCnt)|RptCnt|Rule Name |Execution Time |Object|Triggered
Value | | |Value |
-----|-----|-----|-----|-----|-----|-----|
(Source port: 23/1)
Port Health(176) |3 | defALL_TARGET_PORTSCRC_1|12/18/16 06:03:45|Port1 | 2 |
(Destination port: 3/17)
Port Health(24) |6 | defALL_HOST_PORTSCRC_2 |12/18/16 09:13:24|Port17| 5 |
Port Health(15) |2 | defNON_E_F_PORTS_LF_0 |12/18/16 09:18:24|Port17| 3 |
```

## Switch Specific Data:

```

=====
SwitchDomain: 23 (0x17)
Name: sw0 Model: 180 Uptime: (2 days 2 hrs 37 mins)
FirmwareVersion : v9.1.0
OperationalStatus: MARGINAL Reason : BAD_PWR
RebootReason: Reboot
Fenced Ports : 24
Decommissioned Ports : None
Quarantined Ports : None
Category(RuleCnt) | RptCnt | Rule Name | Execution Time | Object | Triggered |
 | | | | | |

Port Health(176) | 3 | defALL_TARGET_PORTSCRC_1 | 12/18/16 06:03:45 | Port1 | 2 |
 | | defNON_E_F_PORTSLF_0 | 12/17/16 09:13:24 | Port5 | 6 |

```

```

=====
SwitchDomain: 3 (0x3)
Name: switch 3 Model: 121.3 Uptime: (41 days 0 hrs 02 mins)
FirmwareVersion : v9.0.1
OperationalStatus: Healthy
RebootReason: Reboot
Fenced Ports : None
Decommissioned Ports : None
Quarantined Ports : None
Category(RuleCnt) | RptCnt | Rule Name | Execution Time | Object | Triggered |
 | | | | | |

Port Health(24) | 6 | defALL_HOST_PORTSCRC_2 | 12/17/16 09:13:24 | Port17 | 5 |
Port Health(15) | 2 | defNON_E_F_PORTSLF_0 | 12/17/16 09:18:24 | Port17 | 3 |

```

```

=====
SwitchDomain: 25 (0x19)
Name: switch 25 Model: 121.3 Uptime: (40 days 21 hrs 12 mins)
FirmwareVersion : v9.0.1
OperationalStatus: CRITICAL Reason : FAULTY_BLADE
RebootReason: Reboot
Fenced Ports : None
Decommissioned Ports : None
Quarantined Ports : None
Category(RuleCnt) | RptCnt | Rule Name | Execution Time | Object | Triggered |
 | | | | | |

Port Health(16) | No violations on the ports in the flow path

```

```

=====
SwitchDomain: 21 (0x15)
Name: switch 21 Model: 121.3 Uptime: (40 days 03 hrs 10 mins)
FirmwareVersion : v9.0.0
OperationalStatus: Healthy

```

```

RebootReason: Reboot
Fenced Ports : None
Decommissioned Ports : None
Quarantined Ports : None

```

```

No MAPS violations
=====

```

## Viewing Additional MAPS History Data

Using Fabric OS 9.1.x, administrators can specify a domain name as an option of the `fabinfo` feature, for example:

```

flow --show -feature fabinfo -srcdev 0x011100
 -egrport 17 -verbose -domain 3

```

When the domain name is specified, only data corresponding to the domain is displayed. Data for other domains is not displayed. The data displayed includes MAPS history data. Showing statistics for only those ports in the path of the flow definition, MAPS history data also provides statistics about events that have not crossed the MAPS violation threshold but still have non-zero values. Also, the data covers multiple durations.

### Example of Displaying MAPS History Data

The following example shows the results of running the command described above. Be aware that the column headings in the example have been edited slightly to allow it to display clearly.

```

switch:admin> flow --show -feature fabinfo -srcdev 0x011100
 -egrport 17 -verbose -domain 3

Flow Dashboard Information:
=====

Switch Specific Data:

=====

SwitchDomain: 3 (0x3)
Name: sw0 Model: 180 Uptime: (41 days 0 hrs 02 mins)
FirmwareVersion : v9.1.x
OperationalStatus: Healthy
RebootReason: Reboot
Fenced Ports : None
Decommissioned Ports : None
Quarantined Ports : None

MAPS violation:

Port Health(24) |6 | defALL_HOST_PORTSCRC_2|12/17/16 09:13:24|Port17| 5 |
Port Health(15) |2 | defNON_E_F_PORTSLF_0 |12/17/16 09:18:24|Port17| 3 |

MAPS history data:

Stats Current 12/16/20 12/15/20 12/14/20 12/13/20 12/12/20 12/11/20
(Units) Port(val) Port(val) Port(val) Port(val) Port(val) Port(val) Port(val)

```

|     |         |        |           |          |        |       |         |
|-----|---------|--------|-----------|----------|--------|-------|---------|
| CRC | 32 (3)  | 32 (2) | 46 (34)   | 45 (7)   | 32 (1) | -     | 32 (7)  |
|     | 42 (2)  | 42 (2) | 43 (33)   | 44 (6)   | 33 (1) | -     | 33 (7)  |
|     | 12 (1)  | 12 (1) | 42 (31)   | 42 (5)   | 34 (1) | -     | 34 (7)  |
|     | 06 (1)  | 06 (1) | 47 (20)   | 46 (5)   | 35 (1) | -     | 35 (7)  |
| ITW | 32 (11) | 32 (8) | 45 (2300) | 45 (539) | 32 (3) | 0 (1) | 32 (21) |
|     | 33 (3)  | 33 (3) | 40 (2282) | 40 (522) | 33 (1) | -     | 33 (7)  |
|     | 34 (3)  | 34 (3) | 44 (2276) | 41 (519) | 34 (1) | -     | 34 (7)  |
|     | 35 (3)  | 35 (3) | 41 (2269) | 44 (517) | 35 (1) | -     | 35 (7)  |

## Limitations and Restrictions on Fabric Flow Dashboard Flows

The following limitations and restrictions apply to a Fabric Flow Dashboard flow:

- This feature is not available on switches running in Access Gateway mode.
- The Fabric Flow Dashboard flow definition can use only F\_Ports.
- Only controller trunk ports are supported. Agent ports are not supported.
- The maximum number of flows that can be discovered is the maximum number of subflows supported for a flow. This count will not be affected by the maximum number of subflows supported across the chassis.
- Only the following MAPS categories are considered for the MAPS violations shown in the Fabric Flow Dashboard report:
  - Port Health
  - Traffic Performance
  - Fabric Performance Impact
- Target Driven zones are not considered.
- The frame options `-lun` and `-frametype` are not supported for the flow definition.
- The `flowname` and `-sortBy` options are not supported.
- In a fabric that includes versions of Fabric OS earlier than 8.2.x, the Fabric Flow Dashboard functionality is limited. While topology and switch data will be shown for all the switches, MAPS violation details are not available for those switches running earlier versions of Fabric OS.

## Fabric Flow Dashboard Use Cases

The following section provides a use case for the Fabric Flow Dashboard.

### Use Case: Slow-Drain Device

If you suspect that there is a slow-draining device in your fabric as evidenced by a significant number of Class 3 timeouts (C3TXTO) on E\_Port 5, you can quickly set up a flow to check all the flows passing through this E\_Port and view the data in the flow dashboard. In this case you can specify the flow as follows:

```
flow -show -feature fabinfo -srcdev "*" dstdev "*" -egrport 5
```

In this command example, you have specified that you want to view all flows by using "\*" for both the source and destination devices, and you have identified port 5 as the egress port you are interested in.

If you see "Fabric Performance Impact" violations on any target port(s) in the output, this suggests that the E\_Port timeouts are probably because of a slow-draining device attached to the destination F\_Port. You can then look more closely at the device(s) attached to that F\_Port.



## Managing Virtual Fabrics

---

Virtual Fabrics is an architecture to virtualize hardware boundaries. Traditionally, SAN design and management is done at the granularity of a physical switch. Virtual Fabrics allows SAN design and management to be done at the granularity of a port.

Virtual Fabrics is a suite of related features that can be customized based on your needs. The Virtual Fabrics suite consists of the following specific features:

- Logical switch
- Logical fabric
- Device sharing

This chapter describes the logical switch and logical fabric features. For information about device sharing with Virtual Fabrics, see [FC-FC routing and Virtual Fabrics](#).

### NOTE

A note on terminology: *Virtual Fabrics* is the name of the suite of features. A *logical fabric* is a type of fabric that you can create using the Virtual Fabrics suite of features.

### NOTE

SNMPv3 is required to manage Virtual Fabrics. If you use SNMPGET to poll the switches, for example, HiTrack, the SNMPv3 user must be configured based on the logical switch that is being monitored. If SNMP monitoring is for the default switch, you do not need the SNMP user to be configured as a user on the switch. When SNMP monitoring is VF-specific, you must create a user with permission for the corresponding VF and must add the created user into the SNMP database as an SNMPv3 user. When the SNMPv3 user is successfully added, you can use the same user name in management applications for SNMP queries to the corresponding logical switch.

## Logical Switch Overview

Traditionally, each switch and all the ports in the switch act as a single Fibre Channel (FC) switch that participates in a single fabric. The logical switch feature allows you to divide a physical chassis into multiple fabric elements. Each of these fabric elements is referred to as a *logical switch*. Each logical switch functions as an independent self-contained FC switch.

### NOTE

Each chassis can have multiple logical switches.

## Default Logical Switch


To use the Virtual Fabrics features, you must first enable Virtual Fabrics on the switch, which creates a single logical switch in the physical chassis. This logical switch is called the *default logical switch*.

The default logical switch initially contains all the ports in the physical chassis.


The following figure shows a switch before and after enabling Virtual Fabrics. In this example, the switch has 10 ports, labeled P0 through P9.

**Figure 57: Switch before and after Enabling Virtual Fabrics**

Before enabling Virtual Fabrics



After enabling Virtual Fabrics



After you enable Virtual Fabrics, you can create up to 7 or 15 additional logical switches, depending on the switch type.

**NOTE**


The X7 Directors running Fabric OS 9.0.0 or later and X6 Directors running Fabric OS 8.1.0 or later support up to 16 logical switches including the default logical switch. All the other platforms including fixed-port switches support up to eight logical switches.

The following figure shows a Virtual Fabrics-enabled switch before and after it is divided into logical switches. Before you create logical switches, the chassis appears as a single switch (default logical switch). After you create logical switches, the chassis appears as multiple independent logical switches. All the ports continue to belong to the default logical switch until you explicitly move them to other logical switches.

The default logical switch always exists. You can add and delete other logical switches, but you cannot delete the default logical switch unless you disable Virtual Fabrics.

**NOTE**

A FICON logical switch or an existing logical switch that has been modified to become a FICON logical switch cannot serve as the default logical switch.

**Figure 58: Switch before and after Creating Logical Switches**

## Logical Switches and Fabric IDs


When you create a logical switch, you must assign it a fabric ID (FID). The fabric ID uniquely identifies each logical switch within a chassis and indicates to which fabric the logical switch belongs. You cannot define multiple logical switches with the same fabric ID within the chassis.

In the following figure, logical switches 2, 3, 4, and 5 are assigned FIDs of 1, 15, 8, and 20, respectively. These logical switches belong to different fabrics, even though they are in the same physical chassis. For example, you could not assign logical switch 5 a fabric ID of 15, because logical switch 3 is already assigned FID 15 in the chassis.

The default logical switch is initially assigned FID 128. You can change this value later.

### NOTE

Each logical switch is assigned one and only one FID. The FID identifies the logical fabric to which the logical switch belongs.


**Figure 59: Fabric IDs Assigned to Logical Switches**

## Port Assignment in Logical Switches

Initially, all ports belong to the default logical switch. When you create additional logical switches, they are empty, and you must assign ports to those logical switches.

As you assign ports to a logical switch, the ports are moved from the default logical switch to the newly created logical switch. A given port can be in only one logical switch.

In the following figure, the default logical switch initially has 10 ports, labeled P0 through P9. After logical switches are created, the ports are assigned to specific logical switches. Note that ports 0, 1, 7, and 8 have not been assigned to a logical switch and so remain assigned to the default logical switch.

**Figure 60: Assigning Ports to Logical Switches**

A given port is always in one (and only one) logical switch. The following scenarios refer to the chassis after port assignment shown in above figure:

- If you assign P2 to logical switch 2, you cannot assign P2 to any other logical switch.
- If you want to remove a port from a logical switch, you cannot delete it from the logical switch, but must move it to a different logical switch. For example, if you want to remove P4 from logical switch 3, you must assign it to a different logical switch: logical switch 2, logical switch 4, or logical switch 1 (the default logical switch).
- If you assign a port to a logical switch, it is removed automatically from the logical switch it is currently in. If you assign P3 to Logical switch 3, P3 is automatically removed from logical switch 2.
- If you do not assign a port to any logical switch, it remains in the default logical switch, as is the case with ports 0, 1, 7, and 8.

See [Adding and Moving Ports on a Logical Switch](#) for instructions for assigning and moving ports on logical switches.

A logical switch can have as many ports as are available in the chassis. In the above figure, the chassis has 10 ports. You could assign all 10 ports to a single logical switch, such as logical switch 2; if you did this, however, no ports would be available for logical switches 3 and 4.


You can move only F\_Ports and E\_Ports from one logical switch to another. If you want to configure a different type of port, such as a VE\_Port or EX\_Port, you must configure them after you move them. Some types of ports cannot be moved from the default logical switch.

## Logical Switches and Connected Devices


You can connect devices to logical switches, as shown in below figure. In logical switch 2, P2 is an F\_Port that is connected to H1. In logical switch 3, P4 is an F\_Port that is connected to D1. H1 and D1 cannot communicate with each other because they are in different fabrics, even though they are both connected to the same physical chassis.

You can also connect other switches to logical switches. In figure below, P6 is an E\_Port that forms an inter-switch link (ISL) between logical switch 4 and the non-Virtual Fabrics switch. Logical switch 4 is the only logical switch that can communicate with the non-Virtual Fabrics switch and D2, because the other logical switches are in different fabrics.

**Figure 61: Logical Switches Connected to Devices and non-Virtual Fabrics switch**



The figure below shows a logical representation of the physical chassis and devices. As shown in below figure, the devices are isolated into separate fabrics.

**Figure 62: Logical Switches in a Single Chassis belong to Separate Fabrics**

## Management Model for Logical Switches

The operations you can perform on a logical switch depend on the context you are in. Some operations affect only a single logical switch, and some operations affect the entire physical chassis.

All user operations are classified into one of the following:

- Chassis management operations  
These are operations that span logical switch boundaries, such as:
  - Logical switch configuration (creating, deleting, or modifying logical switches)
  - Account management (determining which accounts can access which logical switches)
  - Field-replaceable unit (FRU) management (slot commands, such as `slotShow` )
  - Firmware management (firmware upgrade, HA failover)
- Logical switch operations  
These are operations that are limited to the logical switch, such as displaying or changing port states. Logical switch operations include all operations that are not covered in the chassis management operations.

When you log in, you are assigned an active context, or active logical switch. This context filters the view that you get, and determines which ports you can see. You can change the active context. For example, if you are working with logical switch 1, you can change the context to logical switch 5. When you change the context to logical switch 5, you only see the ports that are assigned to that logical switch. You do not see any of the other ports in the chassis.

The scope of logical switch operations is defined by the active context. When you are in the context of a logical switch, you can perform port, switch, and fabric-level operations, subject to Role-Based Access Control (RBAC) rules.

If you have permission to execute chassis-level commands, you can do so, regardless of which logical switch context you are in.

## Logical Fabric Overview

A *logical fabric* is a fabric that contains at least one logical switch.

The four fabrics shown in the following figures are logical fabrics because they each have at least one logical switch.

You can connect logical switches to non-Virtual Fabrics switches and to other logical switches. You can connect logical switches to non-Virtual Fabrics switches using an ISL.

You connect logical switches to other logical switches in two ways:

- Using ISLs
- Using base switches and extended ISLs (XISLs)

## Logical Fabric and ISLs

The following figure shows two physical chassis divided into logical switches. In the following figure, ISLs are used to connect the logical switches with FID 1 and the logical switches with FID 15. The logical switches with FID 8 are each connected to a non-Virtual Fabrics switch. The two logical switches and the non-Virtual Fabrics switch are all in the same fabric, with FID 8.

**Figure 63: Logical Switches Connected to other Logical Switches through Physical ISLs**







**Figure 64: Logical Switches Connected to form Logical Fabrics**

The ISLs between the logical switches are *dedicated ISLs* because they carry traffic only for a single logical fabric. In the above figure, Fabric 128 has two switches (the default logical switches), but they cannot communicate with each other because they have no ISLs between them and they cannot use the ISLs between the other logical fabrics.

**NOTE**

Only logical switches with the same FID can form a fabric. If you connect two logical switches with different FIDs, the link between the switches segments.

## Base Switch and Extended ISLs

One method to connect logical switches is to use extended ISLs and base switches.

When you divide a chassis into logical switches, you can designate one of the switches to be a base switch. A *base switch* is a special logical switch that is used for interconnecting the physical chassis.

A base switch has the following properties:

- ISLs connected through the base switch can be used for communication among the other logical switches.
- Base switches do not support direct device connectivity. A base switch can have only E\_Ports, VE\_Ports, or EX\_Ports but no F\_Ports.
- The base switch provides a common address space for communication between different logical fabrics.
- A base switch can be configured for the preferred domain ID just like a non-Virtual Fabrics switch.
- You can have only one base switch in a physical chassis.


A base switch can be connected to other base switches through a special ISL, called a *shared ISL* or *extended ISL* (XISL). An extended ISL connects base switches. The XISL shares traffic among different logical fabrics.

**NOTE**

A FICON logical switch or an existing logical switch that has been modified to become a FICON logical switch cannot serve as the base switch.


Fabric formation across an XISL is based on the FIDs of the logical switches.

The following figure shows two physical chassis divided into logical switches. Each chassis has one base switch. An ISL connects the two base switches. The ISL is an extended ISL (XISL) because it connects base switches.

**Figure 65: Base Switches Connected by an XISL**

Traffic between the logical switches can now flow across this XISL. The traffic can flow only between logical switches with the same fabric ID. For example, traffic can flow between logical switch 2 in chassis 1 and logical switch 6 in chassis 2, because they both have FID 1. Traffic cannot flow between logical switch 2 and logical switch 7, because they have different fabric IDs (and are, therefore, in different fabrics).

It is useful to think of the logical switches as being connected with logical ISLs, as shown in the following figure. The logical ISLs are not connected to ports, because they are not physical cables. They are a logical representation of the switch connections that are allowed by the XISL.


**Figure 66: Logical ISLs Connecting Logical Switches**

To use the XISL, the logical switches must be configured to allow XISL use. By default, they are configured to do so; you can change this setting, however, using the procedure described in [Configuring a Logical Switch for XISL Use](#).

#### NOTE

It is a good practice to configure at least two XISLs for redundancy.

You can also connect logical switches using a combination of ISLs and XISLs, as shown in the following figure. In this diagram, traffic between the logical switches in FID 1 can travel over either the ISL or the XISL. Traffic between the other logical switches travels only over the XISL.

**Figure 67: Logical Fabric using ISLs and XISLs**

By default, the physical ISL path is favored over the logical path (over the XISL) because the physical path has a lower cost. This behavior can be changed by configuring the cost of the dedicated physical ISL to match the cost of the logical ISL.

#### ATTENTION

If you disable a base switch, all of the logical ISLs are disconnected and the logical switches cannot communicate with each other unless they are connected by a physical ISL.

## Base Fabric

Base switch ports on different chassis can be connected together to form a fabric, called a *base fabric*. Similar to other logical switches, the base switches must have the same FID to be connected. If the base switches have different FIDs, the link between the switches is disabled.

The base fabric follows normal routing policies. As long as physical connectivity is available, the base fabric maintains connectivity for the logical fabrics.

## Logical Ports

Logical ISLs are formed to connect logical switches. A *logical port* represents the ports at each end of a logical ISL. A logical port is a software construct only and does not correspond to any physical port.

Most port commands are not supported on logical ports. For example, you cannot change the state or configuration of a logical port. However, state change on a logical link is permitted using the `lfcfg --lislEnable` command.

The World Wide Name (WWN) for logical ports is in NAA=5 format, using the following syntax:

`5n:nn:nn:nz:zz:zz:zx:xx`

The NAA=5 syntax uses the following variables:

- `nnnnnn` is the Brocade Organizationally Unique Identifier (OUI).
- `zzzzzz` is the logical fabric serial number.
- `xxx` is the logical port number, in the range 0 through FFF.

## Logical Fabric Formation

Fabric formation is not based on connectivity, but on the FIDs of the logical switches. The basic order of fabric formation is as follows:

1. Base fabric forms.
2. Logical fabrics form when the base fabric is stable.
3. Devices begin recognizing one another.
4. Traffic is initiated between the logical switches.

## Account Management and Virtual Fabrics

When user accounts are created, they are assigned a list of logical fabrics to which they can log in and a home logical fabric (home FID). When you connect to a physical chassis, the home FID defines the logical switch to which you are logged in by default. You can change to a different logical switch context, as described in [Changing the Context to a Different Logical Fabric](#).

When you are logged in to a logical switch, the system prompt changes to display the FID of that switch. The following are example prompts for when you are logged in to the default logical switch (FID = 128) and a user-defined logical switch (FID = 15):

```
switch:FID128:admin>
switch:FID15:admin>
```

Refer to [Managing User Accounts](#) for information about creating user accounts and assigning FIDs to user accounts.

## Setting Up IP Addresses for a Logical Switch

Each physical chassis has one common IP address that is shared by all of the logical switches in the chassis. You can set up individual IPv4 addresses for each logical switch.

IPv4 addresses assigned to individual logical switches are assigned to IP over Fibre Channel (IPFC) network interfaces. In Virtual Fabrics environments, a single chassis can be assigned to multiple fabrics, each of which is logically distinct and separate from one another. Each IPFC point of connection to a given chassis needs a separate IPv4 address and prefix to be accessible to a management host.

For a management host to access logical switches, the host bus adapter (HBA) must be able to connect with the common, shared IP address and the individual IPv4 addresses configured for each logical switch.

### NOTE

IPv6 is not supported when setting the IPFC interface for Virtual Fabrics.

IPFC addresses are not handled by configupload or configdownload. The IPFC address of the default logical switch or a non-VF switch is stored on the WWN card or compact flash. This address does not display in a configshow. Non-default

logical switch IPFC addresses display in a `configshow`. The `ipaddrshow` command displays all switch addresses and IPFC addresses configured in the chassis.

Use the following procedure to set up IP addresses for a logical switch:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `ipAddrSet -ls` command.
  - To add an IPv4 address, use the `--add` parameter. Specify the network information in dotted-decimal notation for the Ethernet IPv4 address with a Classless Inter-Domain Routing (CIDR) prefix.
  - To delete an IPv4 address, use the `--delete` parameter.
3. Enter the `ipaddrshow` command to verify the result.

The following example sets IP addresses with the CIDR prefix for logical switches with FID 1, 2, and 128 (default logical switch).

```
switch:FID128:admin> ipaddrset -ls 1 --add 192.0.2.11/24
IP address is being changed...Done.
switch:FID128:admin> ipaddrset -ls 2 --add 192.0.2.22/24
IP address is being changed...Done.
switch:FID128:admin> ipaddrset -ls 128 --add 192.0.2.0/24
IP address is being changed...Done.

switch:FID128:admin> ipaddrshow

SWITCH
Ethernet IP Address: 198.51.100.0
Ethernet Subnetmask: 255.255.255.0
Gateway IP Address: 198.51.100.1
DHCP: Off
IPFC address for virtual fabric ID 128: 192.0.2.0/124
IPFC address for virtual fabric ID 1: 192.0.2.11/24
IPFC address for virtual fabric ID 2: 192.0.2.22/24
```

The following example deletes the IP address for the logical switch with FID 1.

```
switch:FID128:admin> ipaddrset -ls 1 --delete
```

## Logical Switch Login Context

When you log in to a logical switch with Telnet or SSH, the login shell uses the destination IPFC address to look up and set the corresponding logical switch context. When you log in from the serial console, the login shell always sets your default home logical switch context. The IPFC assigned to a logical switch is managed by the `ipAddrSet` command. When you use Telnet or SSH from a remote host, the IPFC address must share the same network (netmask) as the management interface. The values of FCIP addresses and subnet masks are stored in the WWN card. For example, if the management interface IP address is 10.1.18.183 with netmask 255.255.240.0, you can create the IPFC address using the following command.

```
ipaddrset -ls 8 --add 10.1.18.183/20
```

### NOTE

The logical switch login context supersedes the default FID context of a user. The CIDR mask of the IPFC address must match the CIDR of the assigned management IP address to avoid routing issues.

To delete the IPFC address, you can use the following command.

```
ipaddrset -ls 8 --delete
```

If you log in to a switch with the management interface IP address, the home virtual fabric is set for the context. You can log in to the logical switch context by using the telnet or SSH service with the configured IPFC address. You can manage the permissions using the `userConfig` command.

```
switch> userconfig --show admin
Account name: admin
Description: Administrator
Enabled: Yes
Password Last Change Date: Unknown (UTC)
Password Expiration Date: Not Applicable (UTC)
Locked: No
Home LF Role: admin
Role-LF List: admin: 1-128
Chassis Role: admin
Home LF: 128
Day Time Access: N/A
```

**Role-LF** is the list of logical switch contexts for which you have permission to log in over the IPFC address. **Home LF Role** is the default logical switch context when you have no permission to log in to a particular logical switch context or over management interface. The following output shows the result of the `ipAddrShow` command. Only one IPFC is mapped to FID 8.

```
switch> ipaddrshow
SWITCH
Ethernet IP Address: 10.1.18.183
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.1.18.1
DHCP: Off
IPFC address for virtual fabric ID 8: 10.1.18.183/20
```

The following conditions apply for default users:

- For admin users, the login is successful to the given logical switch with the corresponding logical switch context set (FID 1 to 128). When you telnet to IPFC 10.1.18.183, the context is set to FID 8.

#### NOTE

If the home VF of a user is not present on the switch, then the user will be logged in to default VF if the default VF is part of the user's VFs list. If not, the user will be logged in to the lowest VF that is available on the switch that is part of the user's VFs list.

The following conditions apply for non-default users:

- For non-default users configured in the switch with or without the logical switch permission, if the user logs in with the IPFC address without the permission for the corresponding logical switch, only the default home virtual fabric of the user is set for the context.
- For non-default users configured in the switch with or without the logical switch permission, if the user logs in with the IPFC address with the permission for the corresponding logical switch, login is successful to the given logical switch with the corresponding logical switch context set.

## Supported Platforms for Virtual Fabrics

The following platforms support Virtual Fabrics:

- Brocade G730
- Brocade G720
- Brocade G630
- Brocade G620
- Brocade G610 has limited support for two logical switches with no base switch
- Brocade X7 Directors
- Brocade X6 Directors

Some restrictions apply to the ports, depending on the port type and blade type. The following sections explain these restrictions.

## Supported Port Configurations in Fixed-Port Switches

No restrictions exist on the ports in Brocade G730, G720, G630, or G620 switches.

The following rules apply:

- Any port can belong to any logical switch (including the base switch and default logical switch), with the exception that F\_Ports cannot belong to the base switch.
- The default logical switch can use XISLs.
- The default logical switch can also be a base switch except on X7 Directors and X6 Directors.

## Supported Port Configurations in Brocade Directors

Some of the ports in the Brocade X7 and X6 Directors are not supported on all types of logical switches. The following table lists the blades and ports that are supported on each type of logical switch.

**Table 128: Blade and Port Types Supported on Logical Switches**

| Blade Type          | Default Logical Switch | User-Defined Logical Switch | Base Switch |
|---------------------|------------------------|-----------------------------|-------------|
| FC64-48: FC Port    | Yes (F, E)             | Yes (F, E)                  | Yes (E, EX) |
| FC32-X7-48: FC Port | Yes (F, E)             | Yes (F, E)                  | Yes (E, EX) |
| FC32-64: FC Port    | Yes (F, E)             | Yes (F, E)                  | Yes (E, EX) |
| FC32-64: Eth Port   | Yes                    | Yes                         | Yes         |
| FC32-48             | Yes (F, E)             | Yes (F, E)                  | Yes (E, EX) |
| SX6: FC ports       | Yes (F, E)             | Yes (F, E)                  | Yes (E, EX) |
| GE ports            | Yes (VE)               | Yes (VE)                    | Yes (VE)    |
| ICL ports           | Yes (E)                | Yes (E)                     | Yes (E, EX) |

## Virtual Fabrics Interaction with Other Fabric OS Features

The following table lists some Fabric OS features and considerations that apply when using Virtual Fabrics.



**Table 129: Virtual Fabrics Interaction with Fabric OS Features**

| Fabric OS Feature                 | Virtual Fabrics Interaction                                                                                                                                                                                                                                                                                              |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Access Gateway                    | If the AG mode is enabled, Virtual Fabrics is not supported on a switch.                                                                                                                                                                                                                                                 |
| Configuration upload and download | Virtual Fabrics uses a configuration file that is different from the configuration file that is used to download system configuration parameters. See <a href="#">Maintaining the Switch Configuration File</a> for more information about how Virtual Fabrics affects the configuration file.                           |
| FC-FC routing service             | All EX_Ports must reside in a base switch.<br>You cannot attach EX_Ports to a logical switch that has XISL use enabled. You must use ISLs to connect the logical switches in an edge fabric.<br>See <a href="#">Using FC-FC Routing to Connect Fabrics</a> for more information about Virtual Fabrics and FC-FC routing. |
| FICON                             | Up to two logical switches per chassis can run FICON Management Server (CUP), but the FICON logical switch can use both ISLs and XISLs. For more information on FICON logical switches, see <a href="#">Administering FICON Fabrics</a> .                                                                                |
| ISL R_RDY mode                    | ISL R_RDY mode is supported in both the base switch and the logical switch.                                                                                                                                                                                                                                              |
| Licensing                         | Licenses are applicable for all logical switches in a chassis.                                                                                                                                                                                                                                                           |
| Network Time Protocol (NTP)       | The clock information for a chassis is maintained by the default logical switch and not from the principal or primary FCS switch. See <a href="#">Network Time Protocol</a> for more information about how Virtual Fabrics affects NTP.                                                                                  |
| QoS                               | QoS VCs are maintained across the base fabric. See <a href="#">QoS</a> for more information about using the Adaptive Networking features with Virtual Fabrics.                                                                                                                                                           |

## Limitations and Restrictions of Virtual Fabrics

Before you use the Virtual Fabrics feature, you should be aware of the restrictions and limitations regarding QSFP ports and the maximum number of logical switches per chassis.

In the core blades, the four ports belonging to the same QSFP can be assigned to different logical switches when using Gen 6 QSFPs. However, if one of the four ports belongs to the base switch, the remaining three ports cannot be assigned to any other logical switch.

Gen 7 QSFPs does not support being split across logical switches and all must be part of the same logical switch.

The maximum number of logical switches per chassis varies depending on the switch type. The following table lists the supported platforms and the maximum number of logical switches (including the default logical switch) supported on each platform.

**Table 130: Maximum Number of Logical Switches per Chassis**

| Platform             | Maximum Number of Logical Switches |
|----------------------|------------------------------------|
| Brocade X7 Directors | 16                                 |
| Brocade X6 Directors | 16                                 |

| Platform     | Maximum Number of Logical Switches |
|--------------|------------------------------------|
| Brocade G730 | 4                                  |
| Brocade G720 | 4                                  |
| Brocade G630 | 4                                  |
| Brocade G620 | 4                                  |
| Brocade G610 | 2                                  |

See [Supported Port Configurations in Brocade Directors](#) for restrictions on the default logical switch.

If a blade slot is being decommissioned and has ports configured in logical switches, the logical port assignments should be removed from that blade before removing the blade. This action ensures a seamless transition for any new port or AP blade that might occupy that slot in the future and does not apply if you are simply replacing a blade of the same type.

## Restrictions on Using XISLs

The `Allow XISL Use` option of the `configure` command, allows a logical switch to use XISLs in the base switch as well as any standard ISLs that are connected to that logical switch. To allow or disallow XISL use for a logical switch, refer to [Configuring a Logical Switch for XISL Use](#).

The following restrictions apply when using XISLs. XISL use is not permitted in any of the following scenarios:

- The logical switch has ICL ports.
- The logical switch is the default logical switch in the Brocade X7 Director or X6 Director devices.
- The logical switch is a base switch.
- The logical switch is an edge switch for an FC router.

In this case, if the logical switch is enabled, you cannot allow XISL use. If the logical switch is disabled or has not yet joined the edge fabric, you *can* allow XISL use; however, fabric segmentation occurs when the logical switch is enabled or connected to an edge fabric.

### NOTE

Using XISL and `fmsmode` at the same time is permitted, but this combination will only work in a one-hop topology.

## Restrictions on Moving Ports

The following are restrictions on moving ports among logical switches:

- FC ports cannot be moved if any one of the following features is enabled:
  - Long distance
  - QoS
  - F\_Port buffers
  - F\_Port trunking
  - E\_Port credits
  - SIM\_Port
- Before moving VE\_Ports, you must remove the VE\_Port tunnel configuration.
- VE\_Ports on the SX6 blade can be moved to any logical switch independent of the location of the physical GE port.

## Enabling Virtual Fabrics Mode

A fabric is said to be in *Virtual Fabrics mode* (VF mode) when the Virtual Fabrics feature is enabled. Before you can use the Virtual Fabrics features, such as logical switch and logical fabric, you must enable the VF mode.

VF mode is enabled by default.

**NOTE**

When you enable VF mode, the control processors (CPs) are rebooted and all EX\_Ports are disabled after the reboot.

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use the `fosconfig --show` command to verify whether VF mode is enabled:

```
fosconfig --show
```

3. Use the `fosconfig --enable vf` command to enable VF mode:

```
fosconfig --enable vf
```

4. Enter `y` at the prompt.

The following example verifies the status of VF mode. If VF mode is disabled, enable it using the `fosconfig --enable` command.

```
switch:admin> fosconfig --show
FC Routing service: disabled
Virtual Fabric: disabled

switch:admin> fosconfig --enable vf
WARNING: This is a disruptive operation that requires a reboot to take effect.
All EX ports will be disabled upon reboot.
Would you like to continue [Y/N] y
VF has been enabled. Your system is being rebooted.
```

## Disabling Virtual Fabrics Mode

When you disable VF mode, the following update occurs:

- The CPs are rebooted.
- If F\_Port trunking is enabled on ports in the default switch, the F\_Port trunking information is preserved.

**NOTE**

You must delete all `dapcf` mappings before disabling VF mode.

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use the `fosConfig` command to check whether VF mode is disabled:

```
fosconfig --show
```

3. Move all ports to the default logical switch.

```
lscfg --config 128 -slot slot -port port
```

4. Delete all the non-default logical switches.

```
lscfg --delete fabricID
```

5. Use the `fosConfig` command to disable VF mode:

```
fosconfig --disable vf
```

6. Enter **y** at the prompt.

The following example checks whether VF mode is enabled or disabled and then disables it.

```
switchA:FID128:admin> fosconfig --show

FC Routing service: disabled
Virtual Fabric: enabled

switch:admin> fosconfig --disable vf
WARNING: This is a disruptive operation that requires a reboot to take effect.
Would you like to continue [Y/N] y
```

## Configuring Logical Switches to Use Basic Configuration Values

All switches in the fabric are configured to use the same basic configuration values. When you create logical switches, the logical switches might have different configuration values than the default logical switch. Use the following procedure to ensure that newly created logical switches have the same basic configuration values as the default logical switch.

### NOTE

For most users, you do not need to run this procedure. Contact your switch service provider to determine if you need to use this procedure.

You need to run this procedure only once on each chassis, after you enable Virtual Fabrics but before you create logical switches. The configuration settings are then preserved across reboots and firmware upgrades and downgrades.

Use the following procedure to configure logical switches to use basic configuration values:

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Enter the `configureChassis` command to ensure that newly created logical switches have the same basic configuration values as the default logical switch:

```
configurechassis
```

3. Enter **n** at the prompts to configure system and `cfgload` attributes. Enter **y** at the prompt to configure custom attributes.

```
System (yes, y, no, n): [no] n
cfgload attributes (yes, y, no, n): [no] n
Custom attributes (yes, y, no, n): [no] y
```

4. Enter the appropriate value at the Config Index prompt. Contact your switch service provider to determine the appropriate value.

```
Config Index (0 to ignore): (0..1000) [3]:
```

## Creating a Logical Switch or Base Switch

When a logical switch is created, it is automatically enabled and is empty; that is, it does not have any ports. After creating the logical switch, you must disable the switch to configure it and set the domain ID. You then assign ports to the logical

switch. You can optionally define the logical switch to be a base switch using the `-base` option. Each chassis can have only one base switch. To create a FICON logical switch, see [Creating a FICON Logical Switch](#).

#### NOTE

Domain ID conflicts are detected before fabric ID conflicts. If you have both a domain ID conflict and a fabric ID conflict, only the domain ID conflict is reported.

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Enter the `lsCfg` command to create a logical switch:

```
device#> lscfg --create fabricID [-base | -lisldisable] [-force]
```

The *fabricID* is the fabric ID that is to be associated with the logical switch.

The logical switch is enabled by default and the logical ISLs come online immediately, cause fabric merge, and lead to domain ID conflict if there is a conflict. To avoid this problem, use the `-lisldisable` option to disable the logical ports while creating.

3. Set the context to the new logical switch.

```
device#>setcontext fabricID | switchname
```

The *fabricID* parameter is the FID of the logical switch you just created. The *switchname* parameter is the name assigned to the logical switch. You can only use one parameter at a time.

4. Disable the logical switch.

```
device#>switchdisable
```

5. Configure the switch attributes, including assigning a unique domain ID.

```
device#>configure
```

6. Enable the logical switch.

```
device#>switchenable
```

7. Enable the logical ISLs.

```
device#>lfcfg --lislenable
```

#### Example of Creating a Logical Switch

The following example creates a logical switch with FID 4, and then assigns domain ID 14 to it.

```
device#> lscfg --create 4 -lisldisable
A Logical switch with FID 4 will be created with default configuration.
Would you like to continue [y/n]?y
About to create switch with fid=4. Please wait...
Logical Switch with FID (4) has been successfully created.
Logical Switch has been created with default configurations.
Please configure the Logical Switch with appropriate switch
and protocol settings before activating the Logical Switch.
sw0:FID128:admin> setcontext 4

switch_4:FID4:admin> switchdisable

switch_4:FID4:admin> configure
```

```

Configure...
Fabric parameters (yes, y, no, n): [no] yes
 Domain: (1..239) [1] 100
 Select Addressing Mode:
 (1 = Zero Based Area Assignment,
 2 = Port Based Area Assignment): (1..2) [2] 2
 WWN Based persistent PID (yes, y, no, n): [no]
 (output truncated)
WARNING: The domain ID will be changed. The port level zoning may be affected

switch_4:FID4:admin> switchenable
switch_4:FID4:admin> lfcfg --lislenable

```

After you create the logical switch, you assign ports to it, as described in [Adding and Moving Ports on a Logical Switch](#).

## Executing a Command in a Different Logical Switch Context

If you are in the context of a logical switch, you can execute a command for a different logical switch. You can also execute a command for all of the logical switches in a chassis.

The command is not executed on those logical switches for which you do not have permission.

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Enter one of the following commands:
  - To execute a command in a different logical switch context:

```
fosexec --fid fabricID -cmd " command"
```

- To execute a command on all logical switches:

```
fosexec --fid all -cmd " command"
```

### Executing the switchShow Command in a Different Logical Switch Context

```

sw0:FID128:admin> fosexec --fid 4 -cmd "switchshow"

"switchshow" on FID 4:
switchName: switch_4
switchType: 66.1
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 14
switchId: fffc0e
switchWwn: 10:00:00:05:1e:82:3c:2b
zoning: OFF
switchBeacon: OFF
FC Router: OFF
Fabric Name: Fab4
Allow XISL Use: ON
LS Attributes: [FID: 4, Base Switch: No, Default Switch: No, Address Mode 0]
Index Port Address Media Speed State Proto

```

```
=====
22 22 0e1600 -- N8 No_Module FC Disabled
23 23 0e1700 -- N8 No_Module FC Disabled
```

## Executing the fabricShow Command on All Logical Switches

```
sw0:FID128:admin> foexec --fid all -cmd "fabricshow"

"fabricshow" on FID 128:
Switch ID Worldwide Name Enet IP Addr FC IP Addr Name

97: fffc61 10:00:00:05:1e:82:3c:2a 10.1.18.105 0.0.0.0 >"sw0"

"fabricshow" on FID 4:
Switch ID Worldwide Name Enet IP Addr FC IP Addr Name

14: fffc0e 10:00:00:05:1e:82:3c:2b 10.1.18.105 0.0.0.0 >"switch_4" (output truncated)
```

## Deleting a Logical Switch

The following rules apply to deleting a logical switch:

- You must remove all ports from the logical switch before deleting it.
- You cannot delete the default logical switch.

### NOTE

If you are in the context of the logical switch that you want to delete, you are automatically logged out when the fabric ID changes. To avoid being logged out, make sure that you are in the context of a different logical switch from the one that you are deleting.

1. Connect to the physical chassis and log in using an account with admin permissions.
2. Remove all ports from the logical switch as described in [Adding and Moving Ports on a Logical Switch](#).
3. Enter the `lscfg` command to delete the logical switch:

```
lscfg --delete fabricID
```

The *fabricID* parameter is the fabric ID of the logical switch to be deleted.

### Example of Deleting the Logical Switch with FID 7

```
switch_4:FID4:admin> lscfg --delete 7
A Logical switch with FID 7 will be deleted.
Would you like to continue [y/n]?:y
All active login sessions for FID 7 have been terminated.
Switch successfully deleted.
```

## Adding and Moving Ports on a Logical Switch

You must add ports to a logical switch by moving the ports from one logical switch to another. To remove the port from a logical switch, you must move the port to a different logical switch. You cannot just remove it.

When you move a port from one logical switch to another, the port is automatically disabled.

If you are deploying ICLs in the base switch, all ports associated with those ICLs must be assigned to the base switch. If you are deploying ICLs to connect to default switches (that is, XISL use is not allowed), the ICL ports should be assigned (or left) in the default logical switch.

Refer to [Supported Platforms for Virtual Fabrics](#) for port restrictions.

To move ports from a logical switch, perform the following steps:

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Enter the `lsCfg` command to move ports from one logical switch to another.

```
lsCfg --config fabricID -slot slot -port port
```

The ports are assigned to the logical switch specified by *fabricID* and are removed from the logical switch on which they are currently configured.

If the `-port` option is omitted, all ports on the specified slot are assigned to the logical switch.

3. Enter `y` at the prompt.

The ports are automatically disabled, removed from their current logical switch, and assigned to the logical switch specified by *fabricID*.

### Example of Assigning Ports 18 through 20 to the Logical Switch with FID 5

```
switch:admin> lscfg --config 5 -port 18-20
This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Making this configuration change. Please wait...
Configuration change successful.
Please enable your ports/switch when you are ready to continue.
```

## Displaying the Logical Switch Configuration

Use the following procedure to display the configuration for a logical switch:

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Enter the `lsCfg --show -n` command to display information about all of the logical switches.
3. Enter the `lsCfg --show` command to display a list of all logical switches and the ports assigned to them.

### Example of Displaying Basic Information about All Logical Switches

```
sw0:FID128:admin> lscfg --show -n

Switch Information

FID: 5
SwitchType: BS
DomainID: 3
SwitchName: Basel
FabricName: Fab1

FID: 128
SwitchType: DS
DomainID: 1
SwitchName: sw0
```



```
FabricName: DefFab
```

### Example of Displaying a List of All Logical Switches and the Ports Assigned to Them

```
sw0:FID128:admin> lscfg --show
Created switches: 128(ds) 4 5
Port 0 1 2 3 4 5 6 7 8 9

FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 5 | 5 | (output truncated)
```

## Changing the Fabric ID of a Logical Switch

The fabric ID indicates in which fabric the logical switch participates. By changing the fabric ID, you are moving the logical switch from one fabric to another.

Changing the fabric ID (FID) requires permission for chassis management operations. You cannot change the FID of your own logical switch context.

### NOTE

If you are in the context of the logical switch with the fabric ID you want to change, you are automatically logged out when the fabric ID changes. To avoid being logged out, make sure you are in the context of a different logical switch from the one with the fabric ID you are changing.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `lsCfg` command to change the fabric ID of a logical switch:

```
lsCfg --change fabricID -newfid newFID
```

3. Enter **y** at the prompt.
4. Enable the logical switch.

```
fosexec --fid newFID -cmd "switchenable"
```

### Example of Changing the Fabric ID on the Logical Switch from 5 to 7

```
sw0:FID128:admin> lscfg --change 5 -newfid 7
Changing of a switch fid requires that the switch be disabled.
Would you like to continue [y/n]?: y
Disabling switch...
All active login sessions for FID 5 have been terminated.
Checking and logging message: fid = 5.
Please enable your switch.
sw0:FID128:admin> fosexec --fid 7 -cmd "switchenable"

"switchenable" on FID 7:
```

## Changing a Logical Switch to a Base Switch

Use the following procedure to change a logical switch to a base switch.

1. Connect to the switch and log in using an account with the chassis-role permission.
2. Set the context to the logical switch you want to change, if you are not already in that context.

```
setcontext fabricID (or switchname)
```

The *fabricID* parameter is the FID of the logical switch you want to change to a base switch. The *switchname* parameter is the name assigned to the logical switch. You can only use one parameter at a time.

3. Configure the switch to *not* allow XISL use, as described in [Configuring a Logical Switch for XISL Use](#).
4. Enter the `lsCfg` command to change the logical switch to a base switch:

```
lscfg --change fabricID -base
```

The *fabricID* parameter is the fabric ID of the logical switch with the attributes you want to change.

5. Enable the switch.

```
switchenable
```

### Example of Changing the Logical Switch with FID 7 to a Base Switch

```
switch:admin> setcontext 7
switch:admin> switchshow
switchName: switch_25
switchType: 181.0
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 151
switchId: fffc1e
switchWwn: 10:00:00:05:1e:82:3c:2c
zoning: ON
switchBeacon: OFF
FC Router: OFF
Fabric Name: MktFab7
HIF Mode: ON
Allow XISL Use: ON
```

```
LS Attributes: [FID: 7, Base Switch: No, Default Switch: No, Address Mode 0] (output truncated)
```

```
switch:admin> configure
```

Not all options will be available on an enabled switch.

To disable the switch, use the "switchDisable" command.

Configure...

```
Fabric parameters (yes, y, no, n): [no] y
 WWN Based persistent PID (yes, y, no, n): [no]
 Allow XISL Use (yes, y, no, n): [yes] n
 WARNING!! Disabling this parameter will cause removal of LISLs to
 other logical switches. Do you want to continue? (yes, y, no, n): [no] y
System services (yes, y, no, n): [no]
```

```

switch:admin> lscfg --change 7 -base
Creation of a base switch requires that the proposed new base switch on this system be disabled.
Would you like to continue [y/n]?: y
Disabling the proposed new base switch...
Disabling switch fid 7
Please enable your switches when ready.
switch:admin> switchenable

```

## Configuring a Logical Switch for XISL Use

When you create a logical switch, it is configured to use XISLs by default. However, in some cases XISL use is not supported.

Refer to [Limitations and Restrictions of Virtual Fabrics](#) for restrictions on XISL use.

To configure a logical switch for XISL use, perform the following procedure:

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use the `setContext` command to set the context to the logical switch you want to manage, if you are not already in that context.

```
setContext fabricID | switchname
```

The *fabricID* parameter is the FID of the logical switch you want to switch to and manage. The *switchname* parameter is the name assigned to the logical switch. You can only use one parameter at a time.

3. Use the `switchShow` command and check the value of the `Allow XISL Use` parameter.
4. Enter the `configure` command:
5. Enter **y** after the `Fabric parameters` prompt.

```
Fabric parameters (yes, y, no, n): [no] y
```

6. Enter **y** at the `Allow XISL Use` prompt to allow XISL use; enter **n** at the prompt to disallow XISL use.

```
Allow XISL Use (yes, y, no, n): y
```

7. Respond to the remaining prompts or press **Ctrl-D** to accept the other settings and exit.

## Changing the Context to a Different Logical Fabric

You can change the context to a different logical fabric. Your user account must have permission to access the logical fabric.

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use the `setContext` command to switch to a different logical switch in the chassis:

```
setcontext fabricID | switchname
```

The *fabricID* parameter is the fabric ID of the logical switch you want to switch to and manage. The *switchname* parameter is the name assigned to the logical switch. You can only use one parameter at a time.

### Example of Changing the Context from FID 128 to FID 4

In this example, notice that the prompt changes when you change to a different logical fabric.

```
sw0:FID128:admin> setcontext 4
switch_4:FID4:admin>
```

## Administering Advanced Zoning

Zoning is a fabric-based service that enables you to partition your SAN into logical groups of devices that can access each other.

For example, you can partition your SAN into two zones, *winzone* and *unixzone*, so that your Windows servers and storage do not interact with your UNIX servers and storage. You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions; for example, you can create a temporary zone to back up nonmember devices.


A device in a zone can communicate only with other devices connected to the fabric within the same zone. A device not included in the zone is not available to members of that zone. When zoning is enabled, devices that are not included in *any* zone configuration are inaccessible to all other devices in the fabric.

Zones can be configured dynamically. They can vary in size, depending on the number of fabric-connected devices, and devices can belong to more than one zone.

The following figure shows an FC fabric with three configured zones: Red, Green, and Blue. Within this fabric, the following is true:

- Server 1 can communicate only with the Loop 1 storage devices. (Red zone)
- Server 2 can communicate only with the RAID storage devices. (Blue zone)
- Server 3 can communicate with the RAID and Loop 1 storage devices. (Green zone)
- Loop 2 is not assigned to a zone; no other zoned fabric device can access it.

**Figure 68: Zoning Example**



Refer to [Best practices for zoning](#) for additional information that should be kept in mind when working with zones.

To list the commands associated with zoning, use the `zoneHelp` command. For detailed information on the zoning commands used in the procedures, refer to the *Brocade Fabric OS Command Reference Manual*.

## Approaches to Zoning

The following table lists the various approaches you can take when implementing zoning in a fabric.

**Table 131: Approaches to Fabric-Based Zoning**

| Zoning Approach        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Recommended approach   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Single HBA             | <p>Zoning by single HBA most closely re-creates the original SCSI bus. Each zone created has only one HBA (initiator) in the zone; each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports are added. If the HBA also accesses tape devices, a second zone is created with the HBA and associated tape devices in it. In the case of clustered systems, it could be appropriate to have an HBA from each of the cluster members included in the zone; this is equivalent to having a shared SCSI bus between the cluster members and assumes that the clustering software can manage access to the shared devices.</p> <p>In a large fabric, zoning by single HBA requires the creation of possibly hundreds of zones; however, each zone contains only a few members. Zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change. This zoning philosophy is the preferred method.</p> |
| Alternative approaches |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Application            | <p>Zoning by application typically requires zoning multiple, perhaps incompatible, operating systems into the same zones. This method of zoning creates the possibility that a minor server in the application suite could disrupt a major server (such as a Web server disrupting a data warehouse server). Zoning by application can also result in a zone with a large number of members, meaning that more notifications, such as registered state change notifications (RSCNs), or errors, go out to a larger group than necessary.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Operating system       | <p>Zoning by operating system has issues similar to zoning by application. In a large site, this type of zone can become very large and complex. When zone changes are made, they typically involve applications rather than a particular server type. If members of different operating system clusters can see storage assigned to another cluster, they might attempt to own the other cluster's storage and compromise the stability of the clusters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Port allocation        | <p>Avoid zoning by port allocation unless the administration team has very rigidly enforced processes for port and device allocation in the fabric. Zoning by port allocation does, however, provide some positive features. For instance, when a storage port, server HBA, or tape drive is replaced, the change of WWN for the new device is of no consequence. As long as the new device is connected to the original port, it continues to have the same access rights. The ports on the edge switches can be pre-associated to storage ports, and control of the fan-in ratio (the ratio of the input port to output port) can be established. With this pre-assigning technique, the administrative team cannot overload any one storage port by associating too many servers with it.</p>                                                                                                                                                                                                        |
| Not recommended        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

| Zoning Approach  | Description                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No fabric zoning | Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should be utilized only in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed. |

## Zone Types

Zones enable you to partition your fabric into logical groups of devices that can access each other. These are *regular* or *standard* zones. Unless otherwise specified, all references to zones in this chapter refer to regular zones. Beyond this, Fabric OS has the following types of special zones:

- **LSAN zones**  
Provide device connectivity between fabrics without merging the fabrics. See [LSAN zone configuration](#) for more information. LSAN zones support RFIDs.
- **QoS zones**  
Assign high or low priority to designated traffic flows. These are regular zones with additional QoS attributes specified by adding a QoS prefix to the zone name. See [QoS](#) for more information.
- **Peer zones**  
Allow a principal device port to communicate with other members of the zone. See [Peer Zoning](#) for more information.
- **Target driven zones**  
Allow targets to create or activate peer zones. See [Target Driven Zoning](#) for more information.

## Zone Objects

A *zone object* is any device in a zone, such as:

- Physical port number or port index on the switch
- Node World Wide Name (N-WWN)
- Port World Wide Name (P-WWN)

Zone objects identified by port number or index number are specified as a pair of decimal numbers in the form  $D, I$ , where  $D$  is the domain ID of the switch and  $I$  is the index number on that switch in relation to the port you want to specify.

For example, in directors, 4,30 specifies port 14 in slot number 2 (domain ID 4, port index 30). On fixed-port models, 3,13 specifies port 13 in switch domain ID 3.

The following issues affect zone membership based on the type of zone object:

- When a zone object is the physical port number, all devices connected to that port are in the zone.
- WWNs are specified as 8-byte (16-digit) hexadecimal numbers, separated by colons (:), for example, 10:00:00:90:69:00:00:8a.
- When a zone object is the Node WWN name, only the specified device is in the zone.
- When a zone object is the Port WWN name, only the single port is in the zone.

### NOTE

Using a WWN in the same format as a peer property member (starting with 00) is not allowed.

The types of zone objects used to define a zone can be mixed. For example, a zone defined with the zone objects 2,12; 2,14; 10:00:00:80:33:3f:aa:11 contains the devices connected to domain 2, ports 12 and 14, and a device with the WWN 10:00:00:80:33:3f:aa:11 (either node name or port name) that is connected on the fabric.

## Naming Zone Objects

Zone objects, such as zone configuration name, zone name, and alias name, can have the following characters if all switches in the fabric and both the CPs in a chassis system are running Fabric OS 8.1.0 or later.

- Start with a number or a letter.
- Contain a hyphen (-) other than the first character.
- Contain an underscore (\_) other than the first character.
- Contain a dollar sign (\$) other than the first character.
- Contain a caret (^) other than the first character.

You can use these characters in the zone object names while creating them and subsequently on add, remove, delete, show, enable, copy, rename, and expunge commands related to zone objects.

- The following is an example of using these characters with the `cfgcreate` command.

```
switch:admin> cfgCreate 1-2_cfg, "z1;z2"
switch:admin> zoneshow
Defined configuration:
 cfg: 1-2_cfg z1; z2
Effective configuration:
 no configuration in effect
```

- The following is an example of using these characters with the `zonecreate` command.

```
switch:admin> zonecreate abc_$_def, "8,7;9,9"
switch:admin> zoneshow
Defined configuration:
 zone: abc_$_def 8,7; 9,9
Effective configuration:
 no configuration in effect
```

### NOTE

The \$ sign must be prefixed with a \ while using the command prompt.

- The following is an example of using these characters with the `alicreate` command.

```
switch:admin> alicreate ali^def, "8,7;9,9"
switch:admin> zoneshow
Defined configuration:
 alias: ali^def 8,7; 9,9
Effective configuration:
 no configuration in effect
```

### NOTE

- If you merge a domain running Fabric OS 8.0.1 or lower with a domain running Fabric OS 8.1.0 or later having zone object names starting with numbers, or containing \$, -, or ^, the fabric segments.
- Downgrading from Fabric OS 8.1.0 or later to a version earlier than Fabric OS 8.1.0 is blocked if enhanced zone object names are configured.

## Zoning Schemes

You can establish a zone by identifying zone objects using one or more of the following *zoning schemes* :

- Domain,index (D,I)

All members are specified by *domain ID*, *port number*, or *domain,index number* pairs or aliases.



- **WWN**  
All members are specified only by WWNs or aliases of WWNs. They can be node or port versions of the WWN.
- **Mixed zoning**  
A zone containing members specified by a combination of *domain,port* or *domain,index* or aliases, and WWNs or aliases of WWNs.

In any scheme, you can identify zone objects using aliases.

## Zone Configurations

A *zone configuration* is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is in effect, all zones that are members of that configuration are in effect.

Several zone configurations can reside on a switch at once, and you can quickly alternate between them. For example, you might want to have one configuration enabled during the business hours and another enabled overnight. However, only one zone configuration can be enabled at a time.

The different types of zone configurations are:

- **Defined configuration**  
The complete set of all zone objects defined in the fabric.
- **Effective configuration**  
A single zone configuration that is currently in effect. The effective configuration is built when you enable a specified zone configuration.
- **Saved configuration**  
A copy of the defined configuration plus the name of the effective configuration, which is saved in flash memory. (You can also provide a backup of the zone configuration and restore the zone configuration.) There might be differences between the saved configuration and the defined configuration if you have modified any of the zone definitions and have not saved the configuration.
- **Disabled configuration**  
The effective configuration is removed from flash memory.

If you disable the effective configuration, the Advanced Zoning feature is disabled on the fabric, and all devices within the fabric can communicate with all other devices (unless you previously set up a default zone, as described in [Default Zoning Mode](#)). This does not mean that the zone database is deleted, however, only that no configuration is active in the fabric.

On power-up, the switch automatically reloads the saved configuration. If a configuration was active when it was saved, the same configuration is reinstated on the local switch.

## Zoning Enforcement

Zoning enforcement describes a set of predefined rules that the switch uses to determine where to send incoming data. Fabric OS uses hardware-enforced zoning, which means that each frame is checked by hardware before the frame is delivered to a zone member and discarded if there is a zone mismatch. When hardware-enforced zoning is active, the Fabric OS switch monitors the communications and blocks any frames that do not comply with the effective zone configuration. The switch performs this blocking at the transmit side of the port on which the destination device is located.

The two methods of hardware enforcement:

- **Frame-based hardware enforcement:** All frames are checked by the hardware.
- **Session-based hardware enforcement:** The only frames that are checked by hardware are the ELS frames (such as PLOGI and RNID) used to establish a session.

The hardware-enforcement method that is used depends on how the zones are configured.

A zone can contain all WWN members, or all D,I members, or a combination of WWN and D,I members.

With Gen 6 platforms, frame-based hardware enforcement is in effect if all members of a zone are identified the same way, either using WWNs or D,I notation, but not both. If the zone includes aliases, then the aliases must also be defined the same way as the zone.

With Gen 7 platforms, which support Unified Addressing Mode (UAM), frame-based hardware enforcement is in effect regardless of the way the zone is defined (WWN or D,I notation); It is also applies for associated zone aliases.

With Gen 6 platforms, session-based hardware enforcement is in effect if the zone has a mix of WWN and D,I members. The same occurs if a port is in multiple zones, and it is defined by WWN in one zone and by D,I in another.

With Gen 7 platforms, mixed WWN, and D,I zoning does not result in the session-based hardware enforcement.

With all platforms, session-based hardware enforcement can result if ACL resources in hardware have been exhausted.

## Identifying the Enforced Zone Type

Use the following procedure to identify zones and zone types:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `portZoneShow` command.

## Considerations for Zoning Architecture

The following table lists considerations for zoning architecture.

**Table 132: Considerations for Zoning Architecture**

| Items                                               | Descriptions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type of zoning enforcement: frame- or session-based | If security is a priority, use frame-based hardware enforcement. The best way to do this is to use WWN identification exclusively for all zoning configurations.                                                                                                                                                                                                                                                                                                                                                     |
| Use of aliases                                      | The use of aliases is optional with zoning. Using aliases requires structure when defining zones. Aliases aid administrators of zoned fabrics in understanding the structure and context.                                                                                                                                                                                                                                                                                                                            |
| Effect of changes in a production fabric            | Zone changes in a production fabric can result in a disruption of I/O under conditions when an RSCN is issued because of the zone change and the HBA is unable to process the RSCN fast enough. Although RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, you should perform zone changes only when the resulting behavior is predictable and acceptable. Ensuring that the HBA drivers are current can shorten the response time in relation to the RSCN. |
| Confirming operation                                | After changing or enabling a zone configuration, you should confirm that the nodes and storage can identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.                                                                                                                                                                                                                                                                        |

Zoning can be implemented and administered from any switch in the fabric, although you should use a switch running the latest Fabric OS version.

The zone configuration is managed on a fabric basis. When a change in the configuration is saved, enabled, or disabled according to the transactional model, it is automatically distributed to all switches in the fabric, preventing a single point of failure for zone information.

### NOTE

Zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent commands. For a large fabric, you should wait several minutes between commands.

## Best Practices for Zoning

The following are recommendations for using zoning:

- When using a mixed fabric — that is, a fabric containing two or more switches running different release levels of Fabric OS — you should use the switch with the latest Fabric OS level to perform zoning tasks. Switches with earlier versions of Fabric OS do not have the same capability to view all the functionality that more recent versions of Fabric OS provide, as functionality is backward-compatible but not forward-compatible.
- Zone using the core switch in preference to using an edge switch.
- Zone using a director rather than a switch. Director has more resources to handle zoning changes and implementations.
- When you are adding a switch to the existing fabric, prior to joining the fabric you must set the defzone policy of the switch being added as follows:
  - If the joining switch has locally attached devices that are online, the defzone policy of the switch being added should be set to No Access.
  - If the joining switch has no online locally attached devices, the defzone policy of the switch being added can be set to All Access.

This is done to avoid a transitional state where the All Access policy might lead to excessive RSCN activity; with extreme cases having the potential for additional adverse effects. This is especially important for fabrics having a very high device count.

- Initial WWNs starting with '00' are for engineering use only and are auto-created. Starting with Fabric OS 7.4.0, zones having initial WWNs starting with '00' are treated as peer zones. This means that when you attempt one of the following,
  - A firmware upgrade from firmware earlier than Fabric OS 7.4.0
  - A merge with a switch running firmware earlier than Fabric OS 7.4.0
  - A configuration download using firmware earlier than Fabric OS 7.4.0and the WWN of the first zone member starts with 00, Fabric OS will treat the zone as a peer zone. You must delete such invalid zones before continuing. Refer to [Deleting invalid zones](#) for instructions.
- Whenever possible, peer zones are recommended instead of other zoning schemes such as single-initiator, one-to-many, or flat zoning methods. See [Peer Zoning](#) for more information.

## Zone Aliases

A *zone alias* is a name assigned to a logical group of ports or WWNs. By creating an alias, you can assign a familiar name to a device or group multiple devices into a single name. This simplifies cumbersome data entry and allows an intuitive naming structure (such as using NT\_Hosts to define all NT hosts in the fabric). Using zone aliases eliminates the need for long lists of individual zone member names.

Zone aliases also simplify repetitive entry of zone objects such as port numbers or a WWN. For example, you can use the name Eng as an alias for 10:00:00:80:33:3f:aa:11.

Naming zones for the initiator they contain can also be useful. For example, if you use the alias SRV\_MAILSERVER\_SLT5 to designate a mail server in PCI slot 5, the alias for the associated zone is ZNE\_MAILSERVER\_SLT5. This clearly identifies the server HBA associated with the zone.

Zone configuration naming is flexible. One configuration should be named PROD\_*fabricname*, where *fabricname* is the name that the fabric has been assigned. The purpose of the PROD configuration is to easily identify the configuration that can be implemented and provide the most generic services. If other configurations are used for specialized purposes, names, such as BACKUP\_A, RECOVERY\_2, and TEST\_18jun02, can be used. If you are creating a new alias using `aliCreate w, "1,1"`, and a user in another telnet session executes `cfgEnable` (or `cfgDisable`, or `cfgSave`), the other user's transaction will abort your transaction and you will receive an error message. Creating a new alias while there is a zone merge taking place may also abort your transaction. For more details about zone merging and zone merge conflicts, refer to [Zone Merging](#).

**Virtual Fabrics considerations:** Alias definitions should not include logical port numbers. Zoning is not enforced on logical ports.

**TDZ considerations:** After an alias becomes a member of a TDZ, the alias cannot be edited. The alias can be edited only after the alias is removed from the TDZ by the target device or if the TDZ is deleted.

## Creating an Alias

Use the following procedure to create an alias.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aliCreate` command using the following syntax:

```
alicreate "aliasname", "member[; member...]"
```

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to the nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message displays on the other switches to indicate that the transaction was aborted.

An example follows.

```
switch:admin> alicreate "array1", "2,32; 2,33; 2,34; 4,4"

switch:admin> alicreate "array2", "21:00:00:20:37:0c:66:23; 4,3"

switch:admin> alicreate "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.

Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no]
```

## Adding Members to an Alias

Use the following procedure to add a member to an alias.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aliAdd` command using the following syntax:

```
aliadd "aliasname", "member[; member...]"
```

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to the nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

An example follows.

```
switch:admin> aliadd "array1", "1,2"

switch:admin> aliadd "array2", "21:00:00:20:37:0c:72:51"

switch:admin> aliadd "loop1", "5,6"
```

```
switch:admin> cfgsave
```

```
WARNING!!!
```

```
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
```

```
Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no]
```

## Removing Members from an Alias

Use the following procedure to remove a member from an alias.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aliRemove` command using the following syntax:

```
aliremove "aliasname", "member[; member...]"
```

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to the nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

An example follows.

