

Brocade Fabric OS Administration Guide, 8.2.1

Supporting Fabric OS 8.2.1

Copyright © 2018 Brocade Communications Systems LLC. All Rights Reserved. Brocade and the stylized B logo are among the trademarks of Brocade Communications Systems LLC. Broadcom, the pulse logo, and Connecting everything are among the trademarks of Broadcom. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Brocade, a Broadcom Inc. Company, reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Brocade is believed to be accurate and reliable. However, Brocade 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, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <https://www.broadcom.com/support/fibre-channel-networking/tools/oscd>.

Contents

Introduction.....	18
About This Guide.....	18
What's New in This Document.....	18
Supported Hardware and Software.....	18
Brocade Gen 5 (16Gb/s) Fixed-Port Switches.....	18
Brocade Gen 5 (16Gb/s) Directors.....	19
Brocade Gen 6 (32Gb/s) Fixed-Port Switches.....	19
Brocade Gen 6 (32Gb/s) Directors.....	19
Contacting Brocade Technical Support.....	19
Document Feedback.....	20
Understanding Fibre Channel Services	21
Fibre Channel Services Overview.....	21
Management Server.....	22
Application Server.....	22
Platform Services.....	22
Platform services and Virtual Fabrics.....	23
Enabling Platform Services.....	23
Disabling Platform Services	23
Management Server Database.....	23
Displaying the Management Server ACL.....	24
Adding a Member to the ACL.....	24
Deleting a Member from the ACL	25
Viewing the Contents of the Management Server Database	26
Clearing the Management Server Database.....	26
Topology Discovery.....	27
Displaying Topology Discovery Status.....	27
Enabling Topology Discovery	27
Disabling Topology Discovery.....	27
Device Login.....	28
Principal Switch.....	29
E_Port Login Process	29
Fabric Login process	29
Port Login Process.....	29
RSCNs.....	30
Duplicate Port World Wide Name.....	30
High Availability of Daemon Process.....	30
FL_Port and Arbitrated Loop Support.....	31
Performing Basic Configuration Tasks.....	32
Fabric OS Overview.....	32
Command Line Interface.....	32
Console Sessions Using the Serial Port.....	33
Telnet or SSH Session.....	34
Getting Help on a Command.....	35
Viewing the History of Command Line Entry.....	36
Using Fosexec to Run Commands on Remote Switch or Domain.....	39

Modification of Default Password.....	41
Default Account Password.....	42
The Switch Ethernet Interface	43
Brocade Directors.....	43
Brocade Switches.....	43
Virtual Fabrics and the Ethernet interface.....	44
Management Ethernet Port Bonding.....	44
Displaying the Network Interface Settings.....	45
Resetting the Management Network Interface Error Counter.....	46
Management Port Ingress Rate Limiting.....	47
Static Ethernet Addresses.....	47
DHCP Activation.....	49
Configuring the Static IPv6 Gateway Address.....	54
IPv6 Autoconfiguration.....	54
Setting the Ethernet Interface Mode and Speed.....	55
Date and Time Settings.....	56
Setting the Date and Time.....	56
Time Zone Settings.....	57
Network Time Protocol.....	59
Domain IDs.....	61
Domain ID Issues.....	61
Displaying the Domain IDs.....	61
Setting the Domain ID.....	62
Registered Organization Name	63
Setting and Displaying the Registered Organization Name.....	63
Switch Name.....	63
Customizing the Switch Name.....	64
Switch Type.....	64
Chassis Name.....	65
Customizing Chassis Name.....	66
Fabric Name.....	66
Configuring the Fabric Name.....	66
High Availability Considerations for Fabric Names.....	66
Upgrade and Downgrade Considerations for Fabric Names.....	67
Switch Activation and Deactivation.....	67
Disabling a Switch.....	67
Enabling a Switch.....	67
Disabling a Chassis.....	68
Enabling a Chassis.....	68
Device Shutdown.....	68
Powering Off a Brocade Switch.....	68
Powering Off a Brocade Director.....	69
Basic Connections.....	69
Device Connection.....	69
Switch Connection.....	70
Performing Advanced Configuration Tasks.....	71
Port identifiers and PID binding overview.....	71
Core PID addressing mode.....	71
Fixed addressing mode.....	72
10-bit addressing (mode 0).....	72

256-area addressing (mode 1 and mode 2).....	72
WWN-based PID assignment.....	73
Ports.....	75
Port types.....	75
Director port blades.....	76
Setting port names.....	77
Port identification by slot and port number.....	77
Port identification by port area ID.....	77
Port identification by index.....	77
Dynamic Portname.....	78
Configuring a device-switch connection	83
Swapping port area IDs.....	83
Enabling a port.....	84
Disabling a port.....	85
Port decommissioning.....	86
Port decommissioning without disabling.....	86
Setting port speeds.....	87
Setting all ports on a switch to the same speed.....	87
Setting port speed for a port octet.....	88
Setting maximum auto-negotiated port speed.....	90
Decoding fdmiShow command output.....	91
Blade terminology and compatibility.....	92
CP blades.....	93
Core blades.....	93
Port and application blade compatibility.....	94
Enabling and disabling blades.....	94
Enabling blades.....	94
Disabling blades.....	95
Blade swapping.....	95
How blades are swapped.....	96
Swapping blades.....	100
Disabling switches.....	100
Power management.....	100
Powering off a port blade or core blade.....	100
Powering on a port blade or core blade.....	101
Persistently powering off a port blade or core blade.....	101
Equipment status.....	102
Checking switch operation.....	102
Verifying High Availability features (Backbones and Directors only).....	102
Verifying fabric connectivity.....	103
Verifying device connectivity.....	103
Audit log configuration.....	104
Verifying host syslog prior to configuring the audit log.....	105
Configuring an audit log for specific event classes.....	105
Configuring remote syslog servers.....	106
Hostname support for the syslogAdmin command.....	107
Configuring quiet time for RASLog and audit log messages.....	107
Duplicate PWWN handling during device login.....	110
Setting 0, First login precedence.....	111
Setting 1, Second login precedence.....	111

Setting 2, Mixed precedence.....	111
Setting the behavior for handling duplicate PWWNs.....	112
Forward error correction.....	112
FEC limitations.....	113
Enabling forward error correction.....	113
Disabling forward error correction.....	114
Enabling or disabling FEC for long-distance ports	114
Displaying and clearing the FEC counters.....	115
FEC-via-TTS.....	115
Fabric-Assigned PWWN.....	117
Fabric-Assigned PWWN overview.....	117
User- and auto-assigned FA-PWWN behavior	118
Configuring an FA-PWWN for an HBA connected to an Access Gateway.....	119
Configuring an FA-PWWN for an HBA connected to an edge switch.....	120
Supported switches and configurations for FA-PWWN.....	120
Configuration upload and download considerations for FA-PWWN.....	121
Security considerations for FA-PWWN.....	121
Restrictions of FA-PWWN.....	122
Access Gateway N_Port failover with FA-PWWN.....	122
NPIV.....	123
NPIV overview.....	123
Upgrade considerations.....	124
Fixed addressing mode.....	124
10-bit addressing mode.....	124
Configuring NPIV.....	124
Enabling and disabling NPIV.....	125
Base device logout.....	126
Difference in the device logout behaviors.....	126
Enabling base device logout.....	127
Use cases and dependencies.....	127
Viewing base device logout setting.....	129
Viewing NPIV port configuration information.....	129
Viewing virtual PID login information.....	131
Maintaining the Switch Configuration File.....	133
Configuration settings.....	133
Configuration file format	134
Configuration file backup.....	135
Special characters in FTP server credentials.....	135
Uploading a configuration file in interactive mode.....	136
Configuration file restoration.....	137
Restrictions when downloading a configuration file.....	137
Configuration download without disabling a switch.....	138
Configurations across a fabric.....	140
Downloading a configuration file from one switch to another switch of the same model.....	140
Security considerations.....	140
Configuration management for Virtual Fabrics.....	140
Uploading a configuration file from a switch with Virtual Fabrics enabled.....	140
Restoring a logical switch configuration using configDownload.....	141
Configuration upload and download restrictions when Virtual Fabrics is enabled.....	142

Brocade configuration form.....	142
Managing User Accounts.....	144
User accounts overview	144
Role-Based Access Control.....	145
Management channel.....	146
Managing user-defined roles.....	147
Configuring time-based access.....	148
Local database user accounts.....	148
Default accounts.....	149
Local account passwords.....	150
Local user account database distribution.....	151
Distributing the local user database.....	151
Accepting distributed user databases on the local switch.....	152
Rejecting distributed user databases on the local switch.....	152
Password policies.....	152
Password strength policy.....	152
Password history policy.....	153
Password expiration policy.....	155
Account lockout policy.....	156
Changing the root password without the old password.....	157
Configuring the password hash type.....	157
The boot PROM password.....	158
Setting the boot PROM password for a switch with a recovery string.....	159
Setting the boot PROM password for Backbones and Directors with a recovery string.....	159
Setting the boot PROM password for a switch without a recovery string.....	160
Setting the boot PROM password for a Backbone or Director without a recovery string.....	161
Remote authentication.....	162
Remote authentication configuration.....	162
Setting the switch authentication mode.....	164
Fabric OS user accounts.....	165
Fabric OS users on the RADIUS server.....	166
Setting up a RADIUS server.....	168
LDAP configuration and Microsoft Active Directory.....	176
LDAP configuration and OpenLDAP.....	178
TACACS+ service.....	182
Obfuscation of RADIUS and TACACS+ shared secret.....	185
Remote authentication configuration on the switch.....	186
Configuring local authentication as backup.....	187
Configuring Protocols.....	189
Security Protocols.....	189
Secure Protocol Deployment Requirements.....	190
Security Protocol Scenarios.....	190
Secure Copy.....	190
Setting up SCP for configuration uploads and downloads.....	191
Secure Shell protocol.....	191
SSH public key authentication.....	191
SecCryptoCfg templates.....	194
Template structure.....	195
Certificate terminology.....	195

X509v3 validation modes.....	196
X509v3 certificate validation.....	196
SecCryptoCfg template management.....	197
Default secCryptoCfg templates.....	198
Certificate revocation check enforcement.....	202
Client authentication for TLS client - RADIUS/LDAP/SYSLOG.....	203
Creating a complete chain of CA certificates.....	205
Upgrade and downgrade considerations.....	205
Cryptographic compliance in FIPS mode.....	206
Configuring the ciphers, KEX, and MAC algorithms.....	207
Secure Sockets Layer protocol	207
Browser and Java support.....	207
SSL configuration overview.....	208
Creating a complete chain of CA certificates.....	214
The browser	215
Root certificates for the Java plugin.....	216
Simple Network Management Protocol.....	216
SNMP Manager.....	217
SNMP Agent.....	217
Management Information Base.....	217
Basic SNMP operation.....	217
Configuring SNMP using CLI.....	218
Telnet protocol.....	229
Blocking Telnet.....	229
Unblocking Telnet.....	230
Listener applications.....	230
Ports and applications used by switches.....	231
Port configuration.....	231
Configuring Security Policies.....	233
ACL policies overview.....	233
How the ACL policies are stored.....	233
Policy members.....	234
ACL policy management.....	234
Displaying ACL policies.....	234
Saving changes without activating the policies.....	234
Activating ACL policy changes.....	235
Deleting an ACL policy.....	235
Adding a member to an existing ACL policy.....	235
Removing a member from an ACL policy.....	236
Abandoning unsaved ACL policy changes.....	236
FCS policies.....	236
FCS policy restrictions.....	237
Ensuring the Fabric Domains Share Same Policy	238
Creating an FCS policy.....	238
Modifying the order of FCS switches.....	238
FCS policy distribution.....	239
Device Connection Control policies.....	240
Virtual Fabrics considerations for Device Connection Control policies	240
DCC policy restrictions.....	241
Creating a Device Connection Control policy.....	241

Deleting a DCC policy.....	242
DCC policy behavior with Fabric-Assigned PWWNs.....	242
SCC Policies.....	244
Virtual Fabrics considerations for SCC policies	244
Creating an SCC policy.....	244
Authentication policy for fabric elements.....	245
Virtual Fabrics considerations for authentication policies	246
E_Port authentication.....	246
Device authentication policy.....	248
AUTH policy restrictions.....	249
Authentication protocols.....	249
Secret key pairs for DH-CHAP.....	251
FCAP configuration overview.....	252
Fabric-wide distribution of the authorization policy.....	254
IP Filter policy.....	255
Virtual Fabrics considerations for IP Filter policies	255
Creating an IP Filter policy.....	255
Cloning an IP Filter policy.....	255
Displaying an IP Filter policy.....	255
Saving an IP Filter policy.....	256
Activating an IP Filter policy.....	256
Deleting an IP Filter policy.....	256
IP Filter Policy Rules.....	256
IP Filter policy enforcement.....	260
Adding a rule to an IP Filter policy.....	260
Deleting a rule from an IP Filter policy.....	260
Aborting an IP Filter transaction.....	260
IP Filter policy distribution.....	261
Policy database distribution.....	261
Database distribution settings.....	262
ACL policy distribution to other switches.....	263
Fabric-wide enforcement.....	263
Notes on joining a switch to the fabric.....	265
Management Interface Security.....	267
Configuration examples.....	267
IPsec protocols.....	269
Security associations.....	269
Authentication and encryption algorithms.....	270
IPsec policies.....	271
IKE policies.....	271
Creating the tunnel.....	272
Example of an end-to-end transport tunnel mode.....	274
In-flight Encryption and Compression.....	276
In-flight encryption and compression overview.....	276
Supported Gen 5 and Gen 6 in-flight encryption and compression features.....	277
Supported use cases for Brocade G630 switch.....	279
Supported port speed on encryption or compression-enabled Gen 5 (16-Gbps) ports	284
Additional in-flight encryption and compression restrictions.....	284
Availability considerations for encryption and compression.....	285
Virtual Fabrics considerations for encryption and compression.....	285

How in-flight encryption and compression are enabled.....	285
Authentication and key generation for encryption and compression.....	286
In-flight compression on long-distance ports.....	287
Compression ratios for compression-enabled ports.....	287
Configuring in-flight encryption and compression on an EX_Port.....	288
Configuring in-flight encryption and compression on an E_Port.....	289
Viewing the encryption and compression configuration.....	289
Configuring and enabling authentication for in-flight encryption.....	290
Enabling in-flight encryption.....	292
Enabling in-flight compression.....	293
Disabling in-flight encryption.....	294
Disabling in-flight compression.....	295
Encryption and Trunking Support on FC32-48 Blades.....	295
Trunking Behavior for Encrypted Ports.....	296
Recommendations and Considerations.....	296
Firmware Considerations.....	297
Buffer-to-Buffer Credits and Credit Recovery.....	298
Buffer credit management	298
Buffer-to-buffer flow control.....	298
Optimal buffer credit allocation	299
Fibre Channel gigabit values reference definition.....	299
Buffer credit allocation based on full-size frames.....	300
Allocating buffer credits based on average-size frames.....	303
Configuring buffers for a single port directly.....	303
Configuring buffers using frame size.....	304
Calculating the number of buffers required given the distance, speed, and frame size.....	304
Allocating buffer credits for F_Ports.....	305
Monitoring buffers in a port group.....	306
Buffer credits per switch or blade model.....	308
Maximum configurable distances for Extended Fabrics.....	310
Configuring credits for a single VC.....	312
Buffer credit recovery	312
Buffer credit recovery over an E_Port.....	313
Buffer credit recovery over an F_Port.....	313
Buffer credit recovery over an EX_Port.....	313
Enabling and disabling buffer credit recovery.....	314
Credit loss detection.....	314
Back-end credit loss detection and recovery support on Brocade 6520 switches.....	314
Enabling back-end credit loss detection and recovery for link reset thresholds.....	315
Performing link reset for loss of sync from peer ports.....	316
Back-end credit loss detection and recovery for link faults.....	317
Optimizing Fabric Behavior.....	319
Adaptive Networking overview.....	319
Ingress Rate Limiting.....	320
Virtual Fabrics considerations for ingress rate limiting.....	320
Limiting traffic from a particular device.....	320
Disabling Ingress Rate Limiting.....	320
QoS.....	321
License requirements for QoS.....	321

QoS on E_Ports.....	322
QoS over FC routers.....	322
vTap and QoS High Priority Zoning Compatibility mode.....	323
Virtual Fabrics considerations for QoS zone-based traffic prioritization.....	325
Traffic prioritization based on QoS zones.....	325
CS_CTL-based frame prioritization.....	331
Supported configurations for CS_CTL-based frame prioritization.....	332
High availability considerations for CS_CTL-based frame prioritization.....	332
Enabling CS_CTL-based frame prioritization on ports.....	332
Disabling CS_CTL-based frame prioritization on ports.....	332
Using CS_CTL auto mode at the chassis level.....	333
Considerations for using CS_CTL-based frame prioritization.....	333
Managing Trunking Connections.....	335
Trunking overview.....	335
Types of trunking.....	335
Masterless trunking.....	336
License requirements for trunking.....	336
Port groups for trunking.....	336
Supported platforms for trunking.....	337
Supported configurations for trunking.....	337
ISLs between G630 switches and other devices.....	337
High Availability support for trunking.....	342
Requirements for trunk groups.....	342
Recommendations for trunk groups.....	343
Configuring trunk groups.....	344
Enabling trunking.....	344
Disabling trunking.....	344
Displaying trunking information.....	345
ISL trunking over long-distance fabrics.....	346
EX_Port trunking.....	347
Masterless EX_Port trunking.....	347
Supported configurations and platforms for EX_Port trunking.....	348
Configuring EX_Port trunking.....	348
Displaying EX_Port trunking information.....	348
F_Port trunking.....	349
F_Port trunking for Access Gateway.....	349
F_Port trunking for Brocade adapters.....	352
F_Port trunking considerations.....	353
F_Port trunking in Virtual Fabrics.....	355
Displaying F_Port trunking information.....	355
Disabling F_Port trunking.....	356
Enabling the DCC policy on a trunk area.....	356
Inter-chassis Links.....	357
Inter-chassis links	357
License requirements for ICLs.....	357
Using the QSFPs that support 2 km on ICL ports	358
ICLs between DCX 8510 Backbones.....	359
ICL trunking between Brocade DCX 8510 Backbones.....	360
ICLs between Brocade X6 Directors.....	361

ICL trunking between Brocade X6 Directors.....	364
ICLs between Brocade DCX 8510 Backbones and Brocade X6 Directors.....	366
ICL trunking between DCX 8510 Backbones and X6 Directors.....	367
Virtual Fabrics considerations for ICLs.....	368
Supported topologies for ICL connections.....	369
Mesh topology.....	369
Core-edge topology.....	371
Managing Long-Distance Fabrics.....	373
Long-distance fabrics overview.....	373
Extended Fabrics device limitations.....	373
Long-distance link modes.....	374
Configuring an extended ISL.....	374
Enabling long distance when connecting to TDM devices	376
Forward error correction on long-distance links.....	377
Enabling FEC on a long-distance link.....	377
Disabling FEC on a long-distance link.....	377
Configuring NVMe.....	378
NVMe overview.....	378
NVMe over Fibre Channel.....	379
NVMe Name Server support.....	380
NVMe devices and zoning.....	381
Managing Virtual Fabrics.....	382
Virtual Fabrics overview.....	382
Logical switch overview.....	383
Default logical switch.....	383
Logical switches and fabric IDs.....	385
Port assignment in logical switches.....	386
Logical switches and connected devices.....	388
Management model for logical switches.....	389
Logical fabric overview.....	390
Logical fabric and ISLs.....	390
Base switch and extended ISLs.....	391
Account management and Virtual Fabrics.....	395
Setting up IP addresses for a logical switch.....	395
Logical Switch Login Context.....	396
Supported platforms for Virtual Fabrics.....	397
Supported port configurations in the fixed-port switches.....	398
Supported port configurations in Brocade Backbones and Directors.....	398
Virtual Fabrics interaction with other Fabric OS features.....	399
Limitations and restrictions of Virtual Fabrics.....	400
Restrictions on using XISLs.....	400
Restrictions on moving ports	401
Enabling Virtual Fabrics mode.....	401
Disabling Virtual Fabrics mode.....	402
Configuring logical switches to use basic configuration values.....	403
Creating a logical switch or base switch.....	403
Executing a command in a different logical switch context.....	405
Deleting a logical switch.....	406
Adding and moving ports on a logical switch.....	406

Displaying logical switch configuration.....	407
Changing the fabric ID of a logical switch.....	408
Changing a logical switch to a base switch.....	408
Configuring a logical switch for XISL use	409
Creating a FICON logical switch.....	410
Changing an existing logical switch to a FICON logical switch.....	413
Considerations for FICON logical switch.....	413
Changing the context to a different logical fabric.....	414
Administering Advanced Zoning.....	415
Zone types.....	415
Zoning overview.....	416
Approaches to zoning.....	417
Zone objects.....	418
Zone configurations.....	419
Zoning enforcement.....	420
Considerations for zoning architecture.....	421
Best practices for zoning.....	421
Broadcast zones.....	422
Broadcast zones and FC-FC routing.....	422
High availability considerations with broadcast zones.....	423
Loop devices and broadcast zones.....	423
Broadcast zones and default zoning mode.....	423
Zone aliases.....	423
Creating an alias.....	424
Adding members to an alias.....	424
Removing members from an alias.....	425
Deleting an alias.....	425
Viewing an alias in the defined configuration.....	426
Zone creation and maintenance.....	427
Displaying existing zones	427
Creating a zone.....	427
Adding devices (members) to a zone.....	428
Removing devices (members) from a zone.....	429
Replacing zone members.....	430
Deleting a zone.....	431
Viewing a zone.....	432
Viewing zone configuration names without case distinction.....	434
Validating a zone.....	435
Default zoning mode.....	439
Setting the default zoning mode.....	439
Viewing the current default zone access mode.....	440
Zone database size.....	440
Zone configurations.....	441
Creating a zone configuration.....	442
Adding zones to a zone configuration.....	442
Removing members from a zone configuration.....	443
Enabling a zone configuration.....	443
Disabling a zone configuration.....	444
Validating zone members in the zone configuration.....	444
Deleting a zone configuration.....	446

Abandoning zone configuration changes.....	447
Viewing all zone configuration information.....	448
Viewing selected zone configuration information.....	448
Viewing the zone aliases in the zone configuration.....	449
Viewing the configuration in the effective zone database.....	451
Clearing all zone configurations.....	451
Zone object maintenance.....	452
Copying a zone object.....	452
Deleting a zone object.....	452
Renaming a zone object.....	453
Zone configuration management.....	454
Security and zoning.....	454
Zone merging.....	455
Fabric segmentation and zoning.....	456
Zone merging scenarios.....	457
Concurrent zone transactions.....	462
Viewing zone database transactions.....	463
Peer Zoning.....	464
Peer Zoning compared to other zoning types.....	464
Advantages of Peer Zoning.....	467
Peer Zone connectivity rules.....	467
Firmware upgrade and downgrade considerations for Peer Zoning.....	468
LSAN and QoS Peer Zoning considerations.....	468
Peer Zone configuration.....	468
Alias support for peer zoning.....	470
Displaying the property members of peer zones.....	473
Target Driven Zoning.....	473
Limitations and considerations for Target Driven Zoning.....	474
Target Driven Zoning configuration.....	474
Alias support for target driven zoning.....	479
Supported commands for Peer Zones and Target Driven Peer Zones.....	480
Boot LUN zoning.....	480
Setting up boot LUN zoning.....	482
Traffic Isolation Zoning.....	484
Traffic Isolation Zoning overview.....	484
TI zone failover	485
FSPF routing rules and traffic isolation.....	486
Enhanced TI zones.....	488
Invalid configurations with enhanced TI zones.....	489
General rules for TI zones.....	491
Additional configuration rules for enhanced TI zones.....	492
Limitations and restrictions of Traffic Isolation Zoning.....	493
Creating a TI zone.....	493
Modifying TI zones.....	494
Changing the state of a TI zone.....	495
Deleting a TI zone.....	495
Displaying TI zones.....	496
Enforcing Local TI Filtering.....	497
Traffic Isolation Zoning over FC routers.....	497
TI zones within an edge fabric.....	498

TI zones within a backbone fabric.....	499
Limitations of TI zones over FC routers.....	501
Fabric-Level Traffic Isolation in a backbone fabric.....	501
Fabric-Level TI zones.....	502
Failover behavior for Fabric-Level TI zones.....	503
Creating a separate TI zone for each path.....	503
Creating a single TI zone for all paths.....	504
Setting up TI zones over FCR (sample procedure).....	505
Virtual Fabrics considerations for Traffic Isolation Zoning.....	509
Traffic Isolation Zoning over FC routers with Virtual Fabrics.....	511
Troubleshooting TI zone routing problems.....	513
TI Zone violation handling for trunk ports.....	514
Routing Traffic.....	515
Routing overview.....	515
Paths and route selection.....	515
FSPF.....	516
Fibre Channel NAT.....	517
Inter-switch links.....	517
Buffer credits.....	519
Congestion versus over-subscription.....	519
Virtual channels.....	519
Gateway links.....	520
Configuring a link through a gateway.....	521
Routing policies.....	522
Considerations for routing policies.....	522
Displaying the current routing policy.....	522
Port-based routing.....	523
Exchange-based routing.....	523
Device-based routing.....	523
Dynamic Path Selection.....	523
Route selection.....	524
Dynamic Load Sharing.....	525
Frame order delivery.....	525
Forcing in-order frame delivery across topology changes.....	526
Restoring out-of-order frame delivery across topology changes.....	526
Enabling Frame Viewer.....	526
Using Frame Viewer to understand why frames are dropped.....	527
Lossless Dynamic Load Sharing on ports.....	528
Lossless core.....	530
Configuring Lossless Dynamic Load Sharing.....	530
Lossless Dynamic Load Sharing in Virtual Fabrics.....	530
Two-hop Lossless DLS route update.....	531
Frame Redirection.....	532
Creating a frame redirect zone.....	533
Deleting a frame redirect zone.....	533
Viewing frame redirect zones.....	534
Using FC-FC Routing to Connect Fabrics.....	535
FC-FC routing overview.....	535
License requirements for FC-FC routing.....	535

Supported platforms for FC-FC routing.....	536
Supported configurations for FC-FC routing.....	539
Fibre Channel routing concepts.....	539
Proxy devices.....	544
FC-FC routing topologies.....	545
Phantom domains.....	545
FC router authentication	549
Setting up FC-FC routing.....	549
Verifying the setup for FC-FC routing.....	550
Backbone fabric IDs.....	551
Assigning backbone fabric IDs.....	552
Assigning alias names to fabric IDs.....	552
FCIP tunnel configuration.....	553
Inter-fabric link configuration.....	553
Configuring an IFL for both edge and backbone connections.....	553
Configuring EX_Ports on an ICL.....	557
FC router port cost configuration.....	559
Port cost considerations.....	559
Setting router port cost for an EX_Port.....	560
Shortest IFL cost configuration.....	561
Configuring shortest IFL cost.....	564
EX_Port frame trunking configuration.....	566
LSAN zone configuration.....	566
Zone definition and naming.....	566
LSAN zones and fabric-to-fabric communications.....	567
Controlling device communication with the LSAN.....	567
Configuring interconnectivity between backbone and edge fabrics.....	571
Setting the maximum LSAN count.....	572
HA and downgrade considerations for LSAN zones.....	573
LSAN zone policies using LSAN tagging.....	574
LSAN zone binding.....	579
Location embedded LSAN zones.....	583
Creating location embedded LSAN zones.....	584
Migrating LSAN zones to location embedded LSAN zones in the edge fabric.....	586
Limitations of location embedded LSAN zones.....	586
Peer LSAN zone support.....	587
Proxy PID configuration.....	587
Fabric parameter considerations.....	588
Inter-fabric broadcast frames.....	588
Displaying the current broadcast configuration.....	588
Enabling broadcast frame forwarding.....	588
Disabling broadcast frame forwarding.....	589
Resource monitoring.....	589
FC-FC routing and Virtual Fabrics.....	590
Logical switch configuration for FC routing.....	591
Backbone-to-edge routing with Virtual Fabrics.....	593
Upgrade and downgrade considerations for FC-FC routing.....	594
How replacing port blades affects EX_Port configuration.....	595
Displaying the range of output ports connected to xlate domains.....	595
FC-FC routing scalability limits.....	595

- Port Indexing.....597**
- Switch and Blade Sensors..... 599**
 - Brocade switch sensors..... 599
 - Brocade blade temperature sensors.....600
 - System temperature monitoring..... 600
- Hexadecimal Conversion.....601**
 - Hexadecimal overview..... 601
 - Example conversion of the hexadecimal triplet 0x610600..... 601
 - Decimal-to-hexadecimal conversion table..... 602
- Revision History..... 604**
 - FOS-821-AG101; September 28, 2018..... 604
 - FOS-821-AG100; August 28, 2018..... 604

Introduction

- About This Guide..... 18
- What's New in This Document..... 18
- Supported Hardware and Software..... 18
- Contacting Brocade Technical Support..... 19
- Document Feedback..... 20

About This Guide

This document describes the administration of Brocade Fibre Channel Services (FCS) features and the corresponding configuration tasks.

What's New in This Document

This document includes new and modified information for the Fabric OS 8.2.1 release.

Supported Hardware and Software

Although many different software and hardware configurations are tested and supported by Brocade for Fabric OS 8.2.1, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by Brocade Fabric OS 8.2.1.

Brocade Gen 5 (16Gb/s) Fixed-Port Switches

- Brocade 6505 Switch
- Brocade 6510 Switch
- Brocade 6520 Switch
- Brocade M6505 Blade Server SAN I/O module
- Brocade 6542 Blade Server SAN I/O module
- Brocade 6543 Blade Server SAN I/O module
- Brocade 6545 Blade Server SAN I/O module
- Brocade 6546 Blade Server SAN I/O module
- Brocade 6547 Blade Server SAN I/O module
- Brocade 6548 Blade Server SAN I/O module
- Brocade 6558 Blade Server SAN I/O module
- Brocade 7840 Extension Switch

Brocade Gen 5 (16Gb/s) Directors

For ease of reference, Brocade chassis-based storage systems are standardizing on the term “director”. The legacy term “backbone” can be used interchangeably with the term “director”.

- Brocade DCX 8510-4 Director
- Brocade DCX 8510-8 Director

Brocade Gen 6 (32Gb/s) Fixed-Port Switches

- Brocade G610 Switch
- Brocade G620 Switch
- Brocade G630 Switch
- Brocade 7810 Extension Switch

Brocade Gen 6 (32Gb/s) Directors

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

Contacting Brocade Technical Support

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

Online	Telephone
<p>For nonurgent issues, the preferred method is to go to MyBrocade (my.brocade.com) and then go to one of the following sites:</p> <ul style="list-style-type: none"> • My Cases • Software Downloads • Licensing tools • Knowledge Base 	<p>Required for Severity 1-Critical and Severity 2-High issues:</p> <ul style="list-style-type: none"> • North America: 1-800-752-8061 (Toll-free) • International: 1-669-234-1001 (Not toll-free) <p>Toll-free numbers are available in many countries and are listed at https://www.broadcom.com/support/fibre-channel-networking/.</p>

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

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

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, publication number, topic heading, page number, and as much detail as possible.

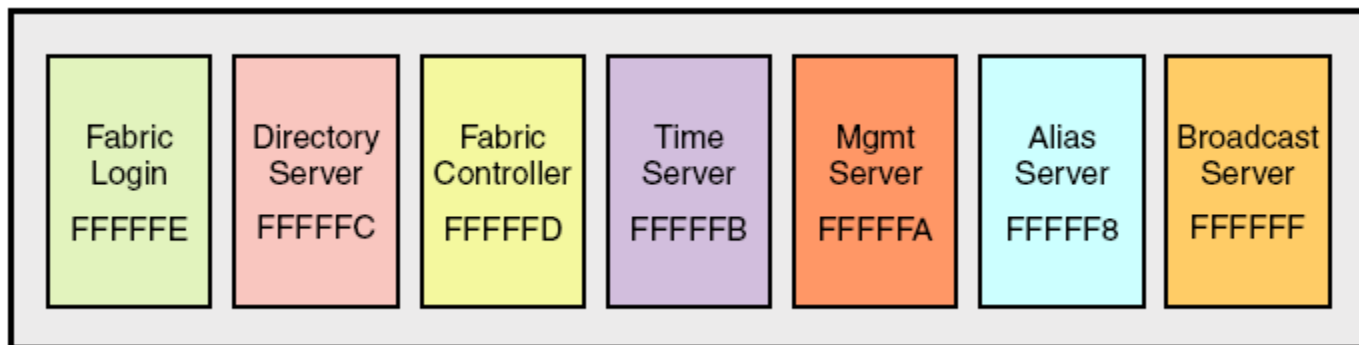
Understanding Fibre Channel Services

- Fibre Channel Services Overview..... 21
- Management Server..... 22
- Platform Services..... 22
- Management Server Database..... 23
- Topology Discovery..... 27
- Device Login..... 28
- High Availability of Daemon Process..... 30
- FL_Port and Arbitrated Loop Support..... 31

Fibre Channel Services Overview

Fibre Channel defines service functions that reside 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 login — 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 (FCS) switch. See [Configuring Security Policies](#) on page 233 for additional information on FCS policies.
- Management server — The management server provides a single point for managing the fabric. This is the only service that users can configure. See [Management Server](#) on page 22 for more details.
- Alias server — The alias server keeps a group of nodes 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 identifying switches can be registered with the management server.

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

- It is accessed by an external Fibre Channel node at the well-known address *FFFFFFAh*, 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.

Application Server

The application server is a key component of the Virtual Machine (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 only supported by Gen6 platforms running Fabric OS 8.1.0 and later. It will not attempt inter-switch communication with Gen4 or Gen5 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* 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.

NOTE

The **msplMgmtActivate** and **msplMgmtDeactivate** commands are allowed only in ADO and AD255 domains.

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 platform service on all logical switches in a Virtual Fabric. The **msPlatShow** command displays all platforms registered in a Virtual Fabric.

Enabling Platform Services

When FCS policy is enabled, the **msplMgmtActivate** 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 **msplMgmtActivate** command, as in the following example.

```
switch:admin> msplmgmtactivate
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 **msplMgmtDeactivate** command.
3. Enter **y** to confirm the deactivation, as in the following example.

```
switch:admin> msplmgmtdeactivate

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

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. 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 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 viewing 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.

NOTE

The **msPIClearDB** command is allowed only in ADO and AD255.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **msplClearDb** command.
3. Enter **y** to confirm the deletion.

The management server platform database is cleared.

Topology Discovery

The topology discovery feature can be displayed, enabled, and disabled; it is disabled by default. The **mstdEnable** and **mstdDisable** commands are allowed only in ADO and AD255 domain.

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 the appropriate following command 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 the appropriate following command 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 management server topology might erase all node ID entries.

Example of disabling discovery

The following is an example of disabling topology discovery.

```
switch:admin> mstdisable
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> mstdisable 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 – 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 – 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 that the new switch is an E_Port 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 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 then it 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 are 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 and establishes the port type and negotiates 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 entered 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** command. For more information, refer to *Brocade Fabric OS Command Reference*.

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 has changed or an IP address has changed.
- Nodes leaving or joining 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](#) on page 110.

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 RAS log is generated.

Device Recovery

To recover devices that have been removed from the Name Server database due to duplicate PWWNs, you must login 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 Daemon Process

Starting non-critical daemons is automatic; you cannot configure the startup process. The following sequence of events occurs when a non-critical 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 reporting 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 non-critical daemon. You can schedule the downtime and reboot the switch at your convenience.

The following table lists the daemons that are considered non-critical and are automatically restarted on failure.

TABLE 1 Automatic Restart Daemons

Daemon	Description
arrd	Asynchronous Response Router, which is used to send management data to hosts when the switch is accessed through the APIs (FA API or SMI-S).
cald	Common Access Layer daemon, which is used by manageability applications.
raslogd	Reliability, Availability, and Supportability daemon logs error detection, reporting, handling, and presentation of data into a format readable by you and management tools.
rpcd	Remote Procedure Call daemon, which is used by the API (Fabric Access API and SMI-S).
snmpd	Simple Network Management Protocol daemon.
npd	Flow Vision daemon.
traced	Trace daemon provides the trace entry date and time translation to the Trace Device at startup and when date/time changed by command. Maintains the trace dump trigger parameters in a Trace Device. Performs the trace Background Dump, trace automatic FTP, and FTP "aliveness check" if auto-FTP is enabled.
trafd	Traffic daemon implements Bottleneck detection.
webd	Webserver daemon used for Web Tools (includes httpd as well).
weblinkerd	Weblinker daemon provides an HTTP interface to manageability applications for switch management and fabric discovery.

FL_Port and Arbitrated Loop Support

The following information applies to FL_Ports and arbitrated loops.

- FL_Ports is supported only on the BrocadeFX8-24 platform.
- FL_Ports are not supported on the BrocadeFC32-64, FC32-48, FC16-32, FC16-48, FC16-64, SX6, Brocade 6505, Brocade 6510, Brocade 6520, Brocade G610, Brocade G620, Brocade G630, or Brocade 7840 platforms.
- VE_Ports on the FX8-24 platform does not support arbitrated loops.

Performing Basic Configuration Tasks

• Fabric OS Overview.....	32
• Command Line Interface.....	32
• Modification of Default Password.....	41
• The Switch Ethernet Interface	43
• Date and Time Settings.....	56
• Domain IDs.....	61
• Registered Organization Name	63
• Switch Name.....	63
• Switch Type.....	64
• Chassis Name.....	65
• Fabric Name.....	66
• Switch Activation and Deactivation.....	67
• Device Shutdown.....	68
• Basic Connections.....	69

Fabric OS Overview

This chapter 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 backbone platform or switch and blades, and then set the IP addresses of those devices. Although this chapter 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 Web Tools Administration Guide*.
- Brocade Network Advisor – For additional information, refer to the *Brocade Network Advisor User Manual*.
- A third-party application using the API – For third-party application procedures, refer to the third-party API documentation.

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, a number of procedures that apply only to variable-port devices are found in [Performing Advanced Configuration Tasks](#) on page 71.

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems LLC, documenting all possible configurations and scenarios is beyond the scope of this document. In some cases, earlier releases are highlighted to present considerations for inter-operating with them.

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 contained in this guide. For additional information, refer to *Brocade Fabric OS Command Reference*.

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](#) on page 145.

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, **BANNERSHOW** and **BannerShow** will not. Refer to the *Brocade Fabric OS Command Reference* for explicit instructions on supported capitalization for each command.
- When command examples in this guide show user input enclosed in quotation marks, the quotation marks are required. For example, **zonecreate** "zonename" requires that the value for *zonename* be in quotation marks.

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 DCX 8510 and X6 Backbone families:** You can connect to CP0 or CP1 using either of the two serial ports.

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 **ttyb** is already in use.

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

- In a Windows environment, enter the following parameters:

TABLE 2 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

Telnet or SSH Session

You can connect to the Fabric OS through a Telnet or SSH 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](#) on page 33.

During the initial configuration login to the device, you must log in using the admin account and change all the default passwords. Once this has been done, you can use one of the following procedures to enable the root account if it is available, as by default the root account (if one is available) is disabled. Not all systems ship with root accounts. If your system does not have a root account, you will not be able to enable it.

NOTE

You must log in to the root account with the manufacturer's default password and change the root account password at the first login.

- If your device has a root account, enable this account using the following command:

```
userconfig --change root -e yes
```

- If you need to access the device remotely using the root account, enable root access over the management interface using the following command. By default, the access is set to console only.

```
rootaccess --set all
```

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. Refer to [DHCP Activation](#) on page 49.

Rules for Telnet Connections

The following rules that are followed for making Telnet connections to your switch:

- Never change the IP address of the switch while Telnet sessions are active; if you do, the existing Telnet sessions will hang and your next attempt to log in will fail. To recover, gain access to the switch by one of these methods:
 - You can use Web Tools to perform a fast boot. When the switch comes up, the Telnet quota is cleared. (For instructions on performing a fast boot with Web Tools, refer to the *Brocade Web Tools Administration 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 processes without disrupting the fabric.
- For accounts with an admin role, Fabric OS limits the number of simultaneous Telnet sessions per switch to four like other users. For more details on session limits, refer to [Managing User Accounts](#) on page 144.
- The **login** command creates nested sessions; therefore, only a maximum of five logins are allowed from the same session, after which the following session limit message is displayed. The **killtelnet** command kills only the latest login but leaves the others open.

```
Switch:admin> admin
Password:
Same session login limit reached. Logout before next login
```

Connecting to Fabric OS Using Telnet

Use the following procedure to connect to the Fabric OS using Telnet.

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 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 Telnet connection using the IP address of the switch to which you want to connect.

The login prompt is displayed when the Telnet connection finds the switch in the network.
5. Enter the account ID at the login prompt.
6. Enter the password.

If you have not changed the system passwords from the default, you are prompted to change them. Enter the new system passwords, or press **Ctrl+C** to skip the password prompts. For more information on system passwords, see [Default Account Password](#) on page 42.

7. Verify the login was successful.

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

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

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 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.

- Enter the **help** command with no specific command and 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 3 Help topic contents

Topic name	Help contents description
diagHelp	Diagnostic help information
fcipHelp	FCIP help information
ficonHelp	FICON help information
mapsHelp	MAPS help information
routeHelp	Routing help information
secHelp	Security help information
zoneHelp	Zoning help information

Viewing the History of Command Line Entry

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 is persistent across reboots and firmware downloads. This command is also supported for 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.

- Connect to the switch and log in.
- Enter the **clihistory** command with the desired argument (refer to the *Brocade Fabric OS Command Reference* for the **clihistory** arguments). Entering no specific argument displays only the command-line history of the currently logged-in user.

NOTE

Commands entered through an interactive SSH login session are recorded in the command-line history. However, if a command is entered as a single line using a SSH login (for example, `ssh admin@10.20.x.x switchshow`), the entire command is not recorded.

CLI History

Entering **cliHistory** with no keywords displays the command line history for the currently logged-in user only. This is true for the root user account as well.

The following example shows the output of the **cliHistory** command for the root login.

```
switch:root> clihistory

CLI history
Date & Time          Message
Thu Sep 27 04:58:00 2012    root, 10.70.12.101, firmwareshow -v
Thu Sep 27 04:58:19 2012    root, 10.70.12.101, telnet 127.1.10.1
Thu Sep 27 05:25:45 2012    root, 10.70.12.101, ipaddrshow]
```

The following example shows the output of the **cliHistory** command for the admin login.

```
switch:admin> clihistory

CLI history
Date & Time          Message
Thu Sep 27 10:14:41 2012    admin, 10.70.12.101, clihistory
Thu Sep 27 10:14:48 2012    admin, 10.70.12.101, clihistory --show
```

The following example shows the output of the **cliHistory** command with Fabric IDs (FID) in virtual fabric mode.

```
switch:admin> clihistory

CLI history
Date & Time          Message
Fri Sep 19 09:41:08 2014    root, FID 128, console, clihistory --clear
Fri Sep 19 09:41:41 2014    root, FID 128, console, lscfg --create 10
Fri Sep 19 09:42:10 2014    root, FID 128, console, lscfg --create 120
Fri Sep 19 09:42:54 2014    root, FID 128, console, lscfg --create 30
Fri Sep 19 09:42:59 2014    root, FID 128, console, setcontext 30
Fri Sep 19 09:43:14 2014    root, FID 30, console, rasdecode -m 65
Fri Sep 19 09:43:32 2014    root, FID 30, console, tracecfg -p -m 138
Fri Sep 19 09:43:42 2014    root, FID 30, console, setcontext 10
```

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 Sep 19 09:43:53 2014    root, FID 10, console, tracedump
Fri Sep 19 09:43:59 2014    root, FID 10, console, coreshow
Fri Sep 19 09:44:21 2014    root, , console, firmwareshow
Fri Sep 19 09:44:25 2014    root, , 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 Root, Admin, and Securityadmin RBAC role permissions. Not all systems ship with root accounts. If your system does not have a root account, you will not be able to enable it.

The following example shows the output of the **clihistory --showuser admin** for the “admin42” user account.

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

CLI history
Date & Time          Message
Thu Sep 27 10:14:41 2012    admin, 10.70.12.101, clihistory
Thu Sep 27 10:14:48 2012    admin, 10.70.12.101, clihistory --show
Thu Sep 27 10:15:00 2012    admin, 10.70.12.101, clihistory
```

cliHistory --showall

Use the **clihistory --showall** command to display the command-line history for all users. For admin and securityadmin users, this also allows them to see the user command history for the root account. This argument is available only to accounts with root, 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, 172.26.3.86, setcontext 20
Mon Oct  5 10:44:17 2015    admin01, FID 20, 172.26.3.86, firmwareshow
Mon Oct  5 10:44:32 2015    admin01, FID 20, 172.26.3.86, islshow
Mon Oct  5 10:44:20 2015    admin01, FID 20, 172.26.3.86, version
Mon Oct  5 10:44:36 2015    admin01, FID 20, 172.26.3.86, switchshow
Mon Oct  5 10:44:40 2015    admin01, FID 20, 172.26.3.86, 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:

- SSH login CLI logs are *not* recorded in the command-line history (as shown in the following example):

```
ssh user@Ipaddr switchshow
```

- 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 10.70.4.109,fvt,/dist,pray4green
Server IP: 10.70.4.109, Protocol IPv4
Checking system settings for firmwaredownload...
Failed to access scp://fvt:*****@10.70.4.109//dist/release.plist

switch:admin> clihistory
```

```
Date & Time          Message
Wed Nov 23 03:39:37 2016    root, console, firmwaredownload
```

From 8.2.0 and later, for security reasons, you can configure the switch and are not allowed to record the CLI usage to the shell CLI buffer. The shell CLI history buffer is not persistent and is removed after you log out and the CLI history is no longer displays the user input.

The following example shows the configuration CLI to disable recording the user input to the shell CLI buffer.

```
switch:admin> clihistory --shelldisable
CLI history shell CLI buffer is disabled.
switch:admin> clihistory --shellenable
CLI history shell CLI buffer is enabled.
switch:admin> clihistory --showshell
CLI history buffer is disabled.
switch:admin> clihistory --enable -shellbuffer
CLI history Shell CLI buffer is enabled.
switch:admin> clihistory --disable -shellbuffer
CLI history shell CLI buffer is disabled.
switch:admin> clihistory --display - shellbuffer
CLI history shell CLI buffer is disabled.
```

Using Fosexec to Run Commands on Remote Switch or Domain

The **fosexec** 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 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 **fosexec** are also captured in the CLI history and audit logs of the remote switch.

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

The **configure** command has the following option under Configure CLI - Fabric Parameters section to enable or disable the remote **fosexec** feature. 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\""
```

In Fabric OS prior to 8.2.0, when executing the **portcfgshow** command with **fosexec** feature for the Access Gateway (AG) configured switch, it ends up displaying the message `RBAC permission denied`. 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:

- **fosexec** requires Fabric OS 7.3.0 or later on both the local and remote domains or switches.
- 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.
- It can be executed only by users with the `fabricadmin`, `admin`, or `root` roles or privileges and RBAC permissions. Not all systems ship with root accounts. If your system does not have a root account, you will not be able to enable it.
- It does not work during remote domain or switch HA failover or reboot.
- Command syntax must match the command syntax supported by the remote switch.
- Commands with parameters such as slot/port specific to a fixed-port switch or Director chassis are not supported with the **all** option.
- It 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.
- It returns a maximum of 64 kilobytes of data from the remote switch. Any data beyond this limit is truncated.
- It is not supported in FIPS and Access Gateway modes.
- When you run **supportsave** through **fosexec**, wait for the execution to complete. To know whether the execution is completed, refer to the `RASLOG [SS-1000]` on all or the specific domain where **fosexec** is being run.
- The **fosexec** feature must be disabled to downgrade from Fabric OS 7.3.0 to any earlier version.

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.

```
sw_85:user9> fosexec --domain all -cmd "islshow"
Domain 19
=====
 1: 80-> 40 10:00:00:05:1e:0f:73:40 53 ls20d53_5100_8_sp:----- bw: 8.000G
 2: 81->289 10:00:00:05:33:0d:7b:02 4 ls20d4_plpl_8_4 sp:----- bw: 8.000G

Domain 23
=====
Remote fosexec feature is disabled.

Domain 53
=====
 1: 40-> 80 10:00:00:05:33:39:0d:da 19 ls20d19_5300_8_sp:----- bw: 8.000G
 2: 41->288 10:00:00:05:33:0d:7b:02 4 ls20d4_plpl_8_4 sp:----- bw: 16.000G

Domain 65
=====
Remote fosexec feature is disabled.
Domain 150
=====
 1: 0-> 38 10:00:00:27:f8:f1:15:c2 1 ls20d1_wedge_8_sp: 16.000G bw: 16.000G TRUNK QOS CR_RECOV FEC

Domain 1
=====
 1: 3-> 10 10:00:00:05:33:7a:4b:76 65 ls20d65_6510_8_47_58 sp: 8.000G bw: 16.000G TRUNK QOS
CR_RECOV
 2: 5-> 13 10:00:00:05:33:e5:56:2c 2 ls20d2_odin_8_39_58 sp: 16.000G bw: 16.000G TRUNK QOS
CR_RECOV FEC
 3: 6-> 42 10:00:00:05:33:0d:7b:02 4 ls20d4_plpl_8_47_30 sp: 16.000G bw: 16.000G TRUNK QOS
CR_RECOV FEC
 4: 22-> 12 10:00:00:05:33:e5:56:2c 2 ls20d2_odin_8_39_58 sp: 16.000G bw: 16.000G TRUNK QOS
CR_RECOV FEC
 5: 38-> 0 10:00:50:eb:1a:64:63:1f 150 ls20d150_amp_38_37_113 sp: 16.000G bw: 16.000G TRUNK QOS
CR_RECOV FEC
 6: 42-> 39 10:00:00:05:33:7a:4b:76 65 ls20d65_6510_8_47_58 sp: 8.000G bw: 8.000G TRUNK QOS
CR_RECOV
```

Example 2

This example shows the output of running **fosexec** only on domain 19.

```
sw_85:user9> fosexec --domain 19 -cmd "islshow"
 1: 80-> 40 10:00:00:05:1e:0f:73:40 53 ls20d53_5100_8_sp:----- bw: 8.000G
 2: 81->289 10:00:00:05:33:0d:7b:02 4 ls20d4_plpl_8_4 sp:----- bw: 8.000G
```

Modification of Default Password

The switch automatically prompts you to change the default account passwords after logging in for the first time. If you do not change the passwords, the switch prompts you at each subsequent login until all the default passwords have been changed.

There is one set of default accounts for the entire chassis: admin, user, and root. 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. The root account is reserved for development and manufacturing or for debugging. Not all systems ship with root accounts. If your system does not have a root account, you will not be able to enable it. For more information on the default accounts, refer to [Default accounts](#) on page 149.

NOTE

The passwords for the default accounts should immediately be changed from their default values when prompted on the initial login; these passwords cannot be changed using the **passwd** command later in the session unless you have done this. If you skip this step, and then later decide to change the passwords, you will need to log out and then log back in.

NOTE

When you log in as admin and the default passwords have not been changed from the manufacturer's defaults, you will be prompted to change the default passwords only for the admin and user accounts. You cannot change the default password for the root account if you log in as admin. The default password for the root account can be changed when you log in as root. When you initially log in to the root account (if one is available), you will be forced to change the default password.

Default Account Password

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

The following table lists the supported characters for default passwords

TABLE 4 Supported Characters for Default Password

Name of the Character	Character
lower-case alphabets	a-z
Numerals	0-9
period	.
underscore	-
space	

Record the passwords exactly as entered and store them in a secure place because recovering passwords requires significant effort and fabric downtime. Although the root account is not meant for general use, change their passwords if prompted to do so and save the passwords in case they are needed for recovery purposes.

For more information, refer to [Password policies](#) on page 152.

Changing the Default Account Password at Login

Use the following procedure to change the default account passwords.

1. Connect to the switch and log in using the default administrative account.
2. At each of the "Enter new password" prompts, either enter a new password or skip the prompt.

To skip a single prompt, press **Enter**. To skip all of the remaining prompts, press **Ctrl-C**.

NOTE

For root login, you cannot skip this step if the default password is present. You must change the default password.

NOTE

For the first time, the default password of the admin user can be changed only from the root login.

The following example shows the output for changing passwords.

```
login: admin

Password:
Please change your passwords now.
Use Control-C to exit or press 'Enter' key to proceed.
Enter new password: <hidden>

Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
(output truncated)
```

The Switch Ethernet Interface

The Ethernet (network) interface provides management access, including direct access to the Fabric OS CLI, and 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 DCX 8510 and X6 Directors, you must set IP addresses for the following components:

- Both Control Processors (CP0 and CP1)
- Chassis management IP

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 (refer to [Static Ethernet Addresses](#) on page 47)
- Activating DHCP (refer to [DHCP Activation](#) on page 49)

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, refer to [Connecting to Fabric OS through the Serial Port](#) on page 33.

Virtual Fabrics and the Ethernet interface

On the Brocade DCX 8510 Backbone 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](#) on page 382.

Management Ethernet Port Bonding

The two external Ethernet ports of a CP (CP8 or CPX6) 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 the Fabric OS kernel modules and applications on the CP blade will 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. Use 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 Backbones and Directors

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

- CP8 blade installed on a Brocade DCX 8510 Backbone
- CPX6 blade installed on a Brocade X6 Director

Setting Up the Second Ethernet Port on a CP8 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. Refer to [Setting the Ethernet Interface Mode and Speed](#) on page 55 for instructions on setting port modes, and [Setting port speeds](#) on page 87 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, refer to [Console Sessions Using the Serial Port](#) on page 33. Otherwise, connect using SSH.

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

2. Enter `ipAddrShow`.

```
ipAddrShow
```

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

```
ecp:admin> ipaddrshow

SWITCH
Ethernet IP Address: 10.1.2.3
Ethernet Subnetmask: 255.255.240.0
CP0
Ethernet IP Address: 10.1.2.3
Ethernet Subnetmask: 255.255.240.0
Host Name: ecp0
Gateway IP Address: 10.1.2.1
CP1
Ethernet IP Address: 10.1.2.4
Ethernet Subnetmask: 255.255.240.0
Host Name: ecp1
Gateway IP Address: 10.1.2.3
IPFC address for virtual fabric ID 123: 11.1.2.3/24
IPFC address for virtual fabric ID 45: 13.1.2.4/20
Slot 7
eth0: 11.1.2.4/24
Gateway: 11.1.2.1
Backplane IP address of CP0 : 10.0.0.5
Backplane IP address of CP1 : 10.0.0.6
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
sw 0 stateless fd00:60:69bc:70:260:69ff:fe00:2/64 preferred
sw 0 stateless fec0:60:69bc:70:260:69ff:fe00:2/64 preferred
cp 0 stateless fd00:60:69bc:70:260:69ff:fe00:197/64 preferred
cp 0 stateless fec0:60:69bc:70:260:69ff:fe00:197/64 preferred
cp 1 stateless fd00:60:69bc:70:260:69ff:fe00:196/64 preferred
cp 1 stateless fec0:60:69bc:70:260:69ff:fe00:196/64 preferred
IPv6 Gateways:
cp 0 fe80:60:69bc:70::3
cp 0 fe80:60:69bc:70::2
cp 0 fe80:60:69bc:70::1
cp 1 fe80:60:69bc:70::3
```

If the Ethernet IP address, subnet mask, and gateway address are displayed, then the network interface is configured. Verify 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 Management Network Interface Error Counter

You can reset the management network interface error counter on the Brocade 7840 Extension Switch and Brocade Gen 6 devices. To reset the error counter, complete the following steps.

1. Enter `ethIf --reseterror interface ID`.

This command clear the counter and display the completion message.

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

- To confirm that the error counter have been cleared, enter **ethif --show interface ID**.

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

```
switch:admin> ethif --show eth0
eth0 interface:
Link mode: negotiated 100baseTx-FD, link ok
MAC Address: 59:AB:AA:B5:75:BC
eth0      Link encap:Ethernet  HWaddr 59:AB:AA:B5:75:BC
          inet addr:10.x.x.x  Bcast:10.x.x.x  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:382496  errors:0  dropped:0  overruns:0  frame:0
          TX packets:393  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0  txqueuelen:1000
```

Management Port Ingress Rate Limiting

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

- Over utilizing the network bandwidth
- Over utilizing the CPU's bandwidth
- Over utilizing 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 directors to prevent the CPU from stalling or overwhelming during a DoS attack.

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

Configuring the Management Port Ingress Rate Limiting

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

NOTE

This is only applicable or available 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, refer to [DHCP Activation](#) on page 49.

If you choose not to use DHCP or to specify an IP address for your switch Ethernet interface, you can do so 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 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 the Static Address for the Ethernet Network Interface

Use the following procedure to set the ethernet network interface static addresses.

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 backbone:
 - If you are setting the IP address for a switch, enter the **ipAddrSet** command.
 - If you are setting the IP address for a backbone, enter the **ipAddrSet** command specifying either CPO or CP1. You must set the IP address for both CPO and CP1.

The following is an example of setting an IPv4 address.

```
switch:admin> ipaddrset cp0

Ethernet IP Address [10.1.2.3]:
Ethernet Subnetmask [255.255.255.0]:
Fibre Channel Subnetmask [255.255.0.0]:
Gateway IP Address [10.1.2.1]:
DHCP [OFF]: off
```

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](#) on page 382.

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. The Fibre Channel prompts are not relevant; you can skip them by pressing **Enter**.
The Fibre Channel IP address is used for management.
6. Enter the **Gateway IP Address**.
7. Enter **off** to disable DHCP.

Setting the Static Address for the Chassis Management IP Interface

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

1. Connect to the switch and log in using an account assigned to the admin role.

2. Enter the **ipAddrSet -chassis** command.

```
switch:admin>> ipaddrset -chassis
DHCP [Off]:
Ethernet IP Address [10.38.18.142]:
Ethernet Subnetmask [255.255.240.0]:
Committing configuration...Done.
```

3. Enter the network information either in dotted-decimal notation for the Ethernet IPv4 address or in semicolon-separated notation for IPv6.
4. Enter the value for **Ethernet Subnet mask**.

DHCP Activation

The Dynamic Host Configuration Protocol (DHCP) is enabled on Brocade switches by default.

TheFabric OS DHCP client conforms to the latest IETF Draft Standard RFCs for IPv4, IPv6, and DHCP. The following parameters are:

- 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 uses a DHCP vendor-class identifier that allows DHCP servers to determine that the discover/request packet is coming from a Brocade switch. The vendor-class identifier is the string "BROCADE" followed by the SWBD model number of the platform. For example, the vendor-class identifier for a request from a Brocade G620 is "BROCADESWBD145."
- 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. If the switch name is changed using the CLI, the DNS record in the DNS server is updated automatically. Dynamic DNS (DDNS) is the method used to automatically update a name server in the DNS.

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, refer to [Static Ethernet Addresses](#) on page 47.

1. Connect to the switch and log in using an account assigned to the admin role.

2. Enter the **ipAddrSet** command.

```
ipaddrset
```

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. You can 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: 10.20.134.219
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.20.128.1
DHCP: On
```

Disabling DHCP for IPv4 in Fixed-Port Devices

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

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

```
ipaddrset
```

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 follow the subsequent steps.

3. Enter **Off** to disable DHCP.

4. You can 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: 10.20.134.219
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.20.128.1
DHCP: Off
```

Enabling DHCP for IPv4 in Directors

To enable DHCP for IPv4 on Brocade DCX 8510 and Brocade X6 Directors, follow these steps.

1. Run the **ipAddrSet -chassis** command.

```
switch:admin> ipaddrset -chassis
DHCP [Off]:on
Info : This feature got enabled on all Management Interfaces of this chassis
Done.
```

2. To display the configured IP address, run the **ipAddrShow** command.

```
switch:admin>ipaddrshow
CHASSIS
Ethernet IP Address: 10.17.46.231
Ethernet Subnetmask: 255.255.240.0

CP0
Ethernet IP Address: 10.17.46.228
Ethernet Subnetmask: 255.255.240.0
Host Name: cp0
Gateway IP Address: 10.17.32.1

CP1
Ethernet IP Address: 10.17.46.229
Ethernet Subnetmask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.17.32.1
```

3. To disable DHCP for IPv4 on Directors, follow these steps.

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

- In interactive mode:

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

- In a single command:

```
switch:admin>ipaddrset -chassis -ipv4 -add -dhcp off
IP address is being changed...
Done.
```

Configuring DHCPv6 for IPv6

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

1. Run the **ipAddrSet -ipv6 -dhcpv6** command.
 2. To display the configured IPv6 address, run the **ipAddrShow** command.
- In case of fixed-port switches, the following output is displayed:

```
switch:admin>ipaddrshow
SWITCH
Ethernet IP Address: 10.17.33.33
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.17.32.1
DHCP: Off
IPv6 Autoconfiguration Enabled: No
Local IPv6 Addresses:
dhcpv6 2620:100:0:f603::f9ff/64 preferred
link local fe80::205:33ff:fe3a:77f9/64
IPv6 Gateways:
fe80::21b:edff:fe0b:9000
DHCPv6: On
```

- In case of directors, the following output is displayed:

```
switch:admin>ipaddrshow
CHASSIS
Local IPv6 Addresses:
dhcpv6 2620:100:0:f603::f912/64 preferred
link local fe80::205:33ff:fe3a:77f9/64

CPO
Local IPv6 Addresses:
dhcpv6 2620:100:0:f603::f9fb/64 preferred
link local fe80::205:33ff:fe3a:77f8/64
IPv6 Gateways:
fe80::21b:edff:fe0b:9000
DHCPv6: On

CP1
Local IPv6 Addresses:
dhcpv6 2620:100:0:f603::f9fa/64 preferred
link local fe80::205:33ff:fe3a:77f7/64
IPv6 Gateways:
fe80::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. Run the **ipAddrSet -ipv6 -auto** command.
2. To display the configured IP address details, run the **ipAddrShow** command.
 - In case of fixed-port switches, the following output is displayed:

```
switch:admin> ipaddrshow
SWITCH
Ethernet IP Address: 10.17.33.33
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.17.32.1
DHCP: Off
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
dhcpv6 2620:100:0:f603::f9ff/64 preferred
stateless 2620:100:0:f602:227:f8ff:fe0:1800/64 preferred
link local fe80::205:33ff:fe3a:77f9/64
IPv6 Gateways:
fe80::21b:edff:fe0b:9000
DHCPv6: On
```

- In case of Directors, the following output is displayed:

```
switch:admin> ipaddrshow
CHASSIS
Ethernet IP Address: 10.17.46.217
Ethernet Subnetmask: 255.255.240.0

CP0
Ethernet IP Address: 10.17.46.240
Ethernet Subnetmask: 255.255.240.0
Host Name: cp0
Gateway IP Address: 10.17.32.1

CP1
Ethernet IP Address: 10.17.46.241
Ethernet Subnetmask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.17.32.1
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
chassis 0 dhcpv6 2620:100:0:f603::f1a2/64 preferred
cp 0 link local fe80::205:1eff:febf:7069/64
chassis 0 stateless 2620:100:0:f603:205:1eff:feb7:3c00/64 preferred
cp 0 static 2620:100:0:f603::f1a3/64 preferred
cp 0 stateless 2620:100:0:f603:205:1eff:febf:7069/64 preferred
cp 0 dhcpv6 2620:100:0:f603::f773/64 preferred
cp 1 link local fe80::205:1eff:febf:7068/64
cp 1 stateless 2620:100:0:f603:205:1eff:febf:7068/64 preferred
cp 1 dhcpv6 2620:100:0:f603::f49f/64 preferred
IPv6 Gateways:
cp 0 fe80::21b:edff:fe0b:3c00
cp 0 fe80::21b:edff:fe0b:9000
DHCPv6: On
```

- 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 = 2620:100:4:e200::10
Name Server IP Address = 2620:100:4:e200::11
```

Configuring the Static IPv6 Gateway Address

You can configure the static IPv6 gateway address using the **ipaddrset -ipv6 -add** command. When the static IPv6 gateway is configured, the 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 always must have the same gateway IP.
- When you delete the static IPv6 gateway address, a chassis reboot is needed for Directors and a reboot for fixed-port switches for the auto-configured 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*.

- Configure the static IPv6 gateway address using the following syntax:

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

- 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 only accessible 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 is 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, though 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 following 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

In Gen 6 device, if this is set to "auto", your manually configured DNS servers will be overwritten by any IPv6 DNS broadcasts after a failover or reboot.

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 backbone, the supported interfaces are eth0 and eth3. On all other platforms, only eth0 is supported.

For dual-CP systems, the `ethlf` 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 `ethlf` command on the active CP; and to set the mode on the standby CP, you must issue the `ethlf` 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 `ethlf` 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 `ethlf` command.

Example of setting the port mode to full auto-negotiate

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

```
switch:admin> ethlf --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 auto-negotiation 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.

In a virtual fabric, there can be a maximum of eight logical switches per backbone. Only the default switch in the chassis can update the hardware clock. When the **date** command is issued from a non-principal pre-Fabric OS v6.2.0 or earlier switch, the **date** command request is dropped by a Fabric OS v6.2.0 and later switch and the pre-Fabric OS v6.2.0 switch or earlier does not receive an error.

Authorization access to set or change the date and time for a switch is role-based. For an understanding of role-based access, refer to [Role-Based Access Control](#) on page 145.

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 "mmddHHMMyy"
```

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 Sep 29 17:01:48 UTC 2007
Stealth200E:admin> date "0204101008"
Mon Feb 4 10:10:00 UTC 2008
```


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. For more information about the **tsTimeZone** command, refer to the *Brocade Fabric OS Command Reference*.

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 backbone.
 - The time zone of the entire backbone 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 only need to set the time zone 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 review or 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/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

switch:admin> tstimezone

Time Zone : US/Central
```

Setting the Time Zone Interactively

Use the following procedure to set the current time zone to PST 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
11) none - I want to specify the time zone using the POSIX TZ format.
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
7) Cameroon         24) Lesotho       41) Spain
8) Central African Rep. 25) Liberia       42) Sudan
9) Chad             26) Libya         43) Swaziland
10) Congo (Dem. Rep.) 27) Malawi        44) Tanzania
11) Congo (Rep.)     28) Mali          45) Togo
12) Cote d'Ivoire    29) Mauritania    46) Tunisia
13) Djibouti        30) Morocco       47) Uganda
14) Egypt           31) Mozambique    48) Western Sahara
15) Equatorial Guinea 32) Namibia       49) Zambia
16) Eritrea         33) Niger         50) 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 Oct 28 15:01:38 CET 2015.
Universal Time is now:  Wed Oct 28 14:01:38 UTC 2015.
```

4. Select a country location at the prompt.

- 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.

```

Is the above information OK?
1) Yes
2) No
Enter number or control-D to quit ?2
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
11) none - I want to specify the time zone using the POSIX TZ format.
Enter number or control-D to quit ?

```

Network Time Protocol

To keep the time in your SAN current, you should synchronize the local time of the principal or primary FCS switch with at least one external Network Time Protocol (NTP) server.

The principal or primary FCS switch connects to the NTP server and broadcasts time service updates to all switches in the fabric. The other switches in the fabric automatically take their time from the principal or primary FCS switch.

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

If virtual fabrics is 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 time does not go out of sync in the chassis. Do not configure LOCL in the NTP server list.
- All default switches in the fabric can query the NTP server. If Virtual Fabrics is not enabled, only the principal or primary FCS switch can query the NTP server.
- The logical switches in a 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 server addresses in IPv4, IPv6, or DNS name formats. When multiple NTP server addresses are passed, **tsClockServer** sets the first obtainable address as the active NTP server. The rest are stored as backup servers that can take over if the active NTP server fails. The principal or primary FCS switch synchronizes its time with the NTP server every 64 seconds.

- Connect to the switch and log in using an account assigned to the admin role.

2. Enter the **tsClockServer** command.

```
switch:admin> tsclockserver ntp1; ntp2
```

In this syntax, *ntp1* is the IP address or DNS name of the first NTP server, which the switch must be able to access. The second variable, *ntp2*, is the second NTP server and is optional. The operand "*ntp1;ntp2*" is optional; by default, this value is LOCL, which uses the local clock of the principal or primary FCS switch as the clock server.

The following example sets the NTP server.

```
switch:admin> tsclockserver
LOCL
switch:admin> tsclockserver "10.1.2.3"
```

The following example displays the NTP server.

```
switch:admin> tsclockserver
10.1.2.3
```

The following example sets up more than one NTP server using a DNS name.

```
switch:admin> tsclockserver "10.1.2.4;10.1.2.5;ntp.localdomain.net"
Updating Clock Server configuration...done.
Updated with the NTP servers
```

Changes to the clock server value on the principal or primary FCS switch are propagated to all switches in the fabric.

NOTE

Running the **reboot** command on a fixed-port switch syncs the NTP time entries but not in a chassis. Instead, the **power -cycle** command syncs the NTP time entries in a chassis-based system.

NTP Configuration Distribution to Access Gateways

Any switch running Fabric OS 7.4.0 or later can distribute the NTP server configuration to core Access Gateway (AG) devices connected to the fabric. The core AG devices must be running Fabric OS 7.4.0 or later to 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.

A switch running Fabric OS 7.4.0 or later can distribute the NTP server configuration to AG devices connected to the fabric. However, AG devices cannot distribute 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. An AG device need not sync with the external NTP server since it can receive the NTP server configuration from the connected Fabric OS switch. 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.
- 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.
- 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, the AG device cannot distribute any configuration to the 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 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's 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 needs to have its domain ID changed to a domain ID not used within the fabric.

NOTE

The default domain ID for Brocade switches is 1. To avoid future conflicts and unexpected changes when additional switches are enabled that might be assigned the default ID, it is recommended that you change the default ID to an insistent domain ID.

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 must be running Fabric OS 7.4.0 or later.

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

- If 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, and TI 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	10.3.220.1	0.0.0.0	"ras001"
3: fffc03	10:00:00:60:69:e0:01:47	10.3.220.2	0.0.0.0	"ras002"
5: fffc05	10:00:00:05:1e:34:01:bd fec0:60:69bc:63:205:1eff:fe34:1bd	10.3.220.5	0.0.0.0	"ras005"
6: fffc06	10:00:00:05:1e:34:02:3e	10.3.220.6	0.0.0.0	>"ras006"
7: fffc07	10:00:00:05:1e:34:02:0c	10.3.220.7	0.0.0.0	"ras007"

(output truncated)
The Fabric has 26 switches

The following table displays the **fabricShow** fields.

TABLE 5 fabricShow Fields

Field	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.
World Wide Name	The switch WWN.
Enet IP Addr	The switch Ethernet IP address for IPv4- 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

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

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 6 Mapping Switch Type to Brocade Product Name

Switch Type	Brocade Product Name	Description
62	Brocade DCX	Chassis with a bandwidth of 4.6Tb/s. Configurable scaling up to 512 ports and supports 512Gb/s UltraScale ICL bandwidth.
64	Brocade 5300	Configurable scaling of 48,64 or 80 ports and supports 640Gb/s bandwidth. It supports 1-, 2-, 4-, and 8Gb/s media types.
66	Brocade 5100	Configurable scaling of 24,32 or 40 ports and supports 680Gb/s bandwidth. It supports 1-, 2-, 4-, and 8Gb/s media types.
67	BrocadeEncryption Switch	Configurable scaling up to 32 ports. It supports 8Gb/s media types.
71	Brocade 300	Configurable scaling up to 239 switches. It supports 4-, 8Gb/s media types.
77	Brocade DCX-4S	Chassis with a bandwidth of 2.3Tb/s. Configurable scaling up to 256 ports and supports 256Gb/s UltraScale ICL bandwidth.
83	Brocade 7800 Extension Switch	It supports 128Gb/s bandwidth and 1GbE media types.
92	Brocade VA-40FC switch	Configurable scaling up to 40 ports and supports 320Gb/s bandwidth. It supports 4-, and 8Gb/s media types.
109	Brocade 6510	Fixed-port switch with configurable scaling from 24 to 48 ports. It supports 16-, 10-, and 8Gb/s media types.

TABLE 6 Mapping Switch Type to Brocade Product Name (continued)

Switch Type	Brocade Product Name	Description
117	Brocade 6547 blade server SAN I/O module	Embedded switch with configurable scaling up to 48 ports and 16Gb/s media types.
118	Brocade 6505	Switch with configurable scaling from 12 to 24 ports. It supports 16 and 8Gb/s media types.
120	Brocade DCX 8510-8	Chassis with a bandwidth of 10.2Tb/s. It supports 512 ports × 16Gb/s + 2.048Tb/s UltraScale ICL bandwidth.
121	Brocade DCX 8510-4	Chassis with a bandwidth of 5.1Tb/s. It supports 256 ports × 16Gb/s data rate + 1.024Tb/s UltraScale ICL bandwidth.
129	Brocade 6548 blade server SAN I/O module	Embedded switch with 16 internal ports and 12 SFP ports. It supports 4-, 8- and 16Gb/s media types.
130	Brocade M6505 Blade Server SAN I/O Module	Embedded switch with configurable scaling up to 24-ports and 16Gb/s media types.
133	Brocade 6520	Fixed-port switch with configurable scaling from 48 to 96 ports. It supports 16-, 10-, and 8Gb/s media types.
148	Brocade 7840 extension switch	Extension switch with configurable scaling of 24 ports and supports 8-,16Gb/s media types. It supports 1 GbE, 8 GbE,10 GbE, and 40 GbE Ethernet interfaces.
149	Brocade 6546 blade server SAN I/O module	Configurable scaling 12 to 24-ports. It supports 4-, 8- and 16Gb/s media types.
150	Brocade 6545 blade server SAN I/O module	Configurable scaling up to 26-ports. It supports 4-, 8- and 16Gb/s media types.
156	Brocade 6543 blade server SAN I/O module	Embedded switch with 16 internal ports and 8 SFP ports.
157	Brocade 6558 blade server SAN I/O module	Embedded switch with configurable scaling from 24 to 48 ports. It supports 4-, 8-, and 16Gb/s media types.
165	Brocade X6-4	Chassis with a bandwidth of 8Tb/s . It supports up to 192 32Gb/s ports or a 256-port equivalent using 16 UltraScale ICL bandwidth ports.
166	Brocade X6-8	Chassis with a bandwidth of 16Tb/s . It supports up to 384 32Gb/s ports or a 512-port equivalent using 128Gb/s, UltraScale ICL ports (32Gb/s × 4 QSFP ports)
162	Brocade G620	Fixed-port switch with configurable scaling from 24 to 64 ports. It supports 128-, 32-, 16-, and 10Gb/s media types.
167	Brocade 6542 blade server SAN I/O module	Embedded switch with configurable scaling from up to 36-ports. It supports 4-, 8-, and 16Gb/s media types.
170	Brocade G610	Fixed-port switch with configurable scaling from 8 to 24 ports. It supports 32-, 16-, and 10Gb/s media types.
173	Brocade G630	Fixed-port switch with configurable scaling from 48 ports to 128 ports. It supports 128-, 32-, 16-, and 10Gb/s media types.
178	Brocade 7810 Extension Switch	Extension switch with configurable scaling of 12 ports and supports 32-, 16-, -8, and 4Gb/s media types. It supports 1 GbE and 10 GbE Ethernet interfaces.

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 supported by Fabric OS v7.0.0 and later allow 31 characters. Chassis names must begin with an alphabetic character and can include alphabetic and numeric characters, and the underscore (_).

Customizing 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 myfabric@1
```

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 quotes.

```
switch:admin> fabricname --set "my new fabric"
```

To set a fabric name that includes special meta-characters or spaces, use the command **fabricName**.

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

To clear the fabric name, use the **fabricName --clear** command.

High Availability Considerations for Fabric Names

Fabric names 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 7 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 logical switch	Enables all ports on physical chassis
Enable chassis	Enables all ports on physical chassis	Not allowed
Disable switch	Disables all ports on logical switch	Disables all ports on physical chassis
Disable chassis	Disables all ports on 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> switcheenable
```

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 have 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, Backbones, 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 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:

```
Broadcast message from root (ttyS0) Wed Jan 25 16:12:09 2006...
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 a Brocade DCX 8510 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:

```
DCX: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....
Broadcast message from root (pts/0) Fri Oct 10 08:36:48 2008...
The system is going down for system halt NOW !!
```

4. Power off the switch.

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, refer to [Performing Advanced Configuration Tasks](#) on page 71.
- For information on configuring the routing of connections, refer to [Routing Traffic](#) on page 515.
- For information on configuring extended inter-switch connections, refer to [Managing Long-Distance Fabrics](#) on page 373.

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

For more information on extended ISL modes, which enable long distance inter-switch links, refer to [Managing Long-Distance Fabrics](#) on page 373.

Performing Advanced Configuration Tasks

• Port identifiers and PID binding overview.....	71
• Ports.....	75
• Blade terminology and compatibility.....	92
• Enabling and disabling blades.....	94
• Blade swapping.....	95
• Disabling switches.....	100
• Power management.....	100
• Equipment status.....	102
• Audit log configuration.....	104
• Duplicate PWWN handling during device login.....	110
• Forward error correction.....	112

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 only select static PID binding if there is a compelling reason, and only after you have evaluated the effect of doing so.

Core PID addressing mode

Core PID is the default PID format for Brocade platforms. It uses the entire 24-bit address space of the domain, 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.

This address does not change unless you choose to swap the address using the **portSwap** command.

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 the Brocade Backbone, fixed addressing mode is used only on the default logical switch.

10-bit addressing (mode 0)

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

NOTE

The default switch in the Brocade Backbones use the fixed addressing mode.

The 10-bit addressing mode utilizes the 8-bit area ID and the borrowed upper two 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 256 NPIV devices. A logical switch can support up to 144 ports that can each support 256 devices. Areas 0x90 through 0xFF use an additional two bits from the AL_PA for the port address. Therefore, these ports support only 64 NPIV devices per port.

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 256 NPIV devices (in fixed addressing mode, only 128 NPIV devices are supported in non-VF mode and 64 NPIV devices in VF mode on a 48-port blade).
- Any port on a 48-port or 64-port blade can support hard port zoning.
- Port index is not guaranteed to be equal to the port area ID.

256-area addressing (mode 1 and mode 2)

The 256-area addressing mode is available only in a logical switch on the Brocade Backbone. 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.

Refer to the *Brocade FICON Administration Guide* 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. This 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.

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 is not supported on the default switch.

The 48-port and 64-port blades have the following restrictions:

- Port-based addressing is not supported on the upper 16 ports of a 64-port blade on the Brocade DCX 8510-4 and X6-4 Backbones.
- Port-based addressing is not supported on the upper 16 ports of a 48-port blade on the Brocade DCX 8510-8 and X6-8 Backbones.
- Port-based addressing is not supported on the upper 32 ports of a 64-port blade on the Brocade DCX 8510-8 and X6-8 Backbones.

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.

Once 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 NPIV devices. The number of point-to-point devices supported depends on the areas available. For example, 448 areas are available on Backbones 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 is supported in the default switch on the Brocade DCX 8510 Backbones and Brocade X6 Directors but 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 FX8-24 and 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 NPIV devices get the same PID across reboots and AL_PAs assigned for the device do not depend on the order in which the devices come up. For more information on NPIV, refer to [NPIV](#) on page 123.

Enabling automatic PID assignment

NOTE

To activate the 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 **wwnAddress -bind** command to assign a 16-bit PID to a given WWN.

Clearing 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 **wwnAddress -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:
wwnAddress --show
 - To display the PID assigned to the specific device WWN, enter the following:
wwnAddress -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. Refer to the "Diagnostic Port" chapter in the *Brocade Fabric OS Troubleshooting and Diagnostics Guide* for more information on this port type.

NOTE

With FOS 8.2.1, we introduce the Brocade 7810 Extension Switch, which supports D-Port capabilities. However, even though the Brocade 7810 switch supports both ETH and FC ports, D_Port is only supported on FC ports.

- **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.
- **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.

- VEX_Port: A virtual EX_Port that connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, a VEX_Port appears as a normal VE_Port. It follows the same Fibre Channel protocol as other VE_Ports. However, the router terminates VEX_Ports rather than allowing different fabrics to merge as would happen on a switch with regular VE_Ports.
- VF_Port: An FCoE enabled ETH port supported on the FC32-64 port blade only.

Director port blades

Because DCX 8510 and X6 Directors contain interchangeable port blades, their procedures differ from those for fixed-port switches. For example, fixed-port models identify ports only by the port number, while Backbones identify ports by *slot/port* notation

NOTE

For detailed information about the Brocade DCX 8510 and Brocade X6 Director families, 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 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 DCX 8510-8 has 12 slots that contain control processor, core, port, and AP blades:

- Slot numbers 6 and 7 contain CPs.
- Slot numbers 5 and 8 contain core blades.
- Slot numbers 1 through 4 and 9 through 12 contain port and AP blades.

The Brocade DCX 8510-4 has 8 slots that contain control processor, core, port, and AP blades:

- Slot numbers 4 and 5 contain CPs.
- Slot numbers 3 and 6 contain core blades.
- Slot numbers 1 and 2, 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 Backbones or 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. [Table 12](#) on page 92 lists the port numbering schemes for the blades.

Setting port names

Perform the following steps to specify a port name. For Backbones, 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 Backbone, 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 port area ID

The relationship between the port number and area ID depends upon the PID format used in the fabric. When Core PID format is in effect, the area ID for port 0 is 0, for port 1 is 1, and so forth.

