

53-1002840-01
22 July 2013



Network OS

Administrator's Guide

Supporting Network OS v4.0.0

BROCADE

Copyright © 2010-2013 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, Brocade Assurance, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, MLX, NetIron, SAN Health, ServerIron, TurboIron, VCS, and VDX are registered trademarks, and AnyIO, Brocade One, CloudPlex, Effortless Networking, ICX, NET Health, OpenScript, and The Effortless Network 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>Network OS Administrator's Guide</i>	53-1002080-01	New document	December 2010
<i>Network OS Administrator's Guide</i>	53-1002339-01	Updated for Network OS v2.1.0	September 2011
<i>Network OS Administrator's Guide</i>	53-1002491-01	Updated for Network OS v2.1.1	December 2011
<i>Network OS Administrator's Guide</i>	53-1002561-01	Updated for Network OS 3.0.0	September 2012
<i>Network OS Administrator's Guide</i>	53-1002561-02	Corrections for Network OS 3.0.0	October 2012
<i>Network OS Administrator's Guide</i>	53-1002795-01	Updated for Network OS 3.0.1	December 2012
<i>Network OS Administrator's Guide</i>	53-1002795-02	Updated for Network OS 3.0.1a	March 2013
<i>Network OS Administrator's Guide</i>	53-1002795-03	Updated for Network OS 3.0.1b	July 2013
<i>Network OS Administrator's Guide</i>	53-1002840-01	Updated for Network OS 4.0.0	July 2013

Contents (High Level)

Section I

Network OS Administration

Chapter 1	Introduction to Network OS and Brocade VCS Fabric Technology	3
Chapter 2	Using the Network OS CLI	17
Chapter 3	Basic Switch Management	27
Chapter 4	Network Time Protocol	67
Chapter 5	Configuration Management	71
Chapter 6	Installing and Maintaining Firmware	81
Chapter 7	Administering Licenses	95
Chapter 8	SNMP	111
Chapter 9	Fabric	123
Chapter 10	Metro VCS	135
Chapter 11	Administering Zones	145
Chapter 12	Configuring Fibre Channel Ports	181
Chapter 13	System Monitor	191
Chapter 14	VMware vCenter	205
Chapter 15	Configuring Remote Monitoring	211

Section II

Network OS Security Configuration

Chapter 16	Managing User Accounts	217
Chapter 17	External server authentication	235
Chapter 18	FIPS Support	265
Chapter 19	Fabric Authentication	279

Section III

Network OS Layer 2 Switch Features

Chapter 20	Administering Edge-Loop Detection	291
Chapter 21	Configuring AMPP	299
Chapter 22	Configuring FCoE Interfaces	313

Chapter 23	Configuring VLANs	331
Chapter 24	Configuring STP-Type Protocols	345
Chapter 25	Configuring UDLD	373
Chapter 26	Configuring Link Aggregation	377
Chapter 27	Configuring LLDP	389
Chapter 28	Configuring ACLs	403
Chapter 29	Configuring QoS	413
Chapter 30	Configuring 802.1x Port Authentication	461
Chapter 31	Configuring sFlow	469
Chapter 32	Configuring Switched Port Analyzer	477
Chapter 33	Configuring SFP breakout mode	485

Section IV

Network OS Layer 3 Routing Features

Chapter 34	In-band Management	495
Chapter 35	IP Route Policy	505
Chapter 36	IP Route Management	509
Chapter 37	Configuring PIM	513
Chapter 38	Configuring OSPF	525
Chapter 39	Configuring VRRP	541
Chapter 40	Configuring VRF	555
Chapter 41	Configuring BGP	559
Chapter 42	Configuring IGMP	587

Section V

Troubleshooting

Chapter 43	Using the Chassis ID (CID) Recovery Tool	595
Chapter 44	Troubleshooting	599
Appendix A	TACACS+ Accounting Exceptions	655
Appendix B	Supported time zones and regions	659

Contents (Detailed)

About This Document

Supported hardware and software	xxxv
What's new in this document	xxxvi
Document conventions	xxxvi
Text formatting	xxxvi
Command syntax conventions	xxxvii
Command examples	xxxviii
Notes, cautions, and warnings	xxxviii
Key terms	xxxviii
Notice to the reader	xxxviii
Additional information	xxxix
Brocade resources	xxxix
Other industry resources	xxxix
Getting technical help	xl
Document feedback	xl

Section I **Network OS Administration**

Chapter 1

Introduction to Network OS and Brocade VCS Fabric Technology

Introduction to Brocade Network OS	3
Brocade VCS Fabric terminology	4
Introduction to Brocade VCS Fabric technology	5
Automation	6
Distributed intelligence	7
Logical chassis	8
Ethernet fabric formation	9
Brocade VCS Fabric technology use cases	10
Classic Ethernet Access and Aggregation use case	10
Large scale server virtualization use case	11
Brocade VCS Fabric connectivity with Fibre Channel SAN	12
Topology and scaling	12
Core-edge topology	13
Ring topology	14
Full mesh topology	14

Chapter 2

Using the Network OS CLI

DCB command line interface	17
----------------------------------	----

Saving your configuration changes	17
Network OS CLI RBAC permissions	18
Default roles	18
Accessing the Network OS CLI through Telnet	18
Network OS CLI command modes	19
Network OS CLI keyboard shortcuts	22
Using the do command as a shortcut	22
Displaying Network OS CLI commands and command syntax.	23
Network OS CLI command completion	24
Network OS CLI command output modifiers.....	24
Show command output information	25

Chapter 3

Basic Switch Management

Connecting to the switch	27
Telnet or SSH connections	28
SSH server key exchange	28
Importing an SSH public key.....	29
Connecting through a Telnet or SSH session.....	29
Connecting through a Telnet or SSH session.....	30
Shutting down Telnet service	30
Starting Telnet service.....	30
Shutting down SSH service.....	30
Starting SSH service	31
Firmware upgrade and downgrade considerations.....	31
Switch attributes.....	31
Setting and displaying the host name	31
Setting and displaying the chassis name.....	32
Switch types.....	33
Disabling and enabling a chassis	33
Rebooting a Brocade switch.....	34
Rebooting a compact switch.....	34
Rebooting a modular chassis	34
Operational modes	35

Logical chassis cluster mode	35
Configuration	36
Creating a logical chassis cluster	37
Precautions to take for mode transitions.....	39
Converting a fabric cluster to a logical chassis cluster	40
Converting a fabric cluster while preserving configuration.....	41
Selecting a principal node for the cluster	43
Converting a logical chassis cluster to a fabric cluster	43
Converting to a fabric cluster while preserving configuration	44
Adding a node to a logical chassis cluster	44
Removing a node from a logical chassis cluster	44
Rejoining a node to the cluster.....	45
Replacing a node in a logical chassis cluster	45
Merging two logical chassis clusters	46
Changing an rbridge-id on a switch within a fabric	46
Examples of global and local configurations	47
Additional operations.....	48
Fabric cluster mode	48
Standalone mode	49
Configuring the Ethernet management interface.....	49
Ethernet interfaces	49
Configuring a static IP address.....	50
Configuring an IP address with DHCP	51
Stateless IPv6 autoconfiguration	52
Setting IPv6 autoconfiguration	52
Displaying the network interface	52
Configuring the management interface speed	53
Outbound Telnet and SSH	54
Establishing a Telnet connection	54
SSH supported features	54
Establishing an SSH connection.....	55
Modular platform basics.....	55
Management module	56
Switch fabric modules	56
Interface modules	56
Supported interface modes	57
Displaying the interfaces.....	57
Slot numbering.....	58
Slot configuration.....	59
Replacing an interface module.....	60
High availability failover.....	61
High availability commands	61
Configuring a switch banner.....	62
Setting and displaying a banner	62

	supportSave data	63
	Uploading supportSave data to an external host	63
	Saving supportSave data to an attached USB device	63
	Displaying the status of a supportSave operation	64
	Configuring autoupload of supportSave data	64
	Displaying the autoupload configuration	64
	Additional supportSave commands	65
	Message logging	65
Chapter 4	Network Time Protocol	
	Date and time settings	67
	Setting the date and time	67
	Time zone settings	68
	Setting the time zone	68
	Displaying the current local clock and time zone	68
	Removing the time zone setting	69
	Network Time Protocol	69
	Synchronizing the local time with an external source	70
	Displaying the active NTP server	70
	Removing an NTP server IP address	70
Chapter 5	Configuration Management	
	Switch configuration overview	71
	Flash file management	71
	Listing the contents of the flash memory	72
	Deleting a file from the flash memory	72
	Renaming a file	72
	Viewing the contents of a file in the flash memory	72
	Configuration file types	73
	Default configuration	73
	Startup configuration	74
	Running configuration	74
	Saving configuration changes	74
	Saving the running configuration	75
	Saving the running configuration to a file	75
	Applying previously saved configuration changes	75
	Configuration backup	76
	Uploading the startup configuration to an external host	76
	Backing up the startup configuration to a USB device	76
	Configuration restoration	76
	Restoring a previous startup configuration from backup	77
	Restoring the default configuration	77
	Configuration management on a modular chassis	78
	Configuration management on interface modules	78
	Configuration management in redundant management modules	78

Configuration management in Brocade VCS Fabric mode	79
Downloading a configuration to multiple switches	80
Automatic distribution of configuration parameters	80

Chapter 6

Installing and Maintaining Firmware

Firmware upgrade overview	81
Obtaining and decompressing firmware	82
Upgrading firmware on a compact switch	82
Upgrading firmware on a modular chassis.	82
Upgrading and downgrading firmware	83
Firmware upgrade in a local switch	84
Preparing for a firmware download.	84
Obtaining the firmware version.	84
Connecting to the switch.	85
Downloading firmware from a remote server.	86
Downloading firmware from a USB device.	87
Activating nodes after an upgrade	88
Firmware upgrade using the manual option	88
Downloading firmware using the manual option.	88
Committing the firmware upgrade	89
Restoring the previous firmware version	89
Verification and error handling	90
Firmware upgrade in Brocade fabric cluster mode.	90
Firmware upgrade in Brocade logical chassis cluster mode	91
Recovering firmware	92
Verification and error handling	92

Chapter 7

Administering Licenses

Licensing overview	95
Permanent licenses	97
Temporary licenses	97
Individual time-based licenses	98
Universal time-based licenses	98
License expiration	98
Extending a license	98
Usage restrictions	98
Managing licenses	99
Displaying the switch license ID	99
Obtaining a license key	99
Installing a license.	99
License Removal	102
Dynamic Ports on Demand.	104
Automatic POD port assignments.	105
Mapping port assignments to a POD port set	105
Activating the Dynamic POD feature.	105
Displaying the Dynamic POD assignments	107
Overriding Dynamic POD assignments.	107

Upgrade and downgrade considerations.	109
Configuration management considerations	109

Chapter 8

SNMP

Simple Network Management Protocol	111
SNMP Manager	111
SNMP Agent	111
Management Information Base (MIB)	112
Basic SNMP operation	112
Understanding MIBs	113
Access to MIB variables.	114
Loading Brocade MIBs	114
Brocade MIB files.	114
Agent Capability MIBs	115
Standard MIBs	115
MIB loading order.	116
SNMP community strings	116
Adding an SNMP community string	117
Changing the access of a read-only community string	117
Removing an SNMP community string	117
SNMP server hosts	118
Setting the SNMP server host.	118
Removing the SNMP server host	119
Configuring the SNMP system group	119
Support for multiple SNMP server contexts	120
Setting the SNMP server context	120
Support for password encryption for SNMPv3 users	120
Displaying SNMP configurations	121

Chapter 9

Fabric

TRILL	123
Brocade VCS Fabric formation	123
How RBridges work	124
Neighbor discovery	125
Brocade trunks.	125
Fabric formation.	125
Fabric routing protocol.	126
Brocade VCS Fabric configuration management	126
Brocade VCS Fabric configuration tasks	127

Fabric interface configuration management.	128
Enabling a Fabric ISL.	128
Disabling a Fabric ISL.	129
Enabling a Fabric trunk.	129
Disabling a Fabric trunk.	129
Broadcast, Unknown Unicast, and Multicast Forwarding.	129
Priorities.	130
Displaying the running configuration.	130
VCS Virtual IP address configuration.	131
Fabric ECMP load balancing.	133

Chapter 10

Metro VCS

Metro VCS overview.	135
Configuring Metro VCS using long distance ISL.	136
Metro VCS using long distance guidelines and restrictions.	137
Configuring Metro VCS using standard ISL.	139
Standard distance Metro VCS guidelines and restrictions.	140
Metro VCS and distributed Ethernet VLAGs.	140
Connecting local VCS Clusters over long distance using vLAG. ...	142

Chapter 11

Administering Zones

Zoning overview.	145
Approaches to zoning.	147
Zone objects.	148
Zoning enforcement.	150
Considerations for zoning architecture.	150
Operational considerations.	150
Supported modes.	150
Supported firmware.	150
Firmware downgrade and upgrade considerations.	151
Zone configuration management.	151
Default zoning access modes.	153
Setting the default zoning mode.	153
Zone database size.	154
Viewing database size information.	154
Zone aliases.	154
Creating an alias.	155
Adding additional members to an existing alias.	155
Removing a member from an alias.	156
Deleting an alias.	157
Zone creation and management.	157
Creating a zone.	157
Adding a member to a zone.	158
Removing a member from a zone.	159
Deleting a zone.	159

	Zone configuration management.	160
	Viewing the defined configuration	160
	Viewing the enabled configuration	161
	Creating a zone configuration	162
	Adding a zone to a zone configuration	163
	Removing a zone from a zone configuration	163
	Enabling a zone configuration	164
	Disabling a zone configuration	165
	Deleting a zone configuration	165
	Clearing changes to a zone configuration	166
	Clearing all zone configurations	166
	Backing up the zone configuration	167
	Restoring a configuration from backup	167
	Zone configuration scenario.	168
	Zone merging.	169
	Fabric segmentation and zoning.	171
	Zone merging scenarios	171
	LSAN Zones	175
	LSAN naming	175
	Managing domain IDs	176
	Configuring LSAN zones—device sharing example	176
Chapter 12	Configuring Fibre Channel Ports	
	Fibre Channel ports overview.	181
	Fibre Channel port activation and deactivation	183
	Enabling a Fibre Channel port	183
	Disabling a Fibre Channel port	184
	Fibre Channel port attributes.	184
	Viewing Fibre Channel port attributes	185
	Setting Fibre Channel port speed	186
	Long-distance operation.	186
	Configuring for long-distance operation	187
	Configuring a Fibre Channel port for trunking.	188
	Monitoring Fibre Channel ports	188
Chapter 13	System Monitor	
	System Monitor overview	191
	Monitored components	191
	Monitored FRUs	192

System Monitor examples	193
Viewing system SFP optical monitoring defaults	193
Setting system thresholds	194
Setting state alerts and actions	194
Displaying the switch health status	194
Configuring e-mail alerts	195
Threshold Monitor overview	195
CPU and memory monitoring	196
SFP monitoring	197
Security monitoring	199
Interface monitoring	199
Threshold Monitor examples	200
Viewing threshold status	200
CPU and memory threshold monitoring examples	201
SFP monitoring example	202
Security monitoring examples	202
Interface monitoring examples	203
Pausing and continuing threshold monitoring	204

Chapter 14

VMware vCenter

vCenter and Network OS integration	205
vCenter properties	205
vCenter guidelines and restrictions	206
vCenter discovery	206
vCenter configuration	207
Step 1: Enabling QoS	207
Step 2: Enabling CDP/LLDP	207
Step 3: Adding and Activating vCenter	207
Discovery timer interval	208
User-triggered vCenter discovery	209
Viewing the discovered virtual assets	209

Chapter 15

Configuring Remote Monitoring

RMON overview	211
RMON configuration and management	211
Default RMON configuration	211
Configuring RMON events	211
Configuring RMON Ethernet group statistics collection	212
Configuring RMON alarm settings	212

Section II

Network OS Security Configuration

Chapter 16

Managing User Accounts

User accounts	217
Default accounts in the local switch user database	217
Creating and modifying a user account	218

Role-based access control (RBAC)	221
Default roles	221
User-defined roles	221
Creating a user-defined role	222
Command access rules	223
Configuration examples	228
Password policies	229
Password strength policy	230
Password encryption policy	230
Account lockout policy	231
Password interaction with remote AAA servers	232
Managing password policies	232
Security event logging	233

Chapter 17

External server authentication

Remote server authentication overview	235
Login authentication mode	236
Conditions for conformance	237
RADIUS	238
Authentication and accounting	238
Authorization	239
Account password changes	239
RADIUS authentication through management interfaces	239
Client-side RADIUS server configuration	239
Server-side RADIUS configuration	241
TACACS+	244
Authorization	245
TACACS+ authentication through management interfaces	245
Supported packages and protocols	245
Client-side TACACS+ server configuration	245
TACACS+ accounting	247
Conditions for conformance	248
Configuring TACACS+ accounting on the client	248
Viewing the TACACS+ accounting logs	249
Firmware downgrade considerations	250
TACACS+ server-side configuration	250
Configuring TACACS+ for a mixed vendor environment	252
LDAP	253
User authentication	253
Server authentication	254
Server authorization	255
FIPS compliance	255
Client-side Active Directory server configuration	255
Active Directory groups	260
Server-side Active Directory configuration	261
Configuring LDAP users on an AD server	262
Clearing sessions	263

Chapter 18	FIPS Support	
	FIPS overview	265
	Zeroization functions	266
	Power-on self-tests	266
	Conditional tests	267
	FIPS-compliant state configuration	267
	Preparing the switch for FIPS	268
	FIPS preparation overview	268
	Enabling the FIPS-compliant state	269
	Zeroizing for FIPS	275
	LDAP in the FIPS-compliant state	275
	Setting up LDAP for the FIPS-compliant state	276

Chapter 19	Fabric Authentication	
	Fabric authentication overview	279
	DH-CHAP	279
	Shared secret keys	280
	SSH server key exchange configuration	281
	Authentication Policy configuration	281
	Configuring device authentication	282
	Switch connection control (SCC) policy	284
	Defined and active SCC policy sets	284

Section III *Network OS Layer 2 Switch Features*

Chapter 20	Administering Edge-Loop Detection	
	Edge-loop detection overview	291
	How ELD detects loops	293
	Configuring edge-loop detection	295
	Setting global ELD parameters for a Brocade VCS Fabric cluster	295
	Setting interface parameters on a port	296
	Edge-loop troubleshooting	296

Chapter 21	Configuring AMPP	
	AMPP overview	299
	AMPP over vLAG	300
	AMPP and Switched Port Analyzer	301
	Scalability	301

Configuring AMPP port-profiles	302
Life of a port-profile	303
Configuring a new port-profile	304
Configuring VLAN profiles	305
Configuring FCoE profiles	306
Configuring QoS profiles	306
Configuring security profiles	307
Deleting a port-profile-port	308
Deleting a port-profile	308
Deleting a sub-profile	309
Monitoring AMPP profiles	309

Chapter 22

Configuring FCoE Interfaces

FCoE overview	313
FCoE terminology	314
End-to-end FCoE	314
FCoE operations	314
FCoE end-to-end forwarding	314
Layer 2 Ethernet overview	317
Layer 2 forwarding	317
VLAN tagging	318
Frame classification (incoming)	319
Congestion control and queuing	319
Access control	321
Trunking	321
Flow control	322
FCoE Initialization Protocol	322
FIP discovery	322
FIP login	323
FIP logout	323
Name server	324
Registered State Change Notification	324
FCoE queuing	325
Configuring FCoE interfaces	325
Assigning an FCoE map onto an interface	325
Assigning an FCoE map onto a LAG member	326
FCoE over LAG	327
Configuration guidelines and restrictions	327
FCoE provisioning on LAGs	327
Logical FCoE ports	328
xSTP reconvergence	329

Chapter 23

Configuring VLANs

VLAN overview	331
Ingress VLAN filtering	331
VLAN configuration guidelines and restrictions	333

Default VLAN configuration	334
VLAN configuration and management	334
Enabling and disabling an interface port	334
Configuring the MTU on an interface port	334
Creating a VLAN	335
Enabling STP on a VLAN	335
Disabling STP on a VLAN	335
Configuring an interface port as a Layer 2 switch port	336
Configuring an interface port as an access interface	336
Configuring an interface port as a trunk interface	337
Disabling a VLAN on a trunk interface	337
Configuring protocol-based VLAN classifier rules	338
Configuring a VLAN classifier rule	338
Configuring MAC address-based VLAN classifier rules	338
Deleting a VLAN classifier rule	339
Creating a VLAN classifier group and adding rules	339
Activating a VLAN classifier group with an interface port	339
Displaying VLAN information	339
Configuring the MAC address table	340
Specifying or disabling the aging time for MAC addresses	340
Adding static addresses to the MAC address table	340
Private VLANs	341
PVLAN configuration guidelines and restrictions	341
Configuring a private VLAN	342
Configuring an interface as a PVLAN promiscuous port	342
Configuring an interface as a PVLAN host port	342
Configuring an interface as a PVLAN trunk port	343
Displaying PVLAN information	343

Chapter 24

Configuring STP-Type Protocols

STP overview	345
Configuring STP	346
Configuration guidelines and restrictions	348
RSTP overview	349
Configuring RSTP	350
MSTP overview	351
Configuring MSTP	352
Overview of PVST+ and Rapid PVST+	353
PVST+ and R-PVST+ guidelines and restrictions	354
Default Spanning Tree Protocol configuration	354

Spanning Tree Protocol configuration and management.	355
Enabling STP, RSTP, MSTP, R-PVST+ or PVST+.	355
Disabling STP, RSTP, or MSTP	355
Shutting down STP, RSTP, or MSTP globally	356
Specifying the bridge priority.	356
Specifying the bridge forward delay	357
Specifying the bridge maximum aging time.	357
Enabling the error disable timeout timer	358
Specifying the error disable timeout interval.	358
Specifying the port-channel path cost	358
Specifying the bridge hello time	359
Specifying the transmit hold count (RSTP, MSTP, and R-PVST+)	360
Enabling Cisco interoperability (MSTP).	360
Disabling Cisco interoperability (MSTP)	360
Mapping a VLAN to an MSTP instance	361
Specifying the maximum number of hops for a BPDU (MSTP)	361
Specifying a name for an MSTP region.	362
Specifying a revision number for MSTP configuration.	362
Clearing spanning tree counters.	362
Clearing spanning tree-detected protocols	363
Displaying STP-related information.	363
Configuring STP, RSTP, or MSTP on DCB interface ports	363
Enabling automatic edge detection	363
Configuring the path cost	364
Enabling a port (interface) as an edge port.	364
Enabling guard root	365
Specifying the STP hello time	366
Specifying restrictions for an MSTP instance	366
Specifying a link type.	367
Enabling port fast (STP).	367
Specifying the port priority	368
Restricting the port from becoming a root port.	369
Restricting the topology change notification	369
Enabling spanning tree	370
Disabling spanning tree.	370
Spanning Tree Protocol and VCS mode.	371
Configuring DiST.	372

Chapter 25

Configuring UDLD

UDLD	373
How UDLD works	373
Configuring UDLD.	374
Other UDLD commands.	375

Chapter 26

Configuring Link Aggregation

Link aggregation overview	377
Link Aggregation Group configuration	378
Link Aggregation Control Protocol	378
Dynamic link aggregation	378
Static link aggregation	378
Brocade-proprietary aggregation	379
LAG distribution process	379
Virtual LAG overview	379
Configuring the vLAG	380
Configuring the vLAG ignore split	381
Configuring load balancing on a remote RBridge	383
LACP configuration guidelines and restrictions	384
Default LACP configuration	384
LACP configuration and management	384
Enabling LACP on a DCB interface	384
Configuring the LACP system priority	385
Configuring the LACP timeout period on a DCB interface	385
Clearing LACP counter statistics on a LAG	386
Clearing LACP counter statistics on all LAG groups	386
Displaying LACP information	386
LACP troubleshooting tips	386

Chapter 27

Configuring LLDP

LLDP overview	389
Layer 2 topology mapping	390
DCBX overview	391
Enhanced Transmission Selection	392
Priority Flow Control	392
LLDP configuration guidelines and restrictions	393
Default LLDP configuration	393
LLDP configuration and management	393
Enabling LLDP globally	393
Disabling and resetting LLDP globally	394
Configuring LLDP global command options	394
Configuring LLDP interface-level command options	401
Clearing LLDP-related information	401
Displaying LLDP-related information	401

Chapter 28

Configuring ACLs

ACL overview	403
Default ACL configuration	404
ACL configuration guidelines and restrictions	404

ACL configuration and management	405
Creating a standard MAC ACL and adding rules	405
Creating an extended MAC ACL and adding rules	406
Applying a MAC ACL to a DCB interface	406
Applying a MAC ACL to a VLAN interface	407
Modifying MAC ACL rules	407
Removing a MAC ACL	408
Reordering the sequence numbers in a MAC ACL	408
IP ACL	408
IP ACL parameters	409
Creating a standard IP ACL	410
Creating an extended IP ACL	411
Applying an IP ACL to a management interface	411
Binding an ACL in standalone mode or fabric cluster mode	412
Displaying the IP ACL configuration	412

Chapter 29

Configuring QoS

Standalone QoS	413
Rewriting	414
Queueing	414
User-priority mapping	414
Traffic class mapping	422
Congestion control	428
Tail drop	428
Configuring CoS thresholds	429
Random Early Discard	430
Ethernet Pause	432
Ethernet Priority Flow Control	434
Multicast rate limiting	435
Creating a receive queue multicast rate-limit	435
BUM storm control	436
Considerations	436
Configuring BUM storm control	436
Scheduling	437
Strict priority scheduling	437
Deficit weighted round robin scheduling	437
Traffic class scheduling policy	438
Multicast queue scheduling	439
Data Center Bridging map configuration	440
Creating a DCB map	442
Defining a priority group table	442
Defining a priority-table map	443
Applying a DCB provisioning map to an interface	443
Verifying the DCB maps	444
Brocade VCS Fabric QoS	444
Configuring Brocade VCS Fabric QoS	444
Restrictions for Layer 3 features in VCS mode	445

	Port-based Policer	445
	Configuring a class map	446
	Configuring a police priority-map	447
	Configuring the policy-map	448
	Attach the mutation to the class	451
	Attach the port-shaper to the class	451
	Attach the scheduler to the class	451
	Attach the priority mapping table to the CEE map [define MQC]	452
	Binding the policy-map to an interface	452
	Policing parameters	453
	Displaying policing settings and policy-maps	456
	Considerations and limitations	457
Chapter 30	Configuring 802.1x Port Authentication	
	802.1x protocol overview	461
	802.1x configuration guidelines and restrictions	461
	802.1x authentication configuration tasks	462
	Configuring authentication	462
	Interface-specific administrative tasks for 802.1x	462
	802.1x readiness check	463
	Configuring 802.1x on specific interface ports	463
	Configuring 802.1x timeouts on specific interface ports	464
	Configuring 802.1x re-authentication on specific interface ports	464
	Configuring 802.1x port-control on specific interface ports	465
	Re-authenticating specific interface ports	465
	Disabling 802.1x on specific interface ports	466
	Disabling 802.1x globally	466
	Checking 802.1x configurations	466
Chapter 31	Configuring sFlow	
	sFlow protocol overview	469
	Interface flow samples	469
	Packet counter samples	470
	Configuring the sFlow protocol globally	470
	Configuring sFlow for interfaces	471
	Enabling and customizing sFlow on specific interfaces	471
	Configuring an sFlow policy map and binding it to an interface	472
	Disabling sFlow on specific interfaces	473
	Hardware support matrix for sFlow	473
	Flow-based sFlow	474
	Configuring flow-based sFlow	474
	Disabling flow-based sFlow on specific interfaces	475
Chapter 32	Configuring Switched Port Analyzer	
	Switched Port Analyzer protocol overview	477
	SPAN guidelines and limitations	477
	Configuring ingress SPAN	478

Configuring egress SPAN	479
Configuring bidirectional SPAN	479
Deleting a SPAN connection from a session.....	480
Deleting a SPAN session.....	480
SPAN in management cluster.....	481
Using RSPAN	482
Mirroring across RSPAN.....	483
Configuring RSPAN.....	483

Chapter 33

Configuring SFP breakout mode

SFP breakout overview	485
Configuring breakout mode (static configuration)	487
For a chassis system	487
For a standalone switch	488
Breakout configuration examples	489
Setting a 40G QSFP port into breakout mode	489
Reserving a 40G QSFP port while in breakout mode	490
Releasing a 40G QSFP port while in breakout mode	491
Breakout limitations	492
High availability considerations	492

Section IV

Network OS Layer 3 Routing Features

Chapter 34

In-band Management

In-band management overview	495
Prerequisites	495
Supported interfaces.....	496
Configuring a standalone in-band management interface	497
Provisioning in standalone mode	497
Configuring an in-band management interface using OSPF	498
Base configuration for a standalone in-band management	499
Base configuration in VCS fabric cluster mode	500

Chapter 35

IP Route Policy

About IP route policy	505
IP prefix-list.....	505
Route-map	506
Configuring IP route policy	506

Chapter 36

IP Route Management

Overview of IP route management.....	509
How IP route management determines best route.....	509

	Configuring static routes.	510
	Specifying the next-hop gateway.	510
	Specifying the egress interface.	510
	Configuring the default route	510
	Other routing commands	511
Chapter 37	Configuring PIM	
	PIM overview	513
	PIM Sparse Mode	514
	Topologies	515
	PIM Sparse device types.	517
	Prerequisites	518
	Standards Conformity	518
	Limitations.	518
	Supportability	519
	Configuration Example	520
Chapter 38	Configuring OSPF	
	Overview of OSPF	525
	OSPF over VRF.	527
	OSPF in a VCS environment	529
	Using designated routers	531
	Performing basic OSPF configuration	532
	Configuration rules	532
	Assigning OSPF areas	532
	Assigning interfaces to an area	537
	Assigning virtual links	537
	Changing other settings	539
	Disabling OSPF on the router.	539
Chapter 39	Configuring VRRP	
	Overview of virtual routers	541
	Guidelines	543
	VRRP/VRRP-E packet behavior	544
	Gratuitous ARP.	544
	VRRP control packets	544
	Source MAC in VRRP control packets.	544
	VRRP basic configuration example	544
	Configuring Router 1 as master for VRRP	545
	Configuring Router 2 as backup for VRRP	545
	VRRP-E differences for basic configuration	546
	Enabling preemption.	546

	Using track ports and track priority with VRRP and VRRP-E	547
	Rules	547
	Track priority example	547
	Using short-path forwarding (VRRP-E only)	547
	Enabling short-path forwarding	549
	Packet routing with short-path forwarding (VRRP-E only)	549
	Multigroup configuration for VRRP/VRRP-E	550
	Configuring a multi-group virtual router cluster	550
	Configuring Router 1 as master for first virtual router group	551
	Configuring Router 1 as backup for second virtual router group . . .	551
	Configuring Router 2 as backup for first virtual router group . . .	552
	Configuring Router 2 as master for second virtual router group . . .	553
Chapter 40	Configuring VRF	
	Overview of VRF	555
	Configuring VRF	556
	Enabling VRRP for VRF	557
	OSPF VRF-Lite for customer-edge routers	558
Chapter 41	Configuring BGP	
	BGP overview	559
	Limitations and considerations	559
	Deployment scenarios	560
	BGP Peering	562
	BGP Attributes	565
	Best-Path Algorithm	565
	Configuration fundamentals	567
	BGP configuration examples	580
	Adjusting defaults to improve routing performance	580
	Using route maps with match and set statements	581
	Clearing configurations	584
Chapter 42	Configuring IGMP	
	IGMP overview	587
	Multicast routing	587
	vLAG and LAG primary port	588
	Configuring IGMP snooping	588
	Configuring IGMP snooping querier	589
	Monitoring IGMP snooping	590
	IGMP scalability	590
	Standalone mode	591
	Brocade VCS Fabric cluster mode	591

Section V Troubleshooting

Chapter 43

Using the Chassis ID (CID) Recovery Tool

Chassis ID card usage	595
Critical SEEPROM data	595
Non-Critical SEEPROM data	596
Automatic auditing and verification of CID card data.....	596
Running the CID recovery tool	596
Data corruption or mismatches	596
CID card failure	597

Chapter 44

Troubleshooting

Troubleshooting overview	599
Gathering troubleshooting information	599
Capturing supportsave data	599
An approach to troubleshooting	600
Understanding troubleshooting hotspots	601
Licensing	602
STP interoperability with Brocade MLX or other switches.....	602
Load balancing distribution.....	603
Static assignment of the routing bridge ID	603
FSPF route change.....	603
VCS Fabric mode and standalone mode	604
vLAG	604
Principal routing bridge availability.....	606
Brocade trunks.....	606
NIC teaming with vLAG	607
Selecting the MTU	607
Avoiding oversubscription	608
ACL limits issues	609

	Troubleshooting procedures	610
	AMPP is not working	610
	Continuous panic reboots	614
	Corrupted CID card	614
	CPU use is unexpectedly high	615
	ECMP not load balancing as expected	616
	ENS functionality check	616
	FCoE devices unable to log in	617
	General debugging for traffic forwarding	619
	ISL does not come up on some ports	620
	License not properly installed	623
	Packets dropped in hardware	624
	Password recovery for the Brocade VDX 8770	630
	Ping failure	633
	QoS configuration causes tail drops	633
	QoS is not marking or treating packets correctly	633
	RBridge ID is duplicated	633
	SNMP MIBs report incorrect values	634
	SNMP traps are missing	634
	Telnet operation into the switch fails	634
	Trunk member not used	635
	Upgrade failure	637
	VCS Fabric cannot be formed	637
	vLAG cannot be formed	638
	Zoning conflict resolution	640
	Zone does not form correctly	640
	Troubleshooting and diagnostic tools	643
	Layer 2 traceroute	644
	show commands	648
	Debug commands	649
	SPAN port and traffic mirroring	650
	Hardware diagnostics	651
	Viewing routing information	652
	Using the packet capture utility	652
Appendix A	TACACS+ Accounting Exceptions	
	Command accounting limitations	655
Appendix B	Supported time zones and regions	
	Africa	659
	America	660
	Antarctica	661
	Arctic	661
	Asia	661
	Atlantic	662
	Australia	662
	Europe	662

Indian	663
Pacific.....	663

Index

Figures

Figure 1	Comparison of classic Ethernet and Brocade VCS Fabric architectures	5
Figure 2	Ethernet fabric with multiple paths	6
Figure 3	Distributed intelligence in an Ethernet fabric	7
Figure 4	Logical chassis in Ethernet fabric	8
Figure 5	Pair of Brocade VDX switches at the top of each server rack	10
Figure 6	Collapsed, flat Layer 3 networks enabling Virtual Machine mobility	11
Figure 7	Brocade VDX 6730 switches deployed as access-level switches	12
Figure 8	Core-edge topology	13
Figure 9	Ring topology	14
Figure 10	Full mesh topology	14
Figure 11	Configuration Database in a logical chassis cluster	37
Figure 12	5-node logical chassis cluster	38
Figure 13	Removal of Node N5 from the logical chassis cluster	45
Figure 14	SNMP structure	112
Figure 15	SNMP query	112
Figure 16	SNMP trap	112
Figure 17	Brocade MIB tree location	113
Figure 18	Metro VCS configuration example	135
Figure 19	Typical deployment topology for Metro VCS using standard ISL	139
Figure 20	Metro VCS and distributed Ethernet fabrics	140
Figure 21	Connecting local VCS Clusters over long-distance using vLAG	142
Figure 22	Zoning	146
Figure 23	LSAN zoning	147
Figure 24	Zone configuration example	168
Figure 25	LSAN zones example	177
Figure 26	FC connection between Network OS fabric and a Fibre Channel SAN	182
Figure 27	Windows server VSA configuration	244
Figure 28	Create new user	262
Figure 29	DH-CHAP authentication	280
Figure 30	Missing LAG causes loop	292
Figure 31	Interconnected Brocade VCS Fabric clusters cause loop	293
Figure 32	Interconnected Brocade VCS Fabric clusters with ELD enabled	294
Figure 33	Port-profile contents	302
Figure 34	FCoE end-to-end header process	315
Figure 35	Multiple switch fabric configuration	317
Figure 36	Ingress VLAN filtering	332

Figure 37	Four-switch example for UDLD.	373
Figure 38	vLAG configuration of the ignore split.	382
Figure 39	Queue depth	428
Figure 40	Strict priority schedule — two queues	437
Figure 41	WRR schedule — two queues.	438
Figure 42	Strict priority and Weighted Round Robin scheduler.	439
Figure 43	A management station with a networked device in standalone mode.	497
Figure 44	In-band management in a VCS fabric with dynamic routes (OSPF)	499
Figure 45	Single VCS deployment	515
Figure 46	Two-tier VCS deployment	516
Figure 47	Example deployment using single VCS	520
Figure 48	OSPF operating in a network	527
Figure 49	OSPF example in a VCS environment	529
Figure 50	Designated and backup router election	531
Figure 51	OSPF network containing an NSSA	535
Figure 52	Defining OSPF virtual links within a network	538
Figure 53	Basic VRRP setup	542
Figure 54	Short path forwarding.	548
Figure 55	Router 1 and 2 configured for dual redundant network access for the host.	550
Figure 56	VRF configuration diagram.	556
Figure 57	Connectivity to the core through an MLX	560
Figure 58	Connectivity to the core without an MLX.	561
Figure 59	BGP with multiple VCS clusters and autonomous systems.	562
Figure 60	BGP route-reflector server and clients	562
Figure 61	IGMP snooping in Brocade VCS Fabric mode.	588
Figure 62	Normal Layer 2 packet traversing a VCS fabric	644
Figure 63	Verifying path continuity with immediate neighbor	645
Figure 64	Verifying path continuity—second hop.	646

Tables

Table 1	Network OS CLI command modes	19
Table 2	Network OS CLI keyboard shortcuts	22
Table 3	CEE CLI command output modifiers	24
Table 4	Mapping switchType to Brocade product names	33
Table 5	Modules supported on the Brocade VDX 8770 platform	56
Table 6	Controlled reload and failover	62
Table 7	Uncontrolled failover	62
Table 8	Standard switch configuration files	73
Table 9	Network OS firmware support by platform	83
Table 10	Licenses requirements by platform	95
Table 11	Brocade licenses for optional Network OS features	96
Table 12	Requirements for activating a license after installation	100
Table 13	Requirements for deactivating a license after removal	102
Table 14	List of available ports when implementing PODs	104
Table 15	Agent Capabilities	115
Table 16	Brocade SNMP MIB dependencies	116
Table 17	Command examples for enabling logical chassis cluster mode	127
Table 18	Command examples for enabling fabric cluster mode	128
Table 19	Command examples when one of the VCS modes is already enabled:	128
Table 20	Virtual IP address configuration scenarios	132
Table 21	ECMP load balancing operands	133
Table 22	Limitations for long distance Metro VCS	136
Table 23	Conditions for long distance Metro VCS	136
Table 24	Long distance Metro VCS port-group schema	137
Table 25	Standard Metro VCS port-group schema	140
Table 26	Approaches to fabric-based zoning	147
Table 27	Considerations for zoning architecture	150
Table 28	Zone merging scenarios: Defined and enabled configurations	172
Table 29	Zone merging scenarios: Different content	173
Table 30	Zone merging scenarios: Different names	174
Table 31	Zone merging scenarios: Default access mode	174
Table 32	Hardware platform default settings for supported switches	192
Table 33	Factory defaults for CPU and memory threshold monitoring	196
Table 34	SFP parameter descriptions	197
Table 35	Factory thresholds for SFP types and monitoring areas	197
Table 36	Security area default settings	199

Table 37	Interface errors that can be monitored on external interfaces	199
Table 38	User account attributes	218
Table 39	Role attributes	222
Table 40	Rule attributes	223
Table 41	Password policy parameters	230
Table 42	RADIUS server parameters	239
Table 43	dictionary.brocade file entries	242
Table 44	TACACS+ server parameters	245
Table 45	AD parameters	256
Table 46	Zeroization behavior	266
Table 47	FIPS-compliant state restrictions	267
Table 48	FIPS-compliant and non-FIPS-compliant states of operation	275
Table 50	Default Ex_Port configuration	280
Table 51	User account attributes	281
Table 52	AMPP scalability values	301
Table 53	AMPP behavior and failure descriptions	303
Table 54	FCoE terminology	314
Table 55	Number of FCoE ports per platform	328
Table 56	Default VLAN configuration	334
Table 57	Private VLAN port types	341
Table 58	STP versus RSTP state comparison	349
Table 59	Default Spanning Tree Protocol configuration	354
Table 60	Default MSTP configuration	354
Table 61	Default 10-gigabit Ethernet DCB interface-specific configuration	354
Table 62	Interface behavior by global and interface configuration	372
Table 63	Load balancing flavors	383
Table 64	Default LACP configuration	384
Table 65	ETS priority grouping of IPC, LAN, and SAN traffic	392
Table 66	Default LLDP configuration	393
Table 67	IP ACL parameters	409
Table 68	Default priority value of untrusted interfaces	415
Table 69	IEEE 802.1Q default priority mapping	415
Table 70	Default DSCP priority mapping	418
Table 71	Default user priority for unicast traffic class mapping	424
Table 72	Default user priority for multicast traffic class mapping	424
Table 73	Pause negotiation results	433
Table 74	Supported scheduling configurations	438
Table 75	Multicast traffic class equivalence mapping	439
Table 76	Default DCB Priority Group Table configuration	441
Table 77	Default DCB priority table	441
Table 78	Policer behavior for L2 and L3 control packets	459
Table 79	sFlow feature support	473

Table 80	Breakout supported platforms.	486
Table 81	SFP breakout values.	486
Table 82	Supported applications for in-band management	495
Table 83	Ports configurable for in-band management	496
Table 84	RFCs supported	518
Table 85	Standalone mode metrics	591
Table 86	Four-node cluster metrics	591
Table 87	Twenty-four-node cluster metrics	592
Table 88	VDX 8770-4 and VDX 8770-8 cluster metrics	592
Table 89	IP multicast metrics	592
Table 90	Load balancing algorithms.	603
Table 91	Port groups	607
Table 92	ACL limits per switch in VCS mode	609
Table 93	Packet header details— Layer 2 packer traverses VCS fabric.	645
Table 94	Packet header details with Layer 2 traceroute—first hop	646
Table 95	Packet header details with Layer 2 traceroute—second hop	647
Table 96	show commands used for troubleshooting	648
Table 97	ASICs and ports.	650
Table 98	Offline diagnostic commands	651
Table 99	Offline diagnostic show commands	652
Table 100	Unsupported commands in privileged EXEC mode	655
Table 101	Unsupported commands in global configuration mode.	657
Table 102	Region/city time zones in Africa	659
Table 103	Region/city time zones in America	660
Table 104	Region/city time zones in Antarctica.	661
Table 105	Region/city time zone in Arctic	661
Table 106	Region/city time zones in Asia.	661
Table 107	Region/city time zones in Atlantic	662
Table 108	Region/city time zones in Australia.	662
Table 109	Region/city time zones in Europe	662
Table 110	Region/city time zones in India	663
Table 111	Region/city time zones in Pacific.	663

About This Document

In this chapter

- [Supported hardware and software](#) xxxv
- [What's new in this document](#) xxxvi
- [Document conventions](#) xxxvi
- [Notice to the reader](#) xxxviii
- [Additional information](#) xxxix
- [Getting technical help](#) xl
- [Document feedback](#) xl

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 Network OS 4.0.0, documenting all possible configurations and scenarios is beyond the scope of this document.

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

- Brocade VDX 6710-54
- Brocade VDX 6720
 - Brocade VDX 6720-24
 - Brocade VDX 6720-60
- Brocade VDX 6730
 - Brocade VDX 6730-32
 - Brocade VDX 6730-76
- Brocade VDX 6740
 - Brocade VDX 6740-48
 - Brocade VDX 6740-64
- Brocade VDX 6740T
 - Brocade VDX 6740T-48
 - Brocade VDX 6740T-64

- Brocade VDX 8770
 - Brocade VDX 8770-4
 - Brocade VDX 8770-8

To obtain information about an OS version other than Network OS v4.0.0, refer to the documentation specific to that OS version.

What's new in this document

This version has been updated to support Network OS v4.0.0, and the new features in this release include:

- Logical Chassis cluster mode
- Metro VCS
- UDLD
- BGP
- High Availability enhancements

For complete information, refer to the Release Notes.

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>code text</code>	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 all lowercase.

Command syntax conventions

Command syntax in this manual follows these conventions:

Convention	Description
[]	Keywords or arguments that appear within square brackets are optional. For example: command [active standby disabled] = One (and only one) of this set of keywords may be used. command [active] [standby] [disabled] = Three independent options, and one or more may be used on the same command line.
{ x y z }	A choice of required keywords appears in braces separated by vertical bars. You must select one. For example: command {active standby disabled} = One (and only one) of this set of keywords/operands must be used.
screen font	Examples of information displayed on the screen.
< >	Nonprinting characters, for example, passwords, appear in angle brackets.
[]	Default responses to system prompts appear in square brackets.
<i>italic text</i>	Identifies variables.
bold text	Identifies literal command options and keywords.

NOTE

In standalone mode, interfaces are identified using *slot/port* notation. In Brocade VCS Fabric mode, interfaces are identified using *rbridge-id/slot/port* notation.

Nesting square brackets and curly brackets

When reading a command entry, optional keywords are surrounded by square brackets and mandatory keywords are surrounded by curly brackets. See [“Command syntax conventions”](#) on page xxxvii for complete details.

In some cases, these brackets can be nested. In this example, rbridge-id is optional as denoted by the square brackets, but if you use it, then you must follow it with either a specific rbridge-id or the word “all.”

Example

```
command [rbridge-id {rbridge-id | all}]
```

However, square brackets can appear within curly brackets, showing that while a keyword is mandatory, supporting operands may be optional, as shown below:

Example

```
command {security [active] [standby] [disabled]}
```

```
command {security [active | standby | disabled]}
```

Command examples

This book describes how to perform configuration tasks using the Network OS command line interface, but does not describe the commands in detail. For complete descriptions of all Network OS commands, including syntax, operand description, and sample output, see the *Network OS Command Reference*.

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 technical glossaries on MyBrocade. See “[Brocade resources](#)” on page xxxix for instructions on accessing MyBrocade.

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
Oracle Corporation	Oracle, Java
Red Hat, Inc.	Red Hat, Red Hat Network, Maximum RPM, Linux Undercover

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> to register at no cost for a user ID and password.

White papers, online demonstrations, and data sheets are available through the Brocade website at:

<http://www.brocade.com/products-solutions/products/index.page>

For additional Brocade documentation, visit the Brocade website:

<http://www.brocade.com>

Release notes are available on the MyBrocade website.

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
- Software name and software version, if applicable
- Error numbers and messages received
- 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 on the switch ID pull-out tab located on the bottom of the port side of the switch.

3. World Wide Name (WWN)

Use the **show license id** command to display the WWN of the chassis.

If you cannot use the **show license id** command because the switch is inoperable, you can get the WWN from the same place as the serial number.

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.

Network OS Administration

This section describes basic Network OS administration features, and includes the following chapters:

- [Introduction to Network OS and Brocade VCS Fabric Technology](#) 3
- [Using the Network OS CLI](#) 17
- [Basic Switch Management](#) 27
- [Network Time Protocol](#) 67
- [Configuration Management](#) 71
- [Installing and Maintaining Firmware](#) 81
- [Administering Licenses](#) 95
- [SNMP](#) 111
- [Fabric](#) 123
- [Metro VCS.](#) 135
- [Configuring PIM](#) 513
- [Administering Zones](#) 145
- [Configuring Fibre Channel Ports.](#) 181
- [System Monitor](#) 191
- [VMware vCenter.](#) 205

Introduction to Network OS and Brocade VCS Fabric Technology

In this chapter

- [Introduction to Brocade Network OS](#) 3
- [Introduction to Brocade VCS Fabric technology](#) 5
- [Brocade VCS Fabric technology use cases](#) 10
- [Topology and scaling](#) 12

Introduction to Brocade Network OS

Brocade Network OS (NOS) is a scalable network operating system available for the Brocade data center switching portfolio products, including the VDX product line. Purpose-built for mission-critical, next generation data centers, Network OS supports the following capabilities:

- Simplified network management

Brocade VCS fabrics are self-forming and self-healing, providing an operationally scalable foundation for very large or dynamic cloud deployments. Multi-node fabrics can be managed as a single logical element, and fabrics can be deployed and easily re-deployed in a variety of configurations optimized to the needs of particular workloads.

Refer to [“Introduction to Brocade VCS Fabric technology”](#) on page 5 for an overview and [Chapter 9, “Fabric,”](#) for detailed information on Brocade VCS Fabric technology.

- High resiliency

Brocade VCS fabrics use hardware-based Inter-Switch Link (ISL) Trunking to provide automatic link failover without traffic interruption.

- Improved network utilization

Transparent Interconnection of Lots of Links (TRILL)-based Layer 2 routing service provides equal-cost multipaths in the network, resulting in improved network utilization. Brocade VCS Fabric technology also delivers multiple active, fully load-balanced Layer 3 gateways to remove constraints on Layer 2 domain growth, eliminate traffic tromboning, and enable inter-VLAN routing within the fabric.

Virtual Router Redundancy Protocol (VRRP) eliminates a single point of failure in a static, default-route environment by dynamically assigning virtual IP routers to participating hosts. The interfaces of all routers in a virtual router must belong to the same IP subnet. There is no restriction against reusing a virtual router ID (VRID) with a different address mapping on different LANs.

Refer to [“TRILL”](#) on page 123 for additional information about TRILL.

Refer to [“Overview of virtual routers”](#) on page 541 for additional information on VRRP/VRRP-E.

- Server virtualization

Automatic Migration of Port Profile (AMPP) functionality provides fabric-wide configuration of network policies, achieves per-port profile forwarding, and enables network-level features to support Virtual Machine (VM) mobility.

Refer to [Chapter 21, “Configuring AMPP”](#) for more information about AMPP.

- Network convergence

Data Center Bridging (DCB)-based lossless Ethernet service provides isolation between IP and storage traffic over a unified network infrastructure. Multi-hop Fibre Channel over Ethernet (FCoE) allows an FCoE initiator to communicate with an FCoE target that is a number of hops away.

Refer to [“End-to-end FCoE”](#) on page 314 for more information about multi-hop FCoE.

In Network OS, all features are configured through a single, industry-standard command line interface (CLI). Refer to the *Network OS Command Reference* for an alphabetical listing and detailed description of all the Network OS commands.

Brocade VCS Fabric terminology

The following terms are used in this document.

Edge ports	In an Ethernet fabric, all switch ports used to connect external equipment, including end stations, switches, and routers.
Ethernet fabric	A topologically flat network of Ethernet switches with shared intelligence, such as the Brocade VCS Fabric.
Fabric ports	The ports on either end of an ISL in an Ethernet fabric.
Inter-Switch Link (ISL)	An interface connected between switches in a VCS fabric. The ports on either end of the interface are called ISL ports or Fabric ports. The ISL can be a single link or a bundle of links forming a Brocade trunk. This trunk can either be created as a proprietary Brocade trunk, or a standard IEEE 802.3ad based link aggregation.
RBridge	A physical switch in a VCS fabric.
RBridge ID	A unique identifier for an RBridge, each switch has a unique RBridge ID. In commands, the RBridge ID is used in referencing all interfaces in the VCS fabric. Refer to “Brocade VCS Fabric configuration management” on page 126 for information about setting the RBridge ID.
VCS ID	A unique identifier for a VCS fabric. The factory default VCS ID is 1. All switches in a VCS fabric must have the same VCS ID.
WWN	World Wide Name. A globally unique ID that is burned into the switch at the factory.

Introduction to Brocade VCS Fabric technology

Brocade VCS Fabric technology is an Ethernet technology that allows you to create flatter, virtualized, and converged data center networks. Brocade VCS Fabric technology is elastic, permitting you to start small, typically at the access layer, and expand your network at your own pace.

Brocade VCS Fabric technology is built upon three core design principles:

- Automation
- Resilience
- Evolutionary design

When two or more Brocade VCS Fabric switches are connected together, they form an Ethernet fabric and exchange information among each other using distributed intelligence. To the rest of the network, the Ethernet fabric appears as a single logical chassis.

Figure 1 shows an example of a data center with a classic hierarchical Ethernet architecture and the same data center with a Brocade VCS Fabric architecture. The Brocade VCS Fabric architecture provides a simpler core-edge topology and is easily scalable as you add more server racks.




FIGURE 1 Comparison of classic Ethernet and Brocade VCS Fabric architectures

Automation

Resilience is a foundational attribute of Brocade Fibre Channel storage networks and resilience is also a requirement in modern data centers with clustered applications and demanding compute Service-Level Agreements (SLAs). In developing its VCS Fabric technology, Brocade naturally carried over this core characteristic to its Ethernet fabric design.

In traditional Ethernet networks running STP, only 50 percent of the links are active; the rest (shown as dotted lines in [Figure 2](#)) act as backups in case the primary connection fails.

When you connect two or more Brocade VCS Fabric mode-enabled switches they form an Ethernet fabric (provided the two switches have a unique RBridgeID and same VCS ID), as shown in [Figure 2](#).




FIGURE 2 Ethernet fabric with multiple paths

The Ethernet fabric has the following characteristics:

- It is a switched network. The Ethernet fabric utilizes an emerging standard called Transparent Interconnection of Lots of Links (TRILL) as the underlying technology.
- All switches automatically know about each other and all connected physical and logical devices.
- All paths in the fabric are available. Traffic is always distributed across equal-cost paths. As shown in [Figure 2](#), traffic from the source to the destination can travel across two paths.
- Traffic travels across the shortest path.
- If a single link fails, traffic is automatically rerouted to other available paths. In [Figure 2](#), if one of the links in Active Path #1 goes down, traffic is seamlessly rerouted across Active Path #2.

- Spanning Tree Protocol (STP) is not necessary because the Ethernet fabric appears as a single logical switch to connected servers, devices, and the rest of the network.
- Traffic can be switched from one Ethernet fabric path to the other Ethernet fabric path.

Distributed intelligence

With Brocade VCS Fabric technology, all relevant information is automatically distributed to each member switch to provide unified fabric functionality, as shown in [Figure 3](#) on page 7.

A Brocade VCS fabric is designed to be managed as a single “logical chassis,” so that each new switch inherits the configuration of the fabric, and the new ports become available immediately. The fabric then appears to the rest of the network as a single switch. This significantly reduces complexity for the management layer, which in turn improves reliability and reduces troubleshooting.

In addition, VCS fabrics provide RESTful and Netconf Application Programming Interfaces (APIs), as well as extensions to OpenStack Quantum to orchestrate both physical and logical networking resources as part of Virtual Machine deployment to support multitiered application topologies.




FIGURE 3 Distributed intelligence in an Ethernet fabric

Distributed intelligence has the following characteristics:

- The fabric is self-forming. When two Brocade VCS Fabric mode-enabled switches are connected, the fabric is automatically created and the switches discover the common fabric configuration.
- The fabric is masterless. No single switch stores configuration information or controls fabric operations. Any switch can fail or be removed without causing disruptive fabric downtime or delayed traffic.
- The fabric is aware of all members, devices, and Virtual Machines (VMs). If the VM moves from one Brocade VCS Fabric port to another Brocade VCS Fabric port in the same fabric, the port-profile is automatically moved to the new port, leveraging Brocade's Auto-Migration Port Profile support.

Logical chassis

All switches in an Ethernet fabric are managed as if they were a single logical chassis. To the rest of the network, the fabric looks no different than any other Layer 2 switch. [Figure 4](#) shows an Ethernet fabric with two switches. The rest of the network is aware of only the edge ports in the fabric, and is unaware of the connections within the fabric.




FIGURE 4 Logical chassis in Ethernet fabric

Each physical switch in the fabric is managed as if it were a blade in a chassis. When a Brocade VCS Fabric mode-enabled switch is connected to the fabric, it inherits the configuration of the fabric and the new ports become available immediately.

Ethernet fabric formation

Brocade VCS Fabric protocols are designed to aid the formation of an Ethernet fabric with minimal user configuration. Refer to [“Brocade VCS Fabric formation”](#) on page 123 for detailed information about the Ethernet fabric formation process.

All supported switches are shipped with Brocade VCS Fabric mode disabled. Refer to [“Brocade VCS Fabric configuration management”](#) on page 126 for information about disabling and enabling Brocade VCS Fabric mode on your switches.

Automatic neighbor discovery

When you connect a switch to a Brocade VCS Fabric mode-enabled switch, the Brocade VCS Fabric mode-enabled switch determines if the neighbor also has Brocade VCS Fabric mode enabled. If the switch has Brocade VCS Fabric mode enabled and the VCS IDs match, the switch joins the Ethernet fabric.

Refer to [“Brocade VCS Fabric configuration management”](#) on page 126 for information about changing the VCS ID.

Automatic ISL formation and hardware-based trunking

When a switch joins an Ethernet fabric, ISLs automatically form between directly connected switches within the fabric.

If more than one ISL exists between two switches, then Brocade ISL trunks can automatically form. All ISLs connected to the same neighboring Brocade switch attempt to form a trunk. The trunks are formed only when the ports belong to the same port group. No user intervention is necessary to form these trunks.

Refer to [“Fabric interface configuration management”](#) on page 128 for information about enabling and disabling ISLs and trunks.

Principal RBridge election

The RBridge with the lowest WWN in the Ethernet fabric is elected as the principal RBridge.

The role of the principal RBridge is to decide whether a new RBridge joining the fabric conflicts with any of the RBridge IDs already present in the fabric. If a conflict arises, the principal RBridge keeps the joining RBridge segmented.

Refer to [“Brocade VCS Fabric configuration management”](#) on page 126 for information about setting the RBridge ID.

Brocade VCS Fabric technology use cases

This section describes the following use cases for Brocade VCS Fabric technology:

- Classic Ethernet
- Large-scale server virtualization

Classic Ethernet Access and Aggregation use case

Brocade VCS Fabric can be deployed in the same fashion as existing top-of-rack switches, as shown in [Figure 5](#). In the right-most two server racks, a two-switch Ethernet fabric replaces the Ethernet switch at the top of each rack.




FIGURE 5 Pair of Brocade VDX switches at the top of each server rack

The servers see a single top-of-rack switch, allowing for active/active connections, end-to-end.

Brocade VCS Fabric technology in this use case provides the following advantages:

- Multiple active-active connections, with increased effective bandwidth
- Preserves existing architecture
- Works with existing core and aggregation networking products
- Co-exists with existing access switches
- Supports 1- and 10-Gbps server connectivity
- Works with server racks or blade servers

Large scale server virtualization use case

Figure 6 shows a logical two-tier architecture with Brocade VCS fabrics at the edge. Each Brocade VCS fabric appears as a single virtual switch to the switches outside the fabric, which results in flattening the network.




FIGURE 6 Collapsed, flat Layer 3 networks enabling Virtual Machine mobility

Brocade VCS Fabric technology in this use case provides the following advantages:

- Optimizes the multipath network (all paths and Layer 3 gateways are active, no single point of failure, and STP is not necessary)
- Increases sphere of Virtual Machine (VM) mobility

Brocade VCS Fabric connectivity with Fibre Channel SAN

Beginning with the Network OS v2.1.1 release, Fibre Channel ports on the Brocade VDX 6730 provide support for connecting a Brocade VCS Fabric to a Fibre Channel SAN. Fibre Channel routers provide the connectivity, which provides access to Fibre Channel devices while preserving isolation between the fabrics. Brocade zoning allows you to determine which FCoE devices can access which storage devices on the Fibre Channel SAN.

Brocade VDX 6730 switches can be deployed into your Brocade VCS Fabric as access-level switches, aggregation-level switches, or as a means of attachment to Brocade VCS Fabric aggregation-level switches. Brocade recommends deployment as access-level switches to minimize congestion issues for storage traffic and isolating FCoE traffic from non-FCoE traffic. [Figure 7](#) shows such a deployment.

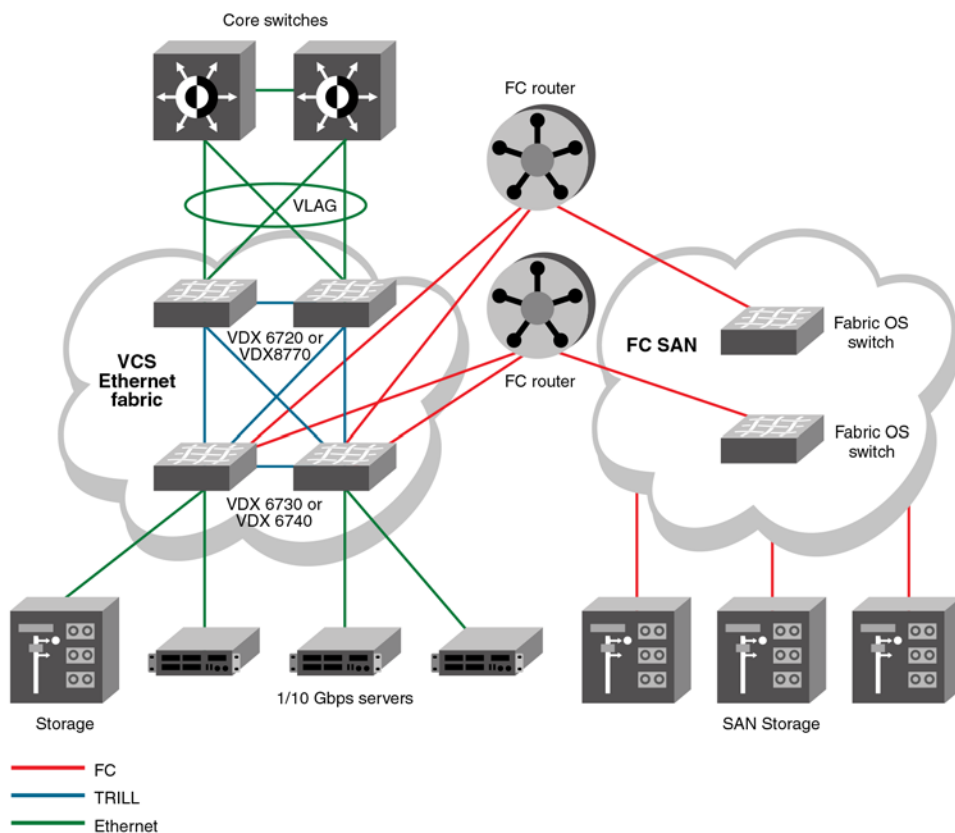


FIGURE 7 Brocade VDX 6730 switches deployed as access-level switches

Topology and scaling

Up to 24 switches can exist in a Brocade VCS Fabric. Although you can use any network topology to build your Brocade VCS Fabric, the following topics discuss the scaling, performance, and availability considerations of topologies more commonly found in data centers:

- [Core-edge topology](#)
- [Ring topology](#)
- [Full mesh topology](#)

Core-edge topology

Core-edge topology, devices connect to edge switches which are connected to each other through core switches. The example shown in [Figure 8](#) uses three core switches. You could use more or fewer switches in the core, depending on whether you need higher availability and greater throughput, or a more efficient use of links and ports.




FIGURE 8 Core-edge topology

This topology is reliable, fast, and scales well. It is reliable because it has multiple core switches. If a core switch or a link to a core switch fails, an alternate path is available. As you increase the number of core switches, you also increase the number of link or core switch failures your cluster can tolerate.

High performance and low latency are assured because throughput is high and the hop count is low. Throughput is high because multiple core switches share the load. Two hops gets you from any edge switch to any other edge switch. If you need greater throughput, simply add another core switch.

Scaling the topology also requires additional core switches and links. However, the number of additional links you need is typically not as great as with, for example, a full mesh topology.

Ring topology

Ring topology connects each node to exactly two other nodes, forming a single continuous pathway. Data travels from node to node, with each node along the path handling every packet of the data.

[Figure 9](#) shows a ring topology.




FIGURE 9 Ring topology

This topology is highly scalable, yet susceptible to failures and traffic congestion. It is highly scalable because of its efficient use of interswitch links and ports; an additional node requires only two ports to connect to the ring. It is susceptible to failures since it provides only one path between any two nodes. Throughput of the fabric is limited by the slowest link or node. Latency can be high because of the potentially high number of hops it takes to communicate between two given switches. This topology is useful where economy of port use is critical, but availability and throughput are less critical.

Full mesh topology

Full mesh topology connects each node to all other cluster nodes. [Figure 10](#) shows a full mesh topology.




FIGURE 10 Full mesh topology

This topology is highly reliable and fast, but does not scale well. It is reliable because it provides many paths through the fabric in case of cable or node failure. It is fast with low latency because you can get to any node in the fabric in just one hop. It does not scale well because each additional node increases the number of fabric links and switch ports exponentially. This topology is suitable for smaller fabrics only.

1 Topology and scaling

Using the Network OS CLI

In this chapter

• DCB command line interface	17
• Saving your configuration changes	17
• Network OS CLI RBAC permissions	18
• Default roles	18
• Accessing the Network OS CLI through Telnet	18
• Network OS CLI command modes	19
• Network OS CLI keyboard shortcuts	22
• Using the do command as a shortcut	22
• Displaying Network OS CLI commands and command syntax	23
• Network OS CLI command completion	24
• Network OS CLI command output modifiers	24
• Show command output information	25

DCB command line interface

The Brocade Data Center Bridging (DCB) CLI is designed to support the management of DCB and Layer 2 Ethernet switching functionality. The Brocade Network OS CLI uses an industry-standard hierarchical shell familiar to Ethernet/IP networking administrators.

The system starts up with the default Network OS configuration and the DCB startup configuration. After logging in, you are in the Network OS shell. For information on accessing the DCB commands from the Network OS shell, see [“Network OS CLI command modes”](#) on page 19.

Saving your configuration changes

Any configuration changes made to the switch are written into the *running-config* file. This is a dynamic file that is lost when the switch reboots. During the boot sequence, the switch resets all configuration settings to the values in the *startup-config* file.

To make your changes permanent, use the **copy** command to commit the *running-config* file to the *startup-config* file, as shown below.

Example of committing the running-config in privileged EXEC mode.

```
switch# copy running-config startup-config
```

Network OS CLI RBAC permissions

Role-Based Access Control (RBAC) defines the capabilities that a user account has based on the role the account has been assigned.

A role is an entity that defines the access privileges of the user accounts on the switch. A user is associated with one role. Refer to [“Role-based access control \(RBAC\)”](#) on page 221 for information about RBAC.

Default roles

Attributes of default roles cannot be modified; however, the default roles can be assigned to non-default user accounts. The following roles are default roles:

- The admin role has the highest privileges. All CLIs are accessible to the user associated with the admin role. By default, the admin role has read and write access.
- The user role has limited privileges that are mostly restricted to show commands in the Privileged EXEC mode. User accounts associated with the user role cannot access configuration CLIs that are in the global configuration mode. By default, the user role has read-only access.

Accessing the Network OS CLI through Telnet

NOTE

While this example uses the **admin** role to log in to the switch, both roles can be used.

The procedure to access the Network OS CLI is the same through either the console interface or through a Telnet session; both access methods bring you to the login prompt.

```
switch login: admin
Password:*****
switch#
```

NOTE

Multiple users can open Telnet sessions and issue commands using the privileged EXEC mode. Network OS v4.0.0 supports up to 32 Telnet sessions with the admin login.

Network OS CLI command modes

Table 1 lists the Network OS CLI command modes and describes how to access them.

NOTE

Use the **pwd** command to view the mode of the current working directory. This command functions in global configuration mode and the modes accessed from global configuration mode.

TABLE 1 Network OS CLI command modes

Command mode	Prompt	How to access the command mode	Description
Privileged EXEC	switch#	This is the default mode for the switch.	Display and change system parameters. Note that this is the administrative mode and includes the basic configuration commands.
Global configuration	switch(config)#	From privileged EXEC mode, enter the configure terminal command.	Configure features that affect the entire switch.
Interface configuration	Port-channel: switch(config-Port-channel-63)# 10-Gigabit Ethernet (DCB port): switch(config-if-te-0/1)# VLAN: switch(config-Vlan-1)#	From global configuration mode, specify an interface by entering one of the following commands: <ul style="list-style-type: none"> interface port-channel interface tengigabitethernet interface VE 	Access and configure individual interfaces.
Protocol configuration	LLDP: switch(config-lldp)# Spanning-tree: switch(config-mstp)# switch(config-rstp)# switch(config-stp)# switch(config-pvst)# switch(config-rpvst)# switch(config-udld)#	From global configuration mode, specify a protocol by entering one of the following commands: <ul style="list-style-type: none"> protocol lldp protocol spanning-tree mstp protocol spanning-tree rstp protocol spanning-tree stp protocol spanning-tree pvst protocol spanning-tree rapid-pvst protocol udld 	Access and configure protocols.
FCoE configuration	FCoE: switch(config-fcoe)# FCoE fabric-map sub-mode: switch(config-fcoe-fabric-map)# FCoE map sub-mode: switch(config-fcoe-map)#	From global configuration mode, use the fcoe command to enter FCoE configuration mode. From FCoE configuration mode, specify an FCoE sub-mode by entering one of the following commands: <ul style="list-style-type: none"> fabric-map default map default 	Access and configure FCoE features.

TABLE 1 Network OS CLI command modes (Continued)

Command mode	Prompt	How to access the command mode	Description
AMPP port-profile mode	AMPP port-profile: switch(config-port-profile-name)# VLAN-profile sub-mode: switch(config-vlan-profile)# QoS-profile sub-mode: switch(config-qos-profile)# FCoE-profile sub-mode: switch(config-fcoe-profile)# Security-profile sub-mode: switch(config-security-profile)#	From global configuration mode, enter the port-profile command to enter port-profile configuration mode. From port-profile configuration mode, specify an AMPP sub-mode by entering one of the following commands: <ul style="list-style-type: none"> • vlan-profile • qos-profile • fcoe-profile • security-profile 	Access and configure AMPP features.
Routing protocol configuration	BGP: switch(config)# switch(config-rbridge-id-1)# switch(config-bgp-router)# BGP route-map configuration mode: switch(config-rbridge-id-1)# switch(config-route-map-myroutemap/permit/1)# BGP address-family IPv4-unicast mode: switch(config-bgp-router)# switch(config-bgp-ipv4u)#	From global configuration mode, specify an RBridge ID to enter RBridge ID configuration mode. From RBridge ID configuration mode, use the router bgp command to enter BGP configuration mode. From RBridge ID configuration mode, use the route-map command with a permit or deny statement and an <i>instance number</i> to enter BGP route-map configuration mode.,. From BGP configuration mode, use the address-family ipv4 unicast command to enter BGP address-family IPv4 unicast configuration mode.	Configure Border Gateway Protocol
Feature configuration	CEE map: switch(config-cee-map-default)# Standard ACL: switch(conf-macl-std)# Extended ACL: switch(conf-macl-ext)#	From global configuration mode, specify a DCB feature by entering one of the following commands: <ul style="list-style-type: none"> • cee-map default • mac access-list standard • mac access-list extended 	Access and configure CEE map features.
RBridge ID configuration mode	RBridge ID: switch (config)# rbridge-id 1 switch (config-rbridge-id-1)#	From global configuration mode, specify a node by entering the rbridge-id rbridge_id command, where <i>rbridge-id</i> is the RBridge ID of the selected node:	Configure features and issue show commands specific to an individual node in a Virtual Cluster Switching (VCS) environment. This includes both fabric cluster and management cluster scenarios.
Hardware configuration mode	switch(config)# hardware	From global configuration mode, specify the hardware node by entering the hardware command.	This mode is a prerequisite for entering connector mode.

TABLE 1 Network OS CLI command modes (Continued)

Command mode	Prompt	How to access the command mode	Description
Connector mode	switch(config-hardware)# connector 2/0/1 switch(config-connector-2/0/1)#	From global configuration mode, enter hardware configuration mode using the following command: switch(config)# hardware	This command is a prerequisite for configuring SFP breakout mode.
DSCP mutation mapping	DSCP Mutation Map: switch(dscp-mutation-mapname)#	From global configuration mode, remap incoming DSCP values using the following command: qos map dscp-mutation mapname	
DSCP to CoS priority mapping	DSCP to CoS Map: switch(dscp-cos-mapname)#	From global configuration mode, create a DSCP to CoS priority map using the following command: qos map dscp-cos mapname	
DSCP to traffic class mapping	DSCP to Traffic Class Map: switch(dscp-traffic-class-mapname)#	From global configuration mode, create a DSCP to traffic class map using the following command: qos map dscp-traffic-class mapname	
QoS Policer configuration	Police Priority Map: switch(config-policemap)# Class Map: switch(config-classmap)# Policy Map: switch(config-policymap)# Policy-class-map submode switch(config-policymap-class)# Policy-class-map-policer attributes submode switch(config-policymap-class-police)#	From global configuration mode, specify a Policer configuration mode by entering one of these command: <ul style="list-style-type: none"> • police-priority-map mapname • class-map mapname • policy-map mapname To enter the policy-class-map sub-mode from the policy-map mode, enter class classmapname . To enter the policy-class-map-policer attributes sub-mode from the policy-map-class mode, enter police followed by the policing attributes.	
Alias configuration	switch(config-alias-config)#	From global configuration mode, enter the alias-config command.	Access configure alias features.
User alias configuration	switch(config-alias-config-user)#	From alias configuration mode, enter the user <name> command.	Access configure user alias features.
Polycymap configuration	switch(config-policymap)#	From global configuration mode, enter the policy-map name command.	
Polycymap class map configuration	switch(config-policymap-class)#	From polycymap mode, enter the class name command.	
Polycymap class police configuration	switch(config-policymap-class-police)#	From polycymap class mode, enter the police cir value command.	
VRF configuration mode	(config-rbridge-l2-vrf-vrf_name)#	From RBridge ID configuration mode, enter the vrf name command.	

NOTE

Pressing **Ctrl+Z** or entering the **end** command in any mode returns you to privileged EXEC mode. Entering **exit** in any mode returns you to the previous mode.

Network OS CLI keyboard shortcuts

Table 2 lists Network OS CLI keyboard shortcuts.

TABLE 2 Network OS CLI keyboard shortcuts

Keystroke	Description
Ctrl+B (or the left arrow key)	Moves the cursor back one character.
Ctrl+F (or the right arrow key)	Moves the cursor forward one character.
Ctrl+A	Moves the cursor to the beginning of the command line.
Ctrl+E	Moves the cursor to the end of the command line.
Esc B	Moves the cursor back one word.
Esc F	Moves the cursor forward one word.
Ctrl+Z	Returns to privileged EXEC mode.
Ctrl+P (or the up arrow key)	Displays commands in the history buffer with the most recent command displayed first.
Ctrl+N (or the down arrow key)	Displays commands in the history buffer with the most recent command displayed last.

NOTE

In privileged EXEC mode, use the **show history** command to list the commands most recently entered. The switch retains the history of the last 1000 commands entered for the current session.

Using the do command as a shortcut

You can use the **do** command to save time when you are working in any configuration mode and you want to run a command in privileged EXEC mode.

For example, if you are configuring LLDP and you want to execute a privileged EXEC mode command, such as the **dir** command, you would first have to exit the LLDP configuration mode. By using the **do** command with the **dir** command, you can ignore the need to change configuration modes, as shown in the following example.

```
switch(conf-lldp)# do dir
Contents of flash://
-rw-r----- 1276 Wed Feb 4 07:08:49 2009 startup_rmon_config
-rw-r----- 1276 Wed Feb 4 07:10:30 2009 rmon_config
-rw-r----- 1276 Wed Feb 4 07:12:33 2009 rmon_configuration
-rw-r----- 1276 Wed Feb 4 10:48:59 2009 starup-config
```

Displaying Network OS CLI commands and command syntax

Enter a question mark (?) in any command mode to display the list of commands available in that mode.

```
switch(conf-lldp)# ?
Possible completions:
advertise      The Advertise TLV configuration.
description    The User description
disable        Disable LLDP
do             Run an operational-mode command
exit           Exit from current mode
hello          The Hello Transmit interval.
help           Provide help information
iscsi-priority Configure the Ethernet priority to advertise for iSCSI
mode           The LLDP mode.
multiplier     The Timeout Multiplier
no             Negate a command or set its defaults
profile        The LLDP Profile table.
pwd            Display current mode path
system-description The System Description.
system-name    The System Name
top            Exit to top level and optionally run command
```

To display a list of commands that start with the same characters, type the characters followed by the question mark (?).

```
switch# e?
Possible completions:
exit  Exit the management session
```

To display the keywords and arguments associated with a command, enter the keyword followed by the question mark (?).

```
switch# terminal ?
Possible completions:
length  Sets Terminal Length for this session
monitor Enables terminal monitoring for this session
no      Sets Terminal Length for this session to default :24.
timeout Sets the interval that the EXEC command interpreter wait for user
        input.
```

If the question mark (?) is typed within an incomplete keyword, and the keyword is the only keyword starting with those characters, the CLI displays help for that keyword only.

```
switch# show d?
Possible completions:
debug  Debug
diag   Show diag related information
dot1x  802.1x configuration
dpod   Provides DPOD license information.
```

If the question mark (?) is typed within an incomplete keyword but the keyword matches several keywords, the CLI displays help for all the matching keywords.

```
switch# show i?
interface Interface status and configuration
ip      Internet Protocol (IP)
```

The Network OS CLI accepts abbreviations for commands. This example is the abbreviation for the **show qos interface all** command.

```
switch# sh q i a
```

If the switch does not recognize a command after **Enter** is pressed, an error message displays.

```
switch# hookup
      ^
syntax error: unknown argument.
```

If an incomplete command is entered, an error message displays.

```
switch# show
      ^
syntax error: unknown argument.
```

Network OS CLI command completion

To automatically complete the spelling of commands or keywords, begin typing the command or keyword and then press **Tab**. For example, at the CLI command prompt, type **te** and press **Tab**:

```
switch# te
```

The CLI displays the following command.

```
switch# terminal
```

If there is more than one command or keyword associated with the characters typed, the Network OS CLI displays all choices. For example, at the CLI command prompt, type **show l** and press **Tab**:

```
switch# show l
```

The CLI displays the following command.

```
Possible completions:
lacp      LACP commands
license   Display license keys installed on the switch.
lldp      Link Layer Discovery Protocol (LLDP).
logging   Show logging
```

Network OS CLI command output modifiers

You can filter the output of the CEE CLI **show** commands using the output modifiers described in [Table 3](#).

TABLE 3 CEE CLI command output modifiers

Output modifier	Description
append	Appends the output to a file.
redirect	Redirects the command output to the specified file.
include	Displays the command output that includes the specified expression.
exclude	Displays the command output that excludes the specified expression.

TABLE 3 CEE CLI command output modifiers (Continued)

Output modifier	Description
begin	Displays the command output that begins with the specified expression.
last	Displays only the last few lines of the command output.
tee	Redirects the command output to the specified file. Note that this modifier also displays the command output.
until <i>string</i>	Ends the output when the output text matches the string.
count	Counts the number of lines in the output.
linnum	Enumerates the lines in the output.
more	Paginates the output.
nomore	Suppresses the pagination of the output.
FLASH	Redirects the output to flash memory.

Show command output information

Network OS contains many versions of the **show** command. The output of the **show** command changes depending on your configuration and situation. However, in general terms the **show** command falls into one of two categories:

- Any **show** commands that are fabric (global configuration) in nature, such as VLAN, MAC Address table, AMPP, Zoning, and so on, should display or clear the information for all nodes in a logical chassis.
- Any **show** commands that are local to a switch, such as Layer 3 or Layer 2 functionality (for example, sFlow, SPAN, and so on), should display the local information by default, and display different switch information specific to an RBridge ID.

2 Show command output information

Basic Switch Management

In this chapter

• Connecting to the switch.....	27
• Switch attributes	31
• Switch types.....	33
• Disabling and enabling a chassis.....	33
• Rebooting a Brocade switch.....	34
• Operational modes	35
• Logical chassis cluster mode	35
• Fabric cluster mode.....	48
• Standalone mode	49
• Configuring the Ethernet management interface	49
• Outbound Telnet and SSH.....	54
• Modular platform basics	55
• Configuring a switch banner.....	62
• supportSave data	63
• Message logging	65

Connecting to the switch

You can connect to your switch through a console session on the serial port, or through a Telnet or Secure Shell (SSH) connection to the management port. You can use any account login present in the local switch database or on a configured authentication, authorization, and accounting (AAA) server for authentication. For initial setup procedures, use the preconfigured administrative account that is part of the default switch configuration.

The switch must 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.

- Refer to the *Brocade VDX Hardware Reference Manuals* for information on connecting through the serial port.
- Refer to “[Configuring the Ethernet management interface](#)” on page 49 for information on configuring the network interface.

Telnet or SSH connections

Telnet and SSH services are enabled by default on the switch. When the Telnet server or SSH server is disabled, access to the switch is not allowed for inbound Telnet or SSH connections, thereby restricting remote access to the switch.

In configuration mode, the CLI can be used to disable Telnet or SSH service on the switch. Doing so will terminate existing inbound Telnet or SSH connections and block any new inbound Telnet or SSH connections to the switch. Additional inbound Telnet or SSH connections will not be allowed until the Telnet server or SSH server is re-enabled. If you have admin privileges, you can re-enable inbound Telnet or SSH connections from configuration mode.

If you are in Logical chassis cluster mode, the command for enabling or disabling Telnet or SSH services is not distributed across the cluster. The RBridge ID of the node should be used to configure the service on individual nodes.

In operational mode, you can use the **show** command to display whether Telnet or SSH is enabled or disabled on the switch.

Note the following:

- Access to the switch is not allowed for both inbound Telnet and SSH connections from both IPv4 and IPv6 addresses when Telnet server and SSH server are disabled.
- Outgoing Telnet or SSH connections from the switch to any remote device is not affected by disabling or enabling the Telnet server or SSH server in the switch.
- No RASLog or auditlog messages are reported when Telnet server or SSH server is disabled or enabled.

SSH server key exchange

SSH allows users to authenticate using public and private key pairs instead of passwords. In password-based authentication, the user must enter a password for authentication purposes. In public-key authentication, the user should have a private key in the local machine and a public key in the remote machine. The user should be logged in to the local machine to be authenticated. If a passphrase is provided while generating the public and private key pair, it must be entered to decrypt the private key while getting authenticated.

SSH key-exchange specifies the method used for generating the one-time session keys for encryption and authentication with the SSH server. A user is allowed to configure the SSH server key-exchange method to DH Group 14. When the SSH server key-exchange method is configured to DH Group 14, the SSH connection from a remote SSH client is allowed only if the key-exchange method at the client is also configured to DH Group 14.

The following steps briefly describe public-key authentication:

1. The user generates a pair of encryption keys in a local machine using the **ssh-keygen** command, along with the public and private key (as shown below). Messages encrypted with the private key can only be decrypted by the public key, and vice-versa.

```
ssh-keygen -t rsa generates RSA public and private keypair
ssh-keygen -t dsa generates DSA public and private keypair
```
2. The user keeps the private key on the local machine, and uploads the public key to the switch.
3. When attempting to log in to the remote host, the user receives an encrypted message from the remote host containing the public key. After decrypting the message in the local host using the private key, the user is authenticated and granted access.