```
switch:admin> aliremove "array1", "1,2"
```

```
switch:admin> aliremove "array2", "21:00:00:20:37:0c:72:51"
```

```
switch:admin> aliremove "loop1", "4,6"
```

```
switch:admin> cfgsave
```

```
WARNING!!!
```

```
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
```

```
Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no]
```

## Deleting an Alias

Use the following procedure to delete an alias.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aliDelete` command using the following syntax.

```
alidelete "aliasname"
```

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

For example,

```
switch:admin> alidelete "array1"

switch:admin> cfgsave
WARNING!!!
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.

Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no]
```

## Viewing an Alias in the Defined Configuration

To view an alias in the defined configuration, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `alishow --ic` command, followed optionally by the "`<pattern>`" and `<mode>` operands, to display the zone configuration information.

### NOTE

If you do not specify any parameters, the entire zone database (both the defined and effective configuration) is displayed. Refer to the *Brocade Fabric OS Command Reference Manual* for more information.

### Displaying All Zone Aliases Regardless of Case

The following example shows all zone aliases beginning with `arr`, regardless of the case:

```
switch:admin> alishow --ic "arr*"
alias: array1 20:e0:00:05:33:11:1f:00
alias: ARRAY2 50:11:00:05:33:c1:37:a2
alias: array3 21:00:00:20:37:0c:66:23
```

### Displaying All Zone Aliases

The following example shows all zone aliases beginning with `arr`, but not if the alias uses these characters in a different case. Here `array1` and `array3` are displayed, but not `ARRAY2`:

```
switch:admin> alishow "arr*"
alias: array1 21:00:00:20:37:0c:76:8c
alias: array3 21:00:00:20:37:0c:66:23
```

## Zone Creation and Maintenance

Fabric OS allows you to create zones to manage devices in a better way. In virtual fabrics, zone definitions must not include logical port numbers. Zoning is not enforced on the logical ports.

## Displaying Existing Zones

Use the following procedure to display a list of existing zones.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgShow` command.

Example of displaying existing zones.

```
switch:admin> cfgshow
Defined configuration:
cfg: sample_cfg
 matt
zone: matt 30:06:00:07:1e:a2:10:20; 3,2
```

```
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2
```

Effective configuration:

```
cfg: sample_cfg
zone: matt 30:06:00:07:1e:a2:10:20
 3,2
```

switch:admin>

## Creating a Zone

### ATTENTION

The `zoneCreate` command adds *all* zone member aliases that match the *aliasname\_pattern* in the zone database to the new zone.

Use the following procedure to create a zone.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `zoneCreate` command using either of the following syntaxes:

```
zonecreate "zonename", "member[; member...]"
zonecreate "zonename", "aliasname_pattern *[, members]"
```

### NOTE

The `zoneCreate` command supports partial pattern matching (wildcards) of zone member aliases, which allows you to add multiple aliases that match the *aliasname\_pattern* in the command line.

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to the nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

4. Enter the `cfgShow` command to view the changes.

Example of creating a zone.

```
switch:admin> zonecreate sloth, "b*; 10:00:00:00:01:1e:20:20"
```

```
switch:admin> cfgadd "sample_cfg", "sloth"
```

```
switch:admin> cfgenable sample_cfg
```

```
switch:admin> cfgshow
```

Defined configuration:

```
cfg: sample_cfg
 matt; sloth
zone: matt 30:06:00:07:1e:a2:10:20; 3,2; bond; jake; jeff; jones; zeus
```

```

zone: sloth 10:00:00:00:01:1e:20:20; bawn; bolt; bond; brain
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2

```

Effective configuration:

```

cfg: sample_cfg
zone: matt 30:06:00:07:1e:a2:10:20
 3,2
 10:00:05:1e:a9:20:00:01
 3,5
 4,7
 8,9
 14,11
 30:00:00:05:1e:a1:cd:02
 40:00:00:05:1e:a1:cd:04
 4,5
 7,3
 6,8
 9,2
zone: sloth 10:00:00:00:01:1e:20:20
 3,5
 4,8
 10:00:00:02:1f:02:00:01
 10:00:05:1e:a9:20:00:01
 11,4
 22,1
 33,6

```

## Adding Devices (Members) to a Zone

Use the following procedure to add members to a zone:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `zoneAdd` command using either of the following syntaxes:

```

zoneadd "zonename", "member[; member...]"
zoneadd "zonename", "aliasname_pattern* [; members]"

```

### NOTE

The `zoneAdd` command supports partial pattern matching ("wildcards") of zone member aliases, which allows you to add multiple aliases that match the `aliasname_pattern` in the command line.

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.



#### 4. Enter the `cfgShow` command to view the changes.

Example of adding members to a zone.

```
switch:admin> zoneadd matt, "ze*; bond*; j*"

switch:admin> cfgenable sample_cfg

switch:admin> cfgshow
Defined configuration:
cfg: sample_cfg
 matt; sloth
zone: matt 30:06:00:07:1e:a2:10:20; 3,2; bond; jake; jeff; jones; zeus
zone: sloth 10:00:00:00:01:1e:20:20; bawn; bolt; bond; brain
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2

Effective configuration:
cfg: sample_cfg
zone: matt 30:06:00:07:1e:a2:10:20
 3,2
 10:00:05:1e:a9:20:00:01
 3,5
 4,7
 8,9
 14,11
 30:00:00:05:1e:a1:cd:02
 40:00:00:05:1e:a1:cd:04
 4,5
 7,3
 6,8
 9,2
zone: sloth 10:00:00:00:01:1e:20:20
 3,5
 4,8
 10:00:00:02:1f:02:00:01
 10:00:05:1e:a9:20:00:01
 11,4
 22,1
 33,6

switch:admin>
```

## Removing Devices (Members) from a Zone

Use the following procedure to remove members from a zone:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `zoneRemove` command using either of the following syntaxes:

```
zoneremove "zonename", "member[; member ...]"
```

```
zoneremove "zonename", "aliasname_pattern* [; members]"
```

### NOTE

This command supports partial pattern matching (wildcards) of zone member aliases, which allows you to remove multiple aliases that match the *aliasname\_pattern* in the command line.

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

4. Enter the `cfgShow` command to view the changes.

Examples of removing members from a zone.

```
switch:admin> cfgshow
Defined configuration:
cfg: sample_cfg
 matt; sloth
zone: matt 30:06:00:07:1e:a2:10:20; 3,2; bond; jake; jeff; jones; zeus
zone: sloth 10:00:00:00:01:1e:20:20; bawn; bolt; bond; brain
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2
```

```
Effective configuration:
cfg: sample_cfg
zone: matt 30:06:00:07:1e:a2:10:20
 3,2
 10:00:05:1e:a9:20:00:01
 3,5
 4,7
 8,9
 14,11
 30:00:00:05:1e:a1:cd:02
 40:00:00:05:1e:a1:cd:04
 4,5
 7,3
```

```

 6,8
 9,2
zone: sloth 10:00:00:00:01:1e:20:20
 3,5
 4,8
 10:00:00:02:1f:02:00:01
 10:00:05:1e:a9:20:00:01
 11,4
 22,1
 33,6

switch:admin>

switch:admin> zoneremove matt,"30:06:00:07:1e:a2:10:20; ja*; 3,2"

switch:admin> cfgsave

switch:admin> cfgenable sample_cfg

switch:admin> cfgshow
Defined configuration:
cfg: sample_cfg
 matt; sloth
zone: matt bond; jeff; jones; zeus
zone: sloth 10:00:00:00:01:1e:20:20; bawn; bolt; bond; brain
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2

Effective configuration:
cfg: sample_cfg
zone: matt 10:00:05:1e:a9:20:00:01
 3,5
 30:00:00:05:1e:a1:cd:02
 40:00:00:05:1e:a1:cd:04
 4,5
 7,3
 4,7
 6,8
 9,2
zone: sloth 10:00:00:00:01:1e:20:20
 3,5
 4,8
 10:00:00:02:1f:02:00:01
 10:00:05:1e:a9:20:00:01
 11,4
 22,1
 33,6

```

```
switch:admin>
```

## Replacing Zone Members

Fabric OS allows you to replace one zone member with another zone member using a CLI command. This command takes two inputs. The first is the member to be replaced, and the second is the new member. These inputs can be formatted only with WWN or D,I zoning schemes.

### Notes and Restrictions

- To make a configuration change effective, issue a `cfgEnable` command after the `zoneObjectReplace` command. Otherwise, the changes are in the transaction buffer but are not committed.
- Only members of regular zones, peer zones (except for TDZs), and aliases (those identified using either D,I or WWN) can be replaced using `zoneobjectreplace`.
- The `zoneobjectreplace` command does not work with aliases. Alias members (that is, members inside an alias) can be replaced using `zoneobjectreplace`, but an alias itself cannot be directly replaced. To achieve the effect of replacement, create an alias (with the desired new name) containing the same members, and then delete the old alias.

Use the following procedure to replace members in a zone.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `zoneObjectReplace` command, using the following syntax:

```
zoneobjectreplace old wwn/D,I new wwn/D,I
```

#### NOTE

The `zoneObjectReplace` command does *not* support partial pattern matching (wildcards) of zone member aliases.

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to the nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

4. Enter the `cfgShow` command to view the changes.

Example replacing zone members.

```
switch:admin> cfgshow
Defined configuration:
cfg: sample_cfg
 matt; sloth
zone: matt bond; jeff; jones; zeus
zone: sloth 10:00:00:00:01:1e:20:20; bawn; bolt; bond; brain
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2
```

Effective configuration:

<output truncated>

```
switch:admin> zoneobjectreplace "3,5" "10,30"
```

```
switch:admin> cfgsave
```

```
switch:admin> cfgenable "sample_cfg"
```

```
switch:admin> cfgshow
```

Defined configuration:

```
cfg: sample_cfg
 matt; sloth
zone: matt bond; jeff; jones; zeus
zone: sloth 10:00:00:00:01:1e:20:20; bawn; bolt; bond; brain
alias: bawn 4,8; 10,30
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 10,30
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2
```

Effective configuration:

<output truncated>



### CAUTION

Executing this command replaces all instances of the older member with the new member in the entire zone database.

## Deleting a Zone

Use the following procedure to delete a zone.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `zoneDelete` command using the following syntax:

```
zonedeldelete "zonename"
```

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

Example of deleting zone members.

```
switch:admin> cfgshow
```

Defined configuration:

```
cfg: sample_cfg
 matt; sloth
zone: matt bond; jeff; jones; zeus
zone: sloth 10:00:00:00:01:1e:20:20; bawn; bolt; bond; brain
alias: bawn 4,8; 10,30
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 10,30
```

```
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2
```

Effective configuration:

<output truncated>

```
switch:admin> zonedel delete sloth
switch:admin> cfmgr remove sample_cfg, sloth
switch:admin> cfmgr save
switch:admin> cfmgr enable sample_cfg
switch:admin> cfmgr show
```

Defined configuration:

```
cfg: sample_cfg
 matt
zone: matt bond; jeff; jones; zeus
alias: bawn 4,8; 10,30
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 10,30
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2
```

Effective configuration:

<output truncated>

```
switch:admin>
```

## Viewing a Zone

The `zoneshow --sort` command displays the defined and effective configurations for a zone. Prior to Fabric OS 7.4.0, the zone members were displayed with sorted D,I pairs. Starting with Fabric OS 7.4.0, the zone members are displayed with D,I pairs and WWNs sorted in ascending order.

Use the following procedure to view a zone in the configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. View the zone information using the `zoneShow` command.

If no parameters are specified, the entire zone database with both the defined and effective configurations is displayed.

### Example of Displaying a Zone in a Defined Configuration

The following example shows all zones beginning with A, B, or C, in ascending order.

```
switch:admin> zoneshow --sort "[A-C]*"

zone: Blue_zone 1,1; array1; 1,2; array2
zone: Bobs_zone 4,5; 4,6; 4,7; 4,8; 4,9
```

## Example of Displaying the Complete Zone Database

The following example shows all zones in the zone database.

```
switch:admin> zoneshow --verbose
Defined configuration:
cfg: cfgnpr npr
cfg: cfgz prz; diz; lsan_bb
zone: diz 1,2; 4,7; 3,2; 1,1
zone: diz1 00:02:00:00:00:02:00:01; 10,4; 10,6; 10,5
zone: lsan_bb 30:12:00:05:1e:61:23:8f; 30:00:00:05:1e:61:23:8f
zone: npr 20:04:00:05:33:88:bb:be; 20:06:00:05:33:88:bb:be;
 20:05:00:05:33:88:bb:be
zone: prz 00:02:00:00:00:03:00:01; 20:04:00:05:33:88:bb:be;
 20:06:00:05:33:88:bb:be; 20:05:00:05:33:88:bb:be
zone: upr 00:02:00:00:00:03:00:02; 20:06:00:05:33:88:bb:be;
 20:05:00:05:33:88:bb:be; 20:04:00:05:33:88:bb:be

Effective configuration:
cfg: cfgz
zone: diz 1,2
 4,7
 3,2
 1,1
zone: lsan_bb 30:12:00:05:1e:61:23:8f
 30:00:00:05:1e:61:23:8f
zone: prz 00:02:00:00:00:03:00:01
 20:04:00:05:33:88:bb:be
 20:06:00:05:33:88:bb:be
 20:05:00:05:33:88:bb:be
```

## Example of Displaying the Sorted Zone Database

The following example shows all zones in the zone database with sorted D,I pairs and WWNs.

```
switch:admin> zoneshow --sort -verbose
Defined configuration:
cfg: cfgnpr npr
cfg: cfgz prz; diz; lsan_bb
zone: diz 1,1; 1,2; 3,2; 4,7
zone: diz1 00:02:00:00:00:02:00:01; 10,4; 10,5; 10,6
zone: lsan_bb 30:00:00:05:1e:61:23:8f; 30:12:00:05:1e:61:23:8f
zone: npr 20:04:00:05:33:88:bb:be; 20:05:00:05:33:88:bb:be;
 20:06:00:05:33:88:bb:be
zone: prz 00:02:00:00:00:03:00:01; 20:04:00:05:33:88:bb:be;
 20:05:00:05:33:88:bb:be; 20:06:00:05:33:88:bb:be
zone: upr 00:02:00:00:00:03:00:02; 20:05:00:05:33:88:bb:be;
 20:06:00:05:33:88:bb:be; 20:04:00:05:33:88:bb:be

Effective configuration:
cfg: cfgz
zone: diz 1,1
 1,2
 3,2
```

```

 4,7
zone: lsan_bb 30:00:00:05:1e:61:23:8f
 30:12:00:05:1e:61:23:8f
zone: prz 00:02:00:00:00:03:00:01
 20:04:00:05:33:88:bb:be
 20:05:00:05:33:88:bb:be
 20:06:00:05:33:88:bb:be

```

## Viewing Zone Configuration Names without Case Distinction

Use the following procedure to view selected zone configuration names for a given pattern without case distinction.

1. Connect to the switch and log in using an account with admin permissions.
2. Use the `zoneShow` command to view configuration names.

```
zonestow[--ic] [" pattern "] [, mode]
```

The following example shows all green zones using pattern search, regardless of the case.

```

switch:admin> zonestow --ic GREEN*

zone: GREEN 44,4; 21:00:00:20:37:0c:71:02; 8,9
zone: green 2,2; 2,3; 21:00:00:20:37:0c:76:8c

```

## Examining Changes in the Zone Database

Fabric OS allows you to check and display any differences between the transaction buffer and the committed database by using the options `--transdiffs` and `--transdiffsonly` in the `zoneShow` and `cfgShow` commands.

The options use the format in the following commands:

```

zoneShow --transdiffs
zoneShow --transdiffsonly
cfgShow --transdiffs
cfgShow --transdiffsonly

```

To reflect the changes made to the zone database (a new zone is added or an existing zone is deleted, or a zone member is added or deleted or any other valid zone database entity is modified), the following notation is used.

- An asterisk (\*) at the start indicates a change in that zone, zone configuration, alias, or any other entity in the zone database.
- A plus sign (+) before any entity (an alias or a zone name or a configuration) indicates that it is a newly added entity.
- A minus sign (-) before any entity indicates that this entity has been deleted. If zone members are added as well as deleted in a zone configuration, a plus sign and a minus sign (+-) are displayed before the member and a \* sign are displayed before the zone name.
- A plus sign (+) before any member of an alias, zone name, or any other entity indicates that this member has been added, and a minus sign (-) indicates that the particular member has been deleted. Notice that the minus sign (-) comes before the port index. For a deleted zone member, it would have been shown as 5,-1.

The following example displays the existing zone database.

```

switch:admin> cfgshow
Defined configuration:
cfg: fabric_cfg Blue_zone

```



```

zone: Blue_zone
1,1; array1; 1,2; array2
zone: green_zone
1,1; 1,2
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28

```

Effective configuration:

```

cfg: fabric_cfg
zone: Blue_zone
1,1
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2

```

The following example adds a new zone red\_zone, deletes 1,1 and adds 6,15 to green\_zone.

```

switch:admin> cfgshow --transdiffs
Defined configuration:
cfg: fabric_cfg Blue_zone
zone: Blue_zone
1,1; array1; 1,2; array2
*zone: green_zone
-1,1; 1,2; +6, 15
*zone: +red_zone
5,1; 4,2
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28

```

Effective configuration:

```

cfg: fabric_cfg
zone: Blue_zone
1,1
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2

```

The following example displays output for `cfgShow --transdiffonly`.

```

switch:admin> cfgshow --transdiffonly
*zone: green_zone -1,1; 1,2; +6,15
*zone: +red_zone 5,1; 4,2

```

## Validating a Zone

Validating a zone evaluates the enforceability of each zone member. Zone members that are defined using worldwide names are flagged as nonexistent if the specified device is not present within the fabric (offline). Zone members that are defined using domain and port index are flagged as nonexistent if either the specified port index is out of bounds, or if the specified domain is not a member of the fabric.

Use the following procedure to validate a zone.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgShow` command to view the zone configuration objects you want to validate.

```
switch:admin> cfgShow

Defined configuration:
cfg: USA_cfg Purple_zone; White_zone; Blue_zone
zone: Blue_zone
1,1; array1; 1,2; array2

zone: Purple_zone
1,0; loop1

zone: White_zone

alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Enter the following command to validate all zones in the zone database in the defined configuration.

```
switch:admin> zone --validate -m 1
Defined configuration:
cfg: cfg1 zone1
cfg: cfg2 zone1; zone2
zone: zone1 1,1; ali1
zone: zone2 1,1; ali2
alias: ali1 10:00:00:05:1e:35:81:7f*; 10:00:00:05:1e:35:81:7d*
alias: ali2 10:00:00:05:1e:35:81:09*; 10:00:00:05:1e:35:81:88*

~ - Invalid configuration
* - Member does not exist
```

The mode flag `-m` specifies the zone database location. Supported mode flag values follow:

- 0 – Zone database from the current transaction buffer
- 1 – Zone database stored from the persistent storage
- 2 – Currently an effective zone database.

If no mode options are given, the validated output of all three buffers is shown.

If the `-f` option is specified, all the zone members that are not enforceable are expunged in the transaction buffer. This pruning operation always occurs on the transaction and defined buffers. You cannot specify a mode option or cannot specify a zone object as an argument with the `-f` option. This mode flag should be used after the zone has been validated.

If the `-i` option is specified, all zone members for a given pattern are listed without the case distinction.

The following example shows how to validate the zone members beginning with *gre*, regardless of the case.

```
switch:admin> zone --validate -i gre*
Defined configuration:
```

```

zone: GREEN 44, 4; 21:00:00:20:37:0c:71:02; 8,9
zone: green 2,2*; 2,3*; 21:00:00:20:37:0c:76:8c*
Effective configuration:
zone: green 2,2*
2,3*
21:00:00:20:37:0c:76:8c*

~ - Invalid configuration
* - Member does not exist

```

## Inconsistencies between the Defined and Effective Configurations

If you edit zone objects in the defined configuration that also exist in the effective configuration and then issue the `cfgSave` command, a warning message is posted stating that a mismatch is observed between the defined and effective configurations, and you are asked to confirm that you want `cfgSave` to continue. If you enter `y`, the updated configuration is saved; if you enter `n`, the updated configuration is discarded.

The following is an example of a warning message.

```

switch: admin> cfgsave
WARNING!!!
The changes you are attempting to save will render the
Effective configuration and the Defined configuration inconsistent.
The inconsistency will result in different Effective Zoning
configurations for switches in the fabric if a zone merge or
HA failover happens. To avoid inconsistency it is recommended to
commit the configurations using the 'cfgenable' command.
Do you want to proceed with saving the Defined
zoning configuration only? (yes, y, no, n): [no] yes

```

If you enter `yes`, and the `cfgSave` operation completes successfully, the following RASLog message is posted.

```

[ZONE-1062], 620/181, FID 128, WARNING, sw0, Defined and Effective zone configurations are
inconsistent, ltime:2012/09/03-23:18:30:983609

```

You can reenable the updated configuration. If there is no impact on the effective configuration with the latest update to the defined configuration, the following message is displayed.

```

You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.

```

```

Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no]

```

The following is an example of the inconsistent defined and effective configuration warning to use.

```

switch: admin> zoneShow
Defined configuration:
cfg: cfg1 zone1; zone2
zone: zone1 10:00:00:00:00:00:00:01; 10:00:00:00:00:00:00:02
zone: zone2 1,1; 1,2

Effective configuration:
cfg: cfg1
zone: zone1 10:00:00:00:00:00:00:01
 10:00:00:00:00:00:00:02

```

```

zone: zone2 1,1; 1,2

switch: admin> zoneadd zone1, 10:00:00:00:00:00:03
switch: admin> cfgsave
WARNING!!!
The changes you are attempting to save will render the
Effective configuration and the Defined configuration inconsistent.
The inconsistency will result in different Effective Zoning
configurations for switches in the fabric if a zone merge or
HA failover happens. To avoid inconsistency it is recommended to
commit the configurations using the 'cfgenable' command.
Do you want to proceed with saving the Defined zoning
configuration only? (yes, y, no, n): [no] y
Updating flash ...

switch:admin> zoneshow
Defined configuration:
cfg: cfg1 zone1; zone2
zone: zone1 10:00:00:00:00:00:00:01; 10:00:00:00:00:00:00:02; 10:00:00:00:00:00:00:03
zone: zone2 1,1; 1,2

Effective configuration:
cfg: cfg1
zone: zone1 10:00:00:00:00:00:00:01; 10:00:00:00:00:00:00:02
zone: zone2 1,1; 1,2

```

## Default Zoning Mode

The default zoning mode controls device access when there is no effective zone configuration. The default zoning mode has two options:

- All Access – All devices within the fabric can communicate with all other devices.
- No Access – Devices in the fabric cannot access any other device in the fabric.

The default zone mode applies to the entire fabric, regardless of switch type.

The default setting is All Access.

Typically, when you disable the zoning configuration in a large fabric with thousands of devices, the name server indicates to all hosts that they can communicate with each other. Each host can receive an enormous list of PIDs, and ultimately cause other hosts to run out of memory or crash. To ensure that all devices in a fabric do not see each other during a configuration disable operation, set the default zoning mode to No Access.

### NOTE

For switches in large fabrics, the default zone mode should be set to No Access. You cannot disable the effective configuration if the default zone mode is All Access and you have more than 120 devices in the fabric.

## Setting the Default Zoning Mode

### NOTE

You should not change the default zone mode from "No Access" to "All Access" if there is no effective zone configuration and more than 120 devices are connected to the fabric.

Use the following procedure to set the default zoning mode.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgActvShow` command to view the current zone configuration.
3. Enter the `defzone --noaccess` or `defzone --allaccess` command. To avoid the user interaction steps in the CLI commands, you can use the `-force` option.

```
switch:admin> defzone --noaccess -force
2021/06/12-12:19:14 (GMT), [ZONE-1043], 929, FID 128, INFO, w84, The Default Zone access mode is set to No
Access.
```

This command initiates a transaction (if one is not already in progress).

4. Enter the `cfgsave`, `cfgenable`, or `cfgdisable` command to commit the change and distribute it to the fabric. The change will not be committed and distributed across the fabric if you do not enter one of these commands.

```
switch:admin> cfgsave

WARNING!!!
The changes you are attempting to save will render the Effective configuration and the Defined
configuration inconsistent. The inconsistency will result in different Effective Zoning configurations
for switches in the fabric if a zone merge or HA failover happens. To avoid inconsistency it is
recommended to commit the configurations using the 'cfgenable' command.
Do you still want to proceed with saving the Defined
zoning configuration only? (yes, y, no, n): [no] y

Updating flash ...
```

## Viewing the Current Default Zone Access Mode

Use the following procedure to view the current default zone access mode.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `defZone --show` command.

### NOTE

If you perform a firmware download of an older release, the current default zone access state appears as it did prior to the download. For example, if the default zoning mode was No Access before the download, it remains as No Access afterward.

## Zone Database Size

The maximum size of a zone database is the upper limit for the defined configuration, and it is determined by the amount of flash memory available for storing the defined configuration.

The zone database size limits are:

- 4MB per fabric and logical switch
- 16MB per chassis (combined size of all zone databases across all logical switches)

The 4MB configuration size is allowed only if all the switches in the fabric are running Fabric OS 9.0.0 version or later. The 16MB zone database limit is chassis wide and hence on virtual-fabric enabled systems, the sum of the zone database size of all the partitions cannot exceed 16 MB. You can create and save a defined configuration up to 4MB in the fabric and the compressed configuration size must not exceed 2 MB during save or enable operations. You can edit the

temporary config and can check the overall size and compressed size of the temporary config by using the `cfgsize` command.

#### NOTE

- Firmware downgrade from Fabric OS 9.0.x or 9.1.x version to a pre-9.0.x version is allowed only if the zone database size is within the limit of the previous release. To downgrade to a pre-9.0.x firmware version, the chassis wide committed zone database size must be less than 2MB for a director, and the per logical switch zone config size must be less than 1MB for a switch.
- The scale of LSAN zones remains at 2MB for the total size of all the enabled LSAN zones in the effective configuration. If a user attempts to add more than 2MB of LSAN zones to the effective configuration and user performs a `cfgenable`, a ZONE-1082 RASLOG message is generated and the `cfgenable` operation fails.

**Zone merging considerations:** Whenever a new zone configuration is merged or pushed to a switch, the total size of the zone configuration in all the partitions of the local and remote domains is recalculated and restricted if the size exceeds the allowed size. If a new switch with the Fabric OS version 9.0.0 tries to join a fabric that has switches with the same firmware version, the zone merge succeeds, and all switches in the fabric have 4MB config. If a new switch with the Fabric OS version 9.0.0 tries to join a fabric that has switches running pre-9.0.0 firmware, and if the committed config is 2MB, the zone merge succeeds.

When a conflict occurs during a zone merge, the output from using the `switchshow` command shows the reason for segmentation, which could be either of the following:

- If a new switch with the Fabric OS version 9.0.0 tries to join a fabric that has switches with the same firmware versions, and if the total size of the committed config is more than 4MB, the zone merge fails and the ISL is segmented. The ZONE-1071 RASLOG is generated with the error message:

```
Fabric Zonedb size exceeded.
```

- If a new switch with the Fabric OS version 9.1.0 tries to join a fabric that has switches running pre-9.0.x firmware, and if the committed config is 4MB, the zone merge fails and the following RASLOGs are generated with the error message:

```
Zone merge failed due to fabric Zonedb size exceeding limits of switch attempting merge.
```

- ZONE-1079 on the 9.0.0 switch
- FABR-1001 on both the switches

**HA out-of-sync considerations:** When the active CP is running 9.0.x or 9.1.x version and the standby CP has a pre-9.0.x version or vice-versa, and if the committed config size on the 9.0.x or 9.1.x CP exceeds the limits of the pre-9.0.0 CP, the HA sync fails. The following error message is displayed in both active and standby CP console:

```
2019/02/04-18:30:41, [ZONE-1028], 4469, SLOT 1 | FID 1, WARNING, AL128_fid1, Commit zone DB size is larger
than the supported limit of the fabric - 4168964
```

```
greater than -4149059.
```

```
2019/02/04-18:30:41, [FSS-1009], 4470, SLOT 1 | FFDC | CHASSIS, ERROR, Allegiance-128, FSS Error: fcswl-zone:
sync-failure: -416.
```

```
2019/02/04-18:30:41, [ZONE-1028], 4471, SLOT 1 | FID 2, WARNING, AL128_fid2, Commit zone DB size is larger
than the supported limit of the fabric - 4168964
```

```
greater than -4149059.
```

```
2019/02/04-18:30:41, [FSS-1009], 4472, SLOT 1 | FFDC | CHASSIS, ERROR, Allegiance-128, FSS Error: fcsw2-zone:
sync-failure: -416.
```

To recover, after the system has stabilized (approximately 5 minutes), reduce the zone database size in the active CP and reboot the standby CP to bring the CPs back in sync with each other.

**Virtual Fabrics considerations:** If Virtual Fabrics is enabled, the sum of the zone database sizes on all the logical fabrics must not exceed the maximum size that is allowed for the chassis. The chassis-wide restriction is displayed in the output as well.

```
switch:admin> cfgsize

Chassis-Wide Max Zone DB size - 16777216 bytes
Chassis-Wide Committed Zone DB size - 1289 bytes
Current Logical Switch Max Zone DB size - 4194304 bytes
Fabric-Wide Max Zone DB size - (Domain 001) 4194304 bytes
Available Zone DB size - 4193015 bytes

Current Logical switch zone config sizes:
 committed - 1289
 transaction - 0
Compressed sizes:
 committed - 44
 transaction - 0
```

#### NOTE

Though Fabric OS does not impose any restriction on the number of members you can have in a zone, it is not advisable to have more than 255 members. Exceeding this limit could lead to the unpredictable behavior.

## Zone Configurations

You can store a number of zones in a zone configuration database. The maximum number of items that can be stored in the zone configuration database depends on the following criteria:

- The lowest maximum size limit supported by the switch having the oldest version of FOS in the fabric.
- The size limits are measured in bytes (or ASCII characters) and limits vary depending on a variety of factors such as the length of each zone object name and the number of members residing in that particular zone object. More density can be achieved if zone objects with shorter names are defined and if the member list of each zone object is defined efficiently (for example, zoning only those devices together that need to communicate instead of zoning everything together in one large zone). The use of peer zones is recommended.

When enabling a new zone configuration, ensure that the size of the defined configuration does not exceed the maximum configuration size supported by all switches in the fabric. This is particularly important if you downgrade to a Fabric OS version that supports a smaller zone database than the current Fabric OS. In this scenario, the zone database in the current Fabric OS must be changed to the smaller zone database before the downgrade.

You can use the `cfgSize` command to check both the maximum available size and the currently saved size on all switches. If you think you are approaching the maximum, you can save a partially completed zone configuration and use the `cfgSize` command to determine the remaining space. The `cfgSize` command reports the maximum available size on the current switch only. It cannot determine the maximum available size on other switches in the fabric.

For important considerations for managing zoning in a fabric, and more details about the maximum zone database size for each version of the Fabric OS, refer to [Zone Database Size](#).

If you create or make changes to a zone configuration, you must enable the configuration for the changes to take effect.

## Creating a Zone Configuration

Use the following procedure to create a zone configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgCreate` command using the following syntax:

```
cfgcreate "cfgname", "member[; member ...]"
```

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

```
switch:admin> cfgcreate "NEW_cfg", "purplezone; bluezone; greenzone"
```

```
switch:admin> cfgsave
```

WARNING!!!

The changes you are attempting to save will render the Effective configuration and the Defined configuration inconsistent. The inconsistency will result in different Effective Zoning configurations for switches in the fabric if a zone merge or HA failover happens. To avoid inconsistency it is recommended to commit the configurations using the 'cfgenable' command.

Do you still want to proceed with saving the Defined zoning configuration only? (yes, y, no, n): [no] y

## Adding Zones to a Zone Configuration

Use the following procedure to add members to a zone configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgAdd` command using the following syntax:

```
cfgadd "cfgname", "member[; member ...]"
```

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed in the other switches to indicate that the transaction was aborted.

```
switch:admin> cfgadd "newcfg", "bluezone"
```

```
switch:admin> cfgsave
```

WARNING!!!

The changes you are attempting to save will render the Effective configuration and the Defined configuration inconsistent. The inconsistency will result in different Effective Zoning configurations for switches in the fabric if a zone merge or HA failover happens. To avoid inconsistency it is recommended to commit the configurations using the 'cfgenable' command.

Do you still want to proceed with saving the Defined zoning configuration only? (yes, y, no, n): [no] y

## Removing Members from a Zone Configuration

Use the following procedure to remove members from a zone configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgRemove` command using the following syntax:



```
cfgremove "cfgname", "member[; member ...]"
```

3. Enter the `cfgsave` command to save the change to the defined configuration.

The `cfgsave` command ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

An example follows.

```
switch:admin> cfgremove "NEW_cfg", "purplezone"
```

```
switch:admin> cfgsave
```

```
WARNING!!!
```

```
The changes you are attempting to save will render the Effective configuration and the Defined
configuration inconsistent. The inconsistency will result in different Effective Zoning configurations
for switches in the fabric if a zone merge or HA failover happens. To avoid inconsistency it is
recommended to commit the configurations using the 'cfgenable' command.
```

```
Do you still want to proceed with saving the Defined zoning configuration only? (yes, y, no, n): [no] y
```

## Enabling a Zone Configuration

The following procedure ends and commits the current zoning transaction buffer to the nonvolatile memory. If a transaction is open on a different switch in the fabric when this procedure is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

Use the following procedure to enable a zone configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgenable` command using the following syntax:

```
cfgenable "cfgname"
```

3. Enter `y` at the prompt.

```
switch:admin> cfgenable "USA_cfg"
```

```
You are about to enable a new zoning configuration.
```

```
This action will replace the old zoning configuration with the current configuration selected.
```

```
Do you want to enable 'USA_cfg' configuration (yes, y, no, n): [no] y
```

```
zone config "USA_cfg" is in effect
```

```
Updating flash ...
```

## Disabling a Zone Configuration

When you disable the current zone configuration, the fabric returns to non-zoning mode. All devices can then access each other or not, depending on the default zone access mode setting.

### NOTE

If the default zoning mode is set to All Access and more than 120 devices are connected to the fabric, you cannot disable the zone configuration because this would enable All Access mode and cause a large number of requests to the switch. In this situation, set the default zoning mode to No Access prior to disabling the zone configuration. Refer to [Default Zoning Mode](#) for information about setting this mode to No Access.

The following procedure ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this procedure is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

Use the following procedure to disable a zone configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgDisable` command.
3. Enter `y` at the prompt.

```
switch:admin> cfgdisable
```

```
You are about to disable zoning configuration. This action will disable any previous zoning configuration enabled.
```

```
Do you want to disable zoning configuration? (yes, y, no, n): [no] y
```

## Validating Zone Members in the Zone Configuration

Use this procedure to validate the zone members present in the zone configurations.

Validating the zone members helps to know the active and inactive devices that are logged into fabric. This procedure adds an indicator "\*" for the zone members that do not exist in the fabric, and "~" for the zone members with the invalid configuration. The zone server updates the details for zone members in the zone database that is based on the information that is received from the name server. Zone members that are defined using world-wide names are flagged as nonexistent if the specified device is not present within the fabric (offline). Zone members that are defined using domain and port index are flagged as nonexistent if either the specified port index is out of bounds, or if the specified domain is not a member of the fabric.

1. Connect to the switch and log in to an account with admin permissions.
2. Enter the `zonestow --validate` command to list the validated zone members in all the zone configurations of the zone database.

Use the pattern option with `zonestow --validate` to validate a specified zone configuration. The pattern must contain the exact zone configuration name or a portion with the wildcard character (\*) at the end to support a range of characters. This operand is optional.

Use the mode option with `zonestow --validate` to validate a specified mode of a zone configuration. This operand is optional. The following modes are supported:

- 0 – To display the contents of the transaction buffer
- 1 – To display the contents of the committed defined buffer
- 2 – To display the contents of the effective configuration

The default mode value is 0.

### Example of Validating All Zone Members

The following example displays the validated zone members of all the zone configurations of the zone database.

```
switch:admin> zonestow --validate
Defined configuration:
cfg: cfg1 zone1; zone10; zone2
zone: zone1 20:1c:00:05:1e:57:b1:c6*; 20:1d:00:05:1e:57:b1:c6
zone: zone10 20:1e:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
zone: zone2 20:03:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
```

```
Effective configuration:
cfg: cfg1
zone: zone1 20:1c:00:05:1e:57:b1:c6*
 20:1d:00:05:1e:57:b1:c6
zone: zone10 20:1e:00:05:1e:57:b1:c6
 20:1f:00:05:1e:57:b1:c6*
zone: zone2 20:03:00:05:1e:57:b1:c6
 20:1f:00:05:1e:57:b1:c6*

~ - Invalid configuration
* - Member does not exist
```

### Example of Validating Zone Members with a Zone Name

The following example displays the validated zone members that contain a specified or a part of the zone name.

```
switch:admin> zoneshow --validate zone1
Defined configuration:
zone: zone1 20:1c:00:05:1e:57:b1:c6*; 20:1d:00:05:1e:57:b1:c6

Effective configuration:
zone: zone1 20:1c:00:05:1e:57:b1:c6*
 20:1d:00:05:1e:57:b1:c6

~ - Invalid configuration
* - Member does not exist
```

```
switch:admin> zoneshow --validate zone*
Defined configuration:
zone: zone1 20:1c:00:05:1e:57:b1:c6*; 20:1d:00:05:1e:57:b1:c6
zone: zone10 20:1e:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
zone: zone2 20:03:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
zone: zone200 20:1d:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*

Effective configuration:
zone: zone1 20:1c:00:05:1e:57:b1:c6*
 20:1d:00:05:1e:57:b1:c6
zone: zone10 20:1e:00:05:1e:57:b1:c6
 20:1f:00:05:1e:57:b1:c6*
zone: zone2 20:03:00:05:1e:57:b1:c6
 20:1f:00:05:1e:57:b1:c6*

~ - Invalid configuration
* - Member does not exist
```

### Example of Validating Zone Members with Mode

The following example displays the validated zone members for a specified mode of a zone configuration.

```
switch:admin> zoneshow --validate zone200, 0
Defined configuration:
zone: zone200 20:1d:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
```

```

~ - Invalid configuration
* - Member does not exist
```

## Deleting a Zone Configuration

Use the following procedure to delete a zone configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgDelete` command using the following syntax:

```
cfgdelete "cfgname"
```

3. Enter the `cfgSave` command to save the change to the defined configuration.

The `cfgSave` command ends and commits the current zoning transaction buffer to the nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message is displayed on the other switches to indicate that the transaction was aborted.

```
switch:admin> cfgdelete "testcfg"
```

```
switch:admin> cfgsave
```

```
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
```

```
Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no]
```

## Deleting Invalid Zones

Invalid zones can be created by a number of events, including firmware upgrades and downgrades.

Use the following steps to remove an invalid zone.

1. Log in to the switch using an account with administrator privileges.
2. Enter `cfgremove "<zone_configuration>", "<invalid_zone>"` to remove the invalid zone from the zone configuration.
3. Enter `zoneddelete "<invalid_zone>"` to delete the invalid zone from the fabric.
4. Enter `cfgsave` to save the changed configuration.
5. Enter `cfgenable "<zone_configuration>"` to reenables the zone configuration.

The following example removes the invalid zone "dead\_zone" from the "baseline" configuration, deletes the zone, then saves the configuration and reenables it.

```
switch:admin> cfgremove "baseline", "dead_zone"
switch:admin> zoneddelete dead_zone
switch:admin> cfgsave
switch:admin> cfgenable "baseline"
```

For detailed information on the zoning commands used in the procedures, refer to the *Brocade Fabric OS Command Reference Manual*.

## Abandoning Zone Configuration Changes

To abandon zone configuration changes, enter the `cfgTransAbort` command.

When this command is executed, all changes since the last save operation (performed with the `cfgSave`, `cfgEnable`, or `cfgDisable` command) are cleared.

Example of assuming that the removal of a member from zone1 was done in error.


```
switch:admin> zoneremove "zone1","3,5"
```

```
switch:admin> cfgtransabort
```


## Zone Fabric Locking

Prior to Fabric OS 9.0.0, the level of zoning transaction protection against multiple concurrent clients only existed at the local switch level. For example, if a user on a switch starts a zone transaction by performing a zone create or edit operation, and if another user tries to perform a zone create or edit or commit operation on the same switch, the second user will be locked out and prevented from performing his operation since the first user opened the transaction first.


**Figure 69: Single Switch Locking Example**



Although there is local switch protection, there is no fabric-wide locking mechanism. For versions prior to Fabric OS 9.0.0, the fabric-level detection support (Concurrent Transaction Detection) consists only of a warning mechanism that warns users if a transaction is open on a remote switch, but does not prevent users from committing changes if concurrent transactions exist across the fabric. Furthermore, this fabric-wide detection is limited in scope to CLI zoning operations only. For example, if a user on one switch opens a zone transaction, and if a user on another switch in the fabric issues a zone edit or commit command, a warning is received that a transaction is occurring in the fabric. However, a zone commit operation can still be issued from either switch, which overwrites any pending zone transaction on other switches.

**Figure 70: Fabric OS 8.2.x and Earlier Fabric-Wide Commit Overwrite Example**

From Fabric OS 9.0.0 version onward, the Zone Fabric Locking feature provides zone transaction protection for a single switch and also fabric wide. This feature is enabled by default. If a zone edit or commit command is occurring in fabric, you cannot perform a zone edit or commit on the same or another switch for a default timeout period of 5 minutes. A lock request is sent at the beginning of a zone edit operation. The Fabric Lock Failsafe Timer is configurable and it is a fabric-wide setting. When a zone fabric lock is active, a failsafe timer is started on all remote switches.

**Figure 71: Fabric OS 9.0.0 Zone Fabric Locking Example**

When the failsafe timer expires, the open zone transaction is not aborted. If the same user attempts to resume the transaction by performing another edit or commit operation after the zone fabric lock has expired, the transaction is allowed and the fabric lock is restarted. If a different user attempts to start a new transaction after the first user's transaction timer has expired, the transaction is allowed and the first user's transaction is aborted before the second user's transaction starts.

**NOTE**

- If a fabric-lock is currently active in the fabric, the Fabric Lock Failsafe Timer timeout value cannot be modified.
- The fabric zone lock is a fabric-wide setting. Thus, the fabric must be stable before performing zone edit, commit, and abort operations.
- If a zone transaction is aborted, the fabric zone lock is cleared.
- When the timer expires, the lock is cleared locally and on each individual switch in the fabric.
- When using virtual fabric functionality, the lock timeout value is independent on each logical switch and can have different values.
- Fabric zone locks can be preempted by zone merges.
- The Zone Fabric Locking feature is supported on the switches with the Fabric OS version 9.0.0 or later. Down-level switches will not participate. However, Fabric OS 9.0.0 switches owning the zone fabric lock will prevent updates from down-level switches.
- On remote switches, an additional 30-second buffer is added to the configured lock timeout value of the lock principal switch.

The fabric-wide locking mechanism ensures that any zone transaction started on any switch by any interface (CLI, REST, Web Tools, SANnav, Target Driven Zone) is not overwritten or aborted by another interface's transaction (except for zone merges) while the lock timer is active.

The following are examples to configure the zone fabric lock failsafe timer.

An example of a successful operation:

```
switch:admin> zoneFabricLock --timeout 15
switch:admin>
```

An example of a failed operation:

```
switch:admin> zoneFabricLock --timeout 10
Error: Zone Fabric Lock Timeout Update failed, can't change timer while Zone Fabric Lock is active (domain 5 owns lock).
```

An example of displaying the locally-configured Zone Fabric Lock timeout value :

```
switch:admin> configshow -pattern zoning.FabLockTimeout
zoning.FabLockTimeout:5
```

An example to view zoneFabricLock details:

```
switch:admin> zoneFabricLock --show
Fabric-Lock State: Active
Lock-Principal Domain: 86
Client: admin/admin/xx.xx.xx.xx/telnet/CLI
Timeout: 30 mins
Time Remaining: 29 mins 47 secs
Fabric Supported: Yes
```

An example of a non-lock principal switch in a fabric with an active lock:

```
switch:admin> cfgtransshow
There is an outstanding zoning transaction in progress on Domain 2

Fabric-Lock State: Active
Time Remaining: 19 mins 22 secs
```

```
Client: admin/admin//console/CLI
Fabric Supported: No
```

### An example of a lock principal switch:

```
switch:admin> cfgtransshow
Current transaction token is 0x142a
It is abortable

Fabric-Lock State: Active
Time Remaining: 18 mins 25 secs
Client: admin/admin//console/CLI
Fabric Supported: No
```

### An example of a switch in a fabric without an active lock:

```
switch:admin> cfgtransshow --fabricLock
Fabric-Lock State: NOT Active
Lock-Principal Domain: n/a
Lock-Principal Transaction Token: n/a
Client: n/a
Timeout: n/a
Time Remaining: n/a
Fabric Supported: Yes
```

## Viewing All Zone Configuration Information

If you do not specify an operand when executing the `cfgShow` command to view zone configurations, then all zone configuration information (both defined and effective) displays. If there is an outstanding transaction, then the newly edited zone configuration that has not yet been saved is displayed. If there are no outstanding transactions, the committed zone configuration is displayed.

Use the following procedure to view all zone configuration information.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgShow` command with no operands.  
An example follows.

```
switch:admin> cfgshow

Defined configuration:
cfg: USA1 Blue_zone
cfg: USA_cfg Purple_zone; Blue_zone
zone: Blue_zone
 1,1; array1; 1,2; array2
zone: Purple_zone
 1,0; loop1
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df

Effective configuration:
cfg: USA_cfg
zone: Blue_zone
 1,1
```



```

21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2
21:00:00:20:37:0c:76:22
21:00:00:20:37:0c:76:28
zone: Purple_zone
1,0
21:00:00:20:37:0c:76:85
21:00:00:20:37:0c:71:df

```

## Viewing Selected Zone Configuration Information

Use the following procedure to view the selected zone configuration information.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgshow` command and specify a pattern.

```
cfgshow [--ic][," pattern"] [, mode]
```

Example displaying all zone configurations that start with "Test".

```

switch:admin> cfgshow "Test*"

cfg: Test1 Blue_zone
cfg: Test_cfg Purple_zone; Blue_zone

```

Example displaying all zone configurations that start with "Test", regardless of the case.

```

switch:admin> cfgshow --ic "Test*"

cfg: Test1 Blue_zone
cfg: Test_2 Red zone; Blue_zone

```

## Viewing the Zone Aliases in the Zone Configuration

The following procedure lists the zone aliases, including the user-defined aliases, in the zone configuration.

1. Log in to the switch.
2. View the zone aliases using the `zonestow --alias [-ic] [pattern]` command.

The `zonestow --alias` command displays all the zones containing aliases matching the pattern. You can use POSIX style regular expression for the alias pattern. The pattern can contain the following:

- A question mark (?) to match any single character.
- An asterisk (\*) to match any string of characters.
- A range of characters to match any character within the range: for example, [0 to 9] or [a to f].

The following example demonstrates using the command to display all the zones containing the alias specified when the complete pattern is specified in the command (case sensitive).

```

switch:admin> zonestow --alias ali1
Defined configuration:
zone: zone1 ali1
zone: zone3 2,2; ali1

```

The following examples demonstrate using the command to display all the zones containing the aliases matching the POSIX-style regular expression specified in each pattern (case sensitive).

```
switch:admin> zoneshow --alias ali1*
Defined configuration:
zone: zone1 ali1
zone: zone3 2,2; ali1
zone: zone4 21,22; ali12

switch:admin> zoneshow --alias "a*"
zone: Zone3 ald1
zone: aliii aldds
zone: h1 al; Ali1; dev3
zone: zone1 ali1
zone: zone4 dev2; aldds

switch:admin> zoneshow --alias "**a*"
zone: Zone3 ald1
zone: aliii aldds
zone: h1 al; Ali1; dev3
zone: zone1 ali1
zone: zone4 dev2; aldds

switch:admin> zoneshow --alias "*1*"
zone: Zone3 ald1
zone: h1 al; Ali1; dev3
zone: zone1 ali1
zone: zone2 Ali1

switch:admin> zoneshow --alias "?l?1"
zone: Zone3 ald1
zone: h1 al; Ali1; dev3
zone: zone1 ali1
zone: zone2 Ali1

switch:admin> zoneshow --alias "al[id]1"
zone: Zone3 ald1
zone: zone1 ali1
```

You can also display a list of zones containing the alias matching the pattern that you specify (not case sensitive).

```
switch:admin> zoneshow
Defined configuration:
cfg: cfg3 Zone2
zone: Zone1 dev1; dev2
zone: Zone2 dev2; DEV3

alias: DEV3 1,8
alias: dev1 1,2
alias: dev2 1,1

Effective configuration:
cfg: cfg3
zone: Zone2 1,1
 1,8

switch:admin> zoneshow --alias -ic dev3

zone: Zone2 dev2; DEV3
```

## Viewing the Configuration in the Effective Zone Database

Use the following procedure to view the configuration in the effective zone database.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgActvShow` command.

```
switch:admin> cfgactvshow
```

```
Effective configuration:
```

```
cfg: NEW_cfg
zone: Blue_zone
1,1
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2
21:00:00:20:37:0c:76:22
21:00:00:20:37:0c:76:28
zone: Purple_zone
1,0
21:00:00:20:37:0c:76:85
21:00:00:20:37:0c:71:df
```

## Clearing All Zone Configurations

Use the following procedure to clear all zone configurations.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgclear [-force/-f]` command to clear all zone information in the transaction buffer.

### ATTENTION

Be careful using the `cfgclear [-force/-f]` command because it deletes the defined configuration.

```
switch:admin> cfgclear
```

```
The Clear All action will clear all aliases, zones, and configurations in the Defined configuration.
Run cfgsave to commit the transaction or cfgtransabort to cancel the transaction.
Do you really want to clear all configurations? (yes, y, no, n): [no]
```

### ATTENTION

The `cfgclear -force` command bypasses the user interaction step and deletes the defined configuration.

```
switch:admin> cfgclear -force
```

```
The Clear All action performed has cleared all aliases, zones, and configurations in the Defined
configuration.
Run cfgsave to commit the transaction or cfgtransabort to cancel the transaction.
```

3. Enter one of the following commands, depending on whether an effective zone configuration exists.

- If no effective zone configuration exists, use the `cfgsave` command.
- If an effective zone configuration exists, use the `cfgdisable` command to disable and clear the zone configuration in nonvolatile memory for all switches in the fabric.

#### NOTE

If the default zone mode is set to `AllAccess` and there are more than 120 devices logged in to the fabric, the `cfgdisable` operation is not allowed. For detailed information, refer to the section [Disabling a Zone Configuration](#).

## Zone Object Group Member Sorting

The members of zone object groups (aliases, zones, and configurations) are sorted in ascending order for zone commit and creation operations. If the members of the zone object groups in a fabric are not listed in the same order (for example, CFG1 in one fabric is Z1, Z3, Z2, and CFG1 in another fabric is Z1, Z2, Z3), the zone members of the zone groups are sorted when you use the following CLI commands:

- `cfgsave`, `cfgdisable`, `cfgenable`
- `zoneadd`, `cfgadd`, `aliadd`, `configdownload`
- `zonecreate`, `cfgcreate`, `alicreate`
- `zoneobjectrename`
- `zoneobjectreplace`

The zone group object member sorting feature ensures that two switches having 9.0.0 and the same configuration objects that are created or updated locally with members in any order merge successfully.

Example of sorting zone object group members.

```
switch:admin> cfgshow
Defined configuration:
cfg: sample_cfg
 matt; sloth
zone: matt 30:06:00:07:1e:a2:10:20; 3,2
zone: sloth 10:00:00:00:01:1e:20:20; bawn; bolt; bond; brain
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; 40:00:00:05:1e:a1:cd:04
alias: jones 4,5; 7,3
alias: zeus 4,7; 6,8; 9,2
```

```
Effective configuration:
cfg: sample_cfg
zone: matt 30:06:00:07:1e:a2:10:20
 3,2
zone: sloth 10:00:00:00:01:1e:20:20
 3,5
 4,8
 10:00:00:02:1f:02:00:01
 10:00:05:1e:a9:20:00:01
 11,4
 22,1
 33,6
```

```
switch:admin>
```

**NOTE**

If a zone group has mixed-member types, the members are sorted in WWN, Domain/Port, and alias order respectively.

**NOTE**

- The zone object group member sorting feature is not applicable for the members of Boot LUN zones.
- Peer zone members are considered for sorting. The principal and non-principal members are sorted separately.
- If a zone object group is created in a pre-9.0.0 switch, and the configuration is pushed to a 9.0.0 switch through merge or upgrade, the members in the zone group object are not sorted.
- Fabric segmentation due to a sort-order conflict may still occur.
- A switch that is upgraded to firmware version 9.0.0 or later from a previous version does not sort the existing unsorted configuration of the pre-9.0.0 switch.

## Zone Object Maintenance

The following procedures describe how to copy, delete, and rename zone objects. Depending on the operation, a zone object can be a zone member, a zone alias, a zone, or a zone configuration.

### Copying a Zone Object

When you copy a zone object, the resulting object has the same name as the original. The zone object can be a zone configuration, a zone alias, or a zone.

Use the following procedure to copy a zone object.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgShow` command to view the zone configuration objects you want to copy.

```
cfgshow ["pattern"] [, mode]
```

For example, to display all zone configuration objects that start with "Test":

```
switch:admin> cfgshow "Test*"
cfg: Test1 Blue_zone
cfg: Test_cfg Purple_zone; Blue_zone
```

3. Enter `zoneObjectCopy` specifying the zone objects you want to copy along with the new object name.

**NOTE**

Zone configuration names are case-sensitive, and blank spaces are ignored.

```
switch:admin> zoneObjectCopy "Test1", "US_Test1"
```

4. Enter the `cfgShow` command to verify the new zone object is present.

```
switch:admin> cfgshow "Test*"
cfg: Test1 Blue_zone
cfg: Test_cfg Purple_zone; Blue_zone
switch:admin> cfgShow "US_Test1"
```

```
cfg: US_Test1 Blue_zone
```

5. If you want the change preserved when the switch reboots, use `cfgSave` to save it to nonvolatile (flash) memory.
6. Enter `cfgEnable` for the appropriate zone configuration to make the change effective.

## Deleting a Zone Object

The following procedure removes all references to a zone object and then deletes the zone object. The zone object can be a zone member, a zone alias, or a zone.

Use the following procedure to delete a zone object.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `cfgshow` command to view the zone configuration objects you want to delete.

```
switch:admin> cfgshow
Defined configuration:
 cfg: USA_cfg Purple_zone; White_zone; Blue_zone
 zone: Blue_zone
 1,1; array1; 1,2; array2
 zone: Purple_zone
 1,0; loop1
 zone: White_zone
 1,3; 1,4
 alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
 alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
 alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
Effective configuration:
 cfg: USA_cfg
 zone: Blue_zone
 1,1
 21:00:00:20:37:0c:76:8c
 21:00:00:20:37:0c:71:02
 1,2
 21:00:00:20:37:0c:76:22
 21:00:00:20:37:0c:76:28
 zone: Purple_zone
 1,0
 21:00:00:20:37:0c:76:85
 21:00:00:20:37:0c:71:df
```

3. Enter `zoneobjectexpunge` to delete the zone object.

### NOTE

Zone configuration names are case-sensitive, and blank spaces are ignored.

```
switch:admin> zoneobjectexpunge "White_zone" -force
You have expunged a configuration or member. This action could result in removing many zoning
configurations recursively. [Removing the last member of a configuration removes the configuration.]
```

### ATTENTION

The `-force` option bypasses the user interaction step and expunges the specified zone object.

4. Enter `cfgshow` to verify that the deleted zone object is no longer present.
5. Use `cfgsave` to save it to the nonvolatile (flash) memory.
6. Enter `cfgenable` for the appropriate zone configuration to make the change effective.

## Renaming a Zone Object

Use the following procedure to rename a zone object.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter `cfgShow` to view the zone configuration objects you want to rename.

```
switch:admin> cfgShow
Defined configuration:
 cfg: USA_cfg Purple_zone; White_zone; Blue_zone
 zone: Blue_zone
 1,1; array1; 1,2; array2
 zone: Purple_zone
 1,0; loop1
 zone: White_zone
 1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Enter `zoneObjectRename` to rename zone configuration objects.

### NOTE

Zone configuration names are case-sensitive and blank spaces are ignored.

```
switch:admin> zoneObjectRename "White_zone", "Purple_zone"
```

4. Enter the `cfgShow` command to verify the renamed zone object is present.
5. If you want the change preserved when the switch reboots, enter the `cfgSave` command to save it to non-volatile (flash) memory.

### NOTE

If the renamed zone object is still used in the effective configuration, running `cfgSave` before `cfgEnable` is not allowed.

6. Enter the `cfgEnable` command for the appropriate zone configuration to make the change effective.

## Zone Configuration Management

You can add, delete, or remove individual elements in an existing zone configuration to create an appropriate configuration for your SAN environment. After the changes have been made, save the configuration to ensure the configuration is permanently saved in the switch and that the configuration is replicated throughout the fabric.

The switch configuration file can also be uploaded to the host for archiving and it can be downloaded from the host to a switch in the fabric. Refer to [Configuration file backup](#), [Configuration file restoration](#), or the `configUpload` and `configDownload` commands in the *Brocade Fabric OS Command Reference Manual* for additional information on uploading and downloading the configuration file.

## Security and Zoning

Zones provide controlled access to fabric segments and establish barriers between operating environments. They isolate systems with different uses, protecting individual systems in a heterogeneous environment; for example, when zoning is in secure mode, no merge operations occur.

Brocade Advanced Zoning is configured on the primary fabric configuration server. The primary FCS switch makes zoning changes and other security-related changes. The primary FCS switch also distributes zoning to all other switches in the secure fabric. All existing interfaces can be used to administer zoning.

You must perform zone management operations from the primary FCS switch using a zone management interface, such as telnet or Web Tools. You can alter a zone database, provided you are connected to the primary FCS switch.

When two secure fabrics join, the traditional zone merge does not occur. Instead, a zone database is downloaded from the primary FCS switch of the merged secure fabric. When E\_Ports are active between two switches, the name of the FCS server and a zoning policy set version identifier are exchanged between the switches. If the views of the two secure fabrics are the same, the fabric's primary FCS server downloads the zone database and security policy sets to each switch in the fabric. If a view conflict exists, the E\_Ports are segmented due to incompatible security data.

All zones should use frame-based hardware enforcement; the best way to do this is to use WWN identification exclusively for all zoning configurations.

## Zone Merging

When a new switch is added to the fabric, it automatically takes on the zone configuration information from the fabric. You can verify the zone configuration on the switch using the procedure that is described in [Viewing the Configuration in the Effective Zone Database](#).

If you are adding a switch that is already configured for zoning, clear the zone configuration on that switch before connecting it to the zoned fabric. Refer to [Clearing All Zone Configurations](#) for instructions.

Adding a new fabric that has no zone configuration information to an existing fabric is similar to adding a new switch. All switches in the new fabric inherit the zone configuration data. If the existing fabric has an effective zone configuration, then the same configuration becomes the effective configuration for the new switches.

### NOTE

From Fabric OS 8.2.0 or later, the admin domain configurations are blocked in zone merging. The ISLs are segmented and admin domain configurations are not pushed into the switches having Fabric OS 8.2.0 or later.

Before the new fabric can merge successfully, it must pass the following criteria:

- Before merging – To facilitate merging, check the following points before merging switches or fabrics:
  - Default zone: The switches must adhere to the default zone merge rules, as described in [Zone Merging Scenarios](#).
  - Effective and defined zone configuration match: Make sure the effective and defined zone configurations match. If they do not match, and you merge with another switch, the merge may be successful but unpredictable zoning and routing behavior can occur.
- Merging and segmentation

When a switch is disabled or enabled, or when a new switch is added, the fabric is checked for segmentation during power-up.

The zone configuration database is stored in a non-volatile memory by the `cfgSave` command. All switches in the fabric have a copy of this database. When a change is made to the defined configuration, the switch where the changes were made must close its transaction for the changes to be propagated throughout the fabric.

If you are implementing default zoning, you must set the switch in the fabric to the same default zone mode setting to avoid the segmentation.
- Merging rules – Observe these rules when merging zones:



- Local and adjacent configurations: If the local and adjacent zone database configurations are the same, they will remain unchanged after the merge.
  - Effective configurations: If there is an effective configuration between two switches, the effective zone configurations must match.
  - Zone object naming: If a zoning object has the same name in both the local and adjacent defined configurations, the object types and member lists must match. When comparing member lists, the content and order of the members are important.
  - Objects in adjacent configurations: If a zoning object appears in an adjacent defined configuration, but not in the local defined configuration, you must add the zoning object to the local defined configuration. The modified zone database must fit in the nonvolatile memory area that is allotted for the zone database.
  - Local configuration modification: If a local defined configuration is modified because of a merge, the new zone database is propagated to other switches within the merge request.
- Merging two fabrics  
Both fabrics have identical zones and configurations that are enabled, including the default zone mode. The two fabrics are joined to make one larger fabric with the same zone configuration across the newly created fabric.
  - Merge conflicts  
When a merge conflict is present, a merge does not take place and the ISL will segment. Use the `switchShow` or `errDump` commands to obtain additional information about possible merge conflicts because many non-zone-related configuration parameters can cause conflicts. Refer to the *Brocade Fabric OS Command Reference Manual* for detailed information about these commands.  
If the fabrics have different zone configuration data, the system attempts to merge the two sets of zone configuration data. If the zones cannot merge, the ISL is segmented.  
If any of the following conditions exist, a merge is not possible:
    - Configuration mismatch: Zoning is enabled in both fabrics, and the zone configurations that are enabled are different in each fabric.
    - Type mismatch: The name of a zone object in one fabric is used for a different type of zone object in the other fabric.
    - Content mismatch: The definition of a zone object in one fabric is different from the definition of the zone object with the same name in the other fabric.
    - Zone database size: The zone database size exceeds the maximum limit of another switch.

**NOTE**

In a large fabric, especially with 1 MB or more, zone configuration takes some amount of time for zone merge, which may cause the host device to not discover the target in the other end of the fabric for a short duration. This problem may not be seen if you follow the guidelines of joining the stable fabric one-by-one with devices offline and ensuring that the zone merge operation is completed before enabling the devices.

**NOTE**

If the zone set members on two switches are not listed in the same order, the configuration is considered a mismatch, resulting in the switches being segmented from the fabric. For example, `cfg1 = z1; z2` is different from `cfg1 = z2; z1`, even though members of the configuration are the same. If zone set members on two switches have the same names that are defined in the configuration, make sure zone set members are listed in the same order.

## Fabric Segmentation and Zoning

If the connections between two fabrics are no longer available, the fabric segments into two separate fabrics. Each new fabric retains the same zone configuration.

If the connections between two fabrics are replaced and no changes have been made to the zone configuration in either of the two fabrics, the two fabrics merge back into one single fabric. If any changes that cause a conflict have been made to either zone configuration, the fabrics may segment.

## Zone Merging Scenarios

The following tables provide information on merging zones and the expected results:

- Defined and effective configurations
- Different content
- Different names
- Default access mode
- Mixed Fabric OS versions

**Table 133: Zone Merging Scenarios: Defined and Effective Configurations**

| Description                                                                                                                               | Switch A                                            | Switch B                                            | Expected Results                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Switch A</b> has a defined configuration.<br><b>Switch B</b> does not have a defined configuration.                                    | defined: cfg1: zone1: ali1; ali2<br>effective: none | defined: none<br>effective: none                    | Configuration from <b>Switch A</b> to propagate throughout the fabric in an inactive state, because the configuration is not enabled. |
| <b>Switch A</b> has a defined and effective configuration.<br><b>Switch B</b> has a defined configuration but no effective configuration. | defined: cfg1 zone1: ali1; ali2<br>effective: cfg1: | defined: cfg1 zone1: ali1; ali2<br>effective: none  | Configuration from <b>Switch A</b> to propagate throughout the fabric. The configuration is enabled after the merge in the fabric.    |
| <b>Switch A</b> and <b>Switch B</b> have the same defined configuration. Neither have an effective configuration.                         | defined: cfg1 zone1: ali1; ali2<br>effective: none  | defined: cfg1 zone1: ali1; ali2<br>effective: none  | No change (clean merge).                                                                                                              |
| <b>Switch A</b> and <b>Switch B</b> have the same defined and effective configuration.                                                    | defined: cfg1 zone1: ali1; ali2<br>effective: cfg1: | defined: cfg1 zone1: ali1; ali2<br>effective: cfg1: | No change (clean merge).                                                                                                              |
| <b>Switch A</b> does not have a defined configuration.<br><b>Switch B</b> has a defined configuration.                                    | defined: none<br>effective: none                    | defined: cfg1 zone1: ali1; ali2<br>effective: none  | <b>Switch A</b> will absorb the configuration from the fabric.                                                                        |
| <b>Switch A</b> does not have a defined configuration.<br><b>Switch B</b> has a defined configuration.                                    | defined: none<br>effective: none                    | defined: cfg1 zone1: ali1; ali2<br>effective: cfg1  | <b>Switch A</b> will absorb the configuration from the fabric, with cfg1 as the effective configuration.                              |
| <b>Switch A</b> and <b>Switch B</b> have the same defined configuration. Only <b>Switch B</b> has an effective configuration.             | defined: cfg1 zone1: ali1; ali2<br>effective: none  | defined: cfg1 zone1: ali1; ali2<br>effective: cfg1  | Clean merge, with cfg1 as the effective configuration.                                                                                |
| <b>Switch A</b> and <b>Switch B</b> have different defined configurations. Neither have an enabled zone configuration.                    | defined: cfg2 zone2: ali3; ali4<br>effective: none  | defined: cfg1 zone1: ali1; ali2<br>effective: none  | Clean merge. The new configuration is a composite of the two defined: cfg1 zone1: ali1; ali2 cfg2: zone2: ali3; ali4 effective: none  |
| <b>Switch A</b> and <b>Switch B</b> have different defined configurations. <b>Switch B</b> has an effective configuration.                | defined: cfg2 zone2: ali3; ali4<br>effective: none  | defined: cfg1 zone1: ali1; ali2<br>effective: cfg1  | Clean merge. The new configuration will be a composite of the two, with cfg1 as the effective configuration.                          |

| Description                                                                                                                                                                                                        | Switch A                         | Switch B                                                                                        | Expected Results                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Switch A</b> does not have a defined configuration.<br><b>Switch B</b> has a defined configuration and an effective configuration, but the effective configuration is different from the defined configuration. | defined: none<br>effective: none | defined: cfg1<br>zone1: ali1; ali2<br>effective: cfg1<br>zone1: ali1; ali2<br>zone2: ali3, ali4 | Clean merge. <b>Switch A</b> absorbs the defined configuration from the fabric, with cfg1 as the effective configuration.<br><br>In this case, however, the effective configurations for <b>Switch A</b> and <b>Switch B</b> are different. Issue a <code>cfgenable</code> command from the switch with the proper effective configuration. |

**Table 134: Zone Merging Scenarios: Different Content**

| Description                       | Switch A                                                             | Switch B                                                             | Expected Results                                        |
|-----------------------------------|----------------------------------------------------------------------|----------------------------------------------------------------------|---------------------------------------------------------|
| Effective configuration mismatch. | defined: cfg1 zone1: ali1; ali2<br>effective: cfg1 zone1: ali1; ali2 | defined: cfg2 zone2: ali3; ali4<br>effective: cfg2 zone2: ali3; ali4 | Fabric segments due to: Zone conflict cfg mismatch.     |
| Configuration content mismatch.   | defined: cfg1 zone1: ali1; ali2<br>effective: irrelevant             | defined: cfg1 zone1: ali3; ali4<br>effective: irrelevant             | Fabric segments due to: Zone conflict content mismatch. |

**Table 135: Zone Merging Scenarios: Different Names**

| Description                                     | Switch A                                                             | Switch B                                                             | Expected Results                                        |
|-------------------------------------------------|----------------------------------------------------------------------|----------------------------------------------------------------------|---------------------------------------------------------|
| Same content, different effective cfg name.     | defined: cfg1 zone1: ali1; ali2<br>effective: cfg1 zone1: ali1; ali2 | defined: cfg2 zone1: ali1; ali2<br>effective: cfg2 zone1: ali1; ali2 | Fabric segments due to: Zone conflict cfg mismatch.     |
| Same content, different zone name.              | defined: cfg1 zone1: ali1; ali2<br>effective: irrelevant             | defined: cfg1 zone2: ali1; ali2<br>effective: irrelevant             | Fabric segments due to: Zone conflict content mismatch. |
| Same content, different alias name.             | defined: cfg1 ali1: A; B effective: irrelevant                       | defined: cfg1 ali2: A; B<br>effective: irrelevant                    | Fabric segments due to: Zone conflict content mismatch. |
| Same alias name, same content, different order. | defined: cfg1 ali1: A; B; C<br>effective: irrelevant                 | defined: cfg1 ali1: B; C; A<br>effective: irrelevant                 | Fabric segments due to: Zone conflict content mismatch. |
| Same name, different types.                     | effective: zone1: MARKETING                                          | effective: cfg1: MARKETING                                           | Fabric segments due to: Zone conflict type mismatch.    |
| Same name, different types.                     | effective: zone1: MARKETING                                          | effective: alias1: MARKETING                                         | Fabric segments due to: Zone conflict type mismatch.    |
| Same name, different types.                     | effective: cfg1: MARKETING                                           | effective: alias1: MARKETING                                         | Fabric segments due to: Zone conflict type mismatch.    |

**Table 136: Zone Merging Scenarios: Default Access Mode**

| Description                                  | Switch A (Initiator) | Switch B (Responder) | Expected Results                                               |
|----------------------------------------------|----------------------|----------------------|----------------------------------------------------------------|
| Different default zone access mode settings. | defzone: allaccess   | defzone: noaccess    | Fabric segments due to zone conflict.                          |
|                                              | defzone: noaccess    | defzone: allaccess   |                                                                |
| Same default zone access mode settings.      | defzone: allaccess   | defzone: allaccess   | Clean merge. defzone configuration is allaccess in the fabric. |

| Description                                 | Switch A (Initiator)                              | Switch B (Responder)                              | Expected Results                                                                                                                                                |
|---------------------------------------------|---------------------------------------------------|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                             | defzone: noaccess                                 | defzone: noaccess                                 | Clean merge. defzone configuration is noaccess in the fabric.                                                                                                   |
| Initiator has effective zone configuration. | effective: cfg1<br>defzone: allaccess             | No effective configuration.<br>defzone: allaccess | Clean merge. effective zone configuration from <b>Switch A</b> propagates to fabric.                                                                            |
|                                             | effective: cfg1<br>defzone: allaccess             | No effective configuration.<br>defzone: noaccess  | Fabric segments because <b>Switch B</b> has a hidden zone configuration (no access) activated and <b>Switch A</b> has an explicit zone configuration activated. |
|                                             | effective: cfg1<br>defzone: noaccess              | No effective configuration.<br>defzone: allaccess | Fabric merges --effective zone configuration from <b>Switch A</b> propagates to fabric. Allaccess on <b>Switch B</b> changes to noaccess.                       |
|                                             | effective: cfg1<br>defzone: noaccess              | No effective configuration.<br>defzone: noaccess  | Fabric merges --effective zone configuration from <b>Switch A</b> propagates to fabric.                                                                         |
| Responder has effective zone configuration. | No effective configuration.<br>defzone: allaccess | effective: cfg2<br>defzone: allaccess             | Fabric merges --effective zone configuration from <b>Switch B</b> propagates to fabric.                                                                         |
|                                             | No effective configuration.<br>defzone: allaccess | effective: cfg2<br>defzone: noaccess              | Fabric merges --effective zone configuration from <b>Switch B</b> propagates to fabric. Allaccess on <b>Switch A</b> changes to noaccess.                       |
|                                             | No effective configuration.<br>defzone: noaccess  | effective: cfg2<br>defzone: allaccess             | Fabric segments because <b>Switch A</b> has a hidden zone configuration (no access) activated and <b>Switch B</b> has an explicit zone configuration activated. |
|                                             | No effective configuration.<br>defzone: noaccess  | effective: cfg2<br>defzone: noaccess              | Fabric merges --effective zone configuration from <b>Switch B</b> propagates to fabric.                                                                         |

## Concurrent Zone Transactions

### NOTE

The Concurrent Zone Transactions feature is supported only on the CLI, and it is mutually exclusive with the Zone Fabric Locking feature.

When updates to the zoning database are being made by multiple users, Fabric OS warns you about the situation to choose which operation prevails.

### NOTE

The concurrent zone transactions are tracked only for Fabric OS CLI users. Other users using the REST API, Web Tools, SANnav, or Brocade Network Advisor or other methods are not tracked.

Here is an example of how users are warned if there is already a pending zoning transaction in the fabric.