For 32-port blades (FC16-32), the numbering is contiguous up to port 15; from port 16, the numbering is still contiguous, but starts with 128. For example, port 15 in slot 1 has a port number and area ID of 15; port 16 has a port number and area ID of 128; port 17 has a port number and area ID of 129.

For 48-port blades (FC16-48, FC32-48), the numbering is the same as for 32-port blades for the first 32 ports on the blade. For ports 32 through 47, area IDs are not unique and port index should be used instead of area ID.

For the 64-port blades (FC16-64), the numbering is the same as for 32-port blades for the first 32 ports on the blade. For ports 32 through 63, area IDs are not unique and port index should be used instead of area ID.

If you perform a port swap operation, the port number and area ID no longer match.

To determine the area ID of a particular port, enter the **switchShow** command. This command displays all ports on the current (logical) switch and their corresponding area IDs.

Port identification by index

With the introduction of 48-port and 64-port blades, indexing was introduced. Unique area IDs are possible for up to 255 areas, but beyond that there needed to be some way to ensure uniqueness.

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.

NOTE

The port area schema does not apply to the Brocade DCX 8510-4 and Brocade X6-4 Backbones.

Dynamic Portname

When the Dynamic Portname 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 Portname feature is disabled by default and can be enabled using the **configure** command.

The FDMI Host Name feature and the Dynamic Portname feature are mutually exclusive. If you try to enable the Dynamic Portname feature along with the FDMI Host Name feature, then the **configure** command fails and exits with an error message. You can configure the Dynamic Portname feature when the switch is online. The switch does not have to be offline to configure the Dynamic Portname feature.

The Dynamic Portname feature is not supported on a switch if AG mode is enabled. The Dynamic Portname feature is not supported in FMS mode.

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 is used to pick 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 is used to configure the Dynamic Portname. The option to configure the Dynamic Portname 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 Portname features are mutually exclusive. If you try to enable the Dynamic Portname 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]
  Disable FID Check (yes, y, no, n): [no]
  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:      66.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:05:1e:82:3c:2a
zoning:          OFF
switchBeacon:    OFF
FC Router:       OFF
Fabric Name:     switch_test
Allow XISL Use:  OFF
LS Attributes:   [FID: 128, Base Switch: No, Default Switch: Yes, Address Mode 0]
```

Index	Port	PortWWN	Name
0	0	20:00:00:05:1e:82:3c:2a	port0
1	1	20:01:00:05:1e:82:3c:2a	port1
2	2	20:02:00:05:1e:82:3c:2a	port2
3	3	20:03:00:05:1e:82:3c:2a	port3
4	4	20:04:00:05:1e:82:3c:2a	port4
5	5	20:05:00:05:1e:82:3c:2a	port5
6	6	20:06:00:05:1e:82:3c:2a	port6
7	7	20:07:00:05:1e:82:3c:2a	port7
8	8	20:08:00:05:1e:82:3c:2a	port8
9	9	20:09:00:05:1e:82:3c:2a	port9
10	10	20:0a:00:05:1e:82:3c:2a	port10
11	11	20:0b:00:05:1e:82:3c:2a	port11
12	12	20:0c:00:05:1e:82:3c:2a	port12
13	13	20:0d:00:05:1e:82:3c:2a	port13
14	14	20:0e:00:05:1e:82:3c:2a	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` command with the `-d` 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 is used to represent 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 8 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 AG and 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) //
```

Configuring a device-switch connection

For 8-Gbps platforms only: To configure an 8 Gbps (and 8 Gbps only) connection between a device and a switch, use the **portCfgFillWord** command.

The **portCfgFillWord** command provides the following configuration options:

- Mode Link Init/Fill Word
- Mode 0 IDLE/IDLE
- Mode 1 ARBF/ARBF
- Mode 2 IDLE/ARBF
- Mode 3 If ARBF/ARBF fails, use IDLE/ARBF

This command not applicable to Gen 5 (16-Gbps) and Gen 6 (32-Gbps) platforms.

ATTENTION

Although this setting only affects devices logged in at 8 Gbps, changing the mode is disruptive regardless of the speed at which the port is operating.

The setting is retained and applied any time an 8 Gbps device logs in. Upgrades from prior releases which supported only Modes 0 and 1 will not change the existing setting, but switches reset to factory defaults with Fabric OS v6.3.1 or later will be configured to Mode 0 by default. The default setting on new units may vary by vendor.

Modes 2 and 3 are compliant with FC-FS-3 specifications (standards specify the IDLE/ARBF behavior of Mode 2, which is used by Mode 3 if ARBF/ARBF fails after three attempts). For most environments, Brocade recommends using Mode 3, as it provides more flexibility and compatibility with a wide range of devices. In the event that the default setting or Mode 3 does not work with a particular device, contact your switch vendor for further assistance.

Swapping port area IDs

If a device that uses port binding is connected to a port that fails, you can use port swapping to make another physical port use the same PID as the failed port. The device can then be plugged into the new port without the need to reboot the device.

If two ports are changed using the **portSwap** command, their respective areas and “P” values are exchanged.

For ports that are numbered above 255, the “P” value is a logical index. The first 256 ports continue to have an index value equal to the area ID assigned to the port. If a switch is using Core PID format, and no port swapping has been done, the port index value for all ports is the same as the physical port numbers. Using **portSwap** on a pair of ports will exchange the area ID and index values of those ports.

Port swapping has the following restrictions:

- Shared area ports cannot be swapped.
- Ports that are part of a trunk group cannot be swapped.
- GbE ports cannot be swapped.
- Ports on a faulty blade cannot be swapped.
- Swapping ports between different logical switches is not supported. The ports on the source and destination blades must be in the same logical switch.
- The **portSwap** command is not supported for ports above 256.
- Port swapping is not supported when TI Zoning is in use.
- Port swapping is not supported on 48-port and 64-port blades.

Use the following procedure to swap the port area IDs of two physical switch ports. The swapped area IDs for the two ports remain persistent across reboots, power cycles, and failovers.

To swap port area IDs, the port swap feature must be enabled, and both switch ports must be disabled.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **portSwapEnable** *port_ID* command to enable the feature.
3. Enter the **portDisable** *port_ID* command on each of the source and destination ports to be swapped.

The following example disables port 1 and port 1/2.

```
switch:admin> portdisable 1
switch:admin> portdisable 1/2
```

4. Enter the **portSwap** *port_ID1 port_ID2* command for the ports you want swapped.

The following example swaps port 1 and port 2.

```
switch:admin> portswap 1 2
switch:admin> portswap 1/1 2/2
```

5. Enter **portSwapShow** to verify that the port area IDs have been swapped.
A table is displayed, showing the physical port numbers and the logical area IDs for any swapped ports.
6. Enter **portSwapDisable** to disable the port swap feature.

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 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 **portCfgPersistence** command while enabling or disabling a port. For more information, refer to the *Fabric OS Command Reference*.

NOTE

The use of the **portCfgDefault** command is not recommended on the Blade Server SAN I/O Modules. For more information refer to the *Fabric OS Command Reference*

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 **portCfgPersistence --set -persistentenable** command.

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

In FMS mode, you cannot use the **portCfgPersistence** 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.

Fabric OS 7.1.0 and later provides F_Port decommissioning and recommissioning using Brocade Network Advisor 12.1.0 and later. Refer to the *Brocade Network Advisor User Manual* for details.

NOTE

All members of a trunk group must have an equal link cost value in order 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:

- The local switch and the remote switch on the other end of the E_Port must both be running Fabric OS 7.0.0 or later.
- 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 there are no other shortest path routes available. An impaired port is automatically removed from trunks. Port impair 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.

NOTE

An impaired TI link is not used if there are any other links available to the same destination regardless of failover enabled or disabled state. To use the port impaired feature with TI zones, ensure that there is always a redundant path for each TI zone so that if one of the links is impaired, there is an alternative path to the same TI zone that can be used.

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.

- 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)
```

- The **portshow** and **isshow** 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.

- Connect to the switch and log in using an account with admin permissions.
- Enter the **portCfgSpeed** command.

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

```
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 --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 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.

- Connect to the switch and log in using an account with admin permissions.
- Enter the **switchCfgSpeed** command.

If the specified speed is not supported on a particular port or blade, then 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 8 Gbps:

```
switch:admin> switchcfgspeed 8
Committing configuration...done.
```

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

```
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.

Fabric OS supports the following per-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 9 Port octet combinations and modes

Port Octet in Combination	Gen 6 octet modes	Gen 5 octet modes
1 (default)	32G 16G 8G 4G	16G 8G 4G 2G
2	10G 8G 4G	8G 4G
3	Not supported	16G 10 G

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, the user needs to disable the port or change the speed to a supported fixed speed in order 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 can disable the online ports.

Example 1: If you have the octet of ports 0-7 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 8-15 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 have to 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 (32-Gbps) and Gen 5 (16-Gbps) switches and ports, except for the core blades, the FC16-64 port blade, and the Brocade 7840 Extension Switch.

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 16 Gbps and 10 Gbps):

```
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 then have to 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 then have to 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 10 Supported speed arguments and command strings for octet combinations

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
	10		10

The final speed achieved also depends on the SFP used.

Example 1: If the octet of ports 0–7 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 8–15 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 maximum auto-negotiated 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 auto-negotiated speed.

For example, if a port is configured with a maximum auto-negotiated speed of 16 Gbps, and the SFP is 32 Gbps, then the attempted speed negotiations are 16 and 8 Gbps.

NOTE

A 32-Gbps optical transceiver in a Gen 6 device can only auto-negotiate to 32 Gbps, 16 Gbps, or 8 Gbps, and a 16-Gbps optical transceiver in a Gen 6 device can only auto-negotiate to 16 Gbps, 8 Gbps, or 4 Gbps; 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` command with the `-m` option.

The following example sets the maximum auto-negotiated speed to 16 Gbps for port 3 on slot 2.

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

Decoding fdmiShow command output

Prior to Fabric OS 7.4.0a, the **fdmiShow** command displays the hexadecimal data for all non-ASCII and non-WWN port attributes. This required you to lookup the hexadecimal conversion data in the standardized specifications to decode the information. Starting with Fabric OS 7.4.0a, 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 can use this command to display Fabric-Device Management Interface (FDMI) information for all Host Bus Adapters (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.

- **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:

```
sw0:admin> fdmishow
Local HBA database contains:
 30:03:00:05:1e:0e:ee:b9
Ports: 1
 30:03:00:05:1e:0e:ee:b9
  Port attributes:
    FC4 Types: FCP
    Supported Speed: 1 2 4 8 Gb/s
    Port Speed: 8 Gb/s
    Max Frame Size: 2048 bytes
    Device Name: /proc/scsi/brcd/edsim
    Host Name: EDSIM-FDMI
    Node Name: 10:00:00:05:1e:0e:ee:b9
    Port Name: 30:03:00:05:1e:0e:ee:b9
    Port Type: N_PORT (0x1)
    Port Symb Name: dsim:fdmi_host
    Class of Service: F, 1
    Fabric Name: 00:00:00:00:00:00:00:00
    FC4 Active Type: FCP
    Port State: 0x0
    Discovered Ports: 0x1
    Port Identifier: 0x000000
HBA attributes:
Node Name: 10:00:00:05:1e:0e:ee:b9
Manufacturer: Brocade Communication Systems
Serial Number: EDSIM003000
Model: BRCD-EDSIM
Model Description: PCI-Express Dual Channel 8Gb Fibre Channel HBA
```

```

Hardware Version:
Driver Version: 8.02.21
Option ROM Version: 2.08
Firmware Version: 4.03.01 [Multi-ID]
OS Name and Version: Linux 2.6.14.2 #1 Wed Mar 4 14:56:56 PST 2015

```

```

Local Port database contains:
30:03:00:05:1e:0e:ee:b9

```

```

Remote HBA database contains:

```

```

Remote Port database contains:

```

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 11 Core and CP blade terminology and platform support

Blade	Blade ID (slotshow)	Supported on:		Definition
		DCX 8510 family	X6 family	
CP8	50	Yes	No	Brocade DCX 8510 Backbone family control processor blade. This CP supports all blades used in the DCX 8510 Backbone families.
CR16-8	98	Yes DCX 8510-8 only	No	A core blade that has 16x4 QSFPs per blade. It can be connected to another CR16-8 or a CR16-4 core blade.
CR16-4	99	Yes DCX 8510-4 only	No	A core blade that has 8x4 QSFPs per blade. It can be connected to another CR16-4 or a CR16-8 core blade.
CPX6	175	No	Yes	Brocade X6 Backbone family control processor blade. This CP supports all blades used in the X6 Backbone families.
CR32-8	177	No	Yes X6-8 only	A core blade that has 16x4 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	A core blade that has 8x4 QSFPs per blade. It can be connected to another CR32-4 or a CR32-8 core blade.

TABLE 12 Port blade terminology, numbering, and platform support

Blade	Blade ID (slotshow)	Supported on:			Definition
		DCX 8510 family	X6 family	Ports	
FC32-48	178	No	Yes	48	A 48-port, 32-Gbps port blade supporting 8, 10, 16, and 32 Gbps port speeds.
FC32-64	204	No	Yes	64	A 16 QSFP-port, 4x32-Gbps port blade supporting 4x8, 4x10, 4x16, and 4x32 Gbps port speeds. The ports can also work in ETH (FCoE) mode supporting 16x40G, 64x25G, and 65x10G port speeds.
SX6	186	No	Yes	24	Extension blade with 32-Gbps Fibre Channel, FCIP, and 10-GbE technology.
FC16-32	97	Yes	No	32	A 32-port, 16-Gbps port blade supporting 2, 4, 8, 10, and 16 Gbps port speeds.

TABLE 12 Port blade terminology, numbering, and platform support (continued)

Blade	Blade ID (slotshow)	Supported on:		Ports	Definition
		DCX 8510 family	X6 family		
					<p>NOTE 10 Gbps speed for FC16- xx blades requires the 10G license.</p> <p>Ports are numbered from 0 through 15 from bottom to top on the left set of ports and 16 through 31 from bottom to top on the right set of ports.</p>
FC16-48	96	Yes	No	48	<p>A 48-port, 16-Gbps port blade supporting 2, 4, 8, 10, and 16 Gbps port speeds.</p> <p>NOTE 10 Gbps speed for FC16- xx blades requires the 10G license.</p> <p>Ports are numbered from 0 through 23 from bottom to top on the left set of ports and 24 through 47 from bottom to top on the right set of ports.</p>
FC16-64	153	Yes	No	64	<p>A 64-port, 16-Gbps port blade supporting 4, 8, and 16 Gbps port speeds. These ports have 16 QSFPs per blade.</p> <p>FC Ports are numbered from 0 through 63 from bottom to top. QSFPs are numbered from 0 through 15 from bottom to top.</p>
FX8-24	75	Yes	No	12 FC 10 1-GbE 2 10-GbE	<p>Extension blade with 8-Gbps Fibre Channel, FCIP, and 10-GbE technology.</p> <p>Port numbering on this blade is as follows.</p> <p>On the left side of the blade going from bottom to top:</p> <ul style="list-style-type: none"> • Six FC ports numbered from 0 through 5 • Two 10-GbE ports numbered xge0 and xge1 • Four 1-GbE ports numbered from ge0 through ge3 <p>On the right side of the blade going from bottom to top:</p> <ul style="list-style-type: none"> • Six FC ports numbered from 6 through 11 • Six 1-GbE ports numbered from ge4 through ge9

CP blades

The control processor (CP) blade provides redundancy and acts as the main controller on the Brocade Backbone. The Brocade DCX 8510 devices support the CP8 blades, and the X6 Backbone devices support the CPX6 blades.

The CP blades in the **Brocade** DCX 8510 and X6 Backbone devices are hot-swappable. The CP8 blades are fully interchangeable among Brocade DCX 8510-4 and DCX 8510-8 Backbones. The CPX6 blades are fully interchangeable among Brocade X6-4 and X6-8 Backbones.

Brocade recommends that 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 DCX 8510 and X6 platforms.

- Brocade DCX 8510-8 supports two CR16-8 core blades.

- Brocade DCX 8510-4 supports two CR16-4 core blades.
- Brocade X6-8 supports two CR32-8 core blades.
- Brocade X6-4 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 intelligent 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 each blade for a specific platform. Software functions are not supported across application blades.

NOTE

The hardware limit is enforced by software.

TABLE 13 Intelligent blade support within Brocade Backbone families

Intelligent blade	DCX 8510-8	DCX 8510-4	X6-8	X6-4
FX8-24	4	4	Not supported	Not supported
SX6	Not supported	Not supported	4	4

Follow these guidelines when using an FX8-24, 7800, 7840 and SX6 platforms

- Brocade FX8-24 blade and 7800 cannot be connected to the 7840 and SX6 platforms. The ports may come online, but they will not communicate with each other.
- If an FX8-24 blade is replaced by another FX8-24 blade, the previous IP configuration data would be applied to the new FX8-24.

Enabling and disabling blades

Port blades are enabled by default. In some cases, you will need to disable a port blade to perform diagnostics. When diagnostics are executed manually (from the Fabric OS command line), many commands require the port blade to be disabled. This ensures that diagnostic activity does not interfere with normal fabric traffic.

If you need to replace an application blade with a different application blade, there may be extra steps you need to take to ensure that the previous configuration is not interfering with your new application blade.

Enabling blades

Use the following procedure to enable a blade.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **bladeEnable** command with the slot number of the port blade you want to enable.

```
ecp:admin> bladeenable 3
Slot 3 is being enabled
```

Exceptions in enabling 48-port and 64-port blades

Because the area IDs are shared with different port IDs, the 48-port and 64-port blades support only F_Ports and E_Ports.

Port swapping is not supported on 48-port and 64-port blades. For the 32-port blades, port swapping is supported on all 32 ports. This means that if you replace a 32-port blade where a port has been swapped with a 48-port blade, the 48-port blade faults. To correct this, reinsert the 32-port blade and issue **portSwap** to restore the original area IDs.

Disabling blades

Use the following procedure to disable a blade.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **bladeDisable** command with the slot number of the port blade you want to disable.

```
ecp:admin> bladedisable 3
Slot 3 is being disabled
```

Blade swapping

Blade swapping allows you to swap one blade with another of the same type; in this way, you can replace a FRU with minimal traffic disruption.

The entire operation is accomplished when the **bladeSwap** command runs on the Fabric OS. Fabric OS then validates each command before implementing the command on the Backbone. If an error is encountered, the blade swap quits without disrupting traffic flowing through the blades. If an unforeseen error does occur during the **bladeSwap** command, an entry will be made in the RASlog and all ports that have been swapped as part of the blade swap operation will be swapped back. On successful completion of the command, the source and destination blades are left in a disabled state, allowing you to complete the cable move.

Blade swapping is based on port swapping and has the same restrictions:

- Shared area ports cannot be swapped.
- Ports that are part of a trunk group cannot be swapped.
- GbE ports cannot be swapped.
- Faulty blades cannot be swapped.
- Swapping ports between different logical switches is not supported. The ports on the source and destination blades must be in the same logical switch.
- Undetermined board types cannot be swapped. For example, a blade swap will fail if the blade type cannot be identified.
- Blade swapping is not supported when swapping to a different model of blade or a different port count. For example, you cannot swap an FC16-32 blade with an FC16-48 port blade.


How blades are swapped

The `bladeSwap` command performs the following operations:

1. Blade selection

The selection process includes selecting the switch and the blades to be affected by the swap operation. The following figure shows the source and destination blades identified to begin the process.

FIGURE 2 Identifying the blades



2. Blade validation

The validation process includes determining the compatibility between the blades selected for the swap operation:

- Blade technology. Both blades must be of compatible technology types (for example, Fibre Channel to Fibre Channel, Ethernet to Ethernet, application to application, and extension to extension).
- Port count. Both blades must support the same number of front ports (for example, 32 ports to 32 ports, 48 ports to 48 ports, and 64 ports to 64 ports).
- Availability. The ports on the destination blade must be available for the swap operation and not attached to any other devices.

ATTENTION


If you have swapped different types of blades by error, or if there is a need to swap two cards of different types, you can do that only by destructively erasing the existing blade configuration and reconfiguring the new blades and ports.

3. Port preparation

The process of preparing ports for a swap operation includes basic operations such as ensuring the source and destination ports are offline, or verifying that none of the destination ports have failed.

The preparation process also includes any special handling of ports associated with logical switches. For example, the following figure shows the source blade has ports in a logical switch or logical fabric, and the corresponding destination ports must be included in the associated logical switch or logical fabric of the source ports.

FIGURE 3 Blade swap with Virtual Fabrics during the swap




4. Port swapping

The swap ports action is an iteration of the **portSwap** command for each port on the source blade to each corresponding port on the destination blade.

As shown in the following figure, the blades can be divided into different logical switches as long as they are divided the same way. If slot 1 and slot 2 ports 0 through 7 are all in the same logical switch, then blade swapping slot 1 to slot 2 will work. The entire blade does not need to be in the same partition.

FIGURE 4 Blade swap with Virtual Fabrics after the swap



Swapping blades

Use the following procedure to swap blades.

1. Connect to the Backbone and log in using an account with admin permissions.
2. Enter the **bladeSwap** command.

If no errors are encountered, the blade swap will complete successfully. If errors are encountered, the command is interrupted and the ports are set back to their original configurations.

3. Once the command completes successfully, move the cables from the source blade to the destination blade.
4. Enter the **bladeEnable** command on the destination blade to enable all user ports.

Disabling switches

Switches are enabled by default. In some cases, you may need to disable a switch to perform diagnostic testing. This 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 will be in a disabled state, and will need to be enabled.

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 required to determine if there will be enough power to operate. If it is predicted to be 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 1 and proceeds to the last slot in the chassis. As power becomes available, slots are powered up in the reverse order. During the initial power up of a chassis, or using the **slotPowerOn** command, or the insertion of a blade, the available power is compared to required power before power is applied to the blade.

NOTE

Some FRUs in the chassis may use significant power, yet cannot be powered off through software.

The **powerOffListShow** command displays the power-off order.

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
```

- Use the following command to display the "(Persistent)" flag for the slot.

```
switch:admin> slotshow
Slot  Blade Type      ID      Status
-----
  1      UNKNOWN              VACANT
  2      UNKNOWN              VACANT
  3      UNKNOWN              VACANT, (Persistent)
  4      SW BLADE      97      INSERTED, NOT POWERED ON (Persistent)
  5      UNKNOWN              VACANT
  6      CP BLADE      50      ENABLED
  7      CP BLADE      50      ENABLED
  8      CORE BLADE    98      ENABLED
  9      SW BLADE      96      ENABLED
 10      SW BLADE      96      ENABLED
 11      UNKNOWN              VACANT
 12      UNKNOWN              VACANT
```

- 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.

- Connect to the switch and log in using an account with admin permissions.
- Enter the **switchShow** command. This command displays a switch summary and a port summary.
- Check that the switch and ports are online.

Verifying High Availability features (Backbones and Directors only)

High Availability (HA) features provide maximum reliability and nondisruptive management of key hardware and software modules.

Use the following procedure to verify High Availability features for a Brocade DCX 8510 Backbones and Brocade X6 Directors.

- Connect to the switch and log in using an account with admin permissions.
- Enter the **chassisShow** command to verify the model of the DCX or X6 and obtain a listing of all field-replaceable units (FRUs).
- Enter the **chassisShow--sort** displays all the field replaceable units (FRUs) in a sorted order.
- Enter the **chassisShow--help** displays 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
```

7. Enter the **fanShow** command to display the current 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 current 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 current status of each slot in the system.

Example of the slot information displayed for a X6-8 Director chassis

```
X6-8:FID165: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	ENABLED
4	SW BLADE	178	FC32-48	ENABLED
5	SW BLADE	178	FC32-48	ENABLED
6	SW BLADE	178	FC32-48	ENABLED
7	CORE BLADE	177	CR32-8	ENABLED
8	CORE BLADE	177	CR32-8	ENABLED
9	SW BLADE	178	FC32-48	ENABLED
10	EXT BLADE	186	SX6	ENABLED
11	SW BLADE	178	FC32-48	ENABLED
12	EXT 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 discussed in [Domain IDs](#) on page 61.

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 devices, hosts, and storage are connected.
3. Optional: Enter the **nsShow** command to verify 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.
- From 8.2.0, and later the last 8192 messages are persistently saved in the audit log, but all audit events are sent to the system message log, which — assuming there are no bottlenecks — will be 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 Backbone.

Switch names are logged for switch components and Backbone names for Backbone components. For example, a Backbone 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* for details on event classes and message formats. For more information on setting up the system error log daemon, refer to the *Brocade Fabric OS Troubleshooting and Diagnostics Guide*.

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* for more information.

Verifying 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 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

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.

```
switch:admin> auditcfg --class 2,4
Audit filter is configured.
```

3. Enter the **auditCfg --enable** command, which enables audit event logging based on the classes configured in step 2.

```
switch:admin> auditcfg --enable
Audit filter is enabled.
```

To disable an audit event configuration, enter the **auditCfg --disable** command.

4. Enter the **auditCfg --show** command to view the filter configuration and confirm that the correct event classes are being audited, and the correct filter state appears (enabled or disabled).

```
switch:admin> auditcfg --show

Audit filter is enabled.
2-SECURITY
4-FIRMWARE
```

5. Enter the **auditDump -s** command to confirm that the audit messages are being generated.

Example of the syslog (system message log) output for audit logging

```
Oct 10 08:52:06 10.3.220.7 raslogd: AUDIT, 2008/10/10-08:20:19 (GMT), [SEC-3020], INFO, SECURITY, admin/
admin/10.3.220.13/telnet/CLI, ad 0/ras007/FID 128, , Event: login, Status: success, Info: Successful login
attempt via REMOTE, IP Addr: 10.3.220.13.
Oct 10 08:52:23 10.3.220.7 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 10.3.220.7 raslogd: AUDIT, 2008/10/10-08:28:16 (GMT), [SEC-3021], INFO, SECURITY, admin/
NONE/10.3.220.13/None/CLI, None/ras007/FID 128, , Event: login, Status: failed, Info: Failed login attempt
via REMOTE, IP Addr: 10.3.220.13.
```

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 UDP protocol to send the error log messages to the syslog server. The default UDP port is 514. 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 and it 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.

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 non-secure syslog server:

```
switch:admin> syslogadmin --set -ip fec0:60:69bc: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 fec0:60:69bc: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 fec0:60:69bc:92:218:8bff:fe40:15c4
```

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 csr and third-party signed certificates to be installed as client certificates.

Hostname 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 10.20.58.113
Syslog IP address 10.20.58.113 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          10.20.58.113
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 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). Once the quiet time is configured, a log message will be generated which 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 14 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 15 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 16 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. 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.

Since FEC is optional at 10-Gbps and 16-Gbps 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 32 Gbps speed based on the standard, regardless of the FEC settings on the switch.
- FEC is supported on E_Ports on 16 Gbps-capable switches at 10/16G 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 only on 16 and 32 Gbps-capable switches (Brocade 6505, 6510, 6520, M6505, 6547, 6548, 7840, G610, G620, G630 and the Brocade DCX 8510 and Brocade X6 Director family). 32 Gbps is supported only with Fabric OS 8.0.0 and later. FEC is mandatory for ports operating at 32 Gbps speed and hence enabled by default.
- For switch-to-adapter connections, FEC is supported only on QLogic BR-1860 and BR-1867 with driver v3.2.x on the Fabric Adapter ports operating in HBA mode connected to 16- and 32- Gbps Brocade switches running Fabric OS 7.1 and later.
- FEC is supported only on link speeds of 10-Gbps, 16-Gbps, and 32-Gbps regardless of whether the platform is FEC-capable.
- To connect between a switch and an HBA at 16 or 32 Gbps, 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 16 or 32 Gbps.
- FEC is not active unless enabled when the HBA port speed changes to less than 16 or 32 Gbps.
- Some DWDMs in transparent mode can support FEC. For details, check with the DWDM vendor.

Enabling forward error correction

Use the following procedure to enable FEC.

ATTENTION

Enabling FEC is disruptive to traffic. FEC can be enabled or disabled only at 16 Gbps, or at 10 Gbps on E_Ports with octet mode 2 or 3 on Gen 5 devices. The FEC is always enabled at 32 Gbps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **portCfgFec --enable** command, specifying the port or range of ports on which FEC is to be enabled.

```
portcfgfec --enable -FEC slot/port
```

3. Enter the **portCfgFec --show** command to display the current FEC configuration.

```
portcfgfec --show slot/port
```

Enabling FEC on a single port

```
switch:admin> portcfgfec --enable -FEC 1
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
```

```
switch:admin> portcfgfec --show 1
```

```
Port: 1
FEC Capable: YES
10G/16G FEC Configured: ON
16 FEC via TTS Configured: OFF
FEC State: active
```

Disabling forward error correction

Use the following procedure to disable FEC.

ATTENTION

Disabling FEC is disruptive to traffic and can be disabled at 10, 16, 32 Gbps. However, FEC is always active (even if disabled) at 32 Gbps.

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
```

Refer to [Managing Long-Distance Fabrics](#) on page 373 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

"fec_cor_detected" is displayed on Gen 5 platforms. Gen 6 platforms display "fec_uncor_detected" and "fec_corrected_rate".

2. Enter the **portstatsclear** command to clear the counters.

FEC-via-TTS

All devices that support FEC-via-TTS have it enabled by default; however, not all devices that support 16 Gbps support FEC-via-TTS. FEC-via-TTS is negotiated during speed negotiation when a link comes up. Any HBA or device connection that supports 16 Gbps but not FEC-via-TTS auto-negotiates to 8 Gbps. Therefore, FEC-via-TTS must only be enabled on switch ports intended for connections to HBAs and devices that support FEC-via-TTS. A Brocade to Brocade ISL connection does not utilize TTS to enable FEC; consequently, FEC-via-TTS must be disabled on all E_Ports. When FEC-via-TTS is enabled on a port, the port will automatically be disabled when connected to another Brocade 16 Gbps Gen 5 switch.

- FEC-via-TTS is disabled by default.
- When enabling FEC-via-TTS, the port is momentarily disabled and therefore is a disruptive action. FEC-via-TTS is active only when both ends of the link have FEC enabled.
- E_Port connectivity is not supported when FEC-via-TTS is enabled on the switch port.
- DWDM devices do not support FEC-via-TTS.
- FEC-via-TTS is supported only on QLE267x and QLE274x with driver v9.1.14.22 and later. Prior to enabling FEC Transmitter Training Signal (TTS) mode, refer to your HBA documentation to confirm FEC-via-TTS functionality and support.
- FEC-via-TTS displays **ON** when control of the FEC state is permitted via TTS by an externally attached host or device even if the peer does not support it. Displays **Active** when FEC has been negotiated between the switch and the attached HBA or device. Displays **(..)** or **OFF** when the external control of FEC is disabled.

The following table shows different speed negotiations starting from the highest capable speed in Gen 5 devices .

TABLE 17 Gen 5: Enabling and disabling FEC via TTS

FEC via TTS Enabled	FEC via TTS Disabled
16-Gbps FEC via TTS	16-Gbps
8-Gbps	8-Gbps
4-Gbps	4-Gbps

When connecting to 32-Gbps capable device, the negotiated speed will be 32-Gbps with FEC is enabled. FEC-via-TTS is only applicable to Gen 5 devices. When connecting a Gen 6 device to a Gen 5 device, FEC-via-TTS can be negotiated, if FEC via TTS is

supported by the attached device. With 16-Gbps in Gen 6 device, the speed negotiation behavior is identical to 16-Gbps in Gen 5 device. FEC via TTS is still disabled by default. With 32-Gbps SFPs in a Gen 6 switch, speed negotiation starting from the highest capable speed, is as follows:

TABLE 18 Gen 6: Enabling and disabling FEC via TTS

FEC via TTS Enabled	FEC via TTS Disabled
32-Gbps FEC	32-Gbps FEC
16-Gbps FEC via TTS	16-Gbps
8-Gbps	8-Gbps

Examples of enabling or disabling FEC-via-TTS

FEC-via-TTS can be enabled or disabled using the CLI on a port range on a fixed-port switch or on a port range within a slot for Directors. FEC-via-TTS is a method to negotiate FEC on Gen 5 and Gen 6 16-Gbps TTS-capable links when FEC is supported. Therefore, FEC must be enabled as well.

- Before executing any commands, remember to set the CLI context to the fabric ID (FID) for the logical switch. For example, to set the context to FID 1:

```
switch:admin> setcontext 1
```

- To enable FEC-via-TTS on all ports of a 32-port blade in slot 1 without confirmation:

```
switch:admin> portcfgfec --enable -tts -f 1/0-31
```

- To enable FEC-via-TTS on port 0 of slot 1 with confirmation:

```
switch:admin> portcfgfec --enable -tts 1/0
```

- To disable FEC-via-TTS on port 0 of slot 1 without confirmation:

```
switch:admin> portcfgfec --disable -tts -f 1/0
```

- To disable FEC on port 0 of slot 1 without confirmation:

```
switch:admin> portcfgfec --disable -fec -f 1/0
```

Fabric-Assigned PWWN

• Fabric-Assigned PWWN overview.....	117
• User- and auto-assigned FA-PWWN behavior	118
• Configuring an FA-PWWN for an HBA connected to an Access Gateway.....	119
• Configuring an FA-PWWN for an HBA connected to an edge switch.....	120
• Supported switches and configurations for FA-PWWN.....	120
• Configuration upload and download considerations for FA-PWWN.....	121
• Security considerations for FA-PWWN.....	121
• Restrictions of FA-PWWN.....	122
• Access Gateway N_Port failover with FA-PWWN.....	122

Fabric-Assigned PWWN overview

Fabric-Assigned PWWN simplifies 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 that multiple administrative teams (for example, server and storage teams) coordinate with each other to perform configuration tasks such as zone creation in the fabric and LUN mapping and masking on the storage device. These tasks must be completed before the server is deployed. Before you can configure WWN zones and LUN masks, you must find out the physical PWWN of the server. This means that 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 gets deployed in an FC SAN.

You can use a fabric-assigned PWWN (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 this 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 5 Fabric-assigned port World Wide Name provisioning scenarios



User- 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. Refer to the *Brocade Adapters Administrator's Guide* for a list of supported HBAs.

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 root.
- 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 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.

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. Refer to the *Brocade Adapters Administrator's Guide* for a list of supported HBAs.

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 root.
- 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 7.0.0 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 6 (32-Gbps) platforms running Fabric OS 8.0.x
 - Brocade G620
 - Brocade X6 Directors
- Switch platforms running Fabric OS v7.0.0 or later:
 - Brocade DCX 8510 Backbones
 - Brocade 6505
 - Brocade 6510
 - Brocade 6520
- Access Gateway platforms running Fabric OS v7.0.0 or later:
 - Brocade 6505
 - Brocade 6510

Refer to the release notes for the supported Brocade HBA or adapter versions.

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

Brocade recommends you 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 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. Refer to [Device authentication policy](#) on page 248 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. Refer to [DCC policy behavior with Fabric-Assigned PWWNs](#) on page 242 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
- Swapped ports (using the *portswap* command)
- 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.

NPIV

- NPIV overview..... 123
- Configuring NPIV..... 124
- Enabling and disabling NPIV..... 125
- Base device logout.....126
- Viewing NPIV port configuration information.....129

NPIV overview

N_Port ID Virtualization (NPIV) 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 is therefore capable of registering with all services of the fabric. This chapter does not discuss the Access Gateway feature. For more information on the Access Gateway feature, refer to the *Brocade Access Gateway Administration 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 emulated by NPIV appear no different than regular devices 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, then a port login (PLOGI) to a non-existent virtual PID is not blocked by the switch; rather, it is delivered to the device attached to the NPIV port. In cases where the device is not capable of handling such unexpected PLOGIs, 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:      switch51
switchType:      71.2
switchState:     Online
switchMode:      Access Gateway Mode
switchWwn:       10:00:00:05:1e:41:49:3d
switchBeacon:    OFF
Index Port Address Media Speed State Proto
=====
0 0 010000 id N4 Online FC F-Port 20:0c:00:05:1e:05:de:e4 0xa06601
1 1 010100 id N4 Online FC F-Port 1 N Port + 4 NPIV public
2 2 010200 id N4 Online FC F-Port 1 N Port + 119 NPIV public
3 3 010300 id N4 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.

Upgrade considerations

The maximum logins per switch decreased with Fabric OS v6.4.0. When upgrading from a release previous to Fabric OS v6.4.0, the configured maximum is carried forward and may exceed the Fabric OS v6.4.0 limit. It is recommended to reconfigure this parameter to be within the range permitted in Fabric OS v6.4.0 and later.

Fixed addressing mode

Fixed addressing mode is the default addressing mode used in all 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

The 10-bit addressing mode is the default mode for all the logical switches created in the Brocade DCX 8510 and Brocade X6. The number of NPIV supported on a port is 64.

In 10-bit addressing mode, the switch can dynamically assign up to 112, 8-bit areas, after that 10-bit areas are assigned. So the 8-bit areas can have up to 255 devices, and 10-bit areas can have up to 63 devices.

The following table shows the number of NPIV devices supported on the Brocade DCX 8510 and Brocade X6.

TABLE 19 Number of supported NPIV devices per port

Platform	VF Disabled	VF-Enabled Default Switch	VF-Enabled Logical Switch	VF-Enabled Base Switch
DCX 8510-4 (or) X6-4	Yes, 255	Yes, 255	Yes, 255*	No NPIV support
DCX 8510-8 (or) X6-8	Yes, 127	Yes, 63	Yes, 255**	No NPIV support

NOTE

- All numbers indicate virtual device limit
- *The first 112 physical NPIV-capable devices connected to a logical switch using 10-bit addressing can log in 255 logical devices. The physical NPIV-capable devices after 112 and higher are limited to 63 logical devices. If port is assigned 10-bit area, the number of NPIV devices supported is 63.
- **If port is assigned 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 is used to specify the maximum number of virtual N_port IDs per port on a switch. It can also be used to enable or disable NPIV. Once NPIV is enabled on the port, you can specify the number of logins per port. If the NPIV feature has been disabled, then 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> portcfgnpiport --setloginlimit 7/0 128
NPIV Limit Set to 128 for Port 128
switch:admin> portcfgshow 7/0
Area Number:                128
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            OFF
LOS TOV enable               OFF
NPIV capability              ON
QOS E_Port                   AE
Port Auto Disable:           OFF
Rate Limit                   OFF
EX Port                       OFF
Mirror Port                   OFF
Credit Recovery              ON
F_Port Buffers               OFF
Fault Delay:                  0(R_A_TOV)
NPIV PP Limit:                128
CSCTL mode:                   OFF
Frame Shooter Port           OFF
D-Port mode:                  OFF
D-Port over DWDM              ..
Compression:                  OFF
Encryption:                   OFF
FEC:                           ON
```

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.

- 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 5100:

```
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. This 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 the 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 which 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 20 Base and NPIV logout behaviors

Scenario	Without base device logout enabled	With base device logout enabled
<ul style="list-style-type: none"> No device is logged in. 	<ul style="list-style-type: none"> Requires a device with a base PID to log in to the F_Port before any NPIV device can log in. 	<ul style="list-style-type: none"> Requires a device with a base PID to log in to the F_Port before any NPIV device can log in.
<ul style="list-style-type: none"> Base device is logged in. 	<ul style="list-style-type: none"> 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). 	<ul style="list-style-type: none"> 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.
<ul style="list-style-type: none"> Base device is logged in. At least one NPIV device is logged in. 	<ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> Allows base device to log out while NPIV devices are still logged in. Indicates that the base device is logged out in the switchShow command output. Does not remove the routing information.
<ul style="list-style-type: none"> Base device is logged out. 	N/A	<ul style="list-style-type: none"> Allows base device to log back in as a base or NPIV device.

TABLE 20 Base and NPIV logout behaviors (continued)

Scenario	Without base device logout enabled	With base device logout enabled
<ul style="list-style-type: none"> At least one NPIV device is logged in. 		<ul style="list-style-type: none"> Allows NPIV devices to log in and log out. Does not remove the routing information.
<ul style="list-style-type: none"> Base device is logged out. All NPIV devices are also logged out. 	<ul style="list-style-type: none"> Requires a device with a base PID to log in to the F_Port before any NPIV device can log in. 	<ul style="list-style-type: none"> 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.

```
portcfgflogilogout --enable -all
```

The following example enables base device logout on only specific ports in the logical switch.

```
portcfgflogilogout --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.

- In Fabric OS versions prior 7.3.0, the details of devices logged in with FLOGI (base devices) with AL_PA 0x40 and 0xC0 are not propagated to NS or FCP unless the devices are registered with NS. Because they cannot be probed, devices with these AL_PAs must initiate NS registration. If they do register with NS and then later log out, there can be multiple LUNs losing access during login time because the logged out base device entry is not removed from NS. Furthermore, the logged out base device cannot login to the fabric again. Fabric OS 7.3.0 and later 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 **fcpProbeShow** 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, VE/GE (E/Ex/LG) ports.
- Base device logout must be disabled on all the ports before downgrading from Fabric OS 7.3.0 to an earlier version.

Viewing 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
Fill Word(On Active)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Fill Word(Current)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
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
NPIV PP Limit	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
QOS E_Port	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
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

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	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**, then 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
switchType:    66.1
switchState:   Online
switchMode:    Native
switchRole:    Principal
switchDomain:   1
switchId:      fffc01
switchWwn:     10:00:00:05:1e:82:3c:2a
zoning:        OFF
switchBeacon:  OFF
FC Router:     OFF
FC Router BB Fabric ID: 128
```

Area	Port	Media	Speed	State	Proto
0	0	id	N1	Online	F-Port 1 Nport + 1 NPIV devices.
1	1	id	N4	No_Light	
2	2	id	N4	Online	F-Port 20:0e:00:05:1e:0a:16:59
3	3	id	N4	No_Light	
4	4	id	N4	No_Light	

(output truncated)

- 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

- Configuration settings..... 133
- Configuration file backup..... 135
- Configuration file restoration..... 137
- Configurations across a fabric..... 140
- Configuration management for Virtual Fabrics..... 140
- Brocade configuration form..... 142

Configuration settings

It is important to maintain consistent configuration settings on all switches in the same fabric, because inconsistent parameters, such as inconsistent PID formats, can cause fabric segmentation. As part of standard configuration maintenance procedures, it is recommended that you back up all important configuration data for every switch on a host computer server as a safety measure.

NOTE

For more information about troubleshooting configuration file uploads and downloads, refer to the *Brocade Fabric OS Troubleshooting and Diagnostics Guide*.

There are two ways 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 switch model 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

Editing of the uploaded file is unsupported and can result in system errors if an edited file is subsequently downloaded.

If your user account has chassis account permissions, you can use any of the following options when uploading or downloading a configuration file:

TABLE 21 Configuration management options

-fid	To upload the specified FID configuration.
-all	To upload all of the system configuration, including the chassis section and all switch sections for all logical switches. NOTE Use this parameter when obtaining a complete capture of the switch configuration in a switch that has Virtual Fabrics mode disabled.
-chassis	To upload only the chassis section of the system configuration file.
-switch	To upload the switch configuration only, if Virtual Fabrics mode is disabled.
-cra	Challenge Response Authentication (CRA) is to be used with SCP. You can run this command without arguments for interactive mode. Only a password is supported as a challenge.

TABLE 21 Configuration management options (continued)

-map	Uploads the port-to-area addressing mode configuration files. This command should be used in 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

There is always at least one switch section for the default switch or a switch that has Virtual Fabrics mode disabled, and there are additional sections corresponding 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

It is recommended that you 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 **configUpload -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
- 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, then 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. Therefore, 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 FTP server credentials

FTP 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. These characters have an alternate meaning or are designated to perform a special instruction. Here is a list of some of the more commonly used special characters and their alternate meanings.

NOTE

This list is not exhaustive, and alternate meanings for some characters are contextual. For detailed information on using special characters in FTP credentials, refer to Linux scripting information available on the Internet.

- **&** is used to put a command in background/batch mode.
- **!** is used to recall the last invocation of the command matching the pattern that follows the character.
- **|** is used to direct (pipe) output to the command that follows the character.
- **;** is used to concatenate multiple commands.
- ***** is used to represent a wildcard character.
- **?** is used as a match for any single character in the specified position.
- **()** are used for integer expansion.
- **<** and **>** are used for redirection; **<** represents input, and **>** represents output.
- **\$** is used to represent a shell variable.
- **`** is used for command substitution or for assigning the output of a command to a variable.
- **"** is used for partial quoting.
- **'** is used for full quoting.
- A space is used as a separation character.
- **#**, when preceded by a space, treats all characters through the end of the line as a comment rather than commands.

These special characters may be used to enhance user credential security. However, in order for these characters to be properly interpreted, you must use one of the following methods:

- "Escape" each instance of the special character by preceding it with the escape character (****).
- Enclose the credentials containing special characters within single quotes (**'**).

If single quotes are themselves part of the credential, precede each instance of the single quote with the escape character (****). Alternately, the string may be enclosed in double quotes (**"**) if a more intricate bash substitution is desired to further strengthen the security measure of the credentials.

You can test the representation of the credentials by logging in with root-level permissions and using the **echo** command.

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 model. Downloading a configuration file from a different switch model 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](#) on page 140 to restore the logical switches.

If a **configDownload** command is issued on a non-FC router, any FC router parameters may be viewed in the downloaded data. This is harmless to the switch and can be ignored.

MAPS configuration is downloaded onto the switch only if MAPS is enabled on the local switch.

You can download the configuration file from a device with Fabric OS 7.2.0 or 7.3.0 to 7.4.0.

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

NOTE

If you modify the configuration attributes of a user in the key file, then the configuration download will be blocked.

NOTE

From Fabric OS 8.2.0 or later, all the admin domain commands are blocked, so you cannot download any admin domain configuration from configuration file.

The following table lists restrictions for some of the options of the **configDownload** command.

TABLE 22 Restrictions for using the options of the configDownload command

Option	Restrictions
-chassis	The number of switches defined in the downloaded configuration file must match the number of switches currently defined on the switch.
-fid FID	The FID must be defined in both the downloaded configuration file and the current system. NOTE Brocade recommends you disable a switch before downloading a configuration file. If you plan to download a configuration file while the switch is enabled, refer to Configuration download without disabling a switch on page 138.
-fid FID -sfid FID	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 the number of switches or FIDs currently defined on the switch. The switches must be disabled first. If they are not, the configDownload 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 configDownload command, verify if any switches must be disabled.

TABLE 22 Restrictions for using the options of the configDownload command (continued)

Option	Restrictions
	<p>If you are performing a configuration download due to a configuration error, it is highly recommended to run the configDefault command before running the configDownload command. Refer to Configuration download without disabling a switch on page 138 for more information on non-disruptive configuration downloads.</p> <p>If Virtual Fabrics mode is enabled, the chassisDisable and chassisEnable commands are used to disable all logical switches on the affected switch. This process bypasses the need to disable and enable each switch individually once the configuration download has completed.</p> <p>Non-Virtual Fabric configuration files downloaded to a Virtual Fabric system have a configuration applied only to the default switch. If there are multiple logical switches 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.</p>

If you need to set up your switch again, run the commands listed in [Table 23](#) and save the output in a file format. Store the files in a safe place for emergency reference.

TABLE 23 CLI commands to display or modify switch configuration information

Command	Displays
configShow	System configuration parameters, settings, and license information.
fcLunQuery	LUN IDs and LUNs for all accessible targets.
fcrRouterPortCost	FC Router route information.
fcrXlateConfig	Translate (xlate) domain's domain ID for both EX_Port-attached fabric and backbone fabric.
fosConfig	Fabric OS features.
ipAddrShow	IP address.
isnscCfg	Configuration state of the iSNS client operation.
licenseShow	License keys installed with more detail than the license information from the configShow command.
portCfgEXPort	EX_Port configuration parameters.
portCfgVEXPort	VEX_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 SNMP, MAPS, Fabric Watch, or ACL parameters. However, if there is any changed parameter that does not belong to SNMP, MAPS, or ACL, then 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 be used to 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 you want to save.

1. Verify that the FTP 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 there are any changed parameters 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, local) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: UserFoo
Path/Filename [<home dir>/config.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 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. 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 model 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 *Fabric OS Command Reference*.

Downloading a configuration file from one switch to another switch of the same model

1. Configure one switch.
2. Use the **configUpload** command to save the configuration information. Refer to [Configuration file backup](#) on page 135 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. Refer to [Configuration file restoration](#) on page 137 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, refer to [Configuring Protocols](#) on page 189.

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. For example, a Virtual Fabrics configuration file for Fabric OS 7.4.x cannot be used on a Fabric OS 7.2.x switch and vice versa.

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. It is recommended not to use config.txt for a file name as this name can be confused with a normal uploaded configuration file.

Example of configUpload on a switch with Virtual Fabrics

```
DCX_80: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]: 5100_vf.txt
configUpload complete: VF config parameters are uploaded
2014/07/20-09:13:40, [LOG-1000], 225, SLOT 7 | CHASSIS, INFO, BrocadeDCX, Previous message repeated 7
time(s)
2014/07/20-10:27:14, [CONF-1001], 226, SLOT 7 | FID 128, INFO, DCX_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.

Wait for the configuration file to download to the switch.

6. Verify the L1SL ports are set up correctly.

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

```
5100: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]: 5100_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 in order 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 7.2.x cannot be used on a Fabric OS 7.1.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 24 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	

Configuration information	Values that you chose
Switch name	
Ethernet IP address	
Ethernet subnet mask	
Total number of local devices (nsShow)	
Total number of devices in fabric (nsAllShow)	
Total number of switches in the fabric (fabricShow)	

Managing User Accounts

• User accounts overview	144
• Local database user accounts.....	148
• Local user account database distribution.....	151
• Password policies.....	152
• The boot PROM password.....	158
• Remote authentication.....	162

User accounts overview

In addition to the default root, admin, 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

For new installations of Fabric OS, the root account is disabled by default. Not all systems ship with root accounts. If your system does not have a root account, you will not be able to enable it.

Each user account is associated with the following:

- Permissions — Associate roles with each user account to determine the functional access levels.
- Virtual Fabric list — Specifies the Virtual Fabric a user account is allowed to log in to.
- 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, then 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, then 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, refer to [Managing Virtual Fabrics](#) on page 382.

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 means of 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 means of the local user database. The local user database is manually synchronized by means of the **distribute** command to push a copy of the switch's local user database to all other switches in the fabric running Fabric OS 5.3.0 and later, but the **distribute** command is blocked if users with user-defined roles exist on the sending switch or on any remote, receiving switch.

NOTE

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.

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.

The following table outlines the Fabric OS predefined (default) roles.

TABLE 25 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

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 26 Permission types

Abbreviation	Definition	Description
O	Observe	The user can run commands by using options that display information only, such as running userConfig --show -a 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

TABLE 26 Permission types (continued)

Abbreviation	Definition	Description
		system, such as running the userConfig --change command with the -r 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.

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
```

You can also use the **classConfig --showcli** command to show the permissions that apply to a specific command.

Management channel

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 27 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

NOTE

The total number of sessions on a switch cannot exceed 32. The root user also can have only a maximum of four sessions.

Managing user-defined roles

Fabric OS provides an extensive toolset for managing user-defined roles:

- The **roleConfig** command is available for defining new roles, deleting created roles, or viewing information about user-defined roles.
- The **classConfig** command is available for displaying 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 can be used to assign a user-defined role to a user account.



CAUTION

You can add only a maximum of 64 user defined roles for Fabric OS and other management applications such as Brocade Network Advisor and Web Tools.

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 4 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 can be used to define 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
  mysecurityrole            O
```

To delete a user-defined role, use the **roleConfig --delete** command.

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*.

Configuring time-based access

You can also restrict the access for a user to a specified time of the day. The restriction applies for Telnet, SSH, console, and Web access for a user on all days.

1. Run the **userConfig --add todUser** command to add time-based users.

```
switch:admin> userconfig --add toduser -r user -l 1-128 -h 128 -c user -d "Time of day User" -at
hh:mm-hh:mm
```

The first instance of hh:mm after the “-at” option indicates the start time and the second instance indicates the end time. To downgrade to a previous Fabric OS version, you need to remove the time-based user configuration.



CAUTION

The time of access must not be configured for switch default accounts such as root or admin user, because it can lock you out of the switch access completely. To provide accessibility to the switch, you must have at least one account without the time of access configured.

2. To remove the time based access configuration, run the following command:

```
switch:admin> userconfig --change toduser -at 00:00-00:00
```

Local database user accounts

User **add**, **change**, and **delete** operations are subject to the *subset* rule: an admin with LFlist 1-10 cannot perform operations on an *admin*, *user*, or any role with LFlist 11-128. The user account being changed must have an LFlist 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 28 Default local user accounts

Account name	Role	Logical Fabric	Description
admin	Admin	LF1-128 home: 128	Most commands have <i>observe/modify</i> permission
root	Root	LF1-128 home: 128	Reserved NOTE Not all systems ship with root accounts. If your system does not have a root account, you will not be able to enable it.
user	User	LF-128 home: 128	Most commands have <i>observe-only</i> permission

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.

For all users, an asterisk (*) is displayed next to the username if the password is still a default password when you run **userconfig -show** command.

If there is a root account available on your device, the following items should be kept in mind:

- The default root password must be changed at the first login.
- Root access is disabled by default and restricted to the console only.
 - To enable the root account, enter **userconfig --change root -e** yes.
 - To change the root access settings, use the **rootaccess --set none/consoleonly/all** command to set the desired access permission.

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, displays a list of users that include that LF in their LF permissions.

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 metoo -l 1-128 -h 128 -r admin -c admin
```

This example creates a user account for the user metoo with the following properties:

- Access to Virtual Fabrics 1 through 128
 - Default home logical switch 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 on the command line.

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

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.

Local account passwords

The following rules apply to changing passwords:

- Users can change their own passwords.
- To change the password for another account requires admin permissions or an account associated with a user-defined role with Modify permissions for the LocalUserEnvironment RBAC class of commands. When changing an admin account password, you must provide the current password.
- An admin with LFlist 1-10 cannot change the password on an account with admin, user, or any permission with an LFlist 11-128. The user account being changed must have an VFlist 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.

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 option.

3. Enter the requested information at the prompts.

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 there are logical switches defined other than the default logical switch, then distributing the password database to switches is not supported. Distributing the password database to switches is not allowed if there are users associated with user-defined roles in either the sending switch or the remote switch

NOTE

The passwords are distributed only to switches running Fabric OS 8.0.1 or later if the receiving switch password hash is MD5. Otherwise, the password hash enforcement is applied and the passwords have to be changed before login.

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.

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 (refer to [Local user account database distribution](#) on page 151).

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, then 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 username 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 username 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 password "passAAAword" 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.

- 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 users from recycling recently used passwords, and is enforced across all user accounts when users are setting their own passwords. The password history policy is enforced only when a new password is defined.

The password history policy is enhanced further to prevent the user from configuring the new password similar to current password. A password similarity check is applied using **-minDiff** option to verify the difference of characters between the new and current password. This ensures that the new password is not a simple replacement of some characters of current password.

Use the following attributes of **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 can be set to the most recent password. The default value is 1, which means the current and one previous password cannot be reused. The value 2 indicates that the current 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 or not a newly specified password is identical to one of the specified number (1-24) of previously used passwords.

The password history is not enforced when an administrator sets a password for another user; instead, the user's password history is preserved and the password set by the administrator is recorded in the user's password history.

- Oldpasswd

This option to enable or disable old password check while changing the root password.

- MinDiff

This defines the number of character difference expected between the current and the new password.

The configuration range is between 0 to 40 and must be always less than the configured **-minlength** value. The default value is 0 and if it is set, then the password similarity check is off.

The password is considered as similar if the new password uses the same set of characters as the current password irrespective of the position. The duplicate occurrence of the same character in the new password is not considered as distinct. A character can occur once in the current password and twice in the new password. This second occurrence of the character is not considered as distinct.

NOTE

1. The MinDiff policy is only enforced when the user is logged in and changing the own password. Administrator, changing the user's password is exempted from this policy.

2. The MinDiff parameter is not distributed as part of PWD database.

The following table outlines the different combinations between the current and the new password and their results.

TABLE 29 Different MinDiff Combinations between Current and New Password

Current Password	New Password	Configured MinDiff Value	Similar and Distinct Characters between Passwords	Result
password	password123	5	8 and 3	Fail
password	password123	3	8 and 3	Pass
password	different	5	2 and 7	Pass
clinteastwood	oldwester	5	8 and 1	Fail
oldwester	clinteastwood	5	9 and 4	Fail
manifesto	maniFESTo	5	5 and 4	Fail

Example of a password history policy

The following example shows a password 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 difference between the current and the new password must be five distinct characters.

```
switch:admin> passwdcfg --set -uppercase 3 -lowercase 4 -digits 2 -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. A warning that password expiration is approaching is displayed when the user logs in. 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 in order 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 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 metoo 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 metoo -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 except the root account 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 except the root account, 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 is enforced across all 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 as the other permissions. 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 is used to validate 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 may be used to create 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. Audit logs should be examined to monitor if such attacks are attempted.

Changing the root password without the old password

In Fabric OS 7.4.0 and later, for environments where switch account passwords are controlled by password management appliances, you can configure an option to not verify the old password when changing the root account password. By default, this option is disabled, which means that you must provide the old password to change the root account password.

NOTE

Not all systems ship with root accounts. If your system does not have a root account, you will not be able to enable it. If you have lost or forgotten the root password, you should contact your service provider for the correct password reset or recovery procedures.

To disable verification of the old password, complete the following steps.

1. Connect to the device and log in using an account with root permissions.
2. Enter `passwdcfg --set -oldpasswd 1 / 0`.
 - A value of '1' enables the option; this means that the old password is not required to change the root password.
 - A value of '0' disables the option; this means that the old password is required to change the root password.

```
switch:admin> passwdcfg --set -oldpasswd 1
```

3. Enter the `passwd` command to change the root password.

NOTE

The `-oldpasswd` configuration is discarded when you downgrade to any version prior to Fabric OS 7.4.0. The `-oldpasswd` configuration is retained when you upload the switch configuration or upgrade to a later version of Fabric OS.

```
switch:admin> passwd
Warning: Access to the Root account may be required for proper support of the switch.
Please ensure the Root Passwords are documented in a secure location.
Recovery of a lost Root password will result in fabric downtime.
Changing password for root
Enter new password:
```

Configuring the password hash type

Versions of Fabric OS prior to 8.0.1 use the MD5 hash type for storing the local user passwords. 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 to manually trigger a password change immediately after the hash type configuration.

After a new password hash configuration has started, 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 the logged-in user's password.
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 for the current user, any user, or 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 has for all the users on the switch.

Sample output:

```
switch(mode)# passwdcfg --hash sha256 -manual
switch(mode)# passwdcfg -showhash -all
<Command output>
```

- 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.
- When you upgrade the switch to Fabric OS 8.0.1 or later, the MD5 hash type configuration is retained.
- When you downgrade the switch from Fabric OS 8.0.1 or later, the downgrade is blocked unless the following two conditions are met:
 - The password hash type for all the users is changed to MD5.
 - The switch hash type is changed to MD5.

If the two conditions are met but the password history consists of non-MD5 hash for any user, then the following message is displayed during the firmware downgrade.

```
" WARNING !!!!! Password hash in the history contains non-MD5 hash for user(s) which must be cleared
to proceed with downgrade. Please confirm with [Y/N] to proceed further, when prompted for.
System settings check passed.
```

```
You can run firmwaredownloadstatus to get the status of this command.
```

```
This command will cause a warm/non-disruptive boot but will require that existing telnet, secure
telnet or SSH sessions
be restarted.
```

```
Do you want to continue (Y/N) [Y]:"
```

- When you do a net install, all the password hash types are set to SHA512.

NOTE

You can use the **passwddefault** command to reset all the password hash types to SHA512 and reset all the passwords to default.

The boot PROM password

The boot PROM password provides an additional layer of security by protecting the boot PROM from unauthorized use. Setting a recovery string for the boot PROM password enables you to recover a lost boot PROM password by contacting your switch service provider. Without the recovery string, a lost boot PROM password cannot be recovered.

Although you can set the boot PROM password without also setting the recovery string, it is strongly recommended that you set both the password and the recovery string. If your site procedures dictate that you set the boot PROM password without the recovery string, refer to [Setting the boot PROM password for a switch without a recovery string](#) on page 160.

To set the boot PROM password with or without a recovery string, refer to the section that applies to your switch or Backbone model.

**CAUTION**

Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. Perform this procedure during a planned downtime.

Setting the boot PROM password for a switch with a recovery string

This procedure applies to the fixed-port switch models. The password recovery instructions provided within this section are only for the switches listed in supported hardware section. If your switch is not listed, contact your switch support provider for instructions.

1. Connect to the serial port interface as described in [Connecting to Fabric OS through the Serial Port](#) on page 33.
2. Reboot the switch.
3. Press **Esc** within four seconds after the message "Press escape within 4 seconds..." is displayed.
4. When prompted, enter **2** to select the recovery password option.

- If no password was previously set, the following message is displayed:

```
Recovery password is NOT set. Please set it now.
```

- If a password was previously set, the following messages is displayed:

```
Send the following string to Customer Support for password recovery:
afHTpyLsDo1Pz0Pk5GzhIw==
Enter the supplied recovery password.
Recovery Password:
```

5. Enter the recovery password (string).

The recovery string must be from 8 through 40 alphanumeric characters in length. A random string that is 15 characters or longer is recommended for higher security. The firmware prompts for this password only once. It is not necessary to remember the recovery string, because it is displayed the next time you enter the command shell.

The following prompt is displayed:

```
New password:
```

6. Enter the boot PROM password, and then re-enter it when prompted. The password must be eight alphanumeric characters long (any additional characters are not recorded). Record this password for future use.

The new password is automatically saved.

7. Reboot the switch by entering the **reset** command at the prompt.

Setting the boot PROM password for Backbones and Directors with a recovery string

This procedure applies to the Brocade DCX 8510 Backbones and X6 Directors. The boot PROM and recovery passwords must be set for each CP blade.

1. Connect to the serial port interface on the standby CP blade, as described in [Connecting to Fabric OS through the Serial Port](#) on page 33.
2. Connect to the active CP blade over a serial or Telnet connection and enter the **haDisable** command to prevent failover during the remaining steps.
3. Reboot the standby CP blade by sliding the On/Off switch on the ejector handle of the standby CP blade to **Off**, and then back to **On**.

4. Press **Esc** within four seconds after the message "Press escape within 4 seconds.." is displayed.
5. When prompted, enter **2** to select the recovery password option.

- If no password was previously set, the following message is displayed:

```
Recovery password is NOT set. Please set it now.
```

- If a password was previously set, the following messages are displayed:

```
Send the following string to Customer Support for password recovery:
afHTpyLsDo1Pz0Pk5GzhIw==
Enter the supplied recovery password.
Recovery Password:
```

6. Enter the recovery password (string).

The recovery string must be from 8 through 40 alphanumeric characters in length. A random string that is 15 characters or longer is recommended for higher security. The firmware only prompts for this password once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

The following prompt is displayed:

```
New password:
```

7. Enter the boot PROM password, and then re-enter it when prompted. The password must be eight alphanumeric characters long (any additional characters are not recorded). Record this password for future use.

The new password is automatically saved (the **saveEnv** command is not required).

8. Connect to the active CP blade over a serial or Telnet connection and enter the **haEnable** command to restore high availability, and then fail over the active CP blade by entering the **haFailover** command.

Traffic flow through the active CP blade resumes when the failover is complete.

9. Connect the serial cable to the serial port on the new standby CP blade (previously the active CP blade).
10. Repeat step 2 through step 7 for the new standby CP blade (each CP blade has a separate boot PROM password).
11. Connect to the active CP blade over a serial or Telnet connection and enter the **haEnable** command to restore high availability.

Although you can set the boot PROM password without also setting the recovery string, it is strongly recommended that you set both the password and the string as described in [Setting the boot PROM password for a switch with a recovery string](#) on page 159. If your site procedures dictate that you must set the boot PROM password without the string, follow the procedure that applies to your switch model.

Setting the boot PROM password for a switch without a recovery string

This procedure applies to fixed-port switches.

The password recovery instructions provided within this section are only for the switches listed in the supported hardware section. If your switch is not listed, contact your switch support provider for instructions.

1. Create a serial connection to the switch as described in [Connecting to Fabric OS through the Serial Port](#) on page 33.
2. Reboot the switch by entering the **reboot** command.
3. Press **Esc** within four seconds after the message "Press escape within 4 seconds.." is displayed.
4. When prompted, enter **3** to enter the command shell.
5. At the shell prompt, enter the **passwd** command.

The **passwd** command only applies to the boot PROM password when it is entered from the boot interface.

6. Enter the boot PROM password at the prompt, and then re-enter it when prompted. The password must be eight alphanumeric characters long (any additional characters are not recorded). Record this password for future use.
7. Enter the **saveEnv** command to save the new password.
8. Reboot the switch by entering the **reset** command.

Setting the boot PROM password for a Backbone or Director without a recovery string

This procedure applies to the Brocade DCX 8510 Backbones and X6 Directors.

On the Brocade DCX 8510 Backbones, set the password on the standby CP blade, fail over, and then set the password on the previously active (now standby) CP blade to minimize disruption to the fabric.

1. Determine the active CP blade by opening a Telnet session to either CP blade, connecting as admin, and entering the **haShow** command.
2. Connect to the active CP blade over a serial or Telnet connection and enter the **haDisable** command to prevent failover during the remaining steps.
3. Create a serial connection to the standby CP blade as described in [Connecting to Fabric OS through the Serial Port](#) on page 33.
4. Reboot the standby CP blade by sliding the On/Off switch on the ejector handle of the standby CP blade to **Off**, and then back to **On**.

This causes the blade to reset.

5. Press **Esc** within four seconds after the message "Press escape within 4 seconds..." is displayed.
6. When prompted, enter **3** to enter the command shell.
7. Enter the **passwd** command at the shell prompt.

The **passwd** command applies only to the boot PROM password when it is entered from the boot interface.

8. Enter the boot PROM password at the prompt, and then re-enter it when prompted. The password must be eight alphanumeric characters long (any additional characters are not recorded). Record this password for future use.
9. Enter the **saveEnv** command to save the new password.
10. Reboot the standby CP blade by entering the **reset** command.
11. Connect to the active CP blade over a serial or Telnet connection and enter the **haEnable** command to restore high availability, and then fail over the active CP blade by entering the **haFailover** command.

Traffic resumes flowing through the newly active CP blade after it has completed rebooting.

12. Connect the serial cable to the serial port on the new standby CP blade (previously the active CP blade).
13. Repeat step 3 through step 10 for the new standby CP blade.
14. Connect to the active CP blade over a serial or Telnet connection and enter the **haEnable** command to restore high availability.

NOTE

To recover lost passwords, refer to the *Brocade Fabric OS Troubleshooting and Diagnostics Guide*.

Remote authentication

Fabric OS supports user authentication through the local user database or one of the following external authentication services:

- Remote authentication dial-in user service (RADIUS)
- Lightweight Directory Access Protocol (LDAP) using Microsoft Active Directory in Windows or OpenLDAP in Linux.
- LDAP is supported with Windows 2008-2012 running LDAP 2003-2008 schema.
- Terminal Access Controller Access-Control System Plus (TACACS+)

Remote authentication configuration

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. 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 Backbone should be included in the authentication server configuration.

NOTE

For systems such as the Brocade DCX Backbone, 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 such systems, make sure that the CP IP addresses are used.

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

By default, the remote authentication services are disabled, so AAA services default to the switch's local database.

To enable remote authentication, it is strongly recommended that you 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 Backbone, 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.

Brocade recommends configuring 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 client certificate for the RADIUS server on the switch. This capability is not available on Gen 5 devices.

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, LDAP, and TACACS+ support, the way a switch authenticates users depends on whether a RADIUS, LDAP, or 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.

Supported LDAP options

The following table summarizes the various LDAP options and Brocade support for each.

TABLE 30 LDAP options

Protocol	Description	Channel type	Default port	URL	Brocade supported?
LDAPv3	LDAP over TCP	Unsecured	389	ldap://	No
LDAPv3 with TLS extension	LDAPv3 over TLS	Secured	389	ldap://	Yes (using starttls)
LDAPv3 with TLS and Certificate	LDAPv3 over TLS channel and authenticated using a certificate	Secured	389	ldap://	Yes (using starttls)
LDAPv2 with SSL ¹	LDAPv2 over SSL. Port 636 is used for SSL. Port 389 is for connecting to LDAP.	Secured	636 and 389	ldaps://	No

Command options

The following table outlines the *aaaConfig* command options used to set the authentication mode.

TABLE 31 Authentication configuration options

aaaConfig options	Description
--authspec "local"	Default setting. Authenticates management connections against the local database only. If the password does not match or the user is not defined, the login fails.

¹ This protocol was deprecated in 2003 when LDAPv3 was standardized.

TABLE 31 Authentication configuration options (continued)

aaaConfig options	Description
--authspec "radius"	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 fails back to the AAA server if configured and then to the local switch.
--authspec "radius;local"	Authenticates management connections to the first reachable RADIUS database. If RADIUS fails <i>for any reason</i> , authenticates against the local user database.
--authspec "radius;local" -backup	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 --backup option directs the service to try the secondary authentication database only if the primary authentication database is not available. For example, if user1 has a wrong password on server1 and server2 has the correct password for user1, authentication does query server2 until server1 is reachable. Only when server1 is not reachable, server2 is queried.
--authspec "ldap"	Authenticates management connections against any LDAP databases only. If LDAP service is not available or the credentials do not match, the login fails.
--authspec "ldap; local"	Authenticates management connections against any LDAP databases first. If LDAP fails for any reason, it then authenticates against the local user database.
--authspec "ldap; local" -backup	Authenticates management connections against any LDAP databases first. If LDAP fails for any reason, it then authenticates against the local user database. The --backup option states to try the secondary authentication database only if the primary authentication database is not available.
--authspec "tacacs+"	Authenticates management connections against any TACACS+ databases only. If TACACS+ service is not available or the credentials do not match, the login fails.
--authspec "tacacs+; local"	Authenticates management connections against any TACACS+ databases first. If TACACS+ fails for any reason, it then authenticates against the local user database.
--authspec "tacacs+; local" -backup	Authenticates management connections against any TACACS+ databases first. If TACACS+ fails for any reason, it then authenticates against the local user database. The --backup option states to try the secondary authentication database only if the primary authentication database is not available.
--authspec <auth-type> -nologout	Prevents users from being logged out when you change authentication. Default behavior is to log out users when you change authentication.
--authspec <auth-type> -logpriauth [yes/no]	Suppresses log messages for authentication failure by primary AAA service if authentication is to be done using secondary AAA service which is local switch database. The default value is set to "Yes" and displays login failure messages from the primary AAA service always. This is allowed for LDAP, RADIUS, and TACACS+.

Setting the switch authentication mode

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **aaaConfig --authspec** command.

Fabric OS user accounts

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.

RADIUS, LDAP, and TACACS+ support all the defined RBAC roles described in [Role-Based Access Control](#) on page 145.

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 in a *Brocade Vendor-Specific Attribute* (VSA). If the response does not have a VSA permissions assignment, the user role is assigned or the authentication is rejected.

You can set a user password expiration date and add a warning for RADIUS login and TACACS+ login. The password expiry date must be specified in UTC and in MM/DD/YYYY format. The password warning specifies the number of days prior to the password expiration that a warning of password expiration notifies the user. You either specify both attributes or none. If you specify a single attribute or there is a 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 your RADIUS server and once in the VSA. If the dates do not match, then the RADIUS server authentication fails.

[Table 32](#) describes the syntax used for assigning VSA-based account switch roles on a RADIUS server.

TABLE 32 Syntax for VSA-based account roles

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 SwitchAdminUser ZoneAdmin
	2	<i>Optional:</i> Specifies the Virtual Fabric member list. For more information on Virtual Fabrics, refer to RADIUS configuration with Virtual Fabrics on page 168. Brocade-AVPairs1
	3	Brocade-AVPairs2
	4	Brocade-AVPairs3
	5	Brocade-AVPairs4
	6	Brocade Password ExpiryDate
	7	Brocade Password ExpiryWarning
Vendor length	2 or higher	1 octet, calculated by server, including vendor-type and vendor-length

TABLE 32 Syntax for VSA-based account roles (continued)

Item	Value	Description
Attribute-specific data	ASCII string	Multiple octet, maximum 253, indicating the name of the assigned role and other supported attribute values.

NOTE

In VF mode, if chassis role is assigned as root to any AAA user, then the user login/authentication fails because login for AAA (RADIUS or TACACS+ or LDAP) user with chassis role as root is not allowed in VF mode. Chassis role of remote user is not applicable in non-VF mode.

Configuring role-based LDAP user attributes

Vendor 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 have to 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 already mapped using the `ldapcfg --maprole` command.

Consider the following when configuring role-based LDAP user attributes:

- If these attributes exist both on the server and the switch, then the attribute values from the server are used.
- If these attributes *do not* exist on either the server or the switch, then default values will be used for the attributes.

The priority of configurations is the following order:

1. Configurations on the server.
 2. Configurations on the switch.
 3. Default values (Role: user, Home LF: Default switch) are used when there are no values on both the server and the switch.
1. Map the LDAP role to the switch role.

```
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.

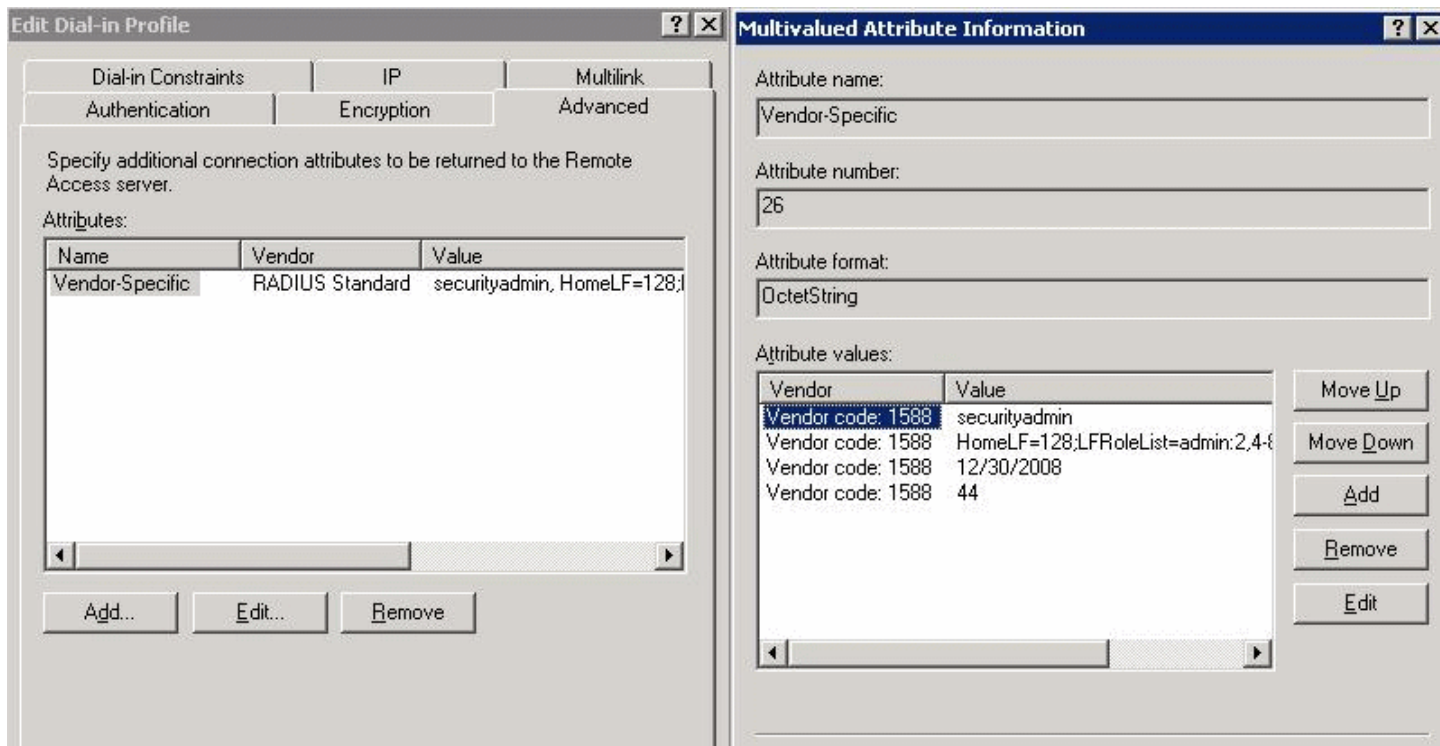
Fabric OS users on the RADIUS server

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 6 Windows 2012 VSA configuration



Linux FreeRADIUS server

For the configuration on a Linux FreeRADIUS server, define the values outlined in Table 33 in a vendor dictionary file called dictionary.brocade.

TABLE 33 Entries in dictionary.brocade file

Include	Key	Value
VENDOR	Brocade	1588
ATTRIBUTE	Brocade-Auth-Role	1 string Brocade
	Brocade-AVPairs1, 2, 3, 4	2, 3, 4, 5 string Virtual Fabric member list
	Brocade-Passwd-ExpiryDate	6 string MM/DD/YYYY in UTC
	Brocade-Passwd-WarnPeriod	7 integer in days

After you have completed the dictionary file, define the permissions for the user in a configuration file. For example, to grant the user admin permissions, you would add the following statement to the configuration file:

```
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"
```

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 Brocade-Auth-Role. The additional attribute values HomeLF and LFRoleList are optional. If they are unspecified, the account can log in with VF128 as its member list and home Virtual Fabric. If there is an error in the LFRoleList or HomeLF specification, the account cannot log in until the Virtual Fabric list is corrected; an error message is displayed.

In the next example, on a Linux FreeRADIUS Server, the user has the "zoneAdmin" permissions, 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"
```

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 Backbones, 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 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.

NOTE

The combination of "peap-mschapv2" and IPv6 when used together, rejects RADIUS authentication. PEAP with IPv4 succeeds. Combination of PEAP-MSCHAPv2 and IPv6 is not supported and hence blocked during configuration. >

When using a fully qualified domain name (FQDN) for RADIUS server, configuring the DNS to return an IPv6 address may be resolved to 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 file `$PREFIX/etc/raddb/dictionary.brocade` 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 file `$PREFIX/etc/raddb/dictionary` in a text editor and add the line:

```
$INCLUDE dictionary.brocade
```

As a result, the file `dictionary.brocade` 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 specified with Brocade-Auth-Role. The valid permissions include admin, switchAdmin, zoneAdmin, securityAdmin, basic SwitchAdmin, fabricAdmin, operator, and user. You must use quotation marks around "password" and "role".

NOTE

Allowing AAA accounts to login to switch with root privileges is a vulnerability and is not permitted.

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, 2008 and a warning period of 30 days:

```
JohnDoe Auth-Type := Local
User-Password == "johnPassword",
Brocade-Auth-Role = "admin",
Brocade-Passwd-ExpiryDate = "05/28/08",
Brocade-Passwd-WarnPeriod = "30"
```

Example of using the 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 Backbones 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/radddb/client.config` 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 10.32.170.59 as a client, the following is displayed:

```
client 10.32.170.59
  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](#) on page 186).

2. Save the file `$PREFIX/etc/raddb/client.config`, 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 or your Microsoft documentation. Confer with your system or network administrator prior to configuration for any special needs that your network environment might have.

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 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, then 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, then 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 will 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 Active Directory Users and Computers tool is available.** If not, select Server Manager > Manager tab > 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 displayed and click Next.

The Features tab is displayed.

- d) From the Features tab, select all the checkboxes starting from the Remote Server Administration Tools checkbox till the AD DS Snap-Ins and Command-Line Tools element checkbox in the navigation tree.
- e) From the Feature tab, select the Group Policy Management checkbox, and then click Next.

The Confirmation tab is displayed.

- f) Select the Restart the destination server automatically if required checkbox 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 will 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 pop-up window.

A remote access policy must be created 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 the **Yes. It conforms** option, and then click **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 box 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 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 box 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 box 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 box 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 box is displayed.
 - i) Select RADIUS Standard from the list and click Add.
The Vendor Specific Attribute Information dialog box is displayed.
 - j) Select RADIUS Standard from the list and choose the Yes, it confirms option. Click OK.
The Configure VSA (RFC Compliant) dialog box 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 box 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 box 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 box 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 box 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 client machine.

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 either something a user has, such as the RSA SecurID token or something uniquely identifies a user such as a biometric.

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.

To learn more about how RSA SecurID works, visit www.rsa.com for more information.

Setting up the RSA RADIUS server

For more information on how to install and configure the RSA Authentication Manager and the RSA RADIUS server, refer to your documentation or visit www.rsa.com.

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 `dictionary.dcm` file located in the following path:

```
C:\Program Files\RSA Security\RSA 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-Passwd-ExpiryDate Brocade-VSA(6,string) r
ATTRIBUTE Brocade-Passwd-WarnPeriod Brocade-VSA(7,integer) r
#####
# brocade.dct -- Brocade Dictionary
#####
```

Example of the `dictionary.dcm` file

```
#####
# dictionary.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 `brocade.dct` file should look like and Example of the `dictiona.dcm` file shows what needs to be modified in the `dictiona.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 **Brocade-Auth-Role** and type the string **Admin**. The string you type equals the role on the switch.
- e) Add the Brocade profile.
- f) In **RSA Authentication Manager**, edit the user records that will be authenticated using RSA SecurID.

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 discusses authentication and authorization using Microsoft Active Directory. For information about authentication and authorization using OpenLDAP, refer to [LDAP configuration and OpenLDAP](#) on page 178.