The **ssh-keygen** command is not distributed across the cluster. The RBridge ID of the node should be used to configure service on individual nodes.

Importing an SSH public key

To import an SSH public key, use the **certutil import sshkey** command.

The following example allows you to import the SSH public key for the user “admin” from a remote host using the credentials shown. The command is performed in Privileged EXEC mode.

Standalone mode:

```
switch# certutil import sshkey user admin host 10.70.4.106 directory
/users/home40/bmeenaks/.ssh file id_rsa.pub login fvt
Password: *****
switch# 2012/11/14-10:28:58, [SEC-3050], 75,, INFO, VDX6720-60, Event: sshutil,
Status: success, Info: Imported SSH public key from 10.70.4.106 for user 'admin'.
switch#
```

VCS mode:

```
switch# certutil import sshkey user admin host 10.70.4.106 directory
/users/home40/bmeenaks/.ssh file id_rsa.pub login fvt rbridge-id 3
Password: *****
switch# 2012/11/14-10:28:58, [SEC-3050], 75,, INFO, VDX6720-60, Event: sshutil,
Status: success, Info: Imported SSH public key from 10.70.4.106 for user 'admin'.
switch#
```

To delete an SSH public key, use the **no certutil sshkey user user [rbridge-id {rbridge-id | all}]** command.

Connecting through a Telnet or SSH session

1. Connect through a serial port to the switch.
2. Verify that the switch’s network interface is configured and that it is connected to the IP network through the RJ-45 Ethernet port.
3. Log off the switch’s serial port.
4. From a management station, open a Telnet or SSH connection using the management IP address of the switch to which you want to connect.

For more information on setting the management IP address, refer to [“Configuring the Ethernet management interface”](#) on page 49.

5. Enter the password.

Brocade recommends that you change the default account password when you log in for the first time. For more information on changing the default password, refer to the Brocade *VDX Hardware Reference* manuals.

6. Verify that the login was successful.

The prompt displays the host name followed by a pound sign (#).

```
switch# login as: admin
admin@10.20.49.112's password:*****
```

3 Connecting to the switch

```
-----  
WARNING: The default password of 'admin' and 'user' accounts have not been  
changed.  
Welcome to the Brocade Network Operating System Software  
admin connected from 10.110.100.92 using ssh on VDX 6720-24
```

Connecting through a Telnet or SSH session

```
switch# login as: admin  
admin@10.20.49.112's password:*****
```

```
-----  
WARNING: The default password of 'admin' and 'user' accounts have not been  
changed.  
Welcome to the Brocade Network Operating System Software  
admin connected from 10.110.100.92 using ssh on VDX 6720-24
```

Shutting down Telnet service

By default, the Telnet server is running on the switch.

While in global configuration mode, use the **telnet server shutdown** command to shut down the Telnet service on the switch.

```
switch(config)# telnet server shutdown  
switch(config)#
```

If you are in VCS mode, enter RBridge ID mode.

```
switch(config)# rbridge-id 3  
switch(config-rbridge-id-3)# telnet server shutdown  
switch(config-rbridge-id-3)#
```

Starting Telnet service

To re-enable Telnet service on a switch, while in global configuration mode, use the **no telnet server shutdown** command.

```
switch(config)# no telnet server shutdown  
switch(config)#
```

If you are in VCS mode, enter RBridge ID mode.

```
switch(config)# rbridge-id 3  
switch(config-rbridge-id-3)# no telnet server shutdown  
switch(config-rbridge-id-3)#
```

Shutting down SSH service

By default, the SSH server is running on the switch.

While in global configuration mode, enter the **ssh server shutdown** command to shut down the SSH service on the switch.

```
switch(config)# ssh server shutdown  
switch(config)#
```

If you are in VCS mode, enter RBridge ID mode.

```
switch(config)# rbridge-id 3  
switch(config-rbridge-id-3)# ssh server shutdown  
switch(config-rbridge-id-3)#
```

Starting SSH service

To re-enable SSH service on a switch, enter the following command while in global configuration mode.

```
switch(config)# no ssh server shutdown  
switch(config)#
```

If you are in VCS mode, enter RBridge ID mode.

```
switch(config)# rbridge-id 3  
switch(config-rbridge-id-3)# no ssh server shutdown  
switch(config-rbridge-id-3)#
```

Firmware upgrade and downgrade considerations

Downgrading to a Network OS version earlier than v4.0.0 is not allowed when the Telnet server or SSH server is disabled. To downgrade to a lower version, both the Telnet Server and SSH Server status must be enabled.

Upgrading to Network OS v4.0.0 or later is automatically allowed because the Telnet server and SSH server status are enabled by default.

For more information, refer to [Chapter 6, “Installing and Maintaining Firmware”](#).

Switch attributes

A switch can be identified by its IP address, World Wide Name (WWN), switch ID or RBridge ID, or by its host name and chassis name. You can customize the host name and chassis name with the **switch-attributes** command.

- A host name can be from 1 through 30 characters long. It must begin with a letter, and can contain letters, numbers, and underscore characters. The default host name is “sw0.” The host name is displayed at the system prompt.
- Brocade recommends that you customize the chassis name for each platform. Some system logs identify the switch by its chassis name; if you assign a meaningful chassis name, logs are more useful. A chassis name can be from 1 through 30 characters long, must begin with a letter, and can contain letters, numbers, and underscore characters. The default chassis name is VDX 6710, VDX 6720-24, VDX 6720-60, VDX 6730-32, or VDX 6730-76, VDX 8770-4, or VDX 8770-8 depending on the switch model.

Setting and displaying the host name

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. If Telnet is not activated on the switch, enter the **no telnet server disable** to activate Telnet.

3. Enter the **switch-attributes** command, followed by a question mark (?) to determine the local RBridge ID.
4. Enter the **switch-attributes** command, followed by the RBridge ID.
5. Enter the **host-name** operand, followed by the host name.
6. Save the configuration changes by using the **do copy running-config startup-config** command.

NOTE

This step is used for switches in standalone mode or fabric cluster mode only. If you are using logical chassis cluster mode, startup configurations are not maintained by the cluster; each node preserves its running configuration. For more information about logical chassis cluster mode, refer to [“Logical chassis cluster mode”](#) on page 35.

7. Verify the configuration with the **do show running-config switch-attributes rbridge-id** command.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no telnet server disable
switch(config)# switch-attributes ?
Possible completions:
<NUMBER:1-239> Specify the rbridge-id 1
switch(config)# switch-attributes 1
switch(config-switch-attributes-1)# host-name lab1_vdx0023
switch(config-switch-attributes-1)# exit
switch(config)# do copy running-config startup-config
switch(config)# do show running-config switch-attributes 1
switch-attributes 1
chassis-name VDX 6720-24
host-name lab1_vdx0023
```

Setting and displaying the chassis name

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. Enter the **switch-attributes** command, followed by a question mark to determine the local RBridge ID.
3. Enter the **switch-attributes** command, followed by the RBridge ID.
4. Enter the **chassis-name** operand, followed by the chassis name.
5. Save the configuration changes using the **do copy running-config startup-config** command.

NOTE

This step is used for switches in standalone mode or fabric cluster mode only. If you are using logical chassis cluster mode, startup configurations are not maintained by the cluster; each node preserves its running configuration. For more information about logical chassis cluster mode, refer to [“Logical chassis cluster mode”](#) on page 35.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# switch-attributes ?
Possible completions:
<NUMBER:1-239> Specify the rbridge-id 1
```



```

switch(config)# switch-attributes 1
switch(config-switch-attributes-1# chassis-name lab1_vdx0023
switch(config)# do copy running-config startup-config
switch(config)# do show running-config switch-attributes 1
switch-attributes 1
  chassis-name lab1_vdx0023
  host-name lab1_vdx0023

```

Switch types

The *switchType* attribute is a unique device model identifier that is displayed when you issue the **show chassis** command. When you are gathering information for your switch support provider, you may be asked for the Brocade product name. Use [Table 4](#) to convert the *switchType* identifier to a Brocade product name.

```
switch# show chassis
```

```

Chassis Family:          VDX 87xx
Chassis Backplane Revision: 1

```

```

switchType: 1000 <=== Use table to convert this parameter
(output truncated)

```

In the example, the number 1000 is the value of the *switchType* attribute. An optional number (.x) indicates the revision of the motherboard.

TABLE 4 Mapping *switchType* to Brocade product names

switchType	Brocade product name	Description
95.1 - 95.2	VDX 6720-24	24 1/10 GbE SFP+ ports
96.2	VDX 6730-32	24 1/10 GbE SFP+ ports and 8 8 Gbps FC ports
97.4	VDX 6720-60	60 1/10 GbE SFP+ ports
107.x	VDX 6730-76	60 1/10 GbE SFP+ ports and 16 8 Gbps FC ports
116.2	VDX 6710-54	48 1 GbE copper and 6 1/10 GbE SFP+ ports
1000.x	VDX 8770-4	4 I/O slot chassis supporting 48x10 GbE, or 12x40 GbE interface modules
1001.x	VDX 8770-8	8 I/O slot chassis supporting 48x10 GbE, 72x1 GbE, or 12x40 GbE interface modules

Disabling and enabling a chassis

The chassis is enabled after power is turned on, and diagnostics and switch initialization routines have finished. All interfaces are online. You can disable and re-enable the chassis as necessary.

- Use the **chassis disable** command if you want to take all interfaces offline. If the switch was part of an Ethernet fabric, the fabric reconfigures.
- Use the **chassis enable** command to bring the interfaces back online. All interfaces that were enabled before the chassis was disabled are expected to come back online. If the switch was part of an Ethernet fabric, it rejoins the fabric.

NOTE

Disabling the chassis is a disruptive operation. Use the **shutdown** command to disable or enable a few selected interfaces only. Refer to the *Network OS Command Reference* for more information on this command.

Rebooting a Brocade switch

Network OS provides several commands to reboot your system: **reload**, **fastboot**, **reload system**, and **ha chassisreboot**

NOTE

All reboot operations are disruptive, and the commands prompt for confirmation before executing. When you reboot a switch connected to a fabric, all traffic to and from that switch stops. All ports on that switch remain inactive until the switch comes back online.

Rebooting a compact switch

- The **reload** command performs a “cold reboot” (power off and restart) of the control processor (CP). If the power-on self-test (POST) is enabled, POST is executed when the system comes back up.
- The **fastboot** command performs a “cold reboot” (power off and restart) of the control processor (CP), bypassing POST when the system comes back up. Bypassing POST can reduce boot time significantly.



CAUTION

Do not perform a reload command between a “chassis disable” command and a “chassis enable” command. Your ports will be closed.

Rebooting a modular chassis

- On a modular chassis, the **reboot** and the **fastboot** commands only reboot the management module on which the command is executed. If you log in to the switch IP address and execute one of these commands, only the active management module reboots and POST is bypassed.
- The **ha chassisreboot** command performs a “cold reboot” (power off and restart) of the entire chassis. If the power-on self-test (POST) is enabled, POST is executed when the system comes back up.

A chassis reboot brings up the system in sequential phases. First, software services are launched on the management modules and brought up to the active state. Then, the interface modules are powered on and initialized. Software services are launched on the interface modules and brought up to the active state. When the interface module initialization reaches the final state, the chassis is ready to accept user commands from the CLI interface.

During the boot process system initialization, configuration data (default or user-defined) are applied to the switch through configuration replay. For more information, refer to [“Configuration management in redundant management modules”](#) in [Chapter 5, “Configuration Management”](#).

Operational modes

Network OS supports three operational modes for Brocade VDX switches:

- Logical chassis cluster mode—One of two types of “VCS” modes for a switch. This mode requires Network OS 4.0.0 or later. In this mode, both the data and configuration paths are distributed. The entire cluster is configured from the principal node. Refer to [“Logical chassis cluster mode”](#) on page 35 for more information.
- Fabric cluster mode—The second of two types of “VCS” modes for a switch. In this mode, the data path for nodes is distributed, but the configuration path is not distributed. Each node keeps its configuration database independently. Refer to [“Fabric cluster mode”](#) on page 48 for more information.
- Standalone mode—Only the Brocade VDX 6710-54, 6720, and 6730 support this mode. Refer to [“Standalone mode”](#) on page 49 for more information.

When a new switch boots up, the switch enters either standalone mode or fabric cluster mode, depending on the switch model.

Important terminology note

The generic term *VCS mode* in this manual applies to both fabric cluster mode and logical chassis cluster mode unless otherwise stated.

Logical chassis cluster mode

Logical chassis cluster mode is defined as a fabric in which both the data and configuration paths are distributed. The entire cluster must be globally configured from the principal node. Logical chassis cluster mode requires NOS 4.0.0 or later.

The following platforms support logical chassis cluster mode and is used in any combination:

- Brocade VDX 6710
- Brocade VDX 6720
- Brocade VDX 6730
- Brocade VDX 6740
- Brocade VDX 6740T
- VDX 8770-4
- VDX 8770-8

Some of the main characteristics of logical chassis cluster mode include:

- The maximum number of nodes supported in a logical chassis cluster is 24.
- Physical connectivity requirements for logical chassis cluster deployment are the same as those for fabric cluster deployment.
- A single global configuration exists across all nodes, while each node can contain its unique local configuration. However, each node contains the local configuration information for all other nodes in the cluster.
- Global and local configurations for the entire logical chassis cluster is performed from one node—the principal node only.

3 Logical chassis cluster mode

- Startup configurations are not maintained by the cluster; each node preserves its running configuration.
- A logical chassis cluster can be transitioned into a fabric cluster while preserving configurations if you follow the steps provided later in this section
- An existing fabric cluster can be transitioned into a logical chassis cluster while preserving configurations if you follow the steps provided later in this section
- A node that is a member of a logical chassis cluster can be transitioned to standalone mode. Platforms that allow standalone mode are the Brocade VDX 6710-54, 6720, and 6730.
- Cluster-wide firmware upgrades can be performed.
- Cluster-wide supportsave can be performed.

Topics covered

The following topics are covered in this section:

- [“Configuration”](#) on page 36
- [“Creating a logical chassis cluster”](#) on page 37
- [“Precautions to take for mode transitions”](#) on page 39
- [“Converting a fabric cluster to a logical chassis cluster”](#) on page 40
- [“Converting a fabric cluster while preserving configuration”](#) on page 41
- [“Selecting a principal node for the cluster”](#) on page 43
- [“Converting a logical chassis cluster to a fabric cluster”](#) on page 43
- [“Converting to a fabric cluster while preserving configuration”](#) on page 44
- [“Adding a node to a logical chassis cluster”](#) on page 44
- [“Removing a node from a logical chassis cluster”](#) on page 44
- [“Rejoining a node to the cluster”](#) on page 45
- [“Replacing a node in a logical chassis cluster”](#) on page 45
- [“Merging two logical chassis clusters”](#) on page 46
- [“Examples of global and local configurations”](#) on page 47
- [“Additional operations”](#) on page 48

Configuration

In logical chassis cluster mode, any operation that results in writing to the configuration database gets automatically distributed. There are no exceptions.

Each node in the logical chassis cluster maintains an individual copy of the configuration to enable high availability of the cluster. [Figure 11](#) represents nodes in a logical chassis cluster. Each node has its own databases, and the databases kept by each node are identical at all times.

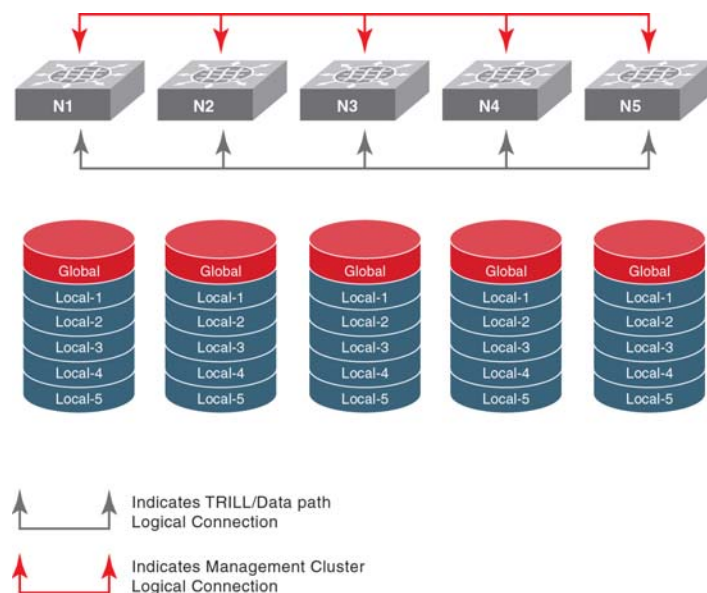


FIGURE 11 Configuration Database in a logical chassis cluster

Network OS switches contain both global and local configuration. In a logical chassis cluster, a single global configuration exists across all cluster members, while each individual member has its own local configuration. (Conversely, in fabric cluster mode, each cluster member can have its own unique global configuration.)

Global configuration is required for cluster-wide operations, whereas local configuration is specific to the operation of an individual node. For more information and examples of each type of configuration, see [“Examples of global and local configurations.”](#)

Creating a logical chassis cluster

This section covers the basic steps to create a logical chassis cluster, with the assumption that all physical connectivity requirements have been met. [Figure 12](#) is a representation of a five-node logical chassis cluster.

To create a logical chassis cluster, follow the steps in the example below:

1. Log into one switch that will be a member of the logical chassis cluster you are creating:
2. In Privileged EXEC mode, enter the **vcs** command with options to set the VCD ID, the RBridge ID and enable logical chassis mode for the switch. The VCS ID and RBridge IDs shown below are chosen for the purposes of this example.

```
switch# vcs vcsid 22 rbridge-id 15 logical-chassis enable
```

3. The switch reboots after you run the **vcs** command. You are asked if you want to apply the default configuration; answer **yes**.

3 Logical chassis cluster mode

4. Repeat the above steps for each node in the cluster, changing only the RBridge ID each time. You must, however, set the VCS ID to the same value on each node that belongs to the cluster.
5. When you have enabled the logical chassis mode on each node in the cluster, run the **show vcs** command to determine which node has been assigned as the cluster principal node. The arrow (>) denotes the principal node. The asterisk (*) denotes the current logged-in node.

```
switch# show vcs
Config Mode      : Distributed
VCS ID           : 1
VCS GUID         : 86024da1-b2c2-4b35-955d-41c27598aaa0
Total Number of Nodes : 2
Rbridge-Id       WWN                               Management IP   Status      HostName
-----
154              >10:00:00:05:33:51:63:42*   10.17.37.154    Online      switch
                2607:f0d0:1002:ff51:ffff:ffff:ffff:fff5
165              10:00:00:05:33:B7:F0:00    10.17.37.165    Online
```

The RBridge ID with the arrow pointing to the WWN is the cluster principal. In this example, RBridge ID 154 is the principal.

6. Set the clock and time zone for the principal node. Time should be consistent across all the nodes. Refer to [Chapter 4, “Network Time Protocol”](#).
7. Log in to the principal cluster and make any desired global and local configuration changes. These changes then are distributed automatically to all nodes in the logical chassis cluster.

NOTE

You can enter the RBridge ID configuration mode for any RBridge in the cluster from the cluster principal node.

NOTE

You can change the principal node by using the **logical-chassis principal priority** and **logical chassis principal switchover** commands. For more information about cluster principal nodes, see [“Selecting a principal node for the cluster”](#) on page 43.

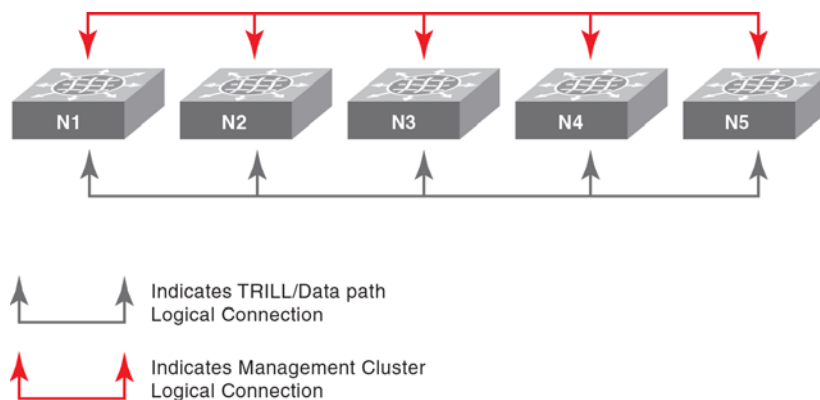


FIGURE 12 5-node logical chassis cluster

Precautions to take for mode transitions

Be sure all nodes that will be transitioned are running the same version of Network OS. Logical chassis cluster mode is supported starting with Network OS release 4.0.0

If you are merging multiple global configuration files to create one new global configuration file, be sure that the same entity name does not exist in the merged file. For example, if mac access-list extended **test1** contains the entries shown in the “Node 1 global configuration” and “Node 2 global configuration” below, when you merge the files you can rename mac access-list extended **test1** from Node 2 to mac access-list extended **test2**, as shown in the “Combined global configuration.”

Node 1 global configuration

```
mac access-list extended test1
seq 10 permit any 1111.2222.333a ffff.ffff.ffff
seq 20 deny any 1111.2222.333b ffff.ffff.ffff
seq 30 deny any 1111.2222.333c ffff.ffff.ffff
seq 40 permit any any
```

Node 2 global configuration

```
mac access-list extended test1
seq 10 permit any 4444.5555.666d ffff.ffff.ffff
seq 20 deny any 4444.5555.666e ffff.ffff.ffff
seq 30 permit any any
```

Combined global configuration

```
mac access-list extended test1
seq 10 permit any 1111.2222.333a ffff.ffff.ffff
seq 20 deny any 1111.2222.333b ffff.ffff.ffff
seq 30 deny any 1111.2222.333c ffff.ffff.ffff
seq 40 permit any any
!
mac access-list extended test2
seq 10 permit any 4444.5555.666d ffff.ffff.ffff
seq 20 deny any 4444.5555.666e ffff.ffff.ffff
seq 30 permit any any
```

The local configuration for Node 2 also needs to be changed accordingly. In this example, one of the local configuration changes would be the interface TenGigabitEthernet. Instead of referencing **test1**, the local configuration file for Node 2 needs to reference **test2** because of the change that was made to the global configuration file. This is shown in the “Node 2 local configuration...” sections below.

Node 2 local configuration before matching the combined global configuration

```
interface TenGigabitEthernet 4/0/3
fabric isl enable
fabric trunk enable
switchport
switchport mode access
switchport access vlan 1
spanning-tree shutdown
mac access-group test1 in
no shutdown
```

Node 2 local configuration after matching the combined global configuration

```
interface TenGigabitEthernet 4/0/3
fabric isl enable
fabric trunk enable
switchport
switchport mode access
switchport access vlan 1
spanning-tree shutdown
mac access-group test2 in
no shutdown
```

Be sure to also take the following precautions:

- Note that the command **copy default-config to startup-config** in logical chassis cluster mode causes a cluster-wide reboot and returns the entire logical chassis cluster to the default configuration. Therefore, use this command only if you want to purge all existing configuration in the logical chassis cluster.
- Make sure that the backup files for global and local configurations are available in a proper scp or ftp location which can be easily retrieved in logical chassis cluster mode during restore. Do not save the files in the local flash because they may not be available on the principal node for replay of local configurations.

Converting a fabric cluster to a logical chassis cluster

This section covers the basic steps to converting an existing fabric cluster to a logical chassis cluster using the default configuration file. Follow the steps in the example below:

1. Be sure all nodes are running the same firmware version. Logical chassis cluster is supported starting on Network OS version 4.0.0.
2. Be sure all the nodes that you intend to transition from a fabric cluster to a logical chassis cluster are online. Run either the **show vcs** or **show vcs detail** command to check the status of the nodes.
3. Log into one switch that you are converting from fabric cluster mode to logical chassis cluster mode.
4. In Privileged EXEC mode, enter the **vcs logical-chassis enable** command with desired options; for example you can convert all R Bridges with one command:

```
switch# vcs logical-chassis enable rbridge-id all default-config
```

NOTE

To convert a specific RBridge from fabric cluster mode to logical chassis mode, use the RBridge ID value in place of the “all” option. You can also specify a range, such as “1,3,4-6”. Refer to the *NOS Command Reference* for details.

The nodes automatically reboot in logical chassis cluster mode. Allow for some down time during the mode transition.

5. Run either the **show vcs** or **show vcs detail** command to check that all nodes are online and now in logical chassis cluster (listed as “Distributed” in the command output) mode.

- The **show vcs** command output can also be used to determine which node has been assigned as the cluster principal node.

```
switch# show vcs
```

R-Bridge	WWN	Switch-MAC	Status
1	>11:22:33:44:55:66:77:81	AA:BB:CC::DD:EE:F1	Online
2	11:22:33:44:55:66:77:82	AA:BB:CC::DD:EE:F2	Online
3	11:22:33:44:55:66:77:83*	AA:BB:CC::DD:EE:F3	Online

The RBridge ID with the arrow pointing to the WWN is the cluster principal. In this example, RBridge ID 1 is the principal.

- Log in to the principal cluster and make any desired global and local configuration changes. These changes then are distributed automatically to all nodes in the logical chassis cluster.

NOTE

You can enter the RBridge ID configuration mode for any RBridge in the cluster from the cluster principal node.

NOTE

You can change the principal node by using the **logical-chassis principal priority** and **logical chassis principal switchover** commands. For more information about cluster principal nodes, see [“Selecting a principal node for the cluster”](#) on page 43.

Converting a fabric cluster while preserving configuration

There is no specific command that can convert a fabric cluster to a logical chassis cluster while preserving current configurations, but you can accomplish this task as follows:

- Be sure all nodes are running the same firmware version. Logical chassis cluster is supported starting on Network OS version 4.0.0.
- Make sure all the nodes that you intend to transition from a fabric cluster to a logical chassis cluster are online. Run either the **show vcs** or **show vcs detail** command to check the status of the nodes.
- Determine which node contains the global configuration you want to use on the logical chassis cluster, and make a back up of this configuration by running the **copy global-running-config** command and saving the configuration to a file on a remote ftp, scp, sftp or usb location:

NOTE

If you need to combine the global configurations of two or more nodes, manually combine the required files into a single file which will be replayed after the transition to logical chassis cluster mode. Refer to the section [“Precautions to take for mode transitions.”](#)

```
copy global-running-config location_configfilename
```

- Back up the local configurations of all individual nodes in the cluster, by running the **copy local-running-config** command on each node and saving the configuration to a file on a remote ftp, scp, sftp, or usb location:

```
copy local-running-config location_configfilename
```

- 5. Perform the mode transition from fabric cluster to logical chassis cluster by running the **vcs logical-chassis enable rbridge-id all default-config** command, as shown in [“Converting a fabric cluster to a logical chassis cluster”](#) on page 40.

The nodes automatically reboot in logical chassis cluster mode. Allow for some down time during the mode transition.

- 6. Run either the **show vcs** or **show vcs detail** command to check that all nodes are online and now in logical chassis cluster (listed as “Distributed” in the command output) mode.
- 7. The **show vcs** command output can also be used to determine which node has been assigned as the cluster principal node.

```
switch# show vcs
```

R-Bridge	WWN	Switch-MAC	Status
1	>11:22:33:44:55:66:77:81	AA:BB:CC::DD:EE:F1	Online
2	11:22:33:44:55:66:77:82	AA:BB:CC::DD:EE:F2	Online
3	11:22:33:44:55:66:77:83*	AA:BB:CC::DD:EE:F3	Online

The RBridge ID with the arrow pointing to the WWN is the cluster principal. In this example, RBridge ID 1 is the principal.

- 8. While logged on to the principal node in the logical chassis cluster, copy the saved global configuration file from the remote location to the principal node as follows:

copy location_configfilename running-config

- 9. Verify that the global configuration is available by running the **show global-running-config** command.
- 10. While logged on to the principal node in the logical chassis cluster, copy each saved local configuration file from the remote location to the principal node as follows:

copy location_configfilename running-config

NOTE
You must run this command for each local configuration file you saved (one for each node).

The configuration file is automatically distributed to all nodes in the logical chassis cluster. Each node will contain the same global configuration after the above steps are performed. Each node will also contain the local configuration information of all the other nodes.

- 11. Verify that the local configurations are available by running the **show local-running-config** command.
- 12. Log in to the principal cluster and make any desired global and local configuration changes. These changes then are distributed automatically to all nodes in the logical chassis cluster.

NOTE
You can enter the RBridge ID configuration mode for any RBridge in the cluster from the cluster principal node.

NOTE

You can change the principal node by using the **logical-chassis principal priority** and **logical chassis principal switchover** commands. For more information about cluster principal nodes, see [“Selecting a principal node for the cluster”](#) on page 43.

Selecting a principal node for the cluster

Logical chassis cluster principal node behavior includes:

- All configuration for the logical chassis cluster must be performed on the principal node.
- By default, the node with the lowest WWN number becomes the principal node.
- You can run the **show vcs** command to determine which node is the principal node. An arrow in the display from this command points to the WWN of the principal node.
- You can select any node in the logical chassis cluster to become the principal by running the **logical chassis principal priority** command, followed by the **logical-chassis principal switchover** command, as shown in the following example (in this example, RBridge ID 5 is being assigned with the highest priority):

```
switch# configure
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# logical-chassis principal-priority 1
switch(config-rbridge-id-5)# end
switch# logical-chassis principal-switchover
```

A lower number means a higher priority. Values range from 1 to 128.

Until you run the **logical-chassis principal switchover** command, the election of the new principal node does not take effect.

Converting a logical chassis cluster to a fabric cluster

To transition all nodes in a logical chassis cluster to a fabric cluster, using default configurations, perform these steps:

1. Make sure all the nodes that you intend to transition from a logical chassis cluster to a fabric cluster are online. Run either the **show vcs** or **show vcs detail** command to check the status of the nodes.
2. Log in to the principal node on the logical chassis cluster.
3. Run the following command to convert all RBridge IDs:

```
no vcs logical-chassis enable rbridge-id all default-config
```

NOTE

To convert just one RBridge ID, specify the ID as shown in the following example:

```
no vcs logical-chassis enable rbridge-id rbridge-id default-config
```

The nodes automatically reboot in fabric cluster mode. Plan for some down time for this transition.

4. Run either the **show vcs** or **show vcs detail** command to check that all nodes are online and now in fabric cluster (listed as “Local-only” in the command output) mode.

Converting to a fabric cluster while preserving configuration

There is no specific command that can convert a logical chassis cluster to a fabric cluster while preserving current configurations, but you can accomplish this task as follows:

1. Make sure all the nodes that you intend to transition from a logical chassis cluster to a fabric cluster are online. Run either the **show vcs** or **show vcs detail** command to check the status of the nodes.
2. Back up the configurations of all nodes in the cluster by running the **copy rbridge-running-config rbridge-id** command on each node and saving the configuration to a file on a remote ftp, scp, sftp, or usb location:

```
copy rbridge-running-config rbridge-id rbridge-id location_configfilename
```

This command copies both the global and local configurations for the specified Rbridge ID.

3. From the principal node of the logical chassis cluster, transition the entire cluster to fabric cluster mode (using the default configuration) by running the following command:

```
no vcs logical-chassis enable rbridge-id all default-config
```

The nodes automatically reboot in fabric cluster mode. Plan for some down time for this transition.

4. Run either the **show vcs** or **show vcs detail** command to check that all nodes are online and now in fabric cluster (listed as “Local-only” in the command output) mode.
 5. Restore the global and local configurations on each individual node for which you backed up these configurations by running the following command on each node:
- ```
copy location_configfilename running-config
```
6. To persist this downloaded configuration for a node, run the **copy running-config startup-config** command.

## Adding a node to a logical chassis cluster

Nodes can be dynamically added to an existing logical chassis cluster. If the proper physical connections exist between the existing logical chassis cluster and the new node, the process is automatic.

Log into the new node and run the **vcs logical-chassis enable** command with the desired options. You must assign the new node the VCS ID of the existing cluster.

You can run the **show vcs** command to verify that the status of the added node is “online.”

## Removing a node from a logical chassis cluster

If the **no vcs enable** command is executed on a switch that is currently in logical chassis cluster mode, the switch boots in standalone mode. (Only the Brocade VDS 6710-54, 6720, and 6730 support standalone mode.) If the **no vcs logical-chassis enable** command is executed on a switch that is currently in logical chassis cluster mode, the switch boots in fabric cluster mode.

Once the node is removed, all configurations corresponding to that node are removed from the cluster configuration database. Similarly, the removed node does not retain any configurations corresponding to the other nodes in the cluster.

Figure 13 depicts the cluster after node N5 has been removed. Nodes N1 through N4 remain in the cluster, and N5 is an island. There is no data path or management path connectivity between the two islands.

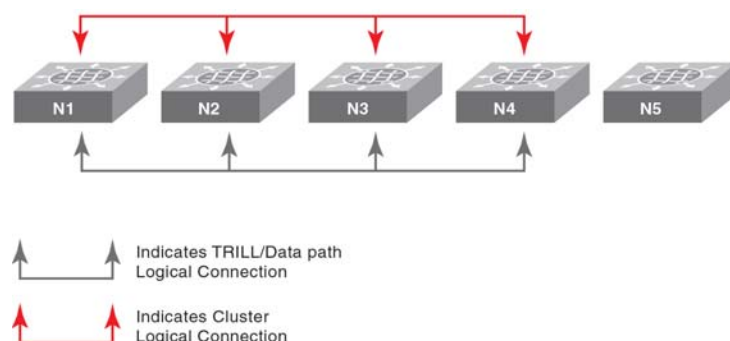


FIGURE 13 Removal of Node N5 from the logical chassis cluster

## Rejoining a node to the cluster

Nodes that are temporarily isolated from a logical chassis cluster can re-join the cluster as long as no configuration or cluster membership changes have taken place on either the deleted node or the cluster. Run the **vcs logical-chassis enable** command with the desired options to rejoin the node to the cluster.

However, if configuration changes have occurred on either the node or cluster since the node was removed, you must reboot the node with its default configuration by issuing **copy default-config startup-config** on the segmented node.

## Replacing a node in a logical chassis cluster

If a node in a logical chassis cluster becomes damaged and no longer be used, a similar node with identical capabilities can be used in its place.

The new node must use the same RBridge ID of the node that is being replaced. When the new node is detected, it joins the cluster as a previously known node instead of being considered a new node.

To replace a node that has an Rbridge ID of 3 and then enter the WWN of the new node, follow the steps shown in the following example:

1. Run the following command on the principal:

```
switch# vcs replace rbridge-id 3
Enter the WWN of the new replacement switch: 11:22:33:44:55:66:77:81
```

2. Assign the RBridge ID of 3 to the new node by running the following command on the new node, assuming the new node is already VCS enabled:

```
switch# vcs rbridge-id 3
```

### NOTE

If the new node is not yet VCS enabled, you can do so at the same time you assign the RBridge ID. Refer to the **vcs** command options in the *Network OS Command Reference*.

## Merging two logical chassis clusters

You can merge two logical chassis clusters that have the same VCS ID. Follow these steps:

1. Make all required physical connections between the two independent clusters.
2. Decide which cluster should retain the configuration after the merge. Only one configuration can be retained.
3. On the cluster whose configuration will not be retained, issue the **copy default-config startup-config** command so that the nodes in this cluster will rebooted with the default configuration.
4. Reboot all nodes in each cluster. The logical chassis cluster whose configuration is being retained recognizes the nodes from the other cluster as new nodes and adds them accordingly.
5. Re-apply the configuration to the cluster whose configuration was not retained.

## Changing an rbridge-id on a switch within a fabric

It may become necessary to change the rbridge-id number on a switch that rebooted and has become orphaned from the cluster.

1. Backup the global configuration before changing the rbridge-id as the local configuration will be reset to default. Refer to [“Configuration backup”](#) on page 76.
2. On the rebooted switch, execute the **chassis disable** command.  
  
switch# **chassis disable**
3. From the fabric principal switch, execute the **no vcs enable rbridge-id rbridge-id** command, where *rbridge-id* is the switch that was orphaned.  
  
switch# **no vcs enable rbridge-id 3**
4. On the rebooted switch, execute the **vcs rbridge-id rbridge-id** command, where *rbridge-id* is the rbridge you want to use.
5. The VCSID should already be set, if it's not set it with the **vcs rbridge-id rbridge-id**.
6. Reboot the orphaned switch.

Note: after the switch reboots

- All interfaces will be in 'shutdown' state. You must perform a **no shutdown** command on ISL interfaces before the switch will rejoin the cluster.
- The original configuration will be lost and the switch will have a default configuration when it rejoins the cluster with the new rbridge-id.

7. Use the **show vcs detail** command to verify the switch is in the fabric.

```
switch# show vcs detail
Config Mode : Local-Only
VCS ID : 1
Total Number of Nodes : 6
Node :1
Serial Number : BKN2501G00R
Condition : Good
Status : Connected to Cluster
VCS Id : 1
Rbridge-Id : 38
```

```
Co-ordinator : NO
WWN : 10:00:00:05:33:52:2A:82
Switch MAC : 00:05:33:52:2A:82
FCF MAC : 0B:20:B0:64:10:27
Switch Type : BR-VDX6720-24-C-24
Internal IP : 127.1.0.38
Management IP : 10.17.10.38
Node :2
Serial Number : BZA0330G00P
```

## Examples of global and local configurations

To view global configuration, run the **show global-running-config** command. Examples of global parameters are:

- Interface vlan
- interface Port-channel
- port-profile
- mac access-list
- ip access-list
- sflow
- snmp-server
- protocol lldp
- protocol udld
- Zoning
- cee-map
- username

To view local configuration, run the **show local-running-config** command. Examples of local parameters are:

- switch-attributes
- interface Management
- interface ve
- diag post
- dpod
- switch-attributes
- fabric route mcast
- rbridge-id
- ip route
- linecard
- router ospf
- router bgp
- protocol vrrp
- vrrp-group
- interface Management

- interface GigabitEthernet
- interface TenGigabitEthernet
- interface FortyGigabitEthernet
- interface Fcoe

## Additional operations

Use the **copy snapshot** commands if you need to upload or download configuration snapshot files to and from an ftp or scp server. You may need to use these commands if you took a snapshot of a configuration on a node that was disconnected from the cluster.

Refer to the *Network OS Command Reference* for information about these and other logical chassis server commands.

## Fabric cluster mode

Fabric cluster mode is defined as a fabric in which the data path for nodes is distributed, but the configuration path is not distributed. Each node keeps its configuration database independently.

By default, the following platforms boot up in fabric cluster mode and will attempt to form interswitch links:

- Brocade VDX 8770-4
- Brocade VDX 8770-8
- Brocade 6740
- Brocade 6740T

If the chassis is not connected to another switch, it forms a “single node VCS fabric.” This means the chassis operates as a standalone system, but the operational mode is always VCS-enabled. You cannot disable the VCS mode on any of the models listed above.

When you issue the **show vcs** command to display the VCS configuration for the chassis, the command output shows a single-node VCS with a VCS ID of 1 and an RBridge ID of 1. Use the **vcs** command to change the default values.

```
switch# show vcs
Config Mode : Local-Only
VCS ID : 1
Total Number of Nodes : 2
Rbridge-Id WWN Management IP Status HostName

1 10:00:00:05:33:15:DE:CC 10.24.82.120 Online dutA1-sw0
 fd00:60:69bc:64:205:33ff:fe15:decc
```



## Standalone mode

All Brocade VDX compact switches (VDX 6710, VDX 6720, and VDX 6730) boot up in standalone (SA) mode. In this restricted mode, the switch supports only legacy features that were available in Network OS v2.1.1, with the exception of IP static routes and in-band management. All other Layer 3 features, or any other features introduced in Network OS v3.0.0, are not available in standalone mode.

---

**NOTE**

The Brocade VDX 8770, 6740, and 6740T do not support standalone mode.

---

When you display the VCS configuration for a compact switch in default mode, it shows that VCS mode is disabled.

```
switch# show vcs
state: Disabled
```

## Configuring the Ethernet management interface

The Ethernet network interface provides management access, including direct access to the Network OS CLI. You must configure at least one IP address using a serial connection to the CLI before you can manage the system with other management interfaces. You can either configure static IP addresses, or you can use a Dynamic Host Configuration Protocol (DHCP) client to acquire IP addresses automatically. For IPv6 addresses, both static IPv6 and stateless IPv6 autoconfiguration are supported.

Setting static IP addresses and using DHCP are mutually exclusive. If DHCP is enabled, remove the DHCP client before you configure a static IP address.

---

**NOTE**

You must connect through the serial port to set the IP address if the network interface is not configured already. Refer to the Brocade VDX *Hardware Reference* manual for your specific product for information on connecting through the serial port.

---

### Ethernet interfaces

The Brocade VDX compact switches have a single configurable Ethernet Interface, Eth0, which can be configured as a management interface.

The modular chassis, the Brocade VDX 8770-8 and the Brocade VDX 8770-4, have two redundant management modules, MM1 and MM2. Each management module can communicate with each of the interface modules (line cards) through an Ethernet connection. Each management module has two Ethernet interfaces, Eth0 and Eth2.

Eth0 is the management interface and can be configured with an IP address. Eth2 provides connectivity to the other management module and the interface modules in the chassis. The Eth2 IP addressing scheme uses default IP addresses to communicate between the modules; these addresses are not user-configurable.

## Configuring a static IP address

Use static Ethernet network interface addresses in environments where the DHCP service is not available. To configure a static IPv4 or IPv6 address, you must first disable DHCP. Refer to [“Configuring an IP address with DHCP”](#) on page 51 for more information.

### *Configuring a static IPv4 Ethernet address*

1. Connect to the switch through the serial console.
2. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
3. Enter the **interface management** *rbridge-id/port* command to configure the management port.

This command enters a management interface configuration mode where you can choose configuration parameters for IPv4 and IPv6 addresses.

- A compact switch has a single management port, and the port number for the management port is always 0.
  - On a modular switch with two redundant management modules, you can configure two management ports. The port numbers are 1 and 2.
4. Enter the **no ip address dhcp** command to disable DHCP.
  5. Enter the **ip address** *IPv4\_address/prefix\_length* command.
  6. Use the **ip route** *0.0.0.0/0 gw-ip* command to configure the gateway address.
  7. Verify the configuration with the **do show running-config interface management** command.

---

#### NOTE

Specifying an IPv4 address with a subnet mask is not supported. Instead, enter a prefix number in Classless Inter-Domain Routing (CIDR) notation. To enter a prefix number for a network mask, type a forward slash (/) and the number of bits in the mask immediately after the IP address. For example, enter, “209.157.22.99/24” for an IP address that has a network mask with 24 leading 1s in the network mask, representing 255.255.255.0.

---

```
switch(config-Management-1/0)# do show running-config interface management
interface Management 1/0
 no ip address dhcp
 ip address 10.24.85.81/20
 r-bridge-id1
 ip route 0.0.0.0/0 10.24.80.1
 no ipv6 address autoconfig
```

8. Apart from the two IP addresses on the management modules, modular switches also supports a chassis virtual IP address. Using this virtual IP address, you can login to the switch. The VCS virtual IP binds to the active MM automatically.

```
switch(config)# rbridge-id 1
switch(config-rbridge-id-1)# chassis virtual-ip 10.24.85.90/20
```

### *Configuring a static IPv6 Ethernet address*

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **interface management rbridge-id/port** command.

This command enters a management interface configuration mode where you can choose configuration parameters for IPv4 and IPv6 addresses.

- A compact switch has a single management port, and the port number for the management port is always 0.
- On a modular switches with two redundant management modules, you can configure two management ports. The port numbers are 1 and 2.

3. Enter the **ipv6 address IPv6\_address/prefix\_length** command.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface management 1/0
switch(config-Management-1/0)# ipv6 address
fd00:60:69bc:832:e61f:13ff:fe67:4b94/64
```

### Configuring an IP address with DHCP

---

#### NOTE

DHCP is not supported for IPv6 addresses.

---

By default, DHCP is disabled. You must explicitly enable the service. Use the **ip address dhcp** command to enable DHCP for IPv4 addresses, and the **ipv6 address dhcp** command to enable DHCP for IPv6 addresses. The Network OS DHCP clients support the following parameters:

- External Ethernet port IP addresses and prefix length
- Default gateway IP address

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

The following example enables DHCP for IPv4 addresses.

```
switch(config)# interface management 1/1
switch(config-Management-1/1)# ip address dhcp
```

The following example enables DHCP for IPv6 addresses.

```
switch(config)# interface management 1/1
switch(config-Management-1/1)# ipv6 address dhcp
```

The **show running-config interface management** command indicates whether DHCP is enabled. The following example shows a switch with DHCP enabled for IPv4 addresses.

```
switch# show running-config interface management
interface Management 2/0
ip address dhcp
ip route 0.0.0.0/0 10.24.80.1
ip address 10.24.73.170/20
```

## 3 Configuring the Ethernet management interface

```
no ipv6 address autoconfig
```

---

### NOTE

Enabling DHCP removes all configured static IP addresses.

---

## Stateless IPv6 autoconfiguration

IPv6 allows the assignment of 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.

When IPv6 autoconfiguration is enabled, the platform will engage in stateless IPv6 autoconfiguration. When IPv6 autoconfiguration is disabled, the platform will relinquish usage of any autoconfigured IPv6 addresses that it may have acquired while IPv6 autoconfiguration was enabled. This same enabled and disabled state also enables or disables the usage of a link local address for each managed entity (though a link local address will continue to be generated for each switch) because those link local addresses are required for router discovery.

The enabled or disabled state of autoconfiguration does not affect any static IPv6 addresses that may have been configured. Stateless IPv6 autoconfiguration and static IPv6 addresses can coexist.

## Setting IPv6 autoconfiguration

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. Take the appropriate action based on whether you want to enable or disable IPv6 autoconfiguration.
  - Enter the **ipv6 address autoconfig** command to enable IPv6 autoconfiguration for all managed entities on the target platform.
  - Enter the **no ipv6 address autoconfig** command to disable IPv6 autoconfiguration for all managed entities on the target platform.

---

### NOTE

On the Brocade VDX 8770, the **autoconfig** command can be issued only on the interface *rbridge-id/1*. However, this operation enables auto-configuration for the entire chassis.

---

## Displaying the network interface

If an IP address has not been assigned to the network interface, you must connect to the Network OS CLI using a console session on the serial port. Otherwise, connect to the switch through Telnet or SSH. Enter the **show interface management** command to display the management interface.

The following example shows the management interface on a Brocade VDX compact switch.

```
switch# show interface management
interface Management 9/0
ip address 10.24.81.65/20
```

```
ip route 0.0.0.0/0 10.24.80.1
ipv6 ipv6-address []
ipv6 ipv6-gateways [fe80::21b:edff:fe0f:bc00 fe80::21b:edff:fe0c:c200]
line-speed actual "1000baseT, Duplex: Full"
line-speed configured Auto
```

The following example shows the management interfaces on a Brocade VDX 8770-4. IPv6 autoconfiguration is enabled for the entire chassis, and, as a result, a stateless IPv6 address is assigned to both management interfaces.

```
switch# show interface management
interface Management 110/1
 ip address 10.20.238.108/21
 ip route 0.0.0.0/0 10.24.80.1
 ipv6 ipv6-address ["stateless fd00:60:69bc:85:205:33ff:fe78:7d88/64 preferred"
]
 ipv6 ipv6-gateways [fe80::21b:edff:fe0b:7800 fe80::21b:edff:fe0b:2400]
 line-speed actual "1000baseT, Duplex: Full"
 line-speed configured Auto
interface Management 110/2
 ip address 10.20.238.109/21
 ip route 0.0.0.0/0 10.24.80.1
 ipv6 ipv6-address ["stateless fd00:60:69bc:85:205:33ff:fe78:be14/64 preferred"
]
 ipv6 ipv6-gateways [fe80::21b:edff:fe0b:7800 fe80::21b:edff:fe0b:2400]
 line-speed actual "1000baseT, Duplex: Full"
 line-speed configured Auto
```

## Configuring the management interface speed

By default, the speed of the interface is set to autoconfiguration, which means the interface speed is optimized dynamically depending on load and other factors. You can override the default with a fixed speed value of 10 Mbps full duplex or 100 Mbps full duplex.

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.

```
switch# configure terminal
Entering configuration mode terminal
```

2. Enter the **interface management** command followed by *rbridge-id/O*.

This command places you in the management interface subconfiguration mode.

```
switch(config)# interface management 1/0
switch(config-Management-1/0)#
```

3. Enter the **speed** command with the selected speed parameter. The valid values are **10**, **100**, and **auto**.

```
switch(config-Management-1/0)# speed auto
```

4. Enter the **do show interface management** command followed by *rbridge-id/O* to display the new settings.

```
switch(config-Management-1/0)# do show interface management 1/0
interface Management 1/0
 ip address 10.24.81.65/20
 ip route 0.0.0.0/0 10.24.80.1
 ipv6 ipv6-address []
 ipv6 ipv6-gateways [fe80::21b:edff:fe0f:bc00 fe80::21b:edff:fe0c:c200]
```

```
line-speed actual "1000baseT, Duplex: Full"
line-speed configured Auto
```

5. Save the configuration changes using the **copy running-config startup-config** command.

```
switch(config-Management-1/0)# do copy running-config startup-config
```

## Outbound Telnet and SSH

Secure Shell (SSH) and Telnet are mechanisms for allowing secure access to management functions on a remote networking device. SSH provides a function similar to Telnet, but unlike Telnet, which offers no security, SSH provides a secure, encrypted connection to the device.

SSH and Telnet support is available in privileged EXEC mode on all Brocade VDX platforms. Both IPv4 and IPv6 addresses are supported.

### Establishing a Telnet connection

To establish a Telnet session, you can use the default settings.

```
switch# telnet 10.17.37.157
Trying 10.17.37.157...
Connected to 10.17.37.157.
Escape character is '^]'.
```

```
Network OS (sw0)
switch login:
```

Telnet connects on port 23. You can override the default port by using the **telnet ip\_address** command with the optional *port* operand (range 0-65535). However, the device must be listening on that port for the connection to succeed.

The following example overrides the default port.

```
switch# telnet 10.17.37.157 87
Trying 10.17.37.157...
Connected to 10.17.37.157.
Escape character is '^]'.
```

```
Network OS (sw0)
switch# login:
```

### Unsupported Features

The following features are not supported with Telnet:

- Displaying Telnet sessions
- Terminating hung Telnet sessions

### SSH supported features

SSHv2 is the supported version of SSH, but not all features typically available with SSHv2 are supported on the Brocade VDX family of switches.

The following encryption algorithms are supported:

- **3des** Triple-DES (default)

- **aes256-cbc**: AES in CBC mode with 256-bit key
- **aes192-cbc**: AES in CBC mode with 192-bit key
- **aes128-cbc**: AES in CBC mode with 128-bit key

The following HMAC (Hash-based Message Authentication Code) message authentication algorithms are supported:

- **hmac-md5**: MD5 encryption algorithm with 128-bit key (default).
- **hmac-md5-96**: MD5 encryption algorithm with 96-bit key.
- **hmac-sha1**: SHA1 encryption algorithm with 160-bit key.
- **hmac-sha1-96**: SHA1 encryption algorithm with 96-bit key.

SSH user authentication is performed with passwords stored on the device or on an external authentication, authorization, and accounting (AAA) server.

### Unsupported Features

The following features are not supported with SSH:

- Displaying SSH sessions
- Deleting stale SSH keys

## Establishing an SSH connection

In privileged EXEC mode, enter the **ssh -l username ip\_address** command to establish an SSH connection with default parameters. Use the **-m** and **-c** options to override the default encryption and hash algorithms. The following example overrides the default settings.

```
switch# ssh -l admin -m hmac-md5 -c aes128-cbc 10.20.51.68
The authenticity of host '10.20.51.68 (10.20.51.68)' can't be established.
RSA key fingerprint is ea:32:38:f7:76:b7:7d:23:dd:a7:25:99:e7:50:87:d0.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.20.51.68' (RSA) to the list of known hosts.
admin@10.20.51.68's password: *****

WARNING: The default password of 'admin' and 'user' accounts have not been
changed.

Welcome to the Brocade Network Operating System Software
admin connected from 10.20.51.66 using ssh on C60_68F
```

## Modular platform basics

The Brocade VDX 8770 platform features two redundant management modules, three or six switch fabric modules, and four or eight interface modules depending on the switch model. The Brocade VDX 8770-4 supports four interface modules and the Brocade VDX 8770-8 supports eight interface modules.

Table 5 shows the modules supported on each platform.

**TABLE 5** Modules supported on the Brocade VDX 8770 platform

| Type     | Module ID  | Slot numbers<br>VDX 8770-4 | Slot numbers<br>VDX 8770-8 | Description                                                |
|----------|------------|----------------------------|----------------------------|------------------------------------------------------------|
| MM       | 0x70 = 112 | M1, M2                     | M1, M2                     | Management module<br>(an 8-core 1.5-GHz Control Processor) |
| SFM      | 0x71 = 113 | S1 - S3                    | S1 - S6                    | Switch fabric module (core blade)                          |
| LC48X10G | 0x72 = 114 | L1 - L4                    | L1 - L8                    | 48-port 10-GbE interface module                            |
| LC12X40G | 0x7F = 127 | L1 - L4                    | L1 - L8                    | 12-port 40-GbE interface module                            |
| LC48X1G  | 0x83 = 131 | L1 - L4                    | L1 - L8                    | 48-port 1-GbE interface module                             |

## Management module

Two management modules provide redundancy and act as the main controller on the Brocade VDX 8770-4 and VDX 8770-8 chassis. The management modules host the distributed Network OS that provides the overall control plane management for the chassis. You can install a redundant management module in slot M1 or M2 in any of the Brocade VDX 8770 chassis. By default, the system considers the module in slot M1 the active management module and the module in slot M2 the redundant, or standby, management module. If the active module becomes unavailable, the standby module automatically takes over management of the system.

Each management module maintains its own copy of the configuration database. The startup configuration is automatically synchronized with the other management module.

Brocade recommends that each management module (primary and secondary partition) should maintain the same firmware version. For more information on maintaining firmware, refer to [Chapter 6, “Installing and Maintaining Firmware”](#).

Each management module has two Ethernet interfaces, Eth0 and Eth2. Eth0 is the management interface and can be configured with an IP address. For more information on configuring the management interface, refer to [“Configuring the Ethernet management interface”](#) on page 49.

## Switch fabric modules

The switch fabric modules play a dual role in the fabric connectivity between interface modules, providing both the data-plane connectivity and the control-plane connectivity needed for end-to-end credit management in each of the interface modules.

In each chassis model, two slots are designated for supporting the control-plane connectivity. In the Brocade VDX 8770-4, the slots S1 and S2 are the designated control-plane slots. In the Brocade VDX 8770-8, the slots S3 and S4 are the designated control-plane slots. At least one of the control-plane slots must be populated to maintain operation. If you remove the switch fabric modules from both the control-plane slots, all interface modules will be faulted and the chassis is no longer operational.

## Interface modules

Three types of interface modules provide I/O ports for network Ethernet protocols, the Brocade LC48x10G, the Brocade LC48x1G, and the Brocade LC12x40G.



The Brocade LC48x1G provides 48 1-GbE/10-GbE SFP+ front ports. The Brocade LC48x10G interface modules provide 48 1-GbE/10-GbE SFP+ front ports.

- The 10-GbE SFP+ interfaces are named “TenGigabitEthernet” or “TE”.
- The 1-GbE SFP+ interfaces are named “GigabitEthernet” or “GE”.

The Brocade LC12x40G interface module provides twelve (12) 40-GbE QSFP front ports.

- The 40-GbE QSFP interfaces are named “FortyGigabitEthernet” or “FO”.

## Supported interface modes

All interfaces in the Brocade VDX 8770 chassis come online as Fabric interswitch links (“Fabric ISLs”) by default and will attempt to form a Brocade VCS fabric. If the ISL formation fails, the interfaces come up as “Edge ports”.

---

### NOTE

The Brocade VDX 8770 chassis always comes up in VCS mode. Standalone mode is not supported on the Brocade VDX 8770 platform.

---

## Displaying the interfaces

Interfaces on the VDX 8770 platform are identified by the RBridge ID, slot number, and port number, separated by forward slashes (/). For example, the notation 9/2/8 indicates port 8 located in slot 2 on a chassis with the RBridge ID of 9.

Enter the **show running-config interface** *interface\_type* command to display the interfaces and their status.

```
switch# show running-config interface tengigabitethernet
interface tengigabitethernet 1/1/1
 fabric isl enable
 fabric trunk enable
 no shutdown
!
interface tengigabitethernet 1/1/2
 fabric isl enable
 fabric trunk enable
 no shutdown
!
interface tengigabitethernet 1/1/3
 fabric isl enable
 fabric trunk enable
 no shutdown
!
interface tengigabitethernet 1/1/4
 fabric isl enable
 fabric trunk enable
 no shutdown
!
interface tengigabitethernet 1/1/5
 fabric isl enable
 fabric trunk enable
 no shutdown
!
interface tengigabitethernet 1/1/6
```

```

fabric isl enable
fabric trunk enable
no shutdown

```

Enter the **show interface interface\_type rbridge\_id/slot/port** command to display the configuration details for the specified interface.

```

switch# show interface tengigabitethernet 1/1/9
tengigabitethernet 1/1/9 is up, line protocol is up (connected)
Hardware is Ethernet, address is 0005.3315.df5a
 Current address is 0005.3315.df5a
Pluggable media present
Interface index (ifindex) is 4702109825
MTU 9216 bytes
LineSpeed Actual : 10000 Mbit
LineSpeed Configured : Auto, Duplex: Full
Flowcontrol rx: off, tx: off
Priority Tag disable
Last clearing of show interface counters: 04:12:03
Queueing strategy: fifo
Receive Statistics:
 1580 packets, 140248 bytes
 Unicasts: 0, Multicasts: 1580, Broadcasts: 0
 64-byte pkts: 0, Over 64-byte pkts: 1561, Over 127-byte pkts: 17
 Over 255-byte pkts: 2, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
 Over 1518-byte pkts(Jumbo): 0
 Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
 Errors: 0, Discards: 0, TrillportCtrlFrames: 1564
Transmit Statistics:
 1583 packets, 140120 bytes
 Unicasts: 0, Multicasts: 1583, Broadcasts: 0
 Underruns: 0
 Errors: 0, Discards: 0, TrillportCtrlFrames: 1583
Rate info (interval 299 seconds):
 Input 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
 Output 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:15:53

```

## Slot numbering

The slot numbering on the Brocade VDX 8770 chassis is based on the module type. The slot numbers for the interface module are numbered L1 through L4 on the Brocade VDX 8770-4, and L1 through L8 on the Brocade VDX 8770-8. The slots for the management modules are numbered M1 and M2. The slots for the switch fabric modules are numbered S1 through S3 on the Brocade VDX 8770-4, and S1 through S6 on the Brocade VDX 8770-8.

### *Displaying slots and module status information*

Use the **show slots** command to display information for all slots in the chassis. The following example shows slot information for the Brocade VDX 8770-8.

```
switch# show slots
```

| Slot | Type | Description          | ID  | Status  |
|------|------|----------------------|-----|---------|
| M1   | MM   | Management Module    | 112 | ENABLED |
| M2   | MM   | Management Module    | 112 | ENABLED |
| S1   | SFM  | Switch Fabric Module | 113 | ENABLED |

```

S2 VACANT@
S3 SFM Switch Fabric Module 113 ENABLED#
S4 SFM Switch Fabric Module 113 ENABLED#
S5 SFM Switch Fabric Module 113 ENABLED
S6 SFM Switch Fabric Module 113 ENABLED
L1 VACANT
L2 VACANT
L3 LC48X10G 48-port 10GE card 114 DIAG RUNNING POST1
L4 LC48X10G 48-port 10GE card 114 ENABLED
L5 VACANT
L6 VACANT
L7 LC48X1G 48-port 1GE card 114 ENABLED
L7 VACANT
L8 VACANT
= At least one enabled SFM in these slots is required.
@ = The SFM Optical Switch is open.

```

Alternatively, you can use the following commands to display slots per module type:

- Use the **show mm** command to display information for the management modules.
- Use the **show sfm** command to display information for the switch fabric modules.
- Use the **show linecard** command to display information for the interface modules.

## Slot configuration

Interface modules are registered with the system by type, and the slot must be configured with the correct type before you can install an interface module in that slot. When you install a new interface module, the system checks whether or not a previous configuration is associated with the slot. The following rules apply when you install or replace an interface module:

- When you install an interface module and boot it up to an online state in a slot that was never occupied or configured, the module type information is automatically detected and saved to the database. No special configuration is required.
- If you install an interface module in a slot that was previously occupied by an interface module of the same type and the slot is configured for that same type, you can hot-swap the modules without powering off the interface modules. No slot configuration changes are required.
- If the slot was previously configured for a different type of interface module, the installation fails and the module is faulted with a “Type mismatch” error. A RASlog error message is generated. You must power off the interface module and clear the slot configuration with the **no linecard** command before you can configure the slot for a new interface module.

The slot configuration persists in the database even after the interface module is physically removed, powered off, or faulted since it first came online. All configuration data associated with the slot is automatically preserved across reboot or hot-swap of the interface module with the same type.

To make the slot configuration persistent across a chassis reboot (which involves reloading the management modules), you must save the configuration persistently by issuing the **copy running-config startup-config** command after the interface module reaches the online state and before the system reboots.

## Replacing an interface module

You can remove an interface module without powering it off. However, doing so will not remove the configuration. When you replace a module with a different type, you must first remove the configuration and then reconfigure the slot for the new interface module type.

Removing the configuration requires the interface module to be powered off.

1. Power off the interface module by issuing the **power-off linecard** command followed by the slot number.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **rbridge-id rbridge-id** command to enter the RBridge sub configuration mode.
4. Enter the **no linecard slot\_number** command to clear the slot configuration.
5. Remove the interface module.
6. Enter the **linecard slot\_number** command followed by a question mark (?) to display the line card menu.
7. Select a line card type and enter the **linecard slot\_number linecard\_type** command.
8. Enter the **exit** command twice to return to privileged EXEC mode.
9. Insert the new interface module into the configured slot.
10. Enter the **power-on linecard** command to power on the interface module.
11. Save the configuration persistently by issuing the **copy running-config startup-config** command after the interface module reaches the online state.
12. Verify the configuration with the **show running-config linecard linecard** command.

```
switch# power-off linecard 4
switch# configure terminal
Entering configuration mode terminal
switch(config)# rbridge-id 1
switch(config-rbridge-id-1)# no linecard 4
switch(config-rbridge-id-1)# linecard 4 ?
Possible completions:
 LC12x40G 12X40G linecard
 LC48x1G 48X1G linecard
 LC48x10G 48X10G linecard
 LC72x1G 72X1G linecard
switch(config-rbridge-id-1)# linecard 4 LC48x10G
Creating new linecard configuration was successful.
switch(config-rbridge-id-1)# exit
switch(config)# exit
switch# copy running-config startup-config
switch# show running-config rbridge-id 4 linecard
rbridge-id 1
 linecard 1 LC48x10G
 linecard 4 LC48x10G
```

## High availability failover

In Network OS v3.0.0, high availability (HA) is cold. When an active management module fails over, the standby management module takes over as the active management module and goes through a cold recovery, during which time, all system components are reset and recovered. All switch fabric modules and interface modules are also reset, resulting in traffic disruption.

With the introduction of Network OS v4.0.0, HA failover is cold, similar to Network OS v3.0.0; however, in fabric cluster and logical chassis mode, HA is warm. In this instance, the following results apply:

- No data path disruption results for Layer 2 (L2) and FCoE traffic.
- All L2 control protocol states are retained.
- The topology state and interface state are retained.
- The last accepted user configuration is retained.
- Data path disruption and control path disruption result for IP.
- IP configuration is replayed post failover.

With the introduction of Network OS 4.0.0, In-service software upgrades (ISSUs) are supported. An ISSU allows a dual management module system to be upgraded non-disruptively and is invoked by entering the firmware download command from the active management module (MM).

ISSU is supported in both fabric cluster mode and logical chassis cluster mode. The HA behavior during ISSU is the same as that of warm failover described above. For more information, refer to [“Upgrading firmware on a modular chassis”](#) on page 82.

## High availability commands

The **ha** commands are available on the switch in privileged EXEC mode.

- Use the **show ha** command to display the management module status.  

```
switch# show ha
Local (M2): Active, Cold Recovered
Remote (M1): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
```
- Use the **ha failover** command to force the active management module to fail over. The standby management module will take over as the active management module.
- Use the **reload system** command to perform a reboot of the entire chassis. This command is supported only on the active management module. This command is not supported on the standby management module. Both management modules must be in sync for the HA reboot operation to succeed.

In logical chassis cluster mode, this command is issued from the principal node to reset one remote node or all of the remote nodes by specifying the rbridge-id parameter.

- Use the **ha sync start** command to enable HA state synchronization after an **ha sync stop** command has been invoked.

---

### NOTE

For additional **ha** and related commands, see the *Network OS Command Reference Supporting Network OS Release 4.0.0* or later.

---

Table 6 and Table 7 identify expected behaviors that result from controlled and uncontrolled reload and failover conditions.

**TABLE 6**      **Controlled reload and failover**

| Command syntax | Behavior in fabric cluster                                                                                                                        | Behavior in logical chassis cluster                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| reload         | <ul style="list-style-type: none"> <li>• Cold failover to standby management module (MM).</li> <li>• Reload system if MMs not in sync.</li> </ul> | <ul style="list-style-type: none"> <li>• Cold failover to standby management module (MM).</li> <li>• Reload system if MMs not in sync.</li> </ul> |
| reload standby | Reboot the standby MM.                                                                                                                            | Reboot the standby MM.                                                                                                                            |
| reload system  | Reboot both MMs. MMs will retain the HA roles.                                                                                                    | Reboot both MMs. MMs will retain the HA roles.                                                                                                    |
| ha failover    | Warm failover to standby MM.                                                                                                                      | Cold failover to standby MM.                                                                                                                      |

**TABLE 7**      **Uncontrolled failover**

| Command syntax | Behavior in fabric cluster                    | Behavior in logical chassis cluster           |
|----------------|-----------------------------------------------|-----------------------------------------------|
| Panic          | Warm failover to standby MM.                  | Cold failover to standby MM.                  |
| MM removal     | Warm failover to standby MM.                  | Cold failover to standby MM.                  |
| Power cycle    | MMs will retain the HA roles upon booting up. | MMs will retain the HA roles upon booting up. |

## Configuring a switch banner

A banner is a text message that displays on the switch console. It can contain information about the switch that an administrator may want users to know when accessing the switch.

The banner can be up to 2048 characters long. To create a multi-line banner, enter the **banner login** command followed by the **Esc-m** keys. Enter **Ctrl-D** to terminate the input.

If you are in logical chassis cluster mode, the configuration is applied to all nodes in the cluster.

### Setting and displaying a banner

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. Enter the **banner login** command and a text message enclosed in double quotation marks ("").
3. Enter the **do show running-config banner** command to display the configured banner.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# banner login "Please do not disturb the setup on this switch"
switch(config)# do show running-config banner
banner login "Please do not disturb the setup on this switch"
```

Use the **no banner login** command to remove the banner.

## supportSave data

If you are troubleshooting a production system, you will have to capture data for further analysis or send the data to your switch service provider. The **copy support** command provides a mechanism for capturing critical system data and uploading the data to an external host or saving the data to an attached USB device.

### Uploading supportSave data to an external host

To upload supportSave data interactively, enter the **copy support-interactive** command and provide input as prompted. Specifying an IPv6 address for the server requires Network OS v3.0.0 or later. For a non-interactive version of the command, refer to the *Network OS Command Reference* for your release.

```
switch# copy support-interactive
Server Name or IP Address: 10.38.33.131
Protocol (ftp, scp): ftp
User: admin
Password: *****
Directory: /home/admin/support
VCS support [y/n]? (y): n
Module timeout multiplier[Range:1 to 5.Default:1]:
copy support start
Saving support information for chassis:sw0, module:RAS...
(output truncated)
```

### Saving supportSave data to an attached USB device

You can use a Brocade-branded USB device to save the support data. The Brocade-branded USB device comes with factory-configured default directories and interacts with the Network OS CLI.

1. Enter the **usb on** command to enable the USB device.
2. Enter the **usb dir** command to display the default directories.
3. Enter the **copy support usb directory support** command.

```
switch# usb on
USB storage enabled
switch# usb dir
firmwarekey\ 0B 2010 Aug 15 15:13
support\ 106MB 2010 Aug 24 05:36
support1034\ 105MB 2010 Aug 23 06:11
config\ 0B 2010 Aug 15 15:13
firmware\ 380MB 2010 Aug 15 15:13
Available space on usbstorage 74%

switch# copy support usb directory support
```

If you are in logical chassis cluster mode, you can use the **rbridge-id all** option to invoke **supportSave** on all nodes at the same time. The **copy support rbridge-id all** command is a blocking command. The Telnet session from which the command is issued will be blocked until **supportSave** is completed on all nodes in the cluster; however, users can again Telnet into the same node or any other nodes in the cluster. When the command is in progress, output messages from all nodes are

shown that include the respective node rbridge-ids. The copy Support command, when executed with USB as the protocol option, will collect support files to the USB device that is connected to the respective nodes. All USB devices connected to each of the nodes should be enabled before the **copy support usb** command is executed.

The following example shows the **copy support** command using the **rbridge-id all** option.

```
switch# copy support ftp host 10.1.2.30 user fvt password pray4green directory
/support rbridge-id all
switch 100: copy support start
switch 117: Saving support information for chassis:sw0, module:RAS...
switch 100: Saving support information for chassis:sw, module:RAS...
switch 117: Saving support information for chassis:sw0, module:CTRACE_OLD...
.....
switch 100: copy support completed
switch 117: copy support completed
2011/04/07-18:03:07, [SS-1000], 2752,, INFO, VDX6720-24, copy support has
uploaded support information to the host with IP address 10.70.4.101.
```

## Displaying the status of a supportSave operation

Enter the **show copy-support status** command.

```
switch# show copy-support status
```

| Slot Name | SS type   | Completion Percentage |
|-----------|-----------|-----------------------|
| # # # # # | # # # # # | # # # # #             |
| M1        | NORMAL    | [100%]                |
| L1/0      | NORMAL    | [100%]                |
| L1/1      | NORMAL    | [100%]                |
| L2/0      | NORMAL    | [100%]                |
| L2/1      | NORMAL    | [100%]                |
| L4/0      | NORMAL    | [100%]                |
| L4/1      | NORMAL    | [100%]                |

## Configuring autoupload of supportSave data

You can configure a switch to upload first-fault data capture (FFDC) and trace data files automatically to a remote server that is specifically set up for collecting information that results from the **supportSave** command. To enable this feature, you must configure a dedicated server, then invoke the **autoupload-param** command to set the parameters, followed by the **support autoupload enable** command to enable the configurations.

```
switch(config)# support autoupload-param hostip 10.31.2.27 username supportadmin
directory /users/support/ffdc_autoupload protocol ftp password
(<string>): *****
```

## Displaying the autoupload configuration

Enter the **show running-config support autoupload-param** command to display the autoupload configuration on the local switch.

```
switch(config)# do show running-config support autoupload-param
support autoupload-param hostip 10.31.2.27 username supportadmin directory
/users/support/ffdc_autoupload protocol ftp password "3iTYxJWEUHp9axZQt2tbvw==\n"
```



## Additional supportSave commands

Use the following commands to configure additional supportSave data collection parameters:

- Use the **show support** command to display a list of core files on the switch.
- Use the **clear support** command to erase support data on the switch.

Refer to the *Network OS Command Reference* for your release for more information on these commands.

## Message logging

Network OS provides several mechanisms for logging error messages including syslog, RASLog, and audit log. The types of message logging available and set up procedures are documented in the “Introduction to Brocade Error Message Logging” chapter of the *Network OS Message Reference*.

## 3 Message logging

# Network Time Protocol

---

## In this chapter

- [Date and time settings](#) ..... 67
- [Time zone settings](#)..... 68
- [Network Time Protocol](#) ..... 69

## Date and time settings

Brocade switches maintain the current date and time inside a battery-backed real-time clock (RTC) circuit. Date and time are used for logging events. Switch operation does not depend on the date and time; a switch with incorrect date and time settings can function correctly. However, because the date and time are used for logging, error detection, and troubleshooting, you should set them correctly.

### Setting the date and time

The **clock set** command sets the local clock date and time. Valid date and time values must be in the range between January 1, 1970 and January 19, 2038. If a time zone is not configured, the time zone defaults to Greenwich Mean Time (GMT). If an active NTP server is configured for the switch, it overrides the local time settings.

Enter the **clock set CCYY-MM-DDTHH:MM:SS** command.

The variables represent the following values:

- **CCYY** specifies the year; the valid range is 1970 through 2038.
- **MM** specifies the month; the valid range is 01 through 12.
- **DD** specifies the day; the valid range is 01 through 31.
- **HH** specifies the hour; the valid range is 00 through 23.
- **MM** specifies the minutes; the valid range is 00 through 59.
- **SS** specifies the seconds; the valid range is 00 through 59.

If you are in VCS mode, setting the time and date is done using the RBridge ID of the node.

#### Example of setting and displaying the date and time in standalone mode

```
switch# clock set 2013-06-06T12:15:00
switch# show clock
rbridge-id 1: 2013-06-06 12:15:05 Etc/GMT+0
```

#### Example of setting and displaying the date and time in VCS mode

```
switch# clock set 2013-06-06T12:15:00 rbridge-id all
switch# show clock
```

```
rbridge-id all: 2013-06-06 12:15:05 Etc/GMT+0
```

# Time zone settings

You can set the time zone by specifying a geographic region and city by name. You can choose one of the following the regions: Africa, America, Pacific, Europe, Antarctica, Arctic, Asia, Australia, Atlantic, and Indian.

The time zone setting has the following characteristics:

- 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 in the Greenwich Mean Time (GMT) time zone (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 will reflect the time zone changes only after the next reboot.
- Time zone settings persist across failover for high availability.
- Time zone settings are not affected by NTP server synchronization.

## Setting the time zone

Use the **clock timezone** command to set the time zone for a switch. You must use the command for all switches for which a 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.

If you are in VCS mode, setting the time and date is done using the rbridge-id of the node.

Refer to [Appendix B, “Supported time zones and regions”](#) for a complete list of configurable regions and cities.

Enter the **clock timezone region/city** command.

**Example for setting and displaying the date and time in standalone mode**

```
switch# clock timezone America/Los_Angeles
```

**Example for setting and displaying the date and time in VCS mode**

```
switch# clock timezone America/Los_Angeles rbridge-id all
```

---

### NOTE

After upgrading your switch firmware, you might need to reconfigure the time zone information.

---

## Displaying the current local clock and time zone

The **show clock** command returns the local time, date, and time zone.

---

### NOTE

This command is currently supported on the local switch.

---

**Example for showing the local switch clock time:**

```
switch# show clock
rbridge-id 1: 2012-05-04 16:01:51 America/Los Angeles
```

**Example for showing clock time for all switches in the cluster (logical chassis cluster mode only):**

```
switch# show clock rbridge-id all
```

**Example for showing clock time for switch with rbridge-id 16:**

```
switch# show clock rbridge-id 16
```

## Removing the time zone setting

Use the **no clock timezone** command to remove the time zone setting for the local clock. This operation returns the local time zone to the default value (GMT). When using the **no** operand, you do not need to reference a timezone setting.

Enter the **no clock timezone** command.

**Example for removing the time zone setting in standalone mode**

```
switch# no clock timezone
```

**Example for removing the time zone setting in VCS mode**

```
switch# no clock timezone rbridge-id 5
```

# Network Time Protocol

Network Time Protocol (NTP) maintains uniform time across all switches in a network. The NTP commands support the configuration of an external time server to maintain synchronization between all local clocks in a network.

To keep the time in your network current, it is recommended that each switch have its time synchronized with at least one external NTP server. External NTP servers should be synchronized among themselves in order to maintain fabric-wide time synchronization.

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

If you are in Standalone mode, Network Time Protocol (NTP) commands must be configured on each individual switch. Network time synchronization is guaranteed only when a common external time server is used by all switches. If you are in VCS mode, when an **ntp server** command is invoked on one switch in a cluster, the configuration is applied to all switches in the cluster.

The **ntp server** command accepts up to five server addresses in IPv4 or IPv6 format. When you configure multiple NTP server addresses, the **ntp server** command sets the first obtainable address as the active NTP server. If there are no reachable time servers, then the local switch time is the default time until a new active time server is configured.

## Synchronizing the local time with an external source

Use the **ntp server** command to synchronize the local switch time with an NTP server. You can configure up to five IP address. At least one IP address in the list must be a reachable, configured NTP server or the request will fail.

Enter the **ntp server ip\_address** command.

```
switch(config)# ntp server 192.168.10.1
```

## Displaying the active NTP server

Use the **show ntp status** command to display the current active NTP server IP address. If an NTP server is not configured or the server is unreachable, the output displays LOCL (for local switch time). Otherwise, the command displays the NTP server IP address. The command displays the local NTP server configuration only.

If the RBridge ID parameter is not provided, status results default to the local switch (LOCL). If **rbridge-id all** is specified, the command displays the status for all switches in the cluster.

**Example for showing the local switch NTP status when an NTP server is not configured:**

```
switch# show ntp status
rbridge-id 1: active ntp server is LOCL
```

**Example for showing the configured NTP server:**

```
switch# show ntp status
active ntp server is 10.31.2.81
```

**Example for showing ntp status for all switches in a cluster.**

```
switch# show ntp status rbridge-id all
rbridge-id 7: active ntp server is LOCL
```

## Removing an NTP server IP address

Use this command to remove an NTP server IP address from a list of server IP addresses. At least one IP address in the remaining list must be a reachable, configured NTP server or the remove request fails.

Enter the **no ntp server** command.

```
switch(config)# no ntp server 192.168.10.1
switch# show ntp status
rbridge-id 1: active ntp server is LOCL
```

# Configuration Management

---

## In this chapter

|                                                             |    |
|-------------------------------------------------------------|----|
| • Switch configuration overview .....                       | 71 |
| • Flash file management .....                               | 71 |
| • Configuration file types .....                            | 73 |
| • Saving configuration changes .....                        | 74 |
| • Configuration backup .....                                | 76 |
| • Configuration restoration .....                           | 76 |
| • Configuration management on a modular chassis .....       | 78 |
| • Configuration management in Brocade VCS Fabric mode ..... | 79 |

## Switch configuration overview

Maintaining consistent configuration settings among switches in the same fabric is an important part of switch management and minimizes fabric disruptions. As part of standard maintenance procedures, it is recommended that you back up all important configuration data for every switch on an external host for emergency reference.

Typical configuration management tasks include the following actions:

- Saving the running configuration to the startup configuration file (“[Saving configuration changes](#)” on page 74).
- Uploading the configuration files to a remote location (“[Configuration backup](#)” on page 76).
- Restoring a configuration file from a remote archive (“[Configuration restoration](#)” on page 76).
- Archiving configuration files for all your switches to a remote location (“[Configuration management in Brocade VCS Fabric mode](#)” on page 79).
- Downloading a configuration file from a remote location to multiple switches (“[Configuration management in Brocade VCS Fabric mode](#)” on page 79).

## Flash file management

Brocade Network OS provides a set of tools for removing, renaming, and displaying files you create in the switch flash memory. You can use the display commands with any file, including the system configuration files. The **rename** and **delete** commands only apply to copies of configuration files you create in the flash memory. You cannot rename or delete any of the system configuration files.

## Listing the contents of the flash memory

To list the contents of the flash memory, enter the **dir** command in the privileged EXEC mode.

```
switch# dir
drwxr-xr-x 2 root sys 4096 Feb 13 00:39 .
drwxr-xr-x 3 root root 4096 Jan 1 1970 ..
-rwxr-xr-x 1 root sys 417 Oct 12 2010 defaultconfig.novcs
-rwxr-xr-x 1 root sys 697 Oct 12 2010 defaultconfig.vcs
-rw-r--r-- 1 root root 6800 Feb 13 00:37 startup-config
```

## Deleting a file from the flash memory

To delete a file from the flash memory, enter the **delete file** command in the privileged EXEC mode.

```
switch# delete myconfig
```

## Renaming a file

To rename a file in the flash memory, enter the **rename source\_file destination\_file** command in the privileged EXEC mode.

```
switch# rename myconfig myconfig_20101010
```

## Viewing the contents of a file in the flash memory

To investigate the contents of a file in the flash memory, enter the **show file file** command in the privileged EXEC mode

```
switch# show file defaultconfig.novcs
!
no protocol spanning-tree
!
vlan dot1q tag native
!
 cee-map default
 remap fabric-priority priority 0
 remap lossless-priority priority 0
 priority-group-table 1 weight 40 pfc on
 priority-group-table 2 weight 60 pfc off
 priority-group-table 15.0 pfc off
 priority-table 2 2 2 1 2 2 2 15.0
!
interface Vlan 1
shutdown
!
port-profile default
vlan-profile
 switchport
 switchport mode trunk
 switchport trunk allowed vlan all
!
protocol lldp
!
end
!
```



**NOTE**

To display the contents of the running configuration, use the **show running-config** command. To display the contents of the startup configuration, use the **show startup-config** command.

## Configuration file types

Brocade Network OS supports three types of configuration files. [Table 8](#) lists the standard configuration files and their functions.

**TABLE 8** Standard switch configuration files

| Configuration file                                                                                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Default configuration <ul style="list-style-type: none"> <li>defaultconfig.novcs</li> <li>defaultconfig.vcs</li> </ul> | Part of the Network OS firmware package. The default configuration is applied, if no customized configuration is available. There are different default configuration files for standalone and Brocade VCS Fabric mode.                                                                                                                                                                                                     |
| Startup configuration <ul style="list-style-type: none"> <li>startup-config</li> </ul>                                 | Configuration effective on startup and after reboot.                                                                                                                                                                                                                                                                                                                                                                        |
| Running configuration <ul style="list-style-type: none"> <li>running-config</li> </ul>                                 | Current configuration active on the switch. Whenever you make a configuration change, it is written to the running configuration. For fabric cluster mode, the running configuration does not persist across reboot, unless you copy it to the startup configuration.<br><br>However, when the switch is in logical chassis cluster mode, the running-config file is saved automatically and it does not need to be copied. |

Configuration management follows a transaction model. When you boot up a switch for the first time, the running configuration is identical to the startup configuration. As you configure the switch, the changes are written to the running configuration. To save the changes, you must save the currently effective configuration (the running configuration) as the startup configuration. When the switch reboots, the configuration changes become effective.

### Default configuration

Network OS provides two different configuration files for switches in standalone and Brocade VCS Fabric mode. When you change from standalone to Brocade VCS Fabric mode, the system chooses the appropriate default configuration based on the mode (Brocade VCS Fabric or standalone). Default configuration files are part of the Network OS firmware package and are automatically applied to the startup configuration under the following conditions:

- When the switch boots up for the first time and no customized configuration is available.
- When you enable or disable Brocade VCS Fabric mode, the appropriate default configuration is applied when the switch reboots.
- When you restore the default configuration.

