

53-1002920-02
9 September 2013



Fabric OS

Administrator's Guide

Supporting Fabric OS 7.2.0

BROCADE

Copyright © 2013 Brocade Communications Systems, Inc. All Rights Reserved.

ADX, AnyIO, Brocade, Brocade Assurance, the B-wing symbol, DCX, Fabric OS, ICX, MLX, MyBrocade, OpenScript, VCS, VDX, and Vyatta are registered trademarks, and HyperEdge, The Effortless Network, and The On-Demand Data Center are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

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 <http://www.brocade.com/support/oscd>.

Brocade Communications Systems, Incorporated

Corporate and Latin American Headquarters
Brocade Communications Systems, Inc.
130 Holger Way
San Jose, CA 95134
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems China HK, Ltd.
No. 1 Guanghua Road
Chao Yang District
Units 2718 and 2818
Beijing 100020, China
Tel: +8610 6588 8888
Fax: +8610 6588 9999
E-mail: china-info@brocade.com

European Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour B - 4ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 5640
Fax: +41 22 799 5641
E-mail: emea-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems Co., Ltd. (Shenzhen WFOE)
Citic Plaza
No. 233 Tian He Road North
Unit 1308 – 13th Floor
Guangzhou, China
Tel: +8620 3891 2000
Fax: +8620 3891 2111
E-mail: china-info@brocade.com

Document History

Title	Publication number	Summary of changes	Date
<i>Fabric OS Administrator's Guide</i>	53-1002920-01	Added Fabric OS v7.2.0 software features and support for embedded switches: Brocade 5431, M6505, and 6547.	July 2013
<i>Fabric OS Administrator's Guide</i>	53-1002920-02	Corrections and additions for the Fabric OS 7.2.0a release.	September 2013

Contents (High Level)

Section I

Standard Features

Chapter 1	Understanding Fibre Channel Services	45
Chapter 2	Performing Basic Configuration Tasks	57
Chapter 3	Performing Advanced Configuration Tasks	83
Chapter 4	Routing Traffic	115
Chapter 5	Buffer-to-Buffer Credits and Credit Recovery	135
Chapter 6	Managing User Accounts	151
Chapter 7	Configuring Protocols	195
Chapter 8	Configuring Security Policies	231
Chapter 9	Maintaining the Switch Configuration File	277
Chapter 10	Installing and Maintaining Firmware	289
Chapter 11	Managing Virtual Fabrics	309
Chapter 12	Administering Advanced Zoning	337
Chapter 13	Traffic Isolation Zoning	379
Chapter 14	Optimizing Fabric Behavior	413
Chapter 15	Bottleneck Detection	427
Chapter 16	In-flight Encryption and Compression	445
Chapter 17	Diagnostic Port	459
Chapter 18	NPIV	473
Chapter 19	Fabric-Assigned PWWN	479
Chapter 20	Managing Administrative Domains	485

Section II

Licensed Features

Chapter 21	Administering Licensing	515
Chapter 22	Inter-chassis Links	543
Chapter 23	Monitoring Fabric Performance	551
Chapter 24	Managing Trunking Connections	569

Chapter 25	Managing Long-Distance Fabrics587
Chapter 26	Using FC-FC Routing to Connect Fabrics593
Appendix A	Port Indexing641
Appendix B	FIPS Support645
Appendix C	Hexadecimal Conversion657

Contents

About This Document

Supported hardware and software	35
What's new in this document.....	36
Document conventions.....	36
Notice to the reader	38
Additional information.....	38
Getting technical help.....	40
Document feedback	41

Section I

Standard Features

Chapter 1

Understanding Fibre Channel Services

Fibre Channel services overview	45
Management server	46
Platform services	46
Platform services and Virtual Fabrics.....	47
Enabling platform services	47
Disabling platform services.....	47
Management server database.....	47
Displaying the management server ACL.....	48
Adding a member to the ACL.....	48
Deleting a member from the ACL.....	49
Viewing the contents of the management server database	50
Clearing the management server database	51
Topology discovery	51
Displaying topology discovery status	51
Enabling topology discovery	51
Disabling topology discovery.....	52
Device login	53
Principal switch	53
E_Port login process	53
Fabric login process.....	54
Port login process	54
RSCNs.....	54
Duplicate Port World Wide Name.....	55
High availability of daemon processes	55

Chapter 2

Performing Basic Configuration Tasks

Fabric OS overview	57
Fabric OS command line interface.	58
Console sessions using the serial port.	58
Telnet or SSH sessions	59
Getting help on a command	60
Viewing a history of command line entries	61
Password modification	63
Default account passwords.	63
The switch Ethernet interface	64
Virtual Fabrics and the Ethernet interface.	65
Management Ethernet port bonding	65
Displaying the network interface settings	66
Static Ethernet addresses.	67
DHCP activation	69
IPv6 autoconfiguration	71
Date and time settings	72
Setting the date and time	72
Time zone settings.	72
Network time protocol	74
Domain IDs	75
Displaying the domain IDs.	75
Setting the domain ID	76
Switch names	76
Customizing the switch name	77
Chassis names	77
Customizing chassis names	77
Fabric name.	77
Configuring the fabric name	78
High availability considerations for fabric names	78
Upgrade and downgrade considerations for fabric names.	78
Switch activation and deactivation	78
Disabling a switch	79
Enabling a switch.	79
Disabling a chassis	79
Enabling a chassis.	80
Switch and Backbone shutdown	80
Powering off a Brocade switch	80
Powering off a Brocade Backbone	81
Basic connections.	81
Device connection	82
Switch connection	82

Chapter 3

Performing Advanced Configuration Tasks

Port identifiers (PIDs) and PID binding overview	83
Core PID addressing mode	84
Fixed addressing mode	84
10-bit addressing mode	84
256-area addressing mode	85
WWN-based PID assignment	86
Ports	88
Port Types	88
Backbone port blades	88
Setting port names	89
Port identification by slot and port number	89
Port identification by port area ID	90
Port identification by index	90
Configuring a device-switch connection	90
Swapping port area IDs	91
Port activation and deactivation	92
Port decommissioning	92
Setting port modes	93
Setting port speeds	94
Setting all ports on a switch to the same speed	94
Setting port speed for a port octet	95
Blade terminology and compatibility	95
CP blades	97
Core blades	97
Port and application blade compatibility	98
FX8-24 compatibility notes	98
Enabling and disabling blades	98
Enabling blades	99
Disabling blades	99
Blade swapping	99
How blades are swapped	100
Swapping blades	102
Disabling switches	102
Power management	103
Powering off a port blade	103
Powering on a port blade	103
Equipment status	104
Checking switch operation	104
Verifying High Availability features (Backbones only)	104
Verifying fabric connectivity	105
Verifying device connectivity	105
Viewing the switch status policy threshold values	105
Setting the switch status policy threshold values	106
Audit log configuration	107
Verifying host syslog prior to configuring the audit log	109
Configuring an audit log for specific event classes	109

Duplicate PWWN handling during device login	110
Setting 0, First login precedence	110
Setting 1, Second login precedence	110
Setting 2, Mixed precedence	110
Setting the behavior for handling duplicate PWWNs	111
Enabling forward error correction	111
FEC Limitations	112
Using the portCfgFec command	112

Chapter 4

Routing Traffic

Routing overview	115
Paths and route selection	116
FSPF	116
Fibre Channel NAT	117
Inter-switch links	118
Buffer credits	119
Congestion versus over-subscription	119
Virtual channels	119
Gateway links	120
Configuring a link through a gateway	121
Routing policies	122
Displaying the current routing policy	122
Port-based routing	123
Exchange-based routing	123
Device-based routing	123
Dynamic Path Selection	124
AP route policies	124
Route selection	125
Dynamic Load Sharing	125
Frame order delivery	126
Forcing in-order frame delivery across topology changes	127
Restoring out-of-order frame delivery across topology changes	127
Using Frame Viewer to understand why frames are dropped	127
Lossless Dynamic Load Sharing on ports	129
Lossless core	130
Configuring Lossless Dynamic Load Sharing	131
Lossless Dynamic Load Sharing in Virtual Fabrics	131
Frame Redirection	132
Creating a frame redirect zone	132
Deleting a frame redirect zone	133
Viewing frame redirect zones	133

Chapter 5

Buffer-to-Buffer Credits and Credit Recovery

Buffer credit management	135
Buffer-to-buffer flow control	135
Optimal buffer credit allocation	136
Fibre Channel gigabit values reference definition.	137
Buffer credit allocation based on full-size frames.	137
Allocating buffer credits based on average-size frames.	140
Configuring buffers for a single port directly	141
Configuring buffers using frame size	141
Calculating the number of buffers required given the distance, speed, and frame size.	142
Allocating buffer credits for F_Ports	142
Monitoring buffers in a port group	142
Buffer credits switch or blade model	143
Maximum configurable distances for Extended Fabrics	144
Downgrade considerations	145
Configuring credits for a single VC	146
Buffer credit recovery	146
Buffer credit recovery over an E_Port.	147
Buffer credit recovery over an F_Port.	147
Buffer credit recovery over an EX_Port.	148
Enabling and disabling buffer credit recovery	148
Credit loss	149
Back-end credit loss detection and recovery support on Brocade 5300 switches	149
Back-end credit loss detection and recovery support on Brocade 6520 switches	149
Enabling back-end credit loss detection and recovery	150

Chapter 6

Managing User Accounts

User accounts overview	151
Role-Based Access Control	152
Management channel	154
Managing user-defined roles	154
Local database user accounts	155
Default accounts	156
Local account passwords	157
Local user account database distribution.	158
Distributing the local user database	158
Accepting distributed user databases on the local switch	158
Rejecting distributed user databases on the local switch	159
Password policies	159
Password strength policy.	159
Password history policy	160
Password expiration policy	161
Account lockout policy	161

The boot PROM password	163
Setting the boot PROM password for a switch with a recovery string	163
Setting the boot PROM password for a Backbone with a recovery string	164
Setting the boot PROM password for a switch without a recovery string	165
Setting the boot PROM password for a Backbone without a recovery string	166
Remote authentication	167
Remote authentication configuration	167
Setting the switch authentication mode	171
Fabric OS user accounts	171
Fabric OS users on the RADIUS server	172
Setting up a RADIUS server	175
LDAP configuration and Microsoft Active Directory	181
LDAP configuration and OpenLDAP	184
TACACS+ service	189
Remote authentication configuration on the switch	192
Configuring local authentication as backup	194

Chapter 7

Configuring Protocols

Security protocols	195
Secure Copy	196
Setting up SCP for configuration uploads and downloads	197
Secure Shell protocol	197
SSH public key authentication	198
Secure Sockets Layer protocol	200
Browser and Java support	200
SSL configuration overview	201
The browser	204
Root certificates for the Java plugin	205
Simple Network Management Protocol	206
SNMP Manager	206
SNMP Agent	206
Management Information Base (MIB)	207
Basic SNMP operation	207
Understanding MIBs	208
Access to MIB variables	208
SNMP support	209
Traps	209
Loading Brocade MIBs	212
Access Gateway and Brocade MIBs	216
Firmware upgrades and enabled traps	216
Support for Administrative Domains	216
Support for Role-Based Access Control	216
Support for IPv6 addressing	217
Support for Virtual Fabric	217
Configuring SNMP using CLI	218

Telnet protocol	226
Blocking Telnet	227
Unblocking Telnet	228
Listener applications	228
Ports and applications used by switches	229
Port configuration	229

Chapter 8

Configuring Security Policies

ACL policies overview	231
How the ACL policies are stored	231
Policy members	232
ACL policy management	232
Displaying ACL policies	233
Saving changes without activating the policies	233
Activating ACL policy changes	233
Deleting an ACL policy	233
Adding a member to an existing ACL policy	234
Removing a member from an ACL policy	234
Abandoning unsaved ACL policy changes	234
FCS policies	235
FCS policy restrictions	235
Ensuring fabric domains share policies	236
Creating an FCS policy	236
Modifying the order of FCS switches	237
FCS policy distribution	238
Device Connection Control policies	238
DCC policy restrictions	239
Creating a DCC policy	239
Deleting a DCC policy	240
DCC policy behavior with Fabric-Assigned PWWNs	241
SCC Policies	242
Creating an SCC policy	243
Authentication policy for fabric elements	243
E_Port authentication	244
Device authentication policy	246
AUTH policy restrictions	247
Authentication protocols	248
Secret key pairs for DH-CHAP	249
FCAP configuration overview	251
Fabric-wide distribution of the authorization policy	253

IP Filter policy	253
Creating an IP Filter policy	254
Cloning an IP Filter policy	254
Displaying an IP Filter policy	254
Saving an IP Filter policy	255
Activating an IP Filter policy	255
Deleting an IP Filter policy	255
IP Filter policy rules	255
IP Filter policy enforcement	258
Adding a rule to an IP Filter policy	259
Deleting a rule from an IP Filter policy	259
Aborting an IP Filter transaction	259
IP Filter policy distribution	260
Policy database distribution	260
Database distribution settings	261
ACL policy distribution to other switches	262
Fabric-wide enforcement	263
Notes on joining a switch to the fabric	264
Management interface security	266
Configuration examples	267
IPsec protocols	269
Security associations	269
Authentication and encryption algorithms	269
IPsec policies	270
IKE policies	271
Creating the tunnel	272
Example of an end-to-end transport tunnel mode	274

Chapter 9

Maintaining the Switch Configuration File

Configuration settings	277
Configuration file format	278
Configuration file backup	279
Uploading a configuration file in interactive mode	279
Configuration file restoration	280
Restrictions	281
Configuration download without disabling a switch	282
Configurations across a fabric	284
Downloading a configuration file from one switch to another switch of the same model	284
Security considerations	284
Configuration management for Virtual Fabrics	285
Uploading a configuration file from a switch with Virtual Fabrics enabled	285
Restoring a logical switch configuration using configDownload	285
Restrictions	286
Brocade configuration form	287

Chapter 10

Installing and Maintaining Firmware

Firmware download process overview	289
Upgrading and downgrading firmware	291
Considerations for FICON CUP environments	291
HA sync state	291
Preparing for a firmware download	292
Obtaining and decompressing firmware	293
Connected switches.	293
Firmware download on switches	294
Switch firmware download process overview.	294
Firmware download on a Backbone.	296
Backbone firmware download process overview.	296
Firmware download from a USB device.	299
Enabling the USB device	299
Viewing the USB file system	299
Downloading from the USB device using the relative path.	300
Downloading from the USB device using the absolute path.	300
FIPS support	300
Public and private key management	300
The firmwareDownload command	301
Power-on firmware checksum test	302
Testing and restoring firmware on switches	302
Testing a different firmware version on a switch.	302
Testing and restoring firmware on Backbones	304
Testing different firmware versions on Backbones	304
Validating a firmware download.	306

Chapter 11

Managing Virtual Fabrics

Virtual Fabrics overview	309
Logical switch overview	310
Default logical switch.	310
Logical switches and fabric IDs.	311
Port assignment in logical switches	312
Logical switches and connected devices.	313
Management model for logical switches.	314
Logical fabric overview	315
Logical fabric and ISLs.	315
Base switch and extended ISLs	316
Account management and Virtual Fabrics	319
Supported platforms for Virtual Fabrics	320
Supported port configurations in the fixed-port switches.	320
Supported port configurations in Brocade Backbones.	321
Virtual Fabrics interaction with other Fabric OS features	322

Limitations and restrictions of Virtual Fabrics	322
Restrictions on XISLs	323
Restrictions on moving ports	324
Enabling Virtual Fabrics mode	324
Disabling Virtual Fabrics mode	325
Configuring logical switches to use basic configuration values.	326
Creating a logical switch or base switch	326
Executing a command in a different logical switch context	328
Deleting a logical switch	329
Adding and moving ports on a logical switch	329
Displaying logical switch configuration	330
Changing the fabric ID of a logical switch	331
Changing a logical switch to a base switch	331
Setting up IP addresses for a logical switch	333
Removing an IP address for a logical switch	333
Configuring a logical switch to use XISLs	333
Changing the context to a different logical fabric	334
Creating a logical fabric using XISLs	334

Chapter 12

Administering Advanced Zoning

Zone types	337
Zoning overview.	338
Approaches to zoning	339
Zone objects.	340
Zone configurations.	341
Zoning enforcement.	342
Considerations for zoning architecture	342
Best practices for zoning.	343
Broadcast zones	343
Broadcast zones and Admin Domains	344
Broadcast zones and FC-FC routing	345
High availability considerations with broadcast zones	346
Loop devices and broadcast zones	346
Broadcast zones and default zoning mode	346
Zone aliases	346
Creating an alias	347
Adding members to an alias	347
Removing members from an alias	348
Deleting an alias	349
Viewing an alias in the defined configuration	349

Zone creation and maintenance	350
Displaying existing zones	350
Creating a zone	350
Adding devices (members) to a zone	351
Removing devices (members) from a zone	352
Replacing zone members	353
Deleting a zone	355
Viewing a zone in the defined configuration	356
Viewing zone configuration names without case distinction	356
Validating a zone	358
Default zoning mode	360
Setting the default zoning mode	361
Viewing the current default zone access mode	361
Zone database size	362
Zone configurations	362
Creating a zone configuration	363
Adding zones to a zone configuration	363
Removing members from a zone configuration	364
Enabling a zone configuration	364
Disabling a zone configuration	365
Deleting a zone configuration	365
Abandoning zone configuration changes	366
Viewing all zone configuration information	366
Viewing selected zone configuration information	367
Viewing the configuration in the effective zone database	367
Clearing all zone configurations	367
Zone object maintenance	368
Copying a zone object	368
Deleting a zone object	369
Renaming a zone object	370
Zone configuration management	370
Security and zoning	371
Zone merging	371
Fabric segmentation and zoning	373
Zone merging scenarios	373
Concurrent zone transactions	376
Viewing zone database transactions	377

Chapter 13

Traffic Isolation Zoning

Traffic Isolation Zoning overview	379
TI zone failover	380
Additional considerations when disabling failover	381
FSPF routing rules and traffic isolation	383
Enhanced TI zones	384
Illegal configurations with enhanced TI zones	385

Traffic Isolation Zoning over FC routers	386
TI zones within an edge fabric	388
TI zones within a backbone fabric	389
Limitations of TI zones over FC routers	390
Fabric-Level Traffic Isolation in a backbone fabric	390
Fabric-Level TI zones	391
Failover behavior for Fabric-Level TI zones	392
Creating a separate TI zone for each path	392
Creating a single TI zone for all paths	393
General rules for TI zones	394
Traffic Isolation Zone violation handling for trunk ports	395
Supported configurations for Traffic Isolation Zoning	396
Additional configuration rules for enhanced TI zones	396
Trunking with TI zones	397
Limitations and restrictions of Traffic Isolation Zoning	398
Admin Domain considerations for Traffic Isolation Zoning	398
Virtual Fabrics considerations for Traffic Isolation Zoning	399
Traffic Isolation Zoning over FC routers with Virtual Fabrics	401
Creating a TI zone	402
Creating a TI zone in a base fabric	404
Modifying TI zones	405
Changing the state of a TI zone	406
Deleting a TI zone	407
Displaying TI zones	407
Troubleshooting TI zone routing problems	408
Setting up TI zones over FCR (sample procedure)	409

Chapter 14

Optimizing Fabric Behavior

Adaptive Networking overview	413
Ingress Rate Limiting	414
Virtual Fabrics considerations	414
Limiting traffic from a particular device	415
Disabling Ingress Rate Limiting	415
QoS	415
License requirements for QoS	416
CS_CTL-based frame prioritization	416
Supported configurations for CS_CTL-based frame prioritization	417
High availability considerations for CS_CTL-based frame prioritization	417
Enabling CS_CTL-based frame prioritization on ports	417
Disabling CS_CTL-based frame prioritization on ports	418
Using CS_CTL auto mode at the chassis level	418
Considerations for using CS_CTL-based frame prioritization	418

QoS zone-based traffic prioritization	419
QoS zones	419
QoS on E_Ports	421
QoS over FC routers	421
Virtual Fabrics considerations for QoS zone-based traffic prioritization	422
High-availability considerations for QoS zone-based traffic prioritization	422
Supported configurations for QoS zone-based traffic prioritization	423
Limitations and restrictions for QoS zone-based traffic prioritization	424
Setting QoS zone-based traffic prioritization	424
Setting QoS zone-based traffic prioritization over FC routers	426
Disabling QoS zone-based traffic prioritization	426

Chapter 15 Bottleneck Detection

Bottleneck detection overview	427
Types of bottlenecks	428
How bottlenecks are reported	428
Supported configurations for bottleneck detection	429
Limitations of bottleneck detection	429
High availability considerations for bottleneck detection	430
Upgrade and downgrade considerations for bottleneck detection	430
Trunking considerations for bottleneck detection	430
Virtual Fabrics considerations for bottleneck detection	430
Access Gateway considerations for bottleneck detection	430
Enabling bottleneck detection on a switch	431
Displaying bottleneck detection configuration details	431
Setting bottleneck detection alerts	433
Setting both a congestion alert and a latency alert	434
Setting a congestion alert only	434
Setting a latency alert only	435
Changing bottleneck detection parameters	435
Examples of applying and changing bottleneck detection parameters	436
Advanced bottleneck detection settings	439
Excluding a port from bottleneck detection	440
Displaying bottleneck statistics	442
Disabling bottleneck detection on a switch	442

Chapter 16

In-flight Encryption and Compression

In-flight encryption and compression overview	445
Supported ports for in-flight encryption and compression	446
In-flight encryption and compression restrictions	446
How in-flight encryption and compression are enabled	448
Authentication and key generation for encryption and compression	448
Availability considerations for encryption and compression	449
Virtual Fabrics considerations for encryption and compression	449
In-flight compression on long-distance ports	450
Compression ratios for compression-enabled ports	450
Configuring in-flight encryption and compression on an EX_Port	450
Configuring in-flight encryption and compression on an E_Port	451
Viewing the encryption and compression configuration	452
Configuring and enabling authentication for in-flight encryption	453
Enabling in-flight encryption	455
Enabling in-flight compression	456
Disabling in-flight encryption	456
Disabling in-flight compression	457

Chapter 17

Diagnostic Port

Diagnostic Port	459
Supported platforms for D_Port	459
Licensing requirements for D_Port	460
Understanding D_Port	460
Advantages of D_Port	461
D_Port configuration mode and nature of test	461
General limitations and considerations for D_Port	462
Supported topologies	463
Topology 1: ISLs	463
Topology 2: ICLs	463
Topology 3: Access Gateways	464
Topology 4: HBA to switch	465
Using D_Port without HBAs	465
Enabling D_Port	465
Disabling D_Port	466
Using D_Port with HBAs	467
Automatic mode configuration	467
Dynamic mode configuration	468
BCU D_Port commands	468
Limitations and considerations for D_Port with HBAs	468
Controlling testing	469

	Example test scenarios and output	469
	Confirming SFP and link status with an HBA	470
	Starting and stopping D_Port testing	470
Chapter 18	NPIV	
	NPIV overview	473
	Upgrade considerations	474
	Fixed addressing mode	474
	10-bit addressing mode	474
	Configuring NPIV	475
	Enabling and disabling NPIV	476
	Viewing NPIV port configuration information	476
	Viewing virtual PID login information	478
Chapter 19	Fabric-Assigned PWWN	
	Fabric-Assigned PWWN overview	479
	User- and auto-assigned FA-PWWN behavior	480
	Configuring an FA-PWWN for an HBA connected to an Access Gateway	481
	Configuring an FA-PWWN for an HBA connected to an edge switch	482
	Supported switches and configurations for FA-PWWN	483
	Configuration upload and download considerations for FA-PWWN	483
	Security considerations for FA-PWWN	483
	Restrictions of FA-PWWN	484
	Access Gateway N_Port failover with FA-PWWN	484
Chapter 20	Managing Administrative Domains	
	Administrative Domains overview	485
	Admin Domain features	487
	Requirements for Admin Domains	487
	Admin Domain access levels	487
	User-defined Admin Domains	488
	System-defined Admin Domains	488
	Home Admin Domains and login	490
	Admin Domain member types	491
	Admin Domains and switch WWNs	492
	Admin Domain compatibility, availability, and merging	494

Admin Domain management for physical fabric administrators	494
Setting the default zoning mode for Admin Domains	495
Creating an Admin Domain	495
User assignments to Admin Domains	496
Removing an Admin Domain from a user account	498
Activating an Admin Domain	498
Deactivating an Admin Domain	499
Adding members to an existing Admin Domain	499
Removing members from an Admin Domain	500
Renaming an Admin Domain	500
Deleting an Admin Domain	501
Deleting all user-defined Admin Domains	502
Deleting all user-defined Admin Domains non-disruptively	502
Validating an Admin Domain member list	506
SAN management with Admin Domains	506
CLI commands in an AD context	507
Executing a command in a different AD context	507
Displaying an Admin Domain configuration	508
Switching to a different Admin Domain context	508
Admin Domain interactions with other Fabric OS features	509
Admin Domains, zones, and zone databases	510
Admin Domains and LSAN zones	511
Configuration upload and download in an AD context	512

Section II Licensed Features

Chapter 21

Administering Licensing

Licensing overview	515
Brocade 7800 Upgrade license	523
ICL licensing	523
ICL 1st POD license	523
ICL 2nd POD license	524
ICL 8-link license	524
ICL 16-link license	524
Enterprise ICL license	524
8G licensing	525
Slot-based licensing	526
Upgrade and downgrade considerations	526
Assigning a license to a slot	526
Removing a license from a slot	527
10G licensing	527
Enabling 10 Gbps operation on an FC port	528
Enabling the 10-GbE ports on an FX8-24 blade	529

Temporary licenses	530
Restrictions on upgrading temporary slot-based licenses . . .	531
Date change restriction	531
Configupload and download considerations	531
Expired licenses	531
Universal temporary licenses	532
Extending a universal temporary license	532
Universal temporary license shelf life.	532
Viewing installed licenses.	532
Activating a license	533
Adding a licensed feature	533
Removing a licensed feature	534
Ports on Demand	535
Displaying installed licenses	536
Activating Ports on Demand	537
Dynamic Ports on Demand	537
Displaying the port license assignments	538
Enabling Dynamic Ports on Demand	538
Disabling Dynamic Ports on Demand	539
Reserving a port license	540
Releasing a port from a POD set.	540

Chapter 22

Inter-chassis Links

Inter-chassis links	543
License requirements for ICLs	544
ICLs for the Brocade DCX 8510 Backbone family.	544
ICL trunking on the Brocade DCX 8510-8 and DCX 8510-4 .	545
ICLs for the Brocade DCX Backbone family.	546
ICL trunking on the Brocade DCX and DCX-4S.	547
Virtual Fabrics considerations for ICLs	547
Supported topologies for ICL connections	547
Mesh topology	547
Core-edge topology	549

Chapter 23

Monitoring Fabric Performance

Advanced Performance Monitoring overview	551
Types of monitors.	551
Restrictions for installing monitors.	552
Virtual Fabrics considerations for Advanced Performance Monitoring	552
Access Gateway considerations for Advanced Performance Monitoring	553

End-to-end performance monitoring	553
Maximum number of EE monitors	553
Supported port configurations for EE monitors	554
Adding EE monitors	554
Setting a mask for an EE monitor	555
Deleting EE monitors	556
Displaying EE monitor counters	557
Clearing EE monitor counters	557
Frame monitoring	558
License requirements for frame monitoring	558
Creating frame types to be monitored	559
Creating a frame monitor	559
Deleting frame types	560
Adding frame monitors to a port	560
Removing frame monitors from a port	560
Saving a frame monitor configuration	560
Displaying frame monitors	561
Clearing frame monitor counters	562
Top Talker monitors	562
Top Talker monitors and FC-FC routing	563
Limitations of Top Talker monitors	565
Adding a Top Talker monitor to a port (port mode)	565
Adding Top Talker monitors on all switches in the fabric (fabric mode)	565
Displaying the top <i>n</i> bandwidth-using flows on a port (port mode)	566
Displaying top talking flows for a given domain ID (fabric mode)	566
Deleting a Top Talker monitor on a port (port mode)	567
Deleting all fabric mode Top Talker monitors	567
Trunk monitoring	567
Trunk monitoring considerations	567
Saving and restoring monitor configurations	567
Performance data collection	568

Chapter 24

Managing Trunking Connections

Trunking overview	569
Types of trunking	570
Masterless trunking	570
License requirements for trunking	571
Port groups for trunking	571
Supported platforms for trunking	571
Supported configurations for trunking	571
High Availability support for trunking	572
Requirements for trunk groups	572
Recommendations for trunk groups	572
Configuring trunk groups	573

Enabling trunking	574
Disabling trunking	574
Displaying trunking information	574
Trunk Area and Admin Domains	576
Example of Trunk Area assignment on port domain, index . . .	576
ISL trunking over long-distance fabrics	576
EX_Port trunking	577
Masterless EX_Port trunking	577
Supported configurations and platforms for EX_Port trunking	578
Configuring EX_Port trunking	578
Displaying EX_Port trunking information	578
F_Port trunking	579
F_Port trunking for Access Gateway	579
F_Port trunking for Brocade adapters	581
F_Port trunking considerations	582
F_Port trunking in Virtual Fabrics	584
Displaying F_Port trunking information	585
Disabling F_Port trunking	585
Enabling the DCC policy on a trunk area	586

Chapter 25

Managing Long-Distance Fabrics

Long-distance fabrics overview	587
Extended Fabrics device limitations	588
Long-distance link modes	588
Configuring an extended ISL	589
Enabling long distance when connecting to TDM devices . . .	590
Forward error correction on long-distance links	591
Enabling FEC on a long-distance link	591
Disabling FEC on a long-distance link	591

Chapter 26

Using FC-FC Routing to Connect Fabrics

FC-FC routing overview	593
License requirements for FC-FC routing	594
Supported platforms for FC-FC routing	594
Supported configurations for FC-FC routing	595
Network OS connectivity limitations	595
Fibre Channel routing concepts	596
Proxy devices	599
FC-FC routing topologies	600
Phantom domains	601
FC router authentication	603
Setting up FC-FC routing	603
Verifying the setup for FC-FC routing	604

Backbone fabric IDs	605
Assigning backbone fabric IDs	606
FCIP tunnel configuration	606
Inter-fabric link configuration	607
Configuring an IFL for both edge and backbone connections	607
Configuring EX_Ports on an ICL	611
FC router port cost configuration	613
Port cost considerations	614
Setting router port cost for an EX_Port.	614
Shortest IFL cost configuration	615
Configuring shortest IFL cost	617
EX_Port frame trunking configuration	619
LSAN zone configuration	620
Use of Admin Domains with LSAN zones and FC-FC routing	620
Zone definition and naming	620
LSAN zones and fabric-to-fabric communications.	621
Controlling device communication with the LSAN.	621
Configuring backbone fabrics for interconnectivity.	623
Setting the maximum LSAN count	624
HA and downgrade considerations for LSAN zones	624
LSAN zone policies using LSAN tagging	624
LSAN zone binding.	628
Proxy PID configuration.	633
Fabric parameter considerations.	633
Inter-fabric broadcast frames.	634
Displaying the current broadcast configuration.	634
Enabling broadcast frame forwarding	634
Disabling broadcast frame forwarding	634
Resource monitoring.	634
FC-FC routing and Virtual Fabrics.	636
Logical switch configuration for FC routing	637
Backbone-to-edge routing with Virtual Fabrics	638
Upgrade and downgrade considerations for FC-FC routing	639
How replacing port blades affects EX_Port configuration.	639
Displaying the range of output ports connected to xlate domains	639

Appendix A Port Indexing

Appendix B FIPS Support

FIPS overview	645
Zeroization functions	645
Power-on self-tests.	647
Conditional tests	647

FIPS mode configuration	647
LDAP in FIPS mode	648
LDAP certificates for FIPS mode	650
Preparing a switch for FIPS.	651
Overview of steps	652
Enabling FIPS mode	652
Zeroizing for FIPS	655
Displaying FIPS configuration	655

Appendix C

Hexadecimal Conversion

Example conversion of the hexadecimal triplet 0x616000 ..	657
Decimal-to-hexadecimal conversion table	658

Index

Figures

Figure 1	Well-known addresses	45
Figure 2	Identifying the blades	100
Figure 3	Blade swap with Virtual Fabrics during the swap	101
Figure 4	Blade swap with Virtual Fabrics after the swap	102
Figure 5	Principal ISLs	116
Figure 6	New switch added to existing fabric	118
Figure 7	Virtual channels on a QoS-enabled ISL	120
Figure 8	Gateway link merging SANs	121
Figure 9	Single host and target	132
Figure 10	Windows 2000 VSA configuration	173
Figure 11	Example of a brocade.dct file	180
Figure 12	Example of the dictiona.dcm file	180
Figure 13	SNMP structure	207
Figure 14	SNMP query	207
Figure 15	SNMP trap	207
Figure 16	Brocade MIB tree location	208
Figure 17	DH-CHAP authentication	244
Figure 18	Protected endpoints configuration	267
Figure 19	Gateway tunnel configuration	268
Figure 20	Endpoint-to-gateway tunnel configuration	268
Figure 21	Switch before and after enabling Virtual Fabrics	310
Figure 22	Switch before and after creating logical switches	311
Figure 23	Fabric IDs assigned to logical switches	312
Figure 24	Assigning ports to logical switches	312
Figure 25	Logical switches connected to devices and non-Virtual Fabrics switch	314
Figure 26	Logical switches in a single chassis belong to separate fabrics	314
Figure 27	Logical switches connected to other logical switches through physical ISLs	316
Figure 28	Logical switches connected to form logical fabrics	316
Figure 29	Base switches connected by an XISL	317
Figure 30	Logical ISLs connecting logical switches	318
Figure 31	Logical fabric using ISLs and XISLs	318
Figure 32	Example of logical fabrics in multiple chassis and XISLs	335
Figure 33	Zoning example	339
Figure 34	Broadcast zones and Admin Domains	345
Figure 35	Traffic Isolation zone creating a dedicated path through the fabric	380
Figure 36	Fabric incorrectly configured for TI zone with failover disabled	382

Figure 37	Dedicated path is the only shortest path	383
Figure 38	Dedicated path is not the shortest path	384
Figure 39	Enhanced TI zones	384
Figure 40	Illegal ETIZ configuration: two paths from one port to two devices on the same remote domain	385
Figure 41	Illegal ETIZ configuration: two paths from one port	386
Figure 42	Traffic Isolation Zoning over FCR	387
Figure 43	TI zone in an edge fabric	388
Figure 44	TI zone in a backbone fabric	389
Figure 45	Fabric-level traffic isolation	391
Figure 46	TI zone misconfiguration	395
Figure 47	Dedicated path with Virtual Fabrics	399
Figure 48	Creating a TI zone in a logical fabric	400
Figure 49	Creating a TI zone in a base fabric	400
Figure 50	Example configuration for TI zones over FC routers in logical fabrics	401
Figure 51	Logical representation of TI zones over FC routers in logical fabrics	401
Figure 52	TI over FCR example	409
Figure 53	QoS traffic prioritization	420
Figure 54	QoS with E_Ports enabled	421
Figure 55	Traffic prioritization in a logical fabric	423
Figure 56	Affected seconds for bottleneck detection	433
Figure 57	Encryption and compression on 16 Gbps ISLs	446
Figure 58	Example of a basic D_Port connection between switches	460
Figure 59	ISLs connecting multiple switches and chassis	463
Figure 60	ICLs connecting chassis blades	463
Figure 61	Single Access Gateway to switch	464
Figure 62	Multiple Access Gateways cascaded to switch	464
Figure 63	Access Gateway to HBA	464
Figure 64	HBA to switch	465
Figure 65	Fabric-assigned port World Wide Name provisioning scenarios	480
Figure 66	Fabric with two Admin Domains	486
Figure 67	Filtered fabric views when using Admin Domains	486
Figure 68	Fabric with ADO and AD255	490
Figure 69	Fabric showing switch and device WWNs	493
Figure 70	Filtered fabric views showing converted switch WWNs	493
Figure 71	ADO and two user-defined Admin Domains, AD1 and AD2	504
Figure 72	ADO with three zones	504
Figure 73	Minimum configuration for 64 Gbps ICLs	545
Figure 74	DCX-4S allowed ICL connections	546
Figure 75	ICL triangular topology with Brocade DCX 8510-8 chassis	548
Figure 76	Full nine-mesh topology	549
Figure 77	64 Gbps ICL core-edge topology	550

Figure 78	Setting end-to-end monitors on a port	554
Figure 79	Mask positions for end-to-end monitors	556
Figure 80	Fabric mode Top Talker monitors on FC router do not monitor any flows	564
Figure 81	Fabric mode Top Talker monitors on FC router monitor flows over the E_Port	564
Figure 82	Port group configuration for the Brocade 5100	571
Figure 83	Switch in Access Gateway mode without F_Port masterless trunking	580
Figure 84	Switch in Access Gateway mode with F_Port masterless trunking	580
Figure 85	A metaSAN with inter-fabric links	596
Figure 86	A metaSAN with edge-to-edge and backbone fabrics and LSAN zones	597
Figure 87	Edge SANs connected through a backbone fabric	599
Figure 88	MetaSAN with imported devices	600
Figure 89	Sample topology (physical topology)	601
Figure 90	EX_Port phantom switch topology	602
Figure 91	Shortest IFL solution	617
Figure 92	Example of setting up Speed LSAN tag	626
Figure 93	LSAN zone binding	629
Figure 94	EX_Ports in a base switch	637
Figure 95	Logical representation of EX_Ports in a base switch	638
Figure 96	Backbone-to-edge routing across base switch using FC router in legacy mode	639

Tables

Table 1	Daemons that are automatically restarted.	55
Table 2	Terminal port parameters	59
Table 3	Help topic contents.	61
Table 4	fabricShow fields	76
Table 5	Ports affected when you enable or disable a switch in VF or non-VF mode.	79
Table 6	Core and CP blade terminology and platform support.	95
Table 7	Port blade terminology, numbering, and platform support	96
Table 8	Blade compatibility within Brocade Backbone families.	98
Table 9	Duplicate PWWN behavior: First login takes precedence over second login	110
Table 10	Duplicate PWWN behavior: Second login overrides first login	110
Table 11	Duplicate PWWN behavior: Port type determines which login takes precedence	111
Table 12	Combinations of routing policy and IOD with Lossless DLS enabled	130
Table 13	Fibre Channel gigabit values	137
Table 14	Fibre Channel data frames	138
Table 15	Total FC ports, ports per port group, and unreserved buffer credits per port group	143
Table 16	Configurable distances for Extended Fabrics.	144
Table 17	Default Fabric OS roles	152
Table 18	Permission types.	153
Table 19	Maximum number of simultaneous sessions.	154
Table 20	Default local user accounts	156
Table 21	LDAP options.	169
Table 22	Authentication configuration options	169
Table 23	Syntax for VSA-based account roles	171
Table 24	Entries in dictionary.brocade file	173
Table 25	Brocade custom TACACS+ attributes	190
Table 26	Secure protocol support.	195
Table 27	Items needed to deploy secure protocols.	196
Table 28	Main security scenarios	196
Table 29	SSL certificate files.	203
Table 30	Brocade SNMP MIB dependencies	214
Table 31	Access Gateway MIB support.	216
Table 32	Security level options	218
Table 33	Blocked listener applications.	228
Table 34	Access defaults.	229

Table 35	Port information	229
Table 36	Valid methods for specifying policy members	232
Table 37	FCS policy states.	235
Table 38	FCS switch operations	236
Table 39	Distribution policy states	238
Table 40	DCC policy states	239
Table 41	DCC policy behavior with FA-PWWN when created using lockdown support . .	241
Table 42	DCC policy behavior when created manually with PWWN	242
Table 43	SCC policy states	242
Table 44	FCAP certificate files.	251
Table 45	Supported services.	256
Table 46	Implicit IP Filter rules	258
Table 47	Default IP policy rules.	258
Table 48	Interaction between fabric-wide consistency policy and distribution settings. .	261
Table 49	Supported policy databases	261
Table 50	Fabric-wide consistency policy settings	263
Table 51	Merging fabrics with matching fabric-wide consistency policies.	265
Table 52	Examples of strict fabric merges	266
Table 53	Fabric merges with tolerant and absent combinations	266
Table 54	Algorithms and associated authentication policies	270
Table 55	CLI commands to display or modify switch configuration information	281
Table 56	Brocade configuration and connection form	287
Table 57	Backbone HA sync states.	292
Table 58	Commands used for validating a firmware download	307
Table 59	Blade and port types supported on logical switches	321
Table 60	Virtual Fabrics interaction with Fabric OS features	322
Table 61	Maximum number of logical switches per chassis.	323
Table 62	Approaches to fabric-based zoning.	339
Table 63	Considerations for zoning architecture.	342
Table 64	Zone merging scenarios: Defined and effective configurations	373
Table 65	Zone merging scenarios: Different content	374
Table 66	Zone merging scenarios: Different names	375
Table 67	Zone merging scenarios: TI zones	375
Table 68	Zone merging scenarios: Default access mode	376
Table 69	Zone merging scenarios: Mixed Fabric OS versions.	376
Table 70	Traffic behavior when failover is enabled or disabled in TI zones	381
Table 71	Comparison between CS_CTL-based and QoS zone-based prioritization.	416
Table 72	Mapping of CS_CTL values to QoS priority for frame prioritization in CS_CTL default mode.	417
Table 73	Mapping of CS_CTL values to QoS priority for frame prioritization in CS_CTL auto mode	417

Table 74	Number of ports supported for in-flight encryption and compression at various port speeds	447
Table 75	Supported platforms for D_Port	460
Table 76	D_Port configuration mode and nature of test.	462
Table 77	Limitation on number of D_Ports for simultaneous tests	469
Table 78	Number of supported NPIV devices	474
Table 79	AD user types	488
Table 80	Ports and devices in CLI output.	507
Table 81	Admin Domain interaction with Fabric OS features.	509
Table 82	Configuration upload and download scenarios in an AD context	512
Table 83	Available Brocade licenses	516
Table 84	License requirements and location name by feature	519
Table 85	Base to Upgrade license comparison	523
Table 86	List of available user ports when implementing PODs	535
Table 87	Number of logical switches that support performance monitors	552
Table 88	Maximum number of frame monitors and offsets per port.	558
Table 89	Predefined values at offset 0.	559
Table 90	Trunking over long distance for the Brocade Backbones and blades	576
Table 91	F_Port masterless trunking considerations	582
Table 92	PWWN format for F_Port and N_Port trunk ports.	584
Table 93	Fabric-wide settings	589
Table 94	LSAN information stored in FC routers, with and without LSAN zone binding .	630
Table 95	Zeroization behavior.	645
Table 96	FIPS mode restrictions.	647
Table 97	FIPS and non-FIPS modes of operation	648
Table 98	Active Directory keys to modify	650
Table 99	Decimal-to-hexadecimal conversion table	658

About This Document

In this chapter

- [Supported hardware and software](#) 35
- [What's new in this document](#) 36
- [Document conventions](#) 36
- [Notice to the reader](#) 38
- [Additional information](#) 38
- [Getting technical help](#) 40
- [Document feedback](#) 41

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some switches but not to others, this guide identifies exactly which switches are supported and which are not.

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for Fabric OS v7.2.0, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by this release of Fabric OS:

- Fixed-port switches:
 - Brocade 300 switch
 - Brocade 5100 switch
 - Brocade 5300 switch
 - Brocade 5410 embedded switch
 - Brocade 5424 embedded switch
 - Brocade 5430 embedded switch
 - Brocade 5431 embedded switch
 - Brocade 5450 embedded switch
 - Brocade 5460 embedded switch
 - Brocade 5470 embedded switch
 - Brocade 5480 embedded switch
 - Brocade M6505 embedded switch
 - Brocade 6505 switch
 - Brocade 6510 switch

- Brocade 6520 switch
- Brocade 6547 embedded switch
- Brocade 7800 extension switch
- Brocade VA-40FC
- Brocade Encryption Switch
- Brocade DCX Backbone family:
 - Brocade DCX
 - Brocade DCX-4S
- Brocade DCX 8510 Backbone family:
 - Brocade DCX 8510-4
 - Brocade DCX 8510-8

What's new in this document

Information that was modified:

- Renamed and moved the section about the two Ethernet ports on the CP blade to [“Management Ethernet port bonding”](#) on page 65.
- Moved the section [“Enabling forward error correction”](#) from the Routing chapter to [Chapter 3, “Performing Advanced Configuration Tasks”](#).
- In [Chapter 17, “Diagnostic Port,”](#) updated [Table 76](#) and added some HA considerations and some considerations for D_Port with HBAs.
- Updated the [“Setting up FC-FC routing”](#) section and the [“Configuring EX_Ports on an ICL”](#) section in [Chapter 26, “Using FC-FC Routing to Connect Fabrics”](#).
- In [Appendix B, “FIPS Support,”](#) updated [Table 96](#) on page 647 with entries for Authentication and FC-FC routing.

Document conventions

This section describes text formatting conventions and important notice formats used in this document.

Text formatting

The narrative-text formatting conventions that are used are as follows:

bold text	Identifies command names
	Identifies the names of user-manipulated GUI elements
	Identifies keywords and operands
	Identifies text to enter at the GUI or CLI

<i>italic text</i>	Provides emphasis Identifies variables Identifies paths and Internet addresses Identifies document titles
code text	Identifies CLI output Identifies command syntax examples

For readability, command names in the narrative portions of this guide are presented in mixed lettercase: for example, switchShow. In actual examples, command lettercase is often all lowercase. Otherwise, this manual specifically notes those cases in which a command is case sensitive.

Command syntax conventions

Command syntax in this manual follows these conventions:

command	Commands are printed in bold.
--option, option	Command options are printed in bold.
-argument, arg	Arguments.
[]	Optional element.
variable	Variables are printed in italics. In the help pages, values are <u>underlined</u> or enclosed in angled brackets < >.
...	Repeat the previous element, for example “member[:member...]”
value	Fixed values following arguments are printed in plain font. For example, --show WWN
	Boolean. Elements are exclusive. Example: --show -mode egress ingress

Notes, cautions, and warnings

The following notices and statements are used in this manual. They are listed below in order of increasing severity of potential hazards.

NOTE

A note provides a tip, guidance or advice, emphasizes important information, or provides a reference to related information.

ATTENTION

An Attention statement indicates potential damage to hardware or data.



CAUTION

A *Caution* statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.



DANGER

A *Danger* statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Key terms

For definitions specific to Brocade and Fibre Channel, see the *Brocade Glossary*.

For definitions of SAN-specific terms, visit the Storage Networking Industry Association online dictionary at:

<http://www.snia.org/education/dictionary>

Notice to the reader

This document may contain references to the trademarks of the following corporations. These trademarks are the properties of their respective companies and corporations.

These references are made for informational purposes only.

Corporation	Referenced Trademarks and Products
Microsoft Corporation	Windows, Windows NT, Internet Explorer
Mozilla Corporation	Mozilla, Firefox
Netscape Communications Corporation	Netscape
Red Hat, Inc.	Red Hat, Red Hat Network, Maximum RPM, Linux Undercover
Sun Microsystems, Inc.	Sun, Solaris

Additional information

This section lists additional Brocade and industry-specific documentation that you might find helpful.

Brocade resources

To get up-to-the-minute information, go to <http://my.brocade.com> and register at no cost for a user ID and password.

For practical discussions about SAN design, implementation, and maintenance, you can obtain *Building SANs with Brocade Fabric Switches* through:

<http://www.amazon.com>

For additional Brocade documentation, visit the Brocade SAN Info Center and click the Resource Library location:

<http://www.brocade.com>

Release notes are available on the My Brocade website and are also bundled with the Fabric OS firmware.

Other industry resources

For additional resource information, visit the Technical Committee T11 website. This website provides interface standards for high-performance and mass storage applications for Fibre Channel, storage management, and other applications:

<http://www.t11.org>

For information about the Fibre Channel industry, visit the Fibre Channel Industry Association website:

<http://www.fibrechannel.org>

Getting technical help


Contact your switch support supplier for hardware, firmware, and software support, including product repairs and part ordering. To expedite your call, have the following information available:

1. General Information

- Switch model
- Switch operating system version
- Error numbers and messages received
- **supportSave** command output
- Detailed description of the problem, including the switch or fabric behavior immediately following the problem, and specific questions
- Description of any troubleshooting steps already performed and the results
- Serial console and Telnet session logs
- syslog message logs

2. switch serial number

The switch serial number and corresponding bar code are provided on the serial number label, as illustrated below.:



The serial number label is located as follows:

- *Brocade 300, 5100, 5300, 6505, M6505, 6510, 6520, 6547, 7800, VA-40FC, and Brocade Encryption Switch*—On the switch ID pull-out tab located inside the chassis on the port side on the left
- *Brocade 5410, 5424, 5430, 5431, 5450, 5460, 5470, 5480*—Serial number label attached to the module
- *Brocade 6510*—On the pull-out tab on the front of the switch
- *Brocade DCX and DCX 8510-8*—On the bottom right on the port side of the chassis
- *Brocade DCX-4S and DCX 8510-4*—On the bottom right on the port side of the chassis, directly above the cable management comb

3. World Wide Name (WWN)

Use the **wwn** command to display the switch WWN.

If you cannot use the **wwn** command because the switch is inoperable, you can get the WWN from the same place as the serial number, except for the Brocade DCX enterprise class platform. For the Brocade DCX enterprise class platform, access the numbers on the WWN cards by removing the Brocade logo plate at the top of the nonport side of the chassis.

For the Brocade 5424 embedded switch: Provide the license ID. Use the **licenseIdShow** command to display the WWN.

Document feedback

Quality is our first concern at Brocade and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback to:

documentation@brocade.com

Provide the title and version number of the document and as much detail as possible about your comment, including the topic heading and page number and your suggestions for improvement.

Standard Features

This section describes standard Fabric OS features, and includes the following chapters:

- [Chapter 1, “Understanding Fibre Channel Services”](#)
- [Chapter 2, “Performing Basic Configuration Tasks”](#)
- [Chapter 3, “Performing Advanced Configuration Tasks”](#)
- [Chapter 4, “Routing Traffic”](#)
- [Chapter 5, “Buffer-to-Buffer Credits and Credit Recovery”](#)
- [Chapter 6, “Managing User Accounts”](#)
- [Chapter 7, “Configuring Protocols”](#)
- [Chapter 8, “Configuring Security Policies”](#)
- [Chapter 9, “Maintaining the Switch Configuration File”](#)
- [Chapter 10, “Installing and Maintaining Firmware”](#)
- [Chapter 11, “Managing Virtual Fabrics”](#)
- [Chapter 12, “Administering Advanced Zoning”](#)
- [Chapter 13, “Traffic Isolation Zoning”](#)
- [Chapter 14, “Optimizing Fabric Behavior”](#)
- [Chapter 15, “Bottleneck Detection”](#)
- [Chapter 16, “In-flight Encryption and Compression”](#)
- [Chapter 17, “Diagnostic Port”](#)
- [Chapter 18, “NPIV”](#)
- [Chapter 19, “Fabric-Assigned PWWN”](#)
- [Chapter 20, “Managing Administrative Domains”](#)

Understanding Fibre Channel Services

In this chapter

- Fibre Channel services overview 45
- Management server 46
- Platform services 46
- Management server database 47
- Topology discovery 51
- Device login 53
- High availability of daemon processes 55

Fibre Channel services overview

Fibre Channel services define service functions that reside at well-known addresses. A *well-known address* is a reserved three-byte address for each service. Services are provided to either nodes or management applications in the fabric.

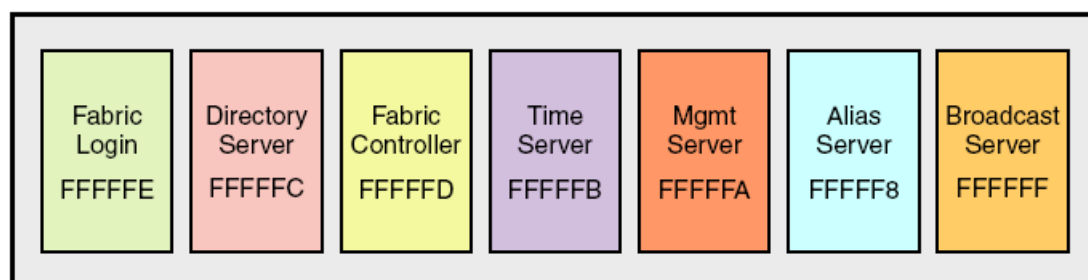


Figure 1 Well-known addresses

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 (0x000000) represents the domain ID, the second node (0x000000) the port area number of the port where the node is attached, and the third node (0x000000) 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.

Refer to [Chapter 8, “Configuring Security Policies,”](#) for additional information on FCS policies.

1 Management server

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

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 commands `msplMgmtActivate` and `msplMgmtDeactivate` are allowed only in ADO and AD255.

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.

The execution of the **msplMgmtActivate** command is subject to Admin Domain restrictions that may be in place.

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.

1 Management server database

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.

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 so you can verify that the WWN you entered was added to the ACL.
6. After verifying that the WWN was 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.

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, that member no longer has access to the management server.

NOTE

If you delete the last member of the ACL, leaving the ACL list is empty, then 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.

1 Management server database

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.

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.

Example of 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 command **msPIClearDB** 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 commands **mstdEnable** and **mstdDisable** are allowed only in ADO and AD255.

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.

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....
```

1 Topology discovery

```
*MS Topology Discovery enabled locally.  
*MS Topology Discovery Enable Operation Complete!!
```

Disabling topology discovery

Use the following procedure to disable 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 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 the Topology Discovery feature.

NOTE

Topology discovery is disabled by default.

ATTENTION

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

If Admin Domains are enabled, you must be in the ADO or AD255 context. Refer to [Chapter 20, "Managing Administrative Domains,"](#) for additional information.

Example of disabling 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 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 ones listed below with the Fabric Controller to access the fabric. Once 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 multiple switches, and one inter-switch link (ISL) exists between any two switches, 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. Once 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 functions to discover the type of device connected and establish the port type and negotiate port speed. See [“Port Types”](#) on page 88 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. This enables private devices that do not explicitly register with the Name Server (NS) to be entered in the NS and receive full fabric access.

A fabric-capable device registers its information with the name server during a FLOGI. These devices typically register information with the name server 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 a switch’s name server, use the **nsShow** or **nsAllShow** command. For more information about these commands, refer to the *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, refer to [“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 Name Server database and a RASlog is generated.

Device recovery

To recover devices that have been removed from the Name Server database due to duplicate PWWNs, the devices must re-login to the fabric. 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 mode.

High availability of daemon processes

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 AUDIT event message are logged.
2. The daemon is automatically started again.
3. If the restart is successful, then 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.

Schedule 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 Daemons that are automatically restarted

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.

1 High availability of daemon processes

Table 1 Daemons that are automatically restarted (Continued)

Daemon	Description
traced	Trace daemon provides trace entry date and time translation to 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.

Performing Basic Configuration Tasks

In this chapter

• Fabric OS overview	57
• Fabric OS command line interface	58
• Password modification	63
• The switch Ethernet interface	64
• Date and time settings	72
• Domain IDs	75
• Switch names	76
• Chassis names	77
• Fabric name	77
• Switch activation and deactivation	78
• Switch and Backbone shutdown	80
• Basic connections	81

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 *Web Tools Administrator's Guide*.

- **Brocade Network Advisor**

For additional information, refer to the *Brocade Network Advisor User Manual* for the version you have.

- **A third-party application using the API**

For third-party application procedures, refer to the third-party API documentation.

Because of 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 [Chapter 3, "Performing Advanced Configuration Tasks"](#).

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc., documenting all possible configurations and scenarios is beyond the scope of this document. In some cases, earlier releases are highlighted to present considerations for interoperating 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 about the commands used in the procedures, refer to the *Fabric OS Command Reference*.

Fabric OS 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 the section [“Role-Based Access Control”](#) on page 152.

Note the following about the command display in this guide:

- Commands are shown and can be entered either in all lower case or using Java-style capitalization. This means that while `bannershow` and `bannerShow` will both work, `BANNERSHOW` and `BannerShow` will not.
- When command examples in this guide show user input enclosed in quotation marks, the quotation marks are required. 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 and DCX 8510 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:

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 on a PC, TERM, TIP, or Kermit in a UNIX environment), and configure the application as follows:

- In a Windows environment enter the following parameters:

TABLE 2 Terminal port parameters

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

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

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

If **tyb** is already in use, use **tya** instead and enter the following string at the prompt:

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

Telnet or SSH sessions

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 58.

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 69.

Rules for Telnet connections

The following rules must be observed when making Telnet connections to your switch:

- Never change the IP address of the switch while two Telnet sessions are active; if you do, your next attempt to log in fails. 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, see the *Web Tools Administrator's 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 kill 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 two. For more details on session limits, refer to [Chapter 6, “Managing User Accounts”](#).

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, refer to [“Default account passwords”](#) on page 63.

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 the switchAdmin roles.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **help [|more]** 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, you can enter **help command |more**, where *command* is the name of the command for which you need specific information.

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
ficonHelp	FICON help information
fwHelp	Fabric Watch help information
iscsiHelp	iSCSI help information
licenseHelp	License help information
perfHelp	Performance Monitoring help information
routeHelp	Routing help information
trackChangesHelp	Track Changes help information
zoneHelp	Zoning help information

Viewing a history of command line entries

The CLI command history log file saves the last 512 commands from all users on a FIFO basis, and this log is persistent across reboots and firmware downloads. This command is also supported for standby CPs.

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

- Timestamp
- Username
- IP address of the telnet session
- Options
- Arguments

Use the following procedure to view the CLI command log:

1. Connect to the switch and log in.
2. Enter the **cliHistory** command with the desired argument (see below for arguments). Entering no specific argument displays only the command line history of the currently logged-in user.

cliHistory

Entering the **cliHistory** command with no arguments displays the command line history for the currently logged-in user only (even for the root user).

Example cliHistory command output from 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]
```

Example cliHistory command output from 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
```

cliHistory --show

Using the "--show" argument displays the same results as entering "cliHistory" without any arguments.

cliHistory --showuser <username>

Using the "--showuser <username>" argument displays the command line history of the named user. This argument is available only to Root, Admin, Factory and Securityadmin RBAC roles.

Example cliHistory command output showing username

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

Using the "--showall" argument displays the command line history for all users. With this option, admin/factory/securityadmin users can see the root user command history. This argument is available only to Root, Admin, Factory and Securityadmin RBAC roles.

Example cliHistory showing history of all users

```
switch:admin> clihistory --showall
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]
Thu Sep 27 05:25:48 2012    root, 10.70.12.101, ipaddrshow
```

cliHistory --help

Using the "-- help" argument displays a list of the available command arguments.

```
swd77: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 --help:
    Displays the command usage
```

Notes:

- SSH login CLI logs are not recorded in the command line history.
- The CLI command log will be collected as part of any “supportsave” operation. The command log record of such an operation will be the equivalent of running “cliHistory --showall”.
- For CLI commands that require a password (Examples: firmwaredownload, configupload/download, supportsave, and so on), only the command (no arguments) is stored (see below for an illustration).

```
sw0:FID128:root> 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
```

```
sw0:FID128:root> clihistory
Date & Time                Message
Wed May 23 03:39:37 2012    root, console, firmwaredownload
```

Password modification

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 after each subsequent login until all the default passwords have been changed.

NOTE

The default account passwords can be changed from their original values only when prompted immediately following the login; the passwords cannot be changed using the **passwd** command later in the session. If you skip the prompt, and then later decide to change the passwords, log out and then back in.

The default accounts on the switch are admin, user, root, and factory. Use the “admin” account to log in to the switch for the first time and to perform the basic configuration tasks. The password for all of these accounts is “password”.

There is only one set of default accounts for the entire chassis. The root and factory default accounts are reserved for development and manufacturing. The user account is primarily used for system monitoring. For more information on default accounts, refer to “[Default accounts](#)” on page 156.

Default account passwords

The change default account passwords prompt is a string that begins with the message “Please change your passwords now”. User-defined passwords can have from 8 through 40 characters. They must begin with an alphabetic character and can include numeric characters, the period (.), and the underscore (_). They are case-sensitive, and they are not displayed when you enter them on the command line.

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 and factory accounts are 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.

Changing the default account passwords 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**.

Example output of changing passwords

```
login: admin
Password:
Please change your passwords now.
Use Control-C to exit or press 'Enter' key to proceed.
for user - root
Changing password for root
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 Backbones

On Brocade Backbones, you must set IP addresses for the following components:

- Both Control Processors (CPO 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 67)
- Activating DHCP (refer to [“DHCP activation”](#) on page 69)

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 58.

Virtual Fabrics and the Ethernet interface

On the Brocade DCX and DCX-4S, 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. On the Brocade DCX and DCX-4S Backbones, 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 on how to set up these IPFC interfaces to your Virtual Fabric, refer to [Chapter 11, “Managing Virtual Fabrics”](#).

Management Ethernet port bonding

The two external Ethernet ports of a CP8 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 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 CP8 blade will use the logical network interface named “bond0” instead of “eth0”.

NOTE

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

The CP8 blade contains multiple Ethernet devices (including eth0 and eth3), which map to the two Ethernet ports on the front of the CP8 blade. Other Ethernet devices on the blade are reserved for use by the operating system.

The CP8 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 even if eth0 comes back up. 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 an HA failover routine.
- Power down the switch and then power it back up again.

ATTENTION

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

Supported devices

Management Ethernet port bonding is available on a CP8 blade when it is installed on a Brocade DCX, Brocade DCX-4S, Brocade DCX 8510-8, or Brocade DCX 8510-4.

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 port modes”](#) on page 93 for instructions on setting port modes, and [“Setting port speeds”](#) on page 94 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

If an IP address has not been assigned to the network interface (Ethernet), you must connect to the Fabric OS CLI using a console session on the serial port. For more information, see [“Console sessions using the serial port”](#) on page 58. Otherwise, connect using SSH.

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

```
ipAddrShow
```

Example output 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: ecpl
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.

Static Ethernet addresses

Use static Ethernet network interface addresses on Brocade DCX and DCX-4S Backbones, and 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 69.

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 addresses 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.

Example of setting an IPv4 address

```
switch:admin> ipaddrset
Ethernet IP Address [10.1.2.3]:
Ethernet Subnetmask [255.255.255.0]:
Fibre Channel IP Address [220.220.220.2]:
Fibre Channel Subnetmask [255.255.0.0]:
Gateway IP Address [10.1.2.1]:
DHCP [OFF]: off
```

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 on setting up an IP address for a Virtual Fabric, refer to [Chapter 11, “Managing Virtual Fabrics”](#).

3. Enter the network information in dotted-decimal notation for the Ethernet IPv4 address or in semicolon-separated notation for IPv6.
4. Enter the **Ethernet Subnetmask** at the prompt.
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 Address** at the prompt.
7. Disable DHCP by entering **off**.

Setting the static addresses 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
Ethernet IP Address [192.168.166.148]:
Ethernet Subnetmask [255.255.255.0]:
Committing configuration...Done.
```

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

DHCP activation

Some Brocade switches have DHCP enabled by default. Fabric OS support for DHCP functionality is only provided for Brocade fixed-port switches. These are listed in the [Preface](#).

NOTE

The Brocade DCX and Brocade DCX-4S Backbones do not support DHCP.

The Fabric OS DHCP client supports the following parameters:

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

The DHCP client uses a DHCP vendor-class identifier that allows DHCP servers to determine that the discover/request packet are 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 5300 is “BROCADESWBD64.”

NOTE

The client conforms to the latest IETF Draft Standard RFCs for IPv4, IPv6, and DHCP. DHCP can obtain stateful IPv6 addresses.

Enabling DHCP for IPv4

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

NOTE

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

Enabling DHCP after the Ethernet information has been configured releases the current Ethernet network interface settings. These include the Ethernet IP address, Ethernet subnet mask, and gateway IP address. The Fibre Channel IP address and subnet mask are static and are not affected by DHCP; for instructions on setting the FC IP address, see “[Static Ethernet addresses](#)” on page 67.

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 complete the following steps.

3. If already set up, you can skip the Ethernet IP address, Ethernet subnet mask, Fibre Channel IP address, and Fibre Channel subnet mask prompts by pressing **Enter**. Otherwise, enter the network information in dotted-decimal notation for the IPv4 address.
4. Enable DHCP by entering **on**.
5. You can confirm that the change has been made using the **ipAddrShow** command.

2 The switch Ethernet interface

Example of enabling DHCP for IPv4 interactively:

```
switch:admin> ipaddrset
Ethernet IP Address [10.1.2.3]:
Ethernet Subnetmask [255.255.255.0]:
Fibre Channel IP Address [220.220.220.2]:
Fibre Channel Subnetmask [255.255.0.0]:
Gateway IP Address [10.1.2.1]:
DHCP [Off]:on
```

Example of enabling 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

When you disable DHCP, enter the static Ethernet IP address and subnet mask of the switch and default gateway address. Otherwise, the Ethernet settings may conflict with other addresses assigned by the DHCP server on the network.

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 complete the following steps.

3. Enter the network information using IPv4 dotted-decimal notation.

NOTE

If a static Ethernet address is not available when you disable DHCP, enter **0.0.0.0** at the Ethernet IP address prompt.

4. You can skip the Fibre Channel prompts by pressing **Enter**.
5. When you are prompted for DHCP[On], disable it by entering **off**.
6. You can confirm that the change has been made using the **ipAddrShow** command.

Example of disabling DHCP for IPv4 interactively:

```
switch:admin> ipaddrset
Ethernet IP Address [10.1.2.3]:
Ethernet Subnetmask [255.255.255.0]:
Gateway IP Address [10.1.2.1]:
DHCP [On]:off
```

Example of disabling 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
```

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* allows an IPv6 host to obtain a unique address using the IEEE 802 MAC address; *stateful* 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.

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 152.

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 "mmdgHHMMyy"
```

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).

Example of showing and setting the 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 *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 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 (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. Enter the **tsTimeZone** command.
 - Use **tsTimeZone** with no parameters to display the current time zone setting.
 - Use **--interactive** to list all of the time zones supported by the firmware.
 - Use **timeZone_fmt** to set the time zone by Country/City or by time zone ID, such as Pacific Standard Time (PST).

Example of displaying and changing 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.

2 Date and time settings

You are prompted to select a general location.

Please identify a location so that time zone rules can be set correctly.

3. Enter the appropriate **number** or press **Ctrl-D** to quit.
4. Select a country location at the prompt.
5. Enter the appropriate number at the prompt to specify the time zone region of **Ctrl-D** to quit.

Network time protocol

You can synchronize the local time of the principal and primary FCS switch to a maximum of eight external Network Time Protocol (NTP) servers. To keep the time in your SAN current, it is recommended that the principal or primary FCS switch has its time synchronized with at least one external NTP server. The other switches in the fabric automatically take their time from the principal or primary FCS switch, as described in [“Synchronizing the local time with an external source.”](#)

All switches in the fabric maintain the current clock server value in nonvolatile memory. By default, this value is the local clock server (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.

If Virtual Fabrics is enabled, all the switches in the fabric must have the same NTP clock server configured. This includes any Fabric OS v6.2.0 or earlier switches in the fabric. This ensures that time does not go out of sync in the logical fabric. It is not recommended to have LOCL in the server list.

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.

NOTE

If Virtual Fabrics is enabled, multiple logical switches can share a single chassis. Therefore, the NTP server list must be the same across all fabrics.

Synchronizing the local time with an external source

The **tsClockServer** command accepts multiple server addresses in IPv4, IPv6, or Domain Name System (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.

1. 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.

Example of setting the NTP server

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

Example of displaying the NTP server

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

Example of setting 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.

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.

The default domain ID for Brocade switches is 1.

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.

Example 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. In this output, the principal switch appears in blue boldface.

```
switch:admin> fabricshow
Switch ID      Worldwide Name      Enet IP Addr      FC IP Addr      Name
```

2 Switch names

```
-----  
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  10.3.220.5  0.0.0.0  "ras005"  
                                     fec0:60:69bc:63:205:1eff:fe34:1bd  
6:  fffc06  10:00:00:05:1e:34:02:3e  10.3.220.6  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

Table 4 displays the **fabricShow** fields.

TABLE 4 fabricShow fields

Field	Description
Switch ID	The switch domain_ID and embedded port D_ID. The numbers are broken down as follows: 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 through 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.

Switch names

Switches can be identified by IP address, domain ID, World Wide Name (WWN), or by customized switch names that are unique and meaningful.

The following considerations apply to switch naming:

- Switch names can be from 1 through 30 characters long.
- All switch names must begin with a letter, and can contain letters, numbers, or the underscore character.
- Switch names must be unique across logical switches.
- Changing the switch name causes a domain address format RSCN to be issued and may be disruptive to the fabric.

Customizing the switch name

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

```
switch:admin> switchname newname
```

The prompt does not change to the new switch name until AFTER you re-login.

3. Record the new switch name for future reference.

```
switch:FID128:# 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>
```

Chassis names

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 names

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.
- All switches in a logical fabric must be running Fabric OS v7.2.0. Switches running earlier versions of the firmware can co-exist 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 as shown here:

```
switch:user> 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, as shown here:

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

To set a fabric name that includes bash special meta-characters or spaces, use the command **fabricName** as shown in the following example:

```
switch:user> 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. [Table 5](#) describes which ports are affected for each type of enable or disable operation.

TABLE 5 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. See [“Disabling a chassis”](#) on page 79.

Enabling a switch

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

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

```
switch:admin> switchenable
```

All Fibre Channel ports that passed 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 is enabled, only the ports allocated to the logical switch are enabled. To enable all of the ports, you must enable the entire chassis. See [“Enabling a 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.

2 Switch and Backbone shutdown

```
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 that passed Power On Self Test (POST) are enabled. If any of the logical switches have inter-switch links (ISLs) to a fabric, it joins the fabric.

Switch and Backbone shutdown

To avoid corrupting your file system, you must perform graceful shutdowns of Brocade switches and Backbones.

Warm reboot (also known as *graceful shutdown*) refers to shutting down the switch or platform by way of the following instructions. *Cold boot* (also known as a *hard boot*) refers to shutting down the switch or platform by suddenly shutting down power and powering on again.

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 Backbone

Use the following procedure to power off a Brocade Backbone device:

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 [Chapter 3, “Performing Advanced Configuration Tasks”](#).
- For information on configuring the routing of connections, refer to [Chapter 4, “Routing Traffic”](#).
- For information on configuring extended inter-switch connections, refer to [Chapter 25, “Managing Long-Distance Fabrics”](#).

Device connection

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

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

Switch connection

See the hardware reference manual of your specific switch for ISL connection and cable management information. The standard or default ISL mode is LO. ISL mode LO 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
- 1 km at 16 Gbps

For more information on extended ISL modes, which enable long distance inter-switch links, refer to [Chapter 25, “Managing Long-Distance Fabrics”](#).

Performing Advanced Configuration Tasks

In this chapter

- Port identifiers (PIDs) and PID binding overview 83
- Ports 88
- Blade terminology and compatibility 95
- Enabling and disabling blades 98
- Blade swapping 99
- Disabling switches 102
- Power management 103
- Equipment status 104
- Audit log configuration 107
- Duplicate PWWN handling during device login 110
- Enabling forward error correction 111

Port identifiers (PIDs) 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 may 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.

For more information on reading and converting hexadecimal, refer to [Appendix C, "Hexadecimal Conversion"](#).

Fixed addressing mode

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. With fixed addressing mode enabled, 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.

10-bit addressing mode

The 10-bit addressing mode is the default mode for all the logical switches created in the Brocade Backbones. This addressing scheme is flexible to support a large number of F_Ports. In the regular 10-bit addressing mode, the **portAddress --auto** command supports addresses from 0x00 to 0x8F.

NOTE

The default switch in the Brocade Backbones uses 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 blade can support loop devices.
- 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

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 in the 256-area addressing mode: zero-based and port-based.

- Zero-based mode 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 mode is a bit more complex:

- Port-based mode is *not* supported on the default switch.
- 48-port cards are supported in port-based addressing mode (mode 2) on both Brocade DCX-4S and 8510-4 devices. However, the upper 16 ports of a 64-port card are *not* supported. The Brocade DCX does not support port-based addressing mode (mode 2) on the FC8-48 blade, but does support zero-based addressing (mode 1).
- The Brocade DCX-4S supports port-based addressing (mode 2) on the FC8-48 blade.
- The Brocade 8510-4 supports port-based addressing (mode 2) on the FC16-48 blade.
- The Brocade 8510-8 does not support port-based addressing (mode 2) on the FC16-48 blade, but does support zero-based addressing (mode 1).

ATTENTION

The Brocade DCX and 8510-8 Backbones have safeguards that disable all 49 port cards if FICON Management Server (FMS) is enabled.

Refer to the *FICON Administrator's Guide* for more details if needed.

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.

ATTENTION

When WWN-based PID assignment is enabled, the area assignment is dynamic and does not guarantee any order in the presence of static WWN-area binding or when the devices are moved around.

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 and DCX 8510 Backbone families. This feature is not supported on application blades such as the FS8-18, FX8-24, and the FCOE10-24. 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 [Chapter 18, "NPIV"](#).

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:
 - **wwnAddress -show** displays the assigned WWN-PID bindings.
 - **wwnAddress -findPID wwn** displays the PID assigned to the specified device 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 may be part of a Brocade device:

- **D_Port** — A diagnostic port lets an administrator isolate the inter-switch link (ISL) to diagnose link level faults. This port runs only specific diagnostics tests and does not carry any fabric traffic. Refer to [Chapter 17, “Diagnostic Port,”](#) for more information on this port type.
- **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 other E_Ports. However, the router terminates EX_Ports rather than allowing different fabrics to merge as would happen on a switch with regular E_Ports. An EX_Port cannot be connected to another EX_Port.
- **F_Port** — A fabric port that is assigned to fabric-capable devices, such as SAN storage devices.
- **G_Port** — A generic port that acts as a transition port for non-loop fabric-capable devices.
- **L_Port or FL_Port** — A loop or fabric loop port that connects loop devices. L_Ports are associated with private loop devices and FL_Ports are associated with public loop devices.
- **M_Port** — A mirror port that is configured to duplicate (mirror) the traffic passing between a specified source port and destination port. This is only supported for pairs of F_Ports. Refer to the *Fabric OS Troubleshooting and Diagnostics Guide* for more information on port mirroring.
- **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.

Backbone port blades

Because Backbones 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 and DCX 8510 Backbone 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 Backbones.
- Port blades are used for host, storage, and interswitch connections.
- Application (AP) blades are used for Fibre Channel Application Services and Routing Services, FCIP, Converged Enhanced Ethernet, and encryption support.

NOTE

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

The Brocade DCX and DCX 8510-8 each have 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-4S and DCX 8510-4 each have 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.

When you have port blades with different port counts in the same Backbone (for example, 16-port blades and 32-port blades, or 16-port blades and 18-port blades with 16 FC ports and 2 GbE ports, or 16-port and 48-port blades), the area IDs no longer match the port numbers.

[Table 7](#) on page 96 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.

```
ecp:admin> portname 1/0 trunk1
```

Port identification by slot and port number

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

To select a specific port in the Backbones, 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.

Example of enabling port 4 on a blade in slot 2

```
ecp: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 (FC8-32, FC8-32E, 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 (FC8-48, FC8-48E, FC16-48), 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 blade (FC8-64), 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. On 48-port blades, port swapping is supported only on ports 0 through 15.

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 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, the ACL DDC, and Admin Domain) 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-4S and DCX 8510-4 Backbones.

Configuring a device-switch connection

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

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 those ports’ area ID and index values.

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.

Use the following procedure to swap the port area IDs of two physical switch ports. To swap port area IDs, the port swap feature must be enabled, and both switch ports must be disabled. The swapped area IDs for the two ports remain persistent across reboots, power cycles, and failovers.

NOTE

On the Brocade DCX and DCX 8510 Backbone families, you can swap only ports 0 through 15 on the FC8-48 port blades. You cannot swap ports 16 through 47.

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

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

4. Enter the **portSwap** command.

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

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

Port activation and deactivation

By default, all licensed ports are enabled. You can disable and re-enable them as necessary. Ports that you activate with the Ports on Demand license must be enabled explicitly, as described in “[Ports on Demand](#)” on page 535.



CAUTION

The fabric will be reconfigured if the port you are enabling or disabling is connected to another switch.

The switch with a port that has been disabled will be segmented from the fabric and all traffic flowing between it and the fabric will be lost.

Enabling a port

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 **portEnable** *[slot/]port* command.
 - To enable a port that is persistently disabled, enter the **portCfgPersistentEnable** *[slot/]port* command.

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

Disabling a port

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** *[slot/]port* command.
 - To disable a port that is persistently enabled, enter the **portCfgPersistentDisable** *[slot/]port* command.

Port decommissioning

Port decommissioning is an Fabric OS 7.0.0 and later 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 links configured for encryption or compression.
- 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.

Setting port modes

Ports can be set to use one of three link operating modes: full duplex, half 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 DCX, DCX-4S, DCX 8510-4, or DCX 8510-8, the supported interfaces are eth0 and eth3. On all other platforms, only eth0 is supported.

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

Enter **y** or **yes** at the prompts to confirm the active link operating mode values. Enter **n** or **no** to deactivate that mode.

Example of setting the port mode to full autonegotiate

The following example sets the mode for eth3 to autonegotiate, and permits both full and half duplex modes to be selected at both 10 and 100 Mbps. Note that the caution shown in this example is not displayed when the command is entered using the serial console port.

```
switch:admin> ifmodeset eth3
Exercise care when using this command. 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 you only use this command from the serial console port.
Are you sure you really want to do this? (yes, y, no, n): [no] y
Proceed with caution.
Auto-negotiate (yes, y, no, n): [no] y
Advertise 100 Mbps / Full Duplex (yes, y, no, n): [yes] y
Advertise 100 Mbps / Half Duplex (yes, y, no, n): [yes] y
Advertise 10 Mbps / Full Duplex (yes, y, no, n): [yes] y
Advertise 10 Mbps / Half Duplex (yes, y, no, n): [yes] y
Committing configuration...done.
```

Example of setting the port mode to 10 Mbps half duplex operation

The following example forces the link for the eth0 interface from autonegotiation to 10 Mbps half-duplex operation:

```
switch:admin> ifmodeset eth0
Auto-negotiate (yes, y, no, n): [yes] n
Force 100 Mbps / Full Duplex (yes, y, no, n): [no] n
Force 100 Mbps / Half Duplex (yes, y, no, n): [no] n
Force 10 Mbps / Full Duplex (yes, y, no, n): [no] n
Force 10 Mbps / Half Duplex (yes, y, no, n): [no] y
Committing configuration...done.
```

Setting port speeds

Use the following procedure to set port speeds.

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

The following example sets the speed for port 3 on slot 2 to 4 Gbps:

```
switch:admin> portcfgspeed 2/3 4
```

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

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

Setting all ports on a switch to the same speed

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

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

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

You can use the `portCfgOctetSpeedCombo` command to configure the speed for a port octet. 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.

Use the following procedure to set the port speed for a port octet.

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

The following example configures the ports in the first octet for combination 3 (support autonegotiated or fixed port speeds of 16 Gbps and 10 Gbps):

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

NOTE

For information on how encryption and compression can affect port speed, refer to [“Port speed on encryption- or compression-enabled ports”](#) on page 447.

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 6](#) includes core and CP blade terminology and descriptions. [Table 7](#) on page 96 includes port blade terminology and descriptions.

TABLE 6 Core and CP blade terminology and platform support

Blade	Blade ID (slotshow)	Supported on:		Definition
		DCX family	DCX 8510 family	
CP8	50	Yes	Yes	Brocade DCX and DCX 8510 Backbone family control processor blade. This CP supports all blades used in the DCX and DCX 8510 Backbone families.
CORE8	52	Yes DCX only	No	A 16-port blade that provides 8 Gbps connectivity between port blades in the Brocade DCX chassis.
CR4S-8	46	Yes DCX-4S only	No	A 16-port blade that provides 8 Gbps connectivity between port blades in the Brocade DCX-4S chassis.
CR16-8	98	No	Yes DCX 8510-8 only	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	No	Yes DCX 8510-4 only	A core blade that has 8x4 QSFPs per blade. It can be connected to another CR16-4 or a CR16-8 core blade.

3 Blade terminology and compatibility

TABLE 7 Port blade terminology, numbering, and platform support

Blade	Blade ID (slotshow)	Supported on:		Ports	Definition
		DCX family	DCX 8510 family		
FC8-16 ¹	21	Yes	No	16	8-Gbps port blade supporting 1, 2, 4, and 8 Gbps port speeds. Ports are numbered from 0 through 15 from bottom to top.
FC8-32 ¹	55	Yes	No	32	8-Gbps port blade supporting 1, 2, 4, and 8 Gbps port speeds. 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.
FC8-32E	125	No	Yes	32	8-Gbps port blade supporting 2, 4, and 8 Gbps port speeds. 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.
FC8-48 ¹	51	Yes	No	48	8-Gbps port blade supporting 1, 2, 4, and 8 Gbps port speeds. 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.
FC8-48E	126	No	Yes	48	8-Gbps port blade supporting 2, 4, and 8 Gbps port speeds. 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.
FC8-64	77	Yes	Yes	64	8-Gbps port blade supporting 2, 4, and 8 Gbps port speeds. The Brocade DCX and Brocade DCX 8510 Backbone families support loop devices on 64-port blades in a Virtual Fabrics-enabled environment. The loop devices can only be attached to ports on a 64-port blade that is not a part of the default logical switch. Ports are numbered from 0 through 31 from bottom to top on the left set of ports and 32 through 63 from bottom to top on the right set of ports.
FC16-32	97	No	Yes	32	A 32-port, 16-Gbps port blade supporting 2, 4, 8, 10, and 16 Gbps port speeds. NOTE: 10 Gbps speed for FC16-xx blades requires the 10G license. 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.
FC16-48	96	No	Yes	48	A 48-port, 16-Gbps port blade supporting 2, 4, 8, 10, and 16 Gbps port speeds. NOTE: 10 Gbps speed for FC16-xx blades requires the 10G license. 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.
FS8-18	68	Yes	Yes	16 FC 2 GbE	Brocade Encryption blade that provides high performance 32-port auto-sensing 8-Gbps Fibre Channel connectivity with data cryptographic (encryption and decryption) and data compression capabilities. Ports are numbered from 0 through 15 from bottom to top. GbE ports are numbered ge0 through ge1 from top to bottom. Going from top to bottom, the 2 GbE ports appear on the top of the blade followed by the 16 FC ports.

TABLE 7 Port blade terminology, numbering, and platform support (Continued)

Blade	Blade ID (slotshow)	Supported on:		Ports	Definition
		DCX family	DCX 8510 family		
FCOE10-24	74	Yes	No	24 10-GbE DCB ports	An application blade that provides Converged Enhanced Ethernet to bridge a Fibre Channel and Ethernet SAN. Ports are numbered from 0 through 11 from bottom to top on the left set of ports and 12 through 23 from bottom to top on the right set of ports.
FX8-24	75	Yes	Yes	12 FC 10 1-GbE 2 10-GbE	Extension blade with 8-Gbps Fibre Channel, FCIP, and 10-GbE technology. Port numbering on this blade is as follows. On the left side of the blade going from bottom to top: <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 On the right side of the blade going from bottom to top: <ul style="list-style-type: none"> • Six FC ports numbered from 6 through 11 • Six 1-GbE ports numbered from ge4 through ge9

1. The Brocade DCX and DCX-4S support loop devices on this blade in a Virtual Fabrics-enabled environment.

CP blades

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

The CP blades in the Brocade DCX and DCX 8510 Backbone families are hot-swappable. The CP8 blades are fully interchangeable among Brocade DCX, DCX-4S, DCX 8510-4, and DCX 8510-8 Backbones.

Brocade recommends that each CP (primary and secondary partition) should maintain the same firmware version.

For more information on maintaining firmware in your Backbone, refer to [Chapter 10, “Installing and Maintaining Firmware”](#).

Core blades

Core blades provide intra-chassis switching and inter-chassis link (ICL) connectivity between DCX/DCX-4S platforms and between DCX 8510 platforms.

- Brocade DCX supports two CORE8 core blades.
- Brocade DCX-4S supports two CR4S-8 core blades.
- Brocade DCX 8510-8 supports two CR16-8 core blades.
- Brocade DCX 8510-4 supports two CR16-4 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 become faulty.

Port and application blade compatibility

Table 7 on page 96 identifies which port and application blades are supported for each Brocade Backbone.

NOTE

During power up of a Brocade DCX or DCX-4S Backbone, if an FCOE10-24 is detected first before any other AP blade, all other AP and FC8-64 blades are faulted. If a non-FCOE10-24 blade is detected first, then any subsequently-detected FCOE10-24 blades are faulted. Blades are powered up starting with slot 1.

The maximum number of intelligent blades supported on a Brocade DCX or DCX 8510-8 is eight.

The maximum number of intelligent blades supported on a Brocade DCX-4S or DCX 8510-4 is four.

Table 8 lists the maximum supported limits of each blade for a specific Fabric OS release. Software functions are not supported across application blades.

TABLE 8 Blade compatibility within Brocade Backbone families

Intelligent blade	Fabric OS v6.3.0		Fabric OS v6.4.0		Fabric OS v7.0.0			
	DCX	DCX-4S	DCX	DCX-4S	DCX	DCX-4S	DCX 8510-8	DCX 8510-4
FS8-18	4	4	4	4	4	4	4	4
FCOE10-24 ¹	2	2	2	2	4	4	0	0
FX8-24 ²	2	4	4	4	4	4	4	4

1. Not compatible with other application blades or with the FC8-64 in the same chassis.
2. The hardware limit is enforced by software.

FX8-24 compatibility notes

Follow these guidelines when using an FX8-24 in the Brocade DCX and DCX-4S Backbones:

- Brocade 7500 GbE ports cannot be connected to either the FX8-24 or Brocade 7800 GbE ports. 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.
- The FX8-24 and FS8-18 blades cannot coexist with the FCOE10-24 blade.