Two operational modes exist in LDAP authentication, FIPS mode and non-FIPS mode. This section discusses LDAP authentication in non-FIPS mode. For more information on LDAP in FIPS mode, refer to [Configuring Security Policies](#) on page 233.

The following are restrictions when using LDAP in non-FIPS mode:

- Passwords cannot be changed through Active Directory.
- 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. Common Name-based authentication 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, in Fabric OS 7.1.0 and later releases, 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, refer to [Role-Based Access Control](#) on page 145.

NOTE

All instructions involving Microsoft Active Directory can be obtained from 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.

Configuring Microsoft Active Directory LDAP service

The following is an overview of the process used to set up LDAP.

1. If your Windows Active Directory server for LDAP needs to be 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 Microsoft Active Directory server.

For instructions on how to create a user, refer to www.microsoft.com or Microsoft documentation to create a user in your Active Directory.

3. Create a group name that uses the switch's role name so that the Active Directory group's name is the same as the switch's role name.

or

Use the `ldapCfg --maprole ldap_role_name switch_role` and `--mapattr` commands to map an LDAP server role to one of the default roles available on the switch.

4. Associate the user to the group by adding the user to the group.
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 entered as a range.

Creating a user

To create a user in Active Directory, refer to www.microsoft.com or Microsoft documentation. There are no special attributes to set. You can use a fully qualified name for logging in; for example, you can log in as "user@domain.com".

Creating a group

To create a group in Active Directory, refer to www.microsoft.com or Microsoft documentation. You must verify that the group has the following attributes:

- The name of the group must match the RBAC role.
- The Group Type must be **Security**.
- The Group Scope must be **Global**.
- 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.
- If the user you created is not a member of the Users OU, then the User Principal Name, in the format of "user@domain", is required to log in.

Assigning the group (role) to the user

To assign the user to a group in Active Directory, refer to www.microsoft.com or Microsoft documentation. If you have a user-defined group, use the `ldapCfg --maprole` command to map LDAP server permissions to one of the default roles available on a switch. Alternatively, update the `memberOf` field with the login permissions (admin, switchAdmin, user, and so on) that the user must use to log in to the switch.

NOTE

Allowing AAA accounts to login to switch with root privileges is a vulnerability and is not permitted.

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** attribute.

The String Attribute Editor dialog box displays.

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, then the lowest 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.

Adding attributes to the Active Directory schema

To create a group in Active Directory, refer to www.microsoft.com or Microsoft documentation. You must:

- Add a new attribute **brcdAdVfData** as Unicode String.
- Add **brcdAdVfData** to the person's properties.

LDAP configuration and OpenLDAP

Fabric OS provides user authentication and authorization by means of OpenLDAP or the Microsoft Active Directory service in conjunction with LDAP on the switch. This section discusses authentication and authorization using OpenLDAP. For information about authentication and authorization using Microsoft Active Directory, refer to [LDAP configuration and Microsoft Active Directory](#) on page 176.

Two operational modes exist in LDAP authentication: FIPS mode and non-FIPS mode. This section discusses LDAP authentication in non-FIPS mode. For information on LDAP in FIPS mode, refer to [Configuring Security Policies](#) on page 233. When using OpenLDAP in non-FIPS mode, you must use the Common-Name for OpenLDAP authentication. User-Principal-Name is not supported in OpenLDAP. OpenLDAP 2.4.23 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 RBAC roles, refer to [Role-Based Access Control](#) on page 145.

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 needs to be verified by the LDAP client (that is, the Brocade switch), then 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 entries (users) in the OpenLDAP Directory.
4. Assign users to groups by using the **member** attribute.
5. Use the **ldapCfg --maprole ldap_role_name switch_role** command to map an LDAP server role to one of the default 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.

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 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 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](#) on page 179 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
```

Assigning the LDAP role to a switch role

Use the `ldapCfg --maprole` command to map LDAP server permissions to one of the default roles available on a switch.

Mapping vendor-specific attributes to LDAP user roles

Use the `ldapCfg --mapattr` command to configure the vendor-specific attributes, such as chassis role, home LF, and LF list to LDAP roles already mapped using the `ldapcfg --maprole` command.

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.
to delete a user attribute

Adding a Virtual Fabric list

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.

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 ( 1.3.6.1.4.1.8412.100
    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 ( 1.3.6.1.4.1.8412.110 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 provided in [Adding a Virtual Fabric list](#).

```
include /usr/local/etc/openldap/schema/local.schema
```

3. Include the **brcdAdvfData** attribute in a user entry in the LDAP directory.
 - If you are using Virtual Fabrics, enter the value of the logical fabrics to which the user has access. Up to three value fields can be specified, separated by an semicolons (;):
 - The HomeLF field specifies the user's home Logical Fabric.
 - The LFRole list field specifies the additional Logical Fabrics to which the user has access and the user's access permissions for those Logical Fabrics. Logical Fabric numbers are separated by commas (,). A hyphen (-) indicates a range.
 - The ChassisRole field designates the permissions that apply to the ChassisRole subset of commands.

Example for adding Virtual Fabrics

In the following example, the logical switch that would be logged in to by default is 10. If 10 is not available, then the lowest FID available will be chosen. The user is given permission to enter logical switches 1 through 128 in an admin role and is also given the chassis role permission of admin.

```
brcdAdvfData: HomeLF=10;LFRoleList=admin:1-128;ChassisRole=admin
```

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.

```
switch:admin> ldapadd -D cn=Sachin,dc=mybrocade,dc=com -x -w secret -f test4.ldif
```

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 or clients. Once configured to use TACACS+, a Brocade switch becomes a Network Access Server (NAS).

The following authentication protocols are supported by the TACACS+ server for user authentication:

- Password Authentication Protocol (PAP)
- Challenge Handshake Authentication Protocol (CHAP)

TACACS+ is not a FIPS-supported protocol, so you cannot configure TACACS+ in FIPS mode. To enable FIPS, any TACACS+ configuration must be removed.

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, provide lists of Virtual Fabrics to which the user should have access. For details, refer to [The tac_plus.cfg file](#) on page 183.

On the Brocade switch, use the **aaaConfig** command to configure the switch to use TACACS+ for authentication. The **aaaConfig** command also allows you to specify up to five TACACS+ servers. When a list of servers is configured, failover from one server to another server happens only if a TACACS+ server fails to respond. It does not happen when user authentication fails.

Failover to another TACACS+ server is achieved by means of a timeout. You can configure a timeout value for each TACACS+ server, so that the next server can be used in case 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 5 attempts. If authentication is rejected or times out, Fabric OS will try again. The retry value can also be customized for each user.

Refer to [Remote authentication configuration on the switch](#) on page 186 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 [The tac_plus.cfg file](#) on page 183 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
```

The tac_plus.cfg file

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. [Table 34](#) lists and defines attributes specific to Brocade.

TABLE 34 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

Adding a user and assigning a role

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.

Refer to [Role-Based Access Control](#) on page 145 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 = exec {
        brcd-role = securityAdmin;
    }
}
```

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

FabricOS 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 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.

Obfuscation of RADIUS and TACACS+ shared secret

You can store and display the RADIUS and TACACS+ server shared secret either as plain text or in an encrypted format. For encrypting the RADIUS server shared secret, both the CPs in a chassis must be running Fabric OS 7.4.0 or later. For encrypting the TACACS+ server shared secret, both the CPs in a chassis must be running Fabric OS 8.1.0 or later.

If you choose to encrypt the shared secret, it is applied during the following operations:

- Set the configuration
- Upload configuration
- Download configuration
- Upload supportSave information
- Download supportSave information
- Upgrade firmware
 - When you upgrade to Fabric OS 7.4.0 or later, the RADIUS shared secret remains as plain text. You must explicitly set the encryption on.
 - When you downgrade to Fabric OS 7.3.0 or earlier, and then back to Fabric OS 7.4.0 or later, the RADIUS shared secret encryption setting is unchanged.
 - When you upgrade to Fabric OS 8.1.0 or later, the TACACS+ shared secret remains as plain text. You must explicitly set the encryption on.
- Downgrade firmware
 - If the RADIUS shared secret encryption is set to on and you downgrade to an earlier version of Fabric OS 7.4.0, then you are prompted to set the encryption to either **none** or **aes256** before downgrading. You are also prompted to remove radius configurations and set the encryption level to none.
 - If the TACACS+ shared secret encryption is set to **aes256** and you downgrade to an earlier version of Fabric OS 8.1.0, then you are prompted to remove the TACACS+ configurations or set the encryption level to none before downgrading.
- The **show** command output

To enable encryption, use the **-e** or **-encr_type** option in the following commands.

```
aaaconfig --add server -conf radius|ldap|tacacs+ [-p port] [-s secret] [-t timeout] [-d domain-name] [-a chap|pap|peap-mschapv2] [-e -encr_type encryption_level]
```

```
--change server -conf radius|ldap|tacacs+ [-p port] [-s secret] [-t timeout] [-a chap|pap|peap-mschapv2] [-d domain-name] [-e -encr_type encryption_level]
```

NOTE

For **-conf ldap**, **-a**, **-s** and **-e** options are not applicable.

Remote authentication configuration on the switch

At least one RADIUS, LDAP, or TACACS+ server must be configured before you can enable a remote authentication service. You can configure the remote authentication service even if it is disabled on the switch. You can configure up to five RADIUS, LDAP, or TACACS+ servers. You must be logged in as `admin` or `switchAdmin` to configure the RADIUS service.

NOTE

On DCX 8510 Backbones and X6 Directors, the switch sends its RADIUS, LDAP, or TACACS+ request using the IP address of the active CP. When adding clients, add both the active and standby CP IP addresses so that users can still log in to the switch in the event of a failover.

RADIUS, LDAP, or TACACS+ configuration is chassis-based configuration data. On platforms containing 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 authentication servers are contacted in the order they are listed, starting from the top of the list and moving to the bottom.

Adding an authentication 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 authentication 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.

Enabling and disabling remote authentication

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `aaaConfig --authspec` command to enable or disable RADIUS, LDAP, or TACACS+.

You must specify the type of service as one of RADIUS, LDAP, or TACACS+. Local is used for local authentication if the user authentication fails on the authentication server.

Example for enabling RADIUS

```
switch:admin> aaaconfig --authspec "radius;local" -backup
```

Example for enabling LDAP

```
switch:admin> aaaconfig --authspec "ldap;local" -backup
```

Example for enabling TACACS+

```
switch:admin> aaaconfig --authspec "tacacs+;local" -backup
```

3. You can configure the `logpriauth` option as `yes` or `no`.

```
switch:admin> aaaconfig --authspec <"aaa1[;aaa2]"> [-backup] [-nologout] [logpriauth <yes/no>]
```

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. Parameters include:

TABLE 35 Current authentication configuration parameters

Parameter	Description
Position	The order in which servers are contacted to provide service.
Server	The server names or IPv4 or IPv6 addresses. IPv6 is not supported when using PEAP authentication.
Port	The server ports.
Secret	The shared secrets.
Timeouts	The length of time servers have to respond before the next server is contacted.
Authentication	The type of authentication being used on servers.

Configuring local authentication as backup

It is useful to enable local authentication, so that the switch can take over authentication locally if the RADIUS or LDAP servers fail to respond because of power outage or network problems.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **aaaConfig --authspec** command to enable or disable RADIUS, LDAP, or TACACS+ with local authentication as a backup authentication mechanism.

You must specify the type of service as one of RADIUS, LDAP, or TACACS+. Local is used for local authentication if the user authentication fails on the authentication server.

Example for RADIUS

Example of enabling local authentication as a backup for RADIUS.

```
switch:admin> aaaconfig --authspec "radius;local" -backup
```

Example for LDAP

```
switch:admin> aaaconfig --authspec "ldap;local" -backup
```

Example for TACACS+

```
switch:admin> aaaconfig --authspec "tacacs+;local" -backup
```

For details about the **aaaConfig** command refer to [Command options](#) on page 163.

When local authentication is enabled and the authentication servers fail to respond, you can log in to the default switch accounts (admin and user) or any user-defined account. You must know the passwords of these accounts.

When the **aaaConfig** command succeeds, the event log indicates that local database authentication is disabled or enabled.

Configuring Protocols

- Security Protocols..... 189
- Secure Copy..... 190
- Secure Shell protocol..... 191
- SecCryptoCfg templates..... 194
- Secure Sockets Layer protocol 207
- Simple Network Management Protocol..... 216
- Telnet protocol..... 229
- Listener applications..... 230
- Ports and applications used by switches..... 231

Security Protocols

Security protocols provide endpoint authentication and communications privacy using cryptography.

Endpoint authentication can be either single authentication or mutual authentication. Typically, a user is authenticated to the switch, which is single authentication; or a TLS client and a server can authenticate each other, which is mutual authentication. This means that the user can be sure which switch they are communicating with.

Fabric OS supports the secure protocols shown in the following table.

TABLE 36 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 used to indicate a secure HTTP connection. Web Tools supports the use of Hypertext Transfer Protocol over SSL (HTTPS).
IPsec	Internet Protocol Security (IPsec) is a framework of open standards for providing confidentiality, authentication and integrity for IP data transmitted over untrusted links or networks.
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 in conjunction with LDAP.
RADIUS	RADIUS uses TLS when configured with PEAP-MSCHAPv2. The CA root certificate for RADIUS can be imported using the seccertmgmt 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 seccertmgmt 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 warrant 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 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 HTTPS. A certificate must be generated and installed on each switch to enable SSL. Supports SSL3, 128-bit encryption by default. Also supports TLS 1.0, TLS 1.1, and TLS 1.2. TLS 1.2 is recommended for stronger encryption.

Fabric OS supports the following features:

- Challenge Response Authentication (CRA) is supported in SCP and SSH sessions.
- RADIUS, LDAP, and Syslog client authentication and chain of certificates validation if connection to an external server is established over secure TLS connection. For more information refer to [#unique_332](#).
- TLS protocol version configuration and Syslog cipher configuration using the `seccryptocfg` command template.

Secure Protocol Deployment Requirements

[Table 37](#) describes additional software or certificates that you must obtain to deploy secure protocols.

TABLE 37 Items needed to deploy secure protocols

Protocol	Host side	Switch side
SSHv2	Secure shell client	None
HTTPS	No requirement on host side except a browser that supports HTTPS	Switch IP certificate for SSL or FQHN of switch
SCP	SSH daemon, SCP server	None
SNMPv3	None	None

Security Protocol Scenarios

Security protocols are designed with the four main use cases, as described in [Table 38](#).

TABLE 38 Main security scenarios

Fabric	Management interfaces	Comments
Nonsecure	Nonsecure	No special setup is needed 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 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, refer to Installing a Switch Certificate on page 213.
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# configure

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
Configure...
System services (yes, y, no, n): [no] n
ssl attributes (yes, y, no, n): [no] n
http attributes (yes, y, no, n): [no] n
snmp attributes (yes, y, no, n): [no] n
rpcd attributes (yes, y, no, n): [no] n
cfgload attributes (yes, y, no, n): [no] y
Enforce secure config Upload/Download (yes, y, no, n): [no]# y
Enforce signature validation for firmware (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. If you start an SSH session, and then use the **login** command to start a nested SSH session, commands that require a secure channel will be rejected.

Fabric OS supports OpenSSH version 6.2p2 and OpenSSL version 1.0.2h with the heartbeat feature disabled .

If you set up a message of the day (MOTD), the MOTD displays either before or after the login prompt, depending on the SSH client implementation. Fabric OS does not control when the message displays.

SSH public key authentication

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 utilize outgoing OpenSSH public key authentication. Any admin user can perform incoming Open SSH public key authentication.

Using OpenSSH RSA, DSA, and 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. ECDSA asynchronous algorithms are FIPS-compliant.

Incoming authentication is used when the remote host needs to authenticate to the switch. Outgoing authentication is used when the switch needs to 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 authentication 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. While creating the key pair, the configured allowed-user can choose a passphrase with which the private key is encrypted. Then the passphrase must always be entered when authenticating to the switch. The allowed-user must have admin permissions to perform OpenSSH public key authentication, import and export keys, generate a key pair for an outgoing connection, and delete public and private keys.

NOTE

There are no special privileges 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 RSA/DSA 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.

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 perform SSH public key authentication, and who can import, 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** command.

You may enter a passphrase for additional security.

Example of generating a key pair on the switch

```
switch:alloweduser> sshutil genkey [-rsa|-dsa|-ecdsa]
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
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** command to export the key.

Example of exporting a public key from the switch

```
switch:alloweduser> sshutil exportpubkey

Enter IP address:192.168.38.244

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` command to delete public keys.

You will be prompted to enter the name of the user whose public keys you want to delete. Enter **all** to delete public keys for all users.

For more information on IP filter policies, refer to [Configuring Security Policies](#) on page 233.

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` command to delete the private key.

For more information on IP filter policies, refer to [Configuring Security Policies](#) on page 233.

Generating and installing hostkeys on a switch

You can generate, install, display, and delete SSH hostkeys on a switch.

1. Log in to the switch as the allowed user.
2. Enter the `sshutil genhostkey [-rsa | -dsa | -ecdsa]` command to generate and install the SSH hostkey.
3. Enter the `sshutil showhostkey` command to display current SSH hostkeys installed on the switch.
4. Enter the `sshutil delhostkey [-rsa | -dsa | -ecdsa]` command to delete the selected SSH hostkeys on the switch.
5. 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
```

6. Enter the `sshutil showrekey` to display the SSH rekey interval.

```
switch:admin> sshutil showrekey
```

SecCryptoCfg templates

SecCryptoCfg templates permit you to predefine standardized sets of values for switch cipher configurations. These can then 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 RADIUS, LDAP, HTTPS, SSH ciphers, Syslog ciphers, SSH key exchange algorithms, and SSH MAC values.

A template can be specific to the requirements of a certification or based on the definition of 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 configurations but can upload the configurations, 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.

The following default templates are provided. Refer to [Default secCryptoCfg templates](#) on page 198 for an example of each template.

- Default Configuration (default_generic)
- Secured configuration (default_strong)
- FIPS configuration (default_fips)

- CC configuration (default_cc)

Template structure

The following table lists the groups within each template and the available options. Refer to [Default secCryptoCfg templates](#) on page 198 for illustrations.

TABLE 39 SecCryptoCfg template groups and options

Group	Option	Function
Version	Number	Version number of the template
SSH	Kex, Mac, Enc	Encryption modes
AAA	RAD_Ciphers, LDAP_Ciphers RAD_Protocol, LDAP_Protocol	Encryption ciphers
LOG	Ciphers, Protocol	Syslog ciphers and protocols
HTTPS	Ciphers, Protocol	HTTPS ciphers and protocol
FIPS	SelfTests, BootProm, Zeroize, Inside	This is present in the FIPS template only. The valid options for these FIPS functions are "on" and "off" for each function.
X509v3	Strict or Basic	X509v3 certificate validation permits Fabric OS to enable or disable certificate validation during certificate import and session establishment. Refer to X509v3 validation modes for details on the modes and the differences between them. User-defined templates may have either basic or strict validation mode.

The following points should be taken into consideration 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 does not include the template files. To import, export, apply, and delete templates, use the template commands described in [SecCryptoCfg template management](#) on page 197.
- If you are downgrading, the script checks for Protocol entries (example: `Syslog_Protocol:TLSv1.2`), and the Log and Validation sections in user-created templates. If you are using a user-created template and any protocol is listed or either of these sections is present, the downgrade will be blocked. To correct this, either apply one of the default templates, or edit the custom template to remove the problem sections and then reapply the template using the `secryptocfg` command.

Certificate terminology

TABLE 40 Certificate terminology

Term	Description
Root CA certificate	A public key certificate that identifies 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, used to authenticate client to server or server to client respectively.
Certificate Chain	A chain of certificates that contain the authority relationship for the certificate authority.
Online Certificate Status Protocol (OCSP)	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 X509v3 certificate validation.

TABLE 41 X509v3 validation modes described

Serial #	Basic Mode	Strict Mode
1	"Certificate Sign" is not mandatory for "KeyUsage" in a CA certificate.	"Certificate Sign" is mandatory for "KeyUsage" in a CA certificate.
2	OCSP validation for identity certificates is not performed.	OCSP validation for identity certificates is performed.
3	"BasicConstraints" field is not mandatory in a CA certificate.	"BasicConstraints" field is mandatory in a CA certificate.
4	"CA:True" value is not mandatory in a CA certificate.	"CA:True" value is mandatory in a CA certificate.
5	"Digital Signature" is not mandatory for "KeyUsage" in a identity certificate.	"Digital Signature" is mandatory for "KeyUsage" in a identity certificate.
6	Audit logs for establishing and terminating TLS session are not printed.	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.
7	In SSH sessions with a switch, the 1GB data transfer limit is not enforced for rekeying.	In SSH sessions with a switch, the 1GB data transfer limit is enforced for rekeying.
8	For RADIUS, LDAP, or Syslog secure authentication using a certificate, a hostname mismatch error is ignored.	For RADIUS, LDAP, or Syslog secure authentication using a certificate, a hostname mismatch error causes authentication failure.
9	For RADIUS, LDAP, or Syslog secure authentication using a certificate, hostname validation is not performed.	For RADIUS, LDAP, or Syslog secure authentication, hostname validation is enforced and a hostname mismatch error causes authentication failure (hostname resolution failure also causes authentication failure).
10	For RADIUS, LDAP, or Syslog secure authentication over TLS using a certificate, if there is a hostname mismatch error, no audit log entry is generated.	For RADIUS, LDAP, or Syslog secure authentication over TLS using a certificate, if there is a hostname 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 42 X509v3 certificate validation

X509 Field	Valid NDcPP 1.0 values	Validation behavior during certificate import	Validation behavior during session establishment
Not Before: Not After:	Certificate is validated for time and date during import and is accepted only if the time and date are within the allowed range.	Certificates that have expired or are not yet valid are not allowed to be imported. Import fails.	Session establishment fails if CA certificate or identity certificate used has expired or is not yet valid. In case of webtools, session can continue if an exception is added to the browser.
CN=	Needs to be FQN device name or IP. Wildcards are allowed only for 1 level of sub-domain and not allowed for the main domain.	Validates identity certificates against the IP or hostname of the host during import. Import fails if common name doesn't match with the IP or hostname or has a wildcard in any field other than the first level of sub-domain. The validation does not apply to commoncert.	Validates identity certificates against IP or hostname of the host. Applies to all identity certificates that are used for session establishment including commoncert. Wildcards are allowed only for the first level of sub-domain.

TABLE 42 X509v3 certificate validation (continued)

X509 Field	Valid NDcPP 1.0 values	Validation behavior during certificate import	Validation behavior during session establishment
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 and the import is successful.	Audit logs are generated for certificates that have keysize less than 2048. The session is established successfully.
CA:	Field must be TRUE for CA certificates.	Validated for CA certificates at import, and import fails if this is not true.	Validated for all CAs used during session establishment.
Key Usage:	Need to have "Certificate Sign" in case of CA certificates and "Digital Signature" in case of identity certificates. NOTE To support RSA-based ciphers for TLS sessions, Key encipherment must be set.	Validated for CA certificates and identity certificates (for "Certificate Sign" and "Digital Signature", respectively, and if validation fails, import fails.	Validated for CA certificates and identity certificates (for "Certificate Sign" and "Digital Signature", respectively, and if validation fails, session establishment fails,
X509v3 Extended Key Usage:	Need to rightly indicate whether it is for use as "server" certificate or "client" certificate. If incorrect, connection is not allowed.	Validated for correct EKU ("server" for identity certificates for servers and "client" for identity certificates of clients) and if incorrect, import fails.	Validated for correct EKU ("server" for identity certificates for servers and "client" for identity certificates of clients) and, if incorrect, session establishment fails.
Authority Information Access:	Must have valid OCSP server respond affirmatively. If this field is not present, OCSP check is not performed.	Not checked during import of certificate.	Checked for OCSP URL in certificate during session establishment. If present, then 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 Method	Not a mandatory attribute. If present, the values stored in this will take priority over the CN in Subject attribute.	Validated during import for all certificates.	Validated during sessions for all certificates.
Basic Constraints Attribute	Attribute must be present and must have CA Field.	Validated for CA certificates at import, and import fails if the attribute does not exist.	Validated for all CAs used during session establishment.

SecCryptoCfg template management

Use the following commands to manage the template files.

- To set a template file, use the **seccryptocfg --apply *template_file*** command.

NOTE

Using the **seccryptocfg --apply *template_file*** command is not an approved method for enabling cryptographic compliance in FIPS mode on a switch. Instead, use the **fipscfg** command. Refer to [Cryptographic compliance in FIPS mode](#) on page 206 for additional information.

- To import a template file, use the **seccryptocfg --import *template_file options*** command to import a template file.

```
switch:admin> seccryptocfg -import <template filename> [-server <ipaddr> -name <user> -proto <ftp|
scp> -file <remotefilename>]
<Interactive Prompt comes up to provide server details>
```

- To export a template file, use the **seccryptocfg --export *template_file options*** command.

```
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, 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.

Default_generic

```

/*****
* Brocade - Generic Template for Security Crypto Configuration
*
* Desc:
*
* Default values for security crypto configurations for operation in
* generic mode
*
*****/

[Ver] 0.1

/*
* 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-md5,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:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
LDAP_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM

/*
* Group : LOG
* Rules : Textual openssl cipherlist (colan,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[LOG]
Syslog_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!3DES:!aNULL
Syslog_Protocol:TLSv1.2

/*
* Group : HTTPS
* Rules : Textual openssl cipherlist (colon,comma, and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[HTTPS]
Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM

/*
* Group : X509v3
* Rules : Textual X509v3 validation options
* Example: Validation:Strict
* Valid options: Strict/Basic
*/
[X509v3]
Validation:Basic

```

Default_strong

Strong configuration consist of TLS 1.2 and SHA above 256.

```

/*****
* Brocade - Template for High Security Crypto Configuration
*
* Desc:
*
* Default values for security crypto configurations for high security
*
*****/

[Ver] 0.1

/*
* 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
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:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!SSLv3
LDAP_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!SSLv3

/*
* Group : LOG
* Rules : Textual openssl cipherlist (colan,comma and space separated)

```

```

* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[LOG]
Syslog_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!3DES:!aNULL
Syslog_Protocol:TLSv1.2

/*
* Group : HTTPS
* Rules : Textual openssl cipherlist (colan,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/

[HTTPS]
Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!SSLv3
/*

* Group : X509v3
* Rules : Textual X509v3 validation options
* Example: Validation:Strict
* Valid options: Strict/Basic
*/
[X509v3]
Validation:Basic

```

Default_fips

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] 0.1

/*
* 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
Mac:hmac-sha1,hmac-sha2-256,hmac-sha2-512

/*
* Group : AAA
* Rules : Textual openssl cipherlist (colan,comma and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: RAD_Ciphers, LDAP_Ciphers
*/
[AAA]
RAD_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
LDAP_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
RAD_Protocol:TLSv1.2
LDAP_Protocol:TLSv1.2

/*
* Group : LOG
* Rules : Textual openssl cipherlist (colan,comma and space separated)

```



```

* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[LOG]
Syslog_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!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:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
Protocol:TLSv1.2

/*
* Group : FIPS
* Rules : Applicable only to FIPS mode configuration. "yes" for configuration enabled and "no" for disabled
* Example: "yes" OR "no" OR "9.xx"(for Inside)
* Valid options: SelfTests, BootProm, Enable, Simulate
* Note: FIPS mode enable via seccryptocfg is not approved in this release.
*/
[FIPS]
SelfTests:yes
BootProm:yes
Zeroize:yes
Inside:yes

/*
* Group : X509v3
* Rules : Textual X509v3 validation options
* Example: Validation:Strict
* Valid options: Ciphers
*/
[X509v3]
Validation:Basic

```

Default_cc

```

/*****
* Brocade - Common Criteria (CC) Template for Security Crypto Configuration
*
* Desc:
*
* Default values for security crypto configurations for CC compliance
*
*****/

[Ver] 0.1

/*
* 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:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!3DES:!aNULL
LDAP_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!3DES:!aNULL
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:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!3DES:!aNULL
Syslog_Protocol:TLSv1.2

/*
* Group : HTTPS
* Rules : Textual openssl cipherlist (colon, comma, and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers
*/
[HTTPS]
Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!3DES:!aNULL
Protocol:TLSv1.2

/*
* Group : X509v3
* Rules : Textual X509v3 validation options
* Example: Validation:Strict
* Valid options: Strict/Basic
*/
[X509v3]
Validation:Strict

```

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 peer (SYSLOG), verify the peer certificate with the OCSP responder for revocation, if the OCSP URI is already present as part of the peer certificate.
- If 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 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. There is no separate OCSP certificate 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

OCSRP Responder Status	TLS session status
Reachable; No Response	Fail
Not reachable	Fail

Client authentication for TLS client - RADIUS/LDAP/SYSLOG

The following steps are applicable to all TLS clients. Only SYSLOG is NDcPP compliant.

- **secCertMgmt** provides options to support 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, then the TLS connection is successfully established only if the client certificate validation passes.
- If client certificate for RADIUS/LDAP/SYSLOG has not been imported on the switch, then the connection proceeds without client certificate validation (provided server does not mandate client certificate validation).
- The CLI enhancements made to support this feature are as shown below:

1. Import 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 client CA certificate, you must import the complete CA chain of certificates that is required to validate the client identity.

2. Import 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".

TLS protocol version configuration option for all TLS based applications

For all TLS-based applications updates have been made to the CLI as described in this section.

- The **seccryptocfg** command template has been enhanced to support the following:
 - TLS protocol version for all TLS based applications
 - Cipher configuration for SYSLOG
- A group "LOG" has been added in all templates to support configurations for SYSLOG
- TLS protocol version key value pairs are supported under the "AAA", "HTTPS" and "LOG" groups, to configure the TLS protocol versions to be used in case of RADIUS, LDAP, HTTPS and SYSLOG respectively.
- Support to configure SYSLOG ciphers has also been added under the "LOG" group in the template.

The template is shown below.

```
[Ver] 0.1

/*
* Group : SSH
* Rules : Comma Seperated
```

```

* 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-md5,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, RAD_protocol, LDAP_protocol
*/
[AAA]
RAD_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
LDAP_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
RAD_protocol:TLSv1
LDAP_protocol:TLSv1.2

/*
* Group : HTTPS
* Rules : Textual openssl cipherlist (colon, comma, and space separated)
* Example: ALL:-MD5:!PSK
* Valid options: Ciphers, Protocol
*/
[HTTPS]
Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
Protocol:TLSv1.2

/*
* Group : LOG
* Rules : Textual openssl cipherlist (colon, comma, and space separated)

* Example: ALL:-MD5:!PSK
* Valid options: Ciphers, Protocol
*/
[LOG]
SYSLOG_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
SYSLOG_Protocol:TLSv1

```

NOTE

The valid values for protocol key value pair are "TLSv1", "TLSv1.2" and "Any".

SSH rekey after specific time interval

The **sshutil** CLI has been enhanced to support rekeying after a specific time interval. The CLI options are shown below:

```

sshutil rekeyinterval <seconds>
sshutil showrekey

```

The **rekeyinterval** option with the parameter value in seconds sets the rekey interval timing. The **showrekey** option displays the value of the parameter.

- The daemon is restarted after configuration, and new sessions for clients rekey based on the updated configuration.
- Rekey is generated after every 1GB of data (incoming + outgoing) of SSH session, when X509v3 flag is set to Strict in the **secCryptoCfg** template.
- If both data and time are configured, rekey happens when either condition is first met, and on rekey, both the timer and the byte counts are reset.

Firmware integrity check

The `firmwarecheck` CLI has been enhanced to support enabling or disabling of firmware integrity check at boot up, in addition to existing functionality of running on-demand firmware integrity check. The CLI format is shown below:

```
firmwarecheck --<enable/disable> -boot
              --show
```

The enable and disable options apply to running the firmware check at boot. The show option displays the current status and the run option causes a firmware integrity check on-demand.

Creating a complete chain of CA certificates

Fabric OS sends the chain of untrusted (intermediary) CA certificates along with the identity certificate whenever there is request for client certificate for client authentication. To comply with this, the complete chain of certificates that is used to sign the identity certificate must be imported using the `seccertmgmt` command with the `-ca` option. Fabric OS validates the chain of certificates when received, as part of server Certificate message. Import the complete hierarchy of certificates by concatenating the certificates. For example, follow this signing hierarchy:

1. `root.pem` is a public key certificate that identifies root certificate authority (CA). Root certificates are self-signed.
2. `Intermediate Level1-ca.pem` is the first level certificate.

NOTE

Intermediate certificate is a sub-ordinate certificates issued by trusted root CA, specifically to issue identity certificates.

3. `Intermediate Level2-ca.pem` is the second level certificate.
4. `Intermediate Level3-ca.pem` is third level certificate.
5. Identity certificate is a client or server certificate used to identify the client or server and used to authenticate the client to server or server to client.

In the above chain of certificates, if the intermediate certificate `Level3-ca.pem` is used to sign the switch or server certificate, chain certificate can be created as below on a Linux host machine:

```
cat Level3-ca.pem Level2-ca.pem Level1-ca.pem root.pem > chain-ca.pem
```

Import `chain-ca.pem` to the switch as the CA certificate.

Upgrade and downgrade considerations

The behavior of certain Fabric OS 8.1.0b features is different than in previous versions. You must pay special consideration to these differences when upgrading to or downgrading from Fabric OS 8.1.0b.

Firmware upgrade

The following considerations are for upgrading to Fabric OS 8.1.0b.

Feature	From Fabric OS version	Details	Behavior
OCSP	Pre-8.1.0b	OCSP URI verification is applied to in-band certificates received from the server. Ex: LDAP.	No impact if the certificate does not have OCSP URI. If the certificate has OCSP URI, and the OCSP responder is not reachable, then connection will fail.
Audit Logs	Pre-8.1.0	Several Audit Logs	No impact.

Feature	From Fabric OS version	Details	Behavior
	Pre-8.1.0b for SYSLOG		
TLS Client Authentication	Pre-8.1.0 Pre-8.1.0b for SYSLOG	Applicable. Client authentication should be seamless.	NOTE See Below for Validity and checks for certificate.
Certificate validation	Pre-8.1.0b	<ol style="list-style-type: none"> 1. Certificates should be imported with all valid attributes as defined for CC. 2. Refer to table for details. 	secCryptocfg uses an attribute to validate either as basic or strict. For CC, strict is the used value. [X509v3] Validation: Strict/Basic.

Firmware downgrade

The following considerations are for downgrading from Fabric OS 8.1.0b.

Feature	From Fabric OS version	Details	Behavior
OCSP	Pre-8.1.0b	OCSP URI verification is applied to in-band certificates received from the server. Ex: LDAP.	No impact if the certificate does not have OCSP URI.
Audit Logs		Several Audit Logs	No impact.
TLS Client Authentication	Pre-8.1.0 Pre-8.1.0b for SYSLOG	Applicable. Client authentication should be seamless.	NOTE See Below for Validity and checks for certificate.
Certificate validation	Pre-8.1.0b	<ol style="list-style-type: none"> 1. Certificates should be imported with all valid attributes as defined for CC. 2. Refer to table for details. 	Downgrade script checks for new attributes: protocol, log, and validation in user-defined templates. If any exist, downgrade is blocked. To downgrade, modify the template to remove the Fabric OS 8.1.0b attributes that are not supported and try again.

Cryptographic compliance in FIPS mode

To ensure that the cryptographic modules on switches configured for FIPS always use FIPS-approved cryptography, you must ensure that only the following supported algorithms are configured; otherwise, the module is in a non-compliant state if it is configure to any other algorithm or protocol.

```
[Ver] 0.1
[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
[AAA]
RAD_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
LDAP_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
[HTTPS]
Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
```

Configuring the ciphers, KEX, and MAC algorithms

Starting with Fabric OS 7.4.0, you can configure the ciphers, key exchange (KEX), and message authentication code (MAC) algorithms dictated by your security policies.

1. To configure the ciphers, 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, 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
```

3. To display the configured algorithm, use the following command.

```
secCryptoCfg --show
```

For crypto configuration, you are recommended to use *seccryptoCfg* templates than manually using the CLI command.

Secure Sockets Layer protocol

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 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 and Java support

Fabric OS supports the following web browsers for SSL connections:

- Internet Explorer v7.0 or later (Microsoft Windows)
- Mozilla Firefox v2.0 or later (Solaris and Red Hat Linux)

NOTE

Review the release notes for the latest information and to verify if your platform and browser are supported.

In countries that allow the use of 128-bit encryption, you should use the latest version of your browser. For example, Internet Explorer 7.0 and later supports 128-bit encryption by default. You can display the encryption support (called "cipher strength") using the Internet

Explorer **Help** > **About** menu option. If you are running an earlier version of Internet Explorer, you may be able to download an encryption patch from the Microsoft website at <http://www.microsoft.com>.

You should upgrade to the Java 1.6.0 plug-in on your management workstation. To find the Java version that is currently running, open the Java console and look at the first line of the window. For more details on levels of browser and Java support, refer to the *Brocade Web Tools Administration Guide*.

Enabling TLS 1.2 for the entire Web Tools session requires you to also enable TLS 1.2 in the Java Control Panel.

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, you must install a certificate in the Java plug-in on the management workstation, and you might need to add a CA certificate to your web browser.

Configuring for SSL involves these main steps, which are shown in detail in the next sections.

1. Choose a certificate authority (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 e–mail (public) or gives access to them on a remote host (private).

5. On each switch, install both the CA and switch certificates. Once the certificate is loaded on the switch, HTTPS starts automatically.
6. If necessary, install the root certificate to the browser on the management workstation.
7. Add the root certificate to the Java plug-in keystore 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 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 generate CSR or certificate, import certificate, export CSR or certificate, delete CSR or certificates, and display the 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 can be used to 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

1. The **-keypair** option works only with **-extn** type.
2. The **sha384** option for **-hash** is supported only with **-extn** type.
3. The **-type ecdsa** is supported only with **-extn** type.

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|mgmtip>
  -ca -client|-server <fcap|commoncert|https|radius|ldap|syslog|extn>
  -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

1. The **-keypair** option works only with **-extn** type.
2. The **-keypair** option is mandatory for **-cert extn** and **-ca extn** options.
3. The **-keypair** option must not be present for **-cert extn**, **-ca extn local**, and **-ca extn** options.
4. The **-certname <certificate_name>** must not contain more than one period (dot) before the extension of the filename.

Exporting the certificates

The following options can be used to 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

1. The **-keypair** option works only with **-extn** type.
2. The **-keypair** option is mandatory for **-cert extn** and **-ca extn** options.
3. The **-keypair** option should not be present for **-cert extn**, **-ca extn local**, and **-ca extn** options.

Deleting the certificates

The following options can be used to delete the certificates:

```
switch:admin> seccertmgmt delete
  -cert <fcap|commoncert|https|radius|ldap|syslog|extn
    (<certificate name> | -keypair_tag <keypair_tag>)|mgmtip <certificate name>|all>
  -ca -client|-server <fcap|commoncert|https|radius|ldap|syslog|extn <certificate name> |all>
  -csr <fcap|commoncert|https|radius|ldap|syslog|extn -keypair_tag <keypair_tag>>
  -all <default|fcap|commoncert|https|radius|ldap|syslog|extn|mgmtip>
  [-f] - no warning
```

NOTE

1. The **-keypair** option works only with **-extn** type.
2. The **-keypair** option is mandatory for **-cert extn** and **-ca extn** options.
3. The **-keypair** option should not be present for **-cert extn**, **-ca extn local**, and **-ca extn** options.
4. The **-cert <certificate_name>** must not contain more than one period (dot) before the extension of the filename.

Example:

```
switch:admin> seccertmgmt delete -cert https
WARNING!!!

About to delete https switch certificate file(s)

Continue (yes, y, no, n): [no] y
```

Displaying the certificates

The following options can be used to display the certificates:

```
switch:admin> seccertmgmt show
name>> -cert <fcap|commoncert|https|radius|ldap|syslog|extn <certificate name>|mgmtip <certificate
name>>
        -ca -client|-server <fcap|commoncert|https|radius|ldap|syslog|extn <certificate name>>
        -csr <fcap|commoncert|https|radius|ldap|syslog|extn <csr name>>
        -keypair_tag <keypair_name>
        -hexdump - Dump the hexadecimal output
        -all
```

NOTE

1. The **-keypair** option works only with **-extn** type.
2. The **-keypair** option is mandatory for **-cert extn** and **-ca extn** options.
3. The **-keypair** option should not be present for **-cert extn remote** and **-ca extn** options.
4. The **-cert <certificate_name>** must not contain more than one period (dot) before the extension of the filename.

Example:

```
switch:admin> seccertmgmt show -all

ssh private key:
  Does not Exist

ssh public keys available for users:
  None

Certificate Files:
-----
Protocol  Client CA          Server CA          SW                CSR                PVT Key
Passphrase
-----
FCAP      Empty              NA                 Empty              Empty              Empty
Empty
RADIUS    Empty              Exist              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
```

The **-keypair_tag** option, the certificate names with **-cacert**, **-certname**, and **-csr** options supports the following characters:

TABLE 43 Supported Characters on Keypair Tags and Certificate Name

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).

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.

```
switch:admin> seccertmgmt generate -csr https
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):Brocade
Organizational Unit Name (eg, department name):FOS
Common Name (Fully qualified Domain Name, or IP address):10.17.37.3
Generating CSR, file name is: 10.37.209.110.csr
Done.
```

Your CA might require specific codes for Country, State or Province, Locality, Organization, and Organizational Unit 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 fully qualified domain name is set on the domain name switch or director. The IP address or FQDN is the switch where the certificate gets installed.

You can also specify the hash type using the **-hash <sha1 | sha256>** option while generating the CSR. For example,

```
switch:admin> seccertmgmt generate -csr https -type rsa -keysize 2048 -hash sha256 -years 1 -f
```

4. Enter **seccertmgmt export** to store the CSR.
5. Enter the requested information. You can use either FTP or SCP.

Select Secure Copy Protocol (SCP) or File Transfer Protocol (FTP) 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 certificates

Once 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 and HTTPS support the following types of files from the Certificate Authority(CA):

- .cer (binary)
- .crt (binary)
- .pem (text)

Typically, the CA provides the certificate files listed in the following table.

TABLE 44 SSL certificate files

Certificate file	Description
<i>name</i> .pem	The switch certificate.
<i>name</i> Root.pem	The root certificate. Typically, this certificate is already installed in the browser, but if not, you must install it.
<i>name</i> CA.pem	The CA root certificate. It must be installed in the browser to verify the validity of the server certificate or server validation fails.

NOTE

You must perform this procedure for each switch.

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](#) on page 212.
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 https** 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 e-mail, 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 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 happens, 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 https** command to import the CA certificate.

For information on CA chain of certificates, refer to [Creating a complete chain of CA certificates](#) on page 205 .

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 https** 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.
6. 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 10.38.41.17 -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 10.38.41.17 -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 enable the HTTPS service. Also, using MAPS, you can configure the time period for notification when certificates expire.

NOTE

Using this feature to generate self-signed certificates for FCAP is not supported.

1. Log in to the device
2. Run the **seccertmgmt generate -cert https** 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.
```

Creating a complete chain of CA certificates

Fabric OS sends the chain of untrusted (intermediary) CA certificates along with the identity certificate whenever there is request for client certificate for client authentication. To comply with this, the complete chain of certificates that is used to sign the identity certificate must be imported using the **seccertmgmt** command with the **-ca** option. Fabric OS validates the chain of certificates when received, as part of server Certificate message. Import the complete hierarchy of certificates by concatenating the certificates. For example, follow this signing hierarchy:

1. `root.pem` is a public key certificate that identifies root certificate authority (CA). Root certificates are self-signed.
2. Intermediate `Level1-ca.pem` is the first level certificate.

NOTE

Intermediate certificate is a sub-ordinate certificates issued by trusted root CA, specifically to issue identity certificates.

3. Intermediate `Level2-ca.pem` is the second level certificate.
4. Intermediate `Level3-ca.pem` is third level certificate.
5. Identity certificate is a client or server certificate used to identify the client or server and used to authenticate the client to server or server to client.

In the above chain of certificates, if the intermediate certificate `Level3-ca.pem` is used to sign the switch or server certificate, chain certificate can be created as below on a Linux host machine:

```
cat Level3-ca.pem Level2-ca.pem Level1-ca.pem root.pem > chain-ca.pem
```

Import `chain-ca.pem` to the switch as the CA certificate.

The browser

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 to Internet Explorer and Mozilla Firefox browsers. For more detailed instructions, refer to the documentation that came with the certificate.

Checking and installing root certificates on Internet Explorer

Use the following procedure to check and install a root security certificate on a switch using IE:

1. Select **Tools** > **Internet Options**.
2. Click the **Content** tab.
3. Click **Certificates**.
4. Click the **Intermediate** or **Trusted Root** tab and scroll the list to see if the root certificate is listed. 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. Follow the instructions in the Certificate Import wizard to import 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 `nameRoot.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 `nameRoot.crt`.

- Click **Open** and follow the instructions to import the certificate.

Root certificates for the Java plugin

For information on Java requirements, refer to [Browser and Java support](#) on page 207.

This procedure is a guide for installing a root certificate to the Java plugin on the management workstation. If the root certificate is not already installed to the plugin, you should install it. For more detailed instructions, refer to the documentation that came with the certificate and to the Oracle Corporation website (www.oracle.com).

Installing a root certificate to the Java plugin

Use the following procedure to install a root certificate to the Java plugin.

- Copy the root certificate file from its location on the FTP server to the Java plugin bin directory. For example, the bin location may be:

```
C: \program files\java\j2re1.6.0\bin
```

- Open a Command Prompt window and change the directory to the Java plugin bin directory.
- Enter the **keyTool** command and respond to the prompts.

Example of installing a root certificate

```
C:\Program Files\Java\j2re1.6.0\bin> keytool -import -alias RootCert -file RootCert.crt -keystore ..\lib
\security\RootCerts

Enter keystore password: changeit

Owner: CN=Brocade, OU=Software, O=Brocade Communications, L=San Jose, ST=California, C=US
Issuer: CN=Brocade, OU=Software, O=Brocade Communications, L=San Jose, ST=California, C=US
Serial number: 0
Valid from: Thu Jan 15 16:27:03 PST 2007 until: Sat Feb 14 16:27:03 PST 2007
Certificate fingerprints:
    MD5: 71:E9:27:44:01:30:48:CC:09:4D:11:80:9D:DE:A5:E3
    SHA1: 06:46:C5:A5:C8:6C:93:9C:FE:6A:C0:EC:66:E9:51:C2:DB:E6:4F:A1
Trust this certificate? [no]: yes

Certificate was added to keystore
```

In the example, **changeit** is the default password and **RootCert** is an example root certificate name.

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 a 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 these 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, which 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 can be used by a network management system 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

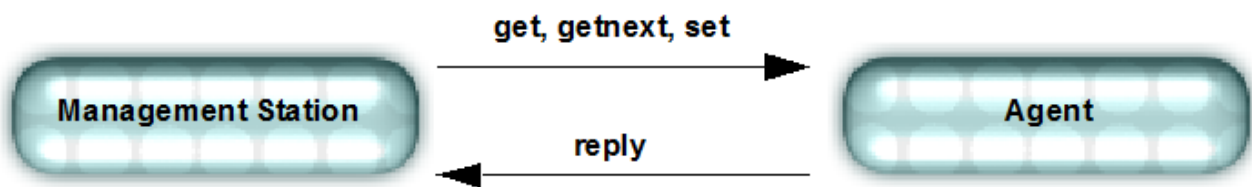
Every Brocade device carries an agent and management information base (MIB). The agent accesses information about a device and makes it available to an SNMP network management station.

FIGURE 7 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 once 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 8 SNMP query



The management station can also receive *traps*, unsolicited messages from the switch agent if an unusual event occurs.

FIGURE 9 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 CLI

For information about Fabric OS commands for configuring SNMP, refer to the Brocade Fabric OS Command Reference.

Configuring 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 that do not have authentication and privacy privileges (noAuthNoPriv).

```
switch:admin> snmpconfig --set secLevel -snmpget 0 -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 45 Security level options

Security level	Protocol	Query behavior	Traps
No security [0] (noAuthNoPriv)	SNMPv1	Allowed.	Sent.
	SNMPv3	Allowed.	Sent.
Authentication only [1] (authNoPriv)	SNMPv1	Allowed.	Sent.
	SNMPv3	All SNMPv3 users allowed except noAuthNoPriv users.	Sent for all SNMPv3 users except noAuthNoPriv users.
Authentication and Privacy [2]	SNMPv1	Not allowed.	Not Sent.

TABLE 45 Security level options (continued)

Security level	Protocol	Query behavior	Traps
(authPriv)	SNMPv3	Only SNMPv3 users with authPriv privilege are allowed.	Sent only for authPriv users.
No Access [3]	SNMPv1 SNMPv3	Not allowed.	Not Sent.

Supported protocol configurations for SNMPv3 users

The following table shows the authentication and privacy protocols supported for configuring SNMPv3 users.

TABLE 46 Supported protocol options

Protocol	Options
Auth protocols	MD5 SHA noAuth
Priv protocols	DES noPriv AES128 AES256

Configuring SNMPv3 user/traps

The following examples list how to configure SNMPv3 users/traps.

NOTE

If you use SNMPGET to poll the switches, for example HiTrack, the SNMPv3 user needs to be configured based on the logical switch that is being monitored. If SNMP monitoring is for the default switch, then you do not need the snmpuser to be configured as a user on the switch. Whereas, when SNMP monitoring is VF specific, then you need to create a user with permission for the corresponding VF and add the created user into SNMP database as SNMPv3 user. Once SNMPv3 user is successfully added, then you can use the same user name in management applications for SNMP queries to the corresponding logical switch.

1. Create a user on the switch in non-VF Context using the **userconfig** command with the required role.

```
switch:admin> userconfig --add fa_adm -r fabricadmin -h0 -a 0-255
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 VF Context using the **userconfig** command with the required role.

```
switch:admin> userconfig --add sa_user -r switchadmin -l 1-128 -h1 -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 --set snmpv3` to create the SNMPv3 user.

```

switch:admin> snmpconfig --set snmpv3

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]
User (rw): [passpass]
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) [4]
New Priv Passwd:
Engine ID: [80:00:05:23:01:ac:1a:1a:ac]
User (rw): [snmpadmin3] password1
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 1
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (1..4) [2] 3
New Priv Passwd:
Verify Priv Passwd:
Engine ID: [00:00:00:00:00:00:00:00:00] 80:00:05:23:01:ac:1a:1a:ac]
User (ro): [snmpuser1]
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: [00:00:00:00:00:00:00:00:00]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [2]
New Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (1..4) [3]
New Priv Passwd:
Engine ID: [00:00:00:00:00:00:00:00:00]
User (ro): [admin]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [80:00:06:34:03]

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]

Notify Type [TRAP(1)/INFORM(2)]: (1..2) [2] 1
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [5]
Trap recipient Port : (0..65535) [1000]
Trap Recipient's IP address : [172.26.26.172]

Notify Type [TRAP(1)/INFORM(2)]: (1..2) [1]
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [4]

SNMPv3 trap recipient configuration:
Trap Recipient's IP address : [0.0.0.0] 10.35.52.33
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [0] 5
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [0.0.0.0] 10.35.52.27
UserIndex: (1..6) [2]

```

```

Trap recipient Severity level : (0..5) [0] 5
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [0.0.0.0]
Trap Recipient's IP address : [0.0.0.0]
Trap Recipient's IP address : [0.0.0.0]
Trap Recipient's IP address : [0.0.0.0]
Committing configuration.....done.

switch:admin> snmpconfig --show snmpv3

SNMP Informs = 1 (ON)

SNMPv3 USM configuration:
User 1 (rw): password
    Auth Protocol: MD5
    Priv Protocol: AES128
    Engine ID: 80:00:05:23:01:ac:1a:1a:ac
User 2 (rw): passpass
    Auth Protocol: MD5
    Priv Protocol: AES256
    Engine ID: 80:00:05:23:01:ac:1a:1a:ac
User 3 (rw): password1
    Auth Protocol: MD5
    Priv Protocol: AES128
    Engine ID: 80:00:05:23:01:ac:1a:1a:ac
User 4 (ro): snmpuser1
    Auth Protocol: MD5
    Priv Protocol: AES128
    Engine ID: 00:00:00:00:00:00:00:00:00
User 5 (ro): snmpuser2
    Auth Protocol: SHA
    Priv Protocol: AES128
    Engine ID: 00:00:00:00:00:00:00:00:00
User 6 (ro): admin
    Auth Protocol: noAuth
    Priv Protocol: noPriv
    Engine ID: 80:00:06:34:03

```

```

SNMPv3 Trap/Informs configuration:
Trap Entry 1:      172.26.26.172
    Trap Port: 162
    Trap User: password
    Trap recipient Severity level: 4
    Notify Type: INFORM(2)
Trap Entry 2:      172.26.26.172
    Trap Port: 1000
    Trap User: password
    Trap recipient Severity level: 5
    Notify Type: TRAP(1)
Trap Entry 3:      172.26.26.172
    Trap Port: 1000
    Trap User: password
    Trap recipient Severity level: 4
    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)

```

An example of the SNMPv3 user trap recipients configured with DNS names and IPv6 addresses:

```

switch:admin> snmpconfig --set snmpv3

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]
User (rw): [passpass]
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) [4]
New Priv Passwd:
Engine ID: [80:00:05:23:01:ac:1a:1a:ac]
User (rw): [password1]
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]
User (ro): [snmpuser1]
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: [00:00:00:00:00:00:00:00:00]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [2]
New Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (1..4) [3]
New Priv Passwd:
Engine ID: [00:00:00:00:00:00:00:00:00]
User (ro): [admin]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [80:00:06:34:03]

SNMPv3 trap/inform recipient configuration:
Trap Recipient's IP address : [172.26.26.172] fe80::224:1dff:fef6:3f98

Notify Type [TRAP(1)/INFORM(2)]: (1..2) [2]
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [172.26.26.172] HCL0389U.corp.brocade.com

Notify Type [TRAP(1)/INFORM(2)]: (1..2) [1]
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [5]
Trap recipient Port : (0..65535) [1000]
Trap Recipient's IP address : [172.26.26.172]

Notify Type [TRAP(1)/INFORM(2)]: (1..2) [1]
Committing configuration.....done.

switch:admin> snmpconfig --show snmpv3

SNMP Informs = 1 (ON)

SNMPv3 USM configuration:
User 1 (rw): password
    Auth Protocol: MD5
    Priv Protocol: AES128
    Engine ID: 80:00:05:23:01:ac:1a:1a:ac
User 2 (rw): passpass
    Auth Protocol: MD5
    Priv Protocol: AES256
    Engine ID: 80:00:05:23:01:ac:1a:1a:ac
User 3 (rw): password1
    Auth Protocol: MD5
    Priv Protocol: AES128
    Engine ID: 80:00:05:23:01:ac:1a:1a:ac
User 4 (ro): snmpuser1
    Auth Protocol: MD5
    Priv Protocol: AES128
    Engine ID: 00:00:00:00:00:00:00:00:00

```

```

User 5 (ro): snmpuser2
    Auth Protocol: SHA
    Priv Protocol: AES128
    Engine ID: 00:00:00:00:00:00:00:00
User 6 (ro): admin
    Auth Protocol: noAuth
    Priv Protocol: noPriv
    Engine ID: 80:00:06:34:03

```

SNMPv3 Trap/Informs configuration:

```

Trap Entry 1: fe80::224:1dff:fef6:3f98
    Trap Port: 162
    Trap User: password
    Trap recipient Severity level: 4
    Notify Type: INFORM(2)
Trap Entry 2: HCL0389U.corp.brocade.com
    Trap Port: 1000
    Trap User: password
    Trap recipient Severity level: 5
    Notify Type: TRAP(1)
Trap Entry 3: 172.26.26.172
    Trap Port: 1000
    Trap User: password
    Trap recipient Severity level: 4
    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:

```

switch:admin> snmptraps --show

# |Mib Name          |Supported Traps
---|-----|-----
001|SW-MIB            |sw-fc-port-scen
    |                  |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|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 all traps to the recipient 10.35.52.33:

```
switch:admin> snmpTraps --send -ip_address 10.35.52.33
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 send the sw-fc-port-scn trap to the recipient 10.35.52.33:

```
switch:admin> snmpTraps --send -trap_name sw-fc-port-scn -ip_address 10.35.52.33
Number of traps sent : 1
```

To unblock port traps on all the ports or on a specific port:

```
switch:admin> snmptraps --unblock -ports 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] 10.32.148.0
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] 10.33.0.0
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 Configure the access permission for **snmpconfig --set snmpv3-groupname ro** .

```
switch:admin> snmpconfig --set snmpv3 -groupname ro
Committing configuration....done.
```

To Configure the access permission for **snmpconfig --set snmpv3-groupname default**.

```
switch:admin> snmpconfig --set snmpv3 -groupname default
Committing configuration....done.
```

NOTE

RO - All entries have read-only permission.

Default - Default configuration (First three entries are RW (read-write) and other three are RO (read-only)).

The following examples display the **mibCapability** configuration.

To enable the `swEventTrap` of the SW-MIB category only (this operation disables all other SNMP traps in this MIB category):

```
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:

```

switch: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
BD-TRAP: YES
                bdTrap: YES
                bdClearTrap: YES
MAPS-TRAP: YES
                mapsTrapAM: 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)
```

- 3.
- 4.
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
```

To set the security level to 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 audit interval.

```
switch:admin> snmpconfig --set auditinterval -interval 31
Committing configuration....done.
switch:admin> snmpconfig --show auditInterval
SNMP Audit Interval (in min): 31
```

Set to 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 Authentication protocol as noAuth, Privacy protocol as noPriv, set the password and set the trap port as 162. (Same values are set as in the switch SNMPv3 configuration.)

Notes:

- SNMPv3 supports AES-128, AES-256, and DES protocols.
- For resolving AES-256 protocol in the USM MIB walk, the eso Consortium MIB has to 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 you will use depends on the mode you are operating in.
 - 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)?"
```

- 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 is used to identify whether the passwords should be encrypted or not for each user.

Changing encrypted SNMPv3 user passwords to unencrypted form

Once 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 can then specify new values for the passwords that will not be encrypted.

NOTE

The SNMPv3 user password decryption option is only available 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 you will use depends on the mode you are operating in.
 - 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. Refer to [Local account passwords](#) on page 150 for instructions on this task.

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, refer to [IP Filter policy](#) on page 255.

ATTENTION

Before blocking Telnet, make sure 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 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. Refer to [Deleting a rule from an IP Filter policy](#) on page 260 for more information on deleting IP filter rules.

6. Save the new IP filter policy by entering the **ipFilter --save** command.
7. Verify 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 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

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.

Refer to [Deleting a rule from an IP Filter policy](#) on page 260 for more information on deleting IP filter rules.

3. To permanently delete the policy, type the **ipfilter --save** command.

ATTENTION

If you deleted the rule to permit 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 47 Blocked listener applications

Listener application	Brocade DCX and Brocade DCX 8510 Backbone 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.

NOTE

If you are using the FC-FC Routing Service, be aware that the **secmodeenable** command is not supported.

TABLE 48 Access defaults

	Access default
Hosts	<ul style="list-style-type: none"> Any host can access the fabric by SNMP. Any host can Telnet to any switch in the fabric. Any host can establish an HTTP connection to any switch in the fabric. Any host can establish an API connection to any switch in the fabric.
Devices	<ul style="list-style-type: none"> All devices can access the management server. Any device can connect to any FC port in the fabric.
Switch access	<ul style="list-style-type: none"> Any switch can join the fabric. All switches in the fabric can be accessed through a serial port.
Zoning	No zoning is enabled.

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 49 Port information

Port	Type	Common use	Comment
22	TCP	SSH, SCP	
23	TCP	Telnet	Use the ipfilter command to block the port.
80	TCP	HTTP	Use the ipfilter command to block the port.

TABLE 49 Port information (continued)

Port	Type	Common use	Comment
123	UDP	NTP	
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 ipfilter command to block the port.
52357	UDP	EZ Server	Use the ipfilter command to block the port.

Configuring Security Policies

• ACL policies overview.....	233
• ACL policy management.....	234
• FCS policies.....	236
• Device Connection Control policies.....	240
• SCC Policies.....	244
• Authentication policy for fabric elements.....	245
• IP Filter policy.....	255
• Policy database distribution.....	261
• Management Interface Security.....	267

ACL policies overview

Each supported Access Control List (ACL) policy 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 (FCS) policy:** Used to restrict which switches can change the configuration of the fabric.
- **Device connection control (DCC) policies:** Used to restrict which Fibre Channel device ports can connect to which Fibre Channel switch ports.
- **Switch connection control (SCC) policy:** Used to restrict which switches can join with a switch.

How the 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 security policy database size is set to 1Mb. The policies are grouped by state and type. A policy can be in either of the following states:

- Active, which means the policy is being enforced by the switch.
- Defined, which means the policy has been set up but is not enforced.

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 1Mb 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 in [Table 50](#).

TABLE 50 Valid methods for specifying policy members

Policy name	Device port WWN 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. It is recommended you make changes from one switch only to prevent multiple transactions from occurring. Each logical switch has its own access control list.

The FCS, SCC, and DCC policies in Secure Fabric OS are not interchangeable with Fabric OS FCS, SCC and DCC policies. Uploading and saving a copy of the Fabric OS configuration after creating policies is strongly recommended. For more information on configuration uploads, refer to [Maintaining the Switch Configuration File](#) on page 133.

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. Refer to [Policy database distribution](#) on page 261 for more information on the database settings and 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 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 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 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 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 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 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.

Example of removing a member

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 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

Fabric configuration server (FCS) 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 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 51 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 normally cannot modify the policy. However, if the Primary FCS switch in the policy list is not reachable, then a backup FCS switch is allowed to modify the policy.

Once 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.

[Table 52](#) shows the commands for switch operations for Primary FCS enforcement.

TABLE 52 FCS switch operations

Allowed on FCS switches	Allowed on all switches
secPolicyAdd (Allowed on all switches for SCC and DCC policies as long as it is not fabric-wide)	secPolicyShow
secPolicyCreate (Allowed on all switches for SCC and DCC policies as long as it is not fabric-wide)	fddCfg --localaccept or fddCfg --localreject
secPolicyDelete (Allowed on all switches for SCC and DCC policies as long as its not fabric-wide)	userconfig, Passwd, Passwdcfg (Fabric-wide distribution is not allowed from a backup or non-FCS switch.)
secPolicyRemove (Allowed on all switches for SCC and DCC policies as long as its not fabric-wide)	secPolicyActivate
fddCfg -- fabwideset	secPolicySave
Any fabric-wide commands	secPolicyAbort
All zoning commands except the show commands	SNMP commands
All AD commands	configupload
	Any local-switch commands

In Fabric OS v7.1.0 and later, to avoid segmentation of ports due to a member-list order mismatch, security policy members are sorted based on WWN. By default, DCC and SCC policy members are sorted based on WWN. Switches running earlier Fabric OS versions will have the member list in the unsorted manner. Any older-version switch with a policy already created in unsorted order will have port segmentation due to order mismatch when attempting to join any switch with Fabric OS v7.1.0 or later. To overcome the order mismatch, you can modify the member list in the switch by using the **-legacy** option in the **secPolicyAdd** and **secPolicyCreate** commands.

Ensuring the Fabric Domains Share Same Policy

To create new FCS policies or manage your current FCS policies, you must follow certain steps to ensure 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 prevents the admin from creating a second FCS master in the fabric.

Creating an FCS 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 **secPolicyCreate "FCS_POLICY"** 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. Once the policy has been activated you can distribute the policy.

NOTE

FCS policy must be consistent across the fabric. If the policy is inconsistent in the fabric, then 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 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](#) on page 262.

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. Since this policy is distributed manually, the command `fddCfg --fabwideset` is used to distribute a fabric-wide consistency policy for FCS policy in an environment consisting of only Fabric OS v6.2.0 and later switches.

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. This is to support 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 is allowed to 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](#) on page 263.

Local switch configuration parameters are needed to control whether a switch accepts or rejects distributions of FCS policy and whether the switch is allowed to 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 53 Distribution policy states

Fabric OS	State
v6.2.0 and later configured to accept	Target switch accepts distribution and fabric state change occurs.
v6.2.0 and later 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 can be used to restrict which device ports can connect to which switch ports. The devices can be initiators, targets, or intermediate devices such as SCSI routers and loop hubs. By default, all device ports are allowed to 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 Access Gateway Administration 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 any DCC policies are allowed to 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.

[Table 54](#) shows the possible DCC policy states.

TABLE 54 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, that device is only allowed 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 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. It is therefore recommended to create policies with the physical PWWN

DCC policies created with the lock down feature result in DCC policies with FA-PWWNs. It is therefore recommended to avoid using the lock down 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 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.

[Table 55](#) lists the behavior of the DCC policy with FA-PWWNs in the fabric when the DCC policy is created using lockdown support.

TABLE 55 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"> FA-PWWN has logged into the switch. DCC policy creation with lockdown (uses FA-PWWN). DCC policy activation. 	FA-PWWN	Traffic is not disrupted. ²	Ports are disabled for security violation. ³
<ul style="list-style-type: none"> DCC policy creation with lockdown (uses physical PWWN). FA-PWWN has logged into the switch. DCC policy activation. 	Physical PWWN	Traffic is not disrupted.	Ports come up without security issues.
<ul style="list-style-type: none"> DCC policy creation with lockdown (uses physical PWWN). DCC policy activation. FA-PWWN has logged into the switch. 	Physical PWWN	Traffic is not disrupted.	Ports come up without security issues.

[Table 56](#) 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.

TABLE 56 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"> FA-PWWN has logged into the switch. DCC policy creation manually with physical PWWN of device. DCC policy activation. 	PWWN	Traffic is not disrupted.	Ports come up without security issues.
<ul style="list-style-type: none"> DCC policy creation manually with physical PWWN. FA-PWWN has logged into the switch. DCC policy activation. 	PWWN	Traffic is not disrupted.	Ports come up without security issues.
<ul style="list-style-type: none"> DCC policy creation manually with physical PWWN. DCC policy activation. 	Physical PWWN	Traffic is not disrupted.	Ports come up without security issues.

² 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. This is done to avoid disturbing any existing management.

³ 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 56 DCC policy behavior when created manually with PWWN (continued)

Configuration	WWN seen on DCC policy list	Behavior when DCC policy activates	Behavior on portDisable and portEnable
<ul style="list-style-type: none"> FA-PWWN has logged into the switch. 			

NOTE

FCoE devices are excluded from DCC control.

SCC Policies

The switch connection control (SCC) policy is used to restrict which switches can join the fabric. Switches are checked against the policy each time an E_Port-to-E_Port connection is made. The policy is named SCC_POLICY and accepts members listed as WWNs, domain IDs, or switch names. Only one SCC policy can be created.

By default, any switch is allowed to join the fabric; the SCC policy does not exist until it is created. When connecting a Fibre Channel router to a fabric or switch that has an active SCC policy, the front domain of the Fibre Channel router must be included in the SCC policy.

SCC policy states are shown in [Table 57](#).

TABLE 57 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, refer to [Managing Virtual Fabrics](#) on page 382.

Creating an SCC policy

- 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.
- Enter the `secPolicyCreate "SCC_POLICY"` command.

3. Save or activate the new policy by entering either the **secPolicySave** 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 v6.2.0 and later 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 have to 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.

FCAP requires the exchange of certificates between two or more switches to authenticate to each other before they form or join a fabric. Beginning with Fabric OS v7.0.0, these certificates are no longer issued by Brocade, but by a third-party which is now the root CA for all of the issued certificates. You can use Brocade and third-party certificates between switches that are Fabric OS v6.4.0, but only Brocade-issued certificates (where Brocade is the root CA) for Fabric OS versions earlier than v6.4.0. The certificates must be in PEM (Privacy Enhanced Mail) encoded format for both root and peer certificates. The switch certificates issued from the third-party vendors can be directly issued from the root CA or from an intermediate CA authority.

When you configure DH-CHAP authentication, you also must define a *pair of shared secrets* known to both switches as a *secret key pair*. [Figure 10](#) 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, refer to [Setting a secret key pair](#) on page 252.

When configured, the *secret key pair* is used for authentication. Authentication occurs whenever there is a state change 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 10 DH-CHAP authentication



If you use DH-CHAP authentication, then 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 switches with Fabric OS v5.3.0 and later. 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_ and F_Port initialization.

ATTENTION

A *secret key pair* has to be installed prior to changing the policy. For more information on setting up secret key pairs, refer to [Setting a secret key pair](#) on page 252.

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. This 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 needs to 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 [Virtual Fabrics considerations for E_Port authentication](#).

Configuring E_Port 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 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 or 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 command **authUtil** 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; this means you can initiate the authentication even if the switch is in PASSIVE mode. This command is used to restart 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 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 Backbones 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, 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 Access Gateway Administration Guide*.

By default the devicepolicy 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 a FLOGI with the 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 a FLOGI accept with the FC-SP bit set, then 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 point-to-point manner and is visible to the entire fabric.

The following are not supported:

- Public loop devices
- Single private devices
- Private loop devices
- Mixed public and private devices in loop
- NPIV devices
- FICON channels
- Configupload and download will not be 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 Access Gateway Administration 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 DH (Diffie-Hellman) 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 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), then switch authentication will fail. If the E_Port is to carry in-flight encrypted traffic, the authentication protocol must be set to DH-CHAP. 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, a shared secret of at least 32 characters is recommended. 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 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 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.

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 8 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 has to 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.

- 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 e-mail (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 58 Certificate files provided by certificate authority

Certificate File	Description
name CA.pem	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 .pem	The switch certificates:switch certificate.

- On each switch, install the CA certificate before installing the switch certificate.
- After the CA certificate is installed, install the switch certificate on each switch.
- Update the switch database for peer switches to use third-party certificates.
- 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 is used to generate the public/private key the CSR, and the passphrase.

- Log in to the switch 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.
- Enter the **seccertmgmt generate -csr fcap** command on the local switch.
- 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](#) on page 253 section and send to a Certificate Authority (CA). The CA will in turn provide two files as outlined in [FCAP configuration overview](#) on page 252.

- Log in to the switch 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.
- Enter the **seccertmgmt export -cert fcap** command.

Importing CA for FCAP

Once you receive the files back from the Certificate Authority with CA root certificate, you will need to install or import them onto the local and remote switches.

- Log in to the switch 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 with the **-ca -client** option and verify the CA certificates are consistent on both local and remote switches.

NOTE

Firmware downgrade from Fabric OS 7.3.0 to an earlier version is blocked if SHA-256 is one of the configured hash types for DH-CHAP or FCAP in at least one of the logical 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 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 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](#) on page 263 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](#) on page 261.

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 will be 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, the most commonly used rules should be specified 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 alpha, 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 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 --create` 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 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 achieve 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 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 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 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`. An active IP Filter policy cannot be deleted.

1. Log in to the switch 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, HTTPS.
- Protocol: The protocol type. Supported types are TCP or UDP.
- Action: The filtering action taken by this rule, either Permit or Deny.

A traffic type and 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 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 keyword **any** is supported to represent any IPv4 address.

For an IPv6 filter policy, the source address has to be a 128-bit IPv6 address, in a format acceptable in RFC 3513. The group prefix has to 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 keyword **any** 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 59](#) lists the supported service names and their corresponding port numbers.

TABLE 59 Supported services

Service name	Port number
echo	7
discard	
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

TABLE 59 Supported services (continued)

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
kerberos4	750

Protocol

TCP and UDP protocols are valid protocol selections. Fabric OS v6.2.0 and later do not support configuration to filter other protocols. Implicitly, ICMP type 0 and type 8 packets are always allowed to support ICMP echo request and reply on commands like ping and traceroute.

Action

For the action, only "permit" and "deny" are valid.

Traffic type and destination IP

The traffic type and destination IP elements allow an IP policy rule to specify filter enforcement for IP forwarding. The INPUT traffic type is the default and restricts rules to manage traffic on IP management interfaces,

The FORWARD traffic type allows management of bidirectional traffic between the external management interface and the inband management interface. In this case, the destination IP element should also be specified.

Implicit filter rules

For every IP Filter policy, the two rules listed in [Table 60](#) 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 60 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. [Table 61](#) lists the rules of the default IP filter policy.

TABLE 61 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 [Table 62](#). Customer can configure the rules from FOS 8.2.1 using `ipfilter` command.

TABLE 62 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	32768-61000	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, then 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 have to be the NAT server address.

Adding a rule to an IP Filter policy

There can be a maximum of 256 rules created 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 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 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 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.

IP Filter policy distribution

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; there is no support for automatic distribution. To distribute the IPFilter policy, refer to [Distributing the local ACL policies](#) on page 263 for instructions.

You can accept or deny IP Filter policy distribution using the commands `fddCfg --localaccept` or `fddCfg --localreject`. Refer to [Policy database distribution](#) on page 261 for more information about distributing the IP Filter policy.

NOTE

Any RPC ports that were allowed in Fabric OS versions earlier than 7.2.0 are removed and ignored in Fabric OS 7.2.0 and later.

NOTE

To distribute the IP Filter policy in a logical fabric using Virtual Fabrics, use the `chassisDistribute` command.

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 or not 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 .

Virtual Fabric considerations: FCS, DCC, SCC, and AUTH databases can be distributed using the `-distribute` command, but the PWD is blocked from distribution. The IPFILTER databases can be distributed using the FID in 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](#) on page 263.
- **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, then the policies are managed on a per-switch basis. For configuration instructions, refer to [Fabric-wide enforcement](#) on page 263.

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 as that of pre-v6.2.0 releases 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 63 Interaction between fabric-wide consistency policy and distribution settings

Distribution setting	Fabric-wide consistency policy		
Absent (default)	Tolerant	Strict	Invalid configuration. ⁴

TABLE 63 Interaction between fabric-wide consistency policy and distribution settings (continued)

Distribution setting	Fabric-wide consistency policy		
Reject	Database is protected, it cannot be overwritten. May not match other databases in the fabric.	Invalid configuration. ⁴	
Accept (default)	Database is not protected, the database can be overwritten. If the switch initiating a distribute command has a strict or tolerant fabric-wide consistency policy, the fabric-wide policy is also overwritten. May not match other databases in the fabric.	Database is not protected. Automatically distributes activated changes to other v6.2.0 or later switches in the fabric. If the fabric-wide consistency is set as "strict" for a particular policy, then the manual distribution is blocked. May not match other databases in the fabric.	Database is not protected. Automatically distributes activated changes to all switches in the fabric. Fabric can only contain switches running Fabric OS v6.2.0 or later. Active database is the same for all switches in the fabric.

NOTE

Starting with Fabric OS 7.3.0, Access Gateways are capable of receiving the password database distributed by native switches and domains. However, the Access Gateways are not capable of distributing the password database to the switches or domains.

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.

[Table 64](#) lists the databases supported in Fabric OS v6.2.0 and later switches.

TABLE 64 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.

⁴ An error is returned indicating that the distribution setting must be Accept before you can set the fabric-wide consistency policy.

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 be running Fabric OS v6.2.0 or later.
- All target switches must accept the database distribution (refer to [Database distribution settings](#) on page 262).
- The fabric must have a tolerant or no (absent) fabric-wide consistency policy (refer to [Fabric-wide enforcement](#) on page 263).

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 be able 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.

NOTE

Starting with Fabric OS 7.3.0, FC routers can join a fabric with a strict fabric-wide consistency policy. FC routers do support the fabric-wide consistency policies.

The following table describes the fabric-wide consistency settings.

TABLE 65 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 Fabric OS v6.2.0 and later 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, there are no ACL merge checks 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.

[Table 66](#) describes the impact of merging fabrics with the same fabric-wide consistency policy that have SCC, DCC, or both policies.

TABLE 66 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.
	None	SCC/DCC	Succeeds	No ACL policies copied.
Tolerant	None	None	Succeeds	No ACL policies copied.

TABLE 66 Merging fabrics with matching fabric-wide consistency policies (continued)

Fabric-wide consistency policy	Fabric A ACL policies	Fabric B ACL policies	Merge results	Database 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. ⁵
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.

Non-matching fabric-wide consistency policies

You may encounter one of the following two scenarios described in [Table 67](#) and [Table 68](#) where you are merging a fabric with a strict policy to a fabric with an absent, tolerant, or non-matching strict policy and the merge fails and the ports are disabled.

[Table 67](#) shows merges that are not supported.

TABLE 67 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 connecting 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	

[Table 68](#) has a matrix of merging fabrics with tolerant and absent policies.

TABLE 68 Fabric merges with tolerant and absent combinations

Fabric-wide consistency policy setting			Expected behavior
	Fabric A	Fabric B	
Tolerant/Absent	SCC;DCC		Error message logged.
	DCC		Run the fddCfg command with the --fabwideset "policy_ID" from any switch with the desired configuration to fix the conflict. The
	SCC;DCC	SCC	
	DCC	SCC	

⁵ 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.

TABLE 68 Fabric merges with tolerant and absent combinations (continued)

Fabric-wide consistency policy setting	Expected behavior
	secPolicyActivate command is blocked until conflict is resolved.

Management Interface Security

You can secure an Ethernet management interface between two Brocade switches or Backbones by implementing IPsec and IKE policies to create a tunnel that protects traffic flows. While the tunnel must have a Brocade switch or Backbone at each end, there may be routers, gateways, and firewalls in between the two ends.

ATTENTION

Enabling secure IPsec tunnels does not provide IPsec protection for traffic flows on the external management interfaces of intelligent blades in a chassis, nor does it support protection of traffic flows on FCIP interfaces. IPsec on Management interface is not supported in tunnel mode.

Internet Protocol security (IPsec) is a framework of open standards that ensures private and secure communications over Internet Protocol (IP) networks through the use of cryptographic security services. The goal of IPsec is to provide the following capabilities:

NOTE

The **ipSecConfig --enable** command is not supported in Fabric OS 8.0.0 in AG mode and supported in Fabric OS 8.0.1 and later with Gen 5 and Gen 6 devices.

- **Authentication** -- Ensures that the sending and receiving end-users and devices are known and trusted by one another.
- **Data Integrity** -- Confirms that the data received was in fact the data transmitted.
- **Data Confidentiality** -- Protects the user data being transmitted, such as utilizing encryption to avoid sending data in clear text.
- **Replay Protection** -- Prevents replay attack in which an attacker resends previously-intercepted packets in an effort to fraudulently authenticate or otherwise masquerade as a valid user.
- **Automated Key Management** --Automates the process, as well as manages the periodic exchange and generation of new keys.

Using the **ipSecConfig** command, you must configure multiple security policies for traffic flows on the Ethernet management interfaces based on IPv4 or IPv6 addresses, a range of IPv4 or IPv6 addresses, the type of application, port numbers, and protocols used (UDP/TCP/ICMP). You must specify the transforms and processing choices for the traffic flow (drop, protect or bypass). Also, you must select and configure the key management protocol using an automatic or manual key.

For more information on IPv4 and IPv6 addressing, refer to [Performing Basic Configuration Tasks](#) on page 32.

Configuration examples


The following examples demonstrate various configurations you can use to implement an IPsec tunnel between two devices. You can configure other scenarios as nested combinations of these configurations.

Endpoint-to-endpoint transport or tunnel

In this scenario, both endpoints of the IP connection implement IPsec, as required of hosts in RFC4301. Transport mode encrypts only the payload while tunnel mode encrypts the entire packet. A single pair of addresses will be negotiated for packets protected by this SA.

It is possible in this scenario that one or both of the protected endpoints will be behind a network address translation (NAT) node, in which case tunneled packets will have to be UDP-encapsulated so that port numbers in the UDP headers can be used to identify individual endpoints behind the NAT.

FIGURE 11 Protected endpoints configuration




A possible drawback of end-to-end security is that various applications that require the ability to inspect or modify a transient packet will fail when end-to-end confidentiality is employed. Various QoS solutions, traffic shaping, and firewalling applications will be unable to determine what type of packet is being transmitted and will be unable to make the decisions that they are supposed to make.

Gateway-to-gateway tunnel

In this scenario, neither endpoint of the IP connection implements IPsec, but the network nodes between them protect traffic for part of the way. Protection is transparent to the endpoints, and depends on ordinary routing to send packets through the tunnel endpoints for processing. Each endpoint would announce the set of addresses behind it, and packets would be sent in tunnel mode where the inner IP header would contain the IP addresses of the actual endpoints.


FIGURE 12 Gateway tunnel configuration



Endpoint-to-gateway tunnel

In this scenario, a protected endpoint (typically a portable computer) connects back to its corporate network through an IPsec-protected tunnel. It might use this tunnel only to access information on the corporate network, or it might tunnel all of its traffic back through the corporate network in order to take advantage of protection provided by a corporate firewall against Internet-based attacks. In either case, the protected endpoint will want an IP address associated with the security gateway so that packets returned to it will go to the security gateway and be tunneled back.

FIGURE 13 Endpoint-to-gateway tunnel configuration



RoadWarrior configuration

In endpoint-to-endpoint security, packets are encrypted and decrypted by the host which produces or consumes the traffic. In the gateway-to-gateway example, a router on the network encrypts and decrypts the packets on behalf of the hosts on a protected network. A combination of the two is referred to as a RoadWarrior configuration where a host on the Internet requires access to a network through a security gateway that is protecting the network.

IPsec protocols

IPsec ensures confidentiality, integrity, and authentication using the following protocols:

- Authentication Header (AH)
- Encapsulating Security Payload (ESP)

IPsec protocols protect IP datagram integrity using hash message authentication codes (HMAC). Using hash algorithms with the contents of the IP datagram and a secret key, the IPsec protocols generate this HMAC and add it to the protocol header. The receiver must have access to the secret key in order to decode the hash.