You cannot remove, rename, or change the default configuration.

### *Displaying the default configuration*

To display the default configuration, enter the **show file** command in the privileged EXEC mode.

```
switch# show file defaultconfig.novcs
switch# show file defaultconfig.vcs
```

### Startup configuration

---

**NOTE**

There is no startup configuration for logical chassis cluster mode. Switches in a logical chassis cluster always preserve their running configuration.

---

The startup configuration is persistent. It is applied when the system reboots.

- When the switch boots up for the first time, the switch uses the default configuration as the startup configuration, depending on the mode.
- The startup configuration always matches the current Brocade VCS Fabric mode. It is deleted when you change modes, unless you make a backup copy.
- When you make configuration changes to the running configuration and save the changes to the startup configuration with the **copy** command, the running configuration becomes the startup configuration.

#### *Displaying the startup configuration*

To display the contents of the startup configuration, enter the **show startup-config** command in the privileged EXEC mode.

```
switch# show startup-config
```

### Running configuration

The configuration currently effective on the switch is referred to as the running configuration. Any configuration change you make while the switch is online is made to the running configuration.

- The running configuration is nonpersistent.
- To save configuration changes, you must copy the running configuration to the startup configuration. If you are not sure about the changes, you can copy the changes to a file, and apply the changes later.

#### *Displaying the running configuration*

To display the contents of the running configuration, enter the **show running-config** command in the privileged EXEC mode.

```
switch# show running-config
```

## Saving configuration changes

Configuration changes are nonpersistent and are lost on reboot unless you save them permanently. You have two options for saving configuration changes:

- Copy the running configuration to the startup configuration. The changes become effective upon reboot.
- Copy the running configuration to a file, and apply it at some later date.

**NOTE**

Always make a backup copy of your running configuration before you upgrade or downgrade the firmware.

## Saving the running configuration

To save the configuration changes you made, copy the running configuration to the startup configuration. The next time the switch reboots, it uses the startup configuration and the changes you made earlier become effective.

**NOTE**

When the switch is in logical chassis cluster mode, the running-config file is saved automatically and it does not need to be copied.

Enter the **copy running-config startup-config** command in the privileged EXEC mode.

```
switch# copy running-config startup-config
copy running-config startup-config
This operation will modify your startup configuration. Do you want to continue?
[Y/N]: y
```

## Saving the running configuration to a file

If you want to save the changes you made to the configuration, but you do not want the changes to take effect when the switch reboots, you can save the running configuration to a file. You can apply the changes at some later time.

1. Enter the **copy running-config file** command in the privileged EXEC mode. Specify the file name as the file URL.

```
switch# copy running-config flash://myconfig
```

2. Verify the transaction by listing the directory contents.

```
switch# dir
total 32
drwxr-xr-x 2 root sys 4096 Feb 17 17:50 .
drwxr-xr-x 3 root root 4096 Jan 1 1970 ..
-rwxr-xr-x 1 root sys 417 Oct 12 2010 defaultconfig.novcs
-rwxr-xr-x 1 root sys 697 Oct 12 2010 defaultconfig.vcs
-rw-r--r-- 1 root root 6777 Feb 17 17:50 myconfig
-rw-r--r-- 1 root root 6800 Feb 13 00:37 startup-config
```

## Applying previously saved configuration changes

When you are ready to apply the configuration changes you previously saved to a file, copy the file (*myconfig* in the example) to the startup configuration. The changes take effect after the switch reboots.

Enter the **copy file startup-config** command in the privileged EXEC mode. Specify the file name as the file URL.

```
switch# copy flash://myconfig startup-config
This operation will modify your startup configuration. Do you want to continue?
[Y/N]: y
```

## Configuration backup

Always keep a backup copy of your configuration files, so you can restore the configuration in the event the configuration is lost or you make unintentional changes. The following recommendations apply:

- Keep backup copies of the startup configuration for all switches in the fabric.
- Upload the configuration backup copies to an external host or to an attached Brocade-branded USB device.
- Avoid copying configuration files from one switch to another. Instead restore the switch configuration files from the backup copy.

### Uploading the startup configuration to an external host

Enter the **copy startup-config destination\_file** command in the privileged EXEC mode.

In the following example, the startup configuration is copied to a file on a remote server using FTP.

```
switch# copy startup-config
ftp://admin:*****@122.34.98.133//archive/startup-config_vdx24-08_20101010
```

### Backing up the startup configuration to a USB device

When you make a backup copy of a configuration file on an attached USB device, the destination file is the file URL on the USB device. You do not need to specify the target directory. The file is automatically recognized as a configuration file and stored in the default configuration directory.

1. Enable the USB device.

```
switch# usb on
USB storage enabled
```

2. Enter the **copy startup-config destination\_file** command in the privileged EXEC mode.

```
switch# copy startup-config usb://startup-config_vdx24-08_20101010
```

## Configuration restoration

Restoring a configuration involves overwriting a given configuration file on the switch by downloading an archived backup copy from an external host or from an attached USB device. There are two typical scenarios.

- [“Restoring a previous startup configuration from backup”](#) on page 77.
- [“Restoring the default configuration”](#) on page 77.

---

#### NOTE

Configuration files that were created using Brocade Network OS 2.x should not be loaded onto a system running Brocade Network OS 3.x or later. The ACL and VLAN configuration information has changed in Brocade Network OS 3.x or later, and the affected lines of configuration are skipped when loading a Brocade Network OS 2.x configuration file.

---

## Restoring a previous startup configuration from backup

- You want to back out of configuration changes you made earlier by overwriting the startup configuration with a modified running configuration.
- You want to take the switch from Brocade VCS Fabric mode back to standalone mode and reapply your original standalone startup configuration.
- This task only applies to Brocade VDX 6720-30 switches.
  1. Disable Brocade VCS Fabric mode and reboot the switch.

The startup configuration associated with Brocade VCS Fabric mode is automatically deleted. The switch boots up in standalone mode and loads the corresponding default configuration.

2. Copy the archived startup configuration file from an FTP server or from an attached USB device to the running configuration.
3. Reboot the switch.

```
switch# no vcs enable
```

The switch automatically reboots at this point.

```
switch# copy ftp://admin:*****@122.34.98.133//archive/\
startup-config_vdx24-08_20101010 running-config
switch# reload system
```

---

### ATTENTION

Make sure that the configuration file you are downloading is the one that belongs to the switch you want to restore. It is a good idea to identify archived configuration files by switch name and date.

---

## Restoring the default configuration

This restoration procedure resets the configuration to the factory defaults. The default configuration files for Brocade VCS Fabric and standalone mode are always present on the switch and can be restored with the **copy** command.

To restore the default configuration, perform the following procedure in privileged EXEC mode

1. Enter the **copy source\_file destination\_file** command to overwrite the startup configuration with the default configuration.

```
switch# copy default-config startup-config
This operation will modify your startup configuration. Do you want to
continue? [Y/N]: y
```

2. Reboot the switch:

```
switch# reload system
```

The configuration restoration operation behaves differently depending on whether the switch is in standalone mode or part of a Brocade VCS Fabric.

In standalone mode, all interfaces are shut down. When the switch comes back up, the restored default configuration is used. The following parameters are unaffected by this command:

- Interface management IP address
- Software feature licenses installed on the switch

In VCS Fabric mode, all interfaces remain online. The following parameters are unaffected by this command:

- Interface management IP address
- Software feature licenses installed on the switch
- Virtual IP address

## Configuration management on a modular chassis

---

### NOTE

When the switch is in logical chassis cluster mode, the running-config file is saved automatically and does not need to be copied. There is no startup configuration for logical chassis cluster mode; therefore the information about startup configuration does not apply to logical chassis cluster mode.

---

The configuration data on a modular chassis is managed in a distributed fashion. The Brocade VDX 8770-4 and VDX 8770-8 chassis maintain two types of configuration data, global configuration parameters and slot configuration parameters. The global configuration, such as the VLAN configuration, applies to the entire chassis. The slot configuration includes specific parameters that apply only to the interface modules.

The startup configuration is maintained at the chassis level and includes both chassis-wide and slot-specific configuration parameters.

## Configuration management on interface modules

When an interface module (line card) boots up in a slot which was never occupied previously or is not configured, the module type is automatically saved in the configuration database. The type configuration associated with a given slot persists in the database even after the interface module is physically removed, powered off, or faulted. This mechanism ensures that all configuration data associated with a given slot is automatically preserved across reboots or hot swaps with the same type of interface module.

If you insert an interface module in a slot that was previously occupied by a module of a different type, the interface module will be faulted with a “type mismatch” error. Before you replace an interface module with a different type, you must clear the existing type configuration from the database. Refer to [“Replacing an interface module”](#) in [Chapter 3, “Basic Switch Management”](#) for more information.

---

### NOTE

The interface module configuration is non-persistent. You must issue the **copy running-config startup-config** command after the interface module comes online. Otherwise, all configuration data associated with the slot along with line module type will be lost after a chassis reboot.

---

## Configuration management in redundant management modules

In modular switches with redundant management modules, the VCS configuration, the startup configuration, and the startup database are synchronized and shared between the two management modules. The initial configuration synchronization occurs when the system boots up. After the initial synchronization has been completed successfully, synchronization can be triggered during the following events:

- When a failover occurs from the active management module to the standby management module. Unsaved configuration changes made on the active management module are lost after a failover. Issue the **copy running-config startup-config** command on the active management module to preserve the running configuration across a management module failover.
- When you insert a standby management module into a chassis after the active management module is already fully initialized.
- When you change the startup configuration by issuing the **copy running-config startup-config** command on the active management module.
- When you restore the default configuration by issuing the **copy default-config startup-config** command on the active management module.
- When you change the VCS configuration (VCS mode, RBridge ID, or VCS ID), the configuration change is synchronized with the standby management module and saved persistently. This event triggers a chassis reboot after the synchronization is complete.
- When you initiate a firmware download. Refer to [Chapter 6, “Installing and Maintaining Firmware”](#) for more information.

## Configuration management in Brocade VCS Fabric mode

With the exception of a few parameters, configuration changes you make to a single switch in a Brocade VCS Fabric are not automatically distributed. When configuring Ethernet fabric parameters and software features on multiple switches, you must configure each switch individually. To simplify the procedure, you can upload a configuration file from one switch and download it to the other switches in the fabric, provided the switches are of the same type.

---

### NOTE

The switches must be of the same model to share a configuration file. For example, downloading a configuration file from a Brocade VDX 6720-24 to a Brocade VDX 6720-60 or to a switch with a different firmware version may cause the switch to misapply the configuration and lead to unpredictable behavior.

---

To determine the switch type, issue the **show system** command. To map the switch type to the Brocade switch model name, refer to [“Switch types”](#) in [Chapter 3, “Basic Switch Management”](#).

If you need to reset affected switches, restore the default configuration as described in [“Restoring the default configuration”](#) on page 77.

### Downloading a configuration to multiple switches

---

**NOTE**

This section does not apply to logical chassis cluster mode because, in that mode, configuration is automatically distributed.

---

1. Configure one switch.
2. Copy the running configuration to the startup configuration as described in [“Saving the running configuration”](#) on page 75.
3. Upload the configuration to an external host ([“Uploading the startup configuration to an external host”](#) on page 76) or to an attached USB device as described in [“Backing up the startup configuration to a USB device”](#) on page 76.
4. Download the configuration file to each of the target switches. Refer to [“Configuration restoration”](#) on page 76 for more information.

### Automatic distribution of configuration parameters

A few configuration parameters are fabric-wide in fabric cluster mode. This means they are automatically distributed to all switches in a VCS fabric when you configure one or more of these parameters on a single RBridge that is part of a VCS fabric. These parameters include the following:

- Zoning configuration
- vCenter parameters
- Virtual IP address

---

**NOTE**

In logical chassis cluster mode, all configuration is automatically distributed.

---

The show running configuration command displays the same configuration for these features on all RBridge in the VCS fabric. Copy operations from any RBridge include all fabric-wide configuration parameters.



# Installing and Maintaining Firmware

---

## In this chapter

- [Firmware upgrade overview](#) ..... 81
- [Firmware upgrade in a local switch](#) ..... 84
- [Firmware upgrade using the manual option](#) ..... 88
- [Firmware upgrade in Brocade fabric cluster mode](#)..... 90
- [Firmware upgrade in Brocade logical chassis cluster mode](#) ..... 91

## Firmware upgrade overview

With the introduction of Network v4.0.0 logical chassis cluster mode, a new VCS mode type, is supported, in which both the data and configuration paths are distributed, and the entire cluster is configured from a principal node. Prior to Network v4.0.0, fabric cluster mode was the only VCS mode type supported, in which the data path for nodes is distributed, but the configuration path is not distributed; each node keeps its configuration database independently. When upgrading firmware, it is important to know which VCS mode is being used, because firmware upgrade instructions will differ depending on the VCS mode type.

Brocade firmware upgrades consist of multiple firmware 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 the firmware packages to be downloaded. These packages are made available periodically to add features or to remedy defects in the firmware.

Starting with Network OS v4.0.0, firmware upgrade is performed incrementally. The **firmware download** command compares the new firmware packages against the current installation and only downloads the packages that contain new features or have been modified.

You can download the firmware from a remote server using the File Transfer Protocol (FTP), SSH File Transfer Protocol (SFTP), or the Secure Copy Protocol (SCP), or you can download the firmware from an attached Brocade-branded USB device.

Brocade Network OS provides a single command line interface (CLI) to download firmware to a compact switch with a single control processor or to a modular chassis with two management modules. If the firmware download process is interrupted by an unexpected reboot, Network OS will make an attempt to recover the previously installed firmware. Success depends on the state of the firmware download. You must wait for the recovery to complete before initiating another firmware download.

---

### ATTENTION

Installing Network OS is disruptive to services, and any unsaved running configuration will be lost during the installation.

---

With the introduction of Network OS v4.0.0, the **logical-chassis firmware download** command allows you to upgrade a single switch or multiple switches of your choice that are connected in logical chassis cluster mode. This command can only be executed from the principal node (coordinator).

The firmware can only be downloaded from the file server through the management Ethernet port, so all nodes must have the management Ethernet port connected. Only one **logical-chassis firmware download** command instance can run at any given time.

If you are in logical chassis cluster mode, after you perform a firmware upgrade, you might find that the switch reverts to its default configurations. To preserve the configurations after an upgrade, back up the configuration using the **copy running-config filename** command before the firmware download. After the upgrade is completed, run the **copy filename running-config** command.

### Obtaining and decompressing firmware

Firmware upgrades are available for customers with support service contracts and for partners on the Brocade website at <http://www.mybrocade.com>.

You must download the firmware package either to an FTP server or to a USB device and decompress the package *before* you can use the **firmware download** command or **firmware download logical-chassis** command (if you are in VCS mode) to upgrade 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 firmware version. When issued with the path to the directory where the firmware is stored, the **firmware download** command or **firmware download logical-chassis** command (if you are in logical chassis cluster mode) performs an automatic search for the correct package file type associated with the device.

### Upgrading firmware on a compact switch

All Brocade compact switches maintain two partitions of nonvolatile storage areas, a primary and a secondary, to store two firmware images. The following steps describe the default behavior after you enter the **firmware download** command (without options) on a compact switch with a single control processor. The procedure applies to all Brocade VDX compact switches.

1. Network OS downloads the firmware to the secondary partition.
2. The switch swaps partitions and performs a reboot. After the system comes back up, the former secondary partition is the primary partition.
3. The system copies the firmware from the primary to the secondary partition and commits the new firmware.

The upgrade process first downloads and then commits the firmware. Use the **show firmwaredownloadstatus** command to monitor the upgrade progress.

### Upgrading firmware on a modular chassis

With the introduction of Network OS 4.0.0, In-service software upgrades (ISSUs) are supported. An ISSU allows a dual management module (MM) system to be upgraded nondisruptively and is invoked by entering the **firmwaredownload** command from the active management module.

On a dual management module system, the **firmwaredownload** command uses the ISSU protocol to upgrade the system. However, before the actual downloading occurs, it runs an ISSU sanity check and determine whether the ISSU will succeed. If the sanity check fails, users are notified to resort to the auto-leveling protocol to upgrade the system.

### *Automatic firmware synchronization*

When you replace a management module (MM) or insert a second management module into a chassis, the active management module automatically synchronizes the hot-plugged standby management module with the same firmware version. The standby management module reboots with the upgraded firmware. The automatic firmware synchronization takes place only if all of the following conditions are met:

- The standby management module is inserted while the chassis is up (hot-plugged insert).
- There was no firmware download process running when the standby management module was inserted.
- The active and standby firmware versions must be different.

---

#### NOTE

Automatic firmware synchronization is intrinsic to Network OS v4.0.0 and no corresponding enable or disable commands are associated with the feature. As a result, the feature cannot be disabled.

---

## Upgrading and downgrading firmware

In most cases, you will be upgrading firmware by installing a more recent firmware version than the one you are currently running. However, some circumstances may require that you downgrade the firmware to an earlier version. The procedures described in the following section assume that you are upgrading firmware, but they work for downgrading as well, provided that the firmware version you are downgrading to is compatible with the version you are currently running. [Table 9](#) displays supported firmware versions by platform.

**TABLE 9 Network OS firmware support by platform**

| Platform               | NOS v2.0.0 | NOS v2.1.0 | NOS v2.1.1 | NOS v3.0.0 | NOS v4.0.0 |
|------------------------|------------|------------|------------|------------|------------|
| Brocade VDX 6710       | no         | yes        | yes        | yes        | yes        |
| Brocade VDX 6720       | yes        | yes        | yes        | yes        | yes        |
| Brocade VDX 6730       | no         | yes        | yes        | yes        | yes        |
| Brocade VDX 6740/6740T | no         | no         | no         | no         | yes        |
| Brocade VDX 8770       | no         | no         | no         | yes        | yes        |

Brocade does not support upgrades from more than one previous release. For example, upgrading from Network OS v3.0.1 to v4.0.0 is supported, but upgrading from Network OS v3.0.0 directly to v4.0.0 is not supported, because there is another release between the two. In other words, upgrading a switch from Network OS v3.0.0 to v4.0.0 is a two-step process; you must first upgrade from v3.0.0 to v3.0.1 and then to v4.0.0 or later. When upgrading to the next patch release, choose the latest patch available for that release (for example, v2.1.0c).

If you are using a Brocade VDX 6740 or 6740T, you cannot download firmware earlier than Network OS v4.0.0.

Always refer to the release notes for compatibility information and take note of restrictions that may exist regarding upgrades and downgrades under particular circumstances.

## Firmware upgrade in a local switch

### Preparing for a firmware download

To prepare for a firmware download, perform the tasks listed in this section. In the unlikely event of a failure or timeout, you will be able to provide your switch support provider the information required to troubleshoot the firmware download.

1. Verify the current firmware version. Refer to [“Obtaining the firmware version”](#) on page 84 for details.
2. Decide on a migration path. Check the connected devices to ensure firmware compatibility and that any older versions are supported. Refer to the Network OS Compatibility section of the *Brocade Network OS Release Notes* for the recommended firmware version.
3. Back up your switch configuration prior to the firmware download. Refer to [Chapter 5, “Configuration Management”](#) for details.
4. *Optional:* For additional support, connect the switch to a computer with a serial console cable. Ensure that all serial consoles and any open network connection sessions, such as Telnet, are logged and included with any trouble reports.
5. Enter the **copy support** command to collect all current core files prior to executing the firmware download. This information helps to troubleshoot the firmware download process in the event of a problem.
6. *Optional:* Enter the **clear logging raslog** command to erase all existing messages in addition to internal messages.

### Obtaining the firmware version

Enter the **show version** command with the **all-partitions** option to obtain the firmware version for both primary and secondary partitions of each module.

```
switch# show version all-partitions
```

```
Network Operating System Software
Network Operating System Version: 4.0.0
Copyright (c) 1995-2013 Brocade Communications Systems, Inc.
Firmware name: NOS_v4.0.0
Build Time: 16:33:57 Jan 31, 2013
Install Time: 10:44:48 Jun 2, 2013
Kernel: 2.6.34.6
BootProm: 1.0.0
Control Processor: e500mc with 7168 MB of memory
```

| Slot | Name | Primary/Secondary Versions | Status  |
|------|------|----------------------------|---------|
| M1   | NOS  | NOS_v4.0.0<br>NOS_v4.0.0   | ACTIVE* |
| M2   | NOS  | NOS_v4.0.0<br>NOS_v4.0.0   | STANDBY |

|      |     |            |         |
|------|-----|------------|---------|
| L1/0 | NOS | NOS_v4.0.0 | ACTIVE  |
|      |     | NOS_v4.0.0 |         |
| L1/1 | NOS | NOS_v4.0.0 | STANDBY |
|      |     | NOS_v4.0.0 |         |

## Connecting to the switch

When you upgrade firmware in default mode, you connect to the switch through the management IP address. Modular switches have one management IP address for the chassis and separate IP addresses for each management module. To upgrade both management modules, you can either connect to the chassis management IP address or to the IP address of the active management module. If you want to upgrade a single management module only, you must connect to the IP address of that management module and run the **firmware download** command in manual mode. In manual mode, only the local management module is upgraded.

Use the **show system** command to display the management IP address for the chassis.

```
switch# show system
Stack MAC : 00:05:33:15:FA:70

-- UNIT 0 --
Unit Name : sw0
Switch Status : Online
Hardware Rev : 1000.0
TengigabitEthernet Port(s) : 56
Up Time : up 8:38
Current Time : 16:39:56 GMT
NOS Version : 3.0.0
Jumbo Capable : yes
Burned In MAC : 00:05:33:15:FA:70
Management IP : 10.24.73.131 <== Chassis Management IP address
Management Port Status : UP
```

Use the **show interface management** command to display the IP addresses for the management modules.

```
switch# show interface management
interface Management 10/1
ip address 10.24.73.130/20
ip gateway-address 10.24.64.1
ipv6 ipv6-address []
ipv6 ipv6-gateways []
line-speed actual "1000baseT, Duplex: Full"
line-speed configured Auto
interface Management 10/2
ip address 10.24.74.23/20
ip gateway-address 10.24.64.1
ipv6 ipv6-address []
ipv6 ipv6-gateways []
line-speed actual "1000baseT, Duplex: Full"
line-speed configured Auto
```

## Downloading firmware from a remote server

Under normal circumstances, Brocade recommends you run the **firmware download** command in default mode. Do not disable the auto-commit mode unless you want to evaluate a firmware upgrade before committing to it. Refer to [“Firmware upgrade using the manual option”](#) on page 88 for details about overriding the auto-commit mode.



### CAUTION

**Do not interrupt the firmware download process. If you encounter a problem, wait for the timeout (30 minutes for network problems) before issuing the firmware download command again. Disrupting the process (for example, by disconnecting the switch from the power source) can render the switch inoperable and may require you to seek help from your switch service provider.**

When upgrading multiple switches, complete the following steps on each switch before you upgrade the next one.

1. Perform the steps described in [“Preparing for a firmware download”](#) on page 84.
2. Verify that the FTP or SSH server is running on the remote server and that you have a valid user ID and password on that server.
3. Download the firmware package from the Brocade website to an FTP server.  
To download the firmware from an attached USB device, refer to [“Downloading firmware from a USB device”](#) on page 87.
4. Decompress the firmware archive.
5. Connect to the switch or management module you are upgrading.  
Refer to [“Connecting to the switch”](#) on page 85 for more information.
6. Issue the **show version** command to determine the current firmware version.
7. Do one of the following:
  - a. Enter the **firmware download interactive** command to download the firmware interactively. When prompted for input, choose defaults whenever possible.
  - b. You can invoke the **firmware download** command with the “noactivate” option. The “noactivate” option allows you to download the firmware to the node without activating it, (which requires a reboot). When you are ready to activate the firmware, you must invoke the **firmware activate** command. The “noactivate” option is available from the command line if a non-interactive option was used during the initial firmware download.

---

### NOTE

To be able to mention the FTP server by name, a Domain Name System (DNS) entry must exist for the server.

---

8. If you invoked the **firmware download** command using “interactive” mode, at the “Do you want to continue [y/n]:” prompt, enter **y**.
9. While the upgrade is proceeding, you can start a separate CLI session on the switch and use the **show firmwaredownloadstatus** command to monitor the upgrade progress.
10. After the switch reboots, enter the **show version** command to verify the firmware upgrade.

The following examples download firmware interactively on a modular switch using default options.

```

switch# firmware download interactive
Server name or IP address: 10.31.2.25
File name: /users/home40/Builds/NOS_v4.0.0
Protocol (ftp, sftp, scp): ftp
User: fvf
Password: *****
Do manual download [y/n]: n
System sanity check passed.

Do you want to continue? [y/n]:y

switch# firmware downloadscp host 10.70.12.110 directory /dist file release.plist
user fvf password pray4green
Performing system sanity check...
This command will use the ISSU protocol to upgrade the system. It will cause a
WARM reboot and will require that existing telnet, secure telnet or SSH sessions
be restarted.
Do you want to continue? [y/n]: y

```

## Downloading firmware from a USB device

The Brocade VDX 6710, VDX 6720, VDX 6730, and VDX 6740 switches, as well as the Brocade VDX 8770, support firmware download from a Brocade-branded USB device. You cannot use a third-party USB device. Before you can access the USB device, you must enable the device and mount it as a file system. The firmware images to be downloaded must be stored in the factory-configured *firmware directory*. Multiple images can be stored under this directory.

1. Ensure that the USB device is connected to the switch.
2. Enter the **usb on** command in privileged EXEC mode.

```

switch# usb on
Trying to enable USB device. Please wait...
USB storage enabled

```

3. *Optional:* Enter the **usb dir** command.

```

switch# usb dir
firmwarekey\ 0B 2013 Jun 15 15:13
support\ 106MB 2013 Jun 24 05:36
config\ 0B 2013 Jun 15 15:13
firmware\ 380MB 2013 Jun 15 15:13
NOS_v4.0.0\ 379MB 2013 Jun 15 15:31
Available space on usbstorage 74%

```

4. Enter the **firmware download usb** command followed by the relative path to the firmware directory.

```

switch# firmware download usb directory NOS_v4.0.0

```

5. *Optional:* Unmount the USB storage device.

```

switch# usb off
Trying to disable USB device. Please wait...
USB storage disabled.

```

## Activating nodes after an upgrade

If you run the **firmware download** command with the “noactivate” option from the command line, after firmware is downloaded to the node, the node will not be rebooted. You must run the **firmware activate** command to activate the firmware in the node. You cannot use the **reload** command. Doing so causes the old firmware to be restored.

The following example shows a request to activate the node after running **firmware download** with the “noactivate” option.

```
switch# firmware activate
```

```
This command will use the ISSU protocol to upgrade the system. It will cause a
WARM reboot and will require that existing telnet, secure telnet or SSH sessions
be restarted.
```

```
Do you want to continue? [y/n]: y
```

```
2010/01/29-23:48:35, [HAM-1004], 226, switchid 1, CHASSIS | VCS, INFO,
Brocade_Elara2, Switch will be rebooted with the new firmware.
```

## Firmware upgrade using the manual option



### CAUTION

Firmware download using the manual option will cause disruption to the traffic. Brocade recommends not to use it unless the release notes indicate otherwise or unless the switch is in certain error conditions in which the default firmwaredownload mode doesn't work.

In a dual management module (MM) system, the manual mode allows you to upgrade the MM on which the firmware download command is issued. Furthermore, it allows you to specify the “noreboot” or “nocommit” options so you have exact control of the MM upgrade sequence.

## Downloading firmware using the manual option

The following procedure applies to a compact switch or a single management module.

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.mybrocade.com> or from your switch support provider and store the file on the FTP, SFTP, or SSH server.
3. Unpack the compressed firmware archive.
4. Enter the **show version** command to view the current firmware version.
5. Enter the **firmware download interactive** command and respond to the prompts.
6. At the “Do Auto-Commit after Reboot [y/n]:” prompt, enter **n**.

```
switch# firmware download interactive
Server name or IP address: 10.31.2.25
File name: /users/home40/Builds/hydra_plat_dev01
Protocol (ftp, scp): ftp
```



```
User: fvt
Password: *****
Do manual download [y/n]: y
Reboot system after download? [y/n]:y
Do Auto-Commit after Reboot? [y/n]:n
```

System sanity check passed.

You are running firmware download on dual MM system with 'manual' option. This will upgrade the firmware only on the local MM.

This command will cause a cold/disruptive reboot and will require that existing telnet, secure telnet or SSH sessions be restarted.

```
Do you want to continue? [y/n]:y
(output truncated)
```

The switch performs a reboot and comes up with the new firmware. Your current CLI session will automatically disconnect.

7. Enter the **show version all-partitions** command to confirm that the primary partitions of the switch contain the new firmware.

## Committing the firmware upgrade

If you decide to keep the firmware upgrade, use the **firmware commit** command to update the secondary partitions with the new firmware. On modular switches, you must run the **firmware commit** command on both management modules. It may take several minutes to complete the commit operation.

1. Enter the **firmware commit** command in privileged EXEC mode.

```
switch# firmware commit
Validating primary partition...
Doing firmwarecommit now.
Please wait ...
Replicating kernel image
.....
FirmwareCommit completes successfully.
```

2. Enter the **show version** command with the **all-partitions** option.

Both partitions on the switch or on the modules should contain the new firmware.

## Restoring the previous firmware version

Use the **firmware restore** command to back out of a firmware upgrade. This option works only if auto-commit mode was disabled during the firmware download. On modular switches, you must run the **firmware restore** command on both management modules.



### CAUTION

The **firmware restore** command is disruptive and reboots the management modules. Therefore, invoking the **firmware restore** command is not recommended.

## Verification and error handling

After a firmware download completes, you should verify that the download has completed properly:

1. Run **show version all-partitions** to verify the management modules (MMs) and all linecard partitions have the correct firmware.
2. Run **show ha all-partitions** to verify the MMs and all linecard partitions are in HA sync.
3. Run **show slots** to verify the MMs and all linecards are in the “enabled” state.

If the MMs are running different firmware, you need to run the **firmware download** command with the “manual” option to update the standby MM to the same level as the active MM.

If a linecard is in the faulty state or the linecard partitions are not in sync, you need to run **power-off linecard** and **power-on linecard** commands to recover the linecard.

## Firmware upgrade in Brocade fabric cluster mode

The **firmware download** command supports local switch upgrades only. To upgrade all switches in a VCS fabric, you must execute the **firmware download** command on each switch separately.

---

### NOTE

For each switch in the fabric, complete the firmware download on the current switch before initiating a firmware download on another switch. This process minimizes traffic disruption.

---

Enter the **show firmwaredownloadstatus** command to verify that the download process is complete, and then move to the next switch.

You can also run the **firmware download** command with the “noactivate” option to download firmware to the nodes without rebooting them. After the firmware is downloaded to all nodes in the cluster, you can run **firmware activate** on each of the nodes to activate the firmware. This allows you to control the reboot sequence of the nodes to avoid traffic disruption in the cluster.

In the following example, a cluster contains four nodes (rbridge-ids 1 through 4). Node 1 is the principal node (coordinator). In the example, the switches need to be rebooted in a specific order, namely node 2, followed by node 3, then node 4, and finally node 1. This can be accomplished by invoking the following sequence of commands from any node.

1. Invoke **firmware download rbridge-id all path path**.
2. Invoke the **firmware activate** command on each node in the following order:
  - a. **firmware activate rbridge-id 2**
  - b. **firmware activate rbridge-id 3**
  - c. **firmware activate rbridge-id 4**
  - d. **firmware activate rbridge-id 1**

## Firmware upgrade in Brocade logical chassis cluster mode

In logical chassis cluster mode, you can upgrade the cluster by logging in to individual nodes and running the **firmware download** command the same as in fabric cluster mode.

Another method for upgrading the logical chassis cluster is by specifying the “logical-chassis” and “rbridge-id” options in the **firmware download** command, which is also referred to as the **firmware download logical-chassis** command. This command allows users to upgrade one or more nodes in the cluster from the principal node. The nodes to be upgraded are specified in the “rbridge-id” option.

After the **firmware download logical-chassis** command is issued, firmware is downloaded to the specified nodes simultaneously through their respective management ports. The download time will not increase based on the number of nodes. By default, after firmware is downloaded to all of the specified nodes, the **firmware download logical-chassis** command will exit. It does *not* reboot the nodes.

You will then need to explicitly issue the **firmware activate rbridge-id** command to activate the new firmware in the nodes. This way, you have exact control over the reboot sequence of the nodes in the cluster. If you want to automatically activate the new firmware, the “auto-activate” option must be specified.



### CAUTION

When the “auto-activate” option is specified, all of the specified nodes in the command will be rebooted at the same time, which can cause traffic disruption. Auto-activate is therefore not recommended.

During the **firmware download logical-chassis** command, you can run the **show firmwaredownloadstatus rbridge-id all** command to verify the download status on the nodes. The status of the nodes after logical-chassis firmware download is completed should display “Ready to activate.” You can also run the **show version rbridge-id all** command to verify the firmware versions on the nodes.

The general procedure for upgrading firmware in logical chassis cluster mode is as follows:

1. Run **firmware download logical-chassis rbridge-id all** to upgrade all nodes in the cluster.

```
switch# firmware download logical-chassis protocol ftp host 10.1.2.30 user fvt
password brocade directory /dist/nos/4.0.0bld01 file release.plist rbridge-id
all
```

| Rbridge-id | Sanity Result        | Current Version |
|------------|----------------------|-----------------|
| 1          | Non-disruptive(ISSU) | 4.0.0           |
| 2          | Disruptive           | 4.0.0           |
| 3          | Disruptive           | 4.0.0           |

2. Run **show firmwaredownloadstatus summary rbridge-id all** and **show version rbridge-id all** to check the firmware download status of the nodes.

```
switch# show firmwaredownloadstatus summary bridge-id all
Rid 1: INSTALLED (Ready for activation)
Rid 2: INSTALLED (Ready for activation)
Rid 3: INSTALLED (Ready for activation)
switch# show version rbridge-id all
rbridge-id 1
```

```

Network Operating System Software
Network Operating System Version: 4.0.0
Copyright (c) 1995-2012 Brocade Communications Systems, Inc.
Firmware name: 4.0.0
Build Time: 19:59:29 Apr 26, 2013
Install Time: 20:36:28 Apr 26, 2013
Kernel: 2.6.34.6
BootProm: 2.2.0
Control Processor: e500v2 with 2048 MB of memory
Appl Primary/Secondary Versions

NOS 4.0.0
 4.0.0bld01

```

3. If any nodes fail to download, issue another **firmware download logical-chassis rbridge-id** *rbridge-id* command to download firmware to the failed nodes and bring all nodes to the same firmware level. Verify that the secondary partition of the nodes has the new firmware before proceeding with the output of this command.
4. Run **firmware activate rbridge-id** *rbridge-id* to activate the nodes in the desired order.

---

### NOTE

All of the nodes specified in the *rbridge-id* parameter in the **firmware activate** command will be rebooted at the same time.

---

```

switch# firmware activate rbridge-id 1-2,3
This command will activate the firmware on the following nodes.

```

```

Rbridge-id Sanity Result

1 Non-disruptive(ISSU)
2 Disruptive
3 Disruptive

```

It will cause these nodes to reboot at the same time.  
Do you want to continue? [y/n]: **y**

## Recovering firmware

The **firmware recover** command is used to discard the firmware that was downloaded using the **firmware download noactive** command.

```

switch# firmware recover rbridge-id 1,2,3
This command will restore the firmware on the following nodes.
Rbridge-id 1, 2, and 3.
Do you want to continue? [y/n]: y

```

## Verification and error handling

Any node failure or principal node failover can cause the logical chassis firmware download to be aborted. After the **firmware download logical-chassis** command is returned, you should verify that the command has been completed successfully:

1. Run **show firmwaredownloadstatus rbridge-id** *rbridge-id* to verify that the nodes have completed the installation and are in the "Ready for activation" state.
2. Run **show version rbridge-id** *rbridge-id* to verify that the nodes have the correct firmware.

If any nodes are not in the “Ready for activation” state or do not have the correct firmware installed, you should issue another **firmware download logical-chassis** command to upgrade the failed nodes before activating the new firmware.

If the logical chassis firmware download keeps failing on any of the nodes, you might need to log in to the failed nodes and recover them. For instructions, refer to [“Firmware upgrade in Brocade logical chassis cluster mode”](#) on page 91.

## 6 Firmware upgrade in Brocade logical chassis cluster mode

# Administering Licenses

## In this chapter

- [Licensing overview](#) ..... 95
- [Permanent licenses](#) ..... 97
- [Temporary licenses](#) ..... 97
- [Managing licenses](#) ..... 99
- [Dynamic Ports on Demand](#) ..... 104
- [Upgrade and downgrade considerations](#) ..... 109

## Licensing overview

The Brocade Network Operating System (Network OS) includes platform support in standalone and Brocade VCS Fabric modes as well as optional features that are enabled by license keys. You can purchase Brocade licenses per product or per feature. Each switch in a fabric must have its own licenses, but universal licenses for multiple switches are available for trial purposes along with individual trial licenses. Licenses may be part of the licensed paperpack supplied with your switch software, or you can purchase them separately from your switch vendor. [Table 10](#) lists the license requirements by platform.

**TABLE 10** Licenses requirements by platform

| Platform               | 40Gb<br>DCB<br>ports | 10Gb<br>DCB<br>ports | 1Gb<br>DCB<br>ports | 8G<br>FC<br>ports | License type                                | POD sets and<br>sizes       | Notes                                                                                                                                         |
|------------------------|----------------------|----------------------|---------------------|-------------------|---------------------------------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Brocade<br>VDX 6720-24 | N/A                  | 24                   | N/A                 | N/A               | POD1, FCoE, VCS<br>Fabric                   | 16 + 8                      | Enabling VCS Fabric mode no longer requires a full POD licenses. Layer 3 features do not require a license.                                   |
| Brocade<br>VDX 6720-60 | N/A                  | 60                   | N/A                 | N/A               | POD1, POD2,<br>FCoE, VCS Fabric             | 40 + 10 + 10                |                                                                                                                                               |
| Brocade<br>VDX 6710    | N/A                  | 6                    | 48                  | N/A               | VCS Fabric                                  | N/A                         | No POD license is required to enable DCB ports. FCoE licensees are not supported on this platform. Layer 3 features do not require a license. |
| Brocade<br>VDX 6730-32 | N/A                  | 24                   | N/A                 | 8                 | POD1, FCoE, VCS<br>Fabric                   | 16 + 8 (CEE<br>ports)       | FC ports require an FCoE license to be enabled (FC ports are not POD-controlled). Layer 3 features do not require a license.                  |
| Brocade<br>VDX 6730-76 | N/A                  | 60                   | N/A                 | 16                | POD1, POD2,<br>FCoE, VCS Fabric,<br>Layer 3 | 40 + 10 + 10<br>(CEE ports) |                                                                                                                                               |

**TABLE 10 Licenses requirements by platform (Continued)**

| Platform              | 40Gb<br>DCB<br>ports | 10Gb<br>DCB<br>ports | 1Gb<br>DCB<br>ports | 8G<br>FC<br>ports | License type                                       | POD sets and<br>sizes | Notes                                                      |
|-----------------------|----------------------|----------------------|---------------------|-------------------|----------------------------------------------------|-----------------------|------------------------------------------------------------|
| Brocade<br>VDX 8770-4 | 48                   | 192                  | 288                 | N/A               | FCoE, VCS Fabric,<br>Layer 3, Advanced<br>Services | N/A                   | There are no POD-enabled<br>ports on the VDX 8770 chassis. |
| Brocade<br>VDX 8770-8 | 96                   | 384                  | 576                 | N/A               | FCoE, VCS Fabric,<br>Layer 3, Advanced<br>Services | N/A                   | There are no POD-enabled<br>ports on the VDX 8770 chassis. |

[Table 11](#) provides descriptive details for each license type.

**TABLE 11 Brocade licenses for optional Network OS features**

| License                                                                                                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Dynamic Ports on Demand: <ul style="list-style-type: none"> <li>PORTS_ON_DEMAND_1</li> <li>PORTS_ON_DEMAND_2</li> </ul> | <p>A Dynamic POD license allows you to instantly scale the fabric by provisioning additional Ethernet ports on the Brocade VDX 6720 and VDX 6730 platforms. Licenses are assigned dynamically from a pool of resources based on either auto-detection of active links or explicit port reservation by the user.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Brocade VCS Fabric <ul style="list-style-type: none"> <li>VCS_FABRIC</li> </ul>                                         | <p>A Brocade VCS Fabric license allows you to provision a Brocade VCS Fabric with up to 24 nodes. You must install a Brocade VCS Fabric license on each node. A Brocade VCS Fabric license is required if your Brocade VCS Fabric includes more than two nodes.</p> <p>A switch with a Brocade VCS Fabric license cannot connect to a switch without such a license. In a two-node VCS, both switches must either have a VCS license, or no VCS license.</p> <p>FC router links do not affect Brocade VCS Fabric licensing. You can connect a switch to an FC router with or without a Brocade VCS Fabric license. FC router links and domains do not count against the two-switch limit for provisioning a VCS Fabric without a Brocade VCS Fabric license.</p>                                                                                                                                                                                                                      |
| FCoE <ul style="list-style-type: none"> <li>FCOE_BASE</li> </ul>                                                        | <p>A Brocade FCoE license is required to enable Fibre Channel over Ethernet functionality. You can use this license on a single a switch, but the FCoE capabilities will be limited to that node only. In order to support multi-hop FCoE traffic, you must enable VCS Fabric mode and install a Brocade FCoE license on each node. A VCS fabric in excess of two nodes requires a VCS Fabric license in addition to the FCoE license to allow FCoE traffic traverse all nodes in the fabric. Without an FCoE license, FCoE logins are not permitted, FCoE traffic does not transit the switch, and most FCoE commands return an error of “No FCoE license present” when executed.</p> <p>Brocade VDX 6730 switches must be in Brocade VCS Fabric mode and have the FCoE license installed to enable the FC ports. FC ports are not POD controlled. The FCoE license is not available on the Brocade VDX 6710. Directly attached FCoE devices are not supported on this platform.</p> |



**TABLE 11 Brocade licenses for optional Network OS features (Continued)**

| License                                                                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Layer 3 <ul style="list-style-type: none"> <li>LAYER_3</li> </ul>                     | <p>On all Brocade VDX compact switches, Layer 3 features are automatically enabled and do not require a license. On modular switches (Brocade VDX 8770-4 and Brocade 8770-8) you must install a license to activate the following Layer 3 features:</p> <ul style="list-style-type: none"> <li>OSFP</li> <li>VRRP</li> <li>PIM-SM</li> <li>Route-maps</li> <li>Prefix-List</li> </ul> <p>No license is required for the following Layer 3 features:</p> <ul style="list-style-type: none"> <li>IGMP Snooping</li> <li>L3 ACL</li> <li>L3 QoS</li> <li>RTM/Static Routes</li> <li>Interface IP Address</li> <li>ARP</li> <li>IP Services (ping/Telnet/FTP)</li> </ul> <p>Brocade Layer 3 licenses are independent of any other licenses or operating modes in the fabric. Brocade Layer 3 licenses are required only on modular switches, and only on those switches in the fabric that have interfaces attached that will use the additional licensed L3 features. Any intermediate switches in the fabric between interfaces do not require the license to be installed. Layer 3 features are only supported in VCS mode.</p> |
| <ul style="list-style-type: none"> <li>Advanced Services</li> </ul> ADVANCED_SERVICES | <p>For the Brocade VDX 8770 platforms, you have the option of purchasing and installing a single license key that will activate FCoE, Layer 3, and VCS features in one convenient bundle. You may also purchase and install each license separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Brocade offers three different models for the Network OS licensed features. They include permanent licenses and two types of temporary trial licenses: individual time-based licenses and universal time-based licenses.

## Permanent licenses

A permanent license (also referred to as a *chassis-wide license*) has no expiration date and is locked to a single switch identified by the switch license ID. The switch license ID is initially the same as the switch World Wide Name (WWN). The switch WWN may change through configuration changes on the product, but the switch license ID remains unchanged. Use only the switch license ID to obtain licenses. Refer to [“Displaying the switch license ID”](#) on page 99 for instructions on how to obtain the license ID of your switch.

## Temporary licenses

A temporary license (also known as a *time-based license*) allows you to evaluate a feature for a limited time prior to buying a permanent license. Brocade offers two types of temporary licenses: individual time-based licenses and universal time-based licenses.

### Individual time-based licenses

An individual time-based license is locked to a single switch and has a fixed expiration date. You cannot install this license on multiple switches.

### Universal time-based licenses

A universal time-based license allows you to use a given feature for a limited trial period defined in days, for example, 30, 60, or 90 days from the date you install the license key on the switch. Each universal license key is valid for a single feature and can be used on any product that supports the feature. In addition, each universal time-based license has an absolute shelf life after which it expires. You cannot install a license with an expired shelf life.

The expiration date is based on the system time at the installation of the license plus the number of days that the universal time-based license is valid. For this reason, you cannot remove and reinstall a universal time-based license.

### License expiration

When a time-based license expires, the feature continues to work while generating warning messages until the switch reboots. Configuration options that require a license will also fail after the license has expired. After the next reboot, the feature will no longer work.

You can display expired licenses with the **show license** command. Expired licenses display a “License has expired” message. RASlog warning messages are generated every hour for licenses that have expired or will expire in the next five days.

### Extending a license

You extend a time-based license by adding another temporary license or by installing a permanent license. Re-installing a temporary license that has expired is not permitted. When you replace an expired license, the warning messages cease.

### Usage restrictions

The following restrictions apply to all time-based licenses:

- Time-based licenses are always retained in the license database and cannot be deleted.
- Once you have installed a time-based license, you cannot change the system date or time.

---

#### NOTE

Other mechanisms for changing date and time, such as Network Time Protocol (NTP), are not blocked. If you are using NTP to synchronize the time between your network devices, including switches or enterprise-class platforms, do not attempt to change the system date and time when a time-based license is installed.

---

## Managing licenses

The following management tasks and associated commands apply to both permanent and temporary licenses.

---

### NOTE

License management in Network OS v3.0.0 is supported only on the local RBridge. You cannot configure or display licenses on remote nodes in the fabric.

---

## Displaying the switch license ID

The switch license ID identifies the switch for which the license is valid. You will need the switch license ID when you activate a license key.

To display the switch license ID, enter the **show license id** command in the privileged ECEC mode.

```
switch# show license id
Rbridge-Id License ID
=====
2 10:00:00:05:33:54:C6:3E
```

## Obtaining a license key

License upgrade orders are fulfilled either through a license activation paperpack, or by an e-mail message containing a transaction key and a link to the Brocade software portal. A device-specific license file is generated in the software portal when you enter the transaction key along with the switch license ID. Use the **show license id** command to obtain the switch license ID.

Follow the instructions in the paperpack or the e-mail message as described for your platform and license type. The transaction key is case-sensitive; you must enter the key exactly as it appears in the paperpack. To lessen the chance of an error, copy and paste the transaction key when you install the license on your switch.

You will receive an e-mail message with the software license keys embedded in an XML file along with installation instructions.

---

### NOTE

Store the license key in a safe place for future reference. The **show license** command does not print out the license key.

---

## Installing a license

1. Open the e-mail message that contains the license key and extract the license key from the XML file. The license key is printed between the XML start `<licKey>` and end `</licKey>` tags. Be sure to copy the entire string, including spaces and non-alphanumeric characters.
2. Enter the **license add licstr** command followed by the license key and, optionally, an RBridge-ID if you are installing the license on a remote switch. If the license key includes spaces, you must enclose the entire string in double quotation marks.
3. Verify that you added the license by entering the **show license** command. The command lists all licensed features currently installed on the switch. If the feature is not listed, enter the **license add licstr** command again.

Depending on the license type, you may be prompted to reload the switch or to disable and re-enable the chassis or specific ports. [Table 12](#) indicates the minimal steps you may need to take to make the installed features fully functional after the license add operation is complete. Take the appropriate action as indicated by the command output.

**TABLE 12** Requirements for activating a license after installation

| License                                | Description                                                                                                                                                                                                                         |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PORTS_ON_DEMAND_1<br>PORTS_ON_DEMAND_2 | One of the following actions may be required depending on the configuration: <ul style="list-style-type: none"> <li>Enabling the ports or the chassis.</li> <li>Disabling and then re-enabling the ports or the chassis.</li> </ul> |
| VCS_FABRIC                             | One of the following actions may be required depending on the configuration: <ul style="list-style-type: none"> <li>Enabling the ports or the chassis.</li> <li>Disabling and then re-enabling the ports or the chassis.</li> </ul> |
| FCOE_BASE                              | You must enable any Fibre Channel (FC) ports that you wish to use (applicable only if the product includes FC ports).                                                                                                               |
| LAYER_3                                | One of the following actions may be required depending on the configuration: <ul style="list-style-type: none"> <li>Enabling the ports or the chassis</li> <li>Disabling and then re-enabling the ports or the chassis</li> </ul>   |
| ADVANCED_SERVICES                      | Combines requirements for FCoE, Layer 3, and VCS licenses.                                                                                                                                                                          |

### *Adding a Dynamic POD license*

The following example adds a second Dynamic POD license on the local switch and verifies the transaction. The command prompts for disabling and then re-enabling the port or the switch.

```
switch# license add licstr "*B
s1SETgzTgeVGUDeQR4WIfRx7mmXODdSwENoRGENAmX3Ca3uHeZgXK0b,
jzxyzfzKLrMsPN8ClSxvDQRRT8VyuULyyKTO0ryU6qm4s1jjiSAeV,COedzCxlv6ycQgnYMeSVp#
"

License Added [*B s1SETgzTgeVGUDeQR4WIfRx7mmXODdSwENoRGENAmX3Ca3uHeZgXK0b,
jzxyzfzKLrMsPN8ClSxvDQRRT8VyuULyyKTO0ryU6qm4s1jjiSAeV,COedzCxlv6ycQgnYMeSVp#
]

For license change to take effect, please disable/enable port or switch...
switch# chassis disable
switch# chassis enable
switch# show license
Rbridge-Id: 2
XX
First Ports on Demand license - additional 10 port upgrade license
Feature name:PORTS_ON_DEMAND_1
XX
Second Ports on Demand license - additional 10 port upgrade license
Feature name:PORTS_ON_DEMAND_2
```

#### **NOTE**

You must install all licenses on the local switch. Remote management of licenses is not supported in Network OS v3.0.0.

### *Adding a Brocade VCS Fabric license*

The following example adds a Brocade VCS Fabric license on the local Brocade VDX 6720 switch and verifies the transaction. The license takes effect immediately after the command is executed. No further action is required.

```
switch# license add licstr "*B
r84pNRtHKdRZujmwAUT63GORXIpBhBZK0ckRq6Bvv13Strvw1:fUjANF
av5W:gWx3hH2:9RsMv3BHfeCRFM2gj9NlkrdIiBPBOa4xfSD2jf,Xx1RwksliX8fH6gpx7,73t#"

Adding license [*B r84pNRtHKdRZujmwAUT63GORXIpBhBZK0ckRq6Bvv13Strvw1:fUjANF
av5W:gWx3hH2:9RsMv3BHfeCRFM2gSLj9NlkrdIiBPBOa4xfSD2jf,Xx1RwksliX8fH6gpx7,73t#
]
switch# show license
Rbridge-Id: 2
XX
First Ports on Demand license - additional 10 port upgrade license
Feature name:PORTS_ON_DEMAND_1
XX
Second Ports on Demand license - additional 10 port upgrade license
Feature name:PORTS_ON_DEMAND_2
XX
VCS Fabric license
Feature name:VCS_FABRIC
```

### *Adding an FCoE license*

The following example adds an FCoE license on the local switch and verifies the transaction.

```
switch# license add licstr "*B
:YFGuJSHxbhlWVwBHjmjfaO20R6QzolkyVR4oqJAU0fqhJRCTioav1A:
HMah2E7uL4d8px4ySTAWSg809etcLwfpLjgXZ1lvWiiKEWcfcZMefx#"

License Added [*B :YFGuJSHxbhlWVwBHjmjfaO20R6QzolkyVR4oqJAU0fqhJRCTioav1A:
HMah2E7uL4d8px4ySTAWSg809etcLwfpLjgXZ1lvWiiKEWcfcZMefx#]

switch# show license
Rbridge-Id: 2
XX
First Ports on Demand license - additional 10 port upgrade license
Feature name:PORTS_ON_DEMAND_1
XX
Second Ports on Demand license - additional 10 port upgrade license
Feature name:PORTS_ON_DEMAND_2
XX
VCS Fabric license
Feature name:VCS_FABRIC
XX
FCoE Base license
Feature name:FCOE_BASE
```

### *Displaying a license*

You display installed licenses with the **show license** command.

The following example displays a Brocade VDX 8770 licensed for a Layer 3 VCS fabric. This configuration does not include FCoE features.

```
switch# show license
```

```

rbridge-id: 60
xx
 Layer 3 license
 Feature name:LAYER_3
xx
 VCS Fabric license
 Feature name:VCS_FABRIC

```

The following example displays a Brocade VDX 8770 licensed for Advanced Services. This configuration enables the use of Layer 3, FCoE, and VCS features.

```

switch# show license
rbridge-id: 60
xx
 Advanced Services license
 Feature name:ADVANCED_SERVICES

```

## License Removal

Depending on the license type, you may be prompted to clear license-dependent configurations, reload the switch, or disable and re-enable the chassis or specific ports. [Table 13](#) indicates the minimal steps you may need to take to remove specific licenses. Take the appropriate action as indicated by the command output.

**TABLE 13** Requirements for deactivating a license after removal

| License           | Description                                                                                                                                                                                                                  |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PORTS_ON_DEMAND_1 | One of the following actions may be required after removing the license: <ul style="list-style-type: none"> <li>Enabling the ports or the chassis.</li> <li>Disabling and then re-enabling the ports the chassis.</li> </ul> |
| PORTS_ON_DEMAND_2 |                                                                                                                                                                                                                              |
| VCS_FABRIC        | Disabling the chassis is required before you can remove the license.                                                                                                                                                         |
| FCOE_BASE         | Disabling all Fibre Channel (FC) ports is required before you can remove the license (applicable only if the product includes FC ports).                                                                                     |
| LAYER_3           | Clearing of Layer 3 configurations on all ports is required before you can remove the license.                                                                                                                               |
| ADVANCED_SERVICES | Combines requirements for FCoE, Layer 3, and VCS licenses.                                                                                                                                                                   |

For some licensed features, you must clear all configurations related to a the feature before you can remove the license for that feature. Some features may require that you reboot the switch and others require you to disable and re-enable selected ports or the entire switch.

For example, removing an FCoE license requires both explicit reset to default of FCoE settings that depend on the license, as well as disabling all FC ports on the platform; all FC ports must be set to a “shut” interface configuration state before the license can be removed.

Refer to the console output for other specific requirements.

### *Removing a license*

1. Enter the **show license** command to display the active licenses.
2. Issue the **license remove** command followed by the license key or the feature name.

The license key is case-sensitive and must be entered exactly as shown. If the license key includes spaces, you must enclose the entire string in double quotation marks.

3. Take the appropriate action as indicated by the command output.

Depending on the license type, you may be prompted to clear license-related features, to reboot the switch, or to disable and re-enable the chassis or specific ports.

4. Enter the **show license** command to verify that the license is removed. If there are no license keys, the command output displays “No licenses.”

---

#### NOTE

You must remember the original license string to use the **license remove** command with the *licenseString* operand. You cannot display the license key with the **show license** command.

---

The following example illustrates the display and removal of an FCoE license by its feature name.

```
switch# show license
Rbridge-Id: 2
xx
 FCoE Base license
 Feature name:FCOE_BASE
xx
 First Ports on Demand license - additional 10 port upgrade license
 Feature name:PORTS_ON_DEMAND_1
xx
 Second Ports on Demand license - additional 10 port upgrade license
 Feature name:PORTS_ON_DEMAND_2

switch# license remove "FCOE_BASE"

License Removed [FCOE_BASE]

For license to take effect, enable the switch and any disabled ports...
```

The remaining licenses are displayed as follows:

```
switch# show license
Rbridge-Id: 2
xx
 First Ports on Demand license - additional 10 port upgrade license
 Feature name:PORTS_ON_DEMAND_1
xx
 Second Ports on Demand license - additional 10 port upgrade license
 Feature name:PORTS_ON_DEMAND_2
```

## Dynamic Ports on Demand

Dynamic Ports on Demand provides a flexible mechanism for allocating port licenses that you can purchase to extend the base functionality of the Brocade VDX 6720 and VDX 6730 switches. The additional ports are enabled after you install the appropriate license keys. The Dynamic POD feature assigns port licenses based on your connectivity choices. Any port on the switch can claim a free assignment from the pool of available POD licenses.

---

### NOTE

DPOD licenses are not supported on the Brocade VDX 6710, VDX 8770, and on the FC ports on the Brocade VDX 6730. When adding a DPOD license to a platform that does not support DPOD, you will be allowed to add the license, but the license display will show that the DPOD license is not supported on the platform.

---

In a Dynamic POD system, each port can be associated with one of three port sets:

- **Base Port Set** – Ports that can be enabled without any POD license.
  - **Single POD License Port Set** – Ports that are assigned first and associated with the existence of a single POD license.
  - **Double POD License Port Set** – Ports that are assigned after the single POD port set is full and are therefore associated with the double POD license.
- 

### NOTE

Licenses are based on the license ID and are not interchangeable between units. For example, if you bought a single POD license for a Brocade VDX 6720-24, you cannot use that license on another Brocade VDX 6720-24, on a Brocade VDX 6720-60, or on a Brocade VDX 6730-76.

---

You can purchase the Brocade VDX 6720 and VDX 6730 switches with the port options indicated in [Table 14](#). You can activate unlicensed ports up to the maximum supported per switch by purchasing and installing additional POD licenses.

**TABLE 14** List of available ports when implementing PODs

| Platform            | Base port set | Single POD set | Double POD set | Total ports |
|---------------------|---------------|----------------|----------------|-------------|
| Brocade VDX 6720-24 | 16            | +8             | N/A            | 24          |
| Brocade VDX 6720-60 | 40            | +10            | +10            | 60          |
| Brocade VDX 6730-32 | 16            | +8             | N/A            | 24          |
| Brocade VDX 6730-76 | 40            | +10            | +10            | 60          |

---

If you purchase both a PORTS\_ON\_DEMAND\_1 (POD1) license and a PORTS\_ON\_DEMAND\_2 (POD2) license for the Brocade VDX 6720-60 or VDX 6730-76, the system has a “double” POD license. If you purchased only one of these features (POD1 or POD2), the system has a “single” POD license. The specific POD license that is installed is not relevant to the port count determination. Only the number of installed POD licenses is relevant.

---

### ATTENTION

The Brocade VDX 6720-24 and VDX 6730-32 support only a single POD license.

---



## Automatic POD port assignments

With the Dynamic POD feature, you can use the base port set plus the number of additional ports you purchased. All ports that do not receive a POD assignment and are trying to come online will go offline. The **show ip interface brief** and **show interface tengigabitethernet rbridge-id/slot/port** commands display the reason for the port disabled status as related to POD licensing.

The Dynamic POD mechanism detects the ports that have active links, and it makes assignments based on the remaining pool of vacancies:

- If the count of assigned ports is below the number of ports in the purchased POD set, additional dynamic assignments can be made at a later time as new links are established. If a port comes online, that port can get assigned if you still have vacancies in your POD set.
- If the number of detected active links is greater than the number of ports in the purchased POD set, port assignments are made in the order in which the ports come online. Because the time it takes for each port to come online varies, the order in which ports are assigned to a given POD set cannot be guaranteed.

If the given assignment order does not align with your intended use of the ports, you can make adjustments using the **dpo rbridge-id/slot/port reserve** or the **dpo rbridge-id/slot/port release** commands. Refer to [“Overriding Dynamic POD assignments”](#) on page 107 for more information.

## Mapping port assignments to a POD port set

Ports are associated with the single or double POD license in the order in which they come online and automatically receive a license assignment from the pool of unassigned ports in the POD set. The first ports that receive a POD assignment are associated with the base port set. When all ports in the base port set are assigned, the next ports that come online receive assignments from the single POD license port set. When this set is full, the remaining port assignments are associated with the double POD license port set.

The association of a specific port to a POD set matters only when you want to remove a POD license from the system. Ports assigned to the double POD license port set will be disabled first. The last-assigned licenses will be released first when you remove POD licenses. Refer to [“Releasing a port from a POD set”](#) on page 108 for more information.

## Activating the Dynamic POD feature

1. Verify the current states of the ports with the **show ip interface brief** command.

The command output indicates whether a port is licensed.

2. Install the Brocade Dynamic POD license.

For instructions on how to install a license, refer to [“Installing a license”](#) on page 99.

3. Use the **shutdown** and **no shutdown** commands to disable and re-enable the ports.

Alternatively, you can disable and re-enable the chassis to activate ports.

4. Use the **show ip interface brief** command to verify the newly activated ports.

5. Use the **show interface tengigabitethernet rbridge-id/slot/port** command to display port details.

The following example shows a Brocade VDX 6720-24 without a Dynamic POD license installed. The 16 ports in the base port set are online and assigned. The remaining 8 ports are unassigned and are down.

```
switch# show ip interface brief
Interface IP-Address Status Protocol
=====
TengigabitEthernet 5/0/1 unassigned up up
TengigabitEthernet 5/0/2 unassigned up up
TengigabitEthernet 5/0/3 unassigned up up
TengigabitEthernet 5/0/4 unassigned up up
TengigabitEthernet 5/0/5 unassigned up up
TengigabitEthernet 5/0/6 unassigned administratively down down (No DPOD
License)
TengigabitEthernet 5/0/7 unassigned administratively down down (No DPOD
License)
TengigabitEthernet 5/0/8 unassigned up up
TengigabitEthernet 5/0/9 unassigned up up
TengigabitEthernet 5/0/10 unassigned up up
TengigabitEthernet 5/0/11 unassigned administratively down down (No DPOD
License)
TengigabitEthernet 5/0/12 unassigned administratively down down (No DPOD
License)
TengigabitEthernet 5/0/13 unassigned up up
TengigabitEthernet 5/0/14 unassigned up up
TengigabitEthernet 5/0/15 unassigned up up
TengigabitEthernet 5/0/16 unassigned up up
TengigabitEthernet 5/0/17 unassigned administratively down down (No DPOD
License)
TengigabitEthernet 5/0/18 unassigned administratively down down (No DPOD
License)
TengigabitEthernet 5/0/19 unassigned administratively down down (No DPOD
License)
TengigabitEthernet 5/0/20 unassigned administratively down down (No DPOD
License)
TengigabitEthernet 5/0/21 unassigned up up
TengigabitEthernet 5/0/22 unassigned up up
TengigabitEthernet 5/0/23 unassigned up up
TengigabitEthernet 5/0/24 unassigned up up
```

The following example displays details for a single port that is offline because it does not have a Dynamic POD license.

```
switch# show interface tengigabitethernet 5/0/6
TengigabitEthernet 5/0/6 is down, line protocol is down (No DPOD License)
Hardware is Ethernet, address is 0005.1eb6.0a25
Current address is 0005.1eb6.0a25
Tracking status: Disabled
Tracked interfaces: None
Pluggable media present, Media type is sfp
Interface index (ifindex) is 1744896001
MTU 2500 bytes
LineSpeed: Auto - 10000 Mbit, Duplex: Full
Flowcontrol rx: on, tx: on
```

## Displaying the Dynamic POD assignments

To display the Dynamic POD assignments, enter the **show dpod** command.

The **show dpod** command provides a summary of POD license status and POD license assignments.

In the following example, all 24 ports are licensed and potentially available. The three unassigned ports are currently persistently disabled and therefore are not assigned to any Dynamic POD license port set.

```
switch# show dpod
rbridge-id: 1
24 ports are available in this switch
 1 POD license is installed
 Dynamic POD method is in use
 24 port assignments are provisioned for use in this switch:
 16 port assignments are provisioned by the base switch license
 8 port assignments are provisioned by the first POD license
 * 0 more assignments are added if the second POD license is installed
 21 ports are assigned to installed licenses:
 16 ports are assigned to the base switch license
 5 ports are assigned to the first POD license
 Ports assigned to the base switch license:
 Te 1/0/1, Te 1/0/10, Te 1/0/11, Te 1/0/12, Te 1/0/13, Te 1/0/14, Te 1/0/15, Te
1/0/16, Te 1/0/17, Te 1/0/18, Te 1/0/19, Te 1/0/20, Te 1/0/21, Te 1/0/22, Te
1/0/23, Te 1/0/24
 Ports assigned to the first POD license:
 Te 1/0/5, Te 1/0/6, Te 1/0/7, Te 1/0/8, Te 1/0/9
 Ports assigned to the second POD license:
 None
 Ports not assigned to a license:
 Te 1/0/2, Te 1/0/3, Te 1/0/4

 3 license reservations are still available for use by unassigned ports
```

## Overriding Dynamic POD assignments

You can override the automatic port license assignments by releasing Dynamic POD assignments from a port and by reserving an assignment for a specific port.

### *Reserving a port assignment*

Reserving an assignment for a port assigns that port to a POD license regardless of whether the port is online or offline. Reserving assignments allocates the POD license to specified ports. This operation overrides automatic port assignments. The reserved assignment will not be available to other ports that come online. To reserve an assignment for a port, a free assignment must be available.

1. Enter the **show dpod** command to determine the unassigned ports.

If all ports are assigned, select a port to release its POD assignment. Follow the instructions in [“Releasing a port from a POD set”](#) on page 108 to release a port from its POD assignment. Once the port is released, you can reuse the assignment for another port.

2. Enter the global configuration mode by issuing the **configure terminal** command.

3. Select the port for which you want to reserve an assignment and enter the **dpod reserve** command.
4. Enter the **exit** command to return to the global configuration mode before you reserve another port.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# dpod 5/0/10 reserve
switch(config-dpod-5/0/10)# exit
switch(config)# dpod 5/0/11 reserve
switch0(config-dpod-5/0/11)# exit
```

---

## NOTE

License reservations or removals do not persist across switch reboots and power cycles. To make them persistent, save the configuration changes by issuing the use **copy running-config startup-config** command before you reboot the switch.

---

5. Save the configuration changes.
6. Reboot the switch.
7. Enter the show running-config dpod command to verify the port is reserved.

```
switch# show running-config dpod 5/0/10
dpod 5/0/10
 reserve
 !
switch# show running-config dpod 5/0/11
dpod 5/0/11
 reserve
 !
```

## *Releasing a port from a POD set*

Once a port has been assigned to a Dynamic POD license port set, it remains licensed (or “reserved”) until you remove the port from the port set. You remove a port from the port set by releasing the port with the **dpod release** command. Releasing a port removes it from the Dynamic POD license port set; the port appears as unassigned until it comes back online.

To prevent a port from coming back online and taking a POD assignment, disable the port and save the running configuration. This action will disable the port persistently.

A port POD assignment can only be released if the port is currently offline. Use the **shutdown** command to disable the port or use the **chassis disable** command to disable the switch if you plan to release multiple ports.

1. Enter the global configuration mode by issuing the **configure terminal** command.
2. Select the interface for the port that you wish to disable using the **interface rbridge-id/slot/port** command.
3. Enter the **shutdown** command to take the port offline.
4. Enter the **exit** command to return to the global configuration mode before you release the port.
5. Enter the **dpod release** command to remove the port from the POD license.

6. Enter the **exit** command to return to the global configuration mode before you reserve another port.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface 1/0/10
switch(config-if-te-1/0/10)# shutdown
switch(config-if-te-1/0/10)# exit
switch(config)# dpod 1/0/10 release
switch(config-dpod-1/0/10)# exit
```

7. Enter **exit** to return to the privileged EXEC mode.
8. Enter the **show dpod** command to verify that the port is no longer assigned to a POD set.
9. Enter the **enable chassis** command to bring the switch back online.
10. Save the configuration changes.

```
switch# copy running-config startup-config
```

---

#### NOTE

Do not release a port unless you plan to disconnect the optical link or disable the port persistently. If you leave the link in a state where the port could be brought online, the POD mechanism will detect this unassigned port and attempt to reassign it to a port set.

---

## Upgrade and downgrade considerations

Downgrading from Network OS v.3.0.0 to a previous version is applicable only to the Brocade VDX 6710, VDX 6720, and VDX 6730. The VDX 8770 chassis only supports Network OS v.3.0.0. There are no downgrade pre-install requirements for any of the licensed features in Network OS v.3.0.0.

You cannot downgrade to Network OS v2.x when the switch is in Brocade VCS Fabric mode without full POD licenses installed. Network OS v2.x requires full POD licenses to be installed when the switch is in Brocade VCS Fabric mode.

If Dynamic POD reserve port configuration data is stored in the running-config file before you downgrade to Network OS v2.x, restoring that running-config file after downgrade does not restore the Dynamic POD reserve port information.

## Configuration management considerations

Licenses are independent of configuration files and are therefore not affected when you make changes to a configuration file or restore the default configuration. The only exceptions are Dynamic POD configurations. When you download a configuration from another switch or you restore the default configuration, the DPOD configuration is moderated and restricted by whatever POD licenses are present on the switch.

For example, if you download a configuration file from a switch that has POD assignments beyond the base number of allowed DPOD ports and there is no POD license installed, then the additional ports beyond the number of ports in the Base DPOD set will not be allowed. When the configuration is played back as part of the copy (or reboot) operation, the licenses are checked for each additional DPOD port that attempts to be assigned a reservation.

## 7 Configuration management considerations

# SNMP

---

## In this chapter

|                                                          |     |
|----------------------------------------------------------|-----|
| • Simple Network Management Protocol .....               | 111 |
| • Basic SNMP operation .....                             | 112 |
| • Understanding MIBs .....                               | 113 |
| • Loading Brocade MIBs .....                             | 114 |
| • SNMP community strings .....                           | 116 |
| • SNMP server hosts .....                                | 118 |
| • Support for multiple SNMP server contexts .....        | 120 |
| • Support for password encryption for SNMPv3 users ..... | 120 |
| • Displaying SNMP configurations .....                   | 121 |

## 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 14](#). The agent accesses information about a device and makes it available to an SNMP network management station.



FIGURE 14 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 15](#)). 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.

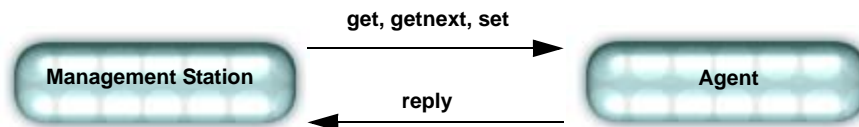


FIGURE 15 SNMP query

The management station can also receive *traps*, unsolicited messages from the switch agent if an unusual event occurs ([Figure 16](#)).



FIGURE 16 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 17](#), the Entity MIB OID is:

```
1.3.6.1.2.1.47
```

The corresponding name is:

```
iso.org.dod.internet.mgmt.mib-2.entityMIB
```

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.

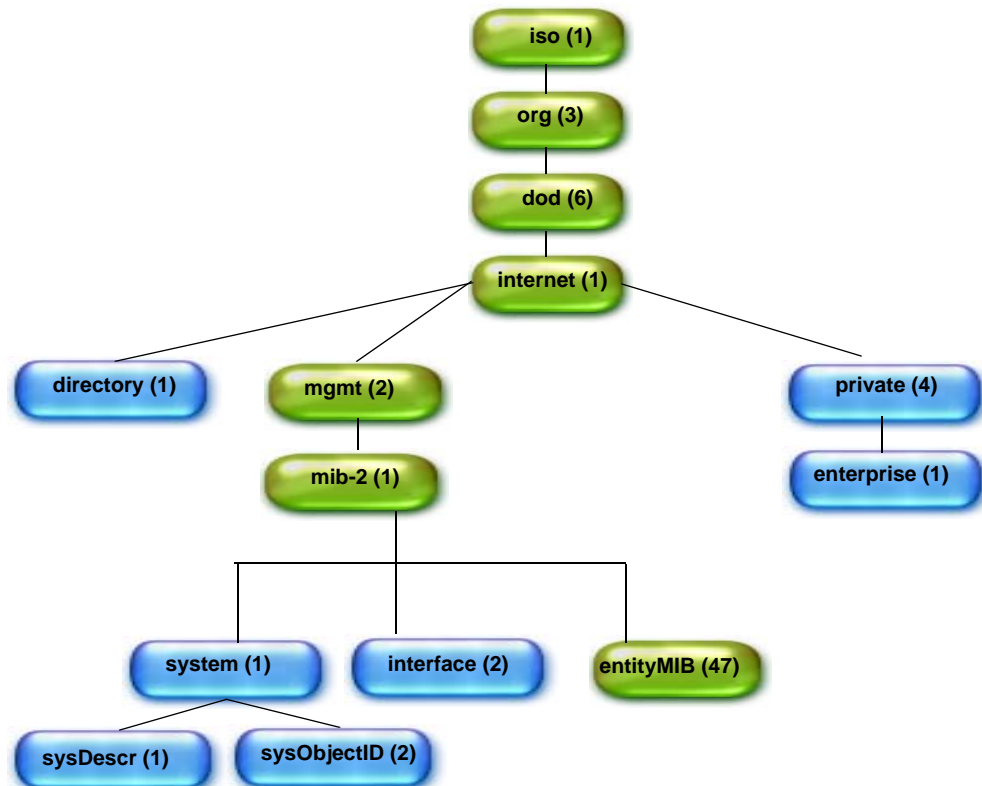


FIGURE 17 Brocade MIB tree location

### 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.

## 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:

- BRCD\_NOS\_PRODUCTS.mib
- BROCADE-PRODUCTS-MIB.mib
- BROCADE-REG-MIB.mib
- BRCD\_TC.mib
- SWBase.mib
- Resource.mib
- System.mib
- FA.mib
- HA.mib

## Agent Capability MIBs

In SNMP, capability MIBs provide the implementation details for the associated MIBs. These MIBs, called AGENT-CAPABILITY MIBs, list supported conformance groups and any deviations from the MIBs as implemented in the associated software version. Table 15 lists the Brocade supported capability MIBs.

**TABLE 15 Agent Capabilities**

| Capability MIBs                      | Description                                                   |
|--------------------------------------|---------------------------------------------------------------|
| BROCADE-IEEE8021-PAE-CAPABILITY-MIB  | Provides the implementation details for the IEEE8021-PAE-MIB  |
| BROCADE-IEEE8023-LAG-CAPABILITY-MIB  | Provides the implementation details for the IEEE8023-LAG-MIB  |
| BROCADE-LLDP-CAPABILITY-MIB          | Provides the implementation details for the LLDP-MIB          |
| BROCADE-LLDP-EXT-DOT3-CAPABILITY-MIB | Provides the implementation details for the LLDP-EXT-DOT3-MIB |

## Standard MIBs

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

- IF-MIB
- LLDP-MIB
- BRIDGE-MIB
- IP-MIB
- LLDP-EXT-DOT3-MIB
- LLDP-EXT-DOT1-MIB
- RSTP-MIB
- RFC1213-MIB
- IEEE8023-LAG-MIB
- Q-BRIDGE-MIB
- IEEE8021-PAE-MIB
- P-BRIDGE-MIB
- RMON-MIB
- SFlow-MIB
- INET-ADDRESS-MIB
- IANAifType-MIB
- IANA-RTPROTO-MIB
- SNMPV2-MIB
- SNMP-FRAMEWORK-MIB
- IANA-ADDRESS-FAMILY-NUMBERS-MIB
- FC-MGMT-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 [Figure 17](#) to ensure any MIB dependencies are loading in the correct order.

### NOTE

Before loading the Brocade MIB files, ensure that you have the correct version of SNMP for the Brocade Network OS. All versions of Network OS support SNMPv1, SNMPv2, and SNMPv3. SNMPv2c is supported in Network OS v2.0.0, but SNMPv2c informs are not supported.

**TABLE 16 Brocade SNMP MIB dependencies**

| MIB Name                 | Dependencies                                                       |
|--------------------------|--------------------------------------------------------------------|
| Brocade-REG-MIB          | RFC1155-SMI                                                        |
| Brocade-TC               | Brocade-REG-MIB<br>SNMPv2-TC<br>SNMPv2-SMI                         |
| BRCD_NOS_PRODUCTS.mib    | SNMPv2-SMI<br>Brocade-REG-MIB                                      |
| BROCADE-PRODUCTS-MIB.mib | SNMPv2-SMI<br>Brocade-REG-MIB                                      |
| SWBase.mib               | SNMPv2-TC<br>SNMPv2-SMI<br>Brocade-REG-MIB                         |
| Resource.mib             | SNMPv2-TC<br>SNMPv2-SMI<br>SWBASE-MIB                              |
| System.mib               | SNMPv2-TC<br>Brocade-TC<br>SWBASE-MIB                              |
| FA.mib                   | RFC1155-SMI<br>RFC-1212<br>RFC1213-MIB<br>RFC-1215                 |
| HA-MIB                   | SNMPv2-SMI<br>Brocade-REG-MIB<br>SW-MIB<br>ENTITY-MIB<br>SNMPv2-TC |

## SNMP community strings

SNMP versions 1 and 2c use community strings to restrict access to the switch. There are six default community strings: three read-write strings and three read-only strings. There is support for a total of 256 SNMP communities, all user-configurable.

The following default community strings are read-write:

- Secret C0de
- OrigEquipMfr
- private

The following default community strings are read-only:

- public
- common
- ConvergedNetwork

## Adding an SNMP community string

The **snmp-server community** command sets the community string and read-write or read-only access for each community, and is applicable to SNMPv1 and SNMPv2c. You can execute this command in global configuration mode.

1. Enter the **configure terminal** command.
2. Enter the **snmp-server community string [ro | rw]** command.

```
switch(config)# snmp-server community private rw
```

- The *string* variable specifies the community string name. The string can be from 2 to 16 characters long.
- The **ro** or **rw** option specifies whether the string is read-only (ro) or read-write (rw).

The command in the example adds the read-write SNMP community string “private” with read-write access.

## Changing the access of a read-only community string

This example changes the access of “user123” from read-only to read-write.

1. Enter the **configure terminal** command.
2. Enter the **snmp-server community string rw** command.

```
switch(config)# snmp-server community user123 rw
```

## Removing an SNMP community string

This example removes the community string “private”.

1. Enter the **configure terminal** command.
2. Enter the **no snmp-server community string [ro | rw]** command.

```
switch(config)# no snmp-server community private
```

## SNMP server hosts

The **snmp-server host** command sets the trap destination IP addresses, SNMP version, community string for SNMPv1 and SNMPv2c, contact information, the destination port for the SNMP server host, and the severity level.

To configure SNMP trap hosts associated with community strings, you must create the community string using the **snmp-server community** command before configuring the host.

The SNMP agent supports six default communities and their associated trap recipients and trap recipient severity levels. The default value for each attribute is as follows: host = 0.0.0.0; udp-port = 162; severity-level = none. The length of the community string must be from 2 through 16 characters. The community strings have the following default values:

- common—read-only
- public—read-only
- ConvergedNetwork—read-only
- OrigEquipMfr—read-write
- private—read-write
- Secret C0de—read-write

### Setting the SNMP server host

You can execute SNMP server commands in global configuration mode.

1. Enter the **configure terminal** command.
2. Enter the **snmp-server host** *ipv4\_host* | *ipv6\_host* | *dns\_host* *community-string* [**version** {**1** | **2c**}] [**udp-port** *port*] [**severity-level** {**none** | **debug** | **info** | **warning** | **error** | **critical**}] command.
  - The *ipv4\_host* | *ipv6\_host* | *dns\_host* variable specifies the IP address of the host.
  - The *community-string* variable sets the community string.
  - The **version** option selects SNMPv1- or SNMPv2c-related configuration parameters. These parameters include the community string. The default SNMP version is 1.
  - The **udp-port** *port* option specifies the UDP port where SNMP traps will be received. The default port is 162. The acceptable range of ports is from 0 through 65535.
  - The **severity** *sev\_level* option provides the ability to filter traps based on severity level on both the host and the v3host. Only RASlog (swEvent) traps can be filtered based on severity level. If the severity level of None is specified, all traps are filtered and no RASLOG traps are received. If the severity level of Critical is specified, no traps are filtered and all traps are received by the host.

Severity level options include None, Debug, Info, Warning, Error, and Critical.

The following example sets up “commaccess” as a read-only community string and sets 10.32.147.6 as a trap recipient with SNMPv2c on target port 162.

```
switch(config)# snmp-server host 10.32.147.6 commaccess version 2c udp-port
162 severity warning
```

## Removing the SNMP server host

The **no snmp-server host** *host community-string string version 2c* command brings version 2c down to version 1.

The **no snmp-server host** *host community-string string* command removes the SNMP server host from the switch configuration altogether.

## Configuring the SNMP system group

These tasks allow you to configure the system contact and system location objects for the SNMP system group.

### *Setting the SNMP server contact*

Use the **snmp-server contact** command to set the SNMP server contact string. The default contact string is "Field Support." The number of characters allowed is from 4 through 255.

1. Enter the **configure terminal** command.
2. Enter the **snmp-server contact** *string* command.

```
switch(config)# snmp-server contact "Operator 12345"
```

The example changes the default contact string to "Operator 12345." You must enclose the text in double quotes if the text contains spaces.

### *Removing the SNMP server contact*

The **no snmp-server contact** *string* command restores the default contact information (Field Support).

### *Setting the SNMP server location*

Use the **snmp-server location** command to set the SNMP server location string. The default SNMP server location string is "End User Premise." The number of characters allowed is from 4 through 255.

1. Enter the **configure terminal** command.
2. Enter the **snmp-server location** *string* command.

```
switch(config)# snmp-server location "Building 3 Room 214"
```

3. Enter the **no snmp-server location** command to remove the location.

The example changes the default location string to "Building 3 Room 214." You must enclose the text in double quotes if the text contains spaces.

### *Setting the SNMP server description*

Use the **snmp-server sys-descr** command to set the SNMP server description string. The default SNMP server description string is "Brocade-VDX-VCS <vcid>." The number of characters allowed is from 4 through 255.

1. Enter the **configure terminal** command.
2. Enter the **snmp-server sys-descr** *string* command.

```
switch(config)# snmp-server sys-descr "Brocade-VDX Test Bed"
```

3. Enter the **no snmp-server sys-descr** command to remove the location.

The example changes the default location string to “Brocade-VDX Test Bed.” You must enclose the text in double quotes if the text contains spaces.

## Support for multiple SNMP server contexts

A single SNMP agent can be supported by the multiple instances of the same MIB module by mapping the context name with the VRF. The context can be mapped to a VRF as described in [“Setting the SNMP server context”](#). The SNMP agent supports 256 contexts to support context to VRF mapping.

### Setting the SNMP server context

Use the **snmp-server context** command to map a context to the name of a VPN routing and forwarding (VRF) instance.