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
```

FC8-48, FC8-48E, FC8-64, and FC16-48 port blade enabling exceptions

Because the area IDs are shared with different port IDs, the FC8-48, FC8-48E, FC8-64, and FC16-48 blades support only F_Ports and E_Ports. They do not support FL_Ports.

Port swapping on an FC8-48, FC8-48E, FC8-64, and FC16-48 is supported only on ports 0 through 15. For the FC8-32, FC8-32E, and FC16-32 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 on ports 16 through 31 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 to ports 16 through 31.

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.

3 Blade swapping

- 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 FC8-32 blade with an FC8-48 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. [Figure 2](#) 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 so on).
- Port count. Both blades must support the same number of front ports (for example, 16 ports to 16 ports, 32 ports to 32 ports, 48 ports to 48 ports, and so on).
- Availability. The ports on the destination blade must be available for the swap operation and not attached to any other devices.

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, [Figure 3](#) 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 [Figure 4](#), 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.

3 Disabling switches




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 diagnostics. This ensures that diagnostic activity does not interfere with normal fabric traffic.

Use the following procedure to disable a switch.

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

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

All blades are powered on by default when the switch chassis is powered on. Blades cannot be powered off when POST or AP initialization is in progress.

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.

NOTE

In the Backbones, the core blades and CP blades cannot be powered off from the CLI. You must manually power off the blades by lowering the slider or removing power from the chassis. If there is no CP up and running, then physical removal or powering off the chassis is required.

Powering off a port blade

Use the following procedure to power off a port blade.

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 you want to power off.

```
ecp:admin> slotpoweroff 3  
Slot 3 is being powered off
```

Powering on a port blade

Use the following procedure to power on a port blade.

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 you want to power on.

```
ecp:admin> slotpoweron 3  
Powering on slot 3
```

Equipment status

You can check the status of switch operation, High Availability features, and fabric connectivity.

Checking switch operation

Use the following procedure to check switch operation.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **switchShow** command. This command displays a switch summary and a port summary.
3. Check that the switch and ports are online.
4. Use the **switchStatusShow** command to further check the status of the switch.

Verifying High Availability features (Backbones 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 Backbone.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **chassisShow** command to verify the model of the DCX and obtain a listing of all field-replaceable units (FRUs).
3. 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.
4. 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.
5. 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.
6. 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 DCX chassis

```
DCX:FID128:admin> slotshow -m
```

Slot	Blade Type	ID	Model Name	Status
1	SW BLADE	55	FC8-32	ENABLED
2	SW BLADE	51	FC8-48	ENABLED
3	SW BLADE	39	FC8-16	ENABLED
4	SW BLADE	51	FC8-48	ENABLED
5	CORE BLADE	52	CORE8	ENABLED
6	CP BLADE	50	CP8	ENABLED
7	CP BLADE	50	CP8	ENABLED
8	CORE BLADE	52	CORE8	ENABLED
9	SW BLADE	37	FC8-16	ENABLED
10	AP BLADE	43	FS8-18	ENABLED
11	SW BLADE	55	FC8-32	ENABLED
12	AP BLADE	24	FS8-18	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 75.

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.

Viewing the switch status policy threshold values

For switches running Fabric Watch, you can view the switch status policy threshold values using the **switchStatusPolicyShow** command

The policy parameter determines the number of failed or inoperable units for each contributor that triggers a status change in the switch. Each parameter can be adjusted so that a specific threshold must be reached before that parameter changes the overall status of a switch to MARGINAL or DOWN. For example, if the FaultyPorts DOWN parameter is set to 3, the status of the switch will change if three ports fail. Only one policy parameter needs to pass the MARGINAL or DOWN threshold to change the overall status of the switch.

For more information about setting policy parameters, refer to the *Fabric Watch Administrator's Guide*.

3 Equipment status

If the switch is running Fabric Watch, you can use the following procedure to view the switch status policy threshold values. If the switch is running MAPS, refer to the *Monitoring and Alerting Policy Suite Administrator's Guide*.

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

Whenever there is a switch change, an error message is logged and an SNMP connUnitStatusChange trap is sent.

The output is similar to the following:

```
ecp:admin> switchstatuspolicyshow
switch:admin> switchstatuspolicyshow
  The current overall switch status policy parameters:
          Down      Marginal
-----
PowerSupplies      0          0
Temperatures       0          0
  Fans              1          0
  WWN               0          0
  CP                0          0
  Blade             0          0
  CoreBlade         0          0
  Flash             0          0
MarginalPorts      0.00%[0]    0.00%[0]
FaultyPorts        0.00%[0]    0.00%[0]
MissingSFPs        0.00%[0]    0.00%[0]
ErrorPorts         0.00%[0]    0.00%[0]
Number of ports: 4
```

Setting the switch status policy threshold values

For switches running Fabric Watch, you can set the switch status policy threshold values using the **switchStatusPolicySet** command

If the switch is running Fabric Watch, you can use the following procedure to set the switch status policy threshold values. If the switch is running MAPS, refer to the *Monitoring and Alerting Policy Suite Administrator's Guide*.

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

The current switch status policy parameter values are displayed. You are prompted to enter values for each DOWN and MARGINAL threshold parameter.

NOTE

By setting the DOWN and MARGINAL values for a parameter to **0,0**, that parameter is no longer used in setting the overall status for the switch.

3. Verify the threshold settings you have configured for each parameter.

Enter the **switchStatusPolicyShow** command to view your current switch status policy configuration.

Example output from a switch

The following example displays what is typically seen from a Brocade switch, but the quantity and types vary by platform.

```
switch:admin> switchstatuspolicyshow

To change the overall switch status policy parameters

The current overall switch status policy parameters:
-----
                Down      Marginal
-----
PowerSupplies   2          1
Temperatures    2          1
  Fans          2          1
  Flash         0          1
MarginalPorts   25.00%[12]    10.00%[5]
FaultyPorts     25.00%[12]    10.00%[5]
MissingSFPS     0.00%[0]      0.00%[0]
ErrorPorts      0.00%[0]      0.00%[0]
Number of ports: 48
```

Note that the value, 0, for a parameter, means that it is NOT used in the calculation.
 ** In addition, if the range of settable values in the prompt is (0..0),
 ** the policy parameter is NOT applicable to the switch.
 ** Simply hit the Return key.

```
The minimum number of
Bad PowerSupplies contributing to DOWN status: (0..2) [2]
Bad PowerSupplies contributing to MARGINAL status: (0..2) [1]
Bad Temperatures contributing to DOWN status: (0..4) [2]1
Bad Temperatures contributing to MARGINAL status: (0..4) [1]2
Bad Fans contributing to DOWN status: (0..2) [2]
Bad Fans contributing to MARGINAL status: (0..2) [1]
(output truncated)
```

NOTE

On the Brocade Backbones, the command output includes parameters related to CP blades.

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.

3 Audit log configuration

Auditable events are generated by the switch and streamed to an external host through a configured system message log daemon (syslog). You specify a filter on the output to select the event classes that are sent through the system message log. The filtered events are streamed chronologically and sent to the system message log on an external host in the specified audit message format. This ensures that they can be easily distinguished from other system message log events that occur in the network. Then, at some regular interval of your choosing, you can review the audit events to look for unexpected changes.

Before you configure audit event logging, familiarize yourself with the following audit event log behaviors and limitations:

- By default, *all event classes* are configured for audit; to create an audit event log *for specific events*, you must explicitly set a filter with the *class* operand and then enable it.
- Audited events are generated specific to a switch and have no negative impact on performance.
- The last 256 events are persistently stored on the switch and are streamed to a system message log.
- 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.

Pushed messages contain the administrative domain of the entity that generated the event. Refer to the *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 *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 *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 **syslogdIpAdd** 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 **syslogdIpAdd** 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.

```
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.
```

3 Duplicate PWWN handling during device login

```
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.
```

Duplicate PWWN handling during device login

If a device attempts to log in with the same port WWN (PWWN) as another device on the switch, you can configure whether the new login or the existing login takes precedence.

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. [Table 9](#) describes the behavior when setting 0 is selected.

TABLE 9 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. [Table 10](#) describes the behavior when setting 1 is selected.

TABLE 10 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.
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.

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 11 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.	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.	New FDISC is rejected.

Setting the behavior for handling duplicate PWWNs

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”.

Enabling 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 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.

The following considerations apply to FEC:

- FEC is supported on E_Ports on 16 Gbps-capable switches.
- 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 on a switch if the device attached supports FEC by using a Brocade host bus adaptor (HBA).
- FEC is enabled by default.

3 Enabling forward error correction

- 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 FEC:

- FEC is configurable only on 16 Gbps-capable switches (Brocade 6505, 6510, 6520, M6505, 6547, and the Brocade DCX 8510 Backbone family).
- For switch to adaptor connections, FEC is supported only on 1860 and 1867 Fabric Adapter ports operating in HBA mode connected to 16 Gbps Brocade switches running Fabric OS 7.1 and later.
- FEC is supported only on link speeds of 10 Gbps and 16 Gbps, regardless of whether the platform is FEC capable.
- FEC is not supported in the following situations:
 - When the HBA port is running on a 16 Gbps link. When the HBA port speed changes to less than this, FEC is disabled.
 - For HBA ports operating in loop mode or in direct-attach configurations.
 - On ports with some DWDM devices.

Using the portCfgFec command

Use the following procedure to enable FEC.

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

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

To enable the FEC feature on a single port and display the configuration, enter the following commands:

```
switch:admin> portcfgfec --enable 1  
switch:admin> portcfgfec --show 1  
Port: 1  
FEC Capable: YES  
FEC Configured: ON
```

Enabling forward error correction

To enable the FEC feature on a port range, enter the **portCfgFec --enable** command. In this example, port 1 already has FEC enabled, and so it remains enabled.

```
switch:admin> portcfgfec --enable 0-8  
Same configuration for port 1
```


Disabling forward error correction

To disable the FEC feature on a port range, enter the **portCfgFec --disable** command.

```
switch:admin> portcfgfec --disable 0-8
```

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 [Chapter 25, “Managing Long-Distance Fabrics”](#) for more details on working with long-distance ports.

Viewing current FEC settings

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

```
switch:admin> portcfgfec --show 1
Port: 1
FEC Capable: YES
FEC Configured: ON
```

3 Enabling forward error correction

Routing Traffic

In this chapter

- Routing overview 115
- Inter-switch links 118
- Gateway links 120
- Routing policies 122
- Route selection 125
- Frame order delivery 126
- Lossless Dynamic Load Sharing on ports 129
- Frame Redirection 132

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.

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. [Figure 5](#) shows the thick red lines as principal ISLs, and thin green lines as regular ISLs.




FIGURE 5 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 [Chapter 1, “Understanding Fibre Channel Services”](#).

You can expand your fabric by connecting new switches to existing switches. [Figure 6](#) shows a new switch being added into an existing fabric. The thick red line is the newly formed ISL.




FIGURE 6 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 *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 **portCfgSLMode** 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 [Chapter 5, “Buffer-to-Buffer Credits and Credit Recovery”](#).

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 [Chapter 14, “Optimizing Fabric Behavior”](#).

Virtual channels are divided into three priority groups. P1 is the highest priority, which is used for Class F, F_RJT, and ACK traffic. P2 is the next highest priority, which is used for data frames. The data virtual channels can be further prioritized to provide higher levels of Quality of Service. P3 is the lowest priority and is used for broadcast and multicast traffic. This example is illustrated in [Figure 7](#).

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 7](#). 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 [Chapter 14, “Optimizing Fabric Behavior”](#).

4 Gateway links




FIGURE 7 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. [Figure 8](#) shows two separate SANs, A-1 and A-2, merged together using a gateway.




FIGURE 8 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 121.
- 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 `portCfgIISLMode` command.
3. Repeat step 1 and step 2 for any additional ports that are connected to the gateway.
4. Repeat this procedure on the switch at the other end of the gateway.

Example of enabling a gateway link on slot 2, port 3

```
ecp:admin> portcfgislmode 2/3, 1
Committing configuration...done.
ISL R_RDY Mode is enabled for port 3. Please make sure the PID
formats are consistent across the entire fabric.
```

Routing policies

By default, all routing protocols place their routes into a routing table. You can control the routes that a protocol places into each table and the routes from that table that the protocol advertises by defining one or more routing policies and then applying them to the specific routing protocol.

The routing policy is responsible for selecting a route based on one of three user-selected routing policies:

- Port-based routing
- Exchange-based routing
- Device-based routing

Notes

Routing is handled by the FSPF protocol and routing policy.

Each switch can have its own routing policy and different policies can exist in the same fabric.

ATTENTION

For most configurations, the default routing policy is optimal and provides the best performance. You should change the routing policy only if there is a significant performance issue, or a particular fabric configuration or application requires it.

Displaying the current routing policy

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **aptPolicy** command with no parameters.

The current policy is displayed, followed by the supported policies for the switch.

In the following example, the current policy is exchange-based routing (3) with the additional AP Dedicated Link policy.

```
switch:admin> aptpolicy
Current Policy: 3
3 : Default Policy
   1: Port Based Routing Policy
   2: Device Based Routing Policy (FICON support only)
   3: Exchange Based Routing Policy
     0: AP Shared Link Policy
     1: AP Dedicated Link 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 *Fibre Channel over IP Administrator's 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 124.

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 124.

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.

AP route policies

Two additional AP policies are supported under exchange-based routing:

- AP Shared Link policy (default)
- AP Dedicated Link policy

NOTE

AP policies are independent of routing policies. Every routing policy supports both AP policies.

The AP Dedicated Link policy relieves internal congestion in an environment in which:

- There is a large amount of traffic going through both directions at the same time.
- There is a reduction of the effect of slow devices on the overall switch performance.

It is recommended that the default AP Shared Link policy be used for most environments. Also, it is recommended that you design a SAN that localizes host-to-target traffic by reducing the amount of traffic through the router.

ATTENTION

Setting either AP route policy is a disruptive process.

Routing in Virtual Fabrics

Virtual Fabrics (VF) supports DPS on all partitions. DPS is limited where multiple paths are available for a logical fabric frame entering a Virtual Fabrics chassis from a base fabric that is sent out using one of the dedicated ISLs in a logical switch.

The AP policy affecting the DPS behavior, whether it is exchange-based, device-based, or port-based, is configured on a per-logical switch basis. In-order delivery (IOD) and DLS settings are set per logical switch as well. IOD and DLS settings for the base switch affect all traffic going over the base fabric including any logical fabric traffic that uses the base fabric.



CAUTION

Setting the routing policy is disruptive to the fabric because it requires that you disable the switch where the routing policy is being changed.

Setting the routing policy

Use the following procedure to set the routing policy.

1. Connect to the VF switch and log in as admin.
2. Enter the **setcontext** [*FID* | *switchname*] command for the correct Fabric ID or switch name.
 - 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.

```
switch:admin> setcontext 20
```
3. Enter the **switchDisable** command to disable the switch.
4. Take the appropriate following action based on the AP route policy you choose to implement:
 - If the exchange-based policy is required, enter the **aptPolicy 3** command.
 - If the port-based policy is required, enter the **aptPolicy 1** command.

Setting the AP route policy

The AP route policy can only be set in the base switches that are using Virtual Fabrics.

Use the following procedure to set the AP route policy.

1. Connect to the base switch and log in as admin.
2. Enter the **switchDisable** command to disable the switch.
3. Take the appropriate following action based on the AP route policy you choose to implement:
 - If the AP Shared Link policy (default) is required, enter the **aptPolicy -ap 0** command.
 - If the AP Dedicated Link policy is required, enter the **aptPolicy -ap 1** command.

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” indicates that DLS is turned on.
- “DLS is not set” 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.
- “DLS is set by default with current routing policy. DLS is set with Lossless enabled.” The current routing policy (exchange-based) requires DLS to be enabled by default. In addition, the Lossless option is enabled. Frame loss is prevented during a load sharing recomputation. If you get this message, you cannot perform [step 3](#), so you are done with this procedure.

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
switch:admin> dlsset
switch:admin> dlsshow
DLS is set
switch:admin> dlsreset
switch:admin> dlsshow
DLS is not set
```

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.

Using Frame Viewer to understand why frames are dropped

When a frame is unable to reach its destination because of a timeout, it is discarded. You can use Frame Viewer to find out which flows contained the dropped frames, which in turn can help you determine which applications might be impacted. Frame Viewer allows you to see the exact time (within one second) that the frames were dropped.

You can view and filter up to 20 discarded frames per chip per second for 1200 seconds using a number of fields with the **frameLog** command.

4 Frame order delivery

Use the following procedure to view frames.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **frameLog --show**.

```
EDCX16_114064:root> frameLog --show
```

```
=====
                Fri Jul 13 23:47:08 UTC 2012
=====
Log
timestamp      TX      RX
                port  port      SID      DID      SFID      DFID      Type      Count
=====
Jul 13 23:47:07 11/45 11/45      0xfffffd  0x40e580  0         0        timeout    2
Jul 13 23:47:07 11/45 11/45      0xfffffc  0x40e580  0         0        timeout    5
Jul 13 23:47:07 11/45 11/45      0xfffffc  0x40e580  0         0        timeout    3
Jul 13 23:47:07 11/45 11/45      0xfffc40  0x40e580  0         0        timeout    2
Jul 13 23:47:07 11/45 11/45      0xfffc40  0x40e580  0         0        timeout    1
```

Using the frameLog --show command

The output of **--show** displays the type of each discard.

The **-type** option of the **frameLog --show** command requires an argument, but only **timeout** is supported at present. The **timeout** argument specifies that only timeout discards are shown.

Filtering results by back-end port in Frame Viewer

The Frame Viewer **--show** command supports specifying that the TX port or RX port of displayed frames should be a back-end port. To filter by TX port or RX port, use a following command:

```
frameLog --show -txport [slot/]port
```

or

```
frameLog --show -rxport [slot/]port
```

or

```
frameLog --show -txport [slot/]port -rxport [slot/]port
```

The **-txport** and **-rxport** options accept the arguments “-1” (for fixed-port switches) or “-1/-1” (for modular switches). These arguments stand for “any back-end port.” Using this notation, you can select specifically those discarded frames that have a back-end port in the TX port or RX port field.

NOTE

Individual back-end ports cannot be specified, only the quality of being a back-end port can be specified.

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 300
- Brocade 5100
- Brocade 5300
- Brocade 6505
- Brocade 6510
- Brocade 6520
- Brocade M6505
- Brocade 6547
- Brocade VA-40FC
- Brocade FC8-16, FC8-32, FC8-48, and FC8-64 port blades
- Brocade DCX 8510 Backbone family and supported blades
- Brocade FC16-32 and FC16-48 port blades
- Brocade FC8-32E and FC8-48E port blades
- Brocade FX8-24 application blades in the Brocade DCX and DCX-4S Backbones

On the Brocade 7800 switch and the FX8-24 application blade, Lossless DLS is supported only on FC-to-FC port flows.

ATTENTION

When you implement Lossless DLS, the switches in the fabric must have either Fabric OS v6.3.0 or Fabric OS v6.4.0 or later installed to guarantee no frame loss.

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 an E_Port
- Adding a slave E_Port
- Removing an E_Port (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. Changes the existing path to a more optimal path.
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.

[Table 12](#) shows the effect of frames when you have a specific routing policy turned on with IOD.

TABLE 12 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.

Lossless core

Lossless core works with the default configuration of the Brocade DCX 8510-8 and DCX 8510-4 hardware 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.

Traffic flow limitations

FA4-18 AP blades, which are supported on the Brocade DCX and DCX-4S devices, may continue to experience frame drops after core blade removal or insertion. The path between an FA4-18 blade and an FX8-24 blade, or vice versa, experiences I/O disruption because the FA4-18 blades do not support this feature.

Configuring Lossless Dynamic Load Sharing

You configure Lossless DLS switch- or chassis-wide by using the **dlsSet** command to specify that no frames are dropped while rebalancing or rerouting traffic.

Use the following procedure to configure Lossless Dynamic Load Sharing.

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
```

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. XISL use must be disabled for Lossless DLS to be enabled.

How DLS affects other logical switches in the fabric

On a Brocade DCX 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 10 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 DSL 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 [Chapter 11, “Managing Virtual Fabrics”](#).

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 9 Single host and target

Figure 9 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 350, and “Enabling a zone configuration” on page 364.

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.
2. Enter the **zone --rddelete** command to remove the base redirect zone object, "red_____base".

NOTE

When the base zone is removed, the redirect zone configuration "r_e_d_i_r_c__fg" is removed as well.

3. Enter the **cfgSave** command to save changes to the defined configuration.

Example of deleting a frame redirect zone

```
switch:admin> zone --rddelete \
red_0917_10_10_10_10_10_10_10_20_20_20_20_20_20_20
```

Viewing frame redirect zones

Use the following procedure to view frame redirect zones.

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

4 Frame Redirection

Buffer-to-Buffer Credits and Credit Recovery

In this chapter

- [Buffer credit management](#) 135
- [Buffer credit recovery](#) 146
- [Credit loss](#) 149

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

Considerations follow for calculating how many ports can be configured for long distance on all Fabric OS v7.x-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 138 and [“Configuring buffers for a single port directly”](#) on page 141 for more information.

Fibre Channel gigabit values reference definition

Use the following Fibre Channel gigabit values to calculate buffer requirements:

[Table 13](#) shows the Fibre Channel gigabit values used to calculate buffer requirements.

TABLE 13 Fibre Channel gigabit values

Gigabit value	Buffer requirements
1 Gbps	1.0625
2 Gbps	2.125
4 Gbps	4.25
8 Gbps	8.5
10 Gbps	10.625
16 Gbps	17

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 140 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 14 describes Fibre Channel data frames.

TABLE 14 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 15 on page 143 and Table 16 on page 144 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 16.

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.

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 15](#) on page 143

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:

$$\text{Maximum Distance } X \text{ (in km)} = (\text{BufferCredits} + 6) * 2 / \text{LinkSpeed}$$

$$498 \text{ km} = (492 + 6 \text{ buffers for Fabric Services}) * 2 / 2 \text{ Gbps}$$

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 137 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 137 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:

$$\text{Total Buffers} = \text{Configured Buffers} + \text{QOS_VC_Credits} + \text{Non-data_VC_Credits}$$

Seven Virtual Channels (VCs) are required for each QoS port. Each VC requires two buffers. Thus, the total number of QoS buffers required for a port is 14 (7*2). An additional 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.

$$\text{Total Buffers} = 100 + 14 + 6 = 120$$

If the configured number of buffers is 100, the total number of buffers allocated for a non-QoS port is 106, as shown in the following example.

$$\text{Total Buffers} = 100 + 6 = 106$$

NOTE

You cannot use the **-buffers** option with the **-distance** option or the **-frameSize** option.

Example

```
switch:admin> portcfglongdistance 2/35 LS 1 -buffers 400
Reserved Buffers =          420
```

Configuring buffers using frame size

You can configure the number of buffers by using the **-frameSize** option of the **portCfgLongDistance** command along with the **-distance** option. Fabric OS calculates the number of buffers from the **-frameSize** option value according to the following formula:

$$\text{buffers_required} = (2048/\text{framesize}) * \text{data_vc_credits}$$

If you enter the average frame size of 1024, Fabric OS 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
```

Allocating buffer credits for F_Ports

The default configured F_Port buffer credit is fixed at eight buffers. You can use the **portCfgFPortBuffers** command to configure a given port with the specified number of buffers.

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

```
switch:admin> portcfgfportbuffers --enable 2/44 12
```

Note that in the sample commands provided in the following procedure, 12 buffers are configured for an F_Port.

To disable the port buffer configuration and return to the default buffer allocation, use the **--disable** option.

```
switch:admin> portcfgfportbuffers --disable 2/44
```

NOTE

The configured number of buffers for a given port is stored in the configuration database and is persistent across reboots. The F_Port buffer feature does not support EX_Port, Port Mirroring, 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 17
```

User Port	Port Type	Lx Mode	Max/Resv Buffers	Avg Buffer Usage & FrameSize		Buffer Usage	Needed Buffers	Link Distance	Remaining Buffers
				Tx	Rx				
64		-	8	- (-)	- (-)	0	-	-	
65		-	8	- (-)	- (-)	0	-	-	
66		-	8	- (-)	- (-)	0	-	-	
67		-	8	- (-)	- (-)	0	-	-	
68	E	LS	806	197 (2012)	201 (2044)	206	206	100km	
69	E	-	8	1 (2016)	1 (2020)	26	26	2km	
70	E	-	8	1 (2012)	1 (2036)	26	26	2km	
71	E	-	8	1 (2008)	2 (2052)	26	26	2km	
192		-	8	- (-)	- (-)	0	-	-	
193		-	8	- (-)	- (-)	0	-	-	
194		-	8	- (-)	- (-)	0	-	-	
195		-	8	- (-)	- (-)	0	-	-	
196		-	8	- (-)	- (-)	0	-	-	
197		-	8	- (-)	- (-)	0	-	-	
198		-	8	- (-)	- (-)	0	-	-	
199		-	8	- (-)	- (-)	0	-	-	4556

Buffer credits switch or blade model

Table 15 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.

TABLE 15 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
300	24	24	484
5100	40	40	1692
5300	80	16	292
5410	12	12	580
5424	24	24	484
5431	16	16	548
5450	26	26	468
5480	24	24	484
M6505	24	24	7904
6505	24	24	7952
6510	48	48	7760
6520	96	48	4256
6547	48	48	7712
7800	16	16	408

TABLE 15 Total FC ports, ports per port group, and unreserved buffer credits per port group (Continued)

Switch/blade model	Total FC ports (per switch/blade)	User port group size	Unreserved buffer credits per port group
VA-40FC	40	40	1692
Brocade Encryption Switch	32	16	1392
FC8-16	16	16	1292/508
FC8-32	32	16	1292/508
FC8-32E	32	16	5456
FC8-48	48	24	1228/716
FC8-48E	48	24	5008
FC8-64	*** Extended Fabrics is not supported on this blade ***		
FC16-32	32	16	5456
FC16-48	48	24	5008
FS8-18	16	8	1604
FX8-24	12	12	1060

For the FC8-x port 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.

Maximum configurable distances for Extended Fabrics

Table 16 shows the maximum supported extended distances (in kilometers) that can be configured for one port on a specific switch or blade at different speeds.

TABLE 16 Configurable distances for 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
300	486	243	121	N/A	N/A
5100	1694	847	423	N/A	N/A
5300	294	147	73	N/A	N/A
5410	582	291	145.5	N/A	N/A
5424	486	243	121.5	N/A	N/A
5431	550	275	137.5	N/A	N/A
5450	470	235	117.5	N/A	N/A
5480	486	243	121.5	N/A	N/A
M6505	N/A	3953	1976	N/A	988
6505	7426	3713	1856	1485	928
6510	6754	3377	1688	1350	844
6520	4064	2032	1016	812	508

TABLE 16 Configurable distances for Extended Fabrics (Continued)

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
6547	7714	3857	1928	1542	964
7800	410	205	102	N/A	N/A
VA-40FC	1694	847	423	N/A	N/A
Brocade Encryption Switch	1392	696	348	N/A	N/A
FC8-16	1294	647	323	N/A	N/A
FC8-32	1294	647	323	N/A	N/A
FC8-32E	5190	2595	1297	1038	648
FC8-48	1230	615	307	N/A	N/A
FC8-48E	4486	2243	1121	897	560
FC8-64	*** Extended Fabrics is not supported on this blade ***				
FC16-32	5190	2595	1297	1038	648
FC16-48	4486	2243	1121	897	560
FS8-18	1604	802	401	N/A	N/A
FX8-24	1062	531	265	N/A	N/A

NOTE

The distances in table 17 assume that QoS is enabled. If QoS is disabled, the maximum supported distances are higher, because QoS requires an additional 20 buffer credits per active port.

Estimated maximum equally distributed distance = 1-port maximum distance/Number of ports

For example, for three ports running at 2 Gbps on a Brocade 300 switch, the maximum equally distributed distance is calculated as $486 / 3 = 164$ km.

Downgrade considerations

When Fabric OS firmware is downgraded from version 7.1 to an earlier version, the effect depends on whether the number of buffer credits for the long-distance port is configured with the **-framesize** and **-distance** options or with the **-buffers** option.

When a port is configured with -framesize and -distance options

In Fabric OS v7.1, if you configure the port by using the **-distance** option alone, the reserved buffers are calculated according to the distance. If you configure both the **-framesize** option and the **-distance** option, more buffers will be reserved, depending on the frame size.

With a firmware downgrade, those ports that were configured with more reserved buffers will keep the reserved buffers as long as the ports remain online. The next time the port is toggled, buffers will again be reserved on the basis of distance only.

When a port is configured with the `-buffers` option

A firmware downgrade is blocked when a port is configured as a long-distance port by means of the `-buffers` option. The following warning message is displayed:

```
Downgrade to selected version is not allowed because few ports are configured with Longdistance -buffers option. Please remove the configuration using portcfglongdistance <slot>/<port> L0 CLI or change the configuration with -distance option on the console.
```

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. This feature is not applicable on ICL_ports.

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 an E_Port.

```
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 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 FC devices and blades (Brocade 6505, 6510, 6520, M6505, 6547, CR16-4, CR16-8, FC8-32E, FC8-48E, FC16-32, FC16-48), you can use the **portCfgCreditRecovery** command to disable or enable buffer credit recovery on a port.

Buffer credit recovery over an E_Port

To support buffer credit recovery, E_Ports must be connected between devices that support 16 Gbps or between devices that support 8 Gbps.

- Devices that support 16 Gbps:
 - Brocade 6505, 6510, 6520, M6505, 6547
 - FC8-32E, FC8-48E, FC16-32, FC16-48
- Devices that support 8 Gbps:
 - Brocade 300, 5100, 5300, 5410, 5424, 5450, 5480, VA-40FC
 - FC8-16, FC8-32, FC8-48

If a device that supports 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.

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 HBA v3.2 firmware or later.
- The adapter must operate at maximum speed.
- 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 both support 16 Gbps or 8 Gbps. Buffer credit recovery is not supported if the EX_Ports do not support the same data rate.
- 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 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, but refer to the *Fabric OS Command Reference* for lists of devices and blades that support 16 Gbps and 8 Gbps.

Enabling and disabling buffer credit recovery

To disable buffer credit recovery on a port, perform the following steps.

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

To enable buffer credit recovery on a port for which it has been disabled, perform the following steps.

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

Fabric OS v7.1 and later supports back-end credit loss detection, back-end ports and core blades, and the Brocade 5300 and 6520 switches, although the support is slightly different on each device. Refer to the following details on these switches, and the *Fabric OS Troubleshooting and Diagnostics Guide* for more general information.

Back-end credit loss detection and recovery support on Brocade 5300 switches

The following credit loss detection methods are supported for Brocade 5300 back-end ports:

- Per-port polling to detect credit loss. If credit loss is detected using this method, the RASlog C2-1012 message is displayed and recorded.
- On-demand VC credit loss detection. If credit loss is detected using this method, the RASlog C2-1027 message is displayed and recorded.
- TX timeout trigger automatic VC credit loss detection. If credit loss is detected using this method, the RASlog C2-1027 message is displayed and recorded.

The following credit loss recovery methods are supported for Brocade 5300 back-end ports:

- For per-port polling and on-demand VC credit loss methods, a link reset will automatically be performed, assuming that this option was enabled. Refer to [“Enabling back-end credit loss detection and recovery”](#) for details on enabling this feature.
- For the TX timeout trigger automatic VC method, a link reset will be automatically performed if complete credit loss on a VC is detected.
- A manual link reset option using the **creditRecovMode** command is also available. Refer to [“Enabling back-end credit loss detection and recovery”](#) for instructions.

NOTE

Whenever a link reset is performed on this switch, the RASlog C2-1014 message is displayed and recorded.

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 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 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

Credit loss detection and recovery is enabled and disabled through the CLI using the **creditRecovMode --cfg** command.

- The execution of this command is subject to Virtual Fabrics or Admin Domain restrictions that may be in place. Refer to the *Fabric OS Troubleshooting and Diagnostics Guide* for more information.
- The bottleneck detection commands are supported on F_Ports, FL_Ports, E_Ports, and EX_Ports.
- The credit recovery commands are supported only on back-end ports of 4 Gbps-, 8 Gbps-, and 16 Gbps-capable FC platforms for blades in the Brocade DCX, DCX-4S, DCX 8510-8, and DCX 8510-4 chassis.

To enable back-end credit loss detection and recovery, perform the following steps.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **creditRecovMode --cfg** command 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 the **creditRecovMode --show** command 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
```

Managing User Accounts

In this chapter

- [User accounts overview](#) 151
- [Local database user accounts](#) 155
- [Local user account database distribution](#) 158
- [Password policies](#) 159
- [The boot PROM password](#) 163
- [Remote authentication](#) 167

User accounts overview

In addition to the default permissions assigned to the roles of root, factory, admin, and user, Fabric OS supports up to 252 additional user accounts on the chassis. These accounts expand your ability to track account access and audit administrative activities.

Each user account is associated with the following:

- Admin Domain list — Specifies the Administrative Domains to which a user account is allowed to log in.
- Home Admin Domain — Specifies the Admin Domain that the user is logged in to by default. The home Admin Domain must be a member of the user's Admin Domain list.
- Permissions — Associate roles with each user account to determine the functional access levels within the bounds of the user's current Admin Domain.
- 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 fabric ID is not available, the next-lower valid fabric ID 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.

NOTE

Admin Domains are mutually exclusive from Virtual Fabrics permissions when you set up user accounts. You will need to set up different user accounts for each feature.

You cannot have Admin Domain mode and Virtual Fabrics mode enabled at the same time.

For more information about Admin Domains, refer to [Chapter 20, "Managing Administrative Domains"](#).

For more information about Virtual Fabrics, refer to [Chapter 11, "Managing Virtual Fabrics"](#).

Fabric OS provides four options for authenticating users: remote RADIUS service, remote LDAP service, remote TACACS+ service, and the local-switch user database. All options allow users to be managed centrally by 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 v5.3.0 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.

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 determines the jobs and tasks that can be performed on a fabric and its associated fabric elements. Fabric OS uses RBAC to determine which commands 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.

[Table 17](#) outlines the Fabric OS predefined (default) roles.

TABLE 17 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 and Admin Domains 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	Nonadministrative use, such as monitoring system activity
ZoneAdmin	Zone administration	Zone management commands only

Admin Domain considerations

Legacy users with no Admin Domain specified and whose current role is admin will have access to ADO through AD255 (physical fabric admin); otherwise, they will have access to ADO only.

If some Admin Domains have been defined for the user and all of them are inactive, the user will not be allowed to log in to any switch in the fabric. If no home domain is specified for a user, the system provides a default home domain.

The default home domain for the predefined account is ADO. For user-defined accounts, the default home domain is the Admin Domain in the user's Admin Domain list with the lowest ID.

Role permissions

Table 18 describes the types of permissions that are assigned to roles.

TABLE 18 Permission types

Abbreviation	Definition	Description
O	Observe	The user can run commands by using options that display information only, such as running <code>userConfig --show -a</code> to show all users on a switch.
M	Modify	The user can run commands by using options that create, change, and delete objects on the system, such as running <code>userConfig --change username -r rolename</code> 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.

```
> classconfig --showroles authentication
Roles that have access to the RBAC Class 'authentication' are:
```

```

Role name          Permission
-----
Admin              OM
Factory            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. [Table 19](#) shows the number of simultaneous login sessions allowed for each role when authenticated locally. The roles are displayed in alphabetic order, which does not reflect their importance. When LDAP, RADIUS, or TACACS+ are used for authentication, the total number of sessions on a switch may not exceed 32.

TABLE 19 Maximum number of simultaneous sessions

Role name	Maximum sessions
Admin	2
BasicSwitchAdmin	4
FabricAdmin	4
Operator	4
SecurityAdmin	4
SwitchAdmin	4
User	4
ZoneAdmin	4

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.

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 on a chassis is 256.

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:

```
> 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.

```
> classConfig --showroles security
Roles that have access to RBAC Class 'security' are:
```

Role Name	Permissions
-----	-----
User	O
Admin	OM
Factory	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

You can assign a user-defined role to a user by using one of the following options of the **userConfig** command:

- **userConfig --add** with the **-r** option to create a new user account and assign a role.
- **userConfig --change** with the **-r** option to add or change a user-defined role for an existing user account.
- **userConfig --add** with the **-c** option to create a new user account and assign a chassis role.
- **userConfig --change** with the **-c** option to add a chassis role to an account.

The following example assigns the mysecurityrole role to the existing anewuser account and adds the admin chassis role:

```
> userConfig --change anewuser -r mysecurityrole -c admin
```

Local database user accounts

User **add**, **change**, and **delete** operations are subject to the *subset* rule: an admin with ADlist 0–10 or LFlist 1–10 cannot perform operations on an *admin*, *user*, or *any* role with ADlist 11–25 or LFlist 11–128. The user account being changed must have an ADlist or 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

Table 20 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 20 Default local user accounts

Account name	Role	Admin Domain	Logical Fabric	Description
admin	Admin	ADO-255 home: 0	LF1-128 home: 128	Most commands have <i>Observe/Modify</i> permission.
factory	Factory	ADO-255 home: 0	LF1-128 home: 128	Reserved
root	Root	ADO-255 home: 0	LF1-128 home: 128	Reserved.
user	User	ADO home: 0	LF-128 home: 128	Most commands have <i>observe-only</i> permission.

Admin Domain and Virtual Fabrics considerations: Administrators can act on other accounts only if that account has an Admin Domain or Logical Fabric list that is a subset of the administrator.

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
 - **userConfig --show username** to show account information for the specified account
 - **userConfig --showad -a adminDomain_ID** to show all accounts permitted to select the specified adminDomain_ID
 - **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.

When changing account parameters, if you change the ADlist for the user account, all of the currently active sessions for that account will be logged out. For more information about changing the Admin Domain on an account, refer to [Chapter 20, “Managing Administrative Domains”](#).

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 ADlist 0–10 or LFlist 1–10 cannot change the password on an account with admin, user, or any permission with an ADlist 11–25 or LFlist 11–128. The user account being changed must have an ADlist 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.
- You cannot change passwords by 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.
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** 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

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 158).

All password policies are enforced during logins to the standby CP. However, you may observe that the password enforcement behavior on the standby CP is inconsistent with prior login activity; this is because password state information from the active CP is automatically synchronized with the standby CP, thereby overwriting any password state information that was previously stored there. Also, password changes are not permitted on the standby CP.

Password authentication policies configured using the `passwdCfg` command are *not* enforced during initial prompts to change default passwords.

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.
- **Repeat**
Specifies the length of repeated character sequences that will be disallowed. For example, if the “repeat” value is set to 3, a password “passAAword” 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.

```
> 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.

Specify the number of past password values that are disallowed when setting a new password. 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 policy 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 policy 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.

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 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.

```
> passwdcfg --setuser metoo -minpasswordage 10 -maxpasswordage 90 -warning 3
```

The following example configures a password expiration policy for all users.

```
> passwdcfg --set -minpasswordage 5 -maxpasswordage 30 -warning 5
```

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 name of the user account that is locked out.

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.

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 165.

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 and software”](#) on page 35. 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 58.
2. Reboot the switch.
3. Press **Esc** within four seconds after the message “Press escape within 4 seconds...” is displayed.

The following options are available:

Option	Description
1 Start system.	Continues the system boot process.
2 Recovery password.	Lets you set the recovery string and the boot PROM password.
3 Enter command shell.	Provides access to boot parameters.

4. Enter **2**.

- 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 a Backbone with a recovery string

This procedure applies to the Brocade DCX, DCX-4S, DCX 8510-4, and DCX 8510-8 Backbones. 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 58.
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.

The following options are available:

Option	Description
1	Start system. Continues the system boot process.
2	Recovery password. Lets you set the recovery string and the boot PROM password.
3	Enter command shell. Provides access to boot parameters.

5. Enter **2**. Take the following appropriate action based on whether you find the password was previously set:
 - 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 163. 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 the fixed-port switch models.

The password recovery instructions provided within this section are only for the switches listed in [“Supported hardware and software”](#) on page 35. 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 58.
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.

The following options are available:

Option	Description
1	Start system. Continues the system boot process.
2	Recovery password. Lets you set the recovery string and the boot PROM password.
3	Enter command shell. Provides access to boot parameters.

4. Enter **3**.
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 without a recovery string

This procedure applies to the Brocade DCX, DCX-4S, DCX 8510-4, and DCX 8510-8 Backbones.

On the Brocade DCX Backbone, 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 58.
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.

The following options are available:

Option	Description
1	Start system. Continues the system boot process.
2	Recovery password. Lets you set the recovery string and the boot PROM password.
3	Enter command shell. Provides access to boot parameters.

6. Enter **3**.
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 *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.
- 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 API. 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 must 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.

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

Table 21 summarizes the various LDAP options and Brocade support for each.

TABLE 21 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
LDAPv3 with TLS and Certificate	LDAPv3 over TLS channel and authenticated using a certificate	Secured	389	ldap://	Yes
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

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

Command options

Table 22 outlines the **aaaConfig** command options used to set the authentication mode.

TABLE 22 Authentication configuration options

aaaConfig options	Description	Equivalent setting in Fabric OS v5.1.0 and earlier	
		--radius	--switchdb ¹
--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.	Off	On
--authspec "radius"	Authenticates management connections against any RADIUS databases only. If the RADIUS service is not available or the credentials do not match, the login fails.	On	Off

TABLE 22 Authentication configuration options (Continued)

aaaConfig options	Description	Equivalent setting in Fabric OS v5.1.0 and earlier	
		--radius	--switchdb ¹
--authspec "radius;local"	Authenticates management connections against any RADIUS databases first. If RADIUS fails for any reason, authenticates against the local user database.	not supported	not supported
--authspec "radius;local" --backup	Authenticates management connections against any RADIUS databases. 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.	On	On
--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.	n/a	n/a
--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.	n/a	On
--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.	n/a	On
--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.	not supported	not supported
--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.	not supported	not supported
--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.	not supported	not supported
--authspec -nologout	Prevents users from being logged out when you change authentication. Default behavior is to log out users when you change authentication.	n/a	n/a

1. Fabric OS v5.1.0 and earlier aaaConfig --switchdb <on | off> setting.

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 ldap_role name switch_role` command to map LDAP server permissions.

RADIUS, LDAP, and TACACS+ support all the defined RBAC roles described in [Table 17](#) on page 152.

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. If no Administrative Domain is assigned, then the user is assigned to the default Admin Domain ADO.

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

TABLE 23 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

TABLE 23 Syntax for VSA-based account roles (Continued)

Item	Value	Description
Vendor type	1	1 octet, Brocade-Auth-Role; valid attributes for the Brocade-Auth-Role are: Admin BasicSwitchAdmin FabricAdmin Operator SecurityAdmin SwitchAdmin User ZoneAdmin
	2	<i>Optional:</i> Specifies the Admin Domain or Virtual Fabric member list. For more information on Admin Domains or Virtual Fabrics, refer to “RADIUS configuration with Admin Domains or Virtual Fabrics” on page 173. 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
Attribute-specific data	ASCII string	Multiple octet, maximum 253, indicating the name of the assigned role and other supported attribute values such as Admin Domain member list.

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 2000 IAS

To configure a Windows 2000 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 [Figure 10](#).




FIGURE 10 Windows 2000 VSA configuration

Linux FreeRADIUS server

For the configuration on a Linux FreeRADIUS server, define the values outlined in Table 24 in a vendor dictionary file called `dictionary.brocade`.

TABLE 24 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 Admin Domain or 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 Admin Domains or Virtual Fabrics

When configuring users with Admin Domains or Virtual Fabrics, you must also include the Admin Domain or 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.

- *HomeAD* is the designated home Admin Domain for the account. The valid range of values is from 0 through 255. The first valid HomeAD key-value pair is accepted by the switch, and any additional HomeAD key-value pairs are ignored.
- *ADList* is a comma-separated list of Administrative Domain numbers of which this account is a member. Valid numbers range from 0 through 255. A dash between two numbers specifies a range. Multiple ADlist key-value pairs within the same or across the different Vendor-Type codes are concatenated. Multiple occurrences of the same Admin Domain number 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 ADList, HomeAD, HomeLF, and LFRoleList are optional. If they are unspecified, the account can log in with ADO as its member list and home Admin Domain or VF128 as its member list and home Virtual Fabric. If there is an error in the ADlist, HomeAD, LFRoleList, or HomeLF specification, the account cannot log in until the AD list or Virtual Fabric list is corrected; an error message is displayed.

For example, on a Linux FreeRADIUS Server, the user (user-za) with the following settings takes the “zoneAdmin” permissions, with AD member list: 1, 2, 4, 5, 6, 7, 8, 9, 12; the Home Admin Domain will be 1.

```
user-za Auth-Type := Local, User-Password == "password"
Brocade-Auth-Role = "ZoneAdmin",
Brocade-AVPairs1 = "ADList=1,2,6,"
Brocade-AVPairs2 = "ADList=4-8;ADList=7,9,12"
```

In the next example, on a Linux FreeRADIUS Server, the user has the “operator” permissions, with ADList 1, 2, 4, 5, 6, 7, 8, 9, 12, 20 and HomeAD 2.

```
user-opr Auth-Type := Local, User-Password == "password"
Brocade-Auth-Role = "operator",
Brocade-AVPairs1 = "ADList=1,2;HomeAD=2",
Brocade-AVPairs2 = "ADList=-4-8,20;ADList=7,9,12"
```

In the next example, on a Linux FreeRADIUS Server, the user has the “zoneAdmin” permissions, with VList 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, one can set these permissions to user, admin, and securityAdmin.

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 root, admin, switchAdmin, zoneAdmin, securityAdmin, basic SwitchAdmin, fabricAdmin, operator, and user. You must use quotation marks around "password" and "role".

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 will 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/raddb/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:

```
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 193).

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 2000

The instructions for setting up RADIUS on a Windows 2000 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 2000 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 RADIUS service on Windows 2000 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 root, admin, factory, switchAdmin, and user, and then add any users whose logins you want to associate to 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.

- a. For the Add RADIUS Client window, provide the following:

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.

- b. 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 (root, admin, factory, switchAdmin, and user) for which you want to use RADIUS. Apply this policy to the user groups that you already created.

- c. In the Vendor-Specific Attribute Information window, enter the vendor code value **1588**. Click the **Yes. It conforms** option, and then click **Configure Attribute**.

- d. 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 (root, admin, switchAdmin, user, and so on) that the user group must use to log in to the switch.

- e. 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.

NOTE

Windows 2008 RADIUS (NPS) support is also available.

RSA RADIUS server

Traditional password-based authentication methods are based on *one-factor* authentication, where you confirm your identity using a memorized password. Two-factor authentication increases the security by using a second factor to corroborate identification. The first factor is either a PIN or password and the second factor is the RSA SecurID token.

RSA SecurID with an RSA RADIUS server is used for user authentication. The Brocade switch does not communicate directly with the RSA Authentication Manager, so the RSA RADIUS server is used in conjunction with the switch to facilitate communication.

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

```
C:\Program Files\RSA Security\RSA RADIUS\Service
```

[Figure 11](#) on page 180 shows what the `brocade.dct` file should look like and [Figure 12](#) on page 180 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

```
#####
# 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
#####
```

FIGURE 11 Example of a brocade.dct file

```
#####
# dictiona.dcm
#####

# Generic Radius

@radius.dct

#
# Specific Implementations (vendor specific)
#
@3comsw.dct
@aat.dct
@acc.dct
@accessbd.dct
@agere.dct
@agns.dct
@airespace.dct
@alcatel.dct
@altiga.dct
@annex.dct
@aptis.dct
@ascend.dct
@ascndvsa.dct
@axc.dct
@bandwagn.dct
@brocade.dct <-----
```

FIGURE 12 Example of the dictiona.dcm file

- 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 184.

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 [Chapter 8, “Configuring Security Policies”](#). The following are restrictions when using LDAP in non-FIPS mode:

- There is no password change through Active Directory.
- There is no automatic migration of newly created users 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).
- Active Directory LDAP 2000, 2003, and 2008 are supported.

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 server -d domain` 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 ldap_role_name switch_role` 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 152.

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) 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` command 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 Administrative Domains or 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 Admin Domains or Virtual Fabrics to the user name. Multiple Admin Domains can be added as a string value separated by the underscore character (`_`). 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 ldap_role_name switch_role` 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 (root, admin, switchAdmin, user, and so on) that the user must use to log in to the switch.

Adding an Admin Domain or 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. Perform the appropriate action based on whether you are using Administrative Domains or Virtual Fabrics:

- If you are using Administrative Domains, enter the values of the Admin Domains separated by an underscore (`_`) into the **Value** field.

Example for adding Admin Domains

```
adlist_0_10_200_endAd
```

Home Admin Domain (homeAD) for the user will be the first value in the *adlist* (Admin Domain list). If a user has no values assigned in the *adlist* attribute, then the homeAD "0" will be the default administrative domain for the user.

- If you are using Virtual Fabrics, 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 181.

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 [Chapter 8, “Configuring Security Policies”](#). 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 ldap_role_name switch_role` command to map LDAP server permissions to one of the default roles available on a switch. For more information on RBAC roles, see [“Role-Based Access Control”](#) on page 152.

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) 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 Admin Domains or 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 Administrative Domains or 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` objectClass. 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

#####
TLSCertificateFile /root/sachin/ldapcert/cacert.pem
TLSCertificateFile /root/sachin/ldapcert/serverCert.pem
TLSCertificateKeyFile /root/sachin/ldapcert/serverKey.pem
TLSVerifyClient never

pidfile      /usr/local/var/run/slapd.pid
argsfile     /usr/local/var/run/slapd.args

database     bdb
suffix       "dc=mybrocade,dc=com"
rootdn       "cn=Manager,dc=mybrocade,dc=com"
rootpw       {SSHA}HL8uT5hPaWyIdcP6yAheMT8n0GoWubr3

# 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](#).

```
> 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 185 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 ldap_role_name switch_role** command to map LDAP server permissions to one of the default roles available on a switch.

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 -f filename** command, where *filename* is the .ldif file you created in [step 1](#).

Example to delete a user attribute

1. Create or edit a .ldif file with an entry similar to the following.

```
#####Deleting an attr
#dn: cn=Sachin,cn=Users,dc=mybrocade,dc=com
```

```
#changetype: modify
#delete: memberof
```

2. Enter the following **ldapmodify** command, where test.ldif is the name of the file you edited in [step 1](#).

```
> ldapmodify -D cn=Sachin,dc=mybrocade,dc=com -x -w secret -f test.ldif
```

Example to add a group member

1. Create or edit a .ldif file with an entry similar to the following.

```
#####Adding an attr value
dn: cn=admin,ou=groups,dc=mybrocade,dc=com
changetype: modify
add: member
member: cn=test1,cn=Users,dc=mybrocade,dc=com
```

2. Enter the following **ldapmodify** command, where test1.ldif is the name of the file you edited in [step 1](#).

```
> ldapmodify -D cn=admin,dc=mybrocade,dc=com -x -w secret -f test1.ldif
```

Example to delete a group member

1. Create or edit a .ldif file with contents similar to the following.

```
#####Deleting an attr value
#dn: cn=admin,ou=groups,dc=mybrocade,dc=com
#changetype: modify
#delete: member
#member: cn=Sachin,cn=Users,dc=mybrocade,dc=com
```

2. Enter the following **ldapmodify** command, where test2.ldif is the name of the file you edited in [step 1](#).

```
> ldapmodify -D cn=admin,dc=mybrocade,dc=com -x -w secret -f test2.ldif
```

Example to change the value of an attribute

1. Create or edit a .ldif file with contents similar to the following.

```
#####Replacing an attribute value
dn: cn=test,cn=Users,dc=mybrocade,dc=com
changetype: modify
replace: uid
uid: test
```

2. Enter the following **ldapmodify** command, where test3.ldif is the name of the file you edited in [step 1](#).

```
> ldapmodify -D cn=admin,dc=mybrocade,dc=com -x -w secret -f test3.ldif
```

The value of the uid attribute is changed to “test”.

Adding an Admin Domain or Virtual Fabric list

If your network uses Admin Domains, you can specify a list of Admin Domain numbers to which the user has access.

Use the **brcdAdvfData** attribute to map a role to a Virtual Fabric or Admin Domain. 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 [step 1](#).

```
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 Administrative Domains, enter the value of each Admin Domain separated by an underscore (_). Each number represents the number of the Admin Domain to which the user has access. The first such number represents the user’s Home domain.

Example for adding Admin Domains

In the following example, the user is granted access to Admin Domains 0, 10, and 200. Admin Domain 0 is the domain that the user initially logs in to.

```
brcdAdvfData: adlist_0_10_200_endAd
```

- 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.

```
> 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, if Admin Domains or Virtual Fabrics are in use, provide lists of Admin Domains or Virtual Fabrics to which the user should have access. For details, refer to [“The tac_plus.cfg file”](#) on page 190.

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 TACAS+ 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 192 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”](#) for details.
3. Run the `tac_plus` daemon to start and enable the TACACS+ service on the server.

```
> 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 25](#) lists and defines attributes specific to Brocade.

TABLE 25 Brocade custom TACACS+ attributes

Attribute	Purpose
<code>brcd-role</code>	Role assigned to the user account
<code>brcd-AV-Pair1</code>	The Admin Domain or Virtual Fabric member list, and chassis role
<code>brcd-AV-Pair2</code>	The Admin Domain or 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 152 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$chap$pswrđ"
    pap = cleartext "pap-password"
    service = exec {
```

```

        brcd-role = securityAdmin;
    }
}

```

Configuring Admin Domain lists

If your network uses Admin Domains, you should create Admin Domain lists for each user to identify the Admin Domains to which the user 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 Admin Domains:

- *HomeAD* is the designated home Admin Domain for the account. The valid range of values is from 0 through 255. The first valid HomeAD key-value pair is accepted by the switch, and any additional HomeAD key-value pairs are ignored.
- *ADList* is a comma-separated list of Administrative Domain numbers of which this account is a member. Valid numbers range from 0 through 255. A - between two numbers specifies a range.

The following example sets the home Admin Domain for the fosuser4 account to 255 and grants the account access to Admin Domains 1, 2, 3, and 200 through 255.

```

user = fosuser4 {
    pap = clear "password"
    chap = clear "password"
    password = clear "password"
    service = shell {
        set brcd-role = securityAdmin
        set brcd-AV-Pair1 = "homeAD=255;ADList=1,2,3";
        set brcd-AV-Pair2 = "ADList=200-255";
    }
}

```

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.

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 dual-CP Backbones (Brocade DCX, DCX-4S, DCX 8510-4, and DCX 8510-8 devices), 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 enabling RADIUS

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

Example enabling LDAP

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

Example enabling TACACS+

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

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+ service is not configured, only the parameter heading line is displayed. Parameters include:

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. 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 [Table 22](#) on page 169.

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

In this chapter

- Security protocols 195
- Secure Copy 196
- Secure Shell protocol 197
- Secure Sockets Layer protocol 200
- Simple Network Management Protocol 206
- Telnet protocol 226
- Listener applications 228
- Ports and applications used by switches 229

Security protocols

Security protocols provide endpoint authentication and communications privacy using cryptography. Typically, you are authenticated to the switch while the switch remains unauthenticated to you. This means that you can be sure with what you are communicating. The next level of security, in which both ends of the conversation are sure with whom they are communicating, is known as two-factor authentication. Two-factor authentication requires public key infrastructure (PKI) deployment to clients.

Fabric OS supports the secure protocols shown in [Table 26](#).

TABLE 26 Secure protocol support

Protocol	Description
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 with TLS uses a certificate authority (CA). By default, LDAP traffic is transmitted unsecured. With the import of signed certificates, you can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL) / Transport Layer Security (TLS) technology in conjunction with LDAP.
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.
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, v2, and v3.

TABLE 26 Secure protocol support (Continued)

Protocol	Description
SSH	Secure Shell (SSH) is a network protocol that allows data to be exchanged over a secure channel between two computers. Encryption provides confidentiality and integrity of data. SSH uses public-key cryptography to authenticate the remote computer and allow the remote computer 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 SSLv3, 128-bit encryption by default.

[Table 27](#) describes additional software or certificates that you must obtain to deploy secure protocols.

TABLE 27 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
SCP	SSH daemon, SCP server	None
SNMPv1, SNMPv2, SNMPv3	None	None

The security protocols are designed with the four main use cases described in [Table 28](#).

TABLE 28 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 Fabric OS versions 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 203.
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** or **yes** at the **cfgload attributes** prompt.
4. Enter **y** or **yes** 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 Blowfish-Cipher block chaining (CBC) 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 protocol v2.0 (ssh2) version 5.2p1. For more information on SSH, refer to the SSH IETF website:

<http://www.ietf.org/ids.by.wg/secsh.html>

You can also refer to *SSH, The Secure Shell: The Definitive Guide* by Daniel J. Barrett, Ph. D., Richard E. Silverman, and Robert G. Byrnes.

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 and DSA, 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. RSA 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.

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
```

Example of RSA/DSA key pair generation

```
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.

Example of adding 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: ouser
Password:
Public key is imported successfully.
```

4. Test the setup by logging in to the switch from a remote device, or by running a command remotely using SSH.

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
```

where 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
```

where the *username* variable is the name of the user who can perform SSH public key authentication, and who can import, export, and delete keys.

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
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: ~ouser/.ssh
Enter login name: ouser
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 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 the public keys you want to delete. Enter **all** to delete public keys for all users.

For more information on IP filter policies, refer to [Chapter 8, “Configuring Security Policies”](#).

Deleting private keys on the switch

Use the following procedure to delete private keys 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 [Chapter 8, “Configuring Security Policies”](#).

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 *Web Tools Administrator's Guide*.

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 may need to add a 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 following items on each switch:
 - a. A public and private key by using the **secCertUtil genkey** command.
 - b. A certificate signing request (CSR) by using the **secCertUtil gencsr** command.
3. Store the CSR on a file server by using the **secCertUtil 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 the certificate. 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.

Generating a public-private key pair

Use the following procedure to generate a public-private key pair.

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 **secCertUtil genkey** command to generate a public/private key pair.
The system reports that this process will disable secure protocols, delete any existing CSR, and delete any existing certificates.
3. Respond to the prompts to continue and select the key size.

Example of generating a key

```
Continue (yes, y, no, n): [no] y
Select key size [1024 or 2048]: 1024
Generating new rsa public/private key pair
Done.
```

Generating and storing a Certificate Signing Request

After generating a public/private key pair, you must generate and store a certificate signing request (CSR).

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **secCertUtil gencsr** command.
3. Enter the requested information.

Example of generating a CSR

```
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):Eng
Common Name (Fully qualified Domain Name, or IP address): 192.1.2.3
Generating CSR, file name is: 192.1.2.3.csr
Done.
```

Your CA may 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/director. The IP address or FQDN is the switch where the certificate gets installed.

4. Enter the **secCertUtil export** command to store the CSR.
5. Enter the requested information. You can use either FTP or SCP.

Example of exporting a CSR

```
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.
```

If you are set up for Secure Copy Protocol (SCP), you can select it; otherwise, select FTP. Enter the IP address of the switch on which you generated the CSR. Enter the remote directory name of the FTP server to which the CSR is to be sent. Enter your account name and password on the server.

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 [Table 29](#).

TABLE 29 SSL certificate files

Certificate file	Description
<i>name.pem</i>	The switch certificate.
<i>nameRoot.pem</i>	The root certificate. Typically, this certificate is already installed in the browser, but if not, you must install it.
<i>nameCA.pem</i>	The CA 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 202.
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. Through a Telnet window, connect to the switch and log in using an account with admin permissions.
4. Enter the **secCertUtil showcsr** 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

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 **secCertUtil import** command.
3. Select a protocol, enter the IP address of the host on which the switch certificate is saved, and enter your login name and password.

Example of installing a switch certificate in interactive mode

```
switch:admin> seccertutil import -config swcert -enable https
Select protocol [ftp or scp]: ftp
Enter IP address: 192.10.11.12
Enter remote directory: path_to_remote_directory
Enter certificate name (must have ".crt", ".cer", \
".pem" or ".psk" suffix): 192.1.2.3.crt
Enter Login Name: your_account
Enter Password: *****
Success: imported certificate [192.1.2.3.crt].
```

Example of installing a switch certificate in noninteractive mode

```
switch:admin> seccertutil import -config swcert -enable https \
-protocol ftp -ipaddr 192.10.11.12 -remotedir path_to_remote_directory \
-certname 192.1.2.3.crt -login your_account -password passwd
Success: imported certificate [192.1.2.3.crt].
Certificate file in configuration has been updated.
Secure http has been enabled.
```

Example of installing a common certificate in non-interactive mode

```
switch:admin> seccertutil import -commonswcert -config swcert -enable https
-protocol scp -ipaddr 192.10.11.12 -remotedir path_to_remote_directory -login
cert -certname 192.1.2.3.pem
```

Important Notes

- 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.

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** tab 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*.
6. 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 200.

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 Sun Microsystems website (www.sun.com).

Installing a root certificate to the Java plugin

Use the following procedure to install a root certificate to the Java plugin.

1. 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
```

2. Open a Command Prompt window and change the directory to the Java plugin bin directory.
3. 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 (MIB)

SNMP Agents in the managed devices store the data about these devices in a database called 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), as shown in [Figure 13](#). The agent accesses information about a device and makes it available to an SNMP network management station.




FIGURE 13 SNMP structure

When active, the management station can **get** information or **set** information when it queries an agent. SNMP commands, such as **get**, **set**, **getnext**, and **getResponse**, are sent from the management station, and the agent replies once the value is obtained or modified ([Figure 14](#)). 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.

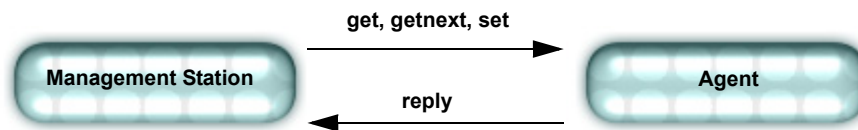


FIGURE 14 SNMP query

The management station can also receive *traps*, unsolicited messages from the switch agent if an unusual event occurs ([Figure 15](#)). For more information, refer to “[Traps](#)” on page 209.



FIGURE 15 SNMP trap

The agent can receive queries from one or more management stations and can send traps to up to six management stations.

Understanding MIBs

The management information base (MIB) is a database of monitored and managed information on a device, in this case a Brocade switch. The MIB structure can be represented by a tree hierarchy. The root splits into three main branches: International Organization for Standardization (ISO), Consultative Committee for International Telegraph and Telephone (CCITT), and joint ISO/CCITT. These branches have short text strings and integers (OIDs) to identify them. Text strings describe *object names*, while integers allow software to create compact, encoded representations of the names.

Each MIB variable is assigned an object identifier (OID). The OID is the sequence of numeric labels on the nodes along a path from the root to the object. For example, as shown in Figure 16, the Brocade SW.MIB OID is:

```
1.3.6.1.4.1.1588.2.1.1.1
```

The corresponding name is:

```
iso.org.dod.internet.private.enterprise.bcsi.commDev.fibreChannel.fcSwitch.sw
```

The other branches are part of the standard MIBs, and the portions relevant to configuring SNMP on a Brocade switch are referenced in the remainder of this reference.




FIGURE 16 Brocade MIB tree location

Access to MIB variables

Use a MIB browser to access the MIB variables: all MIB browsers perform queries and load MIBs.

Once loaded, the MAX-ACCESS provides access levels between the agent and management station. The access levels are as follows:

- not accessible
You cannot read or write to this variable.
- read create
Specifies a tabular object that can be read, modified, or created as a new row in a table.
- read only - *Public*
You can only monitor information.
- read-write - *Private*
You can read or modify this variable.
- accessible-to-notify
You can read this information only through traps.

SNMP support

In addition to the standard MIBs that Brocade devices support, these devices also support Brocade-specific MIBs. Since different vendors vary the information in their private enterprise MIBs, it is necessary to verify their information. The Fibre Channel MIB standards dictate that certain information be included in all MIBs: it is the vendors' responsibility to follow the standards. The standards are as follows:

- FibreAlliance (FA) MIB: Brocade supports v4.4.0 and later releases.
- Fabric Element (FE) MIB: accepted by the Internet Engineering Task Force (IETF).

Brocade supports FE_RFC2837.mib under the MIB-II branch in Fabric OS v7.1.0, v7.0.0, v6.4.1_fcoe, v6.4.0, v6.3.0, v6.2.0, v6.1.2_CEE, v6.1.0, and v6.0.0. The latest version of the FE MIB references the FRAMEWORK.MIB and, based on the MIB browser, it is necessary to load this MIB before the FE.MIB. For more information, refer to [“Loading Brocade MIBs”](#) on page 212.

Traps

An unsolicited message that comes to the management station from the SNMP agent on the device is called a *trap*. Brocade switches send traps out on UDP port 162 and to any configured port. In order to receive traps, the management station IP address and severity level must be configured on the switch. You can define a different message severity level for each recipient so that some recipients receive all trap messages and others receive only the most critical.

NOTE

Due to design limitation, IP address validation cannot be done for trap recipients.

There are two main MIB trap choices:

- FibreAlliance MIB trap - Associated with the Fibre Alliance MIB (FA-MIB), this MIB manages SAN switches and devices from any company that complies with Fibre Alliance specifications.
- Brocade-specific MIB trap - Associated with the Brocade-specific Brocade MIB (SW-MIB), manages Brocade switches only.

There is some overlap in the functionality of these MIBs. If you enable both SW-MIB and FA-MIB traps, you could receive duplicate messages for the switch events that trigger the trap.

You can also use these additional MIBs and their associated traps: HA-MIB; FICON-MIB; and SWEXTTRAP. In Fabric OS v6.4.0, you can use the **snmpConfig-set mibCapability** command to enable or disable all the MIBs.

An event trap (swEventTrap, connUnitEventTrap, or swFabricWatchTrap) is basically an error message (**errShow** output) that is SNMP-formatted and delivered.

FA traps

Consider enabling the FA traps if you want to use SNMP to monitor multiple connectivity units, including Brocade switches.

The **switchStatusPolicySet** command determines the FA-TRAP switch status-related outputs:

- connUnitStatusChange
This trap is generated by Fabric watch such that only the swUnitsStatusChange is controlled by **switchStatusPolicySet** command.
- connUnitSensorStatusChange
This trap is generated by any sensor event.
- connUnitPortStatusChange
This trap sends the instance of connUnitPortName as part of the trap; the instance string is NULL, if the port name is not defined for the specified port.
- connUnitEventTrap
All the external traps gets converted into swEventTrap except for AN-1006, AUTH-3001 to AUTH-3008, FW-3001, SEC-3001 to SEC-3034, and SEC-3044 to SEC-3048 RASlog messages.

Events in the Error Log of a severity at or above the configured threshold will generate SNMP traps.

The Fibre Alliance Trap (FA-TRAP) can be configured to send traps using the **snmpConfig** command. For more information on this command, refer to the *Fabric OS Command Reference*.

HA traps

Consider enabling these traps to monitor field-replaceable unit (FRU) status and control processor (CP) status when you have a Brocade director in your environment:

- fruStatusChanged
This trap is generated by a FRU status change, such as a switch reboot or disabling or enabling a FRU component such as fandisable, fanenable and so on.
- cpStatusChanged
This trap is generated by a change in the status of a CP, including a reboot or firmware download.
- fruHistoryTrap
This trap is generated when a FRU is added or removed. It is not generated when standby CP is removed.

The high availability trap (HA-TRAP) can be configured to send traps using the **snmpConfig** command. For more information on this command, refer to the *Fabric OS Command Reference*.

SW traps

There are fourteen specific traps defined in Brocade SW-TRAP.

- swfault (no longer supported)
- swSensorScn (no longer supported)
- swFCPortScn
This trap is generated by a port state change.
- swEventTrap
This trap is generated by any switch event reported to the system error log.
The desired severity level is introduced to filter a swEvent trap based on the severity level.
- swFabricWatchTrap
This trap is generated when any Fabric Watch threshold is reached.
The desired severity level is introduced to filter a swFabricWatchTrap based on the severity level.
- swTrackChangesTrap
This trap is generated by a login or a logout.
- swIPv6ChangeTrap
This trap is generated when an IPv6 address status change event occurs. It is generated only when IPv6 stateless state changes to the deprecation state and not for address change notification.
- swPmgrEventTrap
This trap is generated when any partition manager change happens.
- swFabricReconfigTrap
The trap to be sent for tracking fabric reconfiguration.
- swFabricSegmentTrap
The trap to be sent for tracking segmentation.
- swExtTrap
The trap adds the SSN binding to the traps if it is enabled.
- swStateChangeTrap
This trap is sent when the switch changes its state to online or offline.
- swPortMoveTrap
This trap is sent when the virtual ports are moved from one switch to another.
- swBrcdGenericTrap
This trap is sent for one of the events, such as fabric change, device change, FAPWWN change, and FDMI events. This trap is for Brocade use.

- swDeviceStatusTrap

This trap is sent whenever a device logs in or logs out.

The Brocade trap (SW-TRAP) can be configured to send traps using the `snmpConfig` command.

FICON traps

- linkRNIDDeviceRegistration
A device registered with the switch.
- linkRNIDDeviceDeRegistration
A device de-registered with the switch.
- linkLIRListenerAdded
A listener for link failure incident is added.
- linkLIRListenerRemoved
A listener for link failure incident is removed.
- linkRLIRFailureIncident
A link failure incident has occurred.

IF traps

- linkDown
A linkDown trap signifies that the SNMPv2 entity, acting in an agent role has detected that the ifOperStatus object for one of its communication links is about to transition into the down state.
- linkUp
A linkUp trap signifies that the SNMPv2 entity, acting in an agent role has detected that the ifOperStatus object for one of its communication links has transitioned out of the down state."

BD traps

- bdTrap
Traps to be sent for bottleneck detection.
- bdClearTrap
Traps to be sent for bottleneck clearance.

Loading Brocade MIBs

The Brocade MIB is a set of variables that are private extensions to the Internet standard MIB-II. The Brocade agents support many other Internet-standard MIBs. These standard MIBs are defined in RFC publications. To find specific MIB information, examine the Brocade proprietary MIB structure and the standard RFC MIBs supported by Brocade.

Brocade MIB files

The Brocade MIB files are as follows:

- bd.mib
- bcCustomOperation.mib
- BRCD_REG.mib
- BRCD_TC.mib
- brcdfcip.mib
- CPQ_HOST.mib
- CPQ_RACK.mib
- FA.mib
- faext.mib
- FICON.mib
- fod.mib
- HA.mib
- IbmBladeCenter.mib
- SW.mib

Standard MIBs

Distribution of standard MIBs has been stopped from Fabric OS v6.4.0. Download the following MIBs from the <http://www.oidview.com/> website:

- BRIDGE-MIB
- ENTITY-MIB
- FC-MGMT-MIB
- FCIP-MGMT-MIB
- FIBRE-CHANNEL-FE-MIB
- IANAifType-MIB
- IEEE 802.1x PAE MIB
- IEEE 802.3 LAG MIB
- IF-MIB
- INET-ADDRESS-MIB
- IP MIB
- LLDP MIB
- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB
- P-BRIDGE MIB
- Q-BRIDGE MIB
- RFC1155-SMI
- RFC1158-MIB
- RFC-1212

- RFC1213-MIB
- RFC-1215
- RMON-MIB
- RSTP-MIB
- SNMP-COMMUNITY-MIB
- SNMP-FRAMEWORK-MIB
- SNMPv2-CONF
- SNMPv2-MIB
- SNMPv2-PARTY-MIB
- SNMPv2-SMI-MIB
- SNMPv2-TC
- SNMP-VIEW-BASED-ACM-MIB
- SNMP-USER-BASED-SM-MIB
- SNMP-TARGET-MIB

MIB loading order

Many MIBs use definitions that are defined in other MIBs. These definitions are listed in the IMPORTS section near the top of the MIB. When loading the Brocade MIBs, refer to [Table 30](#) to ensure any MIB dependencies are loading in the correct order.

NOTE

Before loading Brocade MIB files, ensure that you have the correct version of SNMP for your Fabric OS version. All versions of Fabric OS support SNMPv1 and SNMPv3. SNMPv2 is not supported on Fabric OS v6.1.2_CEE and later versions.

TABLE 30 Brocade SNMP MIB dependencies

MIB Name	Dependencies
BRCD_REG.mib	RFC1155-SMI
BRCD_TC.mib	Brocade-REG-MIB SNMPv2-TC SNMPv2-SMI
FC-MGMT-MIB	SNMPv2-SMI SNMPv2-CONF SNMPv2-MIB IANAifType-MIB SNMPv2-TC IF-MIB SNMP-FRAMEWORK-MIB
FA.mib	RFC1155-SMI RFC1158-MIB RFC-1212 RFC1213-MIB RFC-1215

TABLE 30 Brocade SNMP MIB dependencies

MIB Name	Dependencies
FIBRE-CHANNEL-FE-MIB	SNMPv2-SMI SNMPv2-TC SNMP-FRAMEWORK-MIB SNMPv2-CONF
FCIP-MGMT-MIB	SNMPv2-SMI SNMPv2-TC INET-ADDRESS-MIB FC-MGMT-MIB IF-MIB SNMPv2-CONF SNMP-FRAMEWORK-MIB
ENTITY-MIB	SNMPv2-SMI SNMPv2-TC SNMPv2-CONF SNMP-FRAMEWORK-MIB
SW.mib	SNMPv2-TC SNMPv2-SMI Brocade-TC Brocade-REG-MIB FCMGMT-MIB
bd.mib	SNMPv2-TC SNMPv2-SMI Brocade-TC Brocade-REG-MIB SW-MIB
brcdfcip.mib	SNMPv2-SMI Brocade-REG-MIB SNMPv2-TC INET-ADDRESS-MIB IF-MIB SNMPv2-CONF
faext.mib	SNMPv2-TC SNMPv2-SMI SW-MIB FCMGMT-MIB
FICON.mib	SNMPv2-SMI SNMPv2-TC Brocade-REG-MIB
HA.mib	SNMPv2-SMI Brocade-REG-MIB SW-MIB ENTITY-MIB SNMPv2-TC

Access Gateway and Brocade MIBs

Table 31 shows the MIBs supported by Brocade Access Gateway.

TABLE 31 Access Gateway MIB support

MIB name	Description
MIB-2	Supported in v5.2.1 and later releases.
Entity-MIB	Supported.
HA-MIB	Supported.
SW-MIB	Disabled in Access Gateway because the conventions are specific to fabric switches. In Fabric OS v6.4.0, <code>swConnUnitPortExtensionTable</code> is supported in Access Gateway mode. In Fabric OS v7.0.0, SNMP allows you to access the following tables to support the Advanced Performance Monitoring feature on Access Gateway, even if the SW-MIB is disabled: <ul style="list-style-type: none"> • <code>swBlmPerfEEMntTable</code> • <code>swBlmPerfFitMntTable</code>
FA-MIB	The <code>connUnitSnsTable</code> is not supported because a switch in Access Gateway does support name server services.
CPQ-Rack MIB	Supported on embedded switches only.
IF-MIB	Supported.
BD-MIB	Supported for F-ports.
FA-Ext	Supported.
SNMPv2 MIB	Supported.

Firmware upgrades and enabled traps

You can turn on and off traps individually within a trap group. By default the individual traps are turned off even if the corresponding trap group was enabled before upgrading. You must use the `snmpconfig` command to turn on the individual traps within each trap group.

Support for Administrative Domains

Administrative Domains are supported in Fabric OS v5.3.0 and later releases. An Administrative Domain (AD) is a domain within a fabric. Administrative domains can be used to limit administrator access within a fabric, and to provide service providers with a means to assign portions of a fabric to individual consumers. An AD may contain switches, devices, and ports. An AD may also limit access to a configured set of users.

The following example shows how the `AD:xxx` field is used in the `snmpwalk` command. This command is executed on the host and it walks the entire MIB tree specified (.1).

```
switch# snmpwalk -u admin -v 3 -n AD:4 10.168.176.181.1
```

Support for Role-Based Access Control

Role-Based Access Control (RBAC) is supported in Fabric OS v5.3.0 and later releases. RBAC applies a fixed set of roles that address the access control needs of a majority of customers. Each role is a set of permissions that can be applied to a user that controls the kinds of jobs and tasks the user can perform on a fabric or fabric element.

Support for IPv6 addressing

IPv6 addressing is supported in Fabric OS v5.3.0 and later releases.

Support for Virtual Fabric

Virtual Fabric is supported in Fabric OS v6.2.0 and later releases.

When an SNMPv3 request arrives with a particular user name, it executes in the home Virtual Fabric. From the SNMP manager, all SNMPv3 requests must have a home Virtual Fabric that is specified in the *contextName* field. When the home Virtual Fabric is specified, it will be converted to the corresponding switch ID and the home Virtual Fabric will be set. If the user does not have permission for the specified home Virtual Fabric, this request fails with an error code of *noAccess*.

For an SNMPv3 user to have a home Virtual Fabric, a list of allowed Virtual Fabrics, an RBAC role, and the name of the SNMPv3 user should match that of the Fabric OS user in the local switch database. SNMPv3 users whose names do not match with any of the existing Fabric OS local users have a default RBAC role of admin with the SNMPv3 user access control of read/write. Their SNMPv3 user logs in with an access control of read-only. Both user types will have the default switch as their home Virtual Fabrics.

The *contextName* field should have the format “VF:xxx”, where xxx is the actual VF_ID, for example “VF:1”. If the *contextName* field is empty, then the home Virtual Fabric of the local Fabric OS user with the same name is used. As Virtual Fabrics and Admin Domains are mutually exclusive, this field is considered as Virtual Fabrics context when Virtual Fabrics is enabled. You cannot specify chassis context in the *contextName* field.

The following example shows how the VF:xxx field is used in the **snmpwalk** command. This command is executed on the host and it walks the entire MIB tree specified (.1).

```
switch# snmpwalk -u admin -v 3 -n VF:4 10.168.176.181.1
```

Filtering ports

Each port can belong to only one Virtual Fabric at any time. An SNMP request coming to one Virtual Fabric can only view the port information of the ports belonging to that Virtual Fabric. All port attributes are filtered to allow SNMP to obtain the port information only from within the current Virtual Fabrics context.

Switch and chassis context enforcement

All attributes are classified into one of two categories:

- Chassis-level attributes
- Switch-level attributes

Attributes that are specific to each logical switch belong to the switch category. These attributes are available in the Virtual Fabrics context and not available in the Chassis context.

Attributes that are common across the logical switches belong to the chassis level. These attributes are accessible to users having the chassis-role permission. When a chassis table is queried, the context is set to chassis context, if the user has the chassis-role permission. The context is switched back to the original context after the operation is performed.

Configuring SNMP using CLI

For information about Fabric OS commands for configuring SNMP, refer to the *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
Select SNMP Security Level
(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 =
sxNo Access): (0..3) [0]
Select SNMP SET Security Level
(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 =
No Access): (0..3) [0]
```

Table 32 shows the security level options.

TABLE 32 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] (authPriv)	SNMPv1	Not allowed.	Not Sent.
	SNMPv3	Only SNMPv3 users with authPriv privilege are allowed.	Sent only for authPriv users.
No Access [3]	SNMPv1	Not allowed.	Not Sent.
	SNMPv3		

Configuring SNMPv3 user/traps

The following examples list how to configure SNMPv3 users/traps.

1. Create a user on the switch in non-VF Context using CLI userconfig, 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.
switch:admin>
```

Create a user on the switch in VF Context using CLI userconfig, 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.
switch:admin>
```

2. Create the SNMPv3 user.

```

switch:root> snmpconfig --set snmpv3

SNMP Informs Enabled (true, t, false, f): [false] t

SNMPv3 user configuration(snmp user not configured in FOS user database will
have physical AD and admin role as the default):
User (rw): [snmpadmin1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]
Engine ID: [0:0:0:0:0:0:0:0] 80:00:05:23:01:0A:23:34:21
User (rw): [snmpadmin2]
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)]: (1..3) [2]
New Priv Passwd:
Verify Priv Passwd:
Engine ID: [0:0:0:0:0:0:0:0] 80:00:05:23:01:0A:23:34:1B
User (rw): [snmpadmin3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]
Engine ID: [0:0:0:0:0:0:0:0]
User (ro): [snmpuser1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]
Engine ID: [0:0:0:0:0:0:0:0]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]
Engine ID: [0:0:0:0:0:0:0:0]
User (ro): [snmpuser3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]
Engine ID: [0:0:0:0:0:0:0:0]

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] 4
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] 4
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:root> snmpconfig --show snmpv3

SNMP Informs = 1 (ON)

SNMPv3 USM configuration:
User 1 (rw): snmpadmin1
    Auth Protocol: noAuth

```

```

        Priv Protocol: noPriv
        Engine ID: 80:00:05:23:01:0a:23:34:21
User 2 (rw): snmpadmin2
        Auth Protocol: MD5
        Priv Protocol: DES
        Engine ID: 80:00:05:23:01:0a:23:34:1b
User 3 (rw): snmpadmin3
        Auth Protocol: noAuth
        Priv Protocol: noPriv
        Engine ID: 00:00:00:00:00:00:00:00:00
User 4 (ro): snmpuser1
        Auth Protocol: noAuth
        Priv Protocol: noPriv
        Engine ID: 00:00:00:00:00:00:00:00:00
User 5 (ro): snmpuser2
        Auth Protocol: noAuth
        Priv Protocol: noPriv
        Engine ID: 00:00:00:00:00:00:00:00:00
User 6 (ro): snmpuser3
        Auth Protocol: noAuth
        Priv Protocol: noPriv
        Engine ID: 00:00:00:00:00:00:00:00:00

```

```

SNMPv3 Trap configuration:
Trap Entry 1:      10.35.52.33
    Trap Port: 162
    Trap User: snmpadmin1
    Trap recipient Severity level: 4
Trap Entry 2:      10.35.52.27
    Trap Port: 162
    Trap User: snmpadmin2
    Trap recipient Severity level: 4
Trap Entry 3:      No trap recipient configured yet
Trap Entry 4:      No trap recipient configured yet
Trap Entry 5:      No trap recipient configured yet
Trap Entry 6:      No trap recipient configured yet

```

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): [false]
SNMPv3 user configuration(snmp user not configured in FOS user database will
have physical AD and admin role as the default):
User (rw): [snmpadmin1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]
User (rw): [snmpadmin2]
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)]: (1..3) [2]
User (rw): [snmpadmin3]
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)]: (1..3) [2]
New Priv Passwd:

```

```

Verify Priv Passwd:
User (ro): [snmpuser1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]
User (ro): [snmpuser3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)]: (1..3) [2]

SNMPv3 trap recipient configuration:
Trap Recipient's IP address : [0.0.0.0] 172.26.4.102
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [0] 1
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [0.0.0.0] fe80::224:1dff:fef6:3f98
UserIndex: (1..6) [2]
Trap recipient Severity level : (0..5) [0] 2
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [0.0.0.0]
UserIndex: (1..6) [3]
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]
Committing configuration.....done.
DCX_128:FD128:admin>

switch:admin> snmpconfig --show snmpv3

SNMP Informs = 0 (OFF)

SNMPv3 USM configuration:
User 1 (rw): snmpadmin1
    Auth Protocol: noAuth
    Priv Protocol: noPriv
User 2 (rw): snmpadmin2
    Auth Protocol: MD5
    Priv Protocol: noPriv
User 3 (rw): snmpadmin3
    Auth Protocol: MD5
    Priv Protocol: DES
User 4 (ro): snmpuser1
    Auth Protocol: noAuth
    Priv Protocol: noPriv
User 5 (ro): snmpuser2
    Auth Protocol: noAuth
    Priv Protocol: noPriv
User 6 (ro): snmpuser3
    Auth Protocol: noAuth
    Priv Protocol: noPriv

SNMPv3 Trap configuration:
Trap Entry 1:      172.26.4.102
    Trap Port: 162
    Trap User: snmpadmin1
    Trap recipient Severity level: 1
Trap Entry 2:      fe80::224:1dff:fef6:3f98

```

7 Simple Network Management Protocol

```
Trap Port: 162
Trap User: snmpadmin2
Trap recipient Severity level: 2
Trap Entry 3:      HCL0389U.corp.brocade.com
Trap Port: 162
Trap User: snmpadmin3
Trap recipient Severity level: 5
Trap Entry 4:      No trap recipient configured yet
Trap Entry 5:      No trap recipient configured yet
Trap Entry 6:      No trap recipient configured yet
```

To display the traps and MIBs supported in Fabric OS:

```
switch:root> snmpTraps --show
```

#	Mib Name	Supported Traps
001	SW-MIB	sw-track-changes-trap sw-fabric-watch-trap sw-fc-port-scn ip-v6-change-trap sw-pmgr-event-trap sw-event-trap sw-fabric-reconfig-trap sw-fabric-segment-trap sw-state-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-sensor-status-change conn-unit-port-status-change conn-unit-event-trap
004	RFC1157	cold-restart-trap warm-restart-trap if-link-up-trap if-link-down-trap
005	HA-MIB	snmp-authetication-trap fru-status-change-trap fru-history-trap cp-status-change-trap
006	BD-MIB	bd-trap bd-clear-trap

To send all traps to the configured recipients:

```
switch:root> snmpTraps --send
Number of traps sent : 30
```

To send all traps to the recipient 10.35.52.33:

```
switch:root> 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:root> 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:root> 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
```

Example of 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.
```

Example of mibCapability configuration:

To enable the swFabricWatchTrap non-interactively:

```
switch:admin> snmpconfig --enable mibCapability -mib_name SW-MIB -trap_name
swFabricWatchTrap
Operation succeeded
```

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
swFabricWatchTrap: NO
DesiredSeverity:None
swTrackChangesTrap: NO
swIPv6ChangeTrap: NO
swPmgrEventTrap: NO
swFabricReconfigTrap: NO
swFabricSegmentTrap: NO
swExtTrap: NO
[...]
```

To enable the SW-MIB 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
swFabricWatchTrap: YES
DesiredSeverity:None
swTrackChangesTrap: NO
swIPv6ChangeTrap: NO
swPmgrEventTrap: NO
swFabricReconfigTrap: NO
swFabricSegmentTrap: NO
swExtTrap: NO
[...]
```

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
swFabricWatchTrap: YES
DesiredSeverity:None
swTrackChangesTrap: YES
swIPv6ChangeTrap: YES
swPmgrEventTrap: YES
swFabricReconfigTrap: Yes
swFabricSegmentTrap: Yes
swExtTrap: Yes
[...]
```

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
ISCSI-MIB: YES
IF-MIB: YES
BD-MIB: YES
SW-TRAP: YES
    swFault: YES
    swSensorScn: YES
    swFCPortScn: YES
    swEventTrap: YES
        DesiredSeverity:None
    swFabricWatchTrap: YES
        DesiredSeverity:None
    swTrackChangesTrap: YES
    swIPv6ChangeTrap: YES
    swPmgrEventTrap: YES
    swFabricReconfigTrap: YES
    swFabricSegmentTrap: YES
    swExtTrap: YES
FA-TRAP: YES
    connUnitStatusChange: YES
    connUnitDeletedTrap: YES
    connUnitEventTrap: YES
    connUnitSensorStatusChange: 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
ISCSI-TRAP: YES
    iscsiTgtLoginFailure: YES
    iscsiIntrLoginFailure: YES
    iscsiInstSessionFailure: YES
IF-TRAP: YES
    linkDown: YES
    linkUp: YES
BD-TRAP: YES
    bdTrap: YES
    bdClearTrap: YES

```

To set the system group:

```
DCX_128:FID128:admin> snmpconfig --set systemgroup
```

```

Example of systemGroup configuration (default)
switch:admin> snmpconfig --default systemGroup
*****

```

```

This command will reset the agent's system group configuration back to
factory default
*****

```

```
sysDescr = Fibre Channel Switch
sysLocation = End User Premise
sysContact = Field Support
authTraps = 0 (OFF)
*****
Are you sure? (yes, y, no, n): [no] y
```

3. Set the security level.

```
switch:admin> snmpconfig --set secLevel
Select SNMP GET Security Level
(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 =
No Access): (0..3) [0] 2
Select SNMP SET Security Level
(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 =
No Access): (2..3) [2] 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:

```
DCX_128:FID128:admin> snmpconfig --default seclevel
GET security level = 0, SET level = 0
SNMP GET Security Level: No security
SNMP SET Security Level: No security
SNMP GET Security Level will be set to 'No Security'
SNMP SET Security Level will be set to 'No Security'
Do you want to continue? (yes, y, no, n): [no] y
DCX_128:FID128:admin>

DCX_128:FID128:admin> snmpconfig --show seclevel
GET security level = 0, SET level = 0
SNMP GET Security Level: No security
SNMP SET Security Level: No security
DCX_128:FID128:admin
```

4. 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.)

NOTE

SNMPv3 supports AES-128 and DES protocols.

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 253.

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 typing the **ipFilter --clone** command.