IPsec protocols use a sliding window to assist in flow control. The IPsec protocols also use this sliding window to provide protection against replay attacks in which an attacker attempts a denial of service attack by replaying an old sequence of packets. IPsec protocols assign a sequence number to each packet. The recipient accepts each packet only if its sequence number is within the window. It discards older packets.

Security associations

A security association (SA) is the collection of security parameters and authenticated keys that are negotiated between IPsec peers to protect the IP datagram. A security association database (SADB) is used to store these SAs. Information in these SAs--IP addresses, secret keys, algorithms, and so on--is used by peers to encapsulate and decapsulate the IPsec packets

An IPsec security association is a construct that specifies security properties that are recognized by communicating hosts. The properties of the SA are the security protocol (AH or ESP), destination IP address, and Security Parameter Index (SPI) number. SPI is an arbitrary 32-bit value contained in IPsec protocol headers (AH or ESP) and an IPsec SA is unidirectional. Because most communication is peer-to-peer or client-to-server, two SAs must be present to secure traffic in both directions. An SA specifies the IPsec protocol (AH or ESP), the algorithms used for encryption and authentication, and the expiration definitions used in security associations of the traffic. IKE uses

these values in negotiations to create IPsec SAs. You must create an SA prior to creating an SA-proposal. You cannot modify an SA once it is created. Use the `ipSecConfig --flush manual-sa` command to remove all SA entries from the kernel SADB and re-create the SA.

IPsec proposal

The IPsec sa-proposal defines an SA or an SA bundle. An SA is a set of parameters that define how the traffic is protected using IPsec. These are the IPsec protocols to use for an SA, either AH or ESP, and the encryption and authentication algorithms to use to protect the traffic.

For SA bundles, [AH, ESP] is the supported combination.

Authentication and encryption algorithms

IPsec uses different protocols to ensure the authentication, integrity, and confidentiality of the communication. Encapsulating Security Payload (ESP) provides confidentiality, data integrity and data source authentication of IP packets, and protection against replay attacks. Authentication Header (AH) provides data integrity, data source authentication, and protection against replay attacks, but unlike ESP, AH does not provide confidentiality.

In AH and ESP, `hmac_md5` and `hmac_sha1` are used as authentication algorithms. Only in ESP, `3des_cbc`, `blowfish_cbc`, `aes256_cbc` and `null_enc` are used as encryption algorithms. Use [Table 69](#) when configuring the authentication algorithm.

TABLE 69 Algorithms and associated authentication policies

Algorithm	Encryption Level	Policy	Description
<code>hmac_md5</code>	128-bit	AH, ESP	A stronger MAC because it is a keyed hash inside a keyed hash. When MD5 or SHA-1 is used in the calculation of an HMAC; the resulting MAC algorithm is termed HMAC-MD5 or HMAC-SHA-1 accordingly. NOTE The MD5 hash algorithm is blocked when FIPS mode is enabled
<code>hmac_sha1</code>	160-bit	AH, ESP	
<code>3des_cbc</code>	168-bit	ESP	Triple DES is a more secure variant of DES. It uses three different 56-bit keys to encrypt blocks of 64-bit plain text. The algorithm is FIPS-approved for use by Federal agencies.
<code>blowfish_cbc</code>	64-bit	ESP	Blowfish is a 32-bit to 448-bit keyed, symmetric block cipher.
<code>aes128_cbc</code>	128-bit	ESP	Advanced Encryption Standard is a 128- or 256-bit fixed block size cipher.
<code>aes256_cbc</code>	256-bit	ESP	
<code>null_enc</code>	n/a	ESP	A form of plaintext encryption.

IPsec policies

An IPsec policy determines the security services afforded to a packet and the treatment of a packet in the network. An IPsec policy allows classifying IP packets into different traffic flows and specifies the actions or transformations performed on IP packets on each of the traffic flows. The main components of an IPsec policy are: IP packet filter and selector (IP address, protocol, and port information) and transform set.

IPsec traffic selector

The traffic selector is a traffic filter that defines and identifies the traffic flow between two systems that have IPsec protection. IP addresses, the direction of traffic flow (inbound, outbound) and the upper layer protocol are used to define a filter for traffic (IP datagrams) that is protected using IPsec.

IPsec transform

A *transform set* is a combination of IPsec protocols and cryptographic algorithms that are applied on the packet after it is matched to a selector. The transform set specifies the IPsec protocol, IPsec mode and action to be performed on the IP packet. It specifies the key management policy that is needed for the IPsec connection and the encryption and authentication algorithms to be used in security associations when IKE is used as the key management protocol.

IPsec can protect either the entire IP datagram or only the upper-layer protocols using tunnel mode or transport mode. Tunnel mode uses the IPsec protocol to encapsulate the entire IP datagram. Transport mode handles only the IP datagram payload.

IKE policies

When IKE is used as the key management protocol, IKE policy defines the parameters used in IKE negotiations needed to establish IKE SA and parameters used in negotiations to establish IPsec SAs. These include the authentication and encryption algorithms, and the primary authentication method, such as preshared keys, or a certificate-based method, such as RSA signatures.

Key management

The IPsec key management supports Internet Key Exchange or Manual key/SA entry. The Internet Key Exchange (IKE) protocol handles key management automatically. SAs require keying material for authentication and encryption. The managing of keying material that SAs require is called *key management*.

The IKE protocol secures communication by authenticating peers and exchanging keys. It also creates the SAs and stores them in the SADB.

The manual key/SA entry requires the keys to be generated and managed manually. For the selected authentication or encryption algorithms, the correct keys must be generated using a third party utility on your LINUX system. The key length is determined by the algorithm selected.

Linux IPsec-tools 0.7 provides tools for manual key entry (MKE) and automatic keyed connections. The LINUX **setKey** command can be used for manually keyed connections, which means that all parameters needed for the setup of the connection are provided by you. Based on which protocol, algorithm, and key used for the creation of the security associations, the switch populates the security association database (SAD) accordingly.

Pre-shared keys

A pre-shared key has the .psk extension and is one of the available methods IKE can be configured to use for primary authentication. You can specify the pre-shared keys used in IKE policies; add and delete pre-shared keys (in local database) corresponding to the identity of the IKE peer or group of peers.

The **ipseconfig** command does not support manipulating pre-shared keys corresponding to the identity of the IKE peer or group of peers. Use the **seccertmgmt** command to import, delete, or display the pre-shared keys in the local switch database. For more information on this procedure, refer to [Configuring Protocols](#) on page 189.

Security certificates

A certificate is one of the available methods IKE can be configured to use for primary authentication. You can specify the local public key and peer public key (in X.509 format) to be used in a particular IKE policy.

Use the **seccertmgmt import** command to manage third-party certificates on a switch, including Public Key Infrastructure (PKI)-based certificates, Lightweight Directory Access Protocol (LDAP) certificates, FCAP certificates, and syslog CA certificates. This command also imports or exports Certificate Signing Requests (CSRs) from or to a remote host. This command supports IPv4 and IPv6 addresses.

Use the **seccertmgmt import** command to import public key, and peer-public key (in X.509 PEM format) into the switch database. The size of the key can be 1024, 2048, 4096, or 8192 bits. The greater the value, the more secure is the connection; however, performance degrades with size. The keys are generated only after all existing CSRs and certificates have been deleted. For more information on this procedure, refer to [Configuring Protocols](#) on page 189.

Static Security Associations

Manual Key Entry (MKE) provides the ability to manually add, delete and flush SA entries in the SADB. Manual SA entries may not have an associated IPsec policy in the local policy database. Manual SA entries are persistent across system reboots.

Restrictions when using IKE policies

You should be aware of the following restrictions when using IKE policies:

- For Fabric OS 7.0.0 and later versions, IPsec does not support null encryption (null_enc) for IKE policies.
- IPv6 policies cannot tunnel IMCP traffic.

Creating the tunnel

Each side of the tunnel must be configured in order for the tunnel to activate. After you are logged into the switch, do not log off, because each step requires that you be logged in to the switch. IPsec configuration changes take effect upon execution and are persistent across reboots. Use the following procedure to configure each side of the tunnel:

1. Determine the authentication protocol and algorithm to be used on the tunnel.
Refer to [Table 69](#) on page 270 to determine which algorithm to use in conjunction with a specific authentication protocol.
2. Determine the type of keys to be used on the tunnel.
If you are using CA signed keys, you must generate them prior to setting up your tunnels.
3. Enable IPsec.
 - a) Connect to the switch and log in using an account with admin permissions, or an account associated with the chassis role and having OM permissions for the IPsec RBAC class of commands.
 - b) Enter the **ipSecConfig --enable** command to enable IPsec on the switch.

4. Create an IPsec SA policy on each side of the tunnel using the **ipSecConfig --add** command.

Example of creating an IPsec SA policy

This example creates an IPsec SA policy named "AH01", which uses AH protection with MD5. You would run this command on each switch; on each side of the tunnel so that both sides have the same IPsec SA policy.

```
switch:admin> ipseconfig --add policy ips sa -t AH01 -p ah -auth hmac_md5
```

5. Create an IPsec proposal on each side of the tunnel using the **ipSecConfig --add** command.

Example of creating an IPsec proposal

This example creates an IPsec proposal "IPSEC-AH" to use "AH01" as SA.

```
switch:admin> ipseconfig --add policy ips sa-proposal -t IPSEC-AH -sa AH01
```

6. Import the pre-shared key file.

Refer to [Configuring Protocols](#) on page 189 for information on how to set up pre-shared keys and certificates.

7. Configure the IKE policy using the **ipSecConfig --add** command.

Example of creating an IKE policy

This example creates an IKE policy for the remote peer.

```
switch:admin> ipseconfig --add policy ike -t IKE01 -remote 10.33.74.13
-id 10.33.69.132 -remoteid 10.33.74.13 -enc 3des_cbc -hash hmac_md5
-prf hmac_md5 -auth psk -dh modp1024 -psk ipseckey.psk
```

8. Create an IPsec transform on each switch using the **ipSecConfig --add** command.

Example of creating an IPsec transform

This example creates an IPsec transform TRANSFORM01 to use the transport mode to protect traffic identified for IPsec protection and use IKE01 as key management policy.

```
switch:admin> ipseconfig --add policy ips transform -t TRANSFORM01
-mode transport -sa-proposal IPSEC-AH -action protect -ike IKE01
```

9. Create a traffic selector on each switch using the **ipSecConfig --add** command.

Example of creating a traffic selector

This example creates a traffic selector to select outbound and inbound traffic that needs to be protected.

```
switch:admin> ipseconfig --add policy ips selector -t SELECTOR-OUT -d out
-l 10.33.69.132 -r 10.33.74.13 -transform TRANSFORM01
switch:admin> ipseconfig --add policy ips selector -t SELECTOR-IN -d in
-l 10.33.74.13 -r 10.33.69.132 -t transform TRANSFORM01
```

Inbound and outbound selectors use opposite values for local and remote IP addresses. In this example, notice that the local ("-l") address of SELECTOR-OUT is the same as the remote ("-r") address of SELECTOR-IN. Similarly, the local ("-l") address of SELECTOR-IN is the same as the remote ("-r") address of SELECTOR-OUT. That is, "local" refers to the source IP address of the packet, and "remote" is the destination IP address. Hence inbound packets have opposite source and destination addresses than outbound packets.

NOTE

For Directors, the tunnel must be configured with three IP addresses, one for the chassis and two for the IPs to ensure that all the incoming and outgoing traffic is encrypted. The Directors require three tunnels.

10. Verify traffic is protected.

- a) Initiate a telnet, SSH, or ping session from the two switches.
- b) Verify that IP traffic is encapsulated.
- c) Monitor IPsec SAs created using IKE for the traffic flow:
 - Use the `ipSecConfig --show manual-sa -a` command with the operands specified to display the outbound and inbound SAs in kernel SADB.
 - Use the `ipSecConfig --show policy ips sa -a` command with the specified operands to display all IPsec SA policies.
 - Use the `ipSecConfig --show policy ips sa-proposal -a` command with the specified operands to display IPsec proposals.
 - Use the `ipSecConfig --show policy ips transform -a` command with the specified operands to display IPsec transforms.
 - Use the `ipSecConfig --show policy ips selector -a` command with the specified operands to display IPsec traffic selectors.
 - Use the `ipSecConfig --show policy ike -a` command with the specified operands to display IKE policies.
 - Use the `ipSecConfig --flush manual-sa` command with the specified operands to flush the created SAs in the kernel SADB.

Example of an end-to-end transport tunnel mode

This example illustrates securing traffic between two systems using AH protection with MD5 and configure IKE with pre-shared keys. The two systems are a switch (IPv4 address 10.33.74.13) and an external host (10.33.69.132).

1. On the system console, log in to the switch as Admin.
2. Enable IPsec.
 - a) Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the IPsec RBAC class of commands.
 - b) Enter the `ipSecConfig --enable` command to enable IPsec on the switch.
3. Create an IPsec SA policy named AH01, which uses AH protection with MD5.

```
switch:admin> ipseconfig --add policy ips sa -t AH01 -p ah -auth hmac_md5
```

4. Create an IPsec proposal IPSEC-AH to use AH01 as SA.

```
switch:admin> ipseconfig --add policy ips sa-proposal -t IPSEC-AH -sa AH01
```

5. Configure the SA proposal's lifetime in time units. The maximum lifetime is 86400, or one day.

```
switch:admin> ipseconfig --add policy ips sa-proposal -t IPSEC-AH
-lifetime 86400 -sa AH01
```

6. Import the pre-shared key file using the `seccertmgmt` command. The file name should have a .psk extension.

For more information on importing the pre-shared key file, refer to [Installing a Switch Certificate](#) on page 213.

7. Configure an IKE policy for the remote peer.

```
switch:admin> ipsecconfig --add policy ike -t IKE01 -remote 10.33.69.132
-id 10.33.74.13 -remoteid 10.33.69.132 -enc 3des_cbc -hash hmac_md5
-prf hmac_md5 -auth psk -dh modp1024 -psk ipseckey.psk
```

NOTE

IKE version ('-v' option) needs to be set to 1 (IKEv1) if remote peer is a Windows XP or 2000 Host as Windows XP and 2000 do not support IKEv2.

8. Create an IPsec transform named TRANSFORM01 to use transport mode to protect traffic identified for IPsec protection and use IKE01 as key management policy.

```
switch:admin> ipsecconfig --add policy ips transform -t TRANSFORM01
-mode transport -sa-proposal IPSEC-AH -action protect -ike IKE01
```

9. Create traffic selectors to select the outbound and inbound traffic that needs to be protected.

```
switch:admin> ipsecconfig --add policy ips selector -t SELECTOR-OUT -d out
-l 10.33.74.13 -r 10.33.69.132 -transform TRANSFORM01
```

```
switch:admin> ipsecconfig --add policy ips selector -t SELECTOR-IN -d in
-l 10.33.69.132 -r 10.33.74.13 -transform TRANSFORM01
```

10. Verify the IPsec SAs created with IKE using the **ipSecConfig --show manual-sa -a** command.
11. Perform the equivalent steps on the remote peer to complete the IPsec configuration. Refer to your server administration guide for instructions.
12. Generate IP traffic and verify that it is protected using defined policies.
- Initiate Telnet or SSH or ping the session from the switch to the remote host.
 - Verify that the IP traffic is encapsulated.
 - Monitor IPsec SAs created using IKE for the traffic flow.
 - Use the **ipSecConfig --show manual-sa -a** command with the operands specified to display the outbound and inbound SAs in the kernel SADB.
 - Use the **ipSecConfig --show policy ips sa -a** command with the specified operands to display all IPsec SA policies.
 - Use the **ipSecConfig --show policy ips sa-proposal -a** command with the specified operands to display IPsec proposals.
 - Use the **ipSecConfig --show policy ips transform -a** command with the specified operands to display IPsec transforms.
 - Use the **ipSecConfig --show policy ips selector -a** command with the specified operands to display IPsec traffic selectors.
 - Use the **ipSecConfig --show policy ike -a** command with the specified operands to display IKE policies.
 - Use the **ipSecConfig --flush manual-sa** command with the specified operands to flush the created SAs in the kernel SADB.

ATTENTION

Flushing SAs requires IPsec to be disabled and re-enabled. This operation is disruptive to traffic using the tunnel.

In-flight Encryption and Compression

- In-flight encryption and compression overview..... 276
- Supported Gen 5 and Gen 6 in-flight encryption and compression features..... 277
- How in-flight encryption and compression are enabled..... 285
- Configuring in-flight encryption and compression on an EX_Port..... 288
- Configuring in-flight encryption and compression on an E_Port..... 289
- Viewing the encryption and compression configuration..... 289
- Configuring and enabling authentication for in-flight encryption..... 290
- Enabling in-flight encryption..... 292
- Enabling in-flight compression..... 293
- Disabling in-flight encryption..... 294
- Disabling in-flight compression..... 295
- Encryption and Trunking Support on FC32-48 Blades..... 295

In-flight encryption and compression overview

In-flight encryption provides security for frames while they are in flight between two switches. In-flight compression provides better bandwidth use on the ISLs, especially over long distance.

The in-flight encryption and compression features allow frames to be encrypted or compressed at the egress point of an ISL between two Brocade switches, and then to be decrypted or decompressed at the ingress point of the ISL. 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 in-flight encryption or compression.

NOTE

Both ends of the ISL must terminate in a Gen 5 or Gen 6 switch FC port.

Encryption and compression can be enabled at the same time, or you can enable either encryption or compression selectively. [Figure 1](#) shows an example of 16 Gbps links connecting three Brocade switches. One link is configured with encryption and compression, one with just encryption, and one with just compression.

FIGURE 14 Encryption and compression on 16 Gbps ISLs



Supported Gen 5 and Gen 6 in-flight encryption and compression features

TABLE 70 Supported Gen 6 (32-Gbps) in-flight encryption and compression features

Feature	G630	G620	FC32-48	SX6 (Compression Only)
Number of front-end port groups (ASICs/chips)	3 port groups	1 port group	2 port groups	1 port group
Number of ports per ASIC/chip/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, ISL/IFL between Gen 5 and Gen 6 ports			
Port types that do not support encryption or compression but can exist along the I/O path	VE_Ports, VEX_Port, 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	64 Gbps 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		

TABLE 70 Supported Gen 6 (32-Gbps) in-flight encryption and compression features (continued)

Feature	G630	G620	FC32-48	SX6 (Compression Only)
Other features supported on the in-flight encryption or compression enabled ports	Trunking, QoS, and long distance	QoS and long distance (trunking is NOT supported)	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 mandatorily when in-flight encryption is enabled on all the supported number of ports.	None	SFP+ ports 44-47 (Also, 40-43 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 44-47	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, FC32-64 port blade, CR32-8 and CR32-4 core routing blades, and the 7810 Extension Switch do not support in-flight encryption or compression.
- The encryption trunking on FC32-48 blade is not supported in 8.2.1 release.

TABLE 71 Supported Gen 5 (16-Gbps) in-flight encryption and compression features

Feature	6520	6510	7840	FC16-64	FC16-48	FC16-32	OEM Blade Server SAN I/O modules (refer note below)
Number of front-end port groups (ASICs/chips)	4 port groups	1 port group	1 port group	2 port groups	2 port groups	2 port groups	1 port group
Number of ports per ASIC/chip/port group that can be enabled with encryption or compression	2 ports @ 16/ASN (Auto Speed Negotiation) Gbps, 3 ports @ 10 Gbps and 4 ports @ 2/4/8 Gbps						
Total number of ports per switch/blade that can be enabled with encryption or compression (row 1 x row 2)	8 to 16 ports	2 to 4 ports	2 to 4 ports	4 to 8 ports	4 to 8 ports	4 to 8 ports	2 to 4 ports
Port types that support encryption or compression	E_Ports, EX_Ports, XISL ports in VF mode, ISL/IFL between Gen 5 and Gen 6 ports						
Port types that do not support encryption or compression but can exist along the I/O path	VE_Ports, VEX_Port, GE_Ports, F_Ports, F_Port trunks, ICLs, D_Ports, AE_Ports						
Maximum total bandwidth of the ports online per port group with encryption and/or compression enabled	32 Gbps per port group						
Other features supported on the in-flight encryption or compression enabled ports	trunking, QoS, and long distance						
Minimum Fabric OS version that supports encryption	7.1.0	7.0.0	7.3.0	7.3.0	7.0.0	7.0.0	Varies
Minimum Fabric OS version that supports compression	7.1.0	7.0.0	7.3.0	7.3.0	7.0.0	7.0.0	Varies
Supported flow control modes	R_RDY, VC_RDY, and EXT_VC_RDY						

NOTE

Brocade 6505 switch, any switch in Access Gateway mode, CR16-8 and CR16-4 core routing blades do not support in-flight encryption or compression.

NOTE

The supported OEM Blade Server SAN I/O modules are Brocade 6542 (Inspur), 6543 (Huawei), M6505 (Dell), 6546 (Hitachi), 6547 (Lenovo), 6548 (HPE), and 6558 (HPE) switch modules.

Supported use cases for Brocade G630 switch

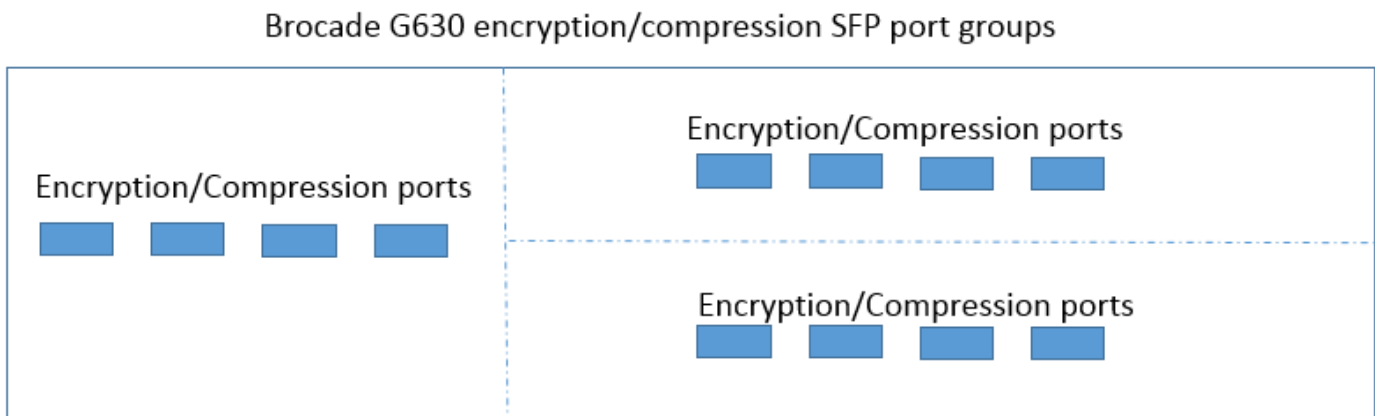
The following are the Brocade G630 switch encryption/compression SFP port groups as you see them in the front panel:

TABLE 72 Brocade G630 encryption/compression SFP port groups

Encryption/compression port group 1 SFP ports 0-15 and 58-63	Encryption/compression port group 2 SFP ports 16-47
	Encryption/compression port group 3 SFP ports 64-95

In this section, the following format is used to represent the Brocade G630 switch with any four ports enabled with encryption and/or compression in each port group:

FIGURE 15 Encryption/compression port groups



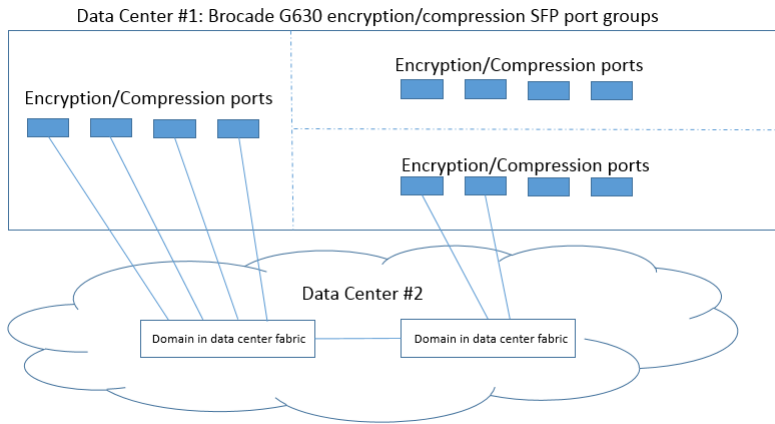
Between two data centers

Only one domain per port group is supported. It is recommended to 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.

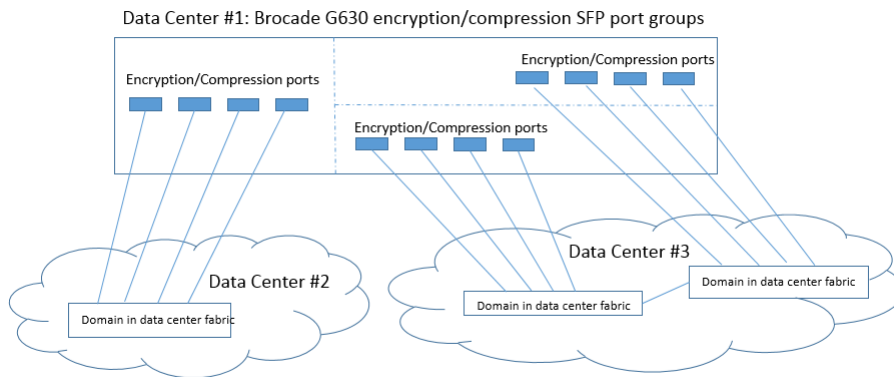
FIGURE 16 Between two data centers



Among three data centers

Only one domain per port group is supported. It is recommended to use all the ports in a port group before going to the next port group to plan for future data center connections.

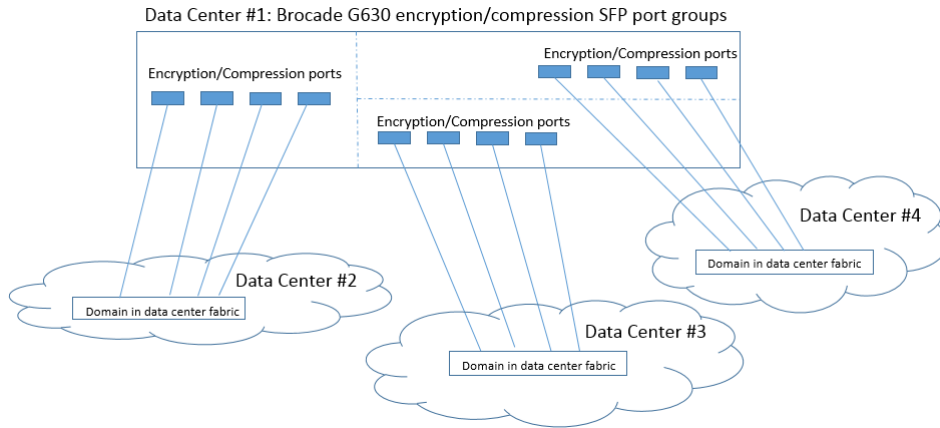
FIGURE 17 Among three data centers



Among four data centers

Only one domain per port group is supported.

FIGURE 18 Among four data centers



With FCR support on 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 LSAN zone on the same port group with EX_Ports in non-VF mode. However, this configuration is supported in case of VF mode.

FIGURE 19 Supported configuration with FCR

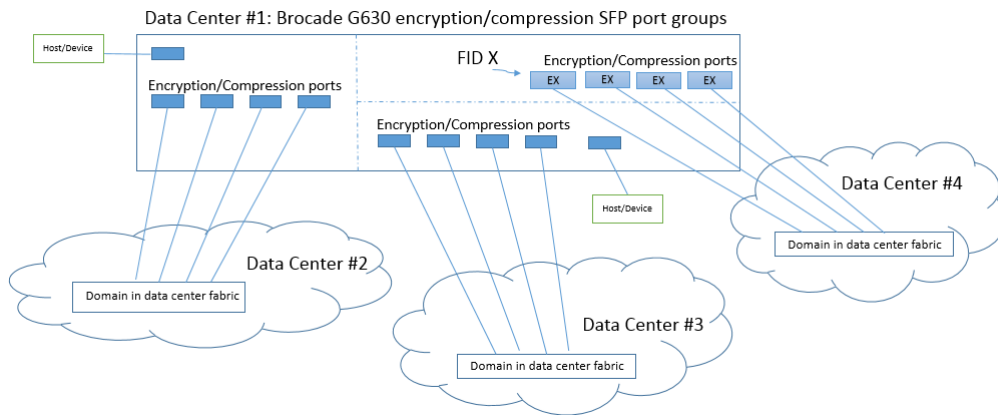
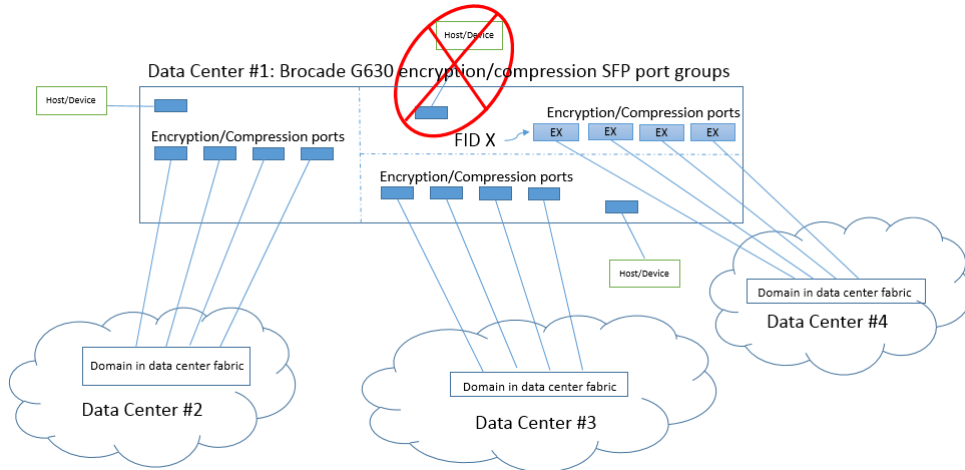


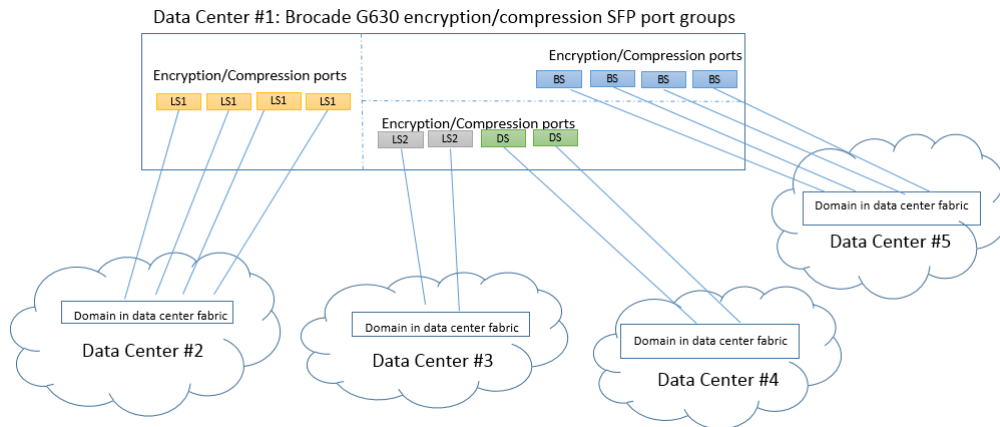
FIGURE 20 Unsupported configuration with FCR



With virtual fabrics support on 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 21 With virtual fabrics support on G630

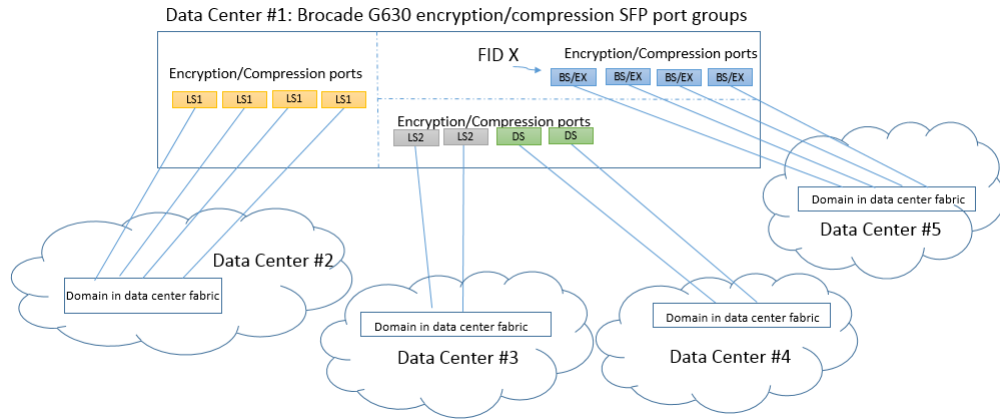


- BS—Base Switch
- DS—Default Switch
- LS1—Logical Switch 1
- LS2—Logical Switch 2

With virtual fabrics and FCR support on G630

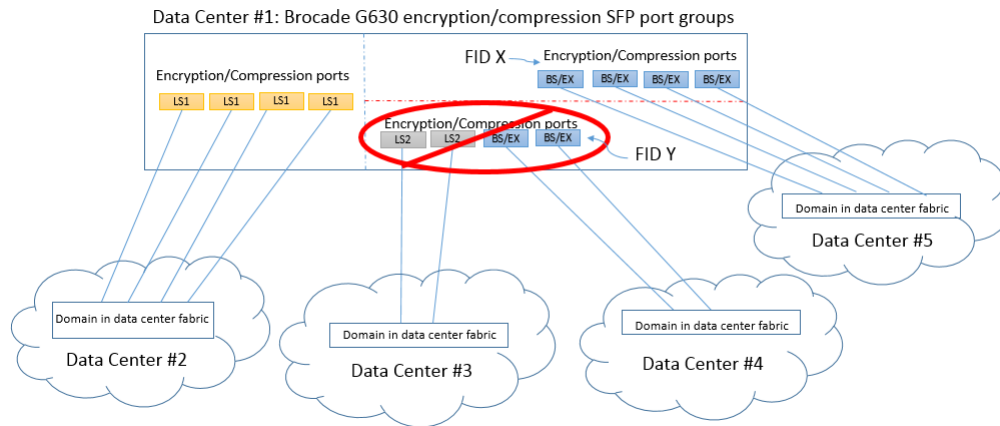
Only one domain per logical switch, the default switch or the base switch is allowed. Only one domain per FID/edge fabric is allowed in a port group. You cannot have a logical switch and the base switch in the same port group.

FIGURE 22 Supported configuration with virtual fabrics and FCR



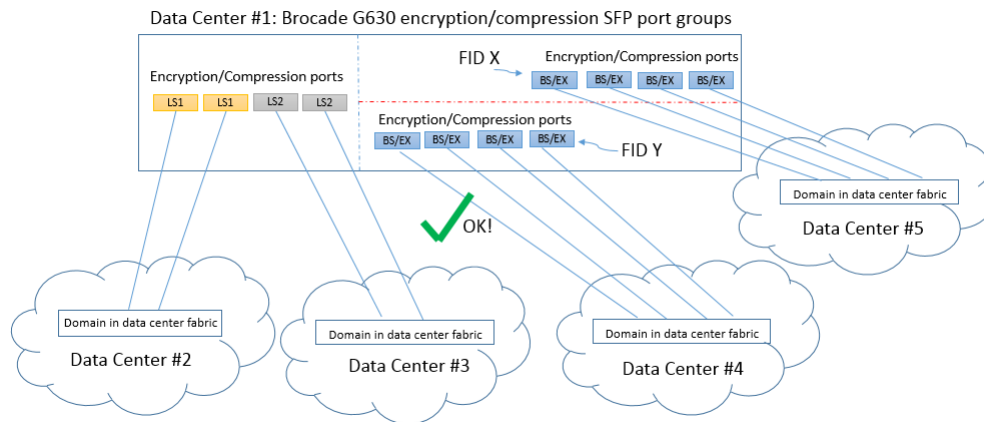
- BS—Base Switch
- DS—Default Switch
- LS1—Logical Switch 1
- LS2—Logical Switch 2

FIGURE 23 Unsupported configuration with virtual fabrics and FCR



- BS—Base Switch
- LS1—Logical Switch 1
- LS2—Logical Switch 2

FIGURE 24 Rectified configuration with virtual fabrics and FCR



- BS—Base Switch
- LS1—Logical Switch 1
- LS2—Logical Switch 2

Supported port speed on encryption or compression-enabled Gen 5 (16-Gbps) ports

The port speed determines the maximum number of ports on a device that can support the in-flight encryption and compression features.

If the port speed is configured as AUTO NEG, the speed of the port is taken as 16 Gbps for calculation purposes. It is recommended that you configure the ports to a specific speed before enabling encryption or compression.

The port speed values can be displayed through several commands, including `portEncCompShow`, `portShow`, and `switchShow`.

You can change the port speed on any port that has encryption or compression enabled with the `portCfgSpeed` command. If the capacity is available, the port is configured with the new speed. If there is not enough capacity available, you cannot increase the port speed.

Refer to [Setting port speeds](#) on page 87 for more information.

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 8.2.1 and later, additional considerations and restrictions apply. Refer to [Encryption and Trunking Support on FC32-48 Blades](#) on page 295 for details.

- FCP data frames and non-FCP data frames except ELS and 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.
- In-flight encryption is not FIPS certified.

Availability considerations for encryption and compression

To provide redundancy in the event of encryption or compression port failures, you should connect each ISL or trunk group to different ASICs on the peer switch. Refer to 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, it is recommended to 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, it is recommended to 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 XISL ports, but not on 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.

For considerations when enabling encrypted ports for trunking on a logical switch when using a FC32-48 blade, refer to [Encryption and Trunking Support on FC32-48 Blades](#) on page 295.

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, then 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, then 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. Refer to [Port decommissioning](#) on page 86 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
 - FIPS mode configuration:
 - > Authentication Type - DH-CHAP
 - > Hash Type - SHA256, SHA1
 - > Group Type - 4
 - > Switch Authentication Policy - Passive

- The in-flight encryption protocol supports the 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 will be generated during authentication phase and it will be used in IKE (Internet Key Exchange) 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. Once 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 pre-requisite.
- 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 master port. Slave ports do not exchange keys. If the master 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, it is recommended to 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 five seconds, and is the ratio between the accumulated original length and the compressed length of the data over the previous five seconds.

When a port is configured for compression, entering **portStatsShow** displays the port's compression ratio. The value shown by **portStatsShow** is a five-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. This is because the compressed length then also includes the encryption header. This overhead affects the ratio calculation. To obtain accurate compression ratio data, it is recommended that you 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 Director.

1. Determine which ports are available for encryption or compression.

Refer to [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. You can use the **fcredgeshow** 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.

Refer to [Configuring and enabling authentication for in-flight encryption](#) on page 290 for instructions.

5. Enable encryption on the port.
Refer to [Enabling in-flight encryption](#) on page 292 for instructions.
6. Enable compression on the port.
Refer to [Enabling in-flight compression](#) on page 293 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.
Refer to [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.
Refer to [Configuring and enabling authentication for in-flight encryption](#) on page 290 for instructions.
3. Enable encryption on the port.
Refer to [Enabling in-flight encryption](#) on page 292 for instructions.
4. Enable compression on the port.
Refer to [Enabling in-flight compression](#) on page 293 for instructions.

Following successful port initialization, the configured features are enabled and active. You can 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. Refer to [Supported Gen 5 and Gen 6 in-flight encryption and compression features](#) on page 277 to determine 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. Refer to [Supported Gen 5 and Gen 6 in-flight encryption and compression features](#) on page 277 to determine ports available per ASIC

```
switch:admin> portenccompshow
```

Index	Encryption Configured	Active	Compression Configured	Active	Config 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 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, then all switches in the fabric must enable DH-CHAP and establish pre-shared secrets. If **fcap** is specified, then 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
```

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 8 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                DId    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, refer to [Authentication policy for fabric elements](#) on page 245.

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. Refer to [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](#) on page 290.

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 44-47.

3. Enter the **portCfgEncrypt --enable** command.

The following example enables encryption on port 15 of an FC16-32 blade in slot 9 of an enterprise class platform:

```
switch:admin> portcfgencrypt --enable 9/15
```

4. Enter the **portEnable** command to enable the port.

After manually enabling the port, the new configuration becomes active.

5. You can verify the configuration using the **portCfgShow** command.

```
switch:admin> portcfgshow 0
Area Number:          0
Octet Speed Combo:   3 (16G,10G)
(output truncated)
D-Port mode:         OFF
D-Port over DWDM:    ..
Compression:         OFF
Encryption:          ON
```

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, it is recommended that you 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. Refer to [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.

The following example enables compression on port 15 in slot 9 of an enterprise class platform:

```
switch:admin> portcfgcompress --enable 9/15
```

4. Enter the **portEnable** command to enable the port.

After enabling the port, the new configuration becomes active.

The following example enables compression on port 0.

```
switch:admin> portdisable 0
switch:admin> portcfgcompress --enable 0
switch:admin> portenable 0
```

You can verify the configuration using the **portCfgShow** command.

```
switch:admin> portcfgshow 0
Area Number:          0
Octet Speed Combo:   3 (16G,10G)
(output truncated)
D-Port mode:         OFF
D-Port over DWDM     ..
Compression:         ON
Encryption:          ON
```

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. Disable encryption on the port using the **portCfgEncrypt --disable** command.

The following example disables encryption on port 15 in slot 9 of an enterprise class platform:

```
switch:admin> portcfgencrypt --disable 9/15
```

4. Enable the port using the **portEnable** command.

The following example disables encryption on port 0.

```
myswitch:admin> portdisable 0
myswitch:admin> portcfgencrypt --disable 0
myswitch:admin> portenable 0
```

You can verify the configuration using the **portCfgShow** command.

```
myswitch:admin> portcfgshow 0
Area Number:          0
Speed Level:         AUTO (SW)
(output truncated)
D-Port mode:         OFF
D-Port over DWDM     ..
Compression:         OFF
Encryption:          OFF
```

Disabling in-flight compression

Disabling compression is an offline event. Ports must be disabled first, and then re-enabled after.

NOTE

Firmware downgrade from Fabric OS 7.3.0 to an earlier version is blocked if in-flight encryption with FCAP protocol is set. Please set the DHCHAP 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.

The following example disables compression on port 15 in slot 9 of an enterprise class platform:

```
switch:admin> portcfgcompress --disable 9/15
```

4. Enable the port using the **portEnable** command.

You can verify the configuration using the **portCfgShow** command.

```
myswitch:admin> portcfgshow 0

Area Number:          0
Speed Level:          AUTO (SW)
(output truncated)
D-Port mode:          OFF
D-Port over DWDM:     ..
Compression:          OFF

Encryption:           OFF
```

The following example disables compression on port 0.

```
myswitch:admin> portdisable 0
myswitch:admin> portcfgcompress --disable 0
myswitch:admin> portenable 0
```

Encryption and Trunking Support on FC32-48 Blades

Fibre Channel ports configured for encryption can be configured in a trunk group on the FC32-48 blade starting with Fabric OS v8.2.1. 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 will be disabled. If you want to reenabling 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, refer to [Managing Trunking Connections](#) on page 335. For requirements for configuring trunking on FC32-48 blade encrypted ports, refer to the following "Recommendations and Considerations" section.

NOTE

The **portenccompshow** command displays ports by index number that are within a specific ASIC/chip/group and identifies those ports that are configured and active for compression and encryption. Refer to [Viewing the encryption and compression configuration](#) on page 289 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 will be 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 will be disabled with an appropriate error.
- When you configure encryption on a port, trunking will be enabled by default. The port will attempt 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 enabled for compression have always been supported for trunking in previous Fabric OS versions. With Fabric OS 8.2.1, ports 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 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 v8.2.1 or greater installed. If one end runs earlier Fabric OS than v8.2.1, only one trunk with single port will come online. In this case, we recommend that you disable trunking on encrypted ports for backward compatibility.
- Downgrade to a Fabric OS release prior to v8.2.1 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 v8.2.1.

Buffer-to-Buffer Credits and Credit Recovery

- Buffer credit management298
- Buffer credit recovery312
- Credit loss detection..... 314

Buffer credit management

Buffer-to-buffer credit management affects performance over distances; therefore, allocating a sufficient number of buffer credits for 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 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 on the basis of this credit system helps prevent the loss of frames and reduces the frequency of entire Fibre Channel sequences needing to be retransmitted across the link.

Because 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 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 is not a recommended practice.

All switch ports provide protection against 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 measured during 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 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 is used to manage 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 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. When QoS is enabled, EXT_VC_RDY flow-control mode allocates VC0 through VC14. 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 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 Fabric OS 7.x and 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. Refer to [Calculating the number of buffers required based on full-size frames](#) on page 301 and [Configuring buffers for a single port directly](#) on page 303 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 73 Fibre Channel gigabit values

Gigabit value	Line rate
1 Gbps	1.0625
2 Gbps	2.125

TABLE 73 Fibre Channel gigabit values (continued)

Gigabit value	Line rate
4 Gbps	4.25
8 Gbps	8.5
10 Gbps	10.625
16 Gbps	14.025
32 Gbps	28.05

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 2,112 bytes (2,148 with headers). Assuming that each payload is 2,112, you need one credit per 1 km of link length at 2 Gbps (smaller payloads require additional buffer credits to maintain link utilization). Refer to [Allocating buffer credits based on average-size frames](#) on page 303 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 2,112 (because $[2,112 - 64] = 2,048$, which is 2 kb of data). The final frame, after it is constructed, is passed through the 8-byte-to-10-byte conversion process.

[Table 74](#) describes Fibre Channel data frames.

TABLE 74 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)	0-2,112 bytes	0-16,896 bits
CRC	4 bytes	32 bits
End of frame	4 bytes	32 bits
Total (number bits/frame)	36-2,148 bytes	288-7,184 bits

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 over LD. The assumption that Fibre Channel payloads are consistently 2,112 bytes is not realistic in practice. To gain the proper number of buffer credits with the LS mode, there must be enough buffer credits available in the pool, because Fabric OS will check before accepting a value.

NOTE

The *desired_distance* parameter of the **portCfgLongDistance** command's is the upper limit of the link distance and is used to calculate buffer availability for other ports in the same port group. When the measured distance exceeds the value of *desired_distance*, this value is used to allocate 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 is used to allocate the buffers required for the port.

Refer to the data in [Table 1](#) and [Table 1](#) 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. The values reflect an estimate, and may differ from the supported values in [Table 1](#).

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:

$$(\text{Reserved_Buffer_for_Distance_Y}) = (X * \text{LinkSpeed} / 2) + 6 + 14$$

- If QoS is not enabled:

$$(\text{Reserved_Buffer_for_Distance_Y}) = (X * \text{LinkSpeed} / 2) + 6$$

The formulas use the following parameters:

X = 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 2 Gbps as the speed of the long-distance connection, insert the numbers into the appropriate formula. The formula should read as follows:

$$(50 \text{ km} * 2 \text{ Gbps} / 2) + 6 = 56 \text{ buffers, which is the number of buffers reserved for distance.}$$

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 1 Gbps, then $(50 \text{ km} * 1 \text{ Gbps} / 2) + 6 = 31$ buffers.
- If you have a distance of 50 km at 2 Gbps, then $(50 \text{ km} * 2 \text{ Gbps} / 2) + 6 = 56$ buffers.
- If you have a distance of 50 km at 4 Gbps, then $(50 \text{ km} * 4 \text{ Gbps} / 2) + 6 = 106$ buffers.
- If you have a distance of 50 km at 8 Gbps, then $(50 \text{ km} * 8 \text{ Gbps} / 2) + 6 = 206$ buffers.
- If you have a distance of 50 km at 10 Gbps, then $(50 \text{ km} * 10 \text{ Gbps} / 2) + 6 = 256$ buffers.
- If you have a distance of 50 km at 16 Gbps, then $(50 \text{ km} * 16 \text{ Gbps} / 2) + 6 = 406$ buffers.
- If you have a distance of 50 km at 32 Gbps, then $(50 \text{ km} * 32 \text{ Gbps} / 2) + 6 = 806$ 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 retrieved from [Table 1](#)

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:

$$\text{Maximum_Distance_X (in km)} = (\text{BufferCredits} + 6) * 2 / \text{LinkSpeed}$$

$$498 \text{ km} = (492 + 6 \text{ buffers for Fabric Services}) * 2 / 2 \text{ Gbps}$$

If you have a distance of 50 km at 8 Gbps, 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 8 Gbps
- 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 1 Gbps, then $484 / (31 - 8) = 21$ ports.

Allocating buffer credits based on average-size frames

In cases where the frame size is average, for example 1,024 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](#) on page 299 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:

$$\text{desired_distance} = \text{roundup} [(real_estimated_distance * 2112) / average_payload_size]$$

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 8-Gbps E_Port. This 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 8 Gbps.
3. Look up the *data_rate* value for the speed of the connection. Refer to [Fibre Channel gigabit values reference definition](#) on page 299 to determine the *data_rate* value.

For example, the *data_rate* is 8.5 for a speed of 8 Gbps.

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 15; 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:

Total Buffers = Configured Buffers + QoS_VC_Credits + 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 6 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, 6 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.

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.

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:

*buffers_required = (2048/framesize) * data_vc_credits*

If you enter the average frame size of 1024, Fabric OS will allocate almost twice as many buffers as for the maximum frame size of 2048.

The **-frameSize** option value is persistent across reboots and HA failover.

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 8-Gbps 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 16 Gbps and 32 Gbps:

```
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 eight buffers for all Gen 5 devices and 20 buffers for all Gen 6 devices, except the Brocade G610, which is fixed at eight buffers. You can use the **portcfgportbuffers** 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 **portcfgportbuffers** command.

```
switch:admin> portcfgportbuffers --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> portcfgportbuffers --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 Usage	Needed Buffers	Link Distance	Remaining Buffers
				Tx	Rx				
0		-	34	- (-)	- (-)	0	-	-	
1	F	-	20	- (-)	- (-)	8	-	-	
2	E	-	20	- (-)	- (-)	26	26	<2km	
3	E	-	20	- (-)	- (-)	34	34	<2km	
4	F	-	48	- (-)	- (-)	48	-	-	
5	F	-	48	- (-)	- (-)	48	-	-	
6	F	-	48	- (-)	- (-)	48	-	-	
7	F	-	48	- (-)	- (-)	48	-	-	
8	E	LS	1634	- (-)	- (-)	1634	1634	100km	
9		-	20	- (-)	- (-)	0	-	-	
10	E	-	34	- (-)	- (-)	8	8	<2km	
11		-	20	- (-)	- (-)	0	-	-	
12		-	20	- (-)	- (-)	0	-	-	
13		-	20	- (-)	- (-)	0	-	-	
14		-	20	- (-)	- (-)	0	-	-	
15		-	20	- (-)	- (-)	0	-	-	
16		-	20	- (-)	- (-)	0	-	-	
17		-	20	- (-)	- (-)	0	-	-	
18		-	20	- (-)	- (-)	0	-	-	
19		-	20	- (-)	- (-)	0	-	-	
20		-	20	- (-)	- (-)	0	-	-	
21		-	20	- (-)	- (-)	0	-	-	
22		-	20	- (-)	- (-)	0	-	-	
23		-	20	- (-)	- (-)	0	-	-	
24		-	20	- (-)	- (-)	0	-	-	
25		-	20	- (-)	- (-)	0	-	-	
26		-	20	- (-)	- (-)	0	-	-	
27		-	20	- (-)	- (-)	0	-	-	
28		-	20	- (-)	- (-)	0	-	-	
29		-	20	- (-)	- (-)	0	-	-	
32		-	20	- (-)	- (-)	0	-	-	
33		-	20	- (-)	- (-)	0	-	-	
34		-	20	- (-)	- (-)	0	-	-	
35		-	20	- (-)	- (-)	0	-	-	
36		-	20	- (-)	- (-)	0	-	-	
37		-	20	- (-)	- (-)	0	-	-	
38		-	20	- (-)	- (-)	0	-	-	
39		-	20	- (-)	- (-)	0	-	-	
40		-	20	- (-)	- (-)	0	-	-	
41		-	20	- (-)	- (-)	0	-	-	
42		-	20	- (-)	- (-)	0	-	-	
43		-	20	- (-)	- (-)	0	-	-	
44		-	20	- (-)	- (-)	0	-	-	
45		-	20	- (-)	- (-)	0	-	-	
46		-	20	- (-)	- (-)	0	-	-	
47	E	-	20	- (-)	- (-)	26	26	<2km	
48		-	20	- (-)	- (-)	0	-	-	
49		-	20	- (-)	- (-)	0	-	-	
50		-	20	- (-)	- (-)	0	-	-	
51		-	20	- (-)	- (-)	0	-	-	
52		-	20	- (-)	- (-)	0	-	-	
53		-	20	- (-)	- (-)	0	-	-	
54		-	20	- (-)	- (-)	0	-	-	
55		-	20	- (-)	- (-)	0	-	-	
56	E	-	20	- (-)	- (-)	70	70	<2km	
57	E	-	20	- (-)	- (-)	70	70	<2km	
58	E	-	20	- (-)	- (-)	70	70	<2km	
59	E	-	20	- (-)	- (-)	70	70	<2km	
60	E	-	20	- (-)	- (-)	70	70	<2km	
61	E	-	20	- (-)	- (-)	70	70	<2km	
62	E	-	20	- (-)	- (-)	105	105	<2km	
63	E	-	20	- (-)	- (-)	105	105	<2km	11522

Buffer credits per switch or blade model

The following tables shows 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 6 32-Gbps devices

Remaining buffers = Total Buffers - [(Total FE Ports * 20) + (Total BE/BI Ports * Buffers for BE/BI Port) + Internal buffer usages (approximately 152)].

TABLE 75 Total FC ports, ports per port group, and unreserved buffer credits per port group

Switch/blade model	Total FC ports (per switch/blade)	User port group size	Unreserved buffer credits per port group (with minimum buffer allocation)	Unreserved buffer credits per port group (with QoS)	Unreserved buffer credits per port group (without QoS)
G610	24	24	1728	960	1296
G620	64	64	13928	8488	10728
G630	128	32	6328	3608	4728
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, then the numbers for that port group are changed as per the following ECB table.

TABLE 76 Total FC ports, ports per port group, and unreserved buffer credits per port group

Switch/blade model	Total FC ports (per switch/blade)	User port group size	Unreserved buffer credits per port group (with minimum buffer allocation)	Unreserved buffer credits per port group (with QoS)	Unreserved buffer credits per port group (without QoS)	Unreserved buffer credits per port group (with ECB port)
7810	12	12	1376	1064	1160	NA
G630	128	32	6328	3608	4728	1816
FC32-64	64	32	7848	5128	6248	NA

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 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 FC32-48 port blade and G620 switch, when a port is configured with in-flight encryption, additionally 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.

- The **portcfgencrypt** command must be used with the **-force** option for the ports 40-47 on the G620 switch and has the following limitations:
 - You cannot have both Octet Combo Mode 2 and in-flight encryption enabled at the same time.
 - You cannot have both SIM port (using the **flow -control -enable** command) and in-flight encryption 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 40-47 cannot be configured to Octet Combo Mode 2.

Formula for Gen 5 16-Gbps and older devices

Remaining buffers = Total Buffers - [(Total FE Ports * 8) + (Total BE/BI Ports * Buffers for BE/BI Port) + EMB Buffers + SAB Buffers.

TABLE 77 Total FC ports, ports per port group, and unreserved buffer credits per port group

Switch/blade model	Total FC ports per switch/blade	User port group size	Unreserved buffer credits per port group (with minimum buffer allocation)	Unreserved buffer credits per port group (with QoS)	Unreserved buffer credits per port group (without QoS)
6505	24	24	7904	7280	7424
M6505	24	24	7904	7280	7424
6510	48	48	7712	6464	6752
6520	96	24	4784	4160	4304
6542	48	48	7712	6464	6752
6543	24	24	7904	7280	7424
6545	26	26	7888	7212	7368
6546	24	24	7904	7280	7424
6547	48	48	7712	6464	6752
6548	28	28	7872	7144	7312
6549	28	28	7904	7280	7424
6558	48	48	7712	6464	6752
6559	48	48	7712	6464	6752
7840	24	24	5024	4400	4544
FC16-32	32	16	5408/2464	4992/2048	5088/2144
FC16-48	48	24	4960/2016	4336/1392	4480/1536
FC16-64	64	32	4288/1344	3456/512	3648/704
CR16-8	64	16	5432/2072	4072/712	4632/1272
CR16-4	32	16			
FX8-24	12	12	988	676	748

For the FC16-x port blades and CR16-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 enabled on the slot. For Gen 5 and older devices, use the **bufOpMode** command to display or change the buffer optimized mode.

Minimum buffers allocated per port type

TABLE 78 Minimum buffers allocated per port type

Port type	Minimum buffers allocated per port
32-Gbps Gen 6 / Condor 4 - F_Port	20
32-Gbps Gen 6 / Condor 4 - E_Port (non QoS) or F_Port trunk without QoS	70
32-Gbps Gen 6 / Condor 4 - E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS	105
16/10/8/4-Gbps Gen 6 / Condor 4 - F_Port	8
16/10/8/4-Gbps Gen 6 / Condor 4 - E_Port (non QoS), or F_Port trunk without QoS	28
16/10/8/4-Gbps Gen 6 / Condor 4 - E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS	34
16/10/8/4/2-Gbps Gen 5 / Condor 3 - F_Port	8
16/10/8/4/2-Gbps Gen 5 / Condor 3 - E_Port (non QoS), or F_Port trunk without QoS	28
16/10/8/4/2-Gbps Gen 5 / Condor 3 - E_Port with QoS, F_Port with AoQ, or F_Port trunk with QoS	34

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 6 32-Gbps 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 79 Configurable distances for Gen 6 Extended Fabrics

Switch/blade model	Maximum distances (km) that can be configured (assuming a 2112-byte frame size)				
	4 Gbps	8 Gbps	10 Gbps	16 Gbps	32 Gbps
7810	*	*	NA	*	*
G610	861 (with min. buffers) 549 (with QoS) 645 (without QoS)	430 (with min. buffers) 274 (with QoS) 322 (without QoS)	N/A	215 (with min. buffers) 137 (with QoS) 161 (without QoS)	107 (with min. buffers) 59 (with QoS) 80 (without QoS)
G620	6068	3034	2427	1517	758
G630	3164 (with min. buffers) 1804 (with QoS) 2364 (without QoS) 908 (with ECB)	1582 (with min. buffers) 902 (with QoS) 1182 (without QoS) 454 (with ECB)	1265 (with min. buffers) 721 (with QoS) 945 (without QoS) 363 (with ECB)	791 (with min. buffers) 451 (with QoS) 591 (without QoS) 227 (with ECB)	395 (with min. buffers) 225 (with QoS) 295 (without QoS) 113 (with ECB)
FC32-48	4073	2036	1629	1018	509
FC32-64	3924 (with min. buffers) 2564 (with QoS) 3124 (without QoS)	1962 (with min. buffers) 1282 (with QoS) 1562 (without QoS)	1569 (with min. buffers) 1025 (with QoS) 1249 (without QoS)	981 (with min. buffers) 641 (with QoS) 781 (without QoS)	490 (with min. buffers) 320 (with QoS) 390 (without QoS)

TABLE 79 Configurable distances for Gen 6 Extended Fabrics (continued)

Switch/blade model	Maximum distances (km) that can be configured (assuming a 2112-byte frame size)				
	4 Gbps	8 Gbps	10 Gbps	16 Gbps	32 Gbps
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

Formula for Gen 5 16-Gbps and older devices

The remaining buffers = (Total FE Ports * 8) + (Total BE/BI Ports * Buffers for BE/BI Port) + EMB Buffers + SAB Buffers

TABLE 80 Configurable distances for Gen 5 Extended Fabrics

Switch/blade model	Maximum distances (km) that can be configured (assuming a 2112-byte frame size)				
	2 Gbps	4 Gbps	8 Gbps	10 Gbps	16 Gbps
M6505	N/A	3949	1974	N/A	987
6505	7898	3949	1974	N/A	987
6510	7706	3853	1926	1541	963
6520	4778	2389	1194	955	597
6542	N/A	3853	1926	N/A	963
6543	N/A	3949	1974	N/A	987
6545	N/A	3941	1970	N/A	985
6546	N/A	3949	1974	N/A	987
6547	N/A	3853	1926	N/A	963
6548	N/A	3933	1966	N/A	983
6549	N/A	3949	1974	N/A	987
6558	N/A	3853	1926	N/A	963
6559	N/A	3853	1926	N/A	963
7840	5018	2509	1254	1003	627
FC16-32	5402/2458	2701/1229	1350/614	1080/491	675/307
FC16-48	4954/2010	2477/1005	1238/502	990/402	619/251
FC16-64	4282/1338	2141/669	1070/334	856/267	535/167
	*** Extended Fabrics is not recommended on this blade due to buffer limitations but can be configured. ***				
FX8-24	982	491	245	N/A	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.

NOTE

For the FC16-x port blades, the first number in the speed column designates the distance (km) supported at that speed without buffer optimized mode; the second number designates the distance (km) supported at that speed with buffer optimized mode enabled on the slot. Use the **bufopmode** command to display or change the buffer optimized mode

For example, for three ports running at 32 Gbps 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 5 devices (16 Gbps-capable blades), the minimum credit allocation is 5 and the maximum is 40. For ICL ports of Gen 5 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, 32 Gbps links are reserved 20 credits per VC). The configured credits will be allocated for each of the medium virtual channels (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** command to allocate credits to an E_Port. In the following example, 14 credits are allocated to each of the medium Virtual Channels (VCs) for non-QoS ports, and to both the medium and high VCs for QoS ports.

```
switch:admin> portcfgeportcredits --enable 12/6 14
Success
```

3. Enter the **portCfgEPortCredits --show** 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> portcfgeportcredits --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 16-Gbps (Gen 5) and 32-Gbps (Gen 6) FC switches and blades:

- You can 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 VCO). If VC init is off, and ISL_RDY mode is set, then VCO will support credit recovery.

Buffer credit recovery over an E_Port

To support buffer credit recovery, E_Ports must be connected between devices that support 32 or 16 Gbps or between devices that support 8 Gbps.

If a device that supports 32/16 Gbps is connected to a device that supports only 8 Gbps, buffer credit recovery is disabled, even if both devices are running 8 Gbps.

There is a limited amount of credit recovery that occurs between 32/16G platforms and 8G platforms. In order for buffer credit recovery to occur between 32 or 16 Gbps devices and 8 Gbps 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 v7.1 or later.
- Fabric OS must support buffer credit recovery at either end of the link.
- If both devices support 16 Gbps, 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 v7.1 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 v7.1 or later.
- The FCR and the switch at either end of the IFL must Gen 5 (16 Gbps) or Gen 6 (32 Gbps).

- Either end of the IFL must support buffer credit recovery.
- If the inter-fabric link (IFL) connects devices that support 8 Gbps only, long-distance mode must also be enabled. Long-distance mode can be enabled or disabled on devices that support 16 and 32 Gbps.
- Virtual Channel flow control (VC_RDY) 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 has been 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 7.1.0 and later supports credit loss detection for back-end ports and core blades, and on the Brocade 5300 and Brocade 6520 switches, although the support is slightly different on each device. Refer to the following topics for information on credit loss detection for these devices; and the *Brocade Fabric OS Troubleshooting and Diagnostics Guide* for more general information on credit loss detection.

Back-end credit loss detection and recovery support on Brocade 6520 switches

The following credit loss detection methods are supported for Brocade 6520 back-end ports:

- Per-port polling to detect credit loss. If credit loss is detected using this method, the RASlog C3-1012 message is displayed and recorded.
- Per-VC credit loss detection. If single-credit loss is detected using this method, it will be automatically recovered and the RASlog C3-1023 message is displayed and recorded. If multi-credit loss is detected using this method, the RASlog C3-1013 message is displayed and recorded. There is no automatic recovery for multi-credit loss.

- Complete VC credit loss detection. If credit loss is detected using this method, the RASlog C3-1011 message is displayed and recorded.

The following credit loss recovery methods are supported for Brocade 6520 back-end ports:

- For all the credit loss methods described previously, a link reset will automatically be performed, assuming that this option was enabled. Refer to [Enabling back-end credit loss detection and recovery for link reset thresholds](#) on page 315 for details on enabling this feature.
- A manual link reset option using the **creditRecovMode** command is also available. Refer to [Enabling back-end credit loss detection and recovery for link reset thresholds](#) on page 315 for instructions.

NOTE

Whenever a link reset is performed on this switch, the RASlog C3-1014 message is displayed and recorded.

Enabling back-end credit loss detection and recovery for link reset thresholds

Credit loss detection and recovery is enabled and disabled using the **creditrecovmode** command, which is only supported on the back-end ports of Brocade 5300, Brocade 6520, and 16 and 32 Gbps-capable Fibre Channel blades in the chassis of Brocade DCX 8510 Backbones and X6 Directors.

ATTENTION

This command might be constrained by Virtual Fabric restrictions. Refer to the *Brocade Fabric OS Troubleshooting and Diagnostics Guide* for more information.

In versions of Fabric OS prior to 7.4.0, the fault option in **creditrecovmode** applied only if a link reset threshold is set. In Fabric OS 7.4.0 and later, **creditrecovmode --fault** makes this fault option applicable to all back-end link failures. Which blade fault is registered due to a back-end link failure is decided by the blade fault option. If credit recovery mode is off and there are any back-end link failures, then Fabric OS will fault the core blade; if there is only one core blade, then Fabric OS will fault the edge blade. Refer to [Back-end credit loss detection and recovery for link faults](#) on page 317 for additional information about using **creditrecovmode --fault**.

To enable back-end credit loss detection and recovery, perform the following steps.

For more information on **creditrecovmode** command options, refer to the *Brocade Fabric OS Command Reference*.

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 back-end ports.

In the following example, back-end 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 back-end port credit recovery configuration. In the following example, back-end 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 link reset for loss of sync from peer ports

You can perform link reset on all the 8-Gbps, 16-Gbps, and 32-Gbps backend ports if there is a loss of sync from peer ports. You can configure either or both of the following link resets.

1. Run the `creditrecovmode --be_crdloss [off|on]` to configure the credit loss detection link reset. By default, it is off.
2. Run the `creditrecovmode --be_losync [off|on]` to configure the loss of sync detection link reset. By default, it is off.

The link reset action is decided based on the `--lrthreshold`, `--crdLoss`, and `--losync` settings as per the following matrix.

TABLE 81 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 command:

```
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 command.

```
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 command.

```
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 command.

```
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 command.

```
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
```

Back-end credit loss detection and recovery for link faults

The following table identifies the blade that is faulted based on the factors.

TABLE 82 Blade fault scenarios

Fault Option	Back-end link failure location	Edge chip fault count	Core chip 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

TABLE 82 Blade fault scenarios (continued)

Fault Option	Back-end link failure location	Edge chip fault count	Core chip fault count	Faulted blade when director has single active core	Faulted blade when director has dual active cores
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

- Adaptive Networking overview.....319
- Ingress Rate Limiting..... 320
- QoS..... 321
- CS_CTL-based frame prioritization..... 331

Adaptive Networking overview

Adaptive Networking is a suite of tools and capabilities that enable you to ensure optimized behavior in the SAN. Under the worst congestion conditions, Adaptive Networking can maximize the fabric behavior and provide necessary bandwidth for high-priority, mission-critical applications and connections.

The Adaptive Networking suite includes the following features:

- **Bottleneck Detection**

The Bottleneck Detection feature identifies devices attached to the fabric that are slowing down traffic. Bottleneck Detection does not require a license.

- **Top Talkers**

The Top Talkers feature provides real-time information about the top n bandwidth-consuming flows passing through a specific port in the network. Top Talkers requires a Fabric Vision license or an Advanced Performance Monitoring license.

- **Traffic Isolation Zoning**

Traffic Isolation Zoning (TI zoning) allows you to control the flow of interswitch traffic by creating a dedicated path for traffic flowing from a specific set of source ports (F_Ports). Traffic Isolation Zoning does not require a license. Refer to [Traffic Isolation Zoning](#) on page 484 for more information about this feature.

- **Ingress Rate Limiting**

Ingress Rate Limiting restricts the speed of traffic from a particular device to the switch port. Ingress Rate Limiting does not require a license. Refer to [Ingress Rate Limiting](#) on page 320 for more information about this feature.

- **Quality of Service (QoS)**

QoS allows you to categorize the traffic flow between a host and target as having a high, medium, or low priority. QoS does not require a license. Refer to [QoS](#) on page 321 for more information about this feature.

You can use the Adaptive Networking features together to optimize the performance of your fabric. For example, you can use the features in the following ways:

- You can use Top Talkers to identify the SID/DID pairs that consume the most bandwidth and can then configure them with certain QoS attributes so they get proper priority.
- If the bottleneck detection feature detects a latency bottleneck, you can use TI zones or QoS to isolate latency device traffic from high-priority application traffic.
- If the bottleneck detection feature detects ISL congestion, you can use Ingress Rate Limiting to slow down low-priority application traffic if it is contributing to the congestion.

Ingress Rate Limiting

Ingress Rate Limiting restricts the speed of traffic from a particular device to the switch port.

Use Ingress Rate Limiting for the following situations:

NOTE

Gen 6 platforms do not support ingress rate limiting.

- To reduce existing congestion in the network or pro-actively avoid congestion.
- To enable you to offer flexible bandwidth-limit services based on requirements.
- To enable more important devices to use the network bandwidth during specific services, such as network backup.

To limit the traffic, you set the maximum speed at which the traffic can flow through a particular F_Port. For example, if you set the rate limit at 4-Gbps, then traffic from a particular device is limited to a maximum of 4-Gbps.

Ingress Rate Limiting enforcement is needed only if the port can run at a speed higher than the rate limit. For example, if the rate limit is 4-Gbps and the port is only a 2-Gbps port, then Ingress Rate Limiting is not enforced. The Ingress Rate Limiting configuration is persistent across reboot.

The following considerations about Ingress Rate Limiting:

- Ingress Rate Limiting is applicable only to F_Ports.
- QoS takes precedence over Ingress Rate Limiting.
- Ingress Rate Limiting is not enforced on trunking ports.
- Ingress Rate Limiting can also be used on simulation ports (SIM ports).

Virtual Fabrics considerations for ingress rate limiting

If Virtual Fabrics is enabled and if a port is configured to have a certain rate limit value, you must first disable the rate limit on the port before moving it to a different logical switch. Ports cannot be moved when they have a rate limit configured on them.

Limiting traffic from a particular device

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

```
portcfgqos --setratelimit [ slot/] port ratelimit
```

Example of setting the rate limit on slot 3, port 9 to 4000 Mbps

```
portcfgqos --setratelimit 3/9 4000
```

Disabling Ingress Rate Limiting

1. Connect to the switch and log in using an account with admin permissions.

2. Enter the `portCfgQos --resetratelimit` command.

```
portcfgqos --resetratelimit [ slot/] port
```

Example of disabling Ingress Rate Limiting on slot 3, port 9

```
portcfgqos --resetratelimit 3/9
```

QoS

Quality of Service (QoS) allows you to categorize the traffic flow between a host and a target as having a high, medium, 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, then 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 shows a basic comparison between CS_CTL-based frame prioritization and QoS zone-based traffic prioritization. Refer to [CS_CTL-based frame prioritization](#) on page 331 and [QoS zone-based traffic prioritization](#) for detailed information about each type of prioritization scheme.

TABLE 83 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 CS_CTL field in frame header	Priority is defined by name of QoS zone
Prioritization is on a frame-basis	Prioritization is on a flow-basis
Setup steps: <ul style="list-style-type: none"> • Enable CS_CTL mode on F_Ports • Ensure that the CS_CTL mode-enabled host and storage are zoned together 	Setup steps: <ul style="list-style-type: none"> • Create QoS zones with host/target members • Add the QoS zones to the zone configuration • Save and then enable the zone configuration • Enable QoS on E_Ports

License requirements for QoS

In Fabric OS 7.2.0 and later, QoS does not require the Adaptive Networking license to be explicitly installed. This license is automatically enabled for new switches and for existing switches that are upgraded to Fabric OS 7.2.0 or later.

If you upgrade to Fabric OS 7.2.0 and you did not previously have an Adaptive Networking license, then all ports that had QoS mode set to AE (automatically enabled) and were not using the QoS feature are automatically set to OFF after the upgrade.


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 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 8-Gbps or higher ports. QoS is disabled by default on all long-distance ports.

FIGURE 25 QoS with E_Ports enabled



You must enable QoS on the E_Ports on both ISLs between domain 3 and domain 4, because either path might be selected to carry the 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, then the path through domain 2 would be used.

To guarantee traffic priority, you should enable QoS on all possible E_Ports. Alternatively, you could use a TI zone to limit the E_Ports that carry the traffic between a "host,target" pair and enable QoS on only those 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," then 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. Refer to [Using FC-FC Routing to Connect Fabrics](#) on page 535 for information about FC routers, phantom switches, and the FC-FC Routing Service.

To establish QoS over FC routers, you must perform the following tasks:

- Define QoS zones in each edge fabric.

- Define LSAN zones in each edge fabric.
- Enable QoS on the E_Ports in each edge fabric.
- Enable QoS on the EX_Ports in the backbone fabric.

Refer to [Setting QoS zone-based traffic prioritization over FC routers](#) on page 329 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 16-Gbps (Gen 5) and 32-Gbps (Gen 6) platforms only (Brocade 6505, 6510, 6520, G620, G630, M6505, 6547, 6548, Brocade DCX 8510 Backbones and X6 Directors), and only if no other platforms are used. 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 both EX_Ports and VEX_Ports.
- The EX_Ports (or VEX_Ports) in the path between the QoS devices must be on switches running Fabric OS v6.3.0 or later.
- QoS zones must use WWN notation only; D,I notation is not supported for QoS over FCRs.

vTap and QoS High Priority Zoning Compatibility mode

The vTap analytics feature utilizes VC 14 for remote flow mirroring, and 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 devices and/or the switches where vTap needs to be configured. You can use the **configureChassis** command to enable the vTap and QoS High Priority Zone Compatibility Mode.

```
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 (i.e., remote flow mirroring) are also active on the local switch, then vTap/QoSH compatibility cannot be disabled. vTap features have to be deactivated on all the logical switches in the chassis prior to disabling vTap/QoSH compatibility.

By default, "vTap and QoS High Priority Zone Compatibility Mode" is disabled. When disabled, the switch assigns QoS high priority zoning VCs in the usual manner.

NOTE

The chassis must be disabled in order to enable the new compatibility mode, but only when the effective zoning configuration has QOSH (default VC assignment) and/or QOSH5 zones.

The **configureChassis** command can be issued from any logical switch partition/context, but affects a chassis-wide scope regardless of switch context.

QoSH5 zones

QoSH/QoSH5 zones created from non-vTap/QoSH compatible switches, can be pushed to CM mode enabled switches as part of zone update. QoSH/QoSH5 zones added to the effective zoning configuration via a **configDownload** are permitted when vTap/QoSH compatibility is enabled. However, creation of new QoSH5 zones using the **zoneCreate** command, Webtools, or BNA version 14.0.1 or higher are not allowed when vTap/QoSH compatibility is enabled. This enforcement is only applied on and for the local switch. For example, if a QoSH5 zone create, rename, or update command is run in a switch that is *not* configured for vTap/QoSH compatibility, the zone update is *accepted* on those switches where vTap/QoSH compatibility is enabled.

NOTE

The presence or absence of QoSH5 zones does not have any effect on fabric build, regardless of vTap and QoSH compatibility settings.

The **zoneCreate** command does not allow creation of QoSH5 zones when vTap and QoSH compatibility is enabled and 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 QoSH5 zone when vTap/QoSH compatibility is enabled and 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 QoSH 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.
```

A best practice to avoid having to disable the chassis, you can perform the following steps:

1. Temporarily enable a zoning configuration that does not have any QOS high priority zones
2. Enable vTap/QoSH compatibility via **configureChassis** on all switches in the fabric
3. Re-enable a zoning configuration having the desired QOS high priority zones, with the recommendation to avoid QOSH5 zones with the new configuration

NOTE

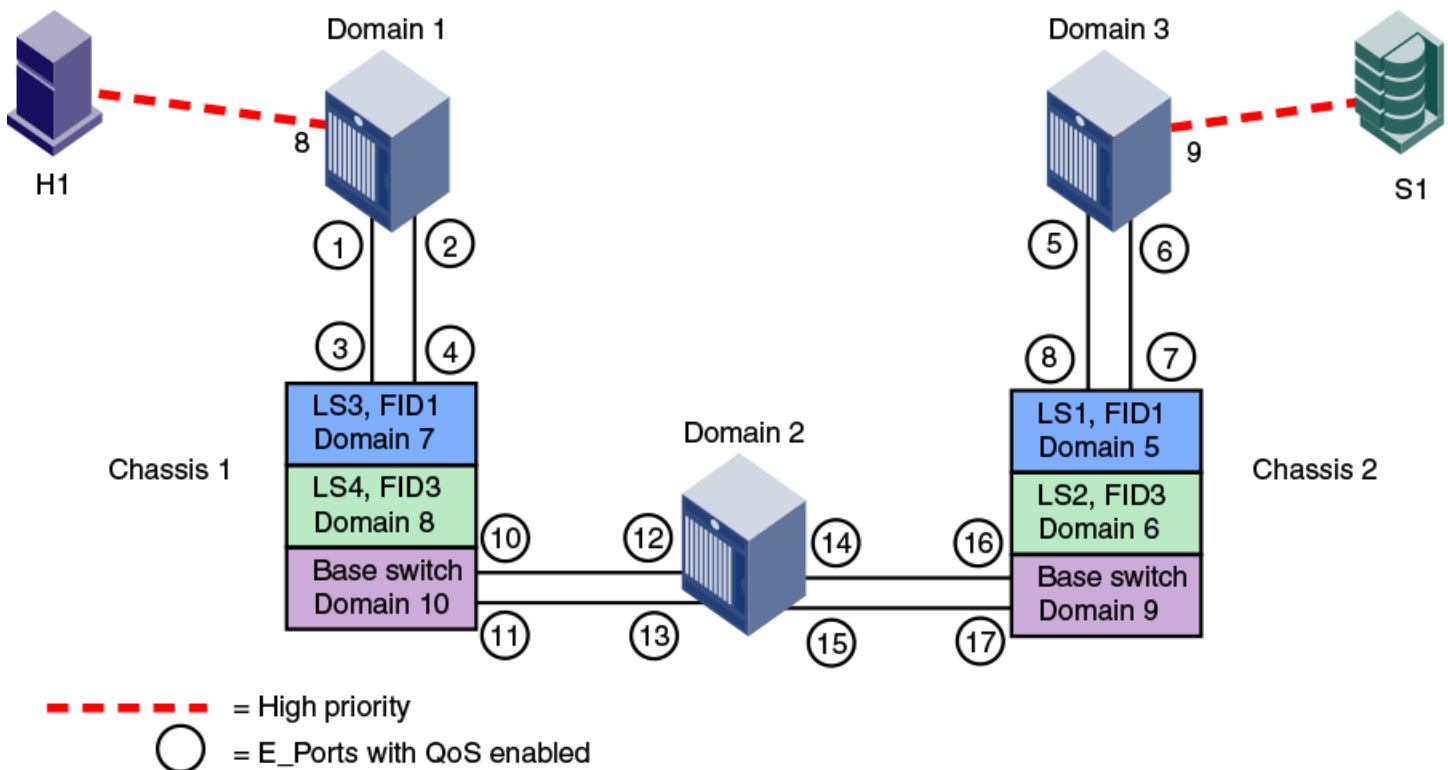
To ensure a non-disruptive Brocade Analytics Monitoring Platform attachment to a fabric, Brocade recommends that you enable QoSH and vTap+QoSH compatibility mode on all switches in the fabric. In addition, you must connect any devices requiring QoSH or QoSH5 services only through switches running versions of Fabric OS that both support this compatibility mode and have it enabled.

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 of the E_Ports shown circled in the figure, including all of the E_Ports in the XISLs (ports 10, 11, 12, 13, 14, 15, 16, and 17).

FIGURE 26 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 high, medium, or low quality of service (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 only belong to one QoS zone. By default, traffic in non-QoS zones is assigned medium priority.

QoS zones can contain WWN members (WWNN or WWPN) or *domain,index* (D,I) members. If you use D,I notation in your QoS zones, refer to [Limitations and restrictions for QoS zone-based traffic prioritization](#) on page 330 for some considerations.

A QoS zone has a special name format to differentiate it from a regular zone, which determines the priority of the traffic flow. The format of the QoS zone name is as follows, and *id* is a flow identifier that designates a specific virtual circuit (VC) for the traffic flow and *xxxxx* 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 virtual channels are allocated by means of a round-robin scheme.

TABLE 84 Name formats for priorities of QoS zones

Priority	Name Format
High	QOSHid_XXXX
Medium	QOSMid_XXXX
Low	QOSLid_XXXX

As examples, "QOSH3_HighPriorityTraffic" and "QOSL1_LowPriorityZone" are both valid QoS zone names.

Each priority level is allocated to different virtual channels (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:

- 1-5 for high-priority traffic, which corresponds to VCs 10-14.
- 1-4 for medium-priority traffic, which corresponds to VCs 2-5. Note, however, that the virtual channels for medium-priority traffic are always allocated by a round-robin scheme, regardless of the *id* value.
- 1-2 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 27 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

The following items should be kept 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 you upgrade to Fabric OS 7.2.0 or later from Fabric OS 7.1.x or earlier, and you did not previously have an Adaptive Networking license, then all ports that had QoS mode set to AE (automatically enabled) and were not using the QoS feature are automatically set to OFF after the upgrade. You must manually configure these ports for QoS.
- 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. Refer to [QoS over FC routers](#) on page 322 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 several medium-priority flows to that same DID, then it is possible that the medium-priority flows would 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

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **zoneCreate** command to create zones for high- and low-priority traffic.
 - For high-priority traffic, use the following syntax:

```
zonecreate "QOSH id_zonename", " member[; member...]"
```

- For low-priority traffic, use the following syntax:

```
zonecreate "QOSL id_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 by means of a round-robin scheme.