1. Enter the **configure terminal** command.
2. Enter the **snmp-server context** *context\_name* **vrf-name** *vrf\_name* command.

```
switch(config)# snmp-server context mycontext vrf-name myvrf
```

The example maps the context name “mycontext” to the VRF name “myvrf.”

## Support for password encryption for SNMPv3 users

For SNMPv3 user, the passwords for **auth-password** and **priv-password** are encrypted. You can configure either with plain text password or encrypted password. In both the cases, the passwords are shown in the **show running-config** command as encrypted.

```
sw0(config)# snmp-server user snmpadmin2 groupname snmpadmin auth md5
auth-password "MVb+360X3kcfBzug5Vo6dQ==" priv DES priv-password
"ckJFoHbzVvhR0xFRPjsMTA==" encrypted
```

---

#### NOTE

This command may not be successful where encrypted passwords are generated by third-party or open-source tools.

---



## Displaying SNMP configurations

Use the **show running-config snmp-server** command to display the current SNMP configurations for the SNMP host, community string, contact, and location, as well as other SNMP configuration options such as SNMPv3 host address, context, and VRF mapping. There are no default configurations for this command. This command can only be executed in Privileged EXEC command mode.

Enter the **show running-config snmp-server** command.

```
switch# show running-config snmp-server
snmp-server contact "Field Support."
snmp-server location "End User Premise."
snmp-server sys-descr "Brocade VDX Switch."
snmp-server community ConvergedNetwork
snmp-server community OrigEquipMfr rw
snmp-server community "Secret C0de" rw
snmp-server community common
snmp-server community private rw
snmp-server community public
snmp-server user snmpadmin1 groupname snmpadmin
snmp-server user snmpadmin2 groupname snmpadmin
snmp-server user snmpadmin3 groupname snmpadmin
snmp-server user snmpuser1
snmp-server user snmpuser2
snmp-server user snmpuser3
```

## 8 Displaying SNMP configurations

# Fabric

---

## In this chapter

- [TRILL](#) ..... 123
- [Brocade VCS Fabric formation](#) ..... 123
- [Brocade VCS Fabric configuration management](#) ..... 126
- [Fabric interface configuration management](#) ..... 128

## TRILL

The Brocade VCS Fabric Ethernet fabric is defined as a group of switches that exchange information between each other to implement distributed intelligence. The Brocade Ethernet fabric uses Transparent Interconnection of Lots of Links (TRILL) protocol, designed for the sole purpose of scaling Ethernet networks by allowing a set of devices, called routing bridges (*RBridges*), to connect with each other.

A link state dynamic routing protocol, rather than Spanning Tree Protocol, determines how the traffic is forwarded between the inter-connected RBridges. Link state routing in Brocade VCS Fabric-based TRILL networks is performed using Fabric Shortest Path First (FSPF) protocol.

TRILL enables Layer 2 networks to behave like routed Layer 3/IP networks. TRILL also defines native support for forwarding both unicast and multicast traffic, and therefore unifies support for both of these different classes of applications over a single transport.

## Brocade VCS Fabric formation

Brocade VCS Fabric technology uses RBridge identifiers (IDs) to discover fabric creation problems, such as duplicate IDs. The RBridge ID of a cluster unit is equal to the domain ID of an FC switch. RBridge ID assignment is implemented by leveraging the domain ID assignment protocols in the FC SANs. Request for Domain ID (RDI) and Domain ID Assignment (DIA) protocols ensure that a single switch (the principal switch) is centrally allocating the domain IDs for every RBridge in the fabric and detecting any domain ID conflicts in the fabric. In case of conflict, the conflicting node is segmented from the fabric. You must take action to resolve the conflict

---

### NOTE

Network OS v4.0.0-based fabrics can have a maximum of 239 RBridges in a single Brocade VCS Fabric. However, Brocade recommends using only 24 RBridges per fabric.

---

The following sequence of events describes the Brocade VCS Fabric formation process:

- Each Brocade VCS Fabric is identified by a VCS ID.
- All Brocade VCS Fabric-capable switches are configured with a factory default VCS ID of 1.

- The switch software searches for the “VCS enable” attribute.

**NOTE**

If the software cannot locate the “VCS enable” attribute, the switch goes into standalone mode and operates like a regular 802.1x switch.

- Assuming the switch is Brocade VCS Fabric-enabled, the switch software invokes a series of protocols:
  - Brocade Link Discovery Protocol (BLDP) attempts to discover if a Brocade VCS Fabric-capable switch is connected to any of the edge ports. See [“Neighbor discovery”](#) on page 125 for more information.
  - BLDP attempts to merge the adjacent Brocade switch into the Brocade VCS Fabric environment at the link level.
- A series of FC fabric formation protocols (RDI, DIA, and FSPF) are initiated once a link level relationship has been established between two neighbor switches. See [“Fabric formation”](#) on page 125 for more information.
- Merge and Join Protocol invokes a merge of switch configuration between the cluster units once the fabric has successfully formed.

## How RBridges work

RBridges find each other by exchanging FSPF Hello frames. Like all TRILL IS-IS frames, Hello frames are transparently forwarded by RBridges and are processed by RBridge Inter-Switch Link (ISL) ports. Using the information exchanged in the Hello frames, the RBridges on each link elect the designated RBridge for that link.

The RBridge link state includes information such as VLAN connectivity, multicast listeners, and multicast router attachment, claimed nicknames, and supported ingress-to-egress options. The designated RBridge specifies the appointed forwarder for each VLAN on the link (which could be itself) and the designated VLAN for inter-RBridge communication. The appointed forwarder handles native frames to and from that link in that VLAN.

The Ingress RBridge function encapsulates frames from the link into a TRILL data frame. The Egress RBridge function decapsulates native frames destined for the link from the TRILL data frames. TRILL data frames with known unicast destinations are forwarded by Rbridge next hop. Multi-destination frames (broadcast, unknown unicast, and multicast) are forwarded on a tree rooted at the multicast root RBridge.

- Unicast forwarding is handled by combining domain routing generated by FSPF and MAC-to-RBridge learning generated by MAC learning and a distributed MAC database.
- Multicast forwarding usually uses one tree that is rooted at the RBridge with the lowest RBridge ID. However, there are several rules for Multicast root tree selection. It is not always the lowest RBridge ID.

If a duplicated RBridge-id is found while the links are still coming up, the links are segmented. Both sides recognize the error and segment the link. If the RBridge-id overlap cannot be found at ISL link bringup time (in the case where a new switch is brought from an offline state into the fabric) it will be found during the fabric build and the conflicting switch is isolated.

An RBridge requests a specific RBridge ID from the coordinator switch. If the coordinator switch detects that this RBridge ID is already used, it returns the next unused RBridge ID. The requesting RBridge is not allowed to take another RBridge ID and it segments itself from the fabric. In this case, you cannot boot the ISLs. The ISLs have to be explicitly disabled and then enabled again in order for the RBridge with the overlapping RBridge ID to be removed.

## Neighbor discovery

Brocade VCS Fabric-capable neighbor discovery involves the following steps:

- Discover whether the neighbor is a Brocade switch.
- Discover whether the Brocade neighbor switch is Brocade VCS Fabric-capable.

Only Brocade VCS Fabric-capable switches with the same VCS ID can form a virtual cluster switch. The default settings for Brocade Network OS switches are Brocade VCS Fabric capable and a VCS ID of “1.”

## Brocade trunks

The Network OS v4.0.0 supports Brocade trunks (hardware-based link aggregation groups, or LAGs). These LAGs are dynamically formed between two adjacent switches. The trunk formation is controlled by the same FC Trunking protocol that controls the trunk formation on FC switches. As such, it does not require user intervention or configuration except enabling or disabling, which instructs the switch software to form a trunk at the global level or not. See [“Enabling a Fabric trunk”](#) on page 129 for instructions.

All ISL ports connected to the same neighbor Brocade switch are attempted to form a trunk. For a successful trunk formation, all ports on the local switch must be part of the same port group and must be configured at the same speed. Rules for these trunks are as follows:

- Eight ports are allowed per trunk group on the Brocade VDX 6720 and VDX 6730.
- Sixteen ports per trunk are allowed on the Brocade VDX 6740.
- On the Brocade VDX 6740, traffic is not broadcasted evenly on all members.
- The trunk is turned on by default.

---

### NOTE

Trunks are not supported between the Brocade 8000 and the Brocade VDX 8770.

---

## Fabric formation

Brocade VCS Fabric technology leverages proven FC Fabric protocols to build a TRILL fabric. The main functions of the fabric formation protocols are to:

- Assign the Brocade VCS Fabric-wide unique RBridge IDs (Domain ID Assignment)
- Create the network topology database using link state routing protocol (Fabric Shortest Path First, or FSPF). FSPF calculates the shortest path routes to a destination RBridge.
- Distribute fabric multicast traffic.

### *Principal switch selection*

Every Brocade VCS Fabric-enabled switch, upon boot-up and after the Fabric port formation, declares itself to be a principal switch and advertises this intent on all fabric ports. The intent includes a priority and its switch WWN. If all switches boot up at the same time, the default priority is the same and all switches will compare their mutual intents. The switch with the lowest Switch WWN becomes the principal switch. The WWN is an industry-standard burnt-in switch identifier, similar to the Bridge-MAC except it is 8 bytes. The role of the principal switch is to decide whether a new RBridge joining the fabric conflicts with any of the RBridge IDs already present in the fabric.

---

#### **NOTE**

Brocade VDX Data Center switches are shipped with factory-programmed world wide names (WWNs) and are unique.

---

At the end of the principal switch selection process, all the switches in the cluster have formed a tree with the principal switch at the root.

---

#### **NOTE**

In a logical chassis cluster, you can select the principle node by using the command line interface. For more information, see the [“Selecting a principal node for the cluster”](#) on page 43.

---

### *RBridge ID allocation*

RBridge ID assignment is implemented by leveraging proven Domain ID assignment protocols from FC SANs. Request for Domain ID (RDI) and Domain ID Assignment (DIA) protocols ensure that a single switch (the principal switch) centrally allocates the domain IDs for every RBridge in the fabric and detects and resolves any domain ID collisions in the fabric. Brocade VCS Fabric supports up to 24 RBridge IDs.

Only the principal switch can allocate RBridge IDs (domain IDs) for all other switches in the fabric. The principal switch starts the allocation process by allocating an RBridge ID for itself (using the ID value supplied by the user), and initiates the DIA messages on all ports.

Other switches, which are now in subordinate mode, upon receiving the DIA frames respond with an RDI message towards the principal switch. The process continues until all the switches in the fabric have been allocated a unique ID.

### **Fabric routing protocol**

After a RBridge ID is assigned to a switch, the Fabric Shortest Path First (FSPF) link state routing protocol begins to form adjacencies and collects topology and inter-connectivity information with its neighbors. Brocade VCS Fabric uses FSPF to calculate and elect a loop-free multicast tree rooted at the multicast root RBridge. The multicast tree is calculated after the unicast routes are computed.

## **Brocade VCS Fabric configuration management**

Complete the following configuration steps to add a new switch into a fabric.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **vcs rbridge-id rbridge-id enable** command.

The switch remembers its RBridge ID once it has been assigned. The **vcs rbridge-id rbridge-id enable** command also sets the insistent RBridge ID property on the switch.

### 3. Reboot the system.

After the required reboot, the switch participates in the RBridge ID allocation protocol which insists that the same value that was manually configured prior to reboot be allocated after reboot.

The switch is not allowed into the fabric if there is a conflict; for example, if another switch with the same ID exists and is operational in the fabric. You have the opportunity to select a new RBridge ID using the same CLI.

Once an ID has been assigned by the fabric protocol, these IDs are then numerically equated to RBridge IDs and are treated as such after that.

Use the **vcs** command to configure the Brocade VCS Fabric parameters, VCS ID, and the switch RBridge ID, and to enable Brocade VCS Fabric mode (also called *VCS mode*).

VCS mode encompasses two mode types:

- Fabric cluster mode—The data path for nodes is distributed, but the configuration path is not distributed. Each node keeps its configuration database independently.
- Logical chassis cluster mode—Both the data and configuration paths are distributed. The entire cluster can be configured from the principal node. Logical chassis cluster mode requires NOS 4.0 or later.

The generic term “VCS mode” in this manual applies to both fabric cluster mode and logical chassis mode unless otherwise stated.

You can set the Brocade VCS Fabric parameters and enable VCS mode at the same time, or you can enable VCS mode and then perform the ID assignments separately. Refer to [Table 17](#), [Table 18](#), and [Table 19](#) for details.

After configuring the Brocade VCS Fabric parameters, the switch applies the changes and reboots.

The switch disable is not saved across a reboot, so if the switch was disabled prior to the reboot, the switch returns to the enabled state when it finishes the boot cycle.

## Brocade VCS Fabric configuration tasks

[Table 17](#) contains command examples for enabling Brocade VCS logical chassis cluster mode:

**TABLE 17** Command examples for enabling logical chassis cluster mode

| Command                                                          | Command Behavior                                                                                                                    |
|------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| switch# <b>vcs logical-chassis enable</b>                        | The VCS ID becomes the default value of 1, the RBridge ID is not changed, and Brocade VCS logical chassis cluster mode is enabled.  |
| switch# <b>vcs vcsid 22 rbridge-id 15 logical-chassis enable</b> | The VCS ID is changed to 22, the RBridge ID is changed to 15, and Brocade VCS logical chassis cluster mode is enabled.              |
| switch# <b>vcs vcsid 11 logical-chassis enable</b>               | The VCS ID is changed to 11, the RBridge ID is not changed, and Brocade VCS logical chassis cluster mode is enabled.                |
| switch# <b>vcs rbridge-id 6 logical-chassis enable</b>           | The VCS ID becomes the default value of 1, the RBridge ID is changed to 6, and Brocade VCS logical chassis cluster mode is enabled. |

Table 18 contains command examples for enabling Brocade VCS fabric cluster mode:

**TABLE 18** Command examples for enabling fabric cluster mode

| Command                                          | Command Behavior                                                                                                         |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| switch# <b>vcs enable</b>                        | The VCS ID becomes 1, the RBridge ID is not changed, and Brocade VCS fabric cluster mode is enabled.                     |
| switch# <b>vcs vcsid 55 rbridge-id 19 enable</b> | The VCS ID is changed to 55, the RBridge ID is changed to 19, and Brocade VCS fabric cluster mode is enabled.            |
| switch# <b>vcs vcsid 73 enable</b>               | The VCS ID is changed to the value of 73, the RBridge ID is not changed, and Brocade VCS fabric cluster mode is enabled. |
| switch# <b>vcs rbridge-id 10 enable</b>          | The VCS ID becomes the default value 1, the RBridge ID is changed to 10, and Brocade VCS fabric cluster mode is enabled. |

Table 19 contains command examples for switches that are already in either fabric cluster mode or logical chassis cluster mode.

**TABLE 19** Command examples when one of the VCS modes is already enabled:

| Command                           | Command Behavior                                                 |
|-----------------------------------|------------------------------------------------------------------|
| <b>vcs vcsid 44 rbridge-id 22</b> | The VCS ID is changed to 44 and the RBridge ID is changed to 22. |
| <b>vcs vcsid 34</b>               | The VCS ID is changed to 34.                                     |

## Fabric interface configuration management

A physical interface in a virtual switch cluster can either be an edge port or a fabric port, but not both. Similar to a switch-port configuration on a physical interface, you can also change a fabric-port configuration on its physical interface by using the **fabric isl enable** and **fabric trunk enable** commands, described below.

### Enabling a Fabric ISL

The **fabric isl enable** command controls whether an ISL should be formed between two cluster members. With the default setting of ISL discovery to **auto** and the ISL formation mode to **enable**, an ISL automatically forms between two cluster switches.

Performing a **fabric isl enable** command on an operational ISL has no effect. However, performing a **no fabric isl enable** command on an interface toggles its link status and subsequently disables ISL formation. In addition, the **no fabric isl enable** command triggers the switch to inform its neighbor that the local interface is ISL disabled. Upon receiving such information, a neighbor switch stops its ISL formation activity regardless of its current interface state.

#### NOTE

After you repair any segmented or disabled ISL ports, toggle the fabric ISL in order to propagate the changes.



---

**NOTE**

A **shutdown** command on an operating ISL interface not only brings down the physical link but also its FSPF adjacency. The main difference between a **shutdown** command and a **no fabric isl enable** command is that the link stays up after a **no fabric isl enable**, while the link stays down after a shutdown.

---

---

**NOTE**

Upon fabric reconvergence due to topology change involving ECMP fabric-isl path, there may be sub-second flooding of known unicast traffic.

---

## Disabling a Fabric ISL

The **no fabric isl enable** command takes this interface out of the trunk group if this interface happens to be currently part of the trunk. If you know and would like to fix the edge and fabric port assignments on a switch, then this command allows you to completely turn off ISL formation logic and shorten any link bring-up delays on edge ports.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **no fabric isl enable** command.

## Enabling a Fabric trunk

---

**NOTE**

Trunks are not supported between the Brocade 8000 and the Brocade VDX 8770.

---

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

## Disabling a Fabric trunk

Fabric trunking is enabled by default. Enter the **no fabric trunk enable** command to revert the ISL back to a standalone adjacency between two Brocade VCS Fabric switch.

## Broadcast, Unknown Unicast, and Multicast Forwarding

All switches in a Brocade VCS Fabric cluster share a single multicast tree rooted at the RBridge with the lowest RBridge ID (domain ID). All broadcast, unknown unicast, and multicast traffic between two edge RBridges is forwarded on this multicast tree inside the Brocade VCS Fabric. The multicast tree includes all RBridges in the Brocade VCS Fabric fabric.

### *Multicast distribution tree-root selection*

Network OS v4.0.0 software supports the following distribution tree behaviors.

- The root of the distribution tree is the switch with the lowest RBridge ID. The automated selection process does not require any user intervention.

- Each switch in the cluster optionally carries a multicast root priority. This priority setting overrides the automatically-selected multicast root. In deployments where a multicast root is required to be a specific switch that does not have the lowest RBridge ID, then the priority setting on that switch can override the root selection. If there are two switches with the same priority, then the switch with the lower RBridge ID prevails.
- A back-up multicast root is pre-selected, which is the switch with the next lowest RBridge ID. The back-up multicast root is automatically selected by all switches should the current multicast root fail.

## Priorities

As stated above, the root of the tree is auto-selected as the switch with the lowest RBridge ID. For example, if you had a cluster with RBridge IDs 5, 6, 7, and 8, then 5 would be the root. If you then added `rbridge-id 1` to this fabric, the tree would be re-calculated with 1 as the root.

In order to avoid this behavior, you can set a priority (default is 1). The highest priority overrides the lowest RBridge ID and becomes the root.

For example, to build a fabric with RBridge ID 7 or 8 as the root, set their priority to something higher than 1 (priority values are 1 through 255). If there is a tie in priority, the lower RBridge is still chosen. If RBridge ID 7 and 8 are both set to priority 1, 7 becomes the root.

### *Changing the priority*

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **`fabric route mcast rbridge-id`** command.

**Example of changing an RBridge multicast priority:**

```
switch(config)# fabric route mcast rbridge-id 12 priority 10
```

## Displaying the running configuration

The **`show running-config fabric route mcast`** command allows you to display fabric route multicast configuration information. The configuration currently effective on the switch is referred to as the running configuration. Any configuration change you make while the switch is online is made to the running configuration. The running configuration is nonpersistent.

---

### NOTE

To save configuration changes, you must save the running-config file to a file, or you can apply the changes by copying the running configuration to the startup configuration.

---

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **`show running-config fabric route mcast`** command.

```
switch# show running-config fabric route mcast priority
fabric route mcast rbridge-id 12 priority 10
```

## VCS Virtual IP address configuration

A virtual IP address is assigned for each VCS cluster. This virtual IP address is tied to the principal switch in the cluster. The management interface of the principal switch can be accessed by means of this virtual IP address. Because the virtual IP address is a property of the fabric cluster and logical chassis cluster, in the event that the principal switch goes down, the next principal switch is assigned this address.

Virtual IP address can be configured by means of the **vcs virtual ip address** command:

```
switch(config)# vcs virtual ip address 10.0.0.23
```

This command can be used in logical chassis cluster and fabric cluster modes only. When the virtual IP address is configured for the first time, the current principal switch in the cluster is assigned this IP address.

Virtual IP configuration is global in nature. All the nodes in the cluster are configured with the same virtual IP address, but address is bound to the current principal switch only. Make sure that the assigned virtual IP address is not a duplicate of an address assigned to any other management port in the cluster or network.

Brocade recommends that you use the same subnet as the IP address of management interface. To see the currently configured virtual IP address, use the **show vcs** command:

```
switch# show vcs virtual-ip
Virtual IP : 10.21.87.2/20
Associated rbridge-id : 2
```

To remove the currently configured virtual IP address, use the **no vcs virtual ip address** command.

```
switch(config)# no vcs virtual ip address
switch# show running-config vcs virtual ip address
% No entries found.
```

---

### NOTE

You should not use the **no vcs virtual ip address** command when logged onto the switch through the virtual IP address. Use the management port IP address of the principal switch, or the serial console connection of the principal switch.

---

If you wish to rebind this virtual IP address to this management interface, remove the currently configured virtual IP address and reconfigure it. This situation can arise when the virtual IP address is not bound to management interface of the principal switch as a result of duplicate address detection.

A separate gateway cannot be configured for virtual IP address. The default gateway is the same as the gateway address for the management port of the same switch.

*Virtual IP address configuration scenarios*

Virtual IP address may be assigned to a switch whenever it is the principal switch in the cluster. The configuration scenarios that may occur are described in [Table 20](#).

**TABLE 20** Virtual IP address configuration scenarios

| Scenario                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| First time cluster formation         | When the cluster is first being formed, and if the virtual IP address is already configured, the principal switch is assigned the Virtual IP address. If no Virtual IP configuration exists, then the principal switch can be accessed using the management port IP address.                                                                                                                                                                                                        |
| Virtual IP configuration             | When you configure the virtual IP address for a cluster the first time, the address is bound to the management interface of the principal switch.                                                                                                                                                                                                                                                                                                                                   |
| Principal switch failover            | If the principal switch becomes a secondary switch while the virtual IP address is assigned to its management interface, then the virtual IP address is reassigned to the new principal switch.                                                                                                                                                                                                                                                                                     |
| Principal switch goes down           | When the principal switch in the cluster goes down, the virtual IP address is released from its management interface. The virtual IP address will be assigned to the next switch that becomes the principal switch.                                                                                                                                                                                                                                                                 |
| Principal switch chassis is disabled | When the <b>chassis disable</b> command is executed on the principal switch, the virtual IP address is released from its management interface. The virtual IP address will be assigned to the next switch that becomes the principal switch.                                                                                                                                                                                                                                        |
| Virtual IP removal                   | If you remove the virtual IP address from the configuration, then the address is unbound from management interface of the principal switch. In this case, the principal switch can still be accessed by using the management port's IP address.                                                                                                                                                                                                                                     |
| Trivial merge                        | In the event that two clusters merge together, the global configuration of the smaller cluster (Cluster A) is overwritten by the larger cluster (Cluster B). During this time, the virtual IP address is unbound from the principal switch of Cluster A. The virtual IP address of Cluster B can now be used to access the principal of new merged cluster. If the virtual IP address of Cluster B is not configured, there will not be a virtual IP address in the merged cluster. |
| Cluster reboot                       | When the cluster reboots, the virtual IP address is persistent and is bound to the new principal switch.                                                                                                                                                                                                                                                                                                                                                                            |
| Cluster Islanding                    | If the ISL link goes down between two or more clusters that are forming, the principal switch in the original cluster retains the virtual IP address. The new principal switch in the second cluster will perform a check to confirm that the virtual IP address is not in use. If it is in use, then the address is not assigned to the switch and an error is logged in RASLog.                                                                                                   |
| Standalone node behavior             | A virtual IP address cannot be configured on a Standalone node in VCS mode.                                                                                                                                                                                                                                                                                                                                                                                                         |
| Virtual MAC address                  | Virtual MAC address are not supported by virtual IP addresses.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Management port primary IPv4 address | For a virtual IP address to work correctly, the management port's IPv4 address should be assigned and functional.                                                                                                                                                                                                                                                                                                                                                                   |

## Fabric ECMP load balancing

Traffic towards ECMP paths are load-balanced using following eight fields as the Key; VlanID, MAC DA/SA, L3\_ULP, L3 DA/SA, L4 Dst/Src. For some pattern of streams, most of the traffic falls into one ECMP path, and rest of the ECMP paths are underutilized. This results in loss of data traffic, even though more ECMP paths are available to offload the traffic. You can configure the ECMP path selection method within the fabric using the **fabric ecmp load-balance** command. The operands for this command are listed in [Table 21](#).

**TABLE 21** ECMP load balancing operands

| Operand                 | Description                                                                       |
|-------------------------|-----------------------------------------------------------------------------------|
| dst-mac-vid             | Destination MAC address and VID-based load balancing                              |
| src-dst-ip              | Source and Destination IP address-based load balancing                            |
| src-dst-ip-mac-vid      | Source and Destination IP and MAC address and VID-based load balancing            |
| src-dst-ip-mac-vid-port | Source and Destination IP, MAC address, VID and TCP/UDP port based load balancing |
| src-dst-ip-port         | Source and Destination IP and TCP/UDP port-based load balancing                   |
| src-dst-mac-vid         | Source and Destination MAC address and VID-based load balancing                   |
| src-mac-vid             | Source MAC address and VID-based load balancing                                   |

Additionally, you can choose to swap adjacent bits of the hash key using the **fabric ecmp load-balance-hash-swap** command. This is useful in cases where a choice of any of the hash key combinations causes the distribution of traffic to not be uniform.

The **fabric ecmp load-balance-hash-swap** command is used to configure the swapping of the input fields before feeding them to the hash function. The integer is interpreted as a bitwise control of the 212-bit key. Each bit controls whether the two adjacent bits of the key are to be swapped. This 32-bit control value is written to all four hash swap control registers. This value is replicated in 32-bit block to form a 106-bit value. A value of 0x0 does not swap any input fields while a value of 0xffffffff swaps all 106 input bit-pairs.

To configure the ECMP load balancing feature, perform the following steps in global configuration mode.

1. Enter RBridge ID configuration mode.

```
switch(config)# rbridge-id 2
switch(config-rbridge-id-2)#
```

2. Execute the **fabric ecmp load-balance** command for the stream you want to favor.

This example uses the Destination MAC address and VID-based load balancing flavor.

```
switch(config-rbridge-id-2)# fabric ecmp load-balance dst-mac-vid
```

3. Optional: Use the **fabric ecmp load-balance-hash-swap** command to swap the input fields before feeding them to the hash function.

```
switch(config-rbridge-id-2)# fabric ecmp load-balance-hash-swap 0x4
```

4. Use the **show fabric ecmp load-balance** command to display the current configuration of hash field selection and hash swap.

```
switch# show fabric ecmp load-balance
Fabric Ecmp Load Balance Information
```

## 9 Fabric interface configuration management

```

Rbridge-Id : 2
Ecmp-Load-Balance Flavor : Destination MAC address and VID based load
balancing
Ecmp-Load-Balance HashSwap : 0x4
```

# Metro VCS

## In this chapter

- [Metro VCS overview](#) ..... 135
- [Configuring Metro VCS using long distance ISL](#) ..... 136
- [Configuring Metro VCS using standard ISL](#) ..... 139
- [Metro VCS and distributed Ethernet VLAGs](#) ..... 140

## Metro VCS overview

Metro VCS allows you to interconnect different locations and form clusters of Data Centers (DC) over long distance in order to provide Disaster Protection/Recovery and load sharing.

There are multiple ways to stretch an Ethernet fabric over distance. By using ISLs (Inter Switch Links) over long distances (more than the standard 1000m), Ethernet fabrics can be distributed across Data Centers located in geographically different locations.

If more complex setups are needed within the different locations, fub-fabrics (also known as local VCS) can be used, as shown in [Figure 18](#).




FIGURE 18 Metro VCS configuration example

If Metro VCS is configured using standard ISL, with distances of up to 1000m, no limitations occur on supported service types or supported number of MAC addresses. Also, no special buffer arrangements (long distance ISL configuration) need to be done. Refer to [“Configuring Metro VCS using long distance ISL”](#) on page 136.

If Metro VCS is configured using long distance ISL, special configurations have to be used and supported fabric functionality is limited. Refer to [“Configuring Metro VCS using standard ISL”](#) on page 139.

## Configuring Metro VCS using long distance ISL

Extending Ethernet Fabrics over distance is accomplished using long distance ISL. The buffer allocation within a single port group is optimized, which extends the supported ISL distance.

Metro VCS supports long-distance ISL ports up to 30 km on the Brocade VDX platforms listed in [Table 22](#). Links up to 10 km are lossless. You can have eight 1-km links forming a Brocade trunk. You can also have mixed length cables forming the ISL. For ECMP purposes, you can have eight 8-link ECMP trunks.

**TABLE 22** Limitations for long distance Metro VCS

| Supported hardware                       | Extended ISL up to 2 km | Extended ISL up to 5 km | Extended ISL up to 10 km | Extended ISL up to 30 km |
|------------------------------------------|-------------------------|-------------------------|--------------------------|--------------------------|
| Brocade VDX 6720-60                      | yes                     | yes                     | yes                      | yes                      |
| Brocade VDX 6730-76                      | yes                     | yes                     | yes                      | yes                      |
| Brocade VDX 6740                         | yes                     | yes                     | yes                      | yes                      |
| Brocade VDX 8770 - VDX LC48x10G linecard | yes                     | yes                     | yes                      | yes                      |

[Table 23](#) displays the limitations on extended ISL for Network OS hardware.

**TABLE 23** Conditions for long distance Metro VCS

| Condition                                                                     | Extended ISL up to 2 km | Extended ISL up to 5 km | Extended ISL up to 10 km | Extended ISL up to 30 km |
|-------------------------------------------------------------------------------|-------------------------|-------------------------|--------------------------|--------------------------|
| Support for lossless FCoE/iSCSI traffic on the Metro VCS port-group           | yes                     | yes                     | yes                      | no                       |
| Layer 2/IP Lossy Traffic support                                              | yes                     | yes                     | yes                      | yes                      |
| Number of Metro VCS long distance ports supported per port group              | 1                       | 1                       | 1                        | 1                        |
| Number of regular ISLs supported on a port group configured for long distance | 1                       | 1                       | 0                        | 0                        |
| Trunking support between multiple LD ISLs                                     | no                      | no                      | no                       | no                       |
| CEE map or FCoE port allowed in same port-group                               | no                      | no                      | no                       | no                       |
| eNS Sync (MAC address table sync)                                             | yes                     | yes                     | yes                      | yes                      |
| Zoning                                                                        | yes                     | yes                     | yes                      | yes                      |
| HA failover                                                                   | yes                     | yes                     | yes                      | yes                      |



**TABLE 23** Conditions for long distance Metro VCS (Continued)

| Condition              | Extended ISL<br>up to 2 km             | Extended ISL<br>up to 5 km             | Extended ISL<br>up to 10 km            | Extended ISL<br>up to 30 km            |
|------------------------|----------------------------------------|----------------------------------------|----------------------------------------|----------------------------------------|
| Node redundancy check  | yes                                    | yes                                    | yes                                    | yes                                    |
| vMotion                | yes                                    | yes                                    | yes                                    | yes                                    |
| Maximum PFCs Supported | 3<br>(2 on the<br>Brocade VDX<br>6740) | 3<br>(2 on the<br>Brocade VDX<br>6740) | 3<br>(2 on the<br>Brocade VDX<br>6740) | 3<br>(2 on the<br>Brocade VDX<br>6740) |

Table 24 displays the port groups and number of port groups available on each platform for long distance Metro VCS.

**TABLE 24** Long distance Metro VCS port-group schema

| Platform                                 | Port groups                             | Number of port groups on platform |
|------------------------------------------|-----------------------------------------|-----------------------------------|
| Brocade VDX 6720-60 (10Gbe)              | 1-10, 11-20, 21-30, 31-40, 41-50, 51-60 | 6                                 |
| Brocade VDX 6730-76 (10Gbe)              | 1-10, 11-20, 21-30, 31-40, 41-50, 51-60 | 6                                 |
| Brocade VDX 6740                         | 1-32, 33-48                             | 2*                                |
| Brocade VDX 8770 (VDX LC48x10G linecard) | 1-8, 9-16, 17-24, 25-32, 33-40, 41-48   | 6 per 10G blade                   |

\*Not a valid deployment scenario at distances longer than 5 km, as no normal ISLs are allowed if both port-groups are configured with long-distance ISLs for 10 km and 30 km.

## Metro VCS using long distance guidelines and restrictions

Consider the following when configuring long distance Metro VCS:

- For fabrics with more than two nodes, the VCS license must be installed on all switches in the Fabric Clusters.
- The Metro VCS fabric is only supported on 10G interface links.
- Brocade trunking is not supported with long distance ISLs, but up to 8-link ECMP trunks can be used.
- Edge ports can not be configured with DCB maps.
- Edge ports and port groups with long distance links can not use the **fcoeport default** command.
- A maximum of three PFCs can be supported on a Metro VCS configured platform. By default, PFC is enabled by class 3 and 7.
- The Brocade VDX 6740 only supports two PFCs.
- All the ports in the portgroup are rebooted when configuring a port for long distance.
- For 2-, 5-, 10 km long distance, use Brocade-supported long-range (LR) optics for direct connectivity.
- For 30 km long distance, use Brocade-supported extended-range (ER) optics for direct connectivity.
- A port-group containing a long-distance port can not have a CEE Map configuration on any edge port.

## 10 Configuring Metro VCS using long distance ISL

To configure a Metro VCS port, perform the following steps in privileged EXEC mode. Each long distance ISL port of a VCS must be connected to a long distance ISL port on the remote VCS.

1. Verify the default standard ISL configuration is correct using the **show running-config** command.

```
switch# show running-config interface tengigabitethernet 51/0/1
interface TenGigabitEthernet 51/0/1
fabric isl enable
fabric trunk enable
no shutdown
```

2. Set the port to support Metro VCS up to 30 km using the **long-distance-isl** command.

```
switch# interface tengigabit 51/0/1
switch(conf-if-te-51/0/1)# long-distance-isl 30000
```

3. Perform the same long-distance ISL configuration on the interface of the peer RBridge on the remote sites of the Metro VCS.

4. Verify the long-distance ISL is correctly formed using the **show fabric isl** and **show fabric islports** command.

```
switch(conf-if-te-51/0/1)# do show fabric isl
Rbridge-id: 51 #ISLs: 1
```

| Src<br>Index | Src<br>Interface | Nbr<br>Index | Nbr<br>Interface | Nbr-WWN                 | BW  |
|--------------|------------------|--------------|------------------|-------------------------|-----|
| 4            | Te 51/0/1        | 4            | Te 53/0/1        | 10:00:00:05:33:65:3B:50 | 10G |
| Yes          | "VCS3-53"        |              |                  |                         |     |

```
switch(conf-if-te-51/0/1)# do show fabric islports
Name: VCS3-51
Type: 131.4
State: Online
Role: Fabric Principal
VCS Id: 3
Config Mode:Local-Only
Rbridge-id: 51
WWN: 10:00:00:05:33:e5:d0:4b
FCF MAC: 00:05:33:e5:d0:cf
```

| Index | Interface  | State | Operational State                                     |
|-------|------------|-------|-------------------------------------------------------|
| 0     | Fo 51/0/49 | Down  |                                                       |
| 1     | Fo 51/0/50 | Down  |                                                       |
| 2     | Fo 51/0/51 | Down  |                                                       |
| 3     | Fo 51/0/52 | Down  |                                                       |
| 4     | Te 51/0/1  | Up    | ISL 10:00:00:05:33:65:3B:50 "VCS3-53" (Trunk Primary) |

<Truncated>

## Configuring Metro VCS using standard ISL

In order to deploy Metro VCS using standard ISL, no configuration is required on the standard ISL. The default configuration on the 10G interface using the **fabric isl enable** and **fabric trunk enable** commands allows ISL formation with other Brocade VDX switches in the same VCS Cluster automatically. BLDP negotiation takes place to form standard ISLs for distances up to 30Km. Refer to [Chapter 9, “Fabric”](#).

Metro VCS using standard ISL is supported on:

- Brocade VDX 6720-60
- Brocade VDX 6730-76
- Brocade VDX 6740
- Brocade VDX 8770 with the VDX LC48x10G line-card

[Figure 19](#) is a typical deployment topology supported for interconnecting Data Centers by extending Brocade VCS Ethernet Fabrics using standard ISL. The local VCS clusters are connected to the Metro VCS Clusters by Brocade vLAGs. In this case, local data-center Ethernet Fabrics from both site are not merged while providing seamless Layer 2 extension. For Metro VCS, Brocade standard ISL Trunking is supported up to a maximum of 8 ISLs to form 80G trunks.




FIGURE 19 Typical deployment topology for Metro VCS using standard ISL

## Standard distance Metro VCS guidelines and restrictions

Consider the following when configuring Metro VCS with standard ISLs:

- Only two data-center and three data-center topologies are supported.
- A maximum of 2 nodes are supported for each site, which provides node redundancy. If more complex local designs are required, a local VCS sub-fabric design must be used.
- Only standard Ethernet services are supported. DCB and FCoE are not supported.
- Brocade trunking with up to (80G) 8x10G links is supported.
- There are no additional limitations on supported overall number of MAC addresses.

Table 25 displays the port groups and number of port groups available on each platform for Metro VCS using standard ISL.

**TABLE 25** Standard Metro VCS port-group schema


| Platform                                 | Port groups                             | Number of port groups on platform |
|------------------------------------------|-----------------------------------------|-----------------------------------|
| Brocade VDX 6720-60 (10Gbe)              | 1-10, 11-20, 21-30, 31-40, 41-50, 51-60 | 6                                 |
| Brocade VDX 6730-76 (10Gbe)              | 1-10, 11-20, 21-30, 31-40, 41-50, 51-60 | 6                                 |
| Brocade VDX 6740                         | 1-16, 17-32, 33-40, 41-48               | 4                                 |
| Brocade VDX 8770 (VDX LC48x10G linecard) | 1-8, 9-16, 17-24, 25-32, 33-40, 41-48   | 6 per 10G blade                   |

## Metro VCS and distributed Ethernet VLAGs

Outside of Metro distances, and whenever bit-transparency may be a problem, Edge to Edge interconnected Fabrics using 10G vLAG over multiple standard Ethernet links can be used.

This allows you to connect separate Ethernet Fabrics which can be located in different data centers, even if the distance between those locations is more than 30km. This is currently supported for distances up to 100km.

As indicated in Figure 20, one side can be a Metro VCS over long distance fabric.



**FIGURE 20** Metro VCS and distributed Ethernet fabrics

In order to connect two distinct VCS Ethernet Fabrics between data-centers, a third Metro VCS Fabric can be formed and the distinct local VCS Ethernet Fabrics can connect to the Metro VCS Fabric by means of Virtual Link Aggregation (vLAG).

Alternatively, the distinct VCS Ethernet Fabrics in respective data-centers can be directly connected to each other using vLAG over xWDM up to a distance of 30 Km. Wherever bit-transparencies is not achievable in xWDM equipment, this solution can be successfully deployed for edge-to-edge interconnectivity (using 10G vLAG over multiple standard Ethernet links). This deployment is referred to as “Distributed Ethernet Fabrics using vLAG”.

This implementation eliminates the need for creation of a separate Metro VCS Fabric to achieve local VCS Cluster isolation while providing Layer 2 connectivity between them. In such a deployment, DCB/FCoE lossless Ethernet traffic is not supported.

---

**NOTE**

When a port-channel from a node in one VCS (or standalone or any switch) spans across multiple R Bridges in other VCS Cluster, a vLAG is formed on the R Bridges in the VCS Cluster that are part of the same port-channel. For Distributed Ethernet Fabrics using vLAG over long distances, only LACP based standard port-channels are supported. For details on how to create port-channels and vLAGs, refer to [Chapter 26, “Configuring Link Aggregation”](#).

---

The following VDX platforms are supported for Distributed Ethernet Fabrics using vLAG:

- Brocade VDX 6720-60
- Brocade VDX 6730-76
- Brocade VDX 6740
- Brocade VDX 8770 with the VDX LC48x10G line-card

## Connecting local VCS Clusters over long distance using vLAG

Following is a typical deployment topology for interconnecting Data Centers by connecting local Brocade VCS Ethernet Fabric through vLAG. A node from the local VCS Cluster is connected via xWDM to the nodes in other VCS to form a vLAG. The other VCS could be spanned across two data-centers using standard or long-distance ISL as depicted in below topology. In this case, the vLAG between two data-centers provide VCS Fabric isolation while providing seamless Layer 2 connectivity (see the first port-channel (Po1) forming the vLAG between VCS1 and VCS2 over long distance through xWDM.)

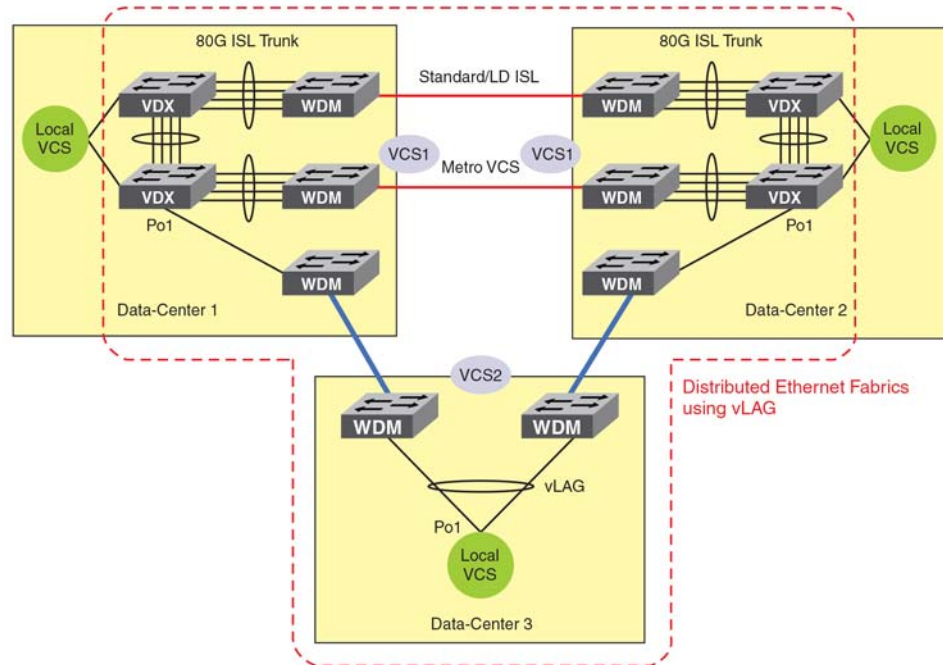


FIGURE 21 Connecting local VCS Clusters over long-distance using vLAG

### Guidelines and restrictions

Note the following guidelines and restrictions when creating a distributed Ethernet VLAG.

- Only dynamic vLAG is supported.
- DCB/FCoE lossless Ethernet traffic is not supported.
- The maximum supported distance is 100km.

### Configuring vLAGs for Distributed Ethernet Fabrics

In order to deploy Distributed Ethernet Fabrics using vLAG, create a port-channel interface on the R Bridges that are to be connected. Then add the member-interfaces to the port-channel and bring them online. Configure switchport and add the vlans that are to be allowed over the port-channel. After the port-channels on all the R Bridges are online, the vLAG forms automatically on the R Bridge that connects to multiple nodes on other VCS. In the deployment topology shown in Figure 21, the vLAG forms on the R Bridges that are part of port-channel (Po) 1 in Data-centers 1 and 2 that forms VCS 1.

This task needs to be configured on R Bridges that connect the two VCS instances. Perform the following task in global configuration mode.

1. Create a port-channel interface on all R Bridges that are directly connected to R Bridges in other VCS instances.

---

#### NOTE

In Logical Chassis mode, the port-channel is created only from the principal node and is created globally.

---

```
switch(config)# interface port-channel 300
```

2. Verify the port-channel is created correctly with the **show running-config** command.

```
switch(config-Port-channel-300)# do show running-config interface Port-channel 300
interface Port-channel 300
 vlag ignore-split
 shutdown
```

3. Configure the port-channel interface for switchport trunk and add vlans to be allowed on the trunk interface with the **switchport** command.

```
switch(config-Port-channel-300)# switchport
switch(config-Port-channel-300)# switchport mode trunk
switch(config-Port-channel-300)# switchport trunk allowed vlan all
```

4. Verify the port-channel interface configuration with the **show running-config** command.

```
switch(config-Port-channel-300)# do show running-config interface Port-channel 300
interface Port-channel 300
 vlag ignore-split
 switchport
 switchport mode trunk
 switchport trunk allowed vlan all
 switchport trunk tag native-vlan
 spanning-tree shutdown
 shutdown
```

5. Add member interfaces to the port-channel interface with the **channel-group** command. Repeat for all interfaces that need to be part of the port-channel.

```
switch(conf-if-te-53/0/31)# channel-group 300 mode active type standard
switch(conf-if-te-53/0/31)# do show running-config interface
TenGigabitEthernet 53/0/31
interface TenGigabitEthernet 53/0/31
 fabric isl enable
 fabric trunk enable
 channel-group 300 mode active type standard
 lacp timeout long
 no shutdown
```

6. Bring the port-channel online in both VCS instances by executing **no shutdown** on the port-channel interface.

```
switch(config-Port-channel-300)# no shutdown
```

7. Verify the port-channel interface configuration with the **show running-config** command.

```
switch(config-Port-channel-300)# do show running-config interface Port-channel 300
interface Port-channel 300
```

```
vlag ignore-split
switchport
switchport mode trunk
switchport trunk allowed vlan all
switchport trunk tag native-vlan
spanning-tree shutdown
no shutdown
```

8. Verify console RASLOGs indicating formation of vLAG with the **no shutdown** command.

```
switch(config-Port-channel-300)# no shutdown
2013/06/17-16:40:53, [NSM-1023], 224126, DCE, INFO, VCS1-51, RBridge ID 51
has joined Port-channel 300. Port-channel is a vLAG with RBridge IDs 52 51.
```

9. Verify port-channel interfaces indicating formation of vLAG with the **show port-channel** command.

```
switch# show port-channel 300
LACP Aggregator: Po 300 (vLAG)
Aggregator type: Standard
Ignore-split is enabled
 Member rbridges:
 rbridge-id: 51 (2)
 rbridge-id: 52 (2)
 Admin Key: 0010 - Oper Key 0010
 Partner System ID - 0x8000,01-e0-52-00-00-02
 Partner Oper Key 0010
Member ports on rbridge-id 51:
 Link: Te 51/0/31 (0x19180E801C) sync: 1
 Link: Te 51/0/32 (0x19180F001D) sync: 1
```



# Administering Zones

---

## In this chapter

|                                       |     |
|---------------------------------------|-----|
| • Zoning overview .....               | 145 |
| • Operational considerations .....    | 150 |
| • Default zoning access modes .....   | 153 |
| • Zone database size .....            | 154 |
| • Zone aliases .....                  | 154 |
| • Zone creation and management .....  | 157 |
| • Zone configuration management ..... | 160 |
| • Zone configuration scenario .....   | 168 |
| • Zone merging .....                  | 169 |
| • LSAN Zones .....                    | 175 |

## Zoning overview

Zoning is a fabric-based service that enables you to partition your network into logical groups of devices that can access each other and prevent access from outside the group. Grouping devices into zones in this manner not only provides security, but also relieves the network from Registered State Change Notification (RSCN) storms that occur when too many native FCoE devices attempt to communicate with one another.

You can use zoning to partition your network in many ways. For example, you can partition your network 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 22](#), 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.

## 11 Zoning overview

- The Storage 3 is not assigned to a zone; no other zoned fabric device can access it.





FIGURE 22 Zoning

Connecting to another network through a Fibre Channel (FC) router, you can create an LSAN zone to include zone objects on other fabrics, including Fabric OS networks. No merging takes place across the FC router when you create an LSAN zone. [Figure 23](#) shows an example in which server1, which is connected to switch in a Brocade VCS Fabric cluster, has access to local storage and to RAID storage on a Fabric OS fabric.



**FIGURE 23** LSAN zoning

**NOTE:** Zoning in Network OS 4.0.0 and later has the following restrictions:

- Zone objects based on physical port number or port ID (D,I ports) are not supported.
- You cannot access a target on a Network OS fabric from a server on the Fabric OS fabric.

## Approaches to zoning

[Table 26](#) lists the various approaches you can take when implementing zoning in a Network OS fabric.

**TABLE 26** Approaches to fabric-based zoning

| Zoning approach             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Recommended approach</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 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 zoning 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. <i>This zoning philosophy is the preferred method.</i> |

**TABLE 26 Approaches to fabric-based zoning (Continued)**

| Zoning approach               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Alternative approaches</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 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 RSCNs, or errors, go out to a larger group than necessary.                                                                                                                                                                                                                                                                      |
| 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. It 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. |
| <b>Not recommended</b>        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| No zoning                     | Using no zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric and causes RSCN storms. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should be used only in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.                                                                                                                                                                                                                                                                                                                                  |

## Zone objects

A zone object can be one of the following types: a zone, a zone member, an alias for one or more zone members, or a zoning configuration.

### *Zones*

A zone is made up of one or more zone members. Each zone member can be a device, a port, or an alias. If the zone member is a device, it must be identified by its Node World Wide Name (node WWN). If it is a port, it must be identified by its Port World Wide Name (port WWN). Port WWNs and node WWNs can be mixed in the same zone. For LSA zones, only port WWNs can be used.

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, 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.

Up to 255 zone member objects are supported for each zone. For LSA zones, this number is further restricted by the FC router which can parse up to 128 entries for each LSA zone.

## *Zone aliases*

A zone alias is a name assigned to a device or a group of devices. By creating an alias, you can assign a familiar name to one or more devices and refer to these devices by that name. Aliases simplify cumbersome data entry by allowing you to create an intuitive naming structure (such as using “NT\_Hosts” to define all NT hosts in the fabric).

As a shortcut for zone members, zone aliases simplify the entry and tracking of zone objects that are defined by their WWNs. 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 kind of naming strategy clearly identifies the server host bus adapter (HBA associated with the zone).

## *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 enabled, all zones that are members of that configuration are enabled.

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.

- Enabled Configuration

A single zone configuration that is currently in effect. The enabled configuration is built when you enable a specified zone configuration.

If you disable the enabled configuration, zoning is disabled on the fabric, and default zoning takes effect. When default zoning takes effect, either all devices within the fabric can communicate with all other devices, or no device communicate with any other device, depending on how default zoning is configured. Disabling the configuration does not mean that the zone database is deleted, however, only that no configuration is active in the fabric.

On power-up, the switch automatically reloads the saved configuration. If a configuration was active when it was saved, the same configuration is reinstated on the local switch.

## *Naming conventions*

Naming zones and zone configurations is flexible. You can devise prefixes to differentiate between zones used for production, backup, recovery, or testing. 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 you want to use other configurations for specific purposes, you can use names such as “BACKUP\_A,” “RECOVERY\_2,” and “TEST\_18jun02”.

## Zoning enforcement

Zone enforcement is by name server. The name server filters queries and RSCNs based on the enabled zoning configuration.

## Considerations for zoning architecture

Table 27 lists considerations for zoning architecture.

**TABLE 27** Considerations for zoning architecture

| 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. |
| Allowing time to propagate changes       | 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.                                                                                                                                                                                                                               |
| 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.                                                                                                                                                                                                                                                                        |
| 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 of zoning.                                                                                                                                                                                                                                                                                                                  |

## Operational considerations

### Supported modes

Zoning is supported only in VCS mode. When you save, enable, or disable a configuration, the changes are automatically distributed to all switches in the VCS Fabric.

Zoning is not supported in standalone mode and all zoning commands are hidden in this mode. The zone database is cleared and disabled when you transition from VCS Fabric mode to standalone mode.

### Supported firmware

Zoning is supported only if all R Bridges in the fabric are running Network OS 2.1 or later.

Connecting an R Bridge running Network OS 2.0 to an R Bridge running Network OS 2.1 or later merges the two networks only if the R Bridge running Network OS 2.1 or later is in Brocade VCS Fabric mode and no zone database elements are defined or enabled.

A switch running Network OS v3.0.0 will segment if it is attached to a switch running Network OS v2.0.0 regardless of zoning configuration. A switch running Network OS v3.0.0 will join the fabric with a 2.1.x switch and zones will be merged, but the cluster will not form, so no further zoning commands will be allowed until all switches are upgraded to the same firmware version and the cluster has formed.

The interswitch links (ISLs) connecting the two R Bridges will segment if the R Bridge running Network OS 2.1 or later has any zone defined or enabled, or the default zone is set to No Access. Any such configuration requires automatic distribution of zoning configuration data, which is not compatible with R Bridges running Network OS 2.0.

## Firmware downgrade and upgrade considerations

A firmware downgrade from Network OS v3.x to v2.1.x is not permitted under the following conditions:

1. One or more zone aliases are configured on the switch. You must remove all references to zone aliases prior to a firmware downgrade. Use the **no zoning defined-configuration alias** command to delete all zone alias objects. Then issue the **zoning enabled-configuration cfg-action {cfg-save | cfg-disable}** command or the **zoning enabled-configuration cfg-name *cfg\_name*** command to commit the operation before re-attempting a firmware download.
2. An open zone transaction in progress. You must either commit or abort the current open transaction before re-attempting a firmware download. Use the **zoning enabled-configuration cfg-action {cfg-save | cfg-disable}** command or the **zoning enabled-configuration cfg-name *cfg\_name*** command to commit the current open transaction. Alternately, use the **zoning enabled-configuration cfg-action *cfg-transaction-abort*** command to abort the open transaction.

You cannot downgrade any switch in a Brocade VCS Fabric to Network OS v2.0 if any zone definition exists in the defined configuration. Any attempt to do so will fail while attempting to download the v2.0 firmware. For the downgrade to succeed, you must clear the defined configuration, disable any active configuration, set the default zoning mode to All Access, and then try again to download the firmware.

When you upgrade from Network OS v2.1.0 to v2.1.1 or v3.0.0 and later, the zone database is cleared.



### CAUTION

Clearing the defined configuration clears the zoning database for the entire fabric. If you want to downgrade just one switch without affecting the rest of the fabric, disconnect the switch from the fabric before deleting the defined configuration.

## Zone configuration management

You can perform zoning operations on any R Bridge in the VCS Fabric, but they are always executed on the principal R Bridge. In Logical Chassis mode, any edits made to the zoning database are allowed only from the principal R Bridge, and you can issue **show** commands from non-principal switches in this mode. In Fabric Cluster mode, you can make edits from any R Bridge.

Automatic distribution of the zoning configuration ensures that the effects of these operations are shared and instantly visible on all switches in the VCS Fabric. However, these operations are not permanent until a transaction commit operation saves them to nonvolatile memory, which holds the master copy of the zoning database. Any user can commit the transaction on any switch, and the commit operation saves the operations performed by all users. Once the zoning configuration is saved in permanent memory, it persists across reboot operations.

A transaction commit occurs when you or another user initiates any of the following zoning operations:

- Saving the database to nonvolatile memory with the **zoning enabled-configuration cfg-action cfg-save** command.
- Enable a specific zone configuration with the **zoning enabled-configuration cfg-name** command.
- Disabling the currently enabled zone configuration with the **no zoning enabled-configuration cfg-name** command.
- Aborting the current transaction with the **zoning enabled-configuration cfg-action cfg-transaction-abort** command. This operation rolls back all zoning operations performed by any user since the last committed transaction.

If the principal RBridge reboots or goes down, Network OS selects a new principal and any pending zoning transaction is rolled back to the last committed transaction, which is the effective zoning configuration saved in nonvolatile memory. Any changes made to the effective configuration prior to an abort operation must be re-entered.

If an RBridge other than the principal reboots or goes down, the ongoing transaction is not backed out. Any zoning operations initiated by the RBridge are still part of the global transaction maintained on the principal RBridge.

If a fabric segments, the newly elected principal RBridge determines whether transaction data is retained. If a segment retains the original principal, it also retains ongoing transaction data. If a segment elects a new principal, the transaction is aborted.

The zone startup configuration is always equal to the running configuration. The running configuration will always be overwritten by the information from the master copy of the zoning database in nonvolatile memory at startup, so you always start up with the previous running configuration. It is not necessary to copy the running configuration to the startup configuration explicitly.

You can save a snapshot of the current running configuration using the **copy running-config file** command. You can add configuration entries from a saved configuration using the **copy file running-config** command. When saving the snapshot you must ensure that the saved running configuration contains no zoning transaction data, otherwise failures will occur when you attempt to restore configuration entries from the saved file. Any transaction data would cause such a failure, including empty configuration definitions or empty zones.

---

### NOTE

When you re-enable the enabled-configuration (using the **zoning enabled-configuration** command) on the principal switch in the cluster, the system propagates the enabled-configuration across the cluster. There is a slight risk of doing this in that the defined-configuration may contain configuration edits that you may not want to enable yet. This feature prevents switches in the cluster having mismatched enabled-configurations.

---

When restoring the running configuration, Brocade recommends copying the file to the running configuration in the absence of any other command line input.



**NOTE**

When you restore a configuration using the **copy** command, the contents of the file are added to the defined configuration; they do not replace the defined configuration. The result is cumulative, as if the input came from the command line.

## Default zoning access modes

The default zoning mode controls device access if zoning is not implemented or if there is no enabled zone configuration. Default zoning has two access modes:

- 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 setting is All Access. Changing the default access mode requires committing the ongoing transaction for the change to take effect.

The default zoning mode takes effect when you disable the effective zone configuration. If your default zone has a large number of devices, to prevent RSCN storms from overloading those devices, you should set the default zoning mode to No Access before attempting to disable the zone configuration. If your default zone includes more than 300 devices, the zoning software prevents you from disabling the zoning configuration if the default zoning mode is All Access.

### Setting the default zoning mode

1. In the privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter one of the following commands, depending on the default access mode you want to configure:
  - To set the default access mode to All Access, enter **zoning enabled-configuration default-zone-access allaccess**.
  - To set the default access mode to No Access, enter **zoning enabled-configuration default-zone-access noaccess**.
3. Enter the **zoning enabled-configuration cfg-action cfg-save** or **zoning enabled-configuration cfg-name** command to commit the ongoing transaction and save the access mode change to nonvolatile memory.
4. Enter the **show running-config zoning enabled-configuration** command to verify the access mode change.

#### Example of setting the default zoning mode to no access

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration default-zone-access noaccess
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)# do show running-config zoning enabled-configuration
zoning enabled-configuration cfg-name cfg1
zoning enabled-configuration default-zone-access noaccess
zoning enabled-configuration cfg-action cfg-save
```

## 11 Zone database size

### Example of viewing enabled zones

```
switch# show zoning enabled-configuration
zoning enabled-configuration cfg-name cfg1

zoning enabled-configuration enabled-zone zone1
member-entry 10:00:00:00:00:00:00:01

zoning enabled-configuration enabled-zone zone2
member-entry 10:00:00:00:00:00:00:02
```

## 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 memory available for storing the master copy of the defined configuration in flash memory.

Use the following information displayed by the **show zoning operation-info** command to determine whether there is enough space to complete outstanding transactions:

- db-max—Theoretical maximum size of the zoning database kept in nonvolatile memory
- db-avail—Theoretical amount of free space available
- db-committed—The size of the defined configuration currently stored in nonvolatile memory
- db-transaction—The amount of memory required to commit the current transaction

The supported maximum zone database size is 100 KB. If the outstanding transaction data (db-transaction field) is less than the remaining supported space (100 KB minus db-committed), enough space exists to commit the transaction.

Note that the db-max field shows a theoretical zone database limit of about 1 MB. However, performance might become unacceptable if the zoning database exceeds 150 KB.

## Viewing database size information

In the privileged EXEC mode, enter the **show zoning operation-info** command.

Database and transaction size information is displayed in bytes.

```
switch# show zoning operation-info
db-max 1045274
db-avail 1043895
db-committed 367
db-transaction 373
transaction-token 1
last-zone-changed-timestamp 2011-11-16 16:54:31 GMT-7:00
last-zone-committed-timestamp 2011-11-16 16:23:44 GMT-7:00
```

## Zone aliases

A zone alias is user-defined name for a logical group of ports or WWNs. You can simplify the process of creating and managing zones by first specifying aliases for zone members. Aliases facilitate tracking and eliminate the need for long lists of individual zone member names. An alias can be a member of a zone, but it cannot be a member of a zoning configuration.

## Creating an alias

1. In the privileged EXEC mode, enter the **show name-server detail** command to list the WWNs of devices and targets available in the Brocade VCS Fabric.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **zoning defined-configuration alias** command followed by a name for the alias.  
A sub-configuration mode prompt appears.
4. Enter the subconfiguration mode **member-entry** command to specify at least one member entry.  
  
The member entry must be specified as a port WWN or a node WWN.  
  
You can add multiple members in one operation by separating each member entry with a semicolon (;). No spaces are allowed after the semicolon.
5. Enter the **exit** command to return to global configuration mode.
6. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the configuration to nonvolatile memory.

### Example of creating an alias with one member node WWN

```
switch# show name-server detail
PID: 013100
Port Name: 20:00:00:00:00:00:01
Node Name: 10:00:00:00:00:00:01
(output truncated)

switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration alias alias1
switch(config-alias-alias1)# member-entry 10:00:00:00:00:00:01
switch(config-alias-alias1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#
```

## Adding additional members to an existing alias

1. In the privileged EXEC mode, enter the **show name-server detail** command to list the WWNs of devices and targets available in the Brocade VCS Fabric.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **zoning defined-configuration alias** command followed the name of an existing zone alias.  
A subconfiguration mode prompt appears.
4. Enter the subconfiguration mode **member-entry** command to specify at least one member entry.  
  
The member entry must be specified as a port WWN or a node WWN.  
  
You can add multiple members in one operation by separating each member entry with a semicolon (;). No spaces are allowed after the semicolon.
5. Enter the **exit** command to return to global configuration mode.

6. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the configuration to nonvolatile memory.

## Example of adding two member node WWNs to an existing alias

```
switch# show name-server detail
PID: 013200
Port Name: 20:00:00:00:00:00:02
Node Name: 10:00:00:00:00:00:02
(output truncated)
PID: 013300
Port Name: 20:00:00:00:00:00:03
Node Name: 10:00:00:00:00:00:03
(output truncated)

switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration alias alias1
switch(config-alias-alias1)# member-entry
10:00:00:00:00:00:02;10:00:00:00:00:00:03
switch(config-alias-alias1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#
```

## Removing a member from an alias

1. In the privileged EXEC mode, enter the **show running-config zoning** command to display the alias and its member WWNs.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **zoning defined-configuration alias** command followed the name of an existing zone alias.  
A subconfiguration mode prompt appears.
4. Enter the subconfiguration mode **no member-entry** command to specify the WWN to be removed from the zone alias.  
You can only remove one member at a time.
5. Enter the **exit** command to return to the global configuration mode.
6. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the configuration to nonvolatile memory.

## Example of removing two members from an alias

```
switch# show running-config zoning
zoning defined-configuration alias alias1
 member-entry 10:00:00:00:00:00:01
 member-entry 10:00:00:00:00:00:02
 member-entry 10:00:00:00:00:00:03
(output truncated)
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration alias alias1
switch(config-alias-alias1)# no member-entry 10:00:00:00:00:00:02
switch(config-alias-alias1)# no member-entry 10:00:00:00:00:00:03
switch(config-alias-alias1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

## Deleting an alias

1. In the privileged EXEC mode, enter the **show running-config zoning** command to display the alias and its member WWNs.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **no zoning defined-configuration alias** command followed by the name of the alias you want to delete.
4. Enter the **show running-config zoning** command to verify the change in the defined configuration (optional).
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the configuration to nonvolatile memory.

### Example of deleting an alias

```
switch# show running-config zoning
zoning defined-configuration alias alias1
 member-entry 10:00:00:00:00:00:01
!
zoning enabled-configuration cfg-name ""
zoning enabled-configuration default-zone-access allaccess
zoning enabled-configuration cfg-action cfg-none
switch# configure terminal
Entering configuration mode terminal
switch(config)# no zoning defined-configuration alias alias1
switch(config)# do show running-config zoning
zoning enabled-configuration cfg-name ""
zoning enabled-configuration default-zone-access allaccess
zoning enabled-configuration cfg-action cfg-none
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#
```

## Zone creation and management

### Creating a zone

A zone cannot persist without any zone members. When you create a new zone, the **zoning defined-configuration zone** command places you in a command subconfiguration mode where you can add the first zone member entry. You can specify multiple members by separating each member from the next by a semicolon (;).

---

#### NOTE

Zones without any zone members cannot exist in volatile memory. They are deleted when the transaction commits successfully.

---

The following procedure adds a new zone to the defined configuration.

1. In the privileged EXEC mode, enter the **show name-server detail** command to obtain the WWNs of servers and targets available in the Brocade VCS Fabric.
2. Enter the **configure terminal** command to enter global configuration mode.

## 11 Zone creation and management

3. Enter the **zoning defined-configuration zone** command and enter a new zone name to add a new zone.  
A subconfiguration mode prompt appears.
4. Enter the subconfiguration mode **member-entry** command to specify at least one member entry.  
The member entry must be specified as a port WWN, a node WWN, or an alias. You can mix WWNs and aliases.  
Add multiple members in one operation by separating each member entry with a semicolon (;). No spaces are allowed after the semicolon.
5. Enter the **exit** command to return to global configuration mode.
6. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

### Example of creating a zone with two members, a WWN and an alias

```
switch# show name-server detail
PID: 012100
Port Name: 10:00:00:05:1E:ED:95:38
Node Name: 20:00:00:05:1E:ED:95:38
(output truncated)

switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration zone zone1
switch(config-zone-zone1)# member-entry 20:00:00:05:1E:ED:95:38;alias2
switch(config-zone-zone1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#
```

## Adding a member to a zone

1. In the privileged EXEC mode, enter the **show name-server detail** command to list the WWNs of devices and targets available on the Brocade VCS Fabric cluster.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **zoning defined-configuration zone** command and enter the name of an existing zone.  
A subconfiguration mode prompt appears.
4. Enter the subconfiguration mode **member-entry** command and specify the member you want to add.  
The new member can be specified by a port WWN, a node WWN, or a zone alias.  
Add multiple members in one operation by separating each member with a semicolon (;).
5. Enter the **exit** command to return to the global configuration mode.
6. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

### Example of adding three members to a zone, two node WWNs and an alias

```
switch# show name-server detail
PID: 012100
Port Name: 50:05:07:61:00:1b:62:ed
```

```

Node Name: 50:05:07:61:00:1b:62:ed
(output truncated)
PID: 012200
Port Name: 50:05:07:61:00:09:20:b4
Node Name: 50:05:07:61:00:09:20:b4
(output truncated)

switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration zone zone1
switch(config-zone-zone1)# member-entry
50:05:07:61:00:1b:62:ed;50:05:07:61:00:09:20:b4;alias3
switch(config-zone-zone1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#

```

## Removing a member from a zone

1. In the privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning defined-configuration zone** command and enter the name of the zone from which you want to remove a member.  
A subconfiguration mode prompt appears.
3. Enter the subconfiguration mode **no member-entry** parameter and specify the WWN or the alias of the member you want to remove.  
You can remove only one member at a time. To remove more than one member, you must issue the **no member-entry** command for each member you want to remove.
4. Enter the **exit** command to return to global configuration mode.
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

---

### NOTE

Saving the configuration to nonvolatile memory also deletes the zone if the member you are removing is the last member in the zone.

---

### Example of removing more than one member from a zone

```

switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration zone zone1
switch(config-zone-zone1)# no member-entry 50:05:07:61:00:09:20:b4
switch(config-zone-zone1)# no member-entry alias3
switch(config-zone-zone1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#

```

## Deleting a zone

Before deleting a zone, ensure that the zone is not a member of any enabled zone configuration. Although the deletion will proceed in RAM, you will not be able to save the configuration to nonvolatile memory if an enabled zone configuration has the deleted zone as a member.

## 11 Zone configuration management

1. In the privileged EXEC mode, enter the **show running-config zoning defined-configuration** command and verify that the zone you want to delete is not a member of an enabled zone configuration. If the zone is a member of an enabled zone configuration, remove it.
2. Enter the **configure terminal** command to enter the global configuration mode.
3. Enter the **no zoning defined-configuration zone** command and enter the name of the zone you want to delete.
4. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

---

### NOTE

Saving the configuration to nonvolatile memory also deletes the zone configuration if the zone you are removing is the last member zone in the configuration.

---

#### Example of removing a zone from the defined configuration

```
switch# show running-config zoning defined-configuration
zoning defined-configuration zone zone1
member-entry 10:00:00:00:00:00:01
!
zoning defined-configuration zone zone2
member-entry 10:00:00:00:00:00:02
!
switch# configure terminal
Entering configuration mode terminal
switch(config)# no zoning defined-configuration zone zone2
switch(config)# zoning enabled-configuration cfg-action cfg-save
Updating flash ...
switch(config)# exit
switch# show running-config zoning defined-configuration
zoning defined-configuration zone zone1
member-entry 10:00:00:00:00:00:01
```

## Zone configuration management

### Viewing the defined configuration

To view the defined configuration, in the privileged EXEC mode, enter the **show running-config zoning defined-configuration** command.

For each configuration, the command lists each member zone. For each zone, the command lists the WWN of each member.

```
switch# show running-config zoning defined-configuration
zoning defined-configuration cfg cfg0
member-zone zone_0_1
member-zone zone_0_2
member-zone zone_0_3
member-zone zone_0_4
member-zone zone_same
!
zoning defined-configuration cfg cfg1
member-zone zone_1_1
member-zone zone_1_2
```



```

member-zone zone_1_3
member-zone zone_1_4
member-zone zone_same
!
zoning defined-configuration cfg cfg2
member-zone zone_2_1
member-zone zone_2_2
member-zone zone_2_3
member-zone zone_2_4
member-zone zone_same
!
zoning defined-configuration cfg cfg4
member-zone zone2
member-zone zone3
!
zoning defined-configuration zone zone0
member-entry 11:22:33:44:55:66:77:80
member-entry 11:22:33:44:55:66:77:81
member-entry 11:22:33:44:55:66:77:82
member-entry 11:22:33:44:55:66:77:83
member-entry 11:22:33:44:55:66:77:84
!
zoning defined-configuration zone zone1
member-entry 11:22:33:44:55:66:77:80
member-entry 11:22:33:44:55:66:77:81
member-entry 11:22:33:44:55:66:77:82
member-entry 11:22:33:44:55:66:77:83
member-entry 11:22:33:44:55:66:77:84
!
zoning defined-configuration zone zone2
member-entry 11:22:33:44:55:66:77:80
member-entry 11:22:33:44:55:66:77:81
member-entry 11:22:33:44:55:66:77:82
member-entry 11:22:33:44:55:66:77:83
member-entry 11:22:33:44:55:66:77:84
!
(output truncated)

```

## Viewing the enabled configuration

To view the enabled configuration, in the privileged EXEC mode, enter the **show zoning enabled-configuration** command. The following information about the enabled configuration is displayed:

- The name of the configuration
- The configuration action
- The mode of the default zone—the mode that will be active if you disable the enabled configuration

---

### NOTE

With NOS 4.0.0, the enabled-zone output is no longer available from the **show zoning enabled-configuration enabled-zone** command. It is now available from the **show zoning enabled-configuration** command.

---

The configuration name has CFG\_MARKER asterisk (\*) appended to it if an outstanding transaction exists; the asterisk is not present if no outstanding transaction exists. Similarly, the configuration action is flagged as “cfg-save” if no outstanding transaction exists; “cfg-none” indicates that an outstanding transaction exists. A CFG\_MARKER flag is appended to the configuration if the enabled configuration does not exactly match the defined configuration. This scenario occurs when you have an enabled configuration and make changes to the defined-configuration, and then, instead of enabling the defined configuration, you issue the **cfg-save** command.



### CAUTION

When edits are made to the defined configuration, and those edits affect a currently enabled zone configuration, issuing a “cfg-save” command makes the enabled configuration effectively stale. Until the enabled configuration is reenabled, the merging of new RBridges into the cluster is not recommended. This merging may cause unpredictable results, with the potential for mismatched enabled-zoning configurations among the RBridges in the cluster.

#### Example of viewing the zoning enabled configuration

```
switch# show zoning enabled-configuration
zoning enabled-configuration cfg-name cfg1

zoning enabled-configuration enabled-zone zone1
 member-entry 10:00:00:00:00:00:00:01

zoning enabled-configuration enabled-zone zone2
 member-entry 10:00:00:00:00:00:00:02
```

Issue the **show zoning** command without any parameters to view both the defined and the enabled configuration parameters.

## Creating a zone configuration

A zone configuration cannot persist without any member zones. When creating a new zone configuration, the **zoning defined-configuration cfg** command places you in a command sub-configuration mode where you must add at least one member zone. While zone configurations without any member zones can exist in volatile memory, they are deleted when the transaction commits successfully.

The following procedure adds a new zone configuration to the defined configuration.

1. In the privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning defined-configuration cfg** command and enter a new configuration name.  
A subconfiguration mode prompt appears.
3. Enter the **member-zone** subconfiguration mode command and specify the name of at least one zone.  
Add multiple zones in one operation by separating each zone name with a semicolon (;).
4. Enter the **exit** command to return to global configuration mode.
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

**Example of creating a zone configuration with one member zone**

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration cfg config1
switch(config-cfg-config1)# member-zone zone1
switch(config-cfg-config1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#
```

**NOTE**

Zone aliases are not valid zone configuration members. Adding an alias to an existing zone configuration will not be blocked. However, the attempt to enable a zone configuration that contains aliases will fail with an appropriate error message.

## Adding a zone to a zone configuration

1. In the privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning defined-configuration cfg** command and enter the name of the configuration to which you want to add zones.

The command prompt changes to indicate a subconfiguration mode.

3. Enter the **member-zone** subconfiguration mode command and specify the name of at least one member zone.

Add multiple zones in one operation by separating each zone name with a semicolon (;).

4. Enter the **exit** command to return to global configuration mode.
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

**Example of adding two zones to config1**

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration cfg config1
switch(config-cfg-config1)# member-zone zone2;zone3
switch(config-cfg-config1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#
```

## Removing a zone from a zone configuration

1. In the privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning defined-configuration cfg** command and enter the name of the configuration from which you want to remove a zone.

The command prompt changes to indicate a subconfiguration mode.

3. Enter the **no member-zone** subconfiguration mode command and specify the name of the zone you want to remove from the configuration.

You can remove only one member at a time. To remove more than one member, you must issue the **no member-zone** command for each member you want to remove.

4. Enter the **exit** command to return to global configuration mode.
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

---

**NOTE**

Saving the configuration to nonvolatile memory deletes the configuration if the zone you are removing is the last member in the configuration.

---

**Example of removing two zones from config1**

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration cfg config1
switch(config-cfg-config1)# no member-zone zone2
switch(config-cfg-config1)# no member-zone zone3
switch(config-cfg-config1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#
```

## Enabling a zone configuration

Only one zone configuration can be enabled in a VCS Fabric. The following procedure selects a configuration from the defined configuration and makes it the enabled configuration. If a zone configuration is currently enabled, the newly enabled configuration replaces the previously enabled configuration.

1. In the privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning enabled-configuration cfg-name** command with the name of the configuration you want to enable.

In addition to enabling the specified configuration, this command also saves any changes made to the zoning database in volatile memory to nonvolatile memory. The saved configuration is persistent.

If the configuration refers to a nonexistent zone or a zone with no members assigned to it, the operation fails and the command returns an error message. The following example enables config1.

**Example of enabling a zone configuration**

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration cfg-name config1
switch(config)#
```

**Example of a failed enable operation**

The enable operation fails because the configuration contains a zone without members.

```
switch(config)# do show running-config zoning
zoning defined-configuration cfg cfg1
member-zone-zone1
member-zone zone2
!
```

```

zoning defined-configuration zone zone1 <-----Zone with no member
!
zoning defined-configuration zone zone2
member-entry 20:03:00:11:0d:bc:76:09
!
zoning enabled-configuration cfg-name ""
zoning enabled-configuration default-zone-access allaccess
zoning enabled-configuration cfg-action cfg-none

switch(config)# zoning enabled-configuration cfg-name cfg1
% Error: Command Failed. Cfg contains empty zone object "zone1"

```

## Disabling a zone configuration

Disabling the currently enabled configuration returns the fabric to no-zoning mode. All devices can then access one another or not at all, depending on the default zone access mode setting.

### NOTE

For fabrics with many devices, Brocade recommends setting the default zone access mode to No Access before disabling a zone configuration to avoid RSCN storms.

1. In the privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **no zoning enabled-configuration cfg-name** command.

In addition to disabling the currently enabled configuration, this command also saves any changes made to the zoning database in volatile memory to nonvolatile memory. The saved configuration is persistent.

### Example of disabling a zoning configuration

```

switch# configure terminal
Entering configuration mode terminal
switch(config)# no zoning enabled-configuration cfg-name
switch(config)#

```

## Deleting a zone configuration

The following procedure deletes a zone configuration from the defined configuration.

1. In the privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **no zoning defined-configuration cfg** command and the name of the zone configuration you want to delete.
3. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified defined configuration to nonvolatile memory.

### Example of deleting a zone configuration

```

switch# configure terminal
Entering configuration mode terminal
switch(config)# no zoning defined-configuration cfg cfg2
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#

```

---

**NOTE**

If you try to delete the enabled configuration from the defined configuration, the **zoning enabled-configuration cfg-action cfg-save** command returns an error. However, if you commit the transaction with the **zoning enabled-configuration cfg-action cfg-disable** command, the operation proceeds without error.

---

### Clearing changes to a zone configuration

The following procedure aborts all pending transactions and removes all uncommitted operations from the database. It returns the configuration in volatile memory to the state it was in when a **zoning enabled-configuration cfg-action cfg-save** or **zoning enabled-configuration cfg-name** command was last executed successfully.

1. In the privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter the **zoning enabled-configuration cfg-action cfg-transaction-abort** command.

**Example of aborting a transaction**

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration cfg-action cfg-transaction-abort
switch(config)#
```

### Clearing all zone configurations

The following procedure clears all zone configurations from the defined configuration and enables the default zone.

---

**NOTE**

For fabrics with many devices, Brocade recommends setting the default access mode to No Access before clearing all zone configurations to avoid RSCN storms.

---

1. In the privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning enabled-configuration cfg-action cfg-clear** command.
3. Enter one of the following commands, depending on whether an enabled zone configuration exists:
  - If no enabled zone configuration exists, enter the **zoning enabled-configuration cfg-action cfg-save** command.
  - If an enabled zone configuration exists, enter the **no zoning enabled-configuration cfg-name** command to disable and clear the zone configuration in nonvolatile memory for all switches in the fabric.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration cfg-action cfg-clear
switch(config)# no zoning enabled-configuration cfg-name
switch(config)#
```

## Backing up the zone configuration

To backup your zoning configuration you copy it to a file and store it on a server or on an attached USB device. You can use the copy to restore the configuration if needed.

---

### NOTE

Ensure that no transaction is pending before you perform the copy operation, otherwise failures will occur when you attempt to restore configuration entries from the saved file. Any transaction data would cause such a failure, including empty configuration definitions or empty zones.

---

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Empty the transaction buffer by either committing the transaction to nonvolatile memory or aborting the transaction.
  - To commit the transaction, enter the **zoning enabled-configuration cfg-action cfg-save** command, the **zoning enabled configuration cfg-name** command, or the **zoning enabled-configuration cfg-action cfg-disable** command.
  - To abort the transaction, enter the **zoning enabled-configuration cfg-action cfg-transaction-abort** command.
3. Enter the **exit** command to return to privileged EXEC mode.
4. Enter the **copy** command. For the source file, use **running-config**. For the destination file, use the file name you want the configuration copied to.

### Example of making a backup copy on a USB device

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)# exit
switch# copy running-config usb://myconfig
```

## Restoring a configuration from backup

When you restore a configuration from backup and add to the running configuration, the zone configuration identified in the backup copy as the enabled configuration becomes the new enabled configuration.

In the privileged EXEC mode, enter the **copy** command. For the source file use the file where the saved configuration is stored. For the destination file, use **running-config**.

This operation updates the defined configuration in RAM.

---

### NOTE


The **copy** command adds to the defined configuration. It does not replace the defined configuration.

---

The following example adds the configuration in the file named myconfig on the attached USB device to the defined configuration.

```
switch# copy usb://myconfig running-config
```

## Zone configuration scenario



**FIGURE 24** Zone configuration example

The following example creates the zone configuration shown in [Figure 24](#). The example assumes that two hosts need access to the same storage device, while each host needs private storage of its own. You create two zones: Zone A contains Host 1, its private storage device, and the shared storage device; Zone B contains Host 2, its private storage device, and the shared storage device. In addition, you create two zone configurations: `cfg1` in which only Zone A is effective; `cfg2`, in which both zones are effective.

1. Log in to any switch in the Brocade VCS Fabric.
2. Enter the **show name-server detail** command to list the available WWNs.
3. Enter the **configure terminal** command to enter global configuration mode.
4. Enter the **zoning defined-configuration zone** command to create Zone A.
5. Enter the **zoning defined-configuration zone** command to create Zone B.
6. Enter the **zoning defined-configuration cfg** command to create the configuration `cfg1` with Zone A as its only member.
7. Enter the **zoning defined-configuration cfg** command to create the configuration `cfg2` with Zone A and Zone B as its members.
8. Enter the **zoning running-config defined-configuration** command to view the defined zone configuration.
9. Enter the **zoning enabled-configuration cfg-name** command to enable `cfg2`.



10. Verify the enabled zoning configuration, by means of the **show zoning enabled-configuration** command.