```
u30:FID128:admin> zonecreate z2, "2,3"
```

WARNING!!

Multiple open transactions are pending in this fabric. Only one transaction can be saved. Please abort all unwanted transactions using the `cfgtransabort` command. Use the `cfgtransshow --opentrans` command to display a list of domains with open transactions

If no other transaction is open in this fabric, no message is displayed.

```
sw0:FID128:admin> zonecreate z7, "4,5;10,3"
```

```
sw0:FID128:admin>
```

Similar messages are shown for `cfgSave` and `cfgEnable` .

```
u30:FID128:admin> cfgenable cfg
```

You are about to enable a new zoning configuration.  
This action will replace the old zoning configuration with the current configuration selected.  
Multiple open transactions are pending in this fabric. Only one transaction can be saved. Please abort all unwanted transactions using the `cfgtransabort` command. Use the `cfgtransshow --opentrans` command to display a list of domains with open transactions  
Do you want to enable 'cfg' configuration (yes, y, no, n): [no]  
]  
u30:FID128:admin> cfgsave

You are about to save the Defined zoning configuration.  
This action will only save the changes on Defined configuration.  
Multiple open transactions are pending in this fabric. Only one transaction can be saved. Please abort all unwanted transactions using the `cfgtransabort` command. Use the `cfgtransshow --opentrans` command to display a list of domains with open transactions  
Do you want to save the Defined zoning configuration only?  
(yes, y, no, n): [no] n

## Viewing Zone Database Transactions

Enter `cfgtransshow` to list the details of the current database transaction token.

Enter `cfgtransshow --opentrans` to display the local open transaction token details and the list of domains with open transactions. Refer to the *Brocade Fabric OS Command Reference Manual* for more details on this command. With the addition of the Zone Fabric Locking feature in Fabric OS 9.0.0, the `--opentrans` example is applicable only in cases where there is an open zone transaction on a pre-9.0.0 down level switch in the fabric, or if there is an expired zone fabric lock on a 9.0.0 switch having an open zone transaction.

The following example shows the output of the `cfgtransshow` command.

```
switch:admin> cfgtransshow
```

```
Current transaction token is 0x571010459
It is abortable
```

The following example shows the output with the `--opentrans` option in the case of an active zone fabric lock on the local switch and an open zone transaction on a downlevel remote domain 3.

```
sw0:FID128:admin> cfgtransshow --opentrans

Current transaction token is 0x5184
It is abortable
Fabric-Lock State: Active
Time Remaining: 4 mins 57 secs
Client: admin/admin/none/console/CLI
Fabric Supported: No
Transactions Detect: Capable
Current Open Transactions Domain List:

3 4
sw0:FID128:admin>
```

## Peer Zoning

Peer zoning allows a *principal* device to communicate with the rest of the devices in the zone. The principal device manages a peer zone. Other *nonprincipal* devices in the zone can communicate with the principal device only; they cannot communicate with each other.

Prior to the introduction of peer zoning, single-initiator zoning was considered a more efficient zoning method in terms of hardware resources and RSCN volume. However, the added storage requirements of defining a unique zone for each host and target rapidly exceed zone database size limits. As the number of zones increases, it is difficult to configure and maintain the zones.


One-to-many zoning defines a zone having one target and other members as initiators. This approach is easy to manage and uses less storage, but zoning all the initiators together in this manner results in the less effective use of hardware resources and greater RSCN traffic.

In a peer zone setup, principal-to-nonprincipal device communication is allowed, but nonprincipal-to-nonprincipal device communication and principal-to-principal device communication is not allowed. This approach establishes zoning connections that provide the efficiency of single-initiator zoning with the simplicity and lower memory characteristics of one-to-many zoning. Typically, a peer zone has a single principal device and one or more nonprincipal devices, but configurations having multiple principal devices are allowed. Peer zones are not mutually exclusive with traditional zones; multiple zoning styles can coexist within the same zoning configuration and fabric.

## Peer Zoning Compared to Other Zoning Types


The following types of zoning methods are used in a SAN:

**Figure 72: Single-Initiator Zoning Example**




**Single-initiator zoning:** Single-initiator zoning eliminates host-to-host visibility, which reduces RSCN traffic. However, the main limitation of single-initiator zoning is that it results in creating a large number of zones, which can lead to exceeding the capacity of the zone database size limits. The Single-Initiator Zoning Example is an example of single-initiator zoning with eight zones; one for each host H1 to H8.


**Figure 73: One-to-Many Zoning Example**



**One-to-many Zoning:** One-to-many zoning creates a zone with multiple hosts and one target. This approach creates fewer zones as compared to single-initiator zoning, but the host-to-host connectivity generates an extraneous amount of RSCN traffic. The one-to-many zoning example is an example of one-to-many zoning where the hosts H1 to H8 are part of the same zone and can communicate with each other.

**Figure 74: Flat Zoning Example**

**Flat Zoning:** Flat zoning consists of a single zone that allows all devices to communicate with each other. This type of zoning supports a large number of hosts. One drawback of this type of zoning is the significant increase in RSCN traffic when compared to other zoning types. The flat zoning example is an example of flat zoning where the hosts H1 to H16 and target T1 and T2, all form a single zone and can communicate with all other devices in the zone.

**Figure 75: Peer Zoning Example**

**Peer Zoning:** Using peer zoning, a peer zone can be created with one device designated as a principal device for that zone. All nonprincipal devices in the peer zone can access only the principal device and cannot communicate with each other. The principal device can communicate with all other nonprincipal devices. The RSCN traffic generated is the same.



as for single-initiator zoning. The zone database size is the same as one-to-many zoning. Peer zoning results in less RSCN traffic on zoning and device changes and it creates fewer zones. The peer zoning example is an example where T1 is the principal member of the peer zone and hosts H1 to H8 are peer devices.

## Advantages of Peer Zoning

Peer zoning has the following advantages when compared to single-initiator zoning and one-to-many zoning:

- The zone database memory usage is the same as one-to-many zoning.
- RSCN and hardware resource efficiencies are the same as single-initiator zoning.

## Peer Zone Connectivity Rules

Peer Zoning adheres to the following connectivity rules:

- Nonprincipal devices can communicate only with the principal device.
  - Nonprincipal devices cannot communicate with other nonprincipal devices unless allowed by some other zone in the active zone set.

### NOTE

For peer zones with D,I members, if there are multiple devices connected to the port index of the switch (for example, NPIV devices), all of those devices will be able to communicate with each other. In this case, it is suggested that the user create a peer zone with WWN members instead to ensure that nonprincipal devices cannot communicate with each other.

- Principal devices cannot communicate with other principal devices unless allowed by some other zone in the active zone set.
- The maximum number of peer zones is determined by the zone database size. See [Zone Database Size](#) for information on zone database sizing.

## Firmware Upgrade and Downgrade Considerations for Peer Zoning

When a switch containing Peer Zones is upgraded to Fabric OS 7.4.x, Peer Zoning connectivity rules are enforced on the devices attached to the upgraded switch. If the device is a member of a Peer Zone, the upgrade can cause disruption.

When a switch containing peer zones is downgraded from Fabric OS 7.4.x to Fabric OS 7.3.x or earlier, the peer zones are treated as regular zones, and the zones follow the regular zoning connectivity rules.

All devices in the peer zone must be attached to switches running Fabric OS 7.4.x, for the peer zoning connectivity rules to be applied. If any device that is part of the peer zone is attached to a switch running an earlier version of Fabric OS, the peer zone is treated as a regular zone.

After downgrading to an earlier version of Fabric OS, Peer Zoning connectivity rules remain on the switch until a port of the Peer Zone device is toggled or a new zoning configuration is committed. In such a case, the regular zoning rules apply to the Peer Zones.

## LSAN and QoS Peer Zoning Considerations

LSAN peer zoning is supported on edge-to-edge and edge-to-backbone topologies. When the zoning configuration is enabled, peer zoning rules are followed

QoS peer zones are supported as long as the zone name has the prefix "QOSH" or "QOSL" in the zone name. See [QoS](#) for more information.

## Peer Zone Configuration

Peer Zones are configured and managed using zone commands with the `--peerzone` option.

In the peer zone, devices must be identified as either WWN or D,I devices. Mixing WWN and D,I device identification within a single peer zone is not allowed.

Aliases are supported as devices for a Peer Zone. For more details, refer to [Alias Support for Peer Zoning](#).

### Creating Peer Zones

To create a Peer Zone configuration, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Create peer zones using the `zonecreate --peerzone` command. The `-members` operand is optional.

The following example creates two peer zones. The first peer zone is called `peerzone_wwn_mbrs`, which contains 10:00:00:00:01:1e:20:20 as the principal device and 10:00:00:02:1f:02:00:01 and 10:00:05:1e:a9:20:00:01 as the non-principal devices. The second peer zone is called `peerzone_di_mbrs`, which contains 10,1 as the principal device and 20,1 and 20,2 as the non-principal devices.

```
switch:admin> zonecreate --peerzone peerzone_wwn_mbrs -principal "10:00:00:00:01:1e:20:20"/
-members "10:00:00:02:1f:02:00:01;10:00:05:1e:a9:20:00:01"
switch:admin> zonecreate --peerzone peerzone_di_mbrs -principal "10,1" -members "20,1;20,2"
```

### Adding Devices to a Peer Zone

To add devices to the Peer Zone, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Add Peer Zone devices using the `zoneadd --peerzone` command.

The following example adds the device "10:00:05:1e:a9:20:00:02" to `peerzone_wwn_mbrs`.

```
switch:admin> zoneadd --peerzone peerzone_wwn_mbrs -members "10:00:05:1e:a9:20:00:02"
```

### Replacing a Device in a Peer Zone

To replace devices in the Peer Zone, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Replace the existing peer zone device with a new device using the `zoneobjectreplace` command.

The `zoneobjectreplace` command restricts replacing the peer zone property member. A WWN can be replaced with a WWN only, and a D,I member can be replaced with a D,I member only.

The following example replaces the device "20:00:00:05:1e:a1:af:b2" with "10:00:00:05:1e:a1:10:c1" in a Peer Zone.

```
switch:admin> zoneobjectreplace "20:00:00:05:1e:a1:af:b2" "10:00:00:05:1e:a1:10:c1"
```

## Removing Devices from a Peer Zone

To remove devices from a Peer Zone, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Remove devices from a peer zone using the `zoneremove --peerzone` command.

Removing the last principal WWN device deletes the peer zone.

The following example removes the device "10:00:05:1e:a9:20:00:02" from the "peerzone\_wwn\_mbrs" peer zone.

```
switch:admin> zoneremove --peerzone peerzone_wwn_mbrs -members "10:00:05:1e:a9:20:00:02"
```

## Deleting Peer Zones

To delete a Peer Zone, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Delete peer zones using the `zonedeldelete` command.

The following example deletes the `peerzone_di_mbrs` peer zone.

```
switch:admin> zonedeldelete peerzone_di_mbrs
```

## Viewing Peer Zones

To view a peer zone configuration, including the devices that are associated with the peer zone, complete the following steps:

1. Connect to the switch, and log in to an account.
2. View the peer zones configuration that is contained in the defined and effective zone databases using the `zoneshow --peerzone` command.
  - Use the `all` option to list all peer zones in the zone database.
  - Use the `user` option to list only the user-created peer zones.
  - Use the `target` option to list all target driven peer zones.
  - Use the `-pattern <pattern_name>` suboption to filter the peer zones using a pattern.

For more information about the commands, refer to the *Brocade Fabric OS Command Reference Manual*.

The following example displays the peer zone configurations using the `all` option to show all peer zones. The `Property Member` field distinguishes a peer zone from a regular zone. The `Peer Member(s)` field lists the non-principal members of the peer zone.

```
switch:admin> zoneshow --peerzone all
Defined configuration:
zone: peer_19_0f_10_SB23
 Property Member: 00:02:00:00:00:03:00:02
 Created by: User
 Principal Member(s):
 20:03:00:11:0d:2d:04:00; 20:06:00:11:0d:41:ae:00
 Peer Member(s):
 10:00:00:10:9b:bd:c0:10; 5d:81:fc:c0:3a:57:d0:01
zone: peer_20_ed_sb23
 Property Member: 00:02:00:00:00:03:00:01
 Created by: User
 Principal Member(s):
```

```

 20:02:00:11:0d:2d:03:00
 Peer Member(s):
 10:00:00:10:9b:8f:2d:ed
zone: peer_25_20244^--_SB23
 Property Member: 00:02:00:00:00:02:00:01
 Created by: User
 Principal Member(s):
 61,5
 Peer Member(s):
 11,33

```

#### Effective configuration:

```

zone: peer_19_0f_10_SB23
 Property Member: 00:02:00:00:00:03:00:02
 Created by: User
 Principal Member(s):
 20:03:00:11:0d:2d:04:00
 20:06:00:11:0d:41:ae:00
 Peer Member(s):
 10:00:00:10:9b:bd:c0:10
 5d:81:fc:c0:3a:57:d0:01
zone: peer_20_ed_sb23
 Property Member: 00:02:00:00:00:03:00:01
 Created by: User
 Principal Member(s):
 20:02:00:11:0d:2d:03:00
 Peer Member(s):
 10:00:00:10:9b:8f:2d:ed
zone: peer_25_20244^--_SB23
 Property Member: 00:02:00:00:00:02:00:01
 Created by: User
 Principal Member(s):
 61,5
 Peer Member(s):
 11,33

```

#### 3 Peer Zones in Eff Cfg

The following example uses the `all -pattern` option to show all peer zones that end with the pattern "ed\_sb23".

```

switch:admin> zoneshow --peerzone all -pattern *ed_sb23
zone: peer_20_ed_sb23
 Property Member: 00:02:00:00:00:03:00:01
 Created by: User
 Principal Member(s):
 20:02:00:11:0d:2d:03:00
 Peer Member(s):
 10:00:00:10:9b:8f:2d:ed

```

The following example uses the `user -pattern` option to show all user-created peer zones that have the pattern "2024" in the zone name.

```

switch:admin> zoneshow --peerzone user -pattern *2024*
zone: peer_25_20244^--_SB23
 Property Member: 00:02:00:00:00:02:00:01
 Created by: User
 Principal Member(s):
 61,5
 Peer Member(s):
 11,3

```

## Alias Support for Peer Zoning

You can also create peer zones using alias names as principal and non-principal members if all the switches in the fabric are running Fabric OS 8.1.0 or a later. If you merge a domain running Fabric OS 8.0.1 or lower with a domain running Fabric OS 8.1.0 or later and have alias peer zones configured, it will result in a zone merge segmentation.

In case of chassis systems, both CPs must be running Fabric OS 8.1.0 or later.

If a standby CP is running a version of Fabric OS that does not support alias peer zones or if the CP is then made to sync with the active CP running a version of Fabric OS that supports and contains alias peer zones, then one of the following results occurs:

- The standby CP will not sync until all alias peer zones are removed from the active CP.
- The standby CP is upgraded to a version of Fabric OS that support alias peer zones.

**Table 137: Zone Merge Results When Using Alias Peer Zoning**

| Existing Fabric Configuration                 | Joining Fabric Configuration                  | Zone Merge Result                                                                     |
|-----------------------------------------------|-----------------------------------------------|---------------------------------------------------------------------------------------|
| Alias peer zoning configured                  | Alias peer zoning configured                  | Merges without any condition.                                                         |
| Alias peer zoning configured                  | Alias peer zoning capable, but not configured | Merges as long as the rest of fabric is alias peer zone capable; otherwise, segments. |
| Alias peer zoning configured                  | Not capable of alias peer zoning              | Segments.                                                                             |
| Alias peer zoning capable, but not configured | Alias peer zoning configured                  | Merges as long as the rest of fabric is alias peer zone capable; otherwise, segments. |
| Alias peer zoning capable, but not configured | Alias peer zoning capable, but not configured | Merges without any condition.                                                         |
| Alias peer zoning capable, but not configured | Not capable fo alias peer zoning              | Merges without any condition.                                                         |
| Not capable of alias peer zoning              | Alias peer zoning configured                  | Segments.                                                                             |
| Not capable of alias peer zoning              | Alias peer zoning capable, but not configured | Merges without any condition.                                                         |
| Not capable of alias peer zoning              | Not capable fo alias peer zoning              | Merges without any condition.                                                         |

The following restrictions are applicable to alias peer zones:

- The members of a peer zone can be either of the following:

- WWN name.
- Alias with only WWN as members.
- Alias with only WWN as members and WWN mixture.
- D,I.
- Alias with only D,I as members.
- Alias with only D,I as members and D,I mixture.
- A peer zone cannot have both WWN and D,I members including the members of aliases included in the peer zone.
- An alias member of a peer zone cannot be both principal as well as non-principal within the same peer zone. It can only be added either as principal (or) peer to a peer zone.
- If an alias is included in a peer zone, it will have a restriction from having mixed type members (that is, D, I versus WWN). Mixed members will be allowed only after removing that alias from all the peer zones of which the alias is a part.
- Aliases in a peer zone are restricted from having duplicate members either in the peer zone or members within the alias itself. The following is an example of an invalid configuration having aliases with duplicate members across two different aliases.

```
Ex: ali_prz_dup 00:02:00:00:02:02:02; ali1; ali2; 4,5
 ali1 7,8
 ali2 7,8
```

- You cannot create a peer zone with aliases if a domain in the fabric is running Fabric OS 8.0.1 or earlier.
- The total number of principal members is restricted to 255 including the members of an alias if the aliases were added as principal members.

```
Ex: p4 00:02:00:00:02:05:09; ali1; 8,7; 9,9; 9,7; ali2; 4,5
 ali1 7,8; 6,7; 9,11
 ali2 5,8; 7,7; 11,11 - total members are 9 including the members of alias
```

- Downgrading the firmware and downloading the configuration to Fabric OS 8.0.1 or lower are blocked if alias peer zones are present.
- Modifying members of an alias to have mixed members is blocked if the alias is part of one or more peer zones. This is applicable for `aliCreate` and `aliAdd` commands.

Use the following commands to configure and display peer zones with alias names as members.

- The `zoneCreate` command supports alias inclusion in principal and member list.
- The `zoneAdd` commands supports aliases as both principal and non-principal members of a peer zone.
- The `zoneRemove` command supports alias inclusion in principal and member list. Deleting the last principal WWN, D,I, or alias name removes the peer zone from the zone database.
- The `zoneShow` commands displays peer zones with aliases.

The following examples show peer zone configuration with alias support.

```
sw0:FID128:admin> zonecreate --peer p1 -p "ali1;8,7;9,9" -m "4,5";

sw0:FID128:admin> zoneshow --verbose
Defined configuration:
 zone: p1 00:02:00:00:00:02:03:05; ali1; 8,7; 9,9; 4,5
 alias: ali1 7,8; 6,7; 9,11
 alias: ali2 5,8; 7,7; 11,11
Effective configuration:
 no configuration in effect

sw0:FID128:admin> cfgcreate cp1,p1
sw0:FID128:admin> cfgenable cp1

sw0:FID128:admin> zoneadd --peer p1 -p "9,7;ali2";
```

```
sw0:FID128:admin> zoneshow --verbose
Defined configuration:
 cfg: cp1 p1
 zone: p1 00:02:00:00:00:02:05:09; ali1; 8,7; 9,9; 9,7; ali2; 4,5
 alias: ali1 7,8; 6,7; 9,11
 alias: ali2 5,8; 7,7; 11,11
Effective configuration:
 cfg: cp1
 zone: p1 00:02:00:00:00:02:03:05
 7,8
 6,7
 9,11
 8,7
 9,9
 4,5

sw0:FID128:admin> zoneremove --peer p1 -p "ali2";
```

```
sw0:FID128:admin> zoneshow --verbose
Defined configuration:
 cfg: cp1 p1
 zone: p1 00:02:00:00:00:02:04:06; ali1; 8,7; 9,9; 9,7; 4,5
 alias: ali1 7,8; 6,7; 9,11
 alias: ali2 5,8; 7,7; 11,11
Effective configuration:
 cfg: cp1
 zone: p1 00:02:00:00:00:02:03:05
 7,8
 6,7
 9,11
 8,7
 9,9
```

```
sw0:FID128:admin> zoneshow --peer all
Defined configuration:
 zone: p1
 Property Member: 00:02:00:00:00:02:04:06
 Created by: User
 Principal Member(s):
 ali1; 8,7; 9,9; 9,7
 Peer Member(s):
 4,5
Effective configuration:
 zone: p1
 Property Member: 00:02:00:00:00:02:03:05
 Created by: User
 Principal Member(s):
 7,8
 6,7
 9,11
 8,7
 9,9
 Peer Member(s):
```

```

4,5
1 Peer Zones in Eff Cfg

sw0:FID128:admin> zoneshow --verbose
Defined configuration:
cfg: c1 p8
zone: p8 00:02:00:00:00:02:02:01; ali5; 86,8

sw0:FID128:admin> alicreate ali5, "30:08:00:05:33:88:e3:f3"
Error: Members of alias "ali5" can either be WWN or D,I: aliases having mixed-type members are not
allowed to be a part of an alias peer zone. To allow mixed members remove this alias from all peer
zones.

sw0:FID128:admin> zoneshow --verbose
Defined configuration:
cfg: c1 p8
zone: p8 00:02:00:00:00:02:02:01; ali5; 86,8
alias: ali5 5,8; 7,7

sw0:FID128:admin> aliadd ali5, "30:08:00:05:33:88:e3:f3"
Error: Members of alias "ali5" can either be WWN or D,I: aliases having mixed-type members are not
allowed to be a part of an alias peer zone. To allow mixed members remove this alias from all peer
zones.

```

## Displaying the Property Members of Peer Zones

By default, the following commands do not display the property members of peer zones. To display the property members of peer zones, you can use the `--verbose` option with these commands except the `zoneshow --peerzone` command, which displays the property members by default.

- `zone`
- `zoneshow`
- `cfgshow`
- `alishow`


For more details, refer to the *Brocade Fabric OS Command Reference Manual*.

## Target-Driven Zoning

Target-driven zoning is a variant of peer zoning. Where a regular peer zone is defined by a user-specified configuration, target-driven peer zone is defined by the principal device. This device is usually a storage device, but not always.

Target-driven zoning manages the zoning using a third-party management interface to manage the device and the switch interactions. To permit a target-driven peer zone, target-driven zoning must be enabled on the F\_Port that connects the principal device to the fabric. Refer to the manual for the device used as the principal device to determine the supported commands and options to construct a target-driven peer zone.



**Figure 76: Target-Driven Zoning Example**

## Limitations and Considerations for Target-Driven Zoning

The following are the general limitations and considerations for target-driven zoning:

- A zone configuration must be created and enabled before a principal device can add or activate a target-driven peer zone. However, in the event where no configuration is enabled and if the zone database is completely empty, the system automatically creates and enables a zone configuration with the name "tdz\_cfg\_auto". As part of auto-creation of "tdz\_cfg\_auto", a zone named "tdz\_cfg\_auto" is also created and added as a member of this configuration to allow the target to remove its last zone member.
- The maximum number of devices allowed in a target-driven peer zone, including a principal device, is 255.
- The target-driven zoning configuration mode must be enabled on the F\_Port to which the principal device is connected.
- Target-driven peer zones cannot be created or modified using Web Tools or SANnav or the CLI; however, they can be displayed or deleted using BNA or SANnav or CLI.
- In target-driven peer zones, devices are identified as WWN devices only.

## Target-Driven Zoning Configuration

Target-driven zoning can be configured and managed at the port level using the `portcfgtdz` command.

To create and commit a TDZ successfully, you must create a zone and add and enable the zone in the active configuration. If no configuration is enabled and the zone DB is completely empty, a "tdz\_cfg\_auto" is automatically created. For procedures to create, add, and enable zones in a zone database, refer to [Zone configurations](#).

You should perform all target-driven zoning tasks using the zone management interface on the principal device.

## Enabling Target-Driven Zoning

Use the management interface on the principal device to create and manage target-driven zoning.

To create and enable target-driven zoning, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions. By default, target-driven zoning configuration is disabled for all ports.
2. Enable target-driven zoning on a port or port range using the `portcfgtdz --enable` command.

The following example enables target-driven zoning. You can enable target-driven zoning for a single port, a range of ports, or all ports on a switch.

```
switch:admin> portcfgtdz --enable 8
```

The following example enables target-driven zoning on a range of ports.

```
switch:admin> portcfgtdz --enable 8-18
```

The following example enables target-driven zoning on all ports.

```
switch:admin> portcfgtdz --enable "*"
```

If the zone-defined database is empty, a zone is created and enabled for a target automatically. If the defined database is not empty and an enabled zone configuration exists, then no action is required. If no enabled zone configuration exists, you can enable a configuration for TDZ using the following steps:

3. Create a zone for target-driven zoning using the `zonecreate` command.

The following example creates a zone for target-driven zoning.

```
switch:admin> zonecreate "targetzone1", "10:00:00:00:c9:2b:c9:0c;
50:05:07:61:00:5b:62:ed;50:05:07:61:00:49:20:b4"
```

4. Create a zone configuration for target-driven zoning in the zone database using the `cfgcreate` command.

The following example creates a zone configuration for target-driven zoning.

```
switch:admin> cfgcreate cfgtdz,targetzone1
```

5. Enable the zone for target-driven zoning using the `cfgenable` command.

The following example enables the zone for target-driven zoning.

```
switch:admin> cfgenable cfgtdz
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'cfgtdz' configuration (yes, y, no, n): [no] y
zone config "cfgtdz" is in effect
Updating flash ...
```

## Disabling Target-Driven Zoning

To disable Target Driven Zoning, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Disable target-driven zoning on a port or port range using the `portcfgtdz --disable` command.

The following example disables target-driven zoning on a port.

```
switch:admin> portcfgtdz --disable 10
```

The following example disables Target Driven Zoning on a range of ports.

```
switch:admin> portcfgtdz --disable 2-10
```

The following example disables target-driven zoning on all ports.

```
switch:admin> portcfgtdz --disable ""
```

You receive a peer zone change RSCN if an active principle device is attached to the specified port.

## Viewing Target-Driven Zoning

To view the target-driven zoning status, complete the following steps.

1. Connect to the switch and log in using an account.
2. View target-driven zoning status on a port or port range.

The following example displays target-driven zoning status on the ports using the `portcfgtdz --show` command.

```
switch:admin> portcfgtdz --show 1-2
Port Mode
=====
1 ON
2 OFF
```

The following example displays the target-driven zoning status on a port using the `portcfgshow` command.

```
switch:admin> portcfgshow 1
Area Number: 1
Octet Speed Combo: 1(16G|8G|4G|2G)
Speed Level: AUTO(SW)
AL_PA Offset 13: OFF
Trunk Port ON
Long Distance OFF
VC Link Init OFF
Locked L_Port OFF
Locked G_Port OFF
Disabled E_Port OFF
Locked E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable ON
LOS TOV mode 0(OFF)
NPIV capability ON
QOS Port AE
Port Auto Disable: OFF
Rate Limit OFF
EX Port OFF
Mirror Port OFF
SIM Port OFF
Credit Recovery ON
F_Port Buffers OFF
E_Port Credits OFF
Fault Delay: 0(R_A_TOV)
NPIV PP Limit: 126
```

```

NPIV FLOGI Logout: OFF
CSCTL mode: OFF
TDZ mode: ON
D-Port mode: OFF
D-Port over DWDM: OFF
Compression: OFF
Encryption: OFF
10G/16G FEC: ON
16G FEC via TTS: OFF
Non-DFE: OFF

```

3. (Optional) View the principal device and the non principal devices in a zone using the `zoneshow --peerzone` command. If the Created by field is User, the zone is a user-created peer zone. If the Created by field is Target, the zone is a target-driven peer zone.

The following example displays the complete peer zone details.

```

switch:admin> zoneshow --peerzone all
Defined configuration:
zone: peerzone_wwn_mbrs
 Property Member: 00:02:00:00:00:03:00:01
 Created by: User
 Principal Member(s):
 10:00:00:00:01:1e:20:20
 Peer Member(s):
 10:00:00:02:1f:02:00:01; 10:00:05:1e:a9:20:00:01
zone: targetzone1
 Property Member: 00:01:00:00:00:03:00:01
 Created by: Target
 Principal Member(s):
 30:04:00:05:1e:61:23:8f
 Peer Member(s):
 10:00:00:02:1f:02:00:05; 10:00:00:02:1f:02:00:12

Effective configuration:
zone: peerzone_wwn_mbrs
 Property Member: 00:02:00:00:00:03:00:01
 Created by: User
 Principal Member(s):
 10:00:00:00:01:1e:20:20
 Peer Member(s):
 10:00:00:02:1f:02:00:01
 10:00:05:1e:a9:20:00:01
zone: targetzone1
 Property Member: 00:01:00:00:00:03:00:01
 Created by: Target
 Principal Member(s):
 30:04:00:05:1e:61:23:8f
 Peer Member(s):
 10:00:00:02:1f:02:00:05
 10:00:00:02:1f:02:00:12

2 Peer Zones in Eff Cfg

```

The following example displays the target-driven peer zones in a zone.

```

switch:admin> zoneshow --peerzone target
Defined configuration:
zone: targetzone1
 Property Member: 00:01:00:00:00:03:00:01
 Created by: Target
 Principal Member(s):
 30:04:00:05:1e:61:23:8f
 Peer Member(s):
 10:00:00:02:1f:02:00:05; 10:00:00:02:1f:02:00:12

Effective configuration:
zone: targetzone1
 Property Member: 00:01:00:00:00:03:00:01
 Created by: Target
 Principal Member(s):
 30:04:00:05:1e:61:23:8f
 Peer Member(s):
 10:00:00:02:1f:02:00:05
 10:00:00:02:1f:02:00:12

```

```
1 Peer Zones in Eff Cfg
```

## Alias Support for Target-Driven Zoning

You can also create target-driven zones containing one or more alias names in the nonprincipal member list in addition to WWN members if all the switches in the fabric are running Fabric OS 8.2.x or later. If you merge a domain running Fabric OS 8.1.x or earlier with a domain running Fabric OS 8.2.0 or later and have alias target driven zones configured, the link is segmented.

In case of chassis systems, both CPs must be running Fabric OS 8.2.0 or later. If a standby CP is running a version of Fabric OS that does not support alias target-driven zones, if the CP is then made to sync with the active CP running a version of Fabric OS that supports alias target-driven zoning and contains alias target-driven zones, the standby CP will not sync until either all alias target driven zones are removed from the active CP or the standby is upgraded to a version of Fabric OS that supports alias target driven zones.

**Table 138: Zone Merge Results When Using Alias Target Driven Zoning**

| Existing Fabric Configuration                          | Joining Fabric Configuration                           | Zone Merge Result                                                                             |
|--------------------------------------------------------|--------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Alias target driven zoning configured                  | Alias target driven zoning configured                  | Merges without any condition.                                                                 |
| Alias target driven zoning configured                  | Alias target driven zoning capable, but not configured | Merges as long as the rest of fabric is alias target driven zone capable; otherwise segments. |
| Alias target driven zoning configured                  | Not capable of alias target driven zoning              | Segments.                                                                                     |
| Alias target driven zoning capable, but not configured | Alias target driven zoning configured                  | Merges as long as the rest of fabric is alias target driven zone capable; otherwise segments. |
| Alias target driven zoning capable, but not configured | Alias target driven zoning capable, but not configured | Merges without any condition.                                                                 |
| Alias target driven zoning capable, but not configured | Not capable of alias target driven zoning              | Merges without any condition.                                                                 |
| Not capable of alias target driven zoning              | Alias target driven zoning configured                  | Segments.                                                                                     |

| Existing Fabric Configuration             | Joining Fabric Configuration                           | Zone Merge Result             |
|-------------------------------------------|--------------------------------------------------------|-------------------------------|
| Not capable of alias target driven zoning | Alias target driven zoning capable, but not configured | Merges without any condition. |
| Not capable of alias target driven zoning | Not capable of alias target driven zoning              | Merges without any condition. |

The following restrictions are applicable to alias target driven zones:

- The nonprincipal members of a target-driven zone can be one of the following:
  - Only WWN members.
  - Aliases with only WWN members.
  - Aliases with WWN members and WWN members.
- The following nonprincipal members for a target driven zone are not supported:
  - Aliases with mixed members (D,I and WWN).
  - D, I members.
  - Aliases with D,I members.
- Switches running Fabric OS 8.2.0 and later do not allow you to edit the aliases that are members of TDZs.
- Although, you can modify the aliases from a device running Fabric OS 8.1.0a, modifying aliases that are members of a target driven zone can lead to the following scenarios:
  - Non-target driven zone administrators can change the target driven zone connectivity using the command line interface or SANnav or the Brocade Network Advisor from devices running Fabric OS 8.1.0a.
  - Peer-zone-change RSCNs are not sent to targets if an alias that is a member of target driven zone is modified from a device running Fabric OS 8.1.0a.
- Downgrading the firmware and downloading configuration to Fabric OS 8.1.x or earlier is blocked if alias target driven zones are present.

## Supported Commands for Peer Zones and Target-Driven Peer Zones

Peer zones are managed using the `--peerzone` operand with regular zone commands for some actions. You can manage target-driven peer zones using the regular zone commands, but not all commands are supported. See the following table for the list of commands supported for peer zones and target-driven peer zones.

**Table 139: Supported Commands for Peer Zones and Target-Driven Peer Zones**

| Command                 | Support for Peer Zones                         | Support for Target-Driven Peer Zones |
|-------------------------|------------------------------------------------|--------------------------------------|
| <code>cfgAdd</code>     | Yes                                            | Yes                                  |
| <code>cfgCreate</code>  | Yes                                            | Yes                                  |
| <code>cfgDelete</code>  | Yes                                            | Yes                                  |
| <code>cfgDisable</code> | Yes                                            | Yes                                  |
| <code>cfgEnable</code>  | Yes                                            | Yes                                  |
| <code>cfgRemove</code>  | Yes                                            | Yes                                  |
| <code>cfgSave</code>    | Yes                                            | Yes                                  |
| <code>cfgShow</code>    | Yes                                            | Yes                                  |
| <code>zone</code>       | No                                             | No                                   |
| <code>zoneAdd</code>    | Yes, using the <code>--peerzone</code> operand | No                                   |
| <code>zoneCreate</code> | Yes, using the <code>--peerzone</code> operand | No                                   |


| Command           | Support for Peer Zones                         | Support for Target-Driven Peer Zones           |
|-------------------|------------------------------------------------|------------------------------------------------|
| zoneDelete        | Yes                                            | Yes                                            |
| zoneObjectCopy    | Yes                                            | No                                             |
| zoneObjectExpunge | Yes                                            | Yes                                            |
| zoneObjectRename  | Yes                                            | No                                             |
| zoneObjectReplace | Yes                                            | No                                             |
| zoneRemove        | Yes, using the <code>--peerzone</code> operand | No                                             |
| zoneShow          | Yes, using the <code>--peerzone</code> operand | Yes, using the <code>--peerzone</code> operand |

## Boot LUN Zoning

Boot from SAN provides rapid server provisioning, rapid recovery from the server failure, improved reliability, and reduced power and cooling. Brocade/QLogic HBAs and Brocade/Emulex HBAs support boot from SAN in all types of environments. Brocade features make boot from SAN easier to deploy, such as putting a label with the Port World Wide Name (PWWN) on the bracket of the HBA and a feature called Fabric-Based Boot LUN Discovery. Fabric-Based Boot LUN Discovery can be deployed in a Brocade SAN using both Brocade/QLogic HBAs and Brocade/Emulex HBAs; thus, eliminating the manual process of boot LUN configuration of each HBA from each server console. This capability dramatically improves configuration and reliability, containing operational costs, and making the adoption of diskless servers a practical alternative for even large data centers.

### NOTE

IBM UEFI and other UEFI-based servers do not support this unique Brocade feature, because of their different discovery processes.

**Figure 77: Boot LUN Zoning**

FOS blocks the creation of any zone with a zone name having the format BFA\_XXXX\_BLUN (that is, starting with BFA\_ and ending in \_BLUN). Zone names with this configuration are considered to indicate an HBA Boot LUN Zone. For peer zoning, restrictions to protect property members starting with 00: were added. This restriction has been extended to exclude any WWNs that start with 00:00:00: because this WWN value indicates a Boot LUN zone member.

Zone configuration rules permit creating a Boot LUN zone using the `bootluncfg` command. To enable a Boot LUN configuration, you must create a basic zone for the host-target pair using the `bootluncfg` command and then add this zone object to the zone configuration to ensure that the HBA has access to the target. After this zone is established, Get Zone Member (GZM) queries from the HBA will handle getting the Boot LUN information from the fabric.

Refer to the *Brocade Fabric OS Command Reference Manual* for additional information on the `bootluncfg` command.


## Setting Up Boot LUN Zoning

You must configure the server, HBA, and Brocade switch for boot LUN zoning to work. You must also know how to flash the HBA BIOS or upload the latest version of Brocade Fabric OS on a Brocade SAN switch or director. Other than that, no additional hardware or extra license or software is needed to deploy Fabric-Based Boot LUN Discovery.

### Configuring the Server and HBA for a BLUN Zone

Fabric OS HBAs are factory-configured to support Fabric-based Boot LUN Discovery in the default BIOS setting. Verify that the BIOS is enabled and that the Boot LUN is set to Auto Discover, as shown in the following figure.



**Figure 78: Boot LUN Zoning BIOS Setting**

## Configuring the Brocade Switch for a BLUN Zone

In addition to zoning and masking, you must configure an additional Boot LUN (BLUN) zone on the switch. At server boot, the Brocade HBA (set to Auto Discover) makes a connection to the switch, retrieves the remote PWWN and the LUN from this BLUN zone, and tries to boot from this LUN.

1. To configure a BLUN zone, use the `bootLunCfg --add` command.

```

switch:admin> bootluncfg --add 11:22:ab:44:44:ff:44:ca \
1b:6c:55:55:55:3a:55:ff 9abc345fa1112410
Operation Successful

```

2. To display existing BLUN mappings, use the `bootLunCfg --show` command.

```

switch:admin> bootluncfg --show
00:11:22:33:44:55:66:77
00:00:00:00:aa:bb:cc:dd;00:00:00:01:ee:ff:11:22; \
00:00:00:02:9a:bc:34:5f;00:00:00:03:a1:11:24:10
aa:aa:aa:aa:aa:aa:aa:aa
00:00:00:00:11:11:11:11;00:00:00:01:11:11:11:11; \
00:00:00:02:9a:bc:34:5f;00:00:00:03:a1:11:24:10
bb:aa:aa:aa:aa:aa:aa:aa
00:00:00:00:11:11:11:11;00:00:00:01:11:11:11:11; \
00:00:00:02:9a:bc:34:5f;00:00:00:03:a1:11:24:10

```

3. To remove an HBA to BLUN mapping, use the `bootLunCfg --delete` command.

```

switch:admin> bootluncfg --delete 11:22:ab:44:44:ff:44:ca \
1b:6c:55:55:55:3a:55:ff 9abc345fa1112410
Operation Successful

```

4. The Brocade Command Line Utility (BCU), loaded with the driver package, provides commands that are comparable to Brocade HCM features for configuring, troubleshooting, and monitoring the HBA and device connection.

### Syntax:

```
switch#> bcu command --operand synopsis
```

**Example:**

```
switch#> bcu adapter --query <ad_id>
```

5. The `bcu boot -blunZone` command helps to create the command, which can then be copied and pasted. Here is the same example but with a different LUN 0xff.

```
switch#> BCU boot --blunZone -c BLUN -p 10:00:00:05:1e:0f:aa:21 -r 50:00:00:e0:d0:44:29:86 -l ff
zonecreate "bfa_100000051e0faa21_blun", "00:00:00:00:50:00:00:e0; 00:00:00:01:d0:44:29:86;
00:00:00:02:00:ff:00:00; 00:00:00:03:00:00:00:00"
```

## Routing Traffic

---

This section describes how the data moves through a fabric from switch to switch and from storage to server along one or more paths that make up a *route*. Routing policies determine the path for each frame of data.

Before the fabric can begin routing traffic, it must discover the route that a packet should take to reach the intended destination. Route tables are lists that indicate the next hop to which packets are directed to reach a destination. Route tables include network addresses, the next address in the data path, and a cost to reach the destination network. Intranet networks have two types of routing protocols: distance vector and link state.

- Distance vector is based on the hop count and it is the number of switches that a frame passes through to get from the source switch to the destination switch.
- Link state is based on a metric value that is based on a cost. The cost could be based on bandwidth, line speed, or round-trip time.

With the link state protocol, switches that discover a route identify the networks to which they are attached, receiving an initial route table from the principal switch. After an initial message is sent out, the switch only notifies the others when changes occur.

No more than seven hops occur between any two switches. This limit is not required or enforced by Fabric Shortest Path First (FSPF). Its purpose is to ensure that a frame is not delivered to a destination after the Resource Allocation TimeOut Value (R\_A\_TOV) has expired.

Fabric OS supports unicast Class 2 and Class 3 traffic, multicast, and broadcast traffic. Broadcast and multicast are supported in Class 3 only.

### NOTE

Routed edge-core-edge EX\_Ports can only be used in the base switch. Other logical switch that is present in a Backbone-routed base fabric cannot be in a default switch even though they are not in the routed traffic path.

## Paths and Route Selection

Paths are possible ways to get from one switch to another. Each inter-switch link (ISL) has a metric cost based on bandwidth. The cumulative cost is based on the sum of all costs of all traversed ISLs.

Route selection is the path that is chosen. Paths that are selected from the routing database are chosen based on the minimal cost.

## FSPF

Fabric Shortest Path First (FSPF) is a link state path selection protocol that directs the traffic along the shortest path between the source and destination that is based on the link cost. FSPF is also referred to as Layer 2 routing. FSPF detects link failures, determines the shortest route for traffic, updates the routing table, provides fixed routing paths within a fabric, and maintains correct ordering of frames. FSPF also tracks the state of the links on all switches in the fabric and associates a cost with each link. The protocol computes paths from a switch to all the other switches in the fabric by adding the cost of all links traversed by the path, and chooses the path that minimizes the costs. This collection of the link states, including costs, of all the switches in the fabric constitutes the topology database or link state database.


When established, FSPF programs the hardware routing tables for all active ports on the switch. FSPF is not involved in frame switching. FSPF uses several frames to perform its functions. Because it may run before fabric routing is set up, FSPF does not use the routing tables to propagate the frames, but it floods the frames throughout the fabric hop-by-hop. Frames are first flooded on all the ISLs; as the protocol progresses, it builds a spanning tree that is rooted on the principal switch. Frames are only sent on the principal ISLs that belong to the spanning tree. When multiple ISLs are between

switches, the first ISL to respond to connection requests becomes the principal ISL. Only one ISL from each switch is used as the principal ISL.

The following figure shows the thick red lines as principal ISLs, and thin green lines as regular ISLs.

**Figure 79: Principal ISLs**

Principal switch



**NOTE**

FSPF only supports 16 routes in a zone.

FSPF makes minimal use of the ISL bandwidth, leaving virtually all of it available for traffic. In a stable fabric, a switch transmits 64 bytes every 20 seconds in each direction. FSPF frames have the highest priority in the fabric, which guarantees that a control frame is not delayed by user data and that FSPF routing decisions occur quickly during convergence.

FSPF guarantees a routing loop-free topology always. It is essential for a fabric to include many physical loops because, without loops, there would not be multiple paths between switches, and therefore no redundancy. Without redundancy, if a link goes down, part of the fabric is isolated. FSPF ensures both that the topology is loop-free and that a frame is never forwarded over the same ISL more than once.

FSPF calculates paths that is based on the destination domain ID. The fabric protocol must complete domain ID assignments before routing can begin. ISLs provide the physical pathway when the Source ID (SID) address has a frame that is destined to a port on a remote switch Destination ID (DID). When an ISL is attached or removed from a switch, FSPF updates the route tables to reflect the addition or deletion of the new routes.

As each host transmits a frame to the switch, the switch reads the SID and DID in the frame header. If the domain ID of the destination address is the same as the switch (intra-switch communications), the frame buffer is copied to the destination port and a credit R\_RDY message is sent to the host. The switch must only read the word zero and the word one of the Fibre Channel frame to perform what is known as *cut-through routing*. A frame may begin to emerge from the

output port before it has been entirely received by the input port. The entire frame does not need to be buffered in the switch.

If the destination domain ID is different from the source domain ID, the switch consults the FSPF route table to identify which local E\_Port provides Fabric Shortest Path First (FSPF) to the remote domain.

## Fibre Channel NAT

Within an edge fabric or across a backbone fabric, the standard Fibre Channel FSPF protocol determines how frames are routed from the source Fibre Channel (FC) device to the destination FC device. The source or destination device can be a proxy device.


Fibre Channel fabrics require that all ports be identified by a unique port identifier (PID). In a single fabric, FC protocol guarantees that domain IDs are unique, and so a PID formed by a domain ID and an area ID is unique within a fabric. However, the domain IDs and PIDs in one fabric may be duplicated within another fabric, just as IP addresses that are unique to one private network are likely to be duplicated within another private network.

In an IP network, a network router can maintain network address translation (NAT) tables to replace private network addresses with public addresses when a packet is routed out of the private network, and to replace public addresses with private addresses when a packet is routed from the public network to the private network. The Fibre Channel routing equivalent to this IP-NAT is Fibre Channel network address translation (FC-NAT). Using FC-NAT, the proxy devices in a fabric can have PIDs that are different from the real devices they represent, allowing the proxy devices to have appropriate PIDs for the address space of their corresponding fabric.

## Inter-Switch Links

An inter-switch link (ISL) is a link between two switches, E\_Port-to-E\_Port. The ports of the two switches automatically come online as E\_Ports after the login process finishes successfully. For more information on the login process, see [Understanding Fibre Channel Services](#).

You can expand your fabric by connecting new switches to existing switches. The following figure shows a new switch being added into an existing fabric. The thick red line is the newly formed ISL.

**Figure 80: New Switch Added to Existing Fabric**

When connecting two switches together, it is best practice that the following parameters are differentiated:

- Domain ID
- Switch name
- Chassis name

You must also verify the following fabric parameters are identical on each switch for a fabric to merge:

- R\_A\_TOV (Resource Allocation TimeOut Value)
- E\_D\_TOV (Error Detect TimeOut Value)
- Data Field Size
- Sequence Level Switching
- Disable Device Probing
- Suppress Class F Traffic
- Per-frame Route Priority

There are non-fabric parameters that must match as well, such as zoning. Some fabric services, such as management server, must match. If the fabric service is enabled in the fabric, the switch you are introducing into the fabric must also have it enabled.

## Buffer Credits

To prevent the dropping of frames in the fabric, a device can never send frames without the receiving device being able to receive them, so an end-to-end flow control is used on the switch. Flow control in Fibre Channel uses buffer-to-buffer credits, which are distributed by the switch. When all buffer-to-buffer credits are utilized, a device waits for a VC\_RDY or an R\_RDY primitive from the destination switch before resuming I/O. The primitive is dependent on whether you have R\_RDYs enabled on your switch using the `portCfgISLMode` command. When a device logs in to a fabric, it typically requests from two to sixteen buffer credits from the switch, depending on the device type, driver version, and configuration. This determines the maximum number of frames the port can transmit before receiving an acknowledgment from the receiving device.


For more information on how to set the buffer-to-buffer credits on an extended link, refer to [Buffer-to-Buffer Credits and Credit Recovery](#).

## Virtual Channels

Virtual channels create multiple logical data paths across a single physical link or connection. They are allocated their own network resources, such as queues and buffer-to-buffer credits. Virtual channel technology is the fundamental building block used to construct Adaptive Networking services. For more information on Adaptive Networking services, see [Optimizing Fabric Behavior](#).

Virtual channels are divided into three priority groups. P1 is the highest priority, which is used for Class F, F\_RJT, and ACK traffic. P2 is the next highest priority, which is used for data frames. The data virtual channels can be further prioritized to provide higher levels of Quality of Service. P3 is the lowest priority and is used for broadcast and multicast traffic. The example is illustrated in the figure below.

Quality of Service (QoS) is a traffic shaping feature available in Fabric OS. QoS allows the prioritization of data traffic based on the SID and DID of each frame. Through the use of QoS zones, traffic can be divided into three priorities: high, medium, and low, as shown in the figure below. The seven data virtual channels (VC8 through VC14) multiplex data frames based upon QoS zones when congestion occurs. For more information on QoS zones, see [QoS](#).


**Figure 81: Virtual Channels on a QoS-Enabled ISL**

## Gateway Links

A gateway merges SANs into a single fabric by establishing point-to-point E\_Port connectivity between two Fibre Channel switches that are separated by a network with a protocol, such as IP or SONET.

Except for link initialization, gateways are transparent to switches; the gateway simply provides E\_Port connectivity from one switch to another. The following figure shows two separate SANs, A-1 and A-2, merged together using a gateway.



**Figure 82: Gateway Link Merging SANs**

By default, switch ports initialize links using the Exchange Link Parameters (ELP) mode 1. However, gateways expect initialization with ELP mode 2, also referred to as ISL R\_RDY mode. Therefore, to enable two switches to link through a gateway, the ports on both switches must be set for ELP mode 2.

Any number of E\_Ports in a fabric can be configured for gateway links, provided these guidelines are followed:

- All switches in the fabric use the core PID format, as described in [Configuring a Link through a Gateway](#).
- The switches connected to both sides of the gateway are included when determining switch-count maximums.
- Extended links (those created using the Extended Fabrics licensed feature) are not supported through gateway links.

## Configuring a Link through a Gateway

1. Connect to the switch at one end of the gateway and log in using an account assigned to the admin role.
2. Enter the `portCfgISLMode` command.
3. Repeat Step 1 and Step 2 for any additional ports that are connected to the gateway.
4. Repeat this procedure on the switch at the other end of the gateway.

Example of enabling a gateway link on slot 2, port 3:

```
eCP:admin> portCfgISLmode 2/3, 1
```

```
Committing configuration...done.
ISL R_RDY Mode is enabled for port 3. Please make sure the PID
formats are consistent across the entire fabric.
```

## Routing Policies

By default, all routing protocols place their routes into a routing table. You can control the routes that a protocol places into each table and the routes from that table that the protocol advertises by defining one or more routing policies and then applying them to the specific routing protocol.

The routing policy is responsible for selecting a route based on one of three user-selected routing policies:

- Port-based routing
- Exchange-based routing
- Device-based routing

## Considerations for Routing Policies

Be aware of the following considerations when working with routing policies:

- Routing is handled by the FSPF protocol and routing policy.
- Each switch can have its own routing policy, and different policies can exist in the same fabric.

### ATTENTION

For most configurations, the default routing policy is optimal and provides the best performance. You should change the routing policy only if there is a significant performance issue, or a particular fabric configuration or application requires it.

## Displaying the Current Routing Policy

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aptPolicy` command with no parameters.

The current policy is displayed, followed by the supported policies for the switch.

In the following example, the current policy is device-based routing (2).

```
switch:admin> aptpolicy

Current Policy: 2
3 : Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy (FICON support only)
3: Exchange Based Routing Policy
```

## Port-Based Routing

The choice of routing path is based only on the incoming port and the destination domain. To optimize port-based routing, Dynamic Load Sharing (DLS) can be enabled to balance the load across the available output ports within a domain.

### NOTE

For FC routers only: When an FC router is in port-based routing mode, the backbone traffic is load-balanced based on the SID and DID. When an FC router is in exchange-based routing mode, the backbone traffic is load-balanced based on the SID, DID, and originator exchange identifier (OXID).

Whatever routing policy a switch is using applies to the VE\_Ports as well. For more information on VE\_Ports, refer to the *Brocade Fabric OS FCIP Administration Guide*.

## Exchange-Based Routing

The choice of routing path is based on the Source ID (SID), Destination ID (DID), and Fibre Channel originator exchange ID (OXID) optimizing path utilization for the best performance. Thus, every exchange can take a different path through the fabric. Exchange-based routing requires the use of the Dynamic Load Sharing (DLS) feature; when this policy is in effect, you cannot disable the DLS feature.

Exchange-based routing is also known as Dynamic Path Selection (DPS). For more information on DPS, see [Dynamic Path Selection](#).

## Device-Based Routing

Device-based routing optimizes routing path selection and utilization based on the Source ID (SID) and Destination ID (DID) of the path source and destination ports. As a result, every distinct flow in the fabric can take a different path through the fabric. Effectively, device-based routing works the same as exchange-based routing but does not use the OXID field, which helps to ensure that the exchanges between a pair of devices stay in order.

### NOTE

Device-based routing requires the use of DLS; when this policy is in effect, you cannot disable the DLS feature.

Device-based routing is also a form of DPS. For more information on DPS refer to [Dynamic Path Selection](#).

### NOTE

Device-based routing is supported in FICON environments, and in open environments only when FICON coexists.

## Dynamic Path Selection

DPS assigns communication paths between end devices in a fabric to egress ports in ratios proportional to the potential bandwidth of the ISL, ICL, or trunk group. When there are multiple paths to a destination, the input traffic is distributed across the different paths in proportion to the bandwidth available on each of the paths. This improves utilization of the available paths, thus reducing possible congestion on the paths. Every time there is a change in the network (which changes the available paths), the input traffic can be redistributed across the available paths. This is an easy and non-disruptive process when the exchange-based routing policy is engaged.

## Displaying a Dynamic Path Selection Group for Reachable Domains

You can use the `portChannelShow` command to display a dynamic path selection group for one or all reachable domains.

- Use the `portChannelShow` command to display the port channels from a domain to all the reachable domains on the switch. For example:

```
switch:admin> portchannelshow

6 domain(s) in the fabric; Local Domain ID: 4

Domain: 1
Name: sw0
WWN: 10:00:00:05:33:c1:26:00
Port Channel: None

Domain: 2
```

```
Name: DCX_35_F_128
WWN: 10:00:00:05:1e:38:e5:23
Port Channel:
 Ports: 384, 385, 386, 387, 400, 401, 402, 403,
 417, 418, 419, 432, 433, 434, 435
```

```
Domain: 3
Name: SW_122_F_128
WWN: 10:00:00:05:1e:9b:10:5b
Port Channel:
 Ports: 111, 248
```

```
Domain: 5
Name: SW_65_F128
WWN: 10:00:00:05:1e:5c:f6:fd
Port Channel:
 Ports: 384, 385, 386, 387, 400, 401, 402, 403,
 417, 418, 419, 432, 433, 434, 435
```

```
Domain: 6
Name: SW_121_F_128
WWN: 10:00:00:05:1e:9c:32:cc
Port Channel: None
```

- Use the `portChannelShow [remote-domain-id | switch-name | switch-wwn]` command to display the port channels from the switch where this command is executed to the specified switch or domain. For example:

```
switch:admin> portchannelshow 2 (OR)
switch:admin> portchannelshow 10:00:00:05:1e:38:e5:23 (OR)
switch:admin> portchannelshow DCX_35_F_128
```

```
Domain: 2
Name: DCX_35_F_128
WWN: 10:00:00:05:1e:38:e5:23
Port Channel:
 Ports: 384, 385, 386, 387, 400, 401, 402, 403,
 417, 418, 419, 432, 433, 434, 435
```

## Route Selection

Selection of specific routes can be dynamic, so that the router can constantly adjust to changing network conditions; or it may be static, so that data packets always follow a predetermined path.

## Dynamic Load Sharing

The Fabric OS Dynamic Load Sharing (DLS) feature for dynamic routing path selection is required by the exchange-based and device-based routing policies. When using these policies, DLS is enabled by default and cannot be disabled. In other words, you cannot enable or disable DLS when the exchange-based routing policy is in effect.

When the port-based policy is in force, you can enable DLS to optimize routing. When DLS is enabled, it shares traffic among multiple equivalent paths between switches. DLS recomputes load sharing when any of the following occurs:

- A switch boots up
- An E\_Port goes offline and online
- An EX\_Port goes offline
- A device goes offline

## Setting DLS

Use the following procedure to set DLS.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `dlsShow` command to view the current DLS setting.

One of the following messages appears:

- DLS is set with Lossless disabled indicates that DLS is turned on.
- DLS is not set with Lossless disabled indicates that DLS is turned off.
- DLS is set with Lossless enabled. DLS is enabled with the Lossless feature. Load sharing is recomputed with every change in the fabric, and existing routes can be moved to maintain optimal balance. In Lossless mode, no frames are lost during this operation.
- It also displays the state of Two-hop lossless setting. For more details, refer to the *Brocade Fabric OS Command Reference Manual*.

3. Enter the `dlsSet` command to enable DLS or enter the `dlsReset` command to disable it.

Example of setting and resetting DLS.