3. Enter the **cfgAdd** command to add the QoS zone to the zone configuration, by using the following syntax:

```
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"
```


- Enter the **portCfgQos** command to enable QoS on the E_Ports, by using the following syntax:

```
portcfgqos --enable [ slot/] port
```

The **portCfgQos** command does not affect QoS prioritization. It only enables or disables the link to pass 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 later use the **portCfgQos** command 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> zoneshow
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> cfigsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled. Until the Effective
configuration is re-enabled, merging new switches into the
fabric is not recommended and may cause unpredictable
results with the potential of mismatched Effective Zoning
configurations.
Do you want to save 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. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
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

- Connect to the switch in the edge fabric and log in using an account with admin permissions.
- Create QoS zones in the edge fabric.

The QoS zones must have WWN members only, and not D,I members. Refer to [Setting QoS zone-based traffic prioritization](#) on page 328 for instructions.

- Create LSAN zones in the edge fabric.

Refer to [Controlling device communication with the LSAN](#) on page 567 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

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:

- All switches in the fabric must be running Fabric OS v6.0.0 or later.

ATTENTION

If QoS traffic crosses an ISL for a switch running a firmware version earlier than Fabric OS v6.0.0, the frames are dropped.

- By default, all devices are assigned medium priority.
 - To be assigned high or low priority, hosts and targets must be connected to a Brocade 8-Gbps or 16-Gbps switch or port blade.
 - To preserve the priority level across ISLs, the switches must be running Fabric OS v6.0.0 or later.
- QoS is enabled by default on 8-Gbps and higher ports. QoS is disabled by default on all 4-Gbps ports and long-distance ports.

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 and medium zones = High priority
 - High and low zones = Low priority
 - Medium and low zones = Low priority
 - High, medium, and low 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.
- 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.
- Traffic prioritization is not supported on a CryptoTarget container (redirection zone). Refer to the *Brocade Fabric OS Encryption Administration Guide* for information about redirection zones.
- Traffic prioritization is not supported in McDATA Fabric Mode (interopmode 2) or Open Fabric Mode (interopmode 3).
- 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 8-Gbps ports in Extended Mode (LE). Refer to [Managing Long-Distance Fabrics](#) on page 373 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 can be used to assign 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 that 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 [Table 85](#). The values are enabled by default to ensure backward compatibility.

TABLE 85 Mapping of CS_CTL values to QoS priority for frame prioritization in CS_CTL default mode

CS_CTL value	Priority
1-8	Low
9-16	Medium
17-24	High

Alternatively, the user can apply CS_CTL auto mode. The CS_CTL auto mode uses only three CS_CTL values, as illustrated in [Table 86](#).

TABLE 86 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, refer to [Using CS_CTL auto mode at the chassis level](#) on page 333, and [Considerations for using CS_CTL-based frame prioritization](#) on page 333.

Supported configurations for CS_CTL-based frame prioritization

CS_CTL-based frame prioritization is supported on all 8-, 16-, and 32-Gbps platforms.

All switches in the fabric should be running Fabric OS v6.3.0 or later.

NOTE

If a switch is running a firmware version earlier than Fabric OS v6.3.0, the outgoing frames from that switch lose their priority.

High availability considerations for CS_CTL-based frame prioritization

If the standby CP is running a Fabric OS version earlier than 6.3.0 and is synchronized with the active CP, then you cannot enable CS_CTL-based frame prioritization on the active CP. If the standby CP is not synchronized or if no standby CP exists, then enabling CS_CTL-based frame prioritization succeeds.

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, so that the frames returned from the destination port for a given exchange always have the same CS_CTL prioritization as the frames originating from the source port.

1. Connect to the switch and log in to an account that has admin permissions.
2. Enable CS_CTL mode:

```
switch:admin> portcfgqos --enable 45 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.
Enabling CSCTL mode will deactivate sys_analytics_vtap on this logical switch.
Do you want to proceed?(y/n)
```

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 86](#) on page 332. Set **CSCTL QoS Mode** to **0** to disable auto mode and revert to default settings, shown in [Table 85](#) on page 331.

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 know 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, proceed with the following steps.

1. Determine whether to use the default mode (refer to [Table 85](#) on page 331) or the auto mode (refer to [Table 86](#) on page 332). No choice 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 `csctl_mode` enabled, as in [Enabling CS_CTL-based frame prioritization on ports](#) on page 332.

3. For a chassis based system, use the **slot poweroff – slot poweron** command pair for all the 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 config and stores the information you provide in the chassis-level persistent data base. When the ASIC is initialized, it retrieves this information to initialize its CSCTL-VC table. You must power cycle all the edge blades in the chassis to have the new information be registered by the ASIC. For fixed-port switches, this requires a reboot of the entire switch. To know 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

• Trunking overview.....	335
• Supported platforms for trunking.....	337
• Supported configurations for trunking.....	337
• Requirements for trunk groups.....	342
• Recommendations for trunk groups.....	343
• Configuring trunk groups.....	344
• Enabling trunking.....	344
• Disabling trunking.....	344
• Displaying trunking information.....	345
• ISL trunking over long-distance fabrics.....	346
• EX_Port trunking.....	347
• F_Port trunking.....	349
• Displaying F_Port trunking information.....	355
• Disabling F_Port trunking.....	356
• Enabling the DCC policy on a trunk area.....	356

Trunking overview

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 link within the trunk group fails.

Trunking is frame-based instead of exchange-based. Because a frame is much smaller than an exchange, this means that frame-based trunks are more granular and better balanced than exchange-based trunks and provide 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 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 DCX 8510 Backbones and/or X6 Directors and is applicable only to ports on the core blades.

Refer to [Inter-chassis Links](#) on page 357 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.

Refer to [EX_Port trunking](#) on page 347 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).

- 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). For more information, refer to [Configuring F_Port trunking for a Brocade adapter](#) on page 352, the *Brocade Access Gateway Administration Guide*, and the *Brocade Adapters Administrators Guide*.

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. This type of trunk might be referred to as N_Port trunking in the *Brocade Access Gateway Administration Guide* or *Brocade Adapters Administrator's Guide*.

Masterless trunking

Masterless trunking means that if the master port goes offline, one of the slave ports automatically becomes the new master port, thus avoiding traffic disruption. The new master port uses the old master port area and the old master port is assigned a new, unused area. In this way, the port identifier (PID) of the trunk does not change if the master port goes offline.

If trunking is not masterless, and if the master port goes offline, traffic disruption can occur because the slave ports in the trunk group go offline to select the new master port and then come back online.

Masterless trunking is supported for most platforms and trunking types:

- All F_Port trunking is masterless.
- ISL and ICL trunking are masterless.
- EX_Port trunking is masterless, except on Brocade DCX 8510 Backbones and Brocade X6 Directors with Virtual Fabrics disabled.

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 re-enable the switch.


Refer to the *Brocade Fabric OS Software Licensing 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 0-7, 8-15, 16-23, 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 28 Port group configuration for the Brocade G620



Supported platforms for trunking

Trunking is supported on the FC ports of all Brocade platforms and blades supported in Fabric OS v7.0.0 and later.

EX_Port trunking is supported only on those platforms that support EX_Ports. Refer to [Supported platforms for FC-FC routing](#) for more information.

Supported configurations for trunking

- Links within a trunk can be 4 Gbps, 8 Gbps, 10 Gbps, 16 Gbps, or 32 Gbps depending on the Brocade platform.

NOTE

32 Gbps 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. Refer to [ICL trunking between Brocade X6 Directors](#) on page 364 for ICL trunking detail.

- 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 128 Gbps (based on a 16-Gbps 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, no more than 30 meters difference is recommended. 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 master and the slave ports. That is, in a given trunk group, if there are some ports with QoS enabled and some with QoS 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 87 G630 QSFP and SFP trunking port groups

Trunk groups	Eight port trunks	Allowed member ports
SFP trunk group 1	Trunk 1	0-7
	Trunk 2	8-15
	Trunk 3	48-55
	Trunk 4	56-63
SFP trunk group 2	Trunk 5	16-23
	Trunk 6	24-31
	Trunk 7	32-39
	Trunk 8	40-47
SFP trunk group 3	Trunk 9	64-71
	Trunk 10	72-79
	Trunk 11	80-87
	Trunk 12	88-95
QSFP trunk group 1	Trunk 13	96-99 and 100-103
	Trunk 14	104-107 and 108-111
QSFP trunk group 2	Trunk 15	112-115 and 116-119
	Trunk 16	120-123 and 124-127


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.

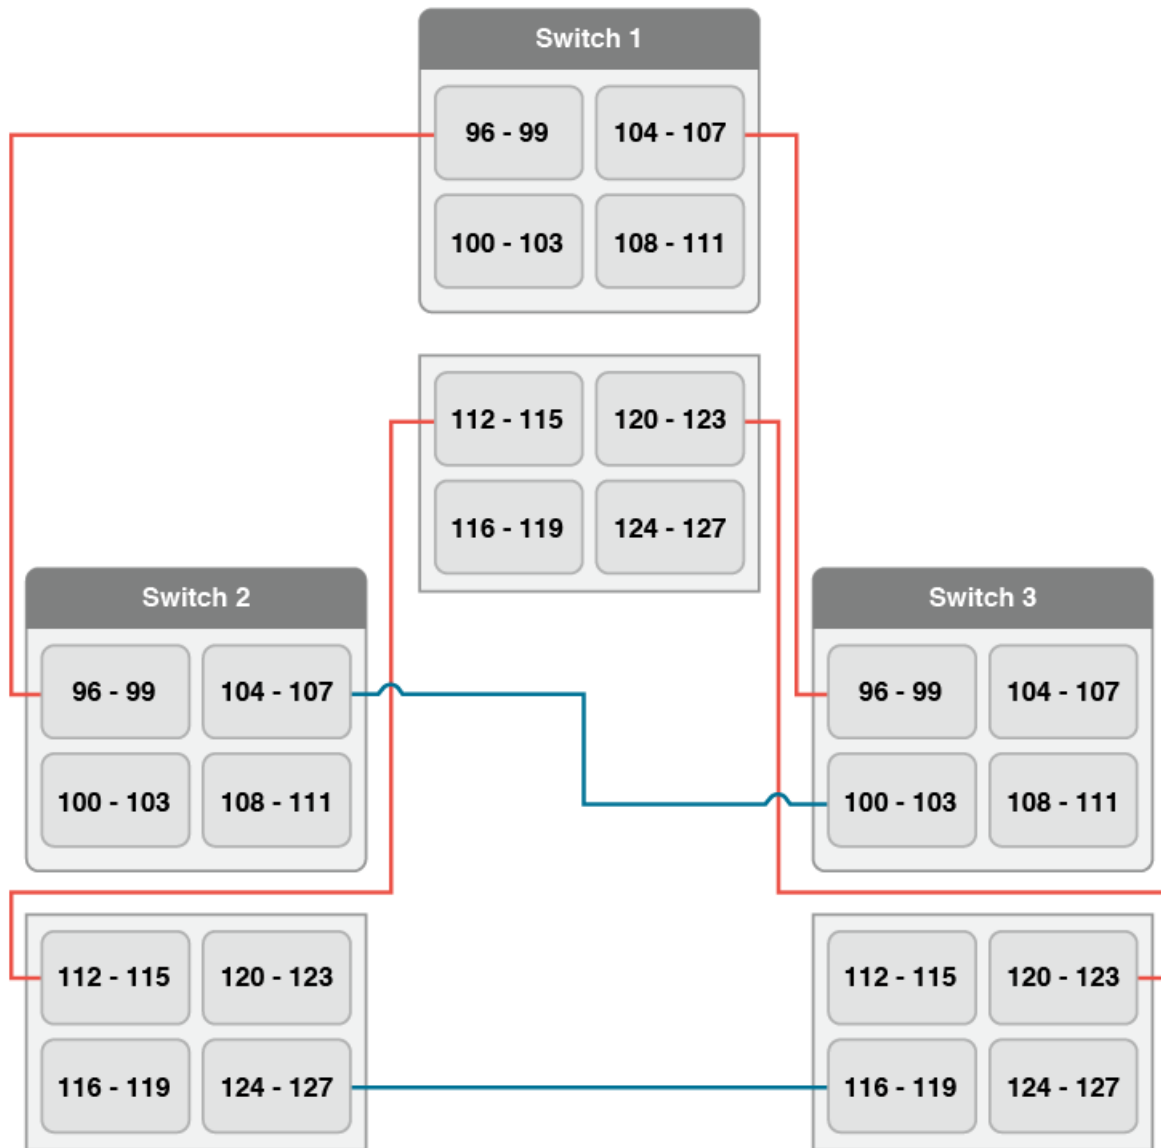
The QSFP ports in a Brocade G630 switch can be used to form ISLs between G630 switches freeing up SFP+ ports for device connectivity.

- 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 256 Gbps 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 port 96-99 (Group 1) and port 112-115 (Group2)
 - Connection 2 between switch 1 and switch 3 using port 104-107 (Group 1) and port 120-123 (Group 2)
 - Connection 3 between switch 2 and switch 3 using port 104-107 (Group 1) and port 124-127 (Group 2)
- QSFP to QSFP connection from a G630 switch to QSFP ports on the DCX 8510 or X6 Director core routing blades are not supported.




QSFP to SFP ISLs

QSFP ports in Brocade G630 can be used to form dense ISLs between G630 switches and Director or any other device SFP+ ports simplifying cabling.

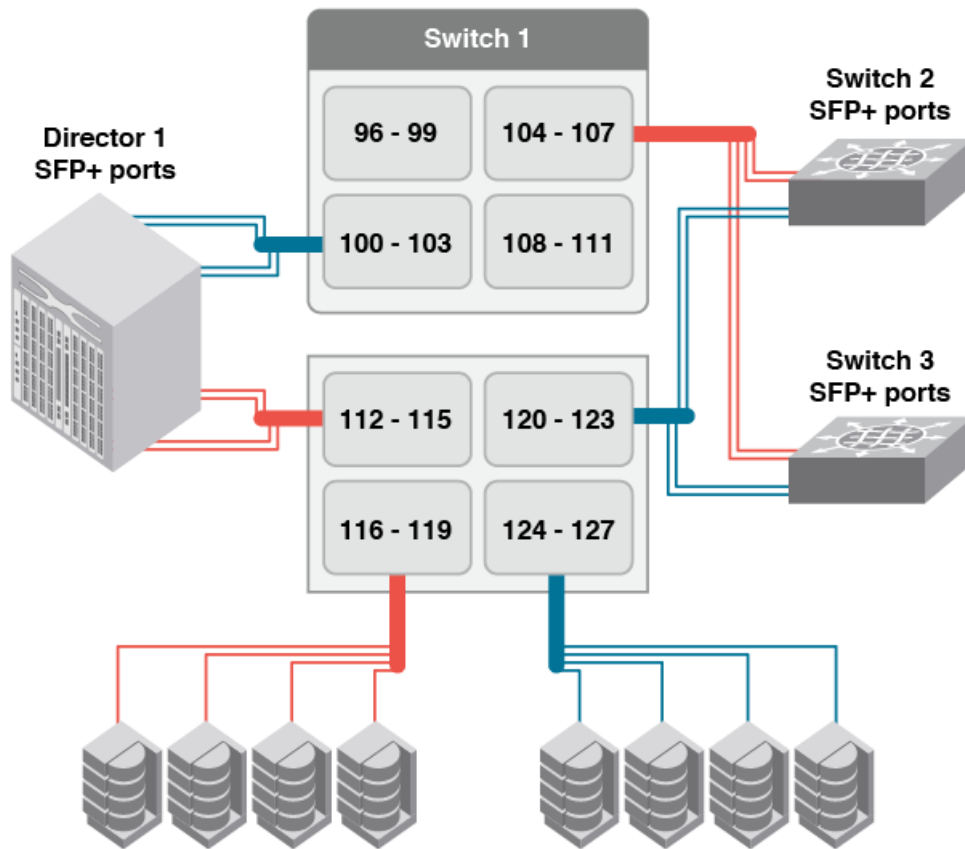
- For all QSFP to SFP connections, you must use only the Brocade qualified break-out cables.
- Up to eight links in 2 x QSFP ports can be trunked to get up to 256 Gbps bandwidth on ISLs.
- When G630 is used as an edge switch, the ISL connections to the Directors should be made observing the trunk boundaries
- Brocade recommends that you buy QSFP POD licenses for the Brocade G630 switch up front and use them for connections to storage ports (F_Ports) for a fully populated G630 switch.
- Brocade recommends that you 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 100-103 (Group 1) and ports 112-115 (Group 2)
 - Connection 2 between switch 1 and switch 2 using ports 104-105 (Group 1) and ports 120-121 (Group 2)
 - Connection 3 between switch 1 and switch 3 using ports 106-107 (Group 1) and ports 122-123 (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

Brocade recommends that you use the SFP ports for connecting hosts and storage devices but 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).
- There must be a direct connection between participating switches.
- Trunking cannot be done 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 or VEX_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 be used to attach other switches or devices.
- Each physical ISL uses two ports that could otherwise be used to 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 master ISL of the group, regardless of the number of ISLs in the group. This allows slave 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.

This helps 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 re-initialize the ports that are to be used in trunk groups so that they recognize that trunking is enabled. This procedure needs to be performed only once, and is required for all types of trunking.

To re-initialize the ports, you can either disable and then re-enable the switch, or disable and then re-enable the affected ports.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **islShow** command to determine which ports are used for ISLs.

```
switch:admin> islshow
1: 0-> 0 10:00:00:05:1e:a2:f9:b1 3 SW_43 sp: 4.000G bw: 4.000G TRUNK
QOS
2: 10-> 1 10:00:00:05:33:88:a1:d0 1 bodc_switch_sw_66_rename_limit sp: 16.000G bw: 16.000G TRUNK
QOS CR_RECOV FEC
3: 12-> 2 10:00:00:05:33:88:a1:d0 1 bodc_switch_sw_66_rename_limit sp: 8.000G bw: 8.000G TRUNK
QOS CR_RECOV
4: 16-> 0 10:00:00:05:1e:a1:99:09 5 bodc_switch_20_chars sp: 8.000G bw: 8.000G QOS
```

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 of 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](#) on page 351 or [Configuring F_Port trunking for a Brocade adapter](#) on page 352 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 re-enables 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 re-enables 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

You can use the **trunkShow** command to view the following information:

- All the trunks and members of a trunk.
- Whether the trunking port connection is the master port connection for the trunk group.
- Whether trunks are formed correctly.
- Trunking information for a switch that is part of an FC router backbone fabric interlinking several edge fabrics.
- Trunking information, including bandwidth and throughput for all the 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 need to adjust the fabric design by adding, removing, or reconfiguring ISLs and trunking groups. For additional information on traffic monitoring, refer to *Brocade Flow Vision Administrator's Guide*.

To view detailed information about F_Port trunking, refer to [Displaying F_Port trunking information](#) on page 355.

Use the following procedure to view trunking information:

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

The following example shows trunking groups 1, 2, and 3; ports 4, 13, and 14 are masters.

```
switch:admin> trunkshow
1: 6-> 4 10:00:00:60:69:51:43:04 99 deskew 15 MASTER
2: 15-> 13 10:00:00:60:69:51:43:04 99 deskew 16 MASTER
    12-> 12 10:00:00:60:69:51:43:04 99 deskew 15
    14-> 14 10:00:00:60:69:51:43:04 99 deskew 17
    13-> 15 10:00:00:60:69:51:43:04 99 deskew 16
3: 24-> 14 10:00:00:60:69:51:42:dd 2 deskew 15 MASTER
```

NOTE

The minimum deskew value for a Gen 6 platform is 1. The minimum deskew value for a Gen 5 or earlier platform is 15.

The following example shows trunking information along with the bandwidth and throughput for all the trunk groups in a switch.

```
switch:admin> trunkshow -perf

1: 2-> 2 10:00:00:05:1e:81:56:8b 1 deskew 15 MASTER
    3-> 3 10:00:00:05:1e:81:56:8b 1 deskew 17
    Tx: Bandwidth 4.00Gbps, Throughput 1.66Gbps (48.45%)
    Rx: Bandwidth 4.00Gbps, Throughput 1.66Gbps (48.44%)
    Tx+Rx: Bandwidth 8.00Gbps, Throughput 3.33Gbps (48.44%)
2: 5->113 10:00:00:05:1e:46:42:01 3 deskew 15 MASTER
    4->112 10:00:00:05:1e:46:42:01 3 deskew 15
    Tx: Bandwidth 16.00Gbps, Throughput 1.67Gbps (12.12%)
    Rx: Bandwidth 16.00Gbps, Throughput 1.67Gbps (12.12%)
    Tx+Rx: Bandwidth 32.00Gbps, Throughput 3.33Gbps (12.12%)
3: 10-> 10 10:00:00:05:1e:81:56:8b 1 deskew 15 MASTER
    11-> 11 10:00:00:05:1e:81:56:8b 1 deskew 15
    Tx: Bandwidth 4.00Gbps, Throughput 1.66Gbps (48.45%)
    Rx: Bandwidth 4.00Gbps, Throughput 1.67Gbps (48.48%)
    Tx+Rx: Bandwidth 8.00Gbps, Throughput 3.33Gbps (48.46%)
4: 12->892 10:00:00:05:1e:46:42:01 3 deskew 15 MASTER
    13->893 10:00:00:05:1e:46:42:01 3 deskew 15
    Tx: Bandwidth 16.00Gbps, Throughput 1.67Gbps (12.12%)
    Rx: Bandwidth 16.00Gbps, Throughput 1.66Gbps (12.11%)
    Tx+Rx: Bandwidth 32.00Gbps, Throughput 3.33Gbps (12.11%)
```

ISL trunking over long-distance fabrics

In long-distance fabrics, if a port speed is set to autonegotiate, then the maximum speed is assumed for reserving buffers for the port. If the port is running at only 2 Gbps, this wastes buffers. For long-distance ports, you should specify the port speed instead of setting it to autonegotiate.

In addition to the criteria listed in [Supported configurations for trunking](#) on page 337, observe the following criteria for trunking over long-distance fabrics:

- Trunking over long-distance fabrics is supported only on switches running Fabric OS 6.1.0 and later.
- 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, refer to [Configuring an extended ISL](#) on page 374.

[Table 88](#) summarizes support for trunking over long distance for the Brocade DCX and Brocade DCX 8510 Backbones and supported blades.

TABLE 88 Trunking over long distance for the Brocade Backbones and blades

Long-distance mode	Distance	Number of 2-Gbps ports	Number of 4-Gbps ports
LE	10 km	48 (six 8-port trunks)	48 (six 8-port trunks)

TABLE 88 Trunking over long distance for the Brocade Backbones and blades (continued)

Long-distance mode	Distance	Number of 2-Gbps ports	Number of 4-Gbps ports
LO	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 LO mode supports up to 5 km at 2 Gbps, up to 2 km at 4 Gbps, and up to 1 km at 8/10 Gbps, 625 m at 16 Gbps and 150 m at 32 Gbps. 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. For more information about setting port speeds, refer to [Performing Advanced Configuration Tasks](#) on page 71.

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 master port. The other ports that come online on the trunk group are considered to be the slave ports. Adding or removing a slave port does not cause frame drop. However, removing a slave port causes the loss of frames in transit.

If router port cost is used with EX_Port trunking, the master port and slave ports share the router port cost of the master 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.

ATTENTION

EX_Port trunking should be enabled only if the entire configuration is running Fabric OS v5.2.0 or later.

Refer to [Using FC-FC Routing to Connect Fabrics](#) on page 535 for more information about EX_Ports and the FC router.

Masterless EX_Port trunking

EX_Port trunking is masterless except for EX_Ports on Brocade DCX and Brocade DCX 8510 Backbones.

For the Backbones, Virtual Fabrics must be enabled for masterless EX_Port trunking to take effect. For the fixed-port switches, Virtual Fabrics can be enabled or disabled.

If masterless EX_Port trunking is not in effect and the master port goes offline, the entire EX_Port-based trunk re-forms and is taken offline for a short period of time. If there are no other links to the edge fabric from the Backbone, the master port going offline may cause a traffic disruption in the backbone fabric.

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 2 Gbps 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.

Masterless EX_Port trunking has additional configuration requirements. Refer to [Masterless EX_Port trunking](#) on page 347 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, then two trunk groups will form: one with QoS enabled and one with QoS disabled.

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. Refer to [Configuring trunk groups](#) on page 344 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 master EX_Port and a slave 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 10:00:00:69:80:1d: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 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 8-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 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. Do 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). Doing so will cause the corresponding N_ports on the Access Gateway module to be disabled.

Refer to the *Brocade Access Gateway Administration Guide* for information about configuring the corresponding N_Port trunking on the Access Gateway.

F_Port trunking for Access Gateway

You can configure trunking between the F_Ports on an edge switch and the N_Ports on an Access Gateway 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 Access Gateway (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.

[Figure 29](#) shows a switch in AG mode without F_Port masterless trunking. [Figure 30](#) shows a switch in AG mode with F_Port masterless trunking.

FIGURE 29 Switch in Access Gateway mode without F_Port masterless trunking

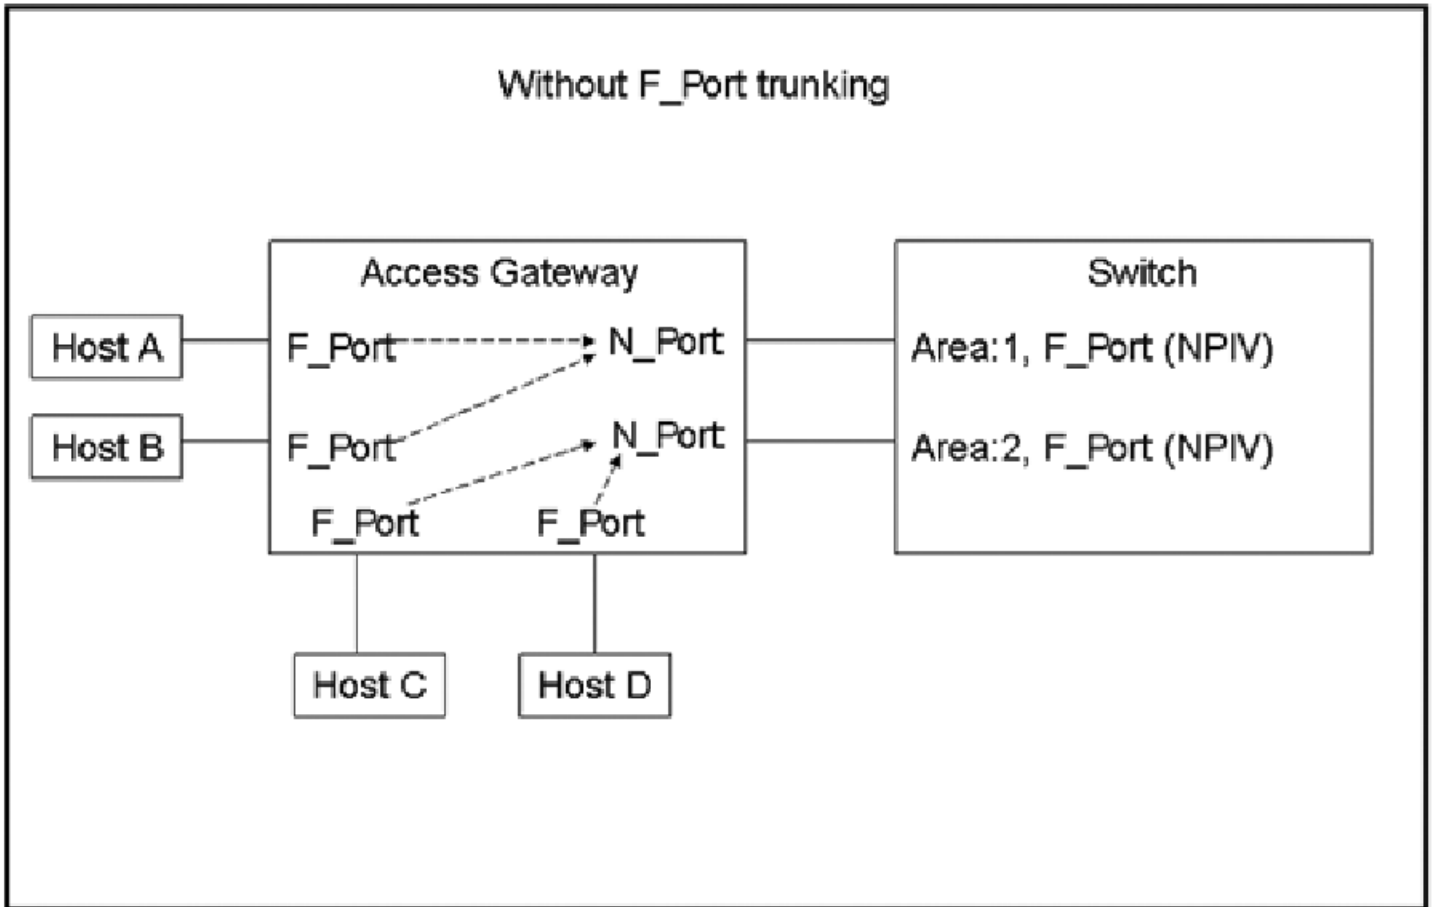
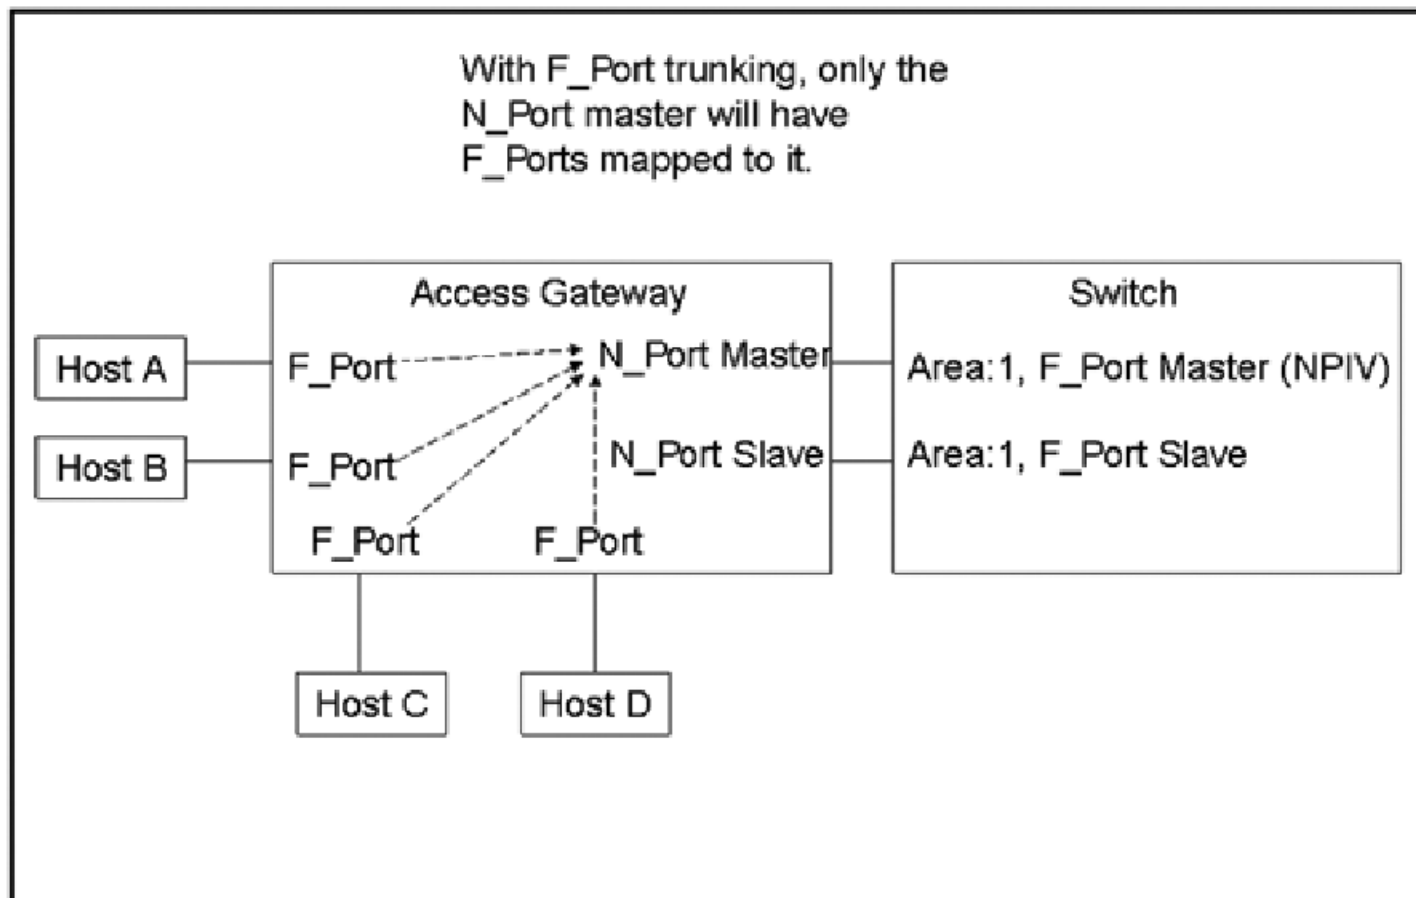


FIGURE 30 Switch in Access Gateway mode with F_Port masterless trunking

**NOTE**

You do not need to map the host to the master port manually because the Access Gateway will perform a cold failover to the master port.

Refer to [Configuring F_Port trunking for an Access Gateway](#) on page 351 for instructions on configuring F_Port trunking.

Requirements for F_Port trunking on an Access Gateway

In addition to the requirements listed in [Requirements for trunk groups](#) on page 342, refer to the Brocade Access Gateway Administration 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 TA.
4. Enter the **portTrunkArea --enable** command to enable the trunk area.

For example, the following command creates a TA for ports 36-39 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 re-enable 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](#) on page 342, 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.

Refer to the *Brocade Adapters Administrator's Guide* for information about configuring the corresponding N_Port trunking on the Access Gateway and the Brocade adapter.

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. Refer to the *Brocade Adapters Administrator's Guide* for a detailed description and requirements of N_Port trunking on the adapters.

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 trunking as described in the *Brocade Adapters Administrator's Guide*.

- 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 masterless trunking considerations.

TABLE 89 F_Port masterless trunking considerations

Category	Description
Area assignment	<p>You statically assign the area within the trunk group on the edge switch. That group is the F_Port trunk.</p> <p>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.</p> <p>10-bit addressing is the default mode for all dynamically created partitions in the Brocade DCX and Brocade DCX 8510-8 platforms.</p>
Authentication	<p>Authentication occurs only on the F_Port trunk master port and only once per the entire trunk. This behavior is the same as E_Port trunk master 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 portShow command.</p> <p>NOTE Switches in Access Gateway (AG) mode support authentication in F_Ports and N_Ports.</p>
configdownload	<p>If you issue the configDownload 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 configDownload.</p> <p>NOTE Long distance, port mirroring, non-CORE_PID, and FastWrite are not compatible with F_Port trunking.</p>
domain,index (D,I)	<p>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 <i>domain,index</i> (D,I), which refers to an index that might have been removed, will no longer be part of the switch.</p> <p>NOTE Be sure to include zoning and DCC when creating a Trunk Area.</p> <p>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.</p> <p>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.</p>
DCC Policy	<p>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 master. 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 master. However, the PWWN of the FDISC will continue to go through DCC policy check.</p>
Default Area	<p>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.</p>
Downgrade	<p>You can have trunking on, but you must disable the trunk ports before performing a firmware downgrade.</p> <p>ATTENTION Removing a Trunk Area on ports running traffic is disruptive because you must disable the port to disable the Trunk Area on the port. Use caution before assigning a Trunk Area if you need to downgrade to a firmware version earlier than Fabric OS 6.2.0.</p>
FastWrite	<p>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.</p>

TABLE 89 F_Port masterless trunking considerations (continued)

Category	Description
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.
HA Sync	If you plug in a standby CP with a firmware version earlier than Fabric OS 6.2.0 and a Trunk Area is present on the switch, the CP blades will become out of sync.
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) ELSs are not supported on F_Port trunks.
NPIV	N_Port ID virtualization (NPIV) is supported on the F_Port master 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. 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, then 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 PWWN is the same across the entire F_Port trunk that has either 0x2f or 0x25 as the first byte of the PWWN. The Trunk Area is part of the PWWN in the format listed in PWWN formats for F_Port trunk ports .
QoS	QoS is supported.
Routing	Routing will route against the F_Port trunk master. Bandwidth information will be modified accordingly as the F_Port trunk forms.
Trunk Master	No more than one trunk master is allowed in a trunk group. The second trunk master will be persistently disabled with the reason "Area has been acquired".
Upgrade	There are no limitations on upgrading to Fabric OS 7.0.0 and later if the F_Port is present on the switch. Upgrading is not disruptive.

PWWN formats for F_Port trunk ports

The following table describes the PWWN formats for F_Port trunk ports. The initial 5 is the default value for an F_Port trunk port in a virtual fabric. The initial 2 is the default value for an F_Port trunk port that is not in a virtual fabric.

TABLE 90 PWWN format for F_Port trunk ports

Fabric type	F_Port trunk port format	Valid range of xx
Virtual fabric (NAA = 5)	5n:nn:nn:nn:nn:n xx	0-4095
Non-virtual fabric (NAA = 2)	25: xx :nn:nn:nn:nn:nn	0-255
	2F: xx :nn:nn:nn:nn:nn	0-255

Refer to [Logical ports](#) on page 394 for more information on PWWN construction.

F_Port trunking in Virtual Fabrics

F_Port trunking functionality performs the same in Virtual Fabrics as it does in non-Virtual Fabrics platforms except for the Brocade DCX 8510-8. Fabric OS uses a 10-bit addressing model.

On the Brocade DCX 8510 platforms and Brocade X6 Directors, 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.

NOTE

Brocade DCX 8510-8 platforms have a maximum of 576 ports. Out of the 1024 10-bit address range, addresses 448–1023 are reserved for the 10-bit address space. Addresses 0–447 are reserved for assigning to NPIV/Loop ports to support 112 (448/4) NPIV/Loop ports in a logical switch with 256 devices each.

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.
- If you swap a port by using the **portSwap** command, you must undo the port swap 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](#) on page 382 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 deskew 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 deskew 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 deskew 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 deskew 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.

In order 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 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](#) on page 233.

Inter-chassis Links

• Inter-chassis links	357
• ICLs between DCX 8510 Backbones.....	359
• ICLs between Brocade X6 Directors.....	361
• ICLs between Brocade DCX 8510 Backbones and Brocade X6 Directors.....	366
• Virtual Fabrics considerations for ICLs.....	368
• Supported topologies for ICL connections.....	369

Inter-chassis links

An inter-chassis link (ICL) is a licensed feature used to interconnect the following chassis platforms. ICL ports in the core blades are used to interconnect the chassis platforms, potentially increasing the number of usable ports in the chassis platforms.

- Two Brocade DCX 8510 Backbones
- Two Brocade X6 Directors
- A Brocade DCX 8510 Backbone with a Brocade X6 Director

NOTE

To connect a DCX 8510 Backbone to an X6 Director, you must fix the port speed on the X6 Director to 16 Gbps and use a 16 Gbps QSFP.

When two chassis platforms are interconnected by ICLs, each chassis requires a unique domain and is managed as a separate switch.

The ICL ports appear as regular ports, with some restrictions. All port parameters associated with 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 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. Refer to [Using FC-FC Routing to Connect Fabrics](#) on page 535 for instructions.

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 meanings and ICL connections, including instructions on how to cable ICLs.

License requirements for ICLs

ICL ports can be used only with an ICL license. An ICL license must be installed on both platforms forming the ICL connection.

All ICL ports must be disabled and then re-enabled for the license to take effect. After the addition or removal of an ICL license, the license enforcement is performed on the ICL ports only when you issue the **portDisable** and **portEnable** commands on the switch for the ports or the **bladeDisable** and **bladeEnable** commands for the core blade.

Enterprise ICL (EICL) license enforcement is applicable for Inter Fabric Links (IFL) via ICL links on Gen 5 platforms and is not applicable for Gen 6 platforms. In a mixed Gen 5 and Gen 6 environment, the edge fabric will additionally count the ICL IFL links for its license enforcement. When a Gen 5 chassis is connected to a Gen 6 chassis, the Gen 5 takes precedence and imposes the EICL license limit on the links. Prior to 8.0.1 release, only ICL ISLs were counted for license enforcement. In case of only Gen 6 platforms, there is no restriction on the number of switches that can be connected via ICL. In case of only Gen 5 platforms, if no EICL license is present, only 3 other direct chassis are allowed to connect (to a total of 4).

In the case of 16-Gbps Gen 5 platforms, if the EICL license is present, a maximum of 9 chassis are allowed to connect (to a total of 10). The limit on the number of chassis depends only on the physical chassis and not on the logical switches.

For more information on how license enforcement occurs, refer to the *Brocade Fabric OS Software Licensing Guide*.

Using the QSFPs that support 2 km on ICL ports

Prior to Fabric OS 7.3.0, all the FE ports and ICL ports used the same buffer credit model. In Fabric OS 7.3.0 and later, 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.

The following points should be considered if you are attempting to support 2 km links on ICL ports:

- Only the QSFPs that have the part number “57-1000310-01” and the serial number starting with “HME114323P00044” support 2 km on ICLs. The second character (‘M’) in this serial number indicates that the QSFP supports 2 km distance. You can also use the **sfpShow** command to identify the QSFPs that support 2 km on ICL ports.
- The new credit model does not affect the HA configuration.
- You cannot downgrade from Fabric OS 7.3.0 to any earlier version if either of the following conditions is true:
 - When you have plugged in the QSFPs that support 2 km on one or more ICL ports.
 - When you have configured buffer credits using the **EportCredit** command on one or more ICL ports.
- The **portCfgEportCredits** configuration cannot have less than 5 buffer credits or more than 16 buffer credits per VC in Gen 5 core blades. Gen 6 core blades cannot have less than 5 buffer credits and more than 160 buffer credits per VC. If there are insufficient buffer credits available, the default configuration is retained and the message *Failed: already exceeds Buffer credits allowed* is displayed.
- Only a maximum of 10 ICL ports can be configured with 2 km QSFPs with 16 buffer credits per VC on Gen 5 core blades.
- Only a maximum of 14 ICL ports can be configured with 2 km QSFPs with 13 buffer credits per VC on Gen 5 core blades.

On the Gen 5 core blades, when you configure the 15th port using **portCfgEportCredit**, the message *Failed: already exceeds Buffer credits allowed* is displayed. The remaining two ports can have only 34 buffer credits. To remedy this, disable some ports and configure the remaining ports using 13 buffer credits per VC.
- On X6-8 core blades, only a maximum of 8 ports can be configured for 2 km link with all frame sizes at line rate. The remaining 8 ports can be configured for 2 km link with full size frames at line rate.
 - For each port that requires 2 km line rate for all frame sizes, assign 96 buffers using the **portCfgEportCredits** command.
 - For each port that requires 2 km line rate for FULL frame sizes, assign 37 buffers using the **portCfgEportCredits** command.
- On X6-4 core blades, only a maximum of 8 ports can be configured for 2 km link with all frame sizes at line rate. The remaining 8 ports can be configured for 2 km link with full size frames at line rate.
 - For each port that requires 2 km line rate for all frame sizes, assign 48 buffers using the **portCfgEportCredits** command.
 - For each port that requires 2 km line rate for FULL frame sizes, assign 18 buffers using the **portCfgEportCredits** command.

- On the Gen 5 core blades, if you are using the `portCfgEportCredit` command, a maximum of 16 buffer credits can be configured on all the ICL ports when all the 16 ICL ports are in the disabled state. After enabling all the ports, until the buffer credits are available, the ports will come up with the configured buffer credits. The remaining ports will come up in degraded mode. In case of remaining QoS-enabled ports, the ports will come up without QoS enabled. If all the ICL ports are QoS enabled, there will only be 448 buffer credits available for 2 km distance support.
- Due to the 2 km QSFP module limitation, the link failure counter is not reliable during module or cable removal or insertion.

ICLs between DCX 8510 Backbones

Each ICL connects the core blades of two Brocade DCX 8510 chassis and provides up to 64 Gbps of throughput within a single cable.

You can have up to 32 QSFP ports in a Brocade DCX 8510-8 chassis or 16 QSFP ports in a Brocade DCX 8510-4 chassis, with up to 2 Tbps ICL bandwidth and support for up to 100 meters on universal optical cables.

The 100-meter ICL is supported beginning in Fabric OS 7.1.0, when using 100-meter-capable QSFPs over OM4 cable only.

The Brocade DCX 8510-8 has four port groups on the CR16-8 core blade. The Brocade DCX 8510-4 has two port groups on the CR16-4 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.

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 91 External port to slotShow port mapping for core blades

External port number	slotShow FC port numbers	External port number	slotShow FC port numbers
0	0-3	8	32-35
1	4-7	9	36-39
2	8-11	10	40-43
3	12-15	11	44-47
4	16-19	12	48-51
5	20-23	13	52-55
6	24-27	14	56-59
7	28-31	15	60-63

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 two Brocade DCX 8510 chassis is four. Additional ICLs should be added in increments of two, one on each core blade.
- For High Availability, you should have at least two ICLs from each core blade. [Figure 31](#) shows two Brocade DCX 8510-8 chassis connected with full redundancy using four ICL connections.

FIGURE 31 Minimum configuration for 64 Gbps ICLs



- The maximum number of ICLs between two Brocade DCX 8510-4 chassis or between a Brocade DCX 8510-8 and a Brocade DCX 8510-4 is 16.

The maximum number of ICLs between two Brocade DCX 8510-8 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

Brocade recommends that you 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 either core blade in the neighboring chassis. Unlike the copper ICLs, the QSFP ICLs do not need to be cross-connected.

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 ICLs and ISLs (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 DCX 8510 Backbones

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 Guide*.

NOTE

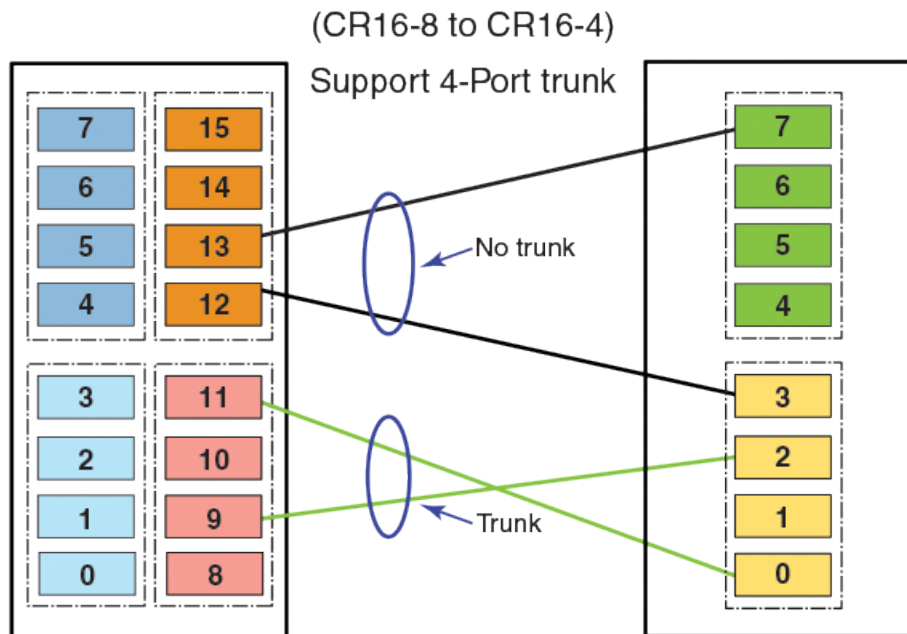
Each QSFP cable has four ports, each terminating on a different ASIC. These ports cannot form a trunk with each other, but can form trunks only with corresponding ports on another QSFP.

Each core blade in the Brocade DCX 8510-8 contains 16 ICL trunk groups. Each core blade in the Brocade DCX 8510-4 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 DCX 8510 Backbone family, the QSFP cables must be in the same trunk group.

The following illustration identifies the 4-port ICL trunking boundaries or port groups for DCX 8510 core routing blades (CR16-8 and CR16-4 blades):

FIGURE 32 ICL trunking boundaries



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 tables summarize the possible trunking connections between DCX 8510 Backbones:

TABLE 92 Trunking connections between DCX 8510 Backbones

Connecting blades	2 to 4-port ICL trunks
CR16-8 and CR16-8	Supported with QSFPs located within the same trunk group on each blade
CR16-8 and CR16-4	Supported with QSFPs located within the same trunk group on each blade
CR16-4 and CR16-4	Supported with QSFPs located within the same trunk group on each blade

Refer to the specific hardware reference manuals for information about port numbering, port trunk groups, and connecting the ICL cables.

ICLs between Brocade X6 Directors

Up to eight directors can be connected through inter-chassis links (ICLs) between 4x32-Gbps QSFPs on core routing blades installed in X6 Directors

Each ICL connects the core blades of two Brocade X6 Director chassis and provides up to 128 Gbps of throughput within a single cable.

You can have up to 32 QSFP ports in a X6-8 chassis or 16 QSFP ports in a X6-4 chassis, with up to 4 Tbps ICL bandwidth and support for up to 100 meters on universal optical cables.

The 100-meter ICL is supported beginning in Fabric OS 8.0.1, when using 100-meter-capable QSFPs over OM4 cable only.

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 93 External port to slotShow port mapping for core blades

External port number	slotShow FC port numbers	External port number	slotShow FC port numbers
0	0-3	8	32-35
1	4-7	9	36-39
2	8-11	10	40-43
3	12-15	11	44-47
4	16-19	12	48-51
5	20-23	13	52-55
6	24-27	14	56-59
7	28-31	15	60-63

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, additional ICL connections should be added in pairs, one on each core blade with each additional pair residing within the same trunk boundary
- For High Availability, you should have at least two ICLs from each core blade. The following illustration shows two X6-8 chassis connected with full redundancy using four ICL connections.

FIGURE 33 Minimum configuration for 128 Gbps 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 a X6-8 and a 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

Brocade recommends that you 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 either core blade in the neighboring chassis. The QSFP ICLs do not need to be cross-connected.

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 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 32 Gbps Fibre Channel ports. Since 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 has to be formed from individual FC ports in different QSFP ports. These QSFP ports 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 Guide*.

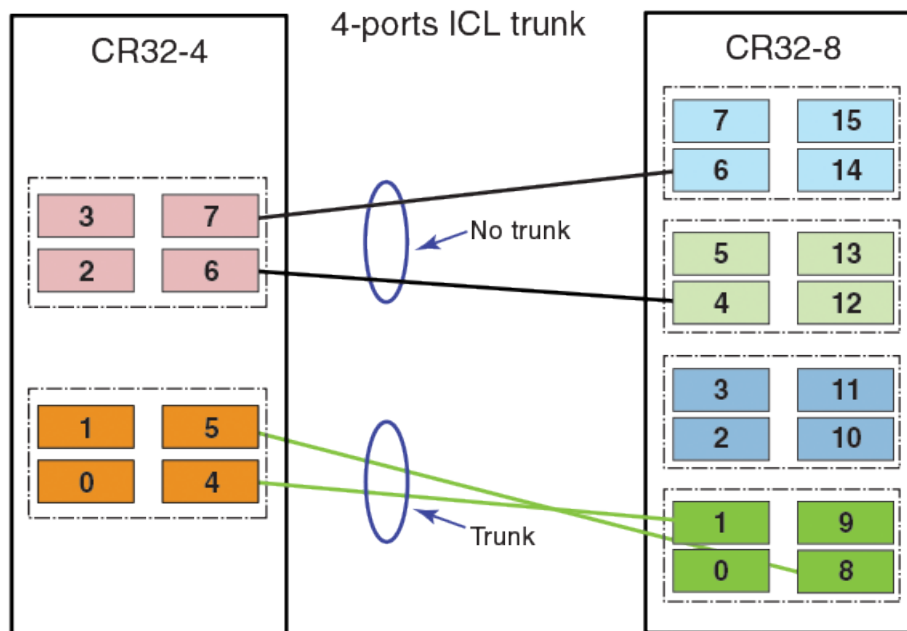
NOTE

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 only with corresponding ports on another QSFP.

Each core blade in the Brocade X6-8 contains 16 ICL trunk groups. Each core blade in a Brocade X6-4 Director 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 X6 Director family, the QSFP cables must be in the same trunk group. Up to four 4x32 Gbps QSFP ports in a trunk group can form a 512 Gbps 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 34 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 are 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 tables summarize the possible trunking connections between Brocade X6 Directors:

TABLE 94 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.

ICLs between Brocade DCX 8510 Backbones and Brocade X6 Directors

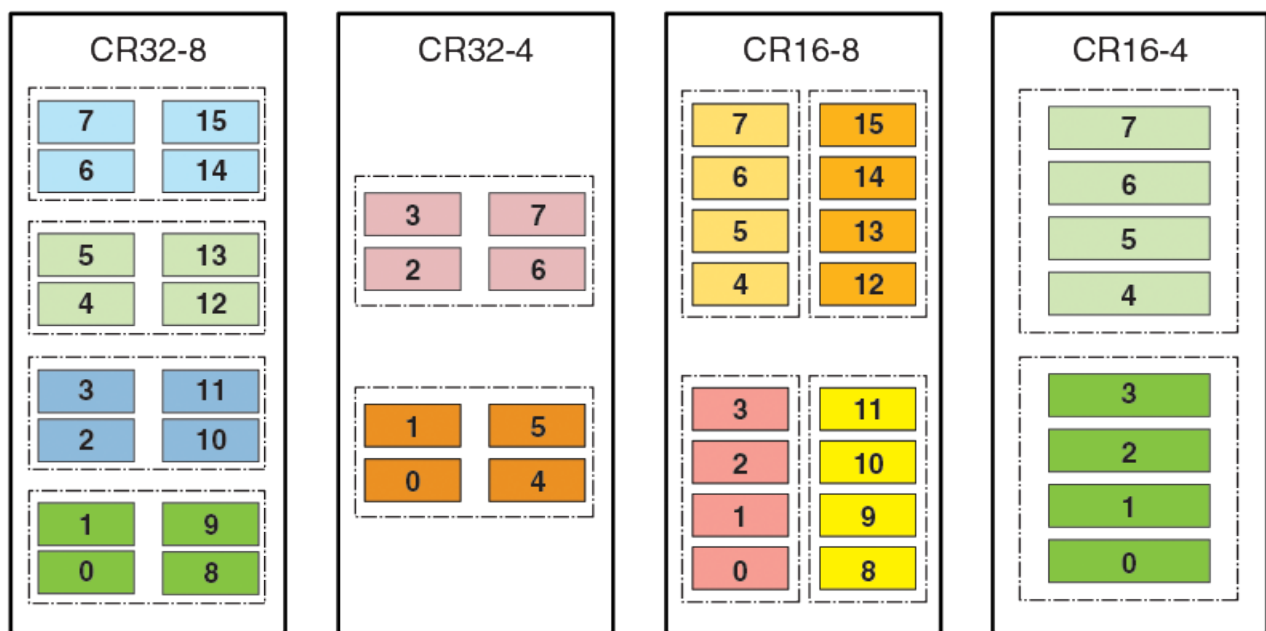
Each ICL connects the core blade of a Brocade DCX 8510 chassis with an X6 Director chassis and provides up to 64 Gbps of throughput within a single cable. You must set the port speed to 16 Gbps and use the 16 Gbps QSFP to interconnect on the X6 Director.

You can have up to 32 QSFP ports in a Brocade DCX 8510-8/Brocade X6-8 chassis, and 16 QSFP ports in a Brocade DCX 8510-4/Brocade X6-4 chassis, with up to 2 Tbps ICL bandwidth and support for up to 100 meters on universal optical cables.

The 100-meter ICL is supported beginning in Fabric OS 8.0.1, when using 100-meter-capable QSFPs over OM4 cable only.

The Brocade DCX 8510-8/Brocade X6-8 has four port groups on the CR16-8/CR32-8 core blade. The Brocade DCX 8510-4/Brocade X6-4 has two port groups on the CR16-4/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 5 vs Gen 6 CR blades, as represented in the following illustration. Each four ports that have the same color representation form a port group.

FIGURE 35 Port groups in Gen 5 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 DCX 8510 chassis and an X6 Director chassis is four. Additional ICLs should be added 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/X6-4 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/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

Brocade recommends that you 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 either core blade in the neighboring chassis. Unlike the copper ICLs, the QSFP ICLs do not need to be cross-connected.

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 16 Gbps 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 Backbone/X6 Director chassis if the ICLs are in a different logical switch than the ISLs.

ICL trunking between DCX 8510 Backbones and X6 Directors

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 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 only with corresponding ports on another QSFP.

Each core blade in the Brocade DCX 8510-8 and X6-8 contains 16 ICL trunk groups. Each core blade in the Brocade DCX 8510-4 and X6-4 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 DCX 8510 Backbone family and X6 Director family, the QSFP cables must be in the same trunk group. Each ICL that connects the core blades between the Gen 5 and Gen 6 platforms provides up to 64 Gbps of throughput within a single cable.

NOTE

When you connect DCX 8510 Backbone ICL ports with X6 Director ICL ports, you must use the Gen 5 4x16 Gbps QSFPs on the Gen 6 (X6 Director) ICL ports and set the port speed to 16 Gbps on the X6 Directors.

The following tables summarize the possible trunking connections between a DCX 8510 Backbone and an X6 Director:

TABLE 95 Trunking connections between a DCX 8510 Backbone and an X6 Director

Connecting blades	2 to 4-port ICL trunks
CR32-8 and CR16-8	Supported with QSFPs located within the same trunk group on each blade
CR32-8 and CR16-4	Supported with QSFPs located within the same trunk group on each blade
CR32-4 and CR16-8	Supported with QSFPs located within the same trunk group on each blade
CR32-4 and CR16-4	Supported with QSFPs located within the same trunk group on each blade

The following figure is an example of ICL trunks between DCX 8510 Backbones and X6 Directors.

FIGURE 36 ICL trunks between DCX 8510 Backbones and X6 Directors



Refer to the specific hardware reference manuals for information about port numbering, port trunk groups, and connecting the ICL cables.

Virtual Fabrics considerations for ICLs

For Virtual Fabrics configured on supported Brocade Gen 5 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, individual QSFP ports can be split across the logical switch, base switch, or default switch.

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 Backbones and/or Directors in a mesh topology and a core-edge topology. A brief description of each follows. (You can also connect two DCX 8510 chassis or X Directors point-to-point.)

The illustrations in this section show sample topologies. Refer to the *Brocade SAN Scalability Guidelines* for details about maximum topology configurations.

Mesh topology

You can connect the Brocade Backbones or Directors in a mesh topology, in which every chassis is connected to every other chassis.


NOTE

A maximum of nine director chassis can be connected to form a mesh topology. The mesh topology can consist of all homogeneous chassis or a mix of Gen 5 and Gen 6 chassis which 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 Backbone or Director chassis. The chassis for each topology must all be from the same family.

- Brocade X6 Director family (Brocade X6-4 or Brocade X6-8)
- Brocade DCX 8510 Backbone family (Brocade DCX 8510-4 or Brocade DCX 8510-8)

FIGURE 37 Three Brocade X6-8 chassis in 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.

For the Brocade DCX Backbone family only: If one ICL is broken but there is a regular ISL, the triangular topology holds given that the ISL cost is lower than the total cost through the ICL linear topology. If a direct ICL link between two chassis is broken, the triangular topology is considered broken when the ISL path between the two switches is a multiple hop. In this case, the triangular topology broken message is posted independently of the cost of the ISL path being lesser or greater than the ICL path between the two switches.

Another form is the full nine-mesh topology shown in [Figure 38](#). This topology is supported by a mix of Brocade X6-4 and Brocade X6-8 Directors or a mix of Gen 5 and Gen 6 Directors and Backbones that have the same slot count.

FIGURE 38 Full nine-mesh topology



Core-edge topology


You can connect Brocade DCX 8510 Backbones or Brocade X6 Directors in a core-edge topology. For example, [Figure 39](#) shows six chassis connected in a core-edge topology (four edges and two cores). Although this illustration shows only the Brocade DCX 8510-8, each chassis can be either a Brocade DCX 8510-4, a Brocade DCX 8510-8, a Brocade X6-4, or a Brocade X6-8. You can have up to eight edges using Brocade DCX 8510-8 or Brocade X6-8 cores or up to four edges using Brocade DCX 8510-4 or Brocade X6-4 cores.

NOTE

A maximum of 12 chassis can be interconnected to form a core-edge topology using ICL links.

Each line in [Figure 39](#) represents four QSFP cables. The cabling scheme should follow the parallel example shown in [Figure 31](#) on page 360.

FIGURE 39 Illustration of ICL core-edge topology using Brocade chassis platforms



Managing Long-Distance Fabrics

- Long-distance fabrics overview.....373
- Extended Fabrics device limitations.....373
- Long-distance link modes..... 374
- Configuring an extended ISL..... 374
- Forward error correction on long-distance links..... 377

Long-distance fabrics overview

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 long-distance dark fiber or wave-length division multiplexing (WDM) links, while the Brocade Extended Fabrics software enables SAN management over long distances.

Brocade Extended Fabrics is an optional licensed feature for 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) that are greater than 10 km.

For details about obtaining and installing licensed features, refer to the *Brocade Fabric OS Software Licensing Guide*.

The Extended Fabrics feature enables the following functionality:

- Fabric interconnectivity over Fibre Channel at longer distances

ISLs 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 be used to increase the capacity of the links. As Fibre Channel speeds increase, the maximum distance decreases for each switch.

The Extended Fabrics feature extends the distance the ISLs 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 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 (E_Ports) are configured with a large pool of buffer credits. The enhanced switch buffers help ensure that data transfer can occur at near-full bandwidth to use the connection over the extended links efficiently. This efficiency ensures the highest possible performance on ISLs.

Extended Fabrics device limitations

Brocade recommends that you 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.

There is a limited number of reserved buffers 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 (LO) — LO 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 2 Gbps, up to 2 km at 4 Gbps, and up to 1 km at 8, 10, and 16 Gbps.
- 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 2 Gbps. This allocation yields the following values for 10 km:
 - 10 buffer credits per port at 2 Gbps
 - 20 buffer credits per port at 4 Gbps
 - 40 buffer credits per port at 8 Gbps
 - 50 buffer credits per port at 10 Gbps
 - 80 buffer credits per port at 16 Gbps
 - 160 buffer credits per port at 32 Gbps
- 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 is used to determine 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](#) on page 335.

- 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.

Table 96 shows the fabric-wide settings that can be set:

TABLE 96 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 $R_A_TOV/4$
MAX_HOPS	Number	7	7 to 19

4. For 8-Gbps platforms only, enter the **portCfgFillword** command to set ARB as the fill word. Refer to the *Brocade Fabric OS Command Reference* for more information on configuring the fill word for a single 8G FC port.

```
portcfgfillword[ slot/] port, mode
```

The *mode* parameter in this command should be set to 3 if the *vc_translation_link_init* parameter in the **portCfgLongDistance** command (in the next step) is set to 1.

5. Enter the **portCfgLongDistance** command.

```
portcfglongdistance[ slot/] port [ distance_level] [ vc_translation_link_init  
[-distance desired_distance]
```