```
switch# show name-server detail
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration zone ZoneA
switch(config-zone-ZoneA)# member-entry
20:00:00:05:1e:ed:95:38;50:05:07:61:00:09:20:b4;50:05:07:61:00:1b:62:ed
switch(config-zone-ZoneA)# exit
switch(config)# zoning defined-configuration zone ZoneB
switch(config-zone-ZoneB)# member-entry
20:00:00:05:1e:ed:18:c3;50:05:07:61:00:22:18:9b;50:05:07:61:00:1b:62:ed
switch(config-zone-ZoneB)# exit
switch(config)# zoning defined-configuration cfg cfg1
switch(config-cfg-cfg1)# member-zone ZoneA
switch(config-cfg-cfg1)# exit
switch(config)# zoning defined-configuration cfg cfg2
switch(config-cfg-cfg2)# member-zone ZoneA;ZoneB
switch(config-cfg-cfg2)# exit
switch(config)# zoning enabled-configuration cfg-name cfg2
switch(config)# exit
switch# show zoning enabled-configuration
zoning enabled-configuration cfg cfg1
 member-zone ZoneA
!
zoning enabled-configuration cfg cfg2
 member-zone ZoneA
 member-zone ZoneB
!
zoning enabled-configuration zone ZoneA
 member-entry 20:00:00:05:1e:ed:95:38
 member-entry 50:05:07:61:00:09:20:b4
 member-entry 50:05:07:61:00:1b:62:ed
!
zoning enabled-configuration zone ZoneB
 member-entry 20:00:00:05:1e:ed:18:c3
 member-entry 50:05:07:61:00:22:18:9b
 member-entry 50:05:07:61:00:1b:62:ed
!
```

## Zone merging

When a new switch is added to a VCS fabric, it automatically inherits the zone configuration information from the fabric. You can verify the zone configuration on any switch by using the procedure described in [“Viewing the defined configuration”](#) on page 160. Take care to avoid mismatched enabled-configuration scenarios.



### CAUTION

When edits are made to the defined configuration, and those edits affect a currently enabled zone configuration, issuing a “cfg-save” command makes the enabled configuration effectively stale. Until the enabled configuration is reenabled, the merging of new R Bridges into the cluster is not recommended. This merging may cause unpredictable results, with the potential for mismatched enabled-zoning configurations among the R Bridges in the cluster.

If you are adding a switch that is already configured for zoning, you must clear the zone configuration on that switch before connecting it to the zoned fabric. Refer to [“Clearing all zone configurations”](#) on page 166 for instructions.

Adding a new fabric that has no zone configuration information to an existing zoned 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 all switches in the added fabric.

---

### NOTE

To prevent an unwanted zone merge, use the **no fabric isl enable** command on ISL interfaces instead of the **shutdown** command on tengigabitethernet ports.

---

Before the new fabric can merge successfully, it must satisfy the following criteria:

- Before merging
  - Ensure that all switches adhere to the default zone merge rules as described in [“Zone merging scenarios”](#) on page 171.
  - Ensure that the enabled and defined zone configurations match. If they do not match and you merge with another switch, the merge might be successful, but unpredictable zoning and routing behavior can occur. See the Caution above and refer to [“Viewing the defined configuration”](#) on page 160.
- Merging and segmentation

The system checks each port as it comes online to determine whether the ports should be segmented. E\_Ports come online on power up, enabling a switch, or adding a new switch, and the system checks the zone database to see if the two database that can be merged safely. Refer to [“Zone merging scenarios”](#) on page 171.
- 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.
  - Enabled configurations: If there is an enabled configuration between two switches, the enabled zone configurations must match.
  - Zone membership: If a zoning object has the same name in both the local and adjacent defined configurations, 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 the other switches within the merge request.
- Merging two fabrics

For best practices, the default-zone access modes should match, although this is not a requirement. Refer to [“Zone merging scenarios”](#) on page 171.

If the two fabrics have conflicting zone configurations, they will not merge. If the two fabrics cannot join, the ISLs between the switches will segment.

The transaction state after the merge depends on which switch is elected as the principal RBridge. The newly elected principal RBridge retains the same transaction information it had before the merge. Transaction data is discarded from any switch that lost its principal status during the merge.

- Merge conflicts

When a merge conflict is present, a merge does not take place and the ISLs will segment.

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 ISLs will be segmented.

A merge is not possible under any of the following conditions:

- Configuration mismatch: Zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
- Zone Database Size: The zone database size exceeds the maximum limit of another switch.

---

#### NOTE

If the zone members on two switches are not listed in the same order, the configuration is considered a mismatch, and the switches will segment from the fabric. For example: `cfg1 = z1; z2` is different from `cfg1 = z2; z1`, even though the members of the configuration are the same. If zone members on two switches have the same names defined in the configuration, make sure the zone 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 previous zone configuration.

If the connections between two fabrics are replaced and no changes have been made to the zone configuration in either of the two fabrics, the two fabrics can merge back into one single fabric. If any changes that cause a conflict have been made to either zone configuration, a fabric merge may fail.

## Zone merging scenarios

The following tables provide information on merging zones and the expected results.

- [Table 28](#) on page 172: Defined and enabled configurations
- [Table 29](#) on page 173: Different content
- [Table 30](#) on page 174: Different names

# 11 Zone merging

- [Table 31](#) on page 174: Default access mode

**TABLE 28** Zone merging scenarios: Defined and enabled configurations

| Description                                                                                                                           | Switch A                                                                                         | Switch B                                                                                         | Expected results                                                                                                                                                                                                                          |
|---------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Switch A</b> has a defined configuration.<br><b>Switch B</b> does not have a defined configuration.                                | defined: cfg1;<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: none | defined: none<br>enabled: none                                                                   | Configuration from <b>Switch A</b> propagates throughout the fabric in an inactive state, because the configuration is not enabled.                                                                                                       |
| <b>Switch A</b> has a defined and enabled configuration.<br><b>Switch B</b> has a defined configuration but no enabled configuration. | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg1; | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: none  | Configuration from <b>Switch A</b> propagates throughout the fabric. The configuration is enabled after the merge in the fabric.                                                                                                          |
| <b>Switch A</b> and <b>Switch B</b> have the same defined configuration. Neither have an enabled configuration.                       | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: none  | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: none  | No change (clean merge).                                                                                                                                                                                                                  |
| <b>Switch A</b> and <b>Switch B</b> have the same defined and enabled configuration.                                                  | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg1; | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg1; | No change (clean merge).                                                                                                                                                                                                                  |
| <b>Switch A</b> does not have a defined configuration.<br><b>Switch B</b> has a defined configuration.                                | defined: none<br>enabled: none                                                                   | defined:cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: none   | <b>Switch A</b> absorbs the configuration from the fabric.                                                                                                                                                                                |
| <b>Switch A</b> does not have a defined configuration.<br><b>Switch B</b> has a defined and enabled configuration.                    | defined: none<br>enabled: none                                                                   | defined:cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg1   | <b>Switch A</b> absorbs the configuration from the fabric, with cfg1 as the enabled configuration.                                                                                                                                        |
| <b>Switch A</b> and <b>Switch B</b> have the same defined configuration. Only <b>Switch B</b> has an enabled configuration.           | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: none  | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg1  | Clean merge, with cfg1 as the enabled configuration.                                                                                                                                                                                      |
| <b>Switch A</b> and <b>Switch B</b> have different defined configurations. Neither have an enabled configuration.                     | defined: cfg2<br>zone2:<br>10:00:00:90:69:00:00:8c;<br>10:00:00:90:69:00:00:8d<br>enabled: none  | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: none  | Clean merge. The new configuration will be a composite of the two.<br>defined: cfg1<br>zone1: 10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>cfg2:<br>zone2: 10:00:00:90:69:00:00:8c;<br>10:00:00:90:69:00:00:8d<br>enabled: none |

TABLE 28 Zone merging scenarios: Defined and enabled configurations (Continued)

| Description                                                                                                                                                                                                 | Switch A                                                                                        | Switch B                                                                                                                                                                                                                                                    | Expected results                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Switch A</b> and <b>Switch B</b> have different defined configurations. <b>Switch B</b> has an enabled configuration.                                                                                    | defined: cfg2<br>zone2:<br>10:00:00:90:69:00:00:8c;<br>10:00:00:90:69:00:00:8d<br>enabled: none | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg1                                                                                                                                                             | Clean merge. The new configuration is a composite of both, with <b>cfg1</b> as the enabled configuration.                                                                                                                                                                                                                                               |
| <b>Switch A</b> does not have a defined configuration. <b>Switch B</b> has a defined configuration and an enabled configuration, but the enabled configuration is different from the defined configuration. | defined: none<br>enabled: none                                                                  | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:0<br>0:8a;<br>10:00:00:90:69:00:0<br>0:8b<br>effective: cfg1<br>zone1:<br>10:00:00:90:69:00:0<br>0:8a;<br>10:00:00:90:69:00:0<br>0:8b<br>zone2:<br>10:00:00:90:69:00:0<br>0:8c,<br>10:00:00:90:69:00:0<br>0:8d | Clean merge. Switch A absorbs the defined configuration from the fabric, with <b>cfg1</b> as the effective configuration. In this case, however, the effective configurations for Switch A and Switch B are different. You should issue a <b>zoning enabled-configuration cfg-name</b> command from the switch with the proper effective configuration. |

TABLE 29 Zone merging scenarios: Different content

| Description                     | Switch A                                                                                                                                                         | Switch B                                                                                                                                                         | Expected results                                       |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| Enabled configuration mismatch. | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b | defined: cfg2<br>zone2:<br>10:00:00:90:69:00:00:8c;<br>10:00:00:90:69:00:00:8d<br>enabled: cfg2<br>zone2:<br>10:00:00:90:69:00:00:8c;<br>10:00:00:90:69:00:00:8d | Fabric segments due to mismatching zone configurations |
| Configuration content mismatch. | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: irrelevant                                                            | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8c;<br>10:00:00:90:69:00:00:8d<br>enabled: irrelevant                                                            | Fabric segments due to mismatching zone content        |

# 11 Zone merging

**TABLE 30 Zone merging scenarios: Different names**

| Description                                         | Switch A                                                                                                                                                         | Switch B                                                                                                                                                        | Expected results                                       |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| Same content, different enabled configuration name. | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b | defined:cfg2<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: cfg2<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b | Fabric segments due to mismatching zone configurations |
| Same content, different zone name.                  | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: irrelevant                                                            | defined: cfg1<br>zone2:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b<br>enabled: irrelevant                                                           | Fabric segments due to mismatching zone content        |
| Same name, same content, different order.           | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8a;<br>10:00:00:90:69:00:00:8b;<br>10:00:00:90:69:00:00:8c<br>enabled: irrelevant                                | defined: cfg1<br>zone1:<br>10:00:00:90:69:00:00:8b;<br>10:00:00:90:69:00:00:8c;<br>10:00:00:90:69:00:00:8a<br>enabled: irrelevant                               | Fabric segments due to mismatching zone content        |
| Same name, different types.                         | effective: zone1:<br>MARKETING                                                                                                                                   | enabled: cfg1:<br>MARKETING                                                                                                                                     | Fabric segments due to mismatching types               |

**TABLE 31 Zone merging scenarios: Default access mode**

| Description                                  | Switch A                                               | Switch B                                               | Expected results                                                                                                                                                |
|----------------------------------------------|--------------------------------------------------------|--------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Different default zone access mode settings. | default zone: All Access                               | default zone: No Access                                | Clean merge — No Access takes precedence and default zone configuration from <b>Switch B</b> propagates to fabric.<br>default zone: No Access                   |
| Same default zone access mode settings.      | default zone: All Access                               | default zone: All Access                               | Clean merge — default zone configuration is All Access in the fabric.                                                                                           |
| Same default zone access mode settings.      | default zone: No Access                                | default zone: No Access                                | Clean merge — default zone configuration is No Access in the fabric.                                                                                            |
| Enabled zone configuration.                  | No enabled configuration.<br>default zone = All Access | enabled: cfg2<br>default zone: All Access or No Access | Clean merge — enabled zone configuration and default zone mode from <b>Switch B</b> propagates to fabric.                                                       |
| Enabled zone configuration.                  | No enabled configuration.<br>default zone = No Access  | enabled: cfg2<br>default zone: All Access              | Fabric segments because <b>Switch A</b> has a hidden zone configuration (No Access) activated and <b>Switch B</b> has an explicit zone configuration activated. |
| Enable zone configuration.                   | enabled: cfg1<br>default zone: No Access               | No enabled configuration.<br>default zone: No Access   | Clean merge — enabled zone configuration from <b>Switch A</b> propagates to fabric.                                                                             |
| Enable zone configuration.                   | enabled: cfg1<br>default zone: All Access              | No enabled configuration.<br>default zone: No Access   | Fabric segments. You can resolve the zone conflict by changing the default zone to No Access on <b>Switch A</b> .                                               |

## LSAN Zones

A Logical SAN (LSAN) consists of zones in two or more edge or backbone fabrics that contain the same devices. LSANs essentially 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. For details of this FC-FC routing service, refer to the *Fabric OS Administrator's Guide*.

---

**NOTE**

A backbone fabric consists of one or more FC switches with configured EX\_Ports. These EX\_Ports in the backbone connect to edge fabric switches through E\_Ports. This type of EX\_Port-to-E\_Port connectivity is called an "Inter-Fabric Link (IFL)".

---

The Brocade VCS Fabric connection to the FC router is an ISL that connects an FC port on a Brocade VDX 6730 to an EX\_Port on the FC router. Similarly, an FC port on the Fabric OS fabric connects to an EX\_Port on the FC router.

You can define and manage LSANs using the same zone management tools as for regular zones. The FC router makes LSAN zoning possible by importing devices in effective zones. For example, consider two devices:

- 11:22:33:44:55:66:77:99 is connected to a switch in a Brocade VCS Fabric cluster.
- 11:22:33:44:55:66:77:88 is connected to a switch in a Fabric OS fabric.

The FC-FC routing service on the FC router that connects the two fabrics presents 11:22:33:44:55:66:77:88 as a phantom device to the Brocade VCS Fabric and also presents 11:22:33:44:55:66:77:99 as a phantom device to the Fabric OS fabric. You can then use the regular zone management tools on the Brocade VCS Fabric cluster to incorporate 11:22:33:44:55:66:77:99 into an LSAN zone on the Brocade VCS Fabric. Similarly, you can use the regular zone management tools in Fabric OS to incorporate 11:22:33:44:55:66:77:88 into an LSAN zone in the Fabric OS fabric. Once both the Brocade VCS Fabric zone and the Fabric OS zone are enabled, the FC router imports devices common to both zones and makes them available to the zones in each fabric.

---

**NOTE**

Each phantom device counts against the maximum supported size of the Brocade VCS Fabric (24 devices).

---

### LSAN naming

Zones that contain hosts and targets that are shared between the two fabrics need to be explicitly coordinated. To share devices between any two fabrics, you must create an LSAN zone in both fabrics containing the WWNs of the devices to be shared. Although an LSAN zone is managed using the same tools as any other zone on the edge fabric, two behaviors distinguish an LSAN zone from a conventional zone:

- A required naming convention. The name of an LSAN zone begins with the prefix "LSAN\_". The LSAN name is case-insensitive; for example, *lsan\_* is equivalent to *LSAN\_*, *Lsan\_*, and so on.
- LSAN zone members in all fabrics must be identified by their WWN. You cannot use the port IDs that are supported only in Fabric OS fabrics.

**NOTE**

The "LSAN\_" prefix must appear at the beginning of the zone name.

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.

## Managing domain IDs

FCoE connectivity across the Fibre Channel link between Brocade VCS Fabric clusters and FC routers uses domain IDs to identify switches. Within a Brocade VCS Fabric cluster, a domain ID is the same as a routing bridge ID. When you connect to a Fibre Channel router, the FC fabric FC router service emulates virtual *phantom* FC domains in the FCoE fabric. Each FCR enabled switch emulates a single "front" phantom domain and each FC fabric is represented by a *translate* phantom domain.

It is important to ensure that front domain IDs and translate domain IDs presented by the FC router do not overlap routing bridge IDs in the FCoE fabric, otherwise the connectivity will fail and the Network OS switch with the overlapping routing bridge ID becomes isolated from the fabric. To prevent potential overlap, use the **portCfgExport -d** Fabric OS command on the FC router to apply a unique front domain ID—one that will not be used in the FCoE fabric. Similarly, use the **fcxlateConfig importedFID exportedFID preferredDomainID** Fabric OS command to set the translate domain ID to a unique value that is also not used as a routing bridge ID.

Refer to the *Fabric OS Command Reference Manual* for details about the **portCfgExport** and **fcxlateConfig** commands.

## Configuring LSAN zones—device sharing example

The following example shows LSANs sharing devices in separate fabrics. The procedure illustrates the creation of two LSAN zones (called *lsan\_zone\_fabric\_02* and *lsan\_zone\_fabric\_01*), which involve the following devices and connections:

- RBridge1 and the host in a Network OS fabric named fabric\_01.
- Switch2, Target A, and Target B in a Fabric OS fabric named fabric\_02.
- RBridge1 is connected by one of its FC\_Ports to an EX\_Port on the FC router.
- 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 RBridge1).
- 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).



Figure 25 shows the connectivity.

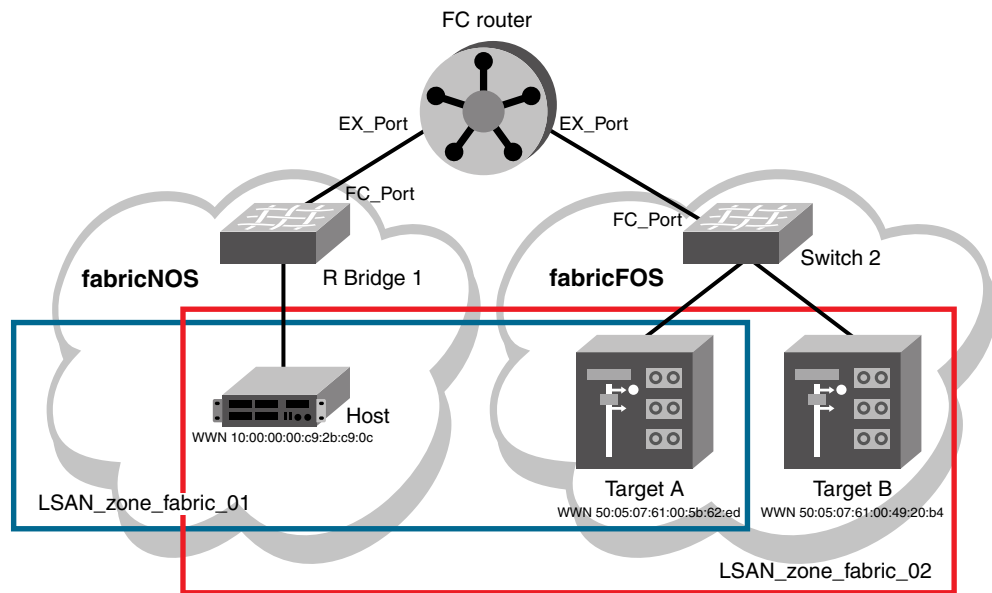


FIGURE 25 LSAN zones example

Obtain the host WWN in fabric\_01:

1. Log in to any switch in fabric\_01.
2. On the fabric\_01 switch, enter the **show name-server detail** command to list the WWN of the host (10:00:00:00:c9:2b:c9:0c).

#### NOTE

The **show name-server detail** output displays both the port WWN and node WWN; the port WWN must be used for LSANs.

```
switch# show name-server detail
PID: 012100
Port Name: 10:00:00:00:c9:2b:c9:0c
Node Name: 20:00:00:00:c9:2b:c9:0c
SCR: 3
FC4s: FCP
PortSymb: [27] "Brocade-1020|2.3.0.0|localhost.localdomain|Red Hat
Enterprise Linux Server release 5.5"
NodeSymb: NULL
Fabric Port Name: 20:21:00:05:1E:CD:79:7A
Permanent Port Name: 10:00:00:00:c9:2b:c9:0c
Device type: Physical Initiator
Interface: Fcoe 1/1/9
Physical Interface: Te 1/0/9
Share Area: No
Redirect: No
```

Obtain the target WWNS in fabric\_02:

3. Log in as admin on switch2 in fabric\_02.
4. On fabric\_02, 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
 The Local Name Server has 2 entries }
```

Create an LSAN zone in the NOS fabric (fabric\_01)

5. In fabric\_01, enter the **zoning defined-configuration zone** command to create the LSAN *lsan\_zone\_fabric\_01*, and include the host.

```
switch# config terminal
switch(config)# zoning defined-configuration zone lsan_zone_fabric_01
switch(config-zone-lsan_zone_fabric_01)# member-entry 10:00:00:00:c9:2b:c9:0c
```

6. In fabric\_01, add Target A to the LSAN.

```
switch(config-zone-lsan_zone_fabric_01)# member-entry 50:05:07:61:00:5b:62:ed
switch(config-zone-lsan_zone_fabric_01)# exit
```

7. In fabric\_01, enter the **zoning defined-configuration cfg** and **zoning enabled-configuration cfg-name** commands to add and enable the LSAN configuration.

```
switch(config)# zoning defined-configuration cfg zone_cfg
switch(config-cfg-zone_cfg)# member-zone lsan_zone_fabric_01
switch(config-cfg-zone_cfg)# exit
switch(config)# zoning enabled-configuration cfg_name zone_cfg
```

Create an LSAN zone in the FOS fabric (fabric\_02)

8. On switch2 (fabric\_02), enter the **zoneCreate** command to create the LSAN *lsan\_zone\_fabric2*, which includes the host (10:00:00:00:c9:2b:c9:0c), Target A (50:05:07:61:00:5b:62:ed), and Target B (50:05:07:61:00:49:20:b4).

```
switch:admin> zonecreate "lsan_zone_fabric_02",
"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. On switch2 (fabric\_02), enter the **cfgShow** command to verify that the zones are correct.

```
switch:admin> cfgshow
Defined configuration:
 zone: lsan_zone_fabric_02
 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. On switch2 (fabric\_02), enter the **cfgAdd** and **cfgEnable** commands to create and enable the LSAN configuration.

```
switch:admin> cfgadd "zone_cfg", "lsan_zone_fabric_02"
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 ...

Display the configuration on the FC router:

11. Log in as an admin and connect to the FC router.
12. On the FC router, 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_fabric_02
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_fabric_01
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
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
in Fabric PID Exists PID

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\_fabric\_02* and *lsan\_zone\_fabric\_01*. However, target B is defined by *lsan\_zone\_fabric\_02* and is not imported because *lsan\_zone\_fabric\_01* does not allow it.



# Configuring Fibre Channel Ports

---

## In this chapter

- Fibre Channel ports overview ..... 181
- Fibre Channel port activation and deactivation ..... 183
- Fibre Channel port attributes ..... 184
- Viewing Fibre Channel port attributes ..... 185
- Setting Fibre Channel port speed ..... 186
- Long-distance operation ..... 186
- Configuring a Fibre Channel port for trunking ..... 188
- Monitoring Fibre Channel ports ..... 188

## Fibre Channel ports overview


Fibre Channel ports provide the ability to connect a Brocade VCS Fabric cluster to a Fabric OS network. These connections provide support for zoning across Network OS and Fabric OS fabric types, which can enable FCoE devices on the Brocade VCS Fabric cluster to access SAN storage and services. See [Chapter 11, “Administering Zones”](#) for information on how to create LSAN zones.

Inter-Switch Links (ISLs) connect Fibre Channel ports on the Network OS switch to EX\_Ports on an FC router, which in turn connects to the Fabric OS network as shown in [Figure 26](#) on page 182.

These connections can be regular or long distance. The following Network OS Fibre Channel port types are supported:

- E\_port
- F\_port
  - Supports Fibre Channel HBA and Fibre Channel target
  - Supports F\_port to FCoE port bidirectional traffic
  - Supports F\_port to E\_port bidirectional traffic
- Auto (G\_port)—This is the default.
- N\_port (for Access Gateway only)

For details of Fibre Channel routing concepts, refer to the *Fabric OS Administrator's Guide*.



**FIGURE 26** FC connection between Network OS fabric and a Fibre Channel SAN

## Example

```
c60-VCS2-1# show fibrechannel login
```

| Interface | Index | PID    | status       | protocol | speed   | PortWWN                 |
|-----------|-------|--------|--------------|----------|---------|-------------------------|
| Fi 1/0/1  | 0     | 0102c3 | up(In_Sync ) | up       | 8G Auto | 2f:0a:00:05:1e:53:10:3f |
| Fi 1/0/1  | 0     | 0102c2 | up(In_Sync ) | up       | 8G Auto | 24:aa:00:05:1e:53:10:3f |
| Fi 1/0/1  | 0     | 0102c1 | up(In_Sync ) | up       | 8G Auto | 10:00:00:05:1e:53:10:3f |
| Fi 1/0/1  | 0     | 0102c0 | up(In_Sync ) | up       | 8G Auto | 20:00:00:05:1e:c7:84:e8 |
| Fi 1/0/4  | 3     | 010100 | up(In_Sync ) | up       | 8G Auto | 10:00:00:05:1e:af:4f:4e |
| Fi 1/0/16 | 15    | 0103c0 | up(In_Sync ) | up       | 8G Auto | 20:21:00:05:1e:c7:84:e8 |

## Example

```
switch(config)# interface FibreChannel 17/0/1
switch(conf-if-fi-17/0/1)#
switch(conf-if-fi-17/0/1)# config-mode
Possible completions:
 auto configure the port as a G-Port (locked)
 eport configure the port as a E-Port (locked)
 fport configure the port as a F-Port (locked)
 nport configure the port as a N-Port (locked)
switch(conf-if-fi-17/0/1)#
```

The mode configured is shown in the command "show running-config"

```
switch(conf-if-fi-17/0/1)# config-mode eport
switch(conf-if-fi-17/0/1)#
switch(config)#
switch# show running-config interface FibreChannel 17/0/1
interface FibreChannel 17/0/1
 desire-distance 0
 no isl-r_rdy
 trunk-enable
 config-mode eport
```

```
no shutdown
!
switch#
```

**NOTE**

Fibre Channel ports can connect F\_ports. You must enable **fcoeport default** for the interface for the Fibre Channel logins to be available.

The Brocade VDX 6730-32 and VDX 6730-76 switches are the only Network OS switches that support Fibre Channel ports. The Brocade VDX 6730-32 switch provides eight 8-Gbps Fibre Channel ports. The Brocade VDX 6730-76 provides sixteen 8-Gbps Fibre Channel ports.

Network OS software provides the following commands for managing Fibre Channel ports:

- **interface FibreChannel**—Global configuration mode command that allows you to enter the interface Fibre Channel configuration submode where you can enter commands to activate and deactivate a Fibre Channel port (**no shutdown** and **shutdown** commands) and to set port attributes (**desire-distance**, **fill-word**, **isl-r\_rdy**, **long-distance**, **speed**, **trunk-enable**, and **vc-link-init** commands).
- **show running-config interface FibreChannel**—A privileged EXEC mode command that displays Fibre Channel port configuration information.
- **show interface FibreChannel**—A privileged EXEC mode command that displays hardware counters that monitor activity and status of a Fibre Channel port.

## Fibre Channel port activation and deactivation

An FCoE license must be installed on a Brocade VDX 6730 switch to allow Fibre Channel port activation. Brocade VCS Fabric mode must be enabled. Once the FCoE license is installed, all Fibre Channel ports are activated by default. Refer to Chapter 7, “Administering Licenses,” for details about installing the FCoE license.

Use the **no shutdown** command to activate a Fibre Channel port. Use the **shutdown** command to deactivate a port.

### Enabling a Fibre Channel port

1. In privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter the **interface FibreChannel rbridge-id/slot/port** command for the Fibre Channel port you want to enable.

A configuration submode prompt appears.

3. Enter the **no shutdown** command.

The following example enables port 1 on routing bridge 8.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface FibreChannel 8/0/1
switch(config-FibreChannel-8/0/1)# no shutdown
```

## Disabling a Fibre Channel port

1. In privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter the **interface FibreChannel** *rbridge-id/slot/port* command for the Fibre Channel port you want to disable.

A configuration submode prompt appears.

3. Enter the **shutdown** command.

The following example disables port 1 on routing bridge 8.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface FibreChannel 8/0/1
switch(conf-FibreChannel-8/0/1)# shutdown
```

## Fibre Channel port attributes

Network OS v2.1.1 allows you to configure and display the following Fibre Channel port attributes for an E\_Port:

- Port speed—Enter the interface Fibre Channel configuration submode **speed** command to set the speed of a Fibre Channel port.
- Fill word—Enter the interface Fibre Channel configuration submode **fill-word** command to configure the link initialization and fill word primitives for an 8-Gbps Fibre Channel port.
- Long distance mode—Enter the interface Fibre Channel configuration submode **long-distance** command to configure the port for long-distance operations.
- VC link init—Enter the interface Fibre Channel configuration submode **vc-link-init** command to configure the fill word for long-distance operation.
- Desired distance—Enter the interface Fibre Channel configuration submode **desire-distance** command to configure manually the distance for a long-distance connection.
- Trunk port—Enter the interface Fibre Channel configuration submode **trunk-enable** command to configure the port for trunking.
- Buffer credit control—Enter the interface Fibre Channel configuration submode **isl-r\_r-rdy** command to enable interswitch link receiver-ready (ISL R\_RDY) mode on the port. Enter the interface Fibre Channel configuration submode **no isl-r\_r-rdy** command to disable ISL R\_RDY mode on the port. If ISL R\_RDY is not set, then interswitch link Virtual Channel ready (ISL VC\_RDY) mode is set by default. We recommend you do *not* set ISL R\_RDY.

The following Fibre Channel port attributes are not supported by Network OS version 2.1.1:

|                 |                    |                    |
|-----------------|--------------------|--------------------|
| AL_PA offset 13 | F_Port buffers     | NPIV capability    |
| Compression     | Fault Delay        | NPIV PP Limit      |
| Credit Recovery | FEC                | Persistent Disable |
| CSCTL mode      | Frame shooter port | Port Auto Disable  |
| D-Port mode     | Locked G_Port      | QoS E_Port         |
| Disabled E_Port | Locked L_Port      | Rate limit         |
| Encryption      | LOS TOV enable     | RSCN suppressed    |
| EX_Port         | Mirror Port        |                    |



## Viewing Fibre Channel port attributes

To view the Fibre Channel port attributes for a single port, in privileged EXEC mode, enter the **show running-config interface FibreChannel *rbridge-id/slot/port*** command for the port you want to view. To view the Fibre Channel port attributes for all Fibre Channel ports in the fabric, enter the **show running-config interface FibreChannel** command without any additional parameters.

Whether you view attributes for a single port or for all ports, the settings for the desire-distance, isl-r\_rdy, trunk-enable, and shutdown attributes are always displayed. The speed, long-distance, vc-link-init, and fill-word attributes are displayed only if they are set to nondefault values.

The following example displays the Fibre Channel port attributes for a single port. In this case, the speed, long-distance, and vc-link-init attributes appear because they have been set to values other than their default values.

```
switch# show running-config interface FibreChannel 8/0/1
interface FibreChannel 8/0/1
 speed 8gbps
 long-distance ld
 vc-link-init arb
 desire-distance 0
 no isl-r_rdy
 trunk-enable
 shutdown
!
```

The following example shows Fibre Channel attributes for all Fibre Channel ports. In this case, the speed, long-distance, vc-link-init, and fill-word attributes are set to their default values for all of the interfaces shown.

```
switch# show running-config interface FibreChannel
interface FibreChannel 3/0/1
 desire-distance 0
 no isl-r_rdy
 trunk-enable
 no shutdown
!
interface FibreChannel 3/0/2
 desire-distance 0
 no isl-r_rdy
 trunk-enable
 no shutdown
!
interface FibreChannel 3/0/3
 desire-distance 0
 no isl-r_rdy
 trunk-enable
 no shutdown
!
(output truncated)
```

To view the setting of a single attribute on a specific port, regardless of whether the attribute is set to its default value, enter the **show running-config interface FibreChannel *rbridge-id/slot/port attribute*** command.

The following example shows the setting of the speed attribute for port 66/0/1:

```
switch# show running-config interface FibreChannel 66/0/1 speed
```

```
interface FibreChannel 66/0/1
 speed auto
!
```

### Setting Fibre Channel port speed

This procedure sets the ports speed to 1, 2, 4, or 8 Gbps, or to autonegotiate (the default value).

1. In privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter the **interface FibreChannel** *rbridge-id/slot/port* command for the port on which you want to set the speed.

A configuration submode prompt appears.

3. Enter the **speed** command and the desired speed in Gbps.

The following example sets the port speed to 4 Gbps.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface FibreChannel 8/0/1
switch(config-FibreChannel-8/0/1)# speed 4
```

### Long-distance operation

Use the **interface FibreChannel long-distance** command to support long-distance links and to allocate enough full-size frame buffers on a specific port. Changes made by this command are persistent across switch reboots and power cycles. This command supports the following long-distance link modes:

- Normal mode (LO) - LO is the normal (default) mode for a port. It configures the port as a regular port. A total of 20 full-size frame buffers are reserved for data traffic, regardless of the port operating speed. The maximum supported link distance is up to 5 km at 2 Gbps, up to 2 km at 4 Gbps, and up to 1 km at 8 Gbps.
- Extended mode (LE) - LE configures an E\_Port distance greater than 5 km and up to 10 km. The baseline for the calculation is one credit per km at 2 Gbps, which yields the following values for 10 km:
  - 5 credits per port at 1 Gbps
  - 10 credits per port at 2 Gbps
  - 20 credits per port at 4 Gbps
  - 40 credits per port at 8 Gbps
- Dynamic Long-Distance mode (LD) - LD calculates buffer-to-buffer (BB) 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 BB credits required in LD extended link mode based on a maximum Fibre Channel payload size of 2,112. You can place an upper limit on the calculation by providing a desired distance value (**desired-distance** command). Network 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 specified desired distance (the smaller value) is used in the calculation.

- Static Long-Distance mode (LS) - LS calculates a static number of BB credits based only on a user-defined desired distance value set using the **desire-distance** command. LS also assumes that all Fibre Channel payloads are 2,112 bytes. Specify LS to configure a static long-distance link with a fixed buffer allocation greater than 10 km. Up to a total of 1,452 full-size frame buffers are reserved for data traffic, depending on the specified desired distance value.

## Configuring for long-distance operation

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 at the same distance level without compromising local switch performance.
- Only qualified Brocade SFP transceivers are used. Only Brocade-branded or certain Brocade-qualified SFPs are supported.

To configure a Fibre Channel port for long-distance operation, follow these steps:

1. In privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter the **interface FibreChannel rbridge-id/slot/port** command for the Fibre Channel port you want to configure.  
  
A configuration submode prompt appears.
3. For 8 Gbps only, enter the **fill-word** command to set the fill word to the same value as for the remote port.
4. Enter the **long-distance** command to set the long distance mode.
5. For LD and LS modes only, enter the **desire-distance** command to set the desired distance.
6. For 8 Gbps only, enter the **vc-link-init** command to set the fill word for the long distance link to the same value as the fill word for the remote port.
7. On the Fabric OS end of the ISL, configure the Fibre Channel port with the same values set in [step 3](#) through [step 4](#) using the Fabric OS **portCfgFillWord** and **portCfgLongDistance** commands.

The following example sets the long distance mode to LS for a distance of 100 km.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface FibreChannel 8/0/1
switch(config-FibreChannel-8/0/1)# fill-word arbff-arbff
switch(config-FibreChannel-8/0/1)# long-distance ls
switch(config-FibreChannel-8/0/1)# desire-distance 100
switch(config-FibreChannel-8/0/1)# vc-link-init arb
switch(config-FibreChannel-8/0/1)# do show running-config interface FibreChannel
8/0/1
interface FibreChannel 8/0/1
 fill-word arbff-arbff
 long-distance ls
 vc-link-init arbff
 desire-distance 100
 no isl-r_rdy-mode
 no shutdown
```

## Configuring a Fibre Channel port for trunking

A link can be configured to be part of a trunk group. Two or more links in a port group form a trunk group when they are configured for the same speed, the same distance level, and their link distances are nearly equal.

1. In privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter the **interface FibreChannel rbridge-id/slot/port** command for the desired port.  
A configuration submode prompt appears.
3. Enter the **trunk-enable** command.

The following example configures the link attached to port 4 on routing bridge 8 to be part of a trunk group.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface FibreChannel 8/0/4
switch(config-FibreChannel-8/0/4)# trunk-enable
```

## Monitoring Fibre Channel ports

To monitor a Fibre Channel port, in privileged EXEC mode, enter the **show interface FibreChannel rbridge-id/slot/port** command for the Fibre Channel port you want to monitor. The command output provides lots of information about the various hardware counters associated with the port.

This command has a basic version and a detail version. The basic version of the command provides general port information such as status, identification, and configuration information, along with interrupt statistics, link status counters, and so on, as shown in the following example:

```
switch# show interface FibreChannel 66/0/1
fibrechannel 66/0/1 is up (No_Light). Protocol state is down.
Pluggable media present
LineSpeed Actual:
PortSpeed: N8Gbps
portDisableReason: None
PortId: 427900
PortIfId: 4302303f
PortWwn: 20:79:00:05:33:67:26:78
Distance: normal

Last clearing of show interface counters: 00:00:00
Interrupts: 0 Link_failure: 0 Frjt: 0
Unknown: 0 Loss_of_sync: 0 Fbsy: 0
Lli: 2 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
```

```

Rate info:
 Bandwidth: 8.00G
 Tx performance: 0 B/sec
 Rx performance: 0 B/sec

```

The detail version of the command, illustrated below, tells you how much traffic has been transmitted or received, and how many times certain error conditions have occurred. Specifically, the `tim_txcrd_z` counters tell you how many times the port was unable to transmit frames because the transmit BB credit was zero. A number greater than zero indicates either congestion on the port or that a device is affected by latency. A bigger number indicates a bigger problem. A sample is taken every 2.5 microseconds.

```

switch# show interface FibreChannel 66/0/1 detail
fibrechannel 66/0/1 is up. Protocol state is up (connected)
Pluggable media present
LineSpeed Actual: 400,800_MB/s
portSpeed: N8Gbps
portDisableReason: None
portId 423100
portIfId: 43020026
portWwn: 20:31:00:05:33:6f:27:57
Distance normal

```

```
Last clearing of show interface counters: 00:00:00
```

#### Rx Statistics:

|             |     |                              |
|-------------|-----|------------------------------|
| stat_wrx    | 118 | 4-byte words received        |
| stat_frx    | 4   | Frames received              |
| stat_c2_frx | 0   | Class 2 frames received      |
| stat_c3_frx | 0   | Class 3 frames received      |
| stat_lc_rx  | 2   | Link control frames received |
| stat_mc_rx  | 0   | Multicast frames received    |

#### Tx Statistics:

|                       |      |                                   |
|-----------------------|------|-----------------------------------|
| stat_wtx              | 282  | 4-byte words transmitted          |
| stat_ftx              | 12   | Frames transmitted                |
| stat_mc_tx            | 0    | Multicast frames transmitted      |
| tim_txcrd_z           | 2881 | Time TX Credit Zero (2.5Us ticks) |
| tim_txcrd_z_vc 0- 3:  | 2881 | 0 0 0                             |
| tim_txcrd_z_vc 4- 7:  | 0    | 0 0 0                             |
| tim_txcrd_z_vc 8-11:  | 0    | 0 0 0                             |
| tim_txcrd_z_vc 12-15: | 0    | 0 0 0                             |

#### Error Statistics

|                    |   |                                                  |
|--------------------|---|--------------------------------------------------|
| er_enc_in          | 0 | Encoding errors inside of frames                 |
| er_crc             | 0 | Frames with CRC errors                           |
| er_trunc           | 0 | Frames shorter than minimum                      |
| er_toolong         | 0 | Frames longer than maximum                       |
| er_bad_eof         | 0 | Frames with bad end-of-frame                     |
| er_enc_out         | 0 | Encoding error outside of frames                 |
| er_bad_os          | 1 | Invalid ordered set                              |
| er_rx_c3_timeout   | 0 | Class 3 receive frames discarded due to timeout  |
| er_tx_c3_timeout   | 0 | Class 3 transmit frames discarded due to timeout |
| er_c3_dest_unreach | 0 | Class 3 frames discarded due to destination      |

#### unreachable

|                  |   |                              |
|------------------|---|------------------------------|
| er_other_discard | 0 | Other discards               |
| er_type1_miss    | 0 | frames with FTB type 1 miss  |
| er_type2_miss    | 0 | frames with FTB type 2 miss  |
| er_type6_miss    | 0 | frames with FTB type 6 miss  |
| er_zone_miss     | 0 | frames with hard zoning miss |
| er_lun_zone_miss | 0 | frames with LUN zoning miss  |
| er_crc_good_eof  | 0 | Crc error with good eof      |
| er_inv_arb       | 0 | Invalid ARB                  |

# 12 Monitoring Fibre Channel ports

```
Port Error Info:
 Loss_of_sync:1
 Loss_of_sig:2
 Frjt:0
 Fbsy:0

Buffer Information:
 Lx Max/Resv Buffer Needed Link Remaining
 Mode Buffers Usage Buffers Distance Buffers
=====
 - 8 0 0 - 924

Rate info:
 Bandwidth: 8.00G
 Tx performance: 0 B/sec
 Rx performance: 0 B/sec
```

# System Monitor

---

## In this chapter

- [System Monitor overview](#) ..... 191
- [System Monitor examples](#)..... 193
- [Threshold Monitor overview](#) ..... 195
- [Threshold Monitor examples](#) ..... 200

## System Monitor overview

System Monitor provides customizable monitoring thresholds, which allow you to monitor the health of each component of a switch. Whenever a switch component exceeds a configured threshold, System Monitor automatically provides notification by means of e-mail or RASlog messages, depending on the configuration.

Because of platform-specific values that vary from platform to platform, it was previously not possible to configure platform-specific thresholds through a global CLI command. Beginning with NOS v4.0.0, it is now possible to monitor not only a stand-alone switch, but also individual switches in a management (or fabric) cluster. This is done in RBridge ID configuration mode, by addressing the rbridge-id of the selected switch. Both monitoring modes—logical chassis cluster and standalone— are illustrated in this chapter.

Threshold and notification configuration procedures are described in the following sections.

### Monitored components

The following FRUs and temperature sensors are monitored on supported switches:

- **LineCard**—Displays the threshold for the line card.
- **MM**—Displays the threshold for the management module.
- **SFM**—Displays the threshold for the switch fabric module device.
- **cid-card**—Displays the threshold for the chassis ID card component.
- **compact-flash**—Displays the threshold for the compact flash device.
- **fan**—Configures fan settings.
- **power**—Configures power supply settings.
- **sfp**—Displays the threshold for the small form-factor pluggable (SFP) device.
- **temp**—Displays the threshold for the temperature sensor component.

**NOTE**

CID cards can be faulted and removed. The system continues to operate normally as long as one CID card is installed. If both CID cards are missing or faulted, the switch will not operate.

## Monitored FRUs

System Monitor monitors the absolute state of the following FRUs:

- Fan
- Power supply
- CID card
- SFP
- Line card

Possible states for all monitored FRUs are removed, inserted, on, off, and faulty. A state of none indicates the switch is not configured. If the FRU is removed, inserted, or goes into a faulty state, System Monitor sends a RASLog message or an e-mail alert, depending on the configuration.

Based on the configured threshold, each component can be in a marginal state or a down state. If a component is in a marginal state or a down state, System Monitor generates a RASLog message to alert the user. It also generates a separate RASLog message for the overall health of the switch.

**NOTE**

For details about each RASLog message, refer to the “RAS System Messages” chapter of the *Network OS Message Reference*.

[Table 32](#) lists the default threshold settings for components monitored by System Monitor on supported switches.

**TABLE 32** Hardware platform default settings for supported switches

| Platform                                        | Hardware component | Default setting | Marginal thresholds | Down thresholds |
|-------------------------------------------------|--------------------|-----------------|---------------------|-----------------|
| Brocade VDX 6710                                | Power supply       | 2               | 1                   | 2               |
|                                                 | Temperature sensor | 3               | 1                   | 2               |
|                                                 | Compact flash      | 1               | 1                   | 0               |
|                                                 | Fan                | 2               | 1                   | 2               |
| Brocade VDX 6720-24 single board 24-port switch | Power supply       | 2               | 1                   | 2               |
|                                                 | Temperature sensor | 3               | 1                   | 2               |
|                                                 | Compact flash      | 1               | 1                   | 0               |
|                                                 | Fan                | 2               | 1                   | 2               |
| Brocade VDX 6720-60 single board 60-port switch | Power supply       | 2               | 1                   | 2               |
|                                                 | Temperature sensor | 6               | 1                   | 2               |
|                                                 | Compact flash      | 1               | 1                   | 0               |
|                                                 | Fan                | 5               | 1                   | 2               |



**TABLE 32** Hardware platform default settings for supported switches (Continued)

| Platform            | Hardware component | Default setting | Marginal thresholds | Down thresholds |
|---------------------|--------------------|-----------------|---------------------|-----------------|
| Brocade VDX 6730-32 | Power supply       | 2               | 1                   | 2               |
|                     | Temperature sensor | 3               | 1                   | 2               |
|                     | Compact flash      | 1               | 1                   | 0               |
|                     | Fan                | 2               | 1                   | 2               |
| Brocade VDX 6730-76 | Power supply       | 2               | 1                   | 2               |
|                     | Temperature sensor | 6               | 1                   | 2               |
|                     | Compact flash      | 1               | 1                   | 0               |
|                     | Fan                | 5               | 1                   | 2               |
| Brocade VDX 8770-4  | Power supply       | 2               | 1                   | 2               |
|                     | Temperature sensor | 3               | 1                   | 2               |
|                     | Compact flash      | 1               | 1                   | 0               |
|                     | Fan                | 2               | 1                   | 2               |
| Brocade VDX 8770-8  | Power supply       | 3               | 6                   | 7               |
|                     | Temperature sensor | 3               | 1                   | 2               |
|                     | Compact flash      | 1               | 1                   | 0               |
|                     | Fan                | 4               | 1                   | 2               |

## System Monitor examples

The following are example basic configurations that illustrate various functions of the **system-monitor** command and related commands:

- [Viewing system SFP optical monitoring defaults](#)
- [Setting system thresholds](#)
- [Setting state alerts and actions](#)
- [Displaying the switch health status](#)
- [Configuring e-mail alerts](#)

### NOTE

For CLI details, see the *Network OS Administrator's Guide Supporting Network OS v.4.0.0*.

## Viewing system SFP optical monitoring defaults

You can view the optical monitoring default values by using the **show defaults threshold** command as follows, and selecting the SFP type:

```
switch# show defaults threshold sfp type 1GLR
```

## Setting system thresholds

Each component can be in one of two states, down or marginal, based on factory-defined or user-configured thresholds. (The default thresholds are listed in [Table 32](#) on page 192.)

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge ID configuration mode, as in the following example.

```
switch(config)# rbridge-id 154
```

Change down-threshold and marginal-threshold values for the SFM:

```
switch(config-rbridge-id-154)# system-monitor sfm threshold down-threshold 3
marginal-threshold 2
```

---

### NOTE

You can disable the monitoring of each component by setting the down threshold and the marginal threshold to 0 (zero).

---

## Setting state alerts and actions

System Monitor generates an alert when there is a change in the state from the default or defined threshold.

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge subconfiguration mode, as in the following example.

```
switch(config)# rbridge-id 154
```

The following are example configurations

To enable a RASLog alert when the power supply is removed:

```
switch(config-rbridge-id-154)# system-monitor power alert state removed action
raslog
```

---

### NOTE

There are no alerts for MM, compact-flash, or temp. There are no alert actions for SFPs.

---

## Displaying the switch health status

Enter the global **show system monitor** command to display the switch health status.

```
switch# show system monitor
```

```
** System Monitor Switch Health Report **
RBridge 154 switch status : MARGINAL
 Time of Report : 2013-03-24 20:51:53
 Power supplies monitor : MARGINAL
 Temperatures monitor : HEALTHY
 Fans monitor : HEALTHY
 Flash monitor : HEALTHY
```

## Configuring e-mail alerts

Use the **system-monitor-mail fru** command to configure e-mail threshold alerts for FRU, SFP, interface, and security monitoring. For an e-mail alert to function correctly, you must add the IP addresses and host names to the domain name server (DNS) in addition to configuring the domain name and name servers. A single email configuration is applicable for all switches in a management cluster. For complete information on the **system-monitor-mail relay host** command, refer to the *Network OS Command Reference*.

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter the following command to enable e-mail alerts and to configure the e-mail address.

```
switch(config)# system-monitor-mail fru enable email-id
```

The following **system-monitor-mail relay host** commands allow the sendmail agent on the switch to resolve the domain name and forward all e-mail messages to a relay server.

To create a mapping:

```
switch(config)# system-monitor-mail relay ip-address 1.2.3.4 domain-name
domain_name1.brocade.com
```

To delete the mapping:

```
switch(config)# no system-monitor-mail relay ip-address 1.2.3.4 domain-name
domain_name1.brocade.com
```

To change the domain name:

---

### NOTE

You must delete the first domain name before you can change it to a new domain name.

---

```
switch(config)# system-monitor-mail relay ip-address 1.2.3.4 domain-name
domain_name2.brocade.com
```

To delete the domain name and return to the default:

```
switch(config)# no system-monitor-mail relay ip-address 1.2.3.4 domain-name
domain_name2.brocade.com
```

## Threshold Monitor overview

With the **threshold-monitor** series of commands you can monitor CPU and memory usage of the system, interface and SFP environmental status, and security status and be alerted when configured thresholds are exceeded. These commands are configured in RBridge ID configuration mode to support Fabric Cluster (FC) and Management Cluster (MC) topologies.

With the **policy** operand (available for interface, SFP, and security monitoring), you can create your own custom policies that have nondefault thresholds, and apply them by means of the **apply** operand. This allows you to toggle between default settings and saved custom configuration settings and to apply actions and thresholds separately. For example, you can choose to use default threshold settings together with a customized subset of available actions, or you can modify some of the threshold settings and use the default action settings. You can also pause monitoring and actions by means of the **pause** operand.

For detailed information on the variables and operands of the **threshold-monitor** series of commands, refer to the *Network OS Command Reference*.

## CPU and memory monitoring

When configuring CPU monitoring, specify a value in the 1–100 range. When the CPU usage exceeds the limit, a threshold monitor alert is triggered. The default CPU limit is 75 percent. With respect to memory, the limit specifies a usage limit as a percentage of available resources.

When used to configure memory or CPU threshold monitoring, the limit value must be greater than the low limit and smaller than the high limit. The alert provided is a RASLog message, with the following options configurable under the **raslog** option of the **threshold-monitor cpu** or the **threshold-monitor memory** commands:

|                   |                                                                                                                                                                                                                                                                                      |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>high-limit</b> | Specifies an upper limit for memory usage as a percentage of available memory. This value must be greater than the value set by <b>limit</b> . When memory usage exceeds this limit, a RASLog CRITICAL message is sent. Values range from 0 through 80 percent.                      |
| <b>limit</b>      | Specifies the baseline memory usage limit as a percentage of available resources. When this value is exceeded, a RASLog WARNING message is sent. When the usage returns below the value set by <b>limit</b> , a RASLog INFO message is sent. Values range from 0 through 80 percent. |
| <b>low-limit</b>  | Specifies a lower limit for memory usage as percentage of available memory. This value must be smaller than the value set by <b>limit</b> . When memory usage exceeds or falls below this limit, a RASLog INFO message is sent.                                                      |
| <b>poll</b>       | Specifies the polling interval in seconds. The range is from 0 through 3600.                                                                                                                                                                                                         |
| <b>retry</b>      | Specifies the number of polling retries before desired action is taken, from 1 through 100.                                                                                                                                                                                          |

### NOTE

For CPU and memory thresholds, the low limit must be the lowest value and the high limit must be the highest value.

Table 33 lists the factory defaults for CPU and memory thresholds.

**TABLE 33** Factory defaults for CPU and memory threshold monitoring

| Operand           | Memory      | CPU         |
|-------------------|-------------|-------------|
| <b>low-limit</b>  | 40%         | N/A         |
| <b>limit</b>      | 60%         | 75%         |
| <b>high-limit</b> | 70%         | N/A         |
| <b>poll</b>       | 120 seconds | 120 seconds |
| <b>retry</b>      | 3           | 3           |

## SFP monitoring

The SFP parameters that can be monitored are shown in [Table 34](#).

**TABLE 34 SFP parameter descriptions**

| SFP parameter        | Description                                                     | Suggested SFP impact                                                                                   |
|----------------------|-----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Temperature          | Measures the temperature of the SFP, in degrees Celsius.        | High temperature suggests the SFP might be damaged.                                                    |
| Receive power (RXP)  | Measures the amount of incoming laser, in $\mu$ Watts.          | Describes the condition of the SFP. If this parameter exceeds the threshold, the SFP is deteriorating. |
| Transmit power (TXP) | Measures the amount of outgoing laser power, in $\mu$ Watts.    | Describes the condition of the SFP. If this parameter exceeds the threshold, the SFP is deteriorating. |
| Current              | Measures the amount of current supplied to the SFP transceiver. | Indicates hardware failures.                                                                           |
| Voltage              | Measures the amount of voltage supplied to the SFP.             | A value higher than the threshold indicates the SFP is deteriorating.                                  |

### *SFP thresholds*

You can customize SFP thresholds or actions by using the **threshold-monitor sfp** command, which enables you to perform the following tasks.

- Customize SFP configurations or accept SFP defaults.
- Manage the actions and thresholds for the Current, Voltage, RXP, TXP, and Temperature areas of the SFP.
- Suspend SFP monitoring.

If you do not provide the SFP type parameters, the default thresholds and actions are used. SFP types, monitoring areas, and default threshold values for the 16-Gbps and QSFP SFPs are listed in [Table 35](#).

**TABLE 35 Factory thresholds for SFP types and monitoring areas**

| SfpType | Area            | Default Value |      |
|---------|-----------------|---------------|------|
| 1 GSR   | Temperature (C) | 100           | -40  |
|         | Voltage (mV)    | 3600          | 3000 |
|         | RXP ( $\mu$ W)  | 1122          | 8    |
|         | TXP ( $\mu$ W)  | 1000          | 60   |
|         | Current (mA)    | 12            | 2    |
| 1 GLR   | Temperature (C) | 90            | -45  |
|         | Voltage (mV)    | 3700          | 2900 |
|         | RXP ( $\mu$ W)  | 501           | 6    |
|         | TXP ( $\mu$ W)  | 794           | 71   |
|         | Current (m)     | 45            | 1    |

**TABLE 35** Factory thresholds for SFP types and monitoring areas (Continued)

| SfpType | Area            | Default Value |      |
|---------|-----------------|---------------|------|
| 10 GSR  | Temperature (C) | 90            | -5   |
|         | Voltage (mVolt) | 3600          | 3000 |
|         | RXP ( $\mu$ W)  | 1000          | 32   |
|         | TXP ( $\mu$ W)  | 794           | 251  |
|         | Current (mA)    | 11            | 4    |
| 10 GLR  | Temperature (C) | 88            | -5   |
|         | Voltage (mV)    | 3600          | 2970 |
|         | RXP ( $\mu$ W)  | 1995          | 16   |
|         | TXP ( $\mu$ W)  | 1585          | 158  |
|         | Current (mA)    | 85            | 15   |
| 10 GUSR | Temperature (C) | 100           | -5   |
|         | Voltage (mV)    | 3600          | 2970 |
|         | RXP ( $\mu$ W)  | 2000          | 32   |
|         | TXP ( $\mu$ W)  | 2000          | 126  |
|         | Current (mA)    | 11            | 3    |
| QSFP    | Temperature (C) | 75            | -5   |
|         | Voltage (mV)    | 3600          | 2970 |
|         | RXP ( $\mu$ W)  | 1995          | 40   |
|         | TXP ( $\mu$ W)  | 0             | 0    |
|         | Current (mA)    | 10            | 1    |

***Threshold values***

High and low threshold values are the values at which potential problems might occur. For example, in configuring a temperature threshold for SFP, you can select the temperatures at which a potential problem can occur because of overheating or freezing.

A combination of high and low threshold settings can cause the following actions to occur:

- Above high threshold—A default or user-configurable action is taken when the current value is above the high threshold.
- Below high threshold—A default or user-configurable action is taken when the current value is between the high and low threshold.
- Below low threshold—A default or user-configurable action is taken when the current value is below the low threshold.

**NOTE**

Above low threshold is not supported.

## Security monitoring

System Monitor monitors all attempts to breach your SAN security, helping you fine-tune your security measures. If there is a security breach, System Monitor sends an email or RASLog alert. The following security areas are monitored:

- Telnet Violation, which occurs when a Telnet connection request reaches a secure switch from an unauthorized IP address.
- Login Violation, which occurs when a secure fabric detects a login failure.

Table 36 lists the factory defaults for security area settings.

**TABLE 36 Security area default settings**

| Area             | High threshold | Low threshold | Buffer | Timebase |
|------------------|----------------|---------------|--------|----------|
| Telnet Violation | 2              | 1             | 0      | Minute   |
| Login Violation  | 2              | 1             | 0      | Minute   |

## Interface monitoring

You can set thresholds for error statistics on all external Gigabit Ethernet interfaces: 1 GbE, 10 GbE, and 40 GbE. When any monitored error crosses the configured high or low threshold, an alert can be generated or a problem interface can be isolated (see “Port Fencing” on page 200).

### Interface error types

Table 37 describes the interface counters that can be monitored on external interfaces.

**TABLE 37 Interface errors that can be monitored on external interfaces**

| Interface area              | Description                                                                                                                                                                                                                                                                                                             | Port Fencing support | Threshold defaults             |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|--------------------------------|
| MissingTerminationCharacter | Number of frames terminated by anything other than the Terminate character; this includes termination due to the Error character.                                                                                                                                                                                       | No                   | Low 12<br>Buffer 0<br>High 300 |
| CRCAlignErrors              | Total number of frames received that had a length (excluding framing bits but including Frame Check Sequence (FCS) octets) of between 64 and 1518 octets. The error indicates either a bad FCS with an integral number of octets (an FCS error) or a bad FCS with a non-integral number of octets (an alignment error). | No                   | Low 12<br>Buffer 0<br>High 300 |
| IFG                         | Minimum-length interframe gap (IFG) between successive frames is violated. A typical IFG is 12 bytes.                                                                                                                                                                                                                   | Yes                  | Low 5<br>Buffer 0<br>High 100  |
| SymbolErrors                | An undefined (invalid) symbol received on the interface. Large symbol errors indicate a bad device, cable, or hardware.                                                                                                                                                                                                 | No                   | Low 0<br>Buffer 0<br>High 5    |

### NOTE

The default setting for above high threshold, above low threshold, below high threshold, and below low threshold actions is “[none].”

### *Port Fencing*

A port that is consistently unstable can harm the responsiveness and stability of the entire fabric and diminish the ability of the management platform to control and monitor the switches within the fabric. Port Fencing is not enabled by default; it disables the interface if a user-defined high threshold is exceeded. When a port that has exceeded its user-defined high threshold is fenced by software, the port is placed in Disabled state and held offline. After a port is disabled, user intervention is required for frame traffic to resume on the port.

---

**NOTE**

Port Fencing is supported for the RX IFG Violated error only.

---

## Threshold Monitor examples

The following are example basic configurations that illustrate various functions of the **threshold-monitor** commands.

- [Viewing threshold status](#)
- [CPU and memory threshold monitoring examples](#)
- [SFP monitoring example](#)
- [Security monitoring examples](#)
- [Interface monitoring examples](#)
- [Pausing and continuing threshold monitoring](#)

---

**NOTE**

For CLI details, see the *Network OS Command Reference*.

---

### Viewing threshold status

To see the status of currently configured thresholds, enter the **show running-config threshold-monitor** command with the RBridge ID, as follows:

```
switch# show running-config rbridge-id rbridge_id threshold-monitor
```

---

**NOTE**

Default values are not displayed under the **show running-config threshold-monitor** command. Only custom values are displayed when a user applies a policy.

---



To display the default values of thresholds and alert options, enter the **show defaults threshold** command, as in the following example for interfaces.

```
switch# show defaults threshold interface type Ethernet
```

Type: GigE-Port

| Area     | High Threshold |                 |                 | Low Threshold |                 |                 | Buffer | Time   |
|----------|----------------|-----------------|-----------------|---------------|-----------------|-----------------|--------|--------|
|          | Value          | Above<br>Action | Below<br>Action | Value         | Above<br>Action | Below<br>Action | Value  | Base   |
| MTC      | 300            | none            | none            | 12            | none            | none            | 0      | minute |
| CRCAlign | 300            | none            | none            | 12            | none            | none            | 0      | minute |
| Symbol   | 5              | none            | none            | 0             | none            | none            | 0      | minute |
| IFG      | 100            | none            | none            | 5             | none            | none            | 0      | minute |

MTC - Missing Termination Character

## CPU and memory threshold monitoring examples

- [Configuring CPU monitoring](#)
- [Configuring memory monitoring](#)

---

### NOTE

Support for the custom **policy** operand is not provided for CPU and memory threshold monitoring.

---

### *Configuring CPU monitoring*

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge ID configuration mode, as in the following example.

```
switch(config)# rbridge-id 154
```

3. Enter the **threshold-monitor cpu ?** command to see available options:

```
switch(config-rbridge-id-154)# threshold-monitor cpu ?
```

The following example changes the thresholds from the default, adjusts polling and retry attempts, and causes a RASLog message to be sent when thresholds are exceeded.

```
switch(config-rbridge-id-154)# threshold-monitor cpu actions raslog limit 65
poll 60 retry 10
```

---

### NOTE

This command does not support **low-limit** or **high-limit** under the **raslog** alert option.

---

## *Configuring memory monitoring*

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge ID configuration mode, as in the following example.

```
switch(config)# rbridge-id 154
```

3. Enter **threshold-monitor memory ?** to see available options:

```
switch(config-rbridge-id-154)# threshold-monitor memory ?
```

The following example changes the thresholds from the default and causes no message to be sent when thresholds are exceeded.

```
switch(config-rbridge-id-1)# threshold-monitor memory actions none high-limit 60 low-limit 40
```

## SFP monitoring example

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge ID configuration mode, as in the following example.

```
switch(config)# rbridge-id 154
```

3. Enter the **threshold-monitor sfp** command and create a custom policy.

```
switch(config-rbridge-id-154)# threshold-monitor sfp policy mypolicy type 1glr area temperature alert above highthresh-action raslog email
```

---

### NOTE

Refer to [“Security monitoring”](#) on page 199.

---

4. Apply the policy.

```
switch(config-rbridge-id-154)# threshold-monitor sfp apply mypolicy
```

## Security monitoring examples

- [Viewing security defaults](#)
- [Configuring security monitoring](#)

### *Viewing security defaults*

To display the default values of security threshold and alert options, enter the **show defaults security area** command with the **login-violation** or **telnet-violation** options:

```
switch# show defaults security area
```

### *Configuring security monitoring*

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge subconfiguration mode, as in the following example.  
`switch(config)# rbridge-id 154`
3. Enter the **threshold-monitor security** command to configure custom login-violation monitoring, as in the following example.
4. `switch(config-rbridge-id-154)# threshold-monitor security policy mypolicy area  
login-violation alert above highthresh-action raslog below  
highthresh-action email lowthresh-action none`
5. Apply the policy.  
`switch(config-rbridge-id-154)# threshold-monitor security apply mypolicy`

## Interface monitoring examples

- [Viewing interface defaults](#)
- [Configuring interface monitoring](#)

### *Viewing interface defaults*

Use the following command to view interface threshold defaults:

```
switch# show defaults threshold interface type Ethernet
```

See “[Viewing threshold status](#)” on page 200 for the results of this command.

### *Configuring interface monitoring*

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge ID configuration mode, as in the following example.  
`switch(config)# rbridge-id 154`
3. Enter the **threshold-monitor interface** command to configure custom interface monitoring, as in the following example.  
`switch(config-rbridge-id-154)# threshold-monitor interface policy mypolicy  
type ethernet area missingterminationcharacter alert above  
lowthresh-action email raslog`
4. Apply the policy.  
`switch(config-rbridge-id-154)# threshold-monitor interface apply mypolicy`

### Pausing and continuing threshold monitoring

By default, threshold monitoring is enabled.

To disable monitoring of a particular type, enter the **threshold-monitor [cpu | interface | memory | security | sfp] pause** command.

To re-enable monitoring, enter the **no version** of the above command.

---

#### NOTE

Not all functions of this command can be disabled. Continue to enter **?** at each level of the command synopsis to confirm which functions can be disabled.

---

# VMware vCenter

---

## In this chapter

- vCenter and Network OS integration ..... 205
- vCenter discovery ..... 206

## vCenter and Network OS integration

The VMware vCenter Server allows for the management of multiple ESX /ESXi servers and virtual machines (VMs) from different ESX servers through a single graphical user interface (GUI). It provides unified management of all the hosts and VMs in the data center, from a single console with an aggregate performance monitoring of clusters, hosts and VMs.

The VMware vCenter and Brocade Network OS integration, supported in Brocade VCS Fabric mode and non-VCS (standalone switch) mode, enables you to discover VMware ESX servers managed by a vCenter server. VMware's server hosts (ESX servers) are connected directly to the physical switches through the switch ports (edge ports in Brocade VCS Fabric mode). The server hosts implement a virtual switch (vSwitch), which is used to provide connections to the VMs. The fundamental requirement for the vCenter and Network OS integration is the IP-level management connectivity of the vCenter Server 4.0 version and later with the Brocade VDX switches.

---

### NOTE

The Network OS integration with vCenter requires vCenter version 4.0, 4.1, and 5.0.

---

You can view virtual switches and virtual machines, their associated MAC addresses, and network policies using the Network OS command line interface (CLI). Refer to the *Brocade Network OS Command Reference* for details about the **vcenter** and **vnetwork** commands.

## vCenter properties

The vCenter manages the VMware ESX/ESXi hosts. The vCenter user interface is provided through a vSphere client on the same management network as the vCenter, and virtual machines (VMs) are created using the vSphere client user interface. In addition to creating the VMs, the server administrator associates the VMs with distributed virtual switches, distributed virtual port groups, standard virtual switches (vSwitches) and standard port groups.

The vCenter automatically generates some of the VM properties (such as the MAC address), and some properties must be configured (such as the VLAN properties). Most of the VM configuration, including network policies, is done using the vCenter's vSphere user interface and is beyond the scope of this document.

For VMWare configuration information, visit the VMware documentation site.

## vCenter guidelines and restrictions

Follow these guidelines and restrictions when configuring vCenter:

- Special characters in the port group names are replaced with the URL-encoded values.
- Standard port groups with the same name that reside in different ESX/ESXi hosts must have identical VLAN settings across all hosts.
- For all vCenter port groups, NOS automatically creates a port profile with the following format: `auto-vcenter_name-datacenter_ID-port-group-name`

User editing of these auto port groups is not supported.

- NOS supports vCenter discovery that is based on events.
- NOS supports LLDP and QoS (IEEE 8021.p) for distributed virtual switches (dvSwitches).
- NOS supports up to 750 port groups in the vCenter.
- CDP/LLDP-receiving interface ports must not have any conflicting configurations (such as switch port and FCoE port configurations) on the interface that prevent them from being in a port-profiled mode.
- Before configuring a vCenter in the fabric, remove all the manually created port profiles that have vCenter inventory MAC associations.
- For NOS 4.0.0 versions, multiple data centers are supported, up to four data centers.
- Duplicate vCenter asset values are not supported, such as duplicate MAC addresses and duplicate Host names.

## vCenter discovery

The Brocade VDX switch connected to VMware ESX/ESXi hosts and virtual machines must be aware of network policies in order to allow or disallow traffic, which requires a discovery process by the VDX switch. During VDX switch configuration, relevant vCenters that exist in its environment and the discovery of virtual assets from the vCenter occurs in the following circumstances:

- When a switch boots up
- When a new vCenter is configured on the VDX switch and activated (activation turns on the timer processing, set to 30-minute intervals)
- When the discovery is explicitly initiated with the CLI

The following assets are discovered from the vCenter:

- Hosts and data centers associated with the vCenter
- Virtual machines (VMs) that have been created on the hosts
- VMware distributed virtual port groups (dvPortGroups)
- Standard port groups, with QoS priority associated with a dvPortGroup
- Standard virtual switches
- Distributed virtual switches

## vCenter configuration

### Step 1: Enabling QoS

You must edit the network resource pool settings and set QoS priorities. Refer to the latest VMware vSphere Networking documentation.

### Step 2: Enabling CDP/LLDP

In order for an Ethernet Fabric to detect the ESX/ESXi hosts, you must first enable Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) on all the virtual switches (vSwitches) and distributed vSwitches (dvSwitches) in the vCenter Inventory.

For more information, refer to the VMware KB article 1003885.

#### *Enabling CDP/LLDP on vSwitches*

Complete the following steps to enable CDP/LLDP on virtual switches (vSwitches).

1. Login as root to the ESX/ESXi Host.
2. Use the following command to verify the current CDP/LLDP settings.

```
[root@server root]# esxconfig-vswitch -b vSwitch1
```

3. Use the following command to enable CDP/LLDP for a given virtual switch. Possible values here are advertise or both.

```
[root@server root]# esxconfig-vswitch -B both vSwitch1
```

#### *Enabling CDP/LLDP on dvSwitches*

Complete the following steps to enable CDP on distributed virtual switches (dvSwitches).

1. Connect to the vCenter server by using the vSphere Client.
2. In the vCenter Server home page, click **Networking**.
3. Right-click the distributed virtual switches (dvSwitches) and click **Edit Settings**.
4. Select **Advanced** under **Properties**.
5. Use the check box and the drop-down list to change the CDP/LLDP settings.

### Step 3: Adding and Activating vCenter

After enabling CDP on all the vSwitches and dvSwitches in the vCenter, the NOS-side configuration is a two step process: adding the vCenter and activating the vCenter.

#### *Adding the vCenter*

You must add the vCenter before initiating any discovery transactions. To authenticate with a specific vCenter, you must first configure the URL, login, and password properties on the VDX switch.

**NOTE**

By default, the vCenter server accepts only HTTPS connection requests.

Enter the **vcenter** command with the name, URL, username, and password of the vCenter.

```
switch(config)# vcenter myvcenter url https://10.2.2.2 username user
password pass
```

**Activating the vCenter**

After adding the vCenter, you must activate the configured vCenter instance.

**NOTE**

In VCS mode, you can configure the vCenter by using any node. Discovery is initiated by the primary node.

1. Enter the **config** command.
2. Enter the **vcenter** command to activate the vCenter.

```
switch(config)# vcenter myvcenter activate
```

Immediately following first-time vCenter activation, the Network OS (NOS) starts the virtual asset discovery process. Use the **show vnetwork vcenter status** command to display the vnetwork status. For example:

```
switch# show vnetwork vcenter status
vCenter Start Elapsed (sec) Status
=====
myvcenter 2011-09-07 14:08:42 10 In progress
```

When the discovery process completes, the status displays as "Success." NOS has performed all the necessary configurations needed for the vCenter Server. NOS is now ready for CDP transmissions from the virtual switches to identify which ESX/ESXi host is connected to which physical interface in the Ethernet Fabric.

**Discovery timer interval**

By default, NOS queries the vCenter updates every three minutes. If any virtual assets are modified (for example, adding or deleting virtual machines (VMs), or changing VLANs), NOS detects those changes and automatically reconfigures the Ethernet Fabric during the next periodic rediscovery attempt.

Use the **vcenter interval** command to manually change the default timer interval value to suit the individual environment needs.

```
switch(config)# vcenter myvcenter interval ?
Possible completions:
<NUMBER:0-1440> Timer Interval in Minutes (default = 30)
```

**NOTE**

Best practice is to keep the discovery timer interval value at the default (30). A value of 0 disables the periodic vCenter discovery.



## User-triggered vCenter discovery

Use the **vnetwork vcenter** command to trigger a vCenter discovery manually.

```
switch# vnetwork vcenter myvcenter discover
```

## Viewing the discovered virtual assets

Enter one of the following **show vnetwork** asset commands:

```
switch# show vnetwork dvpgs datacenter datacenter_name vcenter vcenter_name
switch# show vnetwork dvs datacenter datacenter_name vcenter vcenter_name
switch# show vnetwork hosts datacenter datacenter_name vcenter vcenter_name
switch# show vnetwork pgs datacenter datacenter_name vcenter vcenter_name
switch# show vnetwork vcenter status datacenter datacenter_name
 vcenter vcenter_name
switch# show vnetwork vmpolicy datacenter datacenter_name vcenter vcenter_name
switch# show vnetwork vms datacenter datacenter_name vcenter vcenter_name
switch# show vnetwork vss datacenter datacenter_name vcenter vcenter_name
```

where:

- **dvpgs**—Displays discovered distributed virtual port groups.
- **dvs**—Displays discovered distributed virtual switches.
- **hosts**—Displays discovered hosts.
- **pgs**—Displays discovered standard port groups.
- **vcenter status**—Displays configured vCenter status.
- **vmpolicy**—Displays the following network policies on the Brocade VDX switch: associated media access control (MAC) address, virtual machine, (dv) port group, and the associated port profile.
- **vms**—Displays discovered virtual machines (VMs).
- **vss**—Displays discovered standard virtual switches.

Refer to the *Network OS Command Reference* for detailed information about the **show vnetwork** commands.

## 14 vCenter configuration

# Configuring Remote Monitoring

---

## In this chapter

- [RMON overview . . . . .](#) 211
- [RMON configuration and management . . . . .](#) 211

## RMON overview

Remote monitoring (RMON) is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. The RMON specification defines a set of statistics and functions that can be exchanged between RMON-compliant console managers and network probes. As such, RMON provides you with comprehensive network-fault diagnosis, planning, and performance-tuning information.

## RMON configuration and management

Alarms and events are configurable RMON parameters:

- **Events**—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry, an SNMP trap, or both. You must define the events before an alarm can be configured. If you do not configure the RMON event first, you will receive an error when you configure the alarm settings.
- **Alarms**—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms are paired with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.

### Default RMON configuration

By default, no RMON alarms and events are configured and RMON collection statistics are not enabled.

### Configuring RMON events

You can add or remove an event in the RMON event table that is associated with an RMON alarm number.

To configure RMON events, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
switch# configure terminal
```

2. Configure the RMON event.

```
switch(config)# rmon event 27 description Rising_Threshold log owner
john_smith trap syslog
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Configuring RMON Ethernet group statistics collection

You can collect RMON Ethernet group statistics on an interface. RMON alarms and events must be configured for you to display collection statistics. By default, RMON Ethernet group statistics are not enabled.

To collect RMON Ethernet group statistics on an interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
switch# configure terminal
```

2. Enter the **interface** command to specify the Data Center Bridging (DCB) interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following example format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 0/1
```

3. Enable the DCB interface.

```
switch(conf-if-te-0/1)# no shutdown
```

4. Configure RMON Ethernet group statistics on the interface.

```
switch(conf-if-te-0/1)# rmon collection stats 200 owner john_smith
```

5. Return to privileged EXEC mode.

```
switch(conf-if-te-0/1)# end
```

6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Configuring RMON alarm settings

To configure RMON alarms and events, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
switch# configure terminal
```

2. Configure the RMON alarms.

**Example of an alarm that tests every sample for a rising threshold**

```
switch(config)# rmon alarm 5 1.3.6.1.2.1.16.1.1.1.5.65535 interval 30 absolute
 rising-threshold 95 event 27 owner john_smith
```

**Example of an alarm that tests the delta between samples for a falling threshold**

```
switch(config)# rmon alarm 5 1.3.6.1.2.1.16.1.1.1.5.65535 interval 10 delta
 falling-threshold 65 event 42 owner john_smith
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```



# Network OS Security Configuration

This section describes security features, and includes the following chapters:

- [Managing User Accounts](#) ..... 217
- [External server authentication](#) ..... 235
- [FIPS Support](#) ..... 265
- [Fabric Authentication](#) ..... 279





# Managing User Accounts

---

## In this chapter

- [User accounts](#) ..... 217
- [Role-based access control \(RBAC\)](#) ..... 221
- [Command access rules](#) ..... 223
- [Password policies](#) ..... 229
- [Security event logging](#) ..... 233

## User accounts

A user account allows authorized user access to the switch CLI. A user account must be assigned a role to specify the account's access privileges. A user account can be disabled at any point, preventing the user from logging in to the switch. A user can only be unlocked when the account is auto-locked because the user exceeded the configured threshold for failed login attempts. Only an administrator can create, change, unlock, or delete user accounts.

All modules that pertain to security, for example, user and user roles, RBAC, and password attributes (for example, encryption), are globally configurable data entities. This means that if a switch is in logical chassis cluster mode, all switches in the cluster will have a common configuration for all the above mentioned entities.

### Default accounts in the local switch user database

Network OS comes with two predefined user accounts that are part of the factory-default settings. Brocade recommends that you change the password for all default accounts during the initial installation and configuration for each switch.

The default user accounts are “admin” and “user,” and these accounts are associated with the corresponding admin” and “user” roles in the switch-local user database. Only the “admin” and “user” users can access the CLI and, except for the account password, no other attributes can be changed for the default users “admin” and “user.”

By default, all account information is stored in the switch-local user database. User authentication and tracking of logins to the switch is local by default.

---

#### NOTE

The maximum number of user accounts, including the default accounts, is 64. The maximum number of roles, including the default roles is 64. For any environment requiring more than 64 users, you should adopt an authentication, authorization, and accounting (AAA) service for user management. Refer to [Chapter 17, “External server authentication”](#) for more information. The maximum number of active Telnet or CLI sessions supported per switch is 32.

---

## Creating and modifying a user account

When you create a user account you must specify three mandatory attributes: an account login name, a role, and a password. The remaining attributes are optional.

**TABLE 38** User account attributes

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name                | The name of the account. The user account name is case-sensitive, must not exceed 40 characters, and must begin with a letter. The text string can contain letters, numbers, underscore (_), and periods (.). If the user name specified already exists, the <b>username</b> command modifies the existing role.                                                         |
| role                | The role assigned to the user defines the RBAC access privileges for the account.                                                                                                                                                                                                                                                                                        |
| password            | The account password must satisfy all currently enforced password rules. Refer to the section “ <a href="#">Password policies</a> ” on page 229 for more information.                                                                                                                                                                                                    |
| encryption-level    | The password encryption level. You can choose to encrypt the password (7) or leave it in clear text (0). If you do not specify an encryption level, the default, clear text (0), is the default.                                                                                                                                                                         |
| desc                | A description of the account. The description can be up to 64 characters long, and can include any printable ASCII character, except for the following characters: single quotation marks ('), double quotation marks("), exclamation point (!), colon (:), and semi-colon (;). If the description contains spaces, you must enclose the text in double quotation marks. |
| enable true   false | Indicates whether the account is enabled or disabled. A user whose account is disabled cannot log in. The default account status is enabled.                                                                                                                                                                                                                             |

### *Creating a user account*

The following example creates a new user account with the minimally required attributes: name, role, and password. The account name “brcdUser” has the default user privilege of accessing commands in the privileged EXEC mode

1. In the privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter **username** command with the specified parameters.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# username brcdUser role user password welcome
```

Displaying user account information

The user account information is saved in switch configuration file.

- Use the **show running-config username** command in privileged EXEC mode to display all configured users.

```
switch# show running-config username
username admin password "BwrsDbB+tABWGWpINOVKoQ==\n" encryption-level 7 role
admin desc Administrator
username user password "BwrsDbB+tABWGWpINOVKoQ==\n" encryption-level 7 role
user desc User
```

- Use the **show running-config username username** command in privileged EXEC mode to display a single user.

```
switch# show running-config username admin
username admin password "BwrsDbB+tABWGWpINOVKoQ==\n" encryption-level 7 role
admin desc Administrator
```

- Use the **show running-config username username enable** command in privileged EXEC mode to display whether the account is enabled or disabled.

```
switch# show running-config username admin enable
username admin enable true
```

### *Modifying an existing user account*

The syntax for the account *create* and *modify* operations looks alike. The difference is that there are no mandatory parameters for modifying an existing account. The system internally recognizes whether a new account is created or an existing account is modified operation by checking whether the user account is already present in the configuration database.

The following example adds a description to the previously created “brcdUser” account.

1. In the privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter **username** command with the specified parameters.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# username brcdUser
switch(config-username-brcdUser)# desc "Brocade guest account"
```

### *Disabling a user account*

You can disable a user account by setting the enable parameter to “false”. All active login sessions for a user are terminated when a user account is disabled.

1. In the privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter **username** command with the specified parameters.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# username testUser enable false
```

### *Deleting a user account*

1. In the privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **no username** command.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no username testUser
```

.All active login sessions for a user are terminated when a user account is deleted.

### *Unlocking a user account*

A user account is automatically locked by the system when the configured threshold for repeated failed login attempts has been reached. The account lockout threshold is a configurable parameter. Refer to the section [“Account lockout policy”](#) on page 231 for more information.

If a user account is locked out of a switch, that same user can still try to log in on another switch in the cluster; however, the unlocking is done on the given rbridge-id(s) irrespective of whether the user is not locked on one or more switches.

The following procedure shows the commands used to unlock a user account.

---

**NOTE**

While the **username** and **no username** commands are global configuration commands, the **unlock username** command is a privileged EXEC command.

---

1. Enter the **show users** command in the privileged EXEC mode to display currently active sessions and locked out users.
2. Enter the **unlock username** command in the privileged EXEC mode to unlock the locked user account.
3. Verify that the user has been unlocked. The **show users** command should display “no locked users”.

```
switch# show users
USER SESSIONS
RBridge
ID Username Host Ip Device Time Logged In
2 user 10.70.4.105 vty/0 2012-04-30 01:59:35
1 user 10.70.4.105 vty/0 2012-04-30 01:57:41
1 admin 10.70.4.105 vty/2 2012-04-30 01:58:41
1 user 10.70.4.105 vty/3 2012-09-30 02:04:42
LOCKED USERS
RBridge
ID username
1 testUser
switch# unlock username testUser
Result: Unlocking the user account is successful
switch# show users
USER SESSIONS
RBridge
ID Username Host Ip Device Time Logged In
2 user 10.70.4.105 vty/0 2012-04-30 01:59:35
1 user 10.70.4.105 vty/0 2012-04-30 01:57:41
1 admin 10.70.4.105 vty/2 2012-04-30 01:58:41
1 user 10.70.4.105 vty/3 2012-09-30 02:04:42
LOCKED USERS
RBridge
ID username
no locked users
```

### *Configuring a user alias*

You can specify a global alias and user alias for the switch using the **alias** command. The **alias** command operates in two slightly different ways, depending on which configuration mode you are using; global alias and user-level alias. The global alias is accessible across all users. The user-level alias is accessible only when the respective user logs in.

To set a global alias and a user alias, perform the following task in global configuration mode.

1. Enter alias configuration mode.

```
switch(config)# alias-config
```

2. Set the global user-alias for the switch.

```
switch(config-alias-config)# alias redwood engineering
```

3. Enter user configuration mode.

```
switch(config-alias-config)# user john smith
```

4. Set the user-level alias.

```
switch(config-alias-config-user)# alias manager engineering
```

## Role-based access control (RBAC)

Network OS uses role-based access control (RBAC) as the authorization mechanism. You can create roles dynamically and associate them with rules to define the permissions applicable to a particular role. Every user account must be associated with a role and only a single role can be associated with any given account.

RBAC specifies access rights to resources. When a user executes a command, privileges are evaluated to determine access to the command based on the role of the user.

In Logical chassis cluster mode, the configuration is applied to all nodes in the cluster.

### Default roles

All Brocade VDX switches support two default roles, “user” and “admin.” You cannot modify the attributes of default roles; however, you can assign the default roles to non-default user accounts. The default roles have the following access privileges:

- The user role has limited privileges that are restricted to executing show commands in the Privileged EXEC mode, as well as the following operational commands: **ping**, **ssh**, **telnet**, and **traceroute**. User accounts associated with the user role cannot access configuration commands that are available only in global configuration mode.
- The admin role has the highest privileges. All commands available in privileged EXEC mode and in global configuration mode are accessible to the user associated with the admin role.