```
switch:admin> dlsshow
DLS is not set with Lossless disabled
```

```
switch:admin> dlsset
switch:admin> dlsshow
DLS is set with Lossless disabled
```

```
switch:admin> dlsreset
switch:admin> dlsshow
DLS is not set with Lossless disabled
```

## Frame Order Delivery

The order in which frames are delivered is maintained within a switch and determined by the routing policy in effect. The frame delivery behaviors for each routing policy are:

- Port-based routing  
All frames received on an incoming port destined for a destination domain are guaranteed to exit the switch in the same order in which they were received.
- Exchange-based routing  
All frames received on an incoming port for a given exchange are guaranteed to exit the switch in the same order in which they were received. Because different paths are chosen for different exchanges, this policy does not maintain the order of frames across exchanges.
- Device-based routing  
All frames received on an incoming port for a given pair of devices are guaranteed to exit the switch in the same order in which they were received.

If even one switch in the fabric delivers out-of-order exchanges, then exchanges are delivered to the target out of order, regardless of the policy configured on other switches in the fabric.

**NOTE**

Some devices do not tolerate out-of-order exchanges; in such cases, use the port-based routing policy.

In a stable fabric, frames are always delivered in order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for example, if a link goes down), traffic is rerouted around the failure, and some frames could be delivered out of order. Most destination devices tolerate out-of-order delivery, but some do not.

By default, out-of-order frame-based delivery is allowed to minimize the number of frames dropped. Enabling in-order delivery (IOD) guarantees that frames are either delivered in order or dropped. You should only force in-order frame delivery across topology changes if the fabric contains destination devices that cannot tolerate occasional out-of-order frame delivery.

## Forcing In-Order Frame Delivery across Topology Changes

Use the following procedure to force in-order frame delivery across topology changes.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `iodSet` command.

**NOTE**

The `iodSet` command can cause a delay in the establishment of a new path when a topology change occurs; use it with care.

3. Confirm the in-order delivery has been set by entering the `iodShow` command.

## Restoring Out-of-Order Frame Delivery across Topology Changes

Use the following procedure to restore out-of-order frame delivery across topology changes.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `iodReset` command.

## Enabling Frame Viewer

The Frame Viewer application allows you to view the contents of discarded frames, which can be recorded at up to 40 frames per second per ASIC.

To enable Frame Viewer, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter `frameLog --enable -type type`.

The following table lists the supported values for `-type`. If the `-type all` option is specified, then all seven discard frame type options are enabled and frames will be recorded when they are discarded.

**Table 140: Supported Values for the -type Option.**

| Options                | Reason                                    |
|------------------------|-------------------------------------------|
| <code>timeout</code>   | Due to a timeout.                         |
| <code>du</code>        | Due to the destination being unreachable. |
| <code>unroute</code>   | Due to an unrouteable frame.              |
| <code>type1miss</code> | Due to a type 1 miss.                     |
| <code>type2miss</code> | Due to a type 2 miss.                     |
| <code>type6miss</code> | Due to a type 6 miss.                     |

| Options      | Reason                  |
|--------------|-------------------------|
| hardzonemiss | Due to ACL check miss.  |
| all          | Due to all the reasons. |

3. Use the `frameLog --status` command to view the configuration. The following is a typical example of the output of this command:

```
switch:admin> framelog --status
Service Status: Enabled
Enabled Disc Frame Types: timeout unrout
```

## Using Frame Viewer to Understand Why Frames Are Dropped

When a frame is unable to reach its destination because of a timeout, is unroutable, or has an unreachable destination, or type1miss/type2miss/type6miss, the frame is discarded. You can use Frame Viewer to find out which flows contained the dropped frames, which in turn can help you determine the applications that might be impacted. Frame Viewer allows you to see the exact time (within one second) that the frames were dropped.

With the `frameLog` command, you can view up to 40 discarded frames per ASIC per second for 1200 seconds. You can filter this view output using a number of fields so you are viewing only those frames of interest.

### NOTE

From Fabric OS 9.1.0, Frame Viewer also shows the frames dropped over the FCIP links of an extension switch (orts) for timeout and unroutable reasons.

To display information about discarded frames, complete the following steps. This assumes that the `frameLog` application has been enabled previously.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter `framelog --show`.

Refer to the *Brocade Fabric OS Command Reference Manual* for additional information on using the `frameLog` command, including resetting (clearing) the frame log.

The following example shows a typical output of the `frameLog --show` command:

```
switch:admin> framelog --show
=====
Wed Aug 03 19:54:07 UTC 2016
=====
Timestamp |Tx Port|Rx Port|SID |DID |SFID|DFID|Src Entity Id|Dst Entity Id|Type
Count
Aug 03 2016 19:53:28	--	8	0x520800	0x101010	128	128	vm_2	vm_2	unrout
10
										e
Aug 03 2016 19:52:26	--	8	0x520800	0x101010	128	128	N/A	N/A	unrout
10
										e
```

### NOTE

There may occasionally be frame types listed as *unknown* in the output; these are usually frames that have been corrupted in transit, and so the `frameLog` command cannot identify their type.

## Displaying Discarded Frames by a Backend Port in Frame Viewer

When viewing information about discarded frames, you can filter the results by specifying that the TX port or RX port of displayed frames should be a backend port.

Individual backend ports cannot be specified; only the quality of being a backend port can be specified.

1. Connect to the switch and log in using an account with admin permissions.
2. Use the `frameLog --show` command, followed by `-txport` , `-rxport` , or both, and specify `-1` for the port number.

```
switch:admin> framelog --show -txport -1 -rxport -1
```

Specify `-1` for fixed-port switches and `-1/-1` for Brocade directors. These indicate *any backend port*.

### NOTE

Frame discards can be logged as audit messages using the Fabric OS syslog facility.

## Using Frame Viewer for Fabric Notifications

The Frame Viewer application can also be configured to generate notifications to devices when a SCSI-FCP command frame has been discarded using the FPIN action.

To enable the FPIN action:

1. Connect to the switch and log in as an admin.
2. Enter the `frameLog --action FPIN` command.

Discards of SCSI-FCP command frames will result in an FPIN having a Delivery Descriptor notification to the command originator. The notification is only sent if the discard reason is one of the following types: timeout, du, unroute, and type2miss.

## Lossless Dynamic Load Sharing on Ports

Lossless Dynamic Load Sharing (LDLS) allows you to rebalance port paths without causing I/O failures. For devices where in-order delivery (IOD) of frames is required, you can set IOD separately. You can use this feature with the following hardware:

- Brocade G720
- Brocade G610
- Brocade G620
- Brocade G630
- Brocade X7 Directors with the following blades:
  - Brocade FC64-48 port blade
  - Brocade FC32-X7-48 port blade
- Brocade X6 Directors with the following blades:
  - Brocade FC32-48 port blade
  - Brocade FC32-64 port blade

Lossless DLS must be implemented along the path between the target and the initiator. You can use Lossless DLS on ports connecting switches to perform the following functions:

- Eliminate dropped frames and I/O failures by rebalancing the paths going over the ISLs whenever there is a fabric event that might result in suboptimal utilization of the ISLs.
- Eliminate the frame delay caused by establishing a new path when a topology change occurs.

Lossless mode means no frame loss occurs during a rebalance, and it only takes effect if DLS is enabled. Lossless DLS can be enabled on a fabric topology to have zero frame drops during rebalance operations. If the end device also requires



the order of frames to be maintained during the rebalance operation, IOD must be enabled. However, this combination of Lossless DLS and IOD is supported only in specific topologies, such as in an FICON environment.

You can disable or enable IOD when Lossless DLS is enabled. You can also choose between exchange- or port-based policies with Lossless DLS. The following events cause a rebalance:

- Adding a redundant E\_Port
- Adding a slave E\_Port
- Removing the last E\_Port from a neighbor domain (However, frame loss occurs on traffic flows to this port.)
- Removing an F\_Port (However, frame loss occurs on traffic flows to this port.)

Lossless DLS does the following whenever paths need to be rebalanced:

1. Pauses ingress traffic by not returning credits. Frames that are already in transit are not dropped.
2. Applies the results of the rebalance calculations.
3. *If IOD is enabled*, waits for sufficient time for frames already received to be transmitted. This is needed to maintain IOD.
4. Resumes traffic.

The following table shows the effect of frames when you have a specific routing policy turned on with IOD.

**Table 141: Combinations of routing policy and IOD with Lossless DLS enabled**

| Policy         | IOD      | Rebalance Result with Lossless DLS Enabled                                                             |
|----------------|----------|--------------------------------------------------------------------------------------------------------|
| Port-based     | Disabled | No frame loss, but out-of-order frames may occur.                                                      |
| Port-based     | Enabled  | No frame loss and no out-of-order frames. Topology restrictions apply. Intended for FICON environment. |
| Exchange-based | Disabled | No frame loss, but out-of-order frames may occur.                                                      |
| Exchange-based | Enabled  | No frame loss and no out-of-order frames. Topology restrictions apply. Intended for FICON environment. |
| Device-based   | Disabled | No frame loss, but out-of-order frames may occur.                                                      |
| Device-based   | Enabled  | No frame loss and no out-of-order frames. Topology restrictions apply. Intended for FICON environment. |

#### NOTE

Not only the link cost has influence on DLS, but the order of ports that become online and the speeds of each E\_Port or F\_Port also have an impact on how the switch calculates or chooses the path using DLS.

## Lossless Core

Lossless core works with the default configuration of the Brocade X7 and X6 Directors to prevent frame loss during a core blade removal and insertion. This feature is on by default and cannot be disabled. Lossless core has the following limitations:

- Only supported with IOD disabled, which means lossless core cannot guarantee in-order delivery of exchanges
- ICL limitations
- Traffic flow limitations

## ICL Limitations

If ICL ports are connected during a core blade removal, it is equivalent to removing external E\_Ports, which may cause I/O disruption on the ICL ports that have been removed.

If ICL ports are connected during a core blade insertion, it is equivalent to adding external E\_Ports, which may cause I/O disruption because of reroutes. Lossless DLS, if enabled, takes effect to prevent I/O disruption.

**NOTE**

Before you replace the core blades, enable lossless so that the `portDecom` command can decommission the E\_Ports losslessly and bring back up the E\_Ports losslessly.

## Configuring Lossless Dynamic Load Sharing

You can configure Lossless DLS on either a switch- or chassis-wide basis by using the `dlsset` command to specify that no frames are dropped while rebalancing or rerouting traffic.

**NOTE**

For more information regarding DLS or E\_Port balancing, refer to the `dlsset` command in the *Brocade Fabric OS Command Reference Manual*.

To configure Lossless Dynamic Load Sharing, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the appropriate `dlsset` command to enable or disable Lossless Dynamic Load Sharing.

```
switch:admin> dlsset --enable -lossless
switch:admin> dlsset --disable -lossless
```

**NOTE**

To enable lossless on the whole chassis when multiple logical switches exist, you must enable lossless on each logical switch independently or use the `fosexec --fid all -cmd dlsset` command.

## Lossless Dynamic Load Sharing in Virtual Fabrics

Enabling LDLS is optional on logical switches in Virtual Fabrics. If you enable this feature, it must be on a per-logical switch basis and can affect other logical switches in the fabric.

### How DLS Affects Other Logical Switches in the Fabric

On a Brocade X7 and X6 platform, logical switch 1 consists of ports 0 through 5 in slot 1. Logical switch 2 consists of ports 6 through 10 in slot 1. The Lossless DLS feature is enabled on logical switch 1. Because ports 0 through 5 in slot 1 belong to a logical switch where Lossless DLS is enabled, the traffic in logical switch 2 is affected whenever traffic for logical switch 1 is rebalanced.

**ATTENTION**

Although Lossless DLS is enabled for a specific logical switch, you must have chassis-level permissions to use this feature.

The effect on logical switch 2 is based on the configuration on logical switch 2:

- If logical switch 2 has IOD enabled (`iodSet` only), IOD is enforced.
- If logical switch 2 has Lossless DLS enabled, traffic is paused and resumed.
- If logical switch 2 has no IOD (`iodReset`), traffic is paused and resumed.

To avoid this behavior, define your logical switches as follows:

- Define logical switches that require Lossless DLS at the blade boundary.
- Define logical switches that require Lossless DLS only using supported blades. For example, do not use blades that support IOD, but do not support Lossless DLS.

For more information on Virtual Fabrics and chassis-level permissions, refer to [Managing Virtual Fabrics](#).

## Two-Hop Lossless DLS Route Update

Certain fabric topology changes may affect the Shortest Path First results that lead to rebalance operations, and thus change the paths that are utilized. During these rebalance operations, even if the lossless DLS setting is enabled, route updates that are caused by fabric shortest path changes may result in I/O disruptions and/or out-of-order frame delivery. Route updates are not done globally across all switches at the same time. A new path may be used before the later switches are ready.

Lossless DLS supports rebalance operations after events that affect redundancy within a hop but do not cause fabric shortest path changes. Lossless DLS can only ensure that the local switch updates are performed in the required manner. Lossless DLS cannot ensure that a downstream domain has completed its rebalancing operations prior to a new path being used. The Two-hop enhancement to Lossless DLS extends the lossless route update capabilities to support the addition of a new hop between switches.

When a new link is added between two switches where one was not there previously, the new link is likely to cause changes to the shortest path results. The Two-hop enhancement to Lossless DLS ensures that the route updates occur on the local switches of the new link prior to the rest of the fabric. For devices on those two switches, when they attempt to start using the new link, Two-hop Lossless DLS ensures that the routes are ready for use prior to their route updates.

For a switch that is directly connected to one of the two switches with the new link, when it performs the rebalance operation, its device traffic is shifted to the new path in a lossless manner. If the new path to which traffic is being shifted to is up to two hops (one of which includes the new link), Two-hop Lossless DLS ensures that the routes along that two-hop path are ready for use.

### NOTE

If a fabric shortest path that includes the new link is longer than two hops, when the switches perform their rebalancing operations, the traffic may still experience I/O disruptions, out-of-order frame delivery, or both.

When the new hop has been established, all later route updates use the standard lossless DLS capabilities.

- Two-hop Lossless DLS has the same platform and ISL technology restrictions as the Lossless DLS feature.
- Two-hop Lossless DLS is not supported with any hops to or through an FC router-related domain (front and translate domains).
- Two-hop Lossless DLS is not supported over logical ISLs (LISLs).
- You must manually enable this feature on all switches to get the desired results.

## Configuring Two-Hop Lossless DLS

Use the following commands to configure the Two-hop Lossless DLS feature.

**Table 142: Commands Used to Configure the Two-Hop Lossless DLS Feature**

| Command               | Usage                                                                                                  |
|-----------------------|--------------------------------------------------------------------------------------------------------|
| <code>dlsshow</code>  | Displays the state of the Two-hop Lossless DLS state when it has been enabled.                         |
| <code>dlsreset</code> | Resets the state of DLS in port-based routing.                                                         |
| <code>dlset</code>    | Sets the lossless, two-hop, and eportbal states. Enables or disables the Two-hop Lossless DLS feature. |

The following example shows the use of the `--help` option of the `dlset` command to display the possible uses of the command. The syntax for enabling and disabling the Two-hop Lossless DLS feature are highlighted.

```
switch:admin> dlset --help
dlset usage:
```

```
dlssset
dlssset --enable -lossless
dlssset --disable -lossless
dlssset --enable -twohop
dlssset --disable -twohop
dlssset --enable -eportbal
dlssset --disable -eportbal
dlssset --rebalance
dlssset --rebalance -all
```

The following example shows the use of the `dlssshow` command to display the state of the Two-hop Lossless DLS feature when it has been enabled. If the feature was disabled, the original default output would be displayed.

```
switch:admin> dlssshow
DLS is set with Two-hop Lossless enabled
E_Port Balance Priority is not set
```

## Using FC-FC Routing to Connect Fabrics

---

This section describes the FC-FC routing service, which provides Fibre Channel routing between two or more fabrics without merging those fabrics.

Using FC-FC routing, you can share tape drives across multiple fabrics without the administrative problems, such as change management, network management, scalability, reliability, availability, and serviceability, that might result from merging the fabrics.

The different routing terminologies are:

- *FC routing* is only in a single fabric (Layer 2 routing).
- *FC-FC routing* is routing between two fabrics (Layer 3 routing).

FC-FC routing supports connectivity between the following types of fabrics:

- Fabric OS and Fabric OS

A Fibre Channel router (FC router) is a switch running the FC-FC routing service. The FC-FC routing service can be simultaneously used as an FC router and as a SAN extension over wide area networks (WANs) using FCIP.

You can set up the QoS traffic prioritization over FC routers. See [QoS](#) for information about QoS and instructions for setting traffic prioritization over an FC router.

## License Requirements for FC-FC Routing

The following requirements for FC-FC routing license are:

- A software license might be required for FC-FC routing, depending on the types of fabrics that are connected.
- To enable EX\_Ports on the Brocade 7810 switch, an upgrade combo license is required to install. This adds the Integrated Routing Ports on Demand license on the switch
- The Integrated Routing license or Integrated Routing Port on Demand license is required for FC-FC routing between Fabric OS fabrics.
- To enable EX\_Ports on the Brocade G620 switch or any Gen 7 64G platforms, the Integrated Routing license must be installed.
- To enable EX\_Ports on Gen 6 32G switches other than the Brocade G620 switch, the Integrated Routing Port on Demand License must be installed.
- Enabling the Integrated Routing license and capability does not require a switch reboot.
- To enable EX\_Ports on Gen 7 switches, the Integrated Routing License is required for Gen 7 platforms.

### NOTE

All FC routers in a backbone fabric must either have the Integrated Routing license or no Integrated Routing License. You must not mix licensed and unlicensed FC routers in the backbone fabric.

## Supported Platforms for FC-FC Routing

FC-FC routing is supported on the following Gen 7 (64G) Fabric OS platforms:

- Brocade G730 switch
- Brocade G720 switch
- Brocade X7 Directors:

- 64G port blades (FC64-48)
- ICL ports on the core blades; EX\_Port on ICL is supported only on Virtual fabric enabled FCR switches with a base switch and when the remote port is also a core blade ICL port.
- 32G port blades (FC32-X7-48)

FC-FC routing is supported on the following Gen 6 (32G) Fabric OS platforms:

- Brocade G620 switch
- Brocade G630 switch
- Brocade 7810 Extension Switch, with the following condition:
  - VEX\_Ports are not supported.
- Brocade X6 Directors:
  - FC32-48 port blades
  - FC32-64 port blades
  - SX6 extension blade
    - VEX\_Ports are not supported.
    - EX\_Ports are supported.
  - ICL ports on the core blades; EX\_Port on ICL is supported only on Virtual Fabric enabled FCR switches with a base switch and when the remote port is also a core blade ICL port.

#### NOTE

- Booting from SAN FCR is not a supported feature if devices are not imported in the backbone.
- If a switch is VF enabled, then the EX\_Ports are supported only in the base switch. The 7810 extension switch does not support the base switch.

## Configuring EX\_Ports on QSFP Ports of a Brocade G630 Switch


IFLs can be created between Brocade G630 switches and other supported switches in integrated routing topologies. The trunk port groups on the G630 switch explained in [Table](#) apply for IFLs as well. IFL connections can be formed from QSFP to QSFP ports or from QSFP to SFP ports using Brocade qualified break-out cables.

Guidelines for QSFP IFL connections:

- You must select the QSFP ports from the two separate QSFP port groups of local G630 switch. Consider switch 1 as the FC Router in the backbone fabric. Switch 2 and switch 3 are edge switches in two separate edge fabrics 2 and 3.

#### NOTE

Consider all the three switches in the following figure are Brocade G630 switches.

**Figure 83: QSFP to QSFP IFLs Connectivity**

- IFL connection between switch 1 (FC Router) and switch 2 (edge fabric 2) using ports 96 to 99 (Group 1) and ports 112 to 115 (Group2).
- IFL connection between switch 1 (FC Router) and switch 3 (edge fabric 3) using ports 104 to 107 (Group 1) and ports 120 to 123 (Group 2).
- QSFP to QSFP connection from a G630 switch to QSFP ports on the X7 Director or X6 Director core routing blades are not supported.
- If FC Router connects to an edge fabric using both SFP ports and one of QSFP port group, then the same edge fabric need not be connected to another QSFP port group.
- FC Router detects any inconsistency in the above guidelines during QSFP IFL connectivity and displays a RASLOG message.

## Supported Configurations for FC-FC Routing

Consider the following points for FC-FC routing:

- FC Router connected to a Fabric OS nonsecured edge fabric is supported.
- FC Router connected to a Fabric OS secured edge fabric is supported.
- FC Router connectivity to a Network OS edge fabric is not supported.
- In configurations with two backbone fabrics connected to the same edge fabric, routing is not supported between edge fabrics that are not directly attached to the same backbone fabric.
- Routing over multiple backbone fabrics is a multi-hop topology and is not allowed.

**NOTE**

When multiple EX ports exist between the edge and Backbone fabric, they are considered equal-cost routes. FSPF does not apply from edge to edge. Take this consideration into account when designing FCR fabrics.

**NOTE**


EX\_Port does not support connectivity with a diagnostic port (D\_Port). In case a D\_Port is connected with an EX\_Port, the D\_Port is segmented, and the EX\_Port does not remain online.

## Fibre Channel Routing Concepts


Fibre Channel routing introduces the following concepts:

- Fibre Channel route (FC Router)  
A switch running the FC-FC routing service. Refer to [Supported Platforms for FC-FC Routing](#) for the list of platforms that can be FC routers.
- Backbone fabric  
A backbone fabric is an intermediate network that connects one or more edge fabrics. In a SAN, the backbone fabric consists of at least one FC router and possibly a number of Fabric OS-based Fibre Channel switches.
- Inter-fabric link (IFL)  
The link between an E\_Port and EX\_Port, is called an *inter-fabric link* (IFL). You can configure multiple IFLs from an FC router to an edge fabric.  
The following figure shows a metaSAN consisting of three edge fabrics connected through a Brocade backbone with inter-fabric links.



**Figure 84: A metaSAN with Inter-Fabric Links**

- **EX\_Port**  
An EX\_Port functions similarly to an E\_Port, but terminates at the switch and do not propagate fabric services or routing topology information from one edge fabric to another. Refer to the *Brocade Fabric OS FCIP Administration Guide* for details about VE\_Ports.
- **Edge fabric**  
An edge fabric is a Fibre Channel fabric with targets and initiators connected through the supported platforms.
- **Logical SANs (LSANs)**  
An LSAN is defined by zones in two or more edge or backbone fabrics that contain the same devices. You can create LSANs that span fabrics. These LSANs enable Fibre Channel zones to cross physical SAN boundaries without merging the fabrics while maintaining the access controls of zones.  
An LSAN device can be a *physical device*, meaning that it physically exists in the fabric, or it can be a *proxy device*. The following figure shows a metaSAN with a backbone consisting of one FC router connecting hosts in edge fabrics 1 and 3 with storage in edge fabric 2 and the backbone fabric through the use of LSANs. Three LSAN zones allow device sharing between the backbone fabric and edge fabric 1, between edge fabric 1 and edge fabric 2, and between edge fabric 2 and edge fabric 3.

**Figure 85: A metaSAN with Edge-to-Edge and Backbone Fabrics and LSAN Zones**

- **Proxy device**  
A proxy device is a virtual device imported into a fabric by a Fibre Channel router, and represents a real device on another fabric. It has a name server entry and is assigned a valid port ID. When a proxy device is created in a fabric, the real Fibre Channel device is considered to be imported into this fabric. The presence of a proxy device is required for inter-fabric device communication. See [Proxy Devices](#) for additional information about proxy devices.
- **Proxy PID**  
A proxy PID is the port ID (PID) of the proxy device. The proxy device appears to the fabric as a real Fibre Channel device, has a name server entry, and is assigned a valid port ID. The port ID is relevant only on the fabric in which the proxy device has been created.
- **Fabric ID (FID)**  
Every EX\_Port uses the fabric ID (FID) to identify the fabric at the opposite end of the inter-fabric link. The FID for every edge fabric must be unique from the perspective of each backbone fabric.
  - If multiple EX\_Ports are attached to the same edge fabric, they must be configured with the same FID.
  - If EX\_Ports are attached to different edge fabrics, they must be configured with a unique FID for each edge fabric.

**NOTE**

Backbone fabrics that share connections to the same edge fabrics must have unique backbone fabric IDs.

If two different backbone fabrics are connected to the same edge fabric, the backbone fabric IDs must be different, but the edge fabric IDs must be the same. If you configure the same fabric ID for two backbone fabrics that are connected to the same edge fabric, a RASLog message displays a warning about fabric ID overlap.


You can optionally assign an alias name to the FID.

- **MetaSAN**  
A metaSAN is the collection of all SANs interconnected with Fibre Channel routers.

A simple metaSAN can be constructed using an FC router to connect two or more separate fabrics. Additional FC routers can be used to increase the available bandwidth between fabrics and to provide redundancy.

The following figure shows a metaSAN consisting of a host in Edge SAN 1 connected to storage in Edge SAN 2 through a backbone fabric connecting two FC routers.

**Figure 86: Edge SANs Connected through a Backbone Fabric**



- Phantom domains

A phantom domain is a domain emulated by the FC Router. The FC router can emulate two types of phantom domains: front phantom domains and translate phantom domains. For detailed information about phantom domains, see [Phantom Domains](#).

## Proxy Devices

An FC router achieves inter-fabric device connectivity by creating proxy devices (hosts and targets) in attached fabrics that represent real devices in other fabrics.

For example, a host in Fabric 1 can communicate with a target in Fabric 2 as follows:

- A proxy target in Fabric 1 represents the real target in Fabric 2.
- Likewise, a proxy host in Fabric 2 represents the real host in Fabric 1.

**NOTE**

When the properties of the same physical device are updated on the local edge, the name server properties of a proxy device are updated in the remote edge.


The following NS entries are not seen as part of a proxy FID entry for devices that are imported from a B2E setup or E2E setup:

- Connected through AG: No
- Real device behind AG: No
- FCoE: No
- Connected through FC-LAG: No

The host discovers and sends Fibre Channel frames to the proxy target. The FC router receives these frames, translates them appropriately, and then delivers them to the destination fabric for delivery to the target.

The target responds by sending frames to the proxy host. Hosts and targets are exported from the edge SAN to which they are attached and, correspondingly, imported into the edge SAN reached through Fibre Channel routing. The following figure illustrates this concept.

**Figure 87: MetaSAN with Imported Devices**



## FC-FC Routing Topologies

The FC-FC routing service provides two types of routing:

- *Edge-to-edge*: Occurs when devices in one edge fabric communicate with devices in another edge fabric through one or more FC routers.
- *Backbone-to-edge*: Occurs when devices in the FC routers communicate with devices in an edge fabric—known as a *backbone fabric*—through E\_Ports.

A backbone fabric can be used as a transport fabric that interconnects edge fabrics. FC routers also enable hosts and targets in edge fabrics to communicate with devices in the backbone fabric, known as *backbone-to-edge routing*. From the perspective of the edge fabric, the backbone fabric is like any other edge fabric. For the edge fabric and backbone fabric devices to communicate, the shared devices must be presented to each other's native fabric.

To do so, at least one translate phantom domain is created in the backbone fabric. This translate phantom domain represents the entire edge fabric. The shared physical devices in the edge have corresponding proxy devices on the translate phantom domain.

Each edge fabric has one and only one translate phantom domain to the backbone fabric. The backbone fabric device communicates with the proxy devices whenever it needs to contact the shared physical devices in the edge. The FC-FC routing service receives the frames from the backbone switches destined to the proxy devices, and redirects the frames to the actual physical devices. When connected to edge fabrics, the translate phantom domain can never be the principal switch of the backbone fabric. Front domains are not created, only translate phantom domains are created in the backbone fabric.

Devices are exported from the backbone fabric to one or more edge fabrics using LSANs. See [LSAN Zone Configuration](#) for more information.

## Phantom Domains


A phantom domain is a domain created by the Fibre Channel router. The FC router creates two types of phantom domains: front phantom domains and translates phantom domains.

A *front phantom domain*, or *front domain*, is a domain that is projected from the FC router to the edge fabric. There is one front phantom domain from each FC router to an edge fabric, regardless of the number of EX\_Ports connected from that router to the edge fabric. Another FC router connected to the same edge fabric projects a different front phantom domain.

A *translate phantom domain*, also referred to as *translate domain* or *xlate domain*, is a router virtual domain that represents an entire fabric. The EX\_Ports present xlate domains in edge fabrics as being topologically behind the front domains; if the xlate domain is in a backbone fabric, it is topologically present behind the FC router because there is no front domain in a backbone fabric.


If an FC router is attached to an edge fabric using an EX\_Port, it creates xlate domains in the fabric corresponding to the imported edge fabrics with active LSANs defined. If you import devices into the backbone fabric, an xlate domain is created in the backbone device in addition to the one in the edge fabric.

The following figure shows a sample physical topology with four FC routers in a backbone fabric and four edge fabrics that are connected to the FC routers.

**Figure 88: Sample Topology (Physical Topology)**

The following figure shows a phantom topology for the physical topology shown in [Figure](#). In this figure, the dashed lines and shapes represent the phantom topology from the perspective of Fabric 1. Fabrics 2 and 3 also exhibit phantom topologies, but they are not shown in this example. In this figure, note the following:

- Front domain 1 and Front domain 2 are front domains for EX\_Ports connecting to Fabric 1. There is one front domain for each FC router that is connected to Fabric 1.
- Xlate domain 1 and Xlate domain 2 represent Fabrics 2 and 3, respectively. No xlate domain is created for Fabric 4 because there are no LSAN devices in Fabric 4.
- Target 1', Target 2', and Target 3' are proxy devices for Target 1, Target 2, and Target 3, respectively.

**Figure 89: EX\_Port Phantom Switch Topology**

All EX\_Ports connected to an edge fabric using the same xlate domain ID for an imported edge fabric; this value persists across switch reboots and fabric reconfigurations.

If you lose connectivity to the edge fabric because of link failures or the IFL being disabled, xlate domains remain visible. This prevents unnecessary fabric disruptions caused by xlate domains repeatedly going offline and online due to corresponding IFL failures. To remove the xlate domain from the backbone, see [Identifying and Deleting Stale Xlate Domains](#).

The combination of front domains and xlate domains allows routing around path failures, including path failures through the routers. The multiple paths to an xlate domain provide additional bandwidth and redundancy.

There are some differences in how the xlate domain is presented in the backbone fabric. The backbone xlate domains are topologically connected to FC routers and participate in FC-FC routing protocol in the backbone fabric. Front domains are not needed in the backbone fabric. As in the case of an xlate domain in an edge fabric, backbone fabric xlate domains provide the additional bandwidth and redundancy by being able to present themselves as connected to single or multiple FC routers with each FC router capable of connecting multiple IFLs to edge fabrics.

Use the `fcxlateconfig` command to display or assign a preferred domain ID to a translate domain or, in some scenarios, to prevent the creation of an unnecessary xlate domain.

## Identifying and Deleting Stale Xlate Domains

If a remote edge fabric goes unreachable, the xlate domains created in other edge fabrics for this remote edge fabric are retained and not removed unless there is any disruption in the local edge fabric or unless you remove them manually.

You can use the `fcrlxlateconfig` command to identify and remove these stale xlate domains without disrupting the fabric.

1. Connect to the switch or director and log in using an account with admin permissions.
2. Enter `fcrlxlateconfig --show` to identify any stale xlate domains.
3. Enter `fcrlxlateconfig --del` to delete the stale xlate domains.

```
switch:admin> fcrlxlateconfig --show stalexd
Imported FID Stale XD Owner Domain

 012 002 007 (this FCR)

switch:admin> fcrlxlateconfig --del stalexd 12 2
Xlate domain 2 is deleted
```

## Domain ID Range for Front and Translate Phantom Domains

To avoid FC router phantom domains taking over the real switch domain ID, the principal switch assigns the domain ID based on the following rules:

- The front domain ID for an FC router is assigned within the range 160 through 199.
- The translate domain ID for an FC router is assigned within the range 200 through 239.
- If the preferred domain ID is not configured on an FC router for front or translate domain, the FC router requests 160 for the front domain and 200 for the translate domain. If the preferred domain ID is configured on the FC router for either the front or translates domain, the FC router still requests for the front domain ID of 160 and the translate domain ID of 200 at the discretion of the principal switch in the edge fabric. Depending on, whether there are any domain ID conflicts in the fabric, switches can either be merged with the fabric using a different domain ID (front or translate domain ID) or they are segmented from the fabric.
- If preferred domain ID is configured on an FC router for front domain, translate domain, or both, the FC router requests the preferred domain ID. Phantom domain IDs are stored persistently and used in request domain\_ID (RDI) request. To utilize the new range (160 to 239), do one of the following tasks:

### Option 1:

- a. Convert the EX\_Port to a non-FC router port by removing the existing Xlate configuration using the `fcrlxlateconfig` command.
- b. Apply the default configuration, and then configure the EX\_Ports again and remove the existing Xlate configuration using the `fcrlxlateconfig` command..

### Option 2:

- a. Disable all EX\_Ports in the same edge fabric and persistent XD feature, and then reset the persistent phantom domain using the `fcrlconfigure --resetphantomdomain` command.

```
switch:admin> fcrlconfigure --resetphantomdomain
This operation will reset all the phantom domain to be default range
Do you want to continue (Y/N):y
Phantom Domain IDs were successfully reset to default range
```

Using the `-force` option to force the reset phantom domain.

```
switch:admin> fcrlconfigure --resetphantomdomain -force
Phantom Domain IDs were successfully reset to default range
```

- b. Enable all the EX\_Ports.



**Table 143: Front and Translate Domain IDs**

| Firmware Version in FC Router | Front Domain ID Requested | Translate Domain ID Requested |
|-------------------------------|---------------------------|-------------------------------|
| Fabric OS 8.1.0 and later     | 160                       | 200                           |

## FC Router Authentication

A Brocade FC router is capable of forming a secure link across fabrics. The EX\_Port-enabled router exchanges DH-CHAP information with the edge fabric to enable authentication.

Note that while setting secret keys in the edge switch, use the front phantom WWN as the remote switch WWN in the edge fabric. The front phantom domain's WWN is available through the `portCfgExport` command of the EX\_Port connecting to the edge fabric. The FC router switch should use the edge switch's WWN to configure the secret keys. See [Secret key pairs for DH-CHAP](#) for more details.

FC-FC routing behaves passively to the authentication requests received from edge fabric switches. An FC router never initiates authentication on an EX\_Port and only responds to the edge fabric requests.

### NOTE

Changing the switch authentication policy mode does not affect online EX\_Ports, so it is acceptable to leave the default Passive policy configured on the FC router while the Active or On policy is required on the edge switch.

## Setting Up FC-FC Routing

To set up FC-FC routing, perform the following tasks in the order listed.

1. Verify that you have the proper setup for FC-FC routing. (See [Verifying the Setup for FC-FC Routing](#).)
2. Assign backbone fabric IDs. (See [Backbone Fabric IDs](#).)
3. Configure FCIP tunnels if you are connecting Fibre Channel SANs over IP-based networks.
4. Configure IFLs for edge and backbone fabric connection. (See [Inter-Fabric Link Configuration](#).)
5. Modify port cost for EX\_Ports, if you want to change from the default settings. (See [FC Router Port Cost Configuration](#).)
6. Enable shortest IFL mode if you want to choose a lowest cost IFL path in the backbone fabric. (See [Shortest IFL Cost Configuration](#).)
7. Configure trunking on EX\_Ports that are connected to the same edge fabric. (See [EX\\_Port Trunking Configuration](#).)
8. Configure LSAN zones to enable communication between devices in different fabrics. (See [LSAN Zone Configuration](#).)

## Verifying the Setup for FC-FC Routing

Before configuring a fabric to connect to another fabric, you must perform verification checks on the FC router.

1. Log in to the FC router as admin and enter the `version` command. Verify that Fabric OS v9.0.x is installed on the FC router, as shown in the following example.

```
switch:admin> version
Kernel: 4.1.35rt41
Fabric OS: v9.0.0
Made on: Tue Oct 17 13:21:34 2021
Flash: Fri Oct 17 21:37:48 2021
BootProm: 4.0.19
```

2. Enter the `slotShow -m` command to verify the slot, blade type, ID, model name, and status.

```
switch:admin> slotshow -m
```

| Slot | Blade Type | ID  | Model Name | Status  |
|------|------------|-----|------------|---------|
| 1    | CP BLADE   | 220 | CPX7       | ENABLED |
| 2    | CP BLADE   | 220 | CPX7       | ENABLED |
| 3    | SW BLADE   | 216 | FC64-48    | ENABLED |
| 4    | SW BLADE   | 216 | FC64-48    | ENABLED |
| 5    | SW BLADE   | 216 | FC64-48    | ENABLED |
| 6    | SW BLADE   | 216 | FC64-48    | ENABLED |
| 7    | CORE BLADE | 215 | CR64-8     | ENABLED |
| 8    | CORE BLADE | 215 | CR64-8     | ENABLED |
| 9    | SW BLADE   | 232 | FC64-64    | ENABLED |
| 10   | SW BLADE   | 232 | FC64-64    | ENABLED |
| 11   | SW BLADE   | 204 | FC32-64    | ENABLED |
| 12   | AP BLADE   | 186 | SX6        | ENABLED |

See [Blade Terminology and Capability](#) for a list of blades and their corresponding IDs.

Skip this step if you are using an FOS switch.

3. Verify the FCR is enabled on the chassis or a switch so that the EX\_Ports are set in a core/port blade port or switch port.

```
switch:admin> fosconfig --show
FC Routing service: enabled
Virtual Fabric: enabled
```

4. Enter the `license --show` command to verify that the Integrated Routing license is installed.

```
switch:admin> license --show
```

```
License Id : 10:00:c4:f5:7c:02:57:48
```

```
License 1 :
```

```

License serial number : FOS-86-0-03-11305239
License features : Fabric Vision and IO Insight
Generation date : 09/24/2021
```

```
License 2 :
```

```

License serial number : FOS-86-0-02-11305241
License features : Ports on Demand
License Capacity : 48
Generation date : 09/24/2021
```

```
License 3 :
```

```

License serial number : FOS-86-0-03-11305242
License features : Trunking
Generation date : 09/24/2021
```

```
License 4 :
```

```
License serial number : FOS-86-0-03-11305240
License features : Integrated Routing
Generation date : 09/24/2021
```

If you are connecting to a Fabric OS and the Integrated Routing license is not installed, you must install it, as described in the *Brocade Fabric OS Software Licensing User Guide*.

## Backbone Fabric IDs

If your configuration has only one backbone fabric, then you do not need to assign a backbone fabric ID because the backbone fabric ID in this situation defaults to a value of 128. The default backbone fabric ID is 1 if Virtual Fabrics is disabled.

All switches in a backbone fabric must have the same backbone fabric ID. You can configure the backbone fabric ID using the `fcrConfigure` command. The backbone fabric ID must be unique from the perspective of every attached edge fabric. Fabric ID changes made on a switch are not propagated to other switches in the backbone fabric. The backbone fabric administrator is responsible for making sure that all switches in the backbone have the same fabric ID. Because fabric IDs are used heavily by the routing protocol between the FC Routers, using the wrong fabric ID can affect both edge-to-edge and backbone-to-edge routing.

In addition to ensuring that the backbone fabric IDs are the same within the same backbone, you must make sure that when two different backbones are connected to the same edge fabric, the backbone fabric IDs are different, but the edge fabric ID should be the same. Configuration of two backbones with the same backbone fabric ID that are connected to the same edge is invalid. In this configuration, a RASLog message displays a warning about fabric ID overlap. When two backbone fabrics are *not* connected to the same edge, they can have the same backbone fabric ID.

### ATTENTION

In a multi-switch backbone fabric, modification of the FID within the backbone fabric will cause disruption to local traffic.

## Assigning Backbone Fabric IDs

1. Log in to the switch or backbone.
2. Enter the `switchDisable` command if EX\_Ports are online.
3. Enter the `fosConfig --disable fcr` command to disable the FC-FC routing service.  
The default state for the FC router is disabled.
4. Enter the `fcrConfigure --bbfid` command. At the prompt, enter the fabric ID, or press `Enter` to keep the current fabric ID, which is displayed in brackets.  
The `fcrConfigure --bbfid` command is not supported when VF mode is enabled.
5. Verify that the backbone fabric ID is different from that set for edge fabrics.  
Multiple FC routers attached to the same backbone fabric must have the same backbone fabric ID.
6. Enter the `fosConfig --enable fcr` command.
7. Enter the `switchEnable` command.

```
switch:admin> switchdisable
switch:admin> fosconfig --disable fcr
FC Router service is disabled
switch:admin> fcrconfigure --bbfid
FC Router parameter set. <cr> to skip a parameter
```

```

Please make sure new Backbone Fabric ID does not conflict with any configured EX-Port's Fabric ID
Backbone fabric ID: (1-128)[128]
switch:admin> fosconfig --enable fcr
FC Router service is enabled
switch:admin> switchenable

```

## Assigning Alias Names to Fabric IDs

You can assign an alias name to each fabric ID for easier configuration and manageability.

The alias name is a user-friendly name that you can use instead of a fabric ID (FID) in the `portcfgexport` command.

No relation exists between the FID alias name and the fabric name, which you configure with the `fabricname` command.

1. Log in to the switch or backbone.
2. Set the FID alias using the `fcrconfigure --add` command.

```
fcrconfigure --add -alias alias_name -fid fid
```

3. Verify the configured alias names using the `fcrconfigure --show -alias` command.

```
fcrconfigure --show -alias
```

The following example configures alias names for three fabrics, and then configures an EX\_Port using the FID alias name.

```

switch:admin> fcrconfigure --add -alias Green_fabric -fid 10
switch:admin> fcrconfigure --add -alias Blue_fabric -fid 12
switch:admin> fcrconfigure --add -alias Yellow_fabric -fid 15
switch:admin> fcrconfigure --show -alias
 FID Alias
=====
 10 Green_fabric
 12 Blue_fabric
 15 Yellow_fabric

switch:admin> portcfgexport 3 -a 1 -f Green_fabric; portenable 3
switch:admin> portcfgexport 3
Port 3 info
Admin: enabled
State: OK
Pid format: core(N)
Operate mode: Brocade Native
Edge Fabric ID: 10
Alias Name: Green_fabric
Front Domain ID: 5(output truncated)

```

## FCIP Tunnel Configuration

The optional Fibre Channel over IP (FCIP) Tunneling Service enables you to use tunnels to connect instances of Fibre Channel SANs over IP-based networks to transport all FC ISL and IFL traffic. If you are only using FC\_Ports, then there is no need to configure FCIP tunnels.

If using FCIP in your FC-FC routing configuration, you must first configure FCIP tunnels. When a tunnel is created, it defaults to a disabled state. Then configure the VE\_Port. After the appropriate ports are configured, enable the tunnel.

Refer to the *Brocade Fabric OS Extension User Guide* for instructions on how to configure FCIP tunnels.

## Inter-Fabric Link Configuration

Configuring an inter-fabric link (IFL) involves disabling ports and cabling them to other fabrics, configuring those ports for their intended uses, and then enabling the ports.

Before configuring an IFL, be aware that you cannot configure both IFLs (EX\_Ports) and ISLs (E\_Ports) from a backbone fabric to the same edge fabric.

### ATTENTION

To ensure that fabrics remain isolated, disable the port prior to inserting the cable. If you are configuring an EX\_Port, disable the port prior to making the connection.

## Configuring an IFL for Both Edge and Backbone Connections

1. On the FC router, disable the port that you are configuring as an EX\_Port (the one connected to the Fabric OS switch) by issuing the `portdisable` command.

```
switch:admin> portdisable 7/10
```

You can verify that the port has been disabled by issuing the `portshow` command for the port.

2. Configure each port that connects to an edge fabric as an EX\_Port using the `portcfgexport` command.

### NOTE

The `portcfgexport` (only on the FC ports on the FC router) command works only on ports that are capable of FC-FC routing.

The following example configures an EX\_Port and assigns a fabric ID of 30 to port 10.

```
switch:admin> portcfgexport 7/10 -a 1 -f 30
switch:admin> portcfgexport 7/10
Port 7/10 info
Admin: enabled
State: OK
Pid format: Not Applicable
Operate mode: Brocade Native
Edge Fabric ID: 30
Preferred Domain ID: 160
Front WWN: 50:06:06:9e:20:38:6e:1e
Fabric Parameters: Auto Negotiate
R_A_TOV: Not Applicable
E_D_TOV: Not Applicable
Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

This port can now connect to another switch.

The following example shows configuring a port range:

```
switch:admin> portcfgexport 1/1-2 -a 1
```

```
2014/09/15-07:09:12, [FCR-1071], 11763, SLOT 6 | FID 128, INFO, DCX, Port 1/1 is changed from non FCR port
to FCR port.
2014/09/15-07:09:13, [FCR-1071], 11764, SLOT 6 | FID 128, INFO, DCX, Port 1/2 is changed from non FCR port
to FCR port.
```

For related FC-FC routing commands, refer to `fcledgedshow`, `fcrxlateconfig`, `fcrcconfigure`, and `fcrcproxyconfig` in the *Brocade Fabric OS Command Reference Manual*.

An FC Router can interconnect multiple fabrics. EX\_Ports attached to more than one edge fabric must configure a different fabric ID for each edge fabric.

3. Configure FC router port cost if you want to change the default values. For information about using FC router port cost operations, refer to [FC Router Port Cost Configuration](#).
4. Set up the IFL for EX\_Port trunking. For information on trunking setup, see [Configuring Trunk Groups](#).
5. Physically attach the IFLs from the FC Router to the edge fabric.
6. Enter the `portenable` command to enable the ports that you disabled in Step 1.

```
switch:admin> portenable 7/10
```

7. Enter the `portcfgshow` command to view the ports that are persistently disabled.

```
switch:admin> portcfgshow 7/10
Area Number: 74
Speed Level: AUTO
Trunk Port OFF
Long Distance OFF
VC Link Init OFF
Locked L_Port OFF
Locked G_Port OFF
Disabled E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable OFF
NPIV capability ON
EX Port ON
Mirror Port ON
FC Fastwrite ON
```

8. After identifying the disabled ports, enter `portcfgpersistentenable portID` to enable the port, and then `portcfgshow portID` to verify that the port is enabled.

```
switch:admin> portcfgpersistentenable 7/10
switch:admin> portcfgshow 7/10
Area Number: 74
Speed Level: AUTO
Trunk Port OFF
Long Distance OFF
VC Link Init OFF
Locked L_Port OFF
Locked G_Port OFF
Disabled E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable OFF
NPIV capability ON
EX Port ON
```

```
Mirror Port ON
FC Fastwrite ON
```

## 9. Use the portcfgexport and portshow commands to verify that each port is configured correctly.

```
switch:admin> portcfgexport 7/10
Port 7/10 info
Admin: enabled
State: OK
Pid format: Not Applicable
Operate mode: Brocade Native
Edge Fabric ID: 30
Preferred Domain ID: 160
Front WWN: 50:06:06:9e:20:38:6e:1e
Fabric Parameters: Auto Negotiate
R_A_TOV: Not Applicable
E_D_TOV: Not Applicable
Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A

switch:admin> portshow 7/10

portName:
portHealth: OFFLINE
Authentication: None
EX_Port Mode: Enabled
Fabric ID: 30
Front Phantom: state = Not OK Pref Dom ID: 160
Fabric params: R_A_TOV: 0 E_D_TOV: 0 PID fmt: auto
Authentication Type: None
Hash Algorithm: N/A
DH Group: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
portDisableReason: None
portCFlags: 0x1
portFlags: 0x1 PRESENT U_PORT EX_PORT
portType: 10.0
portState: 2 Offline
portPhys: 2 No_Module
portScn: 0
port generation number: 0
portId: 014a00
portIfId: 4372080f
portWwn: 20:4a:00:60:69:e2:03:86
portWwn of device(s) connected:
Distance: normal
portSpeed: N4Gbps
LE domain: 0
FC Fastwrite: ON
Interrupts: 0 Link_failure: 0 Frjt : 0
```

```

Unknown: 0 Loss_of_sync: 0 Fbsy : 0
Lli: 0 Loss_of_sig: 2
Proc_rqrd: 0 Protocol_err: 0
Timed_out: 0 Invalid_word: 0
Rx_flushed: 0 Invalid_crc: 0
Tx_unavail: 0 Delim_err: 0
Free_buffer: 0 Address_err: 0
Overrun: 0 Lr_in: 0
Suspended: 0 Lr_out: 0
Parity_err: 0 Ols_in: 0
2_parity_err: 0 Ols_out: 0
CMI_bus_err: 0
Port part of other ADs: No

```

10. Enter the `switchshow` command to verify that the EX\_Port, edge fabric ID, and name of the edge fabric switch (containing the E\_Port or VE\_Port) are correct.
11. Enter the `fcrfabricshow` command to view any edge fabric switch names and ensure that links are working as expected.

#### NOTE

The `fcrfabricshow` command displays the static IPv6 addresses for each FC router and each edge fabric switch that is connected to the EX\_Ports.

```

switch:admin> fcrfabricshow
FCR WWN: 10:00:00:05:1e:13:59:00, Dom ID: 2, Info: 10.1.18.52 1080::8:800:200C:1234/64,"Spirit-2"
"fcr_630"
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)

7 10 10:00:00:05:1e:34:11:e5 10.1.18.33 "My630" 1080::8:8FF:FE0C:417A/64
4 116 10:00:00:05:1e:37:00:44 10.1.18.34 "My630"
FCR WWN: 10:00:00:05:1e:12:e0:00, Dom ID: 100, Info:10.1.18.50 1080::8:60F:FE0C:456A/64
"fcr_630"
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)

4 95 10:00:00:05:1e:37:00:45 10.1.18.31 "My630"
FCR WWN: 10:00:00:05:1e:12:e0:00, Dom ID: 100, Info: 10.1.18.50, "fcr_Brocade G630"
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)

4 95 10:00:00:05:1e:37:00:45 10.1.18.31 "Brocade G630"
5 95 10:00:00:05:1e:37:00:45 10.1.18.31 "Brocade G630"
6 95 10:00:00:05:1e:37:00:45 10.1.18.31 "Brocade G630"

```

12. Enter `iflshow` command on the edge fabric switch to display the FC router details and ensure the fabric is functioning correctly.

```

switch:admin> iflshow
 E-Port EX-Port FCR-WWN FCR-FID FCR-Name Speed BW

1 : 350 --> 12 10:00:08:00:88:04:93:94 39 fcr_sw 4G 8G TRUNK

```

## Configuring EX\_Ports on an ICL

The following restrictions apply when configuring EX\_Ports on an ICL:



- An EX\_Port on an ICL is supported in Brocade X7 and X6 Directors when all port blades in the chassis belong to one of the following blade types: FC64-48, FC32-X7-48, FC32-48, FC32-64, or SX6.

**NOTE**

Device discovery will not occur properly with ICL-EX\_Port-connected edge fabrics if the FC router has unsupported blades that are not mentioned in this restriction.

- EX\_Ports on ICLs are allowed only when the chassis is enabled with VF and FCR. Also, the EX\_Ports will be allowed only on the base switch of the chassis. You must ensure that the VF- and FCR-enabled chassis is installed with the FCR and ICL POD licenses.
- F\_Ports are not supported on the base switch.
- You must be in Brocade native mode.
- The ICLs must be in a symmetric topology. For more information on topologies, see [Supported Topologies for ICL Connections](#).
- A bladeless configuration (where the chassis has core blades but no application blades) is not supported.

To configure EX\_Ports on an ICL, perform the following steps:

1. Enter `configure` to configure an 8-bit zero-based Dynamic Area mode or a 10-bit Dynamic Area mode on the FC router to bring up ICL EX\_Ports.

**NOTE**

Gen 7 platforms support only the unified addressing mode (UAM), which simplifies the different addressing mode combinations and features into a single addressing mode. As a result of UAM in Gen 7 platforms, you will not see the configuration of area assignments in the `configure` command. For more information on UAM, see [Unified Addressing Mode](#).

```
switch:FIDxx:admin> switchdisable
switch:FIDxx:admin> configure

Configure...
Fabric parameters (yes, y, no, n):[no] y

Domain: (1..239) [1]
Enable a 256 Area Limit
(0 = No,
 1 = Zero Based Area Assignment,
 2 = Port Based Area Assignment): (0..2) [0]
R_A_TOV: (4000..120000) [10000]
```

The following options are available:

- **0:** 10-bit Dynamic Area Mode. ICL EX\_Ports, F\_Ports, and regular EX\_Ports have 8-bit areas.
- **1:** 8-bit Dynamic Area Mode. ICL\_EX ports have 8-bit areas.
- **2:** Not permitted if ICL ports are present in the logical switch.

2. On the FC router, enter `portdisable portrange` to disable all QSFP ports.

```
switch:FIDxx:admin> portdisable 7/0-3
```

You can verify that all ports have been disabled by issuing the `portshow` command for the ports.

3. Use the `portcfgexport` command to configure the EX\_Ports on the ICL. If you configure EX\_Port on one of the QSFP ports, the configuration is automatically propagated to the other three QSFP ports.

The following example configures EX\_Port on one of the QSFP ports.

```
switch:FIDxx:admin> portcfgexport 7/0-3 -a 1 -f 112
2020/11/26-09:21:54, [FCR-1071], 809, SLOT 1 | FID 50, INFO, Pluto, Port 7/0 is changed from non FCR port
to FCR port.
2020/11/26-09:21:54, [FCR-1071], 810, SLOT 1 | FID 50, INFO, Pluto, Port 7/1 is changed from non FCR port
to FCR port.
2020/11/26-09:21:55, [FCR-1071], 811, SLOT 1 | FID 50, INFO, Pluto, Port 7/2 is changed from non FCR port
to FCR port.
2020/11/26-09:21:55, [FCR-1071], 812, SLOT 1 | FID 50, INFO, Pluto, Port 7/3 is changed from non FCR port
to FCR port.
```

4. Enter `portenable portrange` to enable the QSFP ports that you disabled in Step 2.

```
switch:FIDxx:admin> portenable 7/0-3
```

5. Use the `portcfgexport` or `portshow` commands to verify that each port is configured correctly.

```
switch:FIDxx:admin> portcfgexport 7/0
Port 7/0 info
Admin: enabled
State: OK
Pid format: core(N)
Operate mode: Brocade Native
Edge Fabric ID: 112
Front Domain ID: 160
Front WWN: 50:02:7f:8f:24:d7:1e:70
Principal Switch: 11
Principal WWN: 10:00:00:27:f8:f3:42:02
Fabric Parameters: Auto Negotiate
R_A_TOV: 10000 (N)
E_D_TOV: 2000 (N)
Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Forward Error Correction: ON
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

6. Enter `fcrcdgeshow` to display the FIDs of all configured EX\_Ports.

```
switch:FIDxx:admin> fcrcdgeshow
```

| FID | EX-port | E-port | Neighbor Switch (PWWN, SWWN )                   | Flags |
|-----|---------|--------|-------------------------------------------------|-------|
| 112 | 7/0     | 264    | 2e:08:00:27:f8:f3:42:02 10:00:00:27:f8:f3:42:02 | FEC   |
| 112 | 7/1     | 265    | 2e:09:00:27:f8:f3:42:02 10:00:00:27:f8:f3:42:02 | FEC   |
| 112 | 7/2     | 266    | 2e:0a:00:27:f8:f3:42:02 10:00:00:27:f8:f3:42:02 | FEC   |
| 112 | 7/3     | 267    | 2e:0b:00:27:f8:f3:42:02 10:00:00:27:f8:f3:42:02 | FEC   |
| 112 | 8/0     | 280    | 2e:18:00:27:f8:f3:42:02 10:00:00:27:f8:f3:42:02 | FEC   |
| 112 | 8/1     | 281    | 2e:19:00:27:f8:f3:42:02 10:00:00:27:f8:f3:42:02 | FEC   |
| 112 | 8/2     | 282    | 2e:1a:00:27:f8:f3:42:02 10:00:00:27:f8:f3:42:02 | FEC   |
| 112 | 8/3     | 283    | 2e:1b:00:27:f8:f3:42:02 10:00:00:27:f8:f3:42:02 | FEC   |

7. Enter the `fcrliclpathbwmonitor --show` command to monitor and report path bandwidth imbalances for each edge fabric.

The following example enables the monitoring and reporting of bandwidth balances and imbalances, and it displays the ICL path bandwidth state for each fabric:

```
switch:FIDxx:admin> fcrliclpathbwmonitor --show
ICL Path Bandwidth state :Disabled

switch:FIDxx:admin> fcrliclpathbwmonitor --enable
ICL bandwidth balance functionality is enabled

switch:FIDxx:admin> fcrliclpathbwmonitor --show
ICL Path Bandwidth state :Enabled
```

| FABRIC | SLOT-7 BW | SLOT-8 BW | STATE    |
|--------|-----------|-----------|----------|
| 60     | 200       | 200       | BALANCED |
| 112    | 128       | 128       | BALANCED |

## FC Router Port Cost Configuration

FC routers optimize the usage of the router port links by directing traffic to the link with the smallest router port cost.

The FC router port cost is similar to the link cost setting available on E\_Ports, which allows you to customize traffic flow. The router module chooses the router port path based on the lowest cost for each FID connection. If multiple paths exist where one path costs less than the others, the lowest-cost path is used. If exchange-based routing has not been disabled and multiple paths exist with the same lowest cost, load sharing occurs over these paths.

Every IFL has a default cost. The default router port cost value is 1,000 for an EX\_Port IFL.

The FC router port cost settings are 0, 1000, or 10,000. If the cost is set to 0, the default cost will be used for that IFL. The FC router port cost is persistent and saved in the existing port configuration file, where it is also passed to other routers in the same backbone. Link costs from the front domain to the translate (xlate) domain remain at 10,000. You can use the `lsDbShow` command from the edge fabric to display these link costs.

The FC router port cost is set automatically. You can modify the cost for a port if you want to change the default value.

## Port Cost Considerations

The router port cost has the following considerations:

- Use proper SAN design guidelines to connect the IFLs to different port sets for effective router port cost use. For example, if both a low-speed IFL and a high-speed IFL are going to the same edge fabric, connect the lower router cost IFL to a separate port group (for example, ports 0 to 7) than the higher router cost IFL (for example, ports 8 to 15).

You can connect multiple EX\_Ports to the same edge fabric. The EX\_Ports can all be on the same FC router, or they can be on multiple routers. Multiple EX\_Ports create multiple paths for frame routing. Multiple paths can be used in two different, but compatible, ways:

- Failing over from one path to another.
- Using multiple paths in parallel to increase effective data transmission rates.

EX\_Ports, when connected, are assigned different router port costs and traffic will flow only through the EX\_Ports. Routing failover is automatic, but it can result in frames arriving out of order when frames take different routes. The FC router can force in-order delivery, although frame delivery is delayed immediately after the path failover.

Source EX\_Ports can balance loads across multiple destination EX\_Ports attached to the same edge fabric using exchange IDs from the routed frames as keys to distribute the traffic.

## Setting the Router Port Cost for an EX\_Port

The router port cost value for an EX\_Port is set automatically when the EX\_Port is created. You can modify the cost for that port to force a path to have a higher or lower cost.

You can configure the EX\_Port with values of either 1,000 or 10,000. For example, if you want to differentiate between two EX\_Port links with different speeds, you can assign 1,000 to one link and 10,000 to the other link.

1. Enable EX\_Port mode with the `portcfgexport` command.

```
switch:admin> portcfgexport 7/10 -a 1
```

2. Enter the `fcrrouterportcost` command to display the router port cost for each EX\_Port.

```
switch:admin> fcrrouterportcost
Port Cost

7/3 1000
7/4 1000
7/9 1000
7/10 1000
7/13 1000
10/0 1000
```

You can also use the `fcrrouteshow` command to display the router port cost.

To display the router port cost for a single EX\_Port, enter the `fcrrouterportcost` command with a port and slot number.

```
switch:admin> fcrrouterportcost 7/10
Port Cost

7/10 1000
```

3. Enter the appropriate form of the `fcrrouterportcost` command based on the task you want to perform:
  - To set the router port cost for a single EX\_Port, enter the command with a port and slot number and a specific cost:

```
switch:admin> fcrrouterportcost 7/10 10000
```

- To set the cost of the EX\_Port back to the default, enter a cost value of 0:

```
switch:admin> fcrrouterportcost 7/10 0
```

## Shortest IFL Cost Configuration

You can direct traffic flow to take the shortest path between host and target when multiple FC routers in the backbone are connected to an edge fabric. Shortest inter-fabric link mode is disabled by default. When you enable shortest IFL mode, an FC router can choose a lowest cost ISL path in the backbone fabric. This feature is useful when an FC router has multiple connections to the source edge fabric, and the backbone fabric has multiple FC routers connected through

FCIP links (VE\_Ports) and FC links (E\_Ports). The selection of a low cost path depends on individual ISL link cost settings in the backbone fabric. Traffic originating from a domain in an edge fabric can choose any equal cost path to reach the destination edge fabric. This traffic can be transmitted through high-cost paths, such as FCIP links, within the backbone fabric even though low-cost paths, such as FC links, are present.

When the shortest IFL mode is enabled, each FC router calculates the shortest path for each of its locally-connected source edge fabrics to a remote destination edge fabric that is connected to another FC router in the same backbone fabric. The FC router performs this calculation by identifying all paths in the backbone fabric that can connect the source edge fabric to the destination edge fabric. The cumulative ISL link cost for each path is then calculated:

- For any path for which the cumulative ISL link cost of the path is greater than or equal to 10,000, the FC router sets the link cost from the front domain to translate the domain as 10,001.
- For any path for which the cumulative ISL link cost of the path is less than 10,000, the link cost from front domain to translate domain will remain at 10,000, which is the shortest IFL path.

**NOTE**

The shortest IFL solution is applicable only when the edge fabric has multiple FC router connections and the backbone fabric has at least one available low-cost path. Enable the shortest IFL mode on all the FC Routers in the Backbone fabric for this feature to work correctly.


The following figure shows the shortest IFL solution. In order for Host 12 in Domain 1 in Edge fabric 1 to reach Target 2 in Domain 3 in Edge fabric 3, traffic can choose one of two paths in the backbone fabric:

- Path 1: Edge fabric 1-FCR1 - FCR2 -FCR3 -Edge fabric 3.
- Path 2: Edge fabric 1-FCR4 - FCR3 -Edge fabric 3.

The second path is identified as the low cost path by setting the ISL link cost to 5,000. The first path is considered the high-cost path where the cumulative ISL cost is 10,000.

**NOTE**

The link cost should be set in the direction of traffic flow from the source edge fabric to the destination edge fabric.

**Figure 90: Shortest IFL Solution**

## Configuring the Shortest IFL Cost

1. Enter the `fcFabricShow` command to view the FC routers on the backbone fabric.

```
switch:admin> fcFabricShow
FC Router WWN: 10:00:00:05:1e:58:bd:69, Dom ID: 10,
Info: 10.1.18.59, "DID_10"
EX_Port FID Neighbor Switch Info (enet IP, WWN, name)

 34 1 10.1.18.68 10:00:00:05:1e:61:28:22 "DID_4_1"
switch:admin>fcFabricShow
FC Router WWN: 10:00:00:05:1e:58:be:69, Dom ID: 20,
Info: 10.1.18.60, "DID_20"
EX_Port FID Neighbor Switch Info (enet IP, WWN, name)

 2 2 10.1.18.88 10:00:00:05:33:00:90:48 "DID_1_2"
switch:admin>fcFabricShow
```

```

FC Router WWN: 10:00:00:05:1e:58:be:68, Dom ID: 30,
Info: 10.1.18.61, "DID_30"
EX_Port FID Neighbor Switch Info (enet IP, WWN, name)

 31 3 10.1.18.153 10:00:00:05:33:7e:e8:7c "DID_3_3"
switch:admin>fcrfabricshow
FC Router WWN: 10:00:00:05:1e:58:be:67, Dom ID: 40,
Info: 10.1.18.62, "DID_40"
EX_Port FID Neighbor Switch Info (enet IP, WWN, name)

 34 1 10.1.18.68 10:00:00:05:1e:61:28:22 "DID_4_1"

```

2. Enter the `islshow` command to identify the connections between the FC routers and the destination edge fabric.

```

switch:admin> islshow
1: 10->1010:00:00:05:1e:58:be:69 20 DID_20 sp: 8.000G bw: 8.000G TRUNK QOS

```

3. Identify all paths in the backbone fabric through which traffic can flow from the source edge fabric to the destination edge fabric by examining the output generated in Step 1 and Step 2.

The following example step shows how to identify the number of paths from the source edge fabric 1 to the destination edge fabric 2 in the backbone fabric based on the configuration shown in Step 1 and Step 2:

- The path from FC router Domain ID 10 to FC router Domain ID 20.
- The path from FC router Domain 20: This FC router has no local connection to source edge fabric 1, so it cannot be considered for path identification.
- The path from FC router Domain 30: This FC router has no local connection to source edge fabric 1, so it cannot be considered for path identification.
- The path from FC router Domain ID 40 to FC router Domain ID 30 to FC router Domain ID 20.

Therefore, two available paths exist in the backbone fabric through which traffic can flow from the source edge fabric 1 to the destination edge fabric 2.

4. Enter the `linkcost` command to obtain the cumulative ISL cost for each path identified by adding all individual link costs, and to select one single path to be the only low cost path.

- The following example shows the cumulative ISL cost for the first path identified in Step 3: from FC router ID Domain 10 to FC router ID Domain 20.

```

switch:admin>linkcost 10
Interface10 (E_PORT) Cost 500

```

The cumulative link cost for this path is 500. This path is now known as path 1.

- The following example shows the cumulative ISL cost for the second path identified in Step 3: first from FC router ID Domain 40 to FC router ID Domain 30, and then from FC router Domain ID 30 to FC router ID Domain 20.

```

switch:admin>linkcost 10
Interface10 (E_PORT) Cost 500
switch:admin>linkcost 10
Interface10 (E_PORT) Cost 500

```

The cumulative link cost for this path is 1000. This path is now known as path 2.

Path 1 is selected as the low-cost path.

5. Enter the `linkcost` command to set low-cost values, ensuring that the cumulative ISL cost for the selected path is lower than that of all other paths. A low-cost path should have a cumulative ISL cost of less than 10,000.

- In the following example, the ISL link cost of path 2 from FC router ID Domain 40 to FC router Domain ID 30 is modified.

```
switch:admin>linkcost 10 5000
Interface10 (E_PORT) Cost 5000
```

- In the following example, the ISL link cost of path 2 from FC router Domain 30 to FC router Domain 20 is modified.

```
switch:admin>linkcost 10 5000
Interface10 (E_PORT) Cost 5000
```

The modified cumulative link cost for path 2 is 10,000.

6. Enter the `fcrConfigure --enable -shortestifl` command to enable the shortest IFL mode in all the FC routers in the backbone fabric.

```
switch:admin> fcrconfigure --enable -shortestifl
Shortest IFL path is enabled
```

7. Enter the `fcrConfigure -show` command to verify that the shortest IFL mode is enabled in all the FC routers in the Backbone fabric.

```
switch:admin> fcrconfigure --show
Backbone fabric ID: 128
Shortest IFL feature is enabled
```

8. Verify that the link cost of both the front domain and translate domain in the source edge fabric have been modified using the `lsDbShow` command.