- Repeat step 4 and step 5 for the remote extended ISL port. 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. This example is for an 8-Gbps platform.

```
switch:admin> portcfgfillword 1/2 3
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.

- Connect to the switch and log in using an account assigned to the admin role.
- 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. Refer to the *Brocade Fabric OS Command Reference* for more information on setting fill words.

- 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](#) on page 112.

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 NVMe

• NVMe overview.....	378
• NVMe over Fibre Channel.....	379
• NVMe Name Server support.....	380
• NVMe devices and zoning.....	381

NVMe overview

Flash and Solid State Devices (SSDs) are a type of non-volatile memory (NVM). NVM Express (NVMe) is a logical device interface specification for accessing NVM storage media attached by means of PCI Express (PCIe) bus, which removes SCSI from the I/O stack. By its design, NVMe allows 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 hard disk drive (rotating) media which has a number of challenges that do not exist in SSDs.

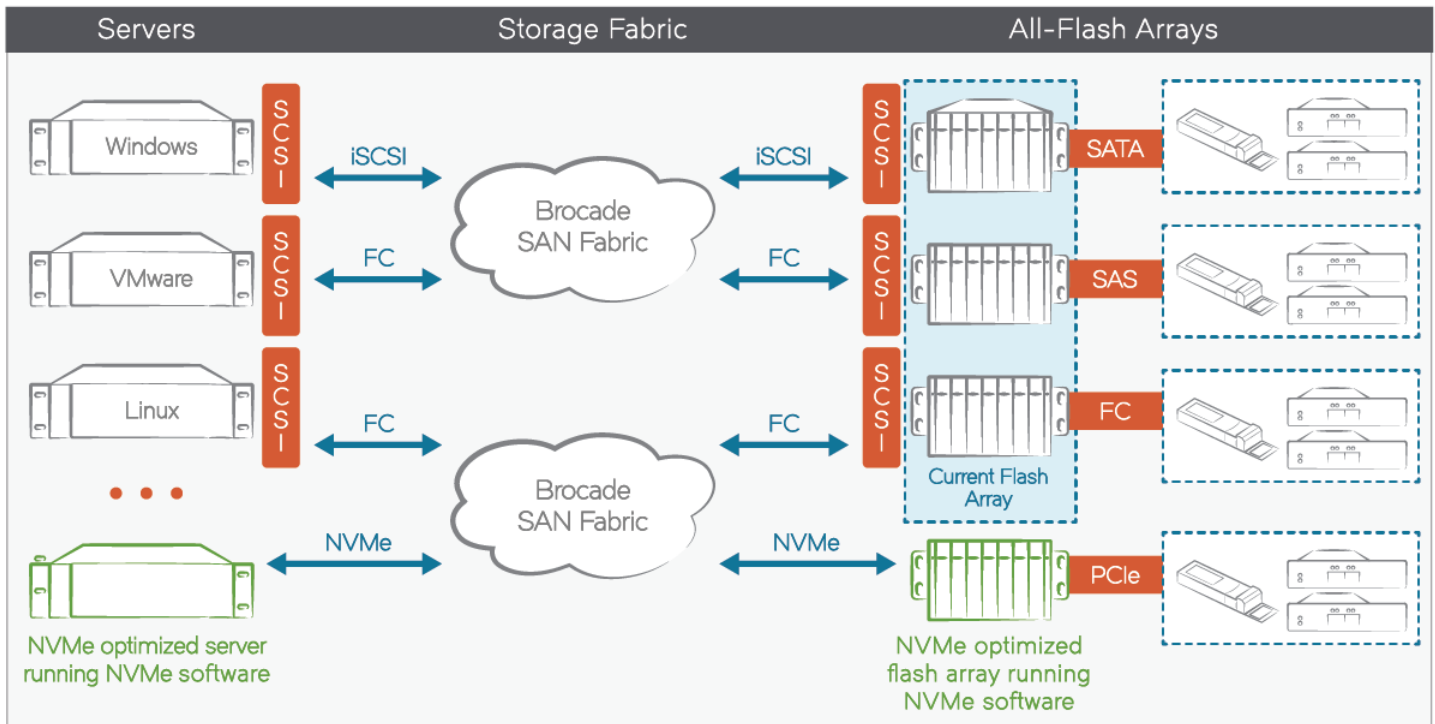
The high performance and reduced cost of SSDs mean that they are increasingly incorporated in data center strategies. SSD media can deliver data with 100 to 200 microseconds of latency; however, the protocol overhead associated with SCSI can significantly reduce the effective speed of SSDs, particularly for latency-sensitive applications that require high I/O and low delay. Next-generation SSD technologies have access latencies of 10 microseconds 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

NVMe over Fibre Channel runs on existing Gen 5 and Gen 6 fabrics and can coexist with 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 40 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 ULP needed by NVMe needs no changes to FC switching logic, which provides backward compatibility with existing FC switches. However, to support NVMe upper layer protocols 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

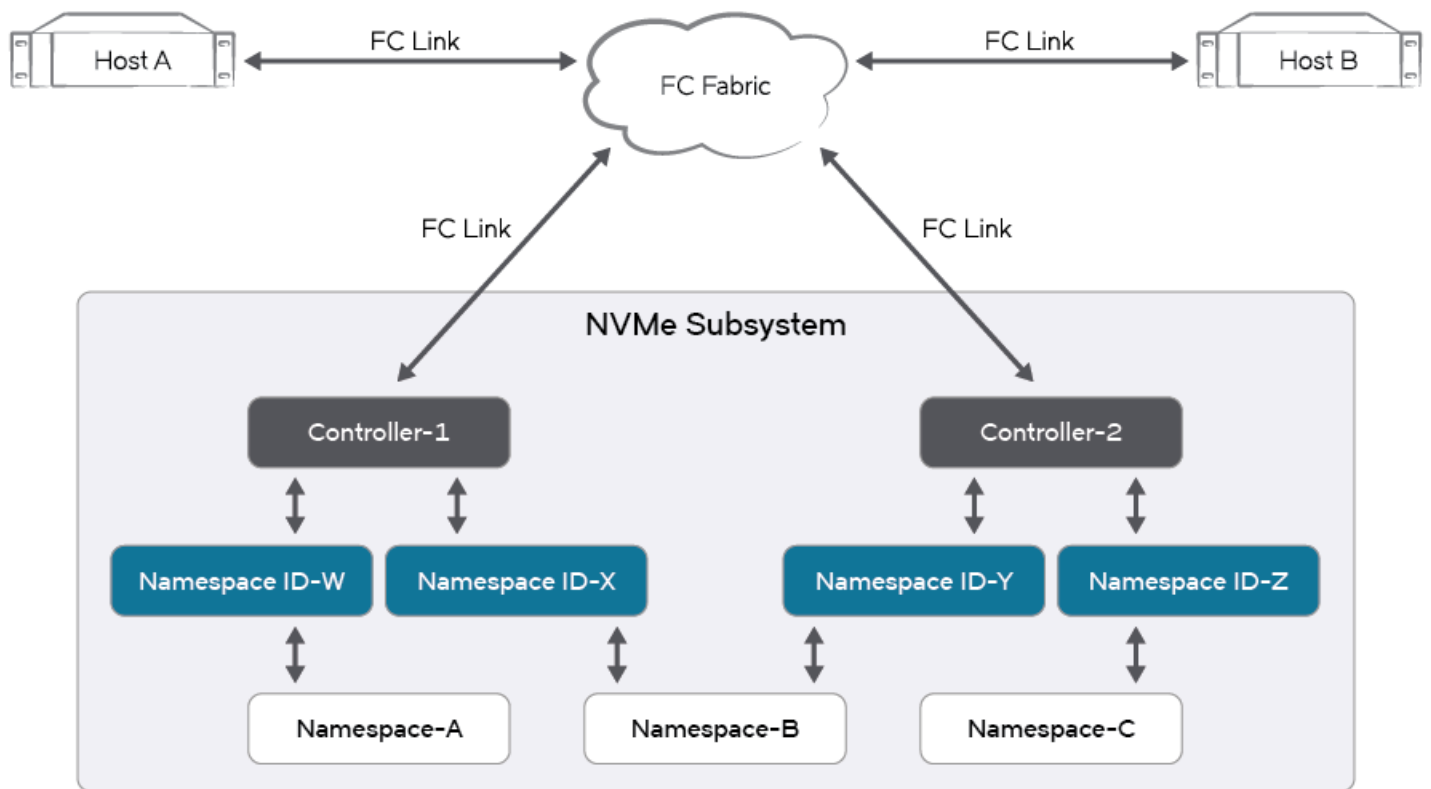
In NVMe over FC, the FC acts as the transport protocol for the NVMe protocol. NVMe frames are encapsulated in FC and transmitted over the SAN. The following table shows the mappings between NVMe protocol and FC protocol. NVMe has a simple set of thirteen commands.

NVMe Entity	FC Mapping
Namespace / NSID	LUN
I/Os	Exchanges
Command	Command RCTL = 6
Status	Status RCTL = 7
Extended/intermediate status	E_Status RCTL = 7
NVMe Subsystem Reset	Unsupported
No Recovery, only retries	SRR/REC
Association Creation	ELS
IO Q creation	ELS

NVMe Entity	FC Mapping
Controller	Target
Namespace	LUN
Rd	Read
Wr	Write

The following illustration 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 the same for both Controller-1 and Controller-2.

FIGURE 41 NVMe Namespace



NVMe Name Server support

When an NVMe device logs in on the fabric, a generic Name Server entry is created. The device registers additional information and the NS entry is updated. 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 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, Web Tools or Brocade Network Advisor. Alternatively, the NVMe device itself can use Target Driven Zoning Mode to create a peer zone. Refer to [Enabling Target Driven Zoning](#) on page 475 for more information.

NOTE

Refer to the Brocade Fabric OS Flow Vision Configuration Guide to view NVMe flow statistics.

Managing Virtual Fabrics

• Virtual Fabrics overview.....	382
• Logical switch overview.....	383
• Management model for logical switches.....	389
• Logical fabric overview.....	390
• Account management and Virtual Fabrics.....	395
• Setting up IP addresses for a logical switch.....	395
• Supported platforms for Virtual Fabrics.....	397
• Virtual Fabrics interaction with other Fabric OS features.....	399
• Limitations and restrictions of Virtual Fabrics.....	400
• Enabling Virtual Fabrics mode.....	401
• Disabling Virtual Fabrics mode.....	402
• Configuring logical switches to use basic configuration values.....	403
• Creating a logical switch or base switch.....	403
• Executing a command in a different logical switch context.....	405
• Deleting a logical switch.....	406
• Adding and moving ports on a logical switch.....	406
• Displaying logical switch configuration.....	407
• Changing the fabric ID of a logical switch.....	408
• Changing a logical switch to a base switch.....	408
• Configuring a logical switch for XISL use	409
• Creating a FICON logical switch.....	410
• Changing the context to a different logical fabric.....	414

Virtual Fabrics overview

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, refer to [FC-FC routing and Virtual Fabrics](#) on page 590.

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 needs to be configured based on the logical switch that is being monitored. If SNMP monitoring is for the default switch, then you do not need the snmpuser to be configured as a user on the switch. Whereas, when SNMP monitoring is VF specific, then you need to create a user with permission for the corresponding VF and add the created user into SNMP database as SNMPv3 user. Once SNMPv3 user is successfully added, then 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. Enabling Virtual Fabrics 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 of the ports in the physical chassis.

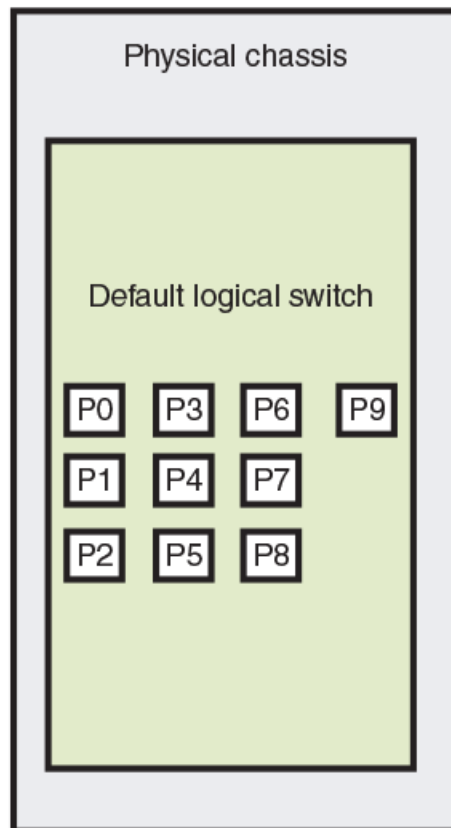
The following figure shows a switch before and after enabling Virtual Fabrics. In this example, the switch has 10 ports, labeled P0 through P9.

FIGURE 42 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 model.

NOTE

The 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 and DCX 8510 Directors support up to 8 logical switches.

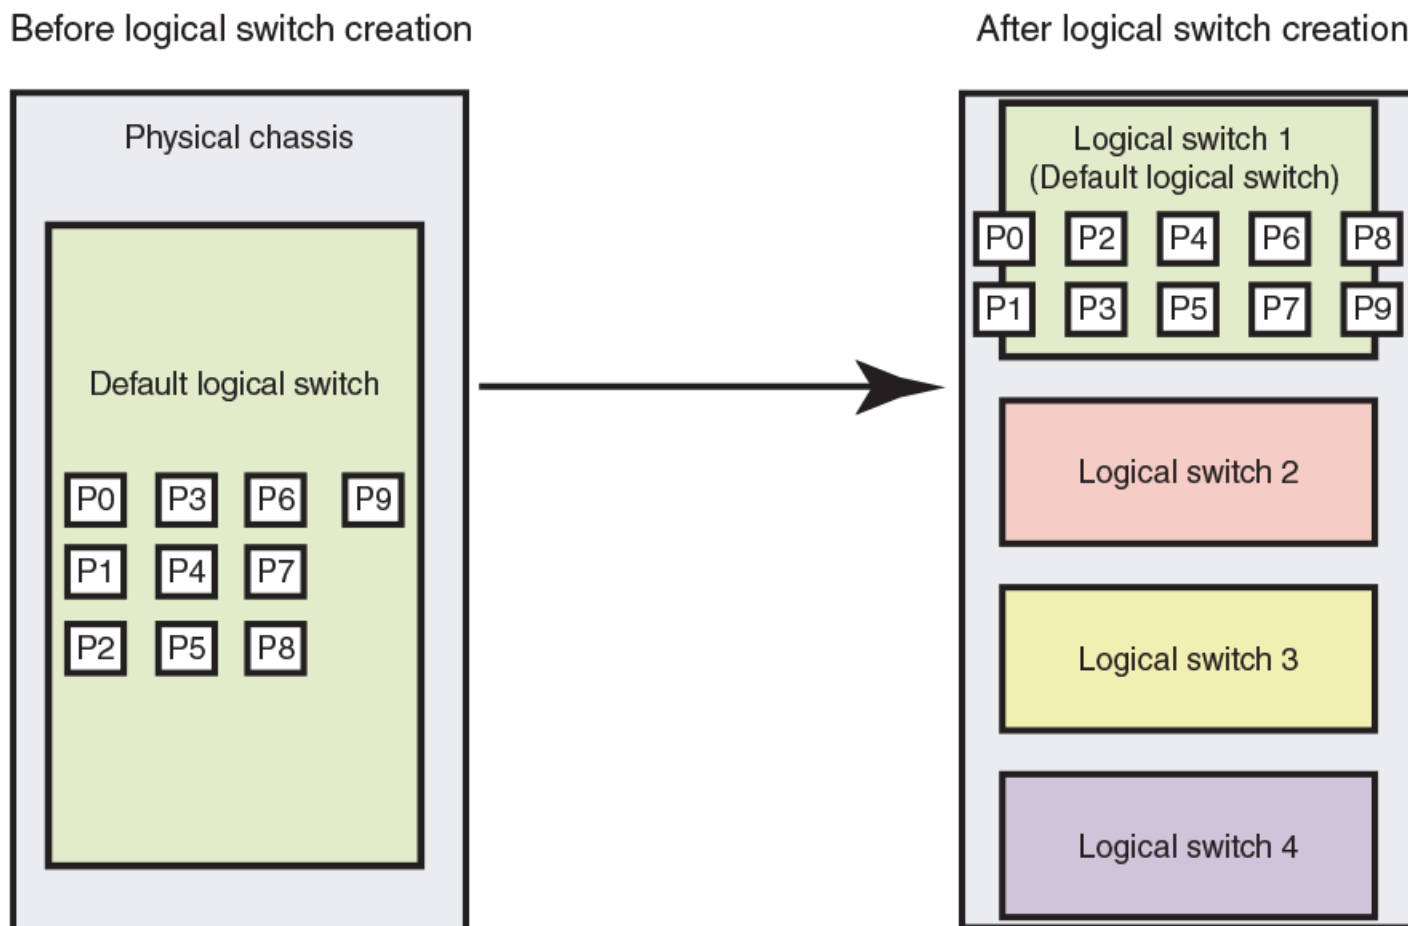
The following figure shows a Virtual Fabrics-enabled switch before and after it is divided into logical switches. Before you create logical switches, the chassis appears as a single switch (default logical switch). After you create logical switches, the chassis appears as multiple independent logical switches. All of 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 43 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.

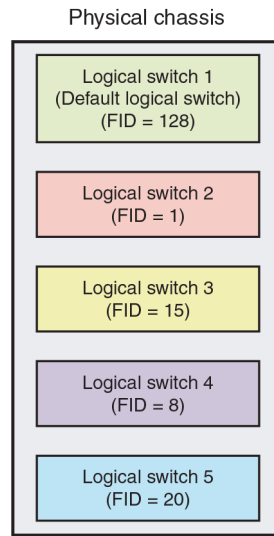
In the following figure, logical switches 2, 3, 4, and 5 are assigned FIDs of 1, 15, 8, and 20, respectively. These logical switches belong to different fabrics, even though they are in the same physical chassis. For example, you could not assign logical switch 5 a fabric ID of 15, because logical switch 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 44 Fabric IDs assigned to logical switches



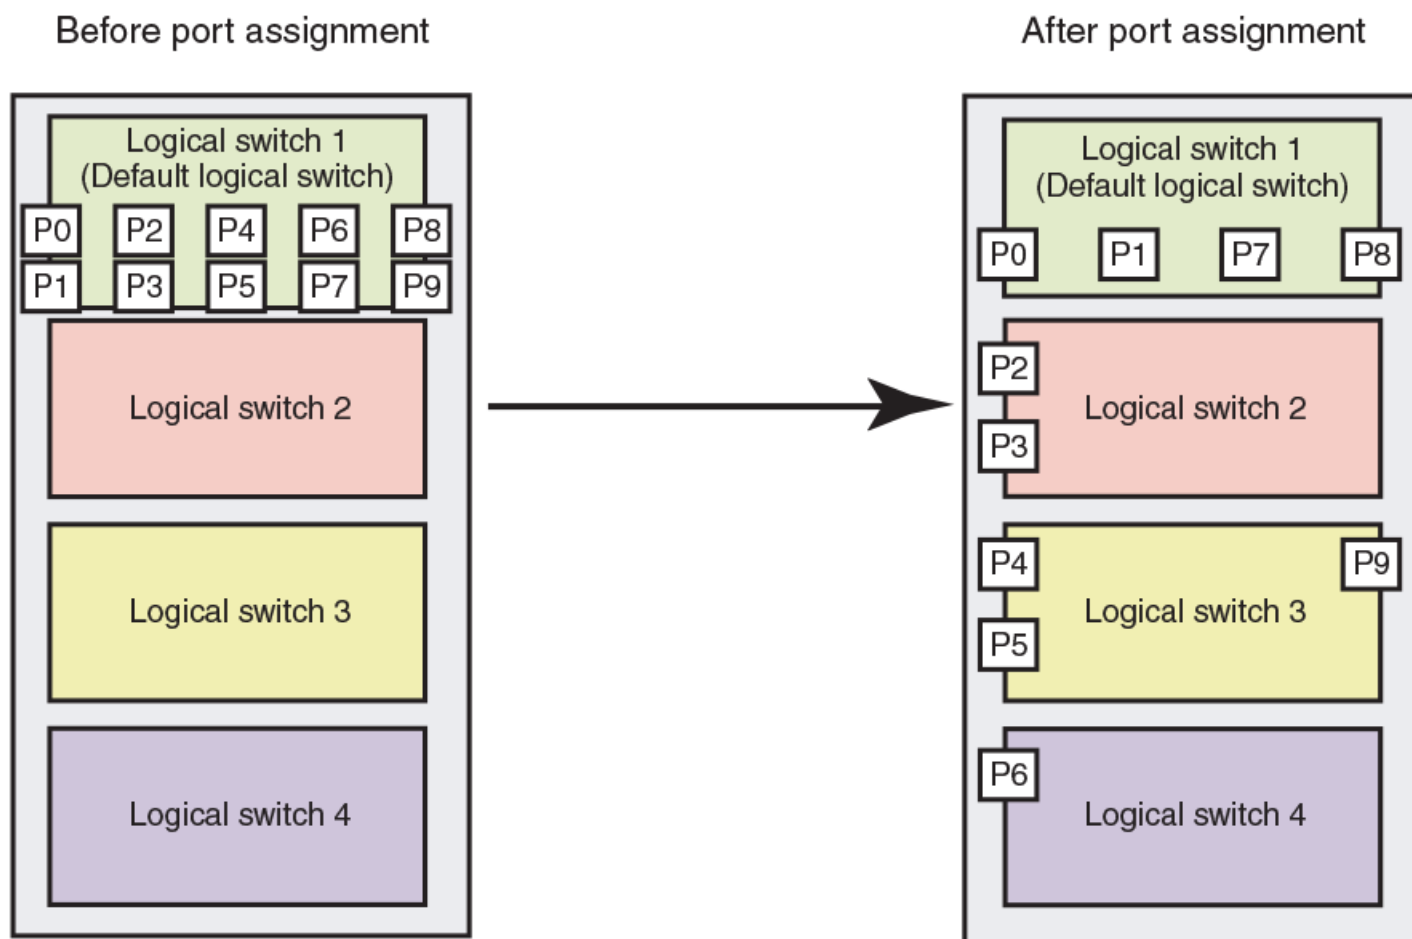
Port assignment in logical switches

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 45 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 in [Figure 45](#):

- 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.

Refer to [Adding and moving ports on a logical switch](#) on page 406 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 [Figure 45](#), 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.

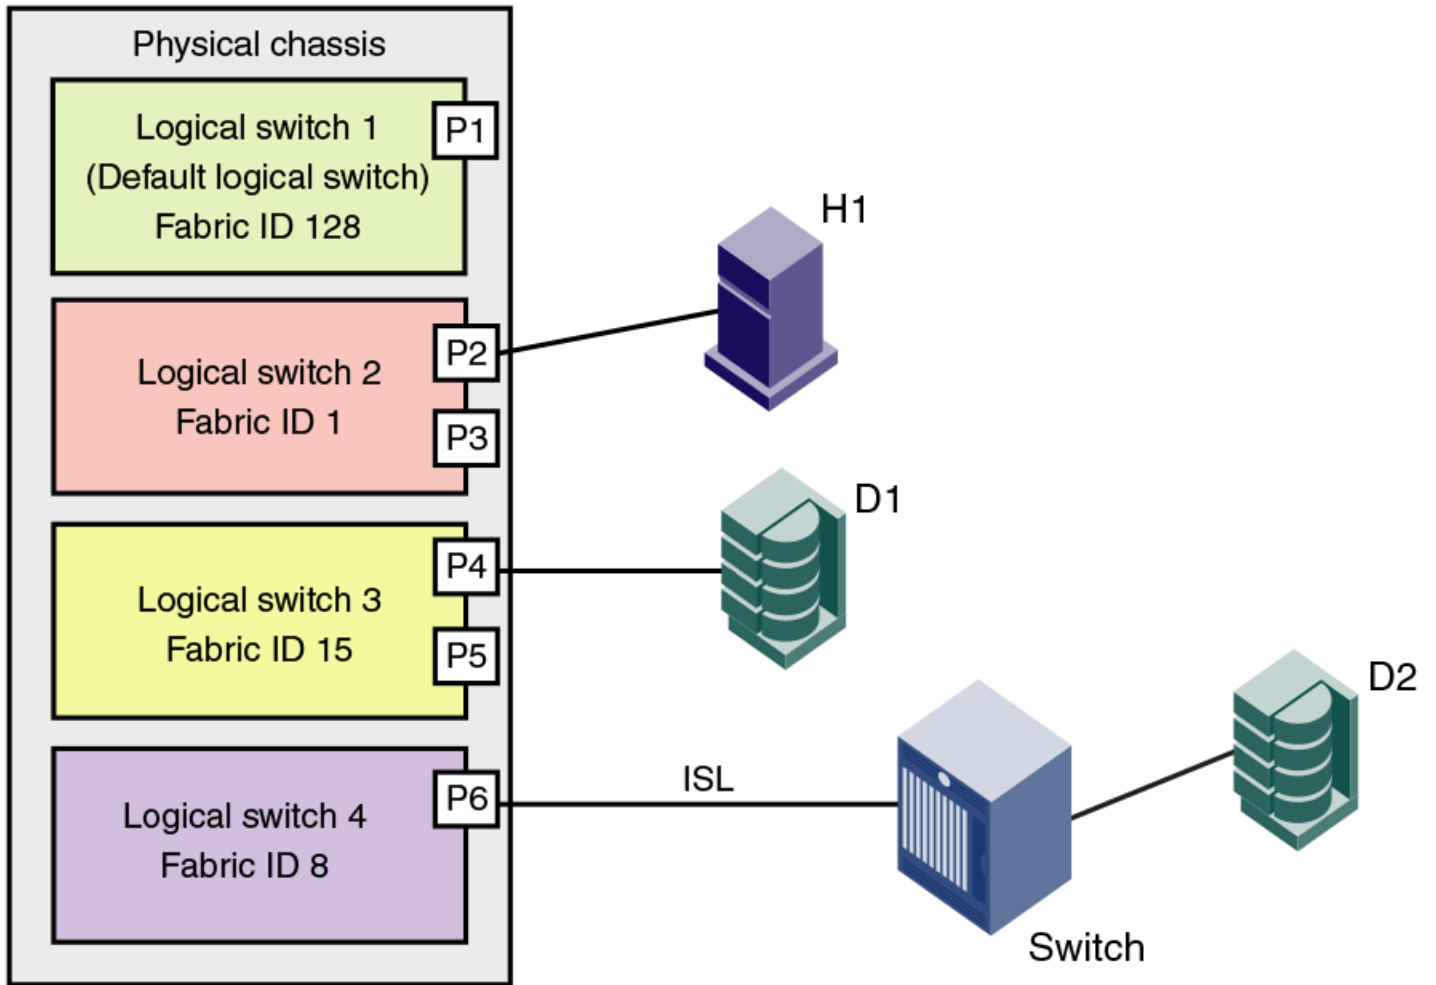
You can move only F_Ports and E_Ports from one logical switch to another. If you want to configure a different type of port, such as a VE_Port or EX_Port, you must configure them after you move them. Some types of ports cannot be moved from the default logical switch. Refer to [Supported platforms for Virtual Fabrics](#) for detailed information about these ports.

Logical switches and connected devices

You can connect devices to logical switches, as shown in [Figure 46](#). 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 46](#), 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 46 Logical switches connected to devices and non-Virtual Fabrics switch



[Figure 47](#) shows a logical representation of the physical chassis and devices in [Figure 46](#). As shown in [Figure 47](#), the devices are isolated into separate fabrics.

FIGURE 47 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 [Figure 46](#) on page 388 and [Figure 47](#) on page 389 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 connect logical switches to non-Virtual Fabrics switches using an ISL, as shown in [Figure 46](#) on page 388.

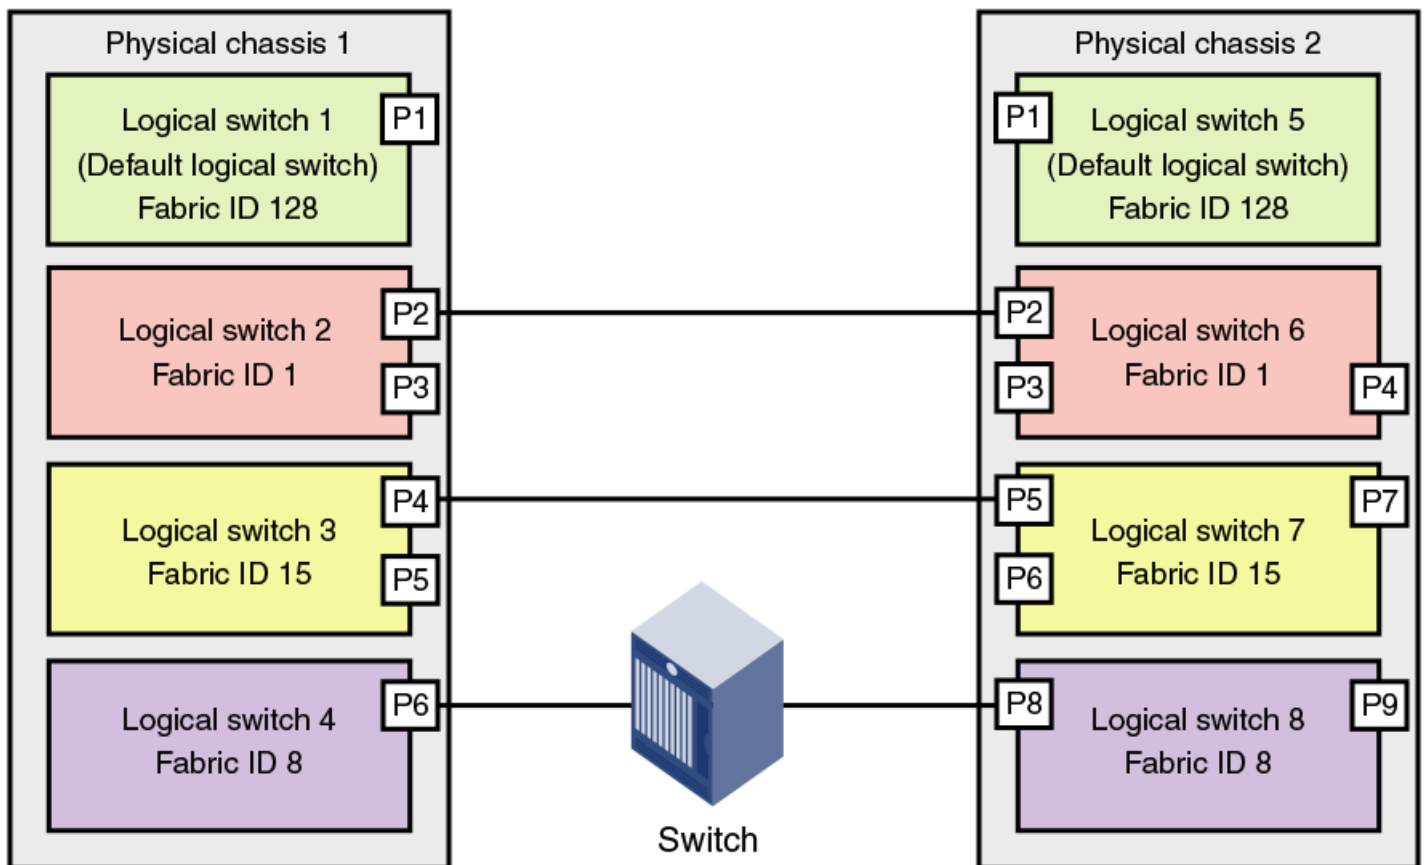
You connect logical switches to other logical switches in two ways:

- Using ISLs
- Using base switches and extended ISLs (XISLs)

Logical fabric and ISLs


[Figure 48](#) shows two physical chassis divided into logical switches. In [Figure 48](#), 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 48 Logical switches connected to other logical switches through physical ISLs



[Figure 49](#) shows a logical representation of the configuration in [Figure 48](#).

FIGURE 49 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 [Figure 48](#), 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 switches.

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, EX_Ports, or VEX_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 is used to share traffic among different logical fabrics.

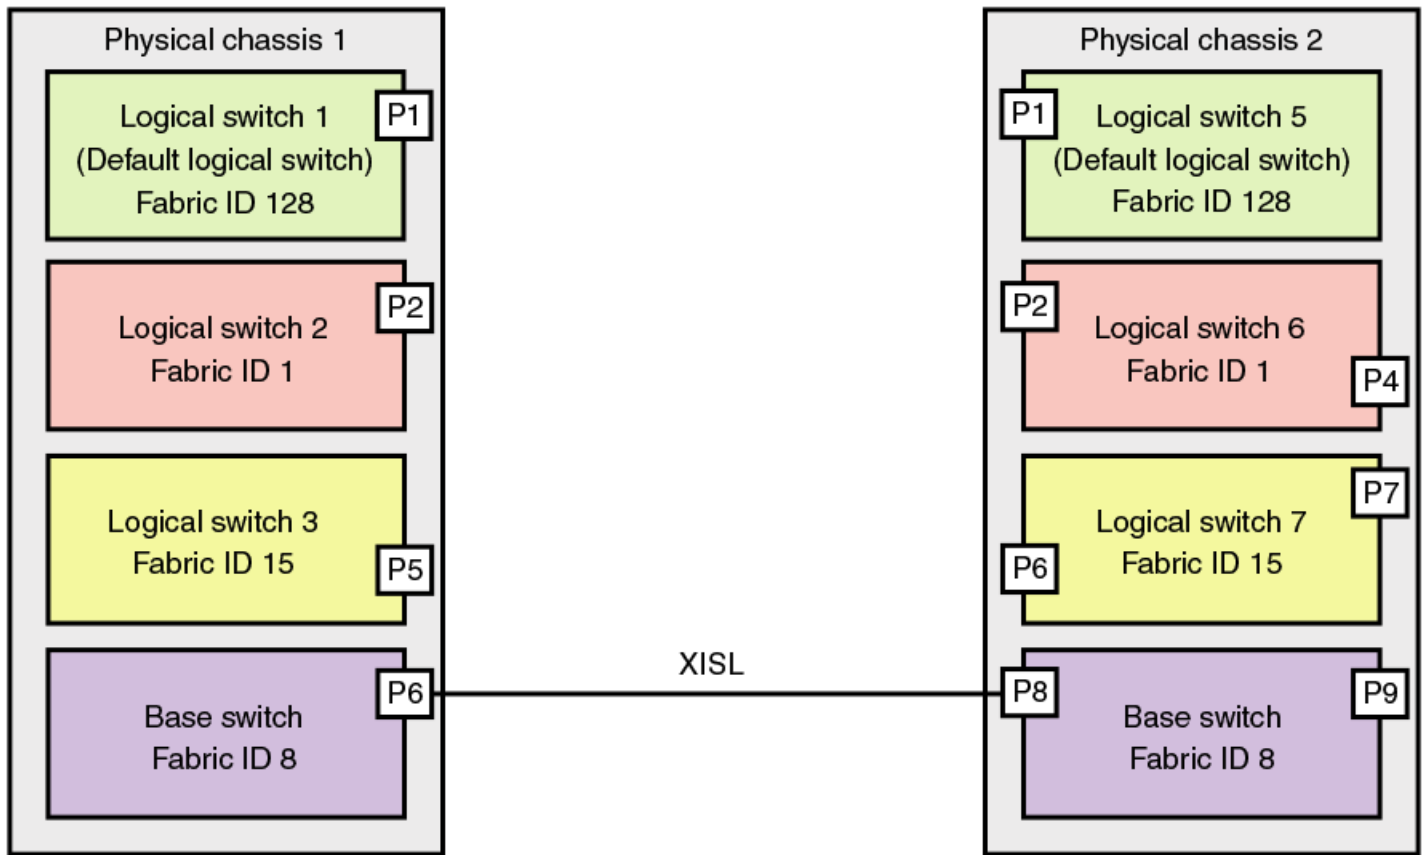
NOTE

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 50 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 51 Logical ISLs connecting logical switches



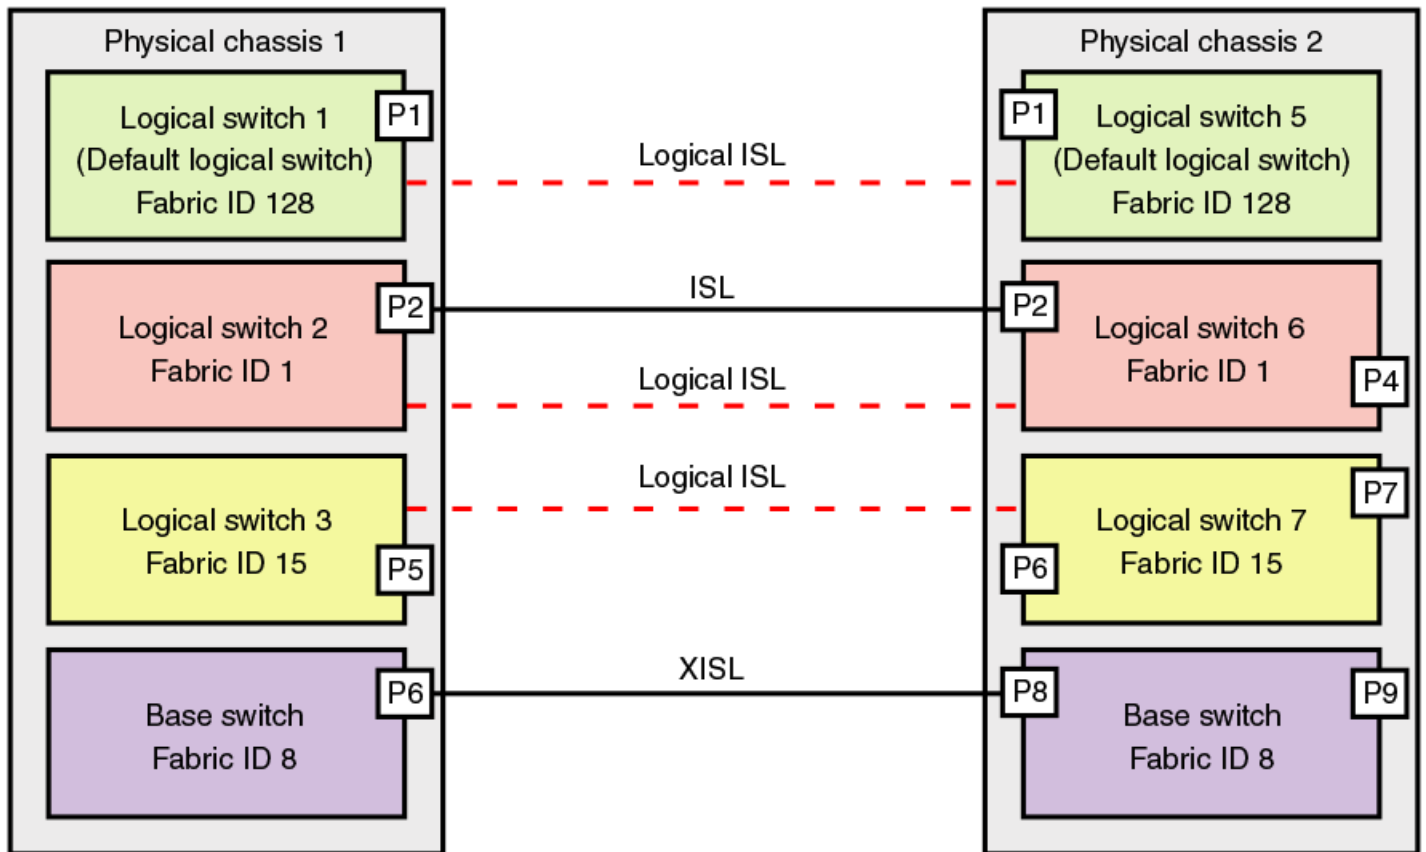
To use the XISL, the logical switches must be configured to allow XISL use. By default, they are configured to do so; you can change this setting, however, using the procedure described in [Configuring a logical switch for XISL use](#) on page 409.

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 52 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

As shown in Figure 52 on page 394, 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 --lis!Enable` 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](#) on page 414.

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](#) on page 144 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 that was supported in releases prior to Fabric OS 7.4.0. 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 10.38.18.183 with netmask 255.255.240.0, you can create the IPFC address using the following command.

```
ipaddrset -ls 8 --add 10.38.18.183/20
```

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. There is only one IPFC mapped to FID 8.

```
switch> ipaddrshow
SWITCH
Ethernet IP Address: 10.38.18.183
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.38.16.1
DHCP: Off
IPFC address for virtual fabric ID 8: 10.38.18.183/20
```

The following conditions apply for default users:

- For admin and root users, the login is successful to the given logical switch with the corresponding logical switch context set (FID 1-128). When you telnet to IPFC 10.38.18.183, the context is set to FID 8.
- For the user privilege, you are allowed to log in to the home virtual fabric (FID 128) only. Direct login is not allowed to the corresponding logical switch context. When you telnet to IPFC 10.38.18.183, the context is set to FID 128.

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, then 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 supports the Virtual Fabrics:

- Brocade 6510
- Brocade 6520
- Brocade G620

- Brocade G630
- Brocade 7840 Extension
- Brocade DCX 8510 family
- Brocade X6 family

Some restrictions apply to the ports, depending on the port type and blade type. The following sections explain these restrictions.

Supported port configurations in the fixed-port switches

There are no restrictions on the ports in Brocade 6510, 6520, G620, or G630 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, *except* on Brocade Backbone family devices.
- The default logical switch can also be a base switch except on Brocade DCX 8510 Backbones and X6 Directors.

For the Brocade 7840, the following rules apply:

- A base switch is supported on the Brocade 7840 starting with Fabric OS 7.4.0.
- XISL is supported on the Brocade 7840 starting with Fabric OS 7.4.0.

Supported port configurations in Brocade Backbones and Directors

Some of the ports in the Brocade DCX 8510 Backbones 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 97 Blade and port types supported on logical switches

Blade type	Default logical switch	User-defined logical switch	Base switch
FC32-64: FC Port	Yes (F, E)	Yes (F, E)	Yes (E, EX)
FC32-64: Eth Port	Yes	Yes	Yes
FC16-32 FC16-48 FC32-48	Yes (F, E)	Yes (F, E)	Yes (E, EX)
FC16-64	Yes (F, E) In the Brocade DCX 8510-8, ports 56-63 of the FC8-64 blade are not supported as E_Ports on the default logical switch. The Brocade DCX 8510-4 does not have this limitation	Yes (F, E)	Yes (E, EX) In the Brocade DCX 8510-8, ports 48-63 of the FC16-64 blades are not supported in the base switch. The Brocade DCX 8510-4 does not have this limitation
SX6: FC ports GE ports	Yes (F, E) Yes (VE)	Yes (F, E) Yes (VE)	Yes (E, EX) Yes (VE)
FX8-24: FC ports GE ports	Yes (F, E) Yes (VE)	Yes (F, E) Yes (VE)	Yes (E, EX) Yes (VE, VEX)
ICL ports	Yes (E)	Yes (E)	Yes (E, EX) EX_Ports on an ICL are supported only in the Brocade DCX 8510 and X6 Director family

Restrictions on Brocade Backbones

The following restrictions apply to Brocade Backbones:

- EX_Ports and VEX_Ports can be in only the base switch.
- ICL ports cannot be in a logical switch that is using XISLs.
- All of the user ports in an ICL cable must be in the same logical switch. Distributing the user ports within the same cable across multiple logical switches is not supported.
- ICL ports that are configured as EX_Ports can be in only the base switch.
- The default logical switch cannot use XISLs.
- The default logical switch cannot be designated as the base switch.
- In Fabric OS v7.0.0 and later, VE_Ports on the FX8-24 blade are supported on a logical switch that is using an XISL, and on the base switch as an XISL.

NOTE

For the FX8-24 blade, if XISL use is enabled it is not recommended that you configure VE_Ports on both the logical switch and the base switch, because FCIP tunnels support only two hops maximum.

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 98 Virtual Fabrics interaction with Fabric OS features

Fabric OS feature	Virtual Fabrics interaction
Access Gateway	Virtual Fabrics is not supported on a switch if AG mode is enabled.
Configuration upload and download	Virtual Fabrics uses a configuration file that is different from the configuration file used to download system configuration parameters. Refer to Maintaining the Switch Configuration File on page 133 for more information about how Virtual Fabrics affects the configuration file.
FC-FC Routing Service	All EX_Ports must reside in a base switch. 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. Refer to Using FC-FC Routing to Connect Fabrics on page 535 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.
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. Refer to Network Time Protocol on page 59 for more information about how Virtual Fabrics affects NTP.
QoS	QoS VCs are maintained across the base fabric. Refer to Optimizing Fabric Behavior on page 319 for more information about using the Adaptive Networking features with Virtual Fabrics.

TABLE 98 Virtual Fabrics interaction with Fabric OS features (continued)

Fabric OS feature	Virtual Fabrics interaction
Traffic Isolation	<p>Traffic Isolation with Failover-Disabled (TI FD) zoning is not supported over LISLs; Traffic Isolation zones with Failover-Enabled (TI FE) is supported. TI FD zoning is supported over DISLs (physical links). TI FD zoning is not supported in the base fabric.</p> <p>Refer to Optimizing Fabric Behavior on page 319 for additional information about using TI Zones with Virtual Fabrics.</p>

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. However, if one of the four ports belongs to the base switch, then the remaining three ports cannot be assigned to any other logical switch.

The maximum number of logical switches per chassis varies depending on the switch model. The following table lists the supported platforms and the maximum number of logical switches (including the default logical switch) supported on each.

TABLE 99 Maximum number of logical switches per chassis

Platform	Maximum number of logical switches
Brocade X6 Directors	16
Brocade DCX 8510 Backbones	8
Brocade G630	4
Brocade G620	4
Brocade 7840	4
Brocade 6520	4
Brocade 6510	4

Refer to [#unique_665](#) for restrictions on the default logical switch.

If a blade slot is being decommissioned and has ports configured in logical switches, it is recommended that the logical port assignments be removed from that blade before removing the blade. This ensures a seamless transition for any new port or AP blade that might occupy that slot in the future. This 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](#) on page 409.

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 DCX 8510 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 FX8-24 or SX6 blade can be moved to any logical switch independent of the location of the physical GE port.
- If you move existing EX_Ports or VEX_Ports to any logical switch other than the base switch, these ports are automatically disabled.

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 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** command to check whether VF mode is enabled:

```
fosconfig --show
```

3. Use the **fosConfig** command to enable VF mode:

```
fosconfig --enable vf
```

4. Enter **y** at the prompt.

The following example checks whether VF mode is enabled or disabled and then enables it.

```
switch:admin> fosconfig --show
FC Routing service:           disabled
iSCSI service:                Service not supported on this Platform
iSNS client service:         Service not supported on this Platform
Virtual Fabric:              disabled
Ethernet Switch Service:     Service not supported on this Platform

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 occurs:

- The CPs are rebooted.
 - If F_Port trunking is enabled on ports in the default switch, the F_Port trunking information is preserved.
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 of 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
iSCSI service:                Service not supported on this Platform
iSNS client service:         Service not supported on this Platform
Virtual Fabric:              enabled
Ethernet Switch Service:     Service not supported on this Platform

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.

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 | -lislisable] [-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, you can use the **-lislisable** 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](#) on page 406.

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> fosexec --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 10.32.79.105 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 10.32.79.105 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 you want to delete, 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 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](#) on page 406.
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 add ports to a logical switch by moving the ports from one logical switch to another. To remove a 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 on 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

NOTE

On the Brocade DCX 8510-8, the `lscfg` command does not allow you to add ports 48-63 of the FC16-64 blade to the base switch. These ports are not supported on the base switch. The Brocade DCX 8510-4 does not have this limitation.

```
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 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 of the logical switches

```
sw0:FID128:admin> lscfg --show -n
-----
Switch Information
-----
FID: 5
SwitchType: BS
DomainID: 3
SwitchName: Base1
FabricName: Fab1
-----
FID: 128
SwitchType: DS
DomainID: 1
SwitchName: sw0
FabricName: DefFab
-----
```

Example of displaying a list of all of the 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](#) on page 409.

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

```
sw0:FID128:admin> setcontext 7
switch_25:FID7:admin> switchshow
switchName:      switch_25
switchType:      66.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     30
switchId:        fffc1e
switchWwn:       10:00:00:05:1e:82:3c:2c
zoning:          OFF
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_25:FID7: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_25:FID7: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_25:FID7: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.

Use the following procedure to allow or disallow the logical switch to use XISLs in the base fabric.

Refer to [Limitations and restrictions of Virtual Fabrics](#) on page 400 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.

- 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.

- Use the **switchShow** command and check the value of the Allow XISL Use parameter.
- Enter the **configure** command:
- Enter **y** after the Fabric Parameters prompt:

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

- 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
```

- Respond to the remaining prompts or press **Ctrl-d** to accept the other settings and exit.

Creating a FICON logical switch

With the FICON logical switch feature, you can create a FICON logical switch. You can also modify an existing logical switch as a FICON logical switch.

The FICON logical switch is supported on all platforms that support Virtual Fabrics in Fabric OS 8.1.0. A FICON logical switch can be created independent of enabling FMS mode. A FICON logical switch cannot be the default switch in the chassis nor can it be the base switch of the chassis.

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
- Routing policy is set to device-based routing
- Area mode is set to zero-based addressing
- SCC security policy is auto-defined with only its own WWN
- The created SCC policy is activated
- The fabric-wide consistency policy is set to strict SCC policy
- High Integrity Fabric (HIF) mode is enabled
- FMS mode is enabled if FICON 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.

Once 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 once it is created.

Use the following procedure to create a FICON logical switch.

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use 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.
```

The command takes approximately 30 seconds to complete.

3. Use the **lscfg** command with the **--show** option to verify the logical switch creation. The output 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. 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 HIF mode is enabled and Address Mode is set to 1 (zero-based address mode).

- Use the **aptpolicy** command to display the routing policy.

```
switch_50:FID50: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
```

- Refer to [Adding and moving ports on a logical switch](#) on page 406 for information on how to move ports into a logical switch.

The following command sequence shows a port migrated into the FICON logical switch.

```
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 (Port not bounded to Address in
FICON Switch)

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
```

8. **NOTE**

For ports that are migrated into the FICON logical switch, you must bind the port by using the **portaddress --bind** command. If the port is not bound after migration, the port enable operations fails and the port is disabled with the information, "Port needs to be bound to area in FICON Switch." You cannot bind ICL ports with the **portaddress --bind** command.

Use the **portaddress --bind** command to bind or unbind ports. 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.

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.
```

Changing an existing logical switch to a FICON logical switch

You can change an existing logical switch to a FICON logical switch.

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. The operation succeeds when the logical switch is already in HIF-enabled mode and switch address mode is set to zero-based addressing. If these conditions are not met, the operation fails.

NOTE

Changing a normal logical switch to a FICON logical switch is non-disruptive if the switch already is in HIF mode, is running zero based addressing, and all online ports are user bound.

To change an existing logical switch to a FICON logical switch, perform the following procedure.

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use the **lscfg** command with the **-ficon** option to change 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.
```

Considerations for FICON logical switch

The following considerations apply to the FICON logical switch:

- Firmware downgrade is blocked to releases prior to Fabric OS 8.1.0. if a FICON logical switch is defined at the time of the downgrade. You are prompted to delete the existing FICON logical switch prior to 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 the FICON mode logical switch creation or modification is attempted while the standby CP is running a software version that does not support 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).
```

- Configupload and configdownload is supported to the extent that the switch-wide attributes can be uploaded and downloaded. However, the FICON switch attribute itself is not part of the configupload/download of a FID and it cannot be saved as part of a configupload of a FID. For example, when a configupload is performed on a FICON logical switch and the logical switch is subsequently made a normal logical switch. A configdownload will not make then normal logical switch a FICON logical switch, only switch related attributes will be downloaded. However, the FICON attribute of the FICON logical switch can be uploaded and downloaded by using the **configupload -vf** command.

Configdownload cannot turn off the IDID mode, HIF mode, and set the address mode to anything other than zero-based addressing in a FICON logical switch.

```
configdownload
Protocol (scp, ftp, sftp, local) [ftp]: scp
:
```

```
.  
Doing configDownload on switch 4...  
High Integrity Fabric mode cannot be turned off in FICON Logical switch.
```

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

• Zone types.....	415
• Zoning overview.....	416
• Broadcast zones.....	422
• Zone aliases.....	423
• Zone creation and maintenance.....	427
• Default zoning mode.....	439
• Zone database size.....	440
• Zone configurations.....	441
• Zone object maintenance.....	452
• Zone configuration management.....	454
• Security and zoning.....	454
• Zone merging.....	455
• Concurrent zone transactions.....	462
• Peer Zoning.....	464
• Target Driven Zoning.....	473
• Supported commands for Peer Zones and Target Driven Peer Zones.....	480
• Boot LUN zoning.....	480

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 these regular zones. Beyond this, Fabric OS has the following types of special zones:

- Broadcast zones
Control which devices receive broadcast frames. A broadcast zone restricts broadcast packets to members of the broadcast zone only. Refer to [Broadcast zones](#) on page 422 for more information.
- Frame redirection zones
Re-route frames between an initiator and a target through a Virtual Initiator and Virtual Target for special processing or functionality, such as for storage virtualization or encryption. Refer to [Frame Redirection](#) on page 532 for more information.
- LSAN zones
Provide device connectivity between fabrics without merging the fabrics. Refer to [LSAN zone configuration](#) on page 566 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. Refer to [QoS](#) on page 321 for more information.
- Traffic Isolation zones (TI zones)
Isolate traffic to a specific, dedicated path through the fabric. Refer to [Traffic Isolation Zoning](#) on page 484 for more information.
- Peer zones
Allow a principal device port to communicate with other members of the zone. Refer to [Peer Zoning](#) on page 464 for more information.
- Target driven zones

Allow targets to create or activate peer zones. Refer to [Target Driven Zoning](#) on page 473 for more information.

Zoning overview

Zoning is a fabric-based service that enables you to partition your storage area network (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.

Consider [Figure 53](#), which shows a 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 53 Zoning example



Refer to [Best practices for zoning](#) on page 421 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*.

Approaches to zoning

Table 100 lists the various approaches you can take when implementing zoning in a fabric.

TABLE 100 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	
No fabric zoning	<p>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</p>

TABLE 100 Approaches to fabric-based zoning (continued)

Zoning approach	Description
	only in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

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 Backbones, "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, then all devices connected to that port are in the zone.
- World Wide Names 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 the 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 **alcreate** command.

```
switch:admin> alcreate 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 8.0.1 or later to a version earlier than Fabric OS 8.0.1 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.
- World Wide Name (WWN)

All members are specified only by World Wide Names (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](#) on page 439). This does not mean that the zone database is deleted, however, only that there is no configuration 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. *Hardware-enforced zoning* means that each frame is checked by hardware (the ASIC) before it is delivered to a zone member and is 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.

There are two methods of hardware enforcement:

- Frame-based hardware enforcement: All frames are checked by the hardware.
- Session-based hardware enforcement: The only frames checked by hardware are the ELS frames (such as PLOGI and RNID) used to establish a session.

The hardware-enforcement method 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.

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.

Session-based hardware enforcement is in effect if the zone has a mix of WWN and D,I members.

If a port is in multiple zones, and is defined by WWN in one zone and by D,I in another, then session-based hardware enforcement is in effect.

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

Table 101 lists considerations for zoning architecture.

TABLE 101 Considerations for zoning architecture

Item	Description
Type of zoning enforcement: frame- or session-based	If security is a priority, frame-based hardware enforcement is recommended. 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.
Testing	Before implementing a new zone, you should run the Zone Analyzer from Web Tools to isolate any possible problems. This is especially useful as fabrics increase in size.
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 it is recommended that you 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 (by closing the transaction) 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 backwards-compatible but not forward-compatible.
- Zone using the core switch in preference to using an edge switch.
- Zone using a Backbone rather than a switch. A Backbone has more resources to handle zoning changes and implementations.
- When you are adding a switch to an 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 7.3.0, 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 7.3.0
 - A merge with a switch running firmware earlier than Fabric OS 7.3.0
 - A configuration download using firmware earlier than Fabric OS 7.3.0

and 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](#) on page 447 for instructions on doing so.

Broadcast zones

Fibre Channel allows sending broadcast frames to all Nx_Ports if the frame is sent to a broadcast well-known address (FFFFFF); however, many target devices and HBAs cannot handle broadcast frames. To control which devices receive broadcast frames, you can create a special zone, called a *broadcast zone*, that restricts broadcast packets to only those devices that are members of the broadcast zone.

If there are no broadcast zones or if a broadcast zone is defined but not enabled, broadcast frames are not forwarded to any F_Ports. If a broadcast zone is enabled, broadcast frames are delivered only to those logged-in Nx_Ports that are members of the broadcast zone and are also in the same zone (regular zone) as the sender of the broadcast packet.

Devices that are not members of the broadcast zone can send broadcast packets, even though they cannot receive them.

A broadcast zone can have *domain,port*, WWN, and alias members.

Broadcast zones do not function in the same way as other zones. A broadcast zone does not allow access within its members in any way. If you want to allow or restrict access between any devices, you must create regular zones for that purpose. If two devices are not part of a regular zone, they cannot exchange broadcast or unicast packets.

To restrict broadcast frames reaching broadcast-incapable devices, create a broadcast zone and populate it with the devices that are capable of handling broadcast packets. Devices that cannot handle broadcast frames must be kept out of the broadcast zone so that they do not receive any broadcast frames.

You create a broadcast zone the same way you create any other zone except that a broadcast zone must have the name "broadcast" (case-sensitive). To enable the broadcast zoning, use **cfgenable** command to add the broadcast zones to the active zone configuration. You set up and manage broadcast zones using the standard zoning commands, described in [Zone creation and maintenance](#) on page 427.

Broadcast zones and FC-FC routing

If you create broadcast zones in a metaSAN consisting of multiple fabrics connected through an FC router, the broadcast zone must include the IP device that exists in the edge or backbone fabric as well as the proxy device in the remote fabric. Refer to [Using FC-FC Routing to Connect Fabrics](#) on page 535 for information about proxy devices and the FC router.

High availability considerations with broadcast zones

If a switch has broadcast zone-capable firmware on the active CP (Fabric OS v5.3.x or later) and broadcast zone-incapable firmware on the standby CP (Fabric OS version earlier than v5.3.0), then you cannot create a broadcast zone because the zoning behavior would not be the same across an HA failover. If the switch failed over, then the broadcast zone would lose its special significance and would be treated as a regular zone.

Loop devices and broadcast zones

Delivery of broadcast packets to individual devices in a loop is not controlled by the switch. Consequently, adding loop devices to a broadcast zone does not have any effect. If a loop device is part of a broadcast zone, then all devices in that loop receive broadcast packets.

Broadcast zones and default zoning mode

The default zoning mode defines the device accessibility behavior if zoning is not implemented or if 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.

If a broadcast zone is active, even if it is the only zone in the effective configuration, the default zone setting is not in effect.

If the effective configuration has only a broadcast zone, then the configuration appears as a No Access configuration. To change this configuration to All Access, you must put all the available devices in a regular zone.

Refer to [Default zoning mode](#) on page 439 for additional information about default zoning.

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, then the alias for the associated zone is ZNE_MAILSERVER_SLT5. This clearly identifies the server host bus adapter (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](#) on page 455.

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. Only after the alias is removed from the TDZ by the Target device or if the TDZ is deleted, the alias can be edited.

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 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.

Example

```
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

WARNING!!!
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
If the update includes changes to one or more traffic isolation
zones, you must issue the 'cfgenable' command for the changes
to take effect.
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 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.

Example

```
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.
If the update includes changes to one or more traffic isolation
zones, you must issue the 'cfgenable' command for the changes
to take effect.
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 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.

Example

```
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.
If the update includes changes to one or more traffic isolation
zones, you must issue the 'cfgenable' command for the changes
to take effect.
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 displays on the other switches to indicate that the transaction was aborted.

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.
If the update includes changes to one or more traffic isolation
zones, you must issue the 'cfgenable' command for the changes
to take effect.
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** command, followed optionally by the **--ic**, *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* for more information on this command.

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 better manage devices.

The following items should be taken into account when creating zones:

- Broadcast zone — To create a broadcast zone, use the reserved name "broadcast". Do *not* give a regular zone this name. Refer to [Broadcast zones](#) on page 422 for additional information about this special type of zone.
- Virtual Fabrics — Zone definitions should not include logical port numbers. Zoning is not enforced on 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:
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 7,3; 4,5
alias: zeus 4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
```

Creating a zone

ATTENTION

The **zoneCreate** command will add *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 " zonenumber ", " member[; member...]"
zonecreate " zonenumber ", " aliasname_pattern *[, members]"
```

NOTE

The **zoneCreate** command supports partial pattern matching ("wildcards") of zone member aliases. This allows you to add multiple aliases that match the "*aliasname_pattern*" in the command line.

To create a broadcast zone, use the reserved name "broadcast".

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 displays on the other switches to indicate that the transaction was aborted.

4. Enter the **cfgShow** command to view the changes.

Example of creating a new zone

```
switch:admin> zonecreate sloth, "b*; 10:00:00:00:01:1e:20:20"

switch:admin> cfgsave

switch:admin> cfgenable

switch:admin> cfgshow

Defined configuration:
zone: matt      30:06:00:07:1e:a2:10:20; 3,2
zone: sloth    bawn; bolt; bond; brain; 10:00:00:00:01:1e:20:20 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   7,3; 4,5
alias: zeus    4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
```

Adding devices (members) to a zone

ATTENTION

The **zoneAdd** command will add *all* zone member aliases that match the "*aliasname_pattern*" in the zone database to the specified 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. This 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 displays on the other switches to indicate that the transaction was aborted.

4. Enter the **cfgShow** command to view the changes.

Example for adding members to a zone

```
switch:admin> zoneadd matt, "ze*; bond*; j*"

switch:admin> cfgsave

switch:admin> cfgshow

Defined configuration:
zone: matt    30:06:00:07:1e:a2:10:20; 3,2; zeus; bond; jake; jeff; jones
zone: sloth  bawn; bolt; bond; brain; 10:00:00:00:01:1e:20:20
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   7,3; 4,5
alias:      zeus    4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
```

Removing devices (members) from a zone

ATTENTION

The **zoneRemove** command will remove *all* zone member aliases that match the "*aliasname_pattern*" in the zone database from the specified 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. This 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 displays on the other switches to indicate that the transaction was aborted.

4. Enter the **cfgShow** command to view the changes.

Examples removing members from a zone

```
switch:admin> cfgshow

Defined configuration:
zone: matt    zeus; bond; jake; jeff; jones; 3,2; 30:06:00:07:1e:a2:10:20
zone: sloth   bawn; bolt; bond; brain; 10:00:00:00:01:1e:20:20
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   7,3; 4,5
alias:        zeus    4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
switch:admin>
switch:admin> zoneremove matt,"30:06:00:07:1e:a2:10:20; ja*; 3,2"

switch:admin> cfgsave

switch:admin> cfgshow

Defined configuration:
zone: matt    zeus; bond; jake; jeff; jones zone: sloth   bawn; bolt; bond; brain; 10:00:00:00:01:1e:20:20
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   7,3; 4,5
alias:        zeus    4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
```

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, a **cfgEnable** command should be issued after the **zoneObjectReplace** command. Otherwise, the changes will be in the transaction buffer but not committed.
- Only members of regular zones and aliases (those identified using either *D,I* or WWN) can be replaced using **zoneObjectReplace**.
- The **zoneObjectReplace** command is not applicable for Frame Redirect (FR) and Traffic Isolation (TI) zones. Only members of regular zones can be replaced using this command.
- 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 a new 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 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.

4. Enter the **cfgShow** command to view the changes.

Example replacing zone members

```
switch:admin> cfgshow

Defined configuration:
zone: matt zeus; bond; jeff;jones; 11, 2
zone: sloth bawn; bolt; bond; brain; brit; bru; 10:00:00:00:01:1e:20:20
alias:      bawn  3,5
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 7,3; 4,5
alias:      zeus  4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
switch:admin>
switch:admin> zoneobjectreplace 11,2 4,8

switch:admin> cfgsave

switch:admin> cfgshow

Defined configuration:
zone: matt zeus; bond; jeff; 4,8
zone: sloth bawn; bolt; bond; brain; 10:00:00:00:01:1e:20:20
alias:      bawn  3,5
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 7,3; 4,5
alias:      zeus  4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
```



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:

```
zonedelete " 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 displays on the other switches to indicate that the transaction was aborted.

Example of deleting zone members

```
switch:admin> cfgshow

Defined configuration:
zone: matt    zeus; bond; jeff; 4,8
zone: sloth   bawn; bolt; bond; brain; brit; bru; 10:00:00:00:01:1e:20:20
alias:        bawn    3,5
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   7,3; 4,5
alias:        zeus    4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
switch:admin>
switch:admin> zonedelete sloth

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

switch:admin>
switch:admin> cfgshow

Defined configuration:
zone: matt    zeus; bond; jeff; 4,8
alias:        bawn    3,5
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   7,3; 4,5
alias:        zeus    4,7; 6,8; 9,2
Effective configuration:
No Effective configuration: (No Access)
```

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.

- 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.

Displaying a zone in a defined configuration example

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
```

Displaying the complete zone database example

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
```

Displaying the sorted zone database example

The following example shows all zones in the zone database with sorted D,l 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.

```
zoneshow[--ic] [" pattern "] [, mode]
```

The following example shows all green zones using pattern search, regardless of the case:

```
switch:admin> zoneshow --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 for and display any differences between the transaction buffer and the committed database by appending the options **--transdiffs** and **--transdiffsonly** to 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, then a plus sign and a minus sign (+-) will be displayed before the member and a * sign will be displayed before the zone name.
- A plus sign (+) before any member of an alias or zone name or any other entity indicates this member has been added, and a minus sign (-) indicates the particular member has been deleted. In the case of TI zones, for inter-fabric links for example, "5,-1" is a valid zone member. Notice that the minus sign (-) comes before the port index. If this was a deleted zone member, it would have been shown as "5,-1". A minus sign (-) before a domain ID indicates that this TI zone member has been deleted.

Example displaying 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
```

Example adding a new zone 'red_zone', deleting "1,1" and adding "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
```

Example cfgShow --transdiffsonly output for the previous example

```
switch:admin> cfgshow --transdiffsonly
*zone: green_zone -1,1; 1,2; +6,15
*zone: +red_zone 5,1; 4,2
switch:admin>
```

Validating a zone

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
      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 the **zone --validate** command to list all zone members that are not part of the current zone enforcement table. Note that zone configuration names are case-sensitive; blank spaces are ignored.

```
switch:admin> zone --validate "White_zone"
```

4. 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** can be used to specify the zone database location. Supported mode flag values are:

- 0 - Zone database from the current transaction buffer
- 1 - Zone database stored from the persistent storage
- 2 - Currently 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 would be expunged in the transaction buffer. This pruning operation always happens on the transaction and defined buffers. You cannot specify a mode option or 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 case distinction

Example validating 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
# - Invalid usage of broadcast zone
```

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 stating that a mismatch is observed between the defined and effective configurations is posted, and you are asked to confirm that you want **cfgSave** to continue. If you enter "y", then the updated configuration will be saved; if you enter "n", then the updated configuration will be discarded.

Example of 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, then the following RASlog message [ZONE-1062] will be 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 then either re-enable the updated configuration or revert to the older configuration. If there is no impact to the effective configuration with the latest update to the defined configuration, then the following message will be displayed.

```
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
If the update includes changes to one or more traffic isolation
zones, you must issue the 'cfgenable' command for the changes
to take effect.
Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no]
```

Example of Inconsistent defined and effective configuration warning to use

```
switch: admin> zoneShow
Defined configuration:
cfg:   cfg1   zone1; zone2
zone:  zone1  10:00:00:00:00:00:01; 10:00:00:00:00:00:02
zone:  zone2  1,1; 1,2

Effective configuration:
cfg:   cfg1
zone:  zone1  10:00:00:00:00:00:01
        10:00:00:00:00:00:02
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  10:00:00:00:00:00:01; 10:00:00:00:00:00:02; 10:00:00:00:00:00:03
zone:  zone2  1,1; 1,2

Effective configuration:
cfg:   cfg1
zone:  zone1  10:00:00:00:00:00:01; 10:00:00:00:00:00:02
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 model.

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** command with one of the following options:

```
defzone --noaccess
```

or

```
defzone --allaccess
```

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> defzone --noaccess

You are about to set the Default Zone access mode to No Access
Do you want to set the Default Zone access mode to No Access ? (yes, y, no, n): [no] y

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, then the current default zone access state will appear as it did prior to the download. For example, if the default zoning mode was No Access before the download, it will remain 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.

To display the zone database size, enter **cfgsize**.

The supported maximum zone database size is 2 MB for a fabric with Brocade DCX 8510 Backbones and Brocade X6 Directors. The presence of any other platform reduces the maximum zone database size to 1 MB. The configuration download and firmware upgrade is blocked when the zone DB size exceeds the allowed size for the fabric in both VF and non-VF modes.

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.

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:

- The merge failed due to fabric limits, resulting in the following output:

```
70 9 6 014640 id N16 Online FC E-Port segmented,10:00:00:27:f8:f0:3a:75
(Fabric Zone DB Size Exceeded)
```


- The merge failed due to chassis limits, resulting in the following output:

```
71 9 7 02ee00 id N16 Online FC E-Port segmented,10:00:00:27:f8:f0:3a:66
(Chassis Zone DB Size Exceeded)
```

HA out-of-sync considerations: Consider the following HA out-of-sync scenarios:

- When the active CP is running 8.0.1 or earlier and the standby CP is running 8.1.0 or later and HA is in sync, if you increase the zone database size of the active CP, the standby CP fails to update and the transaction is aborted with a RASLog message and the CPs will go out of sync.

```
2016/07/08-04:19:58, [ZONE-1028], 201, SLOT 2 FID 8, WARNING, switch_8, Commit zone DB larger than
supported - 2025005 greater than 63614.
```

- When there is no standby CP and the active CP is running Fabric OS 8.0.1 or earlier with the zone database size greater than the maximum supported size, for example, 4 MB (2 MB in LS10 and 2 MB in LS20), if you insert a new standby CP with Fabric OS 8.1.0, the standby CP will be updated with 2 MB on LS10. The zoning transaction on LS20 will be aborted with the following RASLog message. HA will be in sync.

```
2016/07/08-04:19:58, [ZONE-1028], 201, SLOT 2 FID 8, WARNING, switch_8, Commit zone DB larger than
supported - 2025005 greater than 63614.
```

To recover, after `StdbySyncDumpCompleteTimeout` (approximately 15mins), reduce the zone database size in the active CP and reboot the standby CP to bring the HA in sync.

Virtual Fabrics considerations: If Virtual Fabrics is enabled, the sum of the zone database sizes on all of the logical fabrics must not exceed the maximum size allowed for the chassis. The maximum size limit is enforced per-partition and chassis-wide. The chassis-wide restriction is displayed in the output as well.

```
switch:admin> cfgsize
Chassis-Wide Committed Zone DB size - 1086 bytes
Zone DB max size - 2092741 bytes
Available Zone DB size - 2091655 bytes
    committed - 74
    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 unpredictable behavior.

ATTENTION

In a fabric that has a mix of switches running versions of Fabric OS both earlier and later than version 7.0.0, if you run the `cfgsave` or `cfgenable` commands from a switch running a pre-7.0.0 version of Fabric OS, a zone database size of 128 KB might be enforced. Whether this occurs depends on the specific version of pre-Fabric OS 7.0.0 firmware the switches are using. For a detailed list of pre-Fabric OS 7.0.0 firmware compatibility versions, refer to the Zoning Compatibility Note in the Release Notes.

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:

- Number of switches in the fabric.
- Number of bytes for each item name. The number of bytes required for an item name depends on the specifics of the fabric, but cannot exceed 64 bytes for each item.

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 would have to 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.

NOTE

The minimum zone database size is 4 bytes, even if the zone database is empty.

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](#) on page 440.

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 displays 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 displays on 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 displays on the other switches to indicate that the transaction was aborted.

Example

```
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 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 displays 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. If
the update includes changes to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with the traffic isolation zone changes.
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](#) on page 439 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 displays 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 logged on to fabric. This procedure adds an indicator "" for the zone members that do not exist in the fabric, and "~" for the zone members with invalid configuration. The zone server updates the details for zone members in the zone database based on the information received from the Name server.

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 *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 *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.

Validating all zone members example

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
# - Invalid usage of broadcast zone
```

Validating zone members with a zone name example

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
# - Invalid usage of broadcast zone

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
# - Invalid usage of broadcast zone
```

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
# - Invalid usage of broadcast zone
```

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 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.

```
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.
If the update includes changes to one or more traffic isolation
zones, you must issue the 'cfgenable' command for the changes
to take effect.
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 into the switch using an account with administrator privileges.
2. Enter **cfgremove "zone_configuration"** 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 re-enable 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*.

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
```

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, then the committed zone configuration displays.

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.

Example

```
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.

- View the zone aliases using the **zoneshow --alias [-ic] [pattern]** command.

The **zoneshow --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, [0-9] or [a-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> zoneshow --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 "?1?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** command to clear all zone information in the transaction buffer.

ATTENTION

Be careful using the **cfgClear** command because it deletes the defined configuration.

```

switch:admin> cfgclear

The Clear All action will clear all Aliases, Zones, FA 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]

```

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.

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 **zone --copy** 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> zone --copy 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 **zone --expunge** to delete the zone object.

NOTE

Zone configuration names are case-sensitive and blank spaces are ignored.

```
switch:admin> zone --expunge "White_zone"
You are about to expunge one configuration or member. This action could result in removing many
zoning configurations recursively. [Removing the last member of a configuration removes the
configuration.]
Do you want to expunge the member? (yes, y, no, n): [no] yes
```

4. Enter **yes** at the prompt.
5. Enter **cfgShow** to verify the deleted zone object is no longer present.
6. If you want the change preserved when the switch reboots, use **cfgSave** to save it to nonvolatile (flash) memory.
7. 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 nonvolatile (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](#) on page 135, [Configuration file restoration](#) on page 137, or the **configUpload** and **configDownload** commands in the *Brocade Fabric OS Command Reference* 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 (FCS). 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 there is a view conflict, 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 described in [Viewing the configuration in the effective zone database](#) on page 451.

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](#) on page 451 for instructions.

Adding a new fabric that has no zone configuration information to an existing fabric is very 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 will be 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 before merging switches or fabrics:
 - **Defaultzone:** The switches must adhere to the default zone merge rules, as described in [Zone merging scenarios](#) on page 457.
 - **Effective and defined zone configuration match :** Ensure that 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

The fabric is checked for segmentation during power-up, when a switch is disabled or enabled, or when a new switch is added.

The zone configuration database is stored in nonvolatile 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 have implemented default zoning, you must set the switch you are adding into the fabric to the same default zone mode setting as the rest of the fabric to avoid 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, the zoning object is added to the local defined configuration. The modified zone database must fit in the nonvolatile memory area 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 the switches within the merge request.
 - TI zones: If there is an effective configuration between two switches and TI zones are present on either switch, the TI zones are not automatically activated after the merge. Check the TI zone enabled status using the **zone --show** command, and if the status does not match across switches, issue the **cfgEnable** command.
- Merging two fabrics

Both fabrics have identical zones and configurations enabled, including the default zone mode. The two fabrics will join to make one larger fabric with the same zone configuration across the newly created fabric.

If the two fabrics have different zone configurations, they will not be merged. If the two fabrics cannot join, the ISL between the switches will segment.

- Merge conflicts

When a merge conflict is present, a merge will 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* 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 will be segmented.

A merge is not possible if any of the following conditions exist:

- 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

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 defined in the configuration, make sure zone set members are listed in the same order.

NOTE

In a large fabric, especially with 1 MB or more zone configuration takes some amount of time for zone merge. This may cause host device to not to discover the target in other end of the fabric for a short duration. This problem may not be seen if the you follow the guidelines of joining the stable fabric one by one with devices offline and ensure that the zone merge operation is completed before enabling the devices.

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, then the two fabrics merge back into one single fabric. If any changes that cause a conflict have been made to either zone configuration, then the fabrics may segment.

Zone merging scenarios

The following tables provide information on merging zones and the expected results:

- [Table 102](#): Defined and effective configurations
- [Table 103](#): Different content
- [Table 104](#): Different names
- [Table 105](#): TI zones
- [Table 106](#): Default access mode
- [Table 110](#): Mixed Fabric OS versions

TABLE 102 Zone merging scenarios: Defined and effective configurations

Description	Switch A	Switch B	Expected results
Switch A has a defined configuration. Switch B does not have a defined configuration.	defined:cfg1: zone1: ali1; ali2effective: none	defined: none effective: none	Configuration from Switch A to propagate throughout the fabric in an inactive state, because the configuration is not enabled.
Switch A has a defined and effective configuration. Switch B has a defined configuration but no effective configuration.	defined: cfg1zone1: ali1; ali2 effective: cfg1:	defined: cfg1zone1: ali1; ali2 effective: none	Configuration from Switch A to propagate throughout the fabric. The configuration is enabled after the merge in the fabric.
Switch A and Switch B have the same defined configuration. Neither have an effective configuration.	defined: cfg1zone1: ali1; ali2 effective: none	defined: cfg1zone1: ali1; ali2effective: none	No change (clean merge).
Switch A and Switch B have the same defined and effective configuration.	defined: cfg1zone1: ali1; ali2effective: cfg1:	defined: cfg1zone1: ali1; ali2 effective: cfg1:	No change (clean merge).
Switch A does not have a defined configuration. Switch B has a defined configuration.	defined: noneeffective: none	defined:cfg1zone1: ali1; ali2 effective: none	Switch A will absorb the configuration from the fabric.
Switch A does not have a defined configuration. Switch B has a defined configuration.	defined: noneeffective: none	defined:cfg1zone1: ali1; ali2effective: cfg1	Switch A will absorb the configuration from the fabric, with cfg1 as the effective configuration.
Switch A and Switch B have the same defined configuration. Only Switch B has an effective configuration.	defined: cfg1zone1: ali1; ali2effective: none	defined: cfg1zone1: ali1; ali2effective: cfg1	Clean merge, with cfg1 as the effective configuration.
Switch A and Switch B have different defined configurations. Neither have an enabled zone configuration.	defined: cfg2zone2: ali3; ali4 effective: none	defined: cfg1zone1: ali1; ali2effective: none	Clean merge. The new configuration will be a composite of the two.defined: cfg1 zone1: ali1; ali2 cfg2: zone2: ali3; ali4 effective: none

TABLE 102 Zone merging scenarios: Defined and effective configurations (continued)

Description	Switch A	Switch B	Expected results
Switch A and Switch B have different defined configurations. Switch B has an effective configuration.	defined: cfg2zone2: ali3; ali4effective: none	defined: cfg1zone1: ali1; ali2effective: cfg1	Clean merge. The new configuration will be a composite of the two, with cfg1 as the effective configuration.
Switch A does not have a defined configuration. Switch B has a defined configuration and an effective configuration, but the effective configuration is different from the defined configuration.	defined: none effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1 zone1: ali1; ali2 zone2: ali3, ali4	Clean merge. Switch A absorbs the defined configuration from the fabric, with cfg1 as the effective configuration. In this case, however, the effective configurations for Switch A and Switch B are different. You should issue a cfgenable from the switch with the proper effective configuration.

TABLE 103 Zone merging scenarios: Different content

Description	Switch A	Switch B	Expected results
Effective configuration mismatch.	defined: cfg1 zone1: ali1; ali2effective: cfg1 zone1: ali1; ali2	defined: cfg2 zone2: ali3; ali4 effective: cfg2 zone2: ali3; ali4	Fabric segments due to: Zone Conflict cfg mismatch
Configuration content mismatch.	defined: cfg1 zone1: ali1; ali2effective: irrelevant	defined: cfg1 zone1: ali3; ali4 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch

TABLE 104 Zone merging scenarios: Different names

Description	Switch A	Switch B	Expected results
Same content, different effective cfg name.	defined: cfg1 zone1: ali1; ali2 effective: cfg1 zone1: ali1; ali2	defined:cfg2 zone1: ali1; ali2 effective: cfg2 zone1: ali1; ali2	Fabric segments due to: Zone Conflict cfg mismatch
Same content, different zone name.	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone2: ali1; ali2 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same content, different alias name.	defined: cfg1 ali1: A; B effective: irrelevant	defined:cfg1ali2: A; Beffective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same alias name, same content, different order.	defined: cfg1ali1: A; B; Ceffective: irrelevant	defined: cfg1ali1: B; C; Aeffective: 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 105 Zone merging scenarios: TI zones

Description	Switch A	Switch B	Expected results
Switch A does not have Traffic Isolation (TI) zones. Switch B has TI zones.	defined: cfg1 effective: cfg1	defined: cfg1 TI_zone1 effective: cfg1	Clean merge. TI zones are not automatically activated after the merge.
Switch A has TI zones. Switch B has identical TI zones.	defined: cfg1 TI_zone1 effective: cfg1	defined: cfg1 TI_zone1 effective: cfg1	Clean merge. TI zones are not automatically activated after the merge.

TABLE 105 Zone merging scenarios: TI zones (continued)

Description	Switch A	Switch B	Expected results
Switch A has a TI zone. Switch B has a different TI zone.	defined: cfg1 TI_zone1	defined: cfg1 TI_zone2	Fabric segments due to: Zone Conflict cfg mismatch. Cannot merge switches with different TI zone configurations.
Switch A has Enhanced TI zones. Switch B is running Fabric OS v6.4.0 or later.	defined: cfg1 TI_zone1 TI_zone2	defined: none	Clean merge. TI zones are not automatically activated after the merge.
Switch A has Enhanced TI zones. Switch B is running a Fabric OS version earlier than v6.4.0.	defined: cfg1 TI_zone1 TI_zone2	defined: none	Fabric segments because all switches in the fabric must be running Fabric OS v6.4.0 or later to support Enhanced TI zones.

TABLE 106 Zone merging scenarios: Default access mode (pre-Fabric OS 7.3.0)

Description	Switch A	Switch B	Expected results
Different default zone access mode settings.	defzone: allaccess	defzone: noaccess	Clean merge -- noaccess takes precedence and defzone configuration from Switch B propagates to fabric. defzone: noaccess
Same default zone access mode settings.	defzone: allaccess	defzone: allaccess	Clean merge -- defzone configuration is allaccess in the fabric.
	defzone: noaccess	defzone: noaccess	Clean merge -- defzone configuration is noaccess in the fabric.
Effective zone configuration.	No effective configuration.defzone = allaccess	effective: cfg2 defzone: allaccess or noaccess	Clean merge -- effective zone configuration and defzone mode from Switch B propagates to fabric.
Effective zone configuration.	No effective configuration.defzone = noaccess	effective: cfg2 defzone: allaccess	Fabric segments because Switch A has a hidden zone configuration (no access) activated and Switch B has an explicit zone configuration activated.
Effective zone configuration	effective: cfg1 defzone: noaccess	No effective configuration. defzone: noaccess	Clean merge -- effective zone configuration from Switch A propagates to fabric.
Effective zone configuration	effective: cfg1 defzone: allaccess	No effective configuration. defzone: noaccess	Fabric segments. You can resolve the zone conflict by changing defzone to noaccess on Switch 1 .

TABLE 107 Zone merging scenarios: Default access mode (with Fabric OS 7.3.0 or later on initiator and responder)

Description	Switch A (Initiator with FOS 7.3.0)	Switch B (Responder with FOS 7.3.0)	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.

TABLE 107 Zone merging scenarios: Default access mode (with Fabric OS 7.3.0 or later on initiator and responder) (continued)

Description	Switch A (Initiator with FOS 7.3.0)	Switch B (Responder with FOS 7.3.0)	Expected results
	defzone: noaccess	defzone: noaccess	Clean merge -- defzone configuration is noaccess in the fabric.
Initiator has effective zone configuration	effective: cfg1 defzone: allaccess	No effective configuration. defzone: allaccess	Clean merge -- effective zone configuration from Switch A propagates to fabric.
	effective: cfg1 defzone: allaccess	No effective configuration. defzone: noaccess	Fabric segments because Switch B has a hidden zone configuration (no access) activated and Switch A has an explicit zone configuration activated.
	effective: cfg1 defzone: noaccess	No effective configuration. defzone: allaccess	Fabric merges -- effective zone configuration from Switch A propagates to fabric. Allaccess on Switch B changes to noaccess.
	effective: cfg1 defzone: noaccess	No effective configuration. defzone: noaccess	Fabric merges -- effective zone configuration from Switch A propagates to fabric.
Responder has effective zone configuration.	No effective configuration. defzone: allaccess	effective: cfg2 defzone: allaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric.
	No effective configuration. defzone: allaccess	effective: cfg2 defzone: noaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric. Allaccess on Switch A changes to noaccess.
	No effective configuration. defzone: noaccess	effective: cfg2 defzone: allaccess	Fabric segments because Switch A has a hidden zone configuration (no access) activated and Switch B has an explicit zone configuration activated.
	No effective configuration. defzone: noaccess	effective: cfg2 defzone: noaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric.

TABLE 108 Zone merging scenarios: Default access mode (with Fabric OS 7.3.0 or later on initiator and pre-Fabric OS 7.3.0 on responder)

Description	Switch A (Initiator with FOS 7.3.0)	Switch B (Responder with pre-FOS 7.3.0)	Expected results
Different default zone access mode settings.	defzone: allaccess	defzone: noaccess	Fabric merges with allaccess on Switch A changes to noaccess.
	defzone: noaccess	defzone: allaccess	Fabric merges with allaccess on Switch B changes to noaccess.
Same default zone access mode settings.	defzone: allaccess	defzone: allaccess	Clean merge -- defzone configuration is allaccess in the fabric.
	defzone: noaccess	defzone: noaccess	Clean merge -- defzone configuration is noaccess in the fabric.
Initiator has effective zone configuration	effective: cfg1 defzone: allaccess	No effective configuration. defzone: allaccess	Clean merge -- effective zone configuration from Switch A propagates to fabric.

TABLE 108 Zone merging scenarios: Default access mode (with Fabric OS 7.3.0 or later on initiator and pre-Fabric OS 7.3.0 on responder) (continued)

Description	Switch A (Initiator with FOS 7.3.0)	Switch B (Responder with pre-FOS 7.3.0)	Expected results
	effective: cfg1 defzone: allaccess	No effective configuration. defzone: noaccess	Fabric segments because Switch B has a hidden zone configuration (no access) activated and Switch A has an explicit zone configuration activated.
	effective: cfg1 defzone: noaccess	No effective configuration. defzone: allaccess	Fabric merges -- effective zone configuration from Switch A propagates to fabric. Allaccess on Switch B changes to noaccess.
	effective: cfg1 defzone: noaccess	No effective configuration. defzone: noaccess	Fabric merges -- effective zone configuration from Switch A propagates to fabric.
Responder has effective zone configuration.	No effective configuration. defzone: allaccess	effective: cfg2 defzone: allaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric.
	No effective configuration. defzone: allaccess	effective: cfg2 defzone: noaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric. Allaccess on Switch A changes to noaccess.
	No effective configuration. defzone: noaccess	effective: cfg2 defzone: allaccess	Fabric segments because Switch A has a hidden zone configuration (no access) activated and Switch B has an explicit zone configuration activated.
	No effective configuration. defzone: noaccess	effective: cfg2 defzone: noaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric.

TABLE 109 Zone merging scenarios: Default access mode (with pre-Fabric OS 7.3.0 on initiator and Fabric OS 7.3.0 or later on responder)

Description	Switch A (Initiator with pre-FOS 7.3.0)	Switch B (Responder with FOS 7.3.0)	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.
	defzone: noaccess	defzone: noaccess	Clean merge -- defzone configuration is noaccess in the fabric.
Initiator has effective zone configuration	effective: cfg1 defzone: allaccess	No effective configuration. defzone: allaccess	Clean merge -- effective zone configuration from Switch A propagates to fabric.
	effective: cfg1 defzone: allaccess	No effective configuration. defzone: noaccess	Fabric segments because Switch B has a hidden zone configuration (no access) activated and Switch A has an explicit zone configuration activated.
	effective: cfg1	No effective configuration.	Fabric merges -- effective zone configuration from Switch A

TABLE 109 Zone merging scenarios: Default access mode (with pre-Fabric OS 7.3.0 on initiator and Fabric OS 7.3.0 or later on responder) (continued)

Description	Switch A (Initiator with pre-FOS 7.3.0)	Switch B (Responder with FOS 7.3.0)	Expected results
	defzone: noaccess	defzone: allaccess	propagates to fabric. Allaccess on Switch B changes to noaccess.
	effective: cfg1 defzone: noaccess	No effective configuration. defzone: noaccess	Fabric merges -- effective zone configuration from Switch A propagates to fabric.
Responder has effective zone configuration.	No effective configuration. defzone: allaccess	effective: cfg2 defzone: allaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric.
	No effective configuration. defzone: allaccess	effective: cfg2 defzone: noaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric. Allaccess on Switch A changes to noaccess.
	No effective configuration. defzone: noaccess	effective: cfg2 defzone: allaccess	Fabric segments because Switch A has a hidden zone configuration (no access) activated and Switch B has an explicit zone configuration activated.
	No effective configuration. defzone: noaccess	effective: cfg2 defzone: noaccess	Fabric merges -- effective zone configuration from Switch B propagates to fabric.

TABLE 110 Zone merging scenarios: Mixed Fabric OS versions

Description	Switch A	Switch B	Expected results
Switch A is running Fabric OS 7.0.0 or later. Switch B is running a Fabric OS version earlier than 7.0.0.	effective: cfg1 defzone = allaccess	No effective configuration. defzone - noaccess	Fabric segments due to zone conflict.
Switch A is running Fabric OS 7.0.0 or later. Switch B is running a Fabric OS version earlier than 7.0.0.	No effective configuration. defzone = noaccess	effective: cfg2 defzone - allaccess	Fabric segments due to zone conflict.

NOTE

When merging mixed versions of Fabric OS where both sides have default zone mode No Access set, the merge results vary depending on which switch initiates the merge.

Concurrent zone transactions

While working on zone sets, a special workspace is provided to allow you to manipulate the zone sets of your choice. These changes are not put into effect until they are committed to the database. Once they are committed, they will replace the existing active zone sets with the new zone sets or create more zone sets in the defined database. When updates to the zoning database are being made by multiple users, Fabric OS warns you about the situation and allows you to choose which operation will prevail.

NOTE

The concurrent zone transactions are tracked only for Fabric OS CLI users. Other users using the REST API, WebTools, or Brocade Network Advisor or other methods are not tracked.

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 shown.

of what is shown if there is not a pending zoning transaction in the fabric

```
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. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
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 *Fabric OS Command Reference* for more details on this command.

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.

```
switch:admin> cfgtransshow --opentrans

Current transaction token is 0x3109
It is abortable
Transactions Detect: Capable
```

```
Current Open Transactions
Domain List:
-----
1  2  3  4
```

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 "non-principal" 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 exceeds zone database size limits. As the number of zones increase, 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 less effective use of hardware resources and greater RSCN traffic.

In a Peer Zone setup, principal to non-principal device communication is allowed, but non-principal to non-principal device communication and principal to principal device communication are 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 non-principal 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 storage area network (SAN):

FIGURE 54 Single-Initiator Zoning example




Single-Initiator Zoning: Single-Initiator Zoning eliminates the host-to-host visibility, which reduces RSCN traffic. However, the main limitation of Single-Initiator Zoning is that this type of zoning results in creating a large number of zones, which can lead to exceeding the capacity of the zone database size limits. The [Figure 54](#) shows an example of Single-Initiator Zoning with eight zones; one for each host H1 to H8.

FIGURE 55 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 [Figure 55](#) shows 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 56 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 to this type of zoning is the significant increase in RSCN traffic when compared to other zoning types. The [Figure 56](#) shows an example of Flat Zoning where the hosts H1 to H16 and targets T1 and T2, all form a single zone and can communicate with all other devices in the zone.

FIGURE 57 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 non-principal 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 non-principal 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 creates a fewer number of zones. The [Figure 57](#) shows 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 types:

- 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:

- Non-principal devices can communicate only with the principal device.
- Non-principal devices cannot communicate with other non-principal devices, unless allowed by some other zone in the active zone set.
- 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. Refer to [Zone database size](#) on page 440 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 will be 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 and QoS Peer Zoning are supported.

LSAN Peer Zoning is supported on Edge to edge and Edge to backbone topologies. Once 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. Refer to [QoS](#) on page 321 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](#) on page 470.

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"/
-mbrs "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 **zonedelelete** command.

The following example deletes the "peerzone_di_mbrs" Peer Zone.

```
switch:admin> zonedelelete peerzone_di_mbrs
```

Viewing Peer Zones

To view a Peer Zone configuration, including the devices associated with the Peer Zone, complete the following steps.

1. Connect to the switch and log in to an account.

2. View the Peer Zone configuration 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.

For more information about the commands, refer to the *Brocade Fabric OS Command Reference*.

The following example displays the Peer Zone configuration using the **all** option to show all the Peer Zones. Property member field is used to distinguish a Peer Zone from a regular zone. Peer Member (s) field lists the non-principal members of the Peer Zone.

```
switch:admin> zoneshow --peerzone all
Defined configuration:
zone: peerzone_di_mbrs
  Property Member: 00:02:00:00:00:02:00:01
  Created by: User
  Principal Member(s):
    10,1
  Peer Member(s):
    20,1; 20,2
zone: peerzone_wnn_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

Effective configuration:
zone: peerzone_wnn_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

1 Peer Zones in Eff Cfg
```

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 the 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, if the CP is then made to sync with the active CP running a version of Fabric OS that supports alias peer zoning and contain alias peer zones, the standby CP will not sync until either all alias peer zones are removed from the active CP or the standby is upgraded to a version of Fabric OS that supports alias peer zones.

TABLE 111 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.

TABLE 111 Zone merge results when using alias peer zoning (continued)

Existing fabric configuration	Joining fabric configuration	Zone merge result
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 then it will have a restriction from having mixed type members (i.e., D, I versus WWN). Mixed members will be allowed only after removing that alias from all the peer zones the alias is part of.
- 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 lower.
- 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 00:02:00:00:00:02:05:09; 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 configuration to Fabric OS 8.0.1 or lower is blocked if alias peer zones are present.
- Modifying members of an alias to have mixed members are 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.

```
device#> zonecreate --peer p1 -p "ali1;8,7;9,9" -m "4,5";
```

```
device#> 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
```

```
device#> cfgcreate cp1,p1
device#> cfgenable cp1
```

```
device#> zoneadd --peer p1 -p "9,7;ali2";
```

```
device#> 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
```

```
device#> zoneremove --peer p1 -p "ali2";
```

```
device#> 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
```

```
device#> 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

device#> zoneshow --verbose
Defined configuration:
  cfg: c1 p8
  zone: p8 00:02:00:00:00:02:02:01; ali5; 86,8

device#> 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.

device#> 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

device#> 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**

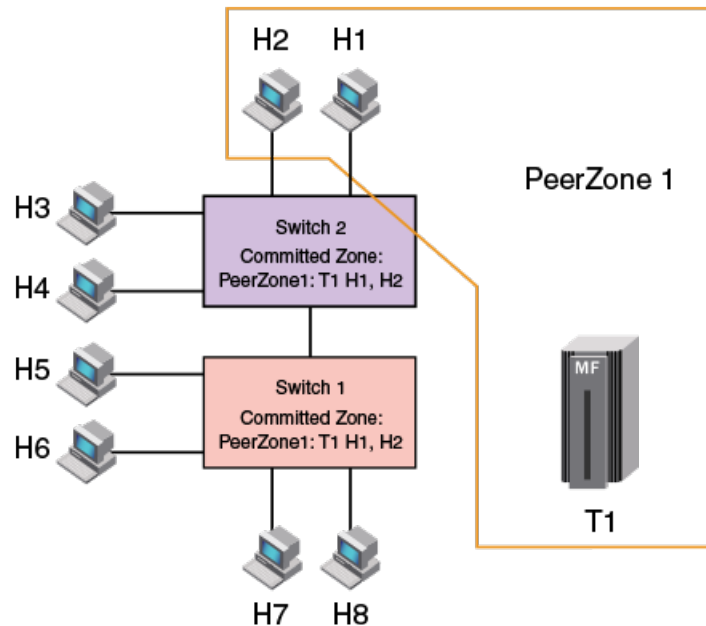
For more details, refer to the *Brocade Fabric OS Command Reference*.

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 58 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 adds or activates a Target Driven Peer Zone. If a zone configuration does not exist or effective zone configuration is disabled, while committing the first set of TDZ commands on a switch automatically creates a default zone configuration with the name "tdz__cfg__auto" if the zone database is empty. If any zone object exists in the defined database, the zone configuration is not automatically created.
 - "tdz__cfg__auto" is created if a AAPZ or RAPZ request is initiated by the target when the zoneDB is completely empty. A dummy zone, "tdz__zone__auto" is created when an "tdz_cfg_auto" is created. This dummy zone allows 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 Brocade Network Advisor (BNA) or the CLI; however, they can be displayed or deleted using BNA or the 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, add and enable the zone in the active configuration. If there is no configuration enabled and the zone DB is completely empty, then a "tdz__cfg__auto" is automatically created. For procedures to create, add, and enable zones in a zone database, refer to [Zone configurations](#) on page 419.

Brocade recommends that all Target Driven Zoning tasks be performed 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 if there is an enabled zone configuration, no action is required. If there is no enabled zone configuration, 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: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. If the update includes changes  
to one or more traffic isolation zones, the update may result in  
localized disruption to traffic on ports associated with  
the traffic isolation zone changes  
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 will get 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
```

- (Optional) View the principal device and the non-principal devices in a zone using the **zonestow --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> zonestow --peerzone all
Defined configuration:
zone: peerzone_wnw_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_wnw_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> zonestow --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 non-principal member list in addition to WWN members if all the switches in the fabric are running Fabric OS 8.1.x or later. If you merge a domain running Fabric OS 8.0.0 or lower with a domain running Fabric OS 8.1.x or later and have alias target driven zones configured, the link gets segmented.

In case of chassis systems, both the CPs must be running Fabric OS 8.1.x 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 contain 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 112 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.
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 non-principal members of a target driven zone can be either of the following:
 - Only WWN members.
 - Aliases with only WWN members.
 - Aliases with WWN members and WWN members.
- The following non-principal 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 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 8.0.0 or lower is blocked if alias target driven zones are present.

Supported commands for Peer Zones and Target Driven Peer Zones

The regular zone commands are used to configure and manage the 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. Refer to the following table for the list of commands supported for Peer Zones and Target Driven Peer Zones.