```
switch:admin> ipfilter --clone BlockTelnet -from default_ipv4
```

3. Save the new policy by typing the **ipFilter --save** command.

```
switch:admin> ipfilter --save BlockTelnet
```

4. Verify the new policy exists by typing the **ipFilter --show** command.

```
switch:admin> ipfilter --show
```

5. Add a rule to the policy, by typing 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 259 for more information on deleting IP filter rules.

6. Save the new IP filter policy by typing the **ipfilter --save** command.

7. Verify the new policy is correct by typing the **ipFilter --show** command.

8. Activate the new IP filter policy by typing 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: 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

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 259 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.

[Table 33](#) lists the listener applications that Brocade switches either block or do not start. Note that RPC ports are blocked.

TABLE 33 Blocked listener applications

Listener application	Brocade DCX and 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

If you are using the FC-FC Routing Service, be aware that the **secModeEnable** command is not supported.

[Table 34](#) lists the defaults for accessing hosts, devices, switches, and zones.

TABLE 34 Access defaults

	Access default
Hosts	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	All devices can access the management server.
	Any device can connect to any FC port in the fabric.
Switch access	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

[Table 35](#) 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 35 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.
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.
512	TCP	exec	
513	TCP	login	
514	TCP	shell	

7 Ports and applications used by switches

Configuring Security Policies

In this chapter

• ACL policies overview	231
• ACL policy management	232
• FCS policies	235
• Device Connection Control policies	238
• SCC Policies	242
• Authentication policy for fabric elements	243
• IP Filter policy	253
• Policy database distribution	260
• Management interface security	266

ACL policies overview

Each supported Access Control List (ACL) policy listed below 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.

NOTE

Run all commands in this chapter by logging in to Administrative Domain (AD) 255 with the suggested permissions. If Administrative Domains have not been implemented, log in to ADO.

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.

Admin Domain considerations: ACL management can be done on AD255 and in ADO only if there are no user-defined Admin Domains. Both ADO (when no other user-defined Admin Domains exist) and AD255 provide an unfiltered view of the fabric.

Virtual Fabric considerations: ACL policies such as DCC, SCC, and FCS can be configured on each logical switch. The limit for 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 36](#).

TABLE 36 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_###	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 will have 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, see [Chapter 9, “Maintaining the Switch Configuration File”](#).

NOTE

All changes, including the creation of new policies, are saved and activated on the local switch only—unless the switch is in a fabric that has a strict or tolerant fabric-wide consistency policy for the ACL policy type for SCC or DCC. See [“Policy database distribution”](#) on page 260 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. Additionally, in a defined policy set, policies created in the same login session also appear but these policies are automatically deleted if the you log out without saving them.

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. Type the **secPolicyShow** command.

```
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. Type **secPolicyDelete** *“policy_name”*.
where *policy_name* is the name of the ACL policy.
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.

Example of adding to an ACL policy

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 37 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 38 shows the commands for switch operations for Primary FCS enforcement.

TABLE 38 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
	Any AD command that does not affect fabric-wide configuration

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 fabric domains share policies

Whether your intention is 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. To distribute the policies, enter the **distribute -p policy_list -d switch_list** command to either send the policies to intended domains, or enter the **distribute -p policy_list -d wild_card (*)** command to send the policies to all switches.

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. Type **secPolicyFCSMove**; then provide the current position of the switch in the list and the desired position at the prompts.

Alternatively, enter **secPolicyFCSMove [From, To]** command. *From* is the current position in the list of the FCS switch and *To* is the desired position in the list for this switch.

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. Type 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 261.

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 262.

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 39 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 *Access Gateway Administrator's 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 40 shows the possible DCC policy states.

TABLE 40 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

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 DCC 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 the **secPolicyCreate** “**DCC_POLICY_***nnn*” command.
 DCC_POLICY_ *nnn* is the name of the DCC policy; *nnn* is a string consisting of up to 19 alphanumeric or underscore characters to differentiate it from any other DCC policies.
 3. To save or activate the new policy, enter the appropriate command:
 - To save the policy, enter the **secPolicySave** command.
 - To save and activate the policy, enter the **secPolicyActivate** command.

If neither of these commands is entered, the changes are lost when the session is logged out.

Example 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 41](#) lists the behavior of the DCC policy with FA-PWWNs in the fabric when the DCC policy is created using lockdown support.

TABLE 41 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 lock down (uses FA-PWWN). DCC policy activation. 	FA-PWWN	Traffic will not be disrupted. ¹	Ports will be disabled for security violation. ²
<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 will not be disrupted.	Ports will 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 will not be disrupted.	Ports will come up without any security issues.

1. 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.

2. 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 login again.

[Table 42](#) 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 42 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 will not be disrupted.	Ports will 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 will not be disrupted.	Ports will come up without security issues.
<ul style="list-style-type: none"> DCC policy creation manually with physical PWWN, DCC policy activation. FA-PWWN has logged into the switch. 	Physical PWWN	Traffic will not be disrupted.	Ports will come up without any security issues.

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 43](#).

TABLE 43 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: 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:

- 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 [Chapter 11, “Managing Virtual Fabrics”](#).

Creating an SCC 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 "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 17](#) 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 250.

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 17 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

If Virtual Fabrics is enabled, all AUTH module parameters such as shared secrets, and shared switch and device policies, are logical switch-wide. That means 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. On logical switch creation, authentication takes 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 250.

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. Refer to [Chapter 16, “In-flight Encryption and Compression,”](#) for details.

Virtual Fabrics considerations

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 which 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 [Chapter 11, “Managing Virtual Fabrics”](#).

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 *Access Gateway Administrator's 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

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 *Access Gateway Administrator's 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. Below are the different 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**.

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

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. See [Chapter 16, “In-flight Encryption and Compression,”](#) for details about in-flight encryption.

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.

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 or at least 32 characters is recommended. See [Chapter 16, “In-flight Encryption and Compression”](#) for details about in-flight encryption.

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	DIId	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.

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 release 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, private key, passphrase and a CSR on each switch.
3. Store the CSR from each switch on a file server.
4. Obtain the certificates from the CA.

You can request a certificate from a CA through a Web browser. After you request a certificate, the CA either sends certificate files by e-mail (public) or gives access to them on a remote host (private). Typically, the CA provides the certificate files listed in [Table 44](#).

ATTENTION

Only the .pem file is supported for FCAP authentication.

TABLE 44 FCAP certificate files

Certificate file	Description
<i>name</i> CA.pem	The CA certificate. It must be installed on the remote and local switch to verify the validity of the switch certificate or switch validation fails.
<i>name</i> .pem	The switch certificate.

5. On each switch, install the CA certificate before installing switch certificate.
6. After the CA certificate is installed, install the switch certificate on each switch.
7. Update the switch database for peer switches to use third-party certificates.
8. Use the newly installed certificates by starting the authentication process.

Generating the key and CSR for FCAP

The public/private key and CSR has to be generated for the local and remote switches that will participate in the authentication. In FCAP, one command is used to generate the public/private key the CSR, and the passphrase.

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 **secCertUtil generate -fcapall -keysize** command on the local switch.

```
switch:admin> seccertutil generate -fcapall -keysize 1024
WARNING!!!

About to create FCAP:
ARE YOU SURE (yes, y, no, n): [no] y
Installing Private Key and Csr...
Switch key pair and CSR generated...
```

3. Repeat [step 2](#) on the remote switch.

Exporting the CSR for FCAP

You will need to export the CSR file created in “[Generating the key and CSR for FCAP](#)” section and send to a Certificate Authority (CA). The CA will in turn provide two files as outlined in “[FCAP configuration overview](#)” on page 251.

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 **secCertUtil export -fcapswcsr** command.

```
switch:admin> seccertutil export -fcapswcsr
Select protocol [ftp or scp]: scp
Enter IP address: 10.1.2.3
Enter remote directory: /myHome/jdoe/OPENSSL
Enter Login Name: jdoe
jdoe@10.1.2.3's password: <hidden text>
Success: exported FCAP CA certificate
```

Importing CA for FCAP

Once you receive the files back from the Certificate Authority, you will need to install or import them onto the local and remote switches.

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 **secCertUtil import -fcapcacert** command and verify the CA certificates are consistent on both local and remote switches.

```
switch:admin> seccertutil import -fcapcacert
Select protocol [ftp or scp]: scp
Enter IP address: 10.1.2.3
Enter remote directory: /myHome/jdoe/OPENSSL
Enter certificate name (must have a ".pem" suffix): CACert.pem
Enter Login Name: jdoe
jdoe@10.1.2.3's password: <hidden text>
Success: imported certificate [CACert.pem].
```

Importing the FCAP switch certificate

ATTENTION

The CA certificates must be installed prior to installing the switch certificate.

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 `secCertUtil import -fcapswcert` command.

```
switch:admin> seccertutil import -fcapswcert
Select protocol [ftp or scp]: scp
Enter IP address: 10.1.2.3
Enter remote directory: /myHome/jdoe/OPENSSL
Enter certificate name (must have ".crt" or ".cer" ".pem" or ".psk"
suffix): 01.pem
Enter Login Name: jdoe
jdoe@10.1.2.3's password: <hidden text>
Success: imported certificate [01.pem].
```

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 and select **active** or **on**, the default is passive. 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 by command; there is no support for automatic distribution. To distribute the AUTH policy, see [“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 may initiate distribution of the policy. To set the local switch configuration parameter, refer to [“Policy database distribution”](#) on page 260.

NOTE

This 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 to go 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. It is not allowed to specify an IPv6 address in the IPv4 filter, or 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: 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 only select port numbers in the well-known port number range, between 0 and 1023, 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 45 lists the supported service names and their corresponding port numbers.

TABLE 45 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

TABLE 45 Supported services (Continued)

Service name	Port number
time	27
name	42
whois	43
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 46](#) 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 46 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. The default IP Filter policy cannot be deleted or changed. When an alternative IP Filter policy is activated, the default IP Filter policy becomes deactivated. [Table 47](#) lists the rules of the default IP Filter policy.

TABLE 47 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
7	Any	443	TCP	Permit
8	Any	161	UDP	Permit
10	Any	123	UDP	Permit
11 ¹	Any	600-1023	TCP	Permit
12 ¹	Any	600-1023	UDP	Permit

1. None of the RPC ports are configurable, even though the action shows "Permit".

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 by 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 may 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 by command; there is no support for automatic distribution. To distribute the IPFilter policy, see [“Distributing the local ACL policies”](#) on page 263 for instructions.

You can accept or deny IP Filter policy distribution through the commands `fddCfg --localaccept` or `fddCfg --localreject`. See [“Policy database distribution”](#) on page 260 for more information on 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.

Virtual Fabrics considerations: To distribute the IPFilter policy in a logical fabric, 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, see [“Database distribution settings”](#) on page 261.

Virtual Fabric considerations: FCS, DCC, SCC, and AUTH databases can be distributed using the `-distribute` command, but the PWD and IPFILTER databases are blocked from distribution.

- **Manually distribute an ACL policy database** — Run the `distribute` command to push the local database of the specified policy type to target switches. [“ACL policy distribution to other switches”](#) on page 262.
- **Fabric-wide consistency policy** — Use 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, see [“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.

[Table 48](#) on page 261 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 distribution command.

TABLE 48 Interaction between fabric-wide consistency policy and distribution settings

Distribution setting	Fabric-wide consistency policy		
	Absent (default)	Tolerant	Strict
Reject	Database is protected, it cannot be overwritten. May not match other databases in the fabric.	Invalid configuration. ¹	Invalid configuration. ¹
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. 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.

1. An error is returned indicating that the distribution setting must be accept before you can set the fabric-wide consistency policy.

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 49](#) lists the databases supported in Fabric OS v6.2.0 and later switches.

TABLE 49 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.

Displaying the database distribution settings

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the FabricDistribution RBAC class of commands.
2. Enter the `fddCfg --showall` command.

Example 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 (see [“Database distribution settings”](#) on page 261).
- The fabric must have a tolerant or no (absent) fabric-wide consistency policy (see [“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

FC routers cannot join a fabric with a strict fabric-wide consistency policy. FC routers do not support the fabric-wide consistency policies.

[Table 50](#) describes the fabric-wide consistency settings.

TABLE 50 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 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.

Example 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.

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 join the fabric, but an error message flags 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 this 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 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 above 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 51](#) describes the impact of merging fabrics with the same fabric-wide consistency policy that have SCC, DCC, or both policies.

TABLE 51 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.
	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 ¹ .
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.

1. 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.

Non-matching fabric-wide consistency policies

You may encounter one of the following two scenarios described in [Table 52](#) and [Table 53](#) 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 52](#) on page 266 shows merges that are not supported.

TABLE 52 Examples of strict fabric merges

Fabric-wide consistency policy setting	Fabric A		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 53 has a matrix of merging fabrics with tolerant and absent policies.

TABLE 53 Fabric merges with tolerant and absent combinations

Fabric-wide consistency policy setting	Fabric A		Expected behavior
	Fabric A	Fabric B	
Tolerant/Absent	SCC;DCC		Error message logged. Run <code>fddCfg --fabwideset "policy_ID"</code> from any switch with the desired configuration to fix the conflict. The <code>secPolicyActivate</code> command is blocked until conflict is resolved.
	DCC		
	SCC;DCC	SCC	
	DCC	SCC	

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.

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:

- **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 [Chapter 2, “Performing Basic Configuration Tasks”](#).

Configuration examples

Below are several examples of 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 18 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 19 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 20 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. For more information on the `ipSecConfig` command, refer to the *Fabric OS Command Reference*.

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 54](#) when configuring the authentication algorithm.

TABLE 54 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 **ipSecConfig** command does not support manipulating pre-shared keys corresponding to the identity of the IKE peer or group of peers. Use the **secCertUtil** command to import, delete, or display the pre-shared keys in the local switch database. For more information on this procedure, refer to [Chapter 7, “Configuring Protocols”](#).

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 private key (in X.509 PEM format) and peer public key (in X.509 format) to be used in a particular IKE policy.

Use the **secCertUtil import** command to import public key, private key and peer-public key (in X.509 PEM format) into the switch database. For more information on this procedure, refer to [Chapter 7, “Configuring Protocols”](#).

ATTENTION

The CA certificate name must have the **IPSECCA.pem** name.

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.

Creating the tunnel

Each side of the tunnel must be configured in order for the tunnel to come up. Once you are logged into the switch, do not log off as each step requires that you be logged in to the switch. IPsec configuration changes take effect upon execution and are persistent across reboots. Configure the following on each side of the tunnel:

1. Determine the authentication protocol and algorithm to be used on the tunnel.
Refer to [Table 54](#) 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> ipsecconfig --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> ipsecconfig --add policy ips sa-proposal -t IPSEC-AH -sa AH01
```

6. Import the pre-shared key file.
Refer to [Chapter 7, "Configuring Protocols"](#) 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> ipsecconfig --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> ipsecconfig --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> ipsecconfig --add policy ips selector -t SELECTOR-OUT -d out
-l 10.33.69.132 -r 10.33.74.13 -transform TRANSFORM01
```

```
switch:admin> ipsecconfig --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.

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 above 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, BROCADE300 (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> ipsecconfig --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> ipsecconfig --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> ipsecconfig --add policy ips sa-proposal -t IPSEC-AH
-lifetime 86400 -sa AH01
```

6. Import the pre-shared key file using the `secCertUtil` 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 203.

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.
 - a. Initiate Telnet or SSH or ping session from BRCD300 to Remote Host.
 - b. Verify that the IP traffic is encapsulated.
 - c. Monitor IPsec SAs created using IKE for the above 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.

**CAUTION**

Flushing SAs requires IPsec to be disabled and re-enabled. This operation is disruptive to traffic using the tunnel.

Notes

- As of Fabric OS 7.0.0, IPsec no longer supports null encryption (null_enc) for IKE policies.
- IPv6 policies cannot tunnel IMCP traffic.

8 Management interface security

Maintaining the Switch Configuration File

In this chapter

- Configuration settings 277
- Configuration file backup 279
- Configuration file restoration 280
- Configurations across a fabric 284
- Configuration management for Virtual Fabrics 285
- Brocade configuration form 287

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, Brocade recommends that you back up all important configuration data for every switch on a host computer server as a safety measure.

NOTE

For information about AD-enabled switches, refer to [Chapter 20, “Managing Administrative Domains”](#).

For more information about troubleshooting configuration file uploads and downloads, refer to the *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:

-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.

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 (slot-based)
- DMM_WWN – Data migration manager World Wide Name configuration
- Licenses – (Feature-based) Licenses configuration
- AGWWN_MAPPING_CONF – Access Gateway WWN mapping configuration
- LicensesLservc – Sentinel License configuration
- GE blade mode – GigE Mode configuration
- FWD CHASSIS CFG – Fabric Watch 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

Brocade recommends keeping 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/download**, **supportsave**, and **auto FFDC/trace upload (supportftp)** commands.

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.

Example of configUpload on a switch without Admin Domains

```
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
```

Example of configUpload on a switch with Admin Domains

NOTE

Administrative domains other than AD255 upload a subset of information. If you want a complete switch configuration, you must use the **configUpload** command while logged in to AD255.

```
switch:AD5:admin> ad --select 5
switch:AD5:admin> configUpload
Protocol (scp or ftp) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: UserFoo
Path/Filename [<home dir>/config.txt]: /pub/configurations/config.txt
Password: <hidden>
configUpload complete: Only zoning parameters are uploaded from ad5.
```

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 285 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.

While it is possible to transfer a Fabric OS 6.4.1 configuration file to a Fabric OS 7.0.0 or later switch, it is not possible to transfer a Fabric OS 7.0.0 or later configuration file to a Fabric OS 6.4.1 switch.

Restrictions

This section lists restrictions for some of the options of the **configDownload** command.

-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 282.

-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.

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 282 for more information on non-disruptive configuration downloads.

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.

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.

If you must set up your switch again, run the commands listed in [Table 55](#) and save the output in a file format. Store the files in a safe place for emergency reference.

TABLE 55 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.
fcRouterPortCost	FC Router route information.

TABLE 55 CLI commands to display or modify switch configuration information (Continued)

Command	Displays
fcxLateConfig	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 is limited in its ability. Edited files can become corrupted and this corruption can lead to switch failures.

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, Fabric Watch, 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.

Example of configDownload without Admin Domains

```
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 configDownload with Admin Domains

```
switch:AD5:admin>configdownload
Protocol (scp or ftp) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: UserFoo
Path/Filename [<home dir>/config.txt]: /pub/configurations/config.txt
```

*** 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>
```

```
Activating configDownload: Switch is disabled
```

```
configDownload complete: Only zoning parameters are downloaded to ad5.
```

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.

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 279 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 280 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 [Chapter 7, “Configuring Protocols”](#).

Configuration management for Virtual Fabrics

You can use the **configUpload -vf** or **configDownload -vf** command to restore configurations to a logical switch. The **-vf** option only restores the Virtual Fabrics configuration information on to a switch of the same model and same release. 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.

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

```
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
```

Example of configUpload on a logical switch configuration

```
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
2009/07/20-09:13:40, [LOG-1000], 225, SLOT 7 | CHASSIS, INFO, BrocadeDCX,
Previous message repeated 7 time(s)
2009/07/20-10:27:14, [CONF-1001], 226, SLOT 7 | FID 128, INFO, DCX_80,
configUpload completed successfully for VF config parameters.
```

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` command.
5. Respond to the prompts.

Wait for the configuration file to download to the switch.

6. Verify the LISL 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)
```

Restrictions

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

Use the form in [Table 56](#) as a hard copy reference for your configuration information. In the hardware reference manuals for the Brocade DCX and DCX-4S Backbones, there is a guide for FC port-setting.

TABLE 56 Brocade configuration and connection form

Brocade configuration settings
IP address
Gateway address
Chassis configuration option
Management connections
Serial cable tag
Ethernet cable tag
Configuration information
Domain ID
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)

9 Brocade configuration form

Installing and Maintaining Firmware

In this chapter

- Firmware download process overview 289
- Preparing for a firmware download 292
- Firmware download on switches 294
- Firmware download on a Backbone 296
- Firmware download from a USB device 299
- FIPS support 300
- Testing and restoring firmware on switches 302
- Testing and restoring firmware on Backbones 304
- Validating a firmware download 306

Firmware download process overview

Fabric OS v7.2.0 provides nondisruptive firmware installation.

This chapter refers to the following specific types of blades inserted into the Brocade DCX and DCX 8510 Backbone families:

- FC blades or port blades that contain only Fibre Channel ports; the Brocade FC8-16, FC8-32, FC8-48, and FC8-64; and the Brocade FC16-32 and FC16-48 blades for 16-Gbps-capable FC blades.
- AP blades contain extra processors and specialized ports: FCOE10-24, FX8-24, and FS8-18 encryption blade.
- CP blades have a control processor (CP) used to control the entire switch; CP blades can be inserted only into slots 6 and 7 on the Brocade DCX or DCX 8510-8, and slots 4 and 5 on the Brocade DCX-4S or DCX 8510-4.
- CR8 and CR4S-8 core blades provide ICL functionality between two Brocade DCX Backbones. CR8 blades can be inserted only into slots 5 and 8 on the Brocade DCX. CR4S-8 blades can be inserted only into slots 3 and 6 on the Brocade DCX-4S.
- CR8 and CR4S-8 core blades provide ICL functionality between two Brocade DCX 8510-8 Backbones. CR4S-8 blades can be inserted only into slots 3 and 6 on the Brocade DCX 8510-4.

NOTE

For more information on troubleshooting a firmware download, refer to the *Fabric OS Troubleshooting and Diagnostics Guide*.

You can download Fabric OS to a Backbone, which is a chassis; and to a nonchassis-based system, also referred to as a fixed-port switch. The difference in the download process is that Backbones have two CPs and fixed-port switches have one CP. Use the **firmwareDownload** command to download the firmware from either an FTP or SSH server by using FTP, SFTP, or SCP to the switch. Or you can use a Brocade-branded USB device.

New firmware consists of multiple files in the form of RPM packages listed in a .plist file. The .plist file contains specific firmware information (time stamp, platform code, version, and so forth) and the names of packages of the firmware to be downloaded. These packages are made available periodically to add features or to remedy defects. Contact your switch support provider to obtain information about available firmware versions.

All systems maintain two partitions (a primary and a secondary) of nonvolatile storage areas to store firmware images. The firmware download process always loads the new image into the secondary partition. It then swaps the secondary partition to be the primary and High Availability (HA) reboots (which is nondisruptive) the system. After the system boots up, the new firmware is activated. The firmware download process then copies the new image from the primary partition to the secondary partition.

In dual-CP systems, the firmware download process, by default, sequentially upgrades the firmware image on both CPs using HA failover to prevent disruption to traffic flowing through the Backbone. This operation depends on the HA status on the Backbone. If the platform does not support HA, you can still upgrade the CPs one at a time.

If you are using a Brocade DCX or DCX 8510 Backbone family platform, with one or more AP blades: Fabric OS automatically detects mismatches between the active CP firmware and the blade's firmware and triggers the autoleveling process. This autoleveling process automatically updates the blade firmware to match the active CP. At the end of the autoleveling process, the active CP and the blade run the same version of the firmware.

If the firmware download process is interrupted by an unexpected reboot, the system automatically repairs and recovers the secondary partition. You must wait for the recovery to complete before issuing another **firmwareDownload** command.

The command supports both non-interactive and interactive modes. If the **firmwareDownload** command is issued without any operands, or if there is any syntax error in the parameters, the command enters an interactive mode, in which you are prompted for input.

ATTENTION

For each switch in your fabric, complete all firmware download changes on the current switch before issuing the **firmwareDownload** command on the next switch. This process ensures nondisruption of traffic between switches in your fabric.

To verify the firmware download process is complete, enter the **firmwareDownloadStatus** command on the switch, verify the process is complete, and then move to the next switch.

Upgrading and downgrading firmware

Upgrading means installing a newer version of firmware. *Downgrading* means installing an older version of firmware.

In most cases, you will be *upgrading* firmware; that is, installing a newer firmware version than the one you are currently running. However, some circumstances may require installing an older version; that is, *downgrading* the firmware. The procedures in this section assume that you are upgrading firmware, but they work for downgrading as well, provided the old and new firmware versions are compatible. Always reference the latest release notes for updates that may exist regarding downgrades under particular circumstances.

For details on Administrative Domains and the firmware download process, refer to [Chapter 20, “Managing Administrative Domains”](#) for more information.

For details about testing and restoring firmware, refer to [“Testing and restoring firmware on Backbones”](#) on page 304.

Passwordless firmware download

You can download firmware without a password using the **sshutil** command for public key authentication when SSH is selected. The switch must be configured to install the private key, and then you must export the public key to the remote host. Before running the **firmwareDownload** command, you must first configure the SSH protocol to permit passwordless logins for outgoing authentication as described in [“Configuring outgoing SSH authentication”](#) on page 199.

Considerations for FICON CUP environments

To prevent channel errors during nondisruptive firmware installation, the switch CUP port must be taken offline from all host systems.

HA sync state

High Availability (HA) synchronization occurs when two CPs in a Backbone are synchronized. This state provides redundancy and a nondisruptive firmware download. In order for a firmware download to successfully occur, the two CPs in a Backbone must be in sync.

If the CPs have mixed versions when you enter the **firmwareDownload** command, the CPs may not be in HA sync. In this case, you must enter the **firmwareDownload -s** command first to upgrade or downgrade the standby CP to the same level as the active CP, and then upgrade the CPs to the desired version of firmware.

NOTE

Do not run mixed firmware levels on CPs.

[Table 57](#) shows the sync state of a Backbone that has different Fabric OS versions installed on the active and standby CPs. Use the table to determine if you need to use the **firmwareDownload -s** command.

TABLE 57 Backbone HA sync states

Active CP Fabric OS version	Standby CP Fabric OS version	HA sync state	Remedy
v6.4.0	v6.4.0	inSync	Run firmwareDownload -s on the standby CP to upgrade it to v7.0.0
v6.4.0	v7.1.0	Not inSync	N/A
v7.0.0	v6.4.0	inSync	Run firmwareDownload -s on the standby CP to upgrade it to v7.0.0
v7.0.0	v7.0.0	inSync	N/A
v7.0.0	v7.1.0	inSync	N/A
v7.0.0	v7.2.0	Not inSync	N/A
v7.1.0	v6.4.0	Not inSync	N/A
v7.1.0	v7.0.0	inSync	N/A
v7.1.0	v7.1.0	inSync	N/A
v7.1.0	v7.2.0	InSync	N/A
v7.2.0	v7.1.0	InSync	N/A
v7.2.0	v7.2.0	InSync	N/A

Preparing for a firmware download

Before executing a firmware download, it is recommended that you perform the tasks listed in this section. In the unlikely event of a failure or timeout, these preparatory tasks enable you to provide your switch support provider the information required to troubleshoot the firmware download.

It is recommended that you use the **configUpload** command to back up the current configuration before you download firmware to a switch. Refer to [“Configuration file backup”](#) on page 279 for details.

1. Read the release notes for the new firmware to find out if there are any updates related to the firmware download process.
2. Connect to the switch and log in using an account with admin permissions. Enter the **firmwareShow** command to verify the current version of Fabric OS.

Brocade does not support non-disruptive upgrades from more than one previous release. Non-disruptive upgrade is supported only from Fabric OS 7.1.x to Fabric OS 7.2.x.

Disruptive upgrade from Fabric OS 7.0.x to Fabric OS 7.2.x is supported.

3. Use the **configUpload** command prior to the firmware download. Save the configuration file on your FTP or SSH server or USB memory device on supported platforms.
4. *Optional:* For additional support, connect the switch to a computer with a serial console cable. Ensure that all serial consoles (both CPs for Backbones) and any open network connection sessions, such as Telnet, are logged and included with any trouble reports.

5. Connect to the switch and log in using an account with admin permissions. Enter the **supportSave** command to retrieve all current core files prior to executing the firmware download. This information helps to troubleshoot the firmware download process if a problem is encountered.
6. *Optional:* Enter the **errClear** command to erase all existing messages in addition to internal messages.

Obtaining and decompressing firmware

Firmware upgrades are available for customers with support service contracts and for partners on the Brocade website at <http://www.brocade.com>.

You must decompress the firmware *before* you can use the **firmwareDownload** command to update the firmware on your equipment. Use the UNIX tar command for .tar files, the gunzip command for all .gz files, or a Windows unzip program for all .zip files

When you unpack the downloaded firmware, it expands into a directory that is named according to the version of Fabric OS it contains. For example, when you download and unzip v7.2.0.zip, it expands into a directory called v7.2.0. When you issue the **firmwareDownload** command, there is an automatic search for the correct package file type associated with the switch. Specify only the path up to and including the v7.2.0 directory.

Connected switches

Before you upgrade the firmware on your switch, you must check the connected switches to ensure compatibility and that any older versions are supported. Refer to the Fabric OS Compatibility section of the *Brocade Fabric OS Release Notes* for the recommended firmware version.

If fixed-port switches are adjacent and you start firmware downloads on them at the same time, there may be traffic disruption.

To determine if you need to upgrade switches connected to the switch you are upgrading, use the following procedure on each connected switch to display firmware information and build dates.

Finding the switch firmware version

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

The following information is displayed:

- **Kernel** displays the version of the switch kernel operating system.
- **Fabric OS** displays the version of the switch Fabric OS.
- **Made on** displays the build date of the firmware running on the switch.
- **Flash** displays the install date of firmware stored in nonvolatile memory.
- **BootProm** displays the version of the firmware stored in the boot PROM.

Firmware download on switches

Brocade fixed-port switches maintain primary and secondary partitions for firmware. The **firmwareDownload** command defaults to an autocommit option that automatically copies the firmware from one partition to the other.

NOTE

This section only applies when upgrading from Fabric OS v7.1.x to v7.2.0, downgrading from v7.2.0 to v7.1.x, or going from v7.2.x to v7.2.x

If you are upgrading from Fabric OS v7.0.x to v7.2.0 or downgrading from v7.2.0 to v7.0.x or earlier, you must enter the **firmwareDownload -s** command as described in [“Testing and restoring firmware on switches”](#) on page 302.

You cannot download firmware if you are going from v7.2.0 to v6.4 (or earlier) or from v6.4 (or earlier) to v7.2.0.

Do not override the autocommit option under normal circumstances; use the default. Refer to [“Testing and restoring firmware on Backbones”](#) on page 304 for details about overriding the autocommit option.

Switch firmware download process overview

The following list describes the default behavior after you enter the **firmwareDownload** command (without options) on Brocade fixed-port switches:

- The Fabric OS downloads the firmware to the secondary partition.
- The system performs a high availability reboot (**haReboot**). After the **haReboot**, the former secondary partition is the primary partition.
- The system replicates the firmware from the primary to the secondary partition.

The upgrade process first downloads and then commits the firmware to the switch. While the upgrade is proceeding, you can start a session on the switch and use the **firmwareDownloadStatus** command to observe the upgrade progress.



CAUTION

After you start the process, do not enter any disruptive commands (such as **reboot**) that interrupt the process. The entire firmware download and commit process takes approximately 17 minutes.

If there is a problem, wait for the timeout (30 minutes for network problems) before issuing the **firmwareDownload** command again. Disrupting the process can render the switch inoperable and require you to seek help from your switch service provider.

Do not disconnect the switch from power during the process. The switch could be inoperable when rebooted.

Upgrading firmware for Brocade fixed-port switches

1. Take the following appropriate action based on what service you are using:
 - If you are using FTP, SFTP, or SCP, verify that the FTP or SSH server is running on the host server and that you have a valid user ID and password on that server.
 - If your platform supports a USB memory device, verify that it is connected and running.
2. Obtain the firmware file from the Brocade website at <http://www.brocade.com> and store the file on the FTP or SSH server or the USB memory device.
3. Unpack the compressed files preserving directory structures.

The firmware is in the form of RPM packages with names defined in a .plist file. The .plist file contains specific firmware information and the names of packages of the firmware to be downloaded.

4. Connect to the switch and log in using an account with admin permissions.
5. Issue the **firmwareShow** command to check the current firmware version on connected switches. Upgrade the firmware on the connected switches, if necessary, before proceeding with upgrading this switch.

Refer to “[Connected switches](#)” on page 293 for details.

6. Enter the **firmwareDownload** command and respond to the prompts.

NOTE

If DNS is enabled and a server name instead of a server IP address is specified in the command line, **firmwareDownload** determines whether IPv4 or IPv6 should be used.

To be able to mention the FTP server by name, you must enter at least one DNS server using the **dnsConfig** command.

7. At the “Do you want to continue [y/n]” prompt, enter **y**.
8. After the HA reboot, connect to the switch and log in again using an account with admin permissions.
9. Enter the **firmwareDownloadStatus** command to determine if the firmware download process has completed.
10. After the firmware commit is completed, which takes several minutes, enter the **firmwareShow** command to verify the firmware level of both partitions is the same.

Example of an interactive firmware download

```
switch:root> firmwaredownload
Server Name or IP Address: 10.31.2.25
User Name: releaseuser
File Name: /home/SAN/fos/v7.2.0/v7.2.0
Network Protocol(1-auto-select, 2-FTP, 3-SCP, 4-SFTP) [1]: 4
Verifying if the public key authentication is available.Please wait ...
The public key authentication is not available.
Password:
Server IP: 10.31.2.25, Protocol IPv4
Checking system settings for firmwaredownload...
```

Firmware download on a Backbone

ATTENTION

To successfully download firmware, you must have an active Ethernet connection on each CP.

You can download firmware to a Backbone without disrupting the overall fabric if the two CP blades are installed and fully synchronized. Use the **haShow** command to verify that the CPs are synchronized prior to beginning the firmware download process. If only one CP blade is inserted, powered on, or plugged into the network, you can run **firmwareDownload -s** to upgrade the CP. If the CPs are not in sync, you can run **firmwareDownload -s** on each of the CPs to upgrade them. These operations are disruptive. If the CPs are not in sync, run the **haSyncStart** command. If the CPs are still not in sync, refer to the *Fabric OS Troubleshooting and Diagnostics Guide*. If the troubleshooting information fails to help resolve the issue, contact your switch service provider.

During the upgrade process, the Backbone fails over to its standby CP blade and the IP address for the Backbone moves to that CP blade's Ethernet port. This may cause informational ARP address reassignment messages to appear on other switches in the fabric. This is normal behavior, because the association between the IP addresses and MAC addresses has changed.

Backbone firmware download process overview

The following summary describes the default behavior of the **firmwareDownload** command (without options) on a Backbone. After you enter the **firmwareDownload** command on the active CP blade the following actions occur.

1. The standby CP blade downloads firmware.
2. The standby CP blade reboots and comes up with the new Fabric OS.
3. The active CP blade synchronizes its state with the standby CP blade.
4. The active CP blade forces a failover and reboots to become the standby CP blade.
5. The new active CP blade synchronizes its state with the new standby CP blade.
6. The *new* standby CP blade (the active CP blade before the failover) downloads firmware.
7. The *new* standby CP blade reboots and comes up with the new Fabric OS.
8. The new active CP blade synchronizes its state with the new standby CP blade.
9. The **firmwareCommit** command runs automatically on both CP blades.



CAUTION

After you start the process, do not enter any disruptive commands (such as **reboot**) that interrupt the process. The entire firmware download and commit process takes approximately 17 minutes.

If there is a problem, wait for the timeout (30 minutes for network problems) before issuing the **firmwareDownload** command again. Disrupting the process can render the switch inoperable and require you to seek help from your switch service provider.

Do not disconnect the switch from power during the process. The switch could be inoperable when rebooted.

Upgrading firmware on Backbones (including blades)

There is only one chassis management IP address for the Brocade Backbones.

NOTE

By default, the **firmwareDownload** command automatically upgrades both the active and the standby CPs and all co-CPs on the CP blades in the Brocade Backbones. It automatically upgrades all AP blades in the Brocade Backbones using autoleveling.

1. Verify that the Ethernet interfaces located on CP0 and CP1 are plugged into your network.
2. Verify that the FTP, SFTP, or SSH server is running on the host server and that you have a user ID on that server.
3. Obtain the firmware file from the Brocade website at <http://www.brocade.com> and store the file on the FTP or SSH server.
4. Unpack the compressed files preserving directory structures.

The firmware is in the form of RPM packages with names defined in a .plist file. The .plist file contains specific firmware information and the names of packages of the firmware to be downloaded.

5. Connect to the chassis IP management interface or active CP and log in using an account with admin permissions.
6. Use the **firmwareShow** command to check the current firmware version on connected switches. Upgrade the firmware, if necessary, before proceeding with upgrading this switch. Refer to “[Connected switches](#)” on page 293.
7. Enter the **haShow** command to confirm that the two CP blades are synchronized.

In the following example, the active CP blade is CP0 and the standby CP blade is CP1:

```
ecp:admin> hashow
Local CP (Slot 5, CP0): Active, Warm Recovered
Remote CP (Slot 6, CP1): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
```

CP blades must be synchronized and running Fabric OS v7.1.0 or later to provide a nondisruptive download. If the two CP blades are not synchronized, enter the **haSyncStart** command to synchronize them. If the CPs still are not synchronized, contact your switch service provider.

For further troubleshooting, refer to the *Fabric OS Troubleshooting and Diagnostics Guide*.

8. Enter the **firmwareDownload** command and respond to the interactive prompts.
9. At the “Do you want to continue [y/n]” prompt, enter **y**.

The firmware is downloaded to one CP blade at a time, beginning with the standby CP blade. During the process, the active CP blade fails over. After the firmware is downloaded, a firmware commit starts on both CP blades. The entire firmware download and commit process takes approximately 17 minutes.

If an AP blade is present: At the point of the failover, an *autoleveling* process is activated. Autoleveling is triggered when the active CP detects a blade that contains a different version of the firmware, regardless of which version is older. Autoleveling downloads firmware to the AP blade, swaps partitions, reboots the blade, and copies the new firmware from the primary partition to the secondary partition. If you have multiple AP blades, they are updated simultaneously; however, the downloads can occur at different rates.

Autoleveling takes place in parallel with the firmware download being performed on the CPs, but does not impact performance. Fibre Channel traffic is not disrupted during autoleveling, but GbE traffic on AP blades may be affected. If there is an active FCIP tunnel on the FX8-24 blade, the FCIP tunnel traffic will be impacted for at least two minutes.

```
ecp:admin> firmwaredownload
Type of Firmware (FOS, SAS, or any application) [FOS]:
Server Name or IP Address: 10.1.2.3
User Name: userfoo
File Name: /home/userfoo/v7.2.0
Network Protocol (1-auto-select, 2-FTP, 3-SCP, 4-SFTP) [1]:
Password: <hidden>
```

```
Checking version compatibility...
Version compatibility check passed.
```

The following AP blades are installed in the system.

Slot Name	Versions	Traffic Disrupted
2 FS8-18	v7.2.0	Encrypted Traffic
8 FX8-24	v7.2.0	GigE

This command will upgrade the firmware on both CPs and all AP blade(s) above. If you want to upgrade firmware on a single CP only, please use `-s` option. You may run `firmwaredownloadstatus` to get the status of this command.

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

```
Do you want to continue [Y]: y
```

```
The firmware is being downloaded to the Standby CP. It may take up to 10 minutes.
```

10. Optionally, after the failover, connect to the switch, and log in again as admin. Using a separate session to connect to the switch, enter the `firmwareDownloadStatus` command to monitor the firmware download status.

```
sw0:FID128:admin> firmwaredownloadstatus
[1]: Mon Jul 22 04:27:21 2013
Slot 7 (CP1, active): Firmware is being downloaded to the switch. This step
may take up to 30 minutes.

[2]: Mon Jul 22 04:34:58 2013
Slot 7 (CP1, active): Relocating an internal firmware image on the CP blade.

[3]: Mon Jul 22 04:35:29 2013
Slot 7 (CP1, active): The internal firmware image is relocated successfully.

[4]: Mon Jul 22 04:35:30 2013
```

```
Slot 7 (CP1, active): Firmware has been downloaded to the secondary partition
of the switch.
```

```
[5]: Mon Jul 22 04:37:24 2013
Slot 7 (CP1, standby): The firmware commit operation has started. This may
take up to 10 minutes.
```

```
[6]: Mon Jul 22 04:41:59 2013
Slot 7 (CP1, standby): The commit operation has completed successfully.
```

```
[7]: Mon Jul 22 04:41:59 2013
Slot 7 (CP1, standby): Firmwaredownload command has completed successfully.
Use firmwareshow to verify the firmware versions.
```

11. Enter the **firmwareShow** command to display the new firmware versions.

Firmware download from a USB device

The Brocade 300, 5100, 5300, 6505, 6510, 6520, 7800, and VA-40FC switches and the Brocade DCX, DCX-4S, or DCX 8510 Backbones support a firmware download from a Brocade-branded USB device attached to the switch or active CP. Before the USB device can be accessed by the **firmwareDownload** command, it must be enabled and mounted as a file system. The firmware images to be downloaded must be stored under the relative path from `/usb/usbstorage/brocade/firmware` or use the absolute path in the USB file system. Multiple images can be stored under this directory. There is a `firmwarekey` directory where the public key signed firmware is stored.

When the **firmwareDownload** command line option, `-U` (uppercase), is specified, the **firmwareDownload** command downloads the specified firmware image from the USB device. When specifying a path to a firmware image in the USB device, you can only specify the relative path to `/firmware` or the absolute path.

Enabling the USB device

1. Log in to the switch using an account assigned to the admin role.
2. Enter the **usbStorage -e** command.

Viewing the USB file system

1. Log in to the switch using an account assigned to the admin role.
2. Enter the **usbStorage -l** command.

```
BrcdDCXB:admin> usbstorage -l
firmware\                381MB    2013 Jul 22 15:33
  v7.2.0\                 381MB    2013 Jul 22 10:39
config\                  0B       2013 Jul 22 15:33
support\                 0B       2013 Jul 22 15:33
firmwarekey\             0B       2013 Jul 22 15:33
Available space on usbstorage 79%
```

Downloading from the USB device using the relative path

1. Log in to the switch using an account assigned to the admin role.
2. Enter the **firmwareDownload -U** command.

```
ecp:admin>firmwaredownload -U v7.2.0
```

Downloading from the USB device using the absolute path

1. Log in to the switch using an account assigned to the admin role.
2. Enter the **firmwareDownload** command with the -U operand.

```
ecp:admin>firmwaredownload -U /usb/usbstorage/brocade/firmware/v7.2.0
```

FIPS support

Federal Information Processing Standards (FIPS) specify the security standards needed to satisfy a cryptographic module utilized within a security system for protecting sensitive information in the computer and telecommunication systems. For more information about FIPS, refer to [Chapter 8, “Configuring Security Policies”](#).

Fabric OS v7.2.0 firmware is digitally signed using the OpenSSL utility to provide FIPS support. To use the digitally signed software, you must configure the switch to enable signed firmware download. If it is not enabled, the firmware download process ignores the firmware signature and performs as before.

If signed firmware download is enabled, and if the validation succeeds, the firmware download process proceeds normally. If the firmware is not signed or if the signature validation fails, firmware download fails.

To enable or disable FIPS mode, refer to [Chapter 8, “Configuring Security Policies”](#).

Public and private key management

For signed firmware, Brocade uses RSA with 1024-bit length key pairs, a private key and a public key. The private key is used to sign the firmware files when the firmware is generated. The public key is packaged in an RPM package as part of the firmware, and is downloaded to the switch. After it is downloaded, it can be used to validate the firmware to be downloaded next time when you run the **firmwareDownload** command.

The public key file on the switch contains only one public key. It is only able to validate firmware signed using one corresponding private key. If the private key changes in future releases, you need to change the public key on the switch by one of the following methods:

- By using the **firmwareDownload** command. When a new firmware is downloaded, firmware download always replaces the public key file on the switch with what is in the new firmware. This allows you to have planned firmware key changes.
- By using the **firmwareKeyUpdate** command. This command retrieves a specified public key file from a specific server location and replaces the one on the switch. The information about firmware versions and their corresponding public key files is documented in the release notes or stored in a known location on the Brocade website. This command allows the customer to handle unplanned firmware key changes.

NOTE

If FIPS mode is enabled, all logins should be handled through SSH or direct serial method, and the transfer protocol should be SCP.

Updating the firmware key

1. Log in to the switch as admin.
2. Enter the **firmwareKeyUpdate** command and respond to the prompts.

The firmwareDownload command

The public key file must be packaged, installed, and run on your switch before you download a signed firmware.

When firmware download installs a firmware file, it must validate the signature of the file. Different scenarios are handled as follows:

- If a firmware file does not have a signature, how it is handled depends on the “signed_firmware” parameter on the switch. If it is enabled, firmware download fails. Otherwise, firmware download displays a warning message and proceeds normally. When downgrading to non-FIPS-compliant firmware, the “signed_firmware” flag must be disabled.
- If the firmware file has a signature but the validation fails, firmware download fails. This means the firmware is not from Brocade, or the contents have been modified.
- If the firmware file has a signature and the validation succeeds, firmware download proceeds normally.

SAS, DMM, and third-party application images are not signed.

Configuring a switch for signed firmware

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **configure** command.
3. Respond to the prompts as follows:

<i>System Service</i>	Press Enter to select default setting; default is no.
<i>ssl attributes</i>	Press Enter to select default setting; default is no.
<i>snmp</i>	
<i>attributes</i>	Press Enter to select default setting; default is no.
<i>rpcd attributes</i>	Press Enter to select default setting; default is no.
<i>cfgload</i>	Select Yes . The following questions are displayed:
<i>attributes</i>	Enforce secure config Upload/Download: Select yes .
	Enforce signed firmware download: Select yes .
<i>Webtools</i>	
<i>attributes</i>	Press Enter to select default setting; default is no.
<i>System</i>	Press Enter to select default setting; default is no.

Power-on firmware checksum test

FIPS requires the checksums of the executables and libraries on the filesystem to be validated before Fabric OS modules are launched. This is to make sure these files have not been changed after they are installed.

When firmware RPM packages are installed during firmware download, the MD5 checksums of the firmware files are stored in the RPM database on the filesystem. The checksums go through all of the files in the RPM database. Every file compares its current checksum with the checksum that is in the RPM database. If they are different, the command displays an output message informing you of the difference.

Because the validation may take up to a few minutes, it is not performed during a hot code load. It is only performed after a cold reboot of the switch.

For more information on FIPS, refer to [Chapter 8, “Configuring Security Policies”](#).

Testing and restoring firmware on switches

NOTE

This section does not apply to SAS or storage applications applied to the FA4-18 AP blade.

Typically, users downgrade firmware after briefly evaluating a newer (or older) version and then restore the original version of the firmware. Testing a new version of firmware in this manner ensures that you do not replace existing firmware because the evaluated version occupies only one partition on the switch.

ATTENTION

When you evaluate new firmware, make sure you disable all features that are not supported by the original firmware before restoring to the original version.

Testing a different firmware version on a switch

1. Verify that the FTP, SFTP, or SSH server is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the Brocade website at <http://www.brocade.com> or the switch support provider and store the file on the FTP or SSH server.
3. Unpack the compressed files preserving directory structures.

The firmware is in the form of RPM packages with names defined in a .plist file that contains specific firmware information and the names of packages of the firmware to be downloaded.

4. Connect to the switch and log in using an account with admin permissions.
5. Enter the **firmwareShow** command to view the current firmware.
6. Enter the **firmwareDownload -s** command to update the firmware, and respond to the prompts.

Example of a firmware download to a single partition

```
ecp:admin> firmwareDownload -s  
Type of Firmware (FOS, SAS, or any application) [FOS]:  
Server Name or IP Address: 10.1.2.3  
Network Protocol (1-auto-select, 2-FTP, 3-SCP, 4-SFTP) [1]:
```

```
User Name: userfoo
File Name: /home/userfoo/v7.2.0
Password: <hidden>
Do Auto-Commit after Reboot [Y]: n
Reboot system after download [N]: y
Firmware is being downloaded to the switch. This step may take up to 30
minutes.
Checking system settings for firmwaredownload...
```

The switch performs a reboot and comes up with the new firmware to be tested. Your current switch session automatically disconnects.

ATTENTION

Downloading firmware to a switch can be disruptive to switch traffic.

7. Connect to the switch, log in as admin, and enter the **firmwareShow** command to confirm that the primary partition of the switch contains the new firmware.

You are now ready to evaluate the new version of firmware.

ATTENTION

Stop! If you want to restore the firmware, stop here and skip ahead to [step 9](#); otherwise, continue to [step 8](#) to commit the firmware on the switch, which completes the firmware download operations.

8. Commit the firmware.
 - a. Enter the **firmwareCommit** command to update the secondary partition with new firmware. Note that it takes several minutes to complete the commit operation.
 - b. Enter the **firmwareShow** command to confirm both partitions on the switch contain the new firmware.

ATTENTION

Stop! If you have completed [step 8](#), then you have committed the firmware on the switch and you have completed the firmware download procedure.

9. Restore the firmware.
 - a. Enter the **firmwareRestore** command. The switch reboots and comes up with the original firmware again.

A firmware commit automatically begins to copy the original firmware from the primary partition to the secondary partition. At the end of the firmware commit process, both partitions have the original firmware. Note that it takes several minutes to complete the commit operation.
 - b. Wait five minutes to ensure that all processes have completed and the switch is fully up and operational.
 - c. Log in to the switch. Enter the **firmwareShow** command and verify that both partitions on the switch have the original firmware.

Testing and restoring firmware on Backbones

This procedure enables you to perform a firmware download on each CP and verify that the procedure was successful before committing to the new firmware. The old firmware is saved in the secondary partition of each CP until you enter the **firmwareCommit** command. If you decide to back out of the installation prior to the firmware commit, you can enter the **firmwareRestore** command to restore the former active Fabric OS firmware image.

The **firmwareRestore** command can only run if autocommit was disabled during the firmware download. This command cannot be used to restore SAS and SA images.

NOTE

Brocade recommends that, under normal operating conditions, you maintain the same firmware version on both CPs, and on both partitions of each CP. This organization enables you to evaluate firmware before you commit. As a standard practice, do not run mixed firmware levels on CPs.

Testing different firmware versions on Backbones

1. Connect to the Brocade Backbone IP address.
2. Enter the **ipAddrShow** command and note the address of CP0 and CP1.
3. Enter the **haShow** command and note which CP is active and which CP is standby. Verify that both CPs are in sync.
4. Enter the **firmwareShow** command and confirm that the current firmware on both partitions on both CPs is listed as expected.
5. Exit the session.
6. Update the firmware on the standby CP.
 - a. Connect to the Backbone and log in as admin to the standby CP.
 - b. Enter the **firmwareDownload -s** command and respond to the prompts.

At this point, the firmware downloads to the standby CP only. When it has completed the download to that CP, reboot it. The current Backbone session is disconnected.
7. Fail over to the standby CP.
 - a. Connect to the Backbone on the active CP.
 - b. Enter the **haShow** command to verify that HA synchronization is complete. It takes a minute or two for the standby CP to reboot and synchronize with the active CP.
 - c. Enter the **firmwareShow** command to confirm that the primary partition of the standby CP contains the new firmware.
 - d. Enter the **haFailover** command. The active CP reboots and the current Backbone session is disconnected.

If an AP blade is present: At the point of the failover, an *autoleveling* process is activated. Refer to [“Backbone firmware download process overview”](#) on page 296 for details about autoleveling.

8. Verify the failover.
 - a. Connect to the Backbone on the active CP, which is the former standby CP.
 - b. Enter the **haShow** command to verify that the HA synchronization is complete. It takes a minute or two for the standby CP, which is the old active CP, to reboot and synchronize with the active CP.

NOTE

If the CPs fail to synchronize, you can still proceed because the version being tested is already present on the active CP, and subsequent steps ensure that the standby CP is updated to the same version as the active CP.

- c. Confirm the evaluation version of firmware is now running on the active CP by entering the **firmwareShow** command.
9. Update firmware on the standby CP.
 - a. Connect to the Backbone on the standby CP, which is the former active CP.
 - b. Enter the **firmwareDownload** command with the **-s -b -n** operands. This ensures that the following steps are successful.

At this point the firmware downloads to the standby CP only and reboots it. The current Backbone session is disconnected.
 - c. Wait one minute for the standby CP to reboot, and then connect to the Backbone and log in as admin.
 - d. Enter the **firmwareShow** command to confirm that *both* primary partitions have the test drive firmware in place.

You are now ready to evaluate the new version of firmware.

ATTENTION

Stop! If you want to *restore* the firmware, stop here and skip ahead to [step 12](#); otherwise, continue to [step 10](#) to commit the firmware on both CPs, which completes the firmware download.

10. Perform a commit on the standby CP.

From the current Backbone session on the standby CP, enter the **firmwareCommit** command to update the secondary partition with new firmware. It takes several minutes to complete the commit operation. Do not do anything on the Backbone while this operation is in process.
11. Perform a commit on the active CP.
 - a. From the current Backbone session on the active CP, enter the **firmwareShow** command and confirm that only the active CP secondary partition contains the old firmware.
 - b. Enter the **firmwareCommit** command to update the secondary partition with the new firmware. It takes several minutes to complete the commit operation. Do not do anything on the Backbone while this operation is in process.
 - c. Upon completion of the **firmwareCommit** command, enter the **firmwareShow** command to confirm both partitions on both CPs contain the new firmware.
 - d. Enter the **haShow** command to confirm that the HA state is in sync.

ATTENTION

Stop! If you have completed [step 11](#), then you have committed the firmware on both CPs and you have completed the firmware download procedure.

12. Restore the firmware on the standby CP.

In the current Backbone session for the standby CP, enter the **firmwareRestore** command. The standby CP reboots and the current Backbone session ends. Both partitions have the same Fabric OS after several minutes.

13. Perform **haFailover** on the active CP.

- a. In the current Backbone session for the active CP, enter the **haShow** command to verify that HA synchronization is complete. It takes a minute or two for the standby CP to reboot and synchronize with the active CP.
- b. Enter the **haFailover** command. The active CP reboots and the current Backbone session ends. The Backbone is now running the original firmware.

14. Restore firmware on the “new” standby CP.

- a. Wait one minute and connect to the Backbone on the new standby CP, which is the former active CP.
- b. Enter the **firmwareRestore** command. The standby CP reboots and the current Backbone session ends. Both partitions have the same Fabric OS after several minutes.
- c. Wait five minutes and log in to the Backbone. Enter the **firmwareShow** command and verify that all partitions have the original firmware.

If an AP blade is present: Blade partitions always contain the same version of the firmware on both partitions. The firmware is stored on the blade’s compact flash card and is always synchronized with the active CP’s firmware. Thus, if you restore the active CP firmware, the blade firmware is automatically downloaded (autoleveled) to become consistent with the new CP firmware (the blade firmware is restored).

Your system is now restored to the original partitions on both CPs. Make sure that servers using the fabric can access their storage devices.

If you want to upgrade a Backbone with only one CP in it, follow the procedures in [“Testing and restoring firmware on switches”](#) on page 302. Be aware that upgrading a Backbone with only one CP is disruptive to switch traffic.

Validating a firmware download

Validate the firmware download by running the following commands: **firmwareShow**, **firmwareDownloadStatus**, **nsShow**, **nsAllShow**, and **fabricShow**.

All of the connected servers, storage devices, and switches should be present in the output of these commands. If there is a discrepancy, it is possible that a device or switch cannot connect to the fabric and further troubleshooting is necessary.

TABLE 58 Commands used for validating a firmware download

Command	Description
firmwareShow	Displays the current firmware level on the switch. For Brocade Backbones, this command displays the firmware loaded on both partitions (primary and secondary) for both CPs and AP blades. Brocade recommends that you maintain the same firmware level on both partitions of each CP within the Brocade Backbone. The firmwareShow command displays the firmware version on each CP.
firmwareDownloadStatus	Displays an event log that records the progress and status of events during Fabric OS, SAS, and SA firmware download. The event log is created by the current firmwareDownload command and is kept until another firmwareDownload command is issued. There is a time stamp associated with each event. When downloading SAS or SA in systems with two control processor (CP) cards, you can only run this command on the active CP. When downloading Fabric OS, the event logs in the two CPs are synchronized. This command can be run from either CP.
nsShow	Displays all devices directly connected to the switch that have logged in to the name server. Make sure the number of attached devices after the firmware download is exactly the same as the number of attached devices prior to the firmware download.
nsAllShow	Displays all devices connected to a fabric. Make sure the number of attached devices after the firmware download is exactly the same as the number of attached devices prior to the firmware download.
fabricShow	Displays all switches in a fabric. Make sure the number of switches in the fabric after the firmware download is exactly the same as the number of attached devices prior to the firmware download.

10 Validating a firmware download

Managing Virtual Fabrics

In this chapter

• Virtual Fabrics overview	309
• Logical switch overview	310
• Logical fabric overview	315
• Management model for logical switches	314
• Account management and Virtual Fabrics	319
• Supported platforms for Virtual Fabrics	320
• Limitations and restrictions of Virtual Fabrics	322
• Enabling Virtual Fabrics mode	324
• Disabling Virtual Fabrics mode	325
• Configuring logical switches to use basic configuration values	326
• Creating a logical switch or base switch	326
• Executing a command in a different logical switch context	328
• Deleting a logical switch	329
• Adding and moving ports on a logical switch	329
• Displaying logical switch configuration	330
• Changing the fabric ID of a logical switch	331
• Changing a logical switch to a base switch	331
• Setting up IP addresses for a logical switch	333
• Removing an IP address for a logical switch	333
• Configuring a logical switch to use XISLs	333
• Changing the context to a different logical fabric	334
• Creating a logical fabric using XISLs	334

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 636.

For information about supported switches and port types, refer to [“Supported platforms for Virtual Fabrics”](#) on page 320.

Virtual Fabrics and Admin Domains are mutually exclusive and are not supported at the same time on a switch.

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.

Logical switch overview

Traditionally, each switch and all the ports in the switch act as a single Fibre Channel switch (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*, and it initially contains all of the ports in the physical chassis.

[Figure 21](#) shows a switch before and after enabling Virtual Fabrics. In this example, the switch has 10 ports, labeled P0 through P9.

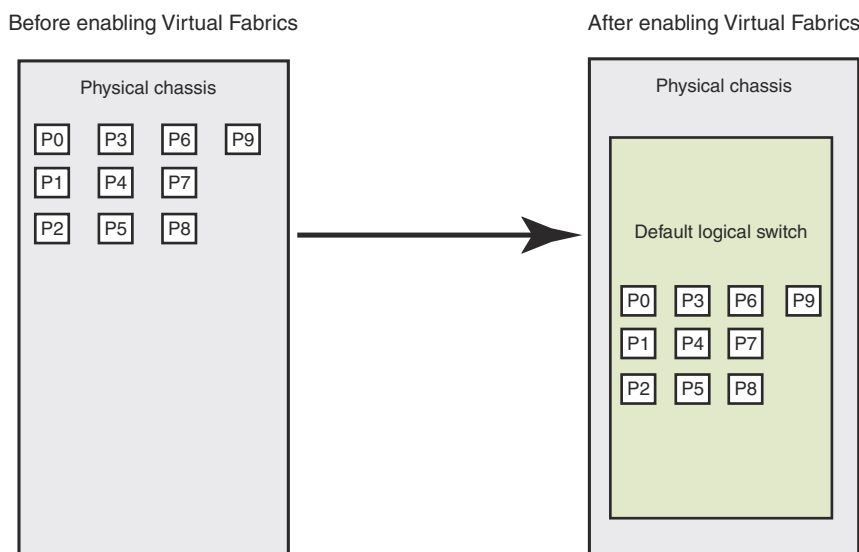


FIGURE 21 Switch before and after enabling Virtual Fabrics

After you enable Virtual Fabrics, you can create up to seven additional logical switches, depending on the switch model.

Figure 22 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.




FIGURE 22 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 Figure 23 on page 312, 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.

11 Logical switch overview




FIGURE 23 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 Figure 24, 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.

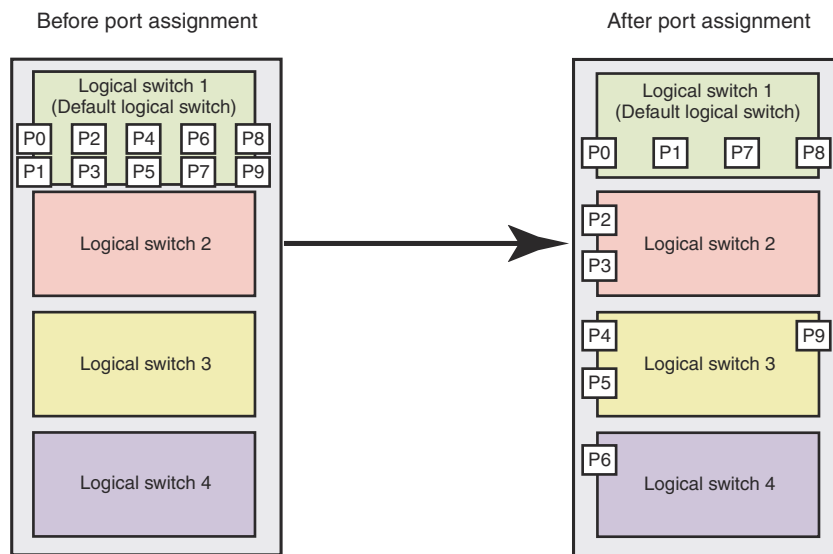


FIGURE 24 Assigning ports to logical switches

A given port is always in one (and only one) logical switch. The following scenarios refer to the chassis after port assignment in [Figure 24](#):

- 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 329 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 24](#), 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”](#) on page 320 for detailed information about these ports.

Logical switches and connected devices

You can connect devices to logical switches, as shown in [Figure 25](#) on page 314. 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 25](#), 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.

11 Management model for logical switches




FIGURE 25 Logical switches connected to devices and non-Virtual Fabrics switch

Figure 26 shows a logical representation of the physical chassis and devices in Figure 25. As shown in Figure 26, the devices are isolated into separate fabrics.




FIGURE 26 Logical switches in a single chassis belong to separate fabrics

For information on allowing device sharing across fabrics in a Virtual Fabrics environment, refer to “FC-FC routing and Virtual Fabrics” on page 636.

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 25](#) and [Figure 26](#) 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 25](#).

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 27](#) shows two physical chassis divided into logical switches. In [Figure 27](#), 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.

11 Logical fabric overview

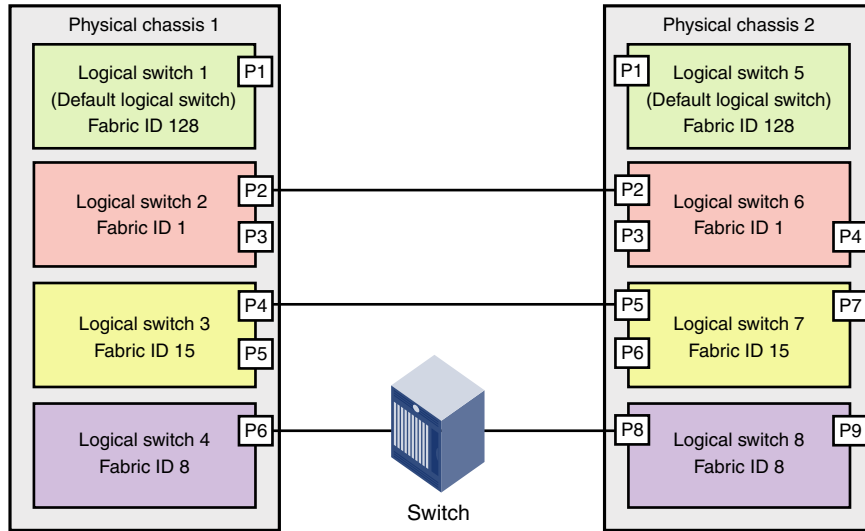


FIGURE 27 Logical switches connected to other logical switches through physical ISLs

Figure 28 shows a logical representation of the configuration in Figure 27.




FIGURE 28 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 27, 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

Another way 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.

Fabric formation across an XISL is based on the FIDs of the logical switches.

Figure 29 shows two physical chassis divided into logical switches. Each chassis has one base switch. An ISL connects the two base switches. This ISL is an extended ISL (XISL) because it connects base switches.




FIGURE 29 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 thus in different fabrics).

Think of the logical switches as being connected with logical ISLs, as shown in Figure 30. In this diagram, 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.

11 Logical fabric overview

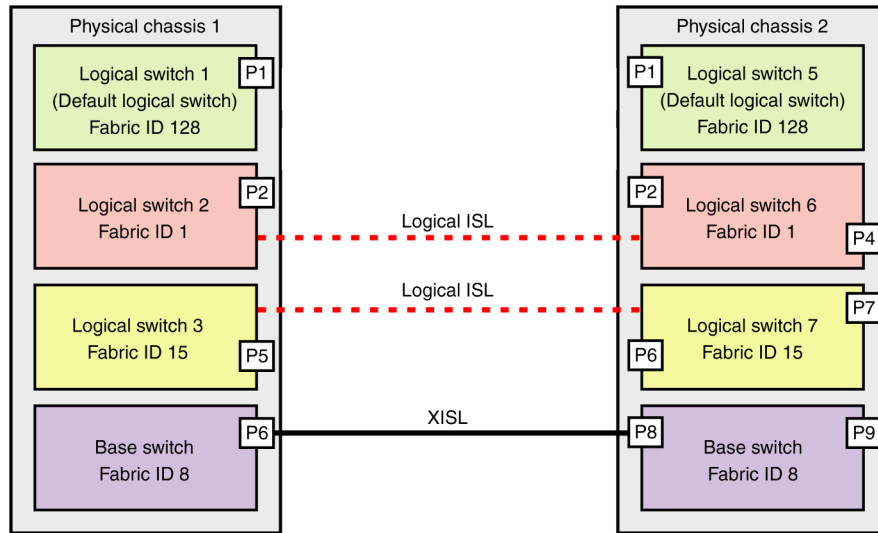


FIGURE 30 Logical ISLs connecting logical switches

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 to use XISLs”](#) on page 333.

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 [Figure 31](#). 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.

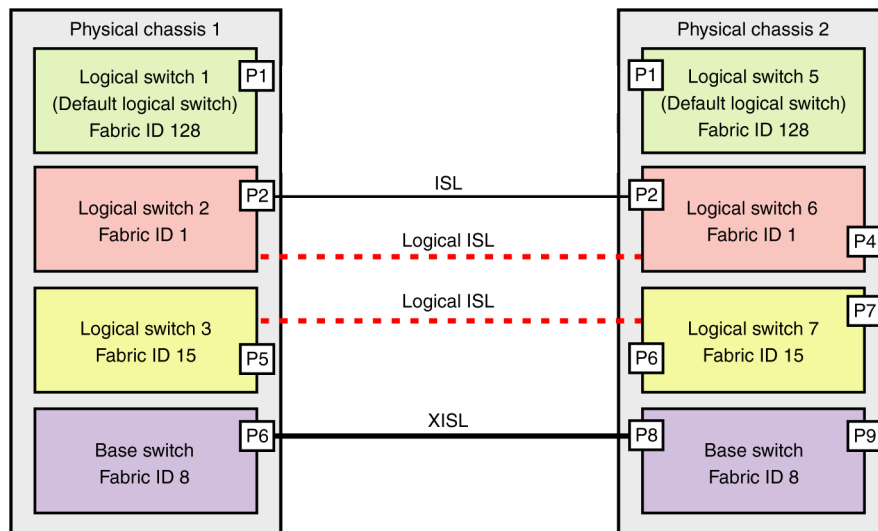


FIGURE 31 Logical fabric using ISLs and XISLs

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 broken 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 31](#), 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.

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. Traffic is initiated between the logical switches.
4. Devices begin recognizing one another.

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 334.

11 Supported platforms for Virtual Fabrics

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 [Chapter 6, “Managing User Accounts,”](#) for information about creating user accounts and assigning FIDs to user accounts.

Supported platforms for Virtual Fabrics

The following platforms are Virtual Fabrics-capable:

- Brocade 5100
- Brocade 5300
- Brocade 6510
- Brocade 6520
- Brocade 7800
- Brocade VA-40FC, in Native mode only
- Brocade DCX
- Brocade DCX-4S
- Brocade DCX 8510 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 the Brocade 5100, 5300, 6510, 6520, and VA-40FC; however, 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 DCX or DCX-4S devices.
- The default logical switch can also be a base switch.

Restrictions on fixed-port switches

Brocade 7800— Although it can be divided into four logical switches, you cannot use an XISL on this switch because a base switch is not supported on this device.

Supported port configurations in Brocade Backbones

Some of the ports in the Brocade DCX and DCX 8510 Backbone families are not supported on all types of logical switches. [Table 59](#) lists the blades and ports that are supported on each type of logical switch.

TABLE 59 Blade and port types supported on logical switches

Blade type	Default logical switch	User-defined logical switch	Base switch
FC8-16 FC8-32 FC8-32E FC8-48 FC8-48E FC16-32 FC16-48	Yes (F, E)	Yes (F, E)	Yes (E, EX)
FC8-64	Yes (F, E) ¹	Yes (F, E)	Yes (E, EX) ²
FS8-18	Yes (F, E)	No	No
FCOE10-24	Yes (F, E)	No	No
FX8-24: FC ports	Yes (F, E)	Yes (F, E,)	Yes (E, EX)
GE ports	Yes (VE)	Yes (VE)	Yes (VE, VEX)
ICL ports	Yes (E)	Yes (E)	Yes (E, EX) ³

1. In the Brocade DCX and 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-4S and DCX 8510-4 do not have this limitation.
2. In the Brocade DCX and DCX 8510-8, ports 48–63 of the FC8-64 blade are not supported in the base switch. The Brocade DCX-4S and DCX 8510-4 do not have this limitation.
3. EX_Ports on an ICL are supported only in the Brocade DCX 8510 Backbone 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

Table 60 lists some Fabric OS features and considerations that apply when using Virtual Fabrics.

TABLE 60 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.
Admin Domains	Virtual Fabrics and Admin Domains are mutually exclusive and are not supported at the same time on a switch. To use Admin Domains, you must first disable Virtual Fabrics; to use Virtual Fabrics, you must first delete all Admin Domains. Refer to “Deleting all user-defined Admin Domains non-disruptively” on page 502 for information on deleting Admin Domains without disrupting device-to-device communication.
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 Chapter 9, “Maintaining the Switch Configuration File,” for more information about how Virtual Fabrics affects the configuration file.
Encryption	Encryption functionality using the FS8-18 blade is available only on the default logical switch.
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 Chapter 26, “Using FC-FC Routing to Connect Fabrics,” 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.
Licensing	Licenses are applicable for all logical switches in a chassis.
Performance monitoring	Performance monitors are supported in a limited number of logical switches, depending on the platform type. Refer to Chapter 23, “Monitoring Fabric Performance,” for more information about performance monitoring when Virtual Fabrics is enabled.
QoS	QoS VCs are maintained across the base fabric. Refer to Chapter 14, “Optimizing Fabric Behavior,” for more information about using the Adaptive Networking features with Virtual Fabrics.
Traffic Isolation	Traffic Isolation zones with failover disabled are not supported in logical fabrics. Refer to Chapter 14, “Optimizing Fabric Behavior,” for additional information about using TI Zones with Virtual Fabrics.

Limitations and restrictions of Virtual Fabrics

Before you use the Virtual Fabrics feature, you should be aware of the restrictions and limitations regarding QSFP ports and the maximum number of logical switches per chassis.

All four ports belonging to a QSFP must be moved to the same logical switch.

The maximum number of logical switches per chassis varies depending on the switch model. [Table 61](#) lists the supported platforms and the maximum number of logical switches (including the default logical switch) supported on each.

TABLE 61 Maximum number of logical switches per chassis

Platform	Maximum number of logical switches
Brocade DCX	8
Brocade DCX-4S	8
Brocade DCX 8510 family	8
Brocade 5300	4
Brocade 5100	3
Brocade 6510	4
Brocade 6520	4
Brocade 7800	4
Brocade VA-40FC	3

Refer to [“Supported port configurations in Brocade Backbones”](#) on page 321 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 XISLs

The **Allow XISL Use** option under 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, see [“Configuring a logical switch to use XISLs”](#) on page 333.

The following restrictions apply to XISL use. 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, DCX-4S, or DCX 8510 family.
- 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 is 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
- Before moving VE_Ports, you must remove the VE_Port tunnel configuration.
- VE_Ports on the FX8-24 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.

Use the following procedure to enable Virtual Fabrics mode:

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use the **fosConfig** command to check whether VF mode is enabled:

```
fosconfig --show
```

3. Delete all Admin Domains, as described in [“Deleting all user-defined Admin Domains non-disruptively”](#) on page 502.
4. Use the **fosConfig** command to enable VF mode:

```
fosconfig --enable vf
```

5. Enter **y** at the prompt.

Example

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 deleted.

ATTENTION

If you want to use Admin Domains in a fabric, you must first disable VF mode.

Use the following procedure to disable Virtual Fabrics mode:

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use the **fosConfig** command to check whether VF mode is disabled:

```
fosconfig --show
```

3. Move all ports to the default logical switch.

```
lscfg --config 128 -slot slot -port port
```

4. Delete all 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.

Example

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 the 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.

Optionally, you can define the logical switch to be a base switch. 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.

Use the following procedure to create a logical switch or a base switch:

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:

```
lsCfg --create fabricID [ -base ] [ -force ]
```

In the command syntax, *fabricID* is the fabric ID that is to be associated with the logical switch.

Specify the **-base** option if the logical switch is to be a base switch.

Specify the **-force** option to execute the command without any user prompts or confirmation.

3. Set the context to the new logical switch.

```
setcontext fabricID (or 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.

```
switchdisable
```

5. Configure the switch attributes, including assigning a unique domain ID.

```
configure
```

6. Enable the logical switch.

```
switchenable
```

7. Assign ports to the logical switch, as described in [“Adding and moving ports on a logical switch”](#) on page 329.

Example

The following example creates a logical switch with FID 4, and then assigns domain ID 14 to it.

```
sw0:FID128:admin> lscfg --create 4
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
```

Executing a command in a different logical switch context

This procedure describes how to execute a command for a logical switch while you are in the context of a different logical switch. You can also execute a command for all the logical switches in a chassis.

The command is not executed on those logical switches for which you do not have permission.

Use the following procedure to execute a command in a different logical switch context:

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 the command on all logical switches:


```
fosexec --fid all -cmd "command"
```

Example 1: 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
```

Example 2: 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.

Use the following procedure to delete a logical switch:

1. Connect to the physical chassis and log in using an account with admin permissions.
2. Remove all ports from the logical switch as described in [“Adding and moving ports on a logical switch.”](#)
3. Enter the **lscfg** command to delete the logical switch:

```
lscfg --delete fabricID
```

The *fabricID* parameter is the fabric ID of the logical switch to be deleted.

Example of deleting the logical switch with FID 7

```
switch_4:FID4:admin> lscfg --delete 7
A Logical switch with FID 7 will be deleted.
Would you like to continue [y/n]?:y
All active login sessions for FID 7 have been terminated.
Switch successfully deleted.
```

Adding and moving ports on a logical switch

This procedure explains how to add and move ports on logical switches.

You add ports to a logical switch by moving the ports from one logical switch to another. See [“Supported platforms for Virtual Fabrics”](#) on page 320 for port restrictions.

If you want to remove a port from a logical switch, you cannot remove it from the logical switch; you must move the port to a different logical switch.

When you move a port from one logical switch to another, the port is automatically disabled. Any performance monitors that were installed on the port are deleted. If monitors are required on the port in its new location, you must manually reinstall them on the port after the move.

Notes

- If the logical switch to which the port is moved has fabric mode Top Talkers enabled, then if the port is an E_Port, fabric mode Top Talker monitors are automatically installed on that port.

11 Displaying logical switch configuration

- 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.

Use the following procedure to add or move ports on a logical switch:

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.

NOTE

On the Brocade DCX and DCX 8510-8, the **lscfg** command does not allow you to add ports 48–63 of the FC8-64 blade to the base switch. These ports are not supported on the base switch. The Brocade DCX-4S and DCX 8510-4 do not have this limitation.

3. Enter **y** at the prompt.

The ports are automatically disabled, then 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

```
sw0:FID128: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** command to display a list of all logical switches and the ports assigned to them:

```
lscfg --show [ -provision ]
```

If the **-provision** option is specified, all ports on all slots are displayed, regardless of the slot status.

Example 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 following procedure describes how you can change the fabric ID of an existing 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 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.

Use the following procedure to change the fabric ID of a logical switch:

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 to use XISLs”](#) on page 333.

11 Changing a logical switch to a base switch

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
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
```

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 also set up individual IPv4 addresses for each logical switch.

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 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.

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. For the **--add** parameter, specify the network information in dotted-decimal notation for the Ethernet IPv4 address with a Classless Inter-Domain Routing (CIDR) prefix.

The following example sets an IP address for a logical switch in a Virtual Fabric with an FID of 123 in non-interactive mode with the CIDR prefix:

```
switch:admin> ipaddrset -ls 123 --add 11.1.2.4/24
```

Removing an IP address for a logical switch

Use the following procedure to delete an IP address for a logical switch:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **ipAddrSet -ls FID -delete** command.

```
switch:admin> ipaddrset -ls 123 -delete
```

Configuring a logical switch to use XISLs

When you create a logical switch, it is configured to use XISLs by default. Use the following procedure to allow or disallow the logical switch to use XISLs in the base fabric.

XISL use is not supported in some cases. See [“Limitations and restrictions of Virtual Fabrics”](#) on page 322 for restrictions on XISL use.

Use the following procedure to configure a logical switch to use XISLs:

1. Connect to the physical chassis and log in using an account with the chassis-role permission.
2. Use the **setContext** command to set the context to the logical switch you want to manage, if you are not already in that context.

11 Changing the context to a different logical fabric

```
setcontext fabricID (or switchname)
```

The *fabricID* parameter is the FID of the logical switch you want to switch to and manage. The *switchname* parameter is the name assigned to the logical switch. You can only use one parameter at a time.

3. Use the **switchShow** command and check the value of the Allow XISL Use parameter.
4. Enter the **configure** command:

```
configure
```

5. Enter **y** after the Fabric Parameters prompt:

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

6. Enter **y** at the Allow XISL Use prompt to allow XISL use; enter **n** at the prompt to disallow XISL use:

```
Allow XISL Use (yes, y, no, n): y
```

7. Respond to the remaining prompts or press **Ctrl-d** to accept the other settings and exit.

Changing the context to a different logical fabric

You can change the context to a different logical fabric. Your user account must have permission to access the logical fabric.

Use the following procedure to change the context to a different 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 (or 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.

The *fabricID* parameter is the fabric ID of the logical switch you want to switch to and manage.

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>
```

Creating a logical fabric using XISLs

This procedure describes how to create a logical fabric using multiple chassis and XISLs and refers to the configuration shown in [Figure 32](#) as an example.




FIGURE 32 Example of logical fabrics in multiple chassis and XISLs

Use the following procedure to create a logical fabric using XISLs:

1. Set up the base switches in each chassis:
 - a. Connect to the physical chassis and log in using an account with the chassis-role permission.
 - b. Enable the Virtual Fabrics feature, if it is not already enabled. See [“Enabling Virtual Fabrics mode”](#) on page 324 for instructions.
 Enabling Virtual Fabrics automatically creates the default logical switch, with FID 128. All ports in the chassis are assigned to the default logical switch.
 - c. Create a base switch and assign it a fabric ID that will become the FID of the base fabric. See [“Creating a logical switch or base switch”](#) on page 326 for instructions on creating a base switch.
 For the example shown in [Figure 32](#), you would create a base switch with fabric ID 8.
 - d. Assign ports to the base switch, as described in [“Adding and moving ports on a logical switch”](#) on page 329.
 - e. Repeat [step a](#) through [step d](#) in all chassis that are to participate in the logical fabric.
2. Physically connect ports in the base switches to form XISLs.
3. Enable all of the base switches. This forms the base fabric.
4. Configure the logical switches in each chassis:
 - a. Connect to the physical chassis and log in using an account with the chassis-role permission.
 - b. Create a logical switch and assign it a fabric ID for the logical fabric. This FID must be different from the FID in the base fabric. See [“Creating a logical switch or base switch”](#) on page 326 for instructions.

11 Creating a logical fabric using XISLs

For the example shown in [Figure 32](#), you would create a logical switch with FID 1 and a logical switch with FID 15.

- c. Assign ports to the logical switch, as described in [“Adding and moving ports on a logical switch”](#) on page 329.
 - d. Physically connect devices and ISLs to these ports on the logical switch.
 - e. (Optional) Configure the logical switch to use XISLs, if it is not already XISL-capable. See [“Configuring a logical switch to use XISLs”](#) on page 333 for instructions.
By default, newly created logical switches are configured to allow XISL use.
 - f. Repeat [step a](#) through [step e](#) in all chassis that are to participate in the logical fabric, using the same fabric ID whenever two switches need to be part of a single logical fabric.
5. Enable all logical switches by entering the **switchEnable** command on each logical switch that you created in [step 4](#) (the base switches are already enabled).

The logical fabric is formed.

The **fabricShow** command displays all logical switches configured with the same fabric ID as the local switch and all non-Virtual Fabrics switches connected through ISLs to these logical switches.

The **switchShow** command displays logical ports as E_Ports, with -1 for the slot and the user port number for the slot port.

Administering Advanced Zoning

In this chapter

• Zone types	337
• Zoning overview	338
• Broadcast zones	343
• Zone aliases	346
• Zone creation and maintenance	350
• Default zoning mode	360
• Zone database size	362
• Zone configurations	362
• Zone object maintenance	368
• Zone configuration management	370
• Security and zoning	371
• Zone merging	371
• Concurrent zone transactions	376

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
Controls which devices receive broadcast frames. A broadcast zone restricts broadcast packets to only those devices that are members of the broadcast zone. Refer to “[Broadcast zones](#)” on page 343 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 132 for more information.
- LSAN zones
Provide device connectivity between fabrics without merging the fabrics. Refer to “[LSAN zone configuration](#)” on page 620 for more information.

- QoS zones
Assign high or low priority to designated traffic flows. QoS zones are regular zones with additional QoS attributes specified by adding a QoS prefix to the zone name. Refer to “QoS” on page 415 for more information.
- Traffic Isolation zones (TI zones)
Isolate traffic to a specific, dedicated path through the fabric. Refer to [Chapter 13, “Traffic Isolation Zoning,”](#) 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 33](#) on page 339, which shows configured zones, Red, Green, and Blue.

- Server 1 can communicate only with the Storage 1 device.
- Server 2 can communicate only with the RAID and Storage 2 devices.
- Server 3 can communicate with the RAID and Storage 1 devices.
- The Storage 3 is not assigned to a zone; no other zoned fabric device can access it.

NOTE

When using a mixed fabric—that is, a fabric containing two or more switches running different release levels of fabric operating systems—you should use the switch with the latest Fabric OS level to perform zoning tasks.

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 *Fabric OS Command Reference*.




FIGURE 33 Zoning example

Approaches to zoning

Table 62 lists the various approaches you can take when implementing zoning in a fabric.

TABLE 62 Approaches to fabric-based zoning

Zoning approach	Description
Recommended approach	
Single HBA	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. 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.
Alternative approaches	
Application	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.

TABLE 62 Approaches to fabric-based zoning (Continued)

Zoning approach	Description
Operating system	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.
Port allocation	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.
Not recommended	
No fabric zoning	Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should be utilized only in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

Zone 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.

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.

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 360). 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 63](#) lists considerations for zoning architecture.

TABLE 63 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.

TABLE 63 Considerations for zoning architecture (Continued)

Item	Description
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:

- Always zone using the switch with the latest Fabric OS release level.
Switches with earlier Fabric OS versions do not have the capability to view all the functionality that a newer Fabric OS version provides, as functionality is backwards-compatible but not forwards-compatible.
- Zone using the core switch versus an edge switch.
- Zone using a Backbone rather than a switch.
A Backbone has more resources to handle zoning changes and implementations.

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). You set up and manage broadcast zones using the standard zoning commands, described in [“Zone creation and maintenance”](#) on page 350.