With a new switch, only the admin user account has access to perform user and role management operations. The admin user can create any roles and configure those roles for access to user and role management operations.

### User-defined roles

In addition to the default roles, Network OS supports the creation of user-defined roles. A user-defined role starts from a basic set of privileges which are then refined by adding special rules. When you have created a role, you can assign a name to the role and then associate the role to one or more user accounts. The following tools are available for managing user-defined roles:

- The **role** command defines new roles and deletes user-defined roles,
- The **rule** command allows you to specify access rules for specific operations and assign these rules to a given role.
- The **username** command associates a given user-defined role with a specific user account.

## Creating a user-defined role

A user-defined role has a mandatory name and an optional description as shown in [Table 39](#).

**TABLE 39** Role attributes

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | The role name must be unique, begin with a letter, and can contain alphanumeric characters and underscores. The length of the role name should be between 4 and 32 characters. The name cannot be same as that of an existing user, an existing default role, or an existing user-defined role.                                                                                                               |
| desc      | An optional description of the role. The description can be up to 64 characters and can include any printable ASCII character, except for the following characters: single quotation marks ('), double quotation marks ("), exclamation point (!), colon (:), and semi-colon (;). If the description contains spaces, you must enclose the text in double quotation marks. If the description contains spaces |

The operation of creating a role must satisfy the following criteria to succeed:

- The maximum number of roles supported on a chassis is 64.
- The command must be run from an account authorized for the operation.
- The **role** command is available in the global configuration mode.
- If the role specified already exists, the **role** command modifies the existing role.

### *Creating or modifying a role*

1. In the privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **role** command with the specified parameters.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# role name VLANAdmin desc "Manages security CLIs"
switch(config)#
```

### *Displaying a role*

In privileged EXEC mode, enter the **show running-config role** command.

```
switch# show running-config role
role name VLANAdmin desc "Manages security CLIs"
role name NetworkAdmin desc "Manages Network CLIs"
role name ClusterAdmin desc "Manages Cluster CLIs"
```

### *Deleting a role*

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **no role** command with the specified parameters.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no role name VLANAdmin
```

## Command access rules

Command authorization is defined in terms of an ordered set of rules that are associated with a role. Rules define and restrict a role to access modes (*read-only* or *read-write* access), and beyond that can define permit or reject on specified command groups or individual commands. You can associate multiple rules with a given user-defined role, but you can only associate one role with any given user account.

To specify a rule, you must specify at least three mandatory attributes: a rule index number, the role to which the rule should apply, and the command that is defined by the rule. [Table 40](#) describes the rule attribute details.

**TABLE 40** Rule attributes

| Parameter | Description                                                                                                                                                                       |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index     | A numeric identifier of the rule in the range between 1 and 512.                                                                                                                  |
| role      | The name of the role for which the rule is defined.                                                                                                                               |
| command   | The command for which access is defined.                                                                                                                                          |
| operation | Optional. Defines the general access mode granted by the rule. Access can be <b>read-only</b> or <b>read-write</b> (default).                                                     |
| action    | Optional. A modifier restricting the general access mode. The specified access is either accepted ( <b>accept</b> ) or rejected ( <b>reject</b> ). The default value is "reject". |

### *Specifying commands with multiple options*

Commands consisting of multiple words indicating command hierarchy are separated by a space, as shown in the following examples.

```
switch(config)# rule 70 action accept operation read-write role NetworkAdmin
command copy running-config

switch(config)# rule 71 action accept operation read-write role NetworkAdmin
command interface management

switch(config)# rule 72 action accept operation read-write role NetworkAdmin
command clear logging
```

#### **NOTE**

Rules cannot be added for commands that are not at the top level of the command hierarchy. For a list of eligible commands, type the help function (?) at the command prompt.

### *Rules for configuration commands*

You can display configuration data for a particular command by using the **show running-config** command. By default, every role can access all the **show running-config** commands. For the non-default roles, even the permission to access the **show running-config** commands can be modified by the authorized user (admin). The user must have the read-write permission for the **configure terminal** command to execute any of the configuration commands.

The following rules govern configuration commands:

- If a role has a rule with a *read-write* operation and the *accept* action for a configuration command, the user associated with this role can execute the command and read the configuration data.
- If a role has a rule with a *read-only* operation and the *accept* action for a configuration command, the user associated with this role can only read the configuration data of the command.
- If a role has a rule with a *read-only* or *read-write* operation and the *reject* action for a configuration command, the user associated with this role cannot execute the command and can read the configuration data of the command.

### *Rules for operational commands*

Rules can be created for the specified operational commands. By default, every role can display all the operational commands but cannot execute them. The show commands can be accessed by all the roles.

The following rules govern operational commands:

- If a role has a rule with a *read-write* operation and the *accept* action for an operational command, the user associated with this role can execute the command.
- If a role has a rule with a *read-only* operation and the *accept* action for an operational command, the user associated with this role can access but cannot execute the command.
- If a role has a rule with a *read-only* or *read-write* operation and the *reject* action for an operational command, the user associated with this role can neither access nor execute the command.

### *Rules for interface key-based commands*

By default, every role has the permission to read the configuration data related to all the instances of the interfaces using the **show running-config interface** *interface\_name* *rbridge-id/slot/port* command.

Rules can be created for a specific instance of the interface-related configuration commands.

The following rules govern interface key-based commands:

- If a role has a rule with a *read-write* operation and the *accept* action for only a particular instance of the interface, the user associated with this role can only modify the attributes of that instance.
- If a role has a rule with a *read-only* operation and the *accept* action for only a particular instance of the interface, the user associated with this role can only read (using the **show running-config** command) the data related to that instance of the interface.
- If a role has a rule with a *read-write* operation and the *reject* action for only a particular instance of the interface, the user associated with this role cannot execute and read the configuration data for that interface instance.

In the following example, the rules are applicable only to a particular instance of the specified interface.

```
switch(config)# rule 60 action accept operation read-write role NetworkAdmin
 command interface tengigabitethernet 0/4

switch(config)# rule 65 action accept operation read-write role NetworkAdmin
 command interface fcoe 0/4
```



```
switch(config)# rule 68 role NetworkAdmin action reject command interface
fortygigabitethernet 1/2/4
```

- If a role has a rule with a *read-only* or *read-write* operation and the *reject* action for an interface or an instance of the interface, the user associated with this role cannot perform clear and show operations related to those interfaces or interface instances. To perform clear and show operations, the user's role must have at least *read-only* and the *accept* permission. By default, every role has the *read-only*, *accept* permission for all interface instances.

In the following example, the user associated with the NetworkAdmin role cannot perform clear and show operations related to all tengigabitethernet instances.

```
switch(config)# rule 30 action accept operation read-write role NetworkAdmin
command interface tengigabitethernet
```

- If a role has a rule with *read-only* or *read-write* operation, and the *reject* action for an interface tengigabitethernet and fcoe instances, the user associated with this role cannot perform clear and show operations related to those instances. To perform clear and show operations related to interface tengigabitethernet and fcoe instances, the user's role should have at least *read-only* and the *accept* permission. By default, every role has the *read-only*, *accept* permission for all interface instances.

In the following example, the user associated with the NetworkAdmin role cannot perform some of the clear and show operations related to all tengigabitethernet instances.

```
switch(config)# rule 30 role NetworkAdmin action reject command interface
tengigabitethernet
```

- A rule created with the **no-operation** command does not enforce any authorization rules. Instead, the **no-operation** instance can be considered as a placeholder for a valid command that will be added later.

For example:

```
switch(config)# rule 75 action reject operation read-write role NetworkAdmin
command no-operation
switch(config)# rule 75 command firmware
```

- The **dot1x** option under the interface instance submode can only be configured if the role has the *read-write* and *accept* permissions for both the **dot1x** command and interface te instances.

In the following example, the user associated with the CfgAdmin role can access and execute the **dot1x** command in the specified tengigabitethernet instance

```
switch(config)# rule 16 action accept operation read-write role cfgadmin
command interface tengigabitethernet
switch(config)# rule 17 action accept operation read-write role cfgadmin
command dot1x
```

- To execute the **no vlan**, and **no spanning-tree** commands under the submode of **interface tengigabitethernet** instances, a user must have *read-write* and *accept* permission for both the **vlan** and the **protocol spanning-tree** commands. If a user has *read-write* and *accept* permission the **vlan**, and **spanning-tree** commands and *read-write* and *accept* permission for at least one interface instance, the user can perform the **no vlan**, and **no spanning-tree** operations on the other interface instances for which the user has only default permissions (*read-only* and *accept*).

### *Configuring a placeholder rule*

A rule created with the **no-operation** command does not enforce any authorization rules. Instead, you can use the **no-operation** instance as a placeholder for a valid command that is added later, as shown in the following example.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **rule** command with the specified parameters and the no-operation placeholder
3. Enter the **rule** command with the specified command to replace the placeholder.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# rule 75 action reject operation read-write role NetworkAdmin
 command no-operation
switch(config)# rule 75 command firmware
```

### *Rule Processing*

When a user executes a command, rules are searched in ascending order by index for a match and the action of the first matching rule is applied. If none of the rules match, command execution is blocked. If there are conflicting permissions for a role in different indices, the rule with lowest index number is applied.

The following exceptions apply:

- When a match is found for a rule with the *read-only* operation, and the *accept* action, the system seeks to determine whether there are any rules with the *read-write* operation and the *accept* action. If such rules are found, the rule with the *read-write* permission is applied.

In the following example, two rules with action *accept* are present and rule 11 is applied.

```
switch(config)# rule 9 operation read-only action accept role NetworkAdmin
 command aaa
switch(config)# rule 11 operation read-write action accept role NetworkAdmin
 command aaa
```

### *Adding a rule*

You add a rule to a role by entering the **rule** command with appropriate options. Any updates to the authorization rules will not apply to the active sessions of the users. The changes will be applied only when users logout from the current session and login to a new session.

The following example creates the rules that authorize the security administrator role to create and manage user accounts:

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Create a rule specifying read-write access to the global configuration mode.
3. Create a second rule specifying read-write access to the **username** command. Enter the **rule** command with the specified parameters

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# rule 150 action accept operation read-write role SecAdminUser
 command config
```

```
switch(config)# rule 155 action accept operation read-write role SecAdminUser
command username
```

4. After creating the rules, the user of the **SecAdminUser** account can log in to the switch and create or modify the user accounts with the **username** command.

```
switch login: SecAdminUser
Password:*****
Welcome to the ConfD CLI
SecAdminUser connected from 127.0.0.1 using console on switch

switch# configure terminal
Entering configuration mode terminal
Current configuration users:
admin console (cli from 127.0.0.1) on since 2010-08-16 18:35:05 terminal mode
switch(config)# username testuser role user password (<string>): *****
```

### *Changing a rule*

The following example changes the previously created rule (index number 155) so that the command “username” is replaced by “role”.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the rule command specifying an existing rule (index 155) and changing the **command** attribute to the **role** command.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# rule 150
switch(config)# rule 155 command role
```

After changing the rule 155, SecAdminUser can log in to the switch and execute the **role** command and not the **username** command.

```
switch login: SecAdminUser
Password:
Welcome to the ConfD CLI
SecAdminUser connected from 127.0.0.1 using console on sw0
switch# configure terminal
Entering configuration mode terminal
Current configuration users:
admin console (cli from 127.0.0.1) on since 2010-08-16 18:35:05 terminal mode
switch(config)# role name NetworkAdmin
```

### *Deleting a rule*

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **no rule** command followed by the index number of the rule you wish to delete.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no rule 155
```

After you deleted rule 155, the SecAdminUser can no longer access the **role** command.

*Displaying a rule*

Enter the **show running-config rule** command in privileged EXEC mode to display all configured rules. You can filter the output by using the command with additional parameters.

```
switch# show running-config rule
rule 30 action accept operation read-write role NetworkSecurityAdmin
rule 30 command role

rule 31 action accept operation read-write role NetworkSecurityAdmin
rule 31 command rule

rule 32 action accept operation read-write role NetworkSecurityAdmin
rule 32 command username

rule 33 action accept operation read-write role NetworkSecurityAdmin
rule 33 command aaa

rule 34 action accept operation read-write role NetworkSecurityAdmin
rule 34 command radius-server

rule 35 action accept operation read-write role NetworkSecurityAdmin
rule 35 command configure

rule 40 action accept operation read-write role FCOEAdmin
rule 40 command "interface fcoe"
```

**Configuration examples**

The following configuration examples illustrate the step-by-step configuration of two frequently used administrative accounts: Brocade VCS Fabric security administrator, and FCoE Fabric administrator.

*Configuring a Brocade VCS Fabric security administrator account*

1. Create a role for a Brocade VCS Fabric security administrator.

```
switch(config)# role name NetworkSecurityAdmin desc "Manages security CLIs"
```

2. Create a user account associated with the newly created role.

```
switch(config)# username SecAdminUser role NetworkSecurityAdmin password
testpassword
```

3. Create the rules to specify the RBAC permissions for the NetworkSecurityAdmin role.

```
switch(config)# rule 30 action accept operation read-write role
NetworkSecurityAdmin command role
switch(config-rule-30)# exit
switch(config)# rule 31 action accept operation read-write role
NetworkSecurityAdmin command rule
switch(config-rule-31)# exit
switch(config)# rule 32 action accept operation read-write role
NetworkSecurityAdmin command username
switch(config-rule-32)# exit
switch(config)# rule 33 action accept operation read-write role
NetworkSecurityAdmin command aaa
switch(config-rule-33)# exit
switch(config)# rule 34 action accept operation read-write role
NetworkSecurityAdmin command radius-server
switch(config-rule-34)# exit
```

```
switch(config)# rule 35 action accept operation read-write role
NetworkSecurityAdmin command config
switch(config-rule-35)# exit
```

The SecAdminUser account has been granted operational access to the configuration-level commands **role**, **rule**, **username**, **aaa**, and **radius-server**. Any account associated with the NetworkSecurityAdmin role can now create and modify user accounts, manage roles, and define rules. In addition, the role permits configuring a RADIUS server and set the login sequence.

### *Configuring a Brocade FCoE administrator account*

1. Create an FCoE administrator role.

```
switch(config)# role name FCOEAdmin desc "Manages FCOE"
```

2. Create an FCoE admin user account.

```
switch(config)# username FCOEAdminUser role FCOEAdmin password testpassword
```

3. Create the rules defining the access permissions for the FCoE administrator role.

```
switch(config)# rule 40 action accept operation read-write role FCOEAdmin
command interface fcoe
```

The FCOEAdminUser account that is associated with the FCOEAdmin role can now perform the FCoE operations.

## Password policies

Password policies define and enforce a set of rules that make passwords more secure by subjecting all new passwords to global restrictions. The password policies described in this section apply to the switch-local user database only. Configured password policies (and all user account attribute and password state information) are synchronized across management modules and remain unchanged after an HA failover.

In Logical chassis cluster mode, the configuration is applied to all the nodes in the cluster.

The following is a list of the configurable password policies:

- [Password strength policy](#)
- [Password encryption policy](#)
- [Account lockout policy](#)

## Password strength policy

Table 41 lists configurable password policy parameters.

**TABLE 41 Password policy parameters**

| Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| character-restriction lower        | Specifies the minimum number of lowercase alphabetic characters that must occur in the password. The maximum value must be less than or equal to the minimum length value. The default value is zero, which means there is no restriction of lowercase characters.                                                                                                                                                                                                                                                                                                                                                                                        |
| character-restriction upper        | Specifies the minimum number of uppercase alphabetic characters that must occur in the password. The maximum value must be less than or equal to the Minimum Length value. The default value is zero, which means there is no restriction of uppercase characters.                                                                                                                                                                                                                                                                                                                                                                                        |
| character-restriction numeric      | Specifies the minimum number of numeric characters that must occur in the password. The maximum value must be less than or equal to the Minimum Length value. The default value is zero, which means there is no restriction of numeric characters.                                                                                                                                                                                                                                                                                                                                                                                                       |
| character-restriction special-char | Specifies the minimum number of punctuation characters that must occur in the password. All printable, non-alphanumeric punctuation characters except the colon (:) are allowed. The value must be less than or equal to the Minimum Length value. The default value is zero, which means there is no restriction of punctuation characters. Characters added after an exclamation point are dropped. For example, if you use the password "first!second", the password will become "first!" Special characters, such as backslash (\) and question mark (?), are not counted as characters in a password unless the password is specified within quotes. |
| min-length                         | Specifies the minimum length of the password. Passwords must be from 8 through 32 characters in length. The default value is 8. The total of the previous four parameters (lowercase, uppercase, digits, and punctuation) must be less than or equal to the Minimum Length value.                                                                                                                                                                                                                                                                                                                                                                         |
| max-retry                          | Specifies the number of failed password logins permitted before a user is locked out. The lockout threshold can range from 0 through 16. The default value is 0. When a password fails more than one of the strength attributes, an error is reported for only one of the attributes at a time.                                                                                                                                                                                                                                                                                                                                                           |

### NOTE

Passwords can be a maximum of 40 characters in length.

## Password encryption policy

Network OS supports encrypting the passwords of all existing user accounts by enabling password encryption at the switch level. By default, the encryption service is disabled and passwords are stored in clear-text. Use the **no service password-encryption** command to enable or disable password encryption. The following rules apply to password encryption:

- When you enable password encryption, all existing clear-text passwords will be encrypted, and any password that are added subsequently in clear-text will be stored in encrypted format

In the following example, the testuser account password is created in clear-text after password encryption has been enabled. The global encryption policy overrides command-level encryption settings. The password is stored as encrypted.

```
switch(config)# service password-encryption
switch(config)# do show running-config service password-encryption
service password-encryption
switch(config)# username testuser role testrole desc "Test User"
```

```

encryption-level 0 password hellothere
switch(config)# do show running-config username
username admin password "BwrsDbB+tABWGWpINOVKoQ==\n" encryptionlevel
7 role admin desc Administrator
username testuser password "cONWlRQ0nTV9Az42/9uCQg==\n"
encryption-level 7 role testrole desc "Test User"
username user password "BwrsDbB+tABWGWpINOVKoQ==\n" encryptionlevel
7 role user desc User

```

- When you disable the password encryption service, any new passwords added in clear-text will be stored as clear-text on the switch. Existing encrypted passwords remain encrypted.

In the following example, the testuser account password is stored in clear-text after password encryption has been disabled. The default accounts, “user” and admin” remain encrypted.

```

switch(config)# no service password-encryption
switch(config)# do show running-config service password-encryption
no service password-encryption
switch(config)# username testuser role testrole desc "Test User"
encryption-level 0 password hellothere enable true.
switch(config)# do show running-config username
username admin password "BwrsDbB+tABWGWpINOVKoQ==\n" encryptionlevel
7 role admin desc Administrator
username testuser password hellothere encryption-level 0 role
testrole desc "Test User"
username user password "BwrsDbB+tABWGWpINOVKoQ==\n" encryptionlevel
7 role user desc User

```

## Account lockout policy

The account lockout policy disables a user account when the user exceeds a configurable number of failed login attempts. A user whose account has been locked cannot log in. SSH login attempts using locked user credentials are denied without notifying the user of the reason for denial.

The account remains locked until explicit administrative action is taken to unlock the account. A user account cannot be locked manually. An account not locked cannot be unlocked.

Failed login attempts are tracked on the local switch only. in VCS mode, the user account is locked only on the switch where the lockout occurred; the same user can still try to log in on another switch in the VCS fabric.

The account lockout policy is enforced across all user accounts except for the root account and accounts with the admin role.

## *Denial of service implications*

The account lockout mechanism may be used to create a denial of service condition by repeatedly attempting to login to an account using an incorrect password. Selected privileged accounts, such as root and admin are exempted from the account lockout policy to prevent them from being locked out by a denial of service attack. However these privileged accounts may then become the target of password guessing attacks. Brocade advises that you periodically examine the Security Audit logs to determine if such attacks are attempted. For information on security audit logging refer to the *Network OS Message Reference*.

### *Configuring the account lockout threshold*

You can configure the lockout threshold with the **password-attributes max-retry** *maxretry* command. The value of the *maxretry* specifies the number of times a user can attempt to log in with an incorrect password before the account is locked. The number of failed login attempts is counted from the last successful login. The *maxretry* can be set to a value from 0 through 16. A value of 0 disables the lockout mechanism (default).

The following example sets the lockout threshold to 5.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **password-attributes** command with the specified parameter.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# password-attributes max-retry 4
```

When a user account is locked, it can be unlocked using the procedure described in [“Unlocking a user account”](#) on page 219.

## Password interaction with remote AAA servers

The password policies apply to local switch authentication only. External AAA servers such as RADIUS, TACACS+, or LDAP provide server-specific password-enforcement mechanisms. The Network OS password management commands operate on the switch-local password database only, even when the switch is configured to use an external AAA service for authentication. When so configured, authentication through remote servers is applied to login only.

When remote AAA server authentication is enabled, an administrator can still perform user and password management functions on the local password database.

For more information on remote AAA server authentication, refer to [Chapter 17, “External server authentication”](#).

## Managing password policies

Use the **password-attributes** command with specified parameters to define or modify existing password policies.

### *Creating a password policy*

The following example defines a password policy that places restrictions on minimum length and enforces character restrictions and account lockout.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **password-attributes** command with the specified parameters.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# password-attributes min-length 8 max-retry 4
character-restriction lower 2 upper 1 numeric 1 special-char 1
```



### *Restoring the default password policy*

The **no** form of the **password-attributes** command resets all password attributes to their default values. When used without operands, the command resets all password attributes.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **password-attributes** command with the specified parameters.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no password-attributes min-length
switch(config)# password-attributes max-retry 4
switch(config)# no password-attributes numeric
```

### *Displaying password attributes*

In privileged EXEC mode, enter the **show running-config password-attributes** command to display configured password attributes:

```
switch(config)# password-attributes max-retry 4
switch(config)# password-attributes character-restriction lower 2
switch(config)# password-attributes character-restriction upper 1 numeric 1
special-char 1
switch(config)# exit
switch# show running-config password-attributes
password-attributes max-retry 4
password-attributes character-restriction upper 1
password-attributes character-restriction lower 2
password-attributes character-restriction numeric 1
password-attributes character-restriction special-char 1
switch# configure terminal
switch(config)# no password-attributes character-restriction lower
switch(config)# no password-attributes character-restriction upper
switch(config)# exit
switch# show running-config password-attributes
password-attributes max-retry 4
password-attributes character-restriction numeric 1
password-attributes character-restriction special-char 1
switch# configure terminal
switch(config)# no password-attributes special-char
switch(config)# exit
switch# show running-config password-attributes
% No entries found.
```

## Security event logging

Security event logging utilizes the RASLog audit infrastructure to record security-related audit events. Any user-initiated security event generates an auditable event. Audited events are generated for all Management interfaces. In Brocade VCS Fabric mode, for cluster-wide events, the audit is generated on all switches of the cluster. Refer to the *Network OS Message Reference* for information on how to configure and monitor security audit logging.



# External server authentication

---

## In this chapter

|                                               |     |
|-----------------------------------------------|-----|
| • Remote server authentication overview ..... | 235 |
| • Login authentication mode .....             | 236 |
| • RADIUS .....                                | 238 |
| • TACACS+ .....                               | 244 |
| • TACACS+ accounting .....                    | 247 |
| • TACACS+ server-side configuration .....     | 250 |
| • LDAP .....                                  | 253 |

## Remote server authentication overview

Network OS supports various protocols to provide external Authentication, Authorization, and Accounting (AAA) services for Brocade devices. Supported protocols include the following:

- RADIUS — Remote authentication dial-in user service
- LDAP/AD — Lightweight directory access protocol using Microsoft Active Directory (AD) in Windows
- TACACS+ — Terminal access controller access-control system plus

When configured to use a remote AAA service, the switch acts as a network access server (NAS) client. The switch sends all authentication, authorization, and accounting (AAA) service requests to the remote RADIUS, LDAP, or TACACS+ server. The remote AAA server receives the request, validates the request, and sends a response back to the switch.

The supported management access channels that integrate with RADIUS, TACACS+, or LDAP include serial port, Telnet, or SSH.

When configured to use a remote RADIUS, TACACS+, or LDAP server for authentication, a switch becomes a RADIUS, TACACS+, or LDAP client. In either of these configurations, authentication records are stored in the remote host server database. Login and logout account name, assigned permissions, and time-accounting records are also stored on the AAA server for each user.

Brocade recommends that you configure at least two remote AAA servers to provide redundancy in the event of failure. For each of the supported AAA protocols, you can configure up to five external servers on the switch. Each switch maintains its own server configuration.

## Login authentication mode

The authentication mode is defined as the order in which AAA services are used on the switch for user authentication during the login process. Network OS supports two sources of authentication: primary and secondary. The secondary source of authentication is used in the event of primary source failover and is optional for configuration. You can configure four possible sources for authentication:

- Local — Use the default switch-local database (default)
- RADIUS — Use an external RADIUS server
- LDAP — Use an external LDAP server
- TACACS+ — Use an external TACACS+ server

By default, external AAA services are disabled, and AAA services default to the switch-local user database. Any environment requiring more than 64 users should adopt AAA servers for user management.

When the authentication, authorization, and accounting (AAA) mode is changed, an appropriate message is broadcast to all logged-in users, and the active login sessions end. If the primary source is set to an external AAA service (RADIUS, LDAP, or TACACS+) and the secondary source is not configured, the following events occur:

- For Telnet-based and SSH connections-based logins, the login authentication fails if none of the configured (primary source) AAA servers respond or if an AAA server rejects the login.
- For a serial port (console) connection-based login, if a user's login fails for any reason with the primary source, failover occurs and the same user credentials are used for login through the local source. This failover is not explicit.
- If the primary source is set to an external AAA service, and the secondary source is configured to be local (for example, **aaa authentication login radius local**), then, if login fails through the primary source either because none of the configured servers is responding or the login is rejected by a server, failover occurs and authentication occurs again through the secondary source (local) for releases earlier than Network OS 4.0.

When local is specified as the secondary authentication service, failover to local does not occur if login is rejected by a server. In addition, when the authentication service is changed, the user sessions are not logged out. If a user wants to log out all connected user sessions, the **clear sessions** command should be used.

- With the introduction of Network OS 4.0, when local is specified as the secondary authentication service, local authentication is tried only when the primary AAA authentication service (TACACS+/Radius/LDAP) is either unreachable or not available. Local authentication will not be attempted if authentication with the primary service fails.

With the introduction of Network OS 4.0, you can specify to use the local switch database if prior authentication methods on a RADIUS or TACACS+ server are not active or if authentication fails. To specify this option, use the **local-auth-fallback** command. In the following example, the local switch database will be used if the RADIUS server is unavailable.

```
switch(config)# aaa authentication login radius local-auth-fallback
```

## Conditions for conformance

- If the first source is specified as *default*, do not specify a second source. A second source signals a request to set the login authentication mode to its default value, which is *local*. If the first source is *local*, the second source cannot be set to any value, because the failover will never occur.
- The source of authentication (except *local*) and the corresponding server type configuration are dependent on each other. Therefore, at least one server should be configured before that server type can be specified as a source.
- If the source is configured to be a server type, you cannot delete a server of that type if it is the only server in the list. For example, if there are no entries in the TACACS+ server list, the authentication mode cannot be set to *tacacs+* or *tacacs+ local*. Similarly, when the authentication mode is *radius* or *radius local*, a RADIUS server cannot be deleted if it is the only one in the list.

## Setting and verifying the login authentication mode

The following procedure configures TACACS+ as the primary source of authentication and the switch-local user database as the secondary source.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter **aaa authentication login** command with the specified parameters.  

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# aaa authentication login tacacs+ local
Broadcast message from root (pts/0) Tue Apr 5 16:34:12 2011...
AAA Server Configuration Change: all accounts will be logged out
```
3. Enter the **do show running-config aaa** command to display the configuration.  

```
switch(config)# do show running-config aaa
aaa authentication login tacacs+ local
```
4. Log in to the switch using an account with TACACS+ only credentials to verify that TACACS+ is being used to authenticate the user.

## Resetting the login authentication mode

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **no aaa authentication login** command to remove the configured authentication sequence and to restore the default value (Local only).  

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no aaa authentication login
```
3. Verify the configuration with the **do show running-config aaa** command.  

```
switch(config)# do show running-config aaa
aaa authentication login local
```
4. Log in to the switch using an account with TACACS+ only credentials. The login should fail with an “access denied” error.

5. Log in to the switch using an account with local only credentials. The login should succeed.

### *Changing the login authentication mode*

You can set the authentication mode with the **aaa authentication login** command, but you cannot change or delete an existing authentication mode with the same command. You can only reset the configuration to the default value using the **no aaa authentication login** command and then reconfigure the authentication sequence to the correct value.

1. Enter the **no aaa authentication login** command to reset the configuration to the default value
2. Enter the **aaa authentication login** command and specify the desired authentication mode.
3. Verify the configuration with the **do show running-config aaa** command.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no aaa authentication login tacacs+ local
switch(config)# aaa authentication login radius local
Broadcast message from root (pts/0) Tue Apr 5 16:34:12 2011...

AAA Server Configuration Change: all accounts will be logged out
switch(config)# do show running-config aaa
aaa authentication login radius local
```

4. Log in to the switch using an account with TACACS+ credentials. The login should fail with an “access denied” error.
5. Log in to the switch using an account with RADIUS credentials. The login should succeed.

## RADIUS

The remote authentication dial-in user service (RADIUS) protocol manages authentication, authorization, and accounting (AAA) services centrally. The supported management access channels that integrate with RADIUS are serial port, Telnet, and SSH.

If you are in Logical chassis cluster mode, the configuration is applied to all nodes in the cluster.

### Authentication and accounting

When a Brocade switch is configured with a set of RADIUS servers to be used for authentication, the switch also sends accounting data to the RADIUS server implicitly. The only accounting events supported on Brocade VDX switches configured to use RADIUS are successful login and logout of the RADIUS user.

During the user authentication process, the switch sends its IP address. When the switch also has a Virtual IP address (in Brocade VCS Fabric mode), it still sends only its unique IP address to the RADIUS server.

---

#### NOTE

If the RADIUS server is not configured to support accounting, the accounting events sent by the switch to the server are dropped.

---

## Authorization

User authorization through the RADIUS protocol is not supported. The access control of RADIUS users is enforced by the Brocade role-based access control (RBAC) protocol at the switch level. A RADIUS user should therefore be assigned a role that is present on the switch using the Vendor Specific Attribute (VSA) “*Brocade-Auth-Role*.” After the successful authentication of the RADIUS user, the role of the user configured on the server is obtained. If the role cannot be obtained or if the obtained role is not present on the switch, the user will assigned “user” role and a session is granted to the user with “user” authorization.

## Account password changes

All existing mechanisms for managing switch-local user accounts and passwords remain functional when the switch is configured to use RADIUS. Changes made to the switch-local database do not propagate to the RADIUS server, nor do the changes affect any account on the RADIUS server; therefore, changes to a RADIUS user password must be done on the RADIUS server.

## RADIUS authentication through management interfaces

You can access the switch through Telnet or SSH from either the Management interface or the data ports (TE interface or in-band). The switch goes through the same RADIUS-based authentication with either access method.

## Client-side RADIUS server configuration

Each Brocade switch client must be individually configured to use RADIUS servers. You use the **radius-server** command to specify the server IP address, authentication protocols, and other parameters. You can configure a maximum of five RADIUS servers on a Brocade switch for AAA service.

The parameters in [Table 42](#) are associated with a RADIUS server that is configured on the switch.

**TABLE 42 RADIUS server parameters**

| Parameter | Description                                                                                                                                                                                                      |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| host      | IP address (IPv4 or IPv6) or host name of the RADIUS server. Host name requires prior DNS configuration. The maximum supported length for the host name is 40 characters.                                        |
| auth-port | The user datagram protocol (UDP) port used to connect the RADIUS server for authentication. The port range is 0 through 65535; the default port is 1812.                                                         |
| protocol  | The authentication protocol to be used. Options include CHAP, PAP, and PEAP. The default protocol is CHAP. IPv6 hosts are not supported if PEAP is the configured protocol.                                      |
| key       | The shared secret between the switch and the RADIUS server. The default value is “sharedsecret.” The key cannot contain spaces and must be from 8 through 40 characters in length. Empty keys are not supported. |
| retries   | The number of attempts permitted to connect to a RADIUS server. The range is 0 through 100, and the default value is 5.                                                                                          |
| timeout   | Time to wait for a server to respond. The range is 1 through 60 seconds. The default value is 5 seconds.                                                                                                         |

**NOTE**

If you do not configure the *key* attribute, the authentication session will not be encrypted. The value of the *key* attribute must match the value configured in the RADIUS configuration file; otherwise, the communication between the server and the switch fails.

***Adding a RADIUS server to the client server list***

You must configure the Domain Name System (DNS) server on the switch prior to adding the RADIUS server with a domain name or a host name. Without the DNS server, name resolution of the RADIUS server fails and therefore the add operation fails. Use the **ip dns** command to configure the DNS server.

**NOTE**

When a list of servers is configured on the switch, failover from one server to another server happens only if a RADIUS server fails to respond; it does not happen when user authentication fails.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter **radius-server** command with the specified parameters.  
Upon execution of the command you are placed into the AAA server configuration submode where you can specify additional parameters.
3. Enter the **exit** command to return to global configuration mode.
4. Enter the **do show running-config radius-server host** command to verify the configuration.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# radius-server host 10.38.37.180 protocol pap key
"new#virgo*secret" timeout 10
switch(config-host-10.38.37.180)# exit
switch# show running-config radius-server host 10.38.37.180
radius-server host 10.38.37.180
protocol pap
key "new# virgo*secret"
timeout 10
```

***Modifying the RADIUS server configuration***

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter **radius-server host** command with the help option (?) to display the configured RADIUS servers.
3. Enter the **radius-server host** command with the IP address of the server you want to modify.  
Upon execution of the command you are placed into the radius-server configuration sub-mode where you can specify the parameters you want to modify.
4. Enter the parameters and values you want to change.
5. Enter the **do show running-config radius-server** command to verify the configuration.

This command does not display default values.

```
switch# configure terminal
Entering configuration mode terminal
```



```

switch(config)# radius-server ?
Possible completions:
<hostname: IP Address or Hostname of this RADIUS server>
10.38.37.180
10.24.65.6
switch(config)# radius-server host 10.38.37.180
switch(config-host-10.38.37.180)# key "changedsec"
switch(config-host-10.38.37.180)# timeout 3
switch(config)# do show running-config radius-server host 10.24.65.6
radius-server host 10.24.65.6
protocol pap
key changedsec
timeout 3

```

The **no radius-server host** command removes the server configuration from the list of configured RADIUS servers. When used with a specified parameter, the command sets the default value of that parameter.

### *Configuring the client to use RADIUS for login authentication*

After you configured the client-side RADIUS server list, you must set the authentication mode so that RADIUS is used as the primary source of authentication. refer to the section “[Login authentication mode](#)” on page 236 for information on how to configure the login authentication mode.

## Server-side RADIUS configuration

With RADIUS servers, you should set up user accounts by their true network-wide identity, rather than by the account names created on a Brocade switch. Along with each account name, you must assign appropriate switch access roles. A user account can exist on a RADIUS server with the same name as a user on the switch at the same time.

When logging in to a switch configured with RADIUS, users enter their assigned RADIUS account names and passwords when prompted. Once the RADIUS server authenticates a user, it responds with the assigned switch role and information associated with the user account information using a Brocade Vendor-Specific Attribute (VSA). An Authentication-Accept response without the role assignment automatically grants the “user” role.

### *Configuring a RADIUS server with Linux*

FreeRADIUS is a open source RADIUS server that runs on Linux (all versions), FreeBSD, NetBSD, and Solaris. Download the package from [www.freeradius.org](http://www.freeradius.org) and follow the installation instructions at the FreeRADIUS website.

You will need the following information to configure Brocade-specific attributes. Refer to the RADIUS product documentation for information on configuring and starting up a RADIUS server.

### *Adding the Brocade attribute to the RADIUS server configuration*

For the configuration on a Linux FreeRadius server, define the values outlined in [Table 43](#) in a vendor dictionary file named `dictionary.brocade`.

**TABLE 43** `dictionary.brocade` file entries

| Include   | Key               | Value            |
|-----------|-------------------|------------------|
| VENDOR    | Brocade           | 1588             |
| ATTRIBUTE | Brocade-Auth-Role | 1 string Brocade |

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.
```

2. Open the master dictionary 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 master configuration directory and loaded for use by the RADIUS server.

### *Configuring a Brocade user account*

1. Open the `$PREFIX/etc/raddb/users` file in a text editor.
2. Add the user name and associated the permissions.

The user will log in using the permissions specified with `Brocade-Auth-Role`. The valid permissions include “user” and “admin”. You must use double quotation marks around the password and role.

The following example configures an account called “jsmith” with admin permissions and a password “jspassword”.

```
jsmith Auth-Type := Local,
 User-Password == "jspassword",
 Brocade-Auth-Role = "admin"
```

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 setting the **pap** option with the **radius-server host** command.

### *Configuring RADIUS server support with Windows server*

Step-by-step instructions for installing and configuring Internet Authentication Service (IAS) with Microsoft Windows server 2008 (or earlier versions, Windows 2003 or 2000) can be obtained from [www.microsoft.com](http://www.microsoft.com) or your Microsoft documentation. Confer with your system or network administrator prior to configuration for any special needs your network environment may have.

Use the following information to configure the Internet Authentication Service for a Brocade switch. This is not a complete presentation of steps.

- In the New RADIUS Client window choose **RADIUS Standard** from the Client-Vendor menu.
- Configure the Dial-in Profile window as follows:
  1. Select the **Advanced** tab.
  2. Scroll to the bottom of the RADIUS Standard list, select **Vendor-Specific**, and click **Add**.  
The Multivalued Attribute Information dialog appears.
  3. Click **Add** in the Multivalued Attribute Information window.  
The Vendor-Specific Attribute Information dialog appears.
  4. Enter the Brocade vendor code value of **1588**.
  5. Select the **Yes. It conforms** radio button and then click **Configure Attribute**.  
The Configure VSA (RFC compliant) dialog appears.
  6. In the Configure VSA (RFC compliant) window, enter the following values and click **OK**.  
Vendor-assigned attribute number—Enter the value **1**.  
Attribute format—Enter **String**.

Attribute value—Enter the **login** role (admin or user).

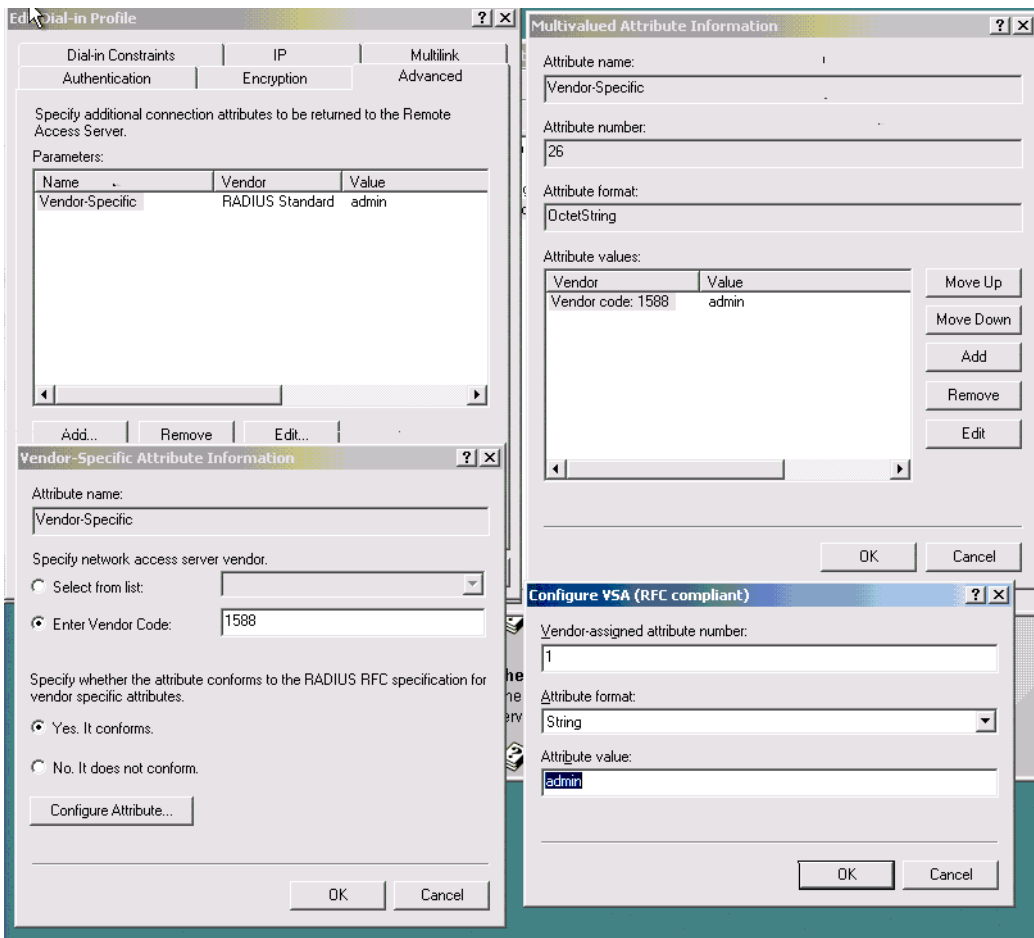


FIGURE 27 Windows server VSA configuration

# TACACS+

The Terminal Access Controller Access-Control System Plus (TACACS+) is an AAA server protocol that uses a centralized authentication server and multiple Network Access Servers (NASs) or clients. With TACACS+ support, management of Brocade switches seamlessly integrates into these environments. Once configured to use TACACS+, a Brocade switch becomes a Network Access Server (NAS).

If you are in Logical chassis cluster mode, the configuration is applied to all nodes in the cluster.

## Authorization

The TACACS+ server is used only for authentication and accounting. Authorization is enforced by the Brocade role-based access control (RBAC) protocol at the switch level. The same role should be assigned to a user configured on the TACACS+ server and configured on the switch. If the switch fails to get the user's role from the TACACS+ server after successful authentication, or if the role does not match any of the roles present on the switch, the *user* role is assigned by default. Thereafter, the role obtained from the TACACS+ server (or the defaulted role) is used for RBAC.

## TACACS+ authentication through management interfaces

You can access the switch through the serial port, or through Telnet or SSH from either the management interface or the data ports (TE interface or in-band). The switch goes through the same TACACS+-based authentication with either access method.

## Supported packages and protocols

Brocade supports the following TACACS+ packages for running the TACACS+ daemon on remote AAA servers.

- Free TACACS+ daemon (tacacs-plus 4.0.4.23-3). You can download this package from [http://www.shrubbery.net/tac\\_plus/](http://www.shrubbery.net/tac_plus/)
- ACS 5.3
- ACS 4.2

The TACACS+ protocol v1.78 is used for AAA services between the Brocade switch client and the TACACS+ server.

The authentication protocols supported for user authentication are Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP).

## Client-side TACACS+ server configuration

Each Brocade switch client must be individually configured to use TACACS+ servers. You use the **tacacs-server** command to specify the server IP address, authentication protocols, and other parameters. You can configure a maximum of five TACACS+ servers on a Brocade switch for AAA service.

The parameters in [Table 44](#) are associated with a TACACS+ server that is configured on the switch.

**TABLE 44 TACACS+ server parameters**

| Parameter | Description                                                                                                                                                                                                      |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| host      | IP address (IPv4 or IPv6) or domain/host name of the TACACS+ server. Host name requires prior DNS configuration. The maximum supported length for the host name is 40 characters.                                |
| port      | The TCP port used to connect the TACACS+ server for authentication. The port range is 1 through 65535; the default port is 49.                                                                                   |
| protocol  | The authentication protocol to be used. Options include CHAP and PAP. The default protocol is CHAP.                                                                                                              |
| key       | The shared secret between the switch and the RADIUS server. The default value is "sharedsecret." The key cannot contain spaces and must be from 8 through 40 characters in length. Empty keys are not supported. |

**TABLE 44** TACACS+ server parameters (Continued)

| Parameter | Description                                                                                                                            |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------|
| retries   | The number of attempts permitted to connect to a RADIUS server. The range is 0 through 100, and the default value is 5.                |
| timeout   | The maximum amount of time to wait for a server to respond. Options are from 1 through 60 seconds, and the default value is 5 seconds. |

**NOTE**

If you do not configure the *key* attribute, the authentication session will not be encrypted. The value of *key* must match with the value configured in the TACACS+ configuration file; otherwise, the communication between the server and the switch fails.

***Adding a TACACS+ server to the client server list***

You must configure the Domain Name System (DNS) server on the switch prior to adding the TACACS+ server with a domain name or a host name. Without the DNS server, name resolution of the TACACS+ server fails and therefore the add operation fails. Use the **ip dns** command to configure the DNS server.

**NOTE**

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.

The following procedure adds a TACACS+ server host in IPv6 format.

1. In the privileged EXEC mode, use the **configure terminal** command to enter the global configuration mode.
2. Enter **tacacs-server** command and specify the server IP address.  
  
Upon execution of the command you are placed into the tacacs-server configuration sub-mode where you can specify additional parameters.
3. Enter the **do show running-config tacacs-server host** command to verify the configuration.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# tacacs-server host fec0:60:69bc:94:211:25ff:fec4:6010
switch(config-tacacs-server-fec0:60:69bc:94:211:25ff:fec4:6010)# protocol \
chap key "new#hercules*secret"
switch(config-tacacs-server-fec0:60:69bc:94:211:25ff:fec4:6010)# exit
switch(config)# do show running-config tacacs-server \
fec0:60:69bc:94:211:25ff:fec4:6010
tacacs-server host fec0:60:69bc:94:211:25ff:fec4:6010
key new# Hercules*secret
```

***Modifying the TACACS+ server configuration***

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter **tacacs-server host** command with the help option (?) to display the configured server IP addresses.
3. Enter **tacacs-server host** command with the IP address of the server you wish to modify.

Upon execution of the command you are placed into the tacacs-server configuration sub-mode where you can specify the parameters you want to modify.

4. Enter the **exit** command to return to the global configuration mode.
5. Enter the **do show running-config tacacs-server host** command to verify the configuration.

This command does not display default values.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# tacacs-server host ?
fec0:60:69bc:94:211:25ff:fec4:6010
switch(config)# tacacs-server host fec0:60:69bc:94:211:25ff:fec4:6010
switch(config-tacacs-server-fec0:60:69bc:94:211:25ff:fec4:6010)# key \
"changedsec" retries 100
switch(config-tacacs-server-fec0:60:69bc:94:211:25ff:fec4:6010)# exit
switch(config)# do show running-config tacacs-server \
fec0:60:69bc:94:211:25ff:fec4:6010
tacacs-server host fec0:60:69bc:94:211:25ff:fec4:6010
key changedesc
retries 100
!
```

The **no tacacs-server host** command removes the server configuration from the list of configured RADIUS servers. If the tacacs-server being deleted is the last one in the list and authentication mode is set to tacacs+, deletion of the server from the switch configuration is denied. When used with a specified parameter, the command sets the default value of that parameter

### *Configuring the client to use TACACS+ for login authentication*

After you configured the client-side TACACS+ server list, you must set the authentication mode so that TACACS+ is used as the primary source of authentication. refer to the section [“Login authentication mode”](#) on page 236 for information on how to configure the login authentication mode.

## TACACS+ accounting

The TACACS+ protocol supports accounting as a function distinctly separate from authentication. You can use TACACS+ for authentication only, for accounting only, or for both. With a TACACS+ server you can track user logins and the commands users execute during a login session by enabling login accounting, command accounting, or both.

If you are in logical chassis cluster mode, the configuration is applied to all nodes in the cluster.

- When *login accounting* is enabled, the switch sends a TACACS+ start accounting packet with relevant attributes to the configured TACACS+ server when the user logs in, and a stop accounting packet when the session terminates.
- When *command accounting* is enabled, the switch sends a TACACS+ stop accounting packet to the server when the command execution completes. No TACACS+ start accounting packet is sent for command accounting. Most configuration commands, show commands and non-configuration commands such as firmware download will be tracked. Commands received through the NetConf interface will also be tracked. For a listing of commands that are not accounted for, refer to [Appendix A, “TACACS+ Accounting Exceptions”](#)

If a TACACS+ server is used for both authentication and accounting, the switch first attempts to connect to the TACACS+ server that was successfully used for authentication when sending accounting packets to the server. If the TACACS+ server cannot be reached, the switch attempts to send the packets to the next server on the list. Note that there is no fail back in this case. When the first TACACS+ server becomes reachable again, the accounting packets continue to be sent to the second TACACS+ server.

If authentication is performed through some other mechanism, such as the switch-local database, a RADIUS, or an LDAP server, the switch will attempt to send the accounting packets to the first configured TACACS+ server. If that server is unreachable, the switch will attempt to send the accounting packets to subsequent servers in the order in which they are configured.

## Conditions for conformance

- Only login and command accounting is supported. System event accounting is not supported.
- You can use a TACACS+ server for accounting regardless of whether authentication is performed through RADIUS, LDAP, TACACS+, or the switch-local user database. The only precondition is the presence of one or more TACACS+ servers configured on the switch.
- No accounting can be performed if authentication fails.
- In command accounting, commands with partial timestamp cannot be logged. For example, a **firmware download** command issued with the reboot option will not be accounted for, because there is no timestamp available for completion of this command.

## Configuring TACACS+ accounting on the client

By default, accounting is disabled on the TACACS+ client (the switch) and you must explicitly enable the feature. Enabling command accounting and login accounting on the TACACS+ client are two distinct operations. To enable login or command accounting, at least one TACACS+ server must be configured. Similarly, if either login or command accounting is enabled, you cannot remove a TACACS+ server, if it is the only server in the list.

### *Enabling login accounting*

The following procedure enables login accounting on a switch where accounting is disabled.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **aaa accounting default exec start-stop tacacs+** command to enable login accounting.
3. Enter the **do show running-config aaa accounting** command to verify the configuration.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# aaa accounting exec default start-stop tacacs+
switch(config)# do show running-config aaa accounting
aaa accounting exec default start-stop tacacs+
aaa accounting commands default start-stop tacacs+
```



### *Enabling command accounting*

The following procedure enables login accounting on a switch where login accounting is enabled and command accounting is disabled.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **aaa accounting default command start-stop tacacs+** command to enable command accounting.
3. Enter the **do show running-config aaa accounting** command to verify the configuration.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# aaa accounting command default start-stop tacacs+
switch(config)# exit
switch# show running-config aaa accounting
aaa accounting exec default start-stop none
aaa accounting commands default start-stop tacacs+
```

### *Disabling accounting*

You have two options to disable accounting, either by using the **aaa accounting** command, with the none option or by using the **no** form of the command. Both variants are functionally equivalent. You must perform the disable operation separately for login accounting and for command accounting. The operation is performed in the global configuration mode.

The following examples show two ways of disabling command accounting. The commands are executed in the global configuration mode.

```
switch(config)# aaa accounting commands default start-stop none
switch(config)# no aaa accounting commands default start-stop
```

The following examples show two ways of disabling login accounting.

```
switch(config)# aaa accounting exec default start-stop none
switch(config)# no aaa accounting exec default start-stop
```

## Viewing the TACACS+ accounting logs

The following excerpts from TACACS+ accounting logs exemplify typical success and failure cases for command and login accounting. These examples were taken from the free TACACS+ server. The order of the attributes may vary depending on the server package, but the values are the same. The location of the accounting logs depend on the server configuration.

### *Command accounting*

The following record shows the successful execution of the **username** command by the admin user.

```
<102> 2012-04-09 15:21:43 4/9/2012 3:21:43 PM NAS_IP=10.17.37.150 Port=0
rem_addr=Console User=admin Flags=Stop task_id=1 timezone=Etc/GMT+0 service=shell
priv-lvl=0 Cmd=username Stop_time=Mon Apr 9 09:43:56 2012
Status=Succeeded
```

The following record shows the failed execution of the **radius-server** command by the admin user due to an invalid host name or server IP address.

```
<102> 2012-04-09 14:19:42 4/9/2012 2:19:42 PM NAS_IP=10.17.37.150 Port=0
rem_addr=Console User=admin Flags=Stop task_id=1 timezone=Etc/GMT+0 service=shell
priv-lvl=0 Cmd=radius-server Stop_time=Mon Apr 9 08:41:56 2012
Status=% Error: Invalid host name or IP address
```

### *Login (EXEC) accounting*

The following record shows the successful login of the trial user.

```
<102> 2012-05-14 11:47:49 5/14/2012 11:47:49 AM NAS_IP=10.17.46.42
Port=/dev/ttyS0 rem_addr=Console User=trial Flags=Start task_id=1
timezone=Asia/Kolkata service=shell
```

The following record shows the successful logout of the trial user.

```
<102>2012-05-14 11:49:52 5/14/2012 11:49:52 AM NAS_IP=10.17.46.42 Port=/dev/ttyS0
rem_addr=console User=trial Flags=Stop task_id=1 timezone=Asia/Kolkata
service=shell elapsed_time=123 reason=admin reset
```

### **Firmware downgrade considerations**

Before downgrading to a version that does not support TACACS+ accounting, you must disable both login and command accounting or the firmware download will fail with an appropriate error message.

## **TACACS+ server-side configuration**

Step-by-step instructions for installing and configuring can be obtained from [www.cisco.com](http://www.cisco.com). Confer with your system or network administrator prior to configuration for any special needs your network environment may have.

### *User account administration*

With TACACS+ servers, you should set up user accounts by their true network-wide identity, rather than by the account names created on a Brocade switch. Along with each account name, you must assign appropriate switch access roles. A user account can exist on TACACS+ server with the same name as a user on the switch at the same time.

When logging in to a switch configured with a TACACS+ server, users enter their assigned TACACS+ account names and passwords when prompted. Once the TACACS+ server authenticates a user, it responds with the assigned switch role and information associated with the user account information using a Brocade Vendor-Specific Attribute (VSA). An Authentication-Accept response without the role assignment automatically grants the “user” role.

User accounts, protocols passwords, and related settings are configured by editing the server configuration files. The following configuration examples are based on the documentation provided by Cisco for its TACACS+ daemon users.

### *Configuring a user account*

The following example assigns the user “Mary” the Brocade role of “vlanadmin” and different passwords depending on whether the CHAP or the PAP protocol is used. In the following example, the `brcd-role` attribute is mandatory, which works in a Brocade-only environment. In a mixed vendor environment, the `brcd-role` attribute must be set to optional. Refer to [“Configuring TACACS+ for a mixed vendor environment”](#) on page 252 for more information.

```
user = Mary {
 chap = cleartext "chap password"
 pap = cleartext "pap password"
 service = exec {
 brcd-role = vlanadmin;
 }
}
```

The following example assigns the user “Agnes” a single password for all types of login authentication.

```
user = Agnes {
 global = cleartext "Agnes global password"
}
```

Alternatively, a user can be authenticated using the `/etc/passwd` file. Configure the account as shown in the following example.

```
user = fred {
 login = file /etc/passwd
}
```

### *Changing a TACACS+ account password*

The change of password for TACACS+ user is done on the server by editing the TACACS+ server configuration file.

### *Setting an account expiration date*

You can set an expiration date for an account by using the “expires” attribute in the TACACS+ server configuration file. The expiration date has the format “MMM DD YYYY”

```
user = Brocade {
 member = admin
 expires = "Jan 1 2011"
 pap = cleartext "pap password"
}
```

### *TACACS+ Server key*

The TACACS+ Server key is the shared secret used to secure the messages exchanged between the Brocade switch and the TACACS+ server. The TACACS+ Server key must be configured on both the TACACS+ server and the client (Brocade switch). Only one key is defined per server in the TACACS+ server configuration file. The key is defined as follows:

```
key = "vcs shared secret"
```

### *Defining a TACACS+ group*

A TACACS+ group or role can contain the same attributes as the users. By inference, all the attributes of a group can be assigned to any user to whom the group is assigned. The TACACS+ group, while functionally similar to the Brocade role concept, has no relation with the value of “brcd-role” attribute.

The following example defines a TACACS+ group.

```
group = admin {
group admin has a cleartext password which all members share
unless they have their own password defined
chap = cleartext "my$parent$chap$password"
}
```

The following example assigns the user “Brocade” with the group “admin”.

```
user = Brocade {
member = admin
pap = cleartext "pap password"
}
```

## Configuring TACACS+ for a mixed vendor environment

Network OS uses Role Based Access Control (RBAC) to authorize access to system objects by authenticated users. In AAA environments users may need to be authorized across Brocade and non-Brocade platforms. You can use TACACS+ to provide centralized AAA services to multiple Network Access Servers (NAS) or clients. To use TACACS+ services in multi-vendor environments, you must configure the Attribute-Value Pair (AVP) argument to be optional as shown in the example.

```
brcd-role=admin
```

The Network OS device sends the optional argument ‘brcd-role’ in the authorization request to the TACACS+ service. Most TACACS+ servers are programmed to return the same argument in response to the authorization request. If ‘brcd-role’ is configured as an optional argument, it is sent in the authorization request and Network OS users are able to successfully authorize with all TACACS+ services in a mixed-vendor environment.

### *Configuring optional arguments in tac\_plus*

The following is a specific example for tac\_plus package. Syntax for other packages may differ.

In the example, the mandatory attribute priv-lvl=15 is set to allow Cisco to authenticate. The optional brcd-role = admin argument is added to the tac\_plus.conf file and allows Brocade VDX switches to authenticate.

The following example configures a user with the optional attribute value pair, brcd-role = admin. A Brocade user must match both the *username* and *usergroup* to authenticate successfully.

```
user = <username> {
 default service = permit
 service = exec {
 priv-lvl=15
 optional brcd-role = admin
 }
}
Or
group = <usergroup> {
```

```

 default service = permit
 service = exec {
 priv-lvl=15
 optional brcd-role = admin
 }
}

user = <username> {
 Member = <usergroup>
}

```

## LDAP

Lightweight Directory Access Protocol (LDAP) is an open-source protocol for accessing distributed directory services that act in accordance with X.500 data and service models. The LDAP protocol assumes that one or more servers jointly provide access to a Directory Information Tree (DIT) where data is stored and organized as entries in a hierarchical fashion. Each entry has a name called the distinguished name that uniquely identifies it.

The LDAP protocol can also be used for centralized authentication through directory service.

Active Directory (AD) is a directory service which supports a number of standardized protocols such as LDAP, Kerberos authentication, and DNS, to provide various network services. AD uses a structured data store as the basis for a logical, hierarchical organization of directory information. AD includes user profiles and groups as the part of directory information, so it can be used as a centralized database for authenticating the third-party resources.

If you are in logical chassis cluster mode, the configuration is applied to all nodes in the cluster.

### User authentication

A Brocade switch can be configured as an LDAP client for authentication with an AD server, supporting authentication with a clear text password over the transport Layer Security (TLS) channel. Optionally, it supports server authentication during the TLS handshake. Only the user principal name from the AD server is supported for LDAP authentication on the Brocade switch. The Common Name-based authentication is not supported. When you log in from the switch, the complete user principal name, including domain, should be entered (for example, "testuser@sec.example.com").

LDAP supports alternative user principal names, such as:

- username
- username@AD.com
- username@ADsuffix.com
- username@newUPN.com

Network OS supports LDAP authentication with the following AD servers:

- Windows 2000
- Windows 2003
- Windows 2008 AD

A Brocade switch configured to perform LDAP-based authentication supports its access through a serial port, Telnet, and SSH. These access channels require that you know the switch IP address or name to connect to the switches.

A maximum of five AD servers can be configured on a Brocade switch.

If you are in logical chassis cluster mode, LDAP server and maprole configurations (except “show certutil” and “certutil”) are applied to all switches in the cluster.

## Server authentication

As a part of user authentication using LDAP, the Brocade switch can be configured to support server certificate authentication. To enable server authentication (server certificate verification), follow these guidelines:

- While configuring the LDAP server, the Fully Qualified Domain Name (FQDN) of the AD server should be added as the host parameter, instead of the IP address. A FQDN is needed to validate the server identity as mentioned in the common name of the server certificate.
- The DNS server must be configured on the switch prior to adding AD server with a domain name or a hostname. Without a DNS server, the name resolution of the server fails, and then the add operation fails. Use the **ip dns** command to configure DNS.
- The CA certificate of the AD server’s certificate should be installed on the switch. Currently, only PEM-formatted CA certificates can be imported into the switch.

If more than one server is configured and an LDAP CA certificate is imported for one server on the switch, the switch performs the server certificate verification on all servers. Thus, either CA certificates for all servers should be imported, or CA certificates should not be imported for any of the servers. After the CA certificate is imported, it is retained even if the switch is set back to its default configuration. If the CA certificate is not required, you should explicitly delete it.

### *Importing an LDAP CA certificate*

The following example imports the LDAP CA certificate from a remote server to a Brocade switch using secure copy (SCP).

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **certutil import ldapca** command with the specified parameters.
3. Verify the import with the **show cert-util ldapcacert** command.

```
switch# configure terminal
Entering configuration mode terminal
switch# certutil import ldapca directory /usr/ldapcacert file cacert.pem
protocol SCP host 10.23.24.56 user admin password *****
switch# show cert-util ldapcacert
List of ldap ca certificate files:

swLdapca.pem
```

### *Deleting CA certificates*

The **no certutil ldapca** command deletes the LDAP CA certificates of all Active Directory servers.

```
switch# no certutil ldapca
Do you want to delete LDAP CA certificate? [y/n]:y
```

## Server authorization

The Active Directory (AD) server is used only for authentication. Command authorization of the AD users is not supported in the AD server. Instead, the access control of AD users is enforced locally by RBAC on the switch.

A user on an AD server should be assigned a nonprimary group, and that group name should be either matched or mapped to one of the existing roles on the switch; otherwise, authentication will fail. After successful authentication, the switch receives the nonprimary group of the user from the AD server and finds the corresponding user role for the group based on the matched or mapped roles.

If the switch fails to get the group from the AD server, or the LDAP user is not a member of any matching AD group, the user authentication fails. Groups that match with the existing switch roles have higher priority than the groups that are mapped with the switch roles. Thereafter, the role obtained from AD server (or default role) is used for RBAC.

If multiple nonprimary groups are associated to the AD user, only one of the groups should be mapped or matched to the switch role. If multiple AD groups of AD users are mapped or matched to the switch roles, authentication of the user is successful, but there is no guarantee as to which role the AD user gets among those multiple roles. After successful authentication, the switch gets the nonprimary group of the user from the AD server and finds the corresponding user role for group based on the matched or mapped roles. Thereafter, the role obtained from the AD server (or default role) will be used for RBAC.

A maximum 16 AD groups can be mapped to the switch roles.

## FIPS compliance

To support FIPS compliance, the CA certificate of the AD server's certificate should be installed on the switch, and the FIPS-compliant TLS ciphers for LDAP should be used.

## Client-side Active Directory server configuration

Each Brocade switch client must be individually configured to use AD servers. You use the **ldap-server** command to specify the host server, authentication protocols, and other parameters. You can configure a maximum of five AD servers on a Brocade switch for AAA service.

The parameters in [Table 45](#) are associated with an AD server that is configured on the switch.

**TABLE 45** AD parameters

| Parameter | Description                                                                                                                                                                          |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| host      | IP address (v4) or Fully Qualified Domain name of the AD server. IPv6 is supported for Windows 2008 AD server only. The maximum supported length for the host name is 40 characters. |
| port      | TCP port used to connect the AD server for authentication. The valid port range is 1024 through 65535. The default port is 389.                                                      |
| timeout   | Time to wait for a server to respond. The range is 1 through 60 seconds. The default value is 5 seconds.                                                                             |
| retries   | Number of unsuccessful attempts to be made to connect to an AD server before quitting. The valid range is 1 through 100. The default value is 5.                                     |
| domain    | Base domain name                                                                                                                                                                     |

A maximum of five LDAP/AD servers can be configured on a Brocade switch for authentication service.

### *Adding an LDAP server to the client server list*

The following procedure configures an LDAP server on an ADAP client (Brocade switch).

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode
2. Use the **ldap-server-host** command to set the parameters for the LDAP server.

This command places you into the ldap-server configuration submenu where you can modify the server default settings.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# ldap-server host 10.24.65.6 basedn sec.brocade.com port 3890
switch(config-ldap-server-10.24.65.6)#
```

3. Modify any settings, such as the domain name or retry limit, in this configuration mode (refer to [Table 45](#)).

```
switch(config-ldap-server 10.24.65.6)# basedn security.brocade.com
switch(config-ldap-server 10.24.65.6)# timeout 8
switch(config-host-10.24.65.6)# retries 3
```

4. Confirm the LDAP settings with the **do show** command.

Attributes holding default values are not displayed.

```
switch(config-ldap-server-10.24.65.6)# do show running-config ldap-server host
10.24.65.6
ldap-server host 10.24.65.6
port 3890
basedn security.brocade.com
retries 3
timeout 8
!
```



5. *Optional:* Use the **no ldap-server** command option to set an attribute back to the default value.

```
switch(config-ldap-server-10.24.65.6)# exit
switch(config)# no ldap-server host 10.24.65.6 retries
```

### *Changing LDAP server parameters*

Changing an LDAP server follows the same procedure as that noted for adding an LDAP server to the client server list. Enter the host IP address or host name, then enter the new values as required. Refer to [“Adding an LDAP server to the client server list”](#) on page 256.

#### **Example:**

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# ldap-server host 10.24.65.6
switch(config-host-10.24.65.6)# domain security.brocade.com
```

### *Removing an LDAP server*

The following example deletes an LDAP server entry from the switch LDAP server list.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode
2. Use the **no ldap-server** command to delete the LDAP server.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no ldap-server host 10.24.65.6
```

### *Importing an LDAP CA 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 role permissions.
2. In privileged EXEC mode, enter the **certutil import ldapca** command. Include the full path to the certificate on the host, specify SCP as the protocol, and include the IP address of the host.

#### **Example: Standalone mode**

```
switch# certutil import ldapca directory /usr/ldapcacert/ file cacert.pem
protocol SCP host 10.23.24.56 user jane password
password:
switch#
```

#### **Example: Logical chassis cluster mode**

```
switch# certutil import ldapca directory /usr/ldapcacert/ file cacert.pem
protocol SCP host 10.23.24.56 user jane password rbridge-id 3
password:
switch#
```

### *Deleting an LDAP CA certificate*

This procedure deletes the LDAP CA certificates of all attached Microsoft Active Directory servers from the switch.

1. Connect to the switch and log in using an account with admin role permissions.
2. In privileged EXEC mode, enter the **no certutil ldapca** command.

```
switch# no certutil ldapca
Do you want to delete LDAP CA certificate? [y/n]:y
```

### *Verifying LDAP CA certificates*

To test whether an LDAP CA certificate has been imported on the switch, in privileged EXEC mode, enter the **no certutil ldapca** command and examine the message returned by the system.

The command returns an error if there is no LDAP CA certificate on the switch. If an LDAP CA certificate exists on the switch, you are prompted to delete it. Enter **no** to retain the certificate.

**Example: When no LDAP CA certificate is present**

```
switch# no certutil ldapcacert
% Error: LDAP CA certificate does not exist.
```

**Example: When an LDAP CA certificate exists on the switch**

```
switch# no certutil ldapcacert
List of swLdapca.pem files:

swLdapca.pem
```

### *Viewing the LDAP CA certificate*

The following procedure allows you to view the LDAP CA certificate that has been imported on the switch.

1. Connect to the switch and log in using an account with admin role permissions.
2. In privileged EXEC mode, enter the **certutil import syslogca** command. Include the full path to the certificate on the host, specify SCP as the protocol, and include the IP address of the host.

**Example: To view the output in standalone mode**

```
switch# show cert-util ldapcacert
```

**Example: To view the output of rbridge-id 3 in logical chassis cluster mode**

```
switch# show cert-util syslogcacert rbridge-id 3
```

### *Importing a syslog CA certificate*

The following procedure imports the syslog CA certificate from the remote host to the switch.

1. Connect to the switch and log in using an account with admin role permissions.
2. In privileged EXEC mode, enter the **certutil import syslogca** command. Include the full path to the certificate on the host, specify SCP as the protocol, and include the IP address of the host.

**Example: Standalone mode**

```
switch# certutil import ldapca directory /usr/ldapcacert/ file cacert.pem
protocol SCP host 10.23.24.56 user jane password
password:
switch#
```

**Example: Logical chassis cluster mode**

```
switch# certutil import syslogca directory /usr/ldapcacert/ file cacert.pem
protocol SCP host 10.23.24.56 user jane password rbridge-id 3
password:
switch#
```

***Deleting a syslog CA certificate***

The following procedure deletes the syslog CA certificates of all attached Microsoft Active Directory servers from the switch.

1. Connect to the switch and log in using an account with admin role permissions.
2. In Privileged EXEC mode, enter the **no certutil syslogca** command.

```
switch# no certutil syslogca
Do you want to delete syslog CA certificate? [y/n]:y
Warning: All the syslogca CA certificates are deleted.
switch#
```

**Example: To view the output in standalone mode**

```
switch# no certutil syslogca
Do you want to delete syslogca certificate? [y/n]:y
Warning: All the syslogca CA certificates are deleted.
switch#
```

**Example: To view the output of rbridge-id 3 in logical chassis cluster mode**

```
switch# no certutil syslogca rbridge-id 3
Do you want to delete syslogca certificate? [y/n]:y
Warning: All the syslog CA certificates are deleted.
switch#
```

***Verifying syslog CA certificates***

To test whether a syslogCA certificate has been imported on the switch, in privileged EXEC mode, enter the **no certutil syslogca** command and examine the message returned by the system. The command returns an error if there is no syslog CA certificate on the switch. If a syslog CA certificate exists on the switch, you are prompted to delete it. Enter **no** to retain the certificate.

**Example: When no syslog CA certificate is present**

```
switch# no certutil syslogcacert
% Error: syslog CA certificate does not exist.
```

**Example: When a syslog CA certificate exists on the switch**

```
switch# no certutil syslogcacert
Do you want to delete syslog CA certificate? [y/n]:n
```

### *Viewing the syslog CA certificate*

The following procedure allows you to view the syslog CA certificate that has been imported on the switch.

1. Connect to the switch and log in using an account with admin role permissions.
2. In privileged EXEC mode, enter the **show cert-util syslogcacert** command.

**Example: To view the output in standalone mode**

```
switch# show cert-util syslogcacert
```

**Example: To view the output of rbridge-id 3 in logical chassis cluster mode**

```
switch# show cert-util syslogcacert rbridge-id 3
```

## Active Directory groups

An Active Directory (AD) group defines access permissions for the LDAP server similar to Brocade roles. You can map an Active Directory group to a Brocade role with the **ldap-server map-role** command. The command confers all access privileges defined by the Active directory group to the Brocade role to which it is mapped.

A user on an AD server should be assigned a nonprimary group, and that group name should be either matched or mapped to one of the existing roles on the switch.

After successful authentication, the user is assigned a role from a nonprimary group (defined on the AD server) based on the matched or mapped switch role.

A user logging in to the switch that is configured to use LDAP and has a valid LDAP user name and password will be assigned LDAP user privileges if the user is not assigned with any nonprimary group.

### *Mapping an Active Directory group to a switch role*

In the following example, a Brocade user with the admin role inherits all privileges associated with the Active Directory (AD) Administrator group.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode

```
switch# config terminal
Entering configuration mode terminal
```

2. Use the **ldap-server** command to set the group information.

A maximum of 16 AD groups can be mapped to the switch roles.

```
switch(config)# ldap-server maprole group Administrator role admin
```

### *Removing the mapping of an Active Directory to a switch role*

The following example removes the mapping between the Brocade admin role and the Active Directory (AD) Administrator group. A Brocade user with the admin role can no longer perform the operations associated with the AD Administrator group.

To unmap an AD group to a switch role, perform the following steps from privileged EXEC mode.

1. Use the **configure terminal** command to enter global configuration mode.

```
switch# configure terminal
Entering configuration mode terminal
```

2. Use the **no ldap-server** command to set the group information.

```
switch(config)# no ldap-server maprole group Administrator
```

### *Configuring the client to use LDAP/AD for login authentication*

After you configured the switch LDAP server list, you must set the authentication mode so that LDAP is used as the primary source of authentication. refer to “[Login authentication mode](#)” on page 236 for information on how to configure the login authentication mode.

## Server-side Active Directory configuration

The following high-level overview of server-side configuration for LDAP/AD servers indicates the steps needed to set up a user account. This overview is provided for your convenience only. All instructions involving Microsoft Active Directory can be obtained from [www.microsoft.com](http://www.microsoft.com) or from your Microsoft documentation. Confer with your system or network administrator prior to configuration for any special needs your network environment may have.

### *Creating a user account on an LDAP/AD server*

1. Create a user on the Microsoft Active Directory server.
2. Create a group. The group should either match with the user's Brocade switch role or you can map the role to the Brocade switch role with the **ldap-server maprole** command.
3. Associate the user with the group by adding the user to the group.

The user account configuration is complete.

### *Verifying the user account on the switch*

1. Log in to the switch as a user with admin privileges
2. Verify that the LDAP/AD server has an entry in the switch LDAP server list.

```
switch# show running-config ldap-server
```

3. In global configuration mode, set the login authentication mode on the switch to use LDAP only and verify the change.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no aaa authentication login
switch(config)# aaa authentication login ldap
```

```
switch(config)# do show running-config aaa
aaa authentication login ldap
```

4. Log in to the switch using an account with valid LDAP/AD only credentials to verify that LDAP/AD is being used to authenticate the user.
5. Log in to the switch using an account with switch-local only credentials. The login should fail with an access denied message.

## Configuring LDAP users on an AD server

1. Create a user.
  - a. Go to **Programs > Administrative Tools > Active directory Users and Computers**.
  - b. Add a user by completing the dialog shown in [Figure 28](#).
  - c. Save the account information.
  - d. From a command prompt, log in using the new user name and enter a password when prompted.

The screenshot shows the 'New Object - User' dialog box. At the top, it says 'Create in: sec.brocade.com/'. Below this, there are several input fields: 'First name:' with 'testuser', 'Initials:' (empty), 'Last name:' (empty), 'Full name:' with 'testuser', 'User logon name:' with 'testuser' and a dropdown menu showing '@sec.brocade.com', and 'User logon name (pre-Windows 2000):' with 'SEC\' and 'testuser'. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

FIGURE 28 Create new user

2. Create a group.
  - a. Go to **Programs > Administrative Tools > Active directory Users and Computers**.
  - b. Add a new group.
  - c. Save the group information.
3. Assign the group to the user.

- a. Click on the user name.
- b. From the **Properties** dialog, click the **Member Of** tab and update the field with the group name. This group should either match the switch role or it must be mapped with the switch role on the Brocade switch. In this instance, Domain Users is the primary group and therefore should not be mapped with the switch role.

## Clearing sessions

With the introduction of Network OS v4.0, you can use the **clear sessions** command to log out user sessions that are connected to a switch. This command is not distributed across a cluster. If you are in VCS mode, you must use the RBridge ID of the node to log out the users connected to the individual nodes.

### Example: In standalone mode

```
switch# clear sessions
This operation will logout all the user sessions. Do you want to continue
(yes/no)? y
```

### Example: In VCS mode

```
switch# clear sessions rbridge-id 3
This operation will logout all the user sessions. Do you want to continue
(yes/no)? y
```





# FIPS Support

---

## In this chapter

|                                            |     |
|--------------------------------------------|-----|
| • FIPS overview .....                      | 265 |
| • Zeroization functions .....              | 266 |
| • FIPS-compliant state configuration ..... | 267 |
| • Preparing the switch for FIPS .....      | 268 |
| • Zeroizing for FIPS .....                 | 275 |
| • LDAP in the FIPS-compliant state .....   | 275 |

## FIPS overview

Federal Information Processing Standards (FIPS) specify the security standards to be satisfied by a cryptographic module utilized in Network OS v4.0.0 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 must be cleared out or *zeroized*.

Before enabling the FIPS-compliant state, 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 nondeterministic random-number generator (non-DRNG). They also verify the consistency of RSA keys with regard to signing and verification and encryption and decryption. These conditional tests also verify that the downloaded firmware is signed.

---

### ATTENTION

Once enabled, the FIPS-compliant state cannot be disabled.

---

FIPS compliance can be applied to switches in standalone and fabric cluster mode. To support FIPS compliance, the CA certificate of the Active Directory server's certificate should be installed on the switch, and FIPS-compliant TLS ciphers for Lightweight Directory Access Protocol (LDAP) should be used.

The Network OS v4.0.0 firmware is signed by means of both SHA1 1024-bit and SHA256 2048-bit keys. Network OS versions earlier than v4.0.0 are signed by means of the SHA1 1024-bit key only.

Firmware signatures are automatically validated during firmware download. When upgrading or downgrading between Network OS v4.0.0 and an earlier firmware version, firmware download uses the SHA1 1024-bit key for firmware signature validation.

## Zeroization functions

Explicit zeroization can be done at the discretion of the security administrator. These functions clear the passwords and the shared secrets. [Table 46](#) lists the various keys used in the system that will be zeroized in a FIPS-compliant Network OS switch.

**TABLE 46** Zeroization behavior

| Keys                   | Zeroization CLI     | Description                                                                                                                                                                                                                                                                                                                                                 |
|------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FCSP CHAP secrets      | fips zeroize        | Automatically zeroized on session termination. All the SFTP sessions gets terminated on zeroization.                                                                                                                                                                                                                                                        |
| Passwords              | fips zeroize        | The <b>fips zeroize</b> command removes user-defined accounts in addition to default passwords for the root, factory, admin, and user default accounts. Only the admin role has permissions for this command which, in addition to removing user accounts and resetting passwords, performs the complete zeroization of the system, and reboots the switch. |
| RADIUS secret          | fips zeroize        | The <b>fips zeroize</b> command zeroizes the secret and deletes a configured 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. All SFTP sessions are terminated on zeroization.                                                                                                                                                                                                                                                             |
| SSH DH private keys    | No command required | Keys will be zeroized within code before they are released from memory.                                                                                                                                                                                                                                                                                     |
| SSH host keys          | No command required | Automatically zeroized on session termination. All the SFTP sessions gets terminated on zeroization.                                                                                                                                                                                                                                                        |
| 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. All SSH sessions terminate on zeroization.                                                                                                                                                                                  |
| 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       | fips zeroize        | 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 the FIPS-compliant state. It 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.

## Conditional tests

The conditional tests are for the random number generators and are executed to verify the randomness of the random number generators. The conditional tests are executed each time before 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 displayed on the local console. This action includes logging both passing and failing results.

## FIPS-compliant state configuration

By default, the switch comes up in the non-FIPS-compliant state. You can bring up the switch in the FIPS-compliant state by enabling the KATs and conditional tests and then rebooting the switch, but you must configure the switch first. The set of prerequisites shown in [Table 47](#) must be satisfied for the system to enter the FIPS-compliant state.

To be FIPS compliant, the switch must be rebooted. KATs are run on the reboot. If the KATs are successful, the switch enters the FIPS-compliant state. If the KATs fail, then the switch reboots until the KATs succeed. If the switch cannot enter the FIPS-compliant state and continues to reboot, you must return the switch to your switch service provider.

When the switch successfully reboots in the FIPS-compliant state, you must follow the restrictions listed in [Table 47](#) to be FIPS compliant. This table lists the Network OS features and their behaviors in the FIPS-compliant and non-FIPS-compliant states.

**TABLE 47 FIPS-compliant state restrictions**

| Features                                              | FIPS-compliant state                                                        | Non-FIPS-compliant state    |
|-------------------------------------------------------|-----------------------------------------------------------------------------|-----------------------------|
| autoupload of FFDC and trace support data             | Not supported                                                               | Supported (FTP)             |
| Configupload/ download/ supportsave/ firmwaredownload | SCP only                                                                    | FTP and SCP                 |
| HTTP/HTTPS access                                     | Disabled                                                                    | HTTP and HTTPS              |
| LDAP CA                                               | CA certificate must be available.<br>Cipher suites: AES256-SHA, AES128-SHA, | CA certificate is optional. |
| Outbound SSH and telnet client                        | Not supported                                                               | Supported                   |
| RADIUS authentication protocols                       | PEAP-MSCHAPv2                                                               | CHAP, PAP, PEAP-MSCHAPv2    |
| Root account                                          | Disabled                                                                    | Enabled                     |
| Signed firmware                                       | Supported                                                                   | Supported                   |
| SSH algorithms                                        | HMAC-SHA1 (MAC)<br>AES128-CBC, AES256-CBC (cipher suites)                   | No restrictions             |
| TACACS+ authentication                                | Not supported                                                               | CHAP and PAP                |
| Telnet/SSH access                                     | Only SSH                                                                    | Telnet and SSH              |
| vCenter                                               | Not supported                                                               | Supported                   |

---

**NOTE**

Although SNMP is not considered to be FIPS compliant, it is not blocked. SNMP is considered to have a plain text interface without any cryptographic content. The few write operations that are supported do not affect the security of the switch. OSPF is considered a plain text interface, and no protection is claimed for protocol data exchange.

---

## Preparing the switch for FIPS

It is important to prepare the switch for the following restrictions that exist in the FIPS-compliant state:

- The root account and all root-only functions are not available.
- Access to the Boot PROM is not available.
- HTTP, HTTPS, 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.
- For USB interfaces, an authorized operator is required to maintain the physical possession (at all times) of the USB token and shall not provide access to unauthorized individuals or entities.

See [Table 47](#) on page 267 for a complete list of restrictions between the FIPS-compliant and non-FIPS-compliant states.

---

**ATTENTION**

You need the admin role permissions to prepare the switch for the FIPS-compliant state.

---

Preparing a switch for FIPS-compliant state operation removes all critical security parameters from the switch. As a consequence, some parameters needed to operate the switch must be applied after enabling the FIPS-compliant state, including the following parameters:

- IP ACL rules used to block HTTP, HTTPS, and Telnet access
- Secret strings used in RADIUS server configuration
- CA certificates used in LDAP authentication

These parameters must be reconfigured after each zeroization of the switch.

### FIPS preparation overview

1. Disable Boot PROM access.
2. *Optional:* Configure an LDAP server for authentication and configure FIPS-compliant ciphers for LDAP.
3. Configure FIPS-compliant ciphers for SSH.
4. Configure FIPS-compliant TLS ciphers for LDAP.
5. Disable root access.
6. Removed configurations of unsupported features vCenter and TACACS+, and disable Dot1x authentication.
7. If any FC-SP authentication policy attributes have been configured, configure all DH-group 0 configuration to groups 1 to 4.

8. Disable auto-upload.
9. Enable the KATs and the conditional tests.
10. Zeroize and reboot the switch into the FIPS-compliant state.
11. Disable the Telnet server.
12. Configure IP ACLS to block HTTP, HTTPS, and Telnet ports.
13. For authentication by a Microsoft Active Directory server, import and install the LDCAP CA certificate for LDAP authentication.
14. *Optional:* For authentication by RADIUS server, configure a RADIUS server.