TABLE 113 Supported commands for Peer Zones and Target Driven Peer Zones

Command	Support for Peer Zones	Support for Target Driven Peer Zones
cfgAdd	Yes	Yes
cfgCreate	Yes	Yes
cfgDelete	Yes	Yes
cfgDisable	Yes	Yes
cfgEnable	Yes	Yes
cfgRemove	Yes	Yes
cfgSave	Yes	Yes
cfgShow	Yes	Yes
zone	No	No
zoneAdd	Yes, using <code>--peerzone</code> operand	No
zoneCreate	Yes, using <code>--peerzone</code> operand	No
zoneDelete	Yes	Yes
zoneObjectCopy	Yes	No
zoneObjectExpunge	No	No
zoneObjectRename	Yes	No
zoneObjectReplace	Yes	No
zoneRemove	Yes, using <code>--peerzone</code> operand	No
zoneShow	Yes, using <code>--peerzone</code> operand	Yes, using <code>--peerzone</code> operand

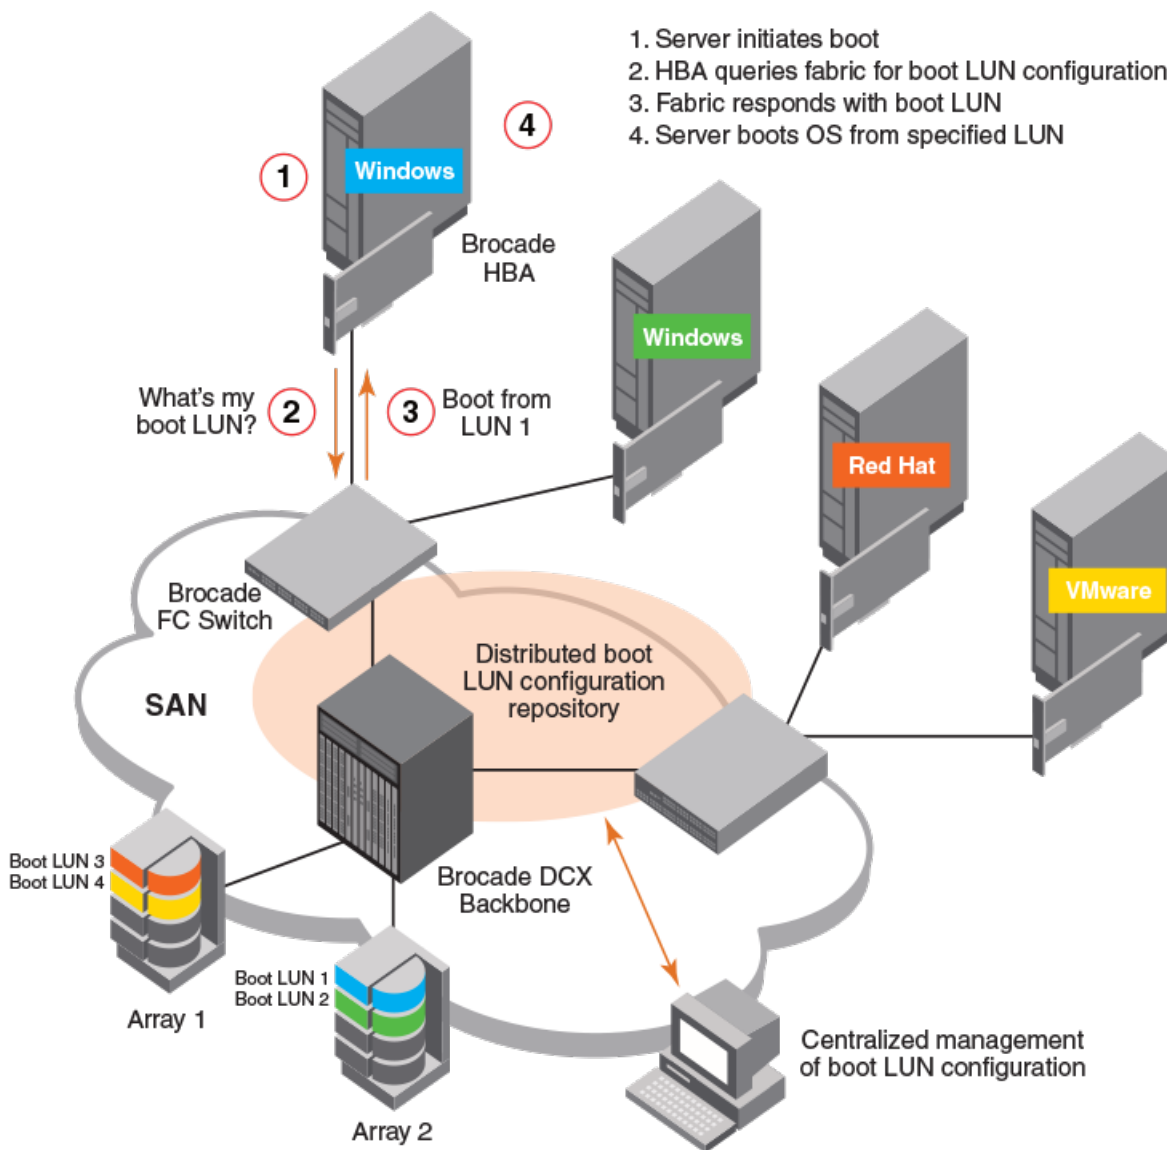
Boot LUN zoning

Boot from SAN provides rapid server provisioning, rapid recovery from server failure, improved reliability, and reduced power and cooling. Brocade/QLLogic Host Bus Adapters (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 only using Brocade/QLLogic HBAs; 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 59 Boot LUN zoning



As Boot LUN zones must not be part of a fabric zone configuration, Fabric OS 7.3.0 and later 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. As part of the support for peer zoning introduced in Fabric OS 7.3.0, restrictions to protect property members starting with "00:" were added. In Fabric OS 8.0.1 and later, this restriction has been extended to exclude any WWNs that start with "00:00:00:" as this WWN value indicates a Boot LUN zone member.

In Fabric OS 8.0.1 and later, zone configuration rules now only permit creating a Boot LUN zone using the **bootluncfg** command. To enable a Boot LUN configuration, you will need to create a basic zone for the host-target pair using the **bootluncfg** command, and then add this zone object to the zone configuration. This will ensure that the HBA has access to the target. Once 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 *Fabric OS Command Reference* 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. Note that Brocade Fabric OS 6.2 or later and HBA BIOS version 1.1 or later are required, and we recommend always using the latest BIOS and FOS version. Other than that, no additional hardware or extra license or software is needed to deploy Fabric-Based Boot LUN Discovery. Use the FOS Command-Line Interface (CLI) for configuration: neither Brocade DCFM nor Brocade Web Tools are required.

Configuring the server and HBA for 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 60 Boot LUN zoning BIOS setting



Configuring the Brocade switch for 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

```

- 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
```

- 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
```

- 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>
```

- 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 Oxff:

```
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"
```

Traffic Isolation Zoning

• Traffic Isolation Zoning overview.....	484
• FSPF routing rules and traffic isolation.....	486
• Enhanced TI zones.....	488
• General rules for TI zones.....	491
• Limitations and restrictions of Traffic Isolation Zoning.....	493
• Creating a TI zone.....	493
• Modifying TI zones.....	494
• Changing the state of a TI zone.....	495
• Deleting a TI zone.....	495
• Displaying TI zones.....	496
• Enforcing Local TI Filtering.....	497
• Fabric-Level Traffic Isolation in a backbone fabric.....	501
• Virtual Fabrics considerations for Traffic Isolation Zoning.....	509
• Traffic Isolation Zoning over FC routers with Virtual Fabrics.....	511
• Troubleshooting TI zone routing problems.....	513

Traffic Isolation Zoning overview

Traffic Isolation Zoning allows you to control the flow of inter-switch traffic by creating a dedicated path for traffic flowing from a specific set of source ports (N_Ports).

You might use Traffic Isolation Zoning for the following scenarios:

- To dedicate an ISL to high priority, host-to-target traffic.
- To force high volume, low priority traffic onto a given ISL to limit the effect on the fabric of this high traffic pattern.
- To ensure that requests and responses of FCIP-based applications such as tape pipelining use the same VE_Port tunnel across a metaSAN.

NOTE

Traffic Isolation Zoning is an advanced feature, but does not require a license. A strong understanding of Fabric Shortest Path First (FSPF) is required.


Isolating traffic requires that you create a special zone, called a *Traffic Isolation zone* (TI zone). A TI zone indicates that a set of N_Ports and E_Ports are to be used only for a specific isolated traffic flow. When a TI zone is activated, the fabric attempts to isolate all inter-switch traffic entering the fabric from a member of the zone to only those E_Ports that have been included in the zone. The fabric also attempts to exclude traffic not in the TI zone from using E_Ports within that TI zone.

The following figure shows a fabric with a TI zone consisting of the following:

- N_Ports: "1,7", "1,8", "4,5", and "4,6"
- E_Ports: "1,1", "3,9", "3,12", and "4,7"

The dotted lines indicate the dedicated path between the initiator in Domain 1 to the target in Domain 4.

FIGURE 61 Traffic Isolation zone creating a dedicated path through the fabric



In the figure, all traffic entering Domain 1 from N_Ports 7 and 8 is routed through E_Port 1. Similarly, traffic entering Domain 3 from E_Port 9 is routed to E_Port 12, and traffic entering Domain 4 from E_Port 7 is routed to the devices through N_Ports 5 and 6. Traffic coming from other ports in Domain 1 would *not* use E_Port 1, but would use E_Port 2 instead.

TI zone failover

A TI zone can have failover enabled or disabled.

Disable failover if you want to guarantee that TI zone traffic uses only the dedicated path, and that no other traffic can use the dedicated path. In other words, an exclusive, members-only path.

Enable failover if you want traffic to have alternate routes when the dedicated path cannot be used, and if you want other traffic to be able to use the dedicated path when the non-dedicated paths are not available.

ATTENTION

If failover is disabled, use care when planning your TI zones so that non-TI zone devices are not isolated. If this feature is not used correctly, it can cause major fabric disruptions that are difficult to resolve. Refer to [Additional considerations when disabling failover](#) on page 486 for additional information about using this feature.

The following table compares the behavior of traffic when failover is enabled and disabled.

TABLE 114 Traffic behavior when failover is enabled or disabled in TI zones

Failover enabled	Failover disabled
If the dedicated path is not the shortest path or if the dedicated path is broken, the TI zone traffic will use a non-dedicated path instead.	If the dedicated path is not the shortest path or if the dedicated path is broken, traffic for that TI zone is halted until the dedicated path is fixed. This condition initiates a RASlog.

TABLE 114 Traffic behavior when failover is enabled or disabled in TI zones (continued)

Failover enabled	Failover disabled
Non-TI zone traffic will use the dedicated path if no other paths through the fabric exist or if the non-dedicated paths are not the shortest paths.	Non-TI zone traffic will never use the dedicated path, even if the dedicated path is the shortest path or if there are no other paths through the fabric. This condition initiates a RASlog.

For example, in [Figure 61](#) on page 485, if the dedicated ISL between Domain 1 and Domain 3 goes offline, then the following occurs, depending on the failover option:

- If failover is disabled for the TI zone, the TI zone traffic is halted until the ISL between Domain 1 and Domain 3 is back online.
- If failover is enabled for the TI zone, the TI zone traffic is routed from Domain 1 to Domain 3 through E_Ports "1,2" and "3,10". Also, when TI zone traffic enters the non-TI path, the TI zone traffic continues to flow through that path. In this example, when the TI zone traffic is routed through E_Ports "1,2" and "3,10", that traffic continues through the non-TI path between domains 3 and 4, even though the TI path between domains 3 and 4 is not broken.

If the non-dedicated ISL between Domain 1 and Domain 3 goes offline, then the following occurs, depending on the failover option:

- If failover is disabled for the TI zone, non-TI zone traffic is halted until the non-dedicated ISL between Domain 1 and Domain 3 is back online.
- If failover is enabled for the TI zone, non-TI zone traffic is routed from Domain 1 to Domain 3 through the dedicated ISL. Also, when non-TI zone traffic enters the TI path, the non-TI zone traffic continues to flow through that path. In this example, when the non-TI zone traffic is routed through E_Ports "1,1" and "3,9", that traffic continues through E_Ports "3,12" and "4,7", even though the non-dedicated ISL between domains 3 and 4 is not broken.

Additional considerations when disabling failover

If failover is disabled, be aware of the following considerations:

- Disabling failover locks the specified route so that only TI zone traffic can use it. Non-TI zone traffic is excluded from using the dedicated path. Ensure that there are also non-dedicated paths through the fabric for all devices that are not in a TI zone.
- If you create a TI zone with E_Ports only, and no N_Ports, failover must be enabled. If failover is disabled, the specified ISLs will not be able to route any traffic.
- It is recommended that TI zone definitions and regular zone definitions match. For example, if four N_Ports are TI zoned together, then these same N_Ports should also be regularly zoned together.
- It is strongly recommended that the insistent Domain ID feature be enabled. TI zone members are defined by Domain ID and Port Index. If a switch changes its active domain ID, the route defined by the TI zone becomes unknown, breaking the dedicated path. Refer to the **configure** command in the *Fabric OS Command Reference* for information about setting insistent Domain ID.

FSPF routing rules and traffic isolation

All traffic must use the lowest cost path. FSPF routing rules take precedence over the TI zones, as described in the following situations.

If the dedicated ISL is not the lowest cost path ISL, then the following rules apply:

- If failover is enabled, the traffic path for the TI zone is broken, and TI zone traffic uses the lowest cost path instead.
- If failover is disabled, the TI zone traffic is blocked.


If the dedicated ISL is the only lowest cost path ISL, then the following rules apply:

- If failover is enabled, non-TI zone traffic as well as TI zone traffic uses the dedicated ISL.

- If failover is disabled, non-TI zone traffic is blocked because it cannot use the dedicated ISL, which is the lowest cost path.


In the following figure, there is a dedicated path between Domain 1 and Domain 3, and another non-dedicated path that passes through Domain 2. If failover is enabled, *all* traffic will use the dedicated path, because the non-dedicated path is not the shortest path. If failover is disabled, non-TI zone traffic is blocked, because the non-dedicated path is not the shortest path.

FIGURE 62 Dedicated path is the only shortest path



In the following figure, a dedicated path between Domain 1 and Domain 4 exists, but is not the shortest path. In this situation, if failover is enabled, the TI zone traffic uses the shortest path, even though the E_Ports are not in the TI zone. If failover is disabled, the TI zone traffic stops until the dedicated path is configured to be the shortest path.

FIGURE 63 Dedicated path is not the shortest path




For information about setting or displaying the FSPF cost of a path, refer to the `linkCost` and `topologyShow` commands in the *Brocade Fabric OS Command Reference*.

Enhanced TI zones

In Fabric OS 6.4.0 and later, ports can be in multiple TI zones at the same time. Zones with overlapping port members are called *enhanced TI zones* (ETIZ). Enhanced TI zones are especially useful in FICON fabrics.

The following figure shows an example of two TI zones. Because these TI zones have an overlapping port (3,8), they are enhanced TI zones.

FIGURE 64 Enhanced TI zones



Refer to the *Brocade FICON Administration Guide* for example topologies using enhanced TI zones.

Refer to [Additional configuration rules for enhanced TI zones](#) on page 492 for more information about enhanced TI zones.

Invalid configurations with enhanced TI zones

When you create TI zones, ensure that all traffic from a port to all destinations on a remote domain have the same path. Do not create separate paths from a local port to two or more ports on the same remote domain. If failover is disabled with such configurations, traffic loss can occur.

A message is sent to the RASlog if this potential error condition is detected in the TI zone configuration. You can also display a report of existing and potential problems with TI zone configurations, as described in [Troubleshooting TI zone routing problems](#) on page 513.

The following cases are examples of invalid enhanced TI zone configurations.


Invalid ETIZ configuration: separate paths from a port to devices on same domain

The following figure shows two enhanced TI zones (failover disabled) that are configured incorrectly because there are two paths from a local port (port 8 on Domain 3) to two or more devices on the same remote domain (ports 1 and 4 on Domain 1).

The TI zones are enhanced TI zones because they have an overlapping member (3,8). Each zone describes a different path from the Target to Domain 1. In this case, traffic is routed correctly from Host 1 and Host 2 to the Target; however, traffic from the Target to the Hosts would be dropped in the following manner.

Traffic from (3,8) destined for Domain 1 cannot go through both port 6 and port 7. That is, two exclusive paths have been defined, but only one path will be chosen. If port 6 is chosen (ETIZ 1 path), then frames destined for the Host at Domain 1, N_Port 4 will be dropped at Domain 1, E_Port 2, because this Host is not a member of the ETIZ 1 path. Alternatively, if port 7 is chosen (ETIZ 2 path), then frames destined for the Host at Domain 1, N_Port 1 will be dropped since it is not a member of the ETIZ 2 path.

FIGURE 65 Invalid ETIZ configuration: two paths from one port to two devices on the same remote domain



The solution is to overlap the E_port members [(1,2), (1,3), (3,6), (3,7)] across both ETIZ zones, in addition to the Target at (3,8).

Invalid ETIZ configuration: separate paths from a single port to the same domain


The following figure shows another example of an invalid ETIZ configuration. In this example, the two hosts are on separate remote domains, but the path to each host goes through the same domain (Domain 1).

This example contains two enhanced TI zones, with port (3,8) as the overlapping member:

- ETIZ 1 contains (1,1), (1,2), (3,6), (3,8)
- ETIZ 2 contains (2,1), (2,2), (1,4), (1,3), (3,7), (3,8)

Both TI zones describe two different paths from the Target on Domain 3 to Hosts on Domain 1 and Domain 2. Traffic from port (3,8) cannot be routed to Domain 1 over both (3,6) and (3,7), only one port/path will be chosen due to the nature of TI zoning. As one example, if E_Port (3,7) is chosen (ETIZ 2 path), frames destined for the Host at (1,1) will be dropped at E_Port (1,3) because said Host is not a member of ETIZ 2. Alternatively, if ETIZ 1 path is chosen, i.e., E_Port (3,6), then traffic to Host 1 at (1,1) can pass, but traffic destined for Host 2 on Domain 2 would be dropped at E_Port (1,2), since (1,4) is not a member of ETIZ 1.

FIGURE 66 Invalid ETIZ configuration: two paths from one port



The solution is to overlap the E_Port members [(1,2),(1,3), (3,6), (3,7)] across both ETIZ zones, in addition to the Target at (3,8); or, overlap all members of both ETIZ zones in this case.

General rules for TI zones


The following general rules apply to TI zones:

- Ports in a TI zone must belong to switches that run Fabric OS v6.0.0 or later. For TI over FCR zones, all switches and FC routers in both edge and backbone fabrics must be running Fabric OS v6.1.0 or later.
- For the FC8-64 and FC16-64 blades in the Brocade DCX and Brocade DCX 8510-8, ports 48-63 can be in a TI zone, only if all switches in that TI zone are running Fabric OS v6.4.0 or later. Ports 48-63 can still be in a failover path for TI traffic. The Brocade DCX-4S and Brocade DCX 8510-4 do not have this limitation.
- VE_Ports can be members of TI zones.
- A TI zone must include E_Ports and N_Ports that form a complete, end-to-end route from initiator to target.
- When an E_Port is a member of a TI zone that E_Port cannot have its index swapped with another port.
- If multiple E_Ports are configured that are on the lowest cost route to a domain, the various source ports for that zone are load-balanced across the specified E_Ports.
- TI zones reside only in the defined configuration and not in the effective configuration. When you make any changes to TI zones, including creating or modifying them, you must enable the effective configuration for the changes to take effect, even if the effective configuration is unchanged.
- A TI zone only provides traffic isolation and is not a "regular" zone.
- Routing rules imposed by TI zones with failover disabled override regular zone definitions. Regular zone definitions should match TI zone definitions.

- FSPF supports a maximum of 16 paths to a given domain. This includes paths in a TI zone.
- To include a trunk group in a TI zone, you must include all ports of the trunk in the TI zone.
- Each TI zone is interpreted by each switch and each switch considers only the routing required for its local ports. No consideration is given to the overall topology and to whether the TI zones accurately provide dedicated paths through the whole fabric.

For example, in the following figure the TI zone was configured incorrectly and E_Port "3,9" was erroneously omitted from the zone. The domain 3 switch assumes that traffic coming from E_Port 9 is *not* part of the TI zone and so that traffic is routed to E_Port 11 instead of E_Port 12, if failover is enabled. If failover is disabled, the route is broken and traffic stops.

FIGURE 67 TI zone misconfiguration



Additional configuration rules for enhanced TI zones

Enhanced TI zones (ETIZ) have the following additional configuration rules:

- Enhanced TI zones are supported only if every switch in the fabric is ETIZ capable. A switch is ETIZ capable if it meets the following qualifications:
 - The switch must be one of the supported platforms, as listed in the Supported Platforms section.
 - The switch must be running Fabric OS v6.4.0 or later.
- If the fabric contains a switch running an earlier version of Fabric OS, you cannot create an enhanced TI zone. You cannot merge a downlevel switch into a fabric containing enhanced TI zones, and you cannot merge a switch with enhanced TI zones defined into a fabric containing switches that do not support ETIZ.
- Overlapping TI zones must have the same failover type. That is, both must be either failover enabled or failover disabled.

NOTE

FC router domains are excluded from the ETIZ platform restrictions. You can create enhanced TI zones with these switches in the fabric.

Limitations and restrictions of Traffic Isolation Zoning

The following limitations and restrictions apply to Traffic Isolation zoning:

- A TI zone name should not be longer than 63 characters.
- A maximum of 255 TI zones can be created in one fabric. A fabric merge resulting in greater than the maximum allowed TI zones results in merge failure and the fabrics are segmented.
- A TI zone can be created using D,I (Domain, Index) notation only, except for TI zones in a backbone fabric, which use port WWNs. Refer to [Traffic Isolation Zoning over FC routers](#) on page 497 for information about TI zones in a backbone fabric.
- To include a trunk group in a TI zone, you must include all ports of the trunk in the TI zone.
- If two N_Ports are online and have the same shared area, and one of them is configured in a TI zone, then they both must be configured in that same TI zone. One of the online shared area N_Ports should not remain outside the TI zone unless it is offline, then it may remain outside the TI zone.
- Port swapping is not permitted in TI zones.
- Ports that are in different TI zones cannot communicate with each other if failover is disabled.
- TI zone members that overlap must have the same TI failover policy across all TI zones to which they belong. That is, if an overlapping member is part of a failover-disabled zone, then it can belong only to other TI zones where the policy is also failover-disabled; the member cannot overlap with failover-enabled TI zones.
- TI zones that have members with port index greater than 511 are not supported with Fabric OS versions earlier than 6.4.0. If such a TI zone and Fabric OS version combination is detected, a warning is issued. These configurations are not prevented, but their behavior is unpredictable.
- When you merge two switches, if there is an effective configuration on the switches and TI zones are present on either switch, the TI zones are not automatically activated after the merge. Check the TI zone enabled status using the **zone --show** command, and if the TI Zone Enabled status does not match across switches, issue the **cfgenable** command.
- Use care when creating TI zones on ICL ports in topologies that span more than two switches connected with ICLs. If a user-defined TI zone breaks the ICL connectivity requirements, a FSPF-1009 RASLog entry and message is generated to notify you of this error condition.
- In a backbone switch, if a TI zone is configured with the remote switch EX_Port even though locally the same ASIC EX_Port exists to reach the destination edge fabric, then PLOGI takes the same local ASIC EX_Port to send the frame out instead of TI Zone EX_Port. Thus, the PLOGI frame is sent out to the edge fabric without translation and gets dropped. To resolve this issue, move the connected destination EX_Port to different ASICs to route based on the configured TI zone.
- When an edge fabric is connected to both a local FC router and a remote FC router, you must configure a TI zone only on the local FC router EX_Port. If you have already configured the FC router TI zone on the remote FC router EX_Port, disable the local FC router EX_Port that is connected to the edge fabric and restart the discovery.
- If an ETH port is configured in the TI zone, then the TI zone must be configured as failover-enabled. Configuring such as TI zone as failover-disabled is not prevented, but its behavior is not predictable.

Creating a TI zone

You can create TI zones using the **zone --create** command.

zone --create -t objtype [-o option list] **name** -p "port list"

The "-t" option specifies the type and should always be specified as "ti" when creating TI zones. When you create a TI zone, you can enable or disable failover mode. By default, failover mode is enabled. You can also set the state of the TI zone to **activated** or **deactivated**. By default the zone state is set to the **activated** state. The following values are used to configure these behaviors using the "-o" option:

TABLE 115 Values used to configure failover mode for TI zones

Possible values for "-o" option list	Description
f	Create a TI zone with failover enabled; mutually exclusive with the 'n' option
n	Create a TI zone with failover disabled; mutually exclusive with the 'f' option
a	Configure the TI zone as activated; mutually exclusive with the 'd' option
d	Configure the TI zone as deactivated; mutually exclusive with the 'a' option

The "name" option specifies the TI zone name.

The "-p" option lets you define the members of the TI zone, typically the N_Ports and E_Ports specified using *Domain* and *Port Index*. For example, given an E_Port at Port Index 12 on Domain ID 5, the resulting TI zone member would be "5,12".

In addition to running **zone --create** to define a TI zone, the effective zoning configuration must also be re-enabled in order to apply the TI zone.

Use the following procedure to create a TI zone:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **zone --create** command, specifying failover mode, activated/deactivated status, TI zone name, and TI zone members (N_Ports and E_Ports)
3. Enter the **cfgEnable** command, specifying the current effective configuration.

TI zone updates are not applied until **cfgEnable** is executed

Example for a TI zone named "bluezone" with failover disabled ('n' option) and an "activated" state ('a' option):

```
switch:admin> zone --create -t ti -o an bluezone -p "10,44; 20,44; 10,38; 20,8"
switch:admin> cfgEnable "myConfiguration"
```

Example for a TI zone named "redzone" with failover enabled ('f' option) and a "deactivated" state ('d' option):

```
switch:admin> zone --create -t ti -o fd redzone -p "1,4; 2,4; 1,8; 2,7"
switch:admin> cfgEnable "myConfiguration"
```

Modifying TI zones

Using the **zone --add** command, you can do the following:

- add ports to an existing TI zone
- change the failover option
- activate or deactivate a TI zone

You must reactivate the effective configuration to apply the changes (**cfgEnable**). Example, changing a TI zone named "greenzone" to failover enabled:

```
switch:admin> zone --add -o f greenzone
switch:admin> cfgEnable "myConfiguration"
```

Using the **zone --remove** command, you can remove ports from existing TI zones. If you remove the last member of a TI zone, the TI zone is deleted. You must reactivate the effective configuration to apply the changes (**cfgEnable**).

Example, removing "5,1" from a TI zone named "silverzone":

```
switch:admin> zone --remove silverzone -p "5,1"
switch:admin> cfgEnable "myConfiguration"
```

NOTE

If you have overlapping TI zones and you want to change the failover option on these zones, you must first remove the overlapping ports from the zones, then change the failover type, and finally re-add the overlapping members.

Changing the state of a TI zone

You can change the configured state of an existing TI zone to activated or deactivated. As is the case with other TI zoning operations, you must enable the current effective configuration to apply the change.

1. Connect to the switch and log in using an account with admin permissions.
2. Perform one of the following actions:

- To activate a TI zone, enter the **zone --activate** command.

```
zone --activate name
```

- To deactivate a TI zone, enter the **zone --deactivate** command.

```
zone --deactivate
name
```

3. Enter the **cfgEnable** command to reactivate your current effective configuration and enforce the TI zones.

```
cfgenable " current_effective_configuration"
```

Deleting a TI zone

Use the **zone --delete** command to delete a TI zone from the defined configuration. This command deletes the entire zone; to only remove port members from a TI zone, use the **zone --remove** command, as described in [Modifying TI zones](#) on page 494.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **zone --delete** command.

```
zone --delete name
```

You can delete multiple zones by separating the zone names with a semicolon and enclosing them in quotation marks.

3. Enter the **cfgEnable** command to reactivate your current effective configuration and apply the TI zone delete operation.

```
cfgenable " current_effective_configuration"
```

Displaying TI zones

Use the **zone --show** command to display information about TI zones. This command displays the following information for each zone:

- Zone name
 - E_Port members
 - N_Port members
 - Configured status (the latest status, which may or may not have been activated by **cfgEnable**)
 - Enabled status (the status that has been activated by **cfgEnable**)
1. Connect to the switch and log in using an account with admin permissions.
 2. Enter the **zone --show** command.

```
zone --show [ name ] [-ascending]
```

Examples of displaying TI zone information

Example: Displaying information about the TI zone purplezone

```
switch:admin> zone --show purplezone

Defined TI zone configuration: TI Zone Name:   redzone:
Port List:      1,2; 1,3; 3,3; 4,5
Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
```

Example: Displaying information about all TI zones in the defined configuration in ascending order

```
switch:admin> zone --show -ascending

Defined TI zone configuration: TI Zone Name:   bluezone:
Port List:      8,3; 8,5; 9,2; 9,3;
Configured Status: Deactivated / Failover-Disabled
Enabled Status: Activated / Failover-Enabled
TI Zone Name:   greenzone:
Port List:      2,2; 3,3; 4,11; 5,3;
Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
TI Zone Name:   purplezone:
Port List:      1,2; 1,3; 3,3; 4,5;
Configured Status: Activated / Failover-Enabled
Enabled Status: Deactivated / Failover-Enabled
```

Example: Displaying members for the zone "ti_red" in ascending order

```
switch:admin> zone --show -ascending ti_red
Defined TI zone configuration:
TI Zone Name:   ti_red
Port List:      3,3; 4,4; 5,5
Configured Status: Activated / Failover-Enabled
Enabled Status: Deactivated
```


Example: Displaying members for the zone "TI_zone", regardless of the case

```
switch:admin> zone --show -ic TI_zone* Defined TI zone configuration:
TI Zone Name:      TI_zone
Port List:        7,8
Configured Status: Activated / Failover-Enabled
Enabled Status:   Deactivated
TI Zone Name:     ti_zone
Port List:        3,3
Configured Status: Activated / Failover-Enabled
Enabled Status:   Deactivated
```

Enforcing Local TI Filtering

Local TI Filtering prevents or restricts unwanted traffic flow by enforcing local-to-local TI zone rules.

Prior to Fabric OS 7.4.0, if regular zoning of local-to-local flows did not reflect the same connectivity as the TI zoning, there were cases where the Name Server reported that the connectivity was allowed, but the connectivity was restricted at the routing layers due to TI zoning rules.

Starting with Fabric OS 7.4.0, the Local TI Filtering feature enforces the TI zone rules on local-to-local flows.

Use the **configure** command to configure Local TI Filtering. By default, this feature is disabled. Disable the switch before configuring Local TI Filtering.

Configuring the Local TI Filtering example

Use Local TI Filtering, listed under Zoning Operation parameters of the **configure** command, to enable or disable the feature.

```
switch:admin> switchdisable
switch:admin> configure

Configure...

Fabric parameters (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no]
D-Port Parameters (yes, y, no, n): [no]
RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24) [1]
Zoning Operation parameters (yes, y, no, n): [no] y

Disable NodeName Zone Checking: (0..1) [0]
Ti Zone Filtering (on, off): [off]
..... <output truncated>
```

Traffic Isolation Zoning over FC routers

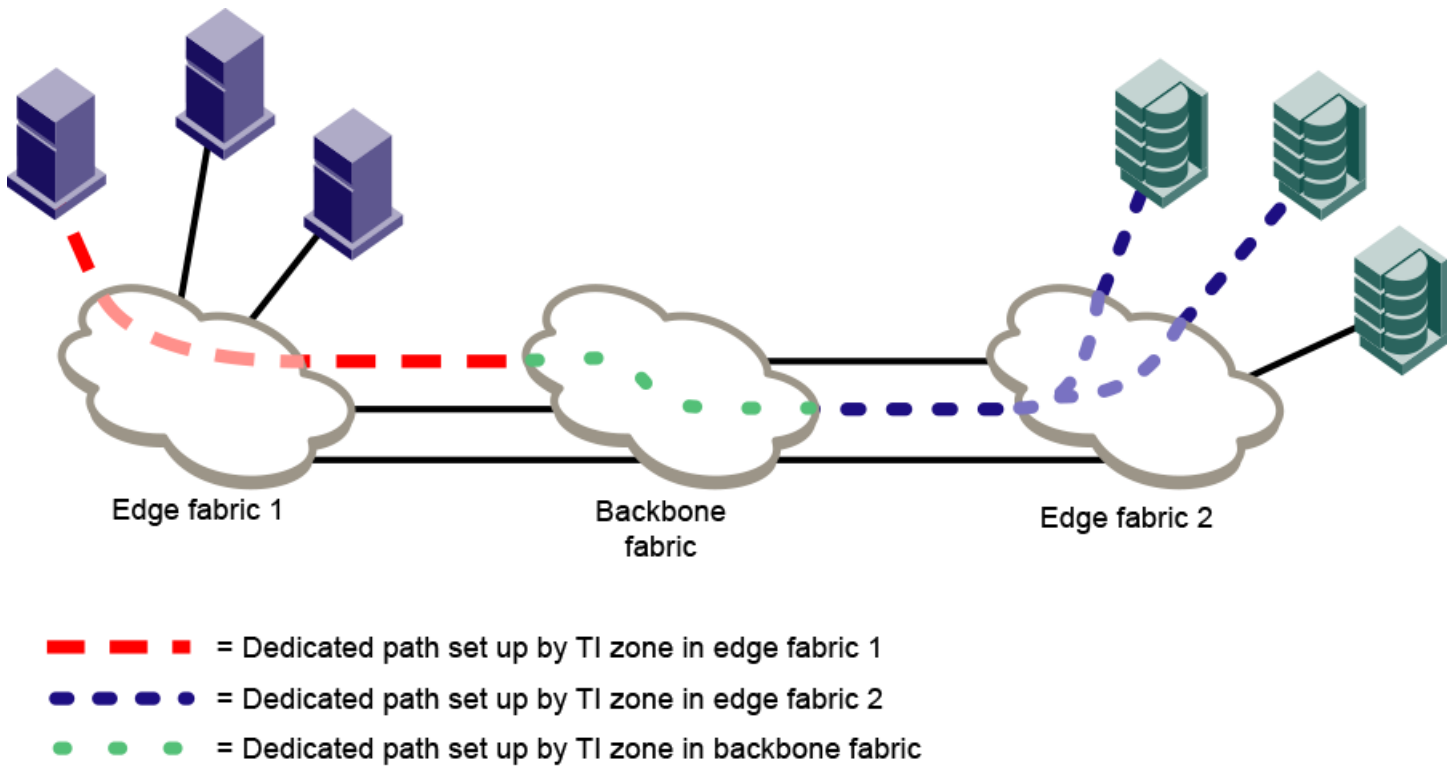
TI zoning can be used with FC routed fabrics with failover enabled. To achieve end-to-end traffic isolation across the metaSAN, TI zones must be configured at each relevant edge fabric, and the backbone fabric, within the metaSAN.

Some VE_Port-based features, such as tape pipelining, require the request and corresponding response traffic to traverse the same VE_Port tunnel across the metaSAN. To ensure that the request and response traverse the same VE_Port tunnel, you must set up Traffic Isolation zones in the edge and backbone fabrics.

- Set up a TI zone in an edge fabric to guarantee that traffic from a specific device in that edge fabric is routed through a particular EX_Port or VEX_Port.
- Set up a TI zone in the backbone fabric to guarantee that traffic between two devices in different fabrics is routed through a particular ISL (VE_Ports or E_Ports) in the backbone.

This combination of TI zones in the backbone and edge fabrics ensures that the traffic between devices in different fabrics traverses the same VE_Port tunnel in a backbone fabric. [Figure 68](#) shows how three TI zones form a dedicated path between devices in different edge fabrics. The backbone fabric can contain one or more FC routers.

FIGURE 68 Traffic Isolation Zoning over FCR



In addition to setting up TI zones, you must also ensure that the devices are in an LSAN zone so that they can communicate with each other.

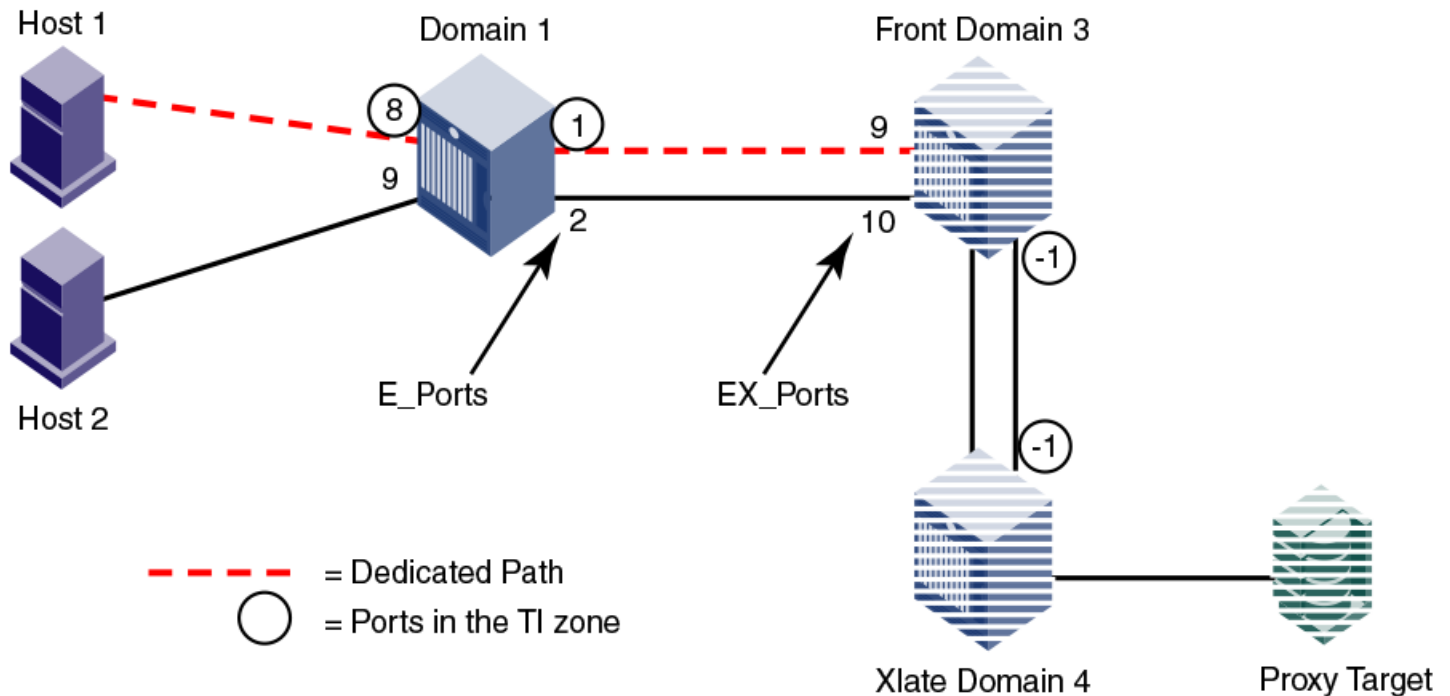
If failover is enabled and the TI path is not available, an alternate path is used. If failover is disabled and the TI path is not available, then devices are not imported.

TI zones within an edge fabric

A TI zone within an edge fabric is used to route traffic between a real device and a proxy device through a particular EX_Port.

For example, in [Figure 69](#), you can set up a TI zone to ensure that traffic between Host 1 and the proxy target is routed through EX_Port 9.

FIGURE 69 TI zone in an edge fabric



In the TI zone, when you designate E_Ports between the front and xlate phantom switches, you must use -1 in place of the "1" in the D,I notation. Both the front and xlate domains must be included in the TI zone.

Using D,I notation, the members of the TI zone in [Figure 69](#) are as follows:

- 1,8
- 1,1
- 3,-1 (E_Port for the front phantom domain)
- 4,-1 (E_Port for the xlate phantom domain)

NOTE

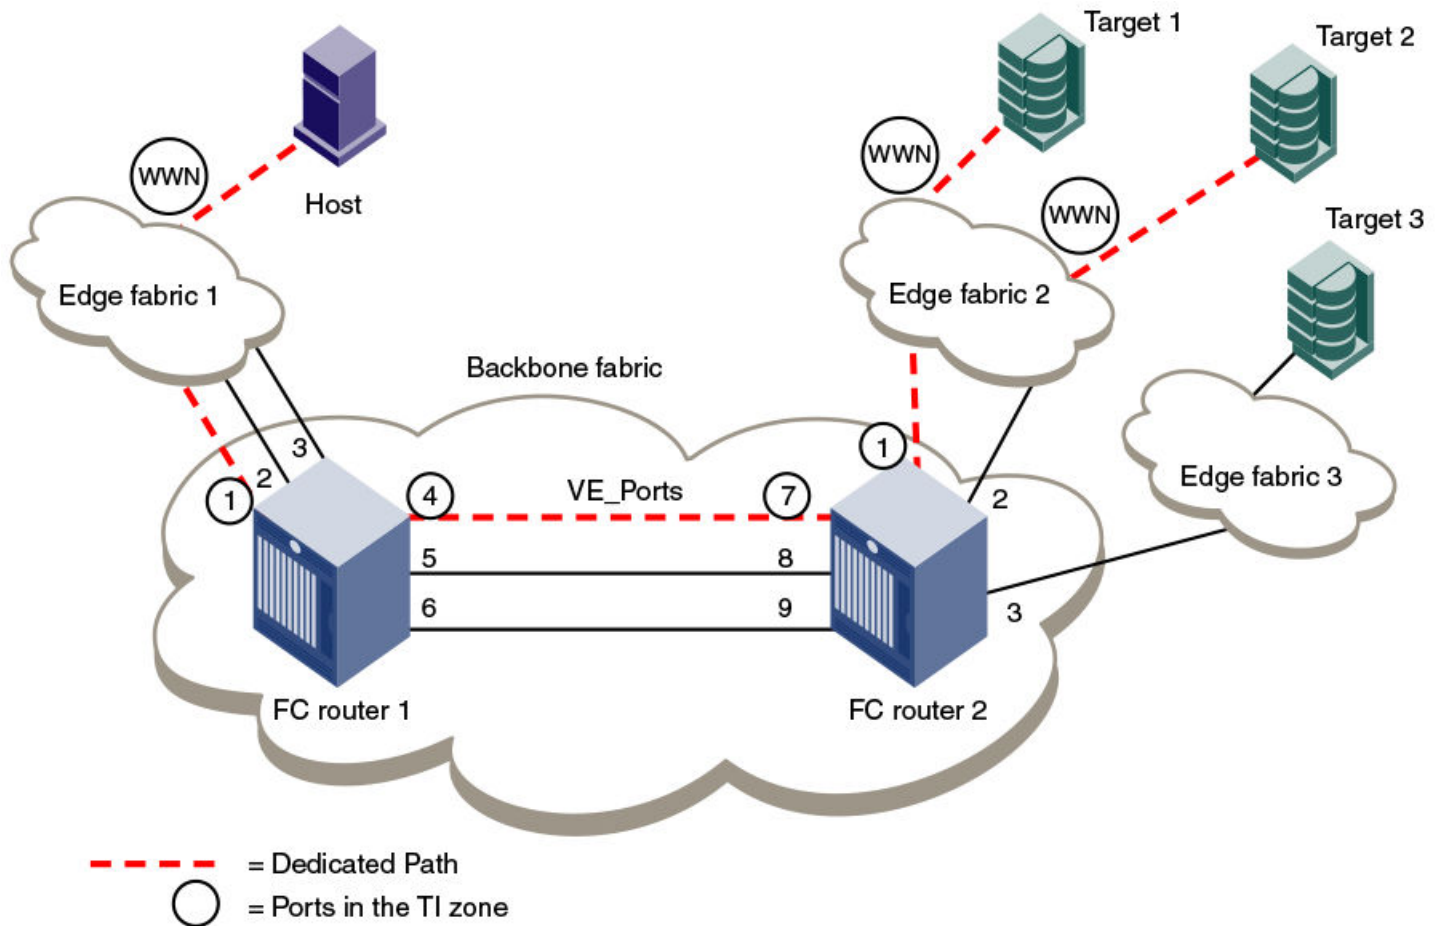
In this configuration the traffic between the front and xlate domains can go through any path between these two domains. The -1 does not identify any specific ISL. To guarantee a specific ISL, you need to set up a TI zone within the backbone fabric.

TI zones within a backbone fabric

A TI zone within a backbone fabric is used to route traffic within the backbone fabric through a particular ISL.

For example, in [Figure 70](#), a TI zone is set up in the backbone fabric to ensure that traffic between EX_Ports "1,1" and "2,1" is routed through VE_Ports "1,4" and "2,7".

FIGURE 70 TI zone in a backbone fabric



TI zones within the backbone fabric use the port WWN instead of D,I notation for devices that are to communicate across fabrics. (You can use the **portShow** command to obtain the port WWN.) Port WWNs should be used only in TI zones within a backbone fabric and should not be used in other TI zones.

Using D,I and port WWN notation, the members of the TI zone in Figure 70 are as follows:

- 1,1 (EX_Port for FC router 1)
- 1,4 (VE_Port for FC router 1)
- 2,7 (VE_Port for FC router 2)
- 2,1 (EX_Port for FC router 2)
- 10:00:00:00:00:01:00:00 (Port WWN for the host)
- 10:00:00:00:00:02:00:00 (Port WWN for target 1)
- 10:00:00:00:00:03:00:00 (Port WWN for target 2)

Limitations of TI zones over FC routers

Be aware of the following when configuring TI zones over FC routers:

- A TI zone defined within the backbone fabric does not guarantee that edge fabric traffic will arrive at a particular EX_Port. You must set up a TI zone in the edge fabric to guarantee this.
- TI zones within the backbone fabric cannot contain more than one destination router port (DRP) per each fabric. This means you cannot define more than one EX_Port to any one edge fabric unless they are part of a trunk.
- Only one egress E_Port or VE_Port connected to the next hop can be defined within TI zones. Only one ISL or trunk can be defined between two backbone switches. The E_Port in the backbone fabric TI zone should be the least cost ISL.
- TI over FCR is supported only from edge fabric to edge fabric. Traffic isolation from backbone to edge is not supported.
- Non-TI data traffic is *not* restricted from going through the TI path in the backbone fabric.
- For TI over FCR, failover must be enabled in the TI zones in the edge fabrics and in the backbone fabric.
- TI over FCR is not supported with FC Fast Write.
- ETIZ over FCR is not supported.
- Configuring backbone TI zone with remote FC router EX_Port is not supported if local chassis has the EX_Port to source and destination edge fabric in the same chip/port group.
- *For the FX8-24 blades only:* If Virtual Fabrics is disabled, two or more shared area EX_Ports connected to the same edge fabric should not be configured in different TI zones. This configuration is not supported.

Fabric-Level Traffic Isolation in a backbone fabric

For Fibre Channel Routed (FCR) environments, you can use TI zoning if you want traffic isolation only at the fabric level and not at the device level.

For example, two fabrics within a MetaSAN need to communicate only with each other. There is no other traffic across the backbone that goes from either of these edge fabrics to any other edge fabric in the MetaSAN. In this case, all of the traffic entering the FCR backbone from one of these edge fabrics will go to the other edge fabric. If these two edge fabrics are connected to two different backbone switches (FC routers), then traffic between these fabrics can be isolated to a specified set of links within the backbone fabric using one of two methods:


- TI over FCR, which includes the PWWN of devices and maintains device level isolation
- TI zoning in the backbone, which provides fabric level isolation

If device-level isolation is needed from one edge fabric to another, then use TI over FCR using Port World Wide Names (PWWNs). However, if there is no need for device-level isolation, but a need for fabric-level isolation, then use Fabric-Level Traffic Isolation, described in this section.

TI over FCR is described in [Traffic Isolation Zoning over FC routers](#) on page 497.

If two edge fabrics are connected to two different backbone switches, then traffic between these fabrics can be isolated to a specified set of links within the backbone fabric using TI zoning in the backbone without including device PWWNs. This is called Fabric-Level Traffic Isolation, as shown in the following figure.

FIGURE 71 Fabric-level traffic isolation



In the figure, there are two links between each edge fabric and the backbone fabric, and there are five links between the two FC routers in the backbone. Fabric ID 1 and Fabric ID 4 communicate only with each other. Two backbone ISLs are dedicated to traffic between FID1 and FID4. These dedicated ISLs are indicated in red and blue.

Fabric-Level TI zones

Fabric-Level Traffic Isolation is accomplished through the use of TI zones. These zones define the dedicated set of paths between the two fabrics. These paths are to be restricted to just the traffic between the two fabrics.

Fabric-Level TI zones are defined in the backbone fabric, and include only EX_Ports and E_Ports in the backbone fabric. The TI zones do not include device PWWNs or ports in the edge fabrics.

The TI zone must include every port in the path from ingress EX_Port to egress EX_Port. The TI zone definitions must include all EX_Ports connected to the two edge fabrics. Unless all possible ingress ports are included, some traffic will not be isolated to the desired paths.

Fabric-Level Traffic Isolation is not enforced on the egress EX_Ports. Any available egress IFL can be used, regardless of whether it is in the Fabric-Level TI zone.

The following rules apply to creating Fabric-Level TI zones:

- Include all EX_Ports connected to the two edge fabrics.

- Include E_Ports for the path between the backbone switches.
- Do not include E_Ports from the edge fabrics.
- Do not include device PWWNs.
- Ensure that failover is enabled.

There are two options for defining the Fabric-Level Traffic Isolation paths within TI zones. The option you select affects the failover behavior of the TI zones.

- Create a separate TI zone for each path
- Combine all of the paths in a single TI zone

Failover behavior for Fabric-Level TI zones

Fabric-Level Traffic Isolation requires the TI zones in the backbone to have failover enabled. The failover behavior differs depending on how you create the TI zones.

If you create a separate TI zone for each path:

- If one of the TI zone paths fails, then traffic on that path is re-routed to a non-dedicated path.

For example, in [Figure 71](#) on page 502, if the BLUE path fails, then traffic associated with the BLUE path is re-routed to a black path.

If you combine all of the paths in a single TI zone:

- If one of the paths in the TI zone fails, then traffic on that path is re-routed to another path in the TI zone.

For example, in [Figure 71](#) on page 502, if the BLUE path fails, then traffic associated with the BLUE path is re-routed to the RED path.

Failover behavior for Fabric-Level TI zones is the same as for other TI zones, as described in [TI zone failover](#) on page 485.

Creating a separate TI zone for each path

Brocade recommends that you create a separate TI zone for each path if you want TI traffic to failover to non-TI zone paths.

This example procedure creates two TI zones in the backbone fabric shown in [Figure 71](#) on page 502.

1. Create TI zones with failover enabled.

Each TI zone must include the ingress and egress EX_Ports, as well as the E_Ports between the two backbone switches. Do not include the edge fabric E_Ports or device PWWNs.

```
switch:admin> zone --create -t ti TI_Zone_Red -p "20,5; 20,3; 30,7; 30,9"
switch:admin> zone --create -t ti TI_Zone_Blue -p "20,6; 20,4; 30,8; 30,10"
```

By default, a new TI zone is configured as "Activated" with failover enabled.

2. Display defined TI zones.

```
switch:admin> zone --show

Defined TI zone configuration:
TI Zone Name:  TI_Zone_Blue
Port List:     20,6; 20,4; 30,8; 30,10
Configured Status: Activated / Failover-Enabled
Enabled Status: Deactivated
TI Zone Name:  TI_Zone_Red
Port List:     20,5; 20,3; 30,7; 30,9
Configured Status: Activated / Failover-Enabled
Enabled Status: Deactivated
```

Be aware that although the configured status is "Activated", the enabled status is "Deactivated".

3. Activate TI zones.

```
switch:admin> cfgactvshow

Effective configuration:
cfg: < current effective configuration
>
...
switch:admin> cfgenable < current effective configuration
>
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
Do you want to enable 'TI_Config' configuration (yes, y, no, n): [no] y

zone config "name" is in effect
Updating flash ...
switch:admin>
Switch:admin> zone --show

Defined TI zone configuration:
TI Zone Name:  TI_Zone_Blue
Port List:     20,6; 20,4; 30,8; 30,10
Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
TI Zone Name:  TI_Zone_Red
Port List:     20,5; 20,3; 30,7; 30,9
Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
```

The enabled status now displays as "Activated".

Creating a single TI zone for all paths

Brocade recommends that you create a single TI zone for all paths if you want TI traffic to failover to other paths in the TI zone.

1. Create a single TI zone with failover enabled.

The TI zone must include the ingress and egress EX_Ports, as well as the E_Ports between the two backbone switches. Do not include the edge fabric E_Ports or device PWWNs.

```
switch:admin> zone --create -t ti TI_Zone_ALL -p "20,3; 20,4; 20,5; 20,6; 30,7; 30,8; 30,9 30,10"
```

By default, a new TI zone is configured as "Activated" with failover enabled.

2. Display defined TI zone.

```
switch:admin> zone --show

Defined TI zone configuration:
TI Zone Name:   TI_Zone_ALL
Port List:     20,3; 20,4 20,5; 20,6; 30,7; 30,8; 30,9; 30,10
Configured Status: Activated / Failover-Enabled
Enabled Status: Deactivated
```

Be aware that although the configured status is "Activated", the enabled status is "Deactivated".

3. Activate the TI zone.

```
switch:admin> cfgactvshow

Effective configuration:
cfg: < current effective configuration
>
...
switch:admin> cfgenable < current effective configuration
>
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
Do you want to enable 'TI_Config' configuration (yes, y, no, n): [no] y

zone config "name" is in effect
Updating flash ...
switch:admin>
Switch:admin> zone --show

Defined TI zone configuration:
TI Zone Name:   TI_Zone_ALL
Port List:     20,3; 20,4; 20,5; 20,6; 30,7; 30,8; 30,9; 30,10
Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
```

Then enabled status now displays as "Activated".

Setting up TI zones over FCR (sample procedure)

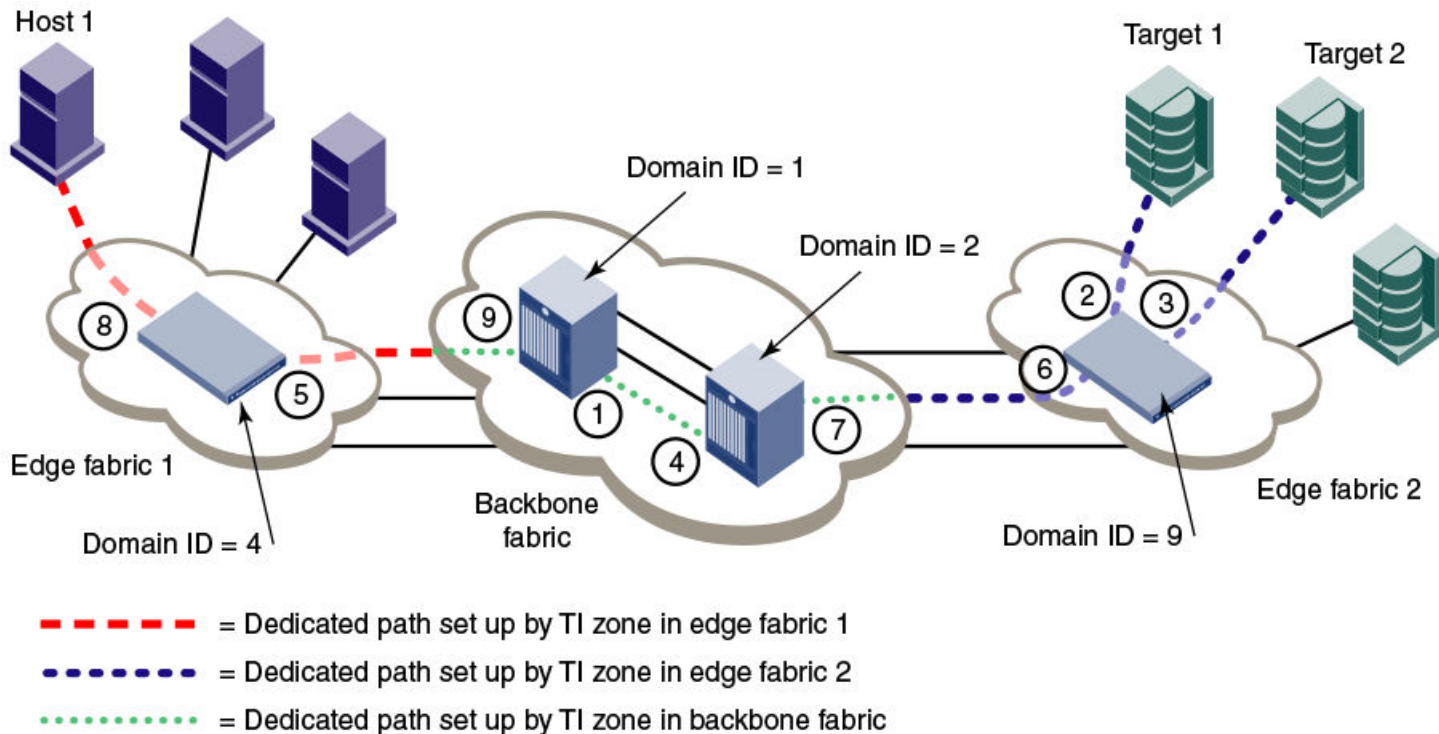
The following example shows how to set up TI zones over FCR to provide a dedicated path, as shown in [Figure 72](#). In this example, three TI zones are created: one in each of the edge fabrics and one in the backbone fabric. The combination of these three TI zones creates a dedicated path for traffic between Host 1 in edge fabric 1 and Targets 1 and 2 in edge fabric 2.

Host 1 has port WWN 10:00:00:00:00:08:00:00

Target 1 has port WWN 10:00:00:00:00:02:00:00

Target 2 has port WWN 10:00:00:00:00:03:00:00

FIGURE 72 TI over FCR example



NOTE

In the following procedure the three TI zones in the edge and backbone fabrics are all given the same name, TI_Zone1. It is not required that the TI zones have the same name, but this is done to avoid confusion. If several dedicated paths are set up across the FC router, the TI zones for each path can have the same name.

1. In each edge fabric, set up an LSAN zone that includes Host 1, Target 1, and Target 2, so these devices can communicate with each other. Refer to [Using FC-FC Routing to Connect Fabrics](#) on page 535 for information about creating LSAN zones.

2. Log in to the edge fabric 1 and set up the TI zone.

- a) Enter the **fabricShow** command to display the switches in the fabric. From the output, you can determine the front and translate domains.

```
Elswitch:admin> fabricshow
```

Switch ID	Worldwide Name	Enet IP Addr	FC IP Addr	Name
1: fffc01	50:00:51:e3:95:36:7e:04	0.0.0.0	0.0.0.0	"fcr_fd_1"
4: fffc04	10:00:00:60:69:80:1d:bc	10.32.72.4	0.0.0.0	>"Elswitch"
6: fffc06	50:00:51:e3:95:48:9f:a0	0.0.0.0	0.0.0.0	"fcr_xd_6_9"

The Fabric has 3 switches

- b) Enter the following commands to create and display a TI zone:

```
Elswitch:admin> zone --create -t ti TI_Zone1 -p "4,8; 4,5, 1,-1; 6,-1"

Elswitch:admin> zone --show

Defined TI zone configuration:
TI Zone Name:   TI_Zone1
Port List:      4,8; 4,5; 1,-1; 6,-1
Status: Activated      Failover: Enabled
```

- c) Enter the following commands to reactivate your current effective configuration and enforce the TI zones.

```
Elswitch:admin> cfgactvshow

Effective configuration:
cfg:   cfg_TI
zone:  lsan_t_i TI_Zone1
      10:00:00:00:00:00:02:00:00
      10:00:00:00:00:00:03:00:00
      10:00:00:00:00:00:08:00:00

Elswitch:admin> cfgenable cfg_TI

You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
If the update includes changes to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with the traffic isolation zone changes
Do you want to enable 'cfg_TI' configuration (yes, y, no, n): [no] y
zone config "cfg_TI" is in effect
Updating flash ...
```

3. Log in to the edge fabric 2 and set up the TI zone.

- a) Enter the **fabricShow** command to display the switches in the fabric. From the output, you can determine the front and translate domains.

```
E2switch:admin> fabricshow
```

Switch ID	Worldwide Name	Enet IP Addr	FC IP Addr	Name
1: fffc01	50:00:51:e3:95:36:7e:09	0.0.0.0	0.0.0.0	"fcr_fd_1"
4: fffc04	50:00:51:e3:95:48:9f:a1	0.0.0.0	0.0.0.0	"fcr_xd_6_9"
9: fffc09	10:00:00:05:1e:40:f0:7d	10.32.72.9	0.0.0.0	>"E2switch"

The Fabric has 3 switches

- b) Enter the following commands to create and display a TI zone:

```
E2switch:admin> zone --create -t ti TI_Zone1 -p "9,2; 9,3; 9,6; 1,-1; 4,-1"
E2switch:admin> zone --show
```

Defined TI zone configuration:
 TI Zone Name: TI_Zone1
 Port List: 9,2; 9,3; 9,6; 1,-1; 4,-1
 Status: Activated Failover: Enabled

- c) Enter the following commands to reactivate your current effective configuration and enforce the TI zones.

```
E2switch:admin> cfgactvshow
```

Effective configuration:
 cfg: cfg_TI
 zone: lsan_t_i TI_Zone1
 10:00:00:00:00:00:02:00:00
 10:00:00:00:00:00:03:00:00
 10:00:00:00:00:00:08:00:00

```
E2switch:admin> cfgenable cfg_TI
```

You are about to enable a new zoning configuration.
 This action will replace the old zoning configuration with the
 current configuration selected.
 If the update includes changes to one or more traffic isolation zones, the update may result in
 localized disruption to traffic on ports associated with the traffic isolation zone changes
 Do you want to enable 'cfg_TI' configuration (yes, y, no, n): [no] y
 zone config "cfg_TI" is in effect
 Updating flash ...

4. Log in to the backbone fabric and set up the TI zone.

```
FCR_Domain_1:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr   FC IP Addr     Name
-----
1: fffc01 10:00:00:05:1e:52:3a:00 10.38.135.15   0.0.0.0        >"FCR_Domain_1"
2: fffc02 10:00:00:27:f8:f1:0b:40 10.38.135.19   0.0.0.0        "FCR_Domain_2"
```

The Fabric has 2 switches

The Backbone domains are represented in "D,I" notation in TI zone configuration.

a) Enter the following commands to create and display a TI zone:

```
FCR_Domain_1:admin> zone --create -t ti TI_Zone1 -p "1,9; 1,1; 2,4; 2,7;
10:00:00:00:00:08:00:00; 10:00:00:00:00:02:00:00; 10:00:00:00:00:03:00:00"
FCR_Domain_1:admin> zone --show
Defined TI zone configuration:
TI Zone Name: TI_Zone1
Port List: 1,9; 1,1; 2,4; 2,7; 10:00:00:00:00:08:00:00;
10:00:00:00:00:02:00:00; 10:00:00:00:00:03:00:00
Status: Activated Failover: Enabled
```

b) Enter the following commands to reactivate your current effective configuration and enforce the TI zones.

```
FCR_Domain_1:admin> cfgactvshow
Effective configuration:
cfg: cfg_TI
zone: lsan_t_i_TI_Zone1
10:00:00:00:00:00:02:00:00
10:00:00:00:00:00:03:00:00
10:00:00:00:00:00:08:00:00
FCR_Domain_1:admin> cfgenable cfg_TI
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
If the update includes changes to one or more traffic isolation zones, the
update may result in localized disruption to traffic on ports associated
with the traffic isolation zone changes
Do you want to enable 'cfg_TI' configuration (yes, y, no, n): [no] y
zone config "cfg_TI" is in effect
Updating flash ...
```

Virtual Fabrics considerations for Traffic Isolation Zoning

TI zones can be created in a logical fabric like in regular fabrics, with the following exceptions:

- The disable failover option is not supported in logical fabrics that use extended ISLs (XISLs).
Although logical switches that use XISLs allow the creation of a TI zone with failover disabled, this is not a supported configuration. Base switches do not allow the creation of a TI zone with failover disabled.
- To create a TI zone for a logical fabric that uses XISLs, you must create two TI zones: one in the logical fabric and one in the base fabric. The combination of TI zones in the base fabric and logical fabric sets the path through the base fabric for logical switches.

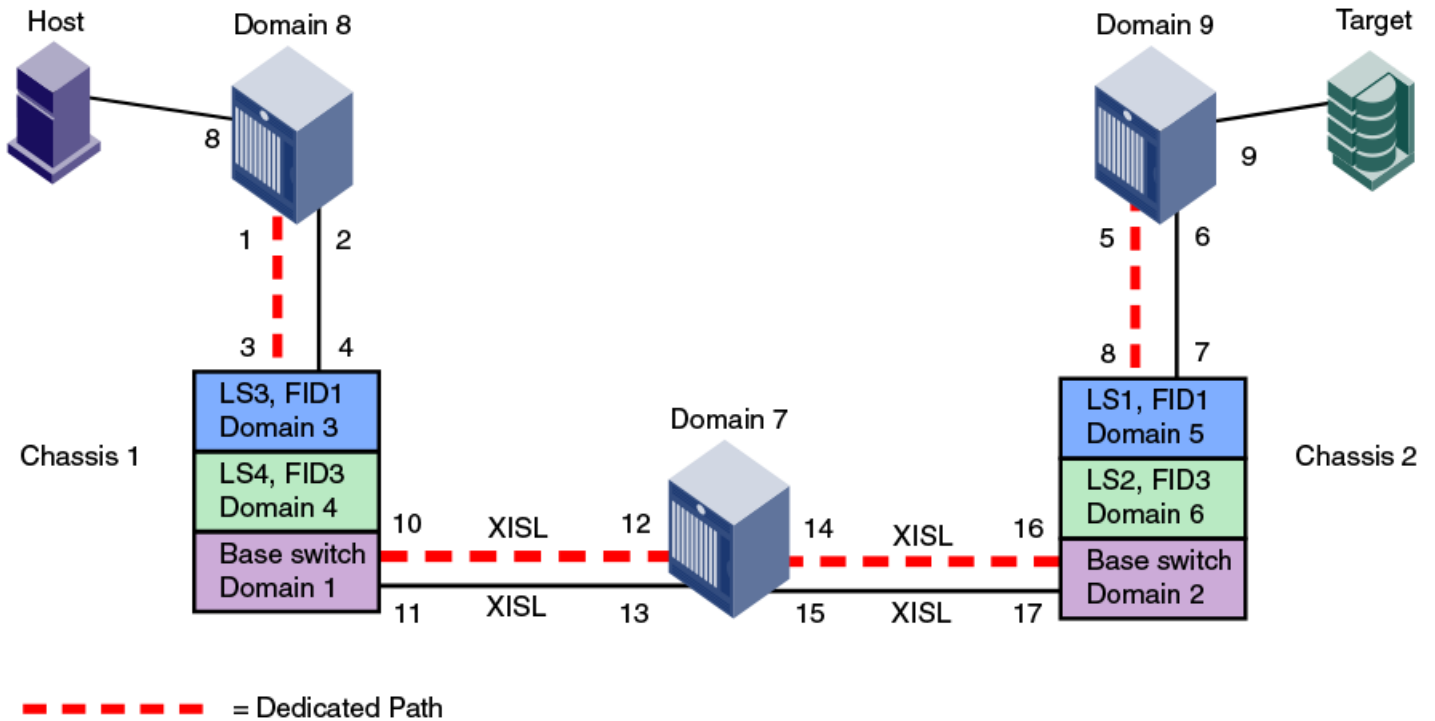
The TI zone in the logical fabric includes the XISL port numbers, as well as the F_Ports and ISLs in the logical fabric.

Because base fabrics do not contain end devices, they normally do not have an effective zone configuration. To activate a TI zone in a base fabric, you must create a "dummy" configuration. TI zones are otherwise created in the same manner as a regular fabric with the exception that TI zone failover cannot be disabled with TI zones in the base fabric.

The TI zone in the base fabric reserves XISLs for a particular logical fabric. The base fabric TI zone should also include ISLs that belong to logical switches participating in the logical fabric.


The following figure shows an initiator and target in a logical fabric (FID1). The dotted line indicates a dedicated path between initiator and target. The dedicated path passes through the base fabric over an XISL. (The figure shows only physical ISLs, not logical ISLs.) To create the TI zones for this dedicated path, you must create a TI zone in the logical fabric (FID 1) and one in the base fabric.

FIGURE 73 Dedicated path with Virtual Fabrics




The following figure shows a logical representation of FID1 in Figure 73. To create the dedicated path, you must create and activate a TI zone in FID1 that includes the circled ports shown in the following figure.

FIGURE 74 Creating a TI zone in a logical fabric



You must also create and activate a TI zone in the base fabric to reserve the XISLs for the dedicated path. In the following figure, the XISLs (indicated by a dotted line) in the base fabric can be reserved for FID1 by defining and activating a base fabric TI zone that consists of ports 10, 12, 14, and 16. You must also include ports 3 and 8, because they belong to logical switches participating in the logical fabric. For the TI zone, it is as though ports 3 and 8 belong to Domains 1 and 2, respectively.

FIGURE 75 Creating a TI zone in a base fabric



Using D,I notation, the port numbers for the TI zones in the logical fabric and base fabric are as follows:

TABLE 116 Port numbers used in TI zones

Port members for the TI zone in logical fabric	Port members for the TI zone in base fabric
8,8 F_Port	1,3 E_Port for ISL in logical switch
8,1 E_Port	1,10 E_Port for XISL
3,3 E_Port	7,12 E_Port for XISL
3,10 E_Port	7,14 E_Port for XISL
5,16 E_Port	2,16 E_Port for XISL
5,8 E_Port	2,8 E_Port for ISL in logical switch
9,5 E_Port	
9,9 F_Port	

Notice that the base fabric zone contains a reference to port 1,3 even though the base switch with domain 1 does not have a port 3 in the switch. This number refers to the port in the *chassis* with port index 3, which actually belongs to LS3 in FID 1.

Traffic Isolation Zoning over FC routers with Virtual Fabrics

This section describes how you can set up TI zones over FC routers in logical fabrics. Figure 76 shows two physical chassis configured into logical switches. The initiator in FID 1 communicates with the target in FID 3 over the EX_Ports in the base switches.

FIGURE 76 Example configuration for TI zones over FC routers in logical fabrics

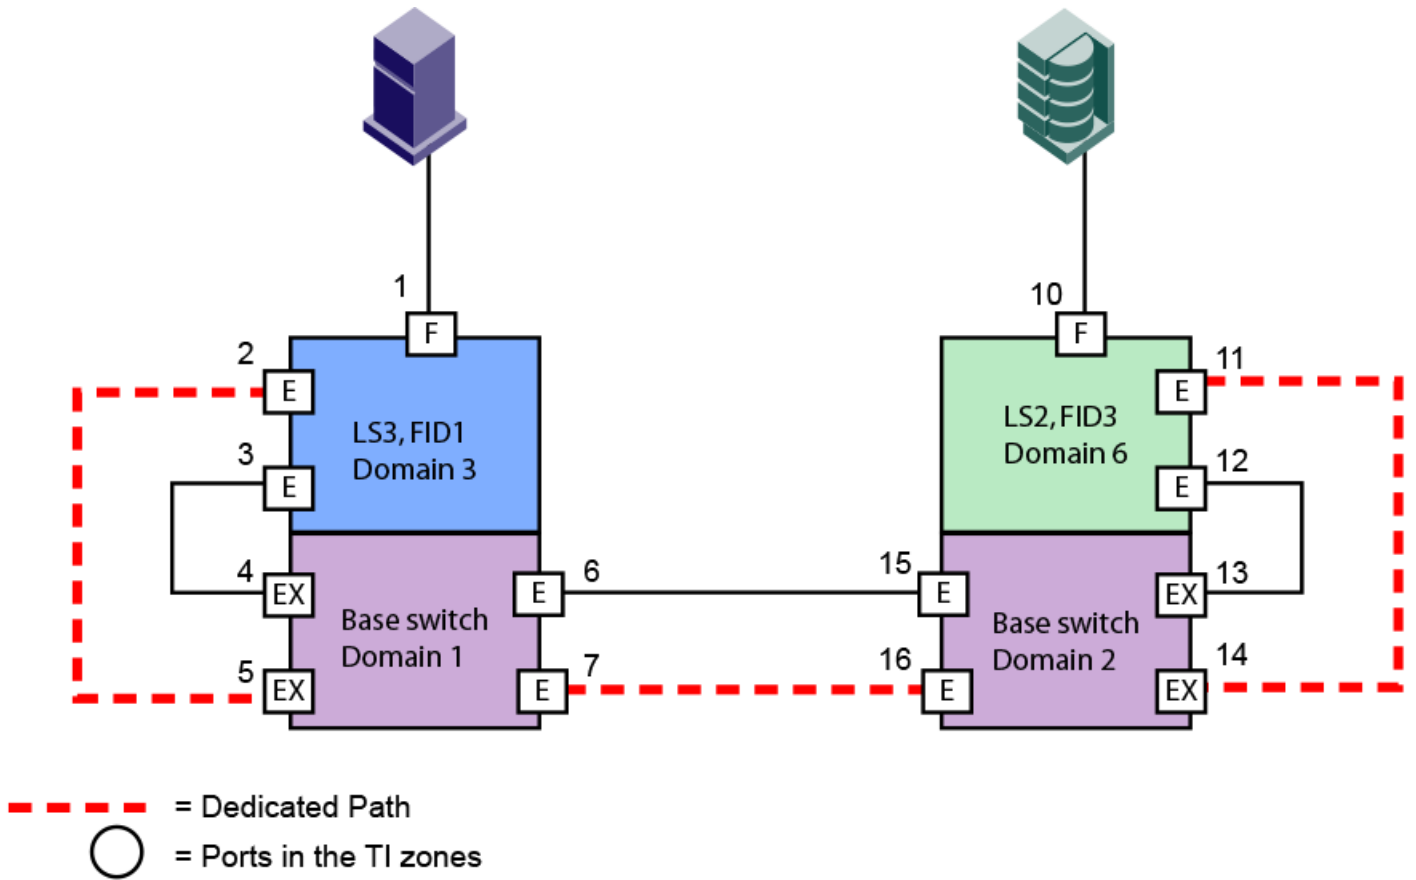



Figure 77 shows a logical representation of the configuration in Figure 76. This SAN is similar to that shown in Figure 68 on page 498 and you would set up the TI zones in the same way as described in Traffic Isolation Zoning over FC routers on page 497.

FIGURE 77 Logical representation of TI zones over FC routers in logical fabrics



Troubleshooting TI zone routing problems

Use the following procedure to generate a report of existing and potential problems with TI zones. The report displays an error type.

- "ERROR" indicates a problem currently exists in the fabric.
- "WARNING" indicates that there is not currently a problem, given the current set of online devices and reachable domains, but given the activated TI zone configuration, parallel exclusive paths between a shared device and a remote domain have been detected, which might cause a problem for devices that join the fabric later.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **zone --showTIerrors** command.

```
zone --showTIerrors
```

Here is an example report that would be generated for the illegal configuration shown in [Figure 65](#) on page 490.

```
switch:admin> zone --showTIerrors

My Domain: 3
Error type:          ERROR
Affected Remote Domain: 1
Affected Local Port: 8
Affected TI Zones:   etiz1, etiz2
Affected Remote Ports: 1, 2, 3, 4
```

TI Zone violation handling for trunk ports

For any trunk group, all the members of the group need to belong to the TI zone to prevent routing issues resulting from changes in the members of the trunk group. This applies to any E_Port or F_Port trunk groups that are included in TI zones using failover disabled mode.

Fabric OS posts a RASLog message (ZONE-1061) if any of the ports part of a trunk group is not added to the TI zone with failover disabled. Also, a CLI (**zone --showTItrunkerrors**) is provided to check if all ports per switch in a TI zone are proper. This helps you identify missing trunk members and take corrective actions.

The following is the RASLog message issued when any port in a trunk group is not in the TI zone

```
SW82:FID128:admin> zone
[ZONE-1061], 620/181, FID 128, WARNING, sw0, Some trunk members are missing from failover disabled active
TI zones.
```

The CLI essentially displays the details of the trunk members present in the TI zone and those not present in the TI zone. These details are displayed per TI Zone basis.

The following is the RASLog message issued when --showTItrunkerrors is added to zone command

```
switch:admin> zone --showTItrunkerrors

TI Zone Name: brackets
E-Port Trunks
Trunk members in TI zone: 16 18
Trunk members not in TI zone: 17
F-Port Trunks
Trunk members in TI zone: 4 5
Trunk members not in TI zone: 6
TI Zone Name: loop
E-Port Trunks
Trunk members in TI zone: 0
Trunk members not in TI zone: 1
TI Zone Name: operand
E-Port Trunks
Trunk members in TI zone: 8
Trunk members not in TI zone: 9 10
E-Port Trunks
Trunk members in TI zone: 16
Trunk members not in TI zone: 17 18
```

Routing Traffic

• Routing overview.....	515
• Inter-switch links.....	517
• Gateway links.....	520
• Routing policies.....	522
• Route selection.....	524
• Frame order delivery.....	525
• Lossless Dynamic Load Sharing on ports.....	528
• Frame Redirection.....	532

Routing overview

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 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. There are two kinds of routing protocols on intranet networks, distance vector and link state.

- Distance vector is based on hop count. This 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 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.

It is recommended that 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 constitutes 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 traffic along the shortest path between the source and destination based upon 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 keeps track of 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.

Once 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 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 rooted on the principal switch. Frames are only sent on the principal ISLs that belong to the spanning tree. When there are multiple ISLs 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 78 Principal ISLs



NOTE

FSPF only supports 16 routes in a zone, including Traffic Isolation Zones.

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. This guarantees that a control frame is not delayed by user data and that FSPF routing decisions occur very quickly during convergence.

FSPF guarantees a routing loop-free topology at all times. It is essential for a fabric to include many physical loops because, without loops, there would not be multiple paths between switches, and consequently 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 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 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 only needs to read word zero and 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, then 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 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 once the login process finishes successfully. For more information on the login process, refer to [Understanding Fibre Channel Services](#) on page 21.

You can expand your fabric by connecting new switches to existing switches. [Figure 79](#) shows a new switch being added into an existing fabric. The thick red line is the newly formed ISL.

FIGURE 79 New switch added to existing fabric



When connecting two switches together, Brocade recommends the 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, then the switch you are introducing into the fabric must also have it enabled. If you experience a segmented fabric, refer to the *Brocade Fabric OS Troubleshooting and Diagnostics Guide* to fix the problem.

Buffer credits

In order 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 acknowledgement 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](#) on page 298.

Congestion versus over-subscription

Congestion occurs when a channel is bottlenecked and fully utilized. This kind of bottleneck is a congestion bottleneck. You should be aware that "over-subscription" does not have the same meaning as "congestion". Over-subscription refers only to the potential for congestion; an over-subscribed link may go through a lifetime of normal operation and never be congested. The term over-subscription is not to be used in place of congestion, which is the actual contention for bandwidth by devices through an ISL.

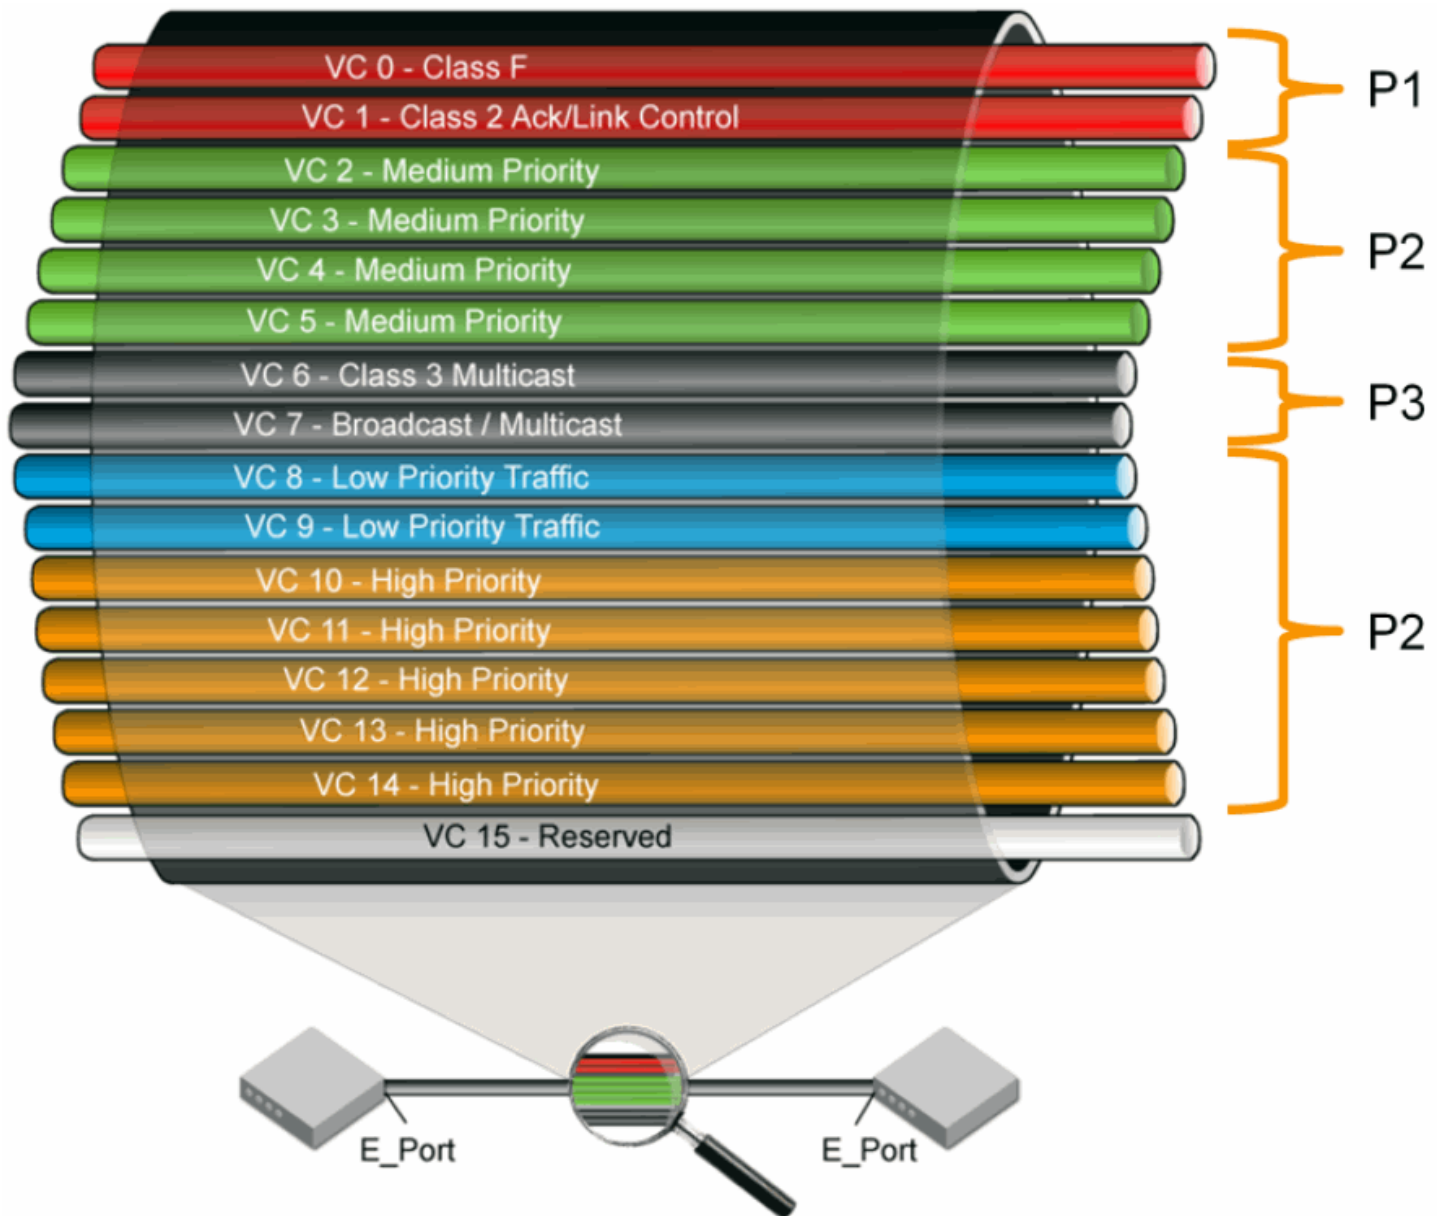
Virtual channels

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, refer to [Optimizing Fabric Behavior](#) on page 319.