```
linkCnt = 2, flags = 0x0
LinkId = 160, out port = 160 rem port = 241, cost = 10001, costCnt = 0, type = 1
LinkId = 161, out port = 161, rem port = 242, cost = 10000, costCnt = 0, type = 1
```

## EX\_Port Trunking Configuration

EX\_Port trunking provides the best utilization and balance of frames transmitted on each link between the FC router and the edge fabric. You should trunk all ports connected to the same edge fabrics.

You configure EX\_Ports to use frame-based trunking just as you do regular for E\_Ports. The restrictions for EX\_Port frame trunking are the same as those for E\_Ports. All the ports must be adjacent to each other using the clearly marked groups on the front of the product.

The FC router front domain has a higher node WWN, derived from the FC router, than that of the edge fabric. Therefore, the FC router front domain initiates the trunking protocol on the EX\_Port.

After initiation, the first port from the trunk group that comes online is designated as the controller port. The other ports that come online on the trunk group are considered the agent ports. Adding or removing a agent port does not cause frame drop. However, removing a agent port causes the loss of frames in transit.

If router port cost is used with EX\_Port trunking, the controller port and agent ports share the router port cost of the controller port.

For information about setting up E\_Port trunking on an edge fabric, see [Managing Trunking Connections](#).



## LSAN Zone Configuration

An LSAN consists of zones in two or more edge or backbone fabrics that contain the same devices. LSANs provide selective device connectivity between fabrics without forcing you to merge those fabrics. FC routers provide multiple mechanisms to manage inter-fabric device connectivity through extensions to existing switch management interfaces. You can define and manage LSANs using Brocade Advanced Zoning.

### NOTE

For performance reasons, Brocade recommends that you do not configure LSANs for device sharing between Fabric OS fabrics until after you activate the Integrated Routing license.

### NOTE

When an FCR front domain is joining a fabric with the Fabric OS version 9.0.0, the list of all LSAN zones in the current config is sent to the FCR switch. The total size of all enabled LSAN zones must not exceed 2 MB. If the total config size for all enabled LSAN zones is more than 2MB, the ZONE-1082 RASLOG is generated whether an FCR domain is present in the fabric or not.

## Zone Definition and Naming

Zones are defined locally on a switch or backbone. Names and memberships, with the exception of hosts and targets exported from one fabric to another, do not need to be coordinated with other fabrics. For example, refer to the [Figure A](#) metaSAN with Edge-to-Edge and Backbone Fabrics and LSAN Zones, when the zones for Edge SAN 1 are defined, you do not need to consider the zones in Edge SAN 2 and vice versa.

Zones that contain hosts and targets that are shared between the two fabrics must be explicitly coordinated. To share devices between any two fabrics, you must create an LSAN zone in both fabrics containing the port WWNs of the devices to be shared. Although an LSAN zone is managed using the same tools as any other zone on the edge fabric, two behaviors distinguish an LSAN from a conventional zone:

- A required naming convention. The name of an LSAN begins with the prefix "LSAN\_". The LSAN name is not case-sensitive; for example, "lsan\_" is equivalent to "LSAN\_", "Lsan\_", and so on.
- Members must be identified by their port WWN because port IDs are not necessarily unique across fabrics. The names of the zones need not be explicitly the same, and membership lists of the zones need not be in the same order.

### NOTE

The "LSAN\_" prefix must appear at the beginning of the zone name. LSAN zones may not be combined with QoS zones. Refer to QoS zones for more information about the naming convention for QoS zones.

To enable device sharing across multiple fabrics, you must create LSAN zones on the edge fabrics (optionally on the backbone fabric if the device sharing occurs between a backbone fabric to an edge fabric), using normal zoning operations to create zones with names that begin with the special prefix "LSAN\_", and adding host and target port WWNs from both local and remote fabrics to each local zone as desired. Zones on the backbone and on multiple edge fabrics that share a common set of devices will be recognized as constituting a single multi-fabric LSAN zone, and the devices that they have in common will be able to communicate with each other across fabric boundaries.

## LSAN Zones and Fabric-to-Fabric Communications

Zoning is enforced by all involved fabrics and any communication from one fabric to another must be allowed by the zoning setup on both fabrics. If the LSANs are under separate administrative control, separate administrators maintain access control.

## Controlling Device Communication with the LSAN

The following procedure illustrates how LSANs control which devices can communicate with each other. The procedure shows the creation of two LSANs (called "lsan\_zone\_fabric75" and "lsan\_zone\_fabric2"), which involve the following devices and connections:

- Switch1 and the host in fabric75
- Switch2, Target A, and Target B in fabric2
- Switch1 is connected to the FC router using an EX\_Port
- Switch2 is connected to the FC router using another EX\_Port
- Host has WWN 10:00:00:00:c9:2b:c9:0c (connected to switch1)
- Target A has WWN 50:05:07:61:00:5b:62:ed (connected to switch2)
- Target B has WWN 50:05:07:61:00:49:20:b4 (connected to switch2)

1. Log in as admin to fabric2.
2. Enter `nsshow` to identify Target A (50:05:07:61:00:5b:62:ed) and Target B (50:05:07:61:00:49:20:b4).

```
switch:admin> nsshow
{
 Type Pid COS PortName NodeName SCR
 N 010200; 3;50:05:07:61:00:5b:62:ed;20:05:07:61:00:5b:62:ed; 0x00000000
 SCR: None
 FC4s: FCP
 PortSymb: [34] "SB V7.4 port10 10.1.18.145 Targ0"
 Fabric Port Name: 20:1b:00:27:f8:f1:4b:41
 Permanent Port Name: 50:05:07:61:00:5b:62:ed
 Device type: Physical Target
 Port Index: 27
 Share Area: No
 Redirect: No
 Partial: No
 LSAN: Yes
 Slow Drain Device: No
 Device link speed: 32G
 Connected through AG: No
 Real device behind AG: No
 FCoE: No
 FC4 Features [FCP]: Target
 N 010200; 3;50:05:07:61:00:49:20:b4;20:05:07:61:00:49:20:b4; 0x00000000
 SCR: None
 FC4s: FCP
 PortSymb: [34] "SB V7.4 port11 10.1.18.145 Targ1"
 Fabric Port Name: 20:1b:00:27:f8:f1:4b:41
 Permanent Port Name: 50:05:07:61:00:49:20:b4
 Device type: Physical Target
 Port Index: 28
 Share Area: No
 Redirect: No
 Partial: No
 LSAN: Yes
 Slow Drain Device: No
 Device link speed: 32G
 Connected through AG: No
 Real device behind AG: No
 FCoE: No
 FC4 Features [FCP]: Target
 The Local Name Server has 2 entries }
```

3. Enter `zonecreate "lsan_zone_fabric2", "10:00:00:00:c9:2b:c9:0c;50:05:07:61:00:5b:62:ed;50:05:07:61:00:49:20:b4"` to create the LSAN "lsan\_zone\_fabric2", which includes the host (10:00:00:00:c9:2b:6a:2c), Target A, and Target B.

```
switch:admin> zonecreate "lsan_zone_fabric2",
"10:00:00:00:c9:2b:c9:0c;50:05:07:61:00:5b:62:ed;50:05:07:61:00:49:20:b4"
```

4. Enter `cfgshow`. This command allows you to verify that the zones are correct.

```
switch:admin> cfgshow
Defined configuration:
zone: lsan_zone_fabric2
 10:00:00:00:c9:2b:c9:0c; 50:05:07:61:00:5b:62:ed;
 50:05:07:61:00:49:20:b4
Effective configuration:
no configuration in effect
```

5. Enter `cfgadd "zone_cfg", "lsan_zone_fabric2"` to create the LSAN configuration.

```
switch:admin> cfgadd "zone_cfg", "lsan_zone_fabric2"
switch:admin> cfgenable "zone_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'zone_cfg' configuration (yes, y, no, n): [no] y
zone config "zone_cfg" is in effect
Updating flash ...
```

6. Enter `cfgenable` to enable the LSAN configuration.

```
switch:admin> cfgenable "zone_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'zone_cfg' configuration (yes, y, no, n): [no] y
zone config "zone_cfg" is in effect
Updating flash ...
```

7. Log in as admin and connect to switch1.

8. Enter `nsshow` to identify the WWN of the host. (Here it is 10:00:00:00:c9:2b:c9:0c.)

The `nsshow` output displays the LSAN zone status of a device, the port WWN, and the node WWN; the port WWN must be used for LSANs.

```
switch:admin> nsshow
{
 Type Pid COS PortName NodeName SCR
 N e50100; 3;10:00:00:00:c9:2b:c9:0c;20:00:00:00:c9:2b:c9:0c; 0x01000003
 SCR: Fabric-Detected Nx-Port-Detected Brocade-Principal-Switch
 FC4s: FCP
 PortSymb: [78] "QLogic-1860-2p | 3.2.6.0 | SQA-141-079 | Windows Server 2012 R2 Standard | N/A"
 NodeSymb: [41] "QLogic-1860-2p | 3.2.6.0 | SQA-141-079 | "
 Fabric Port Name: 20:01:00:05:33:e5:54:ab
 Permanent Port Name: 10:00:00:00:c9:2b:c9:0c
 Device type: Physical Initiator
 Port Index: 1
 Share Area: No
 Redirect: No
 Partial: No
```

```

LSAN: Yes
Slow Drain Device: No
Device link speed: 32G
Connected through AG: No
Real device behind AG: No
FCoE: No
FC4 Features [FCP]: Initiator
The Local Name Server has 1 entry }

```

9. Enter `zonecreate "lsan_zone_fabric75", "10:00:00:00:c9:2b:c9:0c"`. This creates the LSAN "lsan\_zone\_fabric75", which includes the host.

```
switch:admin> zonecreate "lsan_zone_fabric75", "10:00:00:00:c9:2b:c9:0c"
```

10. Enter `zoneadd`. This adds Target A to the LSAN.

```
FID75Domain5:admin> zoneadd "lsan_zone_fabric75", "50:05:07:61:00:5b:62:ed"
```

11. Enter `cfgadd "zone_cfg", "lsan_zone_fabric75"` to add the LSAN configuration. (Alternatively, you can use the `cfgcreate` command.)

```
switch:admin> cfgadd "zone_cfg", "lsan_zone_fabric75"
```

12. Enter `cfgenable "zone_cfg"` to enable the LSAN configuration. You will be asked to confirm this action.

```

switch:admin> cfgenable "zone_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'zone_cfg' configuration (yes, y, no, n): [no] y
zone config "zone_cfg" is in effect
Updating flash ...

```

13. Log in as admin and connect to the FC router.

14. Enter the following commands to display information about the LSANs:

- `lsanzoneshow -d` displays the fabric ID of the LSAN devices and whether the device physically exists in the fabric or is imported from another fabric.

```

switch:admin> lsanzoneshow -d
Fabric ID: 2 Zone Name: lsan_zone_fabric2
 10:00:00:00:c9:2b:c9:0c Imported from FID 75
 50:05:07:61:00:5b:62:ed EXIST in FID 2
 50:05:07:61:00:49:20:b4 EXIST in FID 2
Fabric ID: 75 Zone Name: lsan_zone_fabric75
 10:00:00:00:c9:2b:c9:0c EXIST in FID 75
 50:05:07:61:00:5b:62:ed Imported from FID 2

```

- `lsanzoneshow -o -sort` displays the WWNs in sorted order. The `lsanzoneshow -o` command applies to peer zone rules and displays the WWNs in sorted order if LSAN zones are peer zones.
- `fcrphydevshow` shows the physical devices in the LSAN.

```

switch:admin> fcrphydevshow
 Device WWN Physical
 Exists
 in Fabric

 75 10:00:00:00:c9:2b:c9:0c c70000
 2 50:05:07:61:00:49:20:b4 0100ef
 2 50:05:07:61:00:5b:62:ed 0100e8

```

Total devices displayed: 3

- `fcrproxydevshow` shows the proxy devices in the LSAN.

```
switch:admin> fcrproxydevshow
```

| Proxy<br>Created<br>in Fabric | WWN                     | Proxy<br>PID | Device<br>Exists<br>in Fabric | Physical<br>PID | State    |
|-------------------------------|-------------------------|--------------|-------------------------------|-----------------|----------|
| 75                            | 50:05:07:61:00:5b:62:ed | 01f001       | 2                             | 0100e8          | Imported |
| 2                             | 10:00:00:00:c9:2b:c9:0c | 02f000       | 75                            | c70000          | Imported |

Total devices displayed: 2

On the FC router, the host and Target A are imported, because both are defined by "lsan\_zone\_fabric2" and "lsan\_zone\_fabric75". However, Target B is defined by "lsan\_zone\_fabric2" and is not imported because "lsan\_zone\_fabric75" does not allow it.

When a PLOGI, PDISC, or ADISC arrives at the FC router, the SID and DID of the frame are checked. If they are LSAN-zoned at both SID and DID edge fabrics, the frame is forwarded to the DID. If they are not zoned, only the PLOGI is dropped and the remaining frames zoning enforcement takes place in the edge fabrics.

#### NOTE

If the host-side LSAN zone configuration is enabled prior to the target-side LSAN zone configuration, the host cannot see the target issue and the RAS Log message FCR-1103 reported.

## Configuring Interconnectivity between Backbone and Edge Fabrics

If you want devices in backbone fabrics to communicate with devices in edge fabrics, set up the LSANs as described [Controlling Device Communication with the LSAN](#). However, instead of configuring the LSAN in the second edge fabric, configure the LSAN in the backbone fabric.

## Setting the Maximum LSAN Count

Fabric OS allows you to set the maximum number of LSAN zones, also called the LSAN count, that can be configured on the backbone fabric.

By default, the maximum LSAN count is set to 3,000. You can increase the maximum LSAN count to 7,500 in X6 or X7 director without disabling the switch. The maximum number of LSAN devices supported is 15,000 in X6 or X7 (this includes both physical and proxy devices). If you have 3,000 LSAN zones but have not exceeded the 15,000 device limit, you can increase the LSAN count to 5,000 or 7,500. All FC routers in the same backbone fabric should have the same maximum LSAN count defined, to prevent the FC routers from running into an indefinite state. Asymmetric LSAN configurations due to different maximum LSAN counts may lead to different devices being imported on different FC routers.

With Fabric OS 9.0.x running on a fixed-port switch or both CPs of a chassis system, the scalability limits are increased. To achieve these new FC router scalability numbers, only Brocade Gen 7 and Gen 6 platforms are used as FC router domain in the metaSAN. Downgrading to an older Fabric OS version is blocked when the configuration exceeds the scalability numbers for the earlier version.

- The maximum number of FC routers in the backbone fabric is increased from 12 to 16.
- The maximum number of LSAN zones in a metaSAN is increased from 5000 to 7500. This can be scaled beyond 7500 by using the location-embedded LSAN zones.
- The maximum number of LSAN devices in a metaSAN is increased from 10000 to 15000.

**NOTE**

For Brocade 7810 Extension Switch platform, the maximum LSAN zone count is 3,000.

1. Enter `fcrlsancount` with no parameters to display the current LSAN limit.

```
switch:admin> fcrlsancount
LSAN Zone Limit 3000
```

2. Enter `fcrlsancount` and specify the new LSAN zone limit.

```
switch:admin> fcrlsancount 5000
LSAN Zone Limit 5000
```

For information on how to display the maximum allowed and currently used LSAN zones and devices, see [Resource Monitoring](#).

**NOTE**

Because the maximum number of LSANs is configured for each switch, if there is a different maximum LSAN count on the switches throughout the metaSAN, then the device import and export is not identical on the FC routers. You should enter the same maximum LSAN count for all the FC routers in the same backbone that support this feature. Verify the configured maximum limit against the LSANs configured using the `fcresourceshow` command.

The `lsanzoneshow` running into command is enhanced with following two new options to display which LSAN Zones or Devices need to be removed to downgrade successfully.

```
switch:admin> lsanzoneshow -m
List of LSAN Zones need to be removed before downgrade to pre-FOS v8.1.0:
Fabric IDZone Name

30lsan_5
Total unsupported LSAN Zones: 1
List of LSAN Devices need to be removed before downgrade to pre-FOS v8.1.0:

Fabric IDPortStateImportedZone Name
WWNFabric ID

3030:08:03:05:1e:61:28:22EXIST-lsan_4
3030:0c:03:05:1e:61:28:22Imported60lsan_4
Total unsupported LSAN Devices: 2
switch:admin> lsanzoneshow-r
List of LSAN Zones need to be removed before downgrade to pre-FOS v8.1.0:

----LSAN Zones need to be removed from fabric: 30 ----
cfgremove"<active_cfg_name>","lsan_5"
Total unsupported LSAN Zones: 1
List of LSAN Devices need to be removed before downgrade to pre-FOS v8.1.0:

----LSAN Devices need to be removed from fabric 30 ----
zoneremove"lsan_4","30:08:03:05:1e:61:28:22"
zoneremove"lsan_4","30:0c:03:05:1e:61:28:22"
Total unsupported LSAN Devices: 2
Note: User has to enable the zone configuration in respective edge fabric once devices/zones are removed.
```

## HA and Downgrade Considerations for LSAN Zones

Consider the following points on how LSAN zones impact high availability and firmware downgrades:

- The LSAN zone matrix is synchronized to the standby CP.
- If the feature is enabled on the active CP, introducing a CP with an earlier version of Fabric OS as a standby will cause HA synchronization to fail.
- If the feature is enabled, before downgrading to an earlier Fabric OS version, you will be asked to go back to the default mode.
- This feature does not have any impact on current HA functionality. LSANs will be synchronized as usual after the limit is increased and new LSANs are created.
- In a Backbone-to-edge routing topology, if the "FCR state" in the switch is disabled and the switch is downgraded, and then subsequently the FCR state is enabled, you must re-enable the zone configuration in the backbone fabric if `lsanzoneshow -s` for the backbone fabric ID is not displayed.

## LSAN Zone Policies Using LSAN Tagging

You can create tags for LSAN zones to give them a special meaning.

LSAN zones are zones with names that start with the "lsan\_" prefix. You can specify a tag to append to this prefix that causes the LSAN zone to be treated differently.

You can specify two types of tags:

- Enforce tag - Specifies which LSANs are to be enforced in an FC router.
- Speed tag - Specifies which LSANs are to be imported or exported faster than other LSANs.

The LSAN tags are persistently saved and support `configupload` and `configdownload`.

### Enforce Tag

Use the Enforce tag to achieve better scalability in the FC router by limiting the number of LSAN zones that are enforced in that FC router.


The Enforce tag reduces the resources used in an FC router by limiting the number of LSAN zones that will be enforced in that FC router.

The Enforce tag is useful in the following scenarios:

- Multiple FC routers are connected to the same edge fabric.
- The backbone fabric contains multiple FC routers with multiple edge fabrics.

In these scenarios, often an FC router does not need to know about all edge fabrics and their respective LSAN zones. You can use the Enforce tag to indicate which LSAN zones are imported by which FC routers. Without the Enforce tag, all FC routers import all LSAN zones, even those that are not needed.

For example, the following figure shows a metaSAN with three FC routers in the backbone, and each FC router has one edge fabric connected to it. In this scenario, LSAN zones communicate between edge fabric 1 and edge fabric 2, and between edge fabric 2 and edge fabric 3. There is no communication between edge fabric 1 and edge fabric 3.

**Figure 91: Example of Setting up Enforce LSAN Tag**

FC router 1 does not need to know about the LSAN between edge fabrics 2 and 3. Likewise, FC router 3 does not need to know about the LSAN between edge fabrics 1 and 2.

In this scenario, you could set up two Enforce tags, one for each LSAN. On FC router 2, both Enforce tags would be needed, since FC router 2 uses both LSANs. FC router 1 and FC router 3 each need only one tag, for their respective LSANs.

Setting up Enforce tags in this way keeps resources on FC router 1 and FC router 3 at a minimum. This is a simple example, but it can become more complex as more FIDs are added to each FC Router or as additional FC routers are added to the Backbone.

### **How the Enforce Tag Works**

The FC router automatically accepts all zones with names that start with "lsan\_". You can specify an Enforce tag to indicate that a particular FC router should accept only zones that start with the prefix "lsan\_tag". For example, if you specify an Enforce tag of "abc", the FC router accepts only those LSAN zones that start with "lsan\_abc" and does not import or export any other LSAN zones.

The Enforce tag can be up to eight characters long and can contain only letters and numbers.

The Enforce tag is not case-sensitive; for example, the tag "abc" is equivalent to "ABC" and "Abc".

If you specify "abc", "xyz", and "fab1" as Enforce tags, then the FC router accepts only those LSAN zones with names that start with any of the following:

- lsan\_abc
- lsan\_xyz
- lsan\_fab1

In this example, the following LSAN zones would all be accepted:

- lsan\_abc
- Lsan\_xyz123456
- LSAN\_FAB1\_abc



You can specify up to eight Enforce tags on an FC router.

For example, in the previous figure, you could configure the following Enforce tags on the FC routers:

- For FC router 1, configure one Enforce tag, "21".  
FC router 1 would accept all LSAN zones starting with "LSAN\_21", and so would accept LSAN\_21\_fab, but not LSAN\_23fabrics.
- For FC router 2, configure two Enforce tags, "21" and "23".  
FC router 2 would accept all LSAN zones starting with "LSAN\_21" and "LSAN\_23", and so would accept both "LSAN\_21\_fab" and "LSAN\_23fabrics".  
Alternatively, you could configure only one Enforce tag, "2", which would handle both of the LSANs.
- For FC router 3, configure one Enforce tag, "23".  
FC router 3 would accept all LSAN zones starting with "LSAN\_23", and so would accept LSAN\_23fabrics, but not LSAN\_21\_fab.

## Speed Tag

The Speed tag allows you to speed up the discovery process by importing devices into the remote edge fabrics when the devices come online, regardless of the state of the host.

During target discovery, the FC router process of presenting proxy devices and setting up paths to the proxy devices may cause some sensitive hosts to time out or fail. Using the Speed tag helps sensitive hosts to quickly discover the devices without timing out.

The Speed tag is used primarily when there are boot over SAN devices across an LSAN.

You set the Speed tag on the FC router, and then configure the LSANs in the target edge fabrics with the tag.

The FC router automatically accepts all zones with names that start with "lsan\_". You can specify a Speed tag to indicate that devices in LSAN zones that start with the prefix "lsan\_tag", should be imported into the remote edge fabrics when the devices come online.

Before enabling LSAN zone with the Speed tag, the host device in the edge fabric must be online.

For example, in the following figure, assume that the host, H1, needs fast access to target devices D1 and D2. You could set up the Speed tag as follows:


1. In FC router 1 and FC router 2, configure the Speed tag as "super".
2. In Edge fabric 2, configure two LSANs:
  - lsan\_f2\_f1 (H1, D1)
  - lsan\_f2\_f3 (H1, D2)

The LSAN in the host fabric does not need the tag.

3. In Edge fabric 1, configure the following LSAN:  
lsan\_super\_f1\_f2 (H1, D1)
4. In Edge fabric 3, configure the following LSAN:  
lsan\_super\_f3\_f2 (H1, D2)
5. Choose either the host or target to trigger the fast import process.

The "super" tag is needed only in the LSANs of the target fabrics.

The target proxies D1 and D2 are always present in the host fabric (Edge fabric 2), even if the host is brought down. A target proxy is removed from the host fabric when the target device is offline.

**Figure 92: Example of Setting Up Speed LSAN Tag**

## Rules for LSAN Tagging

Note the following rules for configuring LSAN tags:

- You configure the tags on the FC router, and not on the edge switches. If Virtual Fabrics is enabled, you configure the tags on the base switch on which the EX\_Ports are located. You then must ensure that the LSAN zones in the edge fabrics incorporate the tags correctly.
- If speed tags are set in the LSAN zones in the target edge fabric and in the FC router switches for a Boot-Over-SAN target and an IFL disruption is performed, the fast-import process for the Boot-Over-SAN target is lost. In such cases, you must initiate the discovery of the target by bringing the initiators online.
- The LSAN tags are configured per FC router, not per fabric. If the backbone fabric has multiple FC routers, you should configure the LSAN tags on all of the FC routers.
- The FC router must be disabled before you configure the Enforce tag. Configuring the Speed tag does not require that the FC router be disabled. However, after configuring the Speed tag, you must select the host or target port to trigger the fast-import process.
- The tag is from 1 through 8 alphanumeric characters.
- You can configure only one Speed tag on an FC router, and up to eight Enforce tags on an FC router. The maximum number of tags (Enforce and Speed) on an FC router is eight.
- Up to 500 speed, LSAN tags are supported.

## Configuring an Enforce LSAN Tag

1. Log in to the FC router as admin.
2. Enter the following command to disable the FC router:

```
sw0:admin> switchdisable
```

3. Enter the following command to create an Enforce LSAN tag:

```
sw0:admin> fcrlsan --add -enforce tagname
```

The *tagname* variable is the name of the LSAN tag that you want to create.

4. Enter the following command to enable the FC router:

```
sw0:admin> switchenable
```

5. Change the names of the LSAN zones in the edge fabrics to incorporate the tag in the names.

```
sw0:admin> switchdisable
sw0:admin> fcrlsan --add -enforce enftag1
LSAN tag set successfully
sw0:admin> switchenable
```

## Configuring a Speed LSAN Tag

1. Log in to the FC router as admin.
2. Enter the `fcrlsan --add -speed tagname` command to create a Speed LSAN tag:  
The *tagname* variable is the name of the LSAN tag you want to create.
3. Change the names of the LSAN zones in the edge fabrics to incorporate the tag in the names.
4. Choose the host or target port to trigger the fast-import process.

The following example creates a Speed LSAN tag named `fasttag2`.

```
switch:admin> fcrlsan --add -speed fasttag2
LSAN tag set successfully
```

## Removing an LSAN Tag

When you remove an LSAN tag, the tag is deactivated so that LSAN zones with this tag in the name now behave as regular LSAN zones. Removing an LSAN tag does not remove the LSAN zone.

You must disable the switch before removing an Enforce LSAN tag. You do not need to disable the switch to remove a Speed LSAN tag.

1. Log in to the FC router as admin.
2. Enter the `fcrlsan --remove` command to remove an existing LSAN tag.

If you remove an Enforce LSAN tag, you must disable the switch first.

### Example of Removing an Enforce LSAN Tag

```
sw0:admin> switchdisable
sw0:admin> fcrlsan --remove -enforce enftag1
```

```
LSAN tag removed successfully
sw0:admin> switchenable
```

### Example of Removing a Speed LSAN Tag

```
sw0:admin> fcrlsan --remove -speed fasttag2
LSAN tag removed successfully
```

## Displaying the LSAN Tag Configuration

1. Log in to the FC router as admin.
2. Enter the `fcrlsan --show` command.

```
sw0:admin> fcrlsan --show -enforce
Total LSAN tags : 1
 ENFORCE : enftag1
sw0:admin> fcrlsan --show -speed
Total SPEED tags : 1
 SPEED : fasttag2
sw0:admin> fcrlsan --show -all
Total LSAN tags : 2
 ENFORCE : enftag1
 SPEED : fasttag2
```

## LSAN Zone Binding


LSAN zone binding is an optional, advanced feature that increases the scalability envelope for very large metaSANs.

Without LSAN zone binding, every FC router in the backbone fabric maintains the entire LSAN zone and device state database. The size of this database limits the number of FC routers and devices you can have.

With LSAN zone binding, each FC router in the backbone fabric stores only the LSAN zone entries of the remote edge fabrics that can access its local edge fabrics. The LSAN zone limit supported in the backbone fabric is not limited by the capability of one FC router. In addition, due to the lower LSAN count, the CPU consumption by the FC router is lower. If you configure the metaSAN such that the backbone fabric has two groups of FC routers and there is no LSAN zone sharing and device access between the two groups, the number of FC routers and devices supported in the backbone fabric can be higher.

The following figure shows a sample metaSAN with four FC routers in the backbone fabric. Without LSAN zone binding, each FC router in the backbone fabric would store information about LSAN zones 1, 2, 3, and 4.

Figure 93: LSAN Zone Binding



After you set up LSAN zone binding, each FC router stores information about only those LSAN zones that access its local edge fabrics. The following table shows what LSAN information is stored in each FC router before and after LSAN zone binding is in effect.

Table 144: LSAN Information Stored in FC Routers, with and without LSAN Zone Binding

| Without LSAN Zone Binding |             |             |             | With LSAN Zone Binding |             |             |             |
|---------------------------|-------------|-------------|-------------|------------------------|-------------|-------------|-------------|
| FC router 1               | FC router 2 | FC router 3 | FC router 4 | FC router 1            | FC router 2 | FC router 3 | FC router 4 |
| LSAN1                     | LSAN1       | LSAN1       | LSAN1       | LSAN1                  | LSAN2       | LSAN3       | LSAN4       |
| LSAN2                     | LSAN2       | LSAN2       | LSAN2       | LSAN2                  | —           | LSAN4       | —           |
| LSAN3                     | LSAN3       | LSAN3       | LSAN3       | —                      | —           | —           | —           |

| Without LSAN Zone Binding |       |       |       | With LSAN Zone Binding |   |   |   |
|---------------------------|-------|-------|-------|------------------------|---|---|---|
| LSAN4                     | LSAN4 | LSAN4 | LSAN4 | —                      | — | — | — |

## LSAN Zone Binding Considerations

- Without LSAN zone binding, the maximum number of LSAN devices is 15,000.
- With LSAN zone binding, the metaSAN can import more than 15,000 devices and the backbone fabric can support more FC routers.
- With LSAN zone binding, CPU consumption by an FC router is lower.

## How LSAN Zone Binding Works

LSAN zone binding uses an *FC router matrix*, which specifies pairs of FC routers in the backbone fabric that can access each other, and an *LSAN fabric matrix*, which specifies pairs of edge fabrics that can access each other.

You set up LSAN zone binding using the `fcrLsanMatrix` command. This command has two options: `-fcr` and `-lsan`. The `-fcr` option is for creating and updating the FC router matrix, and the `-lsan` option is for creating and updating the LSAN fabric matrix.

When you set up LSAN zone binding on the local FC router, the resultant matrix database is automatically distributed to all of the FC routers in the backbone fabric.

If a new FC router joins the backbone fabric, the matrix database is automatically distributed to that FC router unless it has a different LSAN fabric matrix or FC router matrix or both defined already.

If you have a dual backbone configuration, where two backbone fabrics share edge fabrics, the LSAN fabric matrix and FC router matrix settings for the shared edge fabrics must be the same on both backbone fabrics. The matrix databases are *not* automatically propagated from one backbone fabric to another, so you must ensure that both backbone fabrics have the same matrix settings.

### NOTE

You can use LSAN zone binding along with LSAN tagging to achieve better scalability and performance. See [LSAN Zone Policies Using LSAN Tagging](#) for information about using the Enforce LSAN tag.

## FC Router Matrix Definition

Depending on the structure of the backbone fabric, you can specify pairs of FC routers that can access each other. For the metaSAN shown in [LSAN Zone Binding](#), the following FC Routers can access each other:

- FC router 1 and FC router 2
- FC router 3 and FC router 4

Because there is no device sharing between the two groups of FC routers, you can use the `fcrLsanMatrix` command with the `-fcr` option to create the corresponding FC Router matrix:

```
fcrLsanmatrix --add -fcr wwn1 wwn2
fcrLsanmatrix --add -fcr wwn3 wwn4
```

The variables `wwn1`, `wwn2`, `wwn3`, and `wwn4` are the WWNs of the four FC routers.

Now edge fabrics 1, 2, 3, 7, and 8 can access each other, and edge fabrics 4, 5, 6, and 9 can access each other; however, edge fabrics in one group cannot access edge fabrics in the other group. The edge fabrics can still communicate with the backbone fabric.

## LSAN Fabric Matrix Definition

With LSAN zone binding, you can specify pairs of fabrics that can access each other. Using the metaSAN shown in [LSAN Zone Binding](#) as an example, the following edge fabrics can access each other:

- Fabric 1 and Fabric 2
- Fabric 2 and Fabric 3
- Fabric 4 and Fabric 5
- Fabric 5 and Fabric 6

You can use the `fcrLsanMatrix` command with the `-lsan` option to create the corresponding LSAN fabric matrix.

```
fcrLsanMatrix --add -lsan 1 2
fcrLsanMatrix --add -lsan 2 3
fcrLsanMatrix --add -lsan 4 5
fcrLsanMatrix --add -lsan 5 6
```

Fabrics that are not specified are part of the default binding and can access other edge fabrics that are not specified. Thus, fabrics 7, 8, and 9 can access each other, but cannot access fabrics 1 through 6.

### ATTENTION

The `fcrLsanMatrix --add -lsan 0 0` command will erase the entire LSAN fabric matrix settings in the cache.

The FC router matrix and the LSAN fabric matrix are used together to determine which fabrics can access each other, with the LSAN fabric matrix providing more specific binding.

## Setting Up LSAN Zone Binding

1. Log in to the FC router as admin.
2. Enter the following command to add a pair of FC Routers that can access each other:

```
FCR:Admin> fcrLsanMatrix --add -fcr wwn1 wwn2
```

The variables `wwn1` and `wwn2` are the WWNs of the FC Routers.

3. Enter the following command to add a pair of edge fabrics that can access each other:

```
FCR:Admin> fcrLsanMatrix --add -lsan fid1 fid2
```

The variables `fid1` and `fid2` are the fabric IDs of the edge fabrics.

4. Enter the following command to apply the changes persistently:

```
FCR:Admin> fcrLsanMatrix --apply -all
```

```
FCR:Admin> fcrLsanMatrix --add -fcr 10:00:00:60:69:c3:12:b2 10:00:00:60:69:c3:12:b3
FCR:Admin> fcrLsanMatrix --add -lsan 4 5
FCR:Admin> fcrLsanMatrix --add -lsan 4 7
FCR:Admin> fcrLsanMatrix --add -lsan 10 19
FCR:Admin> fcrLsanMatrix --apply -all
```

## Viewing the LSAN Zone Binding Matrices

1. Log in to the FC router as admin.
2. Enter `fcrlsanmatrix --fabricview -fcr` to view the FC router matrix.
3. Enter `fcrlsanmatrix --fabricview -lsan` to view the LSAN fabric matrix.

The following example displays both fabric matrices.

```
fcr:admin> fcrlsanmatrix --fabricview -fcr
SAVED FCR PAIRS
=====
FCR FCR

10:00:c4:f5:7c:01:b0:71 (70) 10:00:c4:f5:7c:01:15:81 (230)
```

```
fcr:admin> fcrlsanmatrix --fabricview -lsan
LSAN MATRIX is activated
Fabric ID Fabric ID

 4 5
 4 7
 10 19
```

## Location-Embedded LSAN Zones

The location-embedded LSAN zones help to increase the limit of LSAN zones in the backbone fabric by providing a mechanism to store only the applicable LSAN zones in each FC router. You can configure the LSAN zones appropriately in the edge fabric, at the same time retaining all the FC router features and characteristics such as any-to-any connectivity. You can have more than 10,000 LSAN devices in the entire backbone fabric or metaSAN.

Every FC router switch in the backbone fabric maintains the entire LSAN database having all LSAN zone entries in all edge fabrics subject to the maximum LSAN zone limit. In many cases, the device sharing is done between a few edge fabrics and not with all the edge fabrics that are connected in the backbone. The Location- Embedded LSAN Zones feature allows you to specify or embed the remote fabric ID in the LSAN zone names that share devices. The FC router uses the location information embedded in the LSAN zone name to store only relevant LSAN zone entries. The advantage is that an individual FC router stores less LSAN zone entries and allows you to configure more LSAN zones per FC router which helps to exceed the maximum LSAN zone limit across the backbone fabric.

With location-embedded LSAN zones, the FC Router stores less LSAN zones, resulting in improved zone lookup during pair matching.

## Creating Location-Embedded LSAN Zones

To create location-embedded LSAN zones, perform the following steps:

1. Identify the location and use the `portcfgexport` command to get the remote fabric ID that must be embedded in the LSAN zone.
2. Specify the location in the zone name along with the RFID tag.

```
LSAN_RFID_100
LSAN_LOCAL_SITE_RFID_4
```



- If you want to embed the RFID tag in the middle of the zone name, the underscore (\_) character should be placed before the RFID tag and also appended after the valid fabric ID, as shown in the following example.

```
LSAN_SPEED_NEW_SITE_RFID_20_ZONENAME
```

- Append the RFID tag after the enforce and speed tags in case you need to use the enforce and speed tags along with the RFID tag.

```
LSAN_ENFORCE1_RFID_100
```

```
LSAN_SPEED12_RFID_4_SITENAME
```

#### NOTE

In the FC router, avoid using the RFID tag for configuring enforce or speed tags.

- The priority of processing in the FC router for zones having both enforce and RFID tags are as follows:
  - The FC router processes the LSAN zones based on enforce tags during zone merge. Zones not matching these tags are dropped.
  - The LSAN distribution between FC routers based on RFID is done only on the filtered zones obtained from the previous step.


#### NOTE

Speed tags remain unaffected during this process.

- Rename the LSAN zones in the paired edge fabric.

### Example Configuration of LSAN Zones Environment

**Figure 94: Sample Location-Embedded LSAN Zones Environment**



The following are the LSAN configurations made in the respective fabrics.

```
FID 1: LSAN_A : dev1 dev2 FID 3: LSAN_A : dev3 dev2
FID 1: LSAN_B : dev1 dev4 FID 3: LSAN_B : dev3 dev4
```

```
FID 2: LSAN_A : dev2 dev1 FID 4: LSAN_A : dev4 dev1
FID 2: LSAN_B : dev2 dev3 FID 4: LSAN_B : dev3 dev4
```

After you change the zone name by embedding the remote fabric ID and re-enabling the zone configuration in all the fabrics, LSAN zone configuration changes to the following.

```
FID 1: LSAN_A_RFID_2: dev1 dev2 FID 3: LSAN_RFID_2_A: dev3 dev2
FID 1: LSAN_B_RFID_4: dev1 dev4 FID 3: LSAN_RFID_4_B: dev3 dev4
FID 2: LSAN_RFID_1_A: dev2 dev1 FID 4: LSAN_A_RFID_1: dev4 dev1
FID 2: LSAN_RFID_3_B: dev2 dev3 FID 4: LSAN_B_RFID_3: dev3 dev4
```

FCR 1 has the following zones.

```
FID 1: LSAN_A_RFID_2: dev1 dev2
FID 1: LSAN_B_RFID_4: dev1 dev4
FID 2: LSAN_RFID_1_A: dev2 dev1
FID 4: LSAN_A_RFID_1: dev4 dev1
FID 4: LSAN_B_RFID_3: dev3 dev4
FID 3: LSAN_B_RFID_4: dev3 dev4
```

FCR 2 has the following zones.

```
FID 2: LSAN_RFID_1_A: dev2 dev1
FID 2: LSAN_RFID_3_B: dev2 dev3
FID 1: LSAN_A_RFID_2: dev1 dev2
FID 1: LSAN_B_RFID_4: dev1 dev4
FID 4: LSAN_A_RFID_1: dev4 dev1
FID 4: LSAN_B_RFID_3: dev3 dev4
FID 3: LSAN_RFID_2_A: dev3 dev2
FID 3: LSAN_RFID_4_B: dev3 dev4
```

FCR 3 has the following zones.

```
FID 3: LSAN_RFID_2_A: dev3 dev2
FID 3: LSAN_RFID_4_B: dev3 dev4
FID 2: LSAN_RFID_3_B: dev2 dev3
FID 4: LSAN_A_RFID_1: dev4 dev1
FID 4: LSAN_B_RFID_3: dev3 dev4
FID 1: LSAN_B_RFID_4: dev1 dev4
```

## Migrating LSAN Zones to Location-Embedded LSAN Zones in the Edge Fabric

You must identify a set of LSAN zones for which you want to embed the location. For such LSAN zones, you must reconfigure the LSAN zones in the edge fabric by completing the following steps. Using this procedure, you can scale beyond the maximum limits in the metaSAN.

1. In the FC router running a release prior to Fabric OS 7.4.0, use the `lsanZoneShow -d -f <fid>` command to find the remote fabric ID from which devices are imported for each LSAN zone per configured edge fabric.
2. Rename the LSAN zone name by appending the special `_RFID_` "imported FID" tag.
3. If the LSAN zone has the devices that are imported from more than one edge fabric, split them into multiple LSAN zones based on the number of imported fabrics in `RFID_` "imported FID" tag format and ensure that all newly created LSAN zones have only those imported device pairs.
4. When all the LSAN zones are converted into `RFID_` imported FID format, enable the zone configuration.
5. Repeat the previous steps in all the edge fabrics.

LSAN zones that are not configured with location information will operate as per legacy FC Router behavior.

## Limitations of Location-Embedded LSAN Zones

- The location-embedded LSAN zone feature is not applicable for backbone-to-edge fabric LSAN zones. If you configure an RFID in a backbone fabric LSAN zone, it will not impact anything on resource utilization.
- The entries of `fcProxyDevShow -a`, `fcPhyDevShow -a`, and `lsanZoneShow` may not be same across FC routers in the backbone fabric.
- If you need to change the fabric ID for an EX\_Port whose FID is already being used as an RFID value in other edge fabric LSAN zones, you must modify the LSAN zone configuration to new a RFID value. The steps are as follows:
  - a. Disable the EX\_Ports that need an FID change.
  - b. Configure the EX\_Ports with the new FID value.
  - c. Enable the zone configuration in the respective edge fabrics with the new RFID value.
  - d. Enable the EX\_Ports that were disabled in step a.
- If you configure a Backbone Fabric ID value for the RFID field, the LSAN zone is treated as a regular LSAN zone and will be distributed to other FCRs in the backbone fabric.

## Peer LSAN Zone Support

Starting with Fabric OS 7.4.0, the FC router supports the peer LSAN zones if you have configured them in the edge fabric. Peer zoning rules are applied by the edge fabric switch. The FC router treats peer LSAN zones as normal LSAN zones and imports the devices in the edge fabric as per a pair-matching algorithm. RSCNs of proxy devices present in peer zones generated by the FC Router are distributed to the edge fabric switches and controlled by the edge fabric name server as per peer zoning rules. The FC Router supports peer LSAN zones configured in a backbone fabric.

For more information, see [Peer Zoning](#).

### NOTE

If you are using peer zone over FC Router, both edges or backbones should use peer zone. You cannot combine regular LSAN zone on one end with a peer LSAN zone on another end.

## Proxy PID Configuration

When an FC router is first configured, the PIDs for the proxy devices are automatically assigned. Proxy PIDs (as well as phantom domain IDs) persist across reboots.

The most common situation in which you would set a proxy PID is when you replace a switch. If you replace the switch and want to continue using the old PID assignments, you can configure it to do so; this value remains in the system even if the blade is replaced. To minimize disruption to the edge fabrics, set the proxy PIDs to the same values used with the old hardware.

The `fcProxyConfig` command displays or sets the persistent configuration of proxy devices. Used with the `-s` option, it can also influence the assignment of the `xlata` domain port number (which determines the `Area_ID` field of the PID) and the `Port_ID` field. Like the PIDs in a fabric, a proxy PID must be unique. If the `slot` argument results in a duplicate PID, it is ignored. Proxy PIDs are automatically assigned to devices imported into a fabric, starting at `f001`.

Use the `fcXlataConfig` command to display or assign a preferred domain ID to a translate domain.

## Fabric Parameter Considerations

By default, EX\_Ports detect, autonegotiate, and configure the fabric parameters without user intervention. You can optionally configure these parameters manually.

- To change the fabric parameters on a switch in the edge fabric, use the `configure` command.  
Note that to access all of the fabric parameters controlled by this command, you must disable the switch using the `switchDisable` command. If executed on an enabled switch, only a subset of attributes is configurable.
- To change the fabric parameters of an EX\_Port on the FC router, use the `portCfgEXPort` command.

The backbone fabric PID mode and the edge fabric PID mode do not need to match, but the PID mode for the EX\_Port and the edge fabric to which it is attached must match. You can statically set the PID mode for the fabric by using the `-p` option with the `portCfgEXPort` command. Use the `-t` option to disable the negotiate fabric parameter feature; otherwise, the PID mode is autonegotiated. The various edge fabrics may have different PID modes.

Fabric parameter settings, namely, E\_D\_TOV (error-detect timeout value), R\_A\_TOV (resource-allocation timeout value), and PID format, must be the same on EX\_Ports and on the fabrics to which they are connected. You can set the PID format on an EX\_Port when you configure an inter-fabric link.

The default values for E\_D\_TOV and R\_A\_TOV for an EX\_Port or must match those values on other Fabric OS switches. You do not need to adjust these parameters for an EX\_Port unless you have adjusted them for the edge fabric.

Use the default values for R\_A\_TOV and E\_D\_TOV for all but very large fabrics (ones requiring four or more hops) or high-latency fabrics (such as ones using long-distance FCIP links).

## Inter-Fabric Broadcast Frames

The FC Router can receive and forward broadcast frames between edge fabrics and between the backbone fabric and edge fabrics. By default, broadcast frames are *not* forwarded from the FC router to the edge fabrics.

### Displaying the Current Broadcast Configuration

1. Log in to the FC Router as admin.
2. Enter the following command:

```
fcr:admin> fcrbcastconfig --show
```

This command displays only the FIDs that have the broadcast frame option enabled. The FIDs that are not listed have the broadcast frame option disabled.

### Enabling Broadcast Frame Forwarding

1. Log in to the FC Router as admin.
2. Enter the following command:

```
fcr:admin> fcrbcastconfig --enable -f fabricID
```

The *fabricID* variable is the FID of the edge or backbone fabric on which you want to enable broadcast frame forwarding. Broadcast frame forwarding is enabled by default.

### Disabling Broadcast Frame Forwarding

1. Log in to the FC Router as admin.
2. Enter the following command:

```
fcr:admin> fcrbcastconfig --disable -f fabricID
```

The *fabricID* variable is the FID of the edge or Backbone fabric on which you want to disable broadcast frame forwarding.

## Resource Monitoring

It is possible to exhaust resources, such as proxy PIDs. Whenever a resource is exhausted, Fabric OS generates an error message. The messages are described in the *Brocade Fabric OS Message Reference Manual*.

You can monitor FC router resources using the `fcrresourceshow` command. This command shows FC router resource limits and usage and includes the following:

- **LSAN zones and LSAN devices** —The information shows the maximum versus the currently used zones and device database entries. Each proxy or physical device constitutes an entry. If LSAN zones are defined in two edge fabrics, they are counted as two and not one. One device imported into multiple edge fabrics counts multiple times. The default maximum number of LSAN zones is 3,000. See [Setting the Maximum LSAN Count](#) for information on changing this limit.
- **Proxy device slots** — The physical and proxy devices use the 10,000 device slots. The information shows the maximum pool size for translate phantom node and port WWNs and shows the number of translate node and port WWNs from this pool.
- **Phantom NWWNs**
- **Phantom PWWNs**
- **Max proxy devices**
- **Max NR\_Ports**

The following example shows the use of the `fcrresourceshow` command to display physical port (EX\_Port) resources.

```
switch:admin> fcrresourceshow
Daemon Limits:
```

|                     | Max Allowed | Currently Used |
|---------------------|-------------|----------------|
| LSAN Zones:         | 7500        | 28             |
| LSAN Devices:       | 15000       | 51             |
| Proxy Device Slots: | 15000       | 20             |

|                   | WWN Pool Size | Allocated |
|-------------------|---------------|-----------|
| Phantom Node WWN: | 8192          | 5413      |
| Phantom Port WWN: | 32768         | 16121     |

```
Port Limits:
Max proxy devices: 4000
Max NR_Ports: 1000

Currently Used(column 1: proxy, column 2: NR_Ports):
```

|    |   |    |
|----|---|----|
| 0  | 0 | 34 |
| 1  | 3 | 34 |
| 4  | 0 | 0  |
| 5  | 0 | 0  |
| 6  | 0 | 0  |
| 7  | 0 | 0  |
| 8  | 6 | 34 |
| 9  | 6 | 34 |
| 10 | 6 | 34 |
| 11 | 6 | 34 |
| 12 | 6 | 34 |
| 13 | 6 | 34 |
| 14 | 6 | 34 |
| 15 | 6 | 34 |
| 16 | 8 | 34 |
| 17 | 8 | 34 |
| 18 | 8 | 34 |

|    |   |    |
|----|---|----|
| 19 | 8 | 34 |
| 20 | 8 | 34 |
| 21 | 8 | 34 |
| 22 | 8 | 34 |
| 23 | 8 | 34 |

## FC-FC Routing and Virtual Fabrics

If Virtual Fabrics is not enabled, FC-FC routing behavior is unchanged. If Virtual Fabrics is enabled, then in the FC-FC routing context, a base switch is like a backbone switch and a base fabric is like a backbone fabric.

If Virtual Fabrics is enabled, the following rules apply:

- EX\_Ports can be configured only on the base switch.  
When you enable Virtual Fabrics, the chassis is automatically rebooted. When the switch comes up, only one default logical switch is present, with the default fabric ID (FID) of 128. All previously configured EX\_Ports are persistently disabled with the reason *ExPort in non-base switch*. You must explicitly create a base switch, move the EX\_Ports to the base switch, and then enable the ports.  
If you move existing EX\_Ports to any logical switch other than the base switch, these ports are automatically disabled. If you want to change an EX\_Port on the logical switch to be a non-EX\_Port, you must use the `portCfgDefault` command. You cannot use the `portCfgExPort` command because that command is allowed only on the base switch.
- From FOS v9.0.x onwards, EX\_Ports in a base switch receives a 10-bit dynamic area address instead of a zero bit address. An EX\_Port in the base switch is assigned any 10-bit area from 0x0000 - 0xFFC0, with the following exceptions:
  - Areas 0x8000-8300
  - Areas 0x8040-8340
  - Areas 0x8080-8380
  - Areas 0x80C0-83C0
  - Areas 0xC500, 0xC540, 0xC580 and 0xC5C0
 These areas are reserved for internal lookups.
- EX\_Ports can connect to a logical switch that is in the same chassis or in a different chassis. However, the following configuration rules apply:
  - If the logical switch is on the same chassis, the EX\_Port FID must be set to a different value than the FID of the logical switch to which it is connecting.
  - If the logical switch is on a different chassis, no FID for any logical switch in the FC router backbone fabric can be the same as the FID of the logical switch to which the EX\_Port is connecting.
- EX\_Ports in FC routers and those in a base switch cannot connect to any edge fabric with logical switches configured to use XISLs.  
If you connect an EX\_Port to an edge fabric, you must ensure that there are no logical switches with XISL use enabled in that edge fabric. If any logical switch in the edge fabric allows XISL use, the EX\_Port is disabled. See [Configuring a Logical Switch for XISL Use](#) for instructions on disallowing XISL use.

Because XISL use is disallowed, dedicated links must be configured to route traffic across switches in the same logical fabric.

- Backbone-to-edge routing is not supported in the base switch. See [Figure](#) for information about how to configure legacy FC routers to allow backbone-to-edge routing with Virtual Fabrics.
- All FC router commands can be executed only in the base switch context.
- The `fcconfigure --bbfid` command is not allowed when Virtual Fabrics is enabled. Instead, use the `lscfg` command to configure the FID.
- From FOS v9.0.x onwards, if a port is configured as an EX\_Port and if the administrator moves the port from the base switch to another logical switch, the operation fails. You must remove the EX\_Port configuration before moving it to another logical switch.
- When a VF enabled switch is part of the backbone, only the base switch with active EX\_Ports will be able to route frames between edge fabrics. If there are other base switches that are participating in the fabric and do not have active EX\_Ports, those switches will participate as transient switches to forward frames to the next hop till they can be routed to the intended remote edge fabric. If any other default or logical switch is introduced as a hop between edge fabrics, the newly introduced switch which is not the base switch will not merge with the base switch fabric. This will cause frames to be dropped.

## Logical Switch Configuration for FC Routing

[Figure](#) shows an example of two chassis partitioned into logical switches. This configuration allows the device in Fabric 128 to communicate with the device in Fabric 15 without merging the fabrics.

The following conditions are considered:

- The base switch in Physical chassis 1 serves as an FC router and contains EX\_Ports that connect to logical switches in the two edge fabrics, Fabric 128 and Fabric 15.
- The other logical switches in Fabric 128 and Fabric 15 must be connected with physical ISLs, and they do not use the XISL connection in the base fabric.
- The logical switches in Fabric 1 are configured to allow XISL use. You cannot connect an EX\_Port to these logical switches, so the device in Fabric 1 cannot communicate with the other two devices.

**Figure 95: EX\_Ports in a Base Switch**





Figure shows a logical representation of the physical chassis and devices in Figure. As shown in Figure, Fabric 128 and Fabric 15 are edge fabrics connected to a backbone fabric. Fabric 1 is not connected to the backbone, so the device in Fabric 1 cannot communicate with any of the devices in the other fabrics.

**Figure 96: Logical Representation of EX\_Ports in a Base Switch**



## Backbone-to-Edge Routing with Virtual Fabrics


Backbone-to-edge routing is not supported in the base switch, unless you use a legacy FC Router. A *legacy FC Router* is an FC Router configured on a Brocade 7500 switch.

Base switches can participate in a backbone fabric with legacy FC Routers. You cannot connect devices to the base switch because the base switch does not allow F\_Ports. You can, however, connect devices to the legacy FC Router, thus enabling backbone-to-edge routing.

If you connect a legacy FC Router to a base switch, you must set the backbone FID of the FC Router to be the same as that of the base switch.

In Figure, no devices can be connected to the backbone fabric (Fabric 8) because base switches cannot have F\_Ports. The following figure shows an FC Router in legacy mode connected to a base switch. This FC Router *can* have devices connected to it, and so you can have backbone-to-edge routing through this FC Router. In this figure, Host A in the backbone fabric can communicate with device B in the edge fabric with FID 20; Host A cannot communicate with device C, however, because the base switches do not support backbone-to-edge routing.



**Figure 97: Backbone-to-Edge Routing Across Base Switch Using FC Router in Legacy Mode**

If a backbone fabric has both a Virtual Fabrics-enabled FC Router and a Virtual Fabrics-disabled FC Router, EX\_Ports are not allowed from the base switch of the Virtual Fabrics-enabled FC Router to the same edge fabric that is performing backbone-to-edge routing with the Virtual Fabrics-disabled FC Router.

For example, in the preceding figure, you should not connect EX\_Ports from the base switch (logical switch 4 or logical switch 8) to the edge fabric FID 20, which is performing backbone-to-edge routing with a legacy FC Router.

## Upgrade and Downgrade Considerations for FC-FC Routing

When you upgrade to Fabric OS to a new release in a non-disruptive manner, the EX\_Ports remain functional, and can continue to perform all the FC routing operations on the FCR switch. Similarly, downgrading to the previous available release causes the EX\_Ports to remain functional, and continue to perform all the FC routing operations on the FCR switch.

You should save your FC-FC routing configuration (using the `configUpload` command) before performing any downgrades.

Before upgrading the Brocade G620 switch to 8.1.0 or later, ensure that an Integrated Routing License is installed if EX\_Port is enabled in the switch. This avoids EX\_Port getting disabled after upgrade.

## How Replacing Port Blades Affects EX\_Port Configuration

You must replace the same type of port blades. If you need to swap port blades of different types, you must remove the existing port configurations and then swap the blade types if it is a non-VF backbone switch. In the VF backbone switch, you must move all the ports to the default switch before the blade is replaced.

## Displaying the Range of Output Ports Connected to xlate Domains

The edge fabric detects only one front domain from an FC router connected through multiple output ports. The output port of the front domain is not fixed to 0; the values can be in a range from 129 through 255. The range of the output ports

connected to the xlate domain is from 1 through 128. This range enables the front domain to connect to 127 remote xlate domains.

1. Log in to a switch in the edge fabric.
2. Enter the `lsDbShow` command on the edge fabric.

In the `lsDbShow` output, ports in the range from 129 through 255 are the output ports on the front domain.

### Displaying the Range of Output Ports

The following example shows the range of output ports.

```
linkCnt = 2, flags = 0x0
LinkId = 53, out port = 1, rem port = 35, cost = 500, costCnt = 0, type = 1
LinkId = 57, out port = 129, rem port = 18, cost = 500, costCnt = 0, type = 1
```

### Displaying Ports on the Edge Fabric

The following example shows the use of the `lsDbShow` display on the edge fabric. The front domain, domain 3, has two links representing two EX\_Port connections with output ports 129 and 132.

```
Domain = 3, Link State Database Entry pointer = 0x100bbcc0
.....
linkCnt = 4, flags = 0x0
LinkId = 199, out port = 129, rem port = 2, cost = 10000, costCnt = 0, type = 1
LinkId = 199, out port = 132, rem port = 3, cost = 10000, costCnt = 0, type = 1
LinkId = 2, out port = 1, rem port = 2, cost = 10000, costCnt = 0, type = 1
LinkId = 1, out port = 32, rem port = 2, cost = 10000, costCnt = 0, type = 1
```

## FCR Migration Mode

To non-disruptively migrate the existing FCR traffic to new FCR switches, you must enable the migration mode on an FCR switch. The migration mode allows you to proceed with decommission actions leading to the removal of the FCR switch from the backbone fabric.

When an FCR switch is enabled for migration mode, the traffic that originates from the edge fabric and the traffic that is destined toward remote fabrics get migrated on alternative IFLs in the edge fabric going to the other FCR switches in the backbone fabric. Once the traffic migration is completed, you must decommission the older FCR switch.

### NOTE

The migration mode is supported only in edge-to-edge topologies.

Perform the following steps to enable FCR migration mode on an FCR switch:

1. Identify the FCR switch that must be decommissioned from the topology.
2. Use the `fcrconfigure --enable -migrationmode` command to enable migration mode.

```
switch:admin> fcrconfigure --enable -migrationmode
```

```
switch:admin> fcrconfigure --show
Backbone fabric ID:100
Shortest IFL feature is disabled.
Migration mode is enabled.
```

3. In the edge fabric switches, use the `portperfshow <port_number>` command to confirm that the traffic has migrated away from the IFLs connecting to the migration-mode-enabled FCR switch.
4. In the migration-mode-enabled FCR switch, disable all EX\_Ports connected to the edge fabric where all the traffic was migrated earlier. There are no more IFL paths at the edge fabric to the migration-mode-enabled FCR switch.

5. On the FCR switch, use the `fcrconfigure --disable -migrationmode` command to disable migration mode.

```
switch:admin> fcrconfigure --disable -migrationmode
```

```
switch:admin> fcrconfigure --show
```

```
Backbone fabric ID:100
```

```
Shortest IFL feature is disabled.
```

```
Migration mode is disabled.
```

6. Decommission the FCR switch.

## Administering FICON Fabrics

---

This section describes the IBM Fibre Connection (FICON) protocol that is used between IBM (and compatible) mainframes and storage.

To administer FICON, you must have one of the following roles that is associated with your login name on the switch:

- Admin
- Operator
- SwitchAdmin
- FabricAdmin

The User and BasicSwitchAdmin roles are view-only. The ZoneAdmin and SecurityAdmin roles have no access.

In Virtual Fabrics, the FICON commands apply to the current logical or specified switch only.

## FICON Requirements

The following are requirements for configuring FICON on Brocade switches and directors:

- FICON logical switch (LS)  
You must use a FICON LS in FICON environments. You cannot connect a FICON channel or device to a port in a non-FICON LS. The default switch cannot be changed to a FICON LS, so Virtual Fabrics must be enabled and used.

### NOTE

Converting an existing logical switch to a FICON logical switch (`lscfg -change -ficon` command) is deprecated in Fabric OS 9.0.0. For Gen 6 platforms, you can continue to use the `lscfg -change -ficon` command to convert a logical switch to a FICON logical switch. However, for Gen 7 platforms and subsequent releases of Fabric OS, use the `lscfg -create -ficon` command to create a new logical switch.

- Configuring the domain ID  
Give all switches a unique domain ID.
- Configuring the SCC policy  
An SCC policy is created when you create a FICON LS. The SCC policy initially contains only the current domain. Before a domain can connect to the fabric, you must update the SCC policy list to contain the same switches as the SCC policy in the fabric that you are joining.
- Port addresses  
Ports in a FICON LS must have an address bound to them before they can be enabled. When moving ports into a FICON LS, use the `portAddress --bind` command to bind port addresses before enabling the port.

### NOTE

You cannot bind to port addresses 0xFE and 0xFF in a FICON LS.

## FICON Best Practices

Follow these best practices for configuring FICON on Brocade switches and directors:

- Switch ID  
Make the switch ID in the IOCP or HCD the same as the hexadecimal equivalent of the domain ID. Although not required, this makes it easier to correlate switches to the IOCP. For example, if the domain ID is 32 (0x20), the switch ID should be 20.
- Firmware

- Use only qualified firmware for FICON.
- All versions of firmware for switches that Brocade SANnav Management Portal is managing should be in the SANnav repository.
- All switches in a fabric should be at the same code level.
- When upgrading code more than two levels, plan the upgrades so that no switch in the fabric is more than one release apart.
- Always check the version of firmware on a switch before putting the switch into production.
- If a firmware upgrade is required on a new switch installation, upgrade the firmware before making any configuration changes.
- CHPID and control unit placement

Place CHPID and control unit connections on the same port card when possible to minimize the impact of a card failure. Using the same port card also allows for maximum use of Brocade local switching to minimize the latency and optimize the performance. Consult your Brocade SE for more specifics on maximizing the benefits of local switching with your specific configuration. This should be done in conjunction with your I/O Configuration Data Set (IOCDs) planning.
- Default switch

Leave the default switch FID at 128.
- Port fencing

Use the default MAPS aggressive policy (`dfit_aggressive_policy`) to set the port fencing parameters. For additional information regarding customizing MAPS policies for FICON, refer to [SAN and Fabric Resiliency Best Practices for IBM b-type Products](#).
- WWN-based persistent port area IDs (PIDs)

Do not enable WWN-based persistent PIDs for FICON environments.
- FICON File Access Facility (FAF)

**NOTE**

Fabric OS 9.0.0 was the last release to support FAF.

- If Fabric OS 9.1.0 or later is present in any of the active or standby CPs, do not use PDCM functionalities.
- For Fabric OS 9.1.0 and later, the `configUpload` command does not upload the FAF files.
- For Fabric OS 9.1.0 and later, the `configDownload` command does not update the FAF files (except the IPL file).
- Prohibit Dynamic Connectivity Mask (PDCM)

**NOTE**

Fabric OS 9.0.0 was the last release to support the PDCM used for connectivity management as reflected in the Allow/Prohibit Configuration Matrix option in WebTools. Effectively, the Allow/Prohibit matrix management feature was deprecated in Fabric OS 9.1.0. For more information on managing connectivity for FICON, see [Zoning for FICON](#).

- If Fabric OS 9.1.0 or later is present in any of the active or standby CPs, do not use FAF functionalities.
- For Fabric OS 9.1.0 and later, the `configUpload` command uploads IPL files with the default PDCM configuration.
- For Fabric OS 9.1.0 and later, the `configDownload` command downloads IPL files with the default PDCM configuration.
- Zoning for FICON

Brocade recommends creating zones for FICON that mimic the HCD/IOCDs configuration to manage FICON environments efficiently. Typically, large environments create zones to match the IOCDs configuration or create multiple channel/control unit zones rather than creating lots of 1:1 zones. For smaller environments, a single FICON zone containing all of the FICON ports may be sufficient.

**NOTE**

Default zoning is not supported for FICON when using a Base Fabric.

Engage your local Brocade Representative to ensure proper zoning configuration for your FICON and mainframe environments as the best practices often evolve when new features and functions are introduced.

To configure zoning for your fabric, follow the standard FICON zoning practices that are published by Brocade and IBM (refer to [IBM b-type Gen 7 Installation, Migration, and Best Practices Guide](#)).

## FICON and FEC-via-TTS

All known devices that support Forward Error Correction (FEC) via Transmitter Training Signal (TTS) have it enabled by default; however, not all devices that support 16GFC support FEC-via-TTS. FEC is enabled by default on Brocade switches, but FEC-via-TTS is disabled by default.

Enabling or disabling FEC-via-TTS momentarily disables the port. As with any switch action that disables a port, CHIPIDs connected to that port should be configured offline. Operators should quiesce all traffic or vary offline any link address associated with a port that will be disabled.

The best practice is to always leave FEC enabled; however, a best practice regarding FEC-via-TTS depends on current and future plans regarding 16GFC attachments. If the only plans for 16GFC attachments are with HBAs and devices that will support FEC-via-TTS, then a data center may choose to establish a best practice to enable it during installation. Otherwise, FEC-via-TTS should be enabled only on ports intended for connections to HBAs and on devices that also support FEC-via-TTS.

Because enabling FEC-via-TTS is disruptive, the use of this feature should be considered before putting a port into production.

FEC-via-TTS is supported on the FICON Express16s channel card, which is used for both FICON and FCP.

This feature is not FICON dependent and is supported in FCP environments such as PPRC.

In Gen 6 platforms, 32G F\_Ports are automatically enabled with FEC based on 32G FC standards; there are no `portcfgfec` commands to enter if connecting to a 32G device. For Gen 6 platforms connected to 16G channels and CUs, the 32G switch port will negotiate to 16G and the `portcfgfec` command is required to enable FEC-via-TTS.

In Gen 7 platforms, 64G F\_Ports are automatically enabled with FEC based on 64G FC standards; there are no `portcfgfec` commands to enter if connecting to a 64G device. For Gen 7 platforms connected to 16G channels and CUs, the 64G switch port will negotiate to 16G and the `portcfgfec` command is required to enable FEC-via-TTS.

For further information on FEC limitations and instructions on enablement and disablement, refer to this administration guide and the *Brocade Fabric OS Command Reference Manual*.

## Considerations for a FICON Logical Switch

The following considerations apply to a FICON logical switch:

- In Fabric OS 9.0.x and later, you can only enable FMS mode in a FICON logical switch. Therefore, you cannot download (configDownload) a configuration with FMS mode enabled in a non-FICON logical switch.
- In Fabric OS 9.0.x and later, FMS mode can only be enabled in the FICON logical switch. If FMS mode is enabled in a non-FICON logical switch, the upgrade to Fabric OS 9.0.x or later will be blocked.
- In pre-Fabric OS 9.0.x releases, enabling HIF mode is sufficient for a FICON channel to successfully connect to a logical switch. In Fabric OS 9.0.x, it must be a FICON logical switch for a FICON channel to successfully connect to a logical switch in addition to the HIF mode enabled. If HIF mode is enabled and FMS mode is disabled in a non-FICON LS and a firmware upgrade is attempted, a warning message displays prompting you to cancel or proceed with the upgrade. After the upgrade, FICON channels and devices may continue to be attached and operational. However, if a channel tries to reattach for any reason, it is allowed only to attach to a port in a FICON LS.
- Setting HIF mode locks the IDID, fabric-wide consistency policy, and SCC policy settings to ensure that the fabric is of high integrity so that it can connect with the FICON channel. You cannot change these settings without disabling HIF mode. If HIF configuration requirements have not been met, an error message describes what you must configure for the command to succeed. For example, the following message states that an IDID, SCC policy, or fabric-wide consistency policy has not been configured for the switch. Perform additional configuration if required, and then enable HIF mode.