Broadcast zones and Admin Domains

Each Admin Domain can have only one broadcast zone. However, all of the broadcast zones from all of the Admin Domains are considered as a single consolidated broadcast zone.

Broadcast packets are forwarded to all the ports that are part of the broadcast zone for any Admin Domain, have membership in that Admin Domain, and are zoned together (in a regular zone) with the sender of the broadcast frame.

[Figure 34](#) illustrates how broadcast zones work with Admin Domains. [Figure 34](#) shows a fabric with five devices and two Admin Domains, AD1 and AD2. Each Admin Domain has two devices and a broadcast zone.




FIGURE 34 Broadcast zones and Admin Domains

The dotted box represents the consolidated broadcast zone, which contains all of the devices that can receive broadcast packets. The actual delivery of broadcast packets is also controlled by the Admin Domain and zone enforcement logic. The consolidated broadcast zone is not an actual zone, but is just an abstraction used for explaining the behavior.

- The broadcast zone for AD1 includes member devices "1,1", "3,1" and "5,1"; however, "3,1" and "5,1" are not members of AD1. Consequently, from the AD1 broadcast zone, only "1,1" is added to the consolidated broadcast zone.
- The broadcast zone for AD2 includes member devices "2,1", "3,1", and "4,1". Even though "2,1" is a member of AD1, it is not a member of AD2 and so is not added to the consolidated broadcast zone.
- Device "3,1" is added to the consolidated broadcast zone because of its membership in the AD2 broadcast zone.

When a switch receives a broadcast packet, it forwards the packet only to those devices that are zoned with the sender and are also part of the consolidated broadcast zone.

You can check whether a broadcast zone has any invalid members that cannot be enforced in the current AD context. Refer to ["Validating a zone"](#) on page 358 for complete instructions.

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 [Chapter 26, "Using FC-FC Routing to Connect Fabrics,"](#) 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.

Best practice: All devices in a single loop should have uniform broadcast capability. If all the devices in the loop can handle broadcast frames, then add the FL_Port to the broadcast zone.

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 360 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 371.

Virtual Fabrics considerations: Alias definitions should not include logical port numbers. Zoning is not enforced on logical ports.

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!!!

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 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!!!
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 an alias

Use the following procedure to removing 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!!!
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
```

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!!!
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
```

Viewing an alias in the defined configuration

Use the following procedure to view an alias in the configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **aliShow** command to display zone configuration information.

```
alishow [--ic]["pattern"][, mode]
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

Example

The following example shows all zone aliases beginning with "arr":

```
switch:admin> alishow "arr*"
alias: array1 21:00:00:20:37:0c:76:8c
alias: array2 21:00:00:20:37:0c:66:23
```

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 2f:11:00:05:33:c1:37:a2
```

Zone creation and maintenance

Fabric OS allows you to create zones to better manage devices.

NOTE

Broadcast Zone: To create a broadcast zone, use the reserved name “broadcast”. Do *not* give a regular zone the name of “broadcast”. Refer to “[Broadcast zones](#)” on page 343 for additional information about this special type of zone.

NOTE

Virtual Fabrics considerations: 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 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 "zonename", "member[; member...]"
zonecreate "zonename", "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 creating a new zone

```
switch:admin> zonecreate sloth, "b* ; 10:00:00:00:01:1e:20:20"
switch:admin> cfgsave
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 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.

Example 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; 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:

```
zonedeldelete "zonename"
```

3. Enter the **cfgSave** command to save the change to the defined configuration.

The **cfgSave** command ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message displays on the other switches to indicate that the transaction was aborted.

Example 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> zonedeldelete 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 in the defined configuration

Use the following procedure to view a zone in the configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **zoneShow** command, using the following syntax:

```
zoneShow[--sort] ["pattern"] [, mode]
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

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
```

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]
```

Example

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 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.
Do you want to proceed?” (yes, y, no, n): [no] y
```

Example of inconsistent defined and effective configuration warning to use

```
switch: admin> zoneShow
Defined configuration:
cfg: cfg1 zone1; zone2
zone: zone1 10:00:00:00:00:00:00:01; 10:00:00:00:00:00:00:02
zone: zone2 1,1; 1,2

Effective configuration:
cfg: cfg1
zone: zone1 10:00:00:00:00:00:00:01
10:00:00:00:00:00:00:02
```

12 Default zoning mode

```
zone: zone2 1,1; 1,2

switch: admin> zoneadd zone1, 10:00:00:00:00:00:03
switch: admin> cfigsave
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 'cfigenable' command.
Do you want to proceed with saving the Defined zoning
configuration only? (yes, y, no, n): [no] y
Updating flash ...

switch:admin> zoneShow
Defined configuration:
cfg:   cfg1   zone1; zone2
zone:  zone1  10:00:00:00:00:00:00:01; 10:00:00:00:00:00:00:02;
10:00:00:00:00:00:00:03
zone:  zone2  1,1; 1,2

Effective configuration:
cfg:   cfg1
zone:  zone1  10:00:00:00:00:00:00:01; 10:00:00:00:00:00:00:02
zone:  zone2  1,1; 1,2
```

Default zoning mode

The default zoning mode controls device access 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.

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.

Admin Domain considerations: If you want to use Admin Domains, you must set the default zoning mode to No Access prior to setting up the Admin Domains. You cannot change the default zoning mode to All Access if user-specified Admin Domains are present 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.

Example

```
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.

Use the **cfgSize** command to display the zone database size.

The supported maximum zone database size is 2 MB for systems running only Brocade DCX, DCX-4S, and DCX 8510 platforms. The presence of any other platform reduces the maximum zone database size to 1 MB.

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 (1 MB). The maximum size limit is enforced per-partition, but is not enforced chassis-wide. If the chassis size limit is exceeded, you are not informed of this and unpredictable behavior may occur. It is your responsibility to keep track of the chassis-wide zone database size.

ATTENTION

In a fabric with some switches running Fabric OS 7.0.0 or later and some switches running Fabric OS versions earlier than 7.0.0, if you execute the **cfgSave** or **cfgEnable** command from a pre-7.0.0 switch, a zone database size of 128 KB is enforced.

To avoid this problem, use the switch with the latest Fabric OS version to perform zoning tasks, as described in [“Best practices for zoning”](#) on page 343. Alternatively, make sure that your pre-7.0.0 switches are upgraded with the latest patch release.

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 362.

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.

Example

```
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.

Example

```
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.

Example

```
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 360 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.

Example

```
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
```

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.

Example

```
switch:admin> cfgdelete "testcfg"
switch:admin> cfgsave
```

```
You are about to save the Defined zoning configuration. This action will only
save the changes on Defined configuration.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

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 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 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.

Example

```
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, blank spaces are ignored, and the **zone --copy** command works in any Admin Domain except AD255.

```
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, blank spaces are ignored, and the **zone --expunge** command works in any Admin Domain except AD255.

```
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, blank spaces are ignored, and the **zoneObjectRename** command works in any Admin Domain except AD255.

```
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.
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 279, [“Configuration file restoration”](#) on page 280, or the **configUpload** and **configDownload** commands in the *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 367.

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 367 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.

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 373.
- **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 *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.

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 64](#) on page 373: Defined and effective configurations
- [Table 65](#) on page 374: Different content
- [Table 66](#) on page 375: Different names
- [Table 67](#) on page 375: TI zones
- [Table 68](#) on page 376: Default access mode
- [Table 69](#) on page 376: Mixed Fabric OS versions

TABLE 64 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: <code>cfg1</code> zone1: <code>ali1</code> ; <code>ali2</code> effective: 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: <code>cfg1</code> zone1: <code>ali1</code> ; <code>ali2</code> effective: <code>cfg1</code>	defined: <code>cfg1</code> zone1: <code>ali1</code> ; <code>ali2</code> effective: none	Configuration from Switch A to propagate throughout the fabric. The configuration is enabled after the merge in the fabric.

12 Zone merging

TABLE 64 Zone merging scenarios: Defined and effective configurations (Continued)

Description	Switch A	Switch B	Expected results
Switch A and Switch B have the same defined configuration. Neither have an effective configuration.	defined: cfg1 zone1: ali1; ali2 effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	No change (clean merge).
Switch A and Switch B have the same defined and effective configuration.	defined: cfg1 zone1: ali1; ali2 effective: cfg1:	defined: cfg1 zone1: ali1; ali2 effective: cfg1:	No change (clean merge).
Switch A does not have a defined configuration. Switch B has a defined configuration.	defined: none effective: none	defined:cfg1 zone1: 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: none effective: none	defined:cfg1 zone1: ali1; ali2 effective: 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: cfg1 zone1: ali1; ali2 effective: none	defined: cfg1 zone1: ali1; ali2 effective: 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: cfg2 zone2: ali3; ali4 effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	Clean merge. The new configuration will be a composite of the two. defined: cfg1 zone1: ali1; ali2 cfg2: zone2: ali3; ali4 effective: none
Switch A and Switch B have different defined configurations. Switch B has an effective configuration.	defined: cfg2 zone2: ali3; ali4 effective: none	defined: cfg1 zone1: ali1; ali2 effective: 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 65 Zone merging scenarios: Different content

Description	Switch A	Switch B	Expected results
Effective configuration mismatch.	defined: cfg1 zone1: ali1; ali2 effective: 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; ali2 effective: irrelevant	defined: cfg1 zone1: ali3; ali4 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch

TABLE 66 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:cfg1 ali2: A; B effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same alias name, same content, different order.	defined: cfg1 ali1: A; B; C effective: irrelevant	defined: cfg1 ali1: B; C; A effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same name, different types.	effective: zone1: MARKETING	effective: cfg1: MARKETING	Fabric segments due to: Zone Conflict type mismatch
Same name, different types.	effective: zone1: MARKETING	effective: alias1: MARKETING	Fabric segments due to: Zone Conflict type mismatch
Same name, different types.	effective: cfg1: MARKETING	effective: alias1: MARKETING	Fabric segments due to: Zone Conflict type mismatch

TABLE 67 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.
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.

12 Concurrent zone transactions

TABLE 68 Zone merging scenarios: Default access mode

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.
Same default zone access mode settings.	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 69 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.

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.

Example 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

You can use the `cfgTransShow` command to list all the domains in the fabric with open transactions, as shown in the following syntax:

```
cfgTransShow [ |--opentrans | --help]
```

Example:

```
switch:admin> cfgtransshow
Current transaction token is 0x571010459
It is abortable

switch:admin> cfgtransshow --help
Usage:
cfgTransShow : Displays local open transaction token details
cfgTransShow --openTrans : Displays list of Domains with Open Transactions
cfgTransShow --help : Help

switch:admin> cfgtransshow --opentrans
```

12 Concurrent zone transactions

```
Current transaction token is 0x3109
It is abortable
Transactions Detect: Capable
Current Open Transactions
Domain List:
-----
1 2 3 4
```

Traffic Isolation Zoning

In this chapter

- Traffic Isolation Zoning overview 379
- TI zone failover. 380
- Enhanced TI zones 384
- Traffic Isolation Zoning over FC routers 386
- Fabric-Level Traffic Isolation in a backbone fabric 390
- General rules for TI zones. 394
- Supported configurations for Traffic Isolation Zoning. 396
- Limitations and restrictions of Traffic Isolation Zoning. 398
- Admin Domain considerations for Traffic Isolation Zoning. 398
- Virtual Fabrics considerations for Traffic Isolation Zoning 399
- Traffic Isolation Zoning over FC routers with Virtual Fabrics 401
- Creating a TI zone 402
- Modifying TI zones. 405
- Changing the state of a TI zone 406
- Deleting a TI zone 407
- Displaying TI zones 407
- Troubleshooting TI zone routing problems. 408
- Setting up TI zones over FCR (sample procedure) 409

Traffic Isolation Zoning overview

The Traffic Isolation Zoning feature allows you to control the flow of interswitch traffic by creating a dedicated path for traffic flowing from a specific set of source ports (N_Ports). For example, 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.

Traffic Isolation Zoning does not require a license.

Traffic isolation is implemented using a special zone, called a *Traffic Isolation zone* (TI zone). A TI zone indicates the set of N_Ports and E_Ports to be used for a specific traffic flow. When a TI zone is activated, the fabric attempts to isolate all inter-switch traffic entering 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.

Figure 35 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 line indicates the dedicated path between the initiator in Domain 1 to the target in Domain 4.




FIGURE 35 Traffic Isolation zone creating a dedicated path through the fabric

In Figure 35, 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.

Use the **zone** command to create and manage TI zones.

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.

Enable failover if you want traffic to have alternate routes if the dedicated path cannot be used, and if you want other traffic to be able to use the dedicated path if the non-dedicated paths cannot be used.

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. See [“Additional considerations when disabling failover”](#) on page 381 for additional information about using this feature.

[Table 70](#) compares the behavior of traffic when failover is enabled and disabled.

TABLE 70 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.
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.

For example, in [Figure 35](#) on page 380, 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”.

NOTE

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.

NOTE

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:

- This feature is intended for use in simple linear fabric configurations, such as that shown in [Figure 35](#) on page 380.

- Ensure that there are non-dedicated paths through the fabric for all devices that are not in a TI zone.
- If you create a TI zone with just E_Ports, failover must be enabled. If failover is disabled, the specified ISLs will not be able to route any traffic.
- If the path between devices in a TI zone is broken, no inter-switch RSCNs are generated. Each switch that is part of the TI zone generates RSCNs to locally attached devices that are part of the TI zone and are registered to receive RSCNs.
- Ensure that there are multiple paths between switches.

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.

- You should enable failover-enabled TI zones before enabling failover-disabled TI zones, to avoid dropped frames.

When you issue the **cfgEnable** command to enable the zone configuration, if you have failover disabled zones, do the following:

1. Temporarily change failover-disabled TI zones to failover-enabled.
2. Enable the zones (**cfgEnable**).
3. Reset all the zones you changed in [step 1](#) to failover-disabled.
4. Enable the zones again (**cfgEnable**).

These steps are listed in the procedures in this section.

- It is recommended that TI zone definitions and regular zone definitions match.
- Domain controller frames can use any path between switches. Disabling failover does not affect Domain Controller connectivity.

For example, in [Figure 36](#), if failover is disabled, Domain 2 can continue to send domain controller frames to Domain 3 and 4, even though the path between Domain 1 and Domain 3 is a dedicated path. Domain controller frames include zone updates and name server queries.




FIGURE 36 Fabric incorrectly configured for TI zone with failover disabled

- It is recommended that the insistent Domain ID feature be enabled; if a switch changes its active domain ID, the route is broken. See 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.

For example, in [Figure 37](#), 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 37 Dedicated path is the only shortest path

In [Figure 38](#) on page 384, 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 38 Dedicated path is not the shortest path

NOTE

For information about setting or displaying the FSPF cost of a path, see the `linkCost` and `topologyShow` commands in the *Fabric OS Command Reference*.

Enhanced TI zones

In Fabric OS v6.4.0 and later, ports can be in multiple TI zones at the same time. Zones with overlapping port members are called *enhanced TI zones* (ETIZ).

Figure 39 shows an example of two TI zones. Because these TI zones have an overlapping port (3,8), they are enhanced TI zones.




FIGURE 39 Enhanced TI zones

Enhanced TI zones are especially useful in FICON fabrics. See the *FICON Administrator's Guide* for example topologies using enhanced TI zones.

See “[Additional configuration rules for enhanced TI zones](#)” on page 396 for more information about enhanced TI zones.

Illegal 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 the TI zones are configured with failover disabled, some traffic will be dropped. If the TI zones are configured with failover enabled, all traffic will go through, but half of the traffic will be routed incorrectly according to the TI zone definitions.

A message is sent to the RASlog if a 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 408.

Illegal ETIZ configuration: separate paths from a port to devices on same domain

Figure 40 shows two enhanced TI zones 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. Traffic is routed correctly from Host 1 and Host 2 to the Target; however, traffic from the Target to the Hosts might not be.

Traffic from (3,8) destined for Domain 1 cannot go through both port 6 and port 7, so only one port is chosen. If port 6 is chosen, frames destined for (1,4) will be dropped at Domain 1. If port 7 is chosen, frames destined for (1,1) will be dropped.




FIGURE 40 Illegal ETIZ configuration: two paths from one port to two devices on the same remote domain

Illegal ETIZ configuration: separate paths from a single port to the same domain

Figure 41 shows another example of an illegal 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)

13 Traffic Isolation Zoning over FC routers

In this example traffic from the Target to Domain 2 is routed correctly. Only one TI zone describes a path to Domain 2. However, both TI zones describe different, valid paths from the Target to Domain 1. Only one path will be able to get to (1,1). Traffic from port (3,8) cannot be routed to Domain 1 over both (3,6) and (3,7), so one port will be chosen. If (3,7) is chosen, frames destined for (1,1) will be dropped at Domain 1.




FIGURE 41 Illegal ETIZ configuration: two paths from one port

Traffic Isolation Zoning over FC routers

This section describes how TI zones work with Fibre Channel routing (TI over FCR). See [Chapter 26, "Using FC-FC Routing to Connect Fabrics,"](#) for information about FC routers, phantom switches, and the FC-FC Routing Service.

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 42](#) 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.

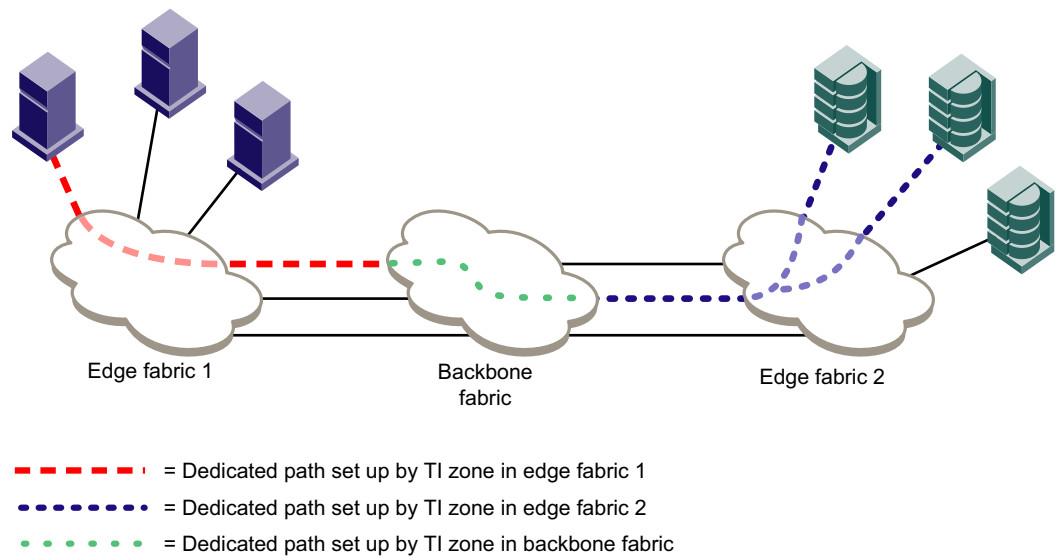


FIGURE 42 Traffic Isolation Zoning over FCR

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.

NOTE

For TI over FCR, all switches in the backbone fabric and in the edge fabrics must be running Fabric OS v6.1.0 or later.

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 43](#), you can set up a TI zone to ensure that traffic between Host 1 and the proxy target is routed through EX_Port 9.

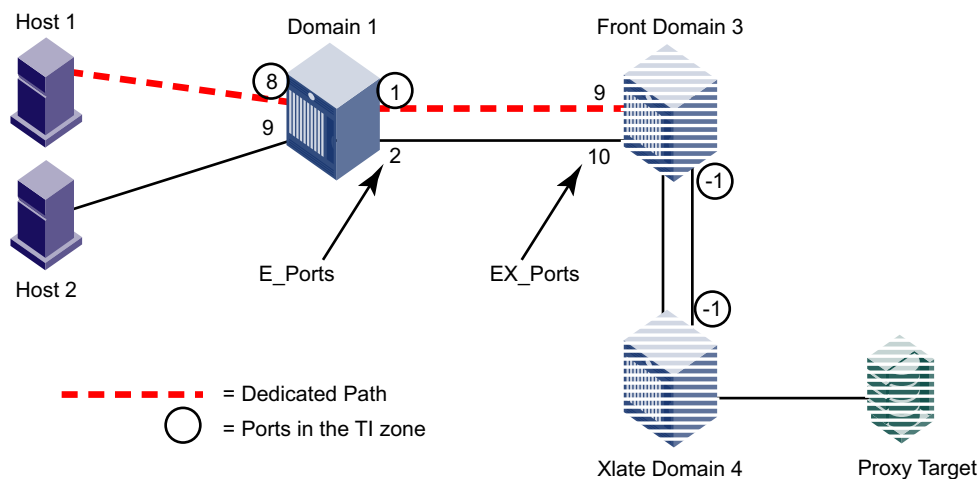


FIGURE 43 TI zone in an edge fabric

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 43](#) are:

- 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 44](#), 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 44 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 44](#) are:

- 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.
- 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.
- *For the FC8-16, FC8-32, FC8-48, FC8-64, and 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 386.

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 [Figure 45](#).




FIGURE 45 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 ISL 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.

Note the following rules for 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.

- Create a separate TI zone for each path
- Combine all of the paths in a single TI zone

The option you select affects the failover behavior of the TI zones.

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 45](#) on page 391, 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 45](#) on page 391, 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 380.

Creating a separate TI zone for each path

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 45](#) on page 391.

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
```

Note 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
```

Then enabled status now displays as “Activated”.

Creating a single TI zone for all paths

Create a single TI zone for all paths if you want TI traffic to failover 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
```

Note 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”.

General rules for TI zones

The following general rules apply to 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 indexed swapped with another port.
- A given E_Port used in a TI zone should not be a member of more than one TI zone.

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.
- 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 [Figure 46](#), 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 46 TI zone misconfiguration

Traffic Isolation 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 will help you identify missing trunk members and take corrective actions.

Example RASlog message 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.

Example RASlog message 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
```

Supported configurations for Traffic Isolation Zoning

The following configuration 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 blade in the Brocade DCX and 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 DCX 8510-4 do not have this limitation.

- VE_Ports are supported in TI zones.
- TI Zoning is not supported in fabrics with switches running firmware versions earlier than Fabric OS v6.0.0. However, the existence of a TI zone in such a fabric is backward-compatible and does not disrupt fabric operation in switches running earlier firmware versions.

TI over FCR is not backward compatible with Fabric OS v6.0.x or earlier. The -1 in the *domain,index* entries causes issues to legacy switches in a zone merge. Firmware downgrade is prevented if TI over FCR zones exist.

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 [“Supported hardware and software”](#) on page 35.
 - 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.

Trunking with TI zones

If you implement trunking and TI zones, you should keep the following points in mind:

- To include a trunk group in a TI zone, you must include all ports of the trunk in the TI zone.
- Trunked ISL ports cannot be members of more than one TI zone.
- The **zone** command includes an option to show TI trunk errors: `zone --showTitrunkerrors`

Description

This command parameter displays the details of the trunk members in the TI zone, separated into present and not present, and displayed per TI Zone basis.

Sample output

```
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
```

Limitations and restrictions of Traffic Isolation Zoning

The following limitations and restrictions apply to Traffic Isolation Zoning:

- For switches running Fabric OS 6.1.0 or later, a maximum of 255 TI zones can be created in one fabric. For switches running Fabric OS 6.0.x, no more than 239 TI zones should be created.
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. See [“Traffic Isolation Zoning over FC routers”](#) on page 386 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. This limitation does not apply to E_Ports that use the same shared area on FC4-48 and FC8-48 port blades.
- 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 v6.4.0. 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.

ATTENTION

Removing a core blade when both ICL connections and lossless dynamic load sharing are enabled may cause frame loss on a number of F_Ports.

Admin Domain considerations for Traffic Isolation Zoning

If you implement Admin Domains and TI zones, you should keep the following points in mind:

- TI zones are applicable only in ADO, and the E_Ports that are members of a TI zone must be in the ADO device list. Because TI zones must use D,I notation, the ADO device list must be declared using D,I notation for ports that are to be used in TI zones.
- A port used in a TI zone should not be a member of multiple Admin Domains.

- Use care if defining TI zones with ports that are shared across Admin Domains because of the limitation that a given port can appear in only one TI zone.
Best practice: Do not use ports that are shared across Admin Domains in a TI zone.

Virtual Fabrics considerations for Traffic Isolation Zoning

This section describes how TI zones work with Virtual Fabrics. See [Chapter 11, “Managing Virtual Fabrics,”](#) for information about the Virtual Fabrics feature, including logical switches and logical fabrics.

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 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 extended XISL (XISL) port numbers, as well as the F_Ports and ISLs in the logical 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.

[Figure 47](#) 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. ([Figure 47](#) 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 47 Dedicated path with Virtual Fabrics

[Figure 48](#) shows a logical representation of FID1 in [Figure 47](#). To create the dedicated path, you must create and activate a TI zone in FID1 that includes the circled ports shown in [Figure 48](#).

13 Virtual Fabrics considerations for Traffic Isolation Zoning




FIGURE 48 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 Figure 49, the XISLs highlighted (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 49 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:

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 50](#) 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.

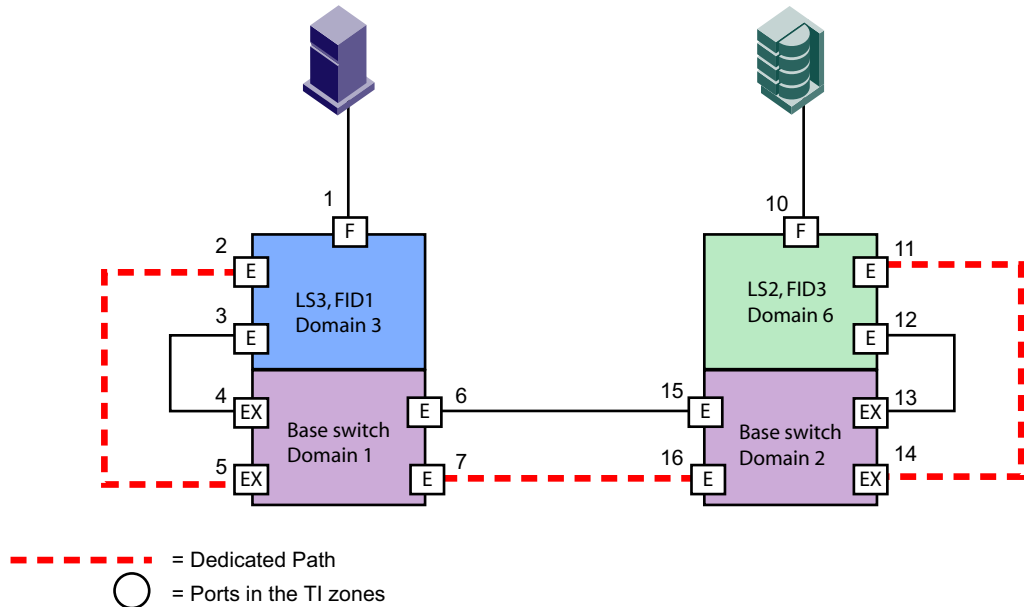


FIGURE 50 Example configuration for TI zones over FC routers in logical fabrics

[Figure 51](#) shows a logical representation of the configuration in [Figure 50](#). This SAN is similar to that shown in [Figure 42](#) on page 387 and you would set up the TI zones in the same way as described in “[Traffic Isolation Zoning over FC routers](#)” on page 386.




FIGURE 51 Logical representation of TI zones over FC routers in logical fabrics

Creating a TI zone

You create and modify TI zones using the **zone** command. Other zoning commands, such as **zoneCreate**, **allCreate**, and **cfgCreate**, cannot be used to manage TI zones.

When you create a TI zone, you can set the state of the zone to activated or deactivated. By default the zone state is set to activated; however, this does not mean that the zone is activated. After you create the TI zone, you must enable the current effective configuration to enforce the new TI zone, which is either activated or deactivated.

Virtual Fabric considerations: 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 should create a “dummy” configuration, as described in [“Creating a TI zone in a base fabric”](#) on page 404.

When you create a TI zone, you can enable or disable failover mode. By default, failover mode is enabled. If you want to change the failover mode after you create the zone, see [“Modifying TI zones”](#) on page 405.

If you are creating a TI zone with failover disabled, note the following:

- Ensure that the E_Ports of the TI zone correspond to valid paths; otherwise, the route might be missing for ports in that TI zone. You can use the **topologyShow** command to verify the paths.
- Ensure that sufficient non-dedicated paths through the fabric exist for all devices that are not in a TI zone; otherwise, these devices might become isolated.

See [“TI zone failover”](#) on page 380 for information about disabling failover mode.

Use the following procedure to create a TI zone. If you are creating a TI zone in a base fabric, use the procedure described in [“Creating a TI zone in a base fabric”](#) on page 404.

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

```
zone --create -t objtype [-o optlist] name -p "portlist"
```

Be aware of the ramifications if you create a TI zone with failover mode disabled. See [“TI zone failover”](#) on page 380 for information about disabling failover mode.

3. Perform the following steps if you have any TI zones with failover disabled. If all of your TI zones are failover-enabled, skip to [step 4](#).

- a. Change the failover option to failover enabled. This is a temporary change to avoid frame loss during the transition.

```
zone --add -o f name
```

- b. Enable the zones.

```
cfgenable "current_effective_configuration"
```

- c. Reset the failover option to failover disabled. Then continue with [step 4](#).

```
zone --add -o n name
```

4. Enter the **cfgEnable** command to reactivate your current effective configuration and enforce the TI zones.

```
cfgenable "current_effective_configuration"
```

Example TI zone creation

The following examples create a TI zone named “bluezone”, which contains E_Ports 1,1 and 2,4 and N_Ports 1,8 and 2,6.

To create a TI zone with failover enabled and in the activated state (default settings):

```
switch:admin> zone --create -t ti bluezone -p "1,1; 2,4; 1,8; 2,6"
```

To create a TI zone with failover enabled (the zone is set to the activated state by default):

```
switch:admin> zone --create -t ti -o f bluezone -p "1,1; 2,4; 1,8; 2,6"
```

To create a TI zone with failover disabled and the state set to activated:

```
switch:admin> zone --create -t ti -o an bluezone -p "1,1; 2,4; 1,8; 2,6"
```

To create a TI zone and set the state to deactivated (failover is enabled by default):

```
switch:admin> zone --create -t ti -o d bluezone -p "1,1; 2,4; 1,8; 2,6"
```

To create a TI zone with failover disabled and the state set to deactivated:

```
switch:admin> zone --create -t ti -o dn bluezone -p "1,1; 2,4; 1,8; 2,6"
```

To create a TI zone in the edge fabric with failover enabled and the state set to activated (default settings):

```
switch:admin> zone --create -t ti bluezone -p "1,1; 1,8; 2,-1; 3,-1"
```

To create a TI zone in the backbone fabric with failover enabled and the state set to activated (default settings):

```
switch:admin> zone --create -t ti backbonezone -p "10:00:00:04:1f:03:16:f2; 1,1; 1,4; 2,7; 2,1; 10:00:00:04:1f:03:18:f1, 10:00:00:04:1f:04:06:e2"
```

To create TI zones in a logical fabric, such as the one shown in [Figure 48](#) on page 400:

Log in to the logical switch FID1, Domain 7 and create a TI zone in the logical fabric with FID=1:

```
LS1> zone --create -t ti -o f "ti_zone1" -p "8,8; 8,1; 3,3; 3,10; 5,16; 5,8; 9,5; 9,9"
```

Then create a TI zone in the base fabric, as described in [“Creating a TI zone in a base fabric”](#).

Remember that your changes are not enforced until you enter the **cfgEnable** command, as shown here:

```
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 ...
```

Creating a TI zone in a base fabric

1. Connect to the switch and log in using an account with admin permissions.
2. Create a “dummy” zone configuration in the base fabric. For example:

```
zone --create "z1", "1,1"
cfgcreate "base_config", z1
```

3. Enter the **zone --create** command to create the TI zone in the base fabric:

```
zone --create -t objtype -o f name -p "portlist"
```

The disable failover option is not supported in base fabrics.

4. Perform the following steps if you have any TI zones with failover disabled. If all of your TI zones are failover-enabled, skip to [step 5](#).
 - a. Change the failover option to failover enabled. This is a temporary change to avoid frame loss during the transition.

```
zone --add -o f name
```

- b. Enable the zones.

```
cfgenable "current_effective_configuration"
```

- c. Reset the failover option to failover disabled. Then continue with [step 4](#).

```
zone --add -o n name
```

5. Enter the **cfgEnable** command to reactivate your current effective configuration and enforce the TI zones.

```
cfgenable "base_config"
```

Example

The following example creates TI zones in the base fabric shown in [Figure 49](#) on page 400:

```
BS_D1> zonecreate "z1", "1,1"
BS_D1> cfgcreate "base_cfg", z1
BS_D1> zone --create -t ti -o f "ti_zone2" -p "1,3; 1,10; 7,12; 7,14; 2,16;
2,8"
BS_D1> cfgenable "base_config"
```

Modifying TI zones

Using the **zone --add** command, you can add ports to an existing TI zone, change the failover option, or both. You can also activate or deactivate the TI zone.

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.

After you modify the TI zone, you must enable the current effective configuration to enforce the changes.

ATTENTION

If failover is disabled, do not allocate all ISLs in TI zones. Make sure sufficient non-dedicated paths exist through the fabric for all devices that are not in a TI zone. See [“TI zone failover”](#) on page 380 for additional information about disabling failover mode.

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.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter one of the following commands, depending on how you want to modify the TI zone.

- Enter the **zone --add** command to add ports or change the failover option for an existing TI zone. You can also activate or deactivate the zone.

```
zone --add [-o optlist] name -p "portlist"
```

```
zone --add -o optlist name [-p "portlist"]
```

- Enter the **zone --remove** command to remove ports from an existing TI zone.

```
zone --remove name -p "portlist"
```

Be aware of the ramifications if you disable failover mode. See [“TI zone failover”](#) on page 380 for information about disabling failover mode.

3. Perform the following steps if you have any TI zones with failover disabled. If all of your TI zones are failover-enabled, skip to [step 4](#).

- a. Change the failover option to failover enabled. This is a temporary change to avoid frame loss during the transition.

```
zone --add -o f name
```

- b. Enable the zones.

```
cfgenable "current_effective_configuration"
```

- c. Reset the failover option to failover disabled. Then continue with [step 4](#).

```
zone --add -o n name
```

4. Enter the **cfgEnable** command to reactivate your current effective configuration and enforce the TI zones.

```
cfgenable "current_effective_configuration"
```

13 Changing the state of a TI zone

Example of modifying a TI zone

To add port members to the existing TI zone bluezone:

```
switch:admin> zone --add bluezone -p "3,4; 3,6"
```

To add port members to the existing TI zone in a backbone fabric:

```
switch:admin> zone --add backbonezone -p "3,4; 3,6; 10:00:00:04:1f:03:16:f2;"
```

To disable failover on the existing TI zone bluezone:

```
switch:admin> zone --add -o n bluezone
```

To enable failover and add ports to TI zone greenzone:

```
switch:admin> zone --add -o f greenzone -p "3,4"
```

To remove ports from the TI zone bluezone:

```
switch:admin> zone --remove bluezone -p "3,4; 3,6"
```

Remember that your changes are not enforced until you enter the **cfgEnable** command.

Changing the state of a TI zone

You can change the state of a TI zone to activated or deactivated. Changing the state does not activate or deactivate the zone. After you change the state of the TI zone, you must enable the current effective configuration to enforce the change.

The TI zone must exist before you can change its state.

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"
```

Example of setting the state of a TI zone

To change the state of the existing TI zone bluezone to activated, type:

```
switch:admin> zone --activate bluezone
```

To change the state of the existing TI zone greenzone to deactivated, type:

```
switch:admin> zone --deactivate greenzone
```

Remember that your changes are not enforced until you enter the **cfgEnable** command.

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 405.

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 enforce the TI zones.

```
cfgenable "current_effective_configuration"
```

Example of deleting a TI zone

To delete the TI zone bluezone, type:

```
switch:admin> zone --delete bluezone
```

Remember that your changes are not enforced until you enter the **cfgEnable** command.

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**)

If you enter the **cfgShow** command to display information about all zones, the TI zones appear in the defined zone configuration only and do not appear in the effective zone configuration.

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

```
zone --show [ name ] [-ascending]
```

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
```

13 Troubleshooting TI zone routing problems

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
```

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 40](#) on page 385.

```
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
```

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 shown in [Figure 52](#). 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

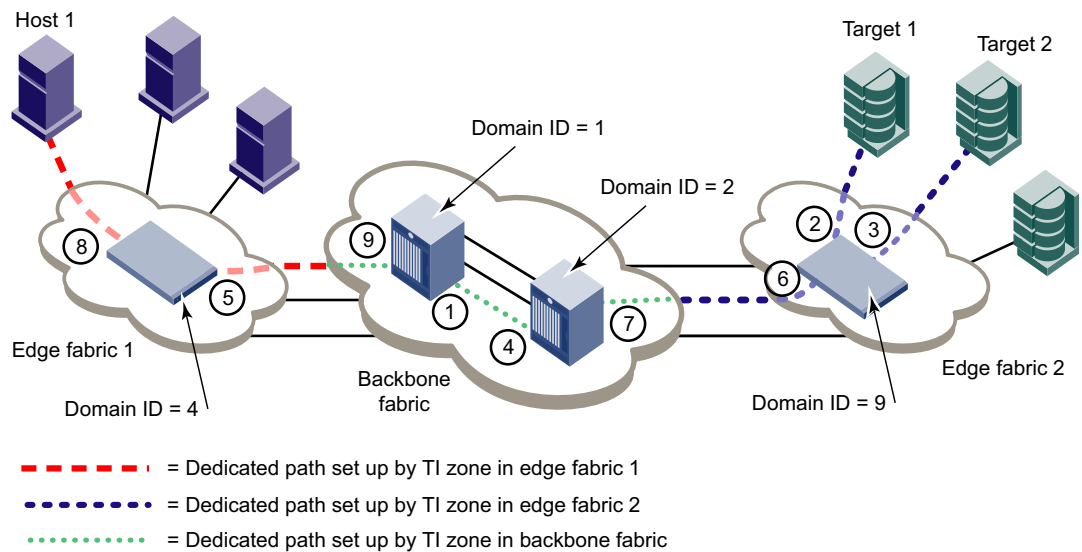


FIGURE 52 TI over FCR example

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. See [Chapter 26, "Using FC-FC Routing to Connect Fabrics,"](#) 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.

- a. Enter the following commands to create and display a TI zone:

```
BB_DCX_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"
BB_DCX_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
```

13 Setting up TI zones over FCR (sample procedure)

- b. Enter the following commands to reactivate your current effective configuration and enforce the TI zones.

```
BB_DCX_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
```

```
BB_DCX_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 ...
```

Optimizing Fabric Behavior

In this chapter

- [Adaptive Networking overview](#) 413
- [Ingress Rate Limiting](#) 414
- [QoS](#) 415
- [CS_CTL-based frame prioritization](#) 416
- [QoS zone-based traffic prioritization](#) 419
- [QoS zones](#) 419
- [Setting QoS zone-based traffic prioritization](#) 424
- [Setting QoS zone-based traffic prioritization over FC routers](#) 426
- [Disabling QoS zone-based traffic prioritization](#) 426

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. Refer to [Chapter 15, “Bottleneck Detection,”](#) for information about this feature.
- **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. Refer to [“Top Talker monitors”](#) on page 562 for more information about this feature.
- **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 [Chapter 13, “Traffic Isolation Zoning,”](#) 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 414 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 415 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:

- To reduce existing congestion in the network or proactively 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 or FL_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 reboots.

You should keep in mind the following considerations about Ingress Rate Limiting:

- Ingress Rate Limiting is applicable only to F_Ports and FL_Ports.
- QoS takes precedence over Ingress Rate Limiting.
- Ingress Rate Limiting is not enforced on trunked ports.

Virtual Fabrics considerations

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 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 or FL_Ports, then QoS zone-based traffic prioritization cannot be used between any devices connected to the F_Ports or FL_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 or FL_Ports that are defined in a QoS zone, CS_CTL-based frame prioritization takes precedence over the QoS zones.

[Table 71](#) 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 416 and [“QoS zone-based traffic prioritization”](#) on page 419 for detailed information about each type of prioritization scheme.

TABLE 71 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 or FL_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 QoS on F_Ports or FL_Ports. • Enable QoS on E_Ports. 	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

Starting in Fabric OS 7.2.0, 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.

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 or 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 72](#). The values are enabled by default to ensure backward compatibility.

TABLE 72 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 73](#).

TABLE 73 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 418, and [“Considerations for using CS_CTL-based frame prioritization”](#) on page 418.

Supported configurations for CS_CTL-based frame prioritization

CS_CTL-based frame prioritization is supported on all 8-Gbps and 16-Gbps platforms.

All switches in the fabric should be running Fabric OS v6.0.0 or later.

NOTE

If a switch is running a firmware version earlier than Fabric OS v6.0.0, 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

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:

```
portcfgqos --enable [slot/]port csctl_mode
```

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:

```
portcfgqos --disable [slot/]port csctl_mode
```

Alternatively, you can disable CS_CTL mode and set QoS to auto-enable (AE):

```
portcfgqos --default [slot/]port csctl_mode
```

Using CS_CTL auto mode at the chassis level

You can use the CS_CTL QoS mode options for the **configureChassis** CLI command to change the chassis-wide default mode (refer to [Table 72](#) on page 417), as in the following example.

```
switch:admin> configurechassis
Configure...
cfgload attributes (yes, y, no, n): [no] y
  Enforce secure config Upload/Download (yes, y, no, n): [no]
  Enforce signature validation for firmware (yes, y, no, n): [no]
  Add Suffix to the uploaded file name (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no] y
  Config Index (0 to ignore): (0..1000) [0]
system attributes (yes, y, no, n): [no] y
  system.blade.bladeFaultOnHwErrMsk: (0x0..0x7fffffff) [0x0]
  system.cpuLoad: (10..121) [121]
  system.i2cTurboCnfg: (0..2) [1]
fos attributes (yes, y, no, n): [no] y
  CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [1]
```

Set **CSCTL QoS Mode** to **1** to enable auto mode, establishing the settings shown in [Table 73](#) on page 417. Set **CSCTL QoS Mode** to **0** to disable auto mode and revert to default settings, shown in [Table 72](#) on page 417.

NOTE

As noted previously, this is a chassis-level configuration. It does not provide options to enable CS_CTL QoS on the ports.

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 72](#) on page 417) or the auto mode (refer to [Table 73](#) on page 417). 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 417.

QoS zone-based traffic prioritization

QoS zone-based traffic prioritization allows you to categorize the traffic flow between a host and a target as having a high, medium, or low priority, depending on the type of zone.

For example, you could assign online transaction processing (OLTP) to high priority and backup traffic to low priority.

All flows without QoS prioritization are considered medium priority.

High-, medium-, and low-priority flows are 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.

NOTE

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.

For new switches, QoS zone-based traffic prioritization is automatically enabled on the E_Ports, except for long-distance E_Ports. For long-distance E_Ports, you must manually enable QoS zone-based traffic prioritization.

If you upgrade to Fabric OS 7.2.0 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.

QoS zones

You assign high, medium, or low priority (QoS level) by configuring a QoS zone. A QoS zone is a special zone that indicates the priority of the traffic flow between a given host/target pair. By default, traffic in non-QoS zones is treated as medium priority.

The members of a QoS zone are the host/target pairs. 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 424 for some considerations.

A QoS zone has a special name to differentiate it from a regular zone. The name of the zone determines the priority of the traffic flow. The format of the QoS zone name is as follows:

For high priority: QOSHid_xxxxx

For medium priority: QOSMid_xxxxx

For low priority: QOSLid_xxxxx

In the QoS zone name format, *id* is a flow identifier that designates a specific virtual channel (VC) for the traffic flow and *xxxxx* is the user-defined portion of the name. For example, the following are valid QoS zone names:

QOSH3_HighPriorityTraffic

QOSL1_LowPriorityZone

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.

For high and low priority traffic, the flow *id* allows you to have control over the VC assignment and control over balancing the flows throughout the fabric. The *id* range is as follows:

- 1 through 5 for high-priority traffic, which corresponds to VCs 10 through 14.
- 1 through 4 for medium-priority traffic, which corresponds to VCs 2 through 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 through 2 for low-priority traffic, 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.

NOTE

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 421 for additional considerations when using QoS to prioritize traffic between device pairs in different edge fabrics.

For example, [Figure 53](#) shows a fabric with two hosts (H1, H2) and three targets (S1, S2, S3). The traffic prioritization is as follows:

- Traffic between H1 and S1 is high priority.
- Traffic between H1 and S3 and between H2 and S3 is low priority.
- All other traffic is medium priority, which is the default.




FIGURE 53 QoS traffic prioritization

For this fabric, you could set up the following QoS zones:

- QOSH_Zone1 Members: H1, S1
- QOSL_Zone3 Members: H1, H2, S3

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 [Figure 54](#), QoS should be enabled on the encircled E_Ports.

NOTE

By default, QoS is enabled on 8-Gbps or higher ports, except for long-distance 8-Gbps ports. QoS is disabled by default on all 4-Gbps ports and long-distance 8-Gbps ports.




FIGURE 54 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 [Figure 54](#) 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 switch 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 [Chapter 26, “Using FC-FC Routing to Connect Fabrics,”](#) for information about FC routers, phantom switches, and the FC-FC Routing Service.

To establish QoS over FC routers, 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 426 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-capable platforms only (Brocade 6505, 6510, 6520, M6505, 6547, and Brocade DCX 8510 Backbone family), 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.

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, [Figure 55](#) 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).

High-availability considerations for QoS zone-based traffic prioritization

If the standby control processor (CP) is running a Fabric OS version earlier than 6.3.0 and is synchronized with the active CP, then QoS zones using D,I notation cannot be created. If the standby CP is not synchronized or if no standby CP exists, then the QoS zone creation succeeds.

If QoS zones using D,I notation exist in either the defined or active configuration and the standby CP tries to synchronize with the active CP, the synchronization fails if the standby CP is running a Fabric OS version earlier than 6.3.0. Synchronization can succeed only if the QoS D,I zones are removed.

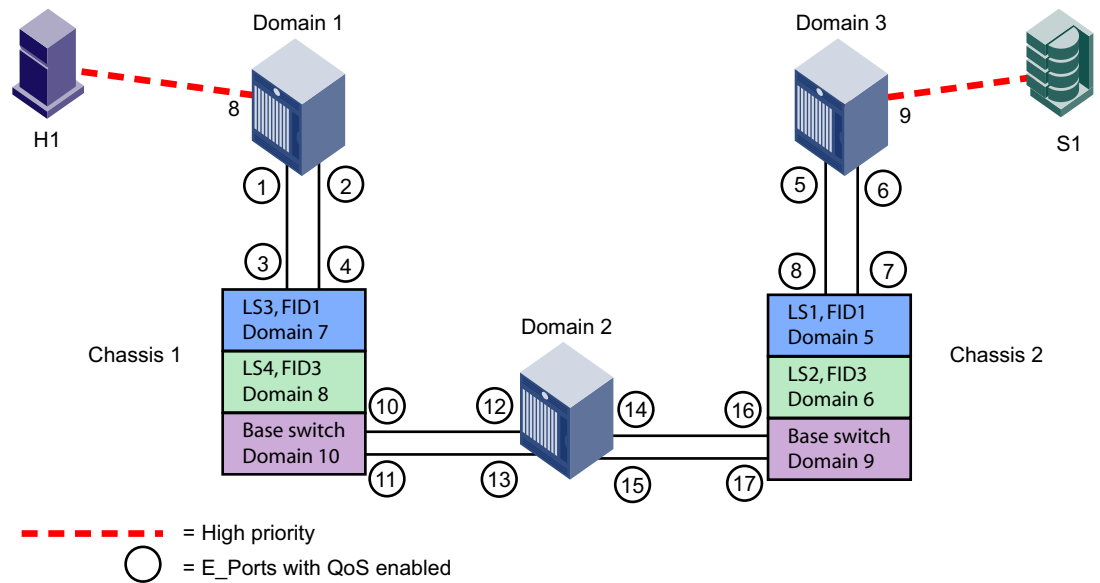


FIGURE 55 Traffic prioritization in a logical fabric

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 10-Gbps ISLs.
- Traffic prioritization is not supported on mirrored ports.
- Traffic prioritization is not supported over LSAN zones. 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 *Fabric OS Encryption Administrator's Guide* for information about redirection zones.
- Traffic prioritization is not supported in McDATA Fabric Mode (interopmode 2) or Open Fabric Mode (interopmode 3).
- You must be running Fabric OS v6.3.0 or later to create QoS zones that use D,I notation.
- 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 [Chapter 25, "Managing Long-Distance Fabrics,"](#) for information about buffer credit allocation in extended fabrics.
- If some ports in a trunk group have QoS enabled and some ports have QoS disabled, then two different trunks are formed, one with QoS enabled and one with QoS disabled.

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 "QOSHid_zonename", "member[; member...]"
```
 - For low-priority traffic, use the following syntax:


```
zonecreate "QOSLid_zonename", "member[; member...]"
```


The *id* range is from 1 through 5 for high-priority traffic, which corresponds to VCs 10 through 14. For low-priority traffic, the *id* range is from 1 through 2, which corresponds to VCs 8 and 9. The *id* is optional; if it is not specified, the virtual channels are allocated 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"
```

6. 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 port again, the port is *not* toggled, because the configuration did not change.

Example

```
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    zone1; 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
zone:  zone1  10:00:00:00:10:00:00:00; 10:00:00:00:20:00:00:00;
      10:00:00:00:30:00:00:00; 10:00:00:00:40:00:00:00

Effective configuration:
No Effective configuration: (No Access)

sw0:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Any changes made on the Effective configuration will not
```

14 Setting QoS zone-based traffic prioritization over FC routers

```
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

1. Connect to the switch in the edge fabric and log in using an account with admin permissions.
2. Create QoS zones in the edge fabric.
The QoS zones must have WWN members only, and not D,I members. Refer to [“Setting QoS zone-based traffic prioritization”](#) on page 424 for instructions.
3. Create LSAN zones in the edge fabric.
Refer to [“Controlling device communication with the LSAN”](#) on page 621 for instructions.
4. Enter the **portCfQoS** command to enable QoS on the E_Ports.
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 **portCfQoS** command to enable QoS on the EX_Ports.

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.
3. Enter the **portCfQoS** command to disable QoS on the E_Ports.

Bottleneck Detection

In this chapter

- Bottleneck detection overview 427
- Supported configurations for bottleneck detection 429
- Enabling bottleneck detection on a switch 431
- Displaying bottleneck detection configuration details 431
- Setting bottleneck detection alerts 433
- Changing bottleneck detection parameters 435
- Advanced bottleneck detection settings 439
- Excluding a port from bottleneck detection 440
- Displaying bottleneck statistics 442
- Disabling bottleneck detection on a switch 442

Bottleneck detection overview

A *bottleneck* is a port in the fabric where frames cannot get through as fast as they should. In other words, a bottleneck is a port where the offered load is greater than the achieved egress throughput. Bottlenecks can cause undesirable degradation in throughput on various links. When a bottleneck occurs at one place, other points in the fabric can experience bottlenecks as the traffic backs up.

Bottleneck detection is configured on a per-switch basis, with optional per-port exclusions. Bottleneck detection is disabled by default. The best practice is to enable bottleneck detection on all switches in the fabric, and leave it on to continuously gather statistics. Bottleneck detection does not require a license.

The bottleneck detection feature enables you to do the following:

- Prevent degradation of throughput in the fabric.
The bottleneck detection feature alerts you to the existence and locations of devices that are causing latency. If you receive alerts for one or more F_Ports, use the CLI to check whether these F_Ports have a history of bottlenecks.
- Reduce the time it takes to troubleshoot network problems.
If you notice one or more applications slowing down, you can determine whether any latency devices are attached to the fabric and where. You can use the CLI to display a history of bottleneck conditions on a port. If the CLI shows above-threshold bottleneck severity, you can narrow the problem down to device latency rather than problems in the fabric.

You can use the bottleneck detection feature with other Adaptive Networking features to optimize the performance of your fabric. For example, you can do the following:

- If the bottleneck detection feature detects a latency bottleneck, you can use TI zones or QoS SID/DID traffic prioritization 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.

Types of bottlenecks

The bottleneck detection feature detects two types of bottlenecks:

- Latency bottleneck
- Congestion bottleneck

A *latency bottleneck* is a port where the offered load exceeds the rate at which the other end of the link can continuously accept traffic, but does not exceed the physical capacity of the link. This condition can be caused by a device attached to the fabric that is slow to process received frames and send back credit returns. A latency bottleneck caused by such a device can spread through the fabric and can slow down unrelated flows that share links with the slow flow.

By default, bottleneck detection detects latency bottlenecks that are severe enough that they cause 98 percent loss of throughput. This default value can be modified to a different percentage.

A *congestion bottleneck* is a port that is unable to transmit frames at the offered rate because the offered rate is greater than the physical data rate of the line. For example, this condition can be caused by trying to transfer data at 8 Gbps over a 4 Gbps ISL.

You can use the **bottleneckMon** command to configure separate alert thresholds for congestion and latency bottlenecks.

Advanced settings allow you to refine the criterion for defining latency bottleneck conditions to allow for more (or less) sensitive monitoring at the sub-second level. For example, you would use the advanced settings to change the default value of 98 percent for loss of throughput. Refer to [“Advanced bottleneck detection settings”](#) on page 439 for specific details.

If a bottleneck is reported, you can investigate and optimize the resource allocation for the fabric. Using the zone setup and Top Talkers, you can also determine which flows are destined to any affected F_Ports.

How bottlenecks are reported

Bottleneck detection uses the concept of an *affected second* when determining whether a bottleneck exists on a port. Each second is marked as being affected or unaffected by a latency or congestion bottleneck, based on certain criteria.

The bottleneck detection feature maintains two histories of affected seconds for each port—one history for latency bottlenecks and another for congestion bottlenecks. A history is maintained for a maximum of three hours for each port. You can view the history using the **bottleneckmon --show** command, as described in [“Displaying bottleneck statistics”](#) on page 442.

Bottlenecks are also reported through RASlog alerts and SNMP traps. These two alerting mechanisms cannot be turned on and off independently.

You can use the **bottleneckMon** command to specify the following alerting parameters:

- Whether alerts are to be sent when a bottleneck condition is detected
- The size of the time window to look at when determining whether to alert
- How many affected seconds are needed to generate the alert
- How long to stay quiet after an alert
- If an enabled alert is for congestion, for latency, or for both

NOTE

Changing alerting parameters affects RASlog alerting as well as SNMP traps.

For more detailed information on the **bottleneckMon** command, refer to the *Fabric OS Command Reference*.

Supported configurations for bottleneck detection

The following configuration rules apply to bottleneck detection:

- Bottleneck detection is supported only on Fibre Channel ports and FCoE F_Ports.
- Bottleneck detection is supported only on the following port types:
 - E_Ports
 - EX_Ports
 - F_Ports
 - FL_Ports
- F_Port and E_Port trunks are supported.
- Long distance E_Ports are supported.
- SIM ports are not supported.
- Bottleneck detection is supported on 4-Gbps, 8-Gbps, and 16-Gbps platforms, including 10-Gbps speeds.
- Bottleneck detection is supported in Access Gateway mode.
- Bottleneck detection is supported whether Virtual Fabrics is enabled or disabled. In VF mode, bottleneck detection is supported on all fabrics, including the base fabric. Refer to [“Virtual Fabrics considerations for bottleneck detection”](#) on page 430 for additional information on using bottleneck detection in VF mode.

Limitations of bottleneck detection

The bottleneck detection feature detects latency bottlenecks only at the point of egress, not ingress. For example, for an E_Port, only the traffic egressing the port is monitored. For FCoE ports at the FC-DCB (Data Center Bridging) boundary, traffic going from FC to DCB is monitored, but traffic going from DCB to FC is not monitored.

ATTENTION

Latency bottleneck detection is *not* recommended for link utilizations above 85 percent.

High availability considerations for bottleneck detection

The bottleneck detection configuration is maintained across a failover or reboot; however, bottleneck statistics collected are lost.

Upgrade and downgrade considerations for bottleneck detection

The bottleneck detection configuration is persistent across firmware upgrades and downgrades.

The sub-second latency criterion parameter settings are not preserved on a downgrade to firmware versions earlier than Fabric OS 7.0.0. If you downgrade and then upgrade back to Fabric OS 7.0.0, the settings revert to their default values.

Trunking considerations for bottleneck detection

A trunk behaves like a single port. Both latency and congestion bottlenecks are reported on the master port only, but apply to the entire trunk.

For masterless trunking, if the master port goes offline, the new master acquires all the configurations and bottleneck history of the old master and continues with bottleneck detection on the trunk.

Virtual Fabrics considerations for bottleneck detection

Bottleneck detection is supported in both VF and non-VF modes.

In VF mode, if a port on which bottleneck detection is enabled is moved out of a logical switch, any per-port configurations are retained by the logical switch. The per-port configuration does not propagate outside of the logical switch. If the port is returned to the logical switch, the previous per-port configurations are automatically set for the port. Refer to [“Changing bottleneck detection parameters”](#) on page 435 for more information about changing per-port configurations.

In logical fabrics, bottleneck detection is not performed on logical ISLs.

Because a base fabric carries traffic from multiple logical fabrics, bottlenecks reported in the base fabric can be caused by a mixture of traffic from multiple logical fabrics or by traffic from a single logical fabric. It is not possible to attribute a base fabric bottleneck to the exact logical fabric causing it. Dedicated ISLs are exclusive to one logical fabric, and any bottleneck on a dedicated ISL E_Port pertains entirely to the traffic of that logical fabric.

Access Gateway considerations for bottleneck detection

If bottleneck detection is enabled on a logical switch with some F_Ports connected to an Access Gateway, you do not get information about which device is causing a bottleneck, because devices are not directly connected to the Access Gateway. To detect bottlenecks on an Access Gateway, enable bottleneck detection on the Access Gateway to which the devices are actually connected.

Enabling bottleneck detection on a switch

Enabling bottleneck detection permits both latency and congestion detection.

Bottleneck detection is enabled on a switch basis. It is recommended that you enable bottleneck detection on every switch in the fabric. If you later add additional switches, including logical switches, to the fabric, be sure to enable bottleneck detection on those switches as well.

When you enable bottleneck detection on a switch, the settings are applied to all eligible ports on that switch. If ineligible ports later become eligible or, in the case of a logical switch, if ports are moved to the logical switch, bottleneck detection is automatically applied to those ports.

You can later override these settings on a per-port basis, as described in [“Changing bottleneck detection parameters”](#) on page 435.

Use the following procedure to enable bottleneck detection.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **bottleneckmon --enable** command to enable bottleneck detection on all eligible ports on the switch.

By default, alerts are not sent unless you specify the **alert** parameter; however, you can view a history of bottleneck conditions for the port as described in [“Displaying bottleneck statistics”](#) on page 442.

3. Repeat [step 1](#) and [step 2](#) on every switch in the fabric.

NOTE

A best practice is to use the default values for the alerting and sub-second latency criterion parameters.

The following example enables bottleneck detection on the switch with alerts using default values for thresholds and time, and is the recommended manner of enabling bottleneck detection

```
switch:admin> bottleneckmon --enable -alert
```

The following example enables bottleneck detection on the switch without alerts. In this case, even though alerts are not delivered, you can still view the bottleneck history using either the CLI or BNA.

```
switch:admin> bottleneckmon --enable
```

Displaying bottleneck detection configuration details

Use the following procedure to display the bottleneck detection configuration details:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **bottleneckmon --status** command to display the details of bottleneck detection configuration for the switch, which includes the following details:
 - Whether the feature is enabled
 - Switch-wide parameters
 - Per-port overrides, if any
 - Excluded ports

15 Displaying bottleneck detection configuration details

The following initials in the “Per-port overrides for alert parameters,” section of the output indicate which alerts have been set:

- C indicates a congestion alert has been set.
- L indicates a latency alert has been set.
- Y indicates both alerts are set.
- N indicates no alerts are set.

The following examples show the status of different bottleneck alerts.

Example of status showing that bottleneck detection is not enabled

```
switch:admin> bottleneckmon --status  
Bottleneck detection - Disabled
```

Refer to [“Enabling bottleneck detection on a switch”](#) on page 431 for instructions on enabling bottleneck detection.

Example of status output showing that congestion and latency alerts are enabled

Example

```
switch:admin> bottleneckmon --status  
Bottleneck detection - Enabled  
=====
```

Switch-wide sub-second latency bottleneck criterion:
=====

Time threshold	- 0.800
Severity threshold	- 50.000

Switch-wide alerting parameters:
=====

Alerts	- Yes
Latency threshold for alert	- 0.100
Congestion threshold for alert	- 0.800
Averaging time for alert	- 300 seconds
Quiet time for alert	- 300 seconds

Example of status output showing that only a congestion alert at the switch level has been set

```
switch:admin> bottleneckmon --status  
Bottleneck detection - Enabled  
=====
```

Switch-wide sub-second latency bottleneck criterion:
=====

Time threshold	- 0.800
Severity threshold	- 50.000

Switch-wide alerting parameters:
=====

Alerts	- Congestion only
Congestion threshold for alert	- 0.800
Averaging time for alert	- 300 seconds
Quiet time for alert	- 300 seconds

Per-port overrides for alert parameters:
=====

Port	Alerts?	LatencyThresh	CongestionThresh	Time (s)	QTime(s)
1	Y	0.100	0.800	300	300
2	C	--	0.800	600	600
3	L	0.100	--	300	300
4	N	--	--	--	--

NOTE
If there are no per-port overrides, “Per-port overrides for alert parameters” section is not displayed.

Setting bottleneck detection alerts

You can configure Fabric OS to log per-port alerts based on the latency and congestion history of the port. Alerts are generated based on the number of affected seconds over a specified period of time. If the number of affected seconds is higher than the threshold, an alert is generated. This evaluation is done independently for latency and congestion.

NOTE
A congestion bottleneck detection alert is generated whenever a frame timeout occurs irrespective of the number of affected seconds in the observation window.

The following **bottleneckmon -alert** parameters determine whether an alert is generated and the reason for the alert.

The **-time** parameter specifies the time window. For this example, **-time** equals 12 seconds.

The **-cthresh** and **-lthresh** parameters specify the thresholds on number of affected seconds that trigger alerts for congestion and latency bottlenecks, respectively.

For example, [Figure 56](#) shows an interval of 12 seconds, in which 6 seconds are affected by a congestion bottleneck and 3 seconds are affected by a latency bottleneck. This example uses the default values for these parameters, where **-cthresh** = 0.8 (80%) and **-lthresh** = 0.1 (10%).

The following command results in the example illustrated in [Figure 56](#):

bottleneckmon -alert -time 12 -cthresh 0.8 -lthresh 0.1




FIGURE 56 Affected seconds for bottleneck detection

For this time window, 50 percent of the seconds (6 out of 12 seconds) are affected by congestion. This is below the threshold of 80 percent, so an alert would not be generated for a congestion bottleneck.

For the same time window, 25 percent of the seconds (3 out of 12 seconds) are affected by latency. This exceeds the threshold of 10 percent, so an alert would be generated for a latency bottleneck.

Setting both a congestion alert and a latency alert

Entering the `bottleneckmon --enable -alert` command enables both alerts using the default alert values.

This example enables both alerts and shows their values.

```
switch:admin> bottleneckmon --enable -alert
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold           - 0.800
Severity threshold       - 50.000

Switch-wide alerting parameters:
=====
Alerts                   - Yes
Latency threshold for alert - 0.100
Congestion threshold for alert - 0.800
Averaging time for alert  - 300 seconds
Quiet time for alert      - 300 seconds
```

Setting a congestion alert only

Entering the `bottleneckmon --enable -alert=congestion` command enables a congestion alert. This example enables a congestion alert and shows its values.

```
switch:admin> bottleneckmon --enable -alert=congestion
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold           - 0.800
Severity threshold       - 50.000

Switch-wide alerting parameters:
=====
Alerts                   - Congestion only
Congestion threshold for alert - 0.800
Averaging time for alert  - 300 seconds
Quiet time for alert      - 300 seconds
```

Setting a latency alert only

Entering the `bottleneckmon --enable -alert=latency` command enables a congestion alert. This example enables a latency alert and shows its values.

```
switch:admin> bottleneckmon --enable -alert=latency
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold           - 0.800
Severity threshold       - 50.000

Switch-wide alerting parameters:
=====
Alerts                   - Latency only
Latency threshold for alert - 0.100
Averaging time for alert  - 300 seconds
Quiet time for alert     - 300 seconds
```

Changing bottleneck detection parameters

When you enable bottleneck detection, you can configure switch-wide or port-specific alerting parameters. The alerting parameters indicate whether alerts are sent, and the threshold, time, and quiet-time options, as well as the sub-second latency criterion for ports.

After you enable bottleneck detection, you can change the alerting parameters for the entire switch or only for individual ports. For example, you can change the latency threshold for only port 47 without affecting any other port. You can also change the parameters on ports that have been excluded from bottleneck detection. For a trunk, you can change the parameters only on the master port.

Alert-related parameters can only be specified with `--config` when `-alert` is specified. This is because `-noalert` is assumed if `-alert` is not specified, and `-noalert` cancels all alert-related parameters. As long as you want alerts, you must include the exact form of alert (`-alert`, `-alert=congestion`, or `-alert=latency`) in every `--config` operation, even if alerts are already enabled.

The retention of settings applies only to the `--config` command, not to `--enable`.

An `--enable` operation behaves as if there is no pre-existing user configuration. If the `--enable` command does not include `-alert`, but does specify alert-related parameters, that command will fail.

NOTE

Entering the `--config` command changes only those settings specified in the command; all others are left alone. The only exceptions are the `-alert` (restores alerts using recorded values) or `-noalert` (disables all alerts) switches. If you want alerts, you must specify what you want as the `-alert` value for every `bottleneckmon --config -alert` command.

Use the following procedure to configure the bottleneck detection parameters.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the `bottleneckmon --config` command to set the alerting and sub-second latency criterion parameters.

Use the **-alert** parameter to enable congestion and latency alerts. Use the **-cthresh** parameter to specify the severity threshold for congestion that triggers an alert. Use the **-lthresh** parameter to specify the severity threshold for latency that triggers an alert. Use the **-time** parameter to specify the time window in seconds over which the percentage of seconds affected by bottleneck conditions is computed and compared with the threshold. Use the **-qtime** parameter to specify the minimum number of seconds between consecutive alerts. Refer to the *Fabric OS Command Reference* for more information.

To remove any port-specific alerting and sub-second latency criterion parameters and revert to the switch-wide parameters, enter the **bottleneckmon --configclear** command. To remove and erase all bottleneck alerts and their criteria, enter **bottleneckmon --disable**. Refer to [“Disabling bottleneck detection on a switch”](#) on page 442 for more details.

Examples of applying and changing bottleneck detection parameters

The following examples show how to change various bottleneck detection parameters, and how the changes made are retained when the next set of changes is made. For each example, after the configuration command is run, the **bottleneckmon --status** command is run to show the new settings, which are **bolded just** for the examples.

Example 1: Setting time window, quiet time, and threshold values for an entire switch

This example sets the time window to 150 seconds, the quiet time to 150 seconds, the congestion threshold to 0.7 (70%) and the latency threshold to 0.2 (20%) for the entire switch.

```
switch:admin> bottleneckmon --config -alert -time 150 -qtime 150 -cthresh 0.7
-lthresh 0.2
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold           - 0.800
Severity threshold       - 50.000

Switch-wide alerting parameters:
=====
Alerts                   - Yes
Latency threshold for alert - 0.200
Congestion threshold for alert - 0.700
Averaging time for alert      - 150 seconds
Quiet time for alert         - 150 seconds
```

Example 2: Changing time window value for an entire switch

This example changes the time window value to 200 seconds for the entire switch.

```
switch:admin> bottleneckmon --config -alert -time 200
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold           - 0.800
Severity threshold       - 50.000
```

```

Switch-wide alerting parameters:
=====
Alerts                - Yes
Latency threshold for alert - 0.200
Congestion threshold for alert - 0.700
Averaging time for alert - 200 seconds
Quiet time for alert   - 150 seconds

```

Example 3: Disabling bottleneck detection alerts for a port

This example disables bottleneck detection alerts for port 46 only.

```

switch:admin> bottleneckmon --config -noalert 46
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

```

```

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold          - 0.800
Severity threshold     - 50.000

```

```

Switch-wide alerting parameters:
=====
Alerts                - Yes
Latency threshold for alert - 0.200
Congestion threshold for alert - 0.700
Averaging time for alert - 200 seconds
Quiet time for alert   - 150 seconds

```

Per-port overrides for alert parameters:

```

=====
Port    Alerts? LatencyThresh  CongestionThresh    Time (s)    QTime (s)
=====
46      N         --                --                --            --

```

Example 4: Selecting latency-only alerts and changing the latency threshold value for a port

This example changes the alerts to latency-only and the latency threshold value to 75 percent, both on port 47 only.

```

switch:admin> bottleneckmon --config -alert=latency -lthresh 0.75 47
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

```

```

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold          - 0.800
Severity threshold     - 50.000

```

```

Switch-wide alerting parameters:
=====
Alerts                - Yes
Latency threshold for alert - 0.200
Congestion threshold for alert - 0.700
Averaging time for alert - 200 seconds
Quiet time for alert   - 150 seconds

```

15 Changing bottleneck detection parameters

```
Per-port overrides for alert parameters:
=====
Port      Alerts? LatencyThresh  CongestionThresh      Time (s)      QTime (s)
=====
46        N        --              --                    --            --
47        L        0.750          --                    200           150
```

Example 5: Changing the latency time value for a port

This example changes the time value to 250 seconds for port 47 only. Notice that the command must include **-alert=latency** to preserve the latency-only alerts configured in the previous example. In general, **-alert** must be specified (with **=latency** or **=congestion** if desired) on every **--config** command when alerts are desired.

```
switch:admin> bottleneckmon --config -alert=latency -time 250 47
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====
```

```
Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold          - 0.800
Severity threshold     - 50.000
```

```
Switch-wide alerting parameters:
=====
Alerts                  - Yes
Latency threshold for alert - 0.200
Congestion threshold for alert - 0.700
Averaging time for alert - 200 seconds
Quiet time for alert    - 150 seconds
```

```
Per-port overrides for alert parameters:
=====
Port      Alerts? LatencyThresh  CongestionThresh      Time (s)      QTime (s)
=====
46        N        --              --                    --            --
47        L        0.750          --                    250           150
```

Example 6: Clearing bottleneck detection override values from ports

This example removes any changed bottleneck detection parameter values from ports 46 and 47. Notice that the “Per-port overrides for alert parameters” section of the output is not displayed because there are no per-port overrides.

```
switch:admin> bottleneckmon --configclear 46-47
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====
```

```
Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold          - 0.800
Severity threshold     - 50.000
```

```
Switch-wide alerting parameters:
=====
Alerts                - Yes
Latency threshold for alert - 0.200
Congestion threshold for alert - 0.700
Averaging time for alert - 200 seconds
Quiet time for alert - 150 seconds
```

Adjusting the frequency of bottleneck alerts

Depending on the circumstances, a problematic switch or port may be triggering alerts more frequently than desired. The **-qtime** parameter can be used to throttle alerts by specifying the minimum number of seconds between consecutive alerts. Thresholds are configured separately for each type of bottleneck and statistical data are collected independently for each condition. This parameter applies individually to each type of bottleneck detection, so there can be one latency alert and one congestion alert in one quiet time.

Example of setting quiet time

This example sets a latency threshold of 0.8 for a time window of 30 seconds, and specifies that an alert should be sent when 80 percent (0.8) of the one-second samples over any period of 30 seconds is affected by latency bottleneck conditions; the system then waits 60 seconds before issuing the next alert (assuming that there is one).

```
switch:admin> bottleneckmon --enable -lthresh 0.8 -time 30 -qtime 60
-alert=latency
```

Advanced bottleneck detection settings

You can use the sub-second latency criterion parameters to refine the criterion for determining whether a second is marked as affected by latency bottlenecks. For example, you may want to use the sub-second latency criterion parameters in the following cases:

- You notice an under-performing application, but do not see any latency bottlenecks detected. You can temporarily increase the sub-second sensitivity of latency bottleneck detection on the specific F_Ports for this application.
- You want greater-than-default (sub-second) latency sensitivity on your fabric, so you set sub-second latency criterion parameters at the time you enable bottleneck detection.
- You want to reduce the number of alerts you are receiving about known latency bottlenecks in the fabric, so you temporarily decrease the sub-second latency sensitivity on these ports.
- You have a latency bottleneck on an ISL that is not at the edge of the fabric.

The sub-second latency criterion parameters are always applicable. These parameters affect alerts and, even if alerting is not enabled, they affect the history of bottleneck statistics.

The following sub-second latency criterion parameters are shown with the default values in parentheses:

- **-subsectimethresh** (0.8) is similar to the **-lthresh** alerting parameter, except on a sub-second level. The default value of 0.8 means that at least 80 percent of a second must be affected by latency for the second to be marked as affected.

15 Excluding a port from bottleneck detection

- **-lsubsecsevthresh** (50) specifies the factor by which throughput must drop in a second for that second to be considered affected by latency. The default value of 50 means that the observed throughput in a second must be no more than 1/50th the capacity of the port for that second to be counted as an affected second. 1/50th of capacity equals 2 percent of capacity, which translates to 98 percent loss of throughput.

Sub-second latency criterion parameters apply only to latency bottlenecks and not congestion bottlenecks.

When you enable bottleneck detection, you can specify switch-wide sub-second latency criterion parameters. After you enable bottleneck detection, you can change the sub-second latency criterion parameters only on a per-port basis. You cannot change them on the entire switch, as you can with alerting parameters, unless you disable and then re-enable bottleneck detection.

Changing the sub-second latency criterion parameters on specific ports causes an interruption in the detection of bottlenecks on those ports, which means the history of bottlenecks is lost on these ports. Also note the following behaviors if you change the sub-second latency criterion parameters:

- Traffic through these ports is not affected.
- History of latency bottlenecks and congestion bottlenecks is lost on these ports. Other ports are not affected, however.
- The interruption occurs whether you set or clear per-port overrides on the sub-second latency criterion parameters.
- Because of the interruption, you can never have an alert for a port such that the alert spans periods of time with different sub-second latency criteria on that port.

Excluding a port from bottleneck detection

When you exclude a port from bottleneck detection, no data is collected from the port and no alerts are generated for the port. All statistics history for the port is discarded.

Alerting parameters for the port are preserved, so if you later include the port for bottleneck detection, the alerting parameters are restored.

Per-port exclusions may be needed if, for example, a long-distance port is known to be a bottleneck because of credit insufficiency. In general, however, per-port exclusions are not recommended.

For trunking, if you exclude a slave port from bottleneck detection, the exclusion has no effect as long as the port is a trunk slave. The exclusion takes effect only if the port becomes a trunk master or leaves the trunk.

Use the following procedure to exclude a port from bottleneck detection.

1. Connect to the switch to which the target port belongs and log in using an account with admin permissions.
2. Enter the **bottleneckmon --exclude** command to exclude the port from bottleneck detection.
To later include the port, enter the **bottleneckmon --include** command.

Example showing how to exclude a single port from bottleneck detection

The following example excludes port 7 only from bottleneck detection. Refer to [“Disabling bottleneck detection on a switch”](#) on page 442 for more information.

NOTE

Excluding the master port excludes the entire trunk, even if individual slave ports are not excluded.

```
switch:admin> bottleneckmon --exclude 7
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold           - 0.800
Severity threshold       - 50.000

Switch-wide alerting parameters:
=====
Alerts                   - Yes
Latency threshold for alert - 0.200
Congestion threshold for alert - 0.700
Averaging time for alert  - 200 seconds
Quiet time for alert     - 150 seconds

Per-port overrides for alert parameters:
=====
Port   Alerts? LatencyThresh  CongestionThresh  Time (s)  QTime (s)
=====
46     N        --                --              --        --
47     L        0.750            --              250       150

Excluded ports:
=====
Port
====
7
```

Example showing how to re-include bottleneck detection for a port

This example restores bottleneck detection for port 7. Notice that the “Excluded ports” section of the output is not displayed because there are no excluded ports.

```
switch:admin> bottleneckmon --include 7
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold           - 0.800
Severity threshold       - 50.000

Switch-wide alerting parameters:
=====
Alerts                   - Yes
Latency threshold for alert - 0.200
Congestion threshold for alert - 0.700
Averaging time for alert  - 200 seconds
Quiet time for alert     - 150 seconds
```

```

Per-port overrides for alert parameters:
=====
Port      Alerts? LatencyThresh  CongestionThresh      Time (s)      QTime (s)
=====
46        N        --              --                    --            --
47        L        0.750          --                    250           150
    
```

Displaying bottleneck statistics

You can use the **bottleneckmon --show** command to display a history of bottleneck conditions, for up to three hours. This command has several display options:

- Display only latency bottlenecks, only congestion bottlenecks, or both combined.
- Display bottleneck statistics for a single port, bottleneck statistics for all ports on the switch, or a list of ports affected by bottleneck conditions.
- Continuously update the displayed data with fresh data.

Use the following procedure to display the bottleneck statistics.

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

Example of displaying the bottleneck history in 5-second windows over a period of 30 seconds

In this example, the definition of *bottlenecked ports* is any port that had a bottleneck occur during any second in the corresponding interval.

```

switch:admin> bottleneckmon --show -interval 5 -span 30

=====
                        Wed Jan 13 18:54:35 UTC 2010
=====
List of bottlenecked ports in most recent interval:
5
=====

From              To              Number of
=====          =====          bottlenecked ports
Jan 13 18:54:05   Jan 13 18:54:10   1
Jan 13 18:54:10   Jan 13 18:54:15   2
Jan 13 18:54:15   Jan 13 18:54:20   1
Jan 13 18:54:20   Jan 13 18:54:25   1
Jan 13 18:54:25   Jan 13 18:54:30   0
Jan 13 18:54:30   Jan 13 18:54:35   0
    
```

Disabling bottleneck detection on a switch

When you disable bottleneck detection on a switch, all bottleneck configuration details are discarded, including the list of excluded ports and non-default values of alerting parameters.

Use the following procedure to disable bottleneck detection.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **bottleneckmon --disable** command to disable bottleneck detection on the switch.

Example of disabling bottleneck detection on a switch

```
switch:admin> bottleneckmon --disable  
switch:admin> bottleneckmon --status  
Bottleneck detection - Disabled
```

15 Disabling bottleneck detection on a switch

In-flight Encryption and Compression

In this chapter

- In-flight encryption and compression overview 445
- Configuring in-flight encryption and compression on an EX_Port 450
- Configuring in-flight encryption and compression on an E_Port 451
- Viewing the encryption and compression configuration 452
- Configuring and enabling authentication for in-flight encryption 453
- Enabling in-flight encryption 455
- Enabling in-flight compression 456
- Disabling in-flight encryption 456
- Disabling in-flight compression 457

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 initially disabled for all ports on a switch.

NOTE

No license is required to configure and enable in-flight encryption or compression.

Both ends of the ISL must terminate in 16 Gbps-capable FC ports.

Encryption and compression can be enabled at the same time, or you can enable either encryption or compression selectively. [Figure 57](#) 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 57 Encryption and compression on 16 Gbps ISLs

Supported ports for in-flight encryption and compression

The in-flight encryption and compression features are supported only on E_Ports and EX_Ports, and only on the Brocade 6510 and 6520 switches, 16 Gbps embedded switches, and the Brocade DCX 8510 Backbone family.

The ports can run at any speed, but must be 16 Gbps-capable.

Encryption and compression are also compatible with the following features:

- E_Ports or EX_Ports with trunking, QoS, or long distance features enabled.
- Flow control modes: R_RDY, VC_RDY, and EXT_VC_RDY.
- XISL ports in VF mode.
- FCP data frames and non-FCP data frames except ELS and BLS frames.

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 and compression restrictions

- Ports must be 16 Gbps-capable, although port speed can be any configurable value.
- Configuration is dynamic based on port speed. Refer to [Table 74](#) on page 447 for specific details about the number of ports supported for encryption and compression.
- The devices at either end of the ISL must run Fabric OS 7.0.0 or later software.
- Only E_Ports, EX_Ports, and XISL ports (in VF mode) support encryption or compression. ICL ports do not support encryption or compression.
- Encryption is not supported in FIPS mode. In-flight encryption is not FIPS-compliant.
- The payload size of a frame is restricted to 2048 bytes.

Bandwidth and port limits for in-flight encryption and compression

Fabric OS supports up to 32 Gbps of data encryption and 32 Gbps of data compression per 16 Gbps-capable FC platform. This limits the number of ports that can have these features enabled at any one time.

The port speed affects the number of supported ports. The slower the speed, the more ports are supported. In general, at 16 Gbps, the number of supported ports is 2 per ASIC or trunk.

The following table shows some examples of how port speed affects the number of supported ports for different implementations.

TABLE 74 Number of ports supported for in-flight encryption and compression at various port speeds

	Blades (FC16-32, FC16-48) ¹		
Port speed	Encryption only	Compression only	Encryption and compression
16 Gbps	4 ports	4 ports	4 ports
10 Gbps	6 ports	6 ports	6 ports
8/4/2 Gbps	8 ports	8 ports	8 ports
Auto-negotiate (AN)	4 ports	4 ports	4 ports
	6510 Fixed-port switches and 16 Gbps embedded switches ²		
16 Gbps	2 ports	2 ports	2 ports
10 Gbps	3 ports	3 ports	3 ports
8/4/2 Gbps	4 ports	4 ports	4 ports
Auto-negotiate (AN)	2 ports	2 ports	2 ports
	6520 Fixed-port switches ³		
16 Gbps	8 ports	8 ports	8 ports
10 Gbps	12 ports	12 ports	12 ports
8/4/2 Gbps	16 ports	16 ports	16 ports
Auto-negotiate (AN)	8 ports	8 ports	8 ports

1. The port blades have two ASICs; the per ASIC limit = numbers above/two
2. The Brocade 6510 switch and 16 Gbps embedded switches have one ASIC; the per ASIC limit = numbers above
3. The Brocade 6520 has four edge ASICs; the per ASIC limit = numbers above/four

This table does not show all the possible combinations of different speeds for the encryption and compression ports; other combinations are also supported. The number of supported ports is automatically calculated based on the speeds chosen.

Port speed on encryption- or compression-enabled 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 change the port speed.

Refer to [“Setting port speeds”](#) on page 94 for more information.

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

Any mismatch in configuration at either end of the IFL or authentication failure results in segmentation or, in rare cases, the port being disabled.

The most common reasons for E_Port or EX_Port segmentation include the following situations:

- Port authentication fails.
- 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.
- 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 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 92 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.
- 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.

- The Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) must be configured along with the DH group 4 for port level authentication as a prerequisite for in-flight encryption. Pre-shared secret keys must be configured on the devices on both ends of the ISL to perform authentication. Authentication secrets greater than 32 characters are recommended for stronger encryption keys. Once the link is authenticated, the keys are generated and exchanged.
- In-flight encryption uses DH-CHAP authentication (SHA-1 algorithm) followed by Internet Key Exchange (IKE) protocol (HMAC-SHA-512 algorithm) to generate the keys.
- 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.

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.

For FC16-32 or FC16-48 blades, if the two ports configured for encryption or compression within the same ASIC are not configured for trunking, it is recommended to connect each ISL to a different ASIC on the peer switch. Similarly, configure the two ports on the other ASIC of the blade. If the ports are configured for trunking, it is recommended to connect each trunk group to different ASICs on the peer switch.

For Brocade 6510 and 6520 switches, and 16 Gbps embedded switches, if the two ports are not configured for trunking, it is recommended that you connect each ISL to different ASICs on the peer switch.

NOTE

If any port on the ASIC with encryption or compression 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.

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 FC router.

1. Determine which ports are available for encryption or compression.
Refer to [“Viewing the encryption and compression configuration”](#) on page 452 for instructions.
2. Obtain the WWN of the edge switch using the **fcrEdgeShow** command.

You 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. 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 453 for instructions.
4. Enable encryption on the port.
Refer to [“Enabling in-flight encryption”](#) on page 455 for instructions.
5. Enable compression on the port.
Refer to [“Enabling in-flight compression”](#) on page 456 for instructions.
6. Obtain the WWN of the front phantom domain using the **portCfExPort** command.
You need this WWN when you set up the secret key on the E_Port on the other end of the IFL.

```
FCR:admin> portCfExPort 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)
```

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.

Next, 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

The following steps summarize how to enable encryption or compression on an E_Port.

To configure in-flight encryption and compression across an IFL, first configure encryption and compression on the EX_Port in the FC router.

Perform the following steps to configure the E_Port in the switch.

1. Determine which ports are available for encryption or compression.
Refer to [“Viewing the encryption and compression configuration”](#) on page 452 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 453 for instructions.
3. Enable encryption on the port.
Refer to [“Enabling in-flight encryption”](#) on page 455 for instructions.
4. Enable compression on the port.

Refer to “[Enabling in-flight compression](#)” on page 456 for instructions.

Following successful port initialization, the configured features are enabled and active. You can use the **isiShow** 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 user ports 348 and 349. Given the limit of two ports per ASIC, ASIC 1 has no more ports available for encryption or compression.

ASIC 0 (above the dashed line) has no ports configured for either encryption or compression and therefore has any two ports available for this purpose.

```
switch:admin> portenccompshow
```

User Port	Encryption		Compression		Config Speed
	Configured	Active	Configured	Active	
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.

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 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 either **dhchap** or **all**. The **dhchap** option explicitly specifies DH-CHAP. Although **all** enables both FCAP and DH-CHAP, the active protocol defaults to DH-CHAP for all ports configured for in-flight encryption.

If DH-CHAP is specified, then all switches in the fabric must enable DH-CHAP and establish pre-shared secrets. 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. Enter the **secAuthSecret --set** command to establish pre-shared secrets 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.

5. Activate DH-CHAP 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.

16 Configuring and enabling authentication for in-flight encryption

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 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 establishing DH-CHAP secrets, refer to [“Secret key pairs for DH-CHAP”](#) on page 249. For additional information about configuring DH-CHAP authentication, refer to [“Authentication policy for fabric elements”](#) on page 243.