## Enabling the FIPS-compliant state

1. Log in to the switch by using an account with the admin role.
2. To enable in standalone mode, enter the **no vcs enable** command in privileged EXEC mode.  
In VCS mode, use **vcs [rbridge-id rbridge-id] [vcsld ID] [enable ID]** command to configure the node.
3. Enter the **unhide** command to provide access to hidden commands. To execute this command, you must enter the password “**fibranne**”.

This step is necessary to gain access to the **prom-access**, **fips root disable**, **fips selftests**, and **fips zeroize** commands.

```
switch# unhide fips
Password: *****
```



### CAUTION

Once access to the Boot PROM has been disabled, you cannot re-enable it.

4. Check the status of prom-access by executing these commands.

```
switch# unhide built-in-self-test
Password: *****
switch#
switch# show prom-access
PROM access Disabled
```

If prom-access is enabled, disable it by running following command, proceed to [step 5](#).

5. Enter the **prom-access disable** command to disable access to the Boot PROM.

```
switch# prom-access disable
You are disabling PROM access. Do you want to continue? [yes/no] (no): yes
PROM access Disabled
```

6. *If LDAP will be used for authentication:*

- a. *Configure FIPS-compliant LDAP ciphers (AES256-SHA, AES128-SHA, DES-CBC3-SHA):*

```
switch# cipheraset ldap
LDAP cipher list configured successfully.
```

- b. Delete any LDAP DSA or RSA 2048 CA certificate that already exists on the switch:

```
switch# no certutil ldapca
```

Do you want to delete LDAP CA certificate? [y/n]:y

---

## NOTE

In the FIPS-compliant state, only RSA 1024 CA certificates are supported. This command deletes all existing LDAP CA certificates on the switch.

---

*For more details about configuring LDAP and the FIPS-compliant LDAP ciphers, refer to [“Setting up LDAP for the FIPS-compliant state”](#) on page 276.*

7. Enter the **cipherset ssh** command to configure the FIPS-compliant ciphers for SSH (HMAC-SHA1 (mac), AES128-CBC, AES256-CBC).

```
switch# cipherset ssh
ssh cipher list configured successfully
switch# show cipherset
LDAP Cipher List : !DH:HIGH:-MD5
SSH Cipher List : aes128-cbc,aes256-cbc
```




---

## CAUTION

Once you have disabled root account access, you cannot re-enable it. To re-enable root account access, you must return your switch to your service provider.

---

8. Enter the **fips root disable** command and enter **yes** at the subsequent prompt to disable access from the root account.

```
switch# fips root disable
This operation disables root account. Do you want to continue? [yes,NO] yes

Network OS (switch)
switch console login: 2011/09/08-17:28:34, [SEC-1197], 19073,, INFO, switch,
Changed account root.
```

---

## NOTE

The **fips root disable** command was exposed by the **unhide** command in [step 3](#). It is normally a hidden command.

---

9. Enter the **show fips** command to confirm the status of fips.

```
switch# show fips
FIPS Selftests: Enabled
Root account: Disabled
```

10. Delete the TACACS+ configuration from the switch by using the following commands.

- a. Enter the **show running-config tacacs-server** command to list the existing TACACS+ configuration.
- b. For each TACACS+ server listed in [step a](#), enter the **no tacacs-server host** command and the IP address or host name to delete the TACACS+ server configuration.

```
switch# show running-config tacacs-server host ?
Description: Configure a TACACS+ Server for AAA
Possible completions:
 10.20.57.13 INETADDRESS;;Domain name or IP Address of this TACACS+
server
| Output modifiers
<cr>
Possible match completions:
```

```

port TCP Port for Authentication (default=49)
protocol Authentication protocol to be used (default=CHAP)
key Secret shared with this server (default='sharedsecret')
retries Number of retries for this server connection (default=5)
timeout Wait time for this server to respond (default=5 sec)
switch# configure terminal
Entering configuration mode terminal
switch(config)# no tacacs-server host 10.10.20.57.13

```

11. Enter the **no dot1x enable** command to disable 802.1x globally.

```

switch# configure terminal
Entering configuration mode terminal
switch(config)# no dot1x enable
switch(config)# exit

```

12. If vCenter is configured, remove the configuration using the following CLI:

```
switch(config)# no vcenter <name>
```

13. DH group 0 is not supported in the fips compliance state of the switch. If DH group 0 or '\*' is configured, execute the following to configure a different group between 1 to 4 (key sizes from 1024 to 2048 bits).

```
switch(config)# fcsp auth group <NUMBER:1-4>
```

14. If autoupload is enabled, disable it.

```
switch# autoupload disable
```



### CAUTION

Once FIPS self-tests are enabled, you cannot disable them. These tests will run on the next reboot and, if successful, will place the switch into the FIPS-compliant state.

15. Enter the **fips selftests** command to enable the FIPS KAT and conditional tests.

```
switch# fips selftests
self tests enabled

```

---

### NOTE

The **fips selftests** command was exposed by the **unhide** command in [step 3](#). It is normally a hidden command.

---

16. Enter the **fips zeroize** command and enter **yes** at the subsequent prompt to clear all passwords and secrets.

The switch reboots and comes up in the FIPS-compliant state.

```

switch# fips zeroize
This operation erases all passwords, shared secrets, private keys etc. on the
switch . Do you want to continue? [yes,NO] yes

```

---

### NOTE

The **fips zeroize** command was exposed by the **unhide** command in [step 3](#). It is normally a hidden command. When the switch reboots, the FIPS commands will be hidden again.

---

On reboot, the switch performs the KATs and conditional tests enabled in [step 15](#). The following sample output indicates successful completion of these tests, after which the switch comes up in the FIPS-compliant state, as shown below:

```
FIPS-mode test application
1. Non-Approved cryptographic operation test...
a. Excluded algorithm (MD5)...successful
b. Included algorithm (D-H)...successful
2. Automatic power-up self test...successful
3. AES-128 CBC encryption/decryption...successful
4. RSA key generation and encryption/decryption...successful
5. TDES-CBC encryption/decryption...successful
6. DSA key generation and signature validation...successful
7a. SHA-1 hash...successful
7b. SHA-256 hash...successful
7c. SHA-512 hash...successful
7d. SHA-1 hash...successful
7e. SHA-224 hash...successful
7f. SHA-256 hash...successful
7g. SHA-384 hash...successful
7h. SHA-512 hash...successful
8. Non-Approved cryptographic operation test...
a. Excluded algorithm (MD5)...Not executed
b. Included algorithm (D-H)...successful as expected
```

---

#### NOTE

If the output shows errors, the switch reboots. If the errors persist, you must return the switch to your service provider for repair.

---

17. Use IP ACLs to block the HTTP, HTTPS, Telnet, and Brocade internal ports. Enter the following commands for IPv4 and IPv6.
  - a. Enter the **ip access-list extended** command and a name for the IP ACL.
  - b. Enter a **seq deny** command to create a rule for blocking the HTTP port (80).
  - c. Enter a **seq deny** command to create a rule for blocking the HTTPS port (443).
  - d. Enter a **seq deny** command to create a rule for blocking the Telnet port (23).
  - e. Enter **seq deny** commands to create rules for blocking the Brocade internal server ports 3016, 4565, 5016, 7013, 7110, 7710, 9013, 9110, 9710, and 9910 through 10110.
  - f. If SSH access is required, enter seq permit commands to allow access on ports 22 and 830.
  - g. If remote access is required, such as through SCP or LDAP, enter seq permit commands to allow UDP and TCP traffic on ports 1024 through 65535. Enter the **interface management rbridge-id/port** command to enter the interface management subconfiguration mode.
  - h. Enter the **ip access-group** command with the ACL name created in [step a](#) to apply the ACL to the management interface.

These commands also disable the non-FIPS-compliant vCenter feature.

For IPv4:

```
switch(conf-ip-ext)# seq 1 deny tcp any any eq www
switch(conf-ip-ext)# seq 2 deny tcp any any eq 443
switch(conf-ip-ext)# seq 3 deny tcp any any eq telnet
switch(conf-ip-ext)# seq 4 deny tcp any any eq 2301
switch(conf-ip-ext)# seq 5 deny tcp any any eq 2401
switch(conf-ip-ext)# seq 6 deny tcp any any eq 3016
```



```

switch(conf-ip-ext)# seq 7 deny tcp any any eq 3516
switch(conf-ip-ext)# seq 8 deny tcp any any eq 4516
switch(conf-ip-ext)# seq 9 deny tcp any any eq 5016
switch(conf-ip-ext)# seq 10 deny tcp any any eq 7013
switch(conf-ip-ext)# seq 11 deny tcp any any eq 7110
switch(conf-ip-ext)# seq 12 deny tcp any any eq 7710
switch(conf-ip-ext)# seq 13 deny tcp any any eq 9013
switch(conf-ip-ext)# seq 14 deny tcp any any eq 9110
switch(conf-ip-ext)# seq 15 deny tcp any any eq 9710
switch(conf-ip-ext)# seq 16 deny tcp any any range 9910 10110
switch(conf-ip-ext)# seq 17 deny udp any any eq 33351
switch(conf-ip-ext)# seq 18 deny udp any any eq 36851
switch(conf-ip-ext)# seq 19 deny udp any any eq 37731
switch(conf-ip-ext)# seq 20 deny udp any any eq 50690
switch(conf-ip-ext)# seq 21 permit tcp any any range 1024 65535
switch(conf-ip-ext)# seq 22 permit udp any any range 1024 65534
switch(conf-ip-ext)# seq 23 permit udp any any eq 65535
switch(conf-ip-ext)# seq 100 permit tcp any any eq 22
switch(conf-ip-ext)# seq 101 permit tcp any any eq 830
switch(conf-ip-ext)# exit
switch(config)#

```

For IPv6:

```

switch(conf-ip-ext)# seq 1 deny tcp any any eq 80
switch(conf-ip-ext)# seq 2 deny tcp any any eq 443
switch(conf-ip-ext)# seq 3 deny tcp any any eq 23
switch(conf-ip-ext)# seq 4 deny tcp any any eq 2301
switch(conf-ip-ext)# seq 5 deny tcp any any eq 2401
switch(conf-ip-ext)# seq 6 deny tcp any any eq 3016
switch(conf-ip-ext)# seq 7 deny tcp any any eq 3516
switch(conf-ip-ext)# seq 8 deny tcp any any eq 4516
switch(conf-ip-ext)# seq 9 deny tcp any any eq 5016
switch(conf-ip-ext)# seq 10 deny tcp any any eq 7013
switch(conf-ip-ext)# seq 11 deny tcp any any eq 7110
switch(conf-ip-ext)# seq 12 deny tcp any any eq 7710
switch(conf-ip-ext)# seq 13 deny tcp any any eq 9013
switch(conf-ip-ext)# seq 14 deny tcp any any eq 9110
switch(conf-ip-ext)# seq 15 deny tcp any any eq 9710
switch(conf-ip-ext)# seq 16 deny tcp any any range 9910 10110
switch(conf-ip-ext)# seq 17 deny udp any any eq 33351
switch(conf-ip-ext)# seq 18 deny udp any any eq 36851
switch(conf-ip-ext)# seq 19 deny udp any any eq 37731
switch(conf-ip-ext)# seq 20 deny udp any any eq 50690
switch(conf-ip-ext)# seq 21 permit tcp any any range 1024 65535
switch(conf-ip-ext)# seq 22 permit udp any any range 1024 65534
switch(conf-ip-ext)# seq 23 permit udp any any eq 65535
switch(conf-ip-ext)# seq 100 permit tcp any any eq 22
switch(conf-ip-ext)# seq 101 permit tcp any any eq 830
switch(conf-ip-ext)# exit
switch(config)#

```

For inband management IPv4 ports, use the following rules:

```

switch(conf-ip-ext)# seq 5 hard-drop tcp any any eq 80
switch(conf-ip-ext)# seq 10 hard-drop tcp any any eq 443
switch(conf-ip-ext)# seq 15 hard-drop tcp any any eq 23
switch(conf-ip-ext)# seq 20 hard-drop tcp any any eq 2301
switch(conf-ip-ext)# seq 25 hard-drop tcp any any eq 2401
switch(conf-ip-ext)# seq 30 hard-drop tcp any any eq 3016
switch(conf-ip-ext)# seq 35 hard-drop tcp any any eq 3516

```

```

switch(conf-ip-ext)# seq 40 hard-drop tcp any any eq 4516
switch(conf-ip-ext)# seq 45 hard-drop tcp any any eq 5016
switch(conf-ip-ext)# seq 50 hard-drop tcp any any eq 7013
switch(conf-ip-ext)# seq 55 hard-drop tcp any any eq 7110
switch(conf-ip-ext)# seq 60 hard-drop tcp any any eq 7710
switch(conf-ip-ext)# seq 65 hard-drop tcp any any eq 9013
switch(conf-ip-ext)# seq 70 hard-drop tcp any any eq 9110
switch(conf-ip-ext)# seq 75 hard-drop tcp any any eq 9710
switch(conf-ip-ext)# seq 80 hard-drop tcp any any range 9910 10110
switch(conf-ip-ext)# seq 85 hard-drop udp any any eq 33351
switch(conf-ip-ext)# seq 90 hard-drop udp any any eq 36851
switch(conf-ip-ext)# seq 95 hard-drop udp any any eq 37731
switch(conf-ip-ext)# seq 100 hard-drop udp any any eq 50690
switch(conf-ip-ext)# seq 105 permit tcp any any range 1024 65535
switch(conf-ip-ext)# seq 110 permit udp any any range 1024 65534
switch(conf-ip-ext)# seq 115 permit udp any any eq 65535
switch(conf-ip-ext)# seq 120 permit tcp any any eq 22
switch(conf-ip-ext)# seq 125 permit tcp any any eq 830
switch(conf-ip-ext)# exit
switch(config)#

```

**NOTE**

For the switch to remain FIPS compliant, the HTTP, HTTPS, Telnet, and Brocade internal server ports (80, 443, 23, 2301, 2401, 3016, 3516, 4516, 5016, 7013, 7110, 7710, 9013, 9110, 9710, 9910 through 10110, 33351, 36851, 37731, and 50690) must be blocked after every zeroization operation.

18. Disable the Telnet server.

```

switch(config)# telnet server shutdown
switch(config)#

```

19. If LDAP authentication is required, in global configuration mode, enter the following command syntax to import the LDAP CA certificate:

**certutil import ldapca directory** *ca-certificate-directory file filename* **protocol** {FTP|SCP} **host** *remote-ip-address* **user** *user-account* **password** *password*

20. Specify SCP for the protocol.

```

switch# certutil import ldapca directory /usr/ldapcacert file cacert.pem
 protocol SCP host 10.23.24.56 user jane password *****

```

21. If RADIUS authentication is required, in global configuration mode, enter the **radius-server host** command with the specified parameters to add RADIUS servers that use only PEAP-MSCHAPv2 ciphers.

```

switch(config)# radius-server host 10.38.37.180 protocol peap-mschap key
"new#Hercules*secret" timeout 10
switch(config-host-10.38.37.180)#

```

22. Enter the **copy running-config startup-config** command to save all settings to the startup configuration file.

```

switch# copy running-config startup-config

```

**NOTE**

After the switch is in the FIPS-compliant state, do not use any non-FIPS-compliant algorithms such as FTP, DHCHAP, MD5. With regards to SCP client on the switch, the remote SCP server must employ RSA host keys with a minimum length of 1024 bits.

## Zeroizing for FIPS

1. Log in to the switch using an account with admin role permissions.
2. In privileged EXEC mode, enter the **fips zeroize** command.

The switch reboots automatically. If the KATs and conditional tests are enabled, then the switch will reboot in the FIPS-compliant state. If the tests are not enabled, the switch comes up in the non-FIPS-compliant state.

---

### NOTE

For the switch to remain FIPS compliant, the HTTP, HTTPS, telnet, and Brocade internal server ports (3016, 4565, 5016, 7013, 7110, 7710, 9013, 9110, 9710, and 9910 through 10110 inclusive) must be blocked after every zeroization operation.

---

## LDAP in the FIPS-compliant state

You can configure your Microsoft Active Directory server to use the Lightweight Directory Access Protocol (LDAP) while in the FIPS-compliant state.

[Table 48](#) lists the differences between the FIPS-compliant and non-FIPS-compliant states of operation.

**TABLE 48 FIPS-compliant and non-FIPS-compliant states of operation**

| FIPS-compliant state                                                                                                                                                                                                      | non-FIPS-compliant state                                                                                                                                          |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The certificate for 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.                                                            |
| 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 the non-FIPS-compliant state, then user authentication will succeed. |

When setting up an LDAP server for FIPS, you will need to perform the following tasks:

- Add a DNS server.
- Configure a Microsoft Active Directory server as the authentication device.
- Import the RSA 1024 LDAP CA certificate from the Microsoft Active Directory server to the switch.

Configuring the DNS server and the Microsoft Active Directory server should be performed before bringing up the switch in the FIPS-compliant state. Any DSA CA certificates must be deleted from the switch.

## Setting up LDAP for the FIPS-compliant state

1. Log in to the switch by using an account with admin role permissions.
2. In privileged EXEC mode, enter the **configureterminal** command to enter global configuration mode.
3. Enter the **ip dns domain-name** and **ip dns name-server** commands to configure DNS on the switch.

Specify the DNS IP address in either IPv4 or IPv6 format. 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.

```
switch# configure
Entering configuration mode terminal
switch(config)# ip dns domain-name sec.brocade.com
switch(config)# ip dns name-server 10.70.20.1
```

4. Enter the **aaa authentication login ldap** command to set the switch authentication mode for LDAP.

```
switch(config)# aaa authentication login ldap local
```

5. Enter the **ldap-server host** command to add your LDAP server. Provide the FQDN of the Microsoft Active Directory server for the host name parameter while configuring LDAP. The maximum supported length for the host name is 40 characters.

```
switch(config)# ldap-server host GEOFF5.ADLAP.LOCAL basedn sec.brocade.com
port 389 retries 3
switch(config-ldap-server-GEOFF5.ADLAP.LOCAL)# exit
switch (config) exit
switch# show running-config ldap-server host GEOFF5.ADLAP.LOCAL
ldap-server host GEOFF5.ADLAP.LOCAL
port 389
domain security.brocade.com
retries 3
!
switch#
```

6. Enter the **cipherset ldap** command to configure the FIPS-compliant ciphers for LDAP operation.

```
switch# cipherset ldap
ldap cipher list configured successfully
```

7. Set up LDAP according to the instructions in “LDAP” in [Chapter 17, “External server authentication”](#) and then perform 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 49](#).

**TABLE 49** 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.



# Fabric Authentication

---

## In this chapter

- [Fabric authentication overview](#) . . . . . 279
- [Switch connection control \(SCC\) policy](#) . . . . . 284

## Fabric authentication overview

When you connect a Brocade VCS Fabric to a Fabric OS fabric, the Network OS Fibre Channel E\_Ports on the Brocade VDX 6730 connect through Inter-Switch Links (ISLs) to EX\_Ports on an FC router, which in turn connects to the Fabric OS network as shown in [Figure 26](#) on page 182.

To ensure that no unauthorized devices can access the fabric, Network OS provides support for security policies and protocols capable of authenticating Network OS devices (E\_Ports) to the EX\_Ports on the FC router (FCR) that provides access to the SAN storage and services.

### DH-CHAP

Network OS use the Diffie Hellman - Challenge Handshake Authentication Protocol (DH-CHAP) to control access between devices. DH-CHAP is a password-based, key exchange authentication protocol that negotiates hash algorithms and Diffie Hellman (DH) groups before performing authentication. It supports both MD5 and SHA-1 hash algorithm-based authentication.

The Fibre Channel Security Protocol (FC-SP) defines the DH groups supported in the DH-CHAP protocol. Following current FC-SP standards, Network OS supports the following DH groups:

- 00 – DH Null option
- 01 – 1024 bit key
- 02 – 1280 bit key
- 03 – 1536 bit key
- 04 – 2048 bit key

To configure DH-CHAP authentication between Network OS switches (E\_Ports) and FC routers (EX\_Ports) you must apply a matching configuration to both sides of the connection. Each device must be configured locally.

---

#### NOTE

The Brocade VDX 6730-32 and VDX 6730-76 are the only platforms that can connect to an FC router providing access to a SAN network of Fabric OS switches.

---

Shared secret keys

When you configure device ports for DH-CHAP authentication, you define a pair of shared secrets known to both devices as a secret key pair. A key pair consists of a local secret and a peer secret. The local secret uniquely identifies the local device. The peer secret uniquely identifies the entity to which the local device may authenticate. Every device may share a secret key pair with any other device or host in a fabric.

Shared secret keys have the following characteristics:

- The shared secrets must be configured locally on every device.
- If shared secrets are not set up for a link, authentication fails. The “Authentication Failed” error is reported for the port.
- The minimum length of a shared secret is 8 bytes and the maximum 40 bytes.




FIGURE 29 DH-CHAP authentication

Figure 29 illustrates how the secrets are configured. Assume two devices, A and B. Each device has a local secret (local secret A and local secret B), and a matching peer secret (Peer secret A and peer secret B). If device B wants to shake hands with A, it will use A's local secret (B's peer secret A) to send the information. In doing so, A authenticates B by confirming its identity through the exchange of matching secret pairs. Conversely, B authenticates A when A sends information to B using B's local secret (A's peer secret B).

On the FC router, the authentication configuration for EX\_Ports is set to fixed default values and cannot be changed. The Fabric OS **authutil** command is applicable only to the E\_Ports on the FC router, not to EX\_Ports. Table 50 shows the default authentication configuration for EX\_Ports:

TABLE 50 Default Ex\_Port configuration

| Operand     | Value             |
|-------------|-------------------|
| Auth-type   | DHCHAP            |
| Auth-Policy | PASSIVE           |
| Auth-Group  | * (0, 1, 2, 3, 4) |
| Auth-Hash   | msd5, sha1        |



## SSH server key exchange configuration

The SSH server key-exchange specifies the method used for generating one-time session keys for encryption and authentication with the SSH server. Currently, you can configure the SSH server key-exchange method to DH Group 14.

If you are in logical chassis cluster mode, the command is not distributed across the cluster. The RBridge ID of the node should be used to configure service on individual nodes.

When the SSH server key exchange method is configured to DH Group 14, SSH connection from a remote SSH client is allowed only if the key-exchange method at the client end is also configured to DH Group 14.

By default, SSH server key-exchange is not configured as DH Group 14. Enter **no ssh server key-exchange dh-group-14** to restore SSH server key-exchange to the default value.

## Authentication Policy configuration

The switch authentication (AUTH) policy initiates DH-CHAP 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 active. You must configure the AUTH policy on all connected fabric entities.

If you are in logical chassis cluster mode, this command is *not* distributed across the cluster. The RBridge ID of the node should be used to configure protocol and policy configurations.

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 requirement if the policy is changed to OFF.

Authentication policy configuration is not distributed across the cluster. The RBridge ID of the node should be used to configure protocol and policy configurations.

You can set the authentication policy to any of the values listed in [Table 51](#). The remaining attributes are optional.

**TABLE 51** User account attributes

| Setting | Description                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ON      | Strict authentication is enforced on all E_Ports. During switch initialization, authentication is initiated on all E_Ports automatically. The authentication handshaking is completed before the switches exchange the fabric parameters (EFP) for E_Port bring-up. If the connecting switch does not support the authentication or the policy is turned off, all ports are disabled and the ISL goes down. |
| ACTIVE  | A switch with an ACTIVE policy is more tolerant and can connect to a device with any type of policy. During switch initialization, authentication is initiated on all E_Ports, but the port is not disabled if the connecting switch does not support authentication, or if the authentication policy is turned off.                                                                                        |

TABLE 51 User account attributes (Continued)

| Setting              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PASSIVE<br>(default) | The switch does not initiate authentication, but participates in authentication if the connecting switch initiates authentication. The switch does not start authentication on E_Ports, but accepts the incoming authentication requests, and will not be disabled if the connecting switch does not support authentication or the policy is turned off.                                                                                                                       |
| OFF                  | The switch does not support authentication, and rejects any authentication negotiation request from a neighbor switch or device. A switch with the policy set to OFF should not be connected to a switch with a policy set to ON. A policy set to ON policy is strict and disables the port if a peer rejects the authentication. DH CHAP shared secrets must be configured on both sides of the connection before you can change the policy from an OFF state to an ON state. |

The behavior of the policy between two adjacent switches is defined as follows:

- If the policy is ON or ACTIVE, the switch sends an Authentication Negotiation request to the connecting device.
- If the connecting device does not support authentication or the policy is OFF, the request is rejected.
- Once the authentication negotiation succeeds, the DH-CHAP authentication is initiated. If DH-CHAP authentication fails, the port is disabled, regardless of the policy settings.

The policy defines the responses of the host if the connecting switch does not support authentication. By default, the policy is set to PASSIVE and you can change the policy with the **fcsp auth** command. This includes starting authentication on all E\_Ports if the policy is changed to ON or ACTIVE, and clearing the authentication if the policy is changed to OFF. Before enabling the policy, you must install the DH-CHAP shared secrets. Refer to [“Configuring DH-CHAP shared secrets”](#) on page 282.

## Configuring device authentication

Configuring a Brocade VDX 6730 switch to access a SAN fabric connected through an FC Router involves the following steps:

1. Configure the matching shared secret pairs on the VDX 6730 and on the FC router.
2. Configure the authentication policy on the VDX 6730 switch (The FC router configuration is fixed).
3. Activate the authentication policy.

Setting up secret keys can quickly become an administrative challenge as your fabric size increases. As a minimum, key pairs need to be installed on all connected fabric entities. However, when connections change, you must install new key pairs to accommodate these changes. If you anticipate this situation, you may install key pairs for all possible connections up front, thus enabling links to change arbitrarily while still maintaining a valid key pair for any new connection.

### *Configuring DH-CHAP shared secrets*

To configure the DH-CHAP shared secrets, execute the **fcsp auth-secret** command in privileged EXEC mode. Provide the following information as shown in the example:

- The world wide name (WWN) of the peer.
- The secret of the peer that authenticates the peer to the local switch.

- The local secret that authenticates the local switch to the peer.

---

**NOTE**

Only the following non-alphanumeric characters are valid for the secret key:

@ \$ % ^ & \* ( ) \_ + - < > { } [ ] ; ' :

---

```
switch# fcsp auth-secret dh-chap node 10:00:00:05:1e:7a:c3:00 peer-secret
12345678 local-secret 87654321
Shared secret is configured successfully.
```

To display the device (WWN) for which the shared secret is configured, use the **show fcsp auth-secret dh-chap** command in privileged EXEC mode.

```
switch# show fcsp auth-secret dh-chap
10:00:00:05:1e:7a:c3:00
```

To remove the shared secrets, use the **no fcsp auth-secret** command in privileged EXEC mode.

```
switch# no fcsp auth-secret dh-chap node 10:00:00:05:1e:7a:c3:00
Shared secret successfully removed
```

### *Setting the authentication policy parameters*

The following procedure configures an authentication policy auth-type DH-CHAP (only option), a DH group of 2, and a hash type of md5. The switch policy is set to OFF until you are ready to explicitly activate the policy.

1. In privileged EXEC mode, use the **configure terminal** command to enter global configuration mode.
2. Enter the **fcsp auth** command with the specified parameters.
3. Enter the **do show running-config fcsp auth** command to verify the configuration.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# fcsp auth auth-type dh-chap hash md5 group 2 switch policy off
switch(config)# do show running-config fcsp auth
fcsp auth group 2
fcsp auth hash md5
fcsp auth policy switch off
```

### *Activating the authentication policy*

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. Enter the **fcsp auth policy active** command to change the policy state from OFF to ON.
3. Enter the **do show running-config fcsp auth** command to verify the configuration.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# fcsp auth auth-type switch policy on
switch(config)# do show running-config fcsp auth
fcsp auth group 2
fcsp auth hash md5
fcsp auth policy switch on
```

## Switch connection control (SCC) policy

The Switch Connection Control (SCC) policy controls access between neighboring devices. The policy defines and restricts which devices can join the fabric. Each time an E\_Port-to-EX\_Port connection is attempted, the devices are checked against the policy and the connection is either accepted or rejected depending on whether the connecting device is listed in the policy. The policy is named SCC\_POLICY and accepts members listed as world wide names (WWNs).

A device configured with an active SCC policy reviews its database whenever a neighboring device tries to establish a connection. If the WWN of the connecting device is found in the SCC active policy database, the connecting device is allowed to join the fabric. If the neighboring device is not specified in the SCC policy active list, both devices are segmented.

By default, any device is allowed to join the fabric; the SCC policy is not enforced until it is created and activated. Creating a policy without any entries blocks access from all devices. The local switch is not required to be included in a switch-local SCC policy.

SCC policy commands are not distributed across the cluster. The RBridge ID of the node should be used to configure policy configurations.

---

### NOTE

The configuration is applicable only to E\_Ports on the Brocade VDX 6730 platforms. All configurations are local to the switch and are not automatically distributed across the fabric

---

## Defined and active SCC policy sets

The Switch Connection Control (SCC) policy maintains two versions, active, and defined, and creating a policy includes two distinct operations

1. Creating the defined SCC policy set
2. Activating the SCC policy

The defined policy includes a list of WWN members and it is configurable. You can create the SCC policy and its members using a single command, **secpolicy defined-policy SCC\_POLICY**. Or you can create the SCC policy first and add the members later. You can modify the defined policy at any time thereafter.

When you create the SCC policy and its defined member set, it remains inactive until you explicitly activate the policy with the **secpolicy activate** command. The SCC policy is enforced on the E\_Ports only after you activate the policy. When the policy is active, only the members included in the activated policy can communicate with each other. If you add additional devices to the defined policy, they remain inactive and access is blocked until you active the defined policy again.

Follow these guidelines and restrictions when configuring SCC policy:

- During the configuration replay operation, the defined and active policies are replayed and the E\_Ports are enabled or disabled based on the SCC policy entries in the active policy list.

During a configuration replay operation, if an E\_port is already disabled due to a violation, it will not come online even if the WWN entry is found in the active policy list. You must explicitly bring up the E\_port to enforce the active policy.

- During execution of the **copy file running-config** command, only the defined policy in the switch is updated with the config file entries; the active policy entries remain unchanged. In this case, you must use the **secpolicy activate** command to activate the defined policy list.

- If an empty policy is created and activated, but not saved, all Fibre Channel (FC) E\_ports will be in the disabled state after a reboot.
- Network OS requires that you invoke the **shutdown** command, followed by the **no shutdown** command to bring up the E\_port. Invoking the **no shutdown** command alone does not enable the port.

### *Creating a defined SCC Policy*

The following procedure creates a Switch Connection Control (SCC) policy, adds two members, and verifies the configuration.

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. Enter the secpolicy **defined-policy SCC\_POLICY** command.  
This command places you into the defined SCC configuration mode where you can add policy member WWNs.
3. Specify a policy member with the **member-entry WWN** command.
4. Specify a second policy member with the **member-entry WWN** command.
5. Exit the defined SCC configuration mode.
6. Enter the **do show running-config secpolicy defined-policy** command to verify the configuration.

#### **Example: Standalone mode**

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# secpolicy defined-policy SCC_POLICY
switch(config-defined-policy-SCC_POLICY)# member-entry \
10:00:00:05:1e:00:69:00
switch(config-defined-policy-SCC_POLICY)# member-entry \
10:00:00:08:2f:00:79:00
switch(config-defined-policy-SCC_POLICY)# exit
switch(config)# do show running-config secpolicy defined-policy
secpolicy defined-policy SCC_POLICY
 member-entry 10:00:00:05:1e:00:69:00
 !
 member-entry 10:00:00:08:2f:00:79:00
 !
```

#### **Example: VCS mode**

```
switch# config
Entering configuration mode terminal
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# secpolicy defined-policy SCC_POLICY
switch(config-defined-policy-SCC_POLICY)# ex
switch(config)#
```

#### **Example: To create SCC policy and add members into the defined policy set in VCS mode**

```
switch# config
Entering configuration mode terminal
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# secpolicy defined-policy SCC_POLICY member-entry
10:00:00:05:1e:00:69:00
switch(config-member-entry-10:00:00:05:1e:00:69:00)# exit
```

```
switch(config-defined-policy-SCC_POLICY)# exit
switch(config-rbridge-id-3)# exit
```

### *Modifying the SCC policy*

The same command sequence that creates the Switch Connection Control (SCC) policy adds additional members. The defined SCC member entries are cumulative. Use the **no** form of the command to remove members from the policy.

The following example adds a member and subsequently removes the same added member:

#### **Example: Standalone mode**

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# secpolicy defined-policy SCC_POLICY
switch(config-defined-policy-SCC_POLICY)# member-entry \
22:22:22:22:22:22:22:22
switch(config-defined-policy-SCC_POLICY)# no member-entry \
22:22:22:22:22:22:22:22
switch(config-defined-policy-SCC_POLICY)# exit
switch(config)# do show running-config secpolicy defined-policy
secpolicy defined-policy SCC_POLICY
member-entry 10:00:00:05:1e:00:69:00
!
member-entry 10:00:00:08:2f:00:79:00
!
```

#### **Example: VCS mode**

```
switch# config
Entering configuration mode terminal
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# secpolicy defined-policy SCC_POLICY
switch(config-defined-policy-SCC_POLICY)# member-entry \
10:00:00:05:1e:00:69:00
switch(config-defined-policy-SCC_POLICY)# no member-entry \
10:00:00:05:1e:00:69:00
switch(config-defined-policy-SCC_POLICY)# exit
switch(config)# do show running-config secpolicy defined-policy
secpolicy defined-policy SCC_POLICY
member-entry 10:00:00:05:1e:00:69:00
!
member-entry 10:00:00:08:2f:00:79:00
```

### *Activating the SCC policy*

1. Define the SCC policy as shown in section [“Creating a defined SCC Policy”](#) on page 285.
2. Enter the **secpolicy activate** in privileged EXEC mode.
3. Enter the **do show running-config secpolicy active -policy** command to verify the configuration.

#### **Example: Standalone mode**

```
switch# secpolicy activate
switch# do show running-config secpolicy active-policy
secpolicy active-policy SCC_POLICY
member-entry 10:00:00:05:1e:00:69:00
!
member-entry 10:00:00:08:2f:00:79:00
```

!

**Example: VCS mode**

```

switch# secpolicy activate rbridge-id 3
switch#
switch# do show running-config rbridge-id 3 secpolicy defined-policy
rbridge-id 3
secpolicy defined-policy SCC_POLICY
 member-entry aa:aa:aa:aa:aa:aa:aa:aa
 !
 member-entry bb:bb:bb:bb:bb:bb:bb:bb
 !
 member-entry cc:cc:cc:cc:cc:cc:cc:cc
 !
!
switch#

```

***Removing the SCC Policy***

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. Enter the secpolicy **no defined-policy SCC\_POLICY** command.
3. Exit global configuration mode.
4. Activate the SCC policy to save the defined policy configuration to the active configuration.
5. Enter the **do show running-config secpolicy active -policy** command to verify that the policy is empty.

**Example: To remove an entry from the policy list in standalone mode**

```

switch# config
Entering configuration mode terminal
switch(config)# no secpolicy defined-policy SCC_POLICY member-entry
10:00:00:05:1e:00:69:00
switch(config)# exit
switch# secpolicy activate
switch# do show running-config secpolicy active-policy
% No entries found.

```

**Example: To remove the SCC\_POLICY entry in standalone mode**

```

switch# config
Entering configuration mode terminal
switch(config)# no secpolicy defined-policy SCC_POLICY
switch(config)# exit
switch# secpolicy activate
switch# do show running-config secpolicy active-policy
% No entries found.

```

**Example: To remove an entry from the policy list of rbridge-id 3 in VCS mode**

```

switch# config
Entering configuration mode terminal
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# no secpolicy defined-policy SCC_POLICY
member-entry 10:00:00:05:1e:00:69:01
switch(config)# exit
switch# do show running-config secpolicy active-policy

```

## 19 Switch connection control (SCC) policy

```
% No entries found.
```

**Example: To remove the SCC\_POLICY entry of rbridge-id 3 in VCS mode**

```
switch# config
Entering configuration mode terminal
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# no secpolicy defined-policySCC_POLICY
switch(config)# exit
switch# do show running-config secpolicy active-policy
% No entries found.
```



## Network OS Layer 2 Switch Features

This section describes the Layer 2 features of Network OS, and includes the following chapters:

- [Administering Edge-Loop Detection . . . . . 291](#)
- [Configuring AMPP . . . . . 299](#)
- [Configuring FCoE Interfaces . . . . . 313](#)
- [Configuring VLANs . . . . . 331](#)
- [Configuring STP-Type Protocols . . . . . 345](#)
- [Configuring UDLD . . . . . 373](#)
- [Configuring Link Aggregation . . . . . 377](#)
- [Configuring LLDP . . . . . 389](#)
- [Configuring ACLs . . . . . 403](#)
- [Configuring QoS . . . . . 413](#)
- [Configuring 802.1x Port Authentication . . . . . 461](#)
- [Configuring sFlow . . . . . 469](#)
- [Configuring Switched Port Analyzer . . . . . 477](#)



# Administering Edge-Loop Detection

---

## In this chapter

- [Edge-loop detection overview . . . . .](#) 291
- [How ELD detects loops . . . . .](#) 293
- [Configuring edge-loop detection. . . . .](#) 295
- [Edge-loop troubleshooting . . . . .](#) 296

## Edge-loop detection overview

Edge-loop detection (ELD) detects and disables Layer 2 loops that would cause broadcast storms. Typically, these loops are caused by misconfigurations.

ELD is configured and enabled on Brocade VCS Fabric clusters. Any topology that includes one or more Brocade VCS Fabric clusters use ELD to detect Layer 2 loops and prevent broadcast storms. Standalone switches can be included in such a cluster, but loop detection takes place on the Brocade VCS Fabric cluster, and not on the standalone switch. You cannot use ELD in a network consisting of standalone switches only.

Specifically, ELD can be used to prevent broadcast storms caused by Layer 2 loops in the following topologies:

- A Brocade VCS Fabric cluster connects to a standalone switch.
- A Brocade VCS Fabric cluster connects to a multiple node network.
- A Brocade VCS Fabric cluster connects to other Brocade VCS Fabric clusters.

Figure 30 shows an example of a misconfiguration between a Brocade VCS Fabric cluster and a standalone switch that could cause a Layer 2 loop. In this case, a VLAG is configured on the edge devices of the Brocade VCS Fabric cluster for the two ISLs that connect the Brocade VCS Fabric cluster to the standalone switch. In this case, a LAG has not been created on the standalone switch at the other end of the ISLs. ELD detects and breaks this potential Layer 2 loop.




FIGURE 30 Missing LAG causes loop

Figure 31 shows another example for which ELD could be used to detect and break a Layer 2 loop. In this case, multiple Brocade VCS Fabric clusters are interconnected in a manner that creates a Layer 2 loop.





FIGURE 31 Interconnected Brocade VCS Fabric clusters cause loop

## How ELD detects loops

ELD works by multicasting Protocol Data Unit (PDU) packets on edge ports. A device recognizes a loop when it receives a PDU that it initiated. Once the device recognizes that a Layer 2 loop exists, it can take action to disable a port and break the Layer 2 loop.

To minimize the number of disabled ports, ELD assigns a priority to each port and a unique receive limit (pdu-rx-limit) to each Brocade VCS Fabric cluster. The port priority determines whether the sending or receiving edge port of the cluster is disabled. The pdu-rx-limit determines on which Brocade VCS Fabric the action takes place. Without these configured values, it is possible that a Layer 2 loop could be detected in multiple clusters at the same time. As a result, multiple ports would be disabled, stopping traffic among the Brocade VCS Fabric clusters.

Figure 32 shows the same interconnections as Figure 31 on page 293, but with ELD enabled on each edge port, and with port priorities and receive limits assigned.



**FIGURE 32** Interconnected Brocade VCS Fabric clusters with ELD enabled

With all ELD enabled edge ports sending PDUs at the same rate, VCS1 reaches its pdu-rx-limit first. Port 2/0/1 has a lower priority (higher priority number) than port 1/0/1, and is therefore selected to be disabled. If both ports have the same priority, the port with the higher port-ID is disabled.

If the port being shutdown by ELD is part of a LAG, all member ports of the LAG are also shutdown. If the port being shutdown is part of a vLAG, all member ports of the vLAG on that RBridge are also shutdown.

Once ELD disables a port, normal operation is for the port to remain disabled until any misconfiguration is repaired. Once the repair is finished, the port can be re-enabled manually.

#### NOTE

When ELD disables a port, the port is operationally down, but administratively still up. If a port is disabled by STP or some other L2 protocol, ELD does not process PDUs for that port.

## Configuring edge-loop detection

Edge-loop detection requires configuration at the global level and at the interface level. For global level configuration, you need to set the number of PDUs that the Brocade VCS Fabric cluster receives on any port before determining that a loop exists. This value is the *pdu-rx-limit*. You must also set the interval between sending PDUs by using the **hello-interval** command. The combination of *pdu-rx-limit* and **hello-interval** timer determines the time it takes for ELD to detect and break a Layer 2 loop.

At the interface level, you must enable ELD on each port you want it to run on and set the port priority. You should also specify a VLAN on which ELD is enabled.

Enter the **pdu-rx-limit** command to set the limit to a different number on each Brocade VCS Fabric cluster so that only one Brocade VCS Fabric cluster disables a port. We recommend setting this value in the increment of two to prevent race conditions which might disable ports on two Brocade VCS Fabric clusters that are incrementally only one apart.

Enter the **hello-interval** command to set the interval between PDUs. This interval must be set to the same value on all Brocade VCS Fabric clusters for which ELD is configured, otherwise the results of edge-loop detection become unpredictable.

Optionally, enter the **shutdown-time** command to configure ports to be re-enabled after a specified period of time (range 10 minutes to 24 hours). A typical use for this feature is in environments in which reconfiguration is common, such as in a typical lab environment. Typical use is to allow the default value of zero, which does not allow ports to be re-enabled automatically.

---

### NOTE

Any change to **shutdown-time** only takes effect for the ports that are disabled by ELD after the configuration change. Any ports that were already disabled by ELD before the **shutdown-time** change continues to follow the old **shutdown-time** value. These ports start to follow the new shutdown time after the currently running timer expires and ELD still detects the loop and shuts down the port again.

---

For each interface on which ELD runs, enter the **edge-loop detection** command to enable ELD. You must also enter the **edge-loop-detection port-priority** command to specify the ELD-port priority.

## Setting global ELD parameters for a Brocade VCS Fabric cluster

Perform this procedure on every Brocade VCS Fabric cluster where you configure ELD.

1. Log in to any switch in a Brocade VCS Fabric cluster.
2. In global configuration mode, enter the **protocol edge-loop-detection** command to enter edge-loop detection configuration mode.
3. Enter the **pdu-rx-limit number** command to set the number of PDUs that will be received before breaking the Layer 2 loop.

The *number* operand must be a value in the range 1 through 5. The default value is 1.

4. Enter the **hello-interval number** command to set the interval between PDUs.

The *number* operand has a unit of 1 ms. It must be in the range from 100 ms to 5000 ms. The default value is 1000 ms.

5. *Optional:* Enter the **edge-loop-detection shutdown-time number** command to set the number of minutes after which the shutdown port is re-enabled.

The *number* operand must be in the range 10 through 1440 (10 minutes through 24 hours). The default value is 0, indicating that the port is not automatically re-enabled.

**Example**

This example configures the Brocade VCS Fabric cluster to detect and break loops on receipt of 5 PDUs. Because the PDU interval is set to 2000 ms (2 seconds), any loop breaks after 10 seconds. The selected port will remain disabled for 24 hours, after which it is automatically re-enabled.

```
switch(config)# protocol edge-loop-detection
switch(config-eld)# edge-loop-detection pdu-rx-limit 5
switch(config-eld)# hello-interval 2000
switch(config-eld)# edge-loop-detection shutdown-time 1440
```

**Setting interface parameters on a port**

Perform this procedure for every port you want to be monitored by ELD.

1. Log in to any switch in a Brocade VCS Fabric cluster.
2. In global configuration mode, enter the **interface** command to select the *rbridge-id/slot/port* for which you want to enable edge-loop detection.
3. In interface configuration mode, enter the **edge-loop-detection vlan** command to specify the VLAN you want ELD to monitor on this port.

If you do not specify a VLAN, the command fails.

4. Enter the **edge-loop-detection port-priority** command to specify the ELD port priority of the specified port for the selected VLAN. However, enabling switching is not mandatory for assigning a port-priority.

**NOTE**

The priority range of values is from 0 through 255. A port with priority 0 means that shutdown for this port is disabled. The default value port priority is 128

**Example**

This example sets the ELD port priority on two port/VLAN pairs: port 1/0/7 VLAN 10 and port 4/0/6 VLAN 10. If both these ports are detected in the same loop, ELD shuts down port 4/0/6 when the pdu-rx-limit for the Brocade VCS Fabric cluster is reached. Port 4/0/6 is chosen for shut down because it has been assigned the lower priority (higher number) then port 1/0/7.

```
switch(config)# interface TenGigabitEthernet 1/0/7
switch(config-if-te-1/0/7)# edge-loop-detection vlan 10
switch(config-if-te-1/0/7)# edge-loop-detection port-priority 5
switch(config-if-te-1/0/7)# top
switch(config)# interface TenGigabitEthernet 4/0/6
switch(config-if-te-1/0/7)# edge-loop-detection vlan 10
switch(config-if-te-1/0/7)# edge-loop-detection port-priority 7
```

**Edge-loop troubleshooting**

Use the edge-loop detection commands to view and correct misconfigurations

1. Log in to any switch in a Brocade VCS Fabric cluster.
2. In the global configuration mode, enter the **show edge-loop-detection** command to display edge-loop detection statistics for the Brocade VCS Fabric cluster.

The command output shows ports disabled by ELD.



3. Correct any misconfigurations detected in [step 2](#).
4. Perform one of the following operations in global configuration mode:
  - To re-enable one port that was disabled by ELD:
    - a. Enter the **shutdown** command on the port disabled by ELD.
    - b. Enter the **no shutdown** command on the port disabled by ELD.

---

**NOTE**

If an edge-port becomes an ISL port because a remote port's VCS ID was changed, a port that was already shutdown by ELD must be cycled with the **shutdown** and **no shutdown** command to be detected as an ISL port.

---

- To re-enable all ports disabled by ELD, enter the **clear edge-loop-detection** command.



# Configuring AMPP

---

## In this chapter

- [AMPP overview](#) ..... 299
- [Configuring AMPP port-profiles](#) ..... 302
- [Monitoring AMPP profiles](#) ..... 309

## AMPP overview

Server virtualization infrastructure associates a server-side Virtual Ethernet Bridge (VEB) port-profile to each Ethernet MAC address used by a virtual machine (VM) to access the network through a VEB port.

If the VM is migrated from one physical server to another, the VEB's port-profile migrates with it, providing automated port-profile migration of the server's VEB ports that are associated with the VM.

For environments where the server's virtualization infrastructure provides sufficient controls, automated port-profile migration approaches are fine. An example of such an environment is a high performance cluster that uses a Layer 2 network that is isolated from external networks through firewalls and security appliances.

However, there is a gap between the access and Quality of Service (QoS) controls supported in external Layer 2 switches and the server virtualization infrastructure. External Layer 2 switches have more advanced controls compared to server VEB implementations.

Some environments prefer the more advanced controls provided by external network switches. An example of such an environment is a multi-tier data center that has several types of applications, each with differing advanced network controls, running over the same Layer 2 network. In this type of environment, the network administrator often prefers the use of advanced access controls available in external switches.

Layer 2 networks do not provide a mechanism for automatically migrating switch access and traffic controls associated with an end-point device when that device migrates from one switch to another. The migration may be physical, such as an operating system image (such as an application, middleware, operating system, and associated state) that is running BareMetal OS on one system and is migrated to another system. The migration may be also be virtual, such as an operating system image that is running over Hypervisor VMware on one system and is migrated to run over Hypervisor VMware on another system.

The Brocade Auto Migrating Port Profile (AMPP) feature provides these advanced controls for maintaining and migrating these port-profile associations when a VM migrates across physical servers.

## AMPP over vLAG

Virtual Link Aggregation Group (vLAG) is the name for Brocade proprietary LAG in which the links to the Brocade VCS Fabric can be connected to one or more physical switches or servers. For redundancy and greater bandwidth, vLAG is a vital component of Brocade VCS Fabric technology. AMPP is supported on vLAG and standard LAG in a manner similar to physical port.

FCoE capability on all port-profiled interfaces can be activated using the fcoe-port configuration in the default port-profile (refer to [“Configuring FCoE profiles”](#) on page 306). This configuration enforces FCoE capability only on physical interfaces, not on the port-channel LAG. Member links of the LAG must be explicitly configured for FCoE capability.

For complete information on vLAG, refer to [Chapter 26, “Configuring Link Aggregation”](#).

The *italic* text in the following example highlights the vLAG information in the port profile:

```
switch# show port-profile status
Port-Profile PPID Activated Associated MAC Interface
auto-dvPortGroup 1 Yes None None
auto-dvPortGroup2 2 Yes None None
auto-dvPortGroup3 3 Yes None None
auto-dvPortGroup_4_0 4 Yes 0050.567e.98b0 None
auto-dvPortGroup_vlag 5 Yes 0050.5678.eaed None
auto-for_iscsi 6 Yes 0050.5673.85f9 None
 0050.5673.fc6d None
 0050.5674.f772 None
 0050.5675.d6e0 Te 234/0/54
 0050.567a.4288 None
auto-VM_Network 9 Yes 000c.2915.4bdc None
 0050.56a0.000d None
 0050.56a0.000e None
 0050.56a0.000f None
 0050.56a0.0010 Po 53
 0050.56a0.0011 Po 53
 0050.56a0.0012 Po 53
 0050.56a0.0013 None
 0050.56a0.0025 None
 0050.56a0.0026 None
 0050.56a0.0027 None
 0050.56a0.0028 None
 0050.56a0.0029 Po 53
 0050.56a0.002a Po 53
 0050.56a0.002b Po 53
 0050.56a0.002c None
 0050.56a0.002d None
 0050.56a0.002e None
 0050.56a0.002f None
 0050.56b3.0001 Po 53
 0050.56b3.0002 Po 53
 0050.56b3.0004 Po 53
 0050.56b3.0005 None
auto-VM_kernel 10 Yes 0050.5671.4d06 None
 0050.5672.862f Po 53
 0050.5678.37ea None
 0050.567a.ddc3 None
auto-VM_NW_1G 11 Yes 0050.56b3.0000 None
 0050.56b3.0003 Po 82
 0050.56b3.0007 None
 0050.56b3.0008 Po 82
 0050.56b3.0009 Po 82
```

```

auto-VMkernel 12 Yes 0050.567a.fdcf Po 82
 0050.567c.c2e3 None
auto-VMkernel_VS 13 Yes 0050.567d.16b9 None
 0050.567e.e25b None
auto-Management+Network 14 Yes 5cf3.fc4d.ca88 None
auto-Virtual+Machine+Network 15 Yes 000c.2941.27e2 None
 000c.2980.335d None

switch# show port-profile int all
Interface Port-Profile
Gi 234/0/1 None
Gi 234/0/13 None
Gi 234/0/25 None
Gi 234/0/26 None
Te 234/0/54 auto-for_iscsi
Po 82 auto-VM_NW_1G
 auto-VMkernel
Po 53 auto-VM_Network
 auto-VM_kernel

```

## AMPP and Switched Port Analyzer

Switched Port Analyzer (SPAN), or Port Mirroring, selects network traffic for analysis by a network analyzer. If you are interested in listening or snooping on traffic that passes through a particular port, Port Mirroring is necessary to artificially copy the packets to a port connected to the analyzer.

AMPP support for a mirrored port provides the support needed to make the mirrored port as profiled-port, and the reverse as well. This does not allow configuring the destination port as the profiled port, or the reverse. SPAN allows the capability to mirror the traffic learnt on the profiled port.

For complete information on SPAN, refer to [Chapter 32, “Configuring Switched Port Analyzer”](#).

## Scalability

[Table 52](#) describes the scalability values supported by Network OS v4.0.0.


**TABLE 52** AMPP scalability values

| Metric                              | Logical chassis mode                     | Fabric cluster mode                      |
|-------------------------------------|------------------------------------------|------------------------------------------|
| Number of profiles                  | 750                                      | 750                                      |
| Number of VLANs in port- profiles   | 2000                                     | 2000                                     |
| QoS profile                         | 1 cee-map<br>1 mutation-map              | 1 cee-map<br>1 cos-mutation-map          |
| Number of ACLs in Security profiles | Same as Layer 2 ACL                      | Same as Layer 2 ACL                      |
| Number of Mac associations          | 8000 (4000 for the<br>Brocade VDX 6710.) | 8000 (4000 for the<br>Brocade VDX 6710.) |

The MAC and VLAN scaling numbers in [Table 52](#) are based on mac-association and vlan-profile scaling without any ACL configuration. Additionally, AMPP is subject to the maximum number of vLAGs and LAGs supported on the switch, which is 750 in this case.

## Configuring AMPP port-profiles

As shown in [Figure 33](#), the default port-profile contains the entire configuration needed for a VM to get access to the LAN and SAN.



**FIGURE 33** Port-profile contents

In addition, all the combinations can be mixed up with some security rules grouped under a security-profile.

---

**NOTE**

A port-profile does not contain some of the interface level configurations, such as LLDP, SPAN, LAG, and so on.

---

A port-profile operates as a self-contained configuration container. In other words, if a port-profile is applied on a completely new switch without any configuration, it is capable of configuring the interface's local configuration and starting to carry traffic. Any changes to the policies are immediately applied to the data plane.

Security profiles are applied to the ACLs based on the profile or PolicyID. Therefore, multiple security profiles can be applied to the same profiled port, except on the Brocade VDX 6730.

---

**NOTE**

The fcoe-profile is only available on the default profile. User-defined port-profiles do not have access to the fcoe-profile. See [“Configuring FCoE profiles”](#) on page 306 for details.

---

However, editing of the port-profile is not allowed once the port-profile is activated. Activation of the port-profile is mandatory when it is applied to a port.

## Life of a port-profile

A port-profile during creation will go through multiple states. The states of a port-profile are as follows:

- *Created*—This state specifies that a port-profile is created or modified, but may not be complete.
- *Activated*—This state specifies that a port-profile is activated and is available for MAC->port-profile association. If the port-profile created is not complete then the activation fails; you must resolve any conflicts or dependencies and reactivate the port-profile.
- *Associated*—This state specifies that one or more MAC addresses have been associated to this port-profile within the fabric.
- *Applied*—This state indicates that the port-profile is applied on the profiled port where the associated MAC address appeared. In the absence of any signaling protocol, the system snoops the packet to detect if the associated MAC address has appeared on the profiled port. Configuration of two different port-profiles can co-exist on a profiled port, but if there is a conflict then the application of the later port-profile fails.

Table 53 describes the AMPP events and the applicable failure behaviors.

**TABLE 53** AMPP behavior and failure descriptions

| AMPP event                                     | Applicable behavior and failures                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create port-profile                            | <ul style="list-style-type: none"> <li>• If the port-profile does not exist, then it is created. If it exists, then it is available for modification (if it is not yet activated).</li> </ul>                                                                                                                                                                                                                 |
| Activate port-profile                          | <ul style="list-style-type: none"> <li>• If the port-profile configuration is not complete, activation fails. Unless the port-profile is activated, it is not applied on any port-profile-port.</li> <li>• If all the dependency validations succeed, the port-profile is in the ACTIVE state and is ready for association.</li> <li>• A vlan-profile is mandatory for all port-profiles.</li> </ul>          |
| De-activate port-profile                       | <ul style="list-style-type: none"> <li>• This event removes the applied port-profile configuration from all the profiled-ports.</li> <li>• De-activation is allowed even if there are MAC addresses associated with the port-profile.</li> </ul>                                                                                                                                                              |
| Modify port-profile                            | <ul style="list-style-type: none"> <li>• Port-profile can be edited only in the pre-activation stage.</li> <li>• The port-profile is set to the INACTIVE state if any conflicting attributes are configured, or some dependent configuration is not completed.</li> <li>• Port-profile state is set as INACTIVE and any attempt to associate the port-profile to a MAC address may not be allowed.</li> </ul> |
| Associate MAC addresses to a port-profile      | <ul style="list-style-type: none"> <li>• If the MAC is already associated with a port-profile, the port-profile to MAC association fails.</li> <li>• Otherwise, if the port-profile to MAC association succeeds, when the same MAC is learned on any of the ports, the port-profile which has the MAC association is applied to the port.</li> </ul>                                                          |
| De-associate MAC addresses from a port-profile | <ul style="list-style-type: none"> <li>• If mapping exists, all the policies configured for a specific MAC address are removed from that port or switch.</li> </ul>                                                                                                                                                                                                                                           |
| Deleting a port-profile                        | <ul style="list-style-type: none"> <li>• An IN USE error is generated if the port-profile is in an activated state. AMPP forces you to de-activate the profile before deleting.</li> <li>• If the port-profile is in an inactive state, then deletion of profile removes all the MAC associations as well.</li> </ul>                                                                                         |

**TABLE 53** AMPP behavior and failure descriptions (Continued)

| AMPP event                                                 | Applicable behavior and failures                                                                                                                                          |
|------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Modifying port-profile content when in an associated state | <ul style="list-style-type: none"> <li>An IN USE error is generated if the port-profile is already activated.</li> </ul>                                                  |
| Moving the VM MAC and notifying the fabric                 | <ul style="list-style-type: none"> <li>All policies associated to the port-profile ID are mapped on the MAC address and applied to the new port in the fabric.</li> </ul> |
| Unused port-profile                                        | <ul style="list-style-type: none"> <li>You must manually remove the port-profile to MAC associations.</li> </ul>                                                          |

## Configuring a new port-profile

To support VM MAC address learning, the default port-profile is employed. The default profile is different from the other user-defined AMPP profiles:

- The port-profile ID (ppid) of the profile cannot be changed.
- The VLAN sub-profile cannot be modified.
- The QoS sub-profile and security-profile cannot be added.
- The default port-profile cannot be activated.

Brocade recommends that you create a new port-profile to accommodate your requirements. To configure a new port-profile, perform the following steps in privileged EXEC mode.

1. Configure the physical interface, LAG, or vLAG as a port-profile port.

```
switch(if-te-2/0/1)# port-profile-port
```

2. Create and configure a new port-profile name.

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# port-profile vml-port-profile
switch(config-port-profile-vml-port-profile)# vlan-profile
switch(config-pp-vlan)# switchport
switch(config-pp-vlan)# switchport mode trunk
switch(config-pp-vlan)# switchport trunk native-vlan 300
switch(config-pp-vlan)# switchport trunk allowed vlan add 300
```

3. Exit VLAN profile configuration mode.

```
switch(config-pp-vlan)# exit
```

4. Activate the profile.

```
switch(config)# port-profile vml-port-profile activate
```

5. Associate the profile to the MAC address for each host.

```
switch(config)# port-profile vml-port-profile static 0050.56bf.0001
switch(config)# port-profile vml-port-profile static 0050.56bf.0002
switch(config)# port-profile vml-port-profile static 0050.56bf.0003
switch(config)# port-profile vml-port-profile static 0050.56bf.0004
switch(config)# port-profile vml-port-profile static 0050.56bf.0005
```



## Configuring VLAN profiles

The VLAN profile defines the VLAN membership of the overall port-profile, which includes both the tagged and untagged VLANs.

---

### NOTE

Private VLAN port mode commands are not available for AMPP VLAN profiles.

---

To configure the VLAN profile, perform the following steps in global configuration mode.

1. AMPP profiles cannot be modified while active. De-activate the port-profile before modifying the VLAN profile.

```
switch(config)# no port-profile vml-port-profile activate
```

2. Enter VLAN profile configuration mode.

```
switch(config)# port-profile vml-port-profile
switch(config-port-profile-vml-port-profile)# vlan-profile
```

3. Use the **switchport** command to change the mode to Layer 2 and set the switching characteristics to the defaults.

```
switch(config-pp-vlan)# switchport
```

4. Access the VLAN profile mode for the correct VLAN.

```
switch(config-pp-vlan)# switchport access vlan 200
```

5. Enter trunk configuration mode.

```
switch(config-pp-vlan)# switchport mode trunk
```

6. Configure the trunk mode for the allowed VLAN IDs.

```
switch(config-pp-vlan)# switchport trunk allowed vlan add 10, 20, 30-40
```

7. Configure the trunk mode to be a native VLAN.

```
switch(config-pp-vlan)# switchport trunk native-vlan 300
```

8. Exit VLAN profile configuration mode.

```
switch(config-pp-vlan)# exit
```

9. Activate the profile.

```
switch(config)# port-profile vml-port-profile activate
```

10. Associate the profile to the MAC address for each host.

```
switch(config)# port-profile vml-port-profile static 0050.56bf.0001
switch(config)# port-profile vml-port-profile static 0050.56bf.0002
switch(config)# port-profile vml-port-profile static 0050.56bf.0003
switch(config)# port-profile vml-port-profile static 0050.56bf.0004
switch(config)# port-profile vml-port-profile static 0050.56bf.0005
```

## Configuring FCoE profiles

Only the FCoE profile of the default profile can be modified. The FCoE profile can only be part of the default profile. When it is part of the default profile, FCoE is enabled globally and all the profiled ports automatically become FCoE ports.

In the absence of the FCoE profile in the default AMPP profile, you can configure FCoE on a per-interface basis, based on the profiled ports. See [“Configuring FCoE interfaces”](#) on page 325 for details.

To globally configure the FCoE profile, perform the following steps in global configuration mode.

1. Enter port-profile configuration mode.

```
switch(config)# port-profile default
```

2. Enter FCoE-profile configuration mode.

```
switch(config-port-profile-default)# fcoe-profile
```

3. Activate the FCoE port profile.

An FCoE map cannot be applied on interfaces that already have a CEE map applied to it.

```
switch(config-fcoe-profile)# fcoeport default
```

## Configuring QoS profiles

QoS profiles define the following values:

- Incoming 802.1p priority is set to internal queue priority. If the port is in QoS untrusted mode, all incoming priorities will be mapped to default best effort priority.
- Incoming priority is set to outgoing priority.
- Mapping of incoming priorities is set to strict or WRR traffic classes.
- Enabling of flow control on a strict or a WRR traffic class.

The QoS profile has two flavors: CEE QoS and Ethernet QoS. The QoS profile may contain either CEE QoS or Ethernet QoS. Server side ports typically are carrying converged traffic.

To configure the QoS profile, perform the following steps in global configuration mode.

1. AMPP profiles cannot be modified while active. Deactivate the port-profile before modifying the VLAN profile.

```
switch(config)# no port-profile vml-port-profile activate
```

2. Enter QoS profile mode.

```
switch(config)# port-profile vml-port-profile
switch(config-port-profile-vml-port-profile)# qos-profile
switch(config-qos-profile)#
```

3. Apply the CEE map.

```
switch(config-qos-profile)# cee default
```

4. Set the default CoS value.

```
switch(config-qos-profile)# qos cos 7
```

5. Set the QoS trust attribute for CoS

```
switch(config-qos-profile)# qos trust cos
```

6. Apply a map to the profile. You can do either of the following:

- Apply the existing CoS-to-CoS mutation map.

```
switch(config-qos-profile)# qos cos-mutation vml-cos2cos-map
```

- Apply the existing CoS-to-Traffic-Class map.

```
switch(config-qos-profile)# qos cos-traffic-class vml-cos2traffic-map
```

7. Enable pause generation. You can do either of the following:

- Without PFC.

```
switch(config-qos-profile)# qos flowcontrol tx on rx on
```

- With PFC for each CoS.

```
switch(config-qos-profile)# qos flowcontrol pfc 1 tx on rx on
```

```
switch(config-qos-profile)# qos flowcontrol pfc 2 tx on rx on
```

8. Exit QoS profile mode.

```
switch(config-qos-profile)# exit
```

9. Activate the profile.

```
switch(config)# port-profile vml-port-profile activate
```

10. Associate the profile to the MAC address for each host.

```
switch(config)# port-profile vml-port-profile static 0050.56bf.0001
```

```
switch(config)# port-profile vml-port-profile static 0050.56bf.0002
```

```
switch(config)# port-profile vml-port-profile static 0050.56bf.0003
```

```
switch(config)# port-profile vml-port-profile static 0050.56bf.0004
```

```
switch(config)# port-profile vml-port-profile static 0050.56bf.0005
```

## Configuring security profiles

A security profile defines all the security rules needed for the server port. A typical security profile contains attributes for MAC-based standard and extended ACLs. Security profiles are applied to the ACLs based on the profile or PolicyID. Therefore, multiple security profiles can be applied to the same profiled port.

To configure the security profile, perform the following steps in global configuration mode.

1. AMPP profiles cannot be modified while active. Deactivate the port-profile before modifying the security profile.

```
switch(config)# no port-profile vml-port-profile activate
```

2. Enter security profile configuration mode.

```
switch(config)# port-profile vml-port-profile
```

```
switch(config-pp)# security-profile
```

```
switch(config-pp-security)#
```

3. Modify the ACL security attributes.

See [Chapter 28, “Configuring ACLs”](#) for details.

4. Apply the ACL to the security profile.

```
switch(config-pp-security)# mac access-group vml-acl in
```

5. Exit security profile configuration mode.

```
switch(config-pp-security)# exit
```

6. Activate the profile.

```
switch(config)# port-profile vml-port-profile activate
```

7. Associate the profile to the MAC address for each host.

```
switch(config)# port-profile vml-port-profile static 0050.56bf.0001
switch(config)# port-profile vml-port-profile static 0050.56bf.0002
switch(config)# port-profile vml-port-profile static 0050.56bf.0003
switch(config)# port-profile vml-port-profile static 0050.56bf.0004
switch(config)# port-profile vml-port-profile static 0050.56bf.0005
```

8. Activate the interface configuration mode for the interface you wish to modify.

The following example activates the mode for the 10-gigabit Ethernet interface in slot 0/port 0.

```
switch(config)# interface tengigabitethernet 1/0/1
```

9. Configure port-profile-port on the physical interface

```
switch(conf-int-te-1/0/1)# port-profile-port
```

### Deleting a port-profile-port

To delete a port-profile-port, perform the following steps in global configuration mode.

1. Activate the interface configuration mode for the interface you wish to modify.

The following example activates the mode for the 10-gigabit Ethernet interface in slot 0/port 0.

```
switch(config)# interface tengigabitethernet 1/0/1
```

2. Unconfigure port-profile-port on the physical interface

```
switch(conf-int-te-1/0/1)# no port-profile-port
switch(conf-int-te-1/0/1)# no shutdown
```

### Deleting a port-profile

To delete a port-profile, perform the following steps in privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#
```

2. Deactivate the port-profile.

```
switch(config)# no port-profile vml-port-profile activate
```

3. Use the **no** form of the **port-profile** command to delete the custom profile.

You cannot delete the default port-profile.

```
switch(config)# no port-profile vml-port-profile
```

## Deleting a sub-profile

To delete a sub-profile, perform the following steps in privileged EXEC mode.

1. Enter global configuration mode

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#
```

2. Deactivate the port-profile.

```
switch(config)# no port-profile vml-port-profile activate
```

3. Enter port-profile mode.

```
switch(conf-vml-port-profile)# port-profile vml-port-profile
```

4. To delete vlan sub-profile:

```
switch(conf-vml-port-profile)# no vlan-profile
```

5. To delete security sub-profile:

```
switch(conf-vml-port-profile)# no security-profile
```

6. To delete fcoe sub-profile under default profile:

```
switch(conf-pp-default)# no fcoe-profile
```

7. To delete qos sub-profile:

```
switch(conf-vml-port-profile)# no qos-profile
```

## Monitoring AMPP profiles

To monitor the AMPP profiles, perform the following steps in privileged EXEC mode.

1. Use the **show** command to display the current MAC details.

```
switch# show mac-address-table port-profile
Legend: Untagged(U), Tagged (T), Not Forwardable(NF) and Conflict(C)
VlanId Mac-address Type State Port-Profile Ports
1 0050.5679.5351 Dynamic Active Profiled(U) Te 111/0/10
1 0050.567b.7030 Dynamic Active Profiled(U) Te 111/0/12
1 005a.8402.0000 Dynamic Active Profiled(T) Te 111/0/24
1 005a.8402.0001 Dynamic Active Profiled(NF) Te 111/0/24
1 005a.8402.0002 Dynamic Active Not Profiled Te 111/0/24
1 005a.8402.0003 Dynamic Active Not Profiled Te 111/0/24
1 005a.8402.0004 Dynamic Active Not Profiled Te 111/0/24
1 005a.8402.0005 Dynamic Active Profiled(NF) Te 111/0/24
1 005a.8402.0006 Dynamic Active Not Profiled Te 111/0/24
1 005a.8402.0007 Dynamic Active Profiled(T) Te 111/0/24
```

## 21 Monitoring AMPP profiles

```
1 005b.8402.0001 Dynamic Active Profiled(T) Te 111/0/24
1 005c.8402.0001 Dynamic Active Profiled(T) Te 111/0/24
100 005a.8402.0000 Dynamic Active Profiled Te 111/0/24
100 005a.8402.0001 Dynamic Active Profiled(NF) Te 111/0/24
100 005a.8402.0003 Dynamic Active Not Profiled Te 111/0/24
100 005a.8402.0005 Dynamic Active Profiled(NF) Te 111/0/24
100 005a.8402.0007 Dynamic Active Profiled Te 111/0/24
Total MAC addresses : 17
```

2. Use the **show running-config** command to display all the available port-profile configurations.

```
switch# show running-config port-profile
port-profile default
vlan-profile
switchport
switchport mode trunk
switchport trunk allowed vlan all
switchport trunk native-vlan 1
!
fcoe-profile
fcoeport default
!
!
port-profile ppl
vlan-profile
switchport
switchport mode access
switchport access vlan 1
!
qos-profile
!
!
port-profile ppl activate
port-profile ppl static 1000.0000.0001
```

3. Use the **show port-profile** command to display the current port-profile configuration.

```
switch# show port-profile
port-profile default
ppid 0
vlan-profile
switchport
switchport mode trunk
switchport trunk allowed vlan all
port-profile vm_kernel
ppid 1
vlan-profile
switchport
switchport mode access
switchport access vlan 1
```

4. Use the **show port-profile status** command to display the current status of all AMPP profiles.

```
switch# show port-profile status applied
Port-Profile PPID Activated Associated MAC Interface
auto-for_iscsi 6 Yes 0050.5675.d6e0 Te 9/0/54
auto-VM_Network 9 Yes 0050.56b3.0001 Te 9/0/53
 0050.56b3.0002 Te 9/0/53
 0050.56b3.0004 Te 9/0/53
 0050.56b3.0014 Te 9/0/53