```
Error: Unable to set HIF Mode. No valid IDID settings,
SCC policy and/or Fabric wide(SCC:S) configuration
```

- If a FICON logical switch is defined at the time of the downgrade, firmware downgrade is blocked to releases earlier than Fabric OS 8.1.0. You are prompted to delete the existing FICON logical switch before proceeding with the downgrade.

```
Downgrade is not allowed due to the presence of FICON logical switch. Please delete
the FICON logical switch using the command "lscfg --delete [-FID]" and retry.
```

- If FICON mode logical switch creation or modification is attempted while the standby CP is running a software version that does not support the FICON logical switch feature, a message is displayed.  
The Standby CP firmware version does not support this feature.
- If a standby CP is brought up with a firmware version that does not support FICON logical switch while the active CP already has a FICON logical switch, HA is forced out of sync, and a RASlog message is displayed.  
[SWCH-1041], HA state out of sync: Standby CP (ver = 11) does not support FICON Logical Switch. (active ver = 12).
- Configuration download is supported to the extent that the switch-wide attributes (FICON LS) can be downloaded. You can use the `configdownload` command to download three types of configuration files: Virtual Fabric configuration parameters, system configuration parameters, and port-to-area addressing mode configuration files. You must download all types of configuration data for the system to behave as expected. You can use the `configdownload -vf` command to download the FICON LS attribute through the virtual configuration files, the `configdownload` command to download system configuration files, and the `configdownload -map` to download the port-to-area addressing mode configuration files.

Similarly, the `configUpload` command uploads configuration data to a file and supports Virtual Fabric configuration parameters and system configuration parameters. Use the `configUpload -vf` command to upload Virtual Fabric configuration parameters, the `configUpload` command to upload the system configuration files, and the `configUpload -map` command to upload the port-to-area addressing mode configuration files. For more information, refer to the *Brocade Fabric OS Command Reference Manual*.

## Preparing a Switch for FICON

Follow these steps to verify and prepare a switch for use in a switched point-to-point FICON environment.

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Verify that the management IP addresses have been set up.
3. Verify that the switches can be seen by your management tool.
4. Verify that the switches have the correct version of Fabric OS.
5. Add feature keys, if applicable.
6. Enter the `configUpload` command to save a baseline of the switch configuration.
7. To create a FICON logical switch, see [Creating a FICON Logical Switch](#).

## Creating a FICON Logical Switch

Fabric OS 9.0.0 and later require a FICON logical switch (LS) for FICON environments. FICON host channels and devices cannot attach to ports in a non-FICON LS. Also, a FICON LS feature provides simplified configuration of logical switches for use in FMS mode. A `-ficon` option is provided for the `lscfg` command, which creates a logical switch with FICON configurations. This includes such configuration attributes as an insistent domain ID (IDID), fabric binding, high-integrity fabric mode, a device-based routing (DBR) policy, and 256-area limited addressing mode.

For Gen 6 platforms, you can use the `lscfg --change -ficon` command to convert a logical switch to a FICON logical switch. For Gen 7 platforms, converting an existing logical switch to a FICON logical switch (using the `lscfg --change -ficon` command) is deprecated in Fabric OS 9.0.x. However, for Gen 7 platforms and subsequent releases



of Fabric OS, use the `lscfg --create -ficon` command to create a new logical switch. FICON logical switches are supported on all platforms that support Virtual Fabrics in Fabric OS 8.1.0 or later. FMS mode is enabled by default when you create a FICON logical switch. A FICON logical switch cannot be the default switch in the chassis, nor can it be the base switch of the chassis.

A FICON logical switch provides auto-configuration of the following, which simplifies the process of establishing a logical switch in the FICON environment:

- Insistent domain ID (IDID) mode is set to on.
- The routing policy is set to device-based routing.
- The area mode is set to zero-based addressing.
- The SCC security policy is autodefined with only its own WWN.
- The created SCC policy is activated.
- The fabric-wide consistency policy is set to the strict SCC policy.
- High Integrity Fabric (HIF) mode is enabled.
- FMS mode is enabled if a FICON control unit port (CUP) license is installed on the chassis.

After the FICON mode logical switch is created, each port migrating into the logical switch must be bound to an address. Use the `portaddress` command to bind addresses. Any port that is not bound is kept disabled in the switch until it is bound to an address.

When a FICON logical switch is created, you cannot disable HIF mode or change the address mode to any mode other than zero-based addressing mode. As a result, the `configure` command will not list the following parameters:

- IDID mode
- 256 area (port-based addressing mode is not provided)
- HIF mode

At the end of FICON logical switch creation, the FICON logical switch is kept offline to allow additional configuration operations. By contrast, a non-FICON logical switch is online and ready for use after it is created.

To create a FICON LS, complete the following steps.

1. Connect to the physical chassis, and log on using an account with the chassis-role permission.
2. Enter the `lscfg` command with the `-ficon` option to create a FICON logical switch.

```
switch_128:FID128:admin> lscfg --create 50 -ficon
A Logical switch with FID 50 will be created with ficon configuration.
Would you like to continue [y/n]?: y
About to create switch with fid=50. Please wait...
Logical Switch with FID (50) has been successfully created.
Logical Switch has been created with ficon configurations.
Please configure the Logical Switch with appropriate switch
and protocol settings before activating the Logical Switch.
```

3. Use the `lscfg` command with the `--show` option to verify the logical switch creation.

The following sample is truncated.

```
switch_128:FID128:admin> lscfg --show

Created switches FIDs(Domain IDs): 128(ds)(120) 50(fs)(1)
Slot 1 2 3 4 5 6 7 8

Port
0 | | | | 128 | 128 | | | 128 |
1 | | | | 128 | 128 | | | 128 |
```



```
2 | | | | 128 | 128 | | | 128 |
```

You can see that logical switch 50 is created as a FICON logical switch, indicated by 50 (fs) .

4. Use the `setcontext` command to set context to the FICON logical switch.

```
switch_128:FID128:admin> setcontext 50
switch_50:FID50:admin>
```

5. Enter the `switchEnable` command to set the switch online. This enables the port to be available online when you enable a port.
6. Use the `switchshow` command to display switch information.

```
switch_50:FID50:admin> switchshow
switchName: switch_50
switchType: 165.0
switchState: Offline
switchMode: Native
switchRole: Disabled
switchDomain: 1 (unconfirmed)
switchId: fffc01
switchWwn: 10:00:00:27:f8:f0:a3:d1
zoning: OFF
switchBeacon: OFF
FC Router: OFF
HIF Mode: ON
Allow XISL Use: ON
LS Attributes: [FID: 50, Base Switch: No, Default Switch: No, Ficon Switch: Yes, Address Mode 1]
```

```

Index Slot Port Address Media Speed State Proto
=====
=====
No ports found in the system!!!

```

You can see that HIF mode is enabled and that Address Mode is set to 1 (zero-based address mode).

7. To configure FICON on the new logical switch, see [Configuring FICON](#) for step-by-step instructions.

## Configuring FICON

This section provides detailed steps and commands to configure a switch for FICON operation. The following steps assume that you used your hardware reference manual to perform the initial setup of the switch and that you performed all the steps from [Preparing a Switch for FICON](#).

For basic steps and commands in a checklist format to quickly configure a switch for fabric and possible FICON operation, see the [Basic Switch Configuration](#) appendix, and then return to this chapter for detailed FICON configuration procedures.

Use the worksheet in the [Configuration Information Record](#) appendix to record your configuration information.

In a FICON environment, logical switches must be created as a FICON logical switch (LS). Creating a FICON LS sets many of the configuration parameters to the correct values; however, you must still set some parameters, such as domain ID.

### NOTE

You cannot configure the default logical switch as a FICON LS.

For a list of FICON-related Fabric OS commands, refer to the *Brocade Fabric OS Command Reference Manual*.



### CAUTION

Configuring the switch for FICON is a disruptive process. The switch must be disabled before you configure switch parameters.

1. Connect to the switch, and log on using an account that is assigned to the admin role.
2. Enter the `switchDisable` command. Disable the switch to access all the switch parameters.
3. Configure the switch and chassis name, if applicable.
4. Use the `aptpolicy` command to display the routing policy.

```
switch_50:FID50:admin> aptpolicy
Current Policy: 1

3 : Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy (FICON support only)
3: Exchange Based Routing Policy
```

5. Set the routing policy by entering the `aptPolicy` command.

```
switch_50:FID50:admin> aptpolicy 2
```

Port-based routing (PBR), device-based routing (DBR), and exchange-based routing (EBR) are qualified for IBM z Systems; however, reference your system qualification letter for current support information. The recommended policy is DBR (`aptPolicy 2`).

If FICON emulation features (IBM z/OS Global Mirror or Tape Pipelining) are enabled on an FCIP tunnel in the switch, set either PBR (`aptPolicy 1`) or DBR (`aptPolicy 2`) in the switch.

6. Configure Dynamic Load Sharing.

The recommended best practice is to enable Dynamic Load Sharing (DLS); however, DLS is supported only when Lossless is enabled.

- To enable Lossless with DLS, use the `dlsSet --enable -lossless` command.
- If you do not enable Lossless, use the `dlsReset` command.

### NOTE

If you do not enable Lossless DLS, the routing policy must be port-based routing (`aptPolicy 1`).

7. Set In-Order Delivery using the `iodSet` command.
8. Configure the switch parameters using the `configure` command, and enter the responses in the following table when prompted.

**Table 145: FICON Switch Parameters**

| Parameter        | Response | Comment                                                                                                                                                                                                                                          |
|------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fabric parameter | Yes      | Prompts for the fabric parameters.                                                                                                                                                                                                               |
| Domain           | –        | The domain ID is the switch address. The recommended best practice is to set the switch address and domain ID to be the same. The domain ID is entered in decimal, so a switch address of hexadecimal 0x22 would have a domain ID of decimal 34. |

| Parameter                | Response | Comment                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| R_A_TOV                  | 10000    | <i>Do not change.</i> The Resource Allocation TimeOut Value (R_A_TOV) is entered in milliseconds. The R_A_TOV works with the variable E_D_TOV to determine switch actions when presented with an error condition.<br>Allocated circuit resources with detected errors are not released until the time value has expired. If the condition is resolved before the timeout, the internal timeout clock resets and waits for the next error condition.                                                                                                                                               |
| E_D_TOV                  | 2000     | The Error Detect TimeOut Value (E_D_TOV) is a timeout value that is entered in milliseconds. This timer is used to flag a potential error condition when an expected response is not received within the set time limit. If the time for an expected response exceeds the set value, an error condition occurs. Leave this value at 2 seconds (2000) unless connected to extension equipment. In some cases, when connecting to extension equipment, it must be set to 5 seconds (5000). This field should be changed only after consulting a Certified Fabric Professional for Brocade products. |
| WAN_TOV                  | 0        | The Wide-Area Network TimeOut Value (WAN_TOV) is set in milliseconds. The WAN_TOV is the maximum frame timeout value for a WAN, if any, interconnecting the Fibre Channel islands.                                                                                                                                                                                                                                                                                                                                                                                                                |
| MAX_HOPS                 | 7        | This parameter should not be confused with the one-hop FICON cascading limitation. There are configurations that are more than one hop because more than two domain IDs are used in the path that fit within the IBM support limitation of "one hop." Also, Linux on IBM z Systems using FCP channels is not restricted to "one hop."                                                                                                                                                                                                                                                             |
| Data field size          | 2112     | <i>Do not change.</i> Specifies the largest possible value, in bytes, for the size of a type 1 (data) frame. The switch advertises this value to other switches in the fabric during construction of the fabric and to other devices when they connect to the fabric. Setting this value smaller than 2112 may result in decreased performance.                                                                                                                                                                                                                                                   |
| Sequence level switching | 0        | <i>Do not change.</i> When set to 1, frames of the same sequence from a particular source are transmitted together as a group. When set to 0, frames are interleaved among multiple sequences when transmitted. Under normal conditions, sequence-level switching should be disabled for better performance.                                                                                                                                                                                                                                                                                      |
| Disable device probing   | 1        | The recommended best practice is to disable device probing. When disable device probing is set to 1, devices that are not registered with the Name Server are not present in the Name Server database.                                                                                                                                                                                                                                                                                                                                                                                            |
| Suppress Class F traffic | 0        | <i>Do not change.</i> By default, the switch can send Class F frames. When this option is turned on (set to 1), Class F traffic is converted to Class 2 traffic before being transmitted to support remote fabrics that involve ATM gateways.                                                                                                                                                                                                                                                                                                                                                     |
| Per-frame route priority | 0        | <i>Do not change.</i> In addition to the virtual channels used in the frame routing priority, when this value is set to 1, support is also available for per-frame-based prioritization.                                                                                                                                                                                                                                                                                                                                                                                                          |

| Parameter                   | Response | Comment                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Long-distance fabric        | 0        | <i>Do not change.</i> The recommended best practice is to configure individual ports for long distance when cascading at extended distances.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Insistent Domain ID Mode    | Yes      | In a FICON environment, you must set the domain ID to be insistent.<br>The default value is set automatically when you create a FICON LS.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| High-Integrity Fabric Mode  | Yes      | Before you set HIF mode, configure the following attributes: <ul style="list-style-type: none"> <li>• An insistent domain ID (IDID)</li> <li>• A fabric-wide consistency policy =&gt; SCC:S (strict mode)</li> <li>• A valid SCC policy (configured and activated)</li> </ul> Setting HIF mode locks the IDID, fabric-wide consistency policy, and SCC policy settings to ensure that the fabric is of high integrity so that it can connect with the FICON channel. You cannot change these settings without disabling HIF mode.<br>This value is set automatically when you create a FICON LS. |
| Virtual Channel parameters  | No       | <i>Do not change.</i> The switch allows fine-tuning of a specific application through virtual channel parameters. The default virtual channel settings have already been optimized for switch performance. Changing these values can improve switch performance, but can also degrade performance.                                                                                                                                                                                                                                                                                               |
| F_Port login parameters     | No       | With Fabric OS v6.1.0 and later, it is not necessary to make changes to F_Port login parameters. The only other reason to answer yes (y) to this prompt is if NPIV is being used and there is a need to change the default number of NPIV logins.                                                                                                                                                                                                                                                                                                                                                |
| Maximum logins per switch   | —        | Change this parameter only if NPIV is being used and the number of fabric logins will exceed the default or there is a need to limit the number of logins.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Maximum logins per port     | —        | Change this parameter only if NPIV is being used and the number of port logins will exceed the default or there is a need to limit the number of logins.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Maximum logins per second   | —        | For Fabric OS v6.1.0b and later, do not change this parameter. Leave it at the default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Login stage interval        | —        | <i>Do not change.</i> Specifies the stage interval in the staged F_Port login, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Zoning operation parameters | No       | <i>Do not change.</i> Disables and enables NodeName Zone Checking.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| RSCN transmission mode      | No       | Configures end-device RSCN Transmission Mode values such as: <ul style="list-style-type: none"> <li>• RSCN contains only a single PID</li> <li>• RSCN contains multiple PIDs (default).</li> <li>• Fabric addresses RSCN</li> </ul>                                                                                                                                                                                                                                                                                                                                                              |
| Arbitrated Loop parameters  | No       | Parameters include Send FAN frames, Enable CLOSE on OPEN received, and Always send RSCN. Refer to the <code>configure</code> command in the <i>Brocade Fabric OS Command Reference Manual</i> for more information.                                                                                                                                                                                                                                                                                                                                                                              |
| System services             | No       | Enables or disables the read link status (RLS) probing performed by the FCP daemon.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Portlog events enable       | No       | Determines whether various types of port events are logged.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Parameter           | Response | Comment                                                                                                                                                                              |
|---------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SSL attributes      | No       | Enables or disables Secure Socket Layer attributes.                                                                                                                                  |
| cfigload attributes | No       | Sets configuration upload and download parameters. Refer to the <code>configureChassis</code> command in the <i>Brocade Fabric OS Command Reference Manual</i> for more information. |

9. Enter the `secPolicyCreate` command to configure the SCC policies on all switches to limit connectivity to only the switches in the selected fabric.

```
switch:admin> secPolicyCreate SCC_POLICY, member;...;member
```

#### NOTE

When adding a new switch to a fabric, you also must add the new switch to the SCC policy in the existing fabric using the `secPolicyAdd` command.

In the command syntax, `member` indicates a switch that is permitted to join the fabric. Specify switches by WWN, domain ID, or switch name. Enter an asterisk (\*) to indicate all the switches in the fabric. To create a policy that includes all switches in the fabric, enter the following command:

```
switch:admin> secPolicyCreate SCC_POLICY "*"
```

10. Save or activate the new policy by entering either the `secPolicySave` command or the `secPolicyActivate` command.

If neither of these commands is entered, the changes are lost when the session is logged out.

11. Move ports into the FICON LS.

For more detailed instructions, see [Adding and Moving Ports on a Logical Switch](#).

The following command sequence shows a port migrated into the FICON logical switch.

#### NOTE

To enable ports migrated to the FICON logical switch successfully, you must bind the port using the `portaddress --bind` command.

```
switch:FID128:admin> lscfg --config 50 -slot 4 -port 5
This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Making this configuration change. Please wait...
Configuration change successful.
Please enable your ports/switch when you are ready to continue.

switch:FID128:admin> setcontext 50
switch_50:FID50:admin> switchshow
switchName: switch_50
.
.
.
Allow XISL Use: ON
LS Attributes: [FID: 50, Base Switch: No, Default Switch: No, Ficon Switch: Yes, Address Mode 1]

Index Slot Port Address Media Speed State Proto
=====
 69 4 5 010000 id N32 No_Sync FC Disabled (Persistent)
```

```

switch_50:FID12:admin> portaddress --bind 4/5 0xab00
switch_50:FID12:admin> portenable 4/5
switch_50:FID12:admin> switchshow
switchName: switch_12
switchType: 166.0
.
.
LS Attributes: [FID: 50, Base Switch: No, Default Switch: No, Ficon Switch: Yes, Address Mode 1]

Index Slot Port Address Media Speed State Proto
=====
101 11 5 01ab00 -- N32 No_Module FC

```

## 12. Enter the `portaddress --bind` command to bind a port address to each port.

You can specify a range of ports to bind to an area. When you specify a range of ports, the first port is bound to the first area and each successive port is bound to the successive area.

### NOTE

Ports in a FICON LS must have an address bound to them before they can be enabled. To enable the ports migrated to the FICON LS successfully, you must bind the port using the `portaddress -- bind` command.

The following examples show a range of ports used in the `portaddress --bind` command.

```

switch:FID128:admin> portaddress --bind 12/0-2 0xa000
switch:FID128:admin> portaddress --unbind 12/0-1

switch:FID128:admin> portaddress --bind 12/3-5 0xa000
portaddress: The operation failed for port 5.
The area 0xa200 is assigned to port 12/1.

```

## 13. Enter the `switchEnable` command to set the switch online.

## 14. Enter the `trackChangesSet` command to enable or disable the track-changes feature and to configure SNMP-TRAP mode.

## 15. Enter the `fabricPrincipal` command on the switch that you want to be the principal switch.

In cascaded fabrics, only one switch should be the principal. In core-edge designs, the director class switch should be the principal.

## 16. Enter the `portCfgSpeed` command to configure port speeds.

By default, all ports are configured to auto-negotiate. Normally, the only time the port configuration is changed is when connecting to 1G ports. Some dense wavelength division multiplexing (DWDM) ports are known not to auto-negotiate well.

### NOTE

- 8G ports auto-negotiate to 4G or 2G. 8G ports do not support or auto-negotiate to 1G.
- 16G ports auto-negotiate to 8G or 4G. 16G ports do not support or auto-negotiate to 2G or 1G.
- 32G ports auto-negotiate to 16G or 8G. 32G ports do not support or auto-negotiate to 4G or less.
- 64G ports auto-negotiate to 32G or 16G. 64G ports do not support or auto-negotiate to 8G or less.

17. Enter the `portCfgLongDistance` command for ports with fiber connections exceeding 10 km.

For ports with fiber connections exceeding 10 km, configuring the port for long-distance mode increases the number of buffer-to-buffer (BB) credits available on that port. If you have any Extended Fabrics links, enable VC translation link initialization (`portCfgLongDistance --vc_translation_link_init`) to stabilize them. See [ISL Trunking over Long-Distance Fabrics](#) for details on this option of the `portCfgLongDistance` command.

18. Enter the `defZone --noaccess` command to disable the default zone.

19. Configure zoning for your fabric.

Follow the standard FICON zoning practices that are published by Brocade and IBM. For a list of the standard FICON zoning documents, see [Zoning for FICON](#) or refer to the user guide for your management tool.

20. Enter the `statsClear` command to clear port statistics.

21. Enter the `switchShow` command to verify that the switch and devices are online.

22. Enter the `ficonShow rnid` command to verify that the FICON devices are registered with the switch. The `ficonShow rnid port` command displays RNID information for a port.

23. Enter the `ficonShow lirr` command to verify that the FICON host channels are registered to listen for link incidents.

24. Enter the `configUpload` command to save your configuration.

#### NOTE

This command backs up the current Fabric OS feature and switch configuration, including any FMS-specific or FICON-specific configuration such as an IPL file, active configuration, mode register, and host data.

To capture all the information necessary for FICON, you must enter the `configUpload` command multiple times with the following command line options:

- `configupload -vf`
- `configupload -all`
- `configupload -map`

25. Enter the command appropriate for your environment:

- If you do not have access to an FTP server, use a Telnet session that can log console output and enter the `supportShow` command to save a FICON baseline.
- If you have access to an FTP server, enter the `supportSave` command to save the FICON baseline of your switch.

26. Take the appropriate action based on your configuration:

- If you have a cascaded configuration, connect and enable channel and control unit (CU) devices. The Query Security Attributes (QSA) response to the channel indicates that fabric binding and the IDID are enabled.
- If you have and plan to use a CUP license, continue configuring your switches using the instructions in [Configuring FICON CUP](#).
- If you have a switched point-to-point configuration, you have completed the switch configuration.

## Converting an Existing Logical Switch to a FICON Logical Switch

Converting an existing logical switch to a FICON logical switch (`lscfg --change -ficon` command) is deprecated in Fabric OS 9.0.0. For Gen 6 platforms, you can continue to use the `lscfg --change -ficon` command to convert a logical switch to a FICON logical switch. When converting an existing logical switch to FICON, you must configure an

insistent domain ID, an SCC policy, and HIF mode. However, for Gen 7 platforms and subsequent releases of Fabric OS, use the `lscfg --create -ficon` command to create a new logical switch.

With Gen 6 platforms, if the existing logical switch has ports assigned to it, the port-to-address binding applies, and only ports that are bound to an address can be enabled. All online ports are restricted to an address. With Gen 6 switches, changing a standard logical switch to a FICON logical switch is disruptive if the switch is under the following conditions:

- Is already in HIF mode.
- Contains insistent domain ID (IDID).
- Contains zero-based addressing.
- Contains DBR routing policy.
- Is in the 256-Area limited addressing mode.
- Has all online ports user-bound.

With Gen 6 chassis systems, changing to a FICON logical switch is non-disruptive on a chassis if the same conditions are met.

#### NOTE

The `lscfg --change -ficon` command fails when it is issued on a switch that is already a FICON mode logical switch. When a logical switch has already been configured to meet the criteria for a FICON switch is changed to be a FICON logical switch, there is no need for an outage. If the logical switch is not configured to meet the criteria for a FICON mode logical switch, the switch is disabled, the changes are applied, and you must re-enable the switch.

To change an existing logical switch to a FICON logical switch, complete the following steps.

1. Connect to the physical chassis, and log on using an account with the chassis-role permission.
2. Perform one of the following actions depending on the generation of your switch:
  - For Gen 6 platforms, use the `lscfg --change -ficon` command to convert an existing logical switch to a FICON logical switch.

```
switch:FID128:admin> lscfg --change 5 -ficon
2016/08/23-15:01:53, [PMGR-1013], 35, SLOT 2 CHASSIS, INFO, Alletor7, Attempt to change switch 5 to
FICON switch succeeded.
Please enable your switches when ready.
```

When converting an existing logical switch to FICON, you may need to first configure an insistent domain ID, an SCC policy, a fabric-wide consistency policy, and HIF mode, as described in the following sections.

- [Enabling Insistent Domain ID](#)
- [Creating and Activating the SCC Policy](#)
- [Enabling the Fabric-Wide Consistency Policy](#)
- [Enabling High-Integrity Fabric Mode](#)
- For Gen 7 platforms and subsequent releases of Fabric OS, use the `lscfg --create -ficon` command to create a new FICON logical switch (see [Configuring FICON](#) for step-by-step instructions).



## Enabling the Insistent Domain ID

To enable the insistent domain ID, complete the following steps for each switch in the fabric.

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Enter the `configure` command, and step through the interactive prompts.
  - a) At the `Fabric parameters` prompt, type `y`.
  - b) At the `Insistent Domain ID Mode` prompt, type `y`.

## Creating and Activating the SCC Policy

Creating a switch connection control (SCC) policy defines the switches allowed in the fabric.

To configure and activate an SCC policy, complete the following steps.

1. Connect to the switch and log on.
2. Perform one of the following actions:
  - Enter the `secpolicycreate` command to add all switches in the fabric, if they are connected.

```
secpolicycreate "SCC_POLICY", ""
```

- Enter the `secpolicyadd` command to add one or more members to an existing policy. The following command is an example of adding a member using device WWNs.

```
secpolicyadd "SCC_POLICY", "wn1;wn2"
```

3. Enter the `secpolicyactivate` command to activate the currently defined SCC policy.

This activates the policy set on the local switch or on all switches in the fabric, depending on the configured fabric-wide consistency policy.

## Enabling the Fabric-Wide Consistency Policy

Enable the fabric-wide consistency policy after all switches have joined the merged fabric. If there are fabric-wide data distribution (FDD) conflicts on any of the ISLs, disable the fabric-wide consistency policy on each switch in the fabric.

Once the fabric has merged successfully (use `fabricShow` to verify), enter the following command.

```
fddcfg --fabwideset "SCC:S"
```

When enabling the fabric-wide consistency policy, SCC:S enforces strict mode, which is required for FICON.

## Enabling High-Integrity Fabric Mode

Setting High-Integrity Fabric (HIF) mode on a switch verifies that the switch meets high-integrity fabric requirements through the channel's Extended Link Service Query Security Attributes (ELS QSA) function.

Setting HIF mode locks the IDID, fabric-wide consistency policy, and SCC policy settings to ensure that the fabric is of high integrity so that it can connect with the FICON channel. You cannot change these settings without disabling HIF mode.

### NOTE

HIF mode must be enabled to enable FMS mode.

To enable HIF mode, complete the following steps.

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Enter the `configure` command and step through the interactive prompts.
  - a) At the `Fabric parameters` prompt, type `y`.
  - b) At the `High Integrity Fabric Mode` prompt, type `y`.

If HIF configuration requirements have not been met, an error message describes what you must configure for the command to succeed. For example, the following message states that an IDID, SCC policy, or fabric-wide consistency policy has not been configured for the switch. Perform additional configuration if required, and then enable HIF mode.

```
Error: Unable to set HIF Mode. No valid IDID settings,
SCC policy and/or Fabric wide(SCC:S) configuration
```

## FCR and FICON Cascading

FICON cascading is allowed in a Fibre Channel Routing (FCR) configuration when front or translate domains are present. These domains can co-exist with a fabric-wide "strict" SCC policy.

To connect edge fabrics to high-integrity backbone fabrics that use FCR, use the following these general steps:

1. Define the FC router `EX_Ports`.
2. Enter the `portcfgexport slot/port` command to obtain the front domain WWN.
3. Add the front domain WWN to the SCC policy lists in the edge switches to be merged.
4. Connect the switches.

## FICON and FICON CUP

For FICON CUP, FICON Management Server (FMS) mode is enabled on the logical switch and not the chassis. For example, on a Brocade switch or a Brocade DCX director, enable CUP on each logical switch.

You can configure up to eight logical switches as CUP switches with FMS mode. Complete instructions for creating a FICON logical switch (LS) using Fabric OS commands are in the *Brocade Fabric OS Administration Guide*. Instructions for creating a FICON LS using Brocade SANnav Management Portal are in the *Brocade SANnav Management Portal User Manual*. Once a FICON LS is configured, CUP must be installed on the switch per the instructions in [Configuring FICON CUP](#).

The following table specifies the maximum number of logical switches and maximum number of FICON LSs with FMS mode enabled for supported Brocade FICON switches and directors.

**Table 146: Maximum Number of Logical Switches for FICON Switches and Directors**

| Switch or Director | Maximum Number of Logical Switches | Maximum Number of Logical Switches with FMS (CUP) Enabled |
|--------------------|------------------------------------|-----------------------------------------------------------|
| G720               | 4                                  | 2                                                         |
| X7-4               | 16                                 | 8                                                         |
| X7-8               | 16                                 | 8                                                         |
| G620               | 4                                  | 2                                                         |
| X6-4               | 16                                 | 8                                                         |
| X6-8               | 16                                 | 8                                                         |

When creating a FICON LS, consider the following:

- FICON is not supported in the default logical switch.
- You can use the Fabric OS `configure` command to allow XISL use.

An extended inter-switch link (XISL) is an inter-switch link (ISL) that connects the logical switch to a base switch and carries traffic for multiple logical fabrics. The XISL is a logical switch feature, not a link feature. It is shared by multiple logical switches and can be used for multiple fabric IDs (FIDs). Creating an XISL is disruptive because it provides a core routing function for the switch. For more information on creating XISLs and restrictions for their use, refer to the `configure` command in the *Brocade Fabric OS Command Reference Manual*.

When using XISLs, the effective routing policy for all logical switches configured to allow XISL use is that of the base switch regardless of how the routing policy is defined. The best practice is to set the routing policy in all logical switches that allow XISL use to have the same routing policy as the base switch to avoid confusion about the routing policy used. A base switch can be used for XISLs only. Therefore, the switch cannot be managed through a CUP port and cannot report statistics for RMF.

- The following routing policies are qualified for IBM z Systems:
  - Port-based routing (PBR)
  - Device-based routing (DBR)
  - Exchange-based routing (EBR)

There are caveats with EBR as to the specific processor type, z/OS level, and storage device support. Refer to the z Systems qualification letter and storage vendor support matrixes for current EBR support details. The recommended policy is DBR or `aptPolicy 2`.

- Note that port addresses 0xFE or 0xFF cannot be assigned to physical ports in a FICON LS.
- Each logical switch has its own configuration data, such as the following:
  - Initial Program Load (IPL) file.
  - Configuration data elements, such as Port Information Block (PIB), port address names, mode register, and key counter.
  - CUP configuration files.
- Ports moved into a logical switch will have persistent states based on the Active=Saved (ASM) bit setting and other factors. Refer to [Disabling and enabling ports - persistent states](#) for more information.

## Blocking and Unblocking Ports

You cannot issue the `portCfgPersistentDisable` command if FMS mode is enabled. When FMS mode is enabled, persistently enabled or disabled configuration states are saved in the CUP's IPL file. When FMS mode is enabled and Active=Saved Mode (ASM) is ON, the user can achieve persistent port configuration states by using the `portdisable` and `portenable` command. IBM CUP documentation refers to these persistent configuration states as blocked and unblocked.

The persistent configuration state depends on the following factors:

- If ASM is set to one (1) and you change the port's blocked state to block/disable or unblock/enable in the CUP Active Configuration, then that setting *will* update the IPL file and *will* be in a persistent configuration state.
- If ASM is set to zero (0) and you change the port's blocked state to block/disable or unblock/enable in the CUP Active Configuration, then that setting *will not* update the IPL file and *will not* be in a persistent configuration state.

The following tables compare the achievement of port persistent enable and disable states with interactions between the ASM mode setting, IPL state, and use of the `portdisable` or `portenable` command.

**Table 147: ASM 1: ON**

| Initial IPL state                                             | Blocked                         | Blocked                          | Unblocked                       | Unblocked                        |
|---------------------------------------------------------------|---------------------------------|----------------------------------|---------------------------------|----------------------------------|
| User action                                                   | <code>portenable</code> command | <code>portdisable</code> command | <code>portenable</code> command | <code>portdisable</code> command |
| Physical state of port                                        | Enabled                         | Disabled                         | Enabled                         | Disabled                         |
| Port state in IPL file                                        | Unblocked                       | Blocked                          | Unblocked                       | Blocked                          |
| <code>portcfgshow</code> command display for persistent state | OFF                             | ON                               | OFF                             | ON                               |

**Table 148: ASM 0: OFF**

| Initial IPL state                                             | Blocked                         | Blocked                          | Unblocked                       | Unblocked                        |
|---------------------------------------------------------------|---------------------------------|----------------------------------|---------------------------------|----------------------------------|
| User action                                                   | <code>portenable</code> command | <code>portdisable</code> command | <code>portenable</code> command | <code>portdisable</code> command |
| Physical state of port                                        | Enabled                         | Disabled                         | Enabled                         | Disabled                         |
| Port state in IPL file                                        | Blocked                         | Blocked                          | Unblocked                       | Unblocked                        |
| <code>portcfgshow</code> command display for persistent state | ON                              | ON                               | OFF                             | OFF                              |

Following are summaries of important points in the tables:

- The only time "Persistent Disable ON" should display with the `portcfgshow` command is when you disable a port and ASM is set to ON.
- The only time the IPL file changes is when ASM is set to ON.
- When FMS mode is enabled, persistence is controlled by the interaction of ASM mode and the IPL file.

## Clearing the FICON Management Database

Perform the following steps to clear Registered Link Incident Report (RLIR) and Request Node Identification Data (RNID) records from the FICON management database.

1. Connect to the switch, and log on as admin.
2. Enter `ficonClear rlir` to remove all RLIR records from the local RLIR database.
3. Enter `ficonClear rnid` to remove all RNID records marked "not current" from the local RNID database.

## Automating CS\_CTL Mapping

The CS\_CTL-based frame prioritization feature allows you to prioritize the frames between a host and control unit (CU) as having high, medium, or low priority, depending on the value of the CS\_CTL field in the FC frame header. High-, medium-, and low-priority frames are allocated to different virtual channels (VCs). High-priority frames receive more VCs than medium-priority frames, which receive more VCs than low-priority frames. The virtual channels are allocated according to the CS\_CTL value.

When QoS is enabled, buffers are reserved for the VCs. Most traffic occurs on the medium-priority channel, which means that a number of available buffers are not being used. The recommended action is to disable QoS on ports unless you have a well-defined plan for establishing priorities in your FICON environment. When QoS is enabled on a port, you risk

depleting buffers and experiencing performance degradation. Be aware that disabling or enabling QoS disrupts the link when it is taken offline and brought back online.

#### NOTE

Check with your host and storage manufacturer to determine if they support Fibre Channel CS\_CTL prioritization on their devices.

The *Brocade Fabric OS Administration Guide* provides additional information on this feature and steps to enable CS\_CTL-based frame prioritization on your switch using the `portCfgQos` command.

The `configureChassis` command allows you to set chassis-level configuration parameters, such as `cfgload`, `custom`, `system`, and `Fabric OS` attributes, to be used by ASIC drivers. The following procedure uses the `Fabric OS` attribute `CS_CTL QoS Mode` to change CS\_CTL values mapped to VCs for priority levels from `default` to `auto` mode. While `default` mode maps CS\_CTL values 1-8 (low), 9-16 (medium), and 17-24 (high) to specific VCs, `auto` mode maps CS\_CTL values 1 (low), 2 (medium), and 3 (high) to VCs based on the ASIC supported by your platform.

To enable `auto` mode for CS\_CTL-to-VC mapping on a FICON switch, use the following steps.

#### NOTE

This process is disruptive as you must reboot the chassis to reset the configuration in the ASIC. You cannot enable CS\_CTL mapping on individual ports until this is done. If `auto` mode is selected, you must select `auto` mode for each chassis in the fabric. Not doing so will mean different virtual channels (VCs) will be used across the fabric with undesirable results.

1. Connect to the switch, and log on to an account that has admin permissions.
2. Launch the `configureChassis` command.
3. Select `no` when prompted to modify `cfgload` attributes, `custom` attributes, and `system` attributes, and then select `yes` to modify `Fabric OS` attributes.

```
switch:admin> configurechassis
Configure...
 cfgload attributes (yes, y, no, n): [no]
 Custom attributes (yes, y, no, n): [no]
 system attributes (yes, y, no, n): [no]
 fos attributes (yes, y, no, n): [no] y
```

4. Select `1` (`auto` mode) when prompted to configure CS\_CTL QoS mode.

```
CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [0] 1
```

5. Reboot the switch or chassis.
6. To verify if `auto` mode or `default` mode is set for the chassis, enter the `configShow --all` command.

```
switch:admin> configshow -all | grep fos.csctlMode
fos.csctlMode:1
```

- If `auto` mode is set, the following should display in the command output:

```
fos.csctlMode:1
```

- If `default` mode is set, the following should display in the command output:

```
fos.csctlMode:0
```

To return to `default` mode, follow the preceding steps, but select `0` (`default` mode) for CS\_CTL QoS mode when modifying `fos` attributes through the `configureChassis` command.

**NOTE**

The `configureChassis` command does not provide options to enable the CS\_CTL QoS mode on individual ports in the chassis. Once you set the mode on the chassis, you must enable individual ports to use this mode with the `portCfgQos` command. For details on the `configureChassis` and `portCfgQos` commands, refer to the *Brocade Fabric OS Command Reference Manual*.

## Latency Guideline

The maximum supported distance for a FICON channel is 300 km (1.5 msec of delay). Synchronous remote data replication for direct-access storage devices (DASD) is generally limited to 100 km (0.5 msec of delay). Asynchronous remote data replication for DASDs and remote tape reads/writes requires that the Brocade Advanced Accelerator for FICON feature be used with Fibre Channel over IP (FCIP). The Advanced Accelerator for FICON feature emulates control unit response to the channel to improve the performance over distances greater than locally attached distances.

## FICON CUP Overview

The Control Unit Port (CUP) provides an in-band management interface, which is defined by IBM, that defines the channel command words (CCWs) that the FICON host can use for managing the switch. The protocol that is used is the IBM version of the ANSI FC-SB4 single-byte command code specification, which defines the protocol that is used to transport CCWs to the switch and for the switch to direct data and status back. The CUP is an optional licensed feature.


The CUP provides an interface for host in-band management and collection of FICON switch performance metrics using the Resource Measurement Facility (RMF) 74-7 record, more commonly known as the FICON Director Activity Report. Host-based management programs manage the FICON switches by sending commands to the switch control unit defined in the I/O Configuration Data Set (IOCDS) and hardware configuration definition (HCD). A FICON switch that supports CUP can be controlled by one or more host-based management programs or switch consoles. Control of the FICON switches can be shared between these options. CUP commands, or CCWs, monitor and control FICON switch functions. CUP commands are oriented toward management of a single switch, even though the use of CUP in a cascaded FICON environment is fully supported.

**NOTE**

The CUP port address is always xxFExx. Port addresses xxFE and xxFF are unavailable to assign to physical ports in a FICON LS.

FICON Management Server mode (FMS mode) must be enabled on the switch to enable CUP management features. For more information, see [FMS Mode and FICON CUP](#).

**Figure 98: FICON CUP**



This figure is a simplified representation of a FICON environment and how CUP fits into that environment. FICON switches with FMS enabled use port address "FE" as the embedded internal port for host connectivity to the CUP, and port address "FF" is reserved. Therefore, addresses "FE" and "FF" cannot be assigned to physical ports in a FICON LS.

In a FICON environment, only one RMF logical partition (LPAR) should be configured to poll the CUP for port statistical counter data to minimize contention for the CUP device. Host contention for access to the CUP can lead to missing interrupts, potentially causing the host to box the device. It is a best practice to have more than one CHPID defined as a path to the CUP to ensure redundancy for host-to-CUP connectivity.

## FICON CUP Restrictions

The following restrictions apply to Brocade FICON switches when FMS mode is enabled and when host applications are used for in-band monitoring and management of the switch:

- FICON Management Server mode (FMS mode) can be enabled only in a FICON logical switch (LS).
- FMS mode must be enabled in each logical switch to enable CUP management features. When this mode is enabled, Fabric OS prevents configuration requests from out-of-band management interfaces from interfering with host in-band configuration requests by serializing access to switch parameters.
- The switch is advertised to the mainframe by the CUP as having a specific number of ports, which is based on the hardware platform.
- The port address name for the internal FE port can be read by the host, and it corresponds to the switch name that is set for the switch. A switch name in Fabric OS must begin with an ASCII alphanumeric character and must contain only alphanumeric characters, an underscore, a dash, or a period. The switch name can be modified by the host and is subject to the following restrictions:
  - If the host-assigned FMS name is all numeric, the first character is converted to an "x" and the rest of the numbers remain. For example, the name "12345" appears as "x2345" in Fabric OS.
  - If the first character of an FMS name is not ASCII alphanumeric, that character is converted to an "x". For example, the name "\_switch\_1" appears as "xswitch\_1" in Fabric OS.
  - Spaces in the middle of an FMS name are converted to underscores. Trailing spaces are removed. For example, the name " switch name 1 " appears as "switch\_name\_1" in Fabric OS.
  - If the host-assigned FMS name is all spaces, the Fabric OS value is set to the format "FID*fabid*\_\_\_\_\_", where *fabid* is the fabric ID of the logical switch and the spaces are converted to underscore characters. For example, an FMS name of 10 spaces on fabric ID 32 appears as "FID32\_\_\_\_\_" in Fabric OS.
  - Any character other than ASCII alphanumeric, underscore, dash, period, and space in the FMS name is converted to a period. For example, the FMS switch name "switch(1)" appears as "switch.1." in Fabric OS.

## CUP Configuration Recommendations

The following are recommendations for setting up the FICON CUP:

- Device Type: 2032  
Use device type 2032 as the CUP "UNIT" type for the CNTLUNIT and IODEVICE statements in the IOCDS or HCD. Although IBM has developed new machine types for switches, all Brocade FICON directors and switches should be configured as 2032.
- Paths  
Define at least two CHPIDs as paths to the CUP in the IOCDS Control Unit statement for the switch to ensure redundancy for host connectivity.
- RMF 74-7 statistics gathering  
So as not to overwhelm the CUP, the recommended best practice in a sysplex environment is to disable FICON switch statistics (RMF polling) on all but one host system that has in-band access to the CUP. If no FICON statement is specified, the system will gather switch statistics by default, so FICON switch statistics should be disabled on all but the system intended to gather statistics.
- Missing Interrupt Handler (MIH)



Set the MIHPTO to three minutes (180 seconds).

**NOTE**

Although 180 seconds should be the default value set on the switch when FMS mode is enabled, this value may not be set when older versions of Fabric OS are upgraded. This value must be set on the host system as well.

## Zoning and PDCM Considerations

**NOTE**

Fabric OS 9.0.0 was the last release to support the PDCM used for connectivity management as reflected in the Allow/Prohibit Configuration Matrix option in WebTools. Effectively, the Allow/Prohibit matrix management feature was deprecated in Fabric OS 9.1.0. For more information on managing connectivity for FICON, see [Zoning for FICON](#).

The FICON PDCMs control whether communication between a pair of ports in the switch is prohibited or allowed. PDCMs are a per-port construct. Each port on the switch has its own PDCM that defines whether communication is allowed between that particular port and each of the other ports in the switch, including itself. The Allow/Prohibit Matrix presents a matrix that allows you to set and clear the PDCMs for all external ports. If there are any differences between restrictions set up with Brocade Advanced Zoning and PDCMs, the most restrictive rules are automatically applied.

As stated, PDCMs affect communication between ports on the same switch only. In a cascaded environment, you need to use zoning to control access between ports on different switches. FMS mode must be enabled for the Allow/Prohibit Matrix to be operational, and there must be an active zoning configuration. Domain, Index (D,I) zoning is recommended for FICON, but is not specifically required. However, hard-port or hard-WWN zoning enforcement is required. PDCMs will not be enforced on ports with session-based enforcement.

If Brocade Advanced Zoning is in use, for more information on managing connectivity for FICON, see [Zoning for FICON](#).

## Port and Switch Naming Standards for FMS Mode

The FICON Management Server (FMS) and the FICON programming interface employ 8-bit EBCDIC characters in port address names and switch names. Fabric OS employs 7-bit ASCII characters. Therefore, when FMS mode is enabled, all characters greater than 0x40 and not equal to 0xFF are allowed in port address names. However, a switch name in Fabric OS must begin with an ASCII alphanumeric character and must contain only alphanumeric characters, an underscore, a dash, or a period. See [FICON CUP Restrictions](#) for additional information about Fabric OS switch and port name handling.

A host can set a name with nonprintable characters or characters other than the set allowed by Fabric OS. If a name contains nonprintable characters, those characters are displayed as dots (...). The following characters are also displayed as dots: semicolon (;), comma (,), equal sign (=), and at sign (@). See [EBCDIC Code Page](#).

The port name for the CUP must be unique within the chassis.

**IMPORTANT**

Change to FMS mode first and then configure port names. If you configure port names **before** changing to FMS mode, you will not see the port names when you do enable FMS mode.

## FICON CUP Fabric OS Commands

The following table summarizes the Fabric OS CLI commands related to managing FICON CUP. For detailed information on these commands, refer to the *Brocade Fabric OS Command Reference Manual*.



**Table 149: Fabric OS Commands Related to FICON CUP**

| Command                                                                           | Description                                                                    |
|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| <code>ficonCupSet fmsmode</code>                                                  | Sets FICON Management Server mode on or off for the switch.                    |
| <code>ficonCupSet crp</code>                                                      | Sets the current reporting path (CRP).                                         |
| <code>ficonCupSet MIHPTO</code>                                                   | Sets the Missing Interrupt Handler Primary Timeout value.                      |
| <code>ficonCupSet modereg</code>                                                  | Sets the mode register bits for the switch.                                    |
| <code>ficonCupSet UALERT_ALL   UALERT_HSC   UALERT_FRU<br/>  UALERT_INVATT</code> | Sets which alert types are enabled for display for the switch.                 |
| <code>ficonCupShow lp</code>                                                      | Displays the host logical path information for the CUP.                        |
| <code>ficonCupShow fmsmode</code>                                                 | Displays the FICON Management Server enabled or disabled state for the switch. |
| <code>ficonCupShow modereg</code>                                                 | Displays the mode register bit settings for the switch.                        |
| <code>ficonCupShow UALERT_MODE</code>                                             | Displays the UALERT mode settings for the switch.                              |

## Configuring FICON CUP

To set up FICON CUP, use the following procedure.

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Enter the `switchShow` command to verify that the logical switch attributes show that the logical switch is a FICON logical switch.

### NOTE

Converting an existing logical switch to a FICON logical switch (`lscfg --change -ficon` command) is deprecated in Fabric OS 9.0.0. For Gen 6 platforms, you can continue to use the `lscfg --change -ficon` command to convert a logical switch to a FICON logical switch. However, for Gen 7 platforms and subsequent releases of Fabric OS, use the `lscfg --create -ficon` command to create a new FICON logical switch.

3. Verify that the switch or director has been set up according to the instructions in [Configuring FICON](#).
4. Enter the `ficonCupShow fmsmode` command to verify if FMS mode is disabled.
5. If FMS mode is enabled, disable it using the `ficonCupSet fmsmode disable` command.

### ATTENTION

Disabling FMS mode may be disruptive to operations because ports that were previously prohibited from communicating will be allowed to do so because prohibitions are no longer enforced. However, you must disable FMS mode to continue setting up CUP.

6. Install a CUP license on the switch. For more information on installing a license, refer to this administration guide.
7. Enable FMS mode using the `ficonCupSet fmsmode enable` command.
8. Enter the `ficonCupSet modereg` command to set the mode register bits.  
See [Mode Register Bit Settings](#) for more information.
9. Verify with the systems programmer that the CUP ports and the switches and directors are defined properly in their IOCDs.

FICON ports are defined as part of the sysgen process. The z/OS systems programmer can verify if the CUP ports are defined properly.

## Determining Physical Port Assignment

The following steps help you identify the physical port from `CUP_PortAddress`.

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Set the context to the appropriate logical switch.
3. Enter the `switchShow` command.
4. Look for the PID in the `Address` column.
5. Read across to the **Slot** and **Port** columns to find the physical port number.

For more information and example output from the `switchShow` command, refer to the *Brocade Fabric OS Command Reference Manual*.

## FMS Mode and FICON CUP

Enabling FICON Management Server (FMS) mode in a Fabric OS environment allows the switch to behave within the context of the FICON architecture and creates a FICON director. The FMS mode setting can be changed whether the switch is disabled or enabled. If FMS mode is changed while the switch is enabled, a device reset is performed for the control device and an RSCN is generated with PID 0xDDFE00, where 0xDD is the domain ID of the switch.

### IMPORTANT

Change to FMS mode first and then configure port names. If you configure port names **before** changing to FMS mode, you do not see the port names when you do enable FMS mode.

The following requirements must be configured on the switch to enable FMS mode:

- For Fabric OS 7.0 and later, a FICON CUP license.
- For Fabric OS 9.0.0 and later, FMS mode can be enabled only in a FICON logical switch.

Changing FMS mode from disabled to enabled triggers the following events:

- Access to switch configuration parameters is serialized.
- The active CUP configuration data, including the following, is initialized according to the IPL file:
  - Port and switch names
  - Port "Block" and "Unblock" values
- Brocade Advanced Zoning, if used, remains enforced.
- An RSCN for the CUP is generated to all mainframe FICON and FCP channels in the fabric.

Changing FMS mode from enabled to disabled triggers the following events:

- A device reset is performed on the control device.
- If a given port was set to "Block" or "Unblock," that port remains disabled or enabled.
- The switch interface for host in-band management is no longer enabled.
- An RSCN for the CUP is generated to all mainframe FICON and FCP channels in the fabric.

## Fabric OS Command Limitations and Considerations

This section describes limitations and considerations for using Fabric OS commands with FMS mode enabled. For a list of Fabric OS commands related to FICON, refer to the *Brocade Fabric OS Command Reference Manual*.

- The following Fabric OS commands return a "switch busy" response if they are issued when the host-based management application is performing a write operation. This prevents interference from out-of-band management applications when a host-based management program is being used to administer the switch.

- bladeEnable
  - portDisable
  - portEnable
  - portName
  - portShow
  - switchCfgPersistentDisable
  - switchDisable
  - switchEnable
  - switchName
  - switchShow
- You cannot use the `portCfgPersistentEnable` and `portCfgPersistentDisable` commands to persistently enable and disable ports. See [Persistently Enabling and Disabling Ports for CUP](#) for instructions.
  - Specifying a range of ports for the `portName` command will fail except for the first port in the range. FMS mode does not permit multiple ports with the same name.

## Displaying FMS Mode

Use the `ficonCupShow fmsmode` command to display the enabled or disabled state of FMS mode for the switch.

## Enabling FMS Mode

FMS mode can be enabled only in a FICON logical switch (LS). During FICON LS creation, all security attributes are enforced to ensure a high-integrity fabric.

### IMPORTANT

Change to FMS mode first and then configure port names. If you configure port names **before** changing to FMS mode, you will not see the port names when you do enable FMS mode.

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Enter `ficonCupSet fmsmode enable`.
3. Enter the `ficonCupShow fmsmode` command to verify that FMS mode is enabled.

## Disabling FMS Mode

To disable FMS mode, perform the following steps.

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Enter `ficonCupSet fmsmode disable`.
3. Enter the `ficonCupShow fmsmode` command to verify that FMS mode is disabled.

## Upgrade Considerations

Take the CUP offline before performing a firmware upgrade or downgrade operation. This means that you should vary the CUP offline at the host. The upgrade can disrupt host communication with the CUP, especially for active RMF sessions.

Consider the following when using FMS mode and upgrading from Fabric OS 8.X:

- From Fabric OS 9.0.0, FMS mode can be enabled only in a FICON logical switch. If FMS mode is enabled in a non-FICON logical switch, the upgrade will fail. See [Administering FICON Fabrics](#) for configuration information.
- In the following conditions, firmware upgrade to Fabric OS 9.1.0 or later is blocked:

- Prohibited connections are detected in the Allow/Prohibit matrix configuration on the FICON logical switch.
- Port addresses greater than the maximum physical port range are bound to a port in any FICON logical switch.

## Mode Register Bit Settings

A mode register controls the behavior of the switch in relation to CUP parameters and in relation to the behavior of other management interfaces.

The FICON CUP mode register bits are described in the following table:

**Table 150: FICON CUP Mode Register Bits**

|      |                                                                                                                                                                                                                                                                                                                                                                                             |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ACP  | Alternate control prohibited (ACP). Because the Fabric OS CLI, Web Tools, and Brocade SANnav Management Portal are considered to be switch consoles, this bit has no effect on their operations. Attempts to set CUP parameters through SNMP are denied when this bit is set to on. The default setting is 1 (on).                                                                          |
| ASM  | Active=Saved mode (ASM). When this bit is set to on, all CUP configuration parameters are persistent. The parameters are saved in nonvolatile storage in the IPL file that is applied upon a cold reboot or a power cycle. The default setting is 1 (on).                                                                                                                                   |
| DCAM | Director clock alert mode (DCAM). When this bit is set to on, a warning is issued when the <code>date</code> , <code>tsClockServer</code> , or <code>tsTimeZone</code> commands are entered to set the time and date on the switch. The default setting is 0 (off).<br>The director clock alert mode function is obsolete and has been replaced by the <a href="#">NTP Server</a> function. |
| HCP  | Host control prohibited (HCP). When this bit is set to on, the host is not allowed to set CUP parameters. The default setting is 0 (off).                                                                                                                                                                                                                                                   |
| POSC | Programmed offline state control (POSC). If this bit is set (=1), the command from the host sets the switch offline. If this bit is off (=0), the command from the host is rejected and the switch does not go offline. The default setting is 1 (on).                                                                                                                                      |
| UAM  | User alert mode (UAM). When this bit is set to on, a warning is issued when an action is attempted that will write CUP parameters on the switch; for example, changing the block/unblock state of a port or the port name by means of the out-of-band interfaces. This bit can only be modified in-band by the host. The default setting is 0 (off).                                        |

When the POSC bit is set to 1, the FICON host is allowed to set the switch offline. The bit is set to 1 by default. If you do not want to allow the host to have permission to set the switch offline, the bit should be turned off. The state of this bit cannot be changed by the FICON host; it can be changed only through the command line.

The `ficonCupShow modereg` command displays the mode register bit settings for the switch. A display of 0 indicates that the mode register bit is set to off; and 1 indicates that the bit is set to on. To display the mode register bit HCP for the switch, enter the command as shown in the following example:

```
switch:admin> ficoncupshow modereg HCP
```

For more information on this command and an example of command output, refer to the *Brocade Fabric OS Command Reference Manual*.

## Considerations for Setting the Mode Register Bits

Consider the following when changing mode register bits:

- The UAM bit can be set only by host programming.
- All mode register bits except UAM are saved across power cycles; the UAM bit is reset to 0 following a power-on.
- Mode register bits can be changed when the switch is offline or online. If the ACP or HCP bits are changed when the switch is online, they will take effect any time between the completion of the current command and the end of the CCW command chain (or the next alternate manager operation).

## Setting the Mode Register Bits

Perform the following steps to set mode register bits:

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Enter the `ficonCupSet modereg` command to set the FICON CUP mode register bits for the local switch.

Example of setting the HCP mode register bit to off:

```
switch:admin> ficoncupset modereg HCP 0
```

Example of setting the ACP mode register bit to on:

```
switch:admin> ficoncupset modereg ACP 1
```

## Setting the MIHPTO Value

The missing interrupt handler primary timeout (MIHPTO) value determines how long the channel will wait before timing out an expected response from the CUP. This value is part of the data returned to the host when it reads the configuration data, but the host is under no obligation to adhere to this value. The actual value that the host will use before detecting a timeout can be set by the system programmer (overriding the value returned by the switch).

Use the `ficonCupSet MIHPTO` command to increase, decrease, or display the MIHPTO value that the CUP returns to the host when configuration data is read. Changing the MIHPTO value requires that the CUP control device be varied offline from the host, as the host will not automatically be notified when the value is changed. The host will normally read the new value using the Read Configuration Data command when the CUP is varied back online.

The MIHPTO setting will persist across switch reboots, power-on resets (PORs), and failovers. Setting this value to the upper end of the time range will allow the CUP more time to process requests during intervals of peak usage, without causing the host to detect timeouts. It is recommended that the value be set to 180 seconds (default).

Complete the following steps to set the MIHPTO value.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `ficonCupSet MIHPTO` command and set the value from 15 through 600.

## Persistently Enabling and Disabling Ports for CUP

When FMS mode is enabled, you cannot use the `portCfgPersistentEnable` and `portCfgPersistentDisable` commands to persistently enable and disable ports. Instead, use the following procedure.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the `ficonCupShow modereg` command to display the mode register bit settings.
3. Verify that the ASM bit is set to on (1).
4. If the ASM bit is set to off (0), enter the `ficonCupSet modereg asm 1` command to set the bit to on.

For more information on this command and an example of command output, refer to the *Brocade Fabric OS Command Reference Manual*.

5. Use the `portEnable` and `portDisable` commands to enable and disable ports as necessary.

The ports remain enabled or disabled after a switch reboot.

```
switch:admin> portenable 1/1
```

See [Blocking and Unblocking Ports](#) for more information on enabling and disabling ports and persistent states.

## Administering FICON Extension Services

This section enables you to configure FICON extension services.

### Platforms Supporting FICON Extension over IP

In Fabric OS 9.0.x, FICON extension over IP is supported on the following platforms:

- Brocade X7-4 with SX6 blade  
The SX6 blade has 16 FC ports, 2 40GbE ports, and 16 1/10 GbE ports.
- Brocade X7-8 with SX6 blade
- Brocade X6-4 with SX6 blade
- Brocade X6-8 with SX6 blade


### FICON Emulation Overview

FICON emulation supports FICON traffic over IP WANs using FCIP as the underlying protocol. FICON emulation can be extended to support performance enhancements for specific applications using the following licensed features:

- IBM z/OS Global Mirror emulation (formerly eXtended Remote Copy or XRC)
- FICON tape emulation (tape read and write pipelining)

#### IBM z/OS Global Mirror Emulation

The IBM z/OS Global Mirror (formerly eXtended Remote Copy or XRC) application is an IBM zSystems mainframe application that implements host-based asynchronous disk mirroring, as supported by the disk hardware architecture and a host software component called the System Data Mover (SDM). Within this application, a specific application channel program invokes a well-constrained type of channel program that is called a Read Record Set (RRS) channel program. This channel program is used to read updates to disk volumes (record sets) only after they have been brought into synchronization, with the aim of writing record updates to a mirrored volume only. The RRS channel program accesses primary volumes from a remote host for the purpose of reading these record sets (updates) supported by IBM z/OS Global Mirror emulation. The Brocade emulation feature thus allows the primary volume to be located at a distance from its mirrored secondary volume without the performance degradation associated with IU pacing.

**Figure 99: IBM z/OS Global Mirror Emulation**

This figure shows how the primary volume and the secondary mirrored volume may be geographically distant across an IP WAN. Updates to the primary disk volumes are completed by the production applications and then staged in a control unit cache to be read by SDM. SDM then writes the updates to the secondary volumes. The latency that is introduced by greater distance creates delays in anticipated responses to certain commands. The FICON IU pacing mechanism may interpret delays as an indication of a large data transfer that could monopolize a shared resource and may react by throttling the I/O. IBM z/OS Global Mirror emulation provides local responses to remote hosts, eliminating distance-related delays. You can use the SX6 blade with FICON emulation.

For information on configuring IBM z/OS Global Mirror emulation, see [Configuring FICON Emulation](#). For information on displaying IBM z/OS Global Mirror emulation status and statistics, see [Displaying FICON Emulation Performance Statistics](#).

For z/OS Global Mirror implementations with cascaded FICON, it is recommended that device-based routing (DBR) be used for the routing policy on the ISLs instead of port-based routing (PBR).

FICON emulation features operate in DBR or PBR mode. FICON emulation can tolerate exchange-based routing (EBR), and WAN consumption is not affected for FICON flows with EBR enabled.

Exchange-based routing, which IBM calls FICON Dynamic Routing (FIDR), is qualified only on the IBM z13-class (and higher) platforms. If you are running IBM z13 fabrics with FDIR, you must isolate those fabrics from the earlier IBM platforms in your fabric. The earlier IBM platforms use only port-based or device-based routing.

It is recommended that the z13 be connected to a Brocade Gen 7 or Gen 6 switch by using FICON Express 16S channels and that the disk storage utilize 16G control unit adapters. IBM Health Checker for z/OS is designed to detect and advise the user if any down-level hardware or drivers are operating in the system. Using logical switches is one approach to provide this isolation.



## Tape Emulation


Tape emulation (also called tape pipelining) refers to the concept of maintaining a series of I/O operations across a host-WAN-device environment and should not be confused with the normal FICON streaming of CCWs and data in a single command chain. Normally, tape access methods can be expected to read data sequentially until they reach the end-of-file delimiters (tape marks) or to write data sequentially until either the data set is closed or an end-of-tape condition occurs (multivolume file). The emulation design strategy attempts to optimize performance for sequential reads and writes, while accommodating any other nonconforming conditions in a lower-performance nonemulating frame shuttle. Because write operations can be expected to comprise the larger percentage of I/O operations for tape devices (for archival purposes), they are addressed first.

For information on configuring tape emulation, see [Configuring FICON Emulation](#). For information on displaying tape emulation status and statistics, see [Displaying FICON emulation Performance Statistics](#).

## Tape Write Pipelining

FICON Tape Write Pipelining improves performance for various applications when writing to tape over extended distances. FICON Tape Write Pipelining locally acknowledges write data records, enabling the host to generate more records while previous records are in transit across the IP WAN. If an exception status is received from the device, the writing of data and emulation is terminated.

**Figure 100: Tape Write Pipelining**






## Tape Read Pipelining

FICON Tape Read Pipelining improves performance for certain applications when reading from FICON tape over extended distances. FICON Tape Read Pipelining reads data on the tape directly from the tape device. Reading of the tape continues until a threshold is reached. The buffered data is forwarded to the host in response to requests from the host. When the host sends the status accept frame indicating that the data was delivered, the read processing on the device side credits the pipeline and requests more data from the tape. If an exception status is received from the device, the reading of data and emulation is terminated.

**Figure 101: Tape Read Pipelining**



## FICON and ESCON Conversion Support

Brocade extension switches and blades can operate with the Prizm FICON to ESCON converter and ESBT Bus/Tag Interface Module for Prizm from Optica Technologies, Inc. This allows FICON tape emulation when connecting to ESCON and to bus and tag devices including 3480, 3490, and 3590 controllers. Discovery of these devices is automatic when FICON emulation is enabled on an FCIP tunnel.

## Teradata Emulation

Teradata emulation reduces latency on links to Teradata Warehouse systems caused by WAN propagation delays and bandwidth restrictions. It accomplishes this by processing selected FICON commands for associated control, data,

and status responses. FICON Teradata emulation is supported only between FICON channels and FICON Teradata controllers.

For write commands, control and status frames are generated for the host side of the WAN in order to pipeline write commands over the same or multiple exchanges. Such pipelined write commands and their data are queued at the device side of the WAN for asynchronous transfer to the device.

For read operations received by the device side of the WAN, a number of anticipatory read commands are autonomously generated and transferred to the device. The data and status associated with such commands are sent to the host side of the WAN and queued in anticipation of host-generated read commands.

FICON emulation for Teradata sequences over an FCIP tunnel is controlled by using FICON emulation feature flags in the `PortCfg fcipunnel` command associated with the FCIP tunnel.

For information on configuring Teradata emulation or tape emulation, see [Configuring FICON emulation](#). For information on displaying Teradata emulation status and statistics, see [Displaying FICON emulation performance statistics](#).

## FCIP Configuration Requirements for FICON Extension

FICON extension uses FCIP for transport. FCIP interfaces and tunnels used for FICON extension must be defined before configuring FICON emulation. Ports should remain persistently disabled until after FICON emulation is configured. Refer to the *Brocade Fabric OS Extension User Guide* for information about configuring FCIP interfaces, tunnels, and circuits.

## Configuration Requirements for Switches and Directors

Consider these configuration issues when an extension switch or blade is connected to a switch or director in a FICON configuration:

- If you are creating a cascaded configuration (connecting two switches or directors with different domain IDs), be aware of IBM requirements for high-integrity fabrics.
- In configurations with ISLs, ensure that ISL paths are properly configured to meet FICON emulation requirements for a determinate path for FICON commands and responses.
- FICON networks with FCIP emulating and nonemulating tunnels supports exchange-based routing (`aptpolicy 3`) configurations.
- When running FICON emulation features over an FCIP tunnel, it is highly recommended that both switches providing the FCIP tunnel use the same Fabric OS release level.

## High-Integrity Fabric Requirements for Cascaded Configurations

You can use extended FICON connections between or through Brocade extension switches or blades to create a cascaded FICON switched configuration. The fabric must be set up to provide a secure, or high-integrity, fabric to enable this configuration to operate in an IBM environment.

## FICON Emulation Requirement for a Determinate Path

FICON emulation processing creates FICON commands and responses on extended CHPIDs, and it intercepts all exchanges between a channel and a CU. For FICON emulation processing to function correctly, all exchanges between a channel and a CU must take a single tunnel path.


There are two ways to ensure a determinate path for FICON commands and responses:

- Define only one FCIP tunnel between sites.
- When multiple FICON emulation FCIP tunnels are required between a pair of chassis, define a new logical switch pair (one logical switch in each chassis) for each tunnel. The logical switches and the FCIP tunnel in those logical switches will provide a determinant path for all FICON sequences between the host and device ports in that logical switch pair.

## FCIP Tunnel between Sites

Brocade extension switches and blades use FCIP trunking features to overcome the limitation of one Ethernet interface, one IP address, and one FCIP tunnel. In Fabric OS v6.3 and later, an FCIP tunnel is created with multiple FCIP circuits over different IP interfaces to provide WAN load balancing and failover recovery in the event of a limited WAN outage. This provides a highly redundant WAN configuration for all FICON or FCP emulation technologies with Fabric OS. The following figure shows that a tunnel can have up to four circuits per GbE port.

**Figure 102: FCIP Tunnel and Circuits**



## Configuring FICON Emulation

Before you configure FICON emulation, you need to understand the available options and whether those options are to be implemented in your installation. For more information, see [FICON Emulation Overview](#).

### ATTENTION

Make sure that both ends of the tunnel match before bringing it up or the tunnel will not form.

1. Verify that the following ports are opened in the IP network for FCIP:

- TCP 3225: FC Class F frames
- TCP 3226: FC Class 2 and 3 frames
- TCP 3227: Broadcom proprietary IP path performance utility (ipPerf)
- IP protocol 6: Transmission Control Protocol
- IP protocol 108: IP Payload Compression Protocol (IPComp)

2. Verify that the correct licenses are installed using the `licenseShow` command. You can verify all slot-based licenses in a chassis using the `licenseSlotCfg --show` command.
3. Verify that the switches have been set up according to the instructions in [Configuring FICON](#).
4. Configure the FCIP tunnel. Refer to the *Brocade Fabric OS Extension User Guide* for more information.
5. Enable compression on the FCIP tunnel.
6. Configure FICON emulation on the tunnel that is solution-specific. See [Configuration Examples](#).
7. If you have more than one tunnel being configured to the same destination and their tunnel commit rates are different, configure the link cost values to be the same.

## Configuration Examples

You can configure FICON emulation with such features as Teradata read, Teradata write, tape write pipelining, tape read pipelining, TIN/TIR, device level ACK, and IBM z/OS Global Mirror (formerly XRC) emulation enabled.

Following are examples for enabling IBM z/OS Global Mirror, Teradata emulation, and tape read and write pipelining.