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”](#) on page 452 for details.

You must also authenticate the port as described in [“Configuring and enabling authentication for in-flight encryption”](#) on page 453.

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.
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.

The following example enables in-flight encryption on port 0.

```
switch:admin> portdisable 0
switch:admin> portcfgencrypt --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: 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”](#) on page 452 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.

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.

The following example disables compression on port 0.

```
myswitch:admin> portdisable 0
myswitch:admin> portcfgcompress --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
```

16 Disabling in-flight compression

Diagnostic Port

In this chapter

- [Diagnostic Port](#) 459
- [Supported platforms for D_Port](#) 459
- [Licensing requirements for D_Port](#). 460
- [Understanding D_Port](#). 460
- [Supported topologies](#) 463
- [Using D_Port without HBAs](#). 465
- [Using D_Port with HBAs](#) 467
- [Controlling testing](#) 469
- [Example test scenarios and output](#) 469

Diagnostic Port

Diagnostic Port (D_Port) mode allows you to convert a Fibre Channel port into a diagnostic port for testing link traffic and running electrical loopback and optical loopback tests. The test results can be very useful in diagnosing a variety of port and link problems.

You can run D_Port tests between the following:

- A pair of switches
- A pair of Access Gateways
- An Access Gateway and a switch
- A switch and a host bus adapter (HBA)

The HBAs can be Brocade or non-Brocade HBAs. The non-Brocade HBAs must have Brocade HBA D_Port support.

Supported platforms for D_Port

D_Port functionality is supported only on 16 Gbps-capable platforms, running Fabric OS 7.0.0 or later. The ports must use 10 Gbps or 16 Gbps Brocade-branded SFP transceivers.

[Table 75](#) lists the Brocade switches and Fabric OS releases that support the D_Port feature.

TABLE 75 Supported platforms for D_Port

Product	Fabric OS release and later
Brocade DCX 8510-4 Backbone	v7.0.0
Brocade DCX 8510-8 Backbone	v7.0.0
Brocade 6505 switch	v7.0.1
Brocade 6510 switch	v7.0.0
Brocade 6520 switch	v7.1.0

D_Port functionality is supported on the following HBAs:

- Brocade 16-Gbps HBA (Brocade Fabric Adapter 1860) ports operating in HBA mode with a 16-Gbps SFP+ on Brocade 16-Gbps switches running Fabric OS version 7.1 or later.
Brocade HBA v3.1 provides limited support for D_Port. Brocade HBA v3.2 provides extensive support for D_Port, including dynamic D_Port mode.
- Non-Brocade 16-Gbps HBA (Must have HBA D_Port support.)

For applicable topologies, refer to “Supported topologies” on page 463.

Licensing requirements for D_Port

The D_Port feature does not require a license if you are running tests between a pair of Brocade devices, whether the devices are switches, Access Gateways, or HBAs.

If you want to run D_Port tests between a switch and a non-Brocade HBA, the Fabric Vision license is required. Also, the HBA vendor must have implemented the Brocade HBA D_Port support.

Understanding D_Port

A port in D_Port mode does not carry any user traffic, and is designed to run only specific diagnostics tests for identifying link-level faults or failures.

To bring up a port in D_Port mode, you must configure both ends of the link between a given pair of switches (or switches configured as Access Gateways), and you must disable the existing port before you can configure it as a D_Port. Figure 58 illustrates an example D_Port connection between a pair of switches through SFP transceivers (port assignments will vary). For all topologies supported, refer to “Supported topologies” on page 463.




FIGURE 58 Example of a basic D_Port connection between switches

Once the ports are configured and enabled as D_Ports, the following basic test suite is executed in the following order, depending on the SFPs installed:

1. Electrical loopback (with 16 Gbps SFP+ only)
2. Optical loopback (with 16 Gbps SFP+ only)
3. Link traffic (with 10 Gbps SFPs, 16 Gbps SFP+, and QSFP+)
4. Link latency and distance measurement (with 10 Gbps SFPs, 16 Gbps SFP+, and QSFP+)

NOTE

Electrical and optical loopback tests are not supported for ICLs.

The following steps are the fundamental parts of D_Port testing:

1. The user configures the desired ports on both ends of the connection.
2. Once both sides are configured, a basic test suite is initiated automatically when the link comes online, conducting diagnostic tests in the following order: (1) electrical loopback, (2) optical loopback, (3) link traffic, and (4) link latency and distance measurement.
3. After the automatic test is complete, the user can view results through the CLI or a GUI and rectify issues (if any) that are reported.
4. The user can also start (and restart) the test manually to verify the link.

Advantages of D_Port

Use the D_Port tests for the following situations:

- Testing a new link before adding it to the fabric
- Testing a trunk member before joining it with the trunk
- Testing long-distance cables and SFP transceivers

D_Port configuration mode and nature of test

D_Port has two modes: static and dynamic. In static mode, you explicitly configure the port to as a D_Port. In dynamic mode, the port is automatically set to a D_Port based on an external request from a remote port.

Dynamic D_Port mode is supported only on connections between a switch and an HBA.

By default, the switch has the capability to support dynamic D_Port mode. You can turn this capability off using the **configure** command, as shown here.

```
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]
```

```
Dynamic D-Port (on, off): [on] off
```

Table 76 summarizes D_Port test initiation modes and test start behavior.

TABLE 76 D_Port configuration mode and nature of test

D_Port mode/nature of test	Description	
Mode	Static	User configures port explicitly. Port remains as D_Port until user removes configuration.
	Dynamic	No user configuration required. D_Port mode is initiated by external request from remote port.
Nature of test	Automatic	Test automatically starts when the port comes online.
	Manual	User starts test from switch side (using the portdporttest -start command) or from HBA side (refer to “BCU D_Port commands” on page 468).

When the tests complete, the port behavior depends on the mode:

- For static D_Ports, you must remove the D_Port configuration at either or both ends of the link to bring up the port as a regular E_Port or F_Port.
- For a switch port in dynamic D_Port mode, the port automatically reverts back to an F_Port or an E_Port, if the port at the other end reverts to a regular port.

General limitations and considerations for D_Port

- The link to be tested must be marginally functional and able to carry a minimal number of frames before it can become a D_Port link. The D_Port feature is useful for diagnosing marginal faults only; it cannot detect the complete failure of any component.
- On QSFP+ ICL ports that are configured as D_Ports, only link traffic can be run. Electrical and optical loopback tests are not supported.
- Brocade recommends that D_Port tests be limited to a maximum of eight D_Ports at once. Otherwise, there is a possibility of false alarms.
- When a large number of D_Ports are configured, the test is run on one port per blade at a time, and other ports wait until the test is completed. No test begins until the fabric is stable.
- Note the following High Availability (HA) considerations:
 - There is no HA support for D_Port test options and results. Consequently, such information is not synchronized with the standby side. Any information from a previous test is lost following a failover.
 - During an HA failover and reboot on one side of the link, the link goes through reinitialization and may restart the test. However, the test cannot proceed if the remote port is not ready to proceed further (the remote port may already be done with the D_Port test and in final state). In such a case, the test will eventually fail with “Remote port not ready”. Restarting the test from either side will recover the port.
- D_Ports on an Access Gateway are supported only when there is no mapping between F_Ports and N_Ports; this includes static mapping and preferred port mapping. In addition, device (WWN) mapping is also not retained when a D_Port is used. If an Access Gateway port to be tested is mapped, the port mapping (including static and preferred port mapping) must be removed before the D_Port can be used. (Refer to “Saving port mappings on an Access Gateway” on page 464.)
- Access Gateway does not support D_Port dynamic mode. If the port on the connected switch or adapter is configured as a D_Port, the Access Gateway port is not automatically enabled as a D_Port. The D_Port must be configured on the Access Gateway port in static mode.

Also refer to [“Limitations and considerations for D_Port with HBAs”](#) on page 468.

Supported topologies

The following supported topologies illustrate at a high level how D_Port functionality can be used:

- [“Topology 1: ISLs”](#) on page 463
- [“Topology 2: ICLs”](#) on page 463
- [“Topology 3: Access Gateways”](#) on page 464
- [“Topology 4: HBA to switch”](#) on page 465

Topology 1: ISLs

[Figure 59](#) illustrates ISLs that connect multiple switches through a pair of chassis. **E** represents E_Ports to be configured as D_Ports.




FIGURE 59 ISLs connecting multiple switches and chassis

For configuration details, refer to [“Using D_Port without HBAs”](#) on page 465.

Topology 2: ICLs

[Figure 60](#) illustrates ICLs between slots 5 and 8 in corresponding chassis. **E** represents E_Ports to be configured as D_Ports.




FIGURE 60 ICLs connecting chassis blades

For configuration details, refer to [“Using D_Port without HBAs”](#) on page 465.

Topology 3: Access Gateways

Figure 61 illustrates a switch configured as a single Access Gateway connected to a fabric switch. **N** and **F** represent, respectively, an N_Port and an F_Port to be configured as D_Ports.

The Access Gateway must be either a Brocade 6505 or 6510.




FIGURE 61 Single Access Gateway to switch

Figure 62 illustrates multiple Access Gateways connected to a switch in a cascaded topology. **N** and **F** represent, respectively, an N_Port and an F_Port to be configured as D_Ports.




FIGURE 62 Multiple Access Gateways cascaded to switch

For configuration details, refer to “Using D_Port without HBAs” on page 465.

NOTE

N_Port-to-F_Port and device (WWN) mappings must be removed from an Access Gateway port before configuring the Access Gateway port as a D_Port. Refer to “Saving port mappings on an Access Gateway” on page 464.

Figure 63 illustrates connectivity between an HBA and an Access Gateway. **F** represents an F_Port to be configured as a D_Port.




FIGURE 63 Access Gateway to HBA

Saving port mappings on an Access Gateway

Before configuring ports as D_Ports on a switch configured as an Access Gateway, you must remove N_Port-to-F_Port and device (WWN) mappings. Fabric OS commands are available to save N_Port mappings. Once you save them, you can display the saved N_Port mappings to reconfigure them after D_Port is disabled. A command is also available to delete saved N_Port mappings.

For more details, refer to Chapter 2, “Configuring Ports in Access Gateway Mode,” in the *Access Gateway Administrator’s Guide*.

Topology 4: HBA to switch

Figure 64 illustrates connectivity between an HBA and a switch. F represents an F_Port to be configured as a D_Port.

This topology supports dynamic D_Port mode on both the switch and the HBA. In dynamic mode, the port does not need to be configured explicitly as a D_Port. It comes up in D_Port mode when it receives a request from the remote port.




FIGURE 64 HBA to switch

For configuration details, refer to “Using D_Port with HBAs” on page 467.

Using D_Port without HBAs

The following sections apply to topologies 1, 2, and 3:

- “Enabling D_Port” on page 465
- “Disabling D_Port” on page 466

The following examples use the command-line interface. Refer also to the latest Brocade Host Connectivity Manager (HCM) and Brocade Network Advisor documentation to use those graphical user interface (GUI) applications to configure D_Port.

Enabling D_Port

Use this procedure to configure a basic D_Port diagnostics session between two switches, as shown in Figure 58 on page 460.

NOTE

The automatic test might fail if you do not follow the sequence of steps exactly.

1. Disable Port 1 on Switch A by using the **portDisable [slot/]port** command.


```
switchA:admin> portdisable 1
```
2. Configure Port 1 on Switch A as a D_Port by using **portCfgDport -enable [slot/]port**.


```
switchA:admin> portcfgdport --enable 1
```
3. Repeat steps 1 and 2 for the corresponding port (in this example, Port 2) on Switch B.


```
switchB:admin> portdisable 2
switchB:admin> portcfgdport --enable 2
```
4. Enable Port 1 on Switch A by using the **portEnable [slot/]port** command.


```
switchA:admin> portenable 1
```

5. Enable Port 2 on Switch B by using the **portEnable [slot/]port** command.

```
switchB:admin> portenable 2
```

The basic test suite starts as soon as both ports are enabled.

6. While the test is running, enter the **portDportTest -show [slot/]port** command to view test results. The following test is successful.

```
switch:admin> portdporttest --show 10/39
```

D-Port Information:

```
=====
Slot:                10
Port:                39
Remote WWNN:         10:00:00:05:33:7e:69:c4
Remote port:         24
Mode:                Manual
No. of test frames:  12 Million
Test frame size:     1024 Bytes
Pattern:             JTSPAT
FEC (enabled/option/active):  Yes/No/No
CR (enabled/option/active):  No/No/No
Start time:          Mon Jan 16 05:57:51 2012
End time:            Mon Jan 16 05:58:56 2012
Status:              PASSED
=====
```

Test	Start time	Result	EST(HH:MM:SS)	Comments
Electrical loopback	05:57:52	PASSED	-----	-----
Optical loopback	05:58:06	PASSED	-----	-----
Link traffic test	05:58:13	PASSED	-----	-----

```
=====
Roundtrip link latency: 934 nano-seconds
Estimated cable distance: 1 meters
Buffers required: 1 (for 1024 byte frames at 16Gbps speed)
=====
```

7. To display a summary of the D_Port, use the **portDportTest -show all** command.

```
switch:admin> portdporttest --show all
Port State   SFP Capabilities Test Result
=====
24  ONLINE   E,O           PASSED
26  ONLINE   E,O           PASSED
33  OFFLINE  ---           FAILED
```

8. *Optional:* If one of the switches reboots, or if the test does not complete on one of the switches, restart the test on both switches. Use the **portDportTest -stop** command and restart the test with the **portDportTest -start** command on both switches.

Disabling D_Port

Use this procedure to disable the D_Port diagnostics session enabled in [“Enabling D_Port.”](#)

1. Disable Port 1 on Switch A by using the **portDisable 1 [slot/]port** command.

```
switchA:admin> portdisable 1
```

2. Disable the D_Port on Port 1 on Switch A by using **portCfgDport -disable 1**.

```
switchA:admin> portcfgdport --disable 1
```

3. Repeat steps 1 and 2 for Port 2 on Switch B.

```
switchB:admin> portdisable 2
switchB:admin> portcfgdport --disable 2
```

4. Enable Port 1 on Switch A by using the **portEnable** [slot/]port command.

```
switchA:admin> portenable 1
```

5. Enable Port 2 on Switch B by using the **portEnable** [slot/]port command.

```
switchB:admin> portenable 2
```

Using D_Port with HBAs

When HBAs are used, D_Port mode initiates electrical loopback, optical loopback, and link traffic diagnostic tests on the link between the HBA and the connected switch port. Results can be viewed from the switch by means of Fabric OS commands and from the adapter by means of the Brocade Command Line Utility (BCU) and Brocade Host Connectivity Manager (HCM) during or after the test. Once in D_Port mode, the adapter port does not participate in fabric operations, log in to a remote device, or run data traffic.

HBAs support testing in dynamic D_Port mode. If a D_Port is enabled on the switch only, it forces the connected adapter port into D_Port mode. The switch initiates and stops tests on the adapter port as specified by the switch configuration. Testing is started by means of BCU commands or HCM options.

In dynamic D_Port mode, you can disable the physical port by using the **bcu port -disable** command to exit D_Port mode. When you enable the port again, the switch will again force the adapter port into D_Port mode if the switch port is still enabled as a D_Port.

The following sections apply to “[Topology 4: HBA to switch](#)” on page 465.

- “[Automatic mode configuration](#)” on page 467
- “[Dynamic mode configuration](#)” on page 468
- “[BCU D_Port commands](#)” on page 468

Automatic mode configuration

This procedure enables a D_Port diagnostic session from the connected switch. After the default test suite is run automatically, you can run specific tests manually to obtain additional detail.

1. Disable the switch port by using the **portDisable** [slot/]port command.
2. Configure the switch port that you want to enable as a D_Port by using the **portCfgDport -enable** [slot/]port command.
3. Disable the adapter port by using the adapter **bcu port -disable** command.
4. Enable the switch port by using the **portEnable** [slot/]port command.
5. Enable the adapter port as a D_Port by using the adapter **bcu diag -dportenable** command and configure test parameters.

For more details on adapter configuration, refer to the *Brocade Fabric Adapters Administrator’s Guide*.

Dynamic mode configuration

This procedure enables a dynamic D_Port diagnostic session from the connected switch to an HBA.

NOTE

D_Port on HBAs is supported only on 16-Gbps SFP transceivers.

1. Disable the switch port by using the **portDisable** [slot/]port command.
2. Enable the switch port as a D_Port by using the **portCfgDport - -enable** [slot/]port command.
3. Enable the switch port by using the **portEnable** [slot/]port command.

To verify whether the port is a D_Port, use the **bcu port - -list** command and look for a “D” in the listing.

For more details on adapter configuration, refer to the *Brocade Fabric Adapters Administrator’s Guide*.

BCU D_Port commands

The following BCU commands can be used for D_Port configuration and control:

- **bcu diag --dportenable** — Enables D_Port on a specific port, sets the test pattern, and sets the frame count for testing.
- **bcu diag --dportdisable** — Disables D_Port on a specific port and sets the port back to an N_Port or NL_Port.
- **bcu diag --dportshow** — Displays test results for a test in progress on a specific port.
- **bcu diag --dportstart** — Restarts a test on a specific port when the test has completed.
- **bcu port --list** — Displays the D_Port enabled or disabled state on the adapter and connected switch.

NOTE

If you stop the test from the switch side, you should disable D_Port on the HBA side.

Use **bcu diag - -dportdisable** in static D_Port mode or **bcu port - -disable** in dynamic D_Port mode.

Limitations and considerations for D_Port with HBAs

Keep in mind the following limitations and considerations for D_Port configurations:

- D_Port is supported only on Brocade 16-Gbps HBA (Brocade Fabric Adapter 1860) ports operating in HBA mode with a 16-Gbps SFP+ on Brocade 16-Gbps switches running Fabric OS version 7.1 or later. In addition, the Brocade adapter must be using driver version 3.2.0 or higher.
- D_Port is supported on non-Brocade 16-Gbps HBAs if you have a Fabric Vision license present on the switch and if the HBA vendor has implemented the Brocade HBA D_Port support.
- D_Ports do not support a loop topology.
- D_Ports are not supported in a configuration of an HBA to another HBA (in target mode).
- D_Ports on the HBA do not support forward error correction (FEC) and credit recovery (CR). If these features are enabled on the switch side, the HBA ignores them.
- D_Port is not supported on adapter ports configured in CNA mode.

- Toggling the port on either side of the link does not restart the test.
- Because of SFP electrical wrap (EWRAP) bleed-through, during the beginning of switch electrical loopback testing, the HBA will receive some broken frames, which cause the port statistic error counter to increase. Examples are “CRC err,” “bad EOF,” and “invalid order set.” Similar results occur for the optical loopback test. You should ignore these port statistics on the HBA.
- The following commands from the switch are not supported by the HBA, and the HBA will drop or reject them:
 - **portdporttest --stop**
 - **portdporttest --restart**
 - **portdporttest --setarg**

Although the adapter supports **portdporttest -start**, options for this command are ignored. With the exception of **-fec** and **-cr**, the **-start** suboptions will work for D_Port on an HBA.

- D_Port is useful for diagnosing marginal faults only. A complete failure of any component cannot be detected.
- D_Port configuration is not supported on mezzanine cards.
- The maximum number of D_Ports on which the tests can run simultaneously depends on the HBA firmware version.

TABLE 77 Limitation on number of D_Ports for simultaneous tests

HBA firmware version	Maximum number of D_Ports on which tests can be run simultaneously
HBA v3.2.0	4
HBA v3.2.3	8

- Sometimes a port may get stuck in G_Port state if one end of the link is in D_Port mode and the other end is not. If this happens, make the D_Port mode compatible on both ends of the link. Either re-configure both ends of the link as D_Ports, or remove the D_Port configuration from both ends of the link.
- Powering off and on or plugging in and out slots containing ports in D_Port mode result in those ports losing the dynamic D_Port state when the slot or port is back up. If this happens you need to re-configure dynamic D_Port mode on the port.

Controlling testing

You can stop and start D_Port testing on a port by using the following commands:

- **portdporttest -stop [slot/]port**
- **portdporttest -start [slot/]port**

Example test scenarios and output

In addition to the examples shown in “[Enabling D_Port](#)” on page 465, other practical scenarios are shown in the following sections.

Confirming SFP and link status with an HBA

The steps in the following example illustrate how the **bcu diag --dportenable** command will fail with an SFP installed but without a connection to the switch.

1. Confirm the initial port status.

```
switch:admin> bcu port --list
```

```
-----
Port#  FN  Type  PWWN/MAC                      FC Addr/  Media  State      Spd
                        Eth dev
-----
1/0    -   fc    10:00:8c:7c:ff:1c:e9:00      160000    sw     Linkup     16G*
      0   fc    10:00:8c:7c:ff:1c:e9:00      160000    sw     Linkup     16G*
1/1    -   fc    10:00:8c:7c:ff:1c:e9:01      --        sw     Linkdown   ---
      1   fc    10:00:8c:7c:ff:1c:e9:01      --        sw     Linkdown   ---
-----
```

2. Disable the port.

```
switch:admin> bcu port --disable 1/0
port disabled
```

3. Remove the connection to the switch and attempt to enable the D_Port.

```
switch:admin> bcu diag --dportenable 1/0
ERROR: Timer expired - Retry if persists contact support
```

4. Install an SFP and attempt to enable the D_Port.

```
switch:admin> bcu diag --dportenable 1/0
ERROR: Switch port is not D_Port capable or D_Port is disabled
```

5. Connect to the HBA without the SFP and disable the native port.

```
switch:admin> bcu port --disable 1/0
port disabled
```

6. Attempt to enable the D_Port.

```
switch:admin> bcu diag --dportenable 1/0
ERROR: SFP is not present.
D-port will be enabled but it will be operational only after inserting a valid
SFP.
```

Starting and stopping D_Port testing

Use the **portDportTest** command to start or stop D_Port testing or show test results.

You can display the complete results from either the responder or the initiator switch. If the initiator switch is running Fabric OS v7.1.x or earlier, the responder displays only the local D_Port results, and you must query the initiator to see the complete results.

The following example shows the D_Port results.

```
switch:admin> portdporttest --show 26
D-Port Information:
=====
Port:                26
Remote WWNN:         10:00:00:05:33:13:2f:b5
```

```

Remote port:      42
Mode:            Automatic
Start time:      Wed Feb  2 01:41:43 2011
End time:        Wed Feb  2 01:43:23 2011
Status:          PASSED
=====
Test              Start time  Result          EST(secs)  Comments
=====
Electrical loopback 01:42:08    PASSED         --          -----
Optical loopback   01:42:16    PASSED         --          -----
Link traffic test  01:43:15    PASSED         --          -----
=====
Roundtrip link latency:      1108 nano-seconds
Estimated cable distance:    20 meters
Buffers required:            1 (for 1024 byte frames at 16Gbps speed)

```

The following example shows the **portdporttest --show** output where the electrical and optical tests pass but the link test fails.

```

switch:admin> portdporttest --show 10/39
D-Port Information:
=====
Slot:                10
Port:                39
Remote WWNN:         10:00:00:05:33:7e:69:c4
Remote port:         24
Mode:                Manual
No. of test frames:  12 Million
Test frame size:     1024 Bytes
Pattern:             JTSPAT
FEC (enabled/option/active):  Yes/No/No
CR (enabled/option/active):  No/No/No
Start time:          Mon Jan 16 05:57:51 2012
End time:            Mon Jan 16 05:58:56 2012
Status:              FAILED
=====
Test              Start time  Result          EST(HH:MM:SS)  Comments
=====
Electrical loopback 05:57:52    PASSED         -----          -----
Optical loopback   05:58:06    PASSED         -----          -----
Link traffic test  05:58:13    FAILED         -----          See failure report
=====
Roundtrip link latency:      934 nano-seconds
Estimated cable distance:    1 meters
Buffers required:            1 (for 1024 byte frames at 16Gbps speed)

Failure report:
Errors detected (local):     CRC, Bad_EOF, Enc_out
Errors detected (remote):    CRC, Bad_EOF

```

Please use **portstatsshow** and **porterrshow** for more details on the above errors. Refer to file `/var/tmp/dport/slot10port39_stats.txt`, for link statistics prior to the port was set to D-Port

Use the **portDportTest --show all** command to display the capabilities and test results of all the D_Ports in a switch.

```

switch:admin> portdporttest --show all
Port State      SFP Capabilities  Test Result

```

17 Example test scenarios and output

```
=====
24   ONLINE   E,O           PASSED
26   ONLINE   E,O           FAILED
33   ONLINE   E,O           PASSED
```

Use the **switchShow** command to see D_Port information.

```
switch:admin> switchshow
switchName:      switch_10
switchType:      109.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:05:33:13:2f:b4
zoning:          OFF
switchBeacon:    OFF
FC Router:       OFF
Allow XISL Use:  ON
LS Attributes:   [FID: 10, Base Switch: No, Default Switch: No, Address Mode 0]
```

```
Index Port Address Media Speed State      Proto
=====
  24  24   010000   id   N16   Online    FC   D-Port Loopback->Port 24
  26  26   010200   id   N16   Online    FC   D-Port segmented,(D-Port mode
mismatch)
  33  33   010300   id    N8   Online    FC   D-Port 10:00:00:05:33:13:2f:b5
```

Use the **portCfgShow** command to see which ports are D_Port-enabled.

```
switch:admin> portcfgshow
Ports of Slot 0      24  26  27
-----+-----+-----+
Octet Speed Combo    1   1   1
Speed                AN  AN  AN
AL_PA Offset 13      ..  ..  ..
Trunk Port            ON  ON  ON
Long Distance        ..  ..  ..
.....
.....
Port Auto Disable    ..  ..  ..
CSCTL mode           ..  ..  ..
D-Port mode          ON  ON  ON
D-Port over DWDM     ..  ..  ..
Compression          ..  ..  ..
Encryption           ..  ..  ..
FEC                  ON  ON  ON
Fault Delay          0   0   0
      where AE:QoSAutoEnable, AN:AutoNegotiate, ..:OFF, -:NotApplicable, ?:INVALID
```


NPIV

In this chapter

- [NPIV overview](#) 473
- [Configuring NPIV](#) 475
- [Enabling and disabling NPIV](#) 476
- [Viewing NPIV port configuration information](#) 476

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 *Access Gateway Administrator's 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 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).

Example of NPIV devices

```
switch:admin> switchshow
switchName:      5100
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:e40xa06601
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
    
```

On the Brocade DCX and DCX-4S with the FC8-64 blade, the base port is not included in the NPIV device count. The following example shows 63 NPIV devices total.

```

Index Slot Port Address Media Speed State Proto
=====
127 12 15 a07f40 id N4 Online FC F-Port 1 N Port + 63 NPIV public
(AoQ)
    
```

Upgrade considerations

The maximum logins per switch has decreased with Fabric OS v6.4.0. When upgrading from a release previous to Fabric OS v6.4.0 and later, 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. When Virtual Fabrics is enabled on the Brocade DCX and DCX-4S, fixed addressing mode is used only on the default logical switch. 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 and DCX-4S Backbones. The number of NPIV or loop devices supported on a port is 64.

[Table 78](#) shows the number of NPIV devices supported on the Brocade DCX and DCX-4S Backbones.

TABLE 78 Number of supported NPIV devices

Platform	Virtual Fabrics	Logical switch type	NPIV support
DCX	Disabled	N/A	Yes, 127 virtual device limit. ¹
DCX	Enabled	Default switch	Yes, 63 virtual device limit. ¹
DCX	Enabled	Logical switch	Yes, 255 virtual device limit. ^{2, 3}
DCX	Enabled	Base switch	No.
DCX-4S	Disabled	N/A	Yes, 255 virtual device limit.
DCX-4S	Enabled	Default switch	Yes, 255 virtual device limit.

TABLE 78 Number of supported NPIV devices (Continued)

Platform	Virtual Fabrics	Logical switch type	NPIV support
DCX-4S	Enabled	Logical switch	Yes, 255 virtual device limit. ³
DCX-4S	Enabled	Base switch	No.

1. Maximum limit support takes precedence if user-configured maximum limit is greater. This applies to shared areas on the FC4-48, FC8-48, and FC8-64 port blades.
2. 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, 113, and higher, are limited to 63 logical devices.
3. Maximum limit of 63 for 10-bit areas connected to third-party (non-Brocade) NPIV HBAs.

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. Perform 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. Press **Enter**.
5. Enter the **portEnable** command to enable the port.

Example of setting the login limit

```
switch:admin> portcfgnpivport --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
```

18 Enabling and disabling NPIV

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.
2. To enable or disable NPIV on a port, enter the **portCfgNPIVPort** command with either the **--enable** or **--disable** option.

The following example shows NPIV being enabled on port 10 of a Brocade 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.

Viewing NPIV port configuration information

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **portCfgShow** command to view the switch ports information.

The following example shows whether a port is configured for NPIV:

```
switch:admin> portcfgshow
Ports of Slot 0   0   1   2   3   4   5   6   7   8   9  10  11  12  13  14  15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed           AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Trunk Port      ON ON ON ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance   . . . . .   . . . . .   . . . . .   . . . . .
VC Link Init    . . . . .   . . . . .   . . . . .   . . . . .
Locked L_Port   . . . . .   . . . . .   . . . . .   . . . . .
Locked G_Port   . . . . .   . . . . .   . . . . .   . . . . .
Disabled E_Port . . . . .   . . . . .   . . . . .   . . . . .
ISL R_RDY Mode . . . . .   . . . . .   . . . . .   . . . . .
RSCN Suppressed . . . . .   . . . . .   . . . . .   . . . . .
Persistent Disable . . . . . . . . . . . . . . . . . . . . . . . .
NPIV capability ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON
```

- Use the **switchShow** and **portShow** commands to view NPIV information for a given port. If a port is an F_Port, and you enter the **switchShow** command, 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 is sample 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** command to view the NPIV attributes and all the N_Port (physical and virtual) port WWNs that are listed under *portWwn of device(s) connected*. The following example is sample output for the **portShow** command:

```
switch:admin> portshow 2
portName: 02
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
```

18 Viewing NPIV port configuration information

```
portFlags: 0x24b03 PRESENT ACTIVE F_PORT G_PORT NPIV LOGICAL_ONLINE LOGIN
NOELP LED ACCEPT
portType: 10.0
portState: 1Online
portPhys: 6In_Sync
portScn: 32F_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

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
```

Fabric-Assigned PWWN

In this chapter

- [Fabric-Assigned PWWN overview](#) 479
- [User- and auto-assigned FA-PWWN behavior](#) 480
- [Configuring an FA-PWWN for an HBA connected to an Access Gateway](#) 481
- [Configuring an FA-PWWN for an HBA connected to an edge switch](#) 482
- [Supported switches and configurations for FA-PWWN](#) 483
- [Configuration upload and download considerations for FA-PWWN](#) 483
- [Security considerations for FA-PWWN](#) 483
- [Restrictions of FA-PWWN](#) 484
- [Access Gateway N_Port failover with FA-PWWN](#) 484

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 Brocade HBA or adapter 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

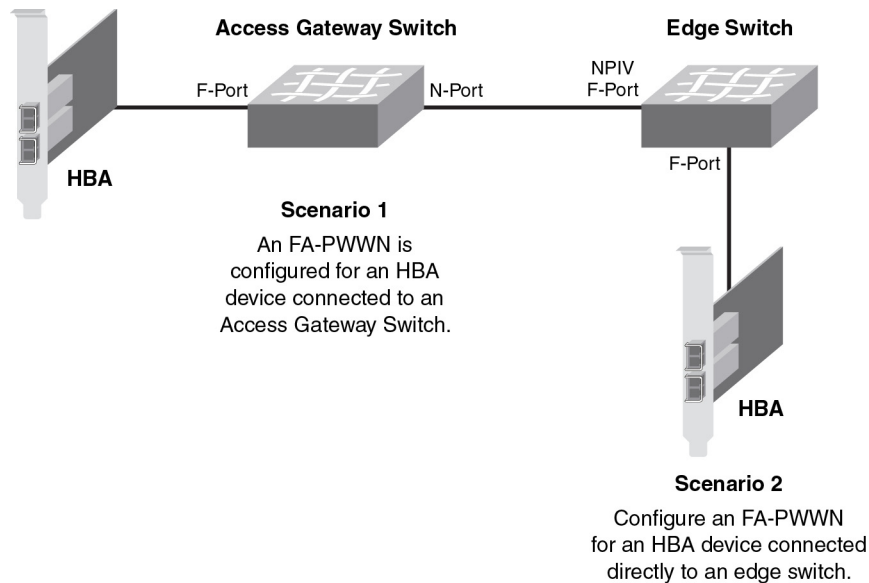


FIGURE 65 Fabric-assigned port World Wide Name provisioning scenarios

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:

- Switch platforms running Fabric OS v7.0.0 or later:
 - Brocade DCX, DCX-4S, and DCX 8510 family
 - Brocade 300
 - Brocade 5100
 - Brocade 5300
 - Brocade 6505
 - Brocade 6510
 - Brocade 6520
 - Brocade 7800
 - Brocade VA-40FC
- Access Gateway platforms running Fabric OS v7.0.0 or later:
 - Brocade 300
 - Brocade 5100
 - 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 246 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 241 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
- FL_Ports
- 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.

Managing Administrative Domains

In this chapter

- [Administrative Domains overview](#) 485
- [Admin Domain management for physical fabric administrators](#) 494
- [SAN management with Admin Domains](#) 506

Administrative Domains overview

An *Administrative Domain* (Admin Domain or AD) is a logical grouping of fabric elements that defines which switches, ports, and devices you can view and modify. An Admin Domain is a filtered administrative view of the fabric.

NOTE

If you do not implement Admin Domains, the feature has no impact on users and you can ignore this chapter.

Admin Domains permit access to a configured set of users. Using Admin Domains, you can partition the fabric into logical groups and allocate administration of these groups to different user accounts. These accounts can manage only the Admin Domains assigned to them and cannot make changes to the rest of the fabric.

For example, you can put all the devices in a particular department in the same Admin Domain for ease of managing those devices. If you have remote sites, you could put the resources in the remote site in an Admin Domain and assign the remote site administrator to manage those resources.

Admin Domains and Virtual Fabrics are mutually exclusive and are not supported at the same time on a switch.

Do not confuse Admin Domains with zones:

- Zones define which devices and hosts can communicate with each other.
- Admin Domains define which users can manage which devices, hosts, and switches.

You can have up to 256 Admin Domains in a fabric (254 user-defined and 2 system-defined), numbered from 0 through 255.

Admin Domains are designated by a name and a number. This document refers to specific Admin Domains using the format “AD n ” where n is a number between 0 and 255.

ATTENTION

The Admin Domain administrator can define up to 254 ADs (AD1 through AD254) in the AD database; however, it is recommended that no more than 16 active Admin Domains run concurrently. More than 16 active Admin Domains might cause performance degradation and unpredictable system behavior.

NOTE

Do not confuse an *Admin Domain number* with the *domain ID* of a switch. They are two different identifiers. The Admin Domain number identifies the Admin Domain and has a range from 0 through 255. The domain ID identifies a switch in the fabric and has a range from 1 through 239.

Figure 66 shows a fabric with two Admin Domains: AD1 and AD2.




FIGURE 66 Fabric with two Admin Domains

Figure 67 shows how users get a filtered view of this fabric, depending on which Admin Domain they are in. As shown in Figure 67, users can see all switches and E_Ports in the fabric, regardless of their Admin Domain; however, the switch ports and end devices are filtered based on Admin Domain membership.




FIGURE 67 Filtered fabric views when using Admin Domains

Admin Domain features

Admin Domains allow you to do the following:

- Define the scope of an Admin Domain to encompass ports and devices within a switch or a fabric.
- Share resources across multiple Admin Domains. For example, you can share array ports and tape drives between multiple departments. In [Figure 66](#) on page 486, one of the storage devices is shared between AD1 and AD2.
- Have a separate zone database for each Admin Domain. Refer to [“Admin Domains, zones, and zone databases”](#) on page 510 for more information.
- Move devices from one Admin Domain to another without traffic disruption, cable reconnects, or discontinuity in zone enforcement.
- Provide strong fault and event isolation between Admin Domains.
- Have visibility of all physical fabric resources. All switches, E_Ports, and FRUs (including blade information) are visible.
- Continue to run existing third-party management applications. Prior and existing versions of third-party management applications continue to work with admin IDs and user IDs.

Requirements for Admin Domains

Implementing Admin Domains in a fabric has the following requirements:

- The default zone mode setting must be set to No Access before you create Admin Domains (refer to [“Setting the default zoning mode for Admin Domains”](#) on page 495 for instructions).
- Virtual Fabrics must be disabled before you create Admin Domains (refer to [“Disabling Virtual Fabrics mode”](#) on page 325 for instructions).
- Gigabit Ethernet (GbE) ports cannot be members of an Admin Domain.
- Traffic Isolation Zoning is supported within Admin Domains, with some restrictions, as described in [“Admin Domain considerations for Traffic Isolation Zoning”](#) on page 398.
- If the fabric includes LSAN zones:
 - The LSAN zone names must not end with “_ADn”.
 - The LSAN zone names must not be longer than 57 characters.

Refer to [Chapter 26, “Using FC-FC Routing to Connect Fabrics,”](#) for information about the FC-FC Routing Service and LSAN zones.

Admin Domain access levels

Admin Domains offer a hierarchy of administrative access. To manage Admin Domains, you must be a *physical fabric administrator*. A physical fabric administrator is a user with admin permissions and access to all Admin Domains (AD0 through AD255). Only a physical fabric administrator can perform Admin Domain configuration and management.

Other administrative access is determined by your defined Role-Based Access Control (RBAC) role and AD membership. Your role determines your access level and permission to perform an operation. Your AD membership determines the fabric resources on which you can operate.

[Table 79](#) lists each Admin Domain user type and describes its administrative access and capabilities.

TABLE 79 AD user types

User type	Description
Physical fabric administrator	User account with admin permissions and with access to all Admin Domains (ADO through AD255). Creates and manages all Admin Domains. Assigns other administrators or users to each Admin Domain. The default admin account is the first physical fabric administrator. Only a physical fabric administrator can create other physical fabric administrators.
Administrative Domain users	Can be assigned to one or more Admin Domains. Manage the resources within their Admin Domains. If their role permits, can create user accounts and assign them to Admin Domains in their list. Cannot view other Admin Domain definitions. They can view only members of their own Admin Domains.

User-defined Admin Domains

AD1 through AD254 are user-defined Admin Domains. These user-defined Admin Domains can be created only by a physical fabric administrator (refer to [“Admin Domain access levels”](#) on page 487 for more information).

In [Figure 66](#) on page 486, AD1 and AD2 are user-defined Admin Domains.

System-defined Admin Domains

ADO and AD255 are system-defined Admin Domains. ADO and AD255 always exist and cannot be deleted or renamed. They are reserved for use in creation and management of Admin Domains.

ADO

ADO is a system-defined Admin Domain. Unlike user-defined Admin Domains, ADO has an implicit and an explicit membership list. User-defined Admin Domains have only an explicit membership list.

- The *implicit membership list* contains all devices, switch ports, and switches that have not been assigned to any other Admin Domain.
Initially, the ADO implicit membership list contains all devices, switch ports, and switches in the fabric. When you create AD1 through AD254, the devices, switch ports, and switches used to create these user-defined Admin Domains disappear from the ADO implicit membership list.
- The *explicit membership list* contains all devices, switch ports, and switches that you explicitly add to ADO and can be used to force device and switch sharing between ADO and other Admin Domains.

ADO is managed like any user-defined Admin Domain. The only difference between ADO and user-defined Admin Domains is the implicit membership list.

The implicit members of ADO change dynamically as the membership of other Admin Domains changes. The explicit members of ADO are not deleted unless you explicitly remove them.

For example, if DeviceA is not a member of any user-defined Admin Domain, then it is an implicit member of ADO.

If you explicitly add DeviceA to ADO, then DeviceA is both an implicit and an explicit member of ADO.

<u>ADO implicit members</u>	<u>ADO explicit members</u>	<u>AD2 members</u>
DeviceA	DeviceA	none

If you add DeviceA to AD2, then DeviceA is deleted from the ADO implicit membership list, but is *not* deleted from the ADO explicit membership list.

<u>ADO implicit members</u>	<u>ADO explicit members</u>	<u>AD2 members</u>
none	DeviceA	DeviceA

If you then remove DeviceA from AD2, DeviceA is added back to the ADO implicit membership list (assuming DeviceA is not in any other Admin Domain).

<u>ADO implicit members</u>	<u>ADO explicit members</u>	<u>AD2 members</u>
DeviceA	DeviceA	none

When a new device is added to the fabric, it automatically becomes an implicit member of ADO until it is explicitly added to an Admin Domain.

ADO is useful when you create Admin Domains because you can see which devices, switch ports, and switches are not yet assigned to any Admin Domains.

ADO owns the root zone database (legacy zone database).

AD255

AD255 is a system-defined Admin Domain that is used for Admin Domain management. AD255 always contains all of the devices in the entire physical fabric. You can use AD255 to get an unfiltered view of the fabric and to view the hierarchical zone databases of ADO through AD254. All Admin Domain management is done in the AD255 context.

AD255 does not have a zone database associated with it; you cannot use AD255 to perform any zoning management tasks (non-read operations such as creating or modifying zones).

[Figure 68](#) on page 490 shows the same fabric from [Figure 66](#) on page 486, but with ADO and AD255 shown. ADO contains the two devices that are not in any of the user-defined Admin Domains (AD1 and AD2). AD255 always encompasses the entire physical fabric.




FIGURE 68 Fabric with AD0 and AD255

Home Admin Domains and login

You are always logged in to an Admin Domain, and you can view and modify only the devices in that Admin Domain.

If you have access to more than one Admin Domain, one of them is designated as your *home Admin Domain*, the one you are automatically logged in to. If your home Admin Domain is deleted or deactivated, then by default you are logged in to the lowest-numbered active Admin Domain in your Admin Domain list. The home Admin Domain, like the Admin Domain list, is a configurable property of a non-default user account. Here is some additional information about AD accounts:

- You can log in to only one Admin Domain at a time. You can later switch to a different Admin Domain (refer to [“Switching to a different Admin Domain context”](#) on page 508 for instructions).
- For default accounts such as admin and user, the home Admin Domain defaults to AD0 and cannot be changed.
- The Admin Domain list for the default admin account is 0 through 255, which gives this account automatic access to any Admin Domain as soon as the domain is created, and makes this account a physical fabric administrator.
- The Admin Domain list for the default user account is AD0 only.

- For user-defined accounts, the home Admin Domain defaults to ADO but an administrator can set the home Admin Domain to any Admin Domain to which the account is given access.
- If you are in any Admin Domain context other than ADO, the Admin Domain number is included in the system prompt displayed during your session. The following are example prompts for when you are in the ADO, AD1, and AD255 contexts, respectively:

```
switch:admin>
switch:AD1:admin>
switch:AD255:admin>
```

Admin Domain member types

You define an Admin Domain by identifying members of that domain. Admin Domain members can be devices, switch ports, or switches. Defining these member types is similar to defining a traditional zone member type. An Admin Domain does not require or have a new domain ID or management IP address linked to it.

Device members

Device members are defined by the device World Wide Name (WWN) and have the following properties:

- A device member can be either a device port WWN or a device node WWN.
- A device member grants view access to the device and zoning rights. View rights are also granted to the switch port to which the device is attached.
- A device member provides a pure virtual view. The cabling and switch port diagnostics and control are done by the physical fabric administrator.

Port control is provided only through switch port membership and is not provided for device members. When you create an Admin Domain, the end device members do not need to be online, even though their WWNs are used in the Admin Domain definition.

You can share device members across multiple Admin Domains. You can also zone shared devices differently in each Admin Domain. A device WWN member does not automatically grant usage of corresponding *domain,index* members in the zone configuration. If you specify a device WWN member in the Admin Domain member list, zone enforcement ignores zones with the corresponding port (the port to which the device is connected) member usage.

Switch port members

Switch port members are defined by switch *domain,index* and have the following properties:

- A switch port member grants port control rights and zoning rights for that switch port.
- A switch port member grants view access and zoning rights to the device connected to that switch port.
- A switch port member allows you to share *domain,index* members across multiple Admin Domains. In each Admin Domain, you can also zone shared devices differently.
- A switch port member implicitly includes all devices connected to the specified *domain,index* members in the Admin Domain membership.
- A switch port member allows you to specify a range of indices as Admin Domain members, for example: 100,5-8. The index range arguments are expanded and stored in the Admin Domain member list.

If a device is a member of an Admin Domain, the switch port to which the device is connected becomes an indirect member of that Admin Domain and the *domain,index* is removed from the ADO implicit membership list.

NOTE

If the switch domain ID changes, the *domain,index* members are invalid (they are not automatically changed). You must then reconfigure the Admin Domain with the current *domain,index* members.

Switch members

Switch members are defined by the switch WWN or domain ID, and have the following properties:

- A switch member grants administrative control to the switch.
- A switch member grants port control for all ports in that switch.
- A switch member allows switch administrative operations such as disabling and enabling a switch, rebooting, and firmware downloads.
- A switch member does not provide zoning rights for the switch ports or devices.

To allow devices to be zoned within Admin Domains, you must specify the port members using *domain,index* or device WWN members.

E_Ports (including VE_Ports, EX_Ports, and VEX_Ports) are implicitly shared across all Admin Domains. An administrator can perform port control operations only if the *domain,index* of the E_Port is part of the Admin Domain.

NOTE

Only the WWN of the switch is saved in the Admin Domain. If you change the domain ID of the switch, the Admin Domain ownership of the switch is not changed.

Admin Domains and switch WWNs

Admin Domains are treated as fabrics. Because switches cannot belong to more than one fabric, switch WWNs are converted so that they appear as unique entities in different Admin Domains (fabrics). This WWN conversion is done only in the AD1 through AD254 context. ADO and AD255 use unconverted switch WWNs.

The switch WWN has the following format:

10:00:nn:nn:nn:nn:nn:nn

In an Admin Domain context, the switch WWN is converted from NAA=1 to NAA=5 format, with the Admin Domain number added, using the following syntax:

5n:nn:nn:nn:nn:nn:n9:xx

In the syntax, xx is the Admin Domain number.

For example, the following switch WWN is in NAA=1 format:

10:00:00:60:69:e4:24:e0

The following switch WWN is the converted WWN for the previous example in AD1:

50:06:06:9e:42:4e:09:01

Figure 69 on page 493 shows an unfiltered view of a fabric with two switches, three devices, and two Admin Domains. The devices are labeled with device WWNs and the switches are labeled with domain IDs and switch WWNs.




FIGURE 69 Fabric showing switch and device WWNs

Figure 70 shows the filtered view of the fabric as seen from AD3 and AD4. The switch WWNs are converted to the NAA=5 syntax; the device WWNs and domain IDs remain the same.

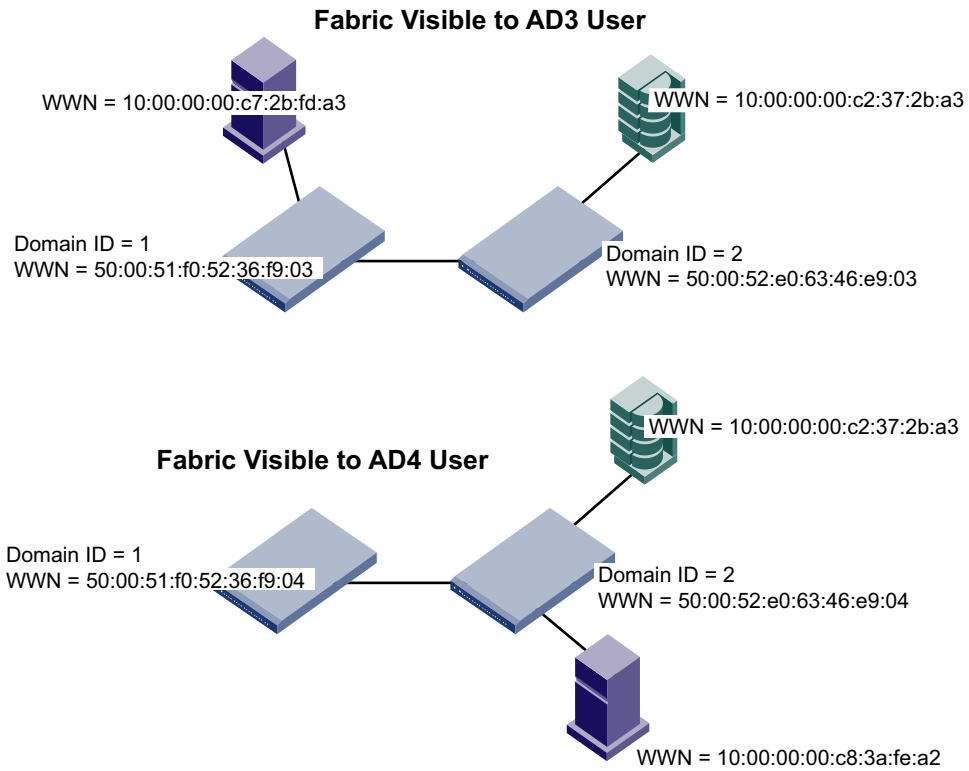


FIGURE 70 Filtered fabric views showing converted switch WWNs

Admin Domain compatibility, availability, and merging

Admin Domains maintain continuity of service for Fabric OS features and operate in mixed-release Fabric OS environments. High availability is supported with some backward compatibility.

When an E_Port comes online, the adjacent switches merge their AD databases. The receiving switch accepts an AD database from the neighboring switch only if the local AD database is empty or if the new AD database exactly matches both the defined and effective configurations of the local AD database. If the AD database merge fails, the E_Port is segmented with an “AD conflict” error code.

Admin Domain management for physical fabric administrators

NOTE

This section is for physical fabric administrators who are managing Admin Domains.

The **ad** command follows a batched-transaction model, which means that changes to the Admin Domain configuration occur in the transaction buffer.

An Admin Domain configuration can exist in several places:

- Effective configuration — The Admin Domain configuration that is currently in effect.
- Defined configuration — The Admin Domain configuration that is saved in flash memory. There might be differences between the effective configuration and the defined configuration.
- Transaction buffer — The Admin Domain configuration that is in the current transaction buffer and has not yet been saved or canceled.

How you end the transaction determines the disposition of the Admin Domain configuration in the transaction buffer. The following commands end the Admin Domain transaction:

- | | |
|------------------------|---|
| ad --save | Saves the changes in the transaction buffer to the defined configuration in persistent storage and propagates the defined configuration to all switches in the fabric. Note that for delete and clear operations, if one or more of the deleted Admin Domains are in the effective configuration, you cannot use --save , but must use --apply instead. |
| ad --apply | Saves the changes to the defined configuration in persistent storage and enforces the defined configuration on all switches in the fabric, replacing the effective configuration. |
| ad --transabort | Aborts the transaction and clears the transaction buffer. The effective and defined configurations remain unchanged. |

You can enter the **ad --transshow** command at any time to display the ID of the current Admin Domain transaction.

Setting the default zoning mode for Admin Domains

To begin implementing an Admin Domain structure within your SAN, you must first set the default zoning mode to No Access. You must be in ADO to change the default zoning mode.

1. Log in to the switch with the appropriate RBAC role.
2. Ensure you are in the ADO context by entering the **ad --show** command to determine the current Admin Domain.
If necessary, switch to the ADO context by entering the **ad --select 0** command.
3. Set the default zoning mode to No Access, as described in [“Setting the default zoning mode”](#) on page 361.

Creating an Admin Domain

To create an Admin Domain, you must specify an Admin Domain name, number, or both:

- If you create an Admin Domain using only a number, the Admin Domain name is automatically assigned to be “AD n ”, where n is the number you specified.
For example, if you specify AD number = 4, then AD name is set to “AD4”.
- If you create an Admin Domain using only a name, the Admin Domain number is automatically assigned and is the lowest available AD number, except if you specify a name in the format “AD n ”, in which case the Admin Domain number is assigned to be n .
For example, if you specify AD name = “blueAD” and the lowest available AD number is 5, then AD name is “blueAD” and AD number is 5.
If you specify AD name = “AD15” and the lowest available AD number is 6, then AD name is “AD15” and AD number is 15. Because the specified name is in the format “AD n ”, the AD number is assigned to be n and *not* the lowest available AD number.

When you create an Admin Domain, you must specify at least one member (switch, switch port, or device). You cannot create an empty Admin Domain. For more information about these member types, refer to [“Admin Domain member types”](#) on page 491.

A newly created Admin Domain has no zoning defined and the default access mode is No Access. This means the devices in the Admin Domain cannot communicate with each other. You must set up zones in the newly created Admin Domain to allow devices to access each other, even if the devices were already zoned together prior to your moving them to the Admin Domain. Refer to [“Admin Domains, zones, and zone databases”](#) on page 510 for additional information about how zones work with Admin Domains.

You create Admin Domains in the transaction buffer. You can either save the newly created Admin Domain to a defined configuration or make it the effective configuration directly.

The following procedure describes the steps for creating Admin Domains.

1. Log in to the switch as the physical fabric administrator.
2. Disable Virtual Fabrics, if necessary, as described in [“Disabling Virtual Fabrics mode”](#) on page 325. Admin Domains and Virtual Fabrics cannot co-exist.
3. Set the default zone mode to No Access, if you have not already done so. Refer to [“Setting the default zoning mode”](#) on page 361 for instructions.

4. Switch to the AD255 context, if you are not already in that context:

```
ad --select 255
```
5. Enter the `ad --create` command using the `-d` option to specify device and switch port members and the `-s` option to specify switch members:

```
ad --create ad_id -d "dev_list" -s "switch_list"
```
6. Enter the appropriate command based on whether you want to save or activate the Admin Domain definition:
 - To save the Admin Domain definition, enter `ad --save`.
 - To save the Admin Domain definition and directly apply the definition to the fabric, enter `ad --apply`.
7. Set up zones in the newly created Admin Domain. Refer to [Chapter 12, “Administering Advanced Zoning,”](#) for instructions.

Example of creating Admin Domains

The following example creates Admin Domain AD1, consisting of two switches, which are designated by domain ID and switch WWN.

```
switch:AD255:admin> ad --create AD1 -s "97; 10:00:00:60:69:80:59:13"
```

The following example creates Admin Domain “blue_ad,” consisting of two switch ports (designated by *domain,index*), one device (designated by device WWN), and two switches (designated by domain ID and switch WWN).

```
switch:AD255:admin> ad --create blue_ad -d "100,5; 1,3; 21:00:00:e0:8b:05:4d:05" -s "97; 10:00:00:60:69:80:59:13"
```

User assignments to Admin Domains

After you create an Admin Domain, you can specify one or more user accounts as the valid accounts that can use that Admin Domain. User accounts have the following characteristics with regard to Admin Domains:

- A user account can have only a single role.
- You can configure a user account to have access to the physical fabric through AD255 and to a list of Admin Domains (AD0 through AD254).
- You can configure a user account to have access to only a subset of your own Admin Domain list. Only a physical fabric administrator can create another physical fabric administrator user account.
- Users capable of using multiple Admin Domains can designate one of these Admin Domains as the home Admin Domain, which is the default Admin Domain context after login.
- If you do not specify one, the home Admin Domain is the lowest valid Admin Domain in the numerically-sorted AD list.
- Users can log in to their Admin Domains and create their own Admin Domain-specific zones and zone configurations.

Creating a new user account for managing Admin Domains

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **userConfig --add** command using the **-r** option to set the role, the **-a** option to provide access to Admin Domains, and the **-h** option to specify the home Admin Domain.

```
userconfig --add username -r role -h home_AD -a "AD_list"
```

Example of creating new user accounts

The following example creates new user account ad1admin with an admin role and assigns one Admin Domain, blue_ad1, to it. This example also assigns blue_ad1 as the user's home Admin Domain.

```
switch:admin> userconfig --add ad1admin -r admin -h blue_ad1 -a "blue_ad1"
```

The following example creates new user account ad2admin with an admin role, access to Admin Domains 1 and 2, and home Admin Domain set to 2.

```
switch:admin> userconfig --add ad2admin -r admin -h 2 -a "1,2"
```

Assigning Admin Domains to an existing user account

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **userConfig --addad** command using the **-a** option to provide access to Admin Domains and the **-h** option to specify the home Admin Domain.

```
userconfig --addad username -h home_AD -a "AD_list"
```

Example

The following example assigns Admin Domain green_ad2 to the existing user account ad1admin.

```
switch:admin> userconfig --addad ad1admin -a "green_ad2"
```

Creating a physical fabric administrator user account

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **userConfig --add** command using the **-r** option to set the role to admin and the **-a** option to provide access to Admin Domains 0 through 255.

```
userconfig --add username -r admin -h home_AD -a "0-255"
```

Example

The following example creates new user account pfa_admin1 with an admin role, access to all Admin Domains (AD0 through AD255), and home Admin Domain set to 255. This user account is now a physical fabric administrator.

```
switch:admin> userconfig --add pfa_admin1 -r admin -h 255 -a "0-255"
```

Removing an Admin Domain from a user account

When you remove an Admin Domain from an account, all of the currently active sessions for that account are logged out.

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

```
userconfig --deletead username [-h admindomain_ID] [-a admindomain_ID_list]
```

If the `-h` argument is not specified, the home Admin Domain either remains as it was or becomes the lowest Admin Domain ID in the remaining list.

Example of removing Admin Domain `green_ad2` from the user account `adm1`

```
switch:admin> userconfig --deletead adm1 -a "green_ad2"
```

```
Broadcast message from root (pts/0) Wed Jan 27 20:57:14 2010...
```

```
Security Policy, Password or Account Attribute Change: adm1 will be logged out  
Ads for account adm1 has been successfully deleted.
```

Activating an Admin Domain

An Admin Domain can be in either an active or inactive state. When you create an Admin Domain, it is automatically in the active state.

1. Connect to the switch and log in using an account with admin permissions.
2. Switch to the AD255 context, if you are not already in that context.

```
ad --select 255
```

3. Enter the `ad --activate` command.

```
ad --activate ad_id
```

You are prompted for confirmation.

By default, after the Admin Domain is activated, the devices specified under that AD are not able to see each other until they are zoned together.

4. Enter the appropriate command based on whether you want to save or activate the Admin Domain definition:
 - To save the Admin Domain definition, enter `ad --save`.
 - To save the Admin Domain definition and directly apply the definition to the fabric, enter `ad --apply`.

Example

The following example activates Admin Domain `AD_B5`.

```
switch:AD255:admin> ad --activate AD_B5  
You are about to activate a new admin domain.  
Do you want to activate 'AD_B5' admin domain (yes, y, no, n): [no]: y  
switch:AD255:admin>
```

Deactivating an Admin Domain

If you deactivate an Admin Domain, the members assigned to the Admin Domain can no longer access their hosts or storage unless those members are part of another Admin Domain.

You cannot log in to an Admin Domain that has been deactivated. You must activate an Admin Domain before you can log in to it.

1. Connect to the switch and log in using an account with admin permissions.
2. Disable the zone configuration under the Admin Domain you want to deactivate.

```
cfgdisable
```

3. Switch to the AD255 context, if you are not already in that context.

```
ad --select 255
```

4. Enter the **ad --deactivate** command.

```
ad --deactivate ad_id
```

You are prompted for confirmation.

5. Enter the appropriate command based on whether you want to save or activate the Admin Domain definition:

- To save the Admin Domain definition, enter **ad --save**.
- To save the Admin Domain definition and directly apply the definition to the fabric, enter **ad --apply**.

All active user sessions associated with the Admin Domain are terminated. The **ad --deactivate** command does not disable ports.

Example of deactivating Admin Domain AD_B4

```
switch:AD255:admin> ad --deactivate AD_B4
You are about to deactivate an AD.
This operation will fail if an effective zone configuration exists in the AD
Do you want to deactivate 'AD_B5' admin domain (yes, y, no, n): [no] y
switch:AD255:admin>
```

Adding members to an existing Admin Domain

1. Connect to the switch and log in using an account with admin permissions.
2. Switch to the AD255 context, if you are not already in that context.

```
ad --select 255
```

3. Enter the **ad --add** command using the **-d** option to specify device and switch port members and the **-s** option to specify switch members.

```
ad --add ad_id -d "dev_list" -s "switch_list"
```

In the syntax, *ad_id* is the Admin Domain name or number, *dev_list* is a list of device WWNs or *domain,index* members, and *switch_list* is a list of switch WWNs or domain IDs.

4. Enter the appropriate command based on whether you want to save or activate the Admin Domain definition:
 - To save the Admin Domain definition, enter **ad --save**.
 - To save the Admin Domain definition and directly apply the definition to the fabric, enter **ad --apply**.

Example of adding two switch ports, designated by *domain,index*, to AD1

```
switch:AD255:admin> ad --add AD1 -d "100,5; 4,1"
```

Removing members from an Admin Domain

If you remove the last member of an Admin Domain, that Admin Domain is automatically deleted.

1. Connect to the switch and log in using an account with admin permissions.
2. Switch to the AD255 context, if you are not already in that context.
3. Enter the **ad --remove** command using the **-d** option to specify device and switch port members and the **-s** option to specify switch members.

```
ad --remove ad_id -d "dev_list" -s "switch_list"
```

Removing the last member element of an Admin Domain deletes the Admin Domain.

4. Enter the appropriate command based on whether you want to save or activate the Admin Domain definition:
 - To save the Admin Domain definition, enter **ad --save**.
 - To save the Admin Domain definition and directly apply the definition to the fabric, enter **ad --apply**.

Example 1

The following example removes port 5 of domain 100 and port 3 of domain 1 from AD1.

```
switch:AD255:admin> ad --remove AD1 -d "100,5; 1,3"
```

Example 2

The following example removes switch 100 from the membership list of AD4.

```
switch:AD255:admin> ad --remove AD4 -s "100"
```

Renaming an Admin Domain

Use this procedure if you want to change the name of an Admin Domain. You can also change auto-assigned names (ADn).

The rename operation does not take effect if the Admin Domain you want to rename is part of the effective configuration.

1. Connect to the switch and log in using an account with admin permissions.
2. Switch to the AD255 context, if you are not already in that context.

```
ad --select 255
```

3. Enter the **ad --rename** command with the present name and the new name.

```
ad --rename present_name new_name
```

4. Enter the appropriate command based on whether you want to save or activate the Admin Domain definition:

- To save the Admin Domain definition, enter **ad --save**.
- To save the Admin Domain definition and directly apply the definition to the fabric, enter **ad --apply**.

The Admin Domain numbers remain unchanged after the operation.

Example of changing the name of Admin Domain Eng_AD to Eng_AD2

```
switch:AD255:admin> ad --rename Eng_AD Eng_AD2
```

Deleting an Admin Domain

When you delete an Admin Domain, its devices no longer have access to the members of the zones with which it was associated.

1. Connect to the switch and log in using an account with admin permissions.
2. Switch to the Admin Domain that you want to delete.

```
ad --select ad_id
```

3. Enter the appropriate command to clear the zone database under the Admin Domain you want to delete.

- To remove the effective configuration, enter **cfgdisable**.
- To remove the defined configuration, enter **cfgclear**.
- To save the changes to nonvolatile memory, enter **cfgsave**.

4. Switch to the AD255 context.

```
ad --select 255
```

5. Enter the **ad --delete** command.

```
ad --delete ad_id
```

The **ad --delete** command prompts you for confirmation before triggering the deletion. The command succeeds whether the Admin Domain is in an activated or deactivated state.

6. Enter the **ad --apply** command to save the Admin Domain definition and directly apply the definition to the fabric.

Example of deleting Admin Domain AD_B3

```
switch:AD255:admin> ad --delete AD_B3
You are about to delete an AD.
This operation will fail if zone configuration exists in the AD
Do you want to delete 'AD_B3' admin domain (yes, y, no, n): [no] y
switch:AD255:admin>
```

Deleting all user-defined Admin Domains

When you clear the Admin Domain configuration, all user-defined Admin Domains are deleted, the explicit membership list of ADO is cleared, and all fabric resources (switches, ports, and devices) are returned to the implicit membership list of ADO.

You cannot clear the Admin Domain configuration if zone configurations exist in any of the user-defined Admin Domains.

If you want to remove all Admin Domains while retaining device connectivity (for example, if you want to enable Virtual Fabrics), use the procedure described in [“Deleting all user-defined Admin Domains non-disruptively.”](#)

1. Clear all individual AD zone databases, in separate transactions, before proceeding with this operation. Refer to [“Clearing all zone configurations”](#) on page 367 for instructions.
2. Connect to the switch and log in using an account with admin permissions.
3. Switch to the AD255 context, if you are not already in that context.

```
ad --select 255
```

4. Enter the **ad --clear** command.

This option prompts you for confirmation before triggering the deletion of all Admin Domains.

5. Enter the **ad --apply** command to save the Admin Domain definition and directly apply the definitions to the fabric.

Example

```
switch:AD255:admin> ad --clear
You are about to delete all ADs definitions.
This operations will fail if zone configurations exists in AD1-AD254
Do you want to clear all admin domains (yes, y, no, n): [no] y
switch:AD255:admin>
```

Deleting all user-defined Admin Domains non-disruptively

To disable Admin Domains non-disruptively, you must do the following before you clear the user-defined ADs:

- Create and activate zone configurations in ADO that are equivalent to the zone configurations in each of the user-defined ADs.
- Define all of the members that are currently in user-defined ADs in ADO.

This will ensure that the devices are able to communicate when they are removed from the user-defined ADs.

You can use this procedure to remove all Admin Domains before enabling Virtual Fabrics.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **cfgshow** command in the AD255 context to display the zone configurations for all Admin Domains.

```
ad --exec 255 "cfgshow"
```

3. Enter the **zone --copy** command to copy the zones from all user-defined Admin Domains to ADO.

```
zone --copy source_AD.source_name dest_name
```

In this syntax, *source_AD* is the name of the user-defined AD from which you are copying the zone, *source_name* is the name of the zone to be copied, and *dest_name* is the name to give to the zone after it is copied to ADO.

4. Copy the newly added zones in ADO to the zone configuration.

```
cfgadd "cfgName", "member[;member]"
```

5. Enable the configuration to complete the transaction.

```
cfgenable cfgName
```

6. Switch to the AD255 context.

```
ad --select 255
```

7. Explicitly add devices that are present in the user-defined ADs to ADO.

```
ad --add AD0 -d "dev_list"
```

8. Enter the **ad --apply** command to save the Admin Domain definition and directly apply the definitions to the fabric.

```
ad --apply
```

At this point, all of the devices in the user-defined ADs are also defined and zoned in ADO.

9. Clear the user-defined ADs.

```
ad --clear -f
```

10. Enter the **ad --apply** command to save the Admin Domain definition and directly apply the definitions to the fabric.

```
ad --apply
```

All user-defined Admin Domains have now been removed, but all device communication that was allowed with the original Admin Domain configuration is still permitted in the context of ADO.

Example

The following example assumes the configuration shown in [Figure 71](#) on page 504:

- Three Admin Domains: AD0, plus two user-defined Admin Domains (AD1 and AD2).
- AD0 has two devices, WWN1 and WWN2, in the AD0_RedZone.
- AD1 has two devices, WWN2 and WWN3, in the AD1_BlueZone.
- AD2 has two devices, WWN4 and WWN5, in the AD2_GreenZone.
- The device WWN2 is in both AD0 and AD1.




FIGURE 71 AD0 and two user-defined Admin Domains, AD1 and AD2

At the conclusion of the procedure, all devices and zones are moved to AD0, and the user-defined Admin Domains are deleted, as shown in Figure 72.




FIGURE 72 AD0 with three zones

```
sw0:admin> ad --exec 255 "cfgshow"

Zone CFG Info for AD_ID: 0 (AD Name: AD0, State: Active) :

Defined configuration:
cfg: AD0_cfg AD0_RedZone
zone: AD0_RedZone
      10:00:00:00:01:00:00:00; 10:00:00:00:02:00:00:00

Effective configuration:
cfg: AD0_cfg
zone: AD0_RedZone
      10:00:00:00:01:00:00:00
      10:00:00:00:02:00:00:00

Zone CFG Info for AD_ID: 1 (AD Name: AD1, State: Active) :

Defined configuration:
cfg: AD1_cfg AD1_BlueZone
zone: AD1_BlueZone
```



```

10:00:00:00:02:00:00:00; 10:00:00:00:03:00:00:00

Effective configuration:
cfg:   AD1_cfg
zone:  AD1_BlueZone
      10:00:00:00:02:00:00:00
      10:00:00:00:03:00:00:00

Zone CFG Info for AD_ID: 2   (AD Name: AD2, State: Active) :

Defined configuration:
cfg:   AD2_cfg AD2_GreenZone
zone:  AD2_GreenZone
      10:00:00:00:04:00:00:00; 10:00:00:00:05:00:00:00

Effective configuration:
cfg:   AD2_cfg
zone:  AD2_GreenZone
      10:00:00:00:04:00:00:00
      10:00:00:00:05:00:00:00

sw0:admin> zone --copy AD1.AD1_BlueZone AD0_BlueZone
sw0:admin> zone --copy AD2.AD2_GreenZone AD0_GreenZone
sw0:admin> cfgadd "AD0_cfg", "AD0_BlueZone; AD0_GreenZone"
sw0:admin> cfgenable AD0_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 'AD0_cfg' configuration (yes, y, no, n): [no] y
zone config "AD0_cfg" is in effect
Updating flash ...

sw0:admin> ad --select 255
sw0:AD255:admin> ad --add AD0 -d "10:00:00:00:03:00:00:00;
10:00:00:00:04:00:00:00; 10:00:00:00:05:00:00:00"
sw0:AD255:admin> ad --apply
You are about to enforce the saved AD configuration.
This action will trigger AD apply to all switches in the fabric
Do you want to apply all admin domains (yes, y, no, n): [no] y

sw0:AD255:admin> ad --clear -f
You are about to delete all ADs definitions and zone databases under them.
This could involve multiple independent zone transactions and
no auto recovery will be done in case of failure in the middle.
Do you want to clear all admin domains (yes, y, no, n): [no] y

sw0:AD255:admin> ad --apply
You are about to enforce the saved AD configuration.
This action will trigger AD apply to all switches in the fabric
Do you want to apply all admin domains (yes, y, no, n): [no] y

```

Validating an Admin Domain member list

You can validate the device and switch member list. You can list non-existing or offline Admin Domain members. You can also identify misconfigurations of the Admin Domain.

The Admin Domain validation process is not applicable for ADO, because ADO implicitly contains all unassigned online switches and their devices.

1. Connect to the switch and log in using an account with admin permissions.
2. Switch to the AD255 context, if you are not already in that context.

```
ad --select 255
```

3. Enter the `ad --validate` command.

```
ad --validate ad_id -m mode
```

If you do not specify any parameters, the entire AD database (transaction buffer, defined configuration, and effective configuration) is displayed.

If you do not specify an Admin Domain, information about all existing Admin Domains is displayed.

The `-m mode` option can be used with the following values:

- 0 to display the Admin Domain configuration in the current transaction buffer.
- 1 to display the Admin Domain configuration stored in the persistent memory (defined configuration).
- 2 to display the currently enforced Admin Domain configuration (effective configuration).

Example of validating the member list of Admin Domain 10 in the current transaction buffer

```
switch:AD255:admin> ad --validate 10 -m 0
Current AD Number: 255  AD Name: AD255

Transaction buffer configuration:
-----
AD Number: 2  AD Name: ad2  State: Active
      Switch port members: 1,1; 1,3; 2,5+; 3,6;
-----
* - Member does not exist
+ - Member is AD Unaware
```

SAN management with Admin Domains

This section is for both users and administrators and describes how Admin Domains affect commands and other Fabric OS features. If you are a physical fabric administrator and you want to create, modify, or otherwise manage Admin Domains, refer to [“Admin Domain management for physical fabric administrators”](#) on page 494.

The Admin Domain looks like a virtual switch or fabric to a user. However, based on the user role and type (User_ID), users are presented with only their relevant AD-based views (refer to [Figure 66](#) on page 486 and [Figure 67](#) on page 486). Any devices and switch ports that are not defined as part of the Admin Domain are not shown and are not available to that AD user.

Each Admin Domain can also have its own zone configurations (defined and effective) with zones and aliases under them.

CLI commands in an AD context

The CLI command input arguments are validated against the AD member list; they do not work with input arguments that specify resources that are not members of the current Admin Domain. All commands present filtered output, showing only the members of the current Admin Domain.

For example, **switchShow** displays details for the list of AD members present in that switch. Note the following about the **switchShow** output:

- Because all E_Ports and EX_Ports are shared across all Admin Domains, they are shown under all Admin Domains.
- Other ports are displayed without any attribute details (with an explanation that they are not part of the current Admin Domain).

A port or device appears in CLI command output or other management tool outputs if any one of the conditions listed in [Table 80](#) is met.

TABLE 80 Ports and devices in CLI output

For	Condition
<i>domain,index</i>	<ul style="list-style-type: none"> • The port is specified in the <i>domain,index</i> member list of the Admin Domain. • One or more WWNs specified in the AD member list is attached to the <i>domain,index</i>.
Device WWN	<ul style="list-style-type: none"> • The device WWN is specified in the AD WWN member list. • The device WWN is attached to one of the <i>domain,index</i> members specified in the AD member list.

RASlog and syslog output is not filtered based on AD membership.

Refer to the *Fabric OS Command Reference* for more detailed information about command syntax and usage and to understand how existing commands behave in an AD context.

Executing a command in a different AD context

You can execute a command in an Admin Domain that is different from your current AD context. The Admin Domain must be one that you can access. This option creates a new shell with the current User_ID, switches to the specified Admin Domain, performs the specified command, and exits the shell.

1. Connect to the switch and log in.
2. Enter the **ad --exec** command, specifying the Admin Domain and the command you want to execute.

```
ad --exec ad_id "command"
```

Example of executing the switchShow command in the AD7 context

```
switch:AD255:admin> ad --exec 7 "switchshow"
```

Displaying an Admin Domain configuration

You can display the membership information and zone database information of a specified Admin Domain. Notice the following differences in the information displayed based on the Admin Domain:

- AD255: If you do not specify the AD name or number, all information about all existing Admin Domains is displayed.
- AD0–AD254: The membership of the current Admin Domain is displayed.
- AD0: The device and switch list members are categorized into implicit and explicit member lists.

1. Connect to the switch and log in as any user type.
2. Enter the **ad --show** command.

```
ad --show
```

If you are in the AD0 context, you can use the **-i** option to display the implicit membership list of AD0; otherwise, only the explicit membership list is displayed.

```
ad --show -i
```

If you are in the AD255 context, all Admin Domain configurations from the transaction buffer, defined configuration, and effective configuration are displayed, unless you use the **-m** option:

```
ad --show ad_id -m mode
```

In the syntax, *ad_id* is the Admin Domain for which you want to display information and *mode* is one of the following values:

- 0 to display the Admin Domain configuration in the current transaction buffer.
- 1 to display the Admin Domain configuration stored in the persistent memory (defined configuration).
- 2 to display the currently enforced Admin Domain configuration (effective configuration).

Example of displaying membership information about AD1

```
switch:AD1:admin> ad --show
Current AD Number: 1 AD Name: TheSwitches

Effective configuration:
-----

AD Number: 1 AD Name: TheSwitches State: Active

Switch WWN members:          50:06:06:99:00:2a:e9:01;
                              50:00:51:e0:23:36:f9:01;
                              50:06:06:98:05:be:99:01;
```

Switching to a different Admin Domain context

You can switch between different Admin Domain contexts. This option creates a new shell with a new Admin Domain context. If the corresponding Admin Domain is not activated, the operation fails.

1. Connect to the switch and log in as any user type.
2. Enter the **ad --select** command and the Admin Domain to which you want to switch.
3. Leave the new Admin Domain context by exiting from the shell.

```
logout
```

You cannot switch to another Admin Domain context from within the shell created by **ad --select**. You must first exit the shell, and then issue the **ad --select** command again.

Example of switching to a different Admin Domain context

The following example switches to the AD12 context and back. Note that the prompt changes to display the Admin Domain.