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. This example is illustrated in [Figure 80](#).

Quality of Service (QoS) is a licensed 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 [Figure 80](#). The seven data virtual channels (VC8 through VC14) are used to multiplex data frames based upon QoS zones when congestion occurs. For more information on QoS zones, refer to [Optimizing Fabric Behavior](#) on page 319.

FIGURE 80 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 81 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 the following guidelines are followed:

- All switches in the fabric use the core PID format, as described in [Configuring a link through a gateway](#) on page 521.
- 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.

- 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

- Connect to the switch and log in using an account with admin permissions.
- 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 SID and DID. When an FC router is in exchange-based routing mode, the backbone traffic is load-balanced based on SID, DID, and 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 refer to [Dynamic Path Selection](#) on page 523.

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. This helps to ensure that the exchanges between a pair of devices stay in order.

NOTE

Device-based routing requires the use of Dynamic Load Sharing (DLS); when this policy is in effect, you cannot disable the DLS feature.

Device-based routing is also a form of Dynamic Path Selection (DPS). For more information on DPS refer to [Dynamic Path Selection](#) on page 523.

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 a very 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*.
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> dlset
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 six discard frame type options are enabled and frames will be recorded when they are discarded.

TABLE 117 Supported values for the **-type** option.

Option	Reason
timeout	Due to a timeout.
du	Due to the destination being unreachable.
unroute	Due to an unroutable frame.
type1miss	Due to a type 1 miss.
type2miss	Due to a type 2 miss.
type6miss	Due to a type 6 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 unroutable
```

NOTE

Frame viewing for unroutable and “destination unreachable” frames is supported only on the following devices:

- Brocade 6505, 6510, 6520, G610, G620, G630, DCX 8510-4, DCX 8510-8, X6-4 and X6-8 switches.
- Brocade CR32-4, CR32-8, CR16-4, CR16-8, FC32-48, FC16-32, FC16-48, and FC16-64 blades. If a chassis has any older blades, only the timeout frames will be captured for those blades.

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 type-1,2, 6 miss, it 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.

You can view up to 40 discarded frames per ASIC per second for 1200 seconds with the **frameLog** command. You filter this view using a number of fields to filter this output so you are viewing only those frames you are concerned with.

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* 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 back-end 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 back-end port.

Individual back-end ports cannot be specified, only the quality of being a back-end 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 specifying **-1** for the port number.

```
switch:admin> framelog --show -txport -1 -rxport -1
```

Specify **-1** for fixed-port switches and **-1/-1** for Brocade Backbones. These indicate "any back-end port".

NOTE

Frame discards can be logged as audit messages using the Fabric OS syslog facility.

Lossless Dynamic Load Sharing on ports

Lossless Dynamic Load Sharing (DLS) allows you to rebalance port paths without causing input/output (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 6505
- Brocade 6510
- Brocade 6520
- Brocade G610
- Brocade G620
- Brocade G630
- Brocade M6505 blade server SAN I/O module
- Brocade 6543 blade server SAN I/O module
- Brocade 6545 blade server SAN I/O module
- Brocade 6546 blade server SAN I/O module
- Brocade 6547 blade server SAN I/O module
- Brocade 6548 blade server SAN I/O module

- Brocade 6558 blade server SAN I/O module
- Brocade DCX 8510 Directors with the following blades:
 - Brocade FC16-32, FC16-48, FC16-64 port blades
- Brocade X6 Directors with the following blades:
 - Brocade FC32-48 port blades

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 during a rebalance and 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, then IOD must be enabled. However, this combination of Lossless DLS and IOD is supported only in specific topologies, such as in a 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 118 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 DCX 8510 Backbones and Brocade 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** can decommission the E_Ports losslessly and bring up the E_Ports back 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*.

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 Lossless Dynamic Load Sharing 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 DCX 8510 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, it is recommended to 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](#) on page 382.

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 utilized. During these rebalance operations, even if the lossless DLS setting is enabled, route updates 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 happen 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 and/or out-of-order frame delivery.

Once 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 TI zones due to non-deterministic ordering of events when multiple ISLs come online at the same time.
- 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 Two-hop Lossless DLS feature.

TABLE 119 Commands used to configure the Two-hop Lossless DLS feature

Command	Usage
dlsshow	Displays the state of the Two-hop Lossless DLS state when it has been enabled.
dlsreset	Resets the state of DLS in port-based routing.
dlssset	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 **dlssset** 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> dlssset --help
dlssset 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 were 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
```

Frame Redirection

Frame Redirection provides a means to redirect traffic flow between a host and a target that use virtualization and encryption applications, such as the Brocade SAS blade and Brocade Data Migration Manager (DMM), so that those applications can perform without having to reconfigure the host and target. You can use this feature if the hosts and targets are not directly attached.

Frame Redirection depends on the wide distribution of the Defined Zone Database. The Defined Zone Database on Fabric OS switches is pushed out to all other Fabric OS switches in the fabric that support Frame Redirection. Redirection zones exist only in the defined configuration and cannot be added to the effective configuration.

NOTE

Fabric OS v7.2.0 is not supported on the Brocade 7600 or Brocade SAS blade. However, this hardware can run in a pre-Fabric OS v7.2.0 system and attach to a Fabric OS v7.2.0 fabric.

Frame Redirection uses a combination of special frame redirection zones and name server changes to spoof the mapping of real device WWNs to virtual PIDs.

FIGURE 82 Single host and target




Figure 82 demonstrates the flow of Frame Redirection traffic. A frame starts at the host with a destination to the target. The port where the appliance is attached to the host switch acts as the virtual initiator and the port where the appliance is attached to the target switch is the virtual target.

Creating a frame redirect zone

The first time the **zone --rdcreate** command is run, the following zone objects are created by default:

- The base zone object, "red_____base".
- The redirect (RD) zone configuration, "r_e_d_i_r_c__fg".

NOTE

Frame redirect zones are not supported with D or I initiator target zones.

ATTENTION

Prior to creating the frame redirect zone, you must create a Layer 2 zone for the Initiator (host) and Target (storage). This zone must be part of the effective configuration and must be defined using the port World Wide Name (WWN). Refer to [Creating a zone](#) on page 427, and [Enabling a zone configuration](#) on page 443.

Use the following procedure to create a frame redirect zone.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **zone --rdcreate** command.
3. Enter the **cfgSave** command to save the frame redirect zones to the defined configuration.

The following example creates a redirect zone, given a host (10:10:10:10:10:10:10:10), target (20:20:20:20:20:20:20:20), virtual initiator (30:30:30:30:30:30:30:30), and virtual target (40:40:40:40:40:40:40:40):

```
switch:admin>zone --rdcreate 10:10:10:10:10:10:10:10 20:20:20:20:20:20:20:20 \
30:30:30:30:30:30:30:30 40:40:40:40:40:40:40:40 restartable noFCR
```

Deleting a frame redirect zone

Use the following procedure to delete a frame redirect zone.

1. Connect to the switch and log in using an account with admin permissions.

Using FC-FC Routing to Connect Fabrics

• FC-FC routing overview.....	535
• Fibre Channel routing concepts.....	539
• Setting up FC-FC routing.....	549
• Backbone fabric IDs.....	551
• Assigning alias names to fabric IDs.....	552
• FCIP tunnel configuration.....	553
• Inter-fabric link configuration.....	553
• FC router port cost configuration.....	559
• Shortest IFL cost configuration.....	561
• EX_Port frame trunking configuration.....	566
• LSAN zone configuration.....	566
• Location embedded LSAN zones.....	583
• Peer LSAN zone support.....	587
• Proxy PID configuration.....	587
• Fabric parameter considerations.....	588
• Inter-fabric broadcast frames.....	588
• Resource monitoring.....	589
• FC-FC routing and Virtual Fabrics.....	590
• Upgrade and downgrade considerations for FC-FC routing.....	594
• Displaying the range of output ports connected to xlate domains.....	595
• FC-FC routing scalability limits.....	595

FC-FC routing overview

The FC-FC routing service provides Fibre Channel routing between two or more fabrics without merging those fabrics.

For example, 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.

Be aware that there are different routing terminologies in use:

- *FC routing* is only in a single fabric (Layer 2 routing). This type of routing is discussed in [Routing Traffic](#) on page 515.
- *FC-FC routing* is routing between two fabrics (Layer 3 routing) and is discussed in this chapter.

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 QoS traffic prioritization over FC routers. Refer to [QoS](#) on page 321 for information about QoS and instructions for setting traffic prioritization over an FC router.

License requirements for FC-FC routing

The following are the License requirements for FC-FC routing:

- A software license might be required for FC-FC routing, depending on the types of fabrics that are connected.

- The Integrated Routing license is required for FC-FC routing between Fabric OS fabrics.
- The Integrated Routing license allows 8-Gbps and 16-Gbps FC ports on Gen 5 platforms to be configured as EX_Ports (or VEX_Ports) supporting FC-FC routing.
- To enable EX_Ports on the Brocade G620 switch, the *Integrated Routing License* needs to be installed. To enable EX_Ports on Gen 6 32-Gbps platform switches excluding the Brocade G620 switch, the *Integrated Routing Port on Demand License* needs to be installed.
- Brocade G630 requires *Integrated Routing Port on Demand License*.
- Enabling the Integrated Routing license and capability does not require a switch reboot.

NOTE

Brocade recommends that all FC routers in a backbone fabric either have the Integrated Routing license or not. You should 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 5 (16 Gbps) Fabric OS platforms:

- Brocade 6510 switch
- Brocade 6520 switch
- Brocade 7840 Extension Switch, with the following conditions:
 - VEX_Ports are not supported.
 - EX_Ports are supported in the base switch.
- Brocade DCX 8510 Backbones:
 - 16-Gbps port blades (FC16-32, FC16-48, FC16-64)
 - FX8-24 extension blade
 - ICL ports on the core blades: EX-Ports on ICL is supported only in DCX 8510-8 and DCX 8510-4 when all the port blades in the chassis belong to one of these blade types: FC16-32, FC16-48, FC16-64.

NOTE

Device discovery will not occur properly with ICL EX_Ports connected edge fabrics if the FC router has unsupported blades that are not mentioned in the previous list.

For the Brocade Backbone families, the backbones have a limit of 128 EX_Ports for each chassis.

FC-FC routing is supported on the following Gen 6 (32 Gbps) Fabric OS platforms:

- Brocade G620
- Brocade G630
- Brocade 7810 Extension Switch, with the following conditions:
 - VEX_Ports are not supported.
 - EX_Ports are supported in the base switch.
- 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 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.

Configuring IFLs using Brocade G630 switches

IFLs can be created between Brocade G630 switches or with FC32-48, FC32-64 port blades on the X6 chassis in integrated routing topologies. The trunk port groups on the G630 switch explained in [Table 87](#) on page 338 applies 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.

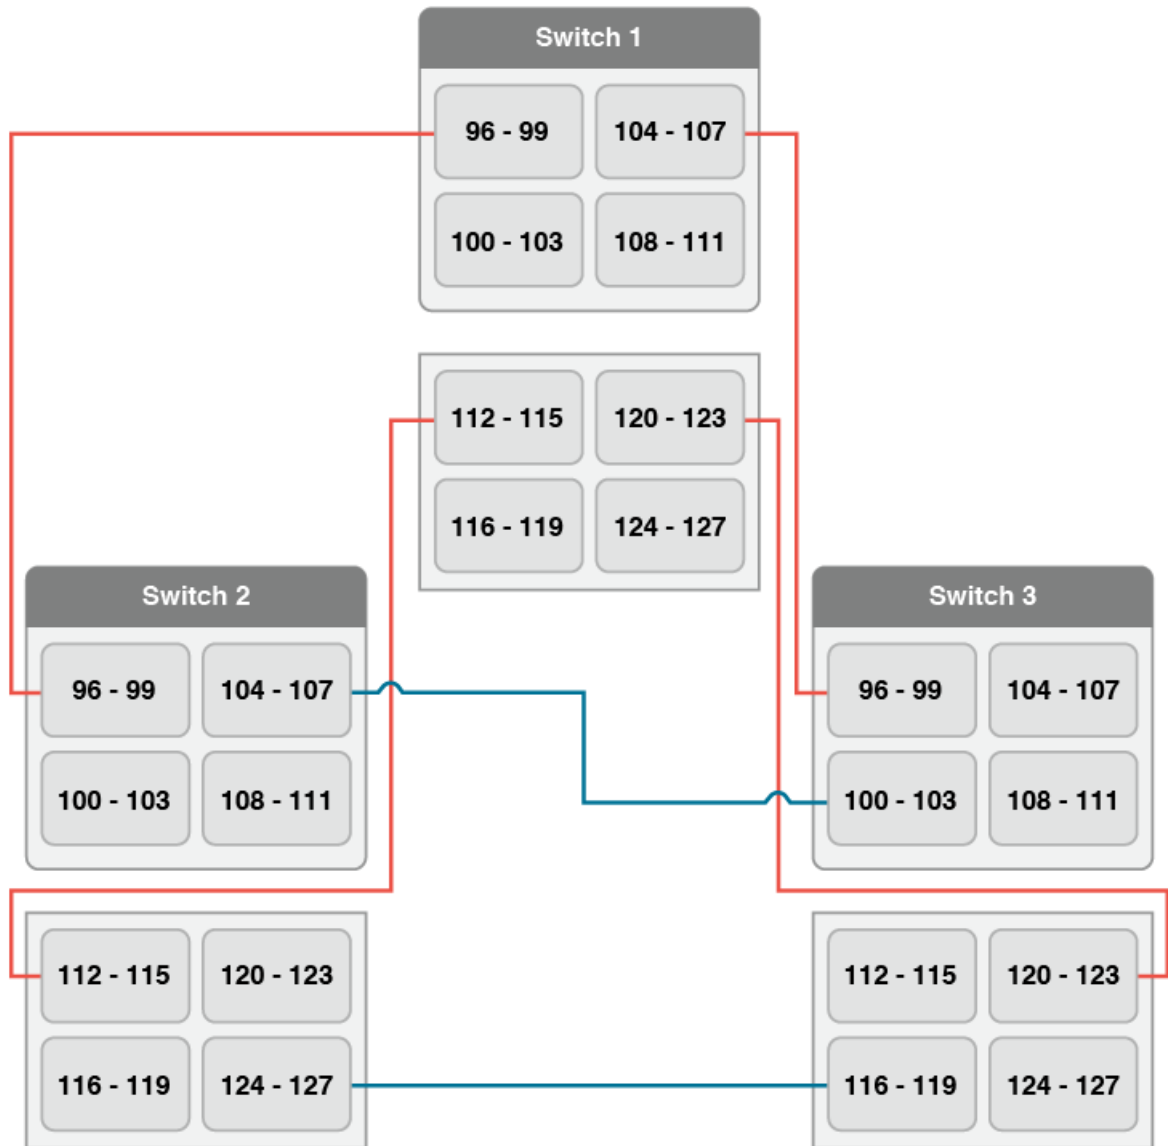
Guidelines for QSFP IFL connections:

- You must select the QSFP ports from the two separate QSFP port groups. Consider switch 1 as 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 port 96-99 (Group 1) and port 112-115 (Group2).
- IFL connection between switch 1 (FC Router) and switch 3 (edge fabric #3) using port 104-107 (Group 1) and port 120-123 (Group 2).
- QSFP to QSFP connection from a G630 switch to QSFP ports on the DCX 8510 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

FC-FC routing supports the following configurations:

- FC router connected to a Fabric OS nonsecured edge fabric.
- FC router connected to a Fabric OS secured edge fabric.
- 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.

VEX edge to VEX edge device sharing is not supported.

NOTE

When there are multiple EX or VEX ports between edge and backbone fabric they will be considered equal cost routes. FSPF does not apply from edge to edge. This needs to be taken into consideration 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 router (FC router)

A switch running the FC-FC routing service. Refer to [Supported platforms for FC-FC routing](#) for a list of platforms that can be FC routers.

- Backbone fabric

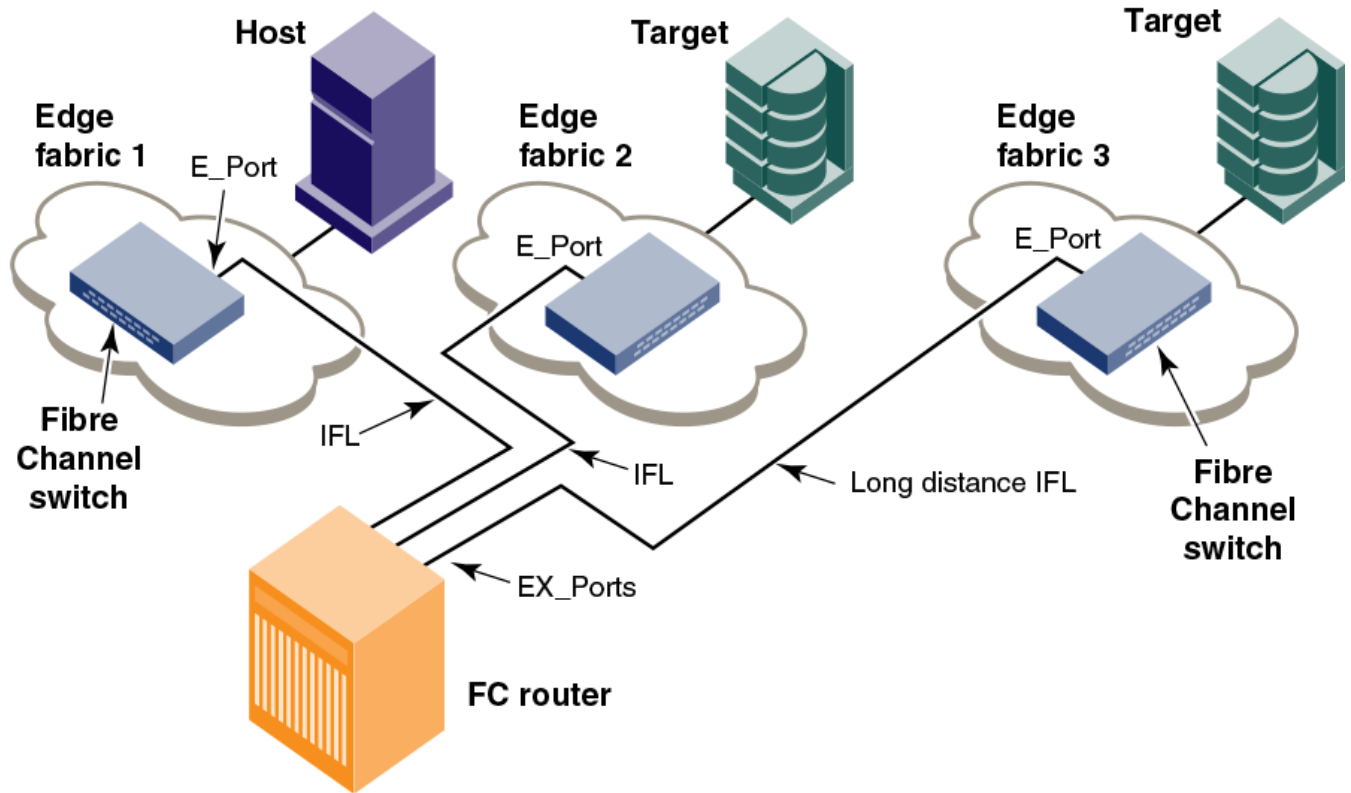
A backbone fabric is an intermediate network that connects one or more edge fabrics. In a SAN, the backbone fabric consists of at least one FC router and possibly a number of Fabric OS-based Fibre Channel switches (refer to [Figure 86](#)).

- Inter-fabric link (IFL)

The link between an E_Port and EX_Port, or VE_Port and VEX_Port, is called an *inter-fabric link* (IFL). You can configure multiple IFLs from an FC router to an edge fabric.

[Figure 84](#) shows a metaSAN consisting of three edge fabrics connected through a Brocade DCX with inter-fabric links.

FIGURE 84 A metaSAN with inter-fabric links



- EX_Port and VEX_Port

An EX_Port and VEX_Port function similarly to an E_Port and VE_Port respectively, but terminate 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.

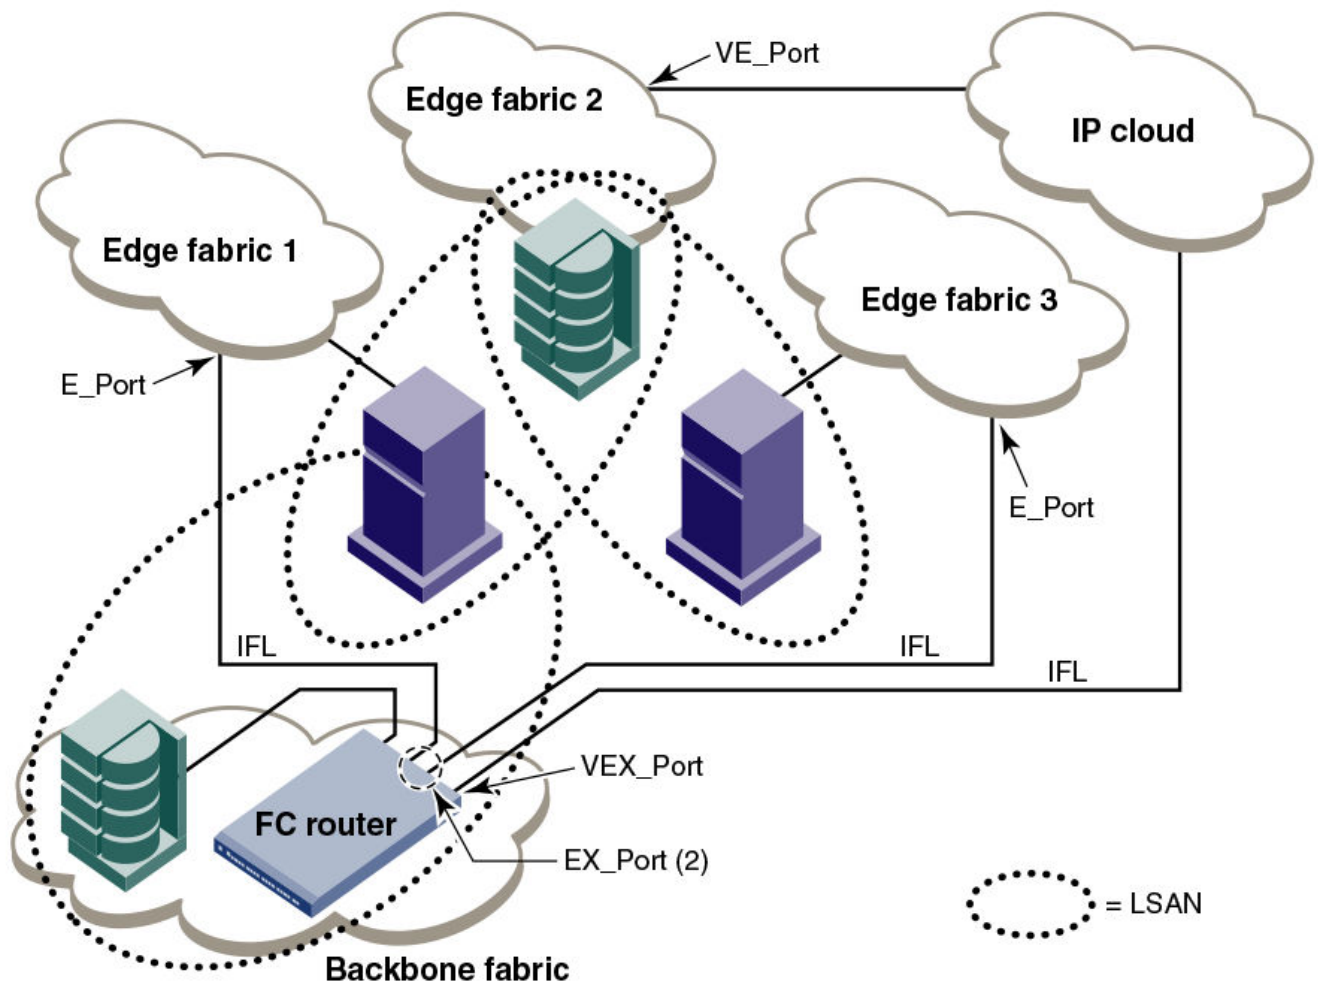
- Logical SANs (LSANs)

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*.

Figure 85 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. Refer to [Proxy devices](#) on page 544 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 and VEX_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 (or multiple VEX_Ports) are attached to the same edge fabric, they must be configured with the same FID.

- If EX_Ports and VEX_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.

[Figure 86](#) 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 Fibre Channel router. The FC router can emulate two types of phantom domains: front phantom domains and translate phantom domains. For detailed information about phantom domains, refer to [Phantom domains](#) on page 545.

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.

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. Refer to [LSAN zone configuration](#) on page 566 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 translate 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, then 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, then 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 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 88. 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 or VEX_Ports connected to an edge fabric use 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, refer to [Identifying and deleting stale xlate domains](#) on page 548.

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 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 **fcrxlateconfig** 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 **fcrxlateconfig** 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 **fcrxlateconfig --show** to identify any stale xlate domains.
3. Enter **fcrxlateconfig --del** to delete the stale xlate domains.

```
switch:admin> fcrxlateconfig --show stalexd
Imported FID      Stale XD      Owner Domain
-----
      012          002          007 ( this FCR )

switch:admin> fcrxlateconfig --del stalexd 12 2
Xlate domain 2 is deleted
```

Domain ID range for front and translate phantom domains

In the releases prior to Fabric OS 7.4.0, the following domain IDs were assigned:

- The FC router front domain is assigned 160 as the domain ID if the insistent domain ID is not configured using the **portCfgExport** command.
- The FC router translate domain is assigned 1 as the domain ID for importing proxies to the edge fabric if the insistent domain ID is not configured using the **fcrXlateConfig** command.

Starting with Fabric OS 7.4.0, 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 translate domain, the FC router will still request 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 and/or translate domain, the FC router requests the preferred domain ID. Phantom domain IDs are stored persistently and used in RDI request. To utilize the new range (160 to 239), do one of the following tasks:

Option 1:

1. Convert the EX_Port to non-FC router port remove existing Xlate configuration using the **fcrXlateConfig** command.
2. Apply the default configuration, and then configure the EX_Ports again and remove existing Xlate configuration using the **fcrXlateConfig** command..

Option 2:

1. Disable all EX_Ports in the same edge fabric and persistent XD feature, and then reset the persistent phantom domain using the **fcrConfigure -resetPhantomDomain** command.

```
switch:admin> fcrconfigure --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> fcrconfigure -resetphantomdomain -force
Phantom Domain IDs were successfully reset to default range
```

2. Enable all the EX_Ports.

TABLE 120 Front and translate domain IDs

Firmware version in FC router	Front domain ID requested	Translate domain ID requested
Prior to Fabric OS 7.4.0	160	1
Fabric OS 7.4.0 and later	160	200

When an FC router is running Fabric OS 7.4.0 and the edge FC router is running Fabric OS 7.3.0, or vice versa, you need to use a different FID to reconstruct the EX_Ports or disable the EX_Port or use the **fcrXlateConfig** command to delete the stale translate domains when moving from Fabric OS 7.3.0 to Fabric OS 7.4.0 as the old translate domains are still in effect.

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, the front phantom WWN should be used 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. Refer to [Secret key pairs for DH-CHAP](#) on page 251 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. (Refer to [Verifying the setup for FC-FC routing](#) on page 550.)
2. Assign backbone fabric IDs. (Refer to [Backbone fabric IDs](#) on page 551.)
3. Configure FCIP tunnels if you are connecting Fibre Channel SANs over IP-based networks.
4. Configure IFLs for edge and backbone fabric connection. (Refer to [Inter-fabric link configuration](#) on page 553.)

5. Modify port cost for EX_Ports, if you want to change from the default settings. (Refer to [FC router port cost configuration](#) on page 559.)
6. Enable shortest IFL mode if you want to choose a lowest cost IFL path in the backbone fabric. (Refer to [Shortest IFL cost configuration](#) on page 561.)
7. Configure trunking on EX_Ports that are connected to the same edge fabric. (Refer to [EX_Port frame trunking configuration](#) on page 566.)
8. Configure LSAN zones to enable communication between devices in different fabrics. (Refer to [LSAN zone configuration](#) on page 566.)

Refer to [Performing Advanced Configuration Tasks](#) on page 71 for more details about configuration options for Brocade Backbones.

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 v7.3.0 or later is installed on the FC router, as shown in the following example.

```
switch:admin> version
Kernel:      2.6.14.2
Fabric OS:   v7.3.0
Made on:     Fri Feb 21 01:15:34 2014
Flash:       Wed Feb 26 20:53:48 2014
BootProm:    1.0.11
```

2. If you are configuring a Brocade DCX 8510 Backbone, enter the **slotShow** command to verify that an FX8-24 blade is present or 16-Gbps port blade is present. The following example shows slots 1, 2, 3, 4, 9, 10, and 12 with 16-Gbps port blades enabled.

```
switch:admin> slotshow -m
Slot  Blade Type      ID   Model Name      Status
-----
  1     SW BLADE           96   FC16-48         ENABLED
  2     SW BLADE           96   FC16-48         ENABLED
  3     SW BLADE           96   FC16-48         ENABLED
  4     SW BLADE           96   FC16-48         ENABLED
  5     CORE BLADE         98   CR16-8          ENABLED
  6     CP BLADE           50   CP8             ENABLED
  7     CP BLADE           50   CP8             ENABLED
  8     CORE BLADE         98   CR16-8          ENABLED
  9     SW BLADE           96   FC16-48         ENABLED
 10     SW BLADE           96   FC16-48         ENABLED
 11     UNKNOWN            96   VACANT          VACANT
 12     SW BLADE           96   FC16-48         ENABLED
```

Refer to [Performing Advanced Configuration Tasks](#) on page 71 for a list of blades and their corresponding IDs.

3. Enter the **licenseShow** command to verify that the Integrated Routing license is installed.

```
switch:admin> licenseshow
S9bddb9SQbTAceeCbzbzRcbcSc0c0SYRyeSzRSscycazft0G:
  Integrated Routing Ports on Demand license
  Capacity 128
```

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 Guide*.

4. For configuring EX_Ports on an ICL, both the Integrated Routing license and the ICL Ports on Demand license are required. The **licenseShow** command displays the license port capacity as shown in the following example.

```
switch:admin> licenseshow
SEEL7KaF9KtSN3HCHLBDWBAQtmgADPtFE3RaTSGAED4A:
  Inter Chassis Link (ICL) license
  Capacity 16 (4 QSFPs per CR blade or 8 QSFPs per chassis)
```

5. If all switches in the fabric are running Fabric OS 7.3.0 or later, skip this step. If any switch is running a Fabric OS version earlier than 7.3.0, then enter the **fddCfg --showall** command to verify that the Fabric-Wide Consistency Policy is not in "strict" mode. When it is in strict mode, an ACL cannot support Fibre Channel routing in the fabric.

```
switch:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
  DATABASE - Accept/Reject
-----
      SCC -      accept
      DCC -      accept
      PWD -      accept
      FCS -      accept
      AUTH -     accept

Fabric-Wide Consistency Policy :- "SCC:S;DCC"
```

If the Fabric-Wide Consistency Policy has the letter "S" in it in the edge fabric or the backbone fabric, do not connect the edge fabric to the FC router. The letter "S" (shown in the preceding sample output) indicates the policy is strict. The Fabric-Wide Consistency Policy must be tolerant before you can connect fabrics to the FC router. Refer to [Configuring Security Policies](#) on page 233 for information about configuring the Fabric-Wide Consistency Policy.

6. For 8-Gbps platforms, delete fabric mode Top Talker monitors, if they are configured.

FC-FC routing and fabric mode Top Talker monitors are not concurrently supported on 8-Gbps platforms.

FC-FC routing and fabric mode Top Talker monitors are concurrently supported only on the Brocade 6510 and 6520 switches, and on the Brocade DCX Backbone family with only 16-Gbps capable ports.

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 Fibre Channel 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.
5. Verify 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** and **portcfgvexport** commands.

There is no relation 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.
3. Verify the configured alias names using the **fcrconfigure --show -alias** command.

```
fcrconfigure --add -alias alias_name -fid fid
```

```
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 Fibre Channel ISL and IFL traffic. FCIP is a prerequisite for configuring VEX_Ports. 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. Once a tunnel is created, it defaults to a disabled state. Then configure the VE_Port or VEX_Port. After the appropriate ports are configured, enable the tunnel.

Refer to the *Brocade Fabric OS FCIP Administration 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 inter-fabric link, be aware that you cannot configure both IFLs (EX_Ports, VEX_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 or VEX_Port using either the **portcfgvexport** or **portcfgexport** command.
 - **portcfgvexport** works only on VE_Ports.
 - **portcfgexport** (only on the FC ports on the FC router) commands work 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 **fcredgeshow**, **fcrxlateconfig**, **fcrconfigure**, and **fcrproxyconfig** in the *Brocade Fabric OS Command Reference*.

A Fibre Channel router can interconnect multiple fabrics. EX_Ports or VEX_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](#) on page 559.
4. Set up ISL or EX_Port trunking. For information on trunking setup, refer to [Configuring EX_Port trunking](#) on page 348.
5. Physically attach ISLs from the Fibre Channel 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 ports that are persistently disabled.

NOTE

FC ports on Brocade 7800 and 7840 switches and FX8-24 blades are configured as persistently disabled by default to avoid inadvertent fabric merges when installing a new FC router.

```
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 such 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 **switchshow** to verify the EX_Port (or VEX_Port), edge fabric ID, and name of the edge fabric switch (containing the E_Port or VE_Port) are correct.

11. Enter **fcrfabricshow** to view any edge fabric switch names and ensure links are working as expected.

NOTE

The **fcrfabricshow** command displays the static IPv6 addresses for each FC router and each edge fabric switch connected to the EX_Ports.

```
switch:admin> fcrfabricshow
FCR WWN: 10:00:00:05:1e:13:59:00, Dom ID: 2, Info: 10.32.156.52 1080::8:800:200C:1234/64,"Spirit-2"
"fcr_5300"
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)
-----
7 10 10:00:00:05:1e:34:11:e5 10.32.156.33 "My5300" 1080::8:8FF:FE0C:417A/64
4 116 10:00:00:05:1e:37:00:44 10.32.156.34 "My5300"
FCR WWN: 10:00:00:05:1e:12:e0:00, Dom ID: 100, Info:10.32.156.50 1080::8:60F:FE0C:456A/64
"fcr_5300"
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)
-----
4 95 10:00:00:05:1e:37:00:45 10.32.156.31 "My5300"
FCR WWN: 10:00:00:05:1e:12:e0:00, Dom ID: 100, Info: 10.32.156.50, "fcr_Brocade 5300"
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)
-----
4 95 10:00:00:05:1e:37:00:45 10.32.156.31 "Brocade 5300"
5 95 10:00:00:05:1e:37:00:45 10.32.156.31 "Brocade 5300"
6 95 10:00:00:05:1e:37:00:45 10.32.156.31 "Brocade 5300"
```

12. Enter **iflshow** 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:

- For DCX 8510 Backbones, both the active and standby CP must be running Fabric OS 7.2.0 or later.
- EX_Port on ICL is supported in DCX 8510 Backbones and Brocade X6 Directors when all the port blades in the chassis belong to one of these blade types: FC16-32, FC16-48, FC16-64, or FC32-48.

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 the previous restriction.

- Switches must have Virtual Fabrics enabled.
- EX_Ports on ICLs are allowed only on the base switch. F_Ports are not supported in the base switch.
- You must be in Brocade native mode.
- You cannot configure a VEX_Port on an ICL.
- The ICLs must be in a symmetric topology. For more information on topologies, refer to the chapter on Inter-Chassis Links (ICLs).
- 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.

```
switch:admin> configure
Configure...
Enable a 256 Area Limit
(0 = No,
1 = Zero Based Area Assignment,
2 = Port Based Area Assignment): (0..2) [0]
```

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:admin> portdisable 6/20-23
```

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 3 QSFP ports.

The following example configures EX_Port on one of the QSFP ports.

```
switch:admin> portcfgexport 6/20 -a 1 -f 45
2013/04/25-21:21:54, [FCR-1071], 29805, SLOT 4 | FID 2, INFO, Pluto, Port 6/20 is changed from non FCR port to FCR port.
2013/04/25-21:21:54, [FCR-1071], 29806, SLOT 4 | FID 2, INFO, Pluto, Port 6/21 is changed from non FCR port to FCR port.
2013/04/25-21:21:55, [FCR-1071], 29807, SLOT 4 | FID 2, INFO, Pluto, Port 6/22 is changed from non FCR port to FCR port.
2013/04/25-21:21:55, [FCR-1071], 29808, SLOT 4 | FID 2, INFO, Pluto, Port 6/23 is changed from non FCR port to FCR port.
```

4. Configure the 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](#) on page 559.
5. Enter **portenable portrange** to enable the QSFP ports that you disabled in step 1.

```
switch:admin> portenable 6/20-23
```

6. Use the **portcfgexport** or **portshow** commands to verify that each port is configured correctly.

```
switch:admin> portcfgexport 6/20
Port 6/20 info
Admin:                enabled
State:                OK
Pid format:           Not Applicable
Operate mode:         Brocade Native
Edge Fabric ID:       50
Preferred Domain ID: 160
Front WWN:            50:00:51:e7:54:c0:0e:32
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
```

7. Enter **switchshow** to verify the EX_Port and edge fabric ID are correct.
8. Enter **fcuredgeshow** to display the FIDs of all configured EX_Ports.

```
switch:admin> fcuredgeshow
  FID EX-port E-port  Neighbor Switch (PWPN, SWWN )
-----
  50  3/12    1180   50:00:51:e4:8f:8f:74:9c  10:00:00:05:1e:48:f8:00
  50  3/13    1181   50:00:51:e4:8f:8f:74:9d  10:00:00:05:1e:48:f8:00
  50  3/14    1182   50:00:51:e4:8f:8f:74:9e  10:00:00:05:1e:48:f8:00
  50  3/15    1183   50:00:51:e4:8f:8f:74:9f  10:00:00:05:1e:48:f8:00
```

9. Enter the **fcriclpathbwmonitor --show** 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 displays the ICL path bandwidth state for each fabric:

```
switch:admin> fcriclpathbwmonitor --show
ICL Path Bandwidth state :Disabled

switch:admin> fcriclpathbwmonitor --enable
ICL bandwidth balance functionality is enabled

switch:admin> fcriclpathbwmonitor --show
ICL Path Bandwidth state :Enabled
FABRIC      SLOT-3 BW      SLOT-6 BW      STATE
=====
    48         128         128      BALANCED
   126         64         128      UNBALANCED
```

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 port link cost values are either 0, 1,000, or 10,000. 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, then the lowest cost path is used. If exchange-based routing has not been disabled and multiple paths exist with the same lowest cost, there will be load sharing over these paths.

Every IFL has a default cost. The default router port cost values are as follows:

- 1,000 for a legacy (v5.1 or XPath FCR) IFL
- 1,000 for an EX_Port IFL
- 10,000 for a VEX_Port IFL

The FC router port cost settings are 0, 1,000, 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 is saved in the existing port configuration file.

FC router port cost is 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 **IsDbShow** 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:

- Router port sets are defined as follows:
 - 0-7 and FCIP Tunnel 16-23

- 8-15 and FCIP Tunnel 24-31
- The router port cost does not help distinguish one IFL (or EX_ and VEX_Port link) from another, if all the IFLs are connected to the same port set. Therefore, if you connect IFL1 and IFL2 to the same edge fabric in port set 0-7 and then configure them to different router port costs, traffic is still balanced across all the IFLs in the same port set.
- 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 0-7) than the higher router cost IFL (for example, ports 8-15). For VEX_Ports, you would use ports in the range of 16-23 or 24-31.

You can connect multiple EX_Ports or VEX_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 and VEX_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 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 or VEX_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. Enter the **portDisable** command to disable any port on which you want to set the router port cost.

```
switch:admin> portdisable 7/10
```

2. Enable EX_Port or VEX_Port mode with the **portCfgEXPort** or **portCfgVEXPort** command.

```
switch:admin> portcfgexport 7/10 -a 1
```


- 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
```

- 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
```

- Enter the **portEnable** command to enable the ports that you disabled in step 1.

```
switch:admin> portenable 7/10
```

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 in order 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. Shortest IFL mode should be enabled on all 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:

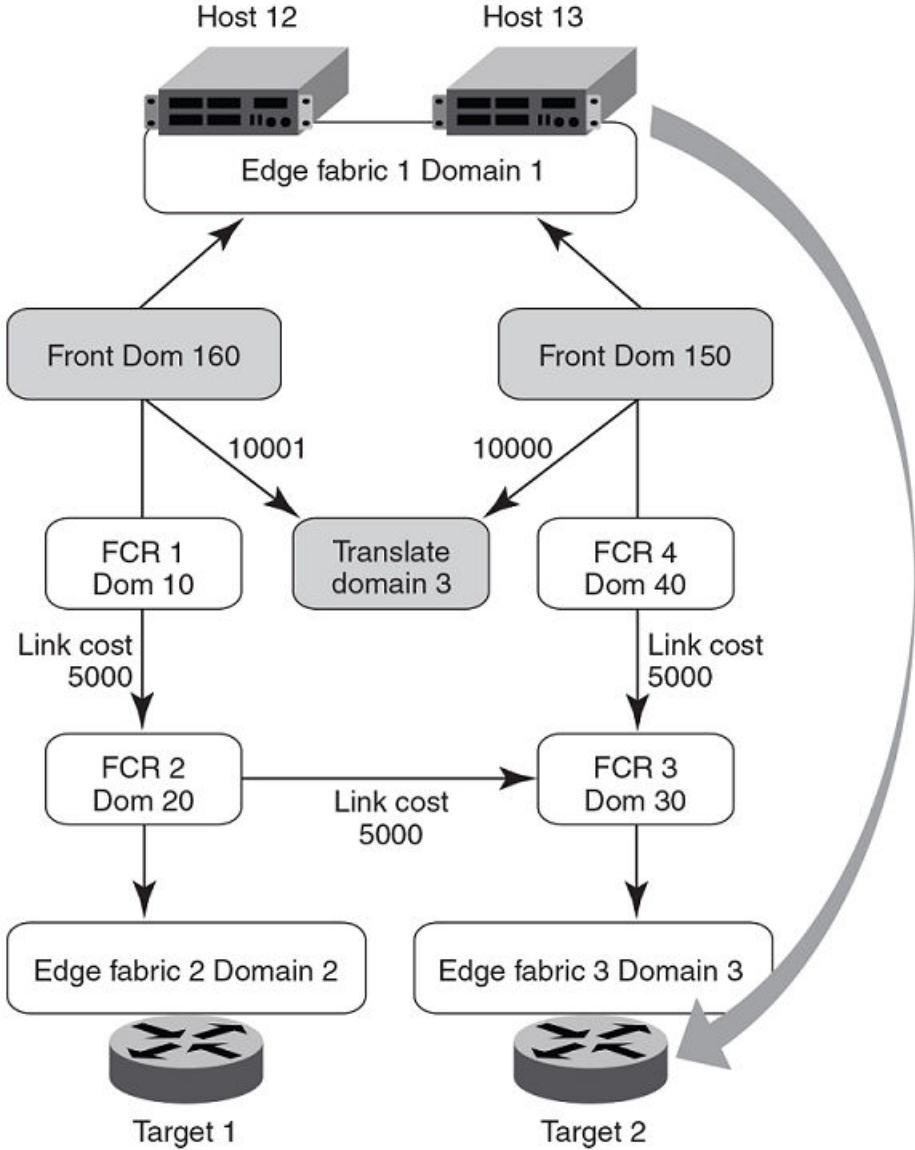
- Path 1: Edge fabric 1-FCR1 - FCR2 -FCR3 -Edge fabric 3.
- Path 2: Edge fabric 1-FCR4 - FCR3 -Edge fabric 3.

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 shortest IFL cost

1. Enter the **fcrFabricShow** command to view the FC routers on the backbone fabric.

```
switch:admin>fcrfabricshow
FC Router WWN: 10:00:00:05:1e:58:bd:69, Dom ID: 10,
Info: 10.17.33.59, "DID_10"
EX_Port   FID   Neighbor Switch Info (enet IP, WWN, name)
-----
    34     1   10.17.33.68     10:00:00:05:1e:61:28:22  "DID_4_1"
switch:admin>fcrfabricshow
FC Router WWN: 10:00:00:05:1e:58:be:69, Dom ID: 20,
Info: 10.17.33.60, "DID_20"
EX_Port   FID   Neighbor Switch Info (enet IP, WWN, name)
-----
    2      2   10.17.33.88     10:00:00:05:33:00:90:48  "DID_1_2"
switch:admin>fcrfabricshow
FC Router WWN: 10:00:00:05:1e:58:be:68, Dom ID: 30,
Info: 10.17.33.61, "DID_30"
EX_Port   FID   Neighbor Switch Info (enet IP, WWN, name)
-----
    31     3   10.17.33.153    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: 10.17.33.62, "DID_40"
EX_Port   FID   Neighbor Switch Info (enet IP, WWN, name)
-----
    34     1   10.17.33.68     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 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 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, there are two available paths 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 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 **IsDbShow** 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 frame trunking configuration

EX_Port frame 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 E_Ports. The restrictions for EX_Port frame trunking are the same as 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 master port. The other ports that come online on the trunk group are considered the slave ports. Adding or removing a slave port does not cause frame drop. However, removing a slave port causes the loss of frames in transit.

If router port cost is used with EX_Port trunking, the master port and slave ports share the router port cost of the master port.

For information about setting up E_Port trunking on an edge fabric, refer to [Managing Trunking Connections](#) on page 335.

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.

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, in [Figure 86](#) on page 543, 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 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 (and optionally on the backbone fabric as well), 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, then 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 or VEX_Port
 - Switch2 is connected to the FC router using another EX_Port or VEX_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.

- 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 10.38.141.145 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 10.38.141.145 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 }

```

- 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"
```

- Enter **cfgshow**. This 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

```


- 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 ...
```

- 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 ...
```

- Log in as admin and connect to switch1.
- Enter `nshow` to identify the WWN of the host. (Here it is 10:00:00:00:c9:2b:c9:0c.)

The `nshow` output displays the LSAN zone status of a device, the port WWN, and the node WWN; the port WWN must be used for LSANs.

NOTE

In the following output **FC4s: FCP APS** means Fibre Channel Protocol Application Services.

```
switch:admin> nshow
{
  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 }
```

- 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"
```

- 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:

- **lsanzone**show -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> lsanzone
```

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

- **lsanzone**show -o -sort displays the WWNs in sorted order. The **lsanzone**show -o command applies to peer zone rules and displays the WWNs in sorted order if LSAN zones are peer zones.
- **fcrphydev**show shows the physical devices in the LSAN.

```
switch:admin> fcrphydev
```

Device	WWN	Physical
Exists		PID
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

- **fcrproxydev**show shows the proxy devices in the LSAN.

```
switch:admin> fcrproxydev
```

Proxy	WWN	Proxy	Device	Physical	State
Created		PID	Exists	PID	
in Fabric			in Fabric		

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](#) on page 567. 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 5,000 without disabling the switch. The maximum number of LSAN devices supported is 10,000 (this includes both physical and proxy devices). If you have 3,000 LSAN zones but have not exceeded the 10,000 device limit, you can increase the LSAN count to 5,000. All FC routers in the same backbone fabric should have the same maximum LSAN count defined, to prevent the FC routers from running into 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 8.1.0 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 DCX 8510, Brocade X6 Directors and Gen 6 fixed-port switches are recommended as FC router domains in the metaSAN. Having only Gen 5 or older fixed-port switches as the FC router domain in the metaSAN is not recommended. Downgrading to an older Fabric OS version is blocked when the configuration exceeds the scalability numbers for the earlier version.

- Maximum number of FC routers in the backbone fabric are increased from 12 to 16.
 - Maximum number of LSAN zones in a metaSAN are increased from 5000 to 7500. This can be scaled beyond 7500 by using the location-embedded LSAN zones.
 - Maximum number of LSAN devices in a metaSAN from 10000 to 15000
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, refer to [Resource monitoring](#) on page 589.

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 will not be 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 **fcresource** command.

Firmware downgrade from Fabric OS 8.1.0 is blocked if the limits are higher than that is supported in Fabric OS 8.0.1 or earlier. The **lsanzoneshow** command is enhanced with following two new options to display which LSAN Zones/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

Be aware of how LSAN zones impact high availability and firmware downgrades:

- The LSAN zone matrix is synchronized to the standby CP.
- On a dual CP switch, both CPs must have Fabric OS v5.3.0 or later.
- 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 to pre-Fabric OS 8.0.1 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. Use the Enforce tag to achieve better scalability in the 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 could 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.


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 and VEX_Ports are located. You then must ensure that the LSAN zones in the edge fabrics incorporate the tags correctly.
- The LSAN tags are configured per FC router, not per fabric. If the backbone fabric has multiple FC routers, it is recommended that you 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:

```
switchdisable
```

3. Enter the following command to create an Enforce LSAN tag:

```
fcrlsan --add -enforce tagname
```

The *tagname* variable is the name of the LSAN tag you want to create.

4. Enter the following command to enable the FC router:

```
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 121 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	FC router 1	FC router 2	FC router 3	FC router 4
LSAN 1	LSAN 1	LSAN 1	LSAN 1	LSAN 1	LSAN 2	LSAN 3	LSAN 4	LSAN 1	LSAN 2	LSAN 3	LSAN 4
LSAN 2	LSAN 2	LSAN 2	LSAN 2	LSAN 2		LSAN 4		LSAN 2			
LSAN 3	LSAN 3	LSAN 3	LSAN 3								
LSAN 4	LSAN 4	LSAN 4	LSAN 4								

LSAN zone binding considerations

- Without LSAN zone binding, the maximum number of LSAN devices is 10,000.
- With LSAN zone binding, the metaSAN can import more than 10,000 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 used for creating and updating the LSAN fabric matrix.

NOTE

Best practice: Use the LSAN zone binding feature in a backbone fabric in which all FC routers are running Fabric OS v6.1.0 or later.

When you set up LSAN zone binding on the local FC router (running Fabric OS v6.1.0 or later), the resultant matrix database is automatically distributed to all of the Fabric OS v6.1.0 or later FC routers in the backbone fabric. You do not need to set up LSAN zone binding on the other FC routers unless those FC routers are running Fabric OS versions earlier than v6.1.0.

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.

Note the following for FC routers running a Fabric OS version earlier than 6.1.0:

- The matrix database is not automatically distributed from this FC router to other FC routers.
- You must manually configure the LSAN fabric matrix on these FC routers to match the other FC routers in the backbone fabric.

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. Refer to [LSAN zone policies using LSAN tagging](#) on page 574 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.

- 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.

- 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

- Log in to the FC router as admin.
- Enter **fcrlsanmatrix --fabricview -fcr** to view the FC router matrix.
- Enter **fcrlsanmatrix --fabricview -lsan** to view the LSAN fabric matrix.

The following example displays both fabric matrices.

```
device:admin> fcrlsanmatrix --fabricview -fcr
SAVED FCR PAIRS
=====
FCR                                FCR
-----
10:00:00:60:69:c3:12:b2 (2)        10:00:00:60:69:c3:12:b3 (unknown)

device:admin> fcrlsanmatrix --fabricview -lsan
LSAN MATRIX is activated
Fabric ID      Fabric ID
-----
      4          5
      4          7
     10         19
```

Location embedded LSAN zones

Starting with Fabric OS 7.4.0, 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.

Prior to Fabric OS 7.4.0, 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.

NOTE

In a backbone fabric having a mix of Fabric OS 7.4.0 and previous releases, there can be mismatch in location embedded LSAN zones in FC routers running Fabric OS pre-v7.4.0. Hence, such mixed combination is not recommended for this feature.

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 needs to 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, then 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
```

- The RFID tag should be appended after the enforce and speed tags in case you need to use the enforce and the 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 is as follows:
 1. The FC router processes the LSAN zones based on enforce tags during zone merge. Zones not matching these tags are dropped.
 2. The LSAN distribution between FC routers based on RFID is done only on the filtered zones obtained from step a.


NOTE

Speed tags remain unaffected during this process.

3. 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. Once 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

- Each FC router has different LSAN zones based on the location embedded in the LSAN zone name. If each FC router has 5000 LSANs, after firmware downgrade to a pre-7.4.0 version, any disruptive action in the metaSAN triggers import or export issues because the pre-7.4.0 FC router accepts only the first 5000 LSAN zones. In such an FC router, you will notice the RAS log FCR-1020 indicates that the FC router LSAN zones are exhausted. Also, the filter will apply depending on where the pre-7.4.0 switch is in the backbone fabric.
- 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 **fcrProxyDevShow -a**, **fcrPhyDevShow -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 recommended steps are as follows:
 1. Disable the EX_Ports that need an FID change.
 2. Configure the EX_Ports with the new FID value.
 3. Enable the zone configuration in the respective edge fabrics with the new RFID value.
 4. Enable the EX_Ports that were disabled in step 1.
- If you configure Backbone Fabric ID value for the RFID field, the LSAN zone is treated as 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 details, refer to [Peer Zoning](#) on page 464.

NOTE

Initial WWNs starting with '00' are for internal use only and are auto-created. From Fabric OS 7.3.0 onwards, zones having such initial WWNs are treated as peer zones. If you upgrade the firmware from pre-7.3.0, or merge with a switch that has pre-7.3.0 firmware, or download configuration from pre-7.3.0 switch, and the first zone member WWN starts with '00', Fabric OS treats the zone as a peer zone. Please delete such invalid zones by first using the **cfgRemove** command, then the **zoneDelete** command and finally the **cfgSave/cfgEnable** commands.

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 **fcrProxyConfig** command displays or sets the persistent configuration of proxy devices. Used with the **-s** option, it can also influence the assignment of the xlate domain port number (which is used to determine 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 will be ignored. Proxy PIDs are automatically assigned to devices imported into a fabric, starting at f001.

Use the **fcrXlateConfig** command to display or assign a preferred domain ID to a translate domain.

Fabric parameter considerations

By default, EX_Ports and VEX_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.
- To change the fabric parameters of a VEX_Port, use the **portCfgVEXPort** 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 or VEX_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 or VEX_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 VEX_Port must match those values on other Fabric OS switches. You do not need to adjust these parameters for an EX_Port or VEX_Port unless you have adjusted them for the edge fabric.

The default values for R_A_TOV and E_D_TOV are the recommended values 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. Many target devices and HBAs cannot handle broadcast frames. In this case, you can set up broadcast zones to control which devices receive broadcast frames. (Refer to [Broadcast zones](#) on page 422 for information about setting up broadcast zones.)

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> fcrbcstconfig --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*.

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. Refer to [Setting the maximum LSAN count](#) on page 572 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 Node WWNs
- Phantom Port WWNs
- 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:      3000              28
LSAN Devices:    10000             51
Proxy Device Slots: 10000             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 and VEX_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 and VEX_Ports are persistently disabled with the reason *Export in non base switch*. You must explicitly create a base switch, move the EX_Ports and VEX_Ports to the base switch, and then enable the ports.

If you move existing EX_Ports or VEX_Ports to any logical switch other than the base switch, these ports are automatically disabled.

If you want to change an EX_Port or VEX_Port on the logical switch to be a non-EX_Port or VEX_Port, you must use the **portCfgDefault** command. You cannot use the **portCfgExport** command because that command is allowed only on the base switch.

- 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 and VEX_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 or VEX_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, then the EX_Port or VEX_Port is disabled. Refer to [Configuring a logical switch for XISL use](#) on page 409 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.

ATTENTION

If you connect an EX_Port or VEX_Port from an FC router running Fabric OS v6.1.x or earlier to a logical switch that allows XISL use, the EX_Port or VEX_Port is *not* disabled. However, this configuration is not supported.

- Backbone-to-edge routing is not supported in the base switch. Refer to [Figure 48](#) on page 390 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 **fcrConfigure** command is not allowed when Virtual Fabrics is enabled. Instead, use the **lsCfg** command to configure the FID.
- Although the Brocade 6510 and Brocade 6520 support up to four logical switches, if you are using FC-FC routing, they can have a maximum of only three logical switches.
- In the Brocade 7840, FC-FC routing is not supported on the base switch.
- When a VF enabled switch is part of backbone, only the base switch will be able to route frames between edge fabrics. If any other (default or logical) switch is included as a hop between edge fabrics, packets will be dropped.

Logical switch configuration for FC routing

[Figure 95](#) 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 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

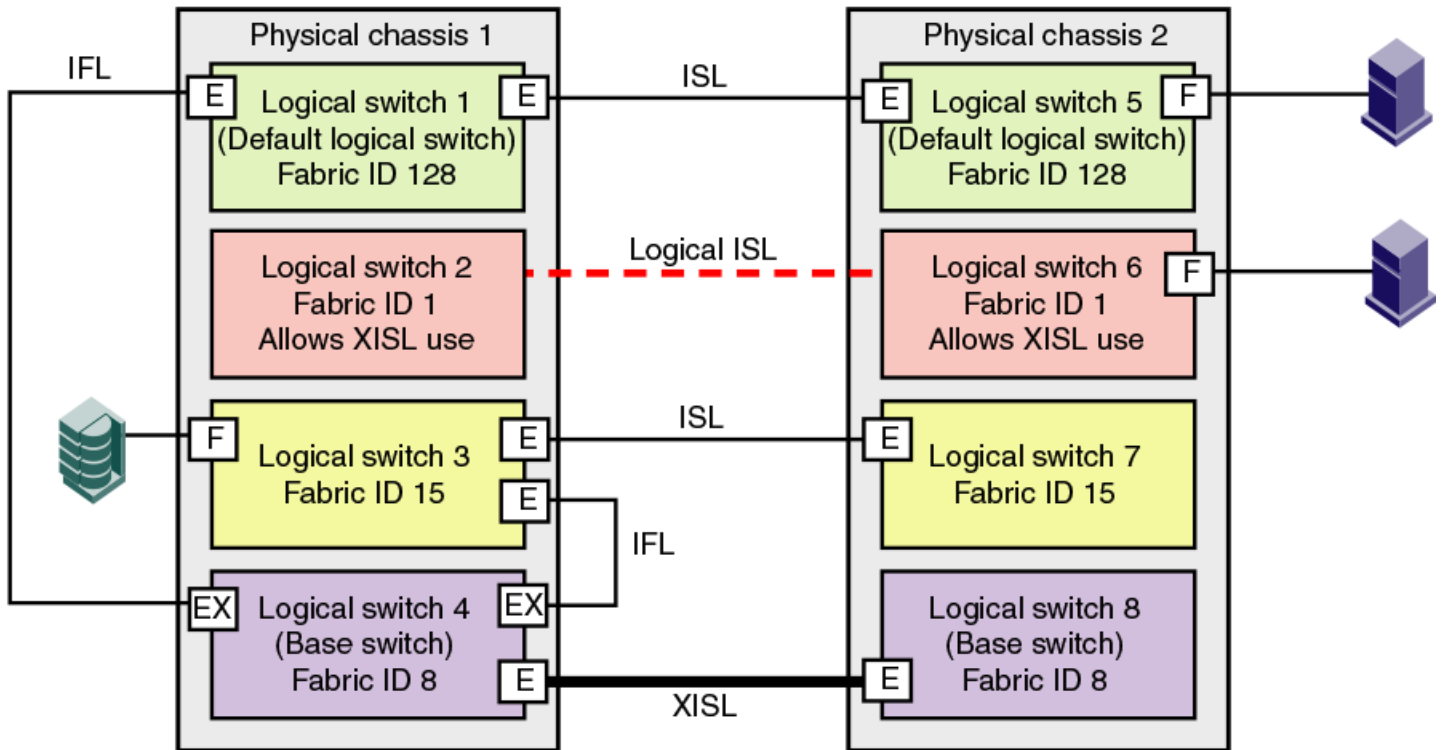



Figure 96 shows a logical representation of the physical chassis and devices in Figure 95. As shown in Figure 96, Fabric 128 and Fabric 15 are edge fabrics connected to a backbone fabric. Fabric 1 is not connected to the backbone, so the device in Fabric 1 cannot communicate with any of the devices in the other fabrics.

FIGURE 96 Logical representation of EX_Ports in a base switch



Backbone-to-edge routing with Virtual Fabrics

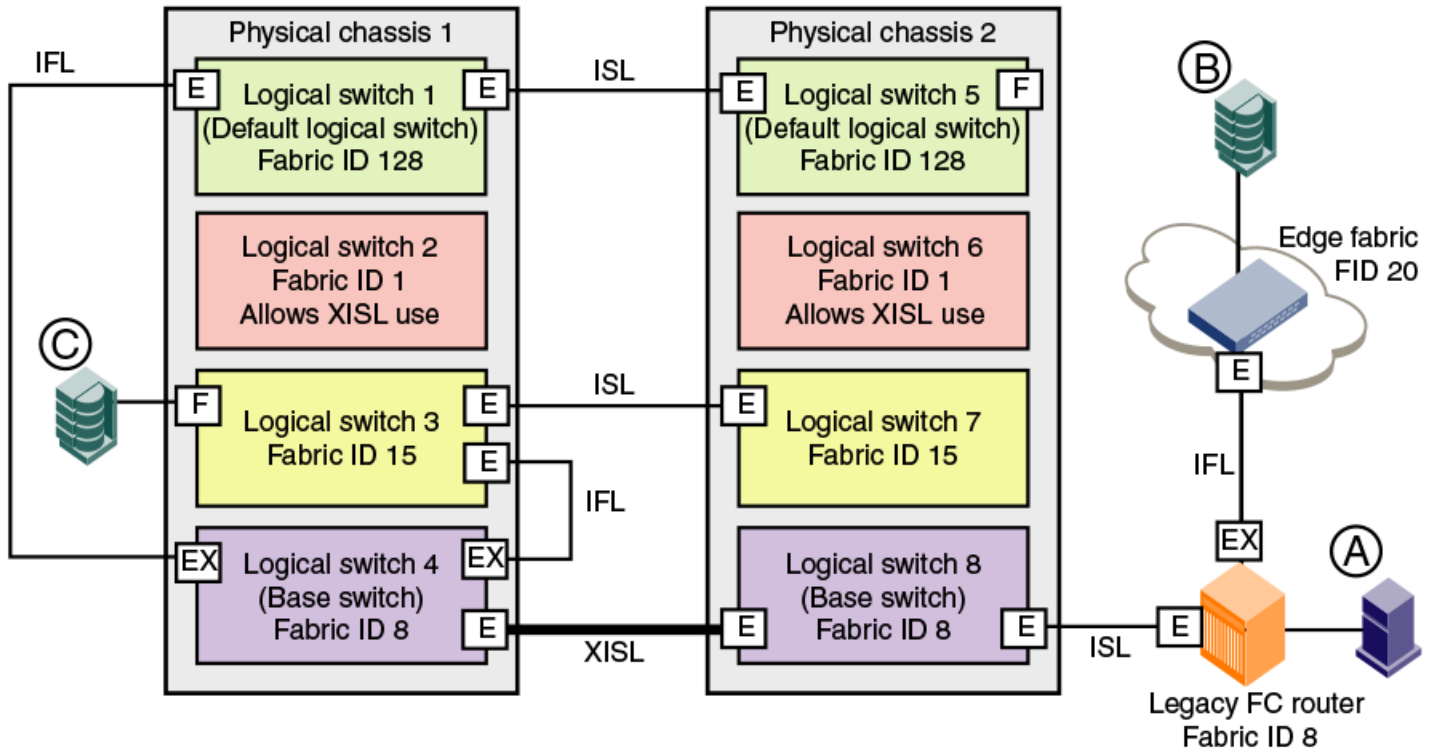
Backbone-to-edge routing is not supported in the base switch, unless you use a legacy FC router. A *legacy FC router* is an FC router configured on a Brocade 7500 switch.

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 95](#) on page 592, no devices can be connected to the backbone fabric (Fabric 8) because base switches cannot have F_Ports. [Figure 97](#) 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 above 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 v7.0.0 or later, EX_Ports remain functional and you can continue to perform all FC router operations on the switch.

Brocade recommends that you save your FC-FC routing configuration (using the `configUpload` command) before performing any downgrades.

Before upgrading the Brocade G620 switch from Fabric OS 8.0.0 to 8.0.1, 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

Consider the following:

- If you replace an FR4-18i blade with an 8-Gbps port blade or FX8-24 blade, the EX_Port configuration remains the same for the first 16 ports on the 8-Gbps port blade (and for the first 12 FC ports on the FX8-24 blade). For all other ports on the blade, the EX_Port configuration is cleared. No ports are persistently disabled.
- If you replace an 8-Gbps port blade with an FX8-24 blade, the EX_Port configuration remains the same for the first 12 FC ports on the FX8-24 blade.
- If you replace an 8-Gbps port blade or FX8-24 blade with another 8-Gbps port blade, the EX_Port configuration remains the same.
- If you replace a 32-Gbps port blade with SX6 blade, the EX_Port configuration remains the same for the first 24 FC ports on the SX6 blade.

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 **IsDbShow** command on the edge fabric.

In the **IsDbShow** 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 **IsDbShow** 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
```

FC-FC routing scalability limits

The following table outlines the scalability limits for FC-FC routing.

TABLE 122 FC-FC routing scalability limits

Scalability feature	Supported limits	Supported limits with enhanced Backbones
Maximum number of Fabric OS edge fabrics per metaSAN (with edge fabrics containing up to 1500 WWNs)	48 (DCX8510-4, DCX8510-8) 24 (6510/520)	8 (DCX8510/5300/6510/6520)
Maximum number of Fabric OS edge fabrics per metaSAN (with edge fabrics containing up to 2000 WWNs)	32 (DCX8510-4, DCX8510-8) 12 (6510/6520)	8 (DCX8510/5300/6510/6520)
Maximum number of edge fabrics per chassis (Edge fabrics per chassis may never exceed maximum number of edge fabrics per metaSAN noted in previous rows - in some cases depending on the number of devices in an edge fabric or type of switches, the maximum number of edge fabrics per metaSAN might be lower than that supported per chassis)	30 (DCX8510-4, DCX8510-8) 12 (6510/6520)	8 (DCX8510/5300/6510/6520)
Maximum number of switches per edge fabric (only Fabric OS switches in edge fabric)	26	26
Maximum number of WWNs per edge fabric (only Fabric OS switches in edge fabric)	4000	4000
Maximum number of imported devices from each edge fabric	4000	4000 (1000 imported from Backbones and Directors)
Maximum number of L2 switches participating in backbone fabric	12	40 (additional 8 XDs from 8 edge fabrics).
Maximum number of FC routers per backbone fabric	16 (DCX8510-4, DCX8510-8, and Gen 6 switches)	2 (These 2 FC routers are 2 of the 40 L2 switches supported in Backbones and Directors and are FC routing capable.)
Maximum number of local WWNs per backbone fabric (not including imported devices)	512	2000
Maximum number of LSAN devices per metaSAN (total number of devices imported from all edge fabrics)	15000 (with only DCX8510-8/ DCX8510-4 and Gen 6 switches as routers) 5000 (with any 6510/ 6520 as routers)	15000 (with only DCX8510-8/ DCX8510-4 and Gen 6 switches as routers) 5000 (with 5300 routers)
Maximum number of LSAN zones per metaSAN	7500*	7500*
Maximum number of devices per LSAN zone	64	64
Maximum number of hops between edge switches	19	19
EX_Ports per chassis with Integrated Routing	128 (DCX8510) To the maximum port count (6510/6520)	128 (DCX8510) To the maximum port count (5300)

NOTE

IPFC over FCR is supported only for edge to edge. FC Fast Write is supported only for edge to edge. The maximum number of LSAN zones in a metaSAN can be scaled beyond 7500 by using location embedded LSAN zone feature. For 4K proxy devices through EX_Port, the DCX family of switches are recommended for better performance.

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 Backbone. 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 CR16-4 core blade in slot 3 of a Brocade DCX 8510-4 Backbone. 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. For a core blade, no PID exists in the Address column.

```
switch:FID128:admin> switchshow -slot 3 -qsfp
```

```
switchName: switch73
switchType: 121.3
switchState: Online
switchMode: Native
switchRole: Subordinate
switchDomain: 75
switchId: fffc4b
switchWwn: 10:00:00:05:1e:4f:eb:00
zoning: ON (zoning name)
switchBeacon: OFF
FC Router: OFF
HIF Mode:      ON
Allow XISL Use: OFF
LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Address Mode 0]
```

Index	Slot	Port	QSFP	Address	Media	Speed	State	Proto
256	3	0	0	-----	id	16G	No_SigDet	FC
257	3	1	0	-----	id	16G	No_SigDet	FC
258	3	2	0	-----	id	16G	No_SigDet	FC
259	3	3	0	-----	id	16G	No_SigDet	FC
260	3	4	1	-----	--	16G	No_Module	FC
261	3	5	1	-----	--	16G	No_Module	FC
262	3	6	1	-----	--	16G	No_Module	FC
263	3	7	1	-----	--	16G	No_Module	FC
264	3	8	2	-----	--	16G	No_Module	FC
265	3	9	2	-----	--	16G	No_Module	FC
266	3	10	2	-----	--	16G	No_Module	FC
267	3	11	2	-----	--	16G	No_Module	FC
268	3	12	3	-----	--	16G	No_Module	FC
269	3	13	3	-----	--	16G	No_Module	FC
270	3	14	3	-----	--	16G	No_Module	FC
271	3	15	3	-----	--	16G	No_Module	FC
736	3	16	4	-----	--	16G	No_Module	FC
737	3	17	4	-----	--	16G	No_Module	FC
738	3	18	4	-----	--	16G	No_Module	FC
739	3	19	4	-----	--	16G	No_Module	FC
740	3	20	5	-----	--	16G	No_Module	FC
741	3	21	5	-----	--	16G	No_Module	FC
742	3	22	5	-----	--	16G	No_Module	FC
743	3	23	5	-----	--	16G	No_Module	FC
744	3	24	6	-----	--	16G	No_Module	FC
745	3	25	6	-----	--	16G	No_Module	FC
746	3	26	6	-----	--	16G	No_Module	FC
747	3	27	6	-----	--	16G	No_Module	FC
748	3	28	7	-----	id	16G	Online	FC E-Port 10:00:00:05:1e:39:e4:5a
trunkmaster name (Trunk master)								
749	3	29	7	-----	id	16G	Online	FC E-Port 10:00:00:05:1e:39:e4:5a
trunkmaster name (Trunk master)								
750	3	30	7	-----	id	16G	Online	FC E-Port 10:00:00:05:1e:39:e4:5a
trunkmaster name (Trunk master)								

```

751      3      31      7      -----      id      16G      Online      FC E-Port 10:00:00:05:1e:39:e4:5a
trunkmaster name (Trunk master)

```

This example shows the truncated output of the **switchShow** command for an FC16-32 port blade in slot 1 of a Brocade DCX 8510-8 Backbone. The Address column shows the PID.

```

switch:FID128:admin> switchshow -slot 1

switchName: DCX8510_8
(output truncated)
LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Address Mode 0]
Index  Slot  Port  Address  Media  Speed  State  Proto
=====
0      1      0      500000  --     N16    No_Module  FC
1      1      1      500100  --     N16    No_Module  FC
2      1      2      500200  --     N16    No_Module  FC
(output truncated)

```

This example shows the truncated **switchShow** output for an FC16-64 port blade on the Brocade DCX 8510 Backbone. The assignment of port index numbers to PIDs will vary depending on blade type, platform type, and slot number.

```

DCX:admin> switchshow

Index  Slot  Port  Address  Media  Speed  State
=====
0      1      0      0a0040  --     N16    No_Module
1      1      1      0a0140  --     N16    No_Module
2      1      2      0a0240  --     N16    No_Module
(output truncated)
768    1      48     0a00c0  --     N16    No_Module
769    1      49     0a01c0  --     N16    No_Module
770    1      50     0a02c0  --     N16    No_Module
(output truncated)
784    1      62     0a0ec0  --     N16    No_Module
783    1      63     0a0fc0  --     N16    No_Module
16     2      0      0a1040  --     N16    No_Module
17     2      1      0a1140  --     N16    No_Module
(output truncated)

```

This example shows the truncated **switchShow** output for an FX8-24 application blade on the Brocade DCX 8510-8 Backbone. The assignment of port index numbers to PIDs will vary depending on blade type, platform type, and slot number.

```

switch:FID128:admin> switchshow -slot 10

switchName: my8510-8
(output truncated)
Slot  Blade  Type  ID  Model  Name  Status
-----
10    AP  BLADE  75  FX8-24  ENABLED
Index  Slot  Port  Address  Media  Speed  State  Proto
=====
80     10    0      505000  id     4G     No_Light  FC
81     10    1      505100  --     4G     No_Module  FC
82     10    2      505200  id     4G     Mod_Inv   FC "Speed Mismatch / Incompatible SFP"
83     10    3      505300  --     4G     No_Module  FC
84     10    4      505400  --     4G     No_Module  FC
(output truncated)
95     10    15     505f00  --     --     Offline   VE
208    10    16     50d000  --     --     Offline   VE
209    10    17     50d100  --     4G     Offline   VE
(output truncated)

```

Switch and Blade Sensors

- Brocade switch sensors..... 599
- Brocade blade temperature sensors..... 600
- System temperature monitoring..... 600

Brocade switch sensors

The following table lists the number of temperature sensors, fans, and PSUs for each Brocade switch model. The *swNum / connUnitNum* column is the sum of the values in the previous three columns.

TABLE 123 Brocade switch sensor counts

Platform	Temperature sensors	Fans	Power supply units	swNum / connUnitNum
Brocade 6505	4	2	2	8
Brocade 6510	4	2	2	8
Brocade 6520	4	2	2	8
Brocade 7840 Extension Switch	4	5	2	11
Brocade DCX 8510-4 Director	35+	2	2	39+
Brocade DCX 8510-8 Director	58+	3	4	65+
Brocade G610	6	4 (a minimum of 2 must be working)	1	11
Brocade G620	7	6	2	13
Brocade G630	15	3 (a minimum of 2 must be working)	2	20
Brocade 7810 Extension Switch	11	6	2	19
Brocade X6-4 Director	35+	2	2	39+
Brocade X6-8 Director	58+	3	4 (3 minimum)	65+
Brocade M6505 Blade Server SAN I/O Module	3	0	0	3
Brocade 6542 Blade Server SAN I/O Module	5	2	0	7
Brocade 6543 Blade Server SAN I/O Module	2	0	0	2
Brocade 6545 Blade Server SAN I/O Module	3	0	0	3
Brocade 6546 Blade Server SAN I/O Module	3	0	0	3
Brocade 6547 Blade Server SAN I/O Module	3	0	0	3
Brocade 6548 Blade Server SAN I/O Module	2	0	0	2
Brocade 6549 Blade Server SAN I/O Module	3	0	0	3
Brocade 6558 Blade Server SAN I/O Module	2	0	0	2
Brocade 6559 Blade Server SAN I/O Module	7	0	0	7

Brocade blade temperature sensors

The following table lists the number of temperature sensors for each Brocade blade model.

TABLE 124 Brocade blade temperature sensors

Model	Temperature sensors
FC16-32	5
FC16-48	5
FC16-64	5
FC32-48	5
FC32-64	5
FX8-24	3
SX6	7

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, as 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.

Hexadecimal Conversion

- Hexadecimal overview.....601

Hexadecimal overview

Hexadecimal, also known as hex, is a numeral system with a base of 16, usually written by means of symbols 0-9 and A-F (or a-f). Its primary purpose is to represent the binary code that computers interpret in a format easier for humans to remember. It 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 *Ox* prefix or a *h* suffix. The address 0xFFFFFA is the same address as FFFFFFFA *h*. This type of address with 6 digits representing 3 bytes is called a hex triplet. Fibre Channel uses hexadecimal notation in hex triplets to specify well-known addresses and port IDs.

Example conversion of the hexadecimal triplet 0x610600

Notice the PID (610600 - bolded) 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 10.38.141.143 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 loop, 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 125 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
Decimal	171	172	173	174	175	176	177	178	179	180
Hex	ab	ac	ad	ae	af	b0	b1	b2	b3	b4
Decimal	181	182	183	184	185	186	187	188	189	190
Hex	b5	b6	b7	b8	b9	ba	bb	bc	bd	be
Decimal	191	192	193	194	195	196	197	198	199	200
Hex	bf	c0	c1	c2	c3	c4	c5	c6	c7	c8

TABLE 125 Decimal-to-hexadecimal conversion table (continued)

Decimal	201	202	203	204	205	206	207	208	209	210
Hex	c9	ca	cb	cc	cd	ce	cf	d0	d1	d2
Decimal	211	212	213	214	215	216	217	218	219	220
Hex	d3	d4	d5	d6	d7	d8	d9	da	db	dc
Decimal	221	222	223	224	225	226	227	228	229	230
Hex	dd	de	df	e0	e1	e2	e3	e4	e5	e6
Decimal	231	232	233	234	235	236	237	238	239	240
Hex	e7	e8	e9	ea	eb	ec	ed	ef	ee	f0
Decimal	241	242	243	244	245	246	247	248	249	250
Hex	f1	f2	f3	f4	f5	f6	f7	f8	f9	fa
Decimal	251	252	253	254	255					
Hex	fb	fc	fd	fe	ff					

Revision History

FOS-821-AG101; September 28, 2018

- Updated [Password history policy](#) on page 153.
- Updated [Encryption and Trunking Support on FC32-48 Blades](#) on page 295.

FOS-821-AG100; August 28, 2018

- Added support for [Management Port Ingress Rate Limiting](#) on page 47.
- Removed "Any port on a 48-port blade can support loop devices" from [10-bit addressing \(mode 0\)](#) on page 72 since FL_Ports are not supported on current 48-port blades.
- Revised the list of restrictions for 48-port and 64-port blades under [Port-based addressing \(mode 2\)](#) on page 73.
- Added support for the Brocade 7810 Extension Switch in [Port types](#) on page 75.
- Added support for the [Password history policy](#) on page 153.
- Added trunking support for the FC32-48 blade in [Supported Gen 5 and Gen 6 in-flight encryption and compression features](#) on page 277.
- Added restrictions for configuring encryption and trunking on Brocade G620 Switch ports to [Additional in-flight encryption and compression restrictions](#) on page 284.
- Added a reference in [Additional in-flight encryption and compression restrictions](#) on page 284 to [Encryption and Trunking Support on FC32-48 Blades](#) on page 295 for restrictions on configuring encryption and trunking on FC32-48 blade ports.
- Updated the recommendations for the FC32-48 blade in [Availability considerations for encryption and compression](#) on page 285.
- Added the new section [Encryption and Trunking Support on FC32-48 Blades](#) on page 295.
- Added support for the Brocade 7810 Extension Switch in [Maximum configurable distances for Extended Fabrics](#) on page 310.
- Modified the table in [Virtual Fabrics interaction with other Fabric OS features](#) on page 399 to support ISL R_RDY mode.
- Added support for the Brocade 7810 Extension Switch in [Supported platforms for FC-FC routing](#) on page 536.
- Added details for the Brocade 7810 Extension Switch in [Brocade switch sensors](#) on page 599.