```
portcfg fciptunnel 23 modify --ficon-xrc 1
portcfg fciptunnel 23 modify --ficon-tera-read 1 1 --ficon-tera-write 1
portcfg fciptunnel 23 modify --ficon-tape-write 1 --ficon-tape-read 1
```

Note that the preceding commands assume that an FCIP tunnel was enabled for FICON operation by using the `-F` or `--ficon` option. Following are two examples:

- To create an empty tunnel as a FICON FCIP tunnel, use the following command:

```
portcfg fciptunnel 1/17 create -F -c 1
```

- To modify a tunnel to be enabled with FICON and hardware compression, use the following command:

```
portcfg fciptunnel 1/17 modify -c 1 -F 1
```

## Displaying FICON Emulation Configuration Values

For Brocade extension switches or blades, use the `portShow fciptunnel` command to display FICON configuration values. Command output shows which emulation features are enabled or disabled and displays related emulation parameters for each tunnel.

For more information, see [Displaying FICON Emulation Performance Statistics](#).

## Modifying FICON Emulation

FICON emulation can be enabled on a per-FCIP-tunnel basis. For Brocade extension switches and blades, use the `portCfg fciptunnel` command to make modifications to the FICON emulation feature.

For details on command usage, refer to the *Brocade Fabric OS Command Reference Manual*.

**ATTENTION**

You must bring the FCIP tunnel down in order to make changes. The following process is disruptive to all traffic on the tunnel.

1. Connect to the switches at both ends of the FCIP tunnel, and log on using an account assigned to the admin role.
2. Disable the FCIP tunnel at both ends where FICON emulation is enabled.
3. Use the `portCfg fcipunnel modify` command for the Brocade extension switches and blades.
4. Make matching FICON emulation changes to both ends of the tunnel.
5. Enable the FCIP tunnel.

**Displaying FICON Emulation Performance Statistics**

You can use the `portShow` command to view performance statistics and monitor the behavior of FICON emulation. You do not need to disable any FICON emulation feature to display statistics.

1. Connect to the switch, and log on using an account assigned to the admin role.
2. Use the `portShow xtun` command to display the statistics for the Brocade extension switches and blades.

See [FICON Emulation Monitoring](#) and [Options for Displaying Statistics](#), and refer to the *Brocade Fabric OS Command Reference Manual* for information on command formats that you can use to display various FICON statistics.

**FICON Emulation Monitoring**

Use the `-emul` option for the `portShow xtun` command to display FICON emulation statistics. The syntax for this command is `portshow xtun [slot/]ve_port -ficon -emul`.

The `-emul` option displays statistics and status for Teradata, tape, and IBM z/OS Global Mirror emulation.

IBM z/OS Global Mirror emulation (formerly eXtended Remote Copy or XRC) and tape statistics are presented differently in output formats for the `-emul` option. The following elements are common to both tape and IBM z/OS Global Mirror emulation output (see [Emulation Statistics Example](#) for an illustration of actual tape and XRC emulation statistics output).

**FDCB Ptr  
Path**

A pointer to the FICON Device Control Block. Support personnel may use this pointer. The device path is in the format VE-HD-HP-DD-DP-LP-CU-DV: where

- VE is the internal VE\_Port number.
- HD is the hex value for the host domain (the entry domain for this host port into the fabric).
- HP is the hex value for the host port (the entry port of this host connection into the fabric).
- DD is the hex value for the device domain (the entry domain for this device into the fabric).
- DP is the hex value for the device port (the entry port for this device connection into the fabric).
- LP is the logical partition (LPAR) value for the host accessing the specific device.
- CU is the control unit number (CUADDR) for the specific FICON connection.
- DV is the Device Number of this FICON connected device.

For example, the value 2463016406050001 breaks down as follows.

| VE | HD | HP | DD | DP | LP | CU | DV |
|----|----|----|----|----|----|----|----|
| 24 | 63 | 01 | 64 | 06 | 05 | 00 | 01 |

**State**

A value of zero indicates that emulation is idle. Any nonzero value indicates that emulation is active.

## Emulation Statistics Examples

The following is a tape statistics output example.

```
admin:switch> portshow xtun 8/16 -ficon -emultape
```

TAPE EMULATION STATS

| FDCB Ptr           | Path             | H | State | Emul | Emul | Rtry | Emulated | Emulated  | RdAvg | Emulated   | WtAvg |
|--------------------|------------------|---|-------|------|------|------|----------|-----------|-------|------------|-------|
| (0x)               | (0x)             | D |       | Pipe | Q'd  | Qd   | Tape Ops | Read CCWs | Size  | Write CCWs | Size  |
| 0x800000028F2A8180 | D0630C6405020001 | H | 0x14  | 0x20 | 0000 | 0000 | 1460     | 0         | 0     | 11680      | 32768 |
| 0x800000028F2AC180 | D0630C640502000C | H | 0x14  | 0x20 | 0000 | 0000 | 1459     | 0         | 0     | 11672      | 32768 |
| 0x800000028F2ADC80 | D0630C640502000E | H | 0x14  | 0x20 | 0000 | 0000 | 1359     | 0         | 0     | 10872      | 32768 |
| 0x800000028F2ACD80 | D0630C640502000F | H | 0x14  | 0x20 | 0000 | 0000 | 1446     | 0         | 0     | 11568      | 32768 |
| 0x800000028F2A3C80 | D0630C6405020101 | H | 0x14  | 0x20 | 0000 | 0000 | 1397     | 0         | 0     | 11176      | 32768 |
| 0x800000028F2A6180 | D0630C6405020105 | H | 0x14  | 0x20 | 0000 | 0000 | 1386     | 0         | 0     | 11088      | 32768 |
| 0x800000028F2A5280 | D0630C6405020106 | H | 0x14  | 0x20 | 0000 | 0000 | 1449     | 0         | 0     | 11592      | 32768 |
| 0x800000028F2A7280 | D0630C640502010B | H | 0x14  | 0x20 | 0000 | 0000 | 1397     | 0         | 0     | 11176      | 32768 |

The following is an IBM z/OS Global Mirror output example.

```
switch:admin> portshow xtun 17 -ficon -emulxrc
```

XRC EMULATION STATS

| FDCB Ptr     | Path             | H | State | Cmds | Cmd  | Data | Data | Emulated | Avg | RRS   | RRS  | Large  |
|--------------|------------------|---|-------|------|------|------|------|----------|-----|-------|------|--------|
| (0x)         | (0x)             | D |       | Qd   | Max  | Qd   | Max  | RRS Ops  | RRS | TLF   | Read | Chains |
| 0x04100D8880 | 1163006400060023 | H | 0x00  | 0000 | 000F | 0000 | 0002 | 1        | 1   | 58864 | 40   | 0      |
| 0x04100DA880 | 116300640006002B | H | 0x00  | 0000 | 000F | 0000 | 0002 | 1        | 1   | 58864 | 40   | 0      |
| 0x041008FB00 | 1163006400060033 | H | 0x00  | 0000 | 0003 | 0000 | 0001 | 1        | 1   | 58864 | 40   | 0      |
| 0x04100ED980 | 116300640006003B | H | 0x00  | 0000 | 0003 | 0000 | 0001 | 1        | 1   | 58864 | 40   | 0      |
| 0x04100DD100 | 1163006400060043 | H | 0x00  | 0000 | 0003 | 0000 | 0001 | 1        | 1   | 58864 | 40   | 0      |
| 0x04100E2000 | 1163006400060047 | H | 0x00  | 0000 | 0003 | 0000 | 0001 | 1        | 1   | 58864 | 40   | 0      |
| 0x04100E7100 | 1163006400060057 | H | 0x00  | 0000 | 0003 | 0000 | 0001 | 1        | 1   | 58864 | 40   | 0      |
| 0x04100ECC80 | 116300640006005F | H | 0x00  | 0000 | 0003 | 0000 | 0001 | 1        | 1   | 58864 | 40   | 0      |

## Options for Displaying Statistics

Use the `portshow xtun [slot/]ve_port -ficon -stats` command to display status, performance statistics, and other statistics for tape, Teradata, and IBM z/OS Global Mirror on Brocade extension switches and blades. See the following table for a description of command formats to achieve various output.

**Table 151: FICON Emulation Statistics Command**

| Command                                                            | Output                                                                                                                                           |
|--------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>portshow xtun slot/ve_port ve_port -ficon -emul</code>       | Statistics and status for Teradata, tape, and IBM z/OS Global Mirror (XRC) emulation.                                                            |
| <code>portshow xtun slot/ve_port -ficon -emulxrc</code>            | FICON IBM z/OS Global Mirror emulation statistics and status.                                                                                    |
| <code>portshow xtun slot/ve_port -ficon -emultape</code>           | FICON tape emulation statistics and status.                                                                                                      |
| <code>portshow xtun slot/ve_port -ficon -emultera</code>           | FICON Teradata emulation statistics and status.                                                                                                  |
| <code>portshow xtun slot/ve_port -ficon -stats</code>              | General statistics, including those for Teradata, XRC, tape emulation, and FICON idle status acceptance.                                         |
| <code>portshow xtun slot/ve_port -ficon -fdcb control_block</code> | Teradata, tape, or IBM z/OS Global Mirror emulation statistics, including the device extension of the FDCB for a specified Device Control Block. |
| <code>portshow xtun ve_port ficon -teraperf</code>                 | Teradata performance statistics.                                                                                                                 |
| <code>portshow xtun ve_port -ficon -tapeperf</code>                | FICON tape pipelining performance statistics.                                                                                                    |
| <code>portshow xtun ve_port -ficon -xrcperf</code>                 | IBM z/OS Global Mirror performance statistics.                                                                                                   |

Refer to the `portShow` command in the *Brocade Fabric OS Command Reference Manual* for more information on these commands.

## Platforms Supporting FICON

This section provides the following support details:

- Brocade products currently supported for FICON for different IBM systems and drivers ([Currently Supported Platforms](#)).
- Brocade blades supported for FICON in various Brocade switch and data center backbone products ([Supported Brocade Blades](#)).

Due to the nature of environments in which the mainframes are used, product qualification for FICON is more rigorous than with other products. As a result, not all product combinations and versions of Fabric OS are qualified for FICON.

### NOTE

Refer to the Fabric OS release notes and the [FICON qualification support letters](#) for additional details about compatibility with IBM systems and drivers. Check that your current versions of IBM systems and drivers are compatible with the Fabric OS firmware that you plan to use.

## Currently Supported Platforms

The following table lists Brocade products that are currently supported for Fabric OS 9.1.0 for various IBM z Systems and driver levels.

### NOTE

Consult your Brocade SE for detailed information on the IBM z14 qualified support. The IBM z14 was introduced on July 17, 2017.

**Table 152: Currently Supported Gen 7 and Gen 6 Platforms**

|                      | Brocade G720<br>Fabric OS | Brocade X7-4<br>Fabric OS | Brocade X7-8<br>Fabric OS | Brocade G620<br>Fabric OS | Brocade X6-4<br>Fabric OS | Brocade X6-8<br>Fabric OS |
|----------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|
| <b>z13 and newer</b> | 9.1.0b                    | 9.1.0b                    | 9.1.0b                    | 9.1.0b                    | 9.1.0b                    | 9.1.0b                    |

## Supported Brocade Blades

The following table lists Brocade blades supported for FICON in various Brocade data center backbone and director products in Fabric OS 9.x.

**Table 153: Currently Supported Blades in Brocade Data Center Backbone and Director Products**

|            | X6-8 | X6-4 | X7-8 | X7-4 |
|------------|------|------|------|------|
| SX6        | Yes  | Yes  | Yes  | Yes  |
| FC32-48    | Yes  | Yes  | Yes  | Yes  |
| CR32-4     | Yes  | Yes  | Yes  | Yes  |
| CR32-8     | Yes  | No   | Yes  | Yes  |
| CR64-4     | No   | No   | Yes  | Yes  |
| CR64-8     | No   | No   | Yes  | Yes  |
| FC64-48    | No   | No   | Yes  | Yes  |
| FC32-X7-48 | No   | No   | Yes  | Yes  |



## Port Indexing

This section shows how to use the `switchshow` command to determine the mapping among the port index, slot or port numbers, and the 24-bit port ID (PID) on any Brocade director.

Enter the `switchShow` command without parameters to show the port index mapping for the entire platform. Enter the `switchShow -slot` command for port-mapping information for the ports on the blade in a specific slot. Include the `--qsfp` option to list also the QSFP number, for slots that contain core blades.

This example shows the output of the `switchShow` command for a FC32-64 port blade in slot 3 of a Brocade X7-8 Director. The leftmost column shows the unique port index. The second and third columns show the corresponding physical slot and port numbers, respectively. The corresponding QSFP number for the port is also shown in the fourth column. For a core blade, no PID exists in the Address column.

```
admin> switchshow -slot 3 -qsfp
switchName: switch1
switchType: 180.0
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 57
switchId: fffc014c
switchWwn: 10:00:c4:f5:7c:94:81:4c
zoning: ON (fos6k)
switchBeacon: OFF
FC Router: OFF
Fabric Name: scale_fid128
HIF Mode: ON
Allow XISL Use: OFF
LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]
Slot Blade Type ID Model Name Status

 3 SW BLADE 204 FC32-64 ENABLED

Index Slot Port QSFP Address Media Speed State Proto
=====
 0 3 0 0 390000 id N32 Online FC E-Port 10:00:c4:f5:7c:01:ae:50
"146152" (downstream) (Trunk master)
 1 3 1 0 390100 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
0)
 2 3 2 0 390200 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
0)
 3 3 3 0 390300 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
0)
 4 3 4 1 390400 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
0)
 5 3 5 1 390500 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
0)
 6 3 6 1 390600 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
0)
```

```

7 3 7 1 390700 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
0)
8 3 8 2 390800 id N32 Online FC E-Port 10:00:c4:f5:7c:01:9f:78
"E_Cassian_146151" (downstream) (Trunk master)
9 3 9 2 390900 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
8)
10 3 10 2 390a00 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
8)
11 3 11 2 390b00 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
8)
12 3 12 3 390c00 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
8)
13 3 13 3 390d00 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
8)
14 3 14 3 390e00 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
8)
15 3 15 3 390f00 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
8)
16 3 16 4 391000 id N16 Online FC E-Port 10:00:00:27:f8:f1:4a:e0
"EG620_146195_LCW" (downstream) (Trunk master)
17 3 17 4 391100 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
16)
18 3 18 4 391200 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
16)
19 3 19 4 391300 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
16)
20 3 20 -- 391400 -- N32 No_Module FC
21 3 21 -- 391500 -- N32 No_Module FC
22 3 22 -- 391600 -- N32 No_Module FC
23 3 23 -- 391700 -- N32 No_Module FC
24 3 24 6 391800 id N16 Online FC E-Port 10:00:c4:f5:7c:ce:93:70
"EG620_146196" (downstream) (Trunk master)
25 3 25 6 391900 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
24)
26 3 26 6 391a00 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
24)
27 3 27 6 391b00 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
24)
28 3 28 7 391c00 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
24)
29 3 29 7 391d00 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
24)
30 3 30 7 391e00 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
24)
31 3 31 7 391f00 id N16 Online FC E-Port (Trunk port, master is Slot 3 Port
24)
32 3 32 8 392000 id N32 Online FC E-Port 10:00:88:94:71:93:4d:b8
"EG620_146194_LCW" (downstream) (Trunk master)
33 3 33 8 392100 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
32)
34 3 34 8 392200 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
32)

```

|                                                   |   |    |    |        |    |     |         |    |        |                                         |
|---------------------------------------------------|---|----|----|--------|----|-----|---------|----|--------|-----------------------------------------|
| 35                                                | 3 | 35 | 8  | 392300 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 32 ) |
| 36                                                | 3 | 36 | 9  | 392400 | id | N16 | Online  | FC | E-Port | 10:00:88:94:71:93:4d:b8                 |
| "EG620_146194_LCW" (Trunk master)                 |   |    |    |        |    |     |         |    |        |                                         |
| 37                                                | 3 | 37 | 9  | 392500 | id | N16 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 36 ) |
| 38                                                | 3 | 38 | 9  | 392600 | id | N16 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 36 ) |
| 39                                                | 3 | 39 | 9  | 392700 | id | N16 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 36 ) |
| 40                                                | 3 | 40 | 10 | 392800 | id | N32 | Online  | FC | E-Port | 10:00:c4:f5:7c:01:20:20                 |
| "Core_G630_114205" (downstream) (Trunk master)    |   |    |    |        |    |     |         |    |        |                                         |
| 41                                                | 3 | 41 | 10 | 392900 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 40 ) |
| 42                                                | 3 | 42 | 10 | 392a00 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 40 ) |
| 43                                                | 3 | 43 | 10 | 392b00 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 40 ) |
| 44                                                | 3 | 44 | 11 | 392c00 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 40 ) |
| 45                                                | 3 | 45 | 11 | 392d00 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 40 ) |
| 46                                                | 3 | 46 | 11 | 392e00 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 40 ) |
| 47                                                | 3 | 47 | 11 | 392f00 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 40 ) |
| 48                                                | 3 | 48 | 12 | 393000 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 52 ) |
| 49                                                | 3 | 49 | 12 | 393100 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 52 ) |
| 50                                                | 3 | 50 | 12 | 393200 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 52 ) |
| 51                                                | 3 | 51 | 12 | 393300 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 52 ) |
| 52                                                | 3 | 52 | 13 | 393400 | id | N32 | Online  | FC | E-Port | 10:00:d8:1f:cc:05:3f:ba                 |
| "E_G630_2_146209_C5c" (downstream) (Trunk master) |   |    |    |        |    |     |         |    |        |                                         |
| 53                                                | 3 | 53 | 13 | 393500 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 52 ) |
| 54                                                | 3 | 54 | 13 | 393600 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 52 ) |
| 55                                                | 3 | 55 | 13 | 393700 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 52 ) |
| 56                                                | 3 | 56 | 14 | 393800 | id | N32 | Online  | FC | E-Port | 10:00:00:27:f8:f1:4a:e0                 |
| "EG620_146195_LCW" (Trunk master)                 |   |    |    |        |    |     |         |    |        |                                         |
| 57                                                | 3 | 57 | 14 | 393900 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 56 ) |
| 58                                                | 3 | 58 | 14 | 393a00 | id | N16 | No_Sync | FC |        |                                         |
| 59                                                | 3 | 59 | 14 | 393b00 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 56 ) |
| 60                                                | 3 | 60 | 15 | 393c00 | id | N32 | Online  | FC | E-Port | 10:00:00:27:f8:f1:4a:e0                 |
| "EG620_146195_LCW" (Trunk master)                 |   |    |    |        |    |     |         |    |        |                                         |
| 61                                                | 3 | 61 | 15 | 393d00 | id | N32 | Online  | FC | E-Port | (Trunk port, master is Slot 3 Port 60 ) |

```

62 3 62 15 393e00 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
60)
63 3 63 15 393f00 id N32 Online FC E-Port (Trunk port, master is Slot 3 Port
60)

```

This example shows the truncated output of the `switchShow` command for an FC32-X7-48 port blade in slot 5 of Brocade X7-8 Director. The Address column shows the PID.

```

switch:admin> switchshow -slot 5 -qsfp
switchName: Core-X7-8_130239
switchType: 180.0
switchState: Online
switchMode: Native
switchRole: Subordinate
switchDomain: 239
switchId: fffcef
switchWwn: 10:00:88:94:71:ba:cd:e1
zoning: ON (fos6k)
switchBeacon: OFF
FC Router: OFF
Fabric Name: scale_fid128
HIF Mode: OFF
Allow XISL Use: OFF
LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]
Slot Blade Type ID Model Name Status

5 SW BLADE 218 FC32-X7-48 ENABLED

Index Slot Port Address Media Speed State Proto
=====
192 5 0 efc000 id N32 Online FC E-Port 10:00:d8:1f:cc:05:3f:ba
"E_G630_2_146209_C5c" (Trunk master)
193 5 1 efc100 id N32 Online FC E-Port 10:00:d8:1f:cc:05:3f:ba "E_G630_2_146209_C5c"
194 5 2 efc200 id N32 Online FC E-Port (Trunk port, master is Slot 5 Port 0)
195 5 3 efc300 id N32 Online FC E-Port (Trunk port, master is Slot 5 Port 0)
196 5 4 efc400 id N16 Online FC E-Port (Trunk port, master is Slot 5 Port 7)
197 5 5 efc500 id N16 Online FC E-Port (Trunk port, master is Slot 5 Port 7)
198 5 6 efc600 id N16 Online FC E-Port (Trunk port, master is Slot 5 Port 7)
199 5 7 efc700 id N16 Online FC E-Port 10:00:50:eb:1a:b5:27:f8
"SW6505_114121" (Trunk master)
200 5 8 efc800 id N16 Online FC E-Port 10:00:50:eb:1a:be:6f:43
"SW6510_114135" (downstream) (Trunk master)
201 5 9 efc900 id N16 Online FC E-Port (Trunk port, master is Slot 5 Port 11)
202 5 10 efca00 id N16 Online FC E-Port (Trunk port, master is Slot 5 Port 11)
203 5 11 efcb00 id N16 Online FC E-Port 10:00:50:eb:1a:be:6f:43
"SW6510_114135" (Trunk master)
204 5 12 efcc00 id N16 Online FC E-Port 10:00:50:eb:1a:b5:27:98
"SW6505_114123" (downstream) (Trunk master)
205 5 13 efcd00 id N16 Online FC E-Port (Trunk port, master is Slot 5 Port 12)
206 5 14 efce00 id N16 Online FC E-Port (Trunk port, master is Slot 5 Port 12)
207 5 15 efcf00 id N16 Online FC E-Port (Trunk port, master is Slot 5 Port 12)
208 5 16 efd000 id N32 No_Light FC

```

|                                             |   |    |        |    |     |           |    |                       |                                         |
|---------------------------------------------|---|----|--------|----|-----|-----------|----|-----------------------|-----------------------------------------|
| 209                                         | 5 | 17 | efd100 | id | N32 | No_Light  | FC |                       |                                         |
| 210                                         | 5 | 18 | efd200 | id | N32 | No_Light  | FC |                       |                                         |
| 211                                         | 5 | 19 | efd300 | id | N32 | No_Light  | FC |                       |                                         |
| 212                                         | 5 | 20 | efd400 | id | N32 | No_Light  | FC |                       |                                         |
| 213                                         | 5 | 21 | efd500 | id | N32 | No_Light  | FC |                       |                                         |
| 214                                         | 5 | 22 | efd600 | id | N32 | No_Light  | FC |                       |                                         |
| 215                                         | 5 | 23 | efd700 | id | N32 | No_Light  | FC |                       |                                         |
| 216                                         | 5 | 24 | efd800 | id | N32 | No_Light  | FC |                       |                                         |
| 217                                         | 5 | 25 | efd900 | id | N32 | No_Light  | FC |                       |                                         |
| 218                                         | 5 | 26 | efda00 | id | N32 | No_Light  | FC |                       |                                         |
| 219                                         | 5 | 27 | efdb00 | id | N32 | No_SigDet | FC | Disabled (Persistent) | (None)                                  |
| 220                                         | 5 | 28 | efdc00 | id | N16 | Online    | FC | E-Port                | (Trunk port, master is Slot 5 Port 29 ) |
| 221                                         | 5 | 29 | efdd00 | id | N16 | Online    | FC | E-Port                | 10:00:50:eb:1a:b7:11:24                 |
| "SW6505_114122" (downstream) (Trunk master) |   |    |        |    |     |           |    |                       |                                         |
| 222                                         | 5 | 30 | efde00 | id | N16 | Online    | FC | E-Port                | (Trunk port, master is Slot 5 Port 29 ) |
| 223                                         | 5 | 31 | efdf00 | id | N16 | Online    | FC | E-Port                | (Trunk port, master is Slot 5 Port 29 ) |
| 224                                         | 5 | 32 | efe000 | id | N32 | No_Light  | FC |                       |                                         |
| 225                                         | 5 | 33 | efe100 | id | N32 | No_Light  | FC |                       |                                         |
| 226                                         | 5 | 34 | efe200 | id | N32 | No_Light  | FC |                       |                                         |
| 227                                         | 5 | 35 | efe300 | id | N32 | No_Light  | FC |                       |                                         |
| 228                                         | 5 | 36 | efe400 | id | N8  | Online    | FC | F-Port                | 20:01:00:11:0d:ac:01:00                 |
| 229                                         | 5 | 37 | efe500 | id | N8  | Online    | FC | F-Port                | 20:00:00:11:0d:ac:00:00                 |
| 230                                         | 5 | 38 | efe600 | id | N8  | Online    | FC | F-Port                | 20:05:00:11:0d:2a:29:00                 |
| 231                                         | 5 | 39 | efe700 | id | N8  | Online    | FC | F-Port                | 20:04:00:11:0d:2a:28:00                 |
| 232                                         | 5 | 40 | efe800 | id | N8  | Online    | FC | F-Port                | 20:03:00:11:0d:2b:01:00                 |
| 233                                         | 5 | 41 | efe900 | id | N8  | Online    | FC | F-Port                | 20:02:00:11:0d:2b:00:00                 |
| 234                                         | 5 | 42 | efea00 | id | N16 | Online    | FC | F-Port                | 1 N Port + 7 NPIV public                |
| 235                                         | 5 | 43 | efeb00 | id | N16 | Online    | FC | F-Port                | 1 N Port + 7 NPIV public                |
| 236                                         | 5 | 44 | efec00 | id | N16 | Online    | FC | F-Port                | 1 N Port + 7 NPIV public                |
| 237                                         | 5 | 45 | efed00 | id | N16 | Online    | FC | F-Port                | 1 N Port + 7 NPIV public                |
| 238                                         | 5 | 46 | efee00 | id | N8  | Online    | FC | F-Port                | 20:07:00:11:0d:63:01:00                 |
| 239                                         | 5 | 47 | efef00 | id | N8  | Online    | FC | F-Port                | 20:06:00:11:0d:63:00:00                 |

This example shows the truncated `switchShow` output for an SX6 application blade on the X7-8 Director. The assignment of port index numbers to PIDs varies depending on the blade type, platform type, and slot number.

```
switch:admin> switchshow -slot 12
switchName: Core-X7-8
switchType: 180.0
switchState: Online
switchMode: Native
switchRole: Subordinate
switchDomain: 239
switchId: fffcef
switchWwn: 10:00:88:94:71:ba:cd:e1
zoning: ON (fos6k)
switchBeacon: OFF
FC Router: OFF
Fabric Name: scale_fid128
HIF Mode: OFF
Allow XISL Use: OFF
```

LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]

| Slot | Blade Type | ID | Model Name | Status |
|------|------------|----|------------|--------|
|------|------------|----|------------|--------|

|    |          |     |     |         |
|----|----------|-----|-----|---------|
| 12 | AP BLADE | 186 | SX6 | ENABLED |
|----|----------|-----|-----|---------|

| Index                             | Slot | Port | Address | Media | Speed | State     | Proto                                              |
|-----------------------------------|------|------|---------|-------|-------|-----------|----------------------------------------------------|
| 672                               | 12   | 0    | efa080  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 673                               | 12   | 1    | efa180  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 674                               | 12   | 2    | efa280  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 675                               | 12   | 3    | efa380  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 676                               | 12   | 4    | efa480  | id    | N32   | No_Sync   | FC Disabled (Persistent) (None)                    |
| 677                               | 12   | 5    | efa580  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 678                               | 12   | 6    | efa680  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 679                               | 12   | 7    | efa780  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 680                               | 12   | 8    | efa880  | id    | N16   | Online    | FC F-Port 20:05:00:11:0d:9a:6b:00                  |
| 681                               | 12   | 9    | efa980  | id    | N32   | No_Sync   | FC Disabled (Persistent) (None)                    |
| 682                               | 12   | 10   | efaa80  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 683                               | 12   | 11   | efab80  | id    | N32   | Online    | FC F-Port 1 N Port + 7 NPIV public                 |
| 684                               | 12   | 12   | efac80  | id    | N16   | Online    | FC E-Port 10:00:00:27:f8:f1:4a:e0                  |
| "EG620_146195_LCW" (Trunk master) |      |      |         |       |       |           |                                                    |
| 685                               | 12   | 13   | efad80  | id    | N16   | Online    | FC E-Port (Trunk port, master is Slot 12 Port 12 ) |
| 686                               | 12   | 14   | efae80  | id    | N16   | Online    | FC E-Port (Trunk port, master is Slot 12 Port 12 ) |
| 687                               | 12   | 15   | efaf80  | id    | N16   | Online    | FC E-Port (Trunk port, master is Slot 12 Port 12 ) |
| 688                               | 12   | 16   | efb080  | --    | --    | Online    | VE VE-Port 10:00:88:94:71:42:cc:a7 "Core_7810"     |
| 689                               | 12   | 17   | efb180  | --    | --    | Offline   | VE                                                 |
| 690                               | 12   | 18   | efb280  | --    | --    | Offline   | VE                                                 |
| 691                               | 12   | 19   | efb380  | --    | --    | Offline   | VE                                                 |
| 692                               | 12   | 20   | efb480  | --    | --    | Offline   | VE                                                 |
| 693                               | 12   | 21   | efb580  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 694                               | 12   | 22   | efb680  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 695                               | 12   | 23   | efb780  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 696                               | 12   | 24   | efb880  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 697                               | 12   | 25   | efb980  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 698                               | 12   | 26   | efba80  | --    | --    | Offline   | VE                                                 |
| 699                               | 12   | 27   | efbb80  | --    | --    | Offline   | VE                                                 |
| 700                               | 12   | 28   | efbc80  | --    | --    | Offline   | VE                                                 |
| 701                               | 12   | 29   | efbd80  | --    | --    | Online    | VE VE-Port 10:00:00:05:33:65:a6:43 "Core_C7810"    |
| 702                               | 12   | 30   | efbe80  | --    | --    | Offline   | VE                                                 |
| 703                               | 12   | 31   | efbf80  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 704                               | 12   | 32   | efc080  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 705                               | 12   | 33   | efc180  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 706                               | 12   | 34   | efc280  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
| 707                               | 12   | 35   | efc380  | --    | --    | Offline   | VE Disabled (10VE Mode)                            |
|                                   | 12   | ge0  |         | id    | 40G   | No_Light  | FCIP                                               |
|                                   | 12   | ge1  |         | id    | 40G   | Online    | FCIP                                               |
|                                   | 12   | ge2  |         | id    | 10G   | Online    | FCIP                                               |
|                                   | 12   | ge3  |         | --    | 10G   | No_Module | FCIP                                               |
|                                   | 12   | ge4  |         | --    | 10G   | No_Module | FCIP                                               |
|                                   | 12   | ge5  |         | --    | 10G   | No_Module | FCIP                                               |
|                                   | 12   | ge6  |         | --    | 10G   | No_Module | FCIP                                               |
|                                   | 12   | ge7  |         | --    | 10G   | No_Module | FCIP                                               |
|                                   | 12   | ge8  |         | --    | 10G   | No_Module | FCIP                                               |

---

|    |      |    |     |           |      |
|----|------|----|-----|-----------|------|
| 12 | ge9  | -- | 10G | No_Module | FCIP |
| 12 | ge10 | -- | 10G | No_Module | FCIP |
| 12 | ge11 | -- | 10G | No_Module | FCIP |
| 12 | ge12 | -- | 10G | No_Module | FCIP |
| 12 | ge13 | -- | 10G | No_Module | FCIP |
| 12 | ge14 | -- | 10G | No_Module | FCIP |
| 12 | ge15 | -- | 10G | No_Module | FCIP |
| 12 | ge16 | -- | 10G | No_Module | FCIP |
| 12 | ge17 | -- | 10G | No_Module | FCIP |

## Brocade Switch and Director Sensors

This section lists the number of sensors, fans, and power supply units (PSUs) for Gen 7 and Gen 6 platforms. The following table lists the number of temperature sensors, fans, and PSUs for each Brocade switch type.

**Table 154: Brocade Switch and Director Sensors and PSUs**

| Platform                      | Temperature Sensors                | Fans                               | Power Supply Units |
|-------------------------------|------------------------------------|------------------------------------|--------------------|
| Brocade G730 Switch           | 39                                 | 3                                  | 2                  |
| Brocade G720 Switch           | 18                                 | 2                                  | 2                  |
| Brocade X7-8 Director         | 58+                                | 3                                  | 4                  |
| Brocade X7-4 Director         | 35+                                | 2                                  | 2                  |
| Brocade G630                  | 10 <sup>a</sup><br>19 <sup>b</sup> | 3 (a minimum of 2 must be working) | 2                  |
| Brocade G620                  | 7 <sup>c</sup><br>12 <sup>d</sup>  | 2                                  | 2                  |
| Brocade G610                  | 9                                  | 4 (a minimum of 2 must be working) | 1                  |
| Brocade 7810 Extension Switch | 12                                 | 2                                  | 2                  |
| Brocade X6-8 Director         | 58+                                | 3                                  | 4 (3 minimum)      |
| Brocade X6-4 Director         | 25+                                | 2                                  | 2                  |

## Brocade Blade Temperature Sensors

This section lists the number of temperature sensors for the blades of Gen 7 and Gen 6 directors. The following table lists the number of temperature sensors for each Brocade blade model.

**Table 155: Brocade Blade Temperature Sensors**

| Model      | Temperature Sensors |
|------------|---------------------|
| FC64-48    | 11                  |
| FC32-X7-48 | 8                   |
| CPX7       | 4                   |
| CR64-8     | 12                  |
| CR64-4     | 9                   |
| FC32-64    | 5                   |
| FC32-48    | 3                   |
| SX6        | 2                   |

a. Amount of temperature sensors for switch type 173

b. Amount of temperature sensors for switch type 184

c. Amount of temperature sensors for switch type 162

d. Amount of temperature sensors for switch type 183



| Model  | Temperature Sensors |
|--------|---------------------|
| CPX6   | 4                   |
| CR32-8 | 5                   |
| CR32-4 | 3                   |

## System Temperature Monitoring

Brocade blades, chassis, and fixed-port switches are continuously monitored for thermal safety. Fabric OS thermal policies are based on a matrix of sensor values particular to each device. Different versions of Fabric OS may also have different thermal policies, because these limits are determined by testing and real-world experience. For all Brocade devices, fan speeds alter to accommodate temperature changes. As a general rule, devices in danger of overheating will be shut down after a two-minute warning. The blades may shutdown sooner if the temperature continues to rise drastically.

## Appendix A: Setting Up the AAA Server Configuration

This section describes the server configuration for RADIUS, LDAP, and TACACS+ server.

### Setting Up a RADIUS Server

#### NOTE

To set up the RADIUS server, you must know the switch IP address (in either IPv4 or IPv6 notation) or the name to connect to switches. Use the `ipAddrShow` command to display a switch IP address.

For Brocade directors, the switch IP addresses are aliases of the physical Ethernet interfaces on the CP blades. When specifying client IP addresses for the logical switches in these systems, make sure that the CP blade IP addresses are used. For accessing both the active and standby CP blades, and for the purpose of HA failover, both of the CP blade IP addresses must be included in the RADIUS server configuration.

User accounts should be set up by their true network-wide identities rather than by the account names created on a Fabric OS switch. Along with each account name, the administrator must assign appropriate switch access permissions. To manage a fabric, set these permissions to user, admin, and securityAdmin.

When using a fully qualified domain name (FQDN) for RADIUS server, configuring the DNS to return an IPv6 address may be resolved to return an IPv6 or IPv4 address.

### Configuring RADIUS Server Support with Linux

The following procedures work for FreeRADIUS on Solaris and Red Hat Linux. FreeRADIUS is a freeware RADIUS server that you can find at the following website:

<http://www.freeradius.org>

Follow the installation instructions at the website. FreeRADIUS runs on Linux (all versions), FreeBSD, NetBSD, and Solaris. If you make a change to any of the files used in this configuration, you must stop the server and restart it for the changes to take effect.

FreeRADIUS installation places the configuration files in `$PREFIX/etc/raddb`. By default, the PREFIX is `/usr/local`.

Configuring RADIUS service on Linux consists of the following tasks:

- Adding the Brocade attributes to the server
- Creating the user
- Enabling clients

#### Adding the Brocade Attributes to the Server

1. Create and save the `$PREFIX/etc/raddb/dictionary.brocade` file with the following information:

```
dictionary.brocade
#
VENDOR Brocade 1588
#
attributes
#
ATTRIBUTE Brocade-Auth-Role 1 string Brocade
ATTRIBUTE Brocade-AVPairs1 2 string Brocade
ATTRIBUTE Brocade-AVPairs2 3 string Brocade
ATTRIBUTE Brocade-AVPairs3 4 string Brocade
```

|           |                           |   |        |         |
|-----------|---------------------------|---|--------|---------|
| ATTRIBUTE | Brocade-AVPairs4          | 5 | string | Brocade |
| ATTRIBUTE | Brocade-Passwd-ExpiryDate | 6 | string | Brocade |
| ATTRIBUTE | Brocade-Passwd-WarnPeriod | 7 | string | Brocade |

This information defines the Brocade vendor ID as 1588, Brocade attribute 1 as Brocade-Auth-Role, Brocade attribute 6 as Brocade-Passwd-ExpiryDate, and Brocade attribute 7 as Brocade-Passwd-WarnPeriod.

2. Open the `$PREFIX/etc/raddb/dictionary` file in a text editor and add the line:

```
$INCLUDE dictionary.brocade
```

As a result, the `dictionary.brocade` file is located in the RADIUS configuration directory and loaded for use by the RADIUS server.

## Creating the User

1. Open the `$PREFIX/etc/raddb/user` file in a text editor.
2. Add the user names and their permissions for users accessing the switch and authenticating through RADIUS.

The user logs in using the permissions that are specified with Brocade-Auth-Role. The valid permissions include the admin, switchadmin, zoneadmin, securityadmin, basicswitchadmin, fabricadmin, operator, and user. Use quotation marks around "password" and "role".

### NOTE

From Fabric OS 9.1.0 release, the underscore "\_" is the only special character that is allowed at the start of the username, whenever you are creating a user on the AAA server.

### Example of Adding a User Name to the RADIUS Authentication

For example, to set up an account called JohnDoe with admin permissions with a password expiry date of May 28, 2016 and a warning period of 30 days:

```
JohnDoe Auth-Type := Local
User-Password == "johnPassword",
Brocade-Auth-Role = "admin",
Brocade-Passwd-ExpiryDate = "05/28/16",
Brocade-Passwd-WarnPeriod = "30"
```

### Example of Using the RADIUS Server Local System Password to Authenticate Users

The following example uses the local system password file to authenticate users.

```
Auth-Type := System
Brocade-Auth-Role = "admin",
Brocade-AVPairs1 = "HomeLF=70",
Brocade-AVPairs2 = "LFRoleList=admin:2,4-8,70,80,128",
Brocade-AVPairs3 = "ChassisRole=switchadmin",
Brocade-Passwd-ExpiryDate = "11/10/2008",
Brocade-Passwd-WarnPeriod = "30"
```

When you use network information service (NIS) for authentication, the only way to enable authentication with the password file is to force the Brocade switch to authenticate using Password Authentication Protocol (PAP); this requires the `-a pap` option with the `aaaConfig` command.

## Enabling Clients

Clients are the switches that use the RADIUS server; each client must be defined. By default, all IP addresses are blocked.

The Brocade Directors send their RADIUS requests using the IP address of the active CP. When adding clients, add both the active and standby CP IP addresses so that, in the event of a failover, users can still log in to the switch.

1. Open the `$PREFIX/etc/raddb/client` file in a text editor, and add the switches that are to be configured as RADIUS clients.

For example, to configure the switch at IP address 10.1.18.59 as a client, the following is displayed:

```
client 10.1.18.59
 secret = Secret
 shortname = Testing Switch
 nastype = other
```

In this example, *shortname* is an alias used to easily identify the client. *Secret* is the shared secret between the client and server. Make sure the shared secret matches that configured on the switch (refer to [Adding an authentication server to the switch configuration](#)).

2. Save the `$PREFIX/etc/raddb/client.config` file, and then start the RADIUS server as follows:

```
$PREFIX/sbin/radiusd
```

## Configuring RADIUS Server Support with Windows 2012

The instructions for setting up RADIUS on a Windows 2012 server are listed here for your convenience but are not guaranteed to be accurate for your network environment. Always check with your system administrator before proceeding with setup.

### NOTE

All instructions involving Microsoft Windows 2012 can be obtained from [www.microsoft.com](http://www.microsoft.com) or your Microsoft documentation. Confer with your system or network administrator prior to configuration for any special needs that your network environment might have.

### NOTE

For RADIUS PEAP-MSCHAPv2 authentication with a Windows server, if the user account name matches with a role/group on the server with administrative privileges such as *admin* or *maintenance*, the server drops the request.

Configuring RADIUS service on Windows 2012 consists of the following steps:

1. **Installing Internet Authentication Service (IAS)**

For more information and instructions on installing IAS, refer to the Microsoft website.

2. **Enabling the Challenge Handshake Authentication Protocol (CHAP)**

If CHAP authentication is required, then Windows must be configured to store passwords with the reversible encryption. Reverse password encryption is not the default behavior; it must be enabled.

### NOTE

If a user is configured prior to enabling reverse password encryption, the user's password is stored and cannot utilize CHAP. To use CHAP, the password must be re-entered after encryption is enabled. If the password is not re-entered, CHAP authentication will not work and the user will be unable to authenticate from the switch. Alternatives to using CHAP are Password Authentication Protocol (PAP) or PEAP-MSCHAPv2.

3. **Configuring a user**

IAS is the Microsoft implementation of a RADIUS server and proxy. IAS uses the Windows native user database to verify user login credentials; it does not list specific users, but instead lists *user groups*. Each user group should be

associated with a specific switch role. For example, you should configure a user group for admin, switchAdmin, and user, and then add any users whose logins you want to associate with the appropriate group.

#### 4. Configuring the server

For more information and instructions on configuring the server, refer to the Microsoft website. You need the following information to configure the RADIUS server for a Brocade switch. A client is the device that uses the RADIUS server; in this case, it is the switch.

#### 5. Check if the Active Directory Users and Computers tool is available. If not, select **Server Manager > Manager > Add Roles and Feature**.

The Add Role and Feature wizard is displayed. Follow these steps to install Active Directory Users and Computers.

- a) Select **Installation Type** from the left pane and choose the **Role-based** or **feature-based Installation** option.
- b) Select **Server Selection** from the left pane and choose the **Select a server** from the server pool option.
- c) Select the server from the list that is displayed and click **Next**.

The Features tab is displayed.

- d) From the **Features** tab, select all the check boxes starting from the **Remote Server Administration Tools** check box until the **AD DS Snap-Ins and Command-Line Tools** element check box in the navigation tree.
- e) From the **Features** tab, select the **Group Policy Management** check box, and then click **Next**.

The **Confirmation** tab is displayed.

- f) Select the **Restart the destination server automatically if required** check box and click **Finish**.

#### 6. Check if the switch is part of RADIUS client. If not, select **Administrative tools > Network policy server > Radius clients and Servers > Radius client** from the menu. Right-click and add the switch IP address as **Radius client**.

Client address (IP or DNS)—Enter the IP address of the switch.

Client-Vendor—Select **RADIUS Standard**.

Shared secret—Provide a password. Shared secret is a password used between the client device and server to prevent IP address spoofing by unwanted clients. Keep your shared secret password in a safe place. You need to enter this password in the switch configuration.

After clicking **Finish**, add a new client for all switches on which RADIUS authentication will be used.

- a) In the Internet Authentication Service window, right-click the Remote Access Policies folder, and then select **New Remote Access Policy** from the popup.

You must create a remote access policy for each group of Brocade login permissions (admin, switchAdmin, and user) for which you want to use RADIUS. Apply this policy to the user groups that you already created.

- b) In the Vendor-Specific Attribute Information window, enter the vendor code value **1588**. Click **Yes**. It conforms option, and then clicks **Configure Attribute**.
- c) In the Configure VSA (RFC compliant) window, enter the following values, and then click **OK**.

Vendor-assigned attribute number -- Enter the value **1**.

Attribute format -- Enter **String**.

Attribute value -- Enter the login role (admin, switchAdmin, user, and so on) that the user group must use to log in to the switch.

- d) After returning to the **Internet Authentication Service** window, add additional policies for all Brocade login types for which you want to use the RADIUS server. After this is done, you can configure the switch.

7. **Add Admin group and CHAP/PAP/PEAP group before adding the user.** To add Admin group in **Active Directory User and Computers**, follow these steps:
  - a) In **Server Manager**, open **AD DS** tab and then open **Active Directory User and Computers**.
  - b) Right-click on the left panel tab and select **New > Group** from the menu.  
The **New Object - Group** dialog is displayed.
  - c) Enter the Group name as *admin*. Select the **Group scope** as **Global** and **Group type** as **Security** and then click **OK**.  
The **admin Properties** window is displayed.
  - d) Enter the Description as *FOS admin role* and click **Apply**.
8. **Add CHAP group in Active Directory Users and Computers.** To add CHAP group, follow these steps:
  - a) In **Server Manager**, open **AD DS** tab and then open **Active Directory User and Computers**.
  - b) Right-click on the left panel tab and select **New > Group** from the menu.  
The **New Object - Group** dialog is displayed.
  - c) Enter the Group name as *chap*. Select the **Group scope** as **Global** and **Group type** as **Security** and then click **OK**.  
The **chap Properties** window displayed.
  - d) Enter the Description as "FOS chap group" and click **Apply**.
9. **Define a CHAP policy for CHAP group users.** To define the CHAP policy, follow these steps:
  - a) Open the **Administrative tools > Network policy server > Policies > Network Policies** from the menu bar.
  - b) Right-click and select **New**.  
The **New Network Policy** dialog is displayed.
  - c) Enter the policy name as *swchap*. Select the Type of network access server as **Unspecified** and click **Next**.
  - d) Click **Add** and then select the **Windows Groups** from the list.  
The **Select Group** dialog is displayed.
  - e) Enter the object name to select as *chap* and then click **OK**. Click **Next**.  
The **Specify Access Permission** window is displayed.
  - f) Select the **Access granted** option and click **Next**.  
The **Configure Authentication Methods** window is displayed.
  - g) Select the **Encrypted authentication (CHAP)** check box and click **Next**.  
The **Configure Settings** window is displayed.
  - h) Select the **Vendor Specific** option from the left pane and click **Add**.  
The **Add Vendor Specific Attribute** dialog is displayed.
  - i) Select **RADIUS Standard** from the list and click **Add**.  
The **Vendor Specific Attribute Information** dialog is displayed.
  - j) Select **RADIUS Standard** from the list and choose **Yes**, it confirms option. Click **OK**.  
The **Configure VSA (RFC Compliant)** dialog is displayed.
  - k) Select the **Vendor assigned attribute number**, **Attribute format** as *String*, and **Attribute value** as *admin* and then click **OK**.
  - l) Click **Finish**.
10. **Add new Radius user with CHAP authentication.** To add a RADIUS user, follow these steps:
  - a) In **Server Manager**, open **AD DS** tab and then open **Active Directory User and Computers**.
  - b) Right-click on the left panel tab and select **New > User** from the menu.  
The **New Object - User** dialog is displayed.

- c) Enter First name and User Logon name as **swchap**, select the server from the list and click **Next**.
- d) Enter a password, confirm the password, and then click **Next**.
- e) Click **Finish**, select the added user from the list.
- f) Right-click and select **Properties**.  
The **swchap Properties** dialog is displayed.
- g) Select the **Account** tab and complete the details.
- h) Select the **Member** of tab, click **Add**, and then add the user Role as *admin*. Apply changes and click **OK**.

**11. Add new Radius user with PAP authentication.** To add a RADIUS user, follow these steps:

- a) In **Server Manager**, open **AD DS** tab and then open **Active Directory User and Computers**.
- b) Right-click on the left panel tab and select **New > User** from the menu.  
The **New Object - User** dialog is displayed.
- c) Enter First name and User Logon name as **swpap**, select the server from the list and click **Next**.
- d) Enter a password, confirm the password, and then click **Next**.
- e) Click **Finish**, select the added user from the list.
- f) Right-click and select **Properties**.  
The **swpap Properties** dialog is displayed.
- g) Select the **Account** tab and complete the details.
- h) Select the **Member** of tab, click **Add**, Add the user Role as *admin* and *PAP group member*. Apply changes and click **OK**.

**12. Configure client for added user server and try to log in with added user from the client machine.**

## LDAP Configuration and Microsoft Active Directory

LDAP provides user authentication and authorization using the Microsoft Active Directory service or using OpenLDAP in conjunction with LDAP on the switch. This section describes the authentication and authorization using Microsoft Active Directory. For information about authentication and authorization using OpenLDAP, see [LDAP Configuration and OpenLDAP](#).

The following are restrictions when using LDAP:

- Migration of newly created users is not automatic from the local switch database to Active Directory. This is a manual process explained later.
- Only IPv4 is supported for LDAP on Windows 2000 and LDAP on Windows Server 2003. For LDAP on Windows Server 2008, both IPv4 and IPv6 are supported.
- LDAP authentication is used on the local switch only and not for the entire fabric.
- You can use the User-Principal-Name and not the Common-Name for AD LDAP authentication.  
To provide backward compatibility, authentication based on the Common Name is still supported for Active Directory LDAP 2000 and 2003. Authentication based on the Common Name is not recommended for new installations.
- A user can belong to multiple groups as long as one of the groups is the primary group. The primary group in the AD server should not be set to the group corresponding to the switch role. You can choose any other group.
- A user can be part of any Organizational Unit (OU).

When authentication is performed by User-Principal-Name, the suffix part of the name (the *@domain-name* part) can be omitted when the user logs in. If the suffix part of the User-Principal-Name name is omitted, the domain name configured for the LDAP server (in the `aaaConfig --add` command) is added and used for authentication purposes.

Roles for Brocade specific users can be added through the Microsoft Management Console. Groups created in Active Directory must correspond directly to the RBAC user roles on the switch. Role assignments can be achieved by including the user in the respective group. A user can be assigned to multiple groups, such as Switch Admin and Security Admin.

For LDAP servers, you can use the `ldapcfg --maprole` command to map LDAP server permissions to one of the default roles available on a switch. For more information on RBAC roles, see [Role Based Access Control](#).

**NOTE**

All instructions involving Microsoft Active Directory can be obtained from [www.microsoft.com](http://www.microsoft.com) or your Microsoft documentation. Confer with your system or network administrator prior to configuration for any special needs your network environment may have.

**NOTE**

From Fabric OS 9.0.1 onwards, the maximum number of characters used in the LDAP configuration of username and domain is increased from 40 characters to 64 characters.

## Configuring Microsoft Active Directory LDAP Service

Perform the following steps to set up an LDAP service:

1. If your Windows Active Directory server for LDAP is verified by the LDAP client (that is, the Brocade switch), then you must install a certificate authority (CA) root certificate on the Windows Active Directory server for LDAP.

Follow Microsoft instructions for generating and installing CA certificates on a Windows server.

2. Create a user in the Microsoft Active Directory server.

From Fabric OS 9.1.0 release, the underscore "\_" is the only special character that is allowed at the start of the username, whenever you are creating a user on the AAA server.

3. Associate the user to the group by adding the user to the group.
4. Use the `ldapCfg --maprole, ldap_role_name, switch_role, and --mapattr` commands to map an LDAP group (including a Nested Group) to one of the available switch roles.
5. Add the user's Virtual Fabrics to the `CN_list` by either editing the `adminDescription` value or adding the `brcdAdVfData` attribute to the existing Active Directory schema.


This action maps the Virtual Fabrics to the user name. Virtual Fabrics are added as a string value separate by a comma (,) and are entered as a range.




## Adding a Brocade VSA in the Windows AD Schema

Perform the following steps to add a Brocade vendor-specific attribute (VSA) in the Windows Active Directory (AD) schema:


1. Press **Windows+R** on your computer to open the **Run** box. Type **cmd** and press **Ctrl+Shift+Enter** to open an administrator command prompt.




2. In the command prompt, enter the `regsvr32 schmmgmt.dll` command.
3. A message similar to the following appears upon the successful addition of the .dll file.




4. To add a snap-in, press **Windows+R** on your computer to open the **Run** box. Type **mmc** and press **Enter**. The microsoft management console opens.



5. Select **File > Add/Remove Snap-In**. Select **Active Directory Schema**, and then select **Add**.



6. Select **OK**. The **Active Directory Schema** option is now available. You are now ready to modify the schema.
7. To create a new attribute, right-click on **Attributes** in the console tree and select **Create New Attribute** from the resulting context menu. A warning message will indicate that you are about to make a permanent change. The new attribute that you are creating can later be disabled, but it cannot be erased. Complete the information for the attribute as appropriate, for example.



**Create New Attribute**

Create a New Attribute Object

Identification

Common Name: Brocade-AD-VF-Data

LDAP Display Name: brocadeAdVFData

Unique X500 Object ID: 1.3.6.1.4.1.1588.6.1

Description: Brocade-AD-VF-Data

Syntax and Range

Syntax: Unicode String


Minimum: 0

Maximum: 1024

☐ Multi-Valued

OK Cancel Help

8. Click **OK**.
9. Right-click on **User** under **Classes** and select **Properties**. Then go to the **Attributes** tab and add the newly created Brocade attribute to the optional list.



10. Click **Apply** and then **OK**. You can ignore any console errors that occur. Now, the `brocdAdvfData` attribute is available in the **Attribute** editor under **Users**.

## Adding a Virtual Fabric List

1. From the Windows Start menu, select **Programs > Administrative Tools > ADSI.msc**.

ADSI is a Microsoft Windows Resource Utility. This utility must be installed to proceed with the rest of the setup. For Windows 2003, this utility comes with Service Pack 1 or you can download this utility from the Microsoft website.

2. Go to **CN=Users**.
3. Select **Properties**. Click the **Attribute Editor** tab.
4. Double-click the **adminDescription** or **brocdAdvfData** attribute.

The String Attribute Editor dialog is displayed.

### NOTE

The attribute can be added to user objects only.

5. Enter the values of the logical fabrics separated by a semi-colon ( ; ) into the **Value** field.

Example for adding Virtual Fabrics:

HomeLF=10;LFRoleList=admin:128,10;ChassisRole=admin

In this example, the logical switch that would be logged in to by default is 10. If 10 is not available, the lowest fabric ID (FID) available will be chosen. You would have permission to enter logical switch 128 and 10 in an admin role, and you would also have the chassis role permission of admin.

### NOTE

You can perform batch operations using the `Ldifde.exe` utility. For more information on importing and exporting schemas, refer to your Microsoft documentation or visit [www.microsoft.com](http://www.microsoft.com).

## LDAP Configuration and OpenLDAP

Fabric OS provides user authentication and authorization by means of OpenLDAP or the Microsoft Active Directory service with LDAP on the switch. This section describes authentication and authorization using OpenLDAP. For information about authentication and authorization using Microsoft Active Directory, see [LDAP Configuration and Microsoft Active Directory](#).

When using OpenLDAP, you must use the Common-Name for OpenLDAP authentication. User-Principal-Name is not supported in OpenLDAP. The v2.4.23 is supported.

When a user is authenticated, the role of the user is obtained from the `memberOf` attribute, which determines group membership. This feature is supported in OpenLDAP through the `memberOf` overlay. You must use this overlay on the OpenLDAP server to assign membership information.

For OpenLDAP servers, you can use the `ldapCfg --maprole` command to map LDAP server permissions to one of the default roles available on a switch. For more information on the RBAC roles, see [Role-Based Access Control](#).

### NOTE

Windows AD specific features such as Global Catalog port queries and `sAMAccountName`-based user logins are not applicable for OpenLDAP.

## Supported LDAP Options

The following table summarizes the various LDAP options and Brocade support for each.

**Table 156: LDAP Options**

| Protocol                        | Description                                                                                      | Channel Type | Default Port                    | URL                   | Brocade Supported?                 |
|---------------------------------|--------------------------------------------------------------------------------------------------|--------------|---------------------------------|-----------------------|------------------------------------|
| LDAPv3                          | LDAP over TCP                                                                                    | Unsecured    | 389                             | <code>ldap://</code>  | No                                 |
| LDAPv3 with TLS extension       | LDAPv3 over TLS                                                                                  | Secured      | 389                             | <code>ldap://</code>  | Yes (using <code>starttls</code> ) |
| LDAPv3 with TLS and Certificate | LDAPv3 over TLS channel and authenticated using a certificate                                    | Secured      | 389                             | <code>ldap://</code>  | Yes (using <code>starttls</code> ) |
| LDAPv3 with SSL                 | LDAPv3 over SSL. Port 636 is the default LDAPS port. Port 3269 is the LDAPS Global Catalog port. | Secured      | 636 (LDAPS) and 3269 (LDAPS GC) | <code>ldaps://</code> | Yes                                |

## OpenLDAP Server Configuration Overview

For complete details about how to install and configure an OpenLDAP server, refer to the OpenLDAP user documentation at <http://www.openldap.org/doc/>. A few key steps for the Brocade environment are outlined here.

1. If your OpenLDAP server is verified by the LDAP client (that is, the Brocade switch), you must install a certificate authority (CA) root certificate on the OpenLDAP server.

Follow OpenLDAP instructions for generating and installing CA certificates on an OpenLDAP server.

2. Enable group membership through the `memberOf` mechanism by including the `memberOf` overlay in the `slapd.conf` file.
3. Create users (entries) in the OpenLDAP directory.  
From Fabric OS 9.1.0 release, the underscore "\_" is the only special character that is allowed at the start of the username, whenever you are creating a user on the AAA server.
4. Assign users to groups by using the `memberOf` attribute.
5. Use the `ldapcfg --maprole` command to map an LDAP server role to one of the roles available on the switch.
6. Add the user's Virtual Fabrics to the user entry.
  - a) Add the `brcdAdvfData` attribute to the existing OpenLDAP schema.
  - b) Add the `brcdAdvfData` attribute to the user entry in the LDAP directory with a value that identifies the Virtual Fabrics with which to associate the user.

### NOTE

To authenticate a user in a switch, create users in OpenLDAP under the **cn=users** group.

## Enabling Group Membership

Group membership in OpenLDAP is specified by an overlay called `memberOf`. Overlays are helpful in customizing the back-end behavior without requiring changes to the back-end code. The `memberOf` overlay updates the `memberOf` attribute whenever changes occur to the membership attribute of the entries of the `groupOfNames` object class. To include this overlay, add `overlay memberOf` to the `slapd.conf` file, as shown in the following example.

```
overlay memberof
Example file:
include /usr/local/etc/openldap/schema/core.schema
include /usr/local/etc/openldap/schema/cosine.schema
include /usr/local/etc/openldap/schema/local.schema
#####
TLSCACertificateFile /root/jsmith/ldapcert/cacert.pem
TLSCertificateFile /root/jsmith/ldapcert/serverCert.pem
TLSCertificateKeyFile /root/jsmith/ldapcert/serverKey.pem
TLSVerifyClient never
pidfile /usr/local/var/run/slapd.pid
argsfile /usr/local/var/run/slapd.args
database bdb
suffix "dc=mycompany,dc=com"
rootdn "cn=Manager,dc=mycompany,dc=com"
rootpw {SSHA}4L8uT5hPa44IdcP6tAheMT8noGoWubr3
The database directory MUST exist prior to running slapd AND
should only be accessible by the slapd and slap tools.
Mode 700 recommended.
directory /usr/local/var/openldap-data
Indices to maintain
index objectClass eq
overlay memberof
```

## Adding Entries to the Directory

To add entries in the OpenLDAP directory, perform the following steps.

1. Using a text editor of your choice, create a `.ldif` file and enter the information for the entry.

The following example defines an organizational role for the Directory Manager in a `.ldif` file for an organization with the domain name `mybrocade.com`.

```
Organization for mybrocade Corporation
dn: dc=mybrocade,dc=com
objectClass: dcObject
objectClass: organization
dc: mybrocade
o: Mybrocade Corporation
description: Mybrocade Corporation
#####
Organizational Role for Directory Manager
dn: cn=Manager,dc=mybrocade,dc=com
objectClass: organizationalRole
cn: Manager
description: Directory Manager
```

2. Enter the `ldapadd` command to add the contents of the `.ldif` file to the Directory, where `test.ldif` is the file you created in the Step 1.

```
switch:admin> ldapadd -D cn=Manager,dc=mybrocade,dc=com -x -w secret -f test.ldif
```

## Assigning a User to a Group

Before you can assign a user to a group, the `memberOf` overlay must be added to the `slapd.conf` file. Refer to [Enabling Group Membership](#) for details.

1. In a `.ldif` file, create a "groupOfNames" objectClass entry with the name of the group, for example, `admin` to create a group.
2. Set a **member** attribute for the group instance to identify the member, as in this example:

```
cn=Sachin,cn=Users,dc=mybrocade,dc=com
```

Automatically, the **memberOf** attribute of the entry Sachin will have the value `cn=admin,ou=groups,dc=mybrocade,dc=com`, which assigns Sachin to the admin group.

3. Enter the `ldapadd` command.

For example, the `.ldif` file might contain information similar to the following:

```
#Groups in organization
dn: ou=groups,dc=mybrocade,dc=com
objectclass:organizationalunit
ou: groups
description: generic groups branch
dn: cn=admin,ou=groups,dc=mybrocade,dc=com
objectclass: groupofnames
cn: admin
description: Members having admin permission
#Add members for admin group
```

```
member: cn=sachin,cn=Users,dc=mybrocade,dc=com
```

## Modifying an Entry

To modify a directory entry, perform the following steps.

1. Create a `.ldif` file containing the information to be modified.
2. Enter the `ldapmodify` command with the `-f` option specifying the `.ldif` file you created in Step 1.

## TACACS+ Service

Fabric OS can authenticate users with a remote server using the Terminal Access Controller Access-Control System Plus (TACACS+) protocol. TACACS+ is a protocol used in AAA server environments consisting of a centralized authentication server and multiple Network Access Servers (NASs) or clients. When configured to use TACACS+, a Brocade switch becomes an NAS.

The following authentication protocols are supported by the TACACS+ server for user authentication:

- Password Authentication Protocol (PAP)
- Challenge Handshake Authentication Protocol (CHAP)

The TACACS+ server can be a Microsoft Windows server or a Linux server. For Linux servers, use TACACS+ 4.0.4 or later from Cisco. For Microsoft Windows servers, use any TACACS+ freeware that uses TACACS+ protocol v1.78 or later.

## TACACS+ Configuration Overview

Configuration is required on both the TACACS+ server and the Brocade switch. On the TACACS+ server, you should assign a role for each user and should provide lists of Virtual Fabrics to which the user should have the access. For details, see [TACACS+ Server Data Configuration](#).

### NOTE

From Fabric OS 9.1.0 release, the underscore "\_" is the only special character that is allowed at the start of the username, whenever you are creating a user on the AAA server.

On the Brocade switch, use the `aaaConfig` command to configure the switch to use TACACS+ for authentication. The `aaaConfig` command allows you to specify up to five TACACS+ servers. When a list of servers is configured, failover from one server to another server occurs only if a TACACS+ server fails to respond. It does not occur when user authentication fails.

Failover to another TACACS+ server is achieved through a timeout. You can configure a timeout value for each TACACS+ server so that the next server can be used if the first server is unreachable. The default timeout value is 5 seconds.

Retry, the number of attempts to authenticate with a TACACS+ server, is also allowed. The default value is three attempts. If the server fails to respond or times out, Fabric OS tries again up to the number of specified retries.

See [Remote Authentication Configuration on the Switch](#) for details about configuring the Brocade switch for authenticating users with a TACACS+ server.

## Configuring the TACACS+ Server on Linux

Fabric OS software supports TACACS+ authentication on a Linux server running the Open Source TACACS+ LINUX package v4.0.4 from Cisco. To install and configure this software, perform the following steps.

1. Download the TACACS+ software from <http://www.cisco.com> and install it.
2. Configure the TACACS+ server by editing the `tac_plus.cfg` file.

Refer to [TACACS+ Server Data Configuration](#) for details.



3. Run the `tac_plus` daemon to start and enable the TACACS+ service on the server.