```
switch:admin> ad --select 12
switch:AD12:admin> logout
switch:admin>
```

Admin Domain interactions with other Fabric OS features

The Admin Domain feature provides interaction with other Fabric OS features and across third-party applications. You can manage Admin Domains with Web Tools as well as the CLI. If the current Admin Domain owns the switch, you can perform Fabric Watch operations.

Admin Domain interactions do not extend to user session tunneling across switches. A user logged in to a switch can control only the local switch ports as specified in the Admin Domain.

When the fabric is in secure mode, the following restrictions apply:

- There is no support for ACL configuration under each Administrative Domain.
- ACL configuration commands are allowed only in ADO and AD255. None of the policy configurations are validated with AD membership.

[Table 81](#) lists some of the Fabric OS features and considerations that apply when using Admin Domains.

TABLE 81 Admin Domain interaction with Fabric OS features

Fabric OS feature	Admin Domain interaction
ACLs	If no user-defined Admin Domains exist, you can run ACL configuration commands in only ADO and AD255. If any user-defined Admin Domains exist, you can run ACL configuration commands only in AD255. You <i>cannot</i> use ACL configuration commands or validate ACL policy configurations against AD membership under each Admin Domain.
Advanced Performance Monitoring (APM)	All APM-related filter setup and statistics viewing is allowed only if the local switch is part of the current Admin Domain.
Configuration upload and download	Refer to “Configuration upload and download in an AD context” on page 512 for details.
Fabric Watch	Fabric Watch configuration operations are allowed only if the local switch is part of the current Admin Domain.
FC-FC Routing Service	You can create LSAN zones as a physical fabric administrator or as an individual AD administrator. The LSAN zone can be part of the root zone database or the AD zone database. FCR collects the LSAN zones from all ADs. If both edge fabrics have matching LSAN zones and both devices are online, FCR triggers a device import. LSAN zone enforcement in the local fabric occurs only if the AD member list contains both of the devices (local and imported devices) specified in the LSAN zone. To support legacy applications, WWNs are reported based on the AD context using NAA=5. As a result, you cannot use the NAA=5 field alone in the WWN to detect an FC router.

TABLE 81 Admin Domain interaction with Fabric OS features (Continued)

Fabric OS feature	Admin Domain interaction
FDMI	FDMI operations are allowed only in ADO and AD255.
FICON	Admin Domains support FICON. However, you must perform additional steps because FICON management requires additional physical control of the ports. You must set up the switch as a physical member of the FICON AD. Device Connection Control (DCC) and Switch Connection Control (SCC) policies are supported only in ADO and AD255, because ACL configurations are supported only in ADO and AD255.
iSCSI	iSCSI operations are supported only in ADO.
LSAN zoning	Refer to “Admin Domains and LSAN zones” on page 511 for details.
Management applications	Management interfaces that access the fabric without a user’s credentials continue to get the physical fabric view. Examples include SNMPv1, Web Tools, HTTP access, unzoned management server query, FAL in-band CT requests from FAL Proxy to FAL Target, and FC-CT-based management applications. Access from applications or hosts using management server calls can be controlled using the management server ACL support provided by the msConfigure command. Note that this is a switch-specific setting and not a fabric-wide setting.
Port swapping and PID formats	Admin Domain port members are specified in <i>domain,index</i> format. Based on the PID format, a <i>domain,index</i> member indicates a slot and port in the switch. The <i>domain,index</i> member is effectively a member of that AD. Port swapping has no effect on AD support as port swapping swaps only the area numbers of two ports and Admin Domains are specified using <i>domain,index</i> members. For detailed information about configuring the PID format, refer to Chapter 3, “Performing Advanced Configuration Tasks” .
RSCN	Admin Domains do not introduce any RSCN changes to devices or hosts.
Virtual Fabrics	Virtual Fabrics and Admin Domains are mutually exclusive and are not supported at the same time on a switch. To use Admin Domains, you must first disable Virtual Fabrics; to use Virtual Fabrics, you must first delete all Admin Domains. If you connect a switch with Admin Domains to a Virtual Fabrics-enabled switch, the link is segmented with the reason “VF AD conflict.”
Zoning	Refer to “Admin Domains, zones, and zone databases” on page 510 for details.

Admin Domains, zones, and zone databases

Admin Domains introduce two types of zone database nomenclature and behavior:

- **Root zone database**
If you do not use Admin Domains, there is only one zone database. This legacy zone database is known as the *root zone database*. If you create Admin Domains, several zone databases exist: the root zone database, which is owned by ADO, and other zone databases, one for each user-defined Admin Domain.
AD-level zone information is merged with the root zone configuration and enforced.
- **AD zone databases**
Each AD (AD1 through AD254) has its own zone database, with the defined and effective zone configurations and all related zone objects (zones, zone aliases, and zone members). Each AD has its own zone transaction buffer. Within an Admin Domain, you can configure zoning only with the devices that are present in that Admin Domain.

The AD zone database also has the following characteristics:

- Each zone database has its own name space. For example, you can define a zone name of `test_z1` in more than one Admin Domain.
- There is no zone database linked to the physical fabric (AD255) and no support for zone database updates. In the physical fabric context (AD255), you can only view the complete hierarchical zone database, which is all of the zone databases in ADO through AD254.
- You can concurrently edit the separate zone databases.
- With AD support, zoning updates are supported selectively at each AD level. For example, a zone change in AD1 results in an update request only for the AD1 zone database.

Zoning operations ignore any resources not in the Admin Domain, even if they are specified in the zone. The behavior functions similarly to specifying offline devices in a zone. All zones from each AD zone configuration are enforced. The enforcement policy encompasses zones in the effective zone configuration of the root zone database and the effective zone configurations of each AD.

NOTE

You must define the Admin Domain members in the same way that they are defined as zone members. That is, if an object is defined by WWN in a zone, it must be defined by WWN in an Admin Domain. If it is defined by *domain,index* in a zone, it must be defined by *domain,index* in the Admin Domain. If both zoning schemes are used, then objects must be defined in the Admin Domain by both WWN and *domain,index*.

Using the **zone --validate** command, you can see all zone members that are not part of the current zone enforcement table but *are* part of the zoning database. A member might not be part of the zone enforcement table for the following reasons:

- The device is offline.
- The device is online but is not part of the current Admin Domain.

Refer to [“Validating a zone”](#) on page 358 for instructions on using the **zone --validate** command.

NOTE

AD zone databases do not have an enforced size limit. The zone database size is calculated by the upper limit of the AD membership definition and the sum of all the zone databases for each AD.

Admin Domains support the default zone mode of No Access only. Before configuring any Admin Domain, you must set the default zone to No Access mode. Admin Domains without effective zone configurations are presented with No Access. Refer to [“Default zoning mode”](#) on page 360 for more information.

If the administrative domain feature is not active (AD1 through AD254 are not configured and no explicit members are added to ADO), ADO supports both All Access and No Access default zone modes.

Admin Domains and LSA zones

Logical storage area networks (LSANs) under each Admin Domain are collated into a single name space and sent out to FCR phantom domains using the following format:

```
<original_LSAN_name>_AD<AD_num>
```

For example, a zone with name `lsan_for_linux_farm` in AD5 is internally converted to `lsan_for_linux_farm_AD005`.

LSAN zone names in ADO are never converted for backward-compatibility reasons.

The auto-converted LSAN zone names might collide with LSAN zone names in ADO (in the example, if ADO contains `lsan_for_linux_farm_ADO05`, this causes a name collision). Fabric OS does not detect or report such name clashes.

LSAN zone names greater than 57 characters are not converted or sent to the FCR phantom domain.

LSAN zones defined within an Admin Domain must contain devices that are applicable to that Admin Domain only. A device must not be included in more than one LSAN zone across multiple Admin Domains. Device discovery problems might occur if LSAN zones in one Admin Domain contain devices that belong to another Admin Domain.

Refer to [Chapter 26, “Using FC-FC Routing to Connect Fabrics,”](#) for information about LSAN zones.

Configuration upload and download in an AD context

The behavior of the `configUpload` and `configDownload` commands varies depending on the AD context and whether the switch is a member of the current Admin Domain. In the AD context, these commands include only the zone configuration of the current Admin Domain. If the switch is a member of the Admin Domain, all switch configuration parameters are saved and the zone database for that Admin Domain is also saved.

[Table 82](#) lists the sections in the configuration file and the Admin Domain contexts in which you can upload and download these sections. Refer to [Chapter 9, “Maintaining the Switch Configuration File,”](#) for additional information about uploading and downloading configurations.

NOTE

You cannot use `configDownload` to restore a single Admin Domain. To restore a single Admin Domain, you must first delete all Admin Domains and then issue `configDownload` to restore them.

TABLE 82 Configuration upload and download scenarios in an AD context

AD contexts	Configuration file sections				
	iSCSI	ACL	Zone	AD headers	Switch configuration and other parameters
AD255: With ADs	Yes	Yes	Yes ¹	Yes	Yes
Without ADs	Yes	Yes	Yes ¹	Yes	Yes
ADO: With ADs and switch membership	Yes	No	Yes ²	No	Yes
With ADs and without switch membership	Yes	No	Yes ²	No	No
Without ADs	Yes	Yes	Yes ²	No	Yes
AD1 – AD254: With switch membership	No	No	Yes ³	No	Yes
Without switch membership	No	No	Yes ³	No	No

1. Zone databases for ADO through AD254.
2. Only zone database for ADO.
3. Only zone database for current AD.

The `configDefault` command does not clear zone or Admin Domain database information. This command is allowed only if the switch is a member of the current Admin Domain.

Licensed Features

This section describes optionally licensed Brocade Fabric OS features and includes the following chapters:

- [Chapter 21, “Administering Licensing”](#)
- [Chapter 22, “Inter-chassis Links”](#)
- [Chapter 23, “Monitoring Fabric Performance”](#)
- [Chapter 24, “Managing Trunking Connections”](#)
- [Chapter 25, “Managing Long-Distance Fabrics”](#)
- [Chapter 26, “Using FC-FC Routing to Connect Fabrics”](#)

Administering Licensing

In this chapter

- [Licensing overview](#) 515
- [Brocade 7800 Upgrade license](#) 523
- [ICL licensing](#) 523
- [8G licensing](#) 525
- [Slot-based licensing](#) 526
- [10G licensing](#) 527
- [Temporary licenses](#) 530
- [Viewing installed licenses](#) 532
- [Activating a license](#) 533
- [Adding a licensed feature](#) 533
- [Removing a licensed feature](#) 534
- [Ports on Demand](#) 535

Licensing overview

Feature licenses are often part of the licensed paper pack supplied with your switch software; if not, they can be purchased separately from your switch vendor, who provides the transaction keys to activate the associated feature or features. Each product, each feature, and each individual switch within a fabric requires its own license key.

Licenses may be associated with a feature version. If a feature has a version-based license, that license is valid only for a particular version of the feature. If you want a newer version of the feature, you must purchase a new license. If a license is not version-based, then it is valid for all versions of the feature. Likewise, if you downgrade Fabric OS to an earlier version, some licenses associated with specific features of the version you are downgrading may not work.

NOTE

To preserve licenses and the functioning of features associated with the licenses installed on your switch, use the **configUpload** command before you upgrade or downgrade Fabric OS.

You can use the **licenseShow** command to display the license keys on a switch. Save the output to a text file in a secure location. If licenses are lost or removed from the switch, you can use the saved output to recover or add the lost licenses.

Fabric OS includes basic switch and fabric support software, and support for optionally licensed software that is enabled using license keys.

Some licenses may display with the text “Obsolete license.” This happens because of changes in licensing requirements of some features that no longer require a license key, yet are still installed on a switch.

ATTENTION

The Adaptive Networking and Server Application Optimization (SAO) licenses are no longer required to be explicitly installed in Fabric OS 7.2.0 and later. These licenses are automatically enabled for new switches and for existing switches that are upgraded to Fabric OS 7.2.0 or later.

If these licenses were installed on the switch prior to upgrading to Fabric OS 7.2.0, they will display as “Obsolete license”. Do not remove these licenses, as they are required if you downgrade and want to maintain QoS functionality.

Some Brocade HBA configuration settings require the Adaptive Networking or SAO license to be installed on the switch. Starting in Fabric OS 7.2.0, these licenses are always installed.

Table 83 lists the optionally licensed features that are available in Fabric OS 7.2.

TABLE 83 Available Brocade licenses

License	Description
10 Gigabit FCIP/Fibre Channel (10G license)	<ul style="list-style-type: none"> Allows 10 Gbps operation of FC ports on the Brocade 6510 or 6520 switches or the FC ports of FC16-32 or FC16-48 port blades installed on a Brocade DCX 8510 Backbone. Enables the two 10-GbE ports on the FX8-24 extension blade when installed on the Brocade DCX, DCX-4S, DCX 8510-4, or DCX 8510-8 Backbone. Allows selection of the following operational modes on the FX8-24 blade: <ul style="list-style-type: none"> 10 1-GbE ports and 1 10-GbE port, or 2 10-GbE ports License is slot-based when applied to a Brocade Backbone. It is chassis-based when applied to a Brocade 6510 or 6520 switch.
7800 Upgrade	<ul style="list-style-type: none"> Enables full hardware capabilities on the Brocade 7800 base switch, increasing the number of Fibre Channel ports from four to sixteen and the number of GbE ports from two to six. Supports up to eight FCIP tunnels instead of two. Supports advanced capabilities such as tape read/write pipelining. <p>NOTE: The Brocade 7800 switch must have the 7800 Upgrade license to add FICON Management Server (CUP) or Advanced FICON Acceleration licenses. Refer to “Brocade 7800 Upgrade license” on page 523 for details.</p>

TABLE 83 Available Brocade licenses (Continued)

License	Description
Advanced Extension	<ul style="list-style-type: none"> Enables two advanced extension features: FCIP Trunking and Adaptive Rate Limiting. FCIP Trunking feature allows all of the following: <ul style="list-style-type: none"> Multiple (up to 4) IP source and destination address pairs (defined as FCIP Circuits) using multiple (up to 4) 1-GbE or 10-GbE interfaces to provide a high bandwidth FCIP tunnel and failover resiliency. Support for up to 4 of the following QoS classes: Class-F, high, medium and low priority, each as a TCP connection. The Adaptive Rate Limiting feature provides a minimum bandwidth guarantee for each tunnel with full usage of available network bandwidth without any negative impact to throughput performance under high traffic load. Available on the Brocade 7800 switch, and the Brocade DCX and DCX-4S and the Brocade DCX 8510 family for the FX8-24 on an individual slot basis.
Advanced FICON Acceleration	<ul style="list-style-type: none"> Allows use of specialized data management techniques and automated intelligence to accelerate FICON tape read and write and IBM Global Mirror data replication operations over distance, while maintaining the integrity of command and acknowledgement sequences. Available on the Brocade 7800 switch, and the Brocade DCX and DCX-4S and the Brocade DCX 8510 family for the FX8-24 on an individual slot basis.
Brocade Advanced Performance Monitoring	<ul style="list-style-type: none"> Enables performance monitoring of networked storage resources. Includes the Top Talkers feature.
Brocade Extended Fabrics	<p>Provides greater than 10 km of switched fabric connectivity at full bandwidth over long distances (depending on the platform, this can be up to 3000 km).</p> <p>NOTE: This license is not required for long distance connectivity using licensed 10G ports.</p>
Brocade Fabric Watch	<ul style="list-style-type: none"> Monitors mission-critical switch operations. Includes Port Fencing capabilities.
Brocade ISL Trunking	<ul style="list-style-type: none"> Provides the ability to aggregate multiple physical links into one logical link for enhanced network performance and fault tolerance. Includes Access Gateway ISL Trunking on those products that support Access Gateway deployment.
Brocade Ports on Demand	<p>Allows you to instantly scale the fabric by provisioning additional ports using license key upgrades.</p> <p>NOTE: Applies to the Brocade 300, 5100, 5300, M6505, 6505, 6510, 6520, 6547, and VA-40FC switches.</p>
DataFort Compatibility	<p>Provides ability to read, write, decrypt, and encrypt the NetApp DataFort-encrypted Disk LUNs and Tapes to all of the following:</p> <ul style="list-style-type: none"> Brocade Encryption Switch Brocade enterprise platforms with FS8-18 blade <p>Includes metadata, encryption and compression algorithms.</p> <p>NOTE: Availability is limited. Contact your vendor for details.</p>

TABLE 83 Available Brocade licenses (Continued)

License	Description
Encryption Performance Upgrade	Provides additional encryption bandwidth on encryption platforms. For the Brocade Encryption Switch, two Encryption Performance Upgrade licenses can be installed to enable the full available bandwidth. On a Brocade enterprise platform, a single Performance License can be installed to enable full bandwidth on all FS8-18 blades installed in the chassis.
Enhanced Group Management	Enables full management of the device in a data center fabric with deeper element management functionality and greater management task aggregation throughout the environment. This license is used in conjunction with Brocade Network Advisor application software. This license is applicable to all of the Brocade 8G and 16G FC platforms. NOTE: This license is enabled by default on all 16G FC platforms, and on DCX and DCX-4S platforms that are running Fabric OS v7.0.0 or later. This license is not included by default on 8G FC fixed port switches (300, 5100, 5300, VA-40FC, and 8G FC embedded switches).
Enterprise ICL	Allows you to connect more than four chassis in a fabric using ICLs. You can connect up to four Brocade DCX 8510 Backbones via ICLs without this license. If the number of interconnected chassis using ICLs exceeds four, then all of the chassis using ICLs require the Enterprise ICL license. You must also have an ICL POD license on each Brocade DCX 8510 to activate the ICL ports. The Enterprise ICL license only allows connection of more than four chassis using ICLs in a fabric; it does not enable the ICL ports on a chassis. NOTE: Applies to the Brocade DCX 8510 Backbone family only.
Fabric Vision (FV)	Allows you to activate the following features: <ul style="list-style-type: none"> • Monitoring and Alerting Policy Suite (MAPS) • Flow Vision • Run D_Port tests between a switch and non-Brocade HBAs This license replaces the Advanced Performance Monitor (APM) and Fabric Watch (FW) licenses. If you have the Fabric Vision license, you can use Advanced Performance Monitoring and Fabric Watch features without the APM and FW licenses.
FCoE	Enables Fibre Channel over Ethernet (FCoE) functions.
FICON Management Server (Also known as Control Unit Port or “CUP”)	Enables host-control of switches in mainframe environments.
High Performance Extension over FCIP/FC (formerly known as “FC-IP Services”)	Includes the IPsec capabilities.
ICL 8-Link	Activates all eight links on ICL ports on a Brocade DCX-4S or half of the ICL bandwidth for each ICL port on the Brocade DCX platform by enabling only eight links out of the sixteen links available. This allows you to purchase half the bandwidth of DCX ICL ports initially and upgrade with an additional ICL 8-Link license to utilize the full ICL bandwidth at a later time. This license is also useful for environments that want to create ICL connections between a DCX and a DCX-4S; the latter cannot support more than eight links on an ICL port. Available on the Brocade DCX and DCX-4S Backbones only.

TABLE 83 Available Brocade licenses (Continued)

License	Description
ICL 16-Link	Activates all 16 links on ICL ports on a Brocade DCX chassis. Each chassis must have the ICL 16-Link license installed in order to enable the full 16-link ICL connections. Available on the Brocade DCX only.
Integrated Routing	Allows any ports in Brocade 5100, 5300, 6510, 6520, 6547, and VA-40FC switches, the Brocade Encryption Switch, or the Brocade DCX, DCX-4S, and DCX 8510 family platforms to be configured as an EX_Port supporting FC-FC routing.
Inter-Chassis Link (1st POD)	Activates half of the ICL bandwidth on a DCX 8510-8, or all the ICL bandwidth on a DCX 8510-4, allowing you to enable only the bandwidth needed, and upgrade to additional bandwidth at a later time. This license is also useful for environments that wish to create ICL connections between a DCX 8510-8 and a DCX 8510-4; the latter platform supports only half the number of ICL links that the former platform supports. Available on the Brocade DCX 8510 Backbones only.
Inter-Chassis Link (2nd POD)	Activates the remaining ICL bandwidth on the Brocade DCX 8510-8 chassis. Each chassis must have this ICL license installed in order to enable all available ICL connections. Available on the Brocade DCX 8510-8 only.

Table 84 lists licensed features, each feature's associated license name, and, if applicable, the location on the local or any connecting switch on which the license must be installed.

TABLE 84 License requirements and location name by feature

Feature	License	Where license should be installed
Adaptive Rate Limiting	Advanced Extension	Local switch.
Administrative Domains	No license required.	N/A
Bottleneck Detection	No license required.	N/A
Brocade Network Advisor	No license required for base use.	See the <i>Brocade Network Advisor User Manual</i> .
Configuration up/download	No license required. NOTE: The <code>configUpload</code> and <code>configDownload</code> commands are provided automatically with Fabric OS on the switch.	N/A
D_Port	No license required for D_Port tests between two switches or between a switch and a Brocade HBA. Fabric Vision license is required for D_Port tests between a switch and a non-Brocade HBA.	If Fabric Vision license is required, it should be installed on the local switch.
Diagnostic tools	No license required.	N/A
Distributed Management Server	No license required.	N/A
Enterprise ICL	Enterprise ICL	Each ICL-connected Brocade DCX 8510 chassis in the fabric when there are five or more such chassis in the fabric.

TABLE 84 License requirements and location name by feature (Continued)

Feature	License	Where license should be installed
Extended Fabrics	Extended Fabrics	Local switch and any attached switches.
Fabric Watch	No license required for baseline monitoring capabilities. Fabric Watch license or Fabric Vision license required for full functionality.	See the <i>Fabric Watch Administrator's Guide</i> .
FCIP	High Performance Extension over FCIP/FC	NOTE: Local and attached switches. License is needed on both sides of tunnel.
FCIP Trunking	Advanced Extension	Local and attached switches.
Fibre Channel Routing/EX_Ports	Integrated Routing	Local switch.
FICON	No license required.	N/A
FICON-CUP	FICON Management Server	Local switch.
FICON Tape Read and Write Emulation over an FCIP Tunnel	<ul style="list-style-type: none"> FICON Tape High Performance Extension over FCIP/FC license or Advanced FICON Acceleration license on Brocade 7800 	Local and attached switches.
FICON XRC Sequence Emulation over an FCIP Tunnel	<ul style="list-style-type: none"> FICON XRC High Performance Extension over FCIP/FC license or Advanced FICON Acceleration license on Brocade 7800 	Local and attached switches.
FIPS	No license required.	N/A
Firmware download	No license required. NOTE: The firmwareDownload command is provided automatically with Fabric OS on the switch.	N/A
Flow Vision: <ul style="list-style-type: none"> Flow Generator Flow Performance Monitor Flow Mirror 	Fabric Vision NOTE: If you have both the Advanced Performance Monitoring and the Fabric Watch licenses installed, you do not need the Fabric Vision license.	Local switch.
Full fabric connectivity	Full Fabric NOTE: Also called the Fabric license (visible in licenseShow output) and the E_Port Upgrade license.	Local switch. May be required on attached switches.
In-flight encryption and compression	No license required.	N/A
Inband Management	No license required.	N/A
Ingress rate limiting	No license required. Adaptive Networking with QoS license is required for switches running Fabric OS versions earlier than 7.2.0. The Brocade 6520 does not require a license regardless of Fabric OS version.	N/A for local switches running Fabric OS 7.2.0 or later. License required on local switches running Fabric OS versions earlier than 7.2.0.

TABLE 84 License requirements and location name by feature (Continued)

Feature	License	Where license should be installed
Inter-chassis link (ICL)	<ul style="list-style-type: none"> ICL 1st POD (Ports on Demand) on the Brocade DCX 8510 Backbone family only. ICL 2nd POD on the Brocade DCX 8510-8 only. ICL 8-link on the Brocade DCX and DCX-4S only. ICL 16-link on the Brocade DCX only. Enterprise ICL on the Brocade DCX 8510 Backbone family only, for topologies with more than four chassis with ICLs. 	Local and attached platforms.
IPSec	No license required.	N/A
IPsec for FCIP tunnels	High Performance Extension over FCIP/FC	NOTE: Local and attached switches. License is needed on both sides of tunnel.
LDAP	No license required.	N/A
Logical fabric	No license required.	N/A
Logical switch	No license required.	N/A
Long distance	Extended Fabrics	Local and attached switches. NOTE: License is needed on both sides of connection.
Monitoring and Alerting Policy Suite (MAPS)	Fabric Vision NOTE: If you have both the Advanced Performance Monitoring and the Fabric Watch licenses installed, you do not need the Fabric Vision license.	Local switch.
NPIV	No license required.	N/A
OpenSSH public key	No license required.	N/A
Performance monitoring	Advanced Performance Monitoring or Fabric Vision license for advanced features. No license required for basic features.	Local switch.
Port fencing	Fabric Watch	Local switch.
Ports	<ul style="list-style-type: none"> Ports on Demand licenses required, applicable to a select set of switches only. 7800 Upgrade license for the 7800 switches to use all ports. 10 Gigabit FCIP/Fibre Channel license to use 10Gb FC ports on FC16-32 blades, FC16-48 blades, and the Brocade 6510 and 6520. 10 Gigabit FCIP/Fibre Channel license to enable 10Gb Ethernet ports on the FX8-24 extension blades. 	Local switch.

TABLE 84 License requirements and location name by feature (Continued)

Feature	License	Where license should be installed
QoS	No license required. Adaptive Networking with QoS license is required for switches running Fabric OS versions earlier than 7.2.0. The Brocade 6520 does not require a license regardless of Fabric OS version.	N/A for local switches running Fabric OS 7.2.0 or later. License required on local and attached switches running Fabric OS versions earlier than 7.2.0.
QoS on HBA	No license required. Adaptive Networking with QoS license is required for switches running Fabric OS versions earlier than 7.2.0. The Brocade 6520 does not require a license regardless of Fabric OS version.	N/A for local switches running Fabric OS 7.2.0 or later. License required on local switches running Fabric OS versions earlier than 7.2.0.
RADIUS	No license required.	N/A
RBAC	No license required.	N/A
Routing traffic	No license required. NOTE: Port-based or exchanged-based routing, static routes, frame-order deliver, and dynamic routes all included.	N/A
Security	No license required. NOTE: DCC, SCC, FCS, IP Filter, and authentication policies all included.	N/A
SNMP	No license required.	N/A
Speed	8 Gbps license needed to support 8 Gbps on the Brocade 300, 5100, 5300, and VA-40FC switches and embedded switches only. NOTE: The 8 Gbps license is installed by default, and you should not remove it. A 10-Gb FCIP/Fibre Channel license is needed to support 10Gb FC ports on FC16-32 blades, FC16-48 blades, and the Brocade 6510 and 6520, as well as to support the 10Gb Ethernet ports on FX8-24 blades. (See the Ports feature above for more information.)	Local switch
SSH public key	No license required.	N/A
TACACS+	No license required.	N/A
Top Talkers	Advanced Performance Monitoring	Local switch and attached switches.
Traffic Isolation	No license required.	N/A
Trunking	<ul style="list-style-type: none"> • ISL Trunking or • ISL Trunking Over Extended Fabrics For ICL trunking, no license is required.	Local and attached switches.
Two-to-four domains in a fabric	Value Line (Two/Four)	Local switch. May be required on attached switches.
USB usage	No license required.	N/A
Virtual Fabrics	No license required.	N/A

TABLE 84 License requirements and location name by feature (Continued)

Feature	License	Where license should be installed
Web Tools	No license required.	Local and any switch you will be managing using Web Tools.
Zoning	No license required.	N/A

Brocade 7800 Upgrade license

The Brocade 7800 has four Fibre Channel (FC) ports and two GbE ports active by default. The number of physical ports active on the Brocade 7800 is fixed. There is one upgrade license to activate the rest of the FC and GbE ports for a total of 16 FC ports and 6 GbE ports. The 7800 Upgrade license activates FC and GbE ports, and also activates additional features outlined in [Table 85](#).

NOTE

You must reboot the Brocade 7800 switch after installing the 7800 Upgrade license.

TABLE 85 Base to Upgrade license comparison

Feature	Base model	7800 Upgrade license
Number of Fibre Channel (FC) ports	4	16
Number of GbE ports	2	6
Number of 10-GbE ports	0	0
Number of FCIP Tunnels	2	8
Tape Pipelining over FCIP Tunnel	No	Yes

ICL licensing

Brocade ICL links operate between the core blades of the DCX 8510 Backbone family, or between the core blades of the DCX and DCX-4S Backbones. Typically, if both core blades are installed, then they are active on the DCX and DCX-4S (or DCX 8510 family) Backbones.

ICL ports on core blades of a DCX 8510-8 can be used only with an ICL (1st or 2nd) POD license. ICL ports on core blades of a DCX 8510-4 can be used only with an ICL 1st POD license.

ICL ports on core blades of a DCX can be used only with an ICL 16-link or ICL 8-link license. ICL ports on core blades of a DCX-4S can be used only with an ICL 8-link license.

After the addition or removal of a license, the license enforcement is performed on the ICL ports only when the **portDisable** and **portEnable** commands are issued on the ports. An ICL license must be installed on the enterprise platforms at both ends of the ICL connection.

ICL 1st POD license

The ICL 1st POD license activates half of the ICL bandwidth on the Brocade DCX 8510-8 platform or all of the ICL bandwidth on the Brocade DCX 8510-4.

On the Brocade DCX 8510-8, this license enables QSFP ports 0–7; QSFP ports 8–15 are disabled. (QSFP ports 0–7 correspond to core blade port numbers 0–31, and QSFP ports 8–15 correspond to core blade port numbers 32–63, as observed in **switchShow** output.)

This license allows you to purchase half the bandwidth of the Brocade DCX 8510-8 ICL ports initially and upgrade with an additional ICL license to use the full ICL bandwidth later. This license is also useful for environments with ICL connections between a Brocade DCX 8510-8 and a DCX 8510-4, as the latter supports half the bandwidth of the DCX 8510-8 on each ICL port.

This license is available on the Brocade DCX 8510-8 and DCX 8510-4 platforms only.

ICL 2nd POD license

The ICL 2nd POD license provides dedicated high-bandwidth links between two Brocade DCX 8510-8 platforms without consuming valuable front-end ports. Each Brocade DCX 8510-8 platform must have the ICL 2nd POD license installed to enable the full number of ICL connections possible.

This license is available for the Brocade DCX 8510-8 only.

ICL 8-link license

The ICL 8-link license activates half of the ICL bandwidth for each ICL port on the Brocade DCX platform by enabling only half of the ICL links available. This allows you to purchase half the bandwidth of the Brocade DCX ICL ports initially and upgrade with an additional ICL license to use the full ICL bandwidth later. This license is also useful for environments with ICL connections between a Brocade DCX and a DCX-4S, as the latter cannot support more than eight links on an ICL port.

This license is available on the Brocade DCX-4S and DCX platforms only.

ICL 16-link license

The ICL 16-link license provides dedicated high-bandwidth links between two Brocade DCX chassis, without consuming valuable front-end ports. Each Brocade DCX chassis must have the ICL 16-link license installed in order to enable the full number of ICL connections possible (16 links in the case of a DCX chassis).

This license is available for the Brocade DCX only.

Enterprise ICL license

The Enterprise ICL (EICL) license allows you to connect more than four Brocade DCX 8510 Backbones through ICLs. This license is available on the Brocade DCX 8510-8 and DCX 8510-4 platforms only.

The EICL license is required in addition to the ICL POD license.

The following requirements apply:

- Connection of four or fewer DCX 8510 Backbones with ICLs does not require the EICL license. However, if you add additional ICL-connected chassis, then all ICL-connected chassis require the EICL license.
- With the EICL license installed, a maximum of 10 chassis are allowed to be connected together via ICLs.

- When Virtual Fabrics are used, the limit on the number of chassis connected together via ICLs depends only on the physical chassis and not on the logical switches.
- If the maximum number of ICL-connected chassis exceeds the allowed limit with or without the EICL license, additional links may either be disabled or segmented. The disabling or segmenting reason code depends on whether the EICL license is installed.
- If ICL links to a chassis become segmented for non-EICL-related reasons, these links are part of the fabric, and the chassis containing these segmented links is included in the maximum chassis count. If the maximum chassis count (with or without the EICL license) is reached with these segmented links, then any additional links will become segmented. Therefore, to add additional links, you first must disable the links that became segmented due to non-EICL reasons. This should reduce the maximum chassis count and allow the new links to join.

Example switchShow output if no Enterprise ICL license is installed

A message such as the following is displayed if a required EICL license is not installed:

```
440    8    24    -----    id    16G        Online        FC    E-Port
segmented,10:00:00:05:33:0d:52:00 (No EICL License) (Trunk master)
441    8    25    -----    id    16G        Online        FC    E-Port
segmented,10:00:00:05:33:0d:52:00 (No EICL License) (Trunk master)
```

Example switchShow output if maximum number of chassis is reached

A message such as the following is displayed if the maximum number of supported chassis is reached:

```
384    5    0    -----    id    16G        Online        FC    E-Port
segmented,10:00:00:05:1e:39:bf:9a (EICL License Limited) (Trunk master)
385    5    1    -----    id    16G        Online        FC    E-Port
segmented,10:00:00:05:1e:39:bf:9a (EICL License Limited) (Trunk master)
```

8G licensing

ATTENTION

The 8G license is installed by default and you should not remove it. Port operation may become disrupted, and ports may be prevented from operating at 8 Gbps when the license is removed.

The 8G license applies to the Brocade 300, 5100, 5300, and VA-40FC switches and the 8 Gbps embedded switches; this license does *not* apply to the Brocade 6505, 6510, or 6520.

The following list describes the basic rules of using, adding, or removing 8G licenses:

- Without an 8G license, even if there is an 8 Gbps SFP plugged into a port in an applicable platform, the port would be enabled to run at a maximum speed of 4 Gbps.
- To obtain an 8G license, only the license ID from the switch is required. When you add the 8G license, you must enter either the **portDisable** and **portEnable** commands on each individual port on the switch, or the **switchDisable** and **switchEnable** commands on the switch, to enable 8 Gbps features.
- When you remove the 8G license, the ports that are online and already running at 8 Gbps are not disturbed until the port goes offline or the switch is rebooted. The switch ports return to their pre-licensed state maximum speed of 4 Gbps.

Slot-based licensing

Slot-based licensing is used on the Brocade DCX and DCX 8510 Backbone families to support the FX8-24 blade, and on the Brocade DCX 8510 Backbone family to support the 16-Gbps FC port blades (FC16-24 and FC16-48). License capacity is equal to the number of slots. These licenses allow you to select the slots that the license will enable up to the capacity purchased and to increase the capacity without disrupting slots that already have licensed features running. Each slot-based license key is for a single feature.

Features utilizing slot-based licenses on the FX8-24 blade include:

- 10 Gigabit FCIP/Fibre Channel
- Advanced Extension
- Advanced FICON Acceleration

NOTE

The 10 GbE feature on the FX8-24 blade and the 10 Gbps FC feature on the 16-Gbps FC blades are both enabled by the same 10 Gigabit FCIP/Fibre Channel license (10G license). This license can also enable the 10 Gbps FC feature on a Brocade 6510 or 6520 switch as a chassis-based license.

Any unassigned slot-based license will be automatically assigned to applicable blades that are detected in the chassis when the license is installed. If you have more applicable blades than available license capacity, then you can manually assign or re-assign the licenses as necessary.

Once a license is assigned to a slot, whether it has been automatically assigned or manually assigned, the assignment will remain until you manually reassign the license to another slot. This design allows for various maintenance operations to occur without having the license move around to other slots.

Use the following procedure to activate a slot-based licensed feature:

1. Install a slot-based license on the platform with sufficient slot count for the number of slots upon which you plan to activate the feature.
2. Configure slots so that the licensed feature is assigned to slots. No more slots can be configured than specified in the license.
3. Configure the application that uses the licensed feature on the blade in the slot. This operation verifies that the previous two steps have been successfully completed.

Once these steps are complete, the feature will work on the blade.

Upgrade and downgrade considerations

When a slot-based license is present on the switch, firmware downgrade to pre-Fabric OS v6.3.0 is allowed, but the slot-based features that were licensed will not be functional.

On upgrade to Fabric OS v7.0.0 or later, any slot-based license that displayed the 10-GbE feature name in the earlier release now appears as “10 Gigabit FCIP/Fibre Channel (FTR_10G) license.”

Assigning a license to a slot

Use the following procedure to assign a license to a slot.

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions in the license class of RBAC commands.
2. Enter the **licenseSlotCfg -add** command to add the license to the appropriate slot.

Removing a license from a slot

Use the following procedure to remove a slot-based license from a blade slot.

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions in the license class of RBAC commands.
2. Deconfigure the application that uses the licensed feature on the blade slot.
3. Enter the **licenseSlotCfg -remove** command to remove the license from the slot.

10G licensing

The 10 Gbps FCIP/Fibre Channel license (10G license) enables the following features:

- 10 Gbps access on the 16-Gbps FC ports on Brocade 6510 or 6520 switches, and FC16-32 and FC16-48 port blades.
- The two 10-GbE ports on the FX8-24 extension blade.

This 10G license is applied as a slot-based license on the FC16-32 and FC16-48 port blades and on the FX8-24 extension blade; generic rules for adding slot-based licenses apply, as described in [“Slot-based licensing”](#) on page 526. When this license is applied to a Brocade 6510 or 6520 switch, it is applied to the whole chassis.

Whether you have a bladed (DCX, DCX-4S, DCX 8510-8, or DCX 8510-4) platform or nonbladed (Brocade 6510 or 6520) switch, you add the 10G license to the chassis using the **licenseAdd** command, as for any license.

For the bladed platforms, you can either allow automatic license assignment, or choose the blades you want the licenses assigned to manually, as for any slot-based license. Automatic assignment is done sequentially by slot number, beginning with the lowest numbered slot with an enabled blade that supports this feature (FX8-24, FC16-32, or FC16-48 blade), and that does not already have the license applied. If the automatic license assignment does not match your needs, you can use the **licenseSlotCfg --remove** and **licenseSlotCfg --add** commands to remove the license manually from a slot and assign it to a different slot with an FX8-24, FC16-32, or FC16-48 blade.

The same multiple slot-based 10G license can be applied to a mixture of 16-Gbps blades and FX8-24 blades. For example, if you have a 10G license for two-slot capacity, and you have an FX8-24 blade in one slot and an FC16-48 blade in a second slot, then the same license can activate the 10GE ports on the FX8-24 blade and enable 10 Gbps operation on the 10-Gbps FC ports on the FC16-48 blade.

After applying a 10G license to the Brocade 6510 or 6520 chassis or to a 16-Gbps FC blade, you must also configure the port octet (**portCfgOctetSpeedCombo** command) with the correct port octet speed group and configure each port to operate at 10 Gbps (**portCfgSpeed** command). It is necessary to configure the port octet because only certain combinations of port speeds are allowed within the port octet. No license is required for the octet group. If the speed configuration operation succeeds and a 10G-capable SFP is inserted in the port connector, the port will allow operation at 10 Gbps when the link becomes active at that speed.

Before removing a 10G license from an entire platform (**licenseRemove** command) or from a specific blade (**licenseSlotCfg --remove** command), you must first deconfigure all affected FC ports to no longer operate at 10 Gbps.

NOTE

An FC port that is operating at 10 Gbps FC speed on a 16-Gbps FC blade or 16-Gbps FC switch does not need an Extended Fabrics license to be used for FC long distance connectivity.

FC ports licensed and configured to operate at 10 Gbps on a Brocade 6510 or 6520 switch or 16-Gbps FC port blade cannot interoperate with 10-Gbps FC ports on the M-6140 platform or the FC10-6 blade. The new FC ports use different protocols and physical connections.

Enabling 10 Gbps operation on an FC port

Use the following procedure to enable 10 Gbps operation on an FC port on a Brocade 6510 or 6520 switch or an FC16-32 or FC16-48 blade:

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the license and SwitchPortConfiguration classes of RBAC commands.
2. Use the **licenseAdd** command to add the 10G license.
3. Use the **licenseShow** command to verify the license.

Bladed platforms only: If the results of the automatic license assignment are not what you intended, use the **licenseSlotCfg** command to reassign the license to the desired blades.

4. Use the **portCfgOctetSpeedCombo** command to set the combination speed for the port octet to a setting that supports 10 Gbps operations. Valid settings for 10 Gbps operations include:
 - 2—Auto-negotiated or fixed port speeds of 10 Gbps, 8 Gbps, 4 Gbps, and 2 Gbps
 - 3—Auto-negotiated or fixed port speeds of 16 Gbps and 10 Gbps
5. Use the **portCfgSpeed** command to set the port speed on each port you want to operate at 10 Gbps.

Example of assigning a 10G license on an FC port blade and enabling 10 Gbps operation on a port

This example assigns a license to slot 4 on a DCX 8510-8 Backbone and enables 10 Gbps operation on port 2 of the port blade in that slot. In this example, the 10G license was first automatically assigned to slot 1.

```
8510-8switch:admin> licenseadd aTFPNFXGLmABANMGtT4LfSBJSDLWTYD3EFrr4WGAEMBA
8510-8switch:admin> licenseshow
aTFPNFXGLmABANMGtT4LfSBJSDLWTYD3EFrr4WGAEMBA
  10 Gigabit FCIP/Fibre Channel (FTR_10G) license
  Capacity 1
  Consumed 1
  Configured Blade Slots 1
8510-8switch:admin> licenseslotcfg -remove FTR_10G 1
8510-8switch:admin> licenseslotcfg -add FTR_10G 4
8510-8switch:admin> licenseshow
aTFPNFXGLmABANMGtT4LfSBJSDLWTYD3EFrr4WGAEMBA
  10 Gigabit FCIP/Fibre Channel (FTR_10G) license
  Capacity 1
  Consumed 1
  Configured Blade Slots 4
8510-8switch:admin> portcfgoctetspeedcombo 4/2 2
```



```
8510-8switch:admin> portcfgspeed 4/2 10
8510-8switch:admin>
```

Example of assigning a 10G license on a Brocade 6510 and enabling 10 Gbps operation on a port

This example assigns a license to a Brocade 6510 switch and enables 10 Gbps operation on port 2.

```
6510-switch:admin> licenseadd aTFPNFXGLmABANMGtT4LfSBJSDDLWTD3EFrr4WGAEMBA
6510-switch:admin> licenseshow
aTFPNFXGLmABANMGtT4LfSBJSDDLWTD3EFrr4WGAEMBA
  10 Gigabit FCIP/Fibre Channel (FTR_10G) license
  Capacity 1
  Consumed 1
6510-switch:admin> portcfgoctetspeedcombo 2 2
6510-switch:admin> portcfgspeed 2 10
```

Enabling the 10-GbE ports on an FX8-24 blade

Use the following procedure to enable the 10-GbE ports on an FX8-24 blade:

1. Connect to the Brocade Backbone and log in using an account with admin permissions, or an account with OM permissions for the license class of RBAC commands.
2. Use the **licenseAdd** command to add the 10G license.
3. Use the **licenseShow** command to check the results of automatic license assignment. If the results are not what you intended, use the **licenseSlotCfg** command to reassign the license to the desired FX8-24 blades.
4. Use the **licenseShow** command to verify the license.
5. Use the **bladeCfgGeMode --set mode** command to configure the GbE port mode for the FX8-24 blade. To enable the 10-GbE ports, set the *mode* parameter to one of the following:
 - **10g**—Enables both 10-GbE ports, disables all ten 1-GbE ports.
 - **dual**—Enables the xge0 port (but not xge1) and also enables all ten 1-GbE ports.

Example of assigning a 10G license on an FX8-24 extension blade and enabling both 10-GbE ports

This example assigns a license to slot 7 on a DCX 8510-4 Backbone and enables both 10-GbE ports on the FX8-24 blade in that slot. In this example, the license was first automatically assigned to slot 1.

```
8510-4switch:admin> licenseadd aTFPNFXGLmABANMGtT4LfSBJSDDLWTD3EFrr4WGAEMBA
8510-4switch:admin> licenseshow
aTFPNFXGLmABANMGtT4LfSBJSDDLWTD3EFrr4WGAEMBA
  10 Gigabit FCIP/Fibre Channel (FTR_10G) license
  Capacity 1
  Consumed 1
  Configured Blade Slots 1
8510-4switch:admin> licenseslotcfg -remove FTR_10G 1
8510-4switch:admin> licenseslotcfg -add FTR_10G 7
8510-4switch:admin> licenseshow
aTFPNFXGLmABANMGtT4LfSBJSDDLWTD3EFrr4WGAEMBA
  10 Gigabit FCIP/Fibre Channel (FTR_10G) license
  Capacity 1
  Consumed 1
  Configured Blade Slots 7
8510-4switch:admin> bladecfggemode --set 10G -slot 7
```

21 Temporary licenses

```
8510-4switch:admin> switchshow -slot 7
...
158  7  30  019e00  --  --  Offline  VE
159  7  31  019f00  --  --  Offline  VE
      7  ge0      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge1      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge2      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge3      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge4      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge5      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge6      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge7      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge8      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  ge9      --  1G  No_Module FCIP  Disabled (10G Mode)
      7  xge0     --  10G No_Module FCIP
      7  xge1     --  10G No_Module FCIP
```

Temporary licenses

A temporary license applies a “try-before-you-buy” approach to certain features so that you can experience the feature and its capabilities prior to buying the license. Once you have installed the license, you are given a time limit to use the feature. A temporary license can be either a regular temporary license or a universal temporary license.

- A regular temporary license is available on a per-switch basis.
- A universal temporary license can be installed on a switch, but can be applied to multiple switches.

The following licenses are available as temporary or universal temporary licenses:

- 10 Gigabit FCIP/Fibre Channel license (slot-based)
- Advanced Extension license (slot-based)
- Advanced FICON Acceleration license (slot-based)
- Advanced Performance Monitoring license
- Enterprise ICL license
- Fabric (E_Port) license
- Fabric Watch license
- FICON Management Server (CUP) license
- Extended Fabrics license
- High Performance Extension over FCIP/FC license
- Integrated Routing license
- ISL Trunking license

Restrictions on upgrading temporary slot-based licenses

If the capacity of the permanent license is equal to or greater than the capacity of the temporary license and you use the same slot assignments, then replacing the temporary license with a permanent license is non-disruptive. If either condition changes, however, then the process is disruptive.

If the permanent license is for fewer slots than the temporary license, you must do the following:

1. Remove the temporary license. The removal process disables the feature.
2. Install the permanent license on the appropriate slots.

If the permanent license is for different slots than the temporary license, you must do the following:

1. Install the permanent license. The temporary license is automatically replaced on the original slots.
2. Deconfigure the application that uses the licensed feature on the original slots.
3. Remove the license from the original slots using the **licenseSlotCfg -remove** command.
4. Add the license to the new slots using the **licenseSlotCfg -add** command.

Date change restriction

Once the temporary license is installed, you cannot change the time of the switch until the temporary license is removed. To change the time, you must remove the license, change the date and time, and then re-install the license on the switch.



CAUTION

If you are using NTP to synchronize the time between your network devices, including switches or Backbones, then do not attempt to change the system date and time when a temporary license is installed.

Configupload and download considerations

The **configDownload** and **configUpload** commands download the legacy, enhanced, consumed capacities, and temporary licenses.

Expired licenses

Once a temporary license has expired, you can view it through the **licenseShow** command. Expired licenses have an output string of "License has expired". RASlog warning messages are generated every hour for licenses present in the database which have expired or are going to expire in the next five days. An expired license may become unusable after a reboot, failover, firmware download, or a port or switch disable or enable operation.

Removing an expired license



CAUTION

This procedure is disruptive to the switch.

Use the following procedure to remove an expired license.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **reboot** command for the expiry to take affect.

Universal temporary licenses

Universal temporary license keys include a duration period. Once installed on a switch, an expiration date is calculated and the duration is decremented until there is no remaining time, at which point it is expired. Because of this, universal temporary licenses should not be installed on a switch until you are ready to use or test the feature, so as not to unnecessarily consume a portion of the temporary-use duration.

The expiration date is based on the system time at the installation of the license plus the number of days for which the universal temporary license is valid. Universal temporary licenses cannot be removed and reinstalled on the same switch.

Universal temporary licenses are always retained in the license database on the product even though they can be explicitly deleted from any user interface.

Extending a universal temporary license

Extending a universal temporary license is done by adding a temporary license with an expiry date after the universal temporary license expiry date, or by adding a permanent license. Re-applying an existing universal temporary license is not allowed.

Universal temporary license shelf life

All universal temporary licenses are encoded with a “shelf life” expiration date. Once this date is reached, the temporary licensed feature can no longer be used on the switch.

Viewing installed licenses

Use the following procedure to view all installed licenses.

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

Activating a license

The transaction key is case-sensitive; it must be entered exactly as it appears in the paperpack. To lessen the chance of error, copy and paste the transaction key. The quotation marks are optional.

Use the following procedure to activate a license.

1. Take the appropriate action based on whether you have a license key:
 - If you have a license key, go to [“Adding a licensed feature”](#).
 - If you do not have a license key and are using a transaction key, launch an Internet browser and go to the Brocade website at <http://www.brocade.com>.

2. Select **Products > Software License Keys**.

The **Software License Keys** instruction page appears.

3. Enter the requested information in the required fields and click **Next**.

A verification screen appears.

4. Verify the information appears correctly.

Click **Submit** if the information displayed is correct. If the information is incorrect, click **Previous**, correct the information, and click **Submit**.

An information screen displays the license keys and you will receive an e-mail with the software license keys and installation instructions.

Adding a licensed feature

To enable a feature, go to the feature’s appropriate section in this manual. Enabling a feature on a switch may be a separate task from adding the license.

For the Brocade Backbones, licenses are effective on both control processor (CP) blades, but are valid only when the CP blade is inserted into a Backbone that has an appropriate license ID stored in the WWN card. If a CP is moved from one Backbone to another, the license works in the new Backbone only if the WWN card is the same in the new Backbone. Otherwise, you must transfer licenses from the old platform to the new platform by obtaining new licenses for the previously licensed features using the new license ID.

For example, if you swap one CP blade at a time, or replace a single CP blade, then the existing CP blade (the active CP blade) propagates the licenses to the new CP blade if the WWN card has been moved to the new platform.

If you move a standby CP from one Backbone to another, then the active CP will propagate its configuration (including license keys) onto that standby CP.

Use the following procedure to add a licensed feature.

1. Connect to the switch and log in using an account with admin permissions.
2. Activate the license using the **licenseAdd** command.
3. Verify the license was added by entering the **licenseShow** command. The licensed features currently installed on the switch are listed. If the feature is not listed, enter the **licenseAdd** command again.

21 Removing a licensed feature

Some features may require additional configuration, or you may need to disable and re-enable the switch to make them operational; see the feature documentation for details.

```
switch:admin> licenseshow
aAYtMJg7tmMZrTZ9JTWBC4SXWLJMY3QfBJYHG:
  Fabric license
  Remote Switch license
  Remote Fabric license
  Extended Fabric license
  Entry Fabric license
  Fabric Watch license
  Performance Monitor license
  Trunking license
  4 Domain Fabric license
  FICON_CUP license
  High-Performance Extension over FCIP/FC license
  Full Ports on Demand license - additional 16 port upgrade license
  2 Domain Fabric license
  Integrated Routing license
  Storage Application Services license
  FICON Tape license
  FICON XRC license
  Inter Chassis Link license
  Enhanced Group Management license
  8 Gig FC license
  DataFort Compatibility license
```

Removing a licensed feature

Use the following procedure to remove a licensed feature.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **licenseShow** command to display the active licenses.
3. Remove the license key using the **licenseRemove** command.

The license key is case-sensitive and must be entered exactly as given. The quotation marks are optional. After removing a license key, the licensed feature is disabled when the switch is rebooted or when a switch disable and enable is performed.

4. Enter the **licenseShow** command to verify the license is disabled.

```
switch:admin> licenseshow
bQebzbRdScRfc0iK:
  Entry Fabric license
  Fabric Watch license
SybbzQQ9edTzcc0X:
  Fabric license
switch:admin> licenseremove "bQebzbRdScRfc0iK"
removing license key "bQebzbRdScRfc0iK"
```

Entering the **licenseShow** command after the **licenseRemove** command displays the remaining licenses.

```
switch:admin> licenseshow
SybbzQQ9edTzcc0X:
  Fabric license
```

If there are no license keys, **licenseShow** displays “No licenses.”

Ports on Demand

The Brocade models in the following list can be purchased with the number of licensed ports indicated. As your needs increase, you can activate unlicensed ports up to a device-constrained maximum by purchasing and installing the optional Ports on Demand licensed product.

- Brocade 300—Can be purchased with 8 ports and no E_Port, 8 ports with full fabric access, or 16 ports with full fabric access. A maximum of 16 ports is allowed; 8-port systems can be upgraded in 4-port increments. An E_Port license upgrade is also available for purchase.
- Brocade 5100—Can be purchased with 24, 32, or 40 licensed ports. A maximum of 40 ports is allowed.
- Brocade 5300—Can be purchased with 48, 64, or 80 licensed ports. A maximum of 80 ports is allowed.
- Brocade M6505—Can be purchased with 12 or 24 licensed ports. A maximum of 24 ports is allowed.
- Brocade 6505—Can be purchased with 12 or 24 licensed ports. A maximum of 24 ports is allowed.
- Brocade 6510—Can be purchased with 24, 36, or 48 licensed ports. A maximum of 48 ports is allowed.
- Brocade 6520—Can be purchased with 48, 72, or 96 licensed ports. A maximum of 96 ports is allowed.
- Brocade 6547—Can be purchased with 12, 24, or 48 licensed ports. A maximum of 48 ports is allowed.
- Brocade VA-40FC—Can be purchased with 24, 32, or 40 licensed ports. A maximum of 40 ports is allowed.

ATTENTION

Licenses are not interchangeable between units. For example, if you bought a POD license for a Brocade 300, you cannot use that license on a Brocade 5100 or VA-40FC. The licenses are based on the switch License Identifiers and are not interchangeable.

[Table 86](#) shows the ports that are enabled by default and the ports that can be enabled after you install the first and second Ports on Demand licenses for each switch type.

TABLE 86 List of available user ports when implementing PODs

Platform	Available user ports, No POD license	Available user ports, POD1 or POD2 present	Available user ports, Both POD licenses present
Brocade 300	0-7	0-15	0-23
Brocade 5100	0-23	0-31	0-39
Brocade 5300	0-47	0-63	0-79
Brocade 5410	0-11	N/A	N/A
Brocade 5424	1-8 and 17-20	POD1: 0, 9-16, and 21-23	0-23
Brocade 5450	1-10 and 19-22	POD1: 0, 11-18, and 23-25	0-25
Brocade 5480	1-8 and 17-20	POD1: 9-12 and 21-22 POD2: 0, 13-16, and 23	0-23
Brocade M6505	1-8 and 17-20	0-23	N/A

TABLE 86 List of available user ports when implementing PODs (Continued)

Platform	Available user ports, No POD license	Available user ports, POD1 or POD2 present	Available user ports, Both POD licenses present
Brocade 6505	0-11	POD 1: 0-23	N/A
Brocade 6510	0-23	0-35	0-47
Brocade 6520	0-47	0-71	0-95
Brocade 6547	0-8 and 29-31	POD1: 0-14 and 29-37 POD2: 0-8, 15-31, and 39-47	0-47
Brocade VA-40FC	0-23	0-31	0-39

Ports on Demand is ready to be unlocked in the switch firmware. Its license key may be part of the licensed paperpack supplied with switch software, or you can purchase the license key separately from your switch vendor. You may need to generate a license key from a transaction key supplied with your purchase. If so, launch an Internet browser and go to the Brocade website at <http://www.brocade.com>. Click **Products > Software Products > Software License Keys** and follow the instructions to generate the key.

Each Ports on Demand license activates the next group of ports in numerical order in either 4-port or 8- or 12-port increments, depending on the model. Before installing a license key, you must insert transceivers in the ports to be activated. Remember to insert the transceivers in the lowest group of inactive port numbers first. For example, if only 16 ports are currently active and you are installing one Ports on Demand license key, make sure to insert the transceivers in ports 16 through 23. If you later install a second license key, insert the transceivers in ports 24 through 31. For details on inserting transceivers, see the switch's hardware reference manual.

Displaying installed licenses

If a single license is installed that enables all Ports on Demand, the license will display as "Full Ports on Demand license - additional X port upgrade license." If there are other individual Ports on Demand licenses installed, these will also be displayed when listing the licenses for a switch, and you will see either "First Ports on Demand license - additional Y port upgrade license" or "Second Ports on Demand license - additional Z port upgrade license." In cases where there are multiple Ports on Demand licenses, the total additional allowed ports will not exceed the total displayed for the "Full Ports on Demand" license.

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

```
switch:admin> licenseshow
SdSSc9SyRSTuTTdz:
    First Ports on Demand license - additional 16 port upgrade license
SdSSc9SyRSTeXTdn:
    Second Ports on Demand license - additional 16 port upgrade license
SdSSc9SyRSTuXTd3:
    Full Ports on Demand license - additional 32 port upgrade license
```

ATTENTION

If you enable or disable an active port, you will disrupt any traffic and potentially lose data flowing on that port.

If the port is connected to another switch, you will segment the switch from the fabric and all traffic flowing between the disabled port and the fabric will be lost.

If you remove a Ports on Demand license, the licensed ports will become disabled after the next platform reboot or the next port deactivation.

Activating Ports on Demand

Use the following procedure to activate Ports on Demand.

1. Connect to the switch and log in using an account with admin permissions.
2. Verify the current states of the ports using the **portShow** command.
In the **portShow** output, the Licensed field indicates whether the port is licensed.
3. Install the Brocade Ports on Demand license.
For instructions on how to install a license, see [“Adding a licensed feature”](#) on page 533.
4. Use the **portEnable** command to enable the ports.
Alternatively, you can disable and re-enable the switch to activate ports.
5. Use the **portShow** command to check the newly activated ports.

Dynamic Ports on Demand

The Dynamic Ports on Demand (POD) feature automatically assigns POD licenses from a pool of available licenses based on the server blade or switch installation.

The following platforms support Dynamic POD:

- Switches:
 - Brocade 6505
 - Brocade 6510
 - Brocade 6520
- Embedded switch modules for bladed servers:
 - Brocade 5410
 - Brocade 5424
 - Brocade 5450
 - Brocade 5460
 - Brocade 5470
 - Brocade 5480
 - Brocade M6505
 - Brocade 6547

For the embedded switch modules, the Dynamic POD feature detects and assigns ports to a POD license only if the server blade is installed with an HBA present. A server blade that does not have a functioning HBA is treated as an inactive link during initial POD port assignment. For the non-server blade switches, the dynamic assignment occurs when an attached Fibre Channel link transitions to the “link active” state.

The Dynamic POD feature assigns the ports to the POD license as they come online. Typically, assignments are sequential, starting with the lowest port number. However, variations in the equipment attached to the ports can cause the ports to take different amounts of time to come online. This means that the port assignment order is not guaranteed.

If the switch detects more active links than allowed by the current POD licenses, then some ports will not be assigned a POD license. Ports that do not receive a POD assignment have a state of *No Sync* or *In Sync*; these ports are not allowed to progress to the online state. Ports that cannot be brought online because of insufficient POD licenses have a state of *(No POD License) Disabled*. You can use the **switchShow** command to display the port states.

Displaying the port license assignments

When you display the available licenses, you can also view the current port assignment of those licenses.

Use the following procedure to display the port license assignments.

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

Example showing manually assigned POD licenses

```
switch:admin> licenseport --show
 24 ports are available in this switch
  Full POD license is installed
  Dynamic POD method is in use
 24 port assignments are provisioned for use in this switch:
    12 port assignments are provisioned by the base switch license
    12 port assignments are provisioned by a full POD license
 24 ports are assigned to installed licenses:
    12 ports are assigned to the base switch license
    12 ports are assigned to the full POD license
Ports assigned to the base switch license:
  1, 2, 3, 4, 5, 6, 7, 8, 17, 18, 19, 20
Ports assigned to the full POD license:
  0, 9, 10, 11, 12, 13, 14, 15, 16, 21, 22, 23
```

Enabling Dynamic Ports on Demand

If the switch is in the static POD mode, then activating the Dynamic POD will erase any prior port license assignments the next time the switch is rebooted. The static POD assignments become the initial Dynamic POD assignments. After the Dynamic POD feature is enabled, you can customize the POD license associations.

Use the following procedure to enable Dynamic Ports on Demand.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **licensePort --method** command with the **dynamic** option to change the license assignment method to dynamic.

```
switch:admin> licenseport --method dynamic
The POD method has been changed to dynamic.
Please reboot the switch now for this change to take effect.
```

3. Enter the **reboot** command to restart the switch.

```
switch:admin> reboot
```

4. Enter the **licensePort --show** command to verify the switch started the Dynamic POD feature.

```
switch:admin> licenseport --show
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use
24 port assignments are provisioned for use in this switch:
    12 port assignments are provisioned by the base switch license
    12 port assignments are provisioned by a full POD license
8 ports are assigned to installed licenses:
    8 ports are assigned to the base switch license
    0 ports are assigned to the full POD license
Ports assigned to the base switch license:
    1, 2, 5, 6, 8*, 21, 22, 23
Ports assigned to the full POD license:
    None
Ports not assigned to a license:
    0, 3, 4, 7, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20

16 license reservations are still available for use by unassigned ports
1 license assignment is held by an offline port (indicated by *)
```

Disabling Dynamic Ports on Demand

Disabling the Dynamic POD feature changes the POD method to static and erases any prior port license associations or assignments the next time the switch is rebooted.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **licensePort --method** command with the **static** option to change the license assignment method to static.

```
switch:admin> licenseport --method static
The POD method has been changed to static.
Please reboot the switch now for this change to take effect.
```

3. Enter the **reboot** command to restart the switch.
4. Enter the **licensePort --show** command to verify the switch changed to static POD.

```
switch:admin> licenseport --show
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use
24 port assignments are provisioned for use in this switch:
12 port assignments are provisioned by the base switch license
12 port assignments are provisioned by a full POD license
24 ports are assigned to installed licenses:
12 ports are assigned to the base switch license
12 ports are assigned to the full POD license
Ports assigned to the base switch license:
1, 2, 3, 4, 5, 6, 7, 8, 17, 18, 19, 20
Ports assigned to the full POD license:
0, 9, 10, 11, 12, 13, 14, 15, 16, 21, 22, 23
```

Reserving a port license

You can allocate licenses by reserving and releasing POD assignments to specific ports. Disabled ports are not candidates for automatic license assignment by the Dynamic POD feature. Persistently disable an otherwise viable port to prevent it from coming online, and thereby preserve a license assignment for another port.

Reserving a license for a port assigns a POD license to that port whether the port is online or offline. That license will not be available to other ports that come online before the specified port.

To allocate licenses to a specific port instead of automatically assigning them as the ports come online, reserve a license for the port. The port receives a POD assignment if any are available.

Use the following procedure to reserve Dynamic Ports on Demand licenses.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **licensePort --show** command to verify there are port reservations available.

```
switch:admin> licenseport --show
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use
24 port assignments are provisioned for use in this switch:
12 port assignments are provisioned by the base switch license
12 port assignments are provisioned by a full POD license
10 ports are assigned to installed licenses:
10 ports are assigned to the base switch license
0 ports are assigned to the full POD license
Ports assigned to the base switch license:
1*, 2*, 3*, 4*, 5*, 6*, 8*, 21, 22, 23
Ports assigned to the full POD license:
None
Ports not assigned to a license:
0, 7, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20
```

3. Take the following appropriate action based on whether port reservations are available:
 - If a port reservation is available, then issue the **licensePort --reserve** command to reserve a license for the port.

```
switch:admin> licenseport --reserve 0
```
 - If all port reservations are assigned, select a port to release its POD license. Follow the instructions in [“Releasing a port from a POD set”](#) to release a port from its POD assignment. Once the port is released, you can reserve it.

Releasing a port from a POD set

Releasing a port removes it from the POD set; the port then appears as “unassigned” until it comes back online. Persistently disabling the port ensures that the port cannot come back online and be automatically assigned to a POD assignment. Before you can re-assign a license, you must disable the port and release the license.

After a port is assigned to the POD set, the port is licensed until it is manually removed from the POD port set. When a port is released from its POD port set (Base, Single, or Double), it creates a vacancy in that port set.

Use the following procedure to release a port from a POD set:

1. Connect to the switch and log in using an account with admin permissions.

2. Enter the **switchDisable** command to take the switch offline.

```
switch:admin> switchdisable
```

3. Enter the **switchShow** command to verify the switch state is offline.

4. Enter the **licensePort --release** command to remove the port from the POD license.

```
switch:admin> licenseport --release 0
```

5. Enter the **licensePort --show** command to verify the port is no longer assigned to a POD set.

```
switch:admin> licenseport --show  
24 ports are available in this switch  
Full POD license is installed  
Dynamic POD method is in use  
24 port assignments are provisioned for use in this switch:  
12 port assignments are provisioned by the base switch license  
12 port assignments are provisioned by a full POD license  
10 ports are assigned to installed licenses:  
10 ports are assigned to the base switch license  
0 ports are assigned to the full POD license  
Ports assigned to the base switch license:  
1*, 2*, 3*, 4*, 5*, 6*, 8*, 21, 22, 23  
Ports assigned to the full POD license:  
None  
Ports not assigned to a license:  
0, 7, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20
```

6. Enter the **switchEnable** command to bring the switch back online.

7. Enter the **switchShow** command to verify the switch state is now online.

Inter-chassis Links

In this chapter

- [Inter-chassis links](#) 543
- [ICLs for the Brocade DCX 8510 Backbone family](#) 544
- [ICLs for the Brocade DCX Backbone family](#) 546
- [Virtual Fabrics considerations for ICLs](#) 547
- [Supported topologies for ICL connections](#) 547

Inter-chassis links

An inter-chassis link (ICL) is a licensed feature used to interconnect two Brocade DCX or DCX 8510 Backbones. ICL ports in the core blades are used to interconnect the Backbones, potentially increasing the number of usable ports in the Backbone chassis.

The Brocade Backbones support two types of ICLs:

- The Brocade DCX 8510 Backbone family supports optical ICL QSFPs.
- The Brocade DCX Backbone family supports proprietary copper ICL connectors.

When two Brocade Backbones are interconnected by ICLs, each chassis requires a unique domain and is managed as a separate switch.

NOTE

You cannot interconnect a Brocade DCX Backbone family chassis to a Brocade DCX 8510 Backbone family chassis.

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. For the Brocade DCX_8510 Backbone family, you can also configure EX_Ports on the ICLs. Refer to [Chapter 26, “Using FC-FC Routing to Connect Fabrics,”](#) for instructions.

When you connect two Brocade Backbones, the following features are supported:

- Trunking
- Buffer-to-buffer credit sharing
- 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.

For more information on how license enforcement occurs, refer to [Chapter 21, “Administering Licensing”](#).

ICLs for the Brocade DCX 8510 Backbone family

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.

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 73](#) shows two Brocade DCX 8510-8 chassis connected with full redundancy using four ICL connections.




FIGURE 73 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 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 Brocade DCX 8510 chassis if the ICLs are in a different logical switch than the ISLs.

ICL trunking on the Brocade DCX 8510-8 and DCX 8510-4

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 [Chapter 21, “Administering Licensing”](#).

Each Quad Small Form-Factor Pluggable (QSFP) cable has four ports, each terminating on a different application-specific integrated circuit (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, as illustrated in [Figure 73](#).

Refer to the specific hardware reference manuals for information about port numbering and connecting the ICL cables.

ICLs for the Brocade DCX Backbone family

The Brocade DCX has two ICL connectors at ports ICL0 and ICL1 on each core blade, each aggregating a set of 16 ports. Thus, each core blade provides 32 ICL ports and there are 64 ICL ports available for the entire Brocade DCX chassis. All the ICL connector ports must be connected to the same two Brocade DCX or DCX-4S chassis.

The Brocade DCX-4S has two ICL connector ports at ICL0 and ICL1, each aggregating a set of 8 ports. Thus, each core blade provides 16 ICL ports and there are 32 ICL ports available for the entire Brocade DCX-4S chassis. All the ICL connector ports must be connected to the same two Brocade DCX or DCX-4S chassis.

Only the following cross-ICL group connections are allowed, as illustrated in [Figure 74](#):

- ICL0 ports on the first chassis connect to ICL1 ports on the second chassis.
- ICL1 ports on the first chassis connect to ICL0 ports on the second chassis.




FIGURE 74 DCX-4S allowed ICL connections

The following ICL connections are not allowed:

- ICL0 ports to ICL0 ports
- ICL1 ports to ICL1 ports

ICL trunking on the Brocade DCX and DCX-4S

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 [Chapter 21, “Administering Licensing”](#). The ICLs are managed the same as ISL trunks.

- On the Brocade DCX, each ICL is managed as two 8-port ISL trunks.
- On the Brocade DCX-4S, each ICL is managed as one 8-port ISL trunk.

Follow the guidelines in the specific hardware reference manuals for connecting the ICL cables.

Virtual Fabrics considerations for ICLs

In Virtual Fabrics, 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.

The following restrictions apply:

- ICL ports cannot be in a logical switch that is using XISLs. The “Allow XISL Use” attribute for the switch must be off.
- 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.

Supported topologies for ICL connections

You can connect the Brocade Backbones in a mesh topology and a core-edge topology. A brief description of each follows. (You can also connect two DCX 8510 chassis 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 in a mesh topology, in which every chassis is connected to every other chassis.

A simple form of the mesh topology is the triangular topology (shown in [Figure 75](#)). The triangular topology is supported by three Brocade Backbone chassis. The chassis for each topology must all be from the same family:

- Brocade DCX Backbone family (DCX or DCX-4S)
- Brocade DCX 8510 Backbone family (DCX 8510-8 or DCX 8510-4)




FIGURE 75 ICL triangular topology with Brocade DCX 8510-8 chassis

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 76](#). This topology is supported by DCX 8510-8 Backbones only. (You can use DCX 8510-4 Backbones for a five-mesh topology.)




FIGURE 76 Full nine-mesh topology

Core-edge topology

You can also connect the Brocade DCX 8510 Backbones in a core-edge topology. For example, [Figure 77](#) shows six chassis connected in a core-edge topology (four edges and two cores).

Although [Figure 77](#) shows only the Brocade DCX 8510-8, each chassis can be either a Brocade DCX 8510-4 or a DCX 8510-8. You can have up to eight edges with DCX 8510-8 cores or up to four edges with DCX 8510-4 cores.

Each line in [Figure 77](#) represents four QSFP cables. The cabling scheme should follow the parallel example shown in [Figure 73](#).




FIGURE 77 64 Gbps ICL core-edge topology

Monitoring Fabric Performance

In this chapter

- [Advanced Performance Monitoring overview](#) 551
- [End-to-end performance monitoring](#) 553
- [Frame monitoring](#) 558
- [Top Talker monitors](#) 562
- [Trunk monitoring](#) 567
- [Saving and restoring monitor configurations](#) 567
- [Performance data collection](#) 568

Advanced Performance Monitoring overview

Advanced Performance Monitoring is a licensed feature that provides a comprehensive tool for monitoring the performance of networked storage resources.

Additional performance monitoring features, such as CRC error reports, are provided through Web Tools and Brocade Network Advisor. Refer to the *Web Tools Administrator's Guide* and *Brocade Network Advisor User Manual* for information about monitoring performance using a graphical interface.

Advanced Performance Monitoring requires either the Brocade Advanced Performance Monitoring license or the Fabric Vision license.

NOTE

Advanced Performance Monitoring features cannot be used if any Flow Vision flow (active or defined) is present on the switch.

Advanced Performance Monitoring commands are available only to users with admin permissions. Use the **perfhel** command to display a list of commands associated with Advanced Performance Monitoring.

NOTE

The command examples in this chapter use the *slot/port* syntax required by Brocade Backbones. For fixed-port switches, use only the port number where needed in the commands.

Types of monitors

Advanced Performance Monitoring provides the following monitors:

- End-to-end monitors (EE monitors) measure the traffic between a host and target pair.

- Frame monitors measure the traffic transmitted through a port with specific values in the first 64 bytes of the frame.
- Top Talker monitors measure the flows that are major consumers of bandwidth on a switch or port.

Restrictions for installing monitors

- Advanced Performance Monitoring is not supported on VE_Ports and EX_Ports. If you issue commands for Advanced Performance Monitoring on VE_Ports or EX_Ports, you will receive error messages.
- All monitor types are allowed only on physical ports.
- Top Talker monitors and EE monitors on E_Ports should be installed only in the ingress direction.

Virtual Fabrics considerations for Advanced Performance Monitoring

In a fabric with Virtual Fabrics enabled, the number of logical switches that can be configured with monitors is restricted. [Table 87](#) lists the platforms that support logical switches and, for each platform, the maximum number of logical switches that can support performance monitors.

TABLE 87 Number of logical switches that support performance monitors

Platform	Maximum number of logical switches supported	Maximum number of logical switches on which monitors are supported
Brocade DCX Brocade DCX-4S Brocade 8510 family	8	4
Brocade 6510	4	4
Brocade 6520	4	4
Brocade 5100 Brocade VA-40FC	3	3
Brocade 5300	4	3

Each logical switch can have its own set of performance monitors. The installation of monitors is restricted to the ports that are present in the respective logical switch.

- Top Talker monitors and EE monitors are supported on the default logical switch, the base switch, and user-defined logical switches.
- Frame monitors are not supported on logical ISLs (LISLs) in user-defined logical switches.

If a port is moved from one logical switch to another, the behavior of monitors installed on that port is as follows:

- **Frame monitor:** Any frame monitors on the port are deleted. To keep the frame monitor, the monitor must be manually installed on the port after the move.
- **Top Talker (fabric mode):** If the fabric mode Top Talkers feature is enabled on the logical switch, a fabric mode Top Talker monitor is automatically installed on the port after it is moved to the logical switch.

- Top Talker (port mode): Any port mode Top Talker monitors on the port are deleted. To keep the port mode Top Talker monitor, the monitor must be manually installed on the port after the move.

Access Gateway considerations for Advanced Performance Monitoring

EE monitors and frame monitors are supported on switches in Access Gateway mode. Top Talker monitors are not supported on these switches.

EE monitors must be installed on F_Ports. Frame monitors can be installed on F_Ports or N_Ports.

Refer to the *Access Gateway Administrator's Guide* for additional information.

End-to-end performance monitoring

Use end-to-end (EE) monitoring when you want to monitor throughput between a pair of devices. End-to-end performance monitoring counts the number of words in Fibre Channel frames for a specified Source ID (SID) and Destination ID (DID) pair.

To enable EE performance monitoring, you must configure an EE monitor on a port, specifying the SID-DID pair (in hexadecimal). The monitor counts only those frames with a matching SID and DID.

Each SID or DID has the following three fields:

- Domain ID (DD)
- Area ID (AA)
- AL_PA

For example, the SID 0x118a0f denotes DD 0x11, AA 0x8a, and AL_PA 0x0f.

An EE monitor includes these counts:

- RX_COUNT - Words in frames received at the port
For frames received at the port with the EE monitor installed, the RX_COUNT is updated if the frame SID is the same as the SID in the monitor and the frame DID is the same as the DID in the monitor.
- TX_COUNT - Words in frames transmitted from the port
For frames transmitted from the port with the EE monitor installed, TX_COUNT is updated if the frame DID is the same as the SID in the monitor and the frame SID is the same as the DID in the monitor.

Maximum number of EE monitors

The maximum number of end-to-end monitors supported varies depending on the switch model:

- The Brocade DCX 8510, 6505, 6510, 6520, M6505, and 6547 models allow up to 512 end-to-end monitors shared by all ports in the same ASIC. Also, these models allow up to 256 end-to-end monitors per port.
- The Brocade DCX, DCX-4S, 5100, VA-40FC, and Brocade Encryption Switch models allow up to 1024 end-to-end monitors shared by all ports in the same ASIC. Also, these models allow up to 256 end-to-end monitors per port.

- The Brocade 300, 5300, 5410, 5424, 5430, 5450, 5460, 5470, 5480, and 7800 models allow up to 768 end-to-end monitors shared by all ports in the same ASIC. Also, these models allow up to 192 end-to-end monitors per port.

The number of interswitch links (ISLs) configured on the switch affects the amount of resources available for end-to-end monitors.

Virtual Fabrics considerations: If Virtual Fabrics is enabled, the Brocade DCX, DCX-4S, DCX 8510, and 5300 models allow up to 256 end-to-end monitors on one logical switch. The Brocade 5100, 6510, 6520, and VA-40FC allow up to 341 end-to-end monitors on one logical switch.

Supported port configurations for EE monitors

You can configure EE monitors on F_Ports and, depending on the switch model, on E_Ports. The following platforms support EE monitors on E_Ports:

- Brocade 6505
- Brocade 6510
- Brocade 6520
- Brocade M6505
- Brocade 6547
- Brocade DCX 8510 family

Identical EE monitors cannot be added to the same port. Two EE monitors are considered identical if they have the same SID and DID values after applying the end-to-end mask.

An EE monitor and a port Top Talker monitor cannot coexist on the same port.

Coexistence of EE monitors and Top Talker monitors on ports belonging to the same ASIC is not recommended because the statistics for the same flow going through ports on the same ASIC may be inaccurate.

Adding EE monitors

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

```
perfaddeemonitor [slotnumber/]portnumber sourceID destID
```

When you add an EE monitor to a port, specify the *sourceID* and *destID* in the ingress direction. For example, [Figure 78](#) shows two devices:

- Host A is connected to domain 1 (0x01), switch area ID 18 (0x12), AL_PA 0x00.
- Dev B is a storage device connected to domain 2 (0x02), switch area ID 30 (0x1e), AL_PA 0x00.




FIGURE 78 Setting end-to-end monitors on a port

End-to-end performance monitoring looks at traffic on SID and DID pairs in any direction. That is, even if the SID is for a remote device, the traffic is monitored in both directions (the Tx and Rx counters are reversed).

Example of monitoring the traffic from Host A to Dev B

On Domain 1, add a monitor to the F_Port, as follows:

```
switch:admin> perfaddeemonitor 2/2 "0x011200" "0x021e00"
```

This monitor (Monitor 1) counts the frames that have an SID of 0x011200 and a DID of 0x021e00. For Monitor 1, RX_COUNT is the number of words from Host A to Dev B, and TX_COUNT is the number of words from Dev B to Host A.

Example of monitoring the traffic from Dev B to Host A

On Domain 2, add a monitor to the F_Port as follows:

```
switch:admin> perfaddeemonitor 2/14 "0x021e00" "0x011200"
```

This monitor (Monitor 4) counts the frames that have an SID of 0x021e00 and a DID of 0x011200. For Monitor 4, RX_COUNT is the number of words from Dev B to Host A, and TX_COUNT is the number of words from Host A to Dev B.

The E_Port monitors are configured similar to the F_Port monitors, but the ingress and egress directions are reversed.

For Monitor 2:

```
switch:admin> perfaddeemonitor 2/3 "0x021e00" "0x011200"
```

For Monitor 3:

```
switch:admin> perfaddeemonitor 2/13 "0x011200" "0x021e00"
```

Setting a mask for an EE monitor

End-to-end monitors count the number of words in Fibre Channel frames that match a specific SID and DID pair. If you want to match only part of the SID or DID, you can set a mask on the port to compare only certain parts of the SID or DID. By default, the frame must match the entire SID and DID to trigger the monitor. By setting a mask, you can choose to have the frame match only one or two of the three fields (domain ID, area ID, and AL_PA) to trigger the monitor.

You specify the masks in the form *dd:aa:pp*, where *dd* is the domain ID mask, *aa* is the area ID mask, and *pp* is the AL_PA mask. The values for *dd*, *aa*, and *pp* are either *ff* (the field must match) or *00* (the field is ignored). The default EE mask value is *ff:ff:ff*.

NOTE

Only one mask per port can be set. When you set a mask, all existing end-to-end monitors are deleted.

ATTENTION

End-to-end masks are supported only on the Brocade Encryption Switch.

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

```
perfsetporteemask [slotnumber/]portnumber "TxSIDMsk" "TxDIDMsk" "RxSIDMsk"
"RxDIDMsk"
```

The **perfSetPortEEMask** command sets the mask for all end-to-end monitors of a port. If any end-to-end monitors are programmed on a port when the **perfSetPortEEMask** command is issued, then a message displays similar to the following example:

```
switch:admin> perfsetporteemask 1/2, "00:ff:ff"
Changing EE mask for this port will cause ALL EE monitors on this port to be
deleted.
Continue? (yes, y, no, n): [no] y
The EE mask on port 2 is set and EE monitors on this port are deleted
```

The **perfSetPortEEMask** command sets a mask for the domain ID, area ID, and AL_PA of the SIDs and DIDs for frames transmitted from and received by the port.

Figure 79 shows the mask positions in the command. A mask (“ff”) is set on slot 1, port 2 to compare the AL_PA fields on the SID and DID in all frames (transmitted and received) on port 2. The frame SID and DID must match only the AL_PA portion of the specified SID and DID pair. Each port can have only one EE mask. The mask is applied to all end-to-end monitors on the port. Individual masks for each monitor on the port cannot be specified.




FIGURE 79 Mask positions for end-to-end monitors

Deleting EE monitors

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **perfMonitorShow** command to list the valid end-to-end monitor numbers for a port.
3. Enter the **perfDelEEMonitor** command to delete a specific monitor.

If you do not specify which monitor number to delete, you are asked if you want to delete all entries.

Example

The following example displays the end-to-end monitors on port 0 (the monitor numbers are listed in the KEY column) and deletes monitor number 2 on port 0:

```
switch:admin> perfmonitorshow --class EE 0

There are 4 end-to-end monitor(s) defined on port 0.

KEY      SID      DID      OWNER_APP      TX_COUNT      RX_COUNT      OWNER_IP_ADDR
-----
0  0x000024  0x000016  WEB_TOOLS      0x0000000000000000  0x0000000000000000  10.106.7.179
1  0x000022  0x000033  WEB_TOOLS      0x0000000000000000  0x0000000000000000  10.106.7.179
2  0x000123  0x000789  WEB_TOOLS      0x0000000000000000  0x0000000000000000  10.106.7.179
3  0x001212  0x003434  WEB_TOOLS      0x0000000000000000  0x0000000000000000  10.106.7.179

switch:admin> perfdeleemonitor 0, 2
End-to-End monitor number 2 deleted
```

Displaying EE monitor counters

You can use this procedure display the end-to-end monitors on a specified port. You can display either the cumulative count of the traffic detected by the monitors or a snapshot of the traffic at specified intervals.

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

```
perfmonitorshow --class monitor_class [slotnumber/]portnumber [interval]
```

Example of displaying an end-to-end monitor on a port at 10-second intervals

```
switch:admin> perfMonitorShow --class EE 4/5 10
Showing EE monitors 4/5 10: Tx/Rx are # of bytes
      0          1          2          3          4
-----
Tx   Rx   Tx   Rx   Tx   Rx   Tx   Rx   Tx   Rx
=====
0    0    0    0    0    0    0    0    0    0
53m  4.9m 53m  4.9m 53m  4.9m 53m  4.9m 53m  0
53m  4.4m 53m  4.4m 53m  4.4m 53m  4.4m 53m  0
53m  4.8m 53m  4.8m 53m  4.8m 53m  4.8m 53m  0
53m  4.6m 53m  4.6m 53m  4.6m 53m  4.6m 53m  0
53m  5.0m 53m  5.0m 53m  5.0m 53m  5.0m 53m  0
53m  4.5m 53m  4.5m 53m  4.5m 53m  4.5m 53m  0
```

Example of displaying EE monitors on a port

```
switch:admin> perfMonitorShow --class EE 4/5
There are 7 end-to-end monitor(s) defined on port 53.
```

KEY	SID	DID	OWNER_APP	TX_COUNT	RX_COUNT	OWNER_IP_ADDR
0	0x58e0f	0x1182ef	TELNET	0x0000000000000000	0x0000000000000000	N/A
0	0x21300	0x21dda	TELNET	0x00000004d0ba9915	0x0000000067229e65	N/A
1	0x21300	0x21ddc	TELNET	0x00000004d0baa754	0x0000000067229e65	N/A
2	0x21300	0x21de0	TELNET	0x00000004d0bab3a5	0x0000000067229e87	N/A
3	0x21300	0x21de1	TELNET	0x00000004d0bac1e4	0x0000000067229e87	N/A
4	0x21300	0x21de2	TELNET	0x00000004d0bad086	0x0000000067229e87	N/A
5	0x11000	0x21fd6	WEB_TOOLS	0x00000004d0bade54	0x0000000067229e87	192.168.169.40
6	0x11000	0x21fe0	WEB_TOOLS	0x00000004d0baed41	0x0000000067229e98	192.168.169.40

Clearing EE monitor counters

The following example clears statistics counters for an end-to-end monitor:

```
switch:admin> perfMonitorClear --class EE 1/2 5
End-to-End monitor number 5 counters are cleared
```

The following example clears statistics counters for all end-to-end monitors on a specific port:

```
switch:admin> perfMonitorClear --class EE 1/2
This will clear ALL EE monitors' counters on port 2, continue?
(yes, y, no, n): [no] y
```

Frame monitoring

Frame monitoring counts the number of times a frame with a particular pattern is transmitted by a port, and generates alerts when thresholds are crossed. Frame monitoring is achieved by defining a filter, or frame type, for a particular purpose. The frame type can be a standard type (for example, an SCSI read command filter that counts the number of SCSI read commands that have been transmitted by the port) or a user-defined frame type customized for your particular use. For a complete list of the standard, predefined frame types, refer to the **fmMonitor** command description in the *Fabric OS Command Reference*.

The maximum number of frame monitors and offsets per port depends on the platform. [Table 88](#) shows the maximum number of frame monitors, in any combination of standard and user-defined frame types, and the maximum number of offsets per port.

TABLE 88 Maximum number of frame monitors and offsets per port

Platform	Maximum number of frame monitors per port	Maximum number of offsets per port
Brocade 300, 5300, 5410, 5424, 5450, 5460, 5470, 5480, and 7800	8	13 ¹
Brocade 5100, 6505, 6510, 6520, M6505, 6547, VA-40FC, DCX, DCX-4S, DCX 8510, and Brocade Encryption Switch	12	25 ²

1. For switches in Access Gateway mode, the maximum number of offsets per port is 7.
2. For switches in Access Gateway mode, the maximum number of offsets per port is 15.

The actual number of frame monitors that can be configured on a port depends on the complexity of the frame types. For trunked ports, the frame monitor is configured on the trunk master.

Static offsets are preset with offset and value combinations. Brocade also supports additional dynamic offsets. When a user-specified offset and value combination matches that already allocated by a Brocade application, resources are shared. The maximum number of configurable offsets by **fmMonitor** may increase or decrease as a result of resource sharing.

Virtual Fabrics considerations: Frame monitors are not supported on logical ISLs (LISLs), but are supported on ISLs and extended ISLs (XISLs).

License requirements for frame monitoring

The Fabric Vision license provides full monitoring capability. If you have this license, you do not need the Advanced Performance Monitoring or the Fabric Watch licenses.

If you do not have the Fabric Vision license, you need the following licenses:

- Advanced Performance Monitoring license is required to use the **fmMonitor** command.
- Fabric Watch license is required to use the monitoring functionality through Fabric Watch.

When you configure actions and alerts through the **fmMonitor** command, Fabric Watch uses these values and generates alerts based on the configuration. If you do not have a Fabric Watch or Fabric Vision license, these values are ignored. Refer to the *Fabric Watch Administrator's Guide* for more information about using Fabric Watch.

Creating frame types to be monitored

In addition to the standard frame types, you can create custom frame types to gather statistics that fit your needs. To define a custom frame type, you must specify a series of offsets, bitmasks, and values. For all transmitted frames, the switch performs the following tasks:

- Locates the byte found in the frame at the specified offset.
- Applies the bitmask to the byte found in the frame.
- Compares the new value with the given value.
- Increments the filter counter if a match is found.

You can specify up to four values to compare against each offset. If more than one offset is required to properly define a filter, the bytes found at each offset must match one of the given values for the filter to increment its counter. If one or more of the given offsets does not match any of the given values, the counter does not increment.

The value of the offset must be between 0 and 63, in decimal format. Byte 0 indicates the first byte of the Start of Frame (SOF), byte 4 is the first byte of the frame header, and byte 28 is the first byte of the payload. Thus, only the SOF, frame header, and first 36 bytes of payload can be selected as part of a filter definition. Offset 0 is a special case, which can be used to monitor the first 4 bytes of the frame (SOF). When the offset is set to 0, the values 0 through 7 that are checked against that offset are predefined, as shown in [Table 89](#).

TABLE 89 Predefined values at offset 0

Value	SOF
0	SOFf
1	SOFc1
2	SOFi1
3	SOFn1
4	SOFi2
5	SOFn2
6	SOFi3
7	SOFn3

Creating a frame monitor

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **fmMonitor --create** command to create a user-defined frame.

Complete details of the **fmMonitor** command parameters are provided in the *Fabric OS Command Reference*. The **highth** and **action** options set values and actions for Fabric Watch, but do not apply to monitoring. To apply the custom values, use the **thconfig --apply** command. Refer to the *Fabric Watch Administrator’s Guide* for more information about using this command.

Example of creating a user-defined frame type

```
switch:admin> fmmonitor --create myframemonitor -pat
"17,0xFF,0x07;7,0x4F,0x01;" -action email
```

Example of creating a user-defined frame type and applying frame monitors to ports 3, 4, and 5

```
switch:admin> fmmonitor --create myframemonitor -pat
"17,0xFF,0x007;7,0x4F,0x01;" -port 3-5
```

Deleting frame types

Deleting a frame type removes the entire configuration, including configured thresholds and associated actions. It also removes any frame monitors of the specified type from all ports.

You can delete only user-defined frame types; you cannot delete the predefined frame types.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **fmMonitor --delete** command to delete a specific frame type.

Example

```
switch:admin> fmmonitor --delete myframemonitor
```

Adding frame monitors to a port

If the switch does not have enough resources to add a frame monitor to a port, then other frame monitors on that port may have to be deleted to free resources.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **fmMonitor --addmonitor** command to add a frame monitor to one or more ports.

The set of ports to be removed from monitoring is automatically saved to the persistent configuration unless you specify the **-nosave** option on the command.

Example

The following example adds a standard SCSI frame type monitor to ports 3 through 12.

```
switch:admin> fmmonitor --addmonitor SCSI -port 3-12
```

Removing frame monitors from a port

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **fmMonitor --delmonitor** command to remove a specific monitor from one or more ports.

The set of ports to be removed from monitoring is automatically saved to the persistent configuration unless you specify the **-nosave** option on the command.

Example

The following example removes the user-defined frame monitor, myframemonitor, from all ports.

```
switch:admin> fmmonitor --delmonitor myframemonitor
```

Saving a frame monitor configuration

When you assign or remove frame monitors on ports, the list of ports to be monitored is automatically saved persistently, unless you specified the **-nosave** option.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **fmMonitor --save** command to save the set of ports on which the frame type is monitored to the persistent configuration.

Example

In the following example, the first command adds a standard SCSI frame type monitor to ports 3 through 12, but does not save the port configuration. The second command saves the port configuration persistently.

```
switch:admin> fmmonitor --addmonitor SCSI -port 3-12 -nosave
switch:admin> fmmonitor --save SCSI
```

Displaying frame monitors

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

Example

The following example displays the existing frame types and associated bit patterns on the switch.

```
switch:admin> fmmonitor --show
FRAME_TYPE      BIT  PATTERN
-----
scsi      12,0xFF,0x08;
          scsiread  12,0xFF,0x08;4,0xFF,0x06;40,0xFF,0x08,0x28;
          scsiwrite 12,0xFF,0x08;4,0xFF,0x06;40,0xFF,0x08,0x28,0x0A,0x2A;
          scsirw   12,0xFF,0x08;4,0xFF,0x06;40,0xFF,0x08,0x28,0x0A,0x2A;
          scsi2reserve 12,0xFF,0x08;4,0xFF,0x06;40,0xFF,0x16,0x56;
          scsi3reserve 12,0xFF,0x08;4,0xFF,0x06;40,0xFF,0x5F;41,0xFF,0x01
          ip      12,0xFF,0x05;
          abts    4,0xFF,0x81;40,0xFF,0x81;12,0xFF,0x0;17,0xFF,0x0;
          baacc   4,0xff,0x84;12,0xff,0x00;17,0xff,00;
```

The following example displays configuration details for the predefined SCSI frame monitor. Notice that in the last entry, the “-” in the Count column indicates that the monitor is configured, but is not installed on the port.

```
switch:admin> fmmonitor --show SCSI
```

Port	Frame Type	Count	HIGH Thres	Actions	TIMEBASE	CFG
000001	scsi	0x0000000000000123	1000	Email	None	saved
000002	scsi	0x0000000000000125	1000	Email	None	saved
000003	scsi	0x0000000000000143	1000	Email	None	saved
000022	scsi	-	0	None	None	saved

The following example displays values for the predefined SCSI frame monitor on port 5 every 5 seconds.

```
switch:admin> fmmonitor --show scsi -port 5 -timeinterval 5
Port|Count
-----
2011-03-21 00:59:50

000005| 48.3k
```

```
2011-03-21 00:59:55
```

```
000005| 48.6k
```

```
(output truncated)
```

Clearing frame monitor counters

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **fmMonitor --clear** command to clear the counters on the ports on which the specified frame type is monitored.

Example

The following example clears the counters for the ABTS monitor from ports 7 through 10.

```
switch:admin> fmmonitor --clear ABTS -port 7-10
```

Top Talker monitors

Top Talker monitors determine the flows (SID and DID pairs) that are the major users of bandwidth (after initial stabilization). Top Talker monitors measure bandwidth usage data in real time and relative to the port on which the monitor is installed.

NOTE

Initial stabilization is the time taken by a flow to reach the maximum bandwidth. This time varies depending on the number of flows in the fabric and other factors. This time can be up to 14 seconds in the Backbones, and up to 82 seconds in the fixed-port switches.

Applications can use Top Talker monitors data to do the following:

- Re-route the traffic through different ports that are less busy, so as not to overload a given port.
- Alert you to the major users of bandwidth (top talking flows) on a port if the total traffic on the port exceeds the acceptable bandwidth consumption.

You can use Top Talker monitors to identify the SID and DID pairs that consume the most bandwidth and can then configure them with certain Quality of Service (QoS) attributes so they get proper priority. Refer to [Chapter 14, “Optimizing Fabric Behavior,”](#) for information on QoS.

The Top Talker monitors are based on SID and DID pairs and not WWNs. Once Top Talker monitors are installed on a switch or port, they remain installed across power cycles.

Top Talker monitors support two modes, port mode and fabric mode:

- Port mode Top Talker monitor

A Top Talker monitor can be installed on a port to measure the traffic originating from the port and flowing to different destinations.

You can configure Top Talker monitors on F_Ports and, depending on the switch model, on E_Ports. The following platforms support Top Talker monitors on E_Ports:

- Brocade 6505
- Brocade 6510
- Brocade 6520
- Brocade M6505
- Brocade 6547
- Brocade DCX 8510 family
- Fabric mode Top Talker monitor

In fabric mode, Top Talker monitors are installed on all E_Ports in the fabric and measure the data rate of all the possible flows in the fabric (ingress E_Port traffic only). In fabric mode, Top Talker monitors can determine the top n bandwidth users on a given switch.

You can install Top Talker monitors in either port mode or fabric mode, but not both.

ATTENTION

A fabric mode Top Talker monitor and an EE monitor cannot be configured on the same fabric. You must delete the EE monitor before you configure the fabric mode Top Talker monitor.

How do Top Talker monitors differ from EE monitors? EE monitors provide counter statistics for traffic flowing between a given SID and DID pair. Top Talker monitors identify all possible SID and DID flow combinations that are possible on a given port and provide a sorted output of the top talking flows. Also, if the number of flows exceeds the hardware resources, existing EE monitors fail to get real-time data for all of them; however, Top Talker monitors can monitor all flows for a given E_Port or F_Port.

Virtual Fabric considerations: All logical switches in the same chassis can use either fabric mode Top Talker monitors or port mode Top Talker monitors and EE monitors. You cannot use fabric mode Top Talker monitors and EE monitors together on the same logical switch.

Admin Domain considerations: Top Talker monitors are always installed in AD255.

NPIV considerations: Top Talker monitors take NPIV devices into consideration when calculating the top talking flows.

Top Talker monitors are not supported on the embedded platforms: Brocade 5410, 5424, 5450, 5460, 5470, and 5480.

Top Talker monitors and FC-FC routing

You can enable Top Talker monitors on a platform that is configured to be an FC router. Top Talker monitors and FC routers are concurrently supported on the following platforms:

- Brocade 6505
- Brocade 6510
- Brocade 6520
- Brocade DCX 8510 Backbone family, with the following blades only: FC16-32, FC16-48

On all other platforms, you can have either Top Talker monitors or FC-FC routing, but not both.

Top Talker monitors are supported on an FC router in both backbone-to-edge and edge-to-edge configurations.

Note the following restrictions:

- An E_Port-attached switch must be connected and merged with the backbone FC router before you can enable Top Talker monitors on the FC router.
- Fabric mode Top Talker monitors do not support requests for domains (either front port domain or xlate domain).
- Fabric mode Top Talker monitors do not monitor flows over EX_Ports.

For example, if a host is connected directly to an FC router and the target is on the edge switch (refer to Figure 80 on page 564), no flows are monitored because none of the flows traverse an E_Port on the FC router.

In Figure 81 on page 564, however, the flows across the E_Port on the FC router are monitored.




FIGURE 80 Fabric mode Top Talker monitors on FC router do not monitor any flows




FIGURE 81 Fabric mode Top Talker monitors on FC router monitor flows over the E_Port

Limitations of Top Talker monitors

Be aware of the following when using Top Talker monitors:

- Top Talker monitors cannot detect transient surges in traffic through a given flow.
- You cannot install a Top Talker monitor on a mirrored port.
- Top Talker monitors can monitor only 10,000 flows at a time.
- Top Talker monitors are not supported on VE_Ports, EX_Ports, and VEX_Ports.
- The maximum number of all port mode Top Talker monitors on an ASIC is 16. If Virtual Fabrics is enabled, the maximum number of all port mode Top Talker monitors on an ASIC is 8.
- If the ingress and egress monitor ports are configured on the same ASIC, port mode Top Talker monitors on F_Ports show the flow from only one of the ports, either the ingress or the egress port, but not both.

Adding a Top Talker monitor to a port (port mode)

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

```
perfttmon --add [egress | ingress] [slotnumber/]port
```

The following example monitors the incoming traffic on port 7.

```
perfttmon --add ingress 7
```

The following example monitors the outgoing traffic on slot 2, port 4 on a Backbone.

```
perfttmon --add egress 2/4
```

Adding Top Talker monitors on all switches in the fabric (fabric mode)

When fabric mode is enabled, you can no longer install Top Talker monitors on an F_Port unless you disable fabric mode.

1. Connect to the switch and log in using an account with admin permissions.
2. Remove any EE monitors in the fabric, as described in [“Deleting EE monitors”](#) on page 556. Fabric mode Top Talker monitors and EE monitors cannot both exist in the fabric.
3. Enter the **perfttmon --add fabricmode** command.

```
perfttmon --add fabricmode
```

The system responds with the following message:

```
Before enabling fabric mode, please remove all EE monitors in the fabric
continue? (yes, y, no, n):
```

4. Enter **y** at the prompt to continue.

Top Talker monitors are added to E_Ports in the fabric and fabric mode is enabled. Any Top Talker monitors that were already installed on F_Ports are automatically uninstalled.

If EE monitors are present on the local switch, the command fails with the following message:

```
Cannot install Fabric Mode Top Talker because EE monitor is already present
```

If EE monitors are present on remote switches, the command succeeds; however, on the remote switches, fabric mode fails and a RASlog message is displayed on those switches.

If a new switch joins the fabric, you must run the **perfttmon --add fabricmode** command on that switch. The Top Talker monitor configuration information is *not* automatically propagated to the new switch.

Displaying the top *n* bandwidth-using flows on a port (port mode)

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

```
perfttmon --show [slotnumber/]port [n] [wwn | pid]
```

The output is sorted based on the data rate of each flow. If you do not specify the number of flows to display, then the command displays the top 8 flows or the total number of flows, whichever is less.

The following example displays the top 5 flows on port 7 in WWN (default) format:

```
perfttmon --show 7 5
```

The following example displays the top flows on slot 2, port 4 on a Backbone in PID format:

```
switch:admin> perfttmon --show 2/4 pid
```

```
=====
Src_PID      Dst_PID      MB/sec
=====
0xa90800     0xa05200     6.926
0xa90800     0xa908ef     6.872
```

Displaying top talking flows for a given domain ID (fabric mode)

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

```
perfttmon --show dom domainid [n] [wwn | pid]
```

Fabric mode must be enabled for this option.

The output is sorted based on the data rate of each flow. If you do not specify the number of flows to display, then the command displays the top 8 flows or the total number of flows, whichever is less. The command can display a maximum of 32 flows.

The following example displays the top 5 flows on domain 1 in WWN (default) format:

```
perfttmon --show dom 1 5
```

The following example displays the top flows on domain 2 in PID format:

```
switch:admin> perfttmon --show dom 2 pid
```

```
=====
Src_PID      Dst_PID      MB/sec      Potential E-Ports
=====
0x03f600     0x011300     121.748     2/0,2/2,2/3
0x03f600     0x011300     121.748     3/14,3/15
```

Deleting a Top Talker monitor on a port (port mode)

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

```
perfttmon --delete [slotnumber/]port
```

The following example deletes the monitor on port 7:

```
perfttmon --delete 7
```

The following example deletes the monitor on slot 2, port 4 on a Backbone:

```
perfttmon --delete 2/4
```

Deleting all fabric mode Top Talker monitors

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

```
perfttmon --delete fabricmode
```

All Top Talker monitors are deleted.

Trunk monitoring

To monitor E_Port (ISL) and F_Port trunks, you can set monitors only on the master port of the trunk. If the master changes, the monitor automatically moves to the new master port.

If a monitor is installed on a port that later becomes a slave port when a trunk comes up, the monitor automatically moves to the master port of the trunk.

Trunk monitoring considerations

- End-to-end (EE) monitors are supported for ISLs only on the Brocade 6505, 6510, 6520, M6505, 6547, and DCX 8510 family.
- If an EE monitor is installed on a trunk group and you disable the trunk, the EE monitor will be installed only on the last master port of that trunk group, which may not be the actual port on which the EE monitor was installed when the trunk was enabled.
- For F_Port trunks, end-to-end masks are allowed only on the F_Port trunk master. Unlike the monitors, if the master changes, the mask does not automatically move to the new master port.
- All platforms support 12 frame monitors for trunks, except for the Brocade 300, which supports 8 frame monitors for trunks.

Saving and restoring monitor configurations

To prevent the switch configuration flash from running out of memory, the number of monitors saved to flash memory is limited as follows:

- The total number of EE monitors per port is limited to 16.

- The total number of frame monitors per port is limited to 16.
- The total number of monitors per switch is limited to 512.

When there are more than 512 monitors in the system, monitors are saved to flash memory in the following order:

- The EE monitors for each port (from 0 to MAX_PORT)
- The frame monitors for each port

EE monitors get preference saving to flash memory when the total number of monitors in a switch exceeds 512. If the total number of monitors per port or switch exceeds the limit, then you will receive an error message indicating the count has been exceeded and that some monitors have been discarded.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter one of the following commands, depending on the action you want to perform:
 - To save the current EE monitor and frame monitor configuration settings into nonvolatile memory, use the **perfCfgSave** command.

```
switch:admin> perfCfgsave
This will overwrite previously saved Performance Monitoring
settings in FLASH. Do you want to continue? (yes, y, no, n): [no] y
Please wait ...
Performance monitoring configuration saved in FLASH.
```

- To restore a saved monitor configuration, use the **perfCfgRestore** command, for example, to restore the original performance monitor configuration after making several changes.

```
switch:admin> perfCfgrestore
This will overwrite current Performance Monitoring settings in RAM. Do you
want to continue? (yes, y, no, n): [no] y
Please wait... Performance monitoring configuration restored from FLASH
ROM.
```

- To clear the previously saved performance monitoring configuration settings from nonvolatile memory, use the **perfCfgClear** command.

```
switch:admin> perfCfgclear
This will clear Performance Monitoring settings in FLASH. The RAM settings
won't change. Do you want to continue? (yes, y, no, n): [no] y
Please wait... Committing configuration...done.
Performance Monitoring configuration cleared from FLASH.
```

Performance data collection

Data collected through Advanced Performance Monitoring is deleted when the switch is rebooted. Using the Brocade Network Advisor Enterprise Edition, you can store performance data persistently. For details on this feature, refer to the *Brocade Network Advisor User Manual*.

Managing Trunking Connections

In this chapter

- Trunking overview 569
- Supported platforms for trunking 571
- Supported configurations for trunking 571
- Requirements for trunk groups 572
- Recommendations for trunk groups 572
- Configuring trunk groups 573
- Enabling trunking 574
- Disabling trunking 574
- Displaying trunking information 574
- Trunk Area and Admin Domains 576
- ISL trunking over long-distance fabrics 576
- EX_Port trunking 577
- F_Port trunking 579
- Displaying F_Port trunking information 585
- Disabling F_Port trunking 585
- Enabling the DCC policy on a trunk area 586

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 [Chapter 21, “Administering Licensing”](#).

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 or DCX 8510 Backbones and is applicable only to ports on the core blades.
Refer to [Chapter 22, “Inter-chassis Links”](#) for detailed information about ICL trunking.
- EX_Port trunking is configured on an inter-fabric link (IFL) between an FC router (EX_Port) and an edge fabric (E_Port). The trunk ports are EX_Ports connected to E_Ports.
Refer to [“EX_Port frame trunking configuration”](#) on page 619 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 581, the *Access Gateway Administrator’s 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 *Access Gateway Administrator’s 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 or DCX 8510 Backbones 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 [Chapter 21, “Administering Licensing,”](#) 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.

[Figure 82](#) shows the port groups for the Brocade 5100.

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 82 Port group configuration for the Brocade 5100

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”](#) on page 594 for more information.

Supported configurations for trunking

- Trunk links can be 2 Gbps, 4 Gbps, 8 Gbps, 10 Gbps, or 16 Gbps, depending on the Brocade platform.
- The maximum number of ports per trunk and 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).

- If in-flight encryption or compression is enabled, you can have a maximum of only two ports per trunk.
- 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. For more information on QoS, refer to “QoS zones” on page 419.

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 encryption, compression, QoS, and FEC settings.
- Trunk groups must be between Brocade switches (or Brocade adapters in the case of F_Port trunking). Brocade trunking is proprietary and is not supported on M-EOS or third-party switches.
- 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 **portCfgIsiMode** 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. Consider these points:
 - 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 **isIShow** command to determine which ports are used for ISLs.

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 580 or [“Configuring F_Port trunking for a Brocade adapter”](#) on page 581 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** *port mode* command to enable trunking on a port.
Enter the **switchCfgTrunk** *mode* command to enable trunking on all ports on the switch.

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** *port mode* command to disable trunking on a port.
Enter the **switchCfgTrunk** *mode* command to disable trunking on all ports on the switch.

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 using the Brocade Advanced Performance Monitor to monitor traffic, refer to [Chapter 23, “Monitoring Fabric Performance”](#).

To view detailed information about F_Port trunking, refer to [“Displaying F_Port trunking information”](#) on page 585.

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
```

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%)
```

Trunk Area and Admin Domains

Ports from different Admin Domains (ADs) are not allowed to join the same Trunk Area (TA) group. The `portTrunkArea` command prevents the different ADs from joining the TA group.

When you assign a TA, the ports within the TA group have the same index. The index that was assigned to the ports is no longer part of the switch. Any domain,index (D,I) AD that was assumed to be part of the domain may no longer exist for that domain because it was removed from the switch.

Example of Trunk Area assignment on port domain,index

If you have AD1: 3,8; 3,9; 4,13; 4,14 and AD2: 3,10; 3,11, and then create a TA with index 8 with ports that have index 8, 9, 10, and 11, then index 9, 10, and 11 are no longer with domain 3. This means that AD2 does not have access to any ports because index 10 and 11 no longer exist on domain 3. This also means that AD1 no longer has 3,9 in effect, because index 9 no longer exists for domain 3. Port 3,8, which is the TA group, can still be seen by AD1 along with 4,13 and 4,14.

If a port within a TA is removed, the index is added back to the switch. For example, the same AD1 and AD2 with TA 8 holds true. If you remove port 9 from the TA, it adds index 9 back to the switch. That means port 3,9 can be seen by AD1 along with 3,8, 4,13, and 4,14.

ISL trunking over long-distance fabrics

In long-distance fabrics, if a port speed is set to autonegotiate, then the maximum speed, which is 16 Gbps, 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 571, observe the following criteria for trunking over long-distance fabrics:

- Trunking over long-distance fabrics is supported only on switches running Fabric OS v6.1.0 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 589.

[Table 90](#) summarizes support for trunking over long distance for the Brocade DCX and DCX 8510 Backbones and supported blades.

TABLE 90 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)
LO	Normal	See note below	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

TABLE 90 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
LD	500 km	0	0
LS	Static	See note below	

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 Gbps. The distance for the LS mode is static. You can specify any distance greater than 10 km.

The distance supported depends on the available buffers, the number of back-end ports, and the number of ports that are offline. For more information on setting port speeds, refer to [Chapter 3, “Performing Advanced Configuration Tasks”](#).

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 [Chapter 26, “Using FC-FC Routing to Connect Fabrics,”](#) 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 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 Fiber 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 a maximum speed of 16 Gbps 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”](#) 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 573 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 10:00:00:60:69:80:1d:bc "MtOlympus_72"
(fabric id = 2 )(Trunk master)

```

F_Port trunking

You can configure F_Port trunking in the following scenarios:

- Between F_Ports on a Fabric OS switch and N_Ports on an Access Gateway module
- Between F_Ports on a Fabric OS switch and N_Ports on a Brocade adapter

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.

Refer to the *Access Gateway Administrator's Guide* and the *Brocade Adapters Administrator's Guide* for information about configuring the corresponding N_Port trunking on the Access Gateway and the Brocade adapter.

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 83](#) shows a switch in AG mode without F_Port masterless trunking. [Figure 84](#) shows a switch in AG mode with F_Port masterless trunking.




FIGURE 83 Switch in Access Gateway mode without F_Port masterless trunking




FIGURE 84 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 580 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 572, refer to the *Access Gateway Administrator’s 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 572, 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*.
3. On the switch side, enable the ports by using the **portEnable** command.

```
switch:admin> portenable 3/40
switch:admin> portenable 3/41
```

F_Port trunking considerations

Table 91 describes the F_Port masterless trunking considerations.

TABLE 91 F_Port masterless trunking considerations

Category	Description
AD	You cannot create a Trunk Area on ports with different Admin Domains. You cannot create a Trunk Area in AD255.
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 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 mode do not perform authentication.</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 (D,I)</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 AD, 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 TA. The PWWN of the FLOGI sent from the AG will be dynamic for the F_Port trunk master. Because you do not know ahead of time what PWWN the AG 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>

TABLE 91 F_Port masterless trunking considerations (Continued)

Category	Description
Default Area	Port X is a port that has its Default Area the same as its Trunk Area. The only time you can remove port X from the trunk group is when the entire trunk group has the Trunk Area disabled.
Downgrade	You can have trunking on, but you must disable the trunk ports before performing a firmware downgrade. 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 v6.2.0.
FastWrite	When you assign a Trunk Area to a trunk group, the trunk group cannot have FastWrite enabled on those ports. If a port is FastWrite-enabled, the port cannot be assigned a Trunk Area.
FICON	FICON is not supported on F_Port trunk ports. However, FICON can still run on ports that are not F_Port trunked within the same switch.
HA Sync	If you plug in a standby CP with a firmware version earlier than Fabric OS v6.2.0 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 F_Port trunk that has 0x2f or 0x25 as the first byte of the PWWN. The TA is part of the PWWN in the format listed in Table 92 on page 584.
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 v7.0.0 and later if the F_Port is present on the switch. Upgrading is not disruptive.

[Table 92](#) describes the PWWN format for F_Port and N_Port trunk ports.

TABLE 92 PWWN format for F_Port and N_Port trunk ports

NAA = 2	2f:xx:nn:nn:nn:nn:nn	Port WWNs for: switch's Fx_Ports.	The valid range of xx is [0-FF], for a maximum of 256.
NAA = 2	25:xx:nn:nn:nn:nn:nn	Port WWNs for: switch's FX_Ports	The valid range of xx is [0-FF], for a maximum of 256.

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 and DCX 8510-8. Fabric OS uses a 10-bit addressing model, which is the default mode for all dynamically created logical switches in the DCX platform.

In the DCX and DCX 8510 platforms, 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

Because the DCX and 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.
- The Port WWN format in a Virtual Fabric is 2z:zz:xx:xx:xx:xx:xx:xx. The z:zz is the logical port number; for example, the logical port 450 will be 1:c2. The xx:xx:xx:xx:xx:xx is based on the logical fabric WWN.
For example, if the logical fabric WWN is 10:00:00:05:1e:39:fa:f3, and logical port number is 450, then the Port WWN of the F_Port trunk will be 21:c2: 00:05:1e:39:fa:f3.
- 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 [Chapter 11, “Managing Virtual Fabrics,”](#) for detailed information about Virtual Fabrics.

Displaying F_Port trunking information

Use the following commands on the edge switch to verify the F_Port trunking configuration.

- Enter the **switchShow** command to display the switch and port information.
- Enter the **portTrunkArea --show enabled** command to display the TA-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 the **portTrunkArea --show trunk** command 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 [Chapter 8, “Configuring Security Policies”](#).

Managing Long-Distance Fabrics

In this chapter

- [Long-distance fabrics overview](#) 587
- [Extended Fabrics device limitations](#) 588
- [Long-distance link modes](#) 588
- [Configuring an extended ISL](#) 589
- [Forward error correction on long-distance links](#) 591

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 [Chapter 21, “Administering Licensing”](#).

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 port blade for long distance because of its limited buffers. This blade does not support long-wavelength (LWL) fiber optics and supports 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; therefore, 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
- 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 [Chapter 24, “Managing Trunking Connections”](#).

- Only qualified Brocade SFP transceivers are used. Only Brocade-branded or certain Brocade-qualified SFP transceivers are supported.
1. Connect to the switch and log in using an account assigned to the admin role.
 2. Enter the **switchDisable** command.
 3. Enter the **configure** command to set the switch fabric-wide configurations.

[Table 93](#) shows the fabric-wide settings that can be set:

TABLE 93 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 *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]
```

6. 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.

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

```
switch:admin> portcfgqos --disable [slot/]port
```

If you do not disable QoS, after the second or third link reset (LR), ARB fill words display. Refer to the *Fabric OS Command Reference* for more for more information on setting fill words.

3. Disable buffer credit recovery; buffer credit recovery is not compatible with the IDLE mode. If you do not disable buffer credit recovery, it continues to perform a link reset.

```
switch:admin> portcfgcreditrecovery --disable [slot/]port
```

4. Configure the port to support long-distance links.

```
switch:admin> portcfglongdistance [slot/]port,LS,0,-distance 100
```

Forward error correction on long-distance links

Forward error correction (FEC) on user ports is supported for LD and LS long-distance modes. Use the **portCfgLongDistance** command with the **-fecEnable** or **-fecDisable** options to enable or disable FEC, respectively, on a user port. Alternatively, you can use the **portCfgFec** command with the **--enable** or **--disable** option as you would for any regular port.

For additional details about FEC, refer to [“Enabling forward error correction”](#) on page 111.

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.

Example

```
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
Forward Error Correction capable:    YES
Forward Error Correction configured:  ON
```

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.

Example

```
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
Forward Error Correction capable:    YES
Forward Error Correction configured:  OFF
```

25 Forward error correction on long-distance links

Using FC-FC Routing to Connect Fabrics

In this chapter

- [FC-FC routing overview](#) 593
- [Fibre Channel routing concepts](#) 596
- [Setting up FC-FC routing](#) 603
- [Backbone fabric IDs](#) 605
- [FCIP tunnel configuration](#) 606
- [Inter-fabric link configuration](#) 607
- [FC router port cost configuration](#) 613
- [Shortest IFL cost configuration](#) 615
- [EX_Port frame trunking configuration](#) 619
- [LSAN zone configuration](#) 620
- [Proxy PID configuration](#) 633
- [Fabric parameter considerations](#) 633
- [Inter-fabric broadcast frames](#) 634
- [Resource monitoring](#) 634
- [FC-FC routing and Virtual Fabrics](#) 636
- [Upgrade and downgrade considerations for FC-FC routing](#) 639
- [Displaying the range of output ports connected to xlate domains](#) 639

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 [Chapter 4, “Routing Traffic”](#).
- *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
- Fabric OS and Brocade Network OS
- Fabric OS and M-EOS

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 415 for information about QoS and instructions for setting traffic prioritization over an FC router.

ATTENTION

FC-FC routing is not supported on a Brocade 7800 that has been enabled for logical switches.

License requirements for FC-FC routing

The Integrated Routing license is required for FC-FC routing between Fabric OS fabrics and between Fabric OS and M-EOS fabrics.

NOTE

The Integrated Routing license is *not* required for connectivity between Fabric OS and Brocade Network OS fabrics or between Brocade Network OS fabrics connected by an FC router.

The Integrated Routing license allows 8-Gbps and 16-Gbps FC ports to be configured as EX_Ports (or VEX_Ports) supporting FC-FC routing.

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. It is not recommended to mix licensed and unlicensed FC routers in the backbone fabric.

Supported platforms for FC-FC routing

FC-FC routing is supported on the following platforms:

- Brocade 5100 switch
- Brocade 5300 switch
- Brocade 6510 switch
- Brocade 6520 switch
- Brocade VA-40FC switch
- Brocade 7800 Extension Switch
- Brocade Encryption Switch
- Brocade DCX and DCX 8510 Backbone families:
 - 8-Gbps port blades (FC8-16, FC8-32, FC8-32E, FC8-48, FC8-48E, FC8-64)
 - 16-Gbps port blades (FC16-32, FC16-48)
 - FX8-24 DCX Extension Blade

- ICL ports on the core blades. EX-Port 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, FC8-32E, FC8-48E).

NOTE

Device discovery will not occur properly with ICL EX-Port connected edge fabrics if the FC Router has unsupported blades which are not mentioned above.

For the Brocade Backbone families, the backbones have a limit of 128 EX_Ports for each chassis. Refer to the *Network OS Administrator's Guide* for supported Network OS platforms.

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 connected to a Brocade Network OS edge fabric (Network OS v2.1.1 or later).
- FC router connected to an M-EOS Open Mode edge fabric.
- FC router connected to an M-EOS Fabric Mode edge fabric.
- FC router interoperating with legacy FC routers (Brocade 7500 switch).

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.

In an edge fabric that contains a mix of administrative domain (AD)-capable switches and switches that are not aware of AD, the FC router must be connected directly to an AD-capable switch. For more information, refer to [“Use of Admin Domains with LSAN zones and FC-FC routing”](#) on page 620.

VEX edge to VEX edge device sharing is not supported.

Network OS connectivity limitations

You should be aware of the following configuration limitations for Network OS connectivity:

- All of the platforms listed in [“Supported platforms for FC-FC routing”](#) on page 594 support FC-FC routing to a Brocade Network OS fabric, except for the Brocade Encryption Switch.
- VEX_Ports do not support Network OS connectivity.
- FCoE devices connected to an FCOE10-24 blade cannot communicate with FCoE devices in the Network OS fabric.
- If Admin Domains are enabled, connectivity between the Fabric OS fabric and a Brocade Network OS fabric is not supported.
- Connectivity between a Brocade Network OS fabric and an M-EOS fabric is not supported.

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](#)” on page 594 for a list of platforms that can be FC routers.
- 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 *Fibre Channel over IP Administrator’s 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 by using an EX_Port or VEX_Port.
- 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 87](#) on page 599).
- 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 85](#) shows a metaSAN consisting of three edge fabrics connected through a Brocade DCX with inter-fabric links.




FIGURE 85 A metaSAN with inter-fabric links

- 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 86 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 86 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 599 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.

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.

NOTE

Backbone fabrics that share connections to the same edge fabrics must have unique backbone fabric IDs.

- 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 87](#) 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 87 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 601.

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. [Figure 88](#) illustrates this concept.




FIGURE 88 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 FC routers connect to a common 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; rather, 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 620 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.

Figure 89 shows a sample physical topology. This figure shows four FC routers in a backbone fabric and four edge fabrics connected to the FC routers.




FIGURE 89 Sample topology (physical topology)

Figure 90 shows a phantom topology for the physical topology shown in Figure 89. In this figure, the dashed lines and shapes represent the phantom topology from the perspective of Fabric 1. Fabrics 2 and 3 also see 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 90 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.”](#)

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.

You can use the `fcrXlateConfig` command to identify and remove these stale xlate domains without disrupting the fabric.

1. Connect to the FC router and log in using an account with admin permissions.
2. Enter the `fcrlateconfig --show` command to identify any stale xlate domains.
3. Enter the `fcrlateconfig --del` command to delete the stale xlate domains.

Example

```
sw0:root> fcrlateconfig --show stalexd
Imported FID      Stale XD      Owner Domain
-----
      012          002          007 ( this FCR )
sw0:root> fcrlateconfig --del stalexd 12 2
Xlate domain 2 is deleted
```

FC router authentication

A Brocade FC router (FCR) 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 port` 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 249 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 604.)
2. Assign backbone fabric IDs. (Refer to [“Backbone fabric IDs”](#) on page 605.)
3. Configure FCIP tunnels if you are connecting Fibre Channel SANs over IP-based networks. (Refer to [“FCIP tunnel configuration”](#) on page 606.)
4. Configure IFLs for edge and backbone fabric connection. (Refer to [“Inter-fabric link configuration”](#) on page 607.)
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 613.)
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 615.)
7. Configure trunking on EX_Ports that are connected to the same edge fabric. (Refer to [“EX_Port frame trunking configuration”](#) on page 619.)

- Configure LSAN zones to enable communication between devices in different fabrics. (Refer to [“LSAN zone configuration”](#) on page 620.)

Refer to [Chapter 3, “Performing Advanced Configuration Tasks,”](#) 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 the following verification checks on the FC router.

- Log in to the switch or backbone as admin and enter the **version** command. Verify that Fabric OS v7.2.0 or later is installed on the FC router, as shown in the following example.

```
switch:admin> version
Kernel:      2.6.14.2
Fabric OS:   v7.2.0
Made on:     Fri Nov 18 01:15:34 2011
Flash:       Mon Nov 21 20:53:48 2011
BootProm:   1.0.9
```

- If you are configuring a backbone, enter the **slotShow** command to verify that an FX8-24 blade is present or an 8-Gbps or 16-Gbps port blade is present. The following example shows slots 1, 2, 3, 4, 9, 10, and 12 with 8-Gbps port blades enabled.

```
switch:admin> slotshow -m
```

Slot	Blade Type	ID	Model Name	Status
1	SW BLADE	37	FC8-16	ENABLED
2	SW BLADE	37	FC8-16	ENABLED
3	SW BLADE	37	FC8-16	ENABLED
4	SW BLADE	39	FC8-16	ENABLED
5	CORE BLADE	52	CORE8	ENABLED
6	CP BLADE	50	CP8	ENABLED
7	CP BLADE	50	CP8	ENABLED
8	CORE BLADE	52	CORE8	ENABLED
9	SW BLADE	37	FC8-16	ENABLED
10	SW BLADE	55	FC8-32	ENABLED
11	UNKNOWN			VACANT
12	SW BLADE	51	FC8-48	ENABLED

Refer to [Chapter 3, “Performing Advanced Configuration Tasks,”](#) for a list of blades and their corresponding IDs.

- Enter the **licenseShow** command to verify that the Integrated Routing license is installed.

```
switch:admin> licenseshow
S9bddb9SQbTAceeC:
    Fabric license
bzbzRcbcSc0c0SY:
    Remote Fabric license
RyeSzRScycazfT0G:
    Integrated Routing license
```

If you are connecting to a Fabric OS or M-EOS fabric and the Integrated Routing license is not installed, you must install it, as described in [Chapter 21, “Administering Licensing”](#). The Integrated Routing license is not required if you are connecting to a Brocade Network OS fabric. For configuring EX_Ports on an ICL, both the Integrated Routing license and the ICL POD license are required.

- Verify that the Fabric-Wide Consistency Policy is not in “strict” mode by issuing the **fddCfg --showall** command. 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
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 policy must be tolerant before you can connect fabrics to the FC router. Refer to [Chapter 8, “Configuring Security Policies”](#) for information about configuring the Fabric-Wide Consistency Policy.

- For 8-Gbps platforms, delete fabric mode Top Talker monitors, if they are configured. Refer to [“Deleting all fabric mode Top Talker monitors”](#) on page 567 for instructions.

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. Rather, 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.

Example

```
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
```

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.

NOTE

FCIP tunnel configuration is applicable only to Fabric OS fabrics and does not apply to Brocade Network OS or M-EOS fabrics.

Refer to the *Fibre Channel over IP Administrator's 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. Note the following:

- **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.

Use the **portCfgEXPort** or **portCfgVEXPort** command to:

- Enable or disable EX_Port or VEX_Port mode.
- Set the fabric ID (avoid using fabric IDs 1 and 128, which are the default IDs for backbone connections).
- Configure an EX_Port to connect to a Brocade Network OS fabric (**portCfgEXPort** only).

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:                NOT 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 configures an EX_Port for connecting to a Brocade Network OS fabric. The **-m 5** option indicates Network OS connectivity.

```
switch:admin> portcfgexport 1/5 -a 1 -f 15 -m 5
switch:admin> portcfgexport 1/5
Port 1/5 info
Admin: enabled
State: NOT OK
Pid format: Not Applicable
Operate mode: Brocade NOS
Edge Fabric ID: 15
(output truncated)
```

For related FC-FC routing commands, refer to **fcrEdgeShow**, **fcrXlateConfig**, **fcrConfigure**, and **fcrProxyConfig** in the *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. (Optional) 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 613.
4. (Optional) Set up ISL or EX_Port trunking. For information on trunking setup, refer to [“Configuring EX_Port trunking”](#) on page 578.
5. Enter the **portEnable** command to enable the ports that you disabled in [step 1](#).

```
switch:admin> portenable 7/10
```

6. Physically attach ISLs from the Fibre Channel router to the edge fabric.
7. Enter the **portCfgShow** command to view ports that are persistently disabled.

FC ports on the Brocade 7800 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 the **portCfgPersistentEnable** command to enable the port, and then the **portCfgShow** command to verify 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. Enter either the **portCfgEXPort** or **portShow** command to verify that each port is configured correctly.

```
switch:admin> portcfgexport 7/10
```

```

Port 7/10 info
Admin:          enabled
State:          NOT 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_06> portshow 7/10
```

```

portName:
portHealth: OFFLINE

Authentication: None

EX_Port Mode:   Enabled
Fabric ID:      30
Front Phantom:  state = Not OK  Pref Dom ID: 160
Fabric params:  R_A_TOV: 0      E_D_TOV: 0      PID fmt: auto

Authentication Type: None
Hash Algorithm: N/A
DH Group: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A

portDisableReason: None
portCFlags: 0x1
portFlags: 0x1  PRESENT U_PORT EX_PORT
portType: 10.0

```

```

portState: 2    Offline
portPhys: 2    No_Module
portScn: 0
port generation number: 0
portId: 014a00
portIfId: 4372080f
portWwn: 20:4a:00:60:69:e2:03:86
portWwn of device(s) connected:

```

```

Distance: normal
portSpeed: N4Gbps

```

```

LE domain: 0

```

```

FC Fastwrite: ON

```

Interrupts:	0	Link_failure:	0	Frjt :	0
Unknown:	0	Loss_of_sync:	0	Fbsy :	0
Lli:	0	Loss_of_sig:	2		
Proc_rqrd:	0	Protocol_err:	0		
Timed_out:	0	Invalid_word:	0		
Rx_flushed:	0	Invalid_crc:	0		
Tx_unavail:	0	Delim_err:	0		
Free_buffer:	0	Address_err:	0		
Overrun:	0	Lr_in:	0		
Suspended:	0	Lr_out:	0		
Parity_err:	0	Ols_in:	0		
2_parity_err:	0	Ols_out:	0		
CMI_bus_err:	0				

```

Port part of other ADs: No

```

10. Enter the **switchShow** command to verify 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 the **fcrFabricShow** command 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 "5300" 1080::8:8FF:FE0C:417A/64
4 116 10:00:00:05:1e:37:00:44 10.32.156.34 "5300"
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 "5300"
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"

```

- Enter the `iflshow` command to display the FC router details and ensure the fabric is functioning correctly.

```

switch:FID128:root> 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:

- Both the active and standby CP must be running Fabric OS 7.2.0 or later.
- EX_Ports on ICLs is supported only on the Brocade DCX 8510-4 and DCX 8510-8.
- EX-Port 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, FC8-32E, FC8-48E).

NOTE

Device discovery will not occur properly with ICL EX-Port connected edge fabrics if the FC Router has unsupported blades which are not mentioned in the above restriction.

- Switches must have Virtual Fabrics enabled.
- EX_Ports on ICLs are allowed only on the base switch. All user ports must be 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 Inter-Chassis Links chapter.
- 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:

- 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.

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, disable all QSFP ports by issuing the **portDisable** command.

```
switch:admin> portdisable 6/20-23
```

You can verify that all ports have been disabled by issuing the **portShow** command for the ports.

3. Configure EX_Ports on the ICL by issuing the **portCfgEXPort** command. 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. (Optional) 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 613.

5. Enter the **portEnable** command to enable the QSFP ports that you disabled in step 1.

```
switch:admin> portenable 6/20-23
```

6. Enter either the **portCfgEXPort** or **portShow** command 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 the **switchShow** command to verify the EX_Port and edge fabric ID are correct.

8. Enter the **fcEdgeShow** command to display the FIDs of all configured EX_Ports.

```
switch:admin> fcEdgeshow
  FID EX-port E-port Neighbor Switch (PWWN, 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** command to monitor and report path bandwidth imbalances for each edge fabric.

The following example enables the monitoring and reporting of bandwidth balances and imbalances, and 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
Pluto:FID2:root> 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

The FC router port cost is set automatically. This section provides information about the router port cost and describes how you can modify the cost for a port if you want to change the default value.

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.

The FC router port cost feature optimizes the usage of the router port links by directing the traffic to a link with a smaller cost.

Every IFL has a default cost. The default router port cost values are:

- 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.

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. However, you can modify the cost for that port. You can configure the EX_Port or VEX_Port with values of either 1,000 or 10,000. 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.

For details about the use of any of the following commands, refer to the *Fabric OS Command Reference*.

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
```

3. Enter the **forRouterPortCost** 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 **forRouteShow** 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
```

4. 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
```

5. 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 FCR 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.

[Figure 91](#) shows a 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

Link cost should be set in the direction of traffic flow from the source edge fabric to the destination edge fabric.

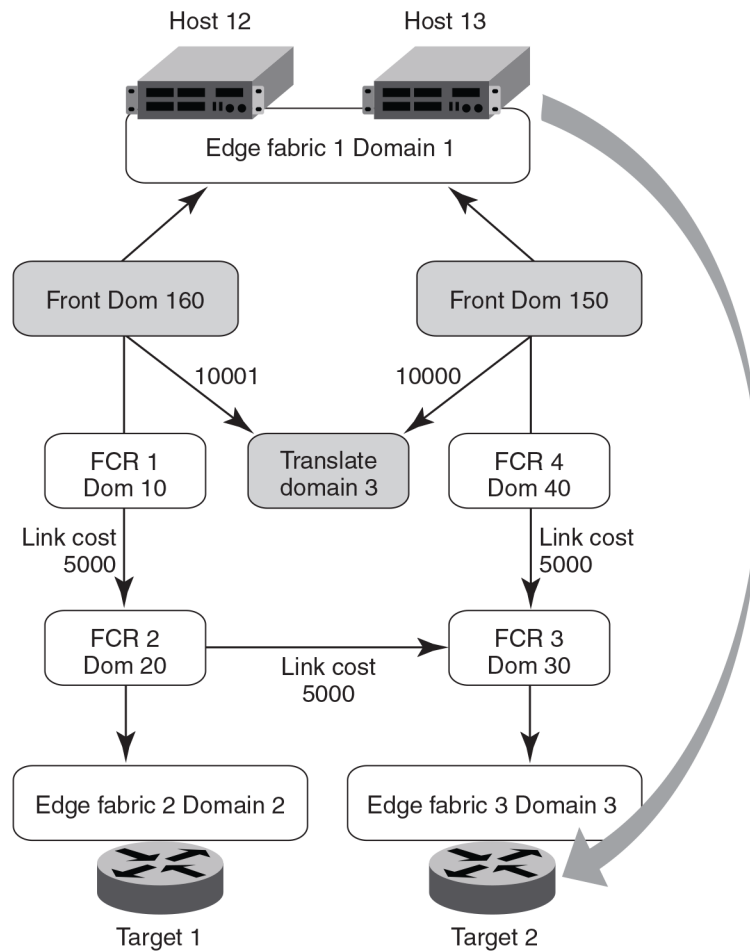


FIGURE 91 Shortest IFL solution

Configuring shortest IFL cost

1. Enter the `fcFabricShow` 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 FIDNeighbor Switch Info (enet IP, WWN, name)
-----
34 110.17.33.68 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 FIDNeighbor Switch Info (enet IP, WWN, name)
-----
2210.17.33.88 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 FIDNeighbor Switch Info (enet IP, WWN, name)
-----
31 310.17.33.153 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 FIDNeighbor Switch Info (enet IP, WWN, name)
-----
      34 110.17.33.68      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 30 is modified.

```
switch:admin>linkcost 10 5000
Interface10 (E_PORT)Cost 5000
```

- In the following example, the ISL link cost of path 2 from FC router Domain 30 to FC router Domain 20 is modified.

```
switch:admin>linkcost 10 5000
Interface10 (E_PORT)Cost 5000
```

The modified cumulative link cost for path 2 is 10,000.

6. Enter the **fcrConfigure -enable -shortestifl** command to enable the shortest IFL mode in all the FC routers in the backbone fabric.

```
switch:admin> fcrconfigure --enable -shortestifl
Shortest IFL path is enabled
```

7. Enter the **fcrConfigure -show** command to verify that the shortest IFL mode is enabled in all the FC routers in the backbone fabric.

```
switch:admin> fcrconfigure --show
Backbone fabric ID: 128
Shortest IFL feature is enabled
```

8. Verify that the link cost of both the front domain and translate domain in the source edge fabric have been modified using the **lsDbShow** command.

```
linkCnt = 2, flags = 0x0
LinkId = 160, out port = 160 rem port = 241, cost = 10001, costCnt = 0, type
= 1
LinkId = 161, out port = 161, rem port = 242, cost = 10000, costCnt = 0, type
= 1
```

EX_Port frame trunking configuration

You can configure EX_Ports to use frame-based trunking just as you do regular E_Ports. EX_Port frame 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 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.

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.

ATTENTION

This feature should be enabled only if the entire configuration is running Fabric OS v5.2.0 or later.

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 [Chapter 24, “Managing Trunking Connections”](#).

LSAN zone configuration

An LSAN consists of zones in two or more edge or backbone fabrics that contain the same devices. LSANs provide selective device connectivity between fabrics without forcing you to merge those fabrics. FC routers provide multiple mechanisms to manage inter-fabric device connectivity through extensions to existing switch management interfaces. You can define and manage LSANs using Brocade Advanced Zoning.

NOTE

For performance reasons, Brocade recommends that you do not configure LSANs for device sharing between Fabric OS fabrics until after you activate the Integrated Routing license.

Use of Admin Domains with LSAN zones and FC-FC routing

You can create LSAN zones as a physical fabric administrator or as an individual Admin Domain (AD) administrator. The LSAN zone can be part of the root zone database or the AD zone database. FC-FC Routing harvests the LSAN zones from all administrative domains. If both edge fabrics have the matching LSAN zones and both devices are online, FC-FC routing triggers a device import. To support legacy applications, WWNs are reported based on the Admin Domain context. As a result, you must not use the network address authority (NAA) field in the WWN to detect an FC router. LSAN zone enforcement in the local fabric occurs only if the Admin Domain member list contains both of the devices (local and imported device) specified in the LSAN zone.

For more information, refer to [Chapter 20, “Managing Administrative Domains”](#).

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 87](#) on page 599, 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” on page 419 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; 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 and connect to switch1.
 2. Enter the **nsShow** command to list the WWN of the host (10:00:00:00:c9:2b:c9:0c).

NOTE

The **nsShow** output displays the LSAN zone status of a device, the port WWN, and the node WWN; the port WWN must be used for LSANs.

```
switch:admin> nsshow
{
  Type Pid      COS      PortName      NodeName
  TTL(sec)
  N   060f00;   2,3;     10:00:00:00:c9:2b:c9:0c;  20:00:00:00:c9:2b:c9:0c; na
  FC4s: FCP
  NodeSymb: [35] "Emulex LP9002 FV3.91A3 DV5-5.20A6 "
  Fabric Port Name: 20:0f:00:05:1e:37:00:44
  Permanent Port Name: 10:00:00:00:c9:2b:c9:0c
  LSAN: Yes
  The Local Name Server has 1 entry }
```

3. Enter the **zoneCreate** command to create the LSAN *lsan_zone_fabric75*, which includes the host.

```
switch:admin> zonecreate "lsan_zone_fabric75", "10:00:00:00:c9:2b:c9:0c"
```

4. Enter the **zoneAdd** command to add Target A to the LSAN.

```
FID75Domain5:admin> zoneadd "lsan_zone_fabric75", "50:05:07:61:00:5b:62:ed"
```

5. Enter the **cfgAdd** or **cfgCreate** and **cfgEnable** commands to add and enable the LSAN configuration.

```
switch:admin> cfgadd "zone_cfg", "lsan_zone_fabric75"
switch:admin> cfgenable "zone_cfg"
```

You are about to enable a new zoning configuration.

This action will replace the old zoning configuration with the current configuration selected.

Do you want to enable 'zone_cfg' configuration (yes, y, no, n): [no] **y**
zone config "zone_cfg" is in effect

Updating flash ...

6. Log in as admin to fabric2.

7. Enter the **nsShow** command to list 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                               TTL(sec)
  NL  0508e8; 3;    50:05:07:61:00:5b:62:ed; 50:05:07:61:00:1b:62:ed; na
      FC4s: FCP [IBM   DNEF-309170   F90F]
      Fabric Port Name: 20:08:00:05:1e:34:11:e5
      Permanent Port Name: 50:05:07:61:00:5b:62:ed
  NL  0508ef; 3;    50:05:07:61:00:49:20:b4; 50:05:07:61:00:09:20:b4; na
      FC4s: FCP [IBM   DNEF-309170   F90F]
      Fabric Port Name: 20:08:00:05:1e:34:11:e5
      Permanent Port Name: 50:05:07:61:00:49:20:b4
  LSAN: Yes
  The Local Name Server has 2 entries }
```

8. Enter the **zoneCreate** command 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"
```

9. Enter the **cfgShow** command 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
```

10. Enter the **cfgAdd** and **cfgEnable** commands to create and enable 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 ...

```

11. Log in as admin and connect to the FC router.

12. Enter the following commands to display information about the LSANs:

- **lsanZoneShow -s** shows the LSAN.

```

switch:admin> lsanzoneshow -s
Fabric ID: 2 Zone Name: lsan_zone_fabric2
    10:00:00:00:c9:2b:c9:0c Imported
    50:05:07:61:00:5b:62:ed EXIST
    50:05:07:61:00:49:20:b4 EXIST
Fabric ID: 75 Zone Name: lsan_zone_fabric75
    10:00:00:00:c9:2b:c9:0c EXIST
    50:05:07:61:00:5b:62:ed Imported

```

- **fcrPhyDevShow** shows the physical devices in the LSAN.

```

switch:admin> fcrphydevshow
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

```

- **fcrProxyDevShow** shows the proxy devices in the LSAN.

```

switch:admin> fcrproxydevshow
Proxy           WWN           Proxy           Device           Physical           State
Created         PID           Exists           PID
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; for the remaining frames zoning enforcement takes place in the edge fabrics.

Configuring backbone fabrics for interconnectivity

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 621. However, instead of configuring the LSAN in the second edge fabric, configure the LSAN in the backbone fabric.

Setting the maximum LSAN count

You can set the maximum number of LSAN zones, or LSAN count, that can be configured on the edge fabrics. 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 could lead to different devices being imported on different FC routers.

1. Enter the **fcrlsancount** command with no parameters to display the current LSAN limit.

```
switch:admin> fcrlsancount
LSAN Zone Limit 3000
```

2. Enter the **fcrlsancount** command 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 634.

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 **fcResourceShow** command.

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.

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

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. This is useful when multiple FC routers are connected to the same edge fabric. Without the Enforce tag, all FC routers import all LSAN zones, even those that are not needed.

Normally, 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 only accept 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.

Speed tag

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. The Speed tag allows you to speed up the discovery process by importing the devices into the remote edge fabrics when the devices come online, regardless of the state of the host. This helps sensitive hosts to quickly discover the devices without timing out.

You set the Speed tag on the FC router, and then configure the LSANs in the target edge fabrics with the tag.

For example, in [Figure 92](#) on page 626, 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.

Example

```
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 following command to create a Speed LSAN tag:

```
fcrlsan --add -speed tagname
```

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.

Example

```
sw0:admin> fcrlsan --add -speed fasttag2  
LSAN tag set successfully
```

Removing an LSAN tag

Use the following procedure to remove an LSAN tag. This procedure does not remove the LSAN zone; it deactivates the tag so that LSAN zones with this tag in the name now behave as regular LSAN zones.

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.

Example

```
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.

NOTE

LSAN zone binding is supported only on FC routers with Fabric OS v5.3.0 and later. The FC router matrix feature is supported only on FC routers with Fabric OS v6.1.0 and later.

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.

Figure 93 on page 629 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. Table 94 shows what LSAN information is stored in each FC router before and after LSAN zone binding is in effect.

TABLE 94 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
LSAN 1	LSAN 1	LSAN 1	LSAN 1	LSAN 1	LSAN 2	LSAN 3	LSAN 4
LSAN 2	LSAN 2	LSAN 2	LSAN 2	LSAN 2		LSAN 4	
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 this 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 624 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 [Figure 93](#), 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 [Figure 93](#) as an example, the following edge fabrics can access each other:

- Fabric 1 and Fabric 2
- Fabric 2 and Fabric 3
- Fabric 4 and Fabric 5
- Fabric 5 and Fabric 6

You can use the **fcrLsanMatrix** command with the **-lsan** option to create the corresponding LSAN fabric matrix:

```
fcrLsanMatrix --add -lsan 1 2
fcrLsanMatrix --add -lsan 2 3
fcrLsanMatrix --add -lsan 4 5
fcrLsanMatrix --add -lsan 5 6
```

Fabrics that are not specified are part of the default binding and can access other edge fabrics that are not specified. Thus, fabrics 7, 8, and 9 can access each other, but cannot access fabrics 1 through 6.

ATTENTION

The **fcrLsanMatrix --add -lsan 0 0** command will erase the entire LSAN fabric matrix settings in the cache.

The FC router matrix and the LSAN fabric matrix are used together to determine which fabrics can access each other, with the LSAN fabric matrix providing more specific binding.

Setting up LSAN zone binding

1. Log in to the FC router as admin.
2. Enter the following command to add a pair of FC routers that can access each other:

```
FCR:Admin> fcrlsanmatrix --add -fcr wwn1 wwn2
```

The variables *wwn1* and *wwn2* are the WWNs of the FC routers.

3. Enter the following command to add a pair of edge fabrics that can access each other:

```
FCR:Admin> fcrlsanmatrix --add -lsan fid1 fid2
```

The variables *fid1* and *fid2* are the fabric IDs of the edge fabrics.

4. Enter the following command to apply the changes persistently:

```
FCR:Admin> fcrlsanmatrix --apply -all
```

Example

```
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 matrixes

1. Log in to the FC router as admin.
2. Enter the following command to view the FC router matrix:

```
fcrlsanmatrix --fabricview -fcr
```

3. Enter the following command to view the LSAN fabric matrix:

```
fcrlsanmatrix --fabricview -lsan
```

Example

```
FCR: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)

FCR:Admin> fcrlsanmatrix --fabricview -lsan
LSAN MATRIX is activated

Fabric ID      Fabric ID
-----
4              5
4              7
10            19
```


Proxy PID configuration

When an FC router is first configured, the PIDs for the proxy devices are automatically assigned. Proxy PIDs (as well as phantom domain IDs) persist across reboots.

The most common situation in which you would set a proxy PID is when you replace a switch. If you replace the switch and want to continue using the old PID assignments, you can configure it to do so; this value remains in the system even if the blade is replaced. To minimize disruption to the edge fabrics, set the proxy PIDs to the same values used with the old hardware.

The **fcProxyConfig** command displays or sets the persistent configuration of proxy devices. Used with the **-s slot** 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. For Proxy IDs projected to an M-EOS edge fabric in McDATA fabric mode, use valid ALPAs (lower 8 bits).

Use the **fcXlateConfig** 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 343 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> fcrbcastconfig --show
```

This command displays only the FIDs that have the broadcast frame option enabled. The FIDs that are not listed have the broadcast frame option disabled.

Enabling broadcast frame forwarding

1. Log in to the FC router as admin.
2. Enter the following command:

```
fcr:admin> fcrbcastconfig --enable -f fabricID
```

The *fabricID* variable is the FID of the edge or backbone fabric on which you want to enable broadcast frame forwarding. Broadcast frame forwarding is enabled by default.

Disabling broadcast frame forwarding

1. Log in to the FC router as admin.
2. Enter the following command:

```
fcr:admin> fcrbcastconfig --disable -f fabricID
```

The *fabricID* variable is the FID of the edge or backbone fabric on which you want to disable broadcast frame forwarding.

Resource monitoring

It is possible to exhaust resources, such as proxy PIDs. Whenever a resource is exhausted, Fabric OS generates an error message. The messages are described in the *Fabric OS Message Reference*.

You can monitor FC router resources using the **fcrResourceShow** command. The **fcrResourceShow** 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 624 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 **forResourceShow** command to display physical port (EX_Port) resources.

```
switch:admin> forresourceshow
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: 2000
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 a different chassis. However, the FID of the EX_Port must be set to a different value than the FID of the logical switch to which it connects.
- EX_Ports and VEX_Ports – those 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 to use XISLs”](#) on page 333 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, as shown in [Figure 27](#) on page 316.

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 [“Backbone-to-edge routing with Virtual Fabrics”](#) on page 638 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 **IsCfg** command to configure the FID.
- Although the Brocade 6510 and 6520 support up to four logical switches, if you are using FC-FC routing, they can have a maximum of only three logical switches.

Logical switch configuration for FC routing

Figure 94 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. Note the following:

- 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 94 EX_Ports in a base switch

Figure 95 shows a logical representation of the physical chassis and devices in Figure 94. As shown in Figure 95, 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 95 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 94](#), no devices can be connected to the backbone fabric (Fabric 8) because base switches cannot have F_Ports. [Figure 96](#) 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 96 Backbone-to-edge routing across base switch using FC router in legacy mode

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.

For further instructions on downgrading, refer to [Chapter 10, “Installing and Maintaining Firmware”](#).

How replacing port blades affects EX_Port configuration

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.

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.

26 Displaying the range of output ports connected to xlate domains

1. Log in to a switch in the edge fabric.
2. Enter the **lsDbShow** command on the edge fabric.

In the **lsDbShow** output, ports in the range from 129 through 255 are the output ports on the front domain.

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
```

The following example also shows the use of the **lsDbShow** display on the edge fabric. The front domain, domain 3, has two links representing two EX_Port connections with output ports 129 and 132.

```
Domain = 3, Link State Database Entry pointer = 0x100bbcc0
.....
linkCnt = 4,    flags = 0x0
LinkId = 199, out port = 129, rem port = 2, cost = 10000, costCnt = 0, type = 1
LinkId = 199, out port = 132, rem port = 3, cost = 10000, costCnt = 0, type = 1
LinkId = 2, out port = 1, rem port = 2, cost = 10000, costCnt = 0, type = 1
LinkId = 1, out port = 32, rem port = 2, cost = 10000, costCnt = 0, type = 1
```


Port Indexing

This appendix shows how to use the **switchShow** command to determine the mapping among the port index, slot/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.

Example of port index mapping on a CR16-4 blade in a DCX 8510-4 Backbone

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: switch name
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
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
```

A Port Indexing

```

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)

```

Example of port index mapping on an FC16-32 blade of a Brocade DCX 8510-8 Backbone

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)

```

Example of port index mapping on an FC8-64 blade on a Brocade DCX Backbone.

This example shows the truncated **switchShow** output for an FC8-64 port blade on the Brocade DCX 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   --     N8     No_Module
1      1      1      0a0140   --     N8     No_Module
2      1      2      0a0240   --     N8     No_Module
(output truncated)
768    1      48      0a00c0   --     N8     No_Module
769    1      49      0a01c0   --     N8     No_Module
770    1      50      0a02c0   --     N8     No_Module
(output truncated)
784    1      62      0a0ec0   --     N8     No_Module
783    1      63      0a0fc0   --     N8     No_Module
16     2      0      0a1040   --     N8     No_Module
17     2      1      0a1140   --     N8     No_Module
(output truncated)

```

Example of port indexing on an FC8-64 blade on a Brocade DCX-4S Backbone.

The Brocade DCX-4S does not need a mapping of ports on port blades because it is a one-to-one mapping. The order is sequential starting at slot 1 port 0 all the way through slot 8 port 255 for the FC8-64 blade. For core blades, the port index mapping for the blade in slot 3 begins with port index 256, and port index mapping for the core blade in slot 6 begins with port index 736. There are no shared areas on the Brocade DCX-4S.

The following example **switchShow** output is from a Brocade DCX-4S. It shows the index and PID addressing. The output has been truncated.

```
DCX-4S:admin> switchshow
  Index  Slot  Port  Address  Media  Speed  State
=====
   0     1    0    0a0000   --    N8     No_Module
   1     1    1    0a0100   --    N8     No_Module
   2     1    2    0a0200   --    N8     No_Module
(output truncated)
  48     1   48    0a3000   --    N8     No_Module
  49     1   49    0a3100   --    N8     No_Module
  50     1   50    0a3200   --    N8     No_Module
(output truncated)
  62     1   62    0a3e00   --    N8     No_Module
  63     1   63    0a3f00   --    N8     No_Module
  64     2    0    0a4000   --    N8     No_Module
(output truncated)
```

Example of port indexing on an FX8-24 blade on a DCX 8510-8 Backbone

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)
```

A Port Indexing

Example of port indexing on an FS8-18 blade on a DCX 8510-8 Backbone

This example shows the truncated **switchShow** output for an FS8-18 encryption 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 2  
switchName: myswitch  
      (output truncated)
```

Slot	Blade Type	ID	Model Name	Status
2	AP BLADE	43	FS8-18	ENABLED

Index	Slot	Port	Address	Media	Speed	State	Proto
16	2	0	501000	--	N8	No_Module	FC
17	2	1	501100	--	N8	No_Module	FC
18	2	2	501200	--	N8	No_Module	FC
19	2	3	501300	--	N8	No_Module	FC
20	2	4	501400	--	N8	No_Module	FC
(output truncated)							
31	2	15	501f00	id	N4	No_Light	FC

FIPS Support

In this appendix

- [FIPS overview](#) 645
- [Zeroization functions](#) 645
- [FIPS mode configuration](#) 647
- [Preparing a switch for FIPS](#) 651

FIPS overview

Federal information processing standards (FIPS) specify the security standards to be satisfied by a cryptographic module utilized in Fabric OS v6.0.0 and later to protect sensitive information in the switch.

As part of FIPS 140-2 level 2, compliance passwords, shared secrets, and the private keys used in SSL, TLS, and system login need to be cleared out or *zeroized*. Before enabling FIPS compliance mode, a power-on self-test (POST) is executed when the switch is powered on to check for the consistency of the algorithms implemented in the switch. Known-answer tests (KATs) are used to exercise various features of the algorithm and their results are displayed on the console for your reference. Conditional tests are performed whenever an RSA key pair is generated. These tests verify the randomness of the deterministic random number generator (DRNG) and the non-deterministic random number generator (non-DRNG). They also verify the consistency of RSA keys with regard to signing and verification and encryption and decryption.

ATTENTION

FIPS mode, when enabled, is a chassis-wide setting that affects all logical switches. Once enabled, FIPS mode cannot be disabled.

Zeroization functions

Zeroization functions can be performed at the discretion of the security administrator. These functions clear the passwords and the shared secrets. Core files and FFDC data are also removed upon FIPS Zeroization. [Table 95](#) lists the various keys used in the system that will be zeroized in a FIPS-compliant Fabric OS module.

TABLE 95 Zeroization behavior

Keys	Zeroization CLI	Description
DH private keys	No command required	Keys will be zeroized within code before they are released from memory.
FCAP private key	<code>secCertUtil delete --fcapall -nowarn</code>	The <code>secCertUtil delete --fcapall -nowarn</code> command removes all FCAP certificates and FCAP private keys.

B Zeroization functions

TABLE 95 Zeroization behavior (Continued)

Keys	Zeroization CLI	Description
FCSP Challenge Handshake Authentication Protocol (CHAP) Secret	secAuthSecret --remove	The secAuthSecret --create command is used to input the keys, and the secAuthSecret --remove command is used to remove and zeroize the keys. All the DH-CHAP and FCAP authenticated ports are disabled after zeroization.
LDAP CA certificate	secCertUtil delete -ldapcert <certname>	The given LDAP certificate file is zeroized and deleted from the module.
Passwords	passwdDefault	<p>The passwdDefault command removes user-defined accounts in addition to default passwords for the root, admin, and user default accounts. However, only the root account has permissions for this command. Users with securityadmin and admin permissions must use fipsCfg --zeroize, which, in addition to removing user accounts and resetting passwords, also performs the complete zeroization of the system.</p> <p>Notes:</p> <ul style="list-style-type: none"> In a dual CP system, executing passwdDefault syncs with the standby. This means that when passwdDefault is executed in the active CP, user-defined accounts are removed from both the active and standby CPs and only the default accounts (root, factory, admin, and user) will be retained. These accounts will have the generic default passwords set. To maintain FIPS 140-2 compliance, passwords for the default accounts (admin and user) must be changed after every zeroization operation.
RADIUS secret	aaaConfig --remove	The aaaConfig --remove command zeroizes the secret and deletes a configured server. The aaaConfig --add command configures the RADIUS server.
RNG seed key	No command required	/dev/urandom is used as the initial source of seed for RNG. The RNG seed key is zeroized on every random number generation.
SFTP session keys	No command required	Automatically zeroized on session termination.
SSH RSA private key	sshUtil delprivkey	Key-based SSH authentication is not used for SSH sessions.
SSH public keys	sshUtil delpubkeys	Zeroizes the SSH public.
SSH session key	No command required	This key is generated for each SSH session that is established with the host. It automatically zeroizes on session termination.
TLS authentication key	No command required	Automatically zeroized on session termination.
TLS pre-master secret	No command required	Automatically zeroized on session termination.
TLS private keys	secCertUtil delkey -all	The secCertUtil delkey -all command is used to zeroize these keys. The secCertUtil genkey command creates the keys. Only RSA keys of size 1024 or 2048 are allowed.
TLS session key	No command required	Automatically zeroized on session termination.

Power-on self-tests

A power-on self-test (POST) is invoked by powering on the switch in FIPS mode and does not require any operator intervention. If any KATs fail, the switch goes into a FIPS Error state, which reboots the system to start the test again. If the switch continues to fail the FIPS POST, you will need to return your switch to your switch service provider for repair. Refer to the *Fabric OS Troubleshooting and Diagnostics Guide* for information about preparing a case for your service provider.

Conditional tests

These tests are for the random number generators and are executed to verify the randomness of the random number generator. The conditional tests are executed each time prior to using the random number provided by the random number generator.

The results of the POST and conditional tests are recorded in the system log or are output to the local console. This action includes logging both passing and failing results. Refer to the *Fabric OS Troubleshooting and Diagnostics Guide* for instructions on how to recover if your system cannot get out of the conditional test mode.

FIPS mode configuration

By default, the switch comes up in non-FIPS mode. You can run the `fipsCfg --enable fips` command to enable FIPS mode, but you must configure the switch first. Self-test mode must be enabled before FIPS mode can be enabled. A set of prerequisites (as shown in [Table 96](#)) must be satisfied for the system to enter FIPS mode. To be FIPS-compliant, the switch must be power-cycled. KATs are run on the reboot. If the KATs are successful, the switch enters FIPS mode. If the KATs fail, then the switch reboots until the KATs succeed. If the switch cannot enter FIPS mode and continues to reboot, you must return the switch to your switch service provider. For information about how to prepare a service provider case, refer to the *Fabric OS Troubleshooting and Diagnostics Guide*.

When the switch successfully reboots in FIPS mode, only FIPS-compliant algorithms are run.

NOTE

RPC is not supported in FIPS mode.

[Table 96](#) lists Fabric OS features and their behaviors in FIPS and non-FIPS mode.

TABLE 96 FIPS mode restrictions

Features	FIPS mode	Non-FIPS mode
Authentication	All ports, including Access Gateway, FC router, and F_Ports, adhere to FIPS guidelines when authentication is enabled.	No restrictions.
Configupload/ download/ supportsave/ firmwaredownload	SCP only	FTP and SCP
DH-CHAP/FCAP hashing algorithms	SHA-1	MD5 and SHA-1
DH-CHAP Shared Secret	Minimum length of 32 bytes for secret used in in-flight encryption	Minimum length of 8 bytes for secret

B FIPS mode configuration

TABLE 96 FIPS mode restrictions (Continued)

Features	FIPS mode	Non-FIPS mode
FC-FC routing	If FIPS is enabled in the FC router and disabled in the edge switch, the EX_Port is disabled if the edge fabric switch has Diffie-Hellman group 0 or hash group MD5.	No restrictions.
HTTP/HTTPS access	HTTPS only	HTTP and HTTPS
HTTPS algorithms	TLS/AES128 cipher suite	TLS AES 128 cipher suite SSL is not supported.
In-flight encryption	Not supported	No restrictions
IPsec	Usage of AES-XCBC, MD5, and DH group 1 are blocked.	No restrictions
LDAP CA	CA certificate must be available.	CA certificate is optional.
Common certificate for FCAP and HTTPS authentication	Not supported	Supported
RADIUS auth protocols	PEAP-MSCHAPv2	CHAP, PAP, PEAP-MSCHAPv2
Root account	Disabled	Enabled
Signed firmware download	Mandatory firmware signature validation (SCP only)	Optional firmware signature validation (FTP and SCP)
SNMP	Read-only operations	Read and write operations
SSH algorithms	HMAC-SHA1 (mac) 3DES-CBC, AES128-CBC, AES192-CBC, AES256-CBC (cipher suites)	No restrictions
SSH public keys	RSA 1024 bit keys and RSA 2048 bit keys	RSA 1024 bit keys, RSA 2048 bit keys, and DSA 1024 bit keys
TACACS+ authentication	Not supported	Supported
Telnet/SSH access	Only SSH	Telnet and SSH

LDAP in FIPS mode

You can configure your Microsoft Active Directory server to use the Lightweight Directory Access Protocol (LDAP) while in FIPS mode. There is no option provided on the switch to configure TLS ciphers for LDAP in FIPS mode. However, the LDAP client checks if FIPS mode is set on the switch and uses the FIPS-compliant TLS ciphers for LDAP. If the FIPS mode is not set and the Microsoft Active Directory server is configured for FIPS ciphers, it uses FIPS-compliant ciphers.

[Table 97](#) lists the differences between FIPS and non-FIPS modes of operation.

TABLE 97 FIPS and non-FIPS modes of operation

FIPS mode	non-FIPS mode
The certificate of the CA that issued the Microsoft Active Directory server certificate must be installed on the switch.	There is no mandatory CA certificate installation on the switch.
Configure FIPS-compliant TLS ciphers [TDES-168, SHA1, and RSA-1024] on the Microsoft Active Directory server. The host needs a reboot for the changes to take effect.	On the Microsoft Active Directory server, there is no configuration of the FIPS-compliant TLS ciphers.

TABLE 97 FIPS and non-FIPS modes of operation (Continued)

FIPS mode	non-FIPS mode
The switch uses FIPS-compliant ciphers regardless of the Microsoft Active Directory server configuration. If the Microsoft Active Directory server is not configured for FIPS ciphers, authentication will still succeed.	The Microsoft Active Directory server certificate is validated if the CA certificate is found on the switch.
The Microsoft Active Directory server certificate is validated by the LDAP client. If the CA certificate is not present on the switch, then user authentication will fail.	If the Microsoft Active Directory server is configured for FIPS ciphers and the switch is in non-FIPS mode, then user authentication will succeed.

Setting up LDAP for FIPS mode

1. Log in to the switch using an account with admin or securityadmin permissions, or an account with OM permissions for the RADIUS and switch configuration RBAC classes of commands.
2. Enter the `dnsConfig` command to configure the DNS on the switch.

Example of setting the DNS

```
switch:admin> dnsconfig
```

```
Enter option
```

```
1 Display Domain Name Service (DNS) configuration
```

```
2 Set DNS configuration
```

```
3 Remove DNS configuration
```

```
4 Quit
```

```
Select an item: (1..4) [4] 2
```

```
Enter Domain Name: [] domain.com
```

```
Enter Name Server IP address in dot notation: [] 123.123.123.123
```

```
Enter Name Server IP address in dot notation: [] 123.123.123.124
```

```
DNS parameters saved successfully
```

```
Enter option
```

```
1 Display Domain Name Service (DNS) configuration
```

```
2 Set DNS configuration
```

```
3 Remove DNS configuration
```

```
4 Quit
```

```
Select an item: (1..4) [4] 4
```

Specify the DNS IP address using either IPv4 or IPv6. This address is needed for the switch to resolve the domain name to the IP address because LDAP initiates a TCP session to connect to your Microsoft Active Directory server. A Fully Qualified Domain Name (FQDN) is needed to validate the server identity as mentioned in the common name of the server certificate.

3. Set the switch authentication mode and add your LDAP server by using the commands shown in the following example. Provide the Fully Qualified Domain Name (FQDN) of the Microsoft Active Directory server for the host name parameter while configuring LDAP.

Example of setting up LDAP for FIPS mode

```
switch:admin> aaaconfig --add GEOFF5.ADLDAP.LOCAL -conf ldap -d adldap.local
-p 389 -t 3
```

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

```
switch:admin> aaaconfig -show
```

```
RADIUS CONFIGURATIONS
```

```
=====
```

```
RADIUS configuration does not exist.
```

B FIPS mode configuration

```
LDAP CONFIGURATIONS
=====

Position           : 1
Server             : GEOFF5.ADLDAP.LOCAL
Port               : 389
Domain             : adldap.local
Timeout(s)         : 3

Primary AAA Service: LDAP
Secondary AAA Service: Switch database
```

4. Set up LDAP according to the instructions in “[LDAP configuration and Microsoft Active Directory](#)” on page 181, and then configure the following additional Microsoft Active Directory settings:
 - a. To support FIPS-compliant TLS cipher suites on the Microsoft Active Directory server, allow the SCHANNEL settings listed in [Table 98](#).

TABLE 98 Active Directory keys to modify

Key	Sub-key
Ciphers	3DES
Hashes	SHA1
Key exchange algorithm	PKCS
Protocols	TLSv1.0

- b. Enable the FIPS algorithm policy on the Microsoft Active Directory.

LDAP certificates for FIPS mode

To utilize the LDAP services for FIPS between the switch and the host, you must generate a certificate signing request (CSR) on the Active Directory server and import and export the CA certificates. To support server certificate validation, it is essential to have the CA certificate installed on the switch and Microsoft Active Directory server. Use the **secCertUtil** command to import the CA certificate to the switch. This command will prompt for the remote IP and login credentials to retrieve the CA certificate. The CA certificate should be in any of the standard certificate formats, “.cer”, “.crt,” or “.pem”.

LDAP CA certificate file names should not contain spaces when using the **secCertUtil** command to import and export the certificate.

Importing an LDAP switch certificate

This procedure imports the LDAP CA certificate from the remote host to the switch.

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the PKI RBAC class of commands.
2. Enter the **secCertUtil import -ldapcacert** command.

Example of importing an LDAP certificate

```
switch:admin> seccertutil import -ldapcacert
Select protocol [ftp or scp]: scp
Enter IP address: 192.168.38.206
```

```

Enter remote directory: /users/aUser/certs
Enter certificate name (must have ".crt" or ".cer" ".pem" suffix):
LDAPTestCa.cer
Enter Login Name: aUser
Password: <hidden>
Success: imported certificate [LDAPTestCa.cer].

```

Exporting an LDAP switch certificate

This procedure exports the LDAP CA certificate from the switch to the remote host.

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the PKI RBAC class of commands.
2. Enter the **secCertUtil export -ldapcacert** command.

Example of exporting an LDAP CA certificate

```

switch:admin> seccertutil export -ldapcacert
Select protocol [ftp or scp]: scp
Enter IP address: 192.168.38.206
Enter remote directory: /users/aUser/certs
Enter Login Name: aUser
Enter LDAP certificate name (must have ".pem" suffix): swLdapca.pem
Password: <hidden>
Success: exported LDAP certificate

```

Deleting an LDAP switch certificate

This procedure deletes the LDAP CA certificate from the switch.

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the PKI RBAC class of commands.
2. Enter the **secCertUtil show -ldapcacert** command to determine the name of the LDAP certificate file.
3. Enter the **secCertUtil delete -ldapcacert file_name** command, where *file_name* is the name of the LDAP certificate on the switch.

Example of deleting an LDAP CA certificate

```

switch:admin> seccertutil delete -ldapcacert swLdapca.pem
WARNING!!!
About to delete certificate: swLdapca.pem
ARE YOU SURE (yes, y, no, n): [no] y
Deleted LDAP certificate successfully

```

Preparing a switch for FIPS

It is important to prepare a switch for the following restrictions that exist in FIPS mode:

- The root account and all root-only functions are not available.
- HTTP, Telnet, and SNMP must be disabled. Once these ports are blocked, you cannot use them to read or write data from and to the switch.
- The **configDownload** and **firmwareDownload** commands using an FTP server are blocked.

B Preparing a switch for FIPS

Refer to [Table 97](#) on page 648 for a complete list of restrictions between FIPS and non-FIPS modes.

ATTENTION

You need both `securityadmin` and `admin` permissions to enable FIPS mode.

Overview of steps

1. Remove legacy OpenSSH DSA keys.
2. *Optional:* Configure the RADIUS server or the LDAP server.
3. *Optional:* Configure any authentication protocols.
4. *For LDAP only:* Install an SSL certificate on the Microsoft Active Directory server and a CA certificate on the switch for using LDAP authentication.
5. Create separate IP filter policies for IPv4 and IPv6 and block access to Telnet (TCP port 23) or HTTP (TCP port 80).
6. Set the SNMP security level to off.
7. Configure the switch for signed firmware.
8. Disable in-flight encryption.
9. Disable IPsec for Ethernet and IPsec for FCIP.
10. Disable in-band management.
11. Disable authspec modes if TACACS + authentication or non-PEAP RADIUS are configured.
12. Disable root access.
13. Enable the KATs and the conditional tests.
14. Disable the boot PROM access.
15. Enable FIPS.
16. Perform zeroization as described in [“Zeroizing for FIPS”](#) on page 655.

Enabling FIPS mode

1. Log in to the switch using an account with `securityadmin` permissions.
2. Enter the `sshutil delpubkeys` and `sshutil delprivkey` commands to remove legacy OpenSSH DSA keys.

These keys, which previously were the default keys, migrate to Fabric OS v7.0.0 but are no longer supported in FIPS mode. You must remove these keys to remain FIPS-compliant.

NOTE

Support for RSA keys is retained. You can implement RSA keys using the `sshutil` command.

3. *Optional:* Select the appropriate authentication method based on your needs:
 - If the switch is set for RADIUS, enter the `aaaConfig --change` or `aaaConfig --remove` command to modify each server to use only PEAP-MSCHAPv2 as the authentication protocol.

The RADIUS server must also be configured to use only PEAP-MSCHAPv2. Note that among the Windows RADIUS servers supported, only Windows 2000-, Windows 2003-, and Windows 2008-based RADIUS servers may be used in a FIPS-compliant configuration.

- If the switch is set for LDAP, refer to the instructions in [“Setting up LDAP for FIPS mode”](#) on page 649.
4. *Optional:* Set the authentication protocols.
 - a. Enter the `authUtil --set -h sha1` command to set the hash type for MD5, which is used in the DH-CHAP and FCAP authentication protocols.
 - b. Enter the `authUtil --set -g n` command (where *n* represents the DH group) to set the DH group to 1, 2, 3, or 4.
 5. Install the LDAP CA certificate on the switch and Microsoft Active Directory server. Refer to [“LDAP certificates for FIPS mode”](#) on page 650.
 6. Enter the `ipFilter --show` command and verify that no active IP filter policy permits access to Telnet or HTTP ports, even if a higher priority policy explicitly denies such access. If an active IP policy does permit any of these ports, you must modify or deactivate the policy. Create separate policies for IPv4 and IPv6, and block access on Telnet and HTTP ports.
 - a. Enter the `ipFilter` command to create IP filter policies for IPv4 and IPv6. Refer to [“Creating an IP Filter policy”](#) on page 254.
 - b. Add rules to each IP filter policy. Refer to [“Adding a rule to an IP Filter policy”](#) on page 259. You can use the following modifications to the rule to block access to Telnet and HTTP ports:


```
ipfilter --addrule polycname -rule rule_number -sip source_IP -dp
dest_port -proto protocol -act deny
```

 - The `-sip` option can be given as *any*.
 - The `-dp` options for the port numbers for Telnet and HTTP are 23 and 80, respectively.
 - The `-proto` option should be set to TCP.
 - c. Activate each IP filter policy. Refer to [“Activating an IP Filter policy”](#) on page 255.
 - d. Save each IP filter policy. Refer to [“Saving an IP Filter policy”](#) on page 255.

Example

```
ipfilter --create http_block_v4 -type ipv4
ipfilter --addrule http_block_v4 -rule 1 -sip any -dp 80 -proto tcp -act deny
ipfilter --activate http_block_v4
```

7. Use the `snmpConfig --set seclevel` command to turn on SNMP security. When prompted to select the SNMP SET Security Level, enter **3**, for no access.

Example

```
switch:FID128:admin> snmpconfig --set seclevel
Select SNMP GET Security Level
(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 =
No Access): (0..3) [0]
Select SNMP SET Security Level
(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 =
No Access): (0..3) [0] 3
```

8. Enter the `fipsCfg --disable bootprom` command to block access to the boot PROM.

NOTE

This command can be entered only from the root account. It must be entered before disabling the root account.

9. Enter the **configure** command and respond to the following prompts to enable signed firmware:
 - System services: No
 - cfgload attributes: Yes
 - Enforce secure config Upload/Download: Press **Enter** to accept the default.
 - Enforce firmware signature validation: Yes

Example

```
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]
...
cfgload attributes (yes, y, no, n): [no] yes
Enforce secure config Upload/Download (yes, y, no, n): [no]
Enforce firmware signature validation (yes, y, no, n): [no] yes
```

10. Enter the **portCfgEncrypt --disable** command to disable in-flight encryption. You must first disable the port.

Example

```
myswitch:root> portdisable 0
myswitch:root> portcfgencrypt --disable 0
myswitch:root> portenable 0
```

11. Enter the **ipSecConfig --disable** command to disable Ethernet IPsec.
12. Disable IPsec for FCIP connections. The procedure depends on the type of extension blade used.

For FX8-24 extension blades, enter the **portCfg fciptunnel [slot/]port modify -ipsec 0** command.
13. Enter the **portCfg --mgmtif delete** command to disable in-band management.
14. Enter the **aaaconfig --authspec local** command to disable to authspec mode if TACACS + authentication, PAP, or CHAP are configured.
15. Enter the **fipsCfg --enable selftests** command to enable KAT and conditional tests on the switch.
16. Enter the **fipsCfg --verify fips** command to verify the switch is FIPS-ready.
17. Enter the **userConfig --change root -e no** command to block access to the root account.

By disabling the root account, RADIUS and LDAP users with root permissions are also blocked in FIPS mode.
18. Enter the **fipsCfg --enable fips** command.
19. Reboot the switch. For a director, reboot both CPs.

Zeroizing for FIPS

1. Log in to the switch using an account with admin or securityadmin permissions, or a user account with OM permissions for the FIPSCfg RBAC class of commands.
2. Enter the **fipsCfg --zeroize** command.

NOTE

Passwords of the default accounts (admin and user) should be changed after every zeroization operation to maintain FIPS 140-2 compliance.

3. Power-cycle the switch.

Displaying FIPS configuration

1. Log in to the switch using an account with admin or securityadmin permissions, or a user account with OM permissions for the FIPSCfg RBAC class of commands.
2. Enter the **fipsCfg --showall** command.

B Preparing a switch for FIPS

Hexadecimal Conversion

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 an *0x* prefix or an *h* suffix. The address 0xFFFFFA is the same address as FFFFFAh. 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 0x616000

Notice the PID (610600 - bolded) in the **nsShow** output is in hexadecimal.