```
switch:admin> tac_plus -d 16 /usr/local/etc/mavis/sample/tac_plus.cfg
```

## Configuring a Windows TACACS+ Server

Fabric OS is compatible with any TACACS+ freeware for Microsoft Windows that uses TACACS+ protocol version v1.78. Refer to the vendor documentation for configuration details.

## Appendix B: Configuring Factory Reset

---

The factory reset feature removes the entire custom configuration data from the device and brings the device to its original manufacturer default settings. Factory Reset avoids installing firmware packages from remote servers and using multiple commands to remove configurations from your device. You can perform a factory reset under the following scenarios:

- Data corruption
- Critical conditions that are irrevocable
- Bringing back the device to the state as it was at the time of shipping

### Supported Switches and Directors

You can perform a factory reset on the following switches and directors with Fabric OS 9.0:

- Brocade Gen 6 switches and directors
- Brocade Gen 6 chassis directors upgraded to Gen 7
- Brocade Gen 6 Blade Server SAN I/O Modules
- Brocade Access Gateway (AG) enabled platforms
- Brocade Gen 7 switches and directors

For more information on the supported devices, see [Supported Hardware and Software](#).

When the `factoryreset` command is applied on AG-capable platforms, the AG mode is converted to the native switch mode except the Brocade MXG610 Blade Server SAN I/O Modules.

### FOS Configurations That Support Fabric Reset

A factory reset removes the following configurations on a switch or a director:

- End-user license agreement (EULA) records
- FC protocol and FCOE related configurations
- bootenv configurations
- User-configurable IP configurations in EPROM
- User accounts other than the default accounts and reset passwords of default accounts
- RASLogs, trace dump, audit logs, and all other logs in the switch
- Extension blade related configurations
- Logical switches

### Resetting a Device to Factory Default Settings

Perform the following steps to reset a switch or director to the factory default settings:

1. Log in to the switch using an account that has admin permissions.
2. Ensure that the custom configuration data is uploaded to a file using the `configupload` command. You can back up the virtual fabric and system configuration-specific files. For more information, refer to the *Brocade Fabric OS Command Reference Manual*.
3. The `factoryreset` command generates an error message if invoked while the switch or chassis is enabled. Enter `switchdisable` to disable switches or `chassisdisable` to disable the virtual-fabric-aware chassis, if you perform the factory reset on a chassis. Enter **Y** in the prompt with a message to disable the chassis as shown in the following example:

```
switch:admin> chassisdisable
```

This command can cause disruption to multiple logical switches.

Are you sure you want to disable all chassis ports now? (yes, y, no, n): [no]: **y**

4. Enter `factoryreset` to start the factory reset. This command displays the following warning message. Enter **y** to continue the factory reset.

#### NOTE

This command should be used with caution on your devices because it can remove the entire configuration data and reset it to the factory default. Once the `factoryreset` command is issued, you cannot revoke it.

```
switch:admin> factoryreset
```

```
This operation will reset all switch configurations to manufacture default,
all customized configurations will be lost.
```

```
Chassis/switch would be rebooted to clear the config.
```

```
Do you want to continue [y/n]?: y
```

```
Do not power off the switch, factory reset in progress...
```

```
...
```

```
...
```

```
Terminating all sessions except factory-reset session
```

```
Factoryreset clear EULA record started...done
```

```
Deleting fips configurations....
```

```
Executing 'secauthsecret --remove -all' for FID 128
```

```
Executing 'secauthsecret --remove -all' for FID 1
```

```
Executing 'secauthsecret --remove -all' for FID 3
```

```
Zeroizing certificates
```

```
Removing all FCAP Keys/Certificates
```

```
Removing all Radius Keys/Certificates
```

```
Removing all LDAP Keys/Certificates
```

```
-----Output truncated-----
```

5. Enter `factoryreset -force` to enter noninteractive mode and start the factory reset.

```
switch:admin> factoryreset -force
```

```
Do not power off the switch, factory reset in progress...
```

```
Removing config database... for fid 128.
```

```
done
```

```
Removing FAWWN data ... for fid 128.
```

```
done
```

```


```

```
Deleting cli and firmware history....done
```

```
Resetting bootenv....
```

```
Factory reset completed successfully!
```

#### NOTE

Do not power off the switch or terminate the session abruptly while the factory reset is in progress. This may cause a failure of the factory reset and put the switch into a unspecified state. If there is an unexpected reboot while the factory reset in progress, review the session output and run the `factoryreset` command again to complete the process.

## 6. Enter `factoryreset -keep ip` to preserve the management interface configurations from the factory reset operations using the interactive mode.

```
switch:admin> factoryreset -keep ip
This operation will reset all switch configurations to manufacture default,
all customized configurations will be lost.
Chassis/switch would be rebooted to clear the config.
Do you want to continue [y/n]?: y
Do not power off the switch, factory reset in progress...
...
...
```

```
Terminating all sessions except factory-reset session
```

```
Deleting fips configurations....
Executing 'secauthsecret --remove -all' for FID 128

Executing 'secauthsecret --remove -all' for FID 10
Zeroizing certificates
Removing all FCAP Keys/Certificates
Removing all Radius Keys/Certificates
Removing all LDAP Keys/Certificates
Removing all Syslog-ng Keys/Certificates
Removing all HTTPS Keys/Certificates
Removing all IP-extension Keys/Certificates
Removing all Kafka Keys/Certificates
Removing all ASC Keys/Certificates
Executing 'passwddefault':
Zeroizing Radius configuration:
Authentication mode already set to local switch database.
RADIUS configuration does not exist.
LDAP configuration does not exist.
TACACS+ configuration does not exist.
Zeroizing IPsec static SA configuration.
Zeroizing SSH key.
private key doesn't exist.
Zeroizing SSH Known Hosts.
Deleting LDAP Map attributes
Resetting auth util data
Removing log files
Deleting crypto templates
Applying default crypto configuration values
Deleting FIPS config files
Disabling Idle session timeout
Zeroize core files in Local.
Zeroize core files in standby.
done
Clearing trace buffer....done
Clearing raslog Audit log cleared.
done
Deleting maps configurations....done
Deleting dns configurations....done
```

```
Resetting switch name, host name and chassis name....done
Resetting diagnostic configurations....done
Resetting BSN configurations....
Removed custom support link configuration.
done
Resetting Banner configurations....done
Resetting motd configurations....done
Resetting security fdd configurations....No Change to distribution parameter(s).
Local Switch Configured to accept policies.
Resetting security fdd config keys....done
Resetting to Native mode if AG mode is enableddone
Deleting VF configurations and switch configurations....done
Deleting cli and firmware history....done
Deleting all license configurations....done
Resetting bootenv....done
Factoryreset starting on Standby CP....
Deleting VF configurations and switch configurations....done
Deleting cli and firmware history....done
Deleting all license configurations....done
Resetting standby bootenv....done
Factoryreset operation is completed successfully.
2020/03/30-07:25:57, [CONF-1033], 9, SLOT 2 | CHASSIS, INFO, sw0, Factoryreset operation is completed
successfully.
```

**NOTE**

If IP is reset, the last set of logs will not be displayed in the `factoryreset` output.

7. Once all the custom configurations are reset to default configurations, your switch or chassis will reboot. After the reboot, the device is the same as it was at the time of shipping with the factory default settings.

## Appendix C: Supported Cipher Suites in Fabric OS

With the upgrade of OpenSSL in Fabric OS 9.0.x, Fabric OS supports additional cipher suites and withdraws supports for certain cipher suites.

The following tables list the cipher suites that are supported in Fabric OS 8.x and 9.x:

**Table 157: Brocade Cipher Suites Supported in Fabric OS 8.x**

| Supported Ciphers             | Supported Ciphers             | Supported Ciphers        |
|-------------------------------|-------------------------------|--------------------------|
| ECDHE-RSA-AES256-GCM-SHA384   | CAMELLIA256-SHA               | ECDH-RSA-AES128-SHA256   |
| ECDHE-ECDSA-AES256-GCM-SHA384 | PSK-AES256-CBC-SHA            | ECDH-ECDSA-AES128-SHA256 |
| ECDHE-RSA-AES256-SHA384       | ECDHE-RSA-AES128-GCM-SHA256   | ECDH-RSA-AES128-SHA      |
| ECDHE-ECDSA-AES256-SHA384     | ECDHE-ECDSA-AES128-GCM-SHA256 | ECDH-ECDSA-AES128-SHA    |
| ECDHE-RSA-AES256-SHA          | ECDHE-RSA-AES128-SHA256       | AES128-GCM-SHA256        |
| ECDHE-ECDSA-AES256-SHA        | ECDHE-ECDSA-AES128-SHA256     | AES128-SHA256            |
| DH-DSS-AES256-GCM-SHA384      | ECDHE-RSA-AES128-SHA          | AES128-SHA               |
| DHE-DSS-AES256-GCM-SHA384     | ECDHE-ECDSA-AES128-SHA        | SEED-SHA                 |
| DH-RSA-AES256-GCM-SHA384      | DH-DSS-AES128-GCM-SHA256      | CAMELLIA128-SHA          |
| DHE-RSA-AES256-GCM-SHA384     | DHE-DSS-AES128-GCM-SHA256     | PSK-AES128-CBC-SHA       |
| DHE-RSA-AES256-SHA256         | DH-RSA-AES128-GCM-SHA256      | ECDHE-RSA-RC4-SHA        |
| DHE-DSS-AES256-SHA256         | DHE-RSA-AES128-GCM-SHA256     | ECDHE-ECDSA-RC4-SHA      |
| DH-RSA-AES256-SHA256          | DHE-RSA-AES128-SHA256         | AECDH-RC4-SHA            |
| DH-DSS-AES256-SHA256          | DHE-DSS-AES128-SHA256         | ADH-RC4-MD5              |
| DHE-RSA-AES256-SHA            | DH-RSA-AES128-SHA256          | ECDH-RSA-RC4-SHA         |
| DHE-DSS-AES256-SHA            | DH-DSS-AES128-SHA256          | ECDH-ECDSA-RC4-SHA       |
| DH-RSA-AES256-SHA             | DHE-RSA-AES128-SHA            | RC4-SHA                  |
| DH-DSS-AES256-SHA             | DHE-DSS-AES128-SHA            | RC4-MD5                  |
| DHE-RSA-CAMELLIA256-SHA       | DH-RSA-AES128-SHA             | PSK-RC4-SHA              |
| DHE-DSS-CAMELLIA256-SHA       | DH-DSS-AES128-SHA             | ECDHE-RSA-DES-CBC3-SHA   |
| DH-RSA-CAMELLIA256-SHA        | DHE-RSA-SEED-SHA              | ECDHE-ECDSA-DES-CBC3-SHA |
| DH-DSS-CAMELLIA256-SHA        | DHE-DSS-SEED-SHA              | EDH-RSA-DES-CBC3-SHA     |
| AECDH-AES256-SHA              | DH-RSA-SEED-SHA               | EDH-DSS-DES-CBC3-SHA     |
| ADH-AES256-GCM-SHA384         | DH-DSS-SEED-SHA               | DH-RSA-DES-CBC3-SHA      |
| ADH-AES256-SHA256             | DHE-RSA-CAMELLIA128-SHA       | DH-DSS-DES-CBC3-SHA      |
| ADH-AES256-SHA                | DHE-DSS-CAMELLIA128-SHA       | AECDH-DES-CBC3-SHA       |
| ADH-CAMELLIA256-SHA           | DH-RSA-CAMELLIA128-SHA        | ADH-DES-CBC3-SHA         |
| ECDH-RSA-AES256-GCM-SHA384    | DH-DSS-CAMELLIA128-SHA        | ECDH-RSA-DES-CBC3-SHA    |
| ECDH-ECDSA-AES256-GCM-SHA384  | AECDH-AES128-SHA              | ECDH-ECDSA-DES-CBC3-SHA  |
| ECDH-RSA-AES256-SHA384        | ADH-AES128-GCM-SHA256         | DES-CBC3-SHA             |
| ECDH-ECDSA-AES256-SHA384      | ADH-AES128-SHA256             | PSK-3DES-EDE-CBC-SHA     |

| Supported Ciphers   | Supported Ciphers            | Supported Ciphers     |
|---------------------|------------------------------|-----------------------|
| ECDH-RSA-AES256-SHA | ADH-AES128-SHA               | ECDH-ECDSA-AES256-SHA |
| ADH-SEED-SHA        | AES256-GCM-SHA384            | ADH-CAMELLIA128-SHA   |
| AES256-SHA256       | ECDH-RSA-AES128-GCM-SHA256   | —                     |
| AES256-SHA          | ECDH-ECDSA-AES128-GCM-SHA256 | —                     |

**Table 158: Brocade Cipher Suites Supported in Fabric OS 9.x**

| Supported Ciphers              | Supported Ciphers              | Supported Ciphers           |
|--------------------------------|--------------------------------|-----------------------------|
| TLS_AES_256_GCM_SHA384         | ECDHE-ECDSA-AES256-SHA384      | ECDHE-RSA-RC4-SHA           |
| TLS_CHACHA20_POLY1305_SHA256   | ECDHE-RSA-AES256-SHA384        | ECDHE-ECDSA-DES-CBC3-SHA    |
| TLS_AES_128_GCM_SHA256         | DHE-RSA-AES256-SHA256          | ECDHE-RSA-DES-CBC3-SHA      |
| ECDHE-ECDSA-AES256-GCM-SHA384  | DHE-DSS-AES256-SHA256          | DHE-RSA-DES-CBC3-SHA        |
| ECDHE-RSA-AES256-GCM-SHA384    | ECDHE-ECDSA-CAMELLIA256-SHA384 | DHE-DSS-DES-CBC3-SHA        |
| DHE-DSS-AES256-GCM-SHA384      | ECDHE-RSA-CAMELLIA256-SHA384   | RSA-PSK-AES256-GCM-SHA384   |
| DHE-RSA-AES256-GCM-SHA384      | DHE-RSA-CAMELLIA256-SHA256     | DHE-PSK-AES256-GCM-SHA384   |
| ECDHE-ECDSA-CHACHA20-POLY1305  | DHE-DSS-CAMELLIA256-SHA256     | RSA-PSK-CHACHA20-POLY1305   |
| ECDHE-RSA-CHACHA20-POLY1305    | ECDHE-ECDSA-AES128-SHA256      | DHE-PSK-CHACHA20-POLY1305   |
| DHE-RSA-CHACHA20-POLY1305      | ECDHE-RSA-AES128-SHA256        | ECDHE-PSK-CHACHA20-POLY1305 |
| ECDHE-ECDSA-AES256-CCM8        | DHE-RSA-AES128-SHA256          | DHE-PSK-AES256-CCM8         |
| ECDHE-ECDSA-AES256-CCM         | DHE-DSS-AES128-SHA256          | DHE-PSK-AES256-CCM          |
| DHE-RSA-AES256-CCM8            | ECDHE-ECDSA-CAMELLIA128-SHA256 | RSA-PSK-ARIA256-GCM-SHA384  |
| DHE-RSA-AES256-CCM             | ECDHE-RSA-CAMELLIA128-SHA256   | DHE-PSK-ARIA256-GCM-SHA384  |
| ECDHE-ECDSA-ARIA256-GCM-SHA384 | DHE-RSA-CAMELLIA128-SHA256     | AES256-GCM-SHA384           |
| ECDHE-ARIA256-GCM-SHA384       | DHE-DSS-CAMELLIA128-SHA256     | AES256-CCM8                 |
| DHE-DSS-ARIA256-GCM-SHA384     | ECDHE-ECDSA-AES256-SHA         | AES256-CCM                  |
| DHE-RSA-ARIA256-GCM-SHA384     | ECDHE-RSA-AES256-SHA           | ARIA256-GCM-SHA384          |
| ECDHE-ECDSA-AES128-GCM-SHA256  | DHE-RSA-AES256-SHA             | PSK-AES256-GCM-SHA384       |
| ECDHE-RSA-AES128-GCM-SHA256    | DHE-DSS-AES256-SHA             | PSK-CHACHA20-POLY1305       |
| DHE-DSS-AES128-GCM-SHA256      | DHE-RSA-CAMELLIA256-SHA        | PSK-AES256-CCM8             |
| DHE-RSA-AES128-GCM-SHA256      | DHE-DSS-CAMELLIA256-SHA        | PSK-AES256-CCM              |
| ECDHE-ECDSA-AES128-CCM8        | ECDHE-ECDSA-AES128-SHA         | PSK-ARIA256-GCM-SHA384      |
| ECDHE-ECDSA-AES128-CCM         | ECDHE-RSA-AES128-SHA           | RSA-PSK-AES128-GCM-SHA256   |
| DHE-RSA-AES128-CCM8            | DHE-RSA-AES128-SHA             | DHE-PSK-AES128-GCM-SHA256   |
| DHE-RSA-AES128-CCM             | DHE-DSS-AES128-SHA             | DHE-PSK-AES128-CCM8         |
| ECDHE-ECDSA-ARIA128-GCM-SHA256 | DHE-RSA-SEED-SHA               | DHE-PSK-AES128-CCM          |
| ECDHE-ARIA128-GCM-SHA256       | DHE-DSS-SEED-SHA               | RSA-PSK-ARIA128-GCM-SHA256  |
| DHE-DSS-ARIA128-GCM-SHA256     | DHE-RSA-CAMELLIA128-SHA        | DHE-PSK-ARIA128-GCM-SHA256  |
| DHE-RSA-ARIA128-GCM-SHA256     | DHE-DSS-CAMELLIA128-SHA        | AES128-GCM-SHA256           |
| AES128-CCM                     | ECDHE-ECDSA-RC4-SHA            | PSK-AES128-CBC-SHA256       |
| ARIA128-GCM-SHA256             | CAMELLIA256-SHA                | PSK-AES128-CBC-SHA          |

| Supported Ciphers            | Supported Ciphers            | Supported Ciphers          |
|------------------------------|------------------------------|----------------------------|
| PSK-AES128-GCM-SHA256        | PSK-AES256-CBC-SHA384        | PSK-CAMELLIA128-SHA256     |
| PSK-AES128-CCM8              | PSK-AES256-CBC-SHA           | ECDHE-PSK-RC4-SHA          |
| PSK-AES128-CCM               | PSK-CAMELLIA256-SHA384       | RSA-PSK-RC4-SHA            |
| PSK-ARIA128-GCM-SHA256       | ECDHE-PSK-AES128-CBC-SHA256  | DHE-PSK-RC4-SHA            |
| AES256-SHA256                | ECDHE-PSK-AES128-CBC-SHA     | RC4-SHA                    |
| CAMELLIA256-SHA256           | RSA-PSK-AES128-CBC-SHA256    | RC4-MD5                    |
| AES128-SHA256                | DHE-PSK-AES128-CBC-SHA256    | PSK-RC4-SHA                |
| CAMELLIA128-SHA256           | RSA-PSK-AES128-CBC-SHA       | ECDHE-PSK-3DES-EDE-CBC-SHA |
| ECDHE-PSK-AES256-CBC-SHA384  | DHE-PSK-AES128-CBC-SHA       | RSA-PSK-3DES-EDE-CBC-SHA   |
| ECDHE-PSK-AES256-CBC-SHA     | ECDHE-PSK-CAMELLIA128-SHA256 | DHE-PSK-3DES-EDE-CBC-SHA   |
| RSA-PSK-AES256-CBC-SHA384    | RSA-PSK-CAMELLIA128-SHA256   | DES-CBC3-SHA               |
| DHE-PSK-AES256-CBC-SHA384    | DHE-PSK-CAMELLIA128-SHA256   | PSK-3DES-EDE-CBC-SHA       |
| RSA-PSK-AES256-CBC-SHA       | AES128-SHA                   | RSA-PSK-CAMELLIA256-SHA384 |
| DHE-PSK-AES256-CBC-SHA       | SEED-SHA                     | AES256-SHA                 |
| ECDHE-PSK-CAMELLIA256-SHA384 | CAMELLIA128-SHA              | AES128-CCM8                |
| DHE-PSK-CAMELLIA256-SHA384   | —                            | —                          |

The following table lists the KEX suites that are supported in Fabric OS 8.x and 9.x:

**Table 159: Brocade KEX Suites Supported in Fabric OS 8.x and 9.0.x**

| Fabric OS 8.x KEX                    | Fabric OS 9.0.x KEX                  |
|--------------------------------------|--------------------------------------|
| ecdh-sha2-nistp256                   | ecdh-sha2-nistp256                   |
| ecdh-sha2-nistp384                   | ecdh-sha2-nistp384                   |
| ecdh-sha2-nistp521                   | ecdh-sha2-nistp521                   |
| diffie-hellman-group-exchange-sha256 | diffie-hellman-group-exchange-sha256 |
| diffie-hellman-group-exchange-sha1   | diffie-hellman-group-exchange-sha1   |
| diffie-hellman-group14-sha1          | diffie-hellman-group14-sha1          |
| diffie-hellman-group1-sha1           | diffie-hellman-group1-sha1           |
| —                                    | curve25519-sha256                    |
| —                                    | diffie-hellman-group16-sha512        |
| —                                    | diffie-hellman-group18-sha512        |
| —                                    | diffie-hellman-group14-sha256        |

The following table lists the supported ciphers that block the downgrade of Fabric OS 8.x and the upgrade of Fabric OS 9.x if configured:

| Supported Ciphers in Fabric OS 9.0.x That Block the Downgrade to Fabric OS 8.x | Supported Ciphers in Fabric OS 8.x That Block the Upgrade to Fabric OS 9.0.x |
|--------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| ECDHE-ECDSA-CHACHA20-POLY1305                                                  | DH-DSS-AES256-GCM-SHA384                                                     |
| ECDHE-RSA-CHACHA20-POLY1305                                                    | DH-RSA-AES256-GCM-SHA384                                                     |
| DHE-RSA-CHACHA20-POLY1305                                                      | DH-RSA-AES256-SHA256                                                         |



| Supported Ciphers in Fabric OS 9.0.x That Block the Downgrade to Fabric OS 8.x | Supported Ciphers in Fabric OS 8.x That Block the Upgrade to Fabric OS 9.0.x |
|--------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| ECDHE-ECDSA-AES256-CCM8                                                        | DH-DSS-AES256-SHA256                                                         |
| ECDHE-ECDSA-AES256-CCM                                                         | DH-RSA-AES256-SHA                                                            |
| DHE-RSA-AES256-CCM8                                                            | DH-DSS-AES256-SHA                                                            |
| DHE-RSA-AES256-CCM                                                             | DH-RSA-CAMELLIA256-SHA                                                       |
| ECDHE-ECDSA-ARIA256-GCM-SHA384                                                 | DH-DSS-CAMELLIA256-SHA                                                       |
| ECDHE-ARIA256-GCM-SHA384                                                       | AECDH-AES256-SHA                                                             |
| DHE-DSS-ARIA256-GCM-SHA384                                                     | ADH-AES256-GCM-SHA384                                                        |
| DHE-RSA-ARIA256-GCM-SHA384                                                     | ADH-AES256-SHA256                                                            |
| ECDHE-ECDSA-AES128-CCM8                                                        | ADH-AES256-SHA                                                               |
| ECDHE-ECDSA-AES128-CCM                                                         | ADH-CAMELLIA256-SHA                                                          |
| DHE-RSA-AES128-CCM8                                                            | ECDH-RSA-AES256-GCM-SHA384                                                   |
| DHE-RSA-AES128-CCM                                                             | ECDH-ECDSA-AES256-GCM-SHA384                                                 |
| ECDHE-ECDSA-ARIA128-GCM-SHA256                                                 | ECDH-RSA-AES256-SHA384                                                       |
| ECDHE-ARIA128-GCM-SHA256                                                       | ECDH-ECDSA-AES256-SHA384                                                     |
| DHE-DSS-ARIA128-GCM-SHA256                                                     | ECDH-RSA-AES256-SHA                                                          |
| DHE-RSA-ARIA128-GCM-SHA256                                                     | ECDH-ECDSA-AES256-SHA                                                        |
| ECDHE-ECDSA-CAMELLIA256-SHA384                                                 | DH-DSS-AES128-GCM-SHA256                                                     |
| ECDHE-RSA-CAMELLIA256-SHA384                                                   | DH-RSA-AES128-GCM-SHA256                                                     |
| DHE-RSA-CAMELLIA256-SHA256                                                     | DH-RSA-AES128-SHA256                                                         |
| DHE-DSS-CAMELLIA256-SHA256                                                     | DH-DSS-AES128-SHA256                                                         |
| ECDHE-ECDSA-CAMELLIA128-SHA256                                                 | DH-RSA-AES128-SHA                                                            |
| ECDHE-RSA-CAMELLIA128-SHA256                                                   | DH-DSS-AES128-SHA                                                            |
| DHE-RSA-CAMELLIA128-SHA256                                                     | DH-RSA-SEED-SHA                                                              |
| DHE-DSS-CAMELLIA128-SHA256                                                     | DH-DSS-SEED-SHA                                                              |
| RSA-PSK-AES256-GCM-SHA384                                                      | DH-RSA-CAMELLIA128-SHA                                                       |
| DHE-PSK-AES256-GCM-SHA384                                                      | DH-DSS-CAMELLIA128-SHA                                                       |
| RSA-PSK-CHACHA20-POLY1305                                                      | AECDH-AES128-SHA                                                             |
| DHE-PSK-CHACHA20-POLY1305                                                      | ADH-AES128-GCM-SHA256                                                        |
| ECDHE-PSK-CHACHA20-POLY1305                                                    | ADH-AES128-SHA256                                                            |
| DHE-PSK-AES256-CCM8                                                            | ADH-AES128-SHA                                                               |
| DHE-PSK-AES256-CCM                                                             | ADH-SEED-SHA                                                                 |
| RSA-PSK-ARIA256-GCM-SHA384                                                     | ADH-CAMELLIA128-SHA                                                          |
| DHE-PSK-ARIA256-GCM-SHA384                                                     | ECDH-RSA-AES128-GCM-SHA256                                                   |
| ES256-CCM8                                                                     | ECDH-ECDSA-AES128-GCM-SHA256                                                 |
| AES256-CCM                                                                     | ECDH-RSA-AES128-SHA256                                                       |
| ARIA256-GCM-SHA384                                                             | ECDH-ECDSA-AES128-SHA256                                                     |
| SK-AES256-GCM-SHA384                                                           | ECDH-RSA-AES128-SHA                                                          |
| PSK-CHACHA20-POLY1305                                                          | ECDH-ECDSA-AES128-SHA                                                        |
| PSK-AES256-CCM8                                                                | AECDH-RC4-SHA                                                                |
| PSK-AES256-CCM                                                                 | ADH-RC4-MD5                                                                  |

| Supported Ciphers in Fabric OS 9.0.x That Block the Downgrade to Fabric OS 8.x | Supported Ciphers in Fabric OS 8.x That Block the Upgrade to Fabric OS 9.0.x |
|--------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| PSK-ARIA256-GCM-SHA384                                                         | ECDH-RSA-RC4-SHA                                                             |
| RSA-PSK-AES128-GCM-SHA256                                                      | ECDH-ECDSA-RC4-SHA                                                           |
| DHE-PSK-AES128-GCM-SHA256                                                      | DH-RSA-DES-CBC3-SHA                                                          |
| DHE-PSK-AES128-CCM8                                                            | DH-DSS-DES-CBC3-SHA                                                          |
| DHE-PSK-AES128-CCM                                                             | AECDH-DES-CBC3-SHA                                                           |
| RSA-PSK-ARIA128-GCM-SHA256                                                     | ADH-DES-CBC3-SHA                                                             |
| DHE-PSK-ARIA128-GCM-SHA256                                                     | ECDH-RSA-DES-CBC3-SHA                                                        |
| AES128-CCM8                                                                    | ECDH-ECDSA-DES-CBC3-SHA                                                      |
| AES128-CCM                                                                     | —                                                                            |
| ARIA128-GCM-SHA256                                                             | —                                                                            |
| PSK-AES128-GCM-SHA256                                                          | —                                                                            |
| PSK-AES128-CCM8                                                                | —                                                                            |
| PSK-AES128-CCM                                                                 | —                                                                            |
| PSK-ARIA128-GCM-SHA256                                                         | —                                                                            |
| CAMELLIA256-SHA256                                                             | —                                                                            |
| CAMELLIA128-SHA256                                                             | —                                                                            |
| ECDHE-PSK-AES256-CBC-SHA384                                                    | —                                                                            |
| ECDHE-PSK-AES256-CBC-SHA                                                       | —                                                                            |
| SA-PSK-AES256-CBC-SHA384                                                       | —                                                                            |
| DHE-PSK-AES256-CBC-SHA384                                                      | —                                                                            |
| RSA-PSK-AES256-CBC-SHA                                                         | —                                                                            |
| DHE-PSK-AES256-CBC-SHA                                                         | —                                                                            |
| ECDHE-PSK-CAMELLIA256-SHA384                                                   | —                                                                            |
| RSA-PSK-CAMELLIA256-SHA384                                                     | —                                                                            |
| DHE-PSK-CAMELLIA256-SHA384                                                     | —                                                                            |
| PSK-AES256-CBC-SHA384                                                          | —                                                                            |
| PSK-CAMELLIA256-SHA384                                                         | —                                                                            |
| ECDHE-PSK-AES128-CBC-SHA256                                                    | —                                                                            |
| ECDHE-PSK-AES128-CBC-SHA                                                       | —                                                                            |
| RSA-PSK-AES128-CBC-SHA256                                                      | —                                                                            |
| DHE-PSK-AES128-CBC-SHA256                                                      | —                                                                            |
| RSA-PSK-AES128-CBC-SHA                                                         | —                                                                            |
| DHE-PSK-AES128-CBC-SHA                                                         | —                                                                            |
| ECDHE-PSK-CAMELLIA128-SHA256                                                   | —                                                                            |
| RSA-PSK-CAMELLIA128-SHA256                                                     | —                                                                            |
| DHE-PSK-CAMELLIA128-SHA256                                                     | —                                                                            |
| PSK-AES128-CBC-SHA256                                                          | —                                                                            |
| PSK-CAMELLIA128-SHA256                                                         | —                                                                            |
| ECDHE-PSK-RC4-SHA                                                              | —                                                                            |
| RSA-PSK-RC4-SHA                                                                | —                                                                            |

| Supported Ciphers in Fabric OS 9.0.x That<br>Block the Downgrade to Fabric OS 8.x | Supported Ciphers in Fabric OS 8.x That<br>Block the Upgrade to Fabric OS 9.0.x |
|-----------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| DHE-PSK-RC4-SHA                                                                   | —                                                                               |
| ECDHE-PSK-3DES-EDE-CBC-SHA                                                        | —                                                                               |
| RSA-PSK-3DES-EDE-CBC-SHA                                                          | —                                                                               |
| DHE-PSK-3DES-EDE-CBC-SHA                                                          | —                                                                               |

## Appendix D: FICON

This appendix provides additional information for FICON including a basic switch configuration, a couple of address binding examples, the configuration information record worksheet, and the Extended Binary Coded Decimal Interchange Code (EBCDIC) page.

### Basic Switch Configuration

This appendix provides basic steps and commands to quickly configure a switch for fabric and possible FICON and cascaded FICON operation.

For detailed concepts, procedures, and additional Fabric OS command options to configure a switch for point-to-point and cascaded FICON operation, see [Administering FICON Fabrics](#).

The following tables are included in this appendix:

- The Switch Configuration Example and Commands table provides the basic steps and commands in a checklist format to quickly configure a switch for fabric and possible FICON operation.
- The Security Example and Commands table provides additional steps and commands to set FICON security parameters.

**Table 160: Switch Configuration Example and Commands**

| Step                                                                                                                                                                        | Relevant Commands                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------|
| Upgrade the code (if necessary).                                                                                                                                            | <code>firmwaredownload</code><br><code>firmwareshow</code>        |
| Add or show feature keys (add licenses).                                                                                                                                    | <code>licenseadd</code><br><code>licenseshow</code>               |
| Create all logical switches, and add all intended ports to each logical switch.<br>The <code>-ficon</code> option sets many of the logical switch attributes for FICON use. | <code>lscfg --create</code><br><code>lscfg --create -ficon</code> |
| Set the switch offline.                                                                                                                                                     | <code>switchdisable</code>                                        |

| Step                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Relevant Commands                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Repeat the following steps to configure switch parameters for each logical switch.</p> <ol style="list-style-type: none"> <li>1. Enter the <code>configure</code> command. You will be asked several questions. Answer questions listed in the Relevant Commands column only. Use defaults for other settings unless instructed by a Brocade certified architect for FICON. Note that the <code>BB_Credit</code> setting should not be changed here. <code>BB_Credit</code> will be set elsewhere in this table.</li> <li>2. Change fabric parameter settings by answering <b>Yes</b>. After answering <b>Yes</b>, you will be asked several questions. Change only the following parameters. Leave all other parameters at the default unless instructed by a Brocade certified architect for FICON. After completing all of the fabric parameter settings, answer <b>No</b> for all other parameters. <ul style="list-style-type: none"> <li>– Domain</li> <li>– Allow XISL Use</li> </ul> </li> </ol> <p>When applicable:</p> <ul style="list-style-type: none"> <li>• Enable a 256 Area Limit: <ul style="list-style-type: none"> <li>– 0 – Not supported for FICON.</li> <li>– 1 – Zero-based addressing. Recommended for FICON.</li> <li>– 2 – Port-based addressing. Not supported for 48-port cards in 8-port slot chassis.</li> </ul> </li> <li>• Disable device probing: Set to 1.</li> <li>• Insistent Domain ID Mode: Answer <b>Yes</b>.</li> </ul> | <pre>configure</pre>                                                                                                                                                                                                                                                                                                  |
| <p>Set the routing policy:</p> <ul style="list-style-type: none"> <li>• 1 – Port-based routing.</li> <li>• 2 – Device-based routing.</li> <li>• 3 – Exchange-based routing.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>To display the routing policy: <code>aptpolicy</code></p> <p>To set the routing policy to port-based routing: <code>aptpolicy 1</code></p> <p>To set the routing policy to device-based routing: <code>aptpolicy 2</code></p> <p>To set the routing policy to exchange-based routing: <code>aptpolicy 3</code></p> |
| Set in-order delivery.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <code>iodset</code>                                                                                                                                                                                                                                                                                                   |
| Enable dynamic load sharing with lossless DLS enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <code>dlsset --enable -lossless</code>                                                                                                                                                                                                                                                                                |
| Set the switch name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <code>switchname</code>                                                                                                                                                                                                                                                                                               |
| Set the chassis name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <code>chassisname</code>                                                                                                                                                                                                                                                                                              |
| Set the switch online.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <code>switchenable</code>                                                                                                                                                                                                                                                                                             |
| Set the principal switch. This can be only one switch in the fabric. In mixed M-Type and B-Type fabrics, the M-Type must be the principal switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <code>fabricprincipal 1</code>                                                                                                                                                                                                                                                                                        |
| Enable FMS if FICON CUP will be used. High-integrity fabric requirements must be met, and HIF mode must be enabled to enable FMS mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <pre>ficoncupset fmsmode enable ficoncupshow fmsmode</pre>                                                                                                                                                                                                                                                            |
| Set the default zone for no access.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <pre>defZone --noaccess cfgsave defZone --show</pre>                                                                                                                                                                                                                                                                  |

| Step                                                                                                                                                                                                              | Relevant Commands                                                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Configure zones. This is much easier with Brocade SANnav or Web Tools. Typically, scripts are used when using Fabric OS commands.                                                                                 | <pre> cfgclear cfgcreate cfgdelete cfgdisable cfgenable cfgremove cfgsave cfgshow </pre>                                              |
| Clear all port statistics and trace information.                                                                                                                                                                  | <pre> supportsave -R errclear statsclear portlogclear diagclearerror -all tracedump -c slotstatsclear errclear ficonclear rlir </pre> |
| Turn on FICON logging.                                                                                                                                                                                            | <pre>supportshowcfgenable ficon</pre>                                                                                                 |
| Set port fencing (refer to the <i>Brocade Fabric OS MAPS User Guide</i> ). Setting port fencing parameters is much easier with Brocade SANnav (refer to the <i>Brocade SANnav Management Portal User Guide</i> ). | <p>To set port fencing:</p> <pre>mapsconfig --actions fence</pre> <p>To define alerts:</p> <pre>mapsConfig --emailcfg</pre>           |

Parameters in the following table provide additional steps and commands to set FICON security parameters.

**Table 161: Security Example and Commands**

| Step                                                                            | Parameter                                                                                                                                                                                    |
|---------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Add members to the SCC policy.                                                  | <pre>secPolicyCreate "SCC_POLICY", "wwn;wwn; ..."</pre> <p>Or, to add all switches already in the fabric, use an asterisk (*) for each WWN:</p> <pre>secPolicyCreate "SCC_POLICY", "*"</pre> |
| Save the policy.                                                                | <pre>secpolicysave</pre>                                                                                                                                                                     |
| Activate the policy.                                                            | <pre>secpolicyactivate</pre>                                                                                                                                                                 |
| Enable the ACL fabric-wide consistency policy, and enforce a strict SCC policy. | <pre>fddcfg --fabwideset "SCC:S;"</pre>                                                                                                                                                      |

## Address Binding Examples

This section provides examples for binding ports sequentially and unbinding multiple ports.

### Sequential Address Binding

This section provides an example that will allow you to configure multiple ports on a blade installed in a Brocade FICON director to use sequential, zero-based, Mode 1 addresses. This type of range binding is the default in a FICON logical switch. For additional information, see [Configuring FICON](#).

Use the `portAddress --bind` command to bind port addresses to ports. Beginning with Fabric OS 9.0.0, you may not need to disable the port before binding an address. For example, if a port is moved into a FICON logical switch and receives a port address of 32, you can bind the port to the address 3200 without disabling the port. However, if you

choose to use a different port address, you must disable the port using the `portDisable` command before binding an address using the `portAddress --bind` command. To change the bound address, you must first unbind the addresses. See [Unbinding Multiple Ports](#) for an easy method to unbind multiple ports at once.

## Binding Ports Sequentially (Mode 1)

The following example uses the `portAddress --bind` command to bind a series of ports on a slot to sequential, zero-based (Mode 1) addresses.

1. Display the list of bound ports.

Note that the output is truncated.

```
switch_128_175:FID128:admin> portaddress --show
Partition Address Mode :1
Index Slot Port Area Mode User_bound
=====
104 4 8 0x2800 8 bit Y
105 4 9 0x2900 8 bit Y
106 4 10 0x2a00 8 bit Y
107 4 11 0x2b00 8 bit Y
108 4 12 0x2c00 8 bit Y
109 4 13 0x2d00 8 bit Y
110 4 14 0x2e00 8 bit Y
111 4 15 0x2f00 8 bit Y
112 4 16 0x1d00 8 bit -
113 4 17 0x1e00 8 bit -
114 4 18 0x1f00 8 bit -
115 4 19 0x3000 8 bit -
116 4 20 0x3100 8 bit -
117 4 21 0x3200 8 bit -
118 4 22 0x3300 8 bit -
119 4 23 0x3400 8 bit -
120 4 24 0x5800 8 bit Y
121 4 25 0x5900 8 bit Y
...
231 7 39 0x2700 8 bit Y
```

2. Bind ports 16 through 23 on slot 4.

```
switch_128_175:FID128:admin> portaddress --bind 4/16-23 3000
```

3. Display the updated list of bound ports.

Note that the output is truncated.

```
switch_128_175:FID128:admin> portaddress --show
Partition Address Mode :1
Index Slot Port Area Mode User_bound
=====
104 4 8 0x2800 8 bit Y
105 4 9 0x2900 8 bit Y
106 4 10 0x2a00 8 bit Y
107 4 11 0x2b00 8 bit Y
108 4 12 0x2c00 8 bit Y
109 4 13 0x2d00 8 bit Y
110 4 14 0x2e00 8 bit Y
111 4 15 0x2f00 8 bit Y
```

|     |   |    |        |       |   |
|-----|---|----|--------|-------|---|
| 112 | 4 | 16 | 0x3000 | 8 bit | Y |
| 113 | 4 | 17 | 0x3100 | 8 bit | Y |
| 114 | 4 | 18 | 0x3200 | 8 bit | Y |
| 115 | 4 | 19 | 0x3300 | 8 bit | Y |
| 116 | 4 | 20 | 0x3400 | 8 bit | Y |
| 117 | 4 | 21 | 0x3500 | 8 bit | Y |
| 118 | 4 | 22 | 0x3600 | 8 bit | Y |
| 119 | 4 | 23 | 0x3700 | 8 bit | Y |
| ... |   |    |        |       |   |
| 231 | 7 | 39 | 0x2700 | 8 bit | Y |

#### 4. Enable the ports.

```
switch_128_175:FID128:admin> portenable 4/16-23
2020/10/01-10:20:33, [FICU-1021], 15, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x30) have been
Unlocked by CLI.
2020/10/01-10:20:34, [FICU-1021], 16, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x31) have been
Unlocked by CLI.
2020/10/01-10:20:34, [FICU-1021], 17, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x32) have been
Unlocked by CLI.
2020/10/01-10:20:34, [FICU-1021], 18, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x33) have been
Unlocked by CLI.
2020/10/01-10:20:34, [FICU-1021], 19, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x34) have been
Unlocked by CLI.
2020/10/01-10:20:34, [FICU-1021], 20, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x35) have been
Unlocked by CLI.
2020/10/01-10:20:35, [FICU-1021], 21, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x36) have been
Unlocked by CLI.
2020/10/01-10:20:35, [FICU-1021], 22, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x37) have been
Unlocked by CLI.
```

## Unbinding Multiple Ports

The following example uses the `portAddress --unbind` command to unbind a series of ports on the same slot.

#### 1. Display the list of bound ports.

Note that the output is truncated.

```
switch_128_175:FID128:admin> portaddress --show
Partition Address Mode :1
Index Slot Port Area Mode User_bound
=====
104 4 8 0x2800 8 bit Y
105 4 9 0x2900 8 bit Y
106 4 10 0x2a00 8 bit Y
107 4 11 0x2b00 8 bit Y
108 4 12 0x2c00 8 bit Y
109 4 13 0x2d00 8 bit Y
110 4 14 0x2e00 8 bit Y
111 4 15 0x2f00 8 bit Y
112 4 16 0x3000 8 bit Y
113 4 17 0x3100 8 bit Y
114 4 18 0x3200 8 bit Y
115 4 19 0x3300 8 bit Y
116 4 20 0x3400 8 bit Y
117 4 21 0x3500 8 bit Y
```



```

118 4 22 0x3600 8 bit Y
119 4 23 0x3700 8 bit Y
...
231 7 39 0x2700 8 bit Y

```

## 2. Disable the ports that you want to unbind (ports 16-23 on slot 4).

```

switch_128_175:FID128:admin> portdisable 4/16-23
2020/10/01-10:18:51, [FICU-1020], 7, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x30) have been
Blocked by CLI.
2020/10/01-10:18:51, [FICU-1020], 8, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x31) have been
Blocked by CLI.
2020/10/01-10:18:51, [FICU-1020], 9, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x32) have been
Blocked by CLI.
2020/10/01-10:18:51, [FICU-1020], 10, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x33) have been
Blocked by CLI.
2020/10/01-10:18:52, [FICU-1020], 11, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x34) have been
Blocked by CLI.
2020/10/01-10:18:52, [FICU-1020], 12, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x35) have been
Blocked by CLI.
2020/10/01-10:18:52, [FICU-1020], 13, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x36) have been
Blocked by CLI.
2020/10/01-10:18:52, [FICU-1020], 14, SLOT 2 | FID 128, INFO, switch_128_175, Port Addrs (0x37) have been
Blocked by CLI.

```

## 3. Unbind ports 16 through 23 on slot 4.

```

switch_128_175:FID128:admin> portaddress --unbind 4/16-23

```

## 4. Display the updated list of bound ports.

Note that the output is truncated.

```

switch_128_175:FID128:admin> portaddress --show
Partition Address Mode :1
Index Slot Port Area Mode User_bound
=====
104 4 8 0x2800 8 bit Y
105 4 9 0x2900 8 bit Y
106 4 10 0x2a00 8 bit Y
107 4 11 0x2b00 8 bit Y
108 4 12 0x2c00 8 bit Y
109 4 13 0x2d00 8 bit Y
110 4 14 0x2e00 8 bit Y
111 4 15 0x2f00 8 bit Y
112 4 16 0x1d00 8 bit -
113 4 17 0x1e00 8 bit -
114 4 18 0x1f00 8 bit -
115 4 19 0x3000 8 bit -
116 4 20 0x3100 8 bit -
117 4 21 0x3200 8 bit -
118 4 22 0x3300 8 bit -
119 4 23 0x3400 8 bit -
120 4 24 0x5800 8 bit Y
121 4 25 0x5900 8 bit Y
...
231 7 39 0x2700 8 bit Y

```



| EBCDIC Code Page 00037 |                   |                   |                   |                   |                       |                   |                   |                   |                   |                   |                     |                   |                   |                   |                   |                       |
|------------------------|-------------------|-------------------|-------------------|-------------------|-----------------------|-------------------|-------------------|-------------------|-------------------|-------------------|---------------------|-------------------|-------------------|-------------------|-------------------|-----------------------|
|                        | -0                | -1                | -2                | -3                | -4                    | -5                | -6                | -7                | -8                | -9                | -A                  | -B                | -C                | -D                | -E                | -F                    |
| 0-                     | NUL<br>0000<br>0  | SOM<br>0001<br>1  | STX<br>0002<br>2  | ETX<br>0003<br>3  | SEL<br>0004<br>4      | HT<br>0005<br>5   | RNL<br>0006<br>6  | DEL<br>0007<br>7  | GE<br>0008<br>8   | SPS<br>0009<br>9  | RPT<br>000A<br>10   | VT<br>000B<br>11  | FF<br>000C<br>12  | CR<br>000D<br>13  | SO<br>000E<br>14  | SI<br>000F<br>15      |
| 1-                     | DLE<br>0010<br>16 | DC1<br>0011<br>17 | DC2<br>0012<br>18 | DC3<br>0013<br>19 | RES ENP<br>0014<br>20 | NL<br>0015<br>21  | SS<br>0016<br>22  | POC<br>0017<br>23 | CAN<br>0018<br>24 | EM<br>0019<br>25  | UBS<br>001A<br>26   | CU1<br>001B<br>27 | IFS<br>001C<br>28 | IGS<br>001D<br>29 | IRS<br>001E<br>30 | IUS ITS<br>001F<br>31 |
| 2-                     | DS<br>0020<br>32  | SOS<br>0021<br>33 | FS<br>0022<br>34  | WUS<br>0023<br>35 | BYP INP<br>0024<br>36 | LF<br>0025<br>37  | ETS<br>0026<br>38 | ESC<br>0027<br>39 | SA<br>0028<br>40  | SFE<br>0029<br>41 | SM SW<br>002A<br>42 | CSP<br>002B<br>43 | MFA<br>002C<br>44 | ENQ<br>002D<br>45 | ACK<br>002E<br>46 | REL<br>002F<br>47     |
| 3-                     |                   |                   | SYN<br>0030<br>48 | IR<br>0031<br>49  | PP<br>0032<br>50      | TRN<br>0033<br>51 | NBS<br>0034<br>52 | EOT<br>0035<br>53 | SSS<br>0036<br>54 | IT<br>0037<br>55  | RFF<br>0038<br>56   | CU2<br>0039<br>57 | DC4<br>003A<br>58 | NAR<br>003B<br>59 |                   | SUB<br>003C<br>60     |
| 4-                     | SP<br>0040<br>64  | RSP<br>0041<br>65 | â<br>0042<br>66   | ä<br>0043<br>67   | à<br>0044<br>68       | á<br>0045<br>69   | ã<br>0046<br>70   | å<br>0047<br>71   | ç<br>0048<br>72   | ñ<br>0049<br>73   | ¢<br>004A<br>74     | .<br>004B<br>75   | <<br>004C<br>76   | (<br>004D<br>77   | +<br>004E<br>78   | <br>004F<br>79        |
| 5-                     | &<br>0050<br>80   | é<br>0051<br>81   | ê<br>0052<br>82   | ë<br>0053<br>83   | è<br>0054<br>84       | í<br>0055<br>85   | î<br>0056<br>86   | ï<br>0057<br>87   | ì<br>0058<br>88   | í<br>0059<br>89   | !<br>005A<br>90     | \$<br>005B<br>91  | *<br>005C<br>92   | )<br>005D<br>93   | ;<br>005E<br>94   | ~<br>005F<br>95       |
| 6-                     | -<br>0060<br>96   | /<br>0061<br>97   | À<br>0062<br>98   | Ä<br>0063<br>99   | À<br>0064<br>100      | Á<br>0065<br>101  | Ã<br>0066<br>102  | Å<br>0067<br>103  | Ç<br>0068<br>104  | Ñ<br>0069<br>105  | ¡<br>006A<br>106    | ,<br>006B<br>107  | %<br>006C<br>108  | _<br>006D<br>109  | ><br>006E<br>110  | ?<br>006F<br>111      |
| 7-                     | ø<br>0070<br>112  | É<br>0071<br>113  | Ê<br>0072<br>114  | Ë<br>0073<br>115  | È<br>0074<br>116      | Í<br>0075<br>117  | Î<br>0076<br>118  | Ï<br>0077<br>119  | Ì<br>0078<br>120  | ~<br>0079<br>121  | :<br>007A<br>122    | #<br>007B<br>123  | @<br>007C<br>124  | '<br>007D<br>125  | =<br>007E<br>126  | "<br>007F<br>127      |
| 8-                     | ø<br>0080<br>128  | a<br>0081<br>129  | b<br>0082<br>130  | c<br>0083<br>131  | d<br>0084<br>132      | e<br>0085<br>133  | f<br>0086<br>134  | g<br>0087<br>135  | h<br>0088<br>136  | i<br>0089<br>137  | «<br>008A<br>138    | »<br>008B<br>139  | ø<br>008C<br>140  | ý<br>008D<br>141  | þ<br>008E<br>142  | ±<br>008F<br>143      |
| 9-                     | °<br>0090<br>144  | j<br>0091<br>145  | k<br>0092<br>146  | l<br>0093<br>147  | m<br>0094<br>148      | n<br>0095<br>149  | o<br>0096<br>150  | p<br>0097<br>151  | q<br>0098<br>152  | r<br>0099<br>153  | ø<br>009A<br>154    | ø<br>009B<br>155  | æ<br>009C<br>156  | .<br>009D<br>157  | Æ<br>009E<br>158  | ¤<br>009F<br>159      |
| A-                     | μ<br>00A0<br>160  | ~<br>00A1<br>161  | s<br>00A2<br>162  | t<br>00A3<br>163  | u<br>00A4<br>164      | v<br>00A5<br>165  | w<br>00A6<br>166  | x<br>00A7<br>167  | y<br>00A8<br>168  | z<br>00A9<br>169  | ı<br>00AA<br>170    | ı<br>00AB<br>171  | Đ<br>00AC<br>172  | Ÿ<br>00AD<br>173  | þ<br>00AE<br>174  | ®<br>00AF<br>175      |
| B-                     | ^<br>00B0<br>176  | £<br>00B1<br>177  | ¥<br>00B2<br>178  | •<br>00B3<br>179  | ©<br>00B4<br>180      | §<br>00B5<br>181  | ¶<br>00B6<br>182  | ¼<br>00B7<br>183  | ½<br>00B8<br>184  | ¾<br>00B9<br>185  | [<br>00BA<br>186    | ]<br>00BB<br>187  | -<br>00BC<br>188  | "<br>00BD<br>189  | ´<br>00BE<br>190  | x<br>00BF<br>191      |
| C-                     | {<br>00C0<br>192  | A<br>00C1<br>193  | B<br>00C2<br>194  | C<br>00C3<br>195  | D<br>00C4<br>196      | E<br>00C5<br>197  | F<br>00C6<br>198  | G<br>00C7<br>199  | H<br>00C8<br>200  | I<br>00C9<br>201  | SHY<br>00CA<br>202  | Ô<br>00CB<br>203  | Ö<br>00CC<br>204  | Ò<br>00CD<br>205  | Ó<br>00CE<br>206  | Õ<br>00CF<br>207      |
| D-                     | }<br>00D0<br>208  | J<br>00D1<br>209  | K<br>00D2<br>210  | L<br>00D3<br>211  | M<br>00D4<br>212      | N<br>00D5<br>213  | O<br>00D6<br>214  | P<br>00D7<br>215  | Q<br>00D8<br>216  | R<br>00D9<br>217  | ¹<br>00DA<br>218    | Û<br>00DB<br>219  | Ü<br>00DC<br>220  | Ù<br>00DD<br>221  | Ú<br>00DE<br>222  | ÿ<br>00DF<br>223      |
| E-                     | \<br>00E0<br>224  | ÷<br>00E1<br>225  | S<br>00E2<br>226  | T<br>00E3<br>227  | U<br>00E4<br>228      | V<br>00E5<br>229  | W<br>00E6<br>230  | X<br>00E7<br>231  | Y<br>00E8<br>232  | Z<br>00E9<br>233  | ²<br>00EA<br>234    | Ô<br>00EB<br>235  | Ö<br>00EC<br>236  | Ò<br>00ED<br>237  | Ó<br>00EE<br>238  | Õ<br>00EF<br>239      |
| F-                     | ø<br>00F0<br>240  | 1<br>00F1<br>241  | 2<br>00F2<br>242  | 3<br>00F3<br>243  | 4<br>00F4<br>244      | 5<br>00F5<br>245  | 6<br>00F6<br>246  | 7<br>00F7<br>247  | 8<br>00F8<br>248  | 9<br>00F9<br>249  | ³<br>00FA<br>250    | Û<br>00FB<br>251  | Ü<br>00FC<br>252  | Ù<br>00FD<br>253  | Ú<br>00FE<br>254  | EO<br>00FF<br>255     |

## Appendix E: Hexadecimal Conversion

This section describes the hexadecimal, also known as hex, which is a numeral system with a base of 16, is written using symbols 0 through 9 and A to F (or a to f).

Its primary purpose is to represent the binary code that computers interpret in a format easier for humans to remember. Hexadecimal acts as a form of shorthand, in which one hexadecimal digit takes the place of four binary bits. For example, the decimal numeral 79, with the binary representation of 01001111, is 4F (or 4f) in hexadecimal where 4 = 0100, and F = 1111.

Hexadecimal numbers can have either a *0x* prefix or an *h* suffix. The address 0xFFFFFA is the same address as FFFFFA *h*. This type of address with six digits representing 3 bytes is called a hex triplet. Fibre Channel uses the hexadecimal notation in hex triplets to specify well-known addresses and port IDs.

### Example Conversion of the Hexadecimal Triplet 0x610600

Notice the PID (**610600** in bold) in the `nsShow` output is in hexadecimal.

```
switch:admin> nsshow
{
 Type Pid COS PortName NodeName SCR
 N
 610600
; 3;20:08:00:11:0d:25:00:00;20:08:00:11:0d:25:00:00; 0x00000000
 SCR: None
 FC4s: FCP
 PortSymb: [33] "SB V7.4 port8 10.1.18.143 Targ0"
 Fabric Port Name: 20:44:00:05:33:0d:84:00
 Permanent Port Name: 20:08:00:11:0d:25:00:00
 Device type: Physical Target
 Port Index: 68
 Share Area: No
 Redirect: No
 Partial: No
 LSAN: Yes
 Slow Drain Device: No
 Device link speed: 8G
 Connected through AG: No
 Real device behind AG: No
 FCoE: No
 FC4 Features [FCP]: Target
```

The Local Name Server has 1 entry }

1. Separate the six digits into three numbers by inserting a space after every two digits: 61 06 00
2. Convert each hexadecimal value to a decimal representation:
  - 61 = Domain ID = 97
  - 06 = Area (port number) = 06

00 = Port (ALPA) = 00 (not used in this instance, but is used in a loop, which are shared areas in PID assignments on blades, NPIV, and Access Gateway devices)

Result: hexadecimal triplet 610600 = decimal triplet 97,06,00

## Decimal-to-Hexadecimal Conversion Table

**Table 162: Decimal-to-Hexadecimal Conversion Table**

| Decimal | 01  | 02  | 03  | 04  | 05  | 06  | 07  | 08  | 09  | 10  |
|---------|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| Hex     | 01  | 02  | 03  | 04  | 05  | 06  | 07  | 08  | 09  | 0a  |
| Decimal | 11  | 12  | 13  | 14  | 15  | 16  | 17  | 18  | 19  | 20  |
| Hex     | 0b  | 0c  | 0d  | 0e  | 0f  | 10  | 11  | 12  | 13  | 14  |
| Decimal | 21  | 22  | 23  | 24  | 25  | 26  | 27  | 28  | 29  | 30  |
| Hex     | 15  | 16  | 17  | 18  | 19  | 1a  | 1b  | 1c  | 1d  | 1e  |
| Decimal | 31  | 32  | 33  | 34  | 35  | 36  | 37  | 38  | 39  | 40  |
| Hex     | 1f  | 20  | 21  | 22  | 23  | 24  | 25  | 26  | 27  | 28  |
| Decimal | 41  | 42  | 43  | 44  | 45  | 46  | 47  | 48  | 49  | 50  |
| Hex     | 29  | 2a  | 2b  | 2c  | 2d  | 2e  | 2f  | 30  | 31  | 32  |
| Decimal | 51  | 52  | 53  | 54  | 55  | 56  | 57  | 58  | 59  | 60  |
| Hex     | 33  | 34  | 35  | 36  | 37  | 38  | 39  | 3a  | 3b  | 3c  |
| Decimal | 61  | 62  | 63  | 64  | 65  | 66  | 67  | 68  | 69  | 70  |
| Hex     | 3d  | 3e  | 3f  | 40  | 41  | 42  | 43  | 44  | 45  | 46  |
| Decimal | 71  | 72  | 73  | 74  | 75  | 76  | 77  | 78  | 79  | 80  |
| Hex     | 47  | 48  | 49  | 4a  | 4b  | 4c  | 4d  | 4e  | 4f  | 50  |
| Decimal | 81  | 82  | 83  | 84  | 85  | 86  | 87  | 88  | 89  | 90  |
| Hex     | 51  | 52  | 53  | 54  | 55  | 56  | 57  | 58  | 59  | 5a  |
| Decimal | 91  | 92  | 93  | 94  | 95  | 96  | 97  | 98  | 99  | 100 |
| Hex     | 5b  | 5c  | 5d  | 5e  | 5f  | 60  | 61  | 62  | 63  | 64  |
| Decimal | 101 | 102 | 103 | 104 | 105 | 106 | 107 | 108 | 109 | 110 |
| Hex     | 65  | 66  | 67  | 68  | 69  | 6a  | 6b  | 6c  | 6d  | 6e  |
| Decimal | 111 | 112 | 113 | 114 | 115 | 116 | 117 | 118 | 119 | 120 |
| Hex     | 6f  | 70  | 71  | 72  | 73  | 74  | 75  | 76  | 77  | 78  |
| Decimal | 121 | 122 | 123 | 124 | 125 | 126 | 127 | 128 | 129 | 130 |
| Hex     | 79  | 7a  | 7b  | 7c  | 7d  | 7e  | 7f  | 80  | 81  | 82  |
| Decimal | 131 | 132 | 133 | 134 | 135 | 136 | 137 | 138 | 139 | 140 |
| Hex     | 83  | 84  | 85  | 86  | 87  | 88  | 89  | 8a  | 8b  | 8c  |
| Decimal | 141 | 142 | 143 | 144 | 145 | 146 | 147 | 148 | 149 | 150 |
| Hex     | 8d  | 8e  | 8f  | 90  | 91  | 92  | 93  | 94  | 95  | 96  |
| Decimal | 151 | 152 | 153 | 154 | 155 | 156 | 157 | 158 | 159 | 160 |
| Hex     | 97  | 98  | 99  | 9a  | 9b  | 9c  | 9d  | 9e  | 9f  | a0  |
| Decimal | 161 | 162 | 163 | 164 | 165 | 166 | 167 | 168 | 169 | 170 |
| Hex     | a1  | a2  | a3  | a4  | a5  | a6  | a7  | a8  | a9  | aa  |



|                |            |            |            |            |            |            |            |            |            |            |
|----------------|------------|------------|------------|------------|------------|------------|------------|------------|------------|------------|
| <b>Decimal</b> | <b>171</b> | <b>172</b> | <b>173</b> | <b>174</b> | <b>175</b> | <b>176</b> | <b>177</b> | <b>178</b> | <b>179</b> | <b>180</b> |
| Hex            | ab         | ac         | ad         | ae         | af         | b0         | b1         | b2         | b3         | b4         |
| <b>Decimal</b> | <b>181</b> | <b>182</b> | <b>183</b> | <b>184</b> | <b>185</b> | <b>186</b> | <b>187</b> | <b>188</b> | <b>189</b> | <b>190</b> |
| Hex            | b5         | b6         | b7         | b8         | b9         | ba         | bb         | bc         | bd         | be         |
| <b>Decimal</b> | <b>191</b> | <b>192</b> | <b>193</b> | <b>194</b> | <b>195</b> | <b>196</b> | <b>197</b> | <b>198</b> | <b>199</b> | <b>200</b> |
| Hex            | bf         | c0         | c1         | c2         | c3         | c4         | c5         | c6         | c7         | c8         |
| <b>Decimal</b> | <b>201</b> | <b>202</b> | <b>203</b> | <b>204</b> | <b>205</b> | <b>206</b> | <b>207</b> | <b>208</b> | <b>209</b> | <b>210</b> |
| Hex            | c9         | ca         | cb         | cc         | cd         | ce         | cf         | d0         | d1         | d2         |
| <b>Decimal</b> | <b>211</b> | <b>212</b> | <b>213</b> | <b>214</b> | <b>215</b> | <b>216</b> | <b>217</b> | <b>218</b> | <b>219</b> | <b>220</b> |
| Hex            | d3         | d4         | d5         | d6         | d7         | d8         | d9         | da         | db         | dc         |
| <b>Decimal</b> | <b>221</b> | <b>222</b> | <b>223</b> | <b>224</b> | <b>225</b> | <b>226</b> | <b>227</b> | <b>228</b> | <b>229</b> | <b>230</b> |
| Hex            | dd         | de         | df         | e0         | e1         | e2         | e3         | e4         | e5         | e6         |
| <b>Decimal</b> | <b>231</b> | <b>232</b> | <b>233</b> | <b>234</b> | <b>235</b> | <b>236</b> | <b>237</b> | <b>238</b> | <b>239</b> | <b>240</b> |
| Hex            | e7         | e8         | e9         | ea         | eb         | ec         | ed         | ef         | ee         | f0         |
| <b>Decimal</b> | <b>241</b> | <b>242</b> | <b>243</b> | <b>244</b> | <b>245</b> | <b>246</b> | <b>247</b> | <b>248</b> | <b>249</b> | <b>250</b> |
| Hex            | f1         | f2         | f3         | f4         | f5         | f6         | f7         | f8         | f9         | fa         |
| <b>Decimal</b> | <b>251</b> | <b>252</b> | <b>253</b> | <b>254</b> | <b>255</b> | —          | —          | —          | —          | —          |
| Hex            | fb         | fc         | fd         | fe         | ff         | —          | —          | —          | —          | —          |

## Revision History

---

The revision history provides a list of the important changes made in each version of the document.

### **FOS-91x-Admin-AG104; January 11, 2024**

Removed references to the *Brocade Fabric OS Troubleshooting and Diagnostics User Guide*.

### **FOS-91x-Admin-AG103; August 31, 2023**

- Added a note in the [Configuring Fibre Channel Link Aggregation](#) section.
- Added a note in the [Configuration File Restoration](#) section.
- Updated the description in the [Creating a Complete Chain of CA Certificates](#) section.
- Added a note in the director clock alert mode (DCAM) of [Mode Register Bit Settings](#) section.
- Updated the sensor values, fan count, and PSUs in the [Brocade Switch and Director Sensors](#) section.

### **FOS-91x-Admin-AG102; June 15, 2022**

- Added support for other (non Read/Write) commands in the [Flow Metrics](#) and [Flow-Level Errors and Exception Metrics](#) sections.

### **FOS-91x-Admin-AG101; March 16, 2022**

- Added a new CPX blade in the [Blade Terminology and Compatibility](#) table.
- Added the [Administering FICON Fabrics](#) and [Appendix D: FICON](#) chapters.

### **FOS-91x-Admin-AG100; December 15, 2021**

- Initial document version.

## Documentation Legal Notice

---

This notice provides copyright and trademark information as well as legal disclaimers.

Copyright © 2024. Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to [www.broadcom.com](http://www.broadcom.com). All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products or to view the licensing terms applicable to the open source software, please download the open source attribution disclosure document in the Broadcom Support Portal. If you do not have a support account or are unable to log in, please contact your support provider for this information.