```
switch:admin> nsshow
{
  Type Pid      COS      PortName                               NodeName                               TTL(sec)
  N      610600;  2,3;10:00:00:00:c9:29:b3:84;20:00:00:00:c9:29:b3:84; na
  FC4s: FCP
  NodeSymb: [36] "Emulex LP9002 FV3.90A7 DV5-5.10A10 "
  Fabric Port Name: 20:08:00:05:1e:01:23:e0
  Permanent Port Name: 10:00:00:00:c9:29:b3:84
  Port Index: 6
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  LSAN: Yes
  The Local Name Server has 1 entry }
```

1. Separate the 6 digits into triplets by inserting a space after every 2 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) = 0 (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 99 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

TABLE 99 Decimal-to-hexadecimal conversion table (Continued)

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
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					

C Hexadecimal Conversion

Index

Numerics

- 10 Gbps operation on an FC port, enabling, 528
- 10-bit addressing mode, 84
- 10G license, 527–530
- 128-bit encryption, in browser, 200
- 16-link ICL license, 524
- 1st POD ICL license, 523
- 256-area addressing mode, 85
- 2nd POD ICL license, 524
- 8G license, 525
- 8-link ICL license, 524

A

- AAA service requests, 168
- aaaConfig command, 169, 171, 189, 193, 194, 652
- accepting distributed user databases locally, 158
- access
 - API, 229
 - browser security support, 200
 - changing account parameters, 157
 - creating accounts, 156
 - deleting accounts, 157
 - HTTP, 229
 - IP address changes, 59
 - log in fails, 59
 - management server, 229
 - NTP, 74
 - password, changing, 64
 - remote access policies, 178
 - secure using SSL, 200
 - serial, 229
 - SNMP, 229
 - switch defaults, 229
 - telnet, 229
 - blocking, 227
 - unblocking, 228
 - using SSL, HTTPS, 200
- Access Control List. *See*: ACL.
- Access Gateway
 - authentication, 248
 - configuring F_Port trunking on, 580

- considerations for Advanced Performance Monitoring, 553
- F_Port trunking for, 579
- F_Port trunking requirements on, 580
- N_Port failover with FA-PWWN, 484
- shared secrets, 250
- accessing
 - devices, 229
 - hosts, 229
 - switches and fabrics, 229
 - zones, 229
- account ID, 60
- account management for Virtual Fabrics, 319
- accounts, 151–194
 - changing parameters, 157
 - creating, 156
 - deleting, 157
 - displaying information, 156
 - local database of users, 155–158
 - lockout policy, 161
 - lockout policy, duration, 162
 - lockout policy, threshold, 162
 - lockouts and denial of service implications, 163
 - managing passwords, 158
 - password policies, 159–163
 - unlocking, 162
 - user-defined, 155
 - See also*: user accounts.
- ACL
 - activating policy changes, 233
 - adding member, 48
 - adding policy member, 234
 - deleting member, 49
 - deleting policy, 233
 - distributing local policies, 263
 - how policies are stored, 231
 - manually distributing policy database, 260
 - policies, 231–235
 - Admin Domain considerations, 232
 - storage, 231
 - viewing, 233
 - Virtual Fabric considerations, 232
 - policy database distribution, 260
 - policy database management, 260
 - policy distribution to other switches, 262

- policy management, 232–235
 - policy members, 232
 - removing policy member, 234
 - resolving conflicting ACL policies, 264
- activating
 - ACL policy changes, 233
 - Admin Domains, 498
 - IP Filter policy, 255
 - licenses, 533
 - ports on demand, 535
 - TI zones, 406
- ad command, 494, 498, 499, 500, 501, 502, 503, 506, 507, 508
- AD0, ACL management, 232
- AD0, and Admin Domains, 488
- AD255, ACL management, 232
- AD255, and Admin Domains, 489
- Adaptive Networking
 - bottleneck detection, 413
 - Ingress Rate Limiting, 414
 - overview, 413–414
 - Quality of Service, 414
 - Top Talkers, 413
 - Traffic Isolation Zoning, 413
- adding
 - a Top Talker monitor to a port (port mode), 565
 - Admin Domain members, 499
 - alias members, 347
 - end-to-end monitors, 553
 - frame monitors to a port, 560
 - licensed features, 533
 - members to a zone configuration, 363
 - ports to logical switches, 329
 - public key to switch, 198
 - rules to an IP Filter policy, 259
 - switch or fabric to a zone, 371
 - switches to a zone, 371
 - Top Talker monitors on all switches in fabric (fabric mode), 565
 - zone members, 351
- address
 - IPv4 filter policy, 256
 - IPv6 filter policy, 256
- addressing mode
 - 10-bit, 84
 - 256-area, 85
 - core PID, 84
 - fixed, 84, 474
- Admin Domain number and domain ID, 486
- Admin Domains
 - about, 485
- access levels, 487
- ACL policy considerations, 232
- activating, 498
- AD list
 - Microsoft Active Directory, 183
 - OpenLDAP, 188
 - RADIUS, 174
 - TACACS+, 191
- AD0, 488
- AD255, 488, 489
- adding members, 499
- assigning to an existing user account, 497
- assigning users to, 496
- compatibility, availability, and merging, 494
- configDownload command, 512
- configUpload command, 512
- configuration, defined, 494
- configuration, displaying, 508
- creating, 495
- deactivating, 499
- defined AD configuration, 494
- deleting, 501, 502
- device members, 491
- effective AD configuration, 494
- features, 487
- home AD, 490
 - Microsoft Active Directory, 183
 - OpenLDAP, 188
 - RADIUS, 174
 - TACACS+, 191
- implementing, 495
- interaction with Fabric OS features, 509
- LDAP server, 183, 187
- logging in to, 490
- LSAN zones, 511
- managing, 485–512
- member types, 491–492
- Microsoft Active Directory service, 183
- numbering, 485
- OpenLDAP server, 187
- physical fabric administrator, 487
- RADIUS configuration, 173
- RADIUS server configuration, 173
- recommended maximum number, 485
- removing from user accounts, 498
- removing members, 500
- renaming, 500
- requirements, 487
- role considerations, 153
- switch members, 492
- switch port members, 491
- switch WWN, 492
- switching context, 508

- system-defined, 488
- TACACS+ service, 191
- TI zone considerations, 398
- transaction model, 494
- trunk area, 576
- user-defined, 488
- using, 506
- validating members, 506
- VF mode and, 325
- Virtual Fabrics permissions, 151
 - zone database, 510
- admin lockout policy, disabling, 162
- admin lockout policy, enabling, 162
- Administrative Domains. *See: Admin Domains.*
- Advanced Performance Monitoring, 551–553
 - Access Gateway considerations, 553
 - monitor types, 551
 - restrictions, 552
 - Virtual Fabrics considerations, 552
- advanced zoning, 337–376
- advanced zoning commands, 338
- alerts
 - bottleneck, 431, 433–440
 - congestion or latency, 433
 - using bottleneckMon command, 433
- alias
 - adding members, 347
 - creating, 347
 - deleting, 349
 - removing members, 348
- Alias server, described, 46
- aliCreate command, 347
- aliDelete command, 349
- aliRemove command, 348
- aliShow command, 349
- AP Dedicated Link policy, 124
- AP policies and routing policies, 124
- AP route
 - policies, 124
 - setting route policy, 125
- AP Shared Link policy, 124
- applications
 - blade compatibility, 98
 - listener applications blocked, 228
 - used by switches, 229
- aptPolicy command, 122, 125
- assigning user-defined roles, 155
- assigning users to Admin Domains, 496
- audit log
 - configuration, 107
 - configuring for specific event classes, 109
- auditCfg command, 109
- auditDump command, 109
- AUTH module, Virtual Fabric considerations, 244
- AUTH policy, 243, 244
 - distributing fabric-wide, 253
- authenticating devices, 246
 - Virtual Fabrics considerations, 247
- authenticating E_Ports, 244
- authentication
 - algorithms used, 269
 - authentication policy restrictions, 247
 - configuring and enabling, 453
 - configuring for devices, 247
 - configuring for E_Ports, 245
 - configuring incoming SSH, 198
 - configuring outgoing SSH, 199
 - configuring policy, 243–253
 - enabling, 171
 - fabric license, 243
 - FCAP, starting, 253
 - joining FC routers and edge fabric switchws, 453
 - key generation, 448
 - protocols, 248
 - re-authenticating an E_Port, 246
 - setting protocols, 248
 - TACACS+ service, 189
 - using with Access Gateway, 248
 - viewing parameter settings, 248
- authentication server, 193
 - adding, 193
 - data, 168
 - deleting, 193
 - reordering, 193
- authentication service
 - configuring, 167–170
 - disabling, 193, 194
 - enabling, 193
 - local, 194
 - Microsoft Active Directory and LDAP, 181
 - modifying, 193
 - OpenLDAP, 184–189
 - RADIUS, 172
 - remote, 167–194
- authorization, fabric-wide distribution of policy, 253
- authUtil command, 245, 246, 247, 248, 253, 453, 653
- auto-assigned FA-PWWN behavior, 480
- auto-leveling, FR4-18i blade, 298, 304
- automatic PID assignment, enabling, 86

B

Backbone

- assigning fabric IDs, 606, 617
- blade compatibility, 98
- fabric ID, 605–606
- fabric, described, 596
- port blades, described, 88
- port configurations supported, 321
- port restrictions, 321
- shutdown, 81
- upgrading firmware, 297

Backbone fabric, and TI zones, 389

Backbone firmware, 296–299

- download, 296
- download process overview, 296
- version testing, 304

Backbone-to-edge routing, 600, 605

backing up a configuration, 279

base fabric, 319

base switch

- about, 316
- creating, 326
- defined, 317
- extended ISLs and, 316

blade

- application compatibility, 98
- Backbone compatibility, 98
- compatibility, 95
- control processor (CP), 97
- core, 97
- disabling, 99
- disabling port blades, 98
- enabling, 99
- FX8-24 compatibility, 98
- limits, 98
- port area ID, 90
- port compatibility, 98
- port identification, 89
- port indexing, 90
- port numbering schemes, 89
- powering off, 103
- powering on, 103
- swapping, 99–102
- terminology and compatibility, 95
- types, 89
- upgrading firmware, 297

bladeCfgGeMode command, 529

bladeDisable command, 99

bladeEnable command, 99

bladeSwap command, 99

blocked listener applications, list, 228

blocking telnet access, 227

bond0 logical network interface, 65

boot PROM password, 163–167

Backbone with recovery string, 164

Backbone without recovery string, 166

switch with recovery string, 163

switch without recovery string, 165

bottleneck, 427–443

access gateway considerations, 430

alert-related parameters, 435

alerts, 433–440

congestion or latency alerts, 433

congestion type, 428

defined, 427

disabling detection, 442

displaying statistics, 442

enabling detection, 431

excluding a port from detection, 440

high availability considerations, 430

latency type, 428

reporting, 428

settings retention, 435

trunking considerations, 430

types, 428

upgrade and downgrade considerations, 430

virtual fabrics considerations, 430

bottleneck detection, 427–443

advanced settings, 439

alert status, 431

configuration retention, 430

configurations, 429

considerations

access gateway, 430

high availability, 430

trunking, 430

upgrades and downgrades, 430

virtual fabrics, 430

disabling, 442

displaying status, 431

enabling, 431

history, 428

history retention time, 428

licensing, 427

limitations, 429

parameters, 435

slave ports, 440

bottleneckMon command, 428, 431, 433, 435, 436, 440, 442

Broadcast server, described, 46

broadcast zones, 337, 343

name restriction, 350

- Brocade 7800, upgrade license, 516, 523
- Brocade 7800, XISL restriction, 320
- Brocade adapters, configuring F_Port trunking for, 581
- Brocade adapters, F_Port trunking for, 581
- Brocade configuration setup form, 287
- Brocade DCX, 519, 543, 547
 - auto-leveling, 290
 - ICLs, 546
- Brocade DCX 8510, 518, 543
 - auto-leveling, 290
 - ICLs, 544
- Brocade DCX 8510-4, 519
- Brocade DCX 8510-8, 519
- Brocade DCX-4S, 547
- Brocade FC16-48 port blade enabling exceptions, 99
- Brocade FC8-48 port blade enabling exceptions, 99
- Brocade FC8-48E port blade enabling exceptions, 99
- Brocade FC8-64 port blade enabling exceptions, 99
- Brocade fixed-port switches, upgrading firmware, 295
- Brocade FX8-24
 - compatibility, 98
 - enabling 10-GbE ports, 529
 - XISL use and VE_Ports, 321
- Brocade MIB files, 213
- Brocade Network Advisor, 57
- Brocade Vendor-Specific Attribute. *See*: VSA.
- browser
 - 128-bit encryption, 200
 - displaying encryption support, 200
 - root certificates in Firefox, 205
 - root certificates in Internet Explorer, 205
 - security certificate configuration, 204
- browser support, 200
- buffer credit
 - management, 135
- buffer credit recovery, 146
- buffer credits, 119
 - allocating, 137–142
 - for average-size frames, 140
 - for F_Ports, 142
 - for full-size frames, 137
 - by switch model, 143
- buffer-to-buffer credits, 119, 135

C

- capitalization in commands, 58
- certificate signing request. *See*: CSR.
- certificates

- browser, configuring, 204
- certificate authorities (CA), 201
- FCAP, 244
 - importing for FCAP, 252
 - installing on switch, 203
 - installing root certificate for Java plugin, 205
- IPsec and, 271
- IPSECCA.pem certificate name, 271
- obtaining, 203
 - private key generation, 201
 - public key generation, 201
- root, 203
- root, configuring, 205
- security, 196
 - signing requests, generating and storing, 202
- SSL, 196, 200, 251
 - SSL certificate files, 203
- switch, 203, 251
 - See also*: root certificates or security certificates.
- cfgActvShow command, 361, 367
- cfgAdd command, 363, 425
- cfgClear command, 367
- cfgCreate command, 363
- cfgDelete command, 365
- cfgDisable command, 365
- cfgEnable command, 364, 382, 406
- cfgRemove command, 364, 426
- cfgSave command, 349
- cfgShow command, 133, 350, 356, 366, 367
- cfgSize command, 362
- cfgTransAbort command, 366
- cfgTransShow command, 377
- Challenge Handshake Authentication Protocol. *See*: CHAP.
- changing
 - authentication server configuration, 193
 - authentication server contact order, 193
 - chassis name, 77
 - logical switch to base switch, 331
 - passwords, 63
- CHAP, 646
 - alternatives, 177
 - password encryption requirement, 177
 - See also*: DH-CHAP.
- chargen listener application, 228
- chassis management IP interface, setting, 68
- chassis names, 77
- chassis, changing name of, 77
- chassisDistribute command, 260, 261
- chassisName command, 77
- ChassisRole

- Microsoft Active Directory, 183
- OpenLDAP, 188
- RADIUS, 173
- TACACS+, 188
- chassisShow command, 104
- CIDR block notation, 67
- class 2 and 3 traffic support, 115
- classConfig command, 153
- classless inter-domain routing. See: CIDR.
- clearing performance monitor counters, 557
- clearing zone configurations, 367
- CLI
 - capitalization in, 58
 - command history, 61
 - commands to display switch configuration, 281
 - commands to modify switch configuration, 281
 - Fabric OS, 58–61
- cliHistory command, 61
- command
 - getting help on, 60
 - aaaConfig, 169, 171, 189, 193, 194, 652
 - ad, 494, 498, 499, 500, 501, 502, 503, 506, 507, 508
 - aliCreate, 347
 - aliDelete, 349
 - aliRemove, 348
 - aliShow, 349
 - aptPolicy, 122, 125
 - auditCfg, 109
 - auditDump, 109
 - authUtil, 245, 246, 247, 248, 253, 453, 653
 - bladeCfgGeMode, 529
 - bladeDisable, 99
 - bladeEnable, 99
 - bladeSwap, 99
 - bottleneckMon, 428, 431, 433, 435, 436, 440, 442
 - capitalization in, 58
 - cfgActvShow, 361, 367
 - cfgAdd, 363, 425
 - cfgClear, 367
 - cfgCreate, 363
 - cfgDelete, 365
 - cfgDisable, 365
 - cfgEnable, 364, 382, 406
 - cfgRemove, 364, 426
 - cfgSave, 349
 - cfgShow, 133, 350, 356, 366, 367
 - cfgSize, 362
 - cfgTransAbort, 366
 - cfgTransShow, 377
 - chassisDistribute, 260, 261
 - chassisName, 77
 - chassisShow, 104
 - classConfig, 153
 - cliHistory, 61
 - configDefault, 284, 512
 - configDownload, 280, 283, 285, 370, 531
 - restrictions, 281
 - Virtual Fabrics mode restrictions, 286
 - configShow, 277
 - configUpload, 277, 279, 284, 285, 292, 370, 531
 - Virtual Fabrics mode restrictions, 286
 - configure, 111, 197, 301, 323, 334, 654
 - configureChassis, 326, 418
 - date, 72
 - defZone, 361
 - distribute, 152, 238, 262, 263, 264
 - dlsReset, 126
 - dlsSet, 126, 131
 - dlsShow, 126
 - dnsConfig, 649
 - errClear, 293
 - fabricName, 78
 - fabricShow, 75, 105
 - described, 307
 - fanShow, 104
 - fcrConfigure, 605, 606, 619
 - fcrXlateConfig, 602
 - fddCfg, 238, 260, 262, 263, 264, 605
 - fipsCfg, 647, 653, 655
 - firmwareCommit, 296, 303, 304, 305
 - firmwareDownload, 290, 291, 293, 294, 295, 296, 297, 299, 300, 301, 302
 - firmwareDownloadStatus, 294, 307
 - firmwareKeyUpdate, 300, 301
 - firmwareRestore, 303, 304
 - firmwareShow, 292, 295, 297, 302, 303
 - described, 307
 - fmMonitor, 558, 559, 560, 561, 562
 - for advanced zoning, 338
 - fosConfig, 324, 325, 606
 - fosExec, 328
 - frameLog, 128
 - haDisable, 164
 - haFailover, 165, 306
 - haShow, 104, 296, 297, 305
 - haSyncStart, 297
 - help, 60
 - ifModeSet, 93
 - iodReset, 127
 - iodSet, 127
 - iodShow, 127
 - ipAddrSet, 68, 69, 70, 259, 333

ipAddrShow, 66, 69, 70
 ipFilter, 227, 228, 254, 255, 259, 653
 ipSecConfig, 267, 269, 271, 272, 274, 654
 islShow, 452, 573
 keyTool, 206
 killTelnet, 59
 ldapAdd, 189
 ldapCfg, 171, 181, 182, 183, 184, 186
 licenseAdd, 528, 529, 533
 licenseIdShow, 40
 licensePort, 538, 539, 540, 541
 licenseRemove, 534
 licenseShow, 528, 529, 531, 532, 533, 534, 604
 licenseSlotCfg, 528, 529, 531
 lsCfg, 325, 326, 329, 330, 331, 332
 msCapabilityShow, 47
 msConfigure, 48, 49
 msPlatShow, 47, 50
 msPIClearDb, 51
 mspIMgmtActivate, 46, 47
 mspIMgmtDeactivate, 46, 47
 mstdDisable, 52
 mstdEnable, 51
 mstdReadConfig, 51
 nsAllShow, 54, 105
 described, 307
 nsShow, 54, 105
 described, 307
 passwd, 63
 passwdCfg, 159, 161
 passwdDefault, 646
 perfAddEeMonitor, 554
 perfCfgClear, 568
 perfCfgRestore, 568
 perfCfgSave, 568
 perfMonitorClear, 557
 perfMonitorShow, 556, 557
 perfSetPortEEMask, 555
 perfTTmon, 565, 566, 567
 portBufferCalc, 450
 portBufferShow, 142, 450
 portCfg, 654
 portCfgCompress, 456, 457
 portCfgEncrypt, 455, 457, 654
 portCfgExport, 603
 portCfgFillWord, 90, 589
 portCfgISLMode, 119, 121
 portCfgLongDistance, 113, 589
 portCfgNpivPort, 475, 476
 portCfgOctetSpeedCombo, 95, 527, 528
 portCfgPersistentEnable, 92
 portCfgQos, 415, 425, 426
 portCfgShow, 476
 portCfgSpeed, 94, 448, 527, 528
 portCfgTrunkPort, 574, 581
 portDecom, 93
 portDisable, 92, 574
 portEnable, 92, 537
 portEncCompShow, 448, 450, 452
 portLoginShow, 478
 portName, 89
 portShow, 477, 537
 portStatsShow, 450
 portSwap, 84, 91
 portSwapEnable, 91
 portTrunkArea, 574, 576, 581
 portZoneShow, 342
 powerOffListSet, 103
 powerOffListShow, 103
 psShow, 104
 roleConfig, 154
 secAuthSecret, 249, 250, 453
 secCertUtil, 201, 202, 203, 252, 253, 271, 650
 secModeEnable, 229
 secPolicyAbort, 234
 secPolicyActivate, 233, 236, 237, 264
 secPolicyAdd, 234
 secPolicyCreate, 236, 240, 243
 secPolicyDelete, 233, 239, 240
 secPolicyFCSMove, 237
 secPolicyRemove, 234
 secPolicySave, 240
 secPolicyShow, 233
 setContext, 125, 333, 334
 slotPowerOff, 103
 slotPowerOn, 103
 slotShow, 104, 604
 snmpConfig, 653
 snmpWalk, 216, 217
 ssh-keygen, 198
 sshUtil, 198, 200, 652
 sshutil, 291
 supportSave, 40
 switchCfgPersistentDisable, 102
 switchCfgSpeed, 94
 switchCfgTrunk, 574
 switchDisable, 79, 111, 125, 541
 switchEnable, 79, 80, 111, 336
 switchName, 77
 switchShow, 90, 104, 105, 334, 336, 452, 473, 477,
 538, 541
 switchShow, 641
 switchStatusPolicySet, 106
 switchStatusPolicyShow, 106

- switchStatusShow, 104
- syslogDlpAdd, 109
- sysShutdown, 80, 81
- tac_plus, 190
- topologyShow, 402
- trunkShow, 574
- tsClockServer, 74
- tsTimeZone, 72, 73
- usbStorage, 299
- userConfig, 155, 497, 498, 654
- version, 604
- wwn, 40
- wwnAddress, 87
- zone, 132, 133, 369, 395, 402, 405, 406, 407, 409
- zoneAdd, 351
- zoneCreate, 350, 424
- zoneDelete, 355
- zoneHelp, 338
- zoneObjectRename, 370
- zoneObjectReplace, 353
- zoneRemove, 352
- zoneShow, 356
- command line interface. See: CLI.
- commands
 - errShow, 210
 - switchStatusPolicySet, 210
- compression
 - configuring, 456
 - disabling, 457
 - enabling, 448, 456
 - enabling on a port, 450, 451
 - in-flight, 445
 - license, 445
 - long distance ports, 450
 - ratios for enc/comp enabled ports, 450
 - restrictions, 446
 - viewing configuration, 452
- concurrent zone transactions, 376
- conditional tests for FIPS, 647
- configDefault command, 284, 512
- configDownload command, 280, 283, 285, 370, 531
 - restrictions, 281
 - Virtual Fabrics mode restrictions, 286
- configShow command, 277
- configUpload command, 277, 279, 284, 285, 292, 370, 531
 - in Admin Domain context, 512
 - Virtual Fabrics mode restrictions, 286
- configuration
 - configDownload command, 280
 - configuration management for Virtual Fabrics, 285
 - FA-PWWN upload and download considerations, 483
 - format of configuration file, 278
 - in fabrics, 284
 - modifying for switches, 281
 - restoring, 282
 - saving for frame monitors, 560
 - security considerations, 284
 - setup form, 287
 - supported for FA-PWWN, 483
 - without disabling a switch, 282
 - zones, 370
- configuration file
 - backing up, 279
 - backup, 279
 - chassis section, 278
 - configDownload command, in Admin Domain context, 512
 - display settings, 277
 - downloading, 512
 - format, 278
 - information not saved, 279
 - restoring, 280
 - save to a host, 277
 - sharing between identical switch models, 284
 - switch section, 278
 - uploading, 512
 - in interactive mode, 279
 - with Virtual Fabrics enabled, 285
- configuration settings, 277–287
- configure command, 111, 197, 301, 323, 334, 654
- configureChassis command, 326, 418
- configuring
 - a switch for signed firmware, 301
 - access methods, Web Tools, 57
 - audit log, 107
 - authentication, 453
 - authentication policy, 243–253
 - browser security certificates, 204
 - compression, 456
 - date and time, 72
 - device authentication, 247
 - device-switch connection, 90
 - DHCP, 69
 - Enforce LSAN tag, 627
 - extended ISLs, 589
 - F_Port trunking on an Access Gateway, 580
 - FA-PWWNs, 481–482
 - FCAP, 251
 - FIPS mode, 647–651
 - FLOGI-time handling of duplicate PWWNs, 110
 - HTTPS access, 200
 - incoming SSH authentication, 198
 - in-flight encryption, 455

- interfabric link, 607
- IPv6 automatically, 71
- links through a gateway, 121
- lossless DLS, 131
- NTP, 74
- outgoing SSH authentication, 199
- remote authentication, 167–170
- remote authentication on switch, 192
- root certificates, 205
- security certificates, 200
- Speed LSA tag, 627
- SSL, 200, 201–205
- TACACS+ service, 189
- zone, rules for, 342
- conflicting ACL policies, resolving, 264
- congestion bottleneck type, 428
- congestions versus over-subscription, 119
- connected devices and logical switches, 313
- connecting
 - device to a switch, 90
 - multiple EX_Ports to an edge fabric, 603
 - switches running different firmware versions, 81
 - to devices, 82
 - to switch, 82
- connection
 - restrictions, 154
 - ssh, 59
 - telnet, 59
- consistency policies, matching fabric-wide, 265
- consistency policies, non-matching fabric-wide, 265
- console session on serial port, 58
- control processor. *See:* CP.
- converting hexadecimal numbers, 657–659
- core blades, 97
- core/edge topology and ISL trunking, 573
- core-edge topology, 549
- CP blades, 97
 - accessing, 175
 - licensed features and, 533
 - standby, 175
 - swapping, 533
- CP8 blade
 - devices supporting dual port, 66
- creating
 - Admin Domains, 495
 - alias, 347
 - base switches, 326
 - DCC policies, 239
 - FCS policies, 236
 - frame monitors, 559
 - frame redirect zones, 132

- frame types to be monitored, 559
- logical switches, 326
- SCC policies, 243
- TI zones, 402
- user-defined roles, 154
- zone configurations, 362, 363
- zones, 350
- CS_CTL auto mode, using at the chassis level, 418
- CS_CTL-based frame prioritization, 416–418
 - considerations, 418
 - default and auto modes, 417
 - disabling, 418
 - enabling, 417
 - high availability considerations, 417
 - supported configurations, 417
- CSR
 - defined, 650
 - exporting for FCAP, 252
 - generating and storing, 202
 - generating for FCAP, 251
 - obtaining certificates, 203
- customizing the switch name, 75
- cut-through routing, 117

D

- D_Port (diagnostic port), 459
 - saving port mappings, 464
 - topologies, 463
 - viewing information, 465
 - with HBAs, 467
 - without HBAs, 465
- D_Port, described, 88
- daemon processes and High Availability, 55
- daemon, tac_plus, 190
- daemons automatically restarted, 55
- date and time, 72
- date change license restriction, 531
- date command, 72
- date settings, 72
- daytime listener application, 228
- DCC
 - creating policy, 239
 - deleting policy, 240
 - policies, 232, 238–242
 - policy member, 232
 - policy restrictions, 239
 - policy, maximum name length, 239
- DCC policies

- for NPIV ports, 241
 - policy behavior with fabric-assigned PWWNs, 241
 - Virtual Fabric considerations, 239
- deactivating
 - Admin Domains, 499
 - TI zones, 406
- decimal to hexadecimal conversion table, 658
- decommissioning ports, 92
- default
 - account passwords, 63
 - accounts, listed, 63
 - Fabric OS roles, 152
 - IP Filter policy names, 254
 - IP Policy Rules, 258
 - logical switch, 310
 - zone access mode, viewing current, 361
 - zone mode, 360, 495
 - zoning mode, setting, 361
- default logical switch
 - base switch restriction, 321
 - XISL restriction, 321
- defZone command, 361
- deleting
 - a Top Talker monitor on a port, 567
 - ACL policy, 233
 - Admin Domains, 501, 502
 - alias, 349
 - all fabric mode Top Talker monitors, 567
 - DCC policy, 240
 - end-to-end monitors, 556
 - frame monitors, 560
 - frame redirect zones, 133
 - IP Filter policy, 255
 - LDAP certificates, 651
 - logical switches, 329
 - private key from switch, 200
 - public key from switch, 200
 - rule from an IP Filter policy, 259
 - TI zones, 407
 - zone configurations, 365
 - zones, 355
- delivery order, forcing for frames, 127
- deploying secure protocols, 196
- device
 - accessing, 229
 - configuring authentication, 247
 - connecting, 82
 - CP8 blade dual port support, 66
 - limiting traffic from, 415
 - login, 53–55
 - proxy devices, 599
 - recovery, 55
 - verifying connectivity, 105
- device authentication policy, 246
 - and Virtual Fabrics considerations, 247
- Device Connection Control. See: DCC.
- device-based routing, 122, 123, 126
 - FICON environments only., 123
- DH-CHAP, 248, 449
 - and F_Ports, 247
 - and FIPS, 647
 - authentication on E_Ports, 244
 - configuring authentication, 243
 - secret key pairs, 249
 - specifying as authentication protocol, 248
 - See also: CHAP.
- DHCP
 - activation, 69
 - configuration, 69
 - disabling, 70
 - enabling, 69
 - stateful IPv6 addresses, 69
- diagnostic port. See: D_Port.
- dictionary.brocade, 173
- Diffie Hellman-Challenge Handshake Authentication Protocol. See: DH-CHAP.
- Directory server, described, 45
- disabled zone configuration, defined, 341
- disabling
 - admin lockout policy, 162
 - bottleneck detection, 442
 - compression, 457
 - CS_CTL-based frame prioritization, 418
 - DHCP, 70
 - F_Port trunking, 585
 - failover in TI zones, considerations, 381
 - in-flight encryption, 456
 - ingress rate limiting, 415
 - ISL trunking, 574
 - local switch protection, 262
 - NPIV, 476
 - port, 92
 - QoS zone-based traffic prioritization, 426
 - remote authentication, 193
 - switches, 79, 102
 - topology discovery, 52
 - Virtual Fabrics, 325
 - zone configurations, 365
- disabling MIBs, 210
- discard listener application, 228
- displaying
 - Admin Domain configuration, 508

- configuration settings, 277
- current routing policy, 122
- domain IDs, 75
- encryption support in browser, 200
- existing zones, 350
- F_Port trunking information, 585
- frame monitors, 561
- logical switch configuration, 330
- LSAN tags, 628
- monitor counters, 557
- network interface settings, 66
- port license assignments, 538
- switch name server contents, 54
- TI zones, 407
- trunking information, 574
- distance vector, in routing, 115
- distribute command, 152, 238, 262, 263, 264
- Distributed Management Server
 - FCS policy, 47
 - management server database, 47
 - topology discovery, 51
 - well-known address, 46, 47
- distributing
 - authorization policy fabric-wide, 253
 - FCS policies, 238
 - IP Filter policy, 260
 - local ACL policies, 263
 - local user account database, 158
- distribution policy states, 238
- DLS
 - computation trigger, 125
 - effect on other logical switches, 131
 - overview, 125
 - rebalancing triggers, 129
 - See also: Dynamic Load Sharing.
- dlsReset command, 126
- dlsSet command, 126, 131
- dlsShow command, 126
- dnsConfig command, 649
- domain ID
 - and Admin Domain number, 486
- domain IDs, 75–76
 - conflicts, 326
 - displaying, 75
 - displaying top talking flows for, 566
 - domain ID 0, 75
 - setting, 76
- downgrading firmware, 291
- download configuration file, 512
- DPS
 - described, 123
 - device-based routing, 123

- support on Virtual Fabrics, 124
- dropped frames, discovering why, 127
- DSA key pair generation, 198
- duplicate
 - F_Port login, 111
 - NPIV port login, 111
 - Port World Wide Name (PWWN), 55
- duplicate PWWNs, handling, 110
- Dynamic Load Sharing. See: DLS.
- Dynamic Path Selection. See: DPS.
- dynamic PID binding, 83

E

- E_Port
 - authentication, 244
 - authentication using DH-CHAP or FCAP, 244
 - configuring authentication, 245
 - described, 88
 - inter-switch link, 118
 - login process, 53
 - re-authenticating, 246
 - Virtual Fabric considerations, 245
- echo listener application, 228
- edge fabric
 - described, 596
 - TI zones and, 388
- edge-to-edge routing, 605
- EE monitors
 - about, 553
 - adding, 554
 - clearing statistic counters, 557
 - defined, 551
 - deleting, 556
 - displaying counters, 557
 - maximum number, 553
 - setting a mask for, 555
 - supported port configurations for, 554
- effective AD configuration, 494
- effective zone configuration, defined, 341
- ELP mode, 121
- enabling
 - 10 Gbps operation on an FC port, 528
 - 10-GbE ports on an FX8-24 blade, 529
 - admin lockout policy, 162
 - authentication, 453
 - bottleneck detection, 431
 - compression, 456
 - CS_CTL-based frame prioritization, 417
 - DCC policy on a trunk area, 586
 - DHCP, 69

- FIPS mode, 652
- FIPS mode, permissions needed, 652
- ISL trunking, 574
- local switch protection, 262
- NPIV, 476
- port, 92
- remote authentication, 193
- switches, 79, 80
- Virtual Fabrics mode, 324
- zone configurations, 364
- Encapsulating Security Payload. *See: ESP.*
- encryption
 - 128-bit in browser, 200
 - algorithms used, 269
 - configuring in-flight, 455
 - disabling in-flight, 456
 - enabling, 448
 - enabling on a port, 450, 451
 - in-flight, 445
 - key generation, 448
 - license, 445
 - restrictions, 446
 - using SSL, 200
 - viewing configuration, 452
- end-to-end (EE) monitoring, 553
- end-to-end monitors
 - deleting, 556
 - restoring configuration, 567
 - saving configuration, 567
 - setting a mask, 555
- end-to-end performance monitoring, 553
- end-to-end transport tunnel mode, example, 274
- enforce LSAN tag, 625
- enforcement of zones, 342
- Enhanced TI zones. *See: ETIZ.*
- ensuring fabric domains share policies, 236
- enterprise ICL license, 524
- equipment status, 104
- errClear command, 293
- errShow command, 210
- ESP, described, 269
- eth0 port on CP8 blade, 65
- eth3 port on CP8 blade, 65
- ethernet address, static, 67
- ethernet interface
 - on switch, 64
 - Virtual Fabrics, 65
- ethernet IP address, setting static, 68
- ETIZ
 - configuration rules, 396
 - defined, 384
 - platform restrictions and FC router domains, 397
 - See also: Enhanced TI zones.*
- events, logging date and time, 72
- EX_Port, 635
 - configuring trunking, 578
 - described, 88, 596
 - displaying trunking information, 578
 - FCR authentication, 603
 - masterless trunking, 577
 - supported trunking configurations and platforms, 578
- EX_Port trunking, 577–579
 - configuring, 578
 - displaying information, 578
 - masterless, 577
 - supported configurations and platforms, 578
- Exchange Link Parameters mode. *See: ELP mode.*
- exchange-based routing, 122, 123, 126
- expired licenses, 531
 - removing, 532
- exporting
 - CSR for FCAP, 252
 - LDAP certificates, 651
 - public key from switch, 199
- extended fabrics
 - about, 588
 - buffer credit management, 135
 - buffer credit recovery, 146
 - buffer requirement calculation, 137
 - buffer-to-buffer credits, 135
 - extended ISLs, 589
 - F_Port buffer credits, 142
 - ISL, 136
 - long-distance mode, 138
 - port buffer credit, 137
 - QoS buffer credit requirements, 145
 - time-division multiplexing, 590
- extended ISL
 - about, 317
 - and base switches, 316
 - and fmsmode, 323
 - logical fabric creation, 334
 - restrictions, 323
 - See also: XISL.*
- extending a universal temporary license, 532

F

- F_Port
 - configuring trunking for Brocade adapters, 581

- configuring trunking on an Access Gateway, 580
- described, 88
- DH-CHAP protocol failure, 247
- disabling trunking, 585
- displaying trunking information, 585
- duplicate logins, 111
- trunking considerations, 582
- trunking for access gateways, 579
- trunking for Brocade adapters, 581
- trunking requirements on an Access Gateway, 580
- F_Port trunking, 579–586
 - and Virtual Fabrics, 584
 - configuring for Brocade adapters, 581
 - considerations, 582
 - for access gateways, 579
 - for Brocade adapters, 581
- fabric
 - access, 229
 - adding Top Talker monitors, 565
 - addresses. *See:* PID.
 - authentication availability, 243
 - authentication license, 243
 - authentication policies, 243–253
 - changing name, 78
 - configurations in, 284
 - connectivity, 105
 - deleting all Top Talker monitors, 567
 - domain policy sharing, 236
 - element authentication, 243
 - fabric login. *See:* FLOGI.
 - fabric-assigned PWWN, 479
 - fabric-assigned PWWNs and DCC policy behavior, 241
 - fabric-wide consistency policy
 - distribution, 260
 - setting, 264
 - viewing, 263
 - fabric-wide distribution of authorization policy, 253
 - fabric-wide policy enforcement, 263
 - joining switches to, 264
 - login command, 53
 - login process, 54
 - matching fabric-wide consistency policies, 265
 - name, 77
 - high availability issues, 78
 - upgrade and downgrade considerations, 78
 - non-matching fabric-wide consistency policies, 265
 - optimizing behavior, 413–426
 - parameters for ISLs, 118
 - segmentation and zoning, 373
- fabric administrator user account, 497
- Fabric configuration server. *See:* FCS.
- Fabric controller, described, 45
- fabric ID
 - assigning for Backbones, 606, 617
 - conflicts, 326
 - described, 598
 - logical switches and, 311
 - See also:* FID.
- Fabric Login server, described, 45
- Fabric mode Top Talker monitor, described, 563
- Fabric OS
 - authentication services, external, 167
 - command line interface, 58, 58–61
 - default roles, 152
 - feature interaction with Virtual Fabrics, 322
 - interaction with Virtual Fabrics, 322
 - policies, 232
 - protocols supported, 196
 - security protocols supported, 195
 - user accounts, 171–172
 - on RADIUS servers, 172–181
 - user accounts through LDAP, 171
 - web server, 204
- Fabric Shortest Path First. *See:* FSPF.
- fabricName command, 78
- fabricShow command, 75, 105
 - described, 307
- fabric-wide consistency policy, 232, 605
- failover, for traffic isolation, 380–382
- fanShow command, 104
- FA-PWWN, 479
 - auto-assigned behavior, 480
 - configuration upload and download considerations, 483
 - configuring, 481–482
 - DCC policy behavior, 241
 - N_Port Access Gateway failover, 484
 - priority, 480
 - restrictions, 484
 - security considerations, 483
 - supported switches and configurations, 483
 - user-assigned behavior, 480
- FC routers
 - setting QoS zone-based traffic prioritization over, 426
- FCAP
 - and FIPS, 647
 - authentication on E_Ports, 244
 - certificates, 244
 - configuration overview, 251
 - CSR export, 252
 - fabric element authentication, 243
 - generating key and CSR, 251
 - importing security certificate, 252

- importing switch certificate, 252
- PKI certificates required, 243
- specifying as authentication protocol, 248
- starting authentication, 253
- FC-FC routing
 - and FCIP, 606
 - and Virtual Fabrics, 636
 - backbone-to-edge, 600
 - configurations supported, 595
 - edge-to-edge, 600
 - fabric mode Top Talker monitors, 605
 - license requirements, 594
 - platforms supported, 594
 - routing service, 593
 - setup, 603–605
 - setup verification, 604
 - Top Talker monitors and, 563
 - topologies, 600
 - See also: FCR and Fibre Channel routing
- FCIP
 - and FC-FC routing, 606
 - tunnel configuration, 606
 - tunnel hop support, 321
- FC-NAT, defined, 117
- FCoE, NPIV required, 476
- FCR
 - and traffic isolation, 386
 - authentication, 603
 - Brocade 7800 logical switches, 594
- fcrConfigure command, 605, 606, 619
- fcrXlateConfig command, 602
- FCS
 - creating policy, 236
 - modifying policy, 235
 - modifying switch order, 237
 - policies, 232, 235, 235–238
 - policy distribution, 238
 - policy member, 232
 - policy restrictions, 235
 - policy states, 235
- fddCfg command, 238, 260, 262, 263, 264, 605
- feature licenses, 515
 - listed, 516
- FEC
 - configuration requirements, 112
 - disabling, 113
 - disabling for long distance ports, 113
 - enabling, 111, 112
 - enabling for long distance ports, 113
 - limitations, 112
 - viewing settings, 113
- Federal Information Processing Standards. See: FIPS.
- Fibre Channel
 - NAT, 117
 - See also: FC.
- Fibre Channel Authentication Protocol. See: FCAP.
- Fibre Channel Common Transport (FC-CT) protocol service, described, 46
- Fibre Channel fabrics, and port ID, 117
- Fibre Channel Over IP service. See: FCIP.
- Fibre Channel port, 88
- Fibre Channel port, enabling 10 Gbps operation, 528
- Fibre Channel protocol auto discovery process, 54
- Fibre Channel router, described, 596
- Fibre Channel routers
 - domains and ETIZ platform restrictions, 397
 - SCC policies, 242
 - TI zone limitations, 390
- Fibre Channel routing, 596
 - concepts, 596–603
 - See also: FCR and FC-FC routing.
- Fibre Channel services, 45–56
 - overview, 45
- FICON CUP environment considerations, 291
- FID. See: fabric ID.
- FIPS
 - conditional tests, 647
 - described, 300
 - DH-CHAP, 647
 - displaying configuration, 655
 - error state, 647
 - FCAP, 647
 - firmware considerations, 300
 - firmwareDownload command, 301
 - LDAP certificates
 - deleting, 651
 - displaying and deleting, 650
 - exporting, 651
 - installing, 650
 - overview, 645
 - permissions needed to enable, 652
 - power-on self tests, 647
 - preparing a switch, 651
 - RADIUS server configuration, 653
 - RPC unsupported, 647
 - support, 645–655
 - zeroization functions, 645
 - zeroizing for, 655
- FIPS mode
 - configuration, 647–651
 - enabling, 652
 - LDAP, 648

- LDAP certificates, 650
- restrictions, 647
- fipsCfg command, 647, 653, 655
- Firefox
 - root certificate installation and verification, 205
 - SSL support, 200
- firmware, 289–307
 - Backbone, 296–299
 - Backbone download process overview, 296
 - Backbone version testing, 304
 - downgrading, 291
 - download process, 289
 - downloading without a password, 291
 - finding version, 293
 - for switches, 294–295
 - obtaining and decompressing, 293
 - power-on checksum test for FIPS, 302
 - signed, 301
 - switch version testing, 302
 - upgrading, 291
 - upgrading for Brocade fixed-port switches, 295
 - upgrading on Backbones, 297
 - upgrading on blades, 297
- firmware download, 290
 - auto-leveling, 304
 - Backbones, 296
 - connected switches, 293
 - FICON CUP considerations, 291
 - FIPS, 300
 - high availability synchronization, 291
 - preparing for download, 292
 - process overview, 294
 - protocol, FTP and SCP, 290
 - switches, 294
 - test and restore on Backbones, 304
 - test and restore on switches, 302
 - testing different firmware versions, 304
 - USB device, 299–300
 - validating, 306
 - verify progress, 290
- firmwareCommit command, 296, 303, 304, 305
- firmwareDownload command, 290, 291, 293, 294, 295, 296, 297, 299, 300, 301, 302
- firmwareDownloadStatus command, 294, 307
- firmwareKeyUpdate command, 300, 301
- firmwareRestore command, 303, 304
- firmwareShow command, 292, 295, 297, 302, 303
 - described, 307
- fixed addressing mode, 84
- fixed-port switches
 - port configurations supported, 320
 - port restrictions, 320
- FL_Port, described, 88
- FLOGI
 - defined, 53
 - FC-SP bit setting, 247
 - process, 54
 - rejected, 247
 - request frame header value, 54
- fmMonitor command, 558, 559, 560, 561, 562
- fmsmode, and XISL, 323
- forcing frame delivery order, 127
- forward error correction. See: FEC.
- fosConfig command, 324, 325, 606
- fosExec command, 328
- frame delivery and topology changes, 127
- frame monitoring, 558
- frame monitors
 - adding to a port, 560
 - clearing statistic counters, 562
 - creating, 559
 - defined, 552
 - deleting, 560
 - displaying, 561
 - removing from a port, 560
 - restoring configuration, 567
 - saving configuration, 560, 567
- frame order, specifying delivery, 126
- frame prioritization, CS_CTL-based, 416–418
- frame redirection, 132
- frame redirection zones, 337
 - creating, 132
 - deleting, 133
 - viewing, 133
- frame types
 - creating to be monitored, 559
 - deleting, 560
- Frame Viewer
 - filtering results, 128
 - using, 127
- frameLog command, 128
- frames
 - creating frame redirect zones, 132
 - deleting frame redirect zones, 133
 - discovering why dropped, 127
 - forcing delivery order, 127
 - restoring unordered delivery order, 127
 - viewing frame redirect zones, 133
- FreeRADIUS
 - clients, enabling, 176
 - configuring, 175
 - Fabric OS user setup, 173

- user, adding, 176
- vendor attributes, 175
- See also: RADIUS and Linux.

FSPF

- described, 116
- number of routes supported, 116
- path calculation, 117
- traffic isolation routing rules, 383

FSPF-1009 RASLOG message, 398

ftp listener application, 228

G

G_Port, described, 88

gateway links, 120

- buffer credits, 588

gateway, configuring a link through, 121

generating

- DSA or RSA key pairs, 198

- key and CSR for FCAP, 251

- PKI key pairs, 199

H

HA. See: High Availability.

haDisable command, 164

haFailover command, 165, 306

haShow command, 104, 296, 297, 305

haSyncStart command, 297

help command, 60

help for commands, 60

hexadecimal conversion, 657–659

hexadecimal to decimal conversion table, 658

High Availability

- daemon processes, 55

- failover and passwords, 159

- failover on RADIUS server, 175

- QoS zone-based traffic prioritization considerations, 422

- support for trunking, 572

- synchronization, 291

- verifying features, 104

history of CLI commands, 61

home Admin Domain, 490

- Microsoft Access Directory, 183

- OpenLDAP, 188

- RADIUS, 174

- TACACS+, 191

home LF

- Microsoft Active Directory, 183

- OpenLDAP, 188

- RADIUS, 174

- TACACS+, 191

host syslog, verifying, 109

hosts, accessing, 229

HTTPS protocol, 200

- described, 195

- secure protocol, 196

I

IAS

- configuring, 178

- remote access policies, 178

ICL

- 16-link license, 524

- 1st POD license, 523

- 2nd POD license, 524

- 8-link license, 524

- about Inter-Chassis Links, 543

- core-edge topology, 549

- enterprise license, 524

- for DCX 8510 family, 544

- for DCX family, 546

- license, 544

- licensing, 523–524

- limitations for lossless core, 130

- maximum numbers of, 545

- mesh topology, 547

- prohibited connections, 546

- topologies supported, 547–550

- triangular topology, 547

- trunking on DCX and DCX-4S, 547

- Virtual Fabric considerations, 547

ICL ports and XISL, 321

identifying ports, 89–90

IE, root certificate installation and verification, 205

IFL

- about, 596

- configuration, 607

- configuring, 607

- described, 596

ifModeSet command, 93

IKE

- policies and IPsec, 271

- policies, null encryption support, 275

implementing Admin Domains, 495

indexing ports, 641–644

in-flight compression and port decommissioning, 448

- in-flight encryption
 - configuring, 455
 - disabling, 456
 - license, 445
 - port decommissioning, 448
 - restrictions, 446
- in-flight encryption and compression, 445
 - overview, 445
- ingress rate limiting, 414–415
 - disabling, 415
 - Virtual Fabrics considerations, 414
- in-order frame delivery, forcing, 127
- installing
 - certificates on switch, 203
 - LDAP certificates, 650
 - root certificate to Java plugin, 205
- Integrated Routing license, 594
- Inter-Chassis Links. See: ICL.
- inter-fabric link See: IFL.
- Internet Explorer and SSL support, 200
- Internet Explorer. See: IE.
- inter-switch link. See: ISL.
- iodReset command, 127
- iodSet command, 127
- iodShow command, 127
- IP addresses
 - configuring for a Virtual Fabric, 333
 - removing from a Virtual Fabric, 333
- IP Filter
 - aborting transaction, 259
 - activating policy, 255
 - adding rule to policy, 259
 - cloning policy, 254
 - creating policy, 254
 - default policy names, 254
 - default policy rules, 258
 - deleting policy, 255
 - deleting rule from policy, 259
 - displaying policy, 254
 - implicit rules, 258
 - IP Filter policy destination port, 256
 - IP Filter policy source address, 256
 - IP Filter policy traffic type and destination IP, 258
 - policy, 253
 - policy and Virtual Fabrics considerations, 254
 - policy distribution, 260
 - policy enforcement, 258
 - policy rules, 255
 - policy rules using service names, 256
 - saving policy, 255
 - supported actions, 257
 - supported protocols, 257
 - supported services and port numbers, 256
- IP interface for chassis management, 68
- ipAddrSet command, 68, 69, 70, 259, 333
- ipAddrShow command, 66, 69, 70
- ipFilter command, 227, 228, 254, 255, 259, 653
- IP-NAT, 117
- IPsec
 - algorithms, 269
 - Authentication Header protocol, 269
 - configuration on the management interface, 266
 - Encapsulating Security Payload protocol, 269
 - flushing security associations, 275
 - IKE policies, 271
 - key management, 271
 - manual key entry, 272
 - null encryption support for IKE policies, 275
 - policies, 270–275
 - policy described, 270
 - pre-shared key, 271
 - protocol, described, 195
 - protocols, 269
 - sa-proposal, 269
 - security association, 269
 - security certificate, 271
 - traffic selector, 270
 - transform set, 270
 - tunnel configurations, 267–268
- IPSECCA.pem, certificate name, 271
- ipSecConfig command, 267, 269, 271, 272, 274, 654
- IPv4 filter policy address, 256
- IPv6
 - autoconfiguration, 71
 - DHCP and stateful IPv6 addresses, 69
 - filter policy address, 256
 - IPv6 policies tunneling ICMP traffic, 275
- ISL, 82
 - best practices, 118
 - configuring extended, 589
 - fabric parameters, 118
 - logical fabrics and, 315
 - maximum distances in LO mode, 82
- ISL R_RDY mode, 121
- ISL trunking
 - disabling, 574
 - enabling, 574
 - over long distance fabrics, 576
- islShow command, 452, 573

J

Java

- installing root certificate in plugin, 205
- installing root certificate to plugin, 205
- support for SSL, 200
- supported version, 201

Java plugin, installing root certificate for, 205

joining a switch to a fabric, 264

K

key

- adding public key to switch, 198
- deleting private from switch, 200
- deleting public from switch, 200
- generating for FCAP, 251
- generation, 201
- key management and IPsec, 271
- key pair generation for RSA or DSA, 198
- manual key entry and IPsec, 272
- PKI key pair generation on switch, 199
- pre-shared, and IPsec, 271
- private key generation, 201
- secret pairs for DH-CHAP, 249
- setting a secret key pair, 250
- viewing list of secret key pairs, 249

key generation, 448

keyTool command, 206

killTelnet command, 59

L

L_Port, described, 88

latency bottleneck type, 428

Layer 2 routing. *See also:* FSPF.

LDAP

- Active Directory LDAP versions supported, 181
- authentication, non-FIPS mode, 184
- certificates for FIPS mode, 650
- configuring for Microsoft Active Directory, 182
- creating Fabric OS user accounts, 171
- deleting certificates, 651
- exporting certificates, 651
- in FIPS mode, 648
- installing certificates, 650
- IPv4 and IPv6 support, 181
- non-FIPS mode restrictions, 181
- role mapping and OpenLDAP, 186

role mapping, and Microsoft Active Directory, 181

secure service, 168

LDAP server

- adding, 193
- deleting, 193
- reordering, 193

LDAP service

- configuration, displaying, 194
- configuring, 181
- configuring for OpenLDAP, 184–189
- disabling, 193, 194
- enabling, 193
- group assignment, 186
- groups, creating, 182
- modifying, 193
- overview, 152
- role, assigning, 183
- users, adding, 182, 185
- vendor attributes, 184

ldapAdd command, 189

ldapCfg command, 171, 181, 182, 183, 184, 186

LDAPS protocol, described, 195

Ldifde.exe utility, using, 183

LF role

- Microsoft Active Directory, 183
- OpenLDAP, 188
- RADIUS, 174
- TACACS+, 191

licence. *See:* license.

license

- 10G, 527–530
- 8G, 525
- activating, 533
- adding features, 533
- Brocade 7800 upgrade, 523
- compression, 445
- date change restriction, 531
- encryption, 445
- enterprise ICL, 524
- expired, 531
- Extended Fabrics, 587
- fabric authentication, 243
- ICL, 523–524, 544
- ICL 16-link, 524
- ICL 1st POD, 523
- ICL 2nd POD, 524
- ICL 8-link, 524
- in-flight encryption, 445
- installation requirements and location, 519
- Integrated Routing, 594
- preserving, 515
- purchasing keys, 536

- removing expired, 532
- removing features, 534
- requirements for SID/DID prioritization, 416
- requirements for trunking, 571
- reserving for POD, 540
- slot-based, 526–527
- temporary, 530–532
- time-based, 530–532
- universal temporary
 - extending, 532
 - shelf life, 532
 - universal temporary, described, 532
 - viewing installed, 532
- licenseAdd command, 528, 529, 533
- licensed features, 515
 - listed, 516
- licenseIdShow command, 40
- licensePort command, 538, 539, 540, 541
- licenseRemove command, 534
- licenseShow command, 528, 529, 531, 532, 533, 534, 604
- licenseSlotCfg command, 528, 529, 531
- licensing, 515–541
 - overview, 515–523
- limiting traffic from a device, 415
- link operating mode, 93
- link state database, 116
- link state, in routing, 115
- link, configuring through a gateway, 121
- Linux
 - FreeRADIUS and Fabric OS user setup, 173
 - LDAP authentication, 184–189
 - RADIUS server support, 175
 - TACACS+ authentication on, 190
- LISL. See: Logical ISL.
- list of port types, 88
- listener application
 - blocked
 - chargen, 228
 - daytime, 228
 - discard, 228
 - echo, 228
 - ftp, 228
 - rexec, 228
 - rlogin, 228
 - rsh, 228
 - rstats, 228
 - rusers, 228
 - time, 228
 - blocked list, 228
 - chargen, 228
 - daytime, 228
 - discard, 228
 - echo, 228
 - ftp, 228
 - rexec, 228
 - rlogin, 228
 - rsh, 228
 - rstats, 228
 - rusers, 228
 - time, 228
- loading brocade MIBs, 212
- local ACL policies, distributing, 263
- local authentication
 - backup, as, 194
 - overview, 152
- local clock, 74
- local database user accounts, 155–158
- local user account database
 - distributing, 158
- local user account database distribution, 158
- local user account passwords, 157
- LOCL, 74
- logging timestamp, 72
- logical fabrics, 315–319
 - about, 315
 - changing context, 334
 - creating using XISLs, 334
 - defined, 310
 - formation, 319
 - ISLs and, 315
- logical ISLs, 317
- logical network interface, bond0, 65
- logical ports
 - zoning, 350
- logical ports in ISL, 319
- Logical SANs, described, 597
- logical switches, 310–314
 - about, 310
 - allowing XISL use, 333
 - basic configuration values, 326
 - changing to a base switch, 331
 - commanding in a different context, 328
 - connected devices and, 313
 - creating, 326
 - deleting, 329
 - displaying configuration, 330
 - DLS effect on, 131
 - fabric IDs and, 311
 - management model, 314
 - moving ports, 329
 - multiple FIDs, 316

- number, 311
- number per chassis, 323
- port assignment, 312
- restoring configuration, 285
- Top Talkers and, 329
- unique names for, 77
- login
 - changing password, 157
 - command for fabric, 53
 - fails, 59
 - process for fabric, 54
 - with Admin Domains, 490
- login sessions, maximum allowed, 154
- long distance fabrics, and ISL trunking, 576
- long distance ports, compression, 450
- lossless core, 130
 - ICL limitations, 130
 - traffic flow limitations, 130
- lossless DLS, 129–131
 - configuring, 131
 - in Virtual Fabrics, 131
- lossless dynamic load sharing. *See: lossless DLS.*
- LSAN, 620
- LSAN tags, 624
- LSAN zone binding, 628
- LSAN zone name length consideration, 512
- LSAN zones, 337
 - in Admin Domains, 511
- IsCfg command, 325, 326, 329, 330, 331, 332

M

- M_Port, described, 88
- making basic connections, 81
- management channel, 154
- management interface
 - IPsec configuration, 266
 - security, 266–275
- management server
 - displaying ACL, 48
 - viewing database, 50
- management server database, 47–51
- Management server, described, 46
- managing
 - Admin Domains, 485–512
 - trunking connections, 569
 - user accounts, 151–194
 - user-defined roles, 154–155
 - zoning configurations in a fabric, 368
- manually distributing ACL policy database, 260

- mask for end-to-end monitors, setting, 555
- masterless EX_Port trunking, 577
- masterless trunking, 570
- matching fabric parameters, 603
- maximum ISL distances in LO mode, 82
- maximum zone database size, 362
- members
 - policy, 232
- merging zones, 362
- mesh topology, 547
- MetaSAN, described, 598
- MIB loading order, 214
- Microsoft Active Directory and LDAP role mapping, 181
- Microsoft Active Directory LDAP service
 - configuring, 182
- Microsoft Active Directory service
 - configuring for LDAP, 181
 - groups, creating, 182
 - role, assigning, 183
 - users, adding, 182
 - vendor attributes, adding to schema, 184
- mirror port. *See also: M_Port.*
- modifying
 - FCS policy, 235
 - FCS switch order, 237
 - TI zones, 405
 - zoning configurations, 362
- monitor configuration, restoring, 567
- monitoring
 - end-to-end performance, 553
 - frames, 558
 - trunks, 567
- Mozilla Firefox. *See: Firefox.*
- msCapabilityShow command, 47
- msConfigure command, 48, 49
- msPlatShow command, 47, 50
- msPIClearDb command, 51
- msPlMgmtActivate command, 46, 47
- msPlMgmtDeactivate command, 46, 47
- mstdDisable command, 52
- mstdEnable command, 51
- mstdReadConfig command, 51

N

- N_Port ID Virtualization. *See: NPIV.*
- N_Port, Access Gateway failover with FA-PWWN, 484
- name
 - chassis, 77

- fabric, 77
 - security certificate name, 271
 - switch, 76
- name server contents, displaying, 54
- naming ports, 89
- NAT, 117
- network address translation, see NAT
- network interface
 - displaying settings, 66
 - logical (bond0), 65
- Network OS connectivity, 593
- Network OS connectivity, unsupported configurations, 595
- network prefix length. See: CIDR.
- network time protocol, 74
- NPIV, 473–478
 - 10-bit addressing mode, 474
 - configuring, 475–476
 - disabling, 476
 - enabling, 476
 - F_Ports, 476
 - FCoE requirement, 476
 - fixed addressing mode, 474
 - overview, 473–474
 - PIDs, 86
 - upgrade considerations, 474
 - viewing PID login information, 478
 - viewing port configuration information, 476
- NPIV ports
 - DCC policy behavior, 241
 - duplicate login, 111
- nsAllShow command, 54, 105
 - described, 307
- nsShow command, 54, 105
 - described, 307
- NTP, 74
- NTP access, 74
- null encryption support for IKE policies, 275

O

- on-demand ports, 535–541
 - activating, 537
 - available ports, 535
 - disabling dynamic, 539
 - displaying installed licenses, 536
 - dynamic, 537
 - enabling dynamic, 538
 - supported devices, 535

- Open LDAP
 - See also: LDAP.
- OpenLDAP
 - configuring, 184–189
 - group assignment, 186
 - group membership, 185
 - group membership, enabling, 185
 - LDAP role mapping and, 186
 - memberOf overlay, 185
 - restrictions, 184
 - server configuration, 184
 - users, adding, 185
 - users, modifying, 186
- over-subscription versus congestions, 119

P

- passwd command, 63
- passwdCfg command, 159, 161
- passwdDefault command, 646
- password, 60
 - changing, 63, 158
 - changing defaults, 64
 - CHAP encryption requirement, 177
 - default for accounts, 63
 - limits, 63
 - recovery string, 165
 - recovery string, boot PROM password, 163
- Password Authentication Protocol (PAP), 177
- password database, distribution restrictions, 158
- password expiration policy, 161
- password history policy, 160
- password policy
 - account lockout, 161
 - strength, 159
- password strength policy, 159
- passwordless firmware download, 291
- passwords
 - boot PROM, 163–167
 - Backbone with recovery string, 164
 - Backbone without recovery string, 166
 - switch with recovery string, 163
 - switch without recovery string, 165
 - local user accounts, 157
 - policies for, 159–163
 - rules, 157
- path calculation using FSPF, 117
- path selection for routing, 116
- paths, defined, 116

- PEAP-MSCHAPv2, 177, 648, 652
- perfAddEeMonitor command, 554
- perfCfgClear command, 568
- perfCfgRestore command, 568
- perfCfgSave command, 568
- perfMonitorClear command, 557
- perfMonitorShow command, 556, 557
- performance data collection, 568
- perfSetPortEEMask command, 555
- perfTTmon command, 565, 566, 567
- permissions assigned to roles, 153
- phantom domains, 601–603
 - described, 599
- physical fabric administrator, 487
- physical fabric administrator user account, creating, 497
- PID, 83–87
 - 10-bit addressing mode, 84
 - assigning static, 87
 - automatic assignment, 86
 - binding overview, 83
 - clearing binding, 87
 - core addressing mode, 84
 - maximum number of assignments, 86
 - showing assignments, 87
 - static and NPIV, 86
 - swapping port area IDs, 91
 - WWN-based assignment, 86
 - WWN-based Virtual Fabrics assignment, 86
- PKI
 - generating DSA or RSA key pairs, 198
 - generating key pairs, 199
 - private key generation, 201
 - public key generation, 201
 - used with SSL, 200
- platform database, 47
- platform services, 46–47
 - disabling, 47
 - enabling, 47
 - Virtual Fabrics, 47
- platforms, FC-FC routing supported, 594
- PLOGI, 54
 - defined, 53
- POD
 - enabling ports, 92
 - releasing a port from a set, 540
 - reserving a port license, 540
 - See *also*: ports on demand.
- policies
 - account lockout, 161
 - account lockout duration, 162
 - account lockout threshold, 162
 - ACL, 231–235
 - ACL and Admin Domain considerations, 232
 - ACL and Virtual Fabric considerations, 232
 - ACL changes, activating, 233
 - ACL, storage, 231
 - ACL, viewing, 233
 - adding member to ACL, 234
 - admin lockout, 162
 - AUTH, 244
 - AUTH, distributing fabric-wide, 253
 - authentication for fabric, 243–253
 - DCC, 232, 238–242
 - DCC and Virtual Fabric considerations, 239
 - distributing local ACL, 263
 - ensuring fabric domains share policies, 236
 - fabric-wide consistency, 232, 605
 - fabric-wide consistency policy distribution, 260
 - fabric-wide policy enforcement, 263
 - FCS, 232, 235–238
 - FCS policy distribution, 238
 - IPsec, 270–275
 - IPv6 tunneling ICMP traffic, 275
 - manually distributing ACL policy database, 260
 - matching fabric-wide consistency policies, 265
 - non-matching fabric-wide consistency policies, 265
 - password, 159–163
 - password expiration, 161
 - password history, 160
 - removing member from ACL, 234
 - resolving conflicting ACL policies, 264
 - routing, 122–125
 - saving changes without activating, 233
 - SCC, 232, 242–243
 - SCC and Virtual Fabric considerations, 242
 - security, 231–275
 - switch database distribution setting, 260
- policy
 - ACL deleting, 233
 - ACL distribution, 262
 - activating IP Filter, 255
 - adding rule to an IP Filter policy, 259
 - authentication restrictions, 247
 - cloning an IP Filter, 254
 - creating DCC, 239
 - creating FCS, 236
 - creating for IP Filter, 254
 - creating SCC, 243
 - DCC deleting, 240
 - DCC restrictions, 239
 - default IP Filter policy rules, 258
 - deleting IP Filter, 255

- deleting rule from an IP Filter policy, 259
- device authentication, 246
- device authentication and Virtual Fabrics
 - considerations, 247
- displaying IP Filter, 254
- enforcing IP Filter, 258
- FCS restrictions, 235
- IP Filter, 253
- IP Filter policy distribution, 260
- management of ACL, 232–235
- members, identifying, 232
- modifying FCS, 235
- password strength, 159
- rules for IP Filter, 255
- saving IP Filter, 255
- using service names in IP Filter rules, 256
- policy database distribution, 260
 - settings, 261
 - viewing settings, 262
 - Virtual Fabric considerations, 260
- policy set, defined, 232
- port, 88–95
 - activating POD, 537
 - activation, 92
 - adding frame monitors to, 560
 - and blade compatibility, 98
 - assignment in logical switches, 312
 - blade enabling exceptions, 99
 - compression for long distance, 450
 - compression ratios
 - for compression-enabled ports, 450
 - for encryption-enabled ports, 450
 - configuration of ports, 229
 - configurations supported for Backbones, 321
 - configurations supported for fixed-port switches, 320
 - configuring E_Port authentication, 245
 - deactivation, 92
 - decommissioning, 92
 - deleting Top Talker monitor on, 567
 - disabling, 92
 - disabling dynamic POD, 539
 - disabling on blades, 98
 - displaying license assignments, 538
 - displaying the top n bandwidth-using flows, 566
 - dynamic POD, 537
 - enabling, 92
 - enabling compression, 450, 451
 - enabling dynamic POD, 538
 - enabling encryption, 450, 451
 - excluding from bottleneck detection, 440
 - FEC disabling for long distance, 113
 - FEC enabling for long distance, 113
 - ICL, 546
 - ID and Fibre Channel fabrics, 117
 - identification
 - by index, 90
 - by port area ID, 90
 - by slot and port number, 89
 - logical and zoning, 350
 - logical in ISL, 319
 - lossless dynamic load sharing, 129–131
 - moving, 313
 - naming, 89
 - port login command, 53
 - port login process, 54
 - port types, 88
 - ports and applications used by switches, 229
 - re-authenticating an E_Port, 246
 - releasing from a POD set, 540
 - removing frame monitors from, 560
 - reserving a POD license, 540
 - restrictions on Backbones, 321
 - restrictions on fixed-port switches, 320
 - restrictions on moving, 324
 - serial connection, 58
 - setting mode, 93
 - setting speed for a port octet, 95
 - slave port bottleneck detection, 440
 - SNMP filtering, 217
 - speed and number of encryption/compression ports, 447
 - Top Talker monitor, adding, 565
- port area ID, 90
- port area IDs, swapping, 91
- port decommissioning
 - on port with in-flight encryption/compression, 448
- port groups for trunking, 571
- port identifier. *See also:* PID.
- port index, 641
- port indexing, 90, 641–644
- port information, viewing, 465
- port mirroring, 88
- Port mode Top Talker monitor, described, 562
- port numbering schemes, 89
- port speeds, setting, 94
- port type
 - D_Port, 88
 - diagnostic, 88
 - E_Port, 88
 - EX_Port, 88
 - F_Port, 88
 - FL_Port, 88
 - G_Port, 88

- L_Port, 88
- M_Port, 88
- mirror, 88
- U_Port, 88
- VE_Port, 88
- VEX_Port, 88
- Port World Wide Name. *See also:* PWWN.
- port-based routing, 122, 123, 126
- portBufferCalc command, 450
- portBufferShow command, 142, 450
- portBufferShow command, 450
- portCfg command, 654
- portCfgCompress command, 456, 457
- portCfgEncrypt command, 455, 457, 654
- portCfgExport command, 603
- portCfgFillWord command, 90, 589
- portCfgISLMode command, 119, 121
- portCfgLongDistance command, 113, 589
- portCfgNpivPort command, 475, 476
- portCfgOctetSpeedCombo command, 95, 527, 528
- portCfgPersistentEnable command, 92
- portCfgQos command, 415, 425, 426
- portCfgShow command, 476
- portCfgSpeed command, 94, 448, 527, 528
- portCfgTrunkPort command, 574, 581
- portDecom command, 93
- portDisable command, 92, 574
- portEnable command, 92, 537
- portEncCompShow command, 448, 450, 452
- portLoginShow command, 478
- portName command, 89
- ports on demand, 535–541
 - activating, 537
 - available ports, 535
 - disabling dynamic, 539
 - displaying installed licenses, 536
 - dynamic, 537
 - enabling dynamic, 538
 - licence restrictions, 535
 - supported devices, 535
 - See also:* POD.
- portShow command, 477, 537
- portStatsShow command, 450
- portSwap command, 84, 91
- portSwapEnable command, 91
- portTrunkArea command, 574, 576, 581
- portZoneShow command, 342
- power management, 103
- powering down, 80
- powerOffListSet command, 103
- powerOffListShow command, 103
- power-on self tests for FIPS, 647
- preparing a switch for FIPS, 651
- preserving licenses, 515
- pre-shared key, and IPsec, 271
- primary FCS, 47
- primary FCS modifying, 237
- Principal ISLs, 116
- principal switch
 - defined, 53
- principal switch, capabilities, 53
- priority groups, for virtual channels, 119
- private key
 - deleting from switch, 200
 - generation, 201
- PRLI, 54
- protocol
 - Fibre Channel Common Transport (FC-CT), described, 46
 - HTTPS, described, 195
 - IPsec, described, 195
 - LDAPS, described, 195
 - SCP, described, 195
 - secure
 - HTTPS, 196
 - SCP, 196
 - SNMPv1, 196
 - SNMPv2, 196
 - SNMPv3, 196
 - SSHv2, 196
 - SNMP, described, 195
 - SSH, described, 196
 - SSL, 200
 - SSL, described, 196
 - telnet, 226
- protocols
 - authentication, 248
 - IPsec, 269
 - items needed to deploy secure, 196
 - secure, 196
 - security, 195
 - setting for authentication, 248
 - supported for IP Filter, 257
- proxy device, 599–600
 - described, 597
- proxy PID, 633
- proxy PID, described, 597
- psShow command, 104
- public key

- adding to switch, 198
- authentication, 198
- deleting from switch, 200
- generation, 201
- public key infrastructure and encryption, 200
- public key infrastructure. *See also:* PKI.

PWWN

- assigned by fabric, 479
- configuring FLOGI-time handling of duplicates, 110
- duplicates, 55
- handling duplicates, 111
- See also:* Port World Wide Name.

Q

QoS, 414

- buffer credit requirement, 145
- described, 119
- enabled by default, 425
- on E_Ports, 421
- over FC routers, 421
- SID/DID traffic prioritization, 415
- traffic prioritization, 421

QoS CS_CTL-based frame prioritization, 416–418

QoS port buffer configuration, 141

QoS zone-based traffic prioritization, 419

- disabling, 426
- High Availability considerations, 422
- limitations and restrictions, 424
- setting, 424
- setting over FC routers, 426
- supported configurations, 423
- trunking considerations, 424
- Virtual Fabrics considerations, 422

QoS zones, 119, 338

- defined, 419
- name prefix specified in an LSAN zone, 420

QSFP ports in DCX 8510 chassis, 544

Quality of Service. *See:* QoS.

R

RADIUS client

- configuration, 176
- enabling, 176

RADIUS server

- adding, 193
- configuration for FIPS, 653
- configuration with Admin Domains or Virtual Fabrics,

173

- configuring support with Linux, 175
- configuring support with Windows 2000, 177
- deleting, 193

- High Availability failover on, 175

- reordering, 193

- RSA setup, 179

- setup, 175–181

RADIUS service

- ADList, 174

- configuration, displaying, 194

- configuring, 175

- ContextRoleList, 174, 191

- disabling, 193, 194

- enabling, 193

- homeAD, 174

- Linux server-based, 175

- modifying, 193

- overview, 152, 172

- user, adding, 176

- Virtual Fabrics HomeContext, 174

- Windows 2008 support, 178

- Windows server-based, 177

RASLOG message

- FSPF-1009, 398

- ZONE-1062, 359

RBAC, 152

- Admin Domain considerations, 153

- and Fabric OS, 58

- role permissions, 153

- recommendations for trunk groups, 572

- recovering a device, 55

- redirecting frames, 132

- Registered State Change Notification, 54

- rejecting distributed user databases locally, 159

- releasing a port from a POD set, 540

- remote access policies, 178

- remote authentication, 167–194

- adding server to the switch configuration, 193

- changing authentication server configuration, 193

- changing authentication server contact order, 193

- deleting server from the switch configuration, 193

- displaying current configuration, 194

- enabling and disabling, 193

- switch configuration, 192

- removing

- Admin Domain members, 500

- Admin Domains from user accounts, 498

- alias members, 348

- frame monitors, 560

- frame monitors from a port, 560

- licensed features, 534

- LSAN tags, 627
- members from a zone configuration, 364
- ports from logical switches, 329
- zone configuration members, 364
- zone members, 352
- renaming Admin Domains, 500
- requirements
 - Admin Domains, 487
 - for F_Port trunking on an Access Gateway, 580
 - for trunk groups, 572
- restoring
 - configuration file, 280
 - logical switch configuration, 285
 - monitor configuration, 567
 - unordered frame delivery, 127
- restrictions
 - authentication policies, 247
 - Backbone ports, 321
 - compression, 446
 - encryption, 446
 - fixed-port switch ports, 320
 - in-flight encryption, 446
 - on FA-PWWN, 484
 - on moving ports, 324
 - QoS zone-based traffic prioritization, 424
 - upgrading temporary slot-based licenses, 531
 - Virtual Fabrics, 322
 - XISLs, 323
- rexec listener application, 228
- rlogin listener application, 228
- Role-Based Access Control. *See:* RBAC.
- roleConfig command, 154
- roles
 - Admin Domain considerations, 153
 - assigning user-defined, 155
 - creating user-defined, 154
 - default, 152
 - managing user-defined, 154–155
 - role permissions, 153
- root certificates
 - in Firefox, 205
 - in Internet Explorer, 205
 - installing in Java plugin, 205
 - installing to the Java plugin, 205
 - See also:* security certificates.
- route selection for routing, 116
- route selection, defined, 116
- routes, number supported using FSPF, 116
- routing
 - AP policies, 124
 - AP route policies, 124

- Backbone-to-edge, 600, 605
- cut-through, 117
- device-based, 122, 123, 126
- displaying current policy, 122
- distance vector, 115
- dynamic load sharing, 125
- exchange-based, 122, 123, 126
- FC-FC setup, 603–605
- frame order delivery, 126
- frame redirection, 132
- link state, 115
- lossless dynamic load sharing, 129
- out-of-order exchanges, 127
- overview, 115–117
- path and route selection, 116
- performance, 122
- port-based, 122, 123, 126
- route selection, 125
- setting AP route policy, 125
- setting policy, 125
- Virtual Fabrics, 124
- routing policies, 122–125
 - VE_Ports, 123
- RPC, unsupported in FIPS mode, 647
- RSA key pair generation, 198
- RSA RADIUS server, 179
- RSA RADIUS server, setup, 179
- RSA SecurID, 179
- RSCN, 77
- RSCN. *See:* Registered State Change Notification.
- rsh listener application, 228
- rstats listener application, 228
- rule
 - adding to an IP Filter policy, 259
 - configuring zones, 342
 - deleting from an IP Filter policy, 259
 - passwords, 157
- rusers listener application, 228

S

- sa-proposal, 269
- saved zone configuration, defined, 341
- saving monitor configuration, 567
- SCC
 - creating policy, 243
 - policies, 232, 242–243
 - Virtual Fabric considerations, 242
 - policy member, 232
- SCC policies, 242

- SCP
 - configuration for uploads and downloads, 197
 - described, 196
 - for certificates, 202
 - protocol, described, 195
 - secure protocol, 196
- SCR, defined, 53
- secAuthSecret command, 249, 250, 453
- secCertUtil command, 201, 202, 203, 252, 253, 271, 650
- secModeEnable command, 229
- secPolicyAbort command, 234
- secPolicyActivate command, 233, 236, 237, 264
- secPolicyAdd command, 234
- secPolicyCreate command, 236, 240, 243
- secPolicyDelete command, 233, 239, 240
- secPolicyFCSMove command, 237
- secPolicyRemove command, 234
- secPolicySave command, 240
- secPolicyShow command, 233
- secret key pair
 - communicating, 249
 - for DH-CHAP, 249
 - length, 249
 - setting, 250
 - viewing list of, 249
- secure copy protocol. *See:* SCP.
- Secure Fabric OS policies, 232
- secure LDAP, 168
- secure protocol
 - HTTPS, 196
 - items needed to deploy, 196
 - SCP, 196
 - SNMPv1, 196
 - SNMPv2, 196
 - SNMPv3, 196
 - SSHv2, 196
- Secure Shell protocol. *See:* SSH.
- Secure Sockets Layer protocol. *See:* SSL.
- security
 - AUTH policy, 243
 - browser support, 200
 - certificates, 196
 - encryption and SSL, 200
 - IAS remote access policies, 178
 - IP policy rules, 258
 - key generation, 201
 - management interface, 266–275
 - obtaining certificates, 203
 - policies, ACL, 231
 - protocols supported, 195
 - public key authentication, 198
 - root certificates in Firefox, 205
 - root certificates in Internet Explorer, 205
 - root certificates in Java plugin, 205
 - secure protocols, supported, 196
 - SSL certificates, 196
 - tunnel creation, 272–275
 - zoning and, 371
- Security Association Database (SADB), described, 269
- security associations, 269
 - flushing, 275
 - IPsec and, 269
- security certificates. *See:* certificates or root certificates.
- security considerations for FA-PWWN, 483
- security policies, 231–275
- security protocols, 195–196
- security scenarios, 196
- serial number, location on switch, 40
- serial port connection, 58
- serial port, console session, 58
- Server Application Optimization. *See:* SAO.
- sessions, maximum allowed, 154
- setContext command, 125, 333, 334
- setting
 - changing passwords, 64
 - chassis configurations, 95
 - chassis management IP interface, 68
 - date, 72
 - default zone mode, 495
 - fabric-wide consistency policy, 264
 - mask for end-to-end monitors, 555
 - port speeds, 94
 - QoS zone-based traffic prioritization, 424
 - QoS zone-based traffic prioritization over FC routers, 426
 - static ethernet IP address, 68
 - switch date and time, 72
 - time, 72
 - time zone, 73
 - time zone interactively, 73
- settings, configuration, 277–287
- shared ISL. *See:* extended ISL.
- shared secrets on Access Gateway, 250
- shelf life of a universal temporary license, 532
- shutdown
 - Backbone, 81
 - switch, 80
- SID/DID traffic prioritization, 415
- signed firmware, 301

Simple Network Management Protocol. See: SNMP.

slapd.conf file, 185

slave port bottleneck detection, 440

slot-based licensing, 526–527

slotPowerOff command, 103

slotPowerOn command, 103

slotShow command, 104, 604

SNMP

configuring, 206

described, 195

filtering ports, 217

host access, 229

password change, 157

SNMPv1

secure protocol, 196

SNMPv2

secure protocol, 196

SNMPv3

secure protocol, 196

switch and chassis context enforcement, 217

snmpConfig command, 653

snmpWalk command, 216, 217

special zones, 337

specification of ACL policy members, 232

Speed LSAN tag, 625

speed, setting for ports, 94

SSH

allowed-user, 198

configuring incoming authentication, 198

configuring outgoing authentication, 199

connection, 59

encrypted sessions, 197

protocol, described, 196

public key authentication, 198

SSHv2 protocol, 196

support, 197

supported OpenSSH protocol version, 197

sshd daemon, 197

ssh-keygen command, 198

sshUtil command, 198, 200, 652

sshutil command, 291

SSL, 251

browser support, 200

certificate files, 203

configuring, 201–205

encryption strength, 200

protocol, 200

protocol, described, 196

security certificates, 200

SSL security certificates, 196

standby CP blade, 175

State Change Registration. See: SCR.

static ethernet address, 67

static ethernet IP address, setting, 68

static PID, assigning, 87

static PIDs, NPIV, 86

statistics, bottleneck, 442

status of equipment, 104

status policy threshold values, setting, 106

status policy threshold values, viewing, 105

supported browsers, 200

supportSave command, 40

swapping blades, 99–102

switch

access, 229

access methods, Web Tools, 57

ACL policy distribution, 262

activation and deactivation, 78

adding public key, 198

applications used, 229

buffer credits by model, 143

certificates, installing, 203

changing name, 77

configuring for signed firmware, 301

configuring without disabling, 282

connecting, 82

connecting to a device, 90

connecting with different firmware, 81

default access, 229

deleting private key, 200

deleting public key, 200

disabling, 79, 102

disabling local switch protection, 262

disabling port, 92

displaying name server contents, 54

enabling, 79, 80

enabling local switch protection, 262

ethernet interface, 64

exporting public key, 199

firmware download, 294

firmware version testing, 302

firmware version, finding, 293

host access, 229

joining to fabric, 264

LDAP certificates

deleting, 651

exporting, 651

installing, 650

modifying FCS order, 237

modifying switch configuration, 281

name limitations, 77

names, 76

- naming, 75
- PKI key pair generation, 199
- ports used, 229
- restoring a configuration, 282
- serial number location, 40
- setting date and time, 72
- setting port speed, 94
- setting status policy threshold values, 106
- shutdown, 80
- switch database distribution setting, 260
- unique names for logical, 77
- user-defined accounts, 155
- viewing status policy threshold values, 105
- switch authentication mode, setting, 171
- switch authentication policy, 244
 - See also: AUTH.
- Switch Connection Control. See: SCC.
- switch firmware, 294–295
- switch WWN in Admin Domains, 492
- switchCfgPersistentDisable command, 102
- switchCfgSpeed command, 94
- switchCfgTrunk command, 574
- switchDisable command, 79, 111, 125, 541
- switchEnable command, 79, 80, 111, 336
- switches supported for FA-PWWN, 483
- switchName command, 77
- switchShow command, 641
- switchShow command, 90, 104, 105, 334, 336, 452, 473, 477, 538, 541
- switchStatusPolicySet command, 210
- switchStatusPolicySet command, 106
- switchStatusPolicyShow command, 106
- switchStatusShow command, 104
- syslogDlpAdd command, 109
- sysShutdown command, 80, 81
- system-defined Admin Domains, 488

T

- tac_plus command, 190
- tac_plus daemon, 190
- tac_plus.cfg file, 190
- TACACS+ server
 - adding, 193
 - deleting, 193
 - failover, 189
 - installing, 190
 - reordering, 193
 - retry, 190

- supported protocols, 189
- timeout, 189
- TACACS+ service
 - ADList, 191
 - Admin Domains, configuring, 191
 - authentication service, 189
 - configuration, 189
 - configuration, displaying, 194
 - disabling, 193, 194
 - enabling, 193
 - home Virtual Fabric, 191
 - homeAD, 191
 - LINUX based, 190
 - modifying, 193
 - overview, 152
 - password expiration, configuring, 192
 - user, adding, 190
 - vendor attributes, 190
 - Virtual Fabrics, configuring, 191
 - Windows server-based, 192
- tags for LSAN zones, 624
- telnet
 - blocking access, 227
 - connection, 59
 - protocol, 226
 - unblocking access, 228
- temporary licenses, 530–532
 - defined, 530
- temporary slot-based licenses, upgrade restrictions, 531
- Terminal Access Controller Access-Control System Plus
 - protocol. See: TACACS+.
- TI zones, 338
 - activating, 406
 - admin domain considerations, 398
 - changing state, 406
 - creating, 402
 - creating in a base fabric, 404
 - deactivating, 406
 - defined, 380
 - deleting, 407
 - displaying, 407
 - failover, 380–382
 - FC router limitations, 390
 - FC routers, 386
 - limitations and restrictions, 398
 - modifying, 405
 - supported configurations, 396
 - trunk port violation handling, 395
 - trunking, 397
 - Virtual Fabric considerations, 402
 - Virtual Fabrics considerations, 399
 - with Virtual Fabrics, 401

- within a Backbone fabric, 389
 - within an edge fabric, 388
- time and date, 72
- time listener application, 228
- Time server, described, 45
- time settings, 72
- time zone
 - setting, 73
 - setting interactively, 73
- time zone settings, 72–74
- time, synchronizing local and external, 74
- time-based licenses, 530–532
- Top Talker monitors
 - adding on all switches in fabric, 565
 - adding to a port (port mode), 565
 - and FC-FC routing, 563
 - defined, 552
 - deleting all in fabric, 567
 - deleting on a port, 567
 - fabric mode, described, 563
 - limitations, 565
 - port mode, described, 562
- Top Talkers, 562
 - logical switches and, 329
- topologies
 - core-edge, 549
 - mesh, 547
 - supported for ICL connections, 547–550
- topology database, 116
- topology discovery, 51–52
 - disabling, 52
 - displaying status, 51
 - enabling, 51
- topologyShow command, 402
- traffic flow limitations for lossless core, 130
- traffic isolation, 379–412
 - enhanced zones, 384–386
 - failover, 380–382
 - FSPF routing rules, 383
 - over FCR, 386
 - over FCR with Virtual Fabrics, 401
 - overview, 379
- traffic isolation zones, 338
- traffic isolation zoning, 379–412
- Traffic Isolation. *See also* TI.
- traffic patterns, planning for, 573
- traffic prioritization, 421
 - QoS zone-based, 419
 - SID/DID, 415
- traffic selector, and IPsec, 270
- traffic support, 115
- traffic, limiting from a device, 415
- transaction model for managing Admin Domains, 494
- transform set, and IPsec, 270
- transform set, defined, 270
- traps, 209
- trunk area and admin domains, 576
- trunk area, enabling DCC policy on, 586
- trunk groups
 - configuring, 573
 - recommendations, 572
 - requirements, 572
- trunk monitoring, 567
- trunk port violation handling, for TI zones, 395
- trunking
 - configuring F_Port for Brocade adapters, 581
 - disabling, 574
 - disabling F_Port trunking, 585
 - disabling ISL, 574
 - displaying F_Port information, 585
 - displaying information, 574
 - enabling, 574
 - enabling ISL, 574
 - EX_Port, 577–579
 - F_Port, 579–586
 - F_Port considerations, 582
 - F_Port for access gateways, 579
 - F_Port for Brocade adapters, 581
 - F_Ports and Virtual Fabrics, 584
 - High Availability support, 572
 - ICL on DCX and DCX-4S, 547
 - ISL over long distance fabrics, 576
 - license requirements, 571
 - managing, 569
 - masterless, 570
 - overview, 569
 - port groups, 571
 - supported configurations, 571
 - supported platforms, 571
 - types, 570
 - with TI zones, 397
- trunkShow command, 574
- tsClockServer command, 74
- tsTimeZone command, 72, 73
- tunnel
 - configurations using IPsec, 267–268
 - creation, 272–275

U

- U_Port, described, 88
- unblocking telnet access, 228
- understanding MIBs, 208
- understanding SNMP bASICS, 207
- universal temporary license
 - defined, 530
 - described, 532
 - extending, 532
 - shelf life, 532
- unlocking an account, 162
- unordered frame delivery, restoring, 127
- upgrading firmware, 291
- upgrading temporary slot-based licenses, restrictions, 531
- uploading AD configuration file, 512
- USB device, 299, 299–300
- usbStorage command, 299
- user account
 - assigning Admin Domains to, 497
 - creating a physical fabric administrator, 497
 - for managing Admin Domains, 497
- user accounts, 151–194
 - distributing local database, 158
 - Fabric OS and LDAP, 171
 - Fabric OS, 171–172
 - Fabric OS and RADIUS servers, 172–181
 - local database distribution, 158
 - local database of, 155–158
 - managing, 151–194
 - overview, 151
 - password policies, 159–163
 - password rules, 157
 - removing Admin Domains, 498
- user authentication, 171
- user databases, 158
 - accepting distributed locally, 158
 - rejecting distributed locally, 159
- user-assigned FA-PWWN behavior, 480
- userConfig command, 155, 497, 498, 654
- user-defined accounts, 155
- user-defined Admin Domains, 488
- user-defined role
 - assigning, 155
 - creating, 154
 - managing, 154–155
- User-Principal-Name, 181
- users
 - assigning to Admin Domains, 496
 - authenticating, 152

- using security certificates, 200

V

- validating a zone, 358
- validating Admin Domain members, 506
- VE_Ports
 - described, 88
 - routing policy, 123
 - XISL and FX8-24, 321
- verification check, 604
- verifying
 - device connectivity, 82, 105
 - High Availability features, 104
 - host syslog, 109
- version command, 604
- VEX_Port
 - described, 88
 - FCIP required, 606
 - in Fibre Channel, 596
- VF mode
 - Admin Domains and, 325
 - definition, 324
 - See also: Virtual Fabrics
- viewing
 - ACL policies, 233
 - alias, 349
 - authentication parameter settings, 248
 - compression configuration, 452
 - current default zone access mode, 361
 - encryption configuration, 452
 - fabric-wide consistency policy, 263
 - frame redirect zones, 133
 - installed licenses, 532
 - list of secret key pairs, 249
 - NPIV port configuration information, 476
 - policy database distribution settings, 262
 - port information, 465
 - virtual PID login information, 478
 - zone configuration information
 - all, 366
 - in effective zone database, 367
 - selected zones, 367
 - zones, 356
- virtual channels, 119
 - priority groups, 119
- Virtual Fabrics
 - account management, 319
 - ACL policy considerations, 232
 - AUTH module considerations, 244

- base switch
 - about, 316
 - creating, 326
- changing logical switch to base switch, 331
- configDownload restrictions, 286
- configUpload restrictions, 286
- configuration management, 285
- considerations
 - for Adv. Perf. Monitoring, 552
 - for WWN-based PID assignment, 86
- considerations for ICLs, 547
- ContextRoleList, 174, 191
- date settings, 72
- DCC policy considerations, 239
- disabling, 325
- DPS support, 124
- E_Port considerations, 245
- enabling, 324
- ethernet interface, 65
- extended ISL (XISL), 317
- F_Port trunking, 584
- FC-FC routing, 636
- home Virtual Fabric, TACCS+, 191
- HomeContext, 174
- ingress rate limiting, 414
- interaction with Fabric OS features, 322
- interaction with other Fabric OS features, 322
- IP address removal, 333
- IP address setup, 333
- IP Filter policy considerations, 254
- LDAP server, 183, 187
- logical fabrics
 - about, 315
 - context change, 334
- logical ISL (LISL), 317
- logical switch
 - creating, 326
 - default, 310
 - deleting, 329
 - displaying configuration, 330
 - overview, 310
- lossless dynamic load sharing, 131
- Microsoft Active Directory service, 183
- OpenLDAP server, 187
- overview, 309
- password database distribution restrictions, 158
- permissions and Admin Domains, 151
- platform services, 47
- policy database distribution considerations, 260
- ports, moving, 329
- QoS zone-based traffic prioritization considerations, 422
- RADIUS configuration, 173

- RADIUS server configuration, 173
- restrictions, 322
- SCC policy considerations, 242
- supported platforms, 320
- TACACS+ service, 191
- TI zone considerations, 399, 402
- with traffic isolation over FCR, 401
- XISL, allowing on logical switches, 333
- zone alias considerations, 347
- zone database size considerations, 362
- virtual PID login information
 - viewing, 478
- VSA
 - permissions assignment, 171
 - VSA-based account role syntax, 171

W

- web server used by Fabric OS, 204
- Web Tools access methods, configuration, 57
- well-known address, described, 45
- Windows 2000
 - IAS and Fabric OS user setup, 172
 - RADIUS authentication on, 177
- Windows 2008 RADIUS (NPS) support, 178
- Windows IAS, Fabric OS user setup, 172
- Windows server
 - LDAP authentication on, 182
 - RADIUS authentication on, 177
 - TACACS+ authentication on, 192
- World Wide Name command. See: wwn command.
- WWN, 533
 - format for logical ports, 319
 - switch WWNs in Admin Domains, 492
- wwn command, 40
- wwnAddress command, 87
- WWN-based PID assignment, 86
 - considerations for Virtual Fabrics, 86

X

- XISL
 - Brocade 7800 restriction, 320
 - default logical switch restriction, 321
 - ICL port restriction, 321
 - on FX8-24, 321
 - See also: extended ISL.
- xlate domain ID, 602
- xlate domains, 601

Z

zeroization functions for FIPS, 645

zeroizing for FIPS, 655

zone

access mode, viewing current, 361

accessing, 229

adding a new switch or fabric, 371

adding members, 351

administering security, 371

alias

adding members, 347

deleting, 349

removing members, 348

viewing, 349

Virtual Fabrics considerations, 347

wildcard usage, 351, 352, 354

all access, 360

broadcast, 337, 343

broadcast (reserved name), 350

concepts, 338

concurrent transactions, 376

configuration management, 370

configurations, 341

adding members, 363

creating and maintaining, 362

managing, 368

configuring rules, 342

creating, 350

creating a configuration, 363

creating frame redirect, 132

creation and maintenance, 350–360

database changes, examining, 356

database configurations, viewing, 367

database size, 362

database size and Virtual Fabric considerations, 362

default zone mode, 360, 495

deleting, 355

configurations, 365

deleting frame redirect, 133

disabled zone configuration, defined, 341

disabling a configuration, 365

effective zone configuration, defined, 341

enabling a configuration, 364

existing, 350

frame redirection, 337

LSAN, 337

maximum database size, 362

merging, 362, 371–376

merging scenarios, 373

no access, 360

objects, 340

optimizing resources, 338

QoS, 338

QoS zones, defined, 419

removing members, 352

from a configuration, 364

replacing member, 353

saved zone configuration, defined, 341

schemes, 341

setting default zoning mode, 361

special, 337

splitting a fabric, 373

terminology, 338

TI, 338

TI zones, defined, 380

traffic isolation, 338

types, 339

viewing

database transactions, 377

frame redirect, 133

zones, 356

viewing configuration information

for all, 366

for selected zones, 367

in effective zone database, 367

zone configuration, defined, 341

zone object, 340

zone types, 337

zone command, 132, 133, 369, 395, 402, 405, 406, 407, 409

zone configuration database, maximum items, 362

zone configurations

clearing, 367

creating, 363

deleting, 365

disabling, 365

enabling, 364

removing members, 364

zone database and Admin Domains, 510

zone database, maximum size, 362

zone object

copying, 368

deleting, 369

maintenance, 368–370

renaming, 370

zoneAdd command, 351

zoneCreate command, 350, 424

zoneDelete command, 355

zoneHelp command, 338

zoneObjectRename command, 370

zoneObjectReplace command, 353

zoneRemove command, 352

zoneShow command, 356

zoning

- advanced, 337–376

- advanced commands, 338

- defined, 338

- enforcement, 342

- on logical ports, 350

- overview, 338