switch# show port-profile status activated
```

| Port-Profile          | PPID | Activated | Associated MAC | Interface |
|-----------------------|------|-----------|----------------|-----------|
| auto-dvPortGroup      | 1    | Yes       | None           | None      |
| auto-dvPortGroup2     | 2    | Yes       | None           | None      |
| auto-dvPortGroup3     | 3    | Yes       | None           | None      |
| auto-dvPortGroup_4_0  | 4    | Yes       | 0050.567e.98b0 | None      |
| auto-dvPortGroup_vlag | 5    | Yes       | 0050.5678.eaed | None      |
| auto-for_iscsi        | 6    | Yes       | 0050.5673.85f9 | None      |

switch# **show port-profile status associated**

| Port-Profile          | PPID | Activated | Associated MAC | Interface |
|-----------------------|------|-----------|----------------|-----------|
| auto-dvPortGroup_4_0  | 4    | Yes       | 0050.567e.98b0 | None      |
| auto-dvPortGroup_vlag | 5    | Yes       | 0050.5678.eaed | None      |
| auto-for_iscsi        | 6    | Yes       | 0050.5673.85f9 | None      |

5. Use show port-profile interface all to display profile and applied interface information

switch# **show port-profile interface all**

| Port-profile    | Interface |
|-----------------|-----------|
| auto-VM_Network | Te 9/0/53 |
| auto-for_iscsi  | Te 9/0/54 |





# Configuring FCoE Interfaces

---

## In this chapter

|                                      |     |
|--------------------------------------|-----|
| • FCoE overview .....                | 313 |
| • End-to-end FCoE .....              | 314 |
| • Layer 2 Ethernet overview .....    | 317 |
| • FCoE Initialization Protocol ..... | 322 |
| • FCoE queuing .....                 | 325 |
| • Configuring FCoE interfaces .....  | 325 |
| • FCoE over LAG .....                | 327 |

## FCoE overview

Fibre Channel over Ethernet (FCoE) enables you to transport FC protocols and frames over Data Center Bridging (DCB) networks. DCB is an enhanced Ethernet network that enables the convergence of various applications in data centers (LAN, SAN, and HPC) onto a single interconnect technology.

FCoE provides a method of encapsulating the Fibre Channel (FC) traffic over a physical Ethernet link. FCoE frames use a unique EtherType (FCoE uses 0x8906 and FIPS uses 0x8914) that enables FCoE SAN traffic and legacy LAN Ethernet traffic to be carried on the same link. FC frames are encapsulated in an Ethernet frame and sent from one FCoE-aware device across an Ethernet network to a second FCoE-aware device. The FCoE-aware devices may be FCoE end nodes (ENodes) such as servers, storage arrays, or tape drives on one end and FCoE Forwarders on the other end. FCoE Forwarders (FCFs) are switches providing SAN fabric services and may also provide FCoE-to-FC bridging.

The motivation behind using DCB networks as a transport mechanism for FC arises from the desire to simplify host protocol stacks and consolidate network interfaces in data center environments. FC standards allow for building highly reliable, high-performance fabrics for shared storage, and these characteristics are what DCB brings to data centers. Therefore, it is logical to consider transporting FC protocols over a reliable DCB network in such a way that it is completely transparent to the applications. The underlying DCB fabric is highly reliable and high performing, the same as the FC SAN.

In FCoE, ENodes discover FCFs and initialize the FCoE connection through the FCoE Initialization Protocol (FIP). The FIP has a separate EtherType from FCoE. The FIP includes a discovery phase in which ENodes discover VLANs supporting FCoE, solicit FCFs on those VLANs, and FCFs respond to the solicitations with advertisements of their own. At this point, the ENodes know enough about the FCFs to log in to them. The virtual link establishment and fabric login (FLOGI/FDISC) for VN-to-VF port connections is also part of the FIP.

## FCoE terminology

Table 54 lists and describes the FCoE terminology used in this document.

**TABLE 54 FCoE terminology**

| Term    | Description                                                             |
|---------|-------------------------------------------------------------------------|
| FCoE    | Fibre Channel over Ethernet                                             |
| DCB     | Data Center Bridging                                                    |
| VN_port | FCoE equivalent of an FC N_port                                         |
| VF_port | FCoE equivalent of an FC F_port                                         |
| ENode   | An FCoE device that supports FCoE VN_ports (servers and target devices) |

## End-to-end FCoE

The Brocade VCS Fabric is a convergence-ready fabric. This means it is capable of providing lossless service and other features expected of a CEE-capable network. Network OS v4.0.0 supports multi-hop FCoE, where an FCoE initiator can communicate with an FCoE target that is a number of hops away.

## FCoE operations

Each switch in the Brocade VCS Fabric cluster acts as a fully functional FCoE Forwarder (FCF). All Fibre Channel (FC) services required to support a Virtual Network (VN) must run on every Brocade VCS Fabric cluster switch, and each switch in the fabric acts as if it were a separate domain in an FC SAN.

For all practical purposes, a Brocade VCS Fabric operates similarly to an FC fabric because all the FCoE initiators and targets are connected to the Brocade VCS Fabric. Each switch in the cluster gets a domain ID, and once the fabric forms, all the FC services (such as Name Server, Login Controller, Domain Controller) are available on each individual cluster switch.

Network OS v4.0.0 supports FCR/LSAN zoning. Network OS v2.1.0 supports only open zoning for FCoE initiators. The fabric device limitation is set to 3000 FCoE devices in a Brocade VCS Fabric cluster, because open zoning floods all the State Change Notifications (SCNs) to every FCoE device. FCoE traffic forwarding across the fabric follows the same equal-cost multi-path (ECMP) routing rules as LAN traffic forwarding.

## FCoE end-to-end forwarding

FCoE frame forwarding between two FCoE devices attached to the Brocade VCS Fabric works similarly to Layer 3 IP routing. The end-node talks to the default gateway's MAC address and the Layer 2 headers are modified hop-by-hop until the frame reaches its final destination. Forwarding decisions are based on the contents of the IP header in the case of IP routing, and the IP header is untouched along the path. FCoE forwarding works the same way.

Figure 34 on page 315 illustrates this process. Assume that VN1 (an FCoE initiator) is trying to access VN2 (an FCoE target).

1. VN1 and VN2 discover VF1 and VF2 through FIP Discovery Protocol and perform a Fabric Login (FLOGI) to their respective VF ports. That is, VN1 performs an FIP FLOGI to VF1 and VN2 performs an FLOGI to VF2. This works like IP in that all communication between the end-station and the network happens to the router's MAC address at Layer 2. This means VN1 is always communicating with VF1 at Layer 2.
2. In a Brocade VCS Fabric implementation, all FC services are available on every cluster unit. This means there is Fibre Channel Network Switch (FCNS) available on both FCF1 and FCF2. The FCNS service functions identically as it does in an FC SAN. As a result, VN1 discovers VN2.
3. VN1 attempts an N\_port Login (PLOGI) to VN2, with the frame information shown at point 1 in Figure 34. The Layer 2 header contains VF1 as the destination MAC address. The Layer 3 header (in this case, the FC header) contains the actual DID and SID of the initiator and the target respectively.

In this example, because VN1 is connected to the FCF with a Domain ID of 1, its PID is 010100. Similarly, because VN2 is connected to FCF3, its FC address is 030100.

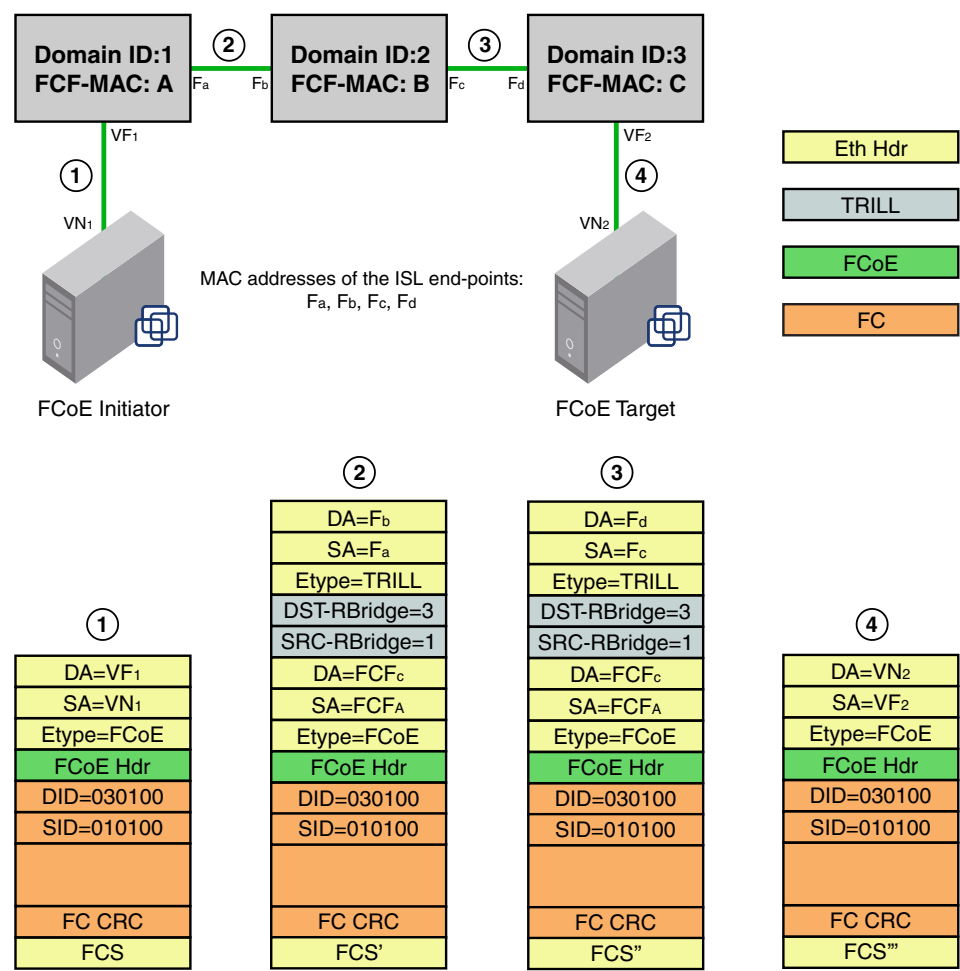


FIGURE 34 FCoE end-to-end header process

4. When FCF-A receives the frame on VF1, it performs a Layer 3 lookup. It looks up the DID in the FC header and determines that the frame is destined to a non-local domain. FCF-A decodes the next hop needed to reach the destination domain of 3, based on Fabric Shortest Path First (FSPF). It is at this point that it does something different than a normal IP router.
5. FCF-A now knows that it needs to reach FCF-C. Each FCF in the Brocade VCS Fabric is assigned an FCF MAC address. FCF-A constructs the Layer 2 header based on this information. So, the original MAC header is now transformed as follows: the DA is changed from VF1 to FCF-C and the SA is changed from VN1 to FCF-A. This occurs at point 2 in [Figure 34](#).
6. The frame gets a Transparent Interconnection of Lots of Links (TRILL) header and traverses across the fabric to reach FCF-C. The TRILL header indicates that the source is RBridge 1 and the destination is RBridge 3. This occurs at point 2 in [Figure 34](#).
7. The outer MAC header is a link level header that gets the frame from FCF-A to FCF-B. FCF-B receives the frame. FCF-B scans the TRILL header, decodes the destination RBridge ID in the frame, and forwards the frame. FCF-B only modifies the Layer 2 header. It neither looks up nor modifies anything in the FC header or the inner MAC header. This occurs at point 3 in [Figure 34](#).
8. FCF-C receives the frame. FCF-C scans the TRILL header and decodes the destination RBridge ID. FCF-C promotes the frame to Layer 3 lookup, because the FCF-C is the DA in the inner MAC header. FCF-C then scans the FC header and does something similar to an area route lookup in FC SAN. This lookup yields the MAC address of VN2 and the VF interface (in this case, VF2) information that it needs to use to forward the frame to VN2. This occurs at point 4 in [Figure 34](#).
9. VN2 receives the PLOGI. The PLOGI response from VN2 traverses back to VN1 in similar fashion.

---


**NOTE**

It is assumed that both VN1 and VN2 are configured to be in the same FCoE VLAN, and FCoE forwarding is enabled on this VLAN in the Brocade VCS Fabric. Network OS v4.0.0 supports only one FCoE VLAN for all FCoE devices connected to the fabric.

---

## Layer 2 Ethernet overview

The Brocade VDX hardware contains DCB ports that support FCoE forwarding. The DCB ports are also backwards-compatible and support classic Layer 2 Ethernet networks (see [Figure 35](#)). In Layer 2 Ethernet operation, a host with a Converged Network Adapter (CNA) can be directly attached to a DCB port on the Brocade VDX hardware. Another host with a classic 10-gigabit Ethernet network interface card (NIC) can be either directly attached to a DCB port, or attached to a classic Layer 2 Ethernet network that is attached to the Brocade VDX hardware.



**FIGURE 35** Multiple switch fabric configuration

## Layer 2 forwarding

Layer 2 Ethernet frames are forwarded on the DCB ports. 802.1Q VLAN support is used to tag incoming frames to specific VLANs, and 802.3ac VLAN tagging support is used to accept VLAN tagged frames from external devices.

Network OS v4.0.0 uses the following 802.1D bridging protocols between Layer 2 switches and to maintain a loop-free network environment:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)
- Per-VLAN Spanning Tree (PVST+)
- Rapid Per-VLAN Spanning Tree (RPVST+)

For detailed information on configuring these protocols, refer to [Chapter 24, “Configuring STP-Type Protocols”](#).

The Brocade VDX hardware handles Ethernet frames as follows:

- When the destination MAC address is not in the lookup table, the frame is flooded on all ports in the same VLAN, except the ingress port.
- When the destination MAC address is present in the lookup table, the frame is switched only to the correct egress port.
- When the destination MAC address is present in the lookup table, and the egress port is the same as the ingress port, the frame is dropped.
- If the Ethernet Frame Check Sequence (FCS) is incorrect, because the switch is in cut-through mode, a correctly formatted Ethernet frame is sent out with an incorrect FCS.
- If the Ethernet frame is too short, the frame is discarded and the error counter is incremented.
- If the Ethernet frame is too long, the frame is truncated and the error counter is incremented. The truncated frame is sent out with an incorrect FCS.
- Frames sent to a broadcast destination MAC address are flooded on all ports in the same VLAN, except the ingress port.
- When MAC address entries in the lookup table time out, they are removed. In this event, frame forwarding changes from unicast to flood.
- An existing MAC address entry in the lookup table is discarded when a device is moved to a new location. When a device is moved, the ingress frame from the new port causes the old lookup table entry to be discarded and the new entry to be inserted into the lookup table. Frame forwarding remains unicast to the new port.
- When the lookup table is full, new entries replace the oldest MAC addresses after the oldest MAC addresses reach a certain age and time out. MAC addresses that still have traffic running are not timed out.

---

**NOTE**

New entries start replacing older entries when the lookup table reaches 90 percent of its 32K capacity.

---

## VLAN tagging

The Layer 2 switch always tags an incoming frame with a VLAN ID. If the incoming frame is untagged, then a tag is added according to the port configuration. A port can classify untagged traffic to a single VLAN or to multiple VLANs. If the incoming frame is already tagged, then the port will either forward or discard the frame according to allowed VLAN rules in the port configuration.

These are three examples of VLAN tagging:

- If the DCB port is configured to tag incoming frames with a single VLAN ID, then incoming frames that are untagged are tagged with the VLAN ID.
- If the DCB port is configured to tag incoming frames with multiple VLAN IDs, then incoming frames that are untagged are tagged with the correct VLAN ID based on the port setting.
- If the DCB port is configured to accept externally tagged frames, then incoming frames that are tagged with a VLAN ID are passed through unchanged.

---

**NOTE**

Only a single switch-wide VLAN is capable of forwarding FCoE traffic.

---

For detailed information on configuring VLANs, refer to [Chapter 23, “Configuring VLANs”](#).

## Frame classification (incoming)

The Brocade VDX hardware is capable of classifying incoming Ethernet frames based on the following criteria:

- Port number
- Protocol
- MAC address

The classified frames can be tagged with a VLAN ID or with 802.1p Ethernet priority. The 802.1p Ethernet priority tagging is done using the Layer 2 Class of Service (CoS). The 802.1p Ethernet priority is used to tag frames in a VLAN with a Layer 2 CoS to prioritize traffic in the VLAN. The Brocade VDX hardware also accepts frames that have been tagged by an external device.

Frame classification options are as follows:

- VLAN ID and Layer 2 CoS by physical port number—With this option, the port is set to classify incoming frames to a preset VLAN ID and the Layer 2 CoS on a physical port on the Brocade VDX hardware.
- VLAN ID and Layer 2 CoS by LAG virtual port number—With this option, the port is set to classify incoming frames to a preset VLAN ID and Layer 2 CoS on a Link Aggregation Group (LAG) virtual port.
- Layer 2 CoS mutation—With this option, the port is set to change the Layer 2 CoS setting by enabling the QoS mutation feature.
- Layer 2 CoS trust—With this option, the port is set to accept the Layer 2 CoS of incoming frames by enabling the QoS trust feature.

For detailed information on configuring QoS, refer to [Chapter 29, “Configuring QoS”](#).

## Congestion control and queuing

The Brocade VDX hardware supports several congestion control and queuing strategies. As an output queue approaches congestion, Random Early Detection (RED) is used to selectively and proactively drop frames to maintain maximum link utilization. Incoming frames are classified into priority queues based on the Layer 2 CoS setting of the incoming frame, or the possible rewriting of the Layer 2 CoS field based on the settings of the DCB port or VLAN.

The Brocade VDX hardware supports a combination of two scheduling strategies to queue frames to the egress ports: Priority queuing, which is also referred to as strict priority, and Deficit Weighted Round Robin (DWRR) queuing.

The scheduling algorithms work on the eight traffic classes as specified in 802.1Qaz Enhanced Transmission Selection (ETS).

Queuing features are described as follows:

- RED—RED increases link utilization. When multiple inbound TCP traffic streams are switched to the same outbound port, and some traffic streams send small frames while other traffic streams send large frames, link utilization will not be able to reach 100 percent. When RED is enabled, link utilization approaches 100 percent.

- Classification—Setting user priority.
  - Inbound frames are tagged with the user priority set for the inbound port. The tag is visible when examining the frames on the outbound port. By default, all frames are tagged to priority zero.
  - Externally tagged Layer 2 frames—When the port is set to accept externally tagged Layer 2 frames, the user priority is set to the Layer 2 CoS of the inbound frames.
- Queuing
  - Input queuing—Input queuing optimizes the traffic flow in the following way. A DCB port has inbound traffic that is tagged with several priority values, and traffic from different priority settings is switched to different outbound ports. Some outbound ports are already congested with background traffic while others are uncongested. With input queuing, the traffic rate of the traffic streams switched to uncongested ports should remain high.
  - Output queuing—Output queuing optimizes the traffic flow in the following way. Several ports carry inbound traffic with different priority settings. Traffic from all ports is switched to the same outbound port. If the inbound ports have different traffic rates, some outbound priority groups will be congested while others can remain uncongested. With output queuing, the traffic rate of the traffic streams that are uncongested should remain high.
  - Multicast rate limit—A typical multicast rate limiting example is where several ports carry multicast inbound traffic that is tagged with several priority values. Traffic with different priority settings is switched to different outbound ports. The multicast rate limit is set so that the total multicast traffic rate on output ports is less than the specified set rate limit.
  - Multicast input queuing—A typical multicast input queuing example is where several ports carry multicast inbound traffic that is tagged with several priority values. Traffic with different priority settings is switched to different outbound ports. Some outbound ports are already congested with background traffic while others are uncongested. The traffic rate of the traffic streams switched to the uncongested ports should remain high. All outbound ports should carry some multicast frames from all inbound ports. This enables multicast traffic distribution relative to the set threshold values.
  - Multicast output queuing—A typical multicast output queuing example is where several ports carry multicast inbound traffic. Each port has a different priority setting. Traffic from all ports is switched to the same outbound port. If the inbound ports have varying traffic rates, some outbound priority groups will be congested while others remain uncongested. The traffic rate of the traffic streams that are uncongested remains high. The outbound ports should carry some multicast frames from all the inbound ports.
- Scheduling—A typical example of scheduling policy (using Strict Priority 0 and Strict Priority 1 modes) is where ports 0 through 7 carry inbound traffic, each port has a unique priority level, port 0 has priority 0, port 1 has priority 1, and so on. All traffic is switched to the same outbound port. In Strict Priority 0 mode, all ports have DWRR scheduling; therefore, the frames per second (FPS) on all ports should correspond to the DWRR settings. In Strict Priority 1 mode, priority 7 traffic uses Strict Priority; therefore, priority 7 can achieve a higher FPS. Frames from input ports with the same priority level should be scheduled in a round robin manner to the output port.

When setting the scheduling policy, each priority group that is using DWRR scheduling can be set to use a percentage of the total bandwidth by setting the PG\_Percentage parameter.

For detailed information on configuring QoS, refer to [Chapter 29, “Configuring QoS”](#).



## Access control

Access Control Lists (ACLs) are used for Layer 2 switching security. Standard ACLs inspect the source address for the inbound ports. Extended ACLs provide filtering by source and destination addresses and protocol. ACLs can be applied to the DCB ports or to VLANs.

ACLs function as follows:

- A standard Ethernet ACL configured on a physical port is used to permit or deny frames based on the source MAC address. The default is to permit all frames.
- An extended Ethernet ACL configured on a physical port is used to permit or deny frames based on the source MAC address, destination MAC address, and EtherType. The default is to permit all frames.
- A standard Ethernet ACL configured on a LAG virtual port is used to permit or deny frames based on the source MAC address. The default is to permit all frames. LAG ACLs apply to all ports in the LAG.
- An extended Ethernet ACL configured on a LAG virtual port is used to permit or deny frames based on the source MAC address, destination MAC address, and EtherType. The default is to permit all frames. LAG ACLs apply to all ports in the LAG.
- A standard Ethernet ACL configured on a VLAN is used to permit or deny frames based on the source MAC address. The default is to permit all frames. VLAN ACLs apply to the Switched Virtual Interface (SVI) for the VLAN.
- An extended Ethernet ACL configured on a VLAN is used to permit or deny frames based on the source MAC address, destination MAC address, and EtherType. The default is to permit all frames. VLAN ACLs apply to the Switched Virtual Interface (SVI) for the VLAN.

For detailed information on configuring ACLs, refer to [Chapter 28, “Configuring ACLs”](#).

## Trunking

---

### NOTE

The term “trunking” in an Ethernet network refers to the use of multiple network links (ports) in parallel to increase the link speed beyond the limits of any one single link or port, and to increase the redundancy for higher availability.

---

802.1ab Link Layer Discovery Protocol (LLDP) is used to detect links to connected switches or hosts. Trunks can then be configured between an adjacent switch or host and the Brocade VDX hardware.

The Data Center Bridging Capability Exchange Protocol (DCBX) extension is used to identify a DCB-capable port on an adjacent switch or host. For detailed information on configuring LLDP and DCBX, refer to [Chapter 27, “Configuring LLDP”](#).

The 802.3ad Link Aggregation Control Protocol (LACP) is used to combine multiple links to create a trunk with the combined bandwidth of all the individual links. For detailed information on configuring LACP, refer to [Chapter 26, “Configuring Link Aggregation”](#).

---

### NOTE

The Brocade software supports a maximum of 24 LAG interfaces.

---

## Flow control

802.3x Ethernet pause and Ethernet Priority-based Flow Control (PFC) are used to prevent dropped frames by slowing traffic at the source end of a link. When a port on a switch or host is not ready to receive more traffic from the source, perhaps due to congestion, it sends pause frames to the source to pause the traffic flow. When the congestion has been cleared, it stops requesting the source to pause traffic flow, and traffic resumes without any frame drop.

When Ethernet pause is enabled, pause frames are sent to the traffic source. Similarly, when PFC is enabled, there is no frame drop; pause frames are sent to the source switch.

For detailed information on configuring Ethernet pause and PFC, refer to [Chapter 29, “Configuring QoS”](#).

# FCoE Initialization Protocol

The FCoE Initialization Protocol (FIP) discovers and establishes virtual links between FCoE-capable entities connected to an Ethernet cloud through a dedicated EtherType, 0x8914, in the Ethernet frame.

## FIP discovery

---

### NOTE

This software version supports ANSI INCITS 462-2010 Fibre Channel – Backbone – 5 (FC-BB-5) / 13-May-2010.

---

The Brocade VDX hardware FIP discovery phase operates as follows:

- The Brocade VDX hardware uses the FCoE Initialization Protocol (FIP). ENodes discover VLANs supporting FCoE, FCFs, and then initialize the FCoE connection through the FIP.
- VF\_port configuration—An FCoE port accepts ENode requests when it is configured as a VF\_port and enabled. An FCoE port does not accept ENode requests when disabled.
- Solicited advertisements—A typical scenario is where a Brocade VDX hardware receives a FIP solicitation from an ENode. Replies to the original FIP solicitation are sent to the MAC address embedded in the original FIP solicitation. After being accepted, the ENode is added to the VN\_port table.
- VLAN 1—The Brocade VDX hardware should not forward FIP frames on VLAN 1 because it is reserved for management traffic only.
- A fabric-provided MAC address is supported. A server-provided MAC address is not supported in the Network OS v4.0.0 release.

---

### NOTE

In the fabric-provided MAC address format, VN\_port MAC addresses are based on a 48-bit fabric-supplied value. The first three bytes of this value are referred to as the FCMAP. The next three bytes are the FC ID, which is assigned by the switch when the ENode logs in to the switch.

---

## FIP login

FIP login operates as follows:

- ENodes can log in to the Brocade VDX hardware using FIP. Fabric login (FLOGI) and fabric discovery (FDISC) are accepted. Brocade VDX hardware in the fabric maintains the MAC address, World Wide Name (WWN), and PID mappings per login. Each ENode port should have a unique MAC address and WWN.
- FIP FLOGI—The Brocade VDX hardware accepts the FIP FLOGI from the ENode. The FIP FLOGI acceptance (ACC) is sent to the ENode if the ENode MAC address or WWN matches the VN\_port table on the Brocade VDX hardware. The FIP FLOGI request is rejected if the ENode MAC address or WWN does not match. The ENode login is added to the VN\_port table. Fabric Provided MAC Addressing (FPMA) is supported.
- FIP FDISC—The Brocade VDX hardware accepts FIP FDISC from the ENode. FIP FDISC acceptance (ACC) is sent to the ENode if the ENode MAC address or WWN matches the VN\_port table on the Brocade VDX hardware. The FIP FDISC request is rejected if the ENode MAC address or WWN does not match. The ENode login is added to the VN\_port table. FPMA is supported.
- Maximum logins per VF\_port—The Brocade VDX hardware supports a maximum of 64 logins per VF\_port. The VF\_port rejects further logins after the maximum is reached.
- Maximum logins per switch—The Brocade VDX hardware accepts a maximum of 1000 logins per switch.

## FIP logout

FIP logout operates as follows:

- ENodes and VN\_ports can log out from the Brocade VDX hardware using FIP. The Brocade VDX hardware in the fabric updates the MAC address, WWN, and PID mappings upon logout. The Brocade VDX hardware also handles scenarios of implicit logout where the ENode has left the fabric without explicitly logging out.
- FIP logout (LOGO)—The Brocade VDX hardware accepts a FIP LOGO from the ENode. The FIP LOGO acceptance (ACC) should be sent to the ENode if the ENode MAC address and the VN\_port MAC address matches the VN\_port table data on the switch. The LOGO is ignored (not rejected) if the ENode MAC address does not match. The ENode logout is updated in the VN\_port table.
- Implicit logout—With the ENode directly connected to a DCB port, if the port that the ENode is attached to goes offline, the Brocade VDX hardware implicitly logs out that ENode. ENode logout is updated in the VN\_port table. The Brocade VDX hardware sends an FIP Clear Virtual Links (CVL) to the ENode.

The FIP Virtual Link Maintenance protocols provide a mechanism to detect reachability loss to an ENode or any VN\_port instantiated on that ENode. This is accomplished by the periodic transmission of FIP Keep-Alive (FKA) messages from the ENode.

If FKA timeouts are enabled on the switch, all VN\_ports associated with an ENode will be implicitly logged out in the event of an ENode FKA timeout.

If FKA timeouts are enabled on the switch, the VN\_port will be implicitly logged out in the event of a VN\_port FKA timeout.

## Name server

The Brocade VDX hardware name server function operates as follows:

- ENode login and logout to and from the Brocade VDX hardware updates the name server in the FC fabric. The Brocade VDX hardware maintains the MAC address to WWN and PID mappings.
- ENode login and logout—When an ENode login occurs through any means (FIP FLOGI, FIP FDISC, FCoE FLOGI, or FCoE FDISC), an entry is added to the name server. When an ENode logout occurs through any means (FIP LOGO, FCoE LOGO, or implicit logout), the entry is removed from the name server.
- ENode data—The Brocade VDX hardware maintains a VN\_port table. The table tracks the ENode MAC address, FIP login parameters for each login from the same ENode, and WWN and PID mappings on the FC side. You can display the VN\_port table with the **show fcoe login** command.

## Registered State Change Notification

The Brocade VDX hardware Registered State Change Notification (RSCN) function operates as follows:

- RSCN events generated in the FC fabric are forwarded to the ENodes. RSCN events generated on the FCoE side are forwarded to the FC devices. DCB is not aware of RSCN events.
- Device RSCN—An RSCN is generated to all registered and affected members when an ENode either logs in or logs out of an FCF through any means. An RSCN is generated when an FC N\_port device either logs in or logs out of the FC fabric.

---

### NOTE

When transmitting an RSCN, zoning rules still apply for FCoE devices as the devices are treated as regular FC N\_ports.

---

- VF\_port RSCN—An RSCN is generated to all registered members when a VF\_port goes online or offline, causing ENode or FC devices to be added or removed.
- Domain RSCN—An RSCN is generated to all registered and affected members when an FC switch port goes online or offline, causing ENode or FC devices to be added or removed. An RSCN is generated when two FC switches merge or segment, causing ENode or FC devices to be added or removed. When FC switches merge or segment, an RSCN is propagated to ENodes.
- Zoning RSCN—An RSCN is generated to all registered and affected members when a zoning exchange occurs in the FC fabric.

## FCoE queuing

The QoS configuration controls the FCoE traffic distribution. Note that changing these settings requires changes on both the Brocade VDX hardware and the Converged Network Adapter (CNA); therefore, the link must be taken offline and put back online after a change is made. Traffic scheduler configuration changes affect FCoE traffic distribution as follows:

- Changing the priority group for a port causes the FCoE traffic distribution to be updated. The priority group and bandwidth are updated.
- Changing the priority table for a port causes the FCoE traffic distribution to be updated. The CoS-to-priority group mapping is updated.
- Changing the class map for a port causes the FCoE traffic distribution to be updated.
- Changing the policy map for a port causes FCoE traffic distribution to be updated.
- Changing the DCB map for a port causes the FCoE traffic distribution to be updated.
- The FCMAP-to-VLAN mapping determines the FCoE VLAN allowed for the FCoE session. Modifying this mapping causes the existing sessions to terminate.

---

**NOTE**

Only one FCoE VLAN is supported in the Network OS v4.0.0 release.

---

## Configuring FCoE interfaces

FCoE maps are used to configure FCoE properties on interfaces. An FCoE map is a placeholder for an FCoE VLAN and a CEE map. You will assign FCoE maps on to physical interfaces using the **fcoeport** command. Once the FCoE map is assigned onto an interface:

- The corresponding FCoE VLAN 1002 is applied to the interface.
- The corresponding CEE map is applied to the interface.
- The FCoE/FIP VLAN classifiers are applied to the interface.

In short, the interface becomes capable of carrying FCoE traffic. The FCoE map can be applied on an interface only if the FCoE map is complete in all aspects. That is, it should have an FCoE VLAN and a CEE map associated with it.

---

**NOTE**

Brocade does not support non-FCoE traffic over the FCoE VLAN. The FCoE VLAN should not carry any mixed traffic.

---

Only a single FCoE map is allowed, which is created automatically with the name “default.” You are not able to delete or rename this map. By default, the FCoE VLAN associated to the FCoE map is FCoE VLAN (1002) and the CEE map associated is the default CEE map (also called “default”).

### Assigning an FCoE map onto an interface

The FCoE map cannot be edited if it is associated with any interfaces.

The FCoE map can be applied, irrespective of whether or not the interface is in “switchport” mode. However, the FCoE map cannot be applied on an interface if the same interface already has a CEE map assigned to it.

To assign the FCoE map onto an interface, perform the following steps in global configuration mode.

1. Activate the interface configuration mode for the interface you wish to modify.

The following example activates the mode for the 10-gigabit Ethernet interface in slot 0/port 1.

```
switch(config)# interface tengigabitethernet 1/0/1
switch(conf-if-te-0/1)#
```

2. Apply the current FCoE profile map to the interface by using the **fcoeport** command.

```
switch(conf-if-te-0/1)# fcoeport default
```

3. Return to the privileged EXEC mode by using the **end** command.

```
switch(conf-if-te-0/1)# end
switch#
```

4. Confirm the changes to the interface by using the **show running-config** command.

```
switch# show running-config interface tengigabitethernet 0/1
interface TenGigabitEthernet 0/1
 fcoeport default
 no shutdown
```

5. Use the **show fcoe fabric-map default** command to confirm the current status of the FCoE map.

```
switch# show fcoe fabric-map default
=====
Fabric-Map VLAN VFID Pri FCMAP FKA Timeout
=====
default 1002[D] 128[D] 3[D] 0xefc00[D] 8000[D] Enabled[D]
Total number of Fabric Maps = 1
```

6. Repeat this procedure for any additional interfaces.

## Assigning an FCoE map onto a LAG member

The **fcoeport default** is a command under interface configuration mode used to provision a port to be an FCoE port. This puts the port in Layer 2 mode, but only for FCoE VLANs. Starting from Network OS v4.0.0, the **fcoeport default** command is supported for LAG member ports where FCoE provisioning is applied to individual 10-gigabit Ethernet ports.

You must apply the **fcoeport default** command on each LAG member interface. Once this command is applied, and if the member port of the LAG is CEE-capable, it carries FCoE-traffic only.

To assign the FCoE map onto a LAG member, perform the following steps in global configuration mode.

1. Activate the interface configuration mode for the interface you wish to modify.

The following example activates the mode for the 10-gigabit Ethernet interface in slot 0/port 0.

```
switch(config)# interface tengigabitethernet 3/0/19
switch(conf-if-te-3/0/19)#
```

2. Activate the channel-group mode.

```
switch(conf-if-te-3/0/19)# channel-group 10 mode active type standard
```

3. Set the LACP timeout to long.

```
switch(conf-if-te-3/0/19)# lacp timeout long
```

4. Apply the current FCoE profile map to the interface by using the **fcoeport** command.

```
switch(conf-if-te-3/0/19)# fcoeport default
```

5. Return to the privileged EXEC mode using the **end** command.

```
switch(conf-if-te-3/0/19)# end
```

6. Confirm the changes to the interface with the **show running-config** command.

```
switch# show running-config interface tengigabitethernet 3/0/19
interface TenGigabitEthernet 3/0/19
 fcoeport default
 shutdown
```

7. Use the **show fcoe interface brief** command to confirm the current status of the FCoE map.

```
switch# show fcoe interface brief
```

8. Repeat this procedure for any additional interfaces.

## FCoE over LAG

Network OS v4.0.0 supports FCoE over LAGs. These are LAGs between the FCoE Forwarder (FCF) and a DCB capable switch. The entire LAG is provisioned for FCoE, so that all member ports are used for FCoE traffic. FCoE traffic is broadcast on all the member links of the LAG.

---

### NOTE

FCoE over LAG supports standard LAGs only. vLAGs are not supported.

---

Additionally, Network OS v4.0.0 supports multiple logins per port. This feature allows multiple ENodes to login to a single 10-gigabit Ethernet port or a LAG.

## Configuration guidelines and restrictions

Follow these configuration guidelines and restrictions when configuring FCoE over LAG:

1. The intermediate switches may or may not be an FSB.
2. All ACLs and FCoE forwarding entries will continue to be on the FCF's ingress ports.
3. It is assumed that the intermediate switch works in "Willing" mode towards the FCF in the DCBX exchange, and accepts the configuration from the FCF and propagates it downstream.
4. The CEE/DCBX configuration is expected to be identical on both FCF and the intermediate switch.
5. Irrespective of item 3 or item 4, the PFC/No-drop behavior from the FCF perspective will be guaranteed only on the links between the FCF and first hop switch. There is no provision in the standard to guarantee this requirement on all paths leading to the Enode.
6. FSBs may or may not be able to forward the FCoE LLS TLV to the ENodes. Hence this TLV may not be present in the LLDP packets sent to the ENodes. The FCF continues to send this TLV in its LLDP packets destined to the intermediate switch.

## FCoE provisioning on LAGs

The existing **fcoeport default** command is extended to the LAG interfaces to support the new feature, as shown in the example below.

```
switch# configure
Entering configuration mode terminal
```

```
switch(config)# interface Port-channel 10
switch(config-Port-channel-10)# fcoeport default
switch(config-Port-channel-10)#
```

This provisions all the member ports of Port-channel 10 for FCoE.

Logical FCoE ports

When the switch boots, a pool of 256 FCoE ports are created. These ports are not bound to any physical ports. The binding are created when an FLOGI is received on the switch. Any free port that is available from the pool is selected and bound to the physical port where the FLOGI is received. The default number of logical ports is 256, and the range of valid values is from 256 though 1000.

TABLE 55 Number of FCoE ports per platform

| Platform            | Number of FCoE ports |
|---------------------|----------------------|
| Brocade VDX 6740    | 256                  |
| Brocade VDX 6740T   | 256                  |
| Brocade VDX 6720-24 | 256                  |
| Brocade VDX 6720-60 | 256                  |
| Brocade VDX 6710    | 256                  |
| Brocade VDX 8770-4  | 256                  |
| Brocade VDX 8770-8  | 256                  |
| Brocade VDX 8770-16 | 256                  |

When the FCoE logical port is automatically bound to a 10-gigabit Ethernet LAG port, it is referred to as dynamic binding. This binding is valid only until the FLOGI session is valid. The binding is automatically removed when CNA logs out. If you want to create a persistent binding between the logical FCoE port and the 10-gigabit or LAG port, use the **bind** command. This is stored in the configuration and retained across reboots.

NOTE

Only one type of binding can be used for each physical port, so the 10-gigabit or LAG binding configuration will overwrite each other.

To create additional logical FCoE ports, perform the following steps in global configuration mode.

1. Enter FCoE configuration mode.  

```
switch(config)# fcoe
```
2. Enter fabric-map configuration mode.  

```
switch(config-fcoe)# fabric-map default
```
3. Enter the **max-enodes** command to set the maximum number of logins allowed on the switch.  

```
switch(config-fcoe)# max-enodes 384
```

Optional: To bind the logical FCoE ports to a physical port, perform the following steps.

1. Exit FCoE configuration mode.  

```
switch(config-fcoe)# exit
```



2. Enter interface configuration mode.

```
switch(config)# interface fcoe 1/1/55
```

3. Bind the logical port to the physical port.

```
switch(config-Fcoe-1/1/55)# bind tengigabitethernet 1/0/1
```

## xSTP reconvergence

For topologies that have redundant LAGs between the intermediate-switch and the VCS (same or different FCFs), one of the LAGs will be in a xSTP-Blocked state. If one LAG fails for some reason, xSTP unblocks the other LAG to restore Layer 2 connectivity.

After a LAG failure, all ENodes are expected to logout and login back again if they discover an alternate path to the FCF. The availability of an alternate path depends on whether the second LAG is configured for FCoE or not. Also, this determines the time-taken for the logout of the ENodes during LAG failures.

If the alternate LAG is not configured for FCoE, the system responds as if the LAG does not exist.

If the alternate LAG is configured for FCoE, then the Unsolicited-advertisements continue to be sent by the FCF and the ENode keep-alives continue to reach the FCF. However, the FCF does not have any login sessions associated with the ENode. Consequently, a CVL is sent to the ENode as soon as the first keep-alive is received by the FCF. This clears the login session in the ENode and forces a re-login. For the default configuration, this happens within one FKA interval, about eight seconds. It may be slightly higher for faster configurations, as some packets (both RX and TX) are lost until STP unblocks the port.

If the two LAGs are connected to different FCFs in the same VCS then the ENodes login to a different domain and the PIDs are updated to reflect the change.

---

### NOTE

Brocade recommends that for faster recovery both the LAGs connecting the intermediate switch and the VCS be configured for FCoE, although only one of them will be operational. It is also recommended to configure faster advertisement intervals, in order to speed up the recovery process.

---

For additional information on xSTP, refer to [Chapter 24, "Configuring STP-Type Protocols"](#).



# Configuring VLANs

---

## In this chapter

|                                                          |     |
|----------------------------------------------------------|-----|
| • VLAN overview .....                                    | 331 |
| • Ingress VLAN filtering .....                           | 331 |
| • VLAN configuration guidelines and restrictions .....   | 333 |
| • Default VLAN configuration .....                       | 334 |
| • VLAN configuration and management .....                | 334 |
| • Configuring protocol-based VLAN classifier rules ..... | 338 |
| • Configuring the MAC address table .....                | 340 |
| • Private VLANs .....                                    | 341 |

## VLAN overview

IEEE 802.1Q Virtual LANs (VLANs) provide the capability to overlay the physical network with multiple virtual networks. VLANs allow you to isolate network traffic between virtual networks and reduce the size of administrative and broadcast domains.

A VLAN contains end stations that have a common set of requirements that are independent of physical location. You can group end stations in a VLAN even if they are not physically located in the same LAN segment. VLANs are typically associated with IP subnetworks and all the end stations in a particular IP subnet belong to the same VLAN. Traffic between VLANs must be routed. VLAN membership is configurable on a per interface basis.

The VLAN used for carrying FCoE traffic needs to be explicitly designated as the FCoE VLAN. FCoE VLANs are configured through the Network OS CLI (see [“Configuring an interface port as a Layer 2 switch port”](#) on page 336).

---

**NOTE**

Currently only one VLAN can be configured as the FCoE VLAN.

---

## Ingress VLAN filtering

A frame arriving at Brocade VDX hardware is either associated with a specific port or with a VLAN, based on whether the frame is tagged or untagged:

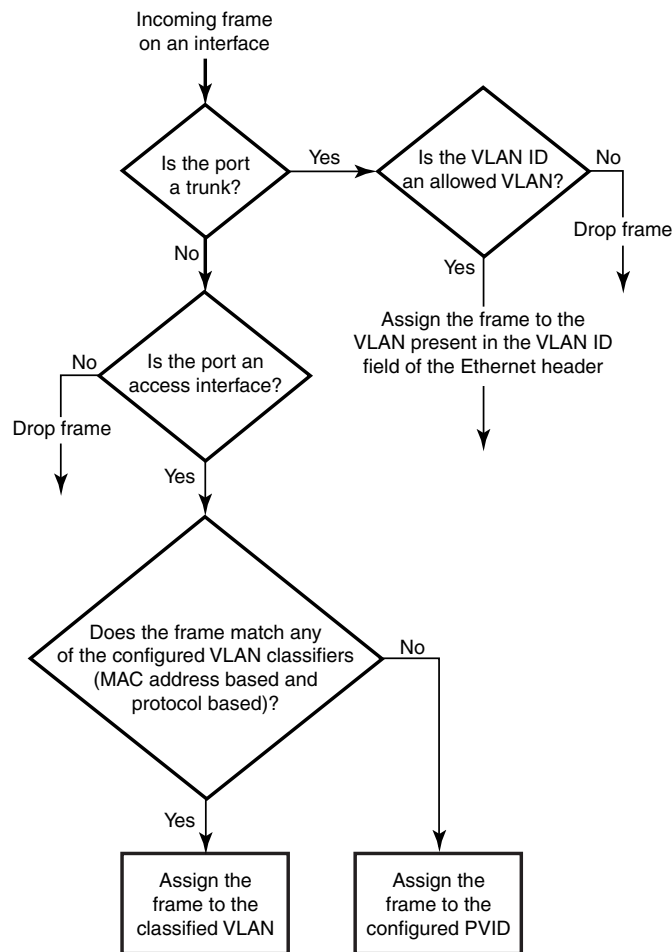
- Admit tagged frames only—The port the frame came in on is assigned to a single VLAN or to multiple VLANs depending on the VLAN ID in the frame’s VLAN tag. This is called trunk mode.
- Admit untagged frames only—These frames are assigned the port VLAN ID (PVID) assigned to the port the frame came in on. This is called access mode.

- Admit VLAN tagged and untagged frames—All tagged and untagged frames would be processed as follows:
  - All untagged frames are classified into native VLANs.
  - If the tengigabitethernet interface port is configured as an fcoeport and is in access mode, untagged Layer 2 or priority-tagged frames are forwarded by the egress port as untagged frames, unless you enable priority-tagging on the tengigabitethernet interface. By default, priority-tagging is disabled.
  - Any tagged frames coming with a VLAN tag equal to the configured native VLAN are processed.
  - For ingress and egress, non-native VLAN tagged frames are processed according to the allowed VLAN user specifications. This is called trunk mode.

**NOTE**

Ingress VLAN filtering is enabled by default on all Layer 2 interfaces. This ensures that VLANs are filtered on the incoming port (depending on the user configuration).

Figure 36 displays the frame processing logic for an incoming frame.



**FIGURE 36** Ingress VLAN filtering

There are important facts you should know about Ingress VLAN filtering:

- Ingress VLAN filtering is based on port VLAN membership.
- Port VLAN membership is configured through the Network OS CLI.
- Dynamic VLAN registration is not supported.
- The Brocade VDX hardware does VLAN filtering at both the ingress and egress ports.
- The VLAN filtering behavior on logical Layer 2 interfaces such as LAG interfaces is the same as on port interfaces.
- The VLAN filtering database (FDB) determines the forwarding of an incoming frame.

Additionally, there are important facts you should know about the VLAN FDB:

- The VLAN FDB contains information that helps determine the forwarding of an arriving frame based on MAC address and VLAN ID data. The FDB contains both statically configured data and dynamic data that is learned by the switch.
- The dynamic updating of FDB entries using learning is supported (if the port state permits).
- Dynamic FDB entries are not created for multicast group addresses.
- Dynamic FDB entries are aged out based on the aging time configured per Brocade VDX hardware. The aging time is between 60 and 1000000 seconds. The default is 300 seconds.
- You can add static MAC address entries specifying a VLAN ID. Static entries are not aged out.
- A static FDB entry overwrites an existing dynamically learned FDB entry and disables learning of the entry going forward.

---

**NOTE**

For more information on frame handling for Brocade VDX hardware, see [“Layer 2 Ethernet overview”](#) on page 317.

---

## VLAN configuration guidelines and restrictions

Follow these guidelines and restrictions when configuring VLANs:

- In an active topology, MAC addresses can be learned, per VLAN, using Independent VLAN Learning (IVL) only.
- A MAC address ACL always overrides a static MAC address entry. In this case, the MAC address is the forwarding address and the forwarding entry can be overwritten by the ACL.
- The Brocade DCB switch supports Ethernet DIX frames and 802.2 LLC SNAP encapsulated frames only.
- You must configure the same native VLAN on both ends of an 802.1q trunk link. Failure to do so can cause bridging loops and VLAN leaks.
- All switches in a fabric cluster or logical chassis cluster must be configured with the same VLAN number.

# Default VLAN configuration

Table 56 lists the default VLAN configuration.

| Parameter                 | Default setting                   |
|---------------------------|-----------------------------------|
| Default VLAN              | VLAN 1                            |
| Interface VLAN assignment | All interfaces assigned to VLAN 1 |
| VLAN state                | Active                            |
| MTU size                  | 2500 bytes                        |

# VLAN configuration and management

**NOTE**  
Enter the `copy running-config startup-config` command to save your configuration changes.

## Enabling and disabling an interface port

**NOTE**  
DCB interfaces are disabled by default in standalone mode, but enabled by default in Brocade VCS Fabric mode.

**NOTE**  
DCB interfaces do not support auto-negotiation of Ethernet link speeds. The DCB interfaces support 10-gigabit Ethernet and gigabit Ethernet.

To enable and disable an interface port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.  
  
The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8 in VCS mode. The prompt for these ports is in the format: `switch(config-if-gi-22/0/1)#`  
  
`switch(config)# interface tengigabitethernet 0/1`

## Configuring the MTU on an interface port

**NOTE**  
The entire fabric acts like a single switch. Therefore, MTU is applicable only on the edge-ports, and not on ISL.

To configure the maximum transmission unit (MTU) on an interface port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the interface port type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8 in VCS mode. The prompt for these ports is in the format: `switch(config-if-gi-22/0/1)#`

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **mtu** command to specify the MTU value on the interface port.

```
switch(conf-if-te-0/1)# mtu 4200
```

## Creating a VLAN

On Brocade VDX hardware, VLANs are treated as interfaces from a configuration point of view.

By default all the DCB ports are assigned to VLAN 1 (VLAN ID equals 1). The *vlan\_ID* value can be 1 through 3963. VLAN IDs 3964 through 4090 are internally-reserved VLAN IDs. However, the **reserved-vlan** command can modify this range. VLANs above 4090 are not configurable. Refer to the *Network OS Command Reference*.

To create a VLAN interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface vlan** command to assign the VLAN interface number.

```
switch(config)# interface vlan 1010
```

## Enabling STP on a VLAN

Once all of the interface ports have been configured for a VLAN, you can enable Spanning Tree Protocol (STP) for all members of the VLAN with a single command. Whichever protocol is currently selected is used by the VLAN. Only one type of STP can be active at a time.

A physical interface port can be a member of multiple VLANs. For example, a physical port can be a member of VLAN 1002 and VLAN 55 simultaneously. In addition, VLAN 1002 can have STP enabled and VLAN 55 can have STP disabled simultaneously.

To enable STP for a VLAN, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol spanning tree** command to select the type of STP for the VLAN.

```
switch(config)# protocol spanning tree mstp
```

3. Enter the **interface** command to select the VLAN interface number.

```
switch(config)# interface vlan 1002
```

4. Enter the **no spanning-tree shutdown** command to enable spanning tree on VLAN 1002.

```
switch(conf-if-vl-1002)# no spanning-tree shutdown
```

## Disabling STP on a VLAN

Once all of the interface ports have been configured for a VLAN, you can disable STP for all members of the VLAN with a single command.

To disable STP for a VLAN, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to select the VLAN interface number.  
`switch(config)# interface vlan 55`
3. Enter the **spanning-tree shutdown** command to disable spanning tree on VLAN 1002.  
`switch(config-if-vl-55)# spanning-tree shutdown`

## Configuring an interface port as a Layer 2 switch port

To configure the interface as a Layer 2 switch port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.  
The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8 in VCS mode. The prompt for these ports is in the format: `switch(config-if-gi-22/0/1)#`  
`switch(config)# interface tengigabitethernet 0/1`
3. Enter the **switchport** command to configure the interface as a Layer 2 switch port.
4. Enter the **do show** command to confirm the status of the DCB interface. For example  
`switch(config-if-te-0/1)# do show interface tengigabitethernet 0/1`
5. Enter the **do show** command to confirm the status of the DCB interface running configuration.  
`switch(config-if-te-0/1)# do show running-config interface tengigabitethernet 0/1`

## Configuring an interface port as an access interface

Each DCB interface port supports admission policies based on whether the frames are untagged or tagged. Access mode admits only untagged and priority-tagged frames.

To configure the interface as an access interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.  
The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8 in VCS mode. The prompt for these ports is in the format: `switch(config-if-gi-22/0/1)#`  
`switch(config)# interface tengigabitethernet 0/1`
3. Enter the **switchport** command to make the interface a Layer 2 switch port.  
`switch(config-if-te-0/1)# switchport`
4. Enter the **switchport** command again to configure the DCB interface as a VLAN.  
`switch(config-if-te-0/1)# switchport access vlan 20`



## Configuring an interface port as a trunk interface

Each DCB interface port supports admission policies based on whether the frames are untagged or tagged. Trunk mode admits only VLAN-tagged frames.

To configure the interface as a trunk interface, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8 in VCS mode. The prompt for these ports is in the format: `switch(config-if-gi-22/0/1)#`

```
switch(config)# interface tengigabitethernet 0/19
```

3. Enter the **switchport** command to place the DCB interface into trunk mode.

```
switch(conf-if-te-0/19)# switchport mode trunk
```

4. Specify whether all, one, or none of the VLAN interfaces are allowed to transmit and receive through the DCB interface. Enter the following command that is appropriate for your needs.

- This example allows the VLAN numbered as 30 to transmit/receive through the DCB interface:

```
switch(conf-if-te-0/19)# switchport trunk allowed vlan add 30
```

- To allow all VLANs to transmit/receive through the DCB interface:

```
switch(conf-if-te-0/19)# switchport trunk allowed vlan all
```

- This example allows all except VLAN 11 to transmit/receive through the DCB interface:

```
switch(conf-if-te-0/19)# switchport trunk allowed vlan except 11
```

- To allow none of the VLANs to transmit/receive through the DCB interface:

```
switch(conf-if-te-0/19)# switchport trunk allowed vlan none
```

## Disabling a VLAN on a trunk interface

To disable a VLAN on a trunk interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8 in VCS mode. The prompt for these ports is in the format: `switch(config-if-gi-22/0/1)#`

```
switch(config)# interface tengigabitethernet 0/10
```

3. Enter the **switchport** command to place the DCB interface into trunk mode.

```
switch(conf-if-te-0/10)# switchport mode trunk none
```

4. Enter the **switchport** command again to remove the VLAN ranges from the trunk port.

```
switch(conf-if-te-0/10)# switchport trunk allowed vlan remove 30
```

## Configuring protocol-based VLAN classifier rules

You can configure VLAN classifier rules to define specific rules for classifying frames to selected VLANs based on protocol and MAC addresses. Sets of rules can be grouped into VLAN classifier groups (see [“Creating a VLAN classifier group and adding rules”](#) on page 339).

VLAN classifier rules (1–256) are a set of configurable rules that reside in one of these categories:

- 802.1Q protocol-based classifier rules
- Source MAC address-based classifier rules
- Encapsulated Ethernet classifier rules

---

### NOTE

Multiple VLAN classifier rules can be applied per interface provided the resulting VLAN IDs are unique for the different rules.

---

802.1Q protocol-based VLANs apply only to untagged frames, or frames with priority tagging.

With both Ethernet-II and 802.2 SNAP encapsulated frames, the following protocol types are supported:

- Ethernet hexadecimal (0x0000 through 0xffff)
- Address Resolution Protocol (ARP)
- Fibre Channel over Ethernet (FCoE)
- FCoE Initialization Protocol (FIP)
- IP version 6 (IPv6)

---

### NOTE

For complete information on all available VLAN classifier rule options, see the *Network OS Command Reference*.

---

## Configuring a VLAN classifier rule

To configure a ARP protocol-based VLAN classifier rule, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **vlan classifier rule** command to configure a protocol-based VLAN classifier rule.

```
switch(config)# vlan classifier rule 1 proto ARP encap ethv2
```

---

### NOTE

See the *Network OS Command Reference* for complete information on all the protocols available for the **vlan classifier rule** command.

---

## Configuring MAC address-based VLAN classifier rules

To configure a MAC address-based VLAN classifier rule, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **vlan classifier rule** command to configure a MAC address-based VLAN classifier rule.

```
switch(config)# vlan classifier rule 5 mac 0008.744c.7fid
```

## Deleting a VLAN classifier rule

VLAN classifier groups (1 through 16) can contain any number of VLAN classifier rules.

To configure a VLAN classifier group and remove a VLAN classifier rule, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Specify a VLAN classifier group and delete a rule.

```
switch(config)# vlan classifier group 1 delete rule 1
```

## Creating a VLAN classifier group and adding rules

VLAN classifier groups (1 through 16) can contain any number of VLAN classifier rules.

To configure a VLAN classifier group and add a VLAN classifier rule, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Create a VLAN classifier group and add a rule.

```
switch(config)# vlan classifier group 1 add rule 1
```

## Activating a VLAN classifier group with an interface port

To associate a VLAN classifier group with an interface port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8 in VCS mode. The prompt for these ports is in the format: `switch(config-if-gi-22/0/1)#`

```
switch(config)# interface tengigabitethernet 0/10
```

3. Enter the **vlan classifier** command to activate and associate it with a VLAN interface (group 1 and VLAN 2 are used in this example).

```
switch(conf-if-te-0/10)# vlan classifier activate group 1 vlan 2
```

---

### NOTE

This example assumes that VLAN 2 was already created.

---

## Displaying VLAN information

---

### NOTE

The command **show vlan brief** displays the VLAN as inactive if there are no member ports associated to that VLAN, or if the ports associated are in an admin down state.

---

To display VLAN information, perform the following steps from privileged EXEC mode.

1. Enter the **show interface** command to display the configuration and status of the specified interface.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8 in VCS mode. The prompt for these ports is in the format: `switch(config-if-gi-22/0/1)#`

```
switch# show interface tengigabitethernet 0/10 port-channel 10 switchport
```

2. Enter the **show vlan** command to display the specified VLAN information. For example, this syntax displays the status of VLAN 20 for all interfaces, including static and dynamic:

```
switch# show vlan 20
```

## Configuring the MAC address table

Each DCB port has a MAC address table. The MAC address table stores a number of unicast and multicast address entries without flooding any frames. Brocade VDX hardware has a configurable aging timer. If a MAC address remains inactive for a specified number of seconds, it is removed from the address table. For detailed information on how the switch handles MAC addresses in a Layer 2 Ethernet environment, see [“Layer 2 Ethernet overview”](#) on page 317.

### Specifying or disabling the aging time for MAC addresses

You can set the length of time that a dynamic entry remains in the MAC address table after the entry is used or updated. Static address entries are never aged or removed from the table. You can also disable the aging time. The default is 300 seconds.

---

#### NOTE

To disable the aging time for MAC addresses, enter an aging time value of 0.

---

To specify an aging time or disable the aging time for MAC addresses, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the appropriate command based on whether you want to specify an aging time or disable the aging time for MAC addresses:

```
switch(config)# mac-address-table aging-time 600
```

### Adding static addresses to the MAC address table

To add a static address to the MAC address table, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Add the static address 0011.2222.3333 to the MAC address table with a packet received on VLAN 100:

```
switch(config)# mac-address-table static 0011.2222.3333 forward
tengigabitethernet 0/1 vlan 100
```

## Private VLANs

A private VLAN (PVLAN) domain is built with at least one pair of VLAN IDs; one (and only one) primary VLAN ID plus one or more secondary VLAN IDs. A primary VLAN is the unique and common VLAN identifier of the whole private VLAN domain and of all its VLAN ID pairs. Secondary VLANs can be configured as one of two types; either isolated VLANs or community VLANs. Only one isolated VLAN can be part of one PVLAN domain.

An isolated VLAN is a secondary VLAN whose distinctive characteristic is that all hosts connected to its ports are isolated at Layer 2. A community VLAN is a secondary VLAN that is associated to a group of ports that connect to a designated community of end devices with mutual trust relationships.

A PVLAN is often used to isolate networks from security attacks, or to simplify IP address assignments.

Within the private VLAN, ports can be assigned port types. A port can be assigned to only one kind of port type at a time. The types of ports available for private VLANs are described in [Table 57](#).

**TABLE 57 Private VLAN port types**

| Port type        | Description                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Isolated port    | An isolated port cannot talk to any other port in the private VLAN domain except for promiscuous ports and traffic ports. If a customer device needs to have access only to a gateway router, then it should be attached to an isolated port.                                                                                                                                               |
| Community port   | A community port is part of a group of ports that have Layer 2 communications with one another, and can also talk to any promiscuous port. For example, if you have two devices that you want to be isolated from other devices, but still be able to communicate between themselves, then community ports should be used. You can not configure multiple community VLANs on a single port. |
| Promiscuous port | A promiscuous port can talk to all other types of ports. A promiscuous port can talk to isolated ports as well as community ports. Layer 3 gateways, DHCP servers, and other trusted devices that need to communicate with the customer endpoints are typically connected using promiscuous ports.                                                                                          |
| Trunk port       | A trunk port is a regular port that connects two switches and carries two or more VLANs.                                                                                                                                                                                                                                                                                                    |

## PVLAN configuration guidelines and restrictions

Follow these guidelines and restrictions when configuring VLANs:

- VxLAN configuration is not supported on private VLAN.
- IGMP is not supported on private VLANs, however you can create an IGMP configuration. The configuration succeeds but the hardware is not programmed.
- For private VLANs, egress ACLs on the primary VLAN are applied only for the traffic that ingresses and egresses from the primary VLAN, and not for the traffic that gets translated from the secondary VLAN to the primary VLAN.
- For private VLANs, egress ACLs on the primary VLAN are also applied to the traffic that gets translated to the secondary VLAN.

## Configuring a private VLAN

This procedure configures the PVLAN and associates the secondary VLAN with the primary VLAN.

1. Configure the VLAN interface.

```
switch(config)#interface vlan 10
```

2. Configure the VLAN as a primary PVLAN.

```
switch(config-if-vl-10)# private-vlan primary
```

3. Configure the secondary VLAN (community).

```
switch(config)#interface vlan 100
switch(config-if-vl-100)# private-vlan community
```

4. Configure the secondary VLAN (isolated).

```
switch(config)#interface vlan 200
switch(config-if-vl-200)# private-vlan isolated
```

5. Associate the secondary VLAN with the primary VLAN. The list can contain one isolated VLAN and multiple community VLAN.

```
switch(config)#interface vlan 10
switch(config-if-vl-10)# private-vlan association add 100
```

6. Exit VLAN configuration mode.

```
switch(config-if-vl-10)# exit
```

## Configuring a interface as a PVLAN promiscuous port

This procedure configures an interface as the PVLAN promiscuous port.

1. Specify the interface.

```
switch(config)#interface tengigabitethernet 0/1
```

2. Mark the interface as switch port

```
switch(config-if-te-0/1)#switchport
```

3. Configure the interface as a PVLAN promiscuous port (untagged).

```
switch(config-if-te-0/1)# switchport mode private-vlan promiscuous
```

4. Configure the interface as a PVLAN promiscuous port (tagged).

```
switch(config-if-te-0/1)# switchport mode private-vlan trunk promiscuous
```

5. Associate the interface with a PVLAN.

```
switch(config-if-te-0/1)# switchport private-vlan mapping add 10 100,200
```

6. Configure a normal VLAN on the PVLAN promiscuous port.

```
switch(config-if-te-0/1)# switchport trunk allowed vlan add 500
```

## Configuring a interface as a PVLAN host port

This procedure configures an interface as the PVLAN host port.

1. Specify the interface.

```
switch(config)#interface tengigabitethernet 0/1
```

2. Mark the interface as a switch port.

```
switch(config-if-te-0/1)# switchport
```

3. Configure the interface as a PVLAN host port that is tagged.

```
switch(config-if-te-0/1)# switchport mode private-vlan trunk host
```

4. Associate the interface with a PVLAN.

```
switch(config-if-te-0/1)# switchport private-vlan host-association 10 100
```

## Configuring an interface as a PVLAN trunk port

This procedure configures an interface as a PVLAN trunk port.

1. Specify the interface.

```
switch(config)#interface tengigabitethernet 0/1
```

2. Mark the interface as switch port

```
switch(config-if-te-0/1)#switchport
```

3. Configure the interface as a PVLAN trunk port.

```
switch(config-if-te-0/1)# switchport mode private-vlan trunk
```

4. Configure the association between primary VLANs and secondary VLANs and the PVLAN trunk port with a PVLAN.

---

### NOTE

Multiple PVLAN pairs can be specified using this command so that a PVLAN trunk port can carry multiple secondary VLANs. If an association is specified for the existing primary VLAN, the existing association is replaced. If there is no trunk association, any packets received on secondary VLANs are dropped.

---

```
switch(config-if-te-0/1)# switchport private-vlan association trunk 10 100
```

5. Configure a normal VLAN on the PVLAN trunk port.

```
switch(config-if-te-0/1)# switchport private-vlan trunk allowed vlan add 100
```

6. Configures a VLAN to which untagged packets (as in IEEE 802.1Q tagging) are assigned on a PVLAN trunk port. If there is no native VLAN configured, all untagged packets are dropped. If the native VLAN is a secondary VLAN and the port does not have the association for the secondary VLAN, the untagged packets are dropped.

```
switch(config-if-te-0/1)# switchport private-vlan trunk native vlan 100
```

## Displaying PVLAN information

To display private VLAN information, use the **show vlan private-vlan** command.





# Configuring STP-Type Protocols

---

## In this chapter

|                                                               |     |
|---------------------------------------------------------------|-----|
| • STP overview .....                                          | 345 |
| • Configuration guidelines and restrictions .....             | 348 |
| • RSTP overview .....                                         | 349 |
| • MSTP overview .....                                         | 351 |
| • Overview of PVST+ and Rapid PVST+ .....                     | 353 |
| • Default Spanning Tree Protocol configuration .....          | 354 |
| • Spanning Tree Protocol configuration and management .....   | 355 |
| • Configuring STP, RSTP, or MSTP on DCB interface ports ..... | 363 |
| • Spanning Tree Protocol and VCS mode .....                   | 371 |

## STP overview

The IEEE 802.1D Spanning Tree Protocol (STP) runs on bridges and switches that are 802.1D-compliant. STP prevents loops in the network by providing redundant links. If a primary link fails, the backup link is activated and network traffic is not affected. Without STP running on the switch or bridge, a link failure can result in a loop.

When the spanning tree algorithm is run, the network switches transform the real network topology into a spanning tree topology in which any LAN in the network can be reached from any other LAN through a unique path. The network switches recalculate a new spanning tree topology whenever there is a change to the network topology.

---

### NOTE

All Brocade VDX switches that are operating in standalone mode need to have some version of xSTP configured in order to avoid VLAN looping issues.

---

For each LAN, the switches that attach to the LAN choose a designated switch that is the closest switch to the root switch. This designated switch is responsible for forwarding all traffic to and from the LAN. The port on the designated switch that connects to the LAN is called the designated port.

The switches decide which of their ports will be part of the spanning tree. A port is included in the spanning tree if it is a root port or a designated port.

With STP, data traffic is allowed only on those ports that are part of the spanning tree topology. Ports that are not part of the spanning tree topology are automatically changed to a blocking (inactive) state. They are kept in the blocking state until there is a break in the spanning tree topology, at which time they are automatically activated to provide a new path.

The STP interface states for every Layer 2 interface running STP are as follows:

- Blocking—The interface does not forward frames.
- Listening—The interface is identified by the spanning tree as one that should participate in frame forwarding. This is a transitional state after the blocking state.
- Learning—The interface prepares to participate in frame forwarding.
- Forwarding—The interface forwards frames.
- Disabled—The interface is not participating in spanning tree because of a shutdown port, no link on the port, or no spanning tree instance running on the port.

A port participating in spanning tree moves through these states:

- From initialization to blocking.
- From blocking to listening or to disabled.
- From listening to learning or to disabled.
- From learning to forwarding, blocking, or disabled.
- From forwarding to disabled.

The following STP features are considered optional features although you might use them in your STP configuration:

- Root guard—For detailed information, see [“Enabling guard root”](#) on page 365.
- Port fast BPDU guard and BPDU filter—For detailed information, see [“Enabling port fast \(STP\)”](#) on page 367.

## Configuring STP

---

### NOTE

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in this format:

```
switch(config-if-gi-22/0/1)#
```

---

The process for configuring STP is as follows.

1. Enter global configuration mode.
2. Enable STP by using the global **protocol spanning-tree** command. For details, see [“Enabling STP, RSTP, MSTP, R-PVST+ or PVST+”](#) on page 355.  
  

```
switch(config)# protocol spanning-tree stp
```
3. Designate the root switch by using the **bridge-priority** command. For details, see [“Specifying the bridge priority”](#) on page 356. The range is 0 through 61440 and the priority values can be set only in increments of 4096.  
  

```
switch(conf-stp)# bridge-priority 28672
```
4. *Optional:* Enable port fast on switch ports by using the **spanning-tree portfast** command. For details, see [“Enabling port fast \(STP\)”](#) on page 367.

---

### NOTE

Port fast only needs to be enabled on ports that connect to workstations or PCs. Repeat these commands for every port connected to workstations or PCs. Do not enable port fast on ports that connect to other switches.

---

**NOTE**

If BPDUs are received on a port fast-enabled interface, the interface loses the edge port status unless it receives a shut/no shut.

**NOTE**

Enabling port fast on ports can cause temporary bridging loops, in both trunking and non-trunking mode.

```
switch(config)# interface tengigabitethernet 0/10
switch(config-if-te-0/10)# spanning-tree portfast
switch(config-if-te-0/10)# exit
switch(config)# interface tengigabitethernet 0/11
switch(config-if-te-0/11)# spanning-tree portfast
switch(config-if-te-0/11)# exit
```

Repeat these commands for every port connected to workstations or PCs.

5. *Optional:* To interoperate with non-Brocade switches in PVST+/R-PVST+ mode, you may need to configure the interface that is connected to that switch by using the following **spanning-tree bpdumac** command.

```
switch(config)# interface tengigabitethernet 0/12
switch(config-if-te-0/12)# spanning-tree bpdumac 0100.0ccc.cccd
```

6. Specify port priorities by using the **spanning-tree priority** command to influence the selection of root/designated ports on the following:
  - All ports of the root switch
  - The root port
  - The designated port
7. *Optional:* Enable the guard root feature with the **spanning-tree guard root** command. The guard root feature provides a way to enforce the root bridge placement in the network. For detailed information, refer to [“Enabling guard root”](#) on page 365.

All other switch ports connect to other switches and bridges are automatically placed in blocking mode.

This does not apply to ports connected to workstations or PCs; these ports remain in the forwarding state.

8. Return to privileged EXEC mode.

```
switch(config-if-te-0/12)# end
```

9. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

When the spanning tree topology is completed, the network switches send and receive data only on the ports that are part of the spanning tree. Data received on ports that are not part of the spanning tree is blocked.

**NOTE**

Brocade recommends leaving other STP variables at their default values.

For more information on STP, see [“Spanning Tree Protocol configuration and management”](#) on page 355.

## Configuration guidelines and restrictions

Follow these configuration guidelines and restrictions when configuring STP:

- You have to disable one form of xSTP before enabling another.
- Packet drops or packet flooding may occur if you do not enable xSTP on all devices connected on both sides of parallel links.
- Network OS switches in logical chassis cluster mode or fabric cluster mode drop all tagged xSTP frames that originate from the Brocade MLX-MCT.
- LAGs are treated as normal links, and by default are enabled for STP.
- You can have 32 MSTP instances and one MSTP region.
- Create VLANs before mapping them to MSTP instances.
- The MSTP force-version option is not supported.
- When a misconfigured local area network running spanning tree has one or more loops, a traffic storm of spanning tree BPDUs can occur. In certain circumstances, VDX can reboot when subjected to an extended period of traffic storm involving spanning tree BPDUs.
- Additionally, when a misconfigured local area network running spanning tree has one or more loops, a traffic storm of spanning tree BPDUs can occur. Edge Loop Detection (ELD) protocol cannot eliminate loops during a traffic storm involving control packets, such as spanning tree BPDUs.
- Do not force an alternate root path through root path cost with PVST+ or R-PVST+ on legacy Foundry equipment, such as the Brocade NetIron MLX or Brocade Turbolron. This can cause traffic issues on the network.
- For load balancing across redundant paths in the network to work, all VLAN-to-instance mapping assignments must match; otherwise, all traffic flows on a single link.
- When you enable MSTP by using the global **protocol spanning-tree mstp** command, RSTP is automatically enabled.
- For two or more switches to be in the same MSTP region, they must have the same VLAN-to-instance map, the same configuration revision number, and the same name.
- Spanning tree topologies must not be enabled on any direct server connections to the front-end 10-gigabit Ethernet ports that may run FCoE traffic. This may result in lost or dropped FCoE logins.
- The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:  

```
switch(config-if-gi-22/0/1)#
```

## RSTP overview

### NOTE

RSTP is designed to be compatible and interoperate with STP. However, the advantages of the RSTP fast convergence are lost when it interoperates with switches running STP.

The IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) standard is an evolution of the 802.1D STP standard. It provides rapid convergence following the failure of a switch, a switch port, or a LAN. It provides rapid convergence of edge ports, new root ports, and ports connected through point-to-point links.

The RSTP interface states for every Layer 2 interface running RSTP are as follows:

- Learning—The interface prepares to participate in frame forwarding.
- Forwarding—The interface forwards frames.
- Discarding—The interface discards frames. Note that the 802.1D disabled, blocking, and listening states are merged into the RSTP discarding state. Ports in the discarding state do not take part in the active topology and do not learn MAC addresses.

Table 58 lists the interface state changes between STP and RSTP.

**TABLE 58 STP versus RSTP state comparison**

| STP interface state | RSTP interface state | Is the interface included in the active topology? | Is the interface learning MAC addresses? |
|---------------------|----------------------|---------------------------------------------------|------------------------------------------|
| Disabled            | Discarding           | No                                                | No                                       |
| Blocking            | Discarding           | No                                                | No                                       |
| Listening           | Discarding           | Yes                                               | No                                       |
| Learning            | Learning             | Yes                                               | Yes                                      |
| Forwarding          | Forwarding           | Yes                                               | Yes                                      |

With RSTP, the port roles for the new interface states are also different. RSTP differentiates explicitly between the state of the port and the role it plays in the topology. RSTP uses the root port and designated port roles defined by STP, but splits the blocked port role into backup port and alternate port roles:

- Backup port—Provides a backup for the designated port and can only exist where two or more ports of the switch are connected to the same LAN; the LAN where the bridge serves as a designated switch.
- Alternate port—Serves as an alternate port for the root port providing a redundant path towards the root bridge.

Only the root port and the designated ports are part of the active topology; the alternate and backup ports do not participate in it.

When the network is stable, the root and the designated ports are in the forwarding state, while the alternate and backup ports are in the discarding state. When there is a topology change, the new RSTP port roles allow a faster transition of an alternate port into the forwarding state.

For more information on RSTP, see [“Spanning Tree Protocol configuration and management”](#) on page 355.

## Configuring RSTP

The basic process for configuring RSTP is as follows.

1. Enter global configuration mode.
2. Enable RSTP by using the global **protocol spanning-tree** command.  
For details, see [“Enabling STP, RSTP, MSTP, R-PVST+ or PVST+”](#) on page 355.

```
switch(config)# protocol spanning-tree rstp
```

3. Designate the root switch by using the **bridge-priority** command.  
For more details, see [“Specifying the bridge priority”](#) on page 356. The range is 0 through 61440 and the priority values can be set only in increments of 4096.

```
switch(config-stp)# bridge-priority 28582
```

4. Configure the **bridge forward delay** value.  
For more details, see [“Specifying the bridge forward delay”](#) on page 357.

```
switch(config-stp)# forward-delay 20
```

5. Configure the **bridge maximum aging time** value.  
For more details, see [“Specifying the bridge maximum aging time”](#) on page 357.

```
switch(config-stp)# max-age 25
```

6. Enable the **error disable timeout timer** value.  
For more details, see [“Enabling the error disable timeout timer”](#) on page 358.

```
switch(config-stp)# error-disable-timeout enable
```

7. Configure the **error-disable-timeout** interval value.  
For more details, see [“Specifying the error disable timeout interval”](#) on page 358.

```
switch(config-stp)# error-disable-timeout interval 60
```

8. Configure the port-channel path cost.  
For more details, see [“Specifying the port-channel path cost”](#) on page 358.

```
switch(config-stp)# port-channel path-cost custom
```

9. Configure the bridge hello-time value.  
For more details, see [“Specifying the bridge hello time”](#) on page 359.

```
switch(config-stp)# hello-time 5
```

10. *Optional:* Enable port fast on switch ports by using the **spanning-tree portfast** command.  
For more details, see [“Enabling port fast \(STP\)”](#) on page 367.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

---

### NOTE

Port fast only needs to be enabled on ports that connect to workstations or PCs. Repeat these commands for every port connected to workstations or PCs. Do not enable port fast on ports that connect to other switches.

---

**NOTE**

Enabling port fast on ports can cause temporary bridging loops, in both trunking and nontrunking mode.

```
switch(config)# interface tengigabitethernet 0/10
switch(config-if-te-0/10)# spanning-tree portfast
switch(config-if-te-0/10)# exit
switch(config)# interface tengigabitethernet 0/11
switch(config-if-te-0/11)# spanning-tree portfast
switch(config-if-te-0/11)# exit
switch(config)#
```

Repeat these commands for every port connected to workstations or PCs.

11. Specify port priorities by using the **spanning-tree priority** command to influence the selection of root/designated ports on the following:

- All ports of the root switch
- The root port
- The designated port

For details, see [“Specifying the port priority”](#) on page 368.

12. *Optional:* Enable the guard root feature with the **spanning-tree guard root** command. The guard root feature provides a way to enforce the root bridge placement in the network. For detailed information, refer to [“Enabling guard root”](#) on page 365.

All other switch ports connected to other switches and bridges are automatically placed in blocking mode.

This does not apply to ports connected to workstations or PCs; these ports remain in the forwarding state.

13. Return to privileged EXEC mode.

```
switch(config)# end
```

14. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## MSTP overview

The IEEE 802.1s Multiple STP (MSTP) helps create multiple loop-free active topologies on a single physical topology. MSTP enables multiple VLANs to be mapped to the same spanning tree instance (forwarding path), which reduces the number of spanning tree instances needed to support a large number of VLANs. Each MSTP instance has a spanning tree topology independent of other spanning tree instances. With MSTP you can have multiple forwarding paths for data traffic. A failure in one instance does not affect other instances. With MSTP, you are able to more effectively utilize the physical resources present in the network and achieve better load balancing of VLAN traffic.

**NOTE**

In MSTP mode, RSTP is automatically enabled to provide rapid convergence.

Multiple switches must be configured consistently with the same MSTP configuration to participate in multiple spanning tree instances. A group of interconnected switches that have the same MSTP configuration is called an MSTP region.

**NOTE**

Brocade supports 32 MSTP instances and one MSTP region.

MSTP introduces a hierarchical way of managing switch domains using regions. Switches that share common MSTP configuration attributes belong to a region. The MSTP configuration determines the MSTP region where each switch resides. The common MSTP configuration attributes are as follows:

- Alphanumeric configuration name (32 bytes)
- Configuration revision number (2 bytes)
- 4096-element table that maps each of the VLANs to an MSTP instance

Region boundaries are determined based on the above attributes. A multiple spanning tree instance is an RSTP instance that operates inside an MSTP region and determines the active topology for the set of VLANs mapping to that instance. Every region has a common internal spanning tree (CIST) that forms a single spanning tree instance that includes all the switches in the region. The difference between the CIST instance and the MSTP instance is that the CIST instance operates across the MSTP region and forms a loop-free topology across regions, while the MSTP instance operates only within a region. The CIST instance can operate using RSTP if all the switches across the regions support RSTP. However, if any of the switches operate using 802.1D STP, the CIST instance reverts to 802.1D. Each region is viewed logically as a single STP/RSTP bridge to other regions.

## Configuring MSTP

The basic process for configuring MSTP is as follows.

1. Enter global configuration mode.
2. Enable MSTP by using the global **protocol spanning-tree** command.  
For more details see [“Enabling STP, RSTP, MSTP, R-PVST+ or PVST+”](#) on page 355.  

```
switch(config)# protocol spanning-tree mstp
```
3. Specify the region name by using the **region** *region\_name* command.  
For more details see [“Specifying a name for an MSTP region”](#) on page 362.  

```
switch(config-mstp)# region brocade1
```
4. Specify the revision number by using the **revision** command.  
For more details see [“Specifying a revision number for MSTP configuration”](#) on page 362.  

```
switch(config-mstp)# revision 1
```
5. Map a VLAN to an MSTP instance by using the **instance** command.  
For more details see [“Mapping a VLAN to an MSTP instance”](#) on page 361.  

```
switch(config-mstp)# instance 1 vlan 2, 3
switch(config-mstp)# instance 2 vlan 4-6
switch(config-mstp)# instance 1 priority 4096
```
6. Specify the maximum hops for a BPDU to prevent the messages from looping indefinitely on the interface by using the **max-hops** *hop\_count* command. For more details see [“Specifying the maximum number of hops for a BPDU \(MSTP\)”](#) on page 361.  

```
switch(config-mstp)# max-hops 25
```
7. Return to privileged EXEC mode.  

```
switch(config)# end
```
8. Enter the **copy** command to save the *running-config* file to the *startup-config* file.



```
switch# copy running-config startup-config
```

For more information on MSTP, see [“Spanning Tree Protocol configuration and management”](#) on page 355.

## Overview of PVST+ and Rapid PVST+

A network topology of bridges typically contains redundant connections to provide alternate paths in case of link failures. However, because there is no concept of TTL in Ethernet frames, this could result in permanent circulation of frames if there are loops in the network. To prevent loops, a spanning tree connecting all the bridges is formed in real time. The redundant ports are put in a blocking (nonforwarding) state. They are enabled when required.

In order to build a spanning tree for the bridge topology, the bridges must exchange control frames (BPDUs – Bridge Protocol Data Units). The protocols define the semantics of the BPDUs and the required state machine. The first Spanning Tree Protocol (STP) became part of the IEEE 802.1d standard.

Because the convergence time of STP is 50 seconds in the case of link failures, this delay soon became increasingly unacceptable. Keeping the main skeleton of STP the same, the state machine was changed to speed up the convergence time as part of the Rapid Spanning Tree protocol (RSTP). RSTP became part of the standard IEEE 802.1w.

Both STP and RSTP build a single logical topology. A typical network has multiple VLANs. A single logical topology does not efficiently utilize the availability of redundant paths for multiple VLANs. If a port is set to “blocked/discarding” for one VLAN (under STP/RSTP), it is the same for all other VLANs too.

Per-VLAN Spanning Tree Plus (PVST+) protocol runs a spanning tree instance for each VLAN in the network. The version of PVST+ that uses the RSTP state machine is called Rapid-PVST Plus (R-PVST+). R-PVST+ has one instance of spanning tree for each VLAN on the switch.

PVST+ is not a scalable model when there are many VLANs in the network, as it consumes a lot of CPU power. A reasonable compromise between the two extremes of RSTP and R-PVST+ is the Multiple Spanning Tree protocol (MSTP), which was standardized as IEEE 802.1s and later incorporated into the IEEE 802.1Q-2003 standard. MSTP runs multiple instances of spanning tree that are independent of VLANs. It then maps a set of VLANs to each instance.

---

### NOTE

Brocade Network OS v4.0.0 supports PVST+ and R-PVST+ only. The PVST and R-PVST protocols are proprietary to Cisco and are not supported.

---

To configure PVST+ or R-PVST+, use the **protocol spanning-tree pvst** and **protocol spanning-tree rpvt** commands. See the *Network OS Command Reference* for details.

For example, the script below sets up PVST+ for VLAN 10:

```
switch(config)# protocol spanning-tree pvst
switch(conf-pvst)# bridge-priority 4096
switch(conf-pvst)# forward-delay 4
switch(conf-pvst)# hello-time 2
switch(conf-pvst)# max-age 7
```

## PVST+ and R-PVST+ guidelines and restrictions

Consider the following items when configuring PVST+ and R-PVST+:

- Disabling the tagging of native VLANs is required on STP/RSTP/MSTP switches in standalone mode, otherwise PVST+/R-PVST+ does not converge and forms a loop on the native VLAN. The tagged native VLAN data traffic is ignored. The native VLAN untagged data is forwarded.
- Disabling the tagging of native VLANs is required on edge ports in fabric cluster mode; otherwise, PVST+/R-PVST+ does not converge and forms a loop on the native VLAN. The tagged native VLAN data traffic is ignored. The native VLAN untagged data is forwarded.
- If a VLAN is configured with tagged ports that do not have PVST+ mode enabled on the interface and are connected to the VDXs, and RSTP is enabled under the VLAN (PVST+), then BPDUs from the tagged ports that are received by the VDX are dropped.

## Default Spanning Tree Protocol configuration

Table 59 lists the default Spanning Tree Protocol configuration.

**TABLE 59** Default Spanning Tree Protocol configuration

| Parameter                      | Default setting                              |
|--------------------------------|----------------------------------------------|
| Spanning-tree mode             | By default, STP, RSTP, and MSTP are disabled |
| Bridge priority                | 32768                                        |
| Bridge forward delay           | 15 seconds                                   |
| Bridge maximum aging time      | 20 seconds                                   |
| Error disable timeout timer    | Disabled                                     |
| Error disable timeout interval | 300 seconds                                  |
| Port-channel path cost         | Standard                                     |
| Bridge hello time              | 2 seconds                                    |

Table 60 lists the switch defaults that apply only to MSTP configurations.

**TABLE 60** Default MSTP configuration

| Parameter                                                 | Default setting |
|-----------------------------------------------------------|-----------------|
| Cisco interoperability                                    | Disabled        |
| Switch priority (when mapping a VLAN to an MSTP instance) | 32768           |
| Maximum hops                                              | 20 hops         |
| Revision number                                           | 0               |

Table 61 lists the switch defaults for the 10-gigabit Ethernet DCB interface-specific configuration.

**TABLE 61** Default 10-gigabit Ethernet DCB interface-specific configuration

| Parameter                | Default setting           |
|--------------------------|---------------------------|
| Spanning tree            | Disabled on the interface |
| Automatic edge detection | Disabled                  |

**TABLE 61** Default 10-gigabit Ethernet DCB interface-specific configuration (Continued)

| Parameter                      | Default setting                                |
|--------------------------------|------------------------------------------------|
| Path cost                      | 2000                                           |
| Edge port                      | Disabled                                       |
| Guard root                     | Disabled                                       |
| Hello time                     | 2 seconds                                      |
| Link type                      | Point-to-point                                 |
| Port fast                      | Disabled                                       |
| Port priority                  | 128                                            |
| DCB interface root port        | Allow the DCB interface to become a root port. |
| DCB interface BPDU restriction | Restriction is disabled.                       |

## Spanning Tree Protocol configuration and management

### NOTE

Enter the **copy running-config startup-config** command to save your configuration changes.

### Enabling STP, RSTP, MSTP, R-PVST+ or PVST+

You enable STP to detect or avoid loops. STP is not required in a loop-free topology. You must turn off one form of STP before turning on another form. By default, STP, RSTP, and MSTP are not enabled.

Perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, MSTP, PVST+, or R-PVST+.

```
switch(config)# protocol spanning-tree rstp
```

### Disabling STP, RSTP, or MSTP

### NOTE

Using the **no protocol spanning-tree** command deletes the context and all the configurations defined within the context or protocol for the interface.

To disable STP, RSTP, or MSTP, perform the following steps from privileged EXEC mode. By default, STP, RSTP, and MSTP are not enabled.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to disable STP, RSTP, MSTP, PVST+, or R-PVST+.

```
switch(config)# no protocol spanning-tree
```

## Shutting down STP, RSTP, or MSTP globally

To shut down STP, RSTP, or MSTP globally, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **shutdown** command to globally shutdown STP, RSTP, MSTP, PVST+, or R-PVST+. The **shutdown** command below works in all three modes.

```
switch(config-mstp)# shutdown
```

## Specifying the bridge priority

In any mode (STP, RSTP, or MSTP), use the **bridge-priority** command to specify the priority of the switch. After you decide on the root switch, set the appropriate values to designate the switch as the root switch. If a switch has a bridge priority that is lower than that of all the other switches, the other switches automatically select the switch as the root switch.

The root switch should be centrally located and not in a “disruptive” location. Backbone switches typically serve as the root switch because they often do not connect to end stations. All other decisions in the network, such as which port to block and which port to put in forwarding mode, are made from the perspective of the root switch.

Bridge Protocol Data Units (BPDUs) carry the information exchanged between switches. When all the switches in the network are powered up, they start the process of selecting the root switch. Each switch transmits a BPDU to directly connected switches on a per-VLAN basis. Each switch compares the received BPDU to the BPDU that the switch sent. In the root switch selection process, if switch 1 advertises a root ID that is a lower number than the root ID that switch 2 advertises, switch 2 stops the advertisement of its root ID, and accepts the root ID of switch 1. The switch with the lowest bridge priority becomes the root switch.

Additionally, you may specify the bridge priority for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

---

### NOTE

On the Brocade VDX 8770, the VLAN value can be 1 through 4086. VLAN IDs 4087 through 4094 are internally reserved VLAN IDs. On all other Brocade VDX switches, the VLAN value can be 1 through 3962, as 3963 through 4094 are reserved.

---

To specify the bridge priority, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, MSTP, PVST+, or R-PVST+.
3. Specify the bridge priority. The range is 0 through 61440 and the priority values can be set only in increments of 4096. The default priority is 32678.

```
switch(conf-stp)# bridge-priority 20480
```

4. Optional: Specify the bridge priority for a specific VLAN.

```
switch(conf-stp)# bridge-priority 20480 vlan 10
```

## Specifying the bridge forward delay

In any mode (STP, RSTP, or MSTP), use this command to specify how long an interface remains in the listening and learning states before the interface begins forwarding all spanning tree instances.

The range is 4 through 30 seconds. The default is 15 seconds. The following relationship should be kept:

$$2 * (\text{forward\_delay} - 1) \geq \text{max\_age} \geq 2 * (\text{hello\_time} + 1)$$

Additionally, you may specify the forward delay for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

On the Brocade VDX 8770, the VLAN value can be 1 through 4086. VLAN IDs 4087 through 4094 are internally-reserved VLAN IDs. On all other Brocade VDX switches, the VLAN value can be 1 through 3962, as 3963 through 4094 are reserved.

To specify the bridge forward delay, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, MSTP, PVST+, or R-PVST+.

```
switch(config)# protocol spanning-tree stp
```

3. Specify the bridge forward delay.

```
switch(conf-stp)# forward-delay 20
```

4. Optional: Specify the bridge forward delay for a specific VLAN.

```
switch(conf-stp)# forward-delay 20 vlan 10
```

## Specifying the bridge maximum aging time

In any mode (STP, RSTP, or MSTP), use this command to control the maximum length of time that passes before an interface saves its BPDU configuration information.

When configuring the maximum aging time, you must set the max-age to be greater than the hello time. The range is 6 through 40 seconds. The default is 20 seconds. The following relationship should be kept:

$$2 * (\text{forward\_delay} - 1) \geq \text{max\_age} \geq 2 * (\text{hello\_time} + 1)$$

Additionally, you may specify the maximum aging for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

On the Brocade VDX 8770, the VLAN value can be 1 through 4086. VLAN IDs 4087 through 4094 are internally reserved VLAN IDs. On all other Brocade VDX switches, the VLAN value can be 1 through 3962, as 3963 through 4094 are reserved.

To specify the bridge maximum aging time, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, MSTP, PVST+, or R-PVST+.

```
switch(config)# protocol spanning-tree stp
```

3. Specify the bridge maximum aging time.

```
switch(config-stp)# max-age 25
```

4. Optional: Specify the bridge maximum aging time for a specific VLAN.

```
switch(config-stp)# max-age 25 vlan 10
```

### Enabling the error disable timeout timer

In any mode (STP, RSTP, or MSTP), use this command to enable the timer to bring a port out of the disabled state. When the STP BPDU guard disables a port, the port remains in the disabled state unless the port is enabled manually. This command allows you to enable the port from the disabled state. For details on configuring the error disable timeout interval, see [“Specifying the error disable timeout interval”](#) on page 358.

To enable the error disable timeout timer, perform the following steps from privileged EXEC mode. By default, the timeout feature is disabled.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, MSTP, PVST+, or R-PVST+.

```
switch(config)# protocol spanning-tree stp
```

3. Enable the error disable timeout timer.

```
switch(config-stp)# error-disable-timeout enable
```

### Specifying the error disable timeout interval

In any mode (STP, RSTP, or MSTP), use this command to specify the time in seconds it takes for an interface to time out. The range is 10 through 1000000 seconds. The default is 300 seconds. By default, the timeout feature is disabled.

To specify the time in seconds it takes for an interface to time out, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, MSTP, PVST+, or R-PVST+.

```
switch(config)# protocol spanning-tree stp
```

3. Specify the time in seconds it takes for an interface to time out.

```
switch(config-stp)# error-disable-timeout interval 60
```

### Specifying the port-channel path cost

In any mode (STP, RSTP, or MSTP), use this command to specify the port-channel path cost. The default port cost is **standard**. The path cost options are as follows:

- **custom**—Specifies that the path cost changes according to the port-channel's bandwidth.
- **standard**—Specifies that the path cost does not change according to the port-channel's bandwidth.

To specify the port-channel path cost, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, MSTP, PVST+, or R-PVST+.
 

```
switch(config)# protocol spanning-tree stp
```
3. Specify the port-channel path cost.
 

```
switch(conf-stp)# port-channel path-cost custom
```
4. Return to privileged EXEC mode.
 

```
switch(config)# end
```
5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.
 

```
switch# copy running-config startup-config
```

## Specifying the bridge hello time

In STP or RSTP mode, use this command to configure the bridge hello time. The hello time determines how often the switch interface broadcasts hello BPDUs to other devices. The range is 1 through 10 seconds. The default is 2 seconds.

When configuring the hello-time, you must set the max-age must be greater than the hello time. The following relationship should be kept:

$$2 * (\text{forward\_delay} - 1) \geq \text{max\_age} \geq 2 * (\text{hello\_time} + 1)$$

Additionally, you may specify the hello time for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

On the Brocade VDX 8770, the VLAN value can be 1 through 4086. VLAN IDs 4087 through 4094 are internally-reserved VLAN IDs. On all other Brocade VDX switches, the VLAN value can be 1 through 3962, as 3963 through 4094 are reserved.

To specify the bridge hello time, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, MSTP, PVST+, or R-PVST+.
 

```
switch(config)# protocol spanning-tree stp
```
3. Specify the time range in seconds for the interval between the hello BPDUs sent on an interface.
 

```
switch(conf-stp)# hello-time 5
```
4. Optional: Specify the time range in seconds for the interval between the hello BPDUs sent on an interface for a specific VLAN.
 

```
switch(conf-stp)# hello-time 5 vlan 10
```
5. Return to privileged EXEC mode.
 

```
switch(config)# end
```
6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.
 

```
switch# copy running-config startup-config
```

## Specifying the transmit hold count (RSTP, MSTP, and R-PVST+)

In RSTP and MSTP mode, use this command to configure the BPDU burst size by specifying the transmit hold count value. The command configures the maximum number of BPDUs transmitted per second for RSTP and MSTP before pausing for 1 second. The range is 1 through 10. The default is 6.

To specify the transmit hold count, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Specify the transmit hold count.

```
switch(config-mstp)# transmit-holdcount 5
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Enabling Cisco interoperability (MSTP)

In MSTP mode, use the **cisco-interoperability** command to enable or disable the ability to interoperate with certain legacy Cisco switches. If Cisco interoperability is required on any switch in the network, then all switches in the network must be compatible, and therefore enabled by means of this command. The default is Cisco interoperability is disabled.

---

### NOTE

This command is necessary because the “version 3 length” field in the MSTP BPDU on some legacy Cisco switches does not conform to current standards.

---

To enable interoperability with certain legacy Cisco switches, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.

```
switch(config)# protocol spanning-tree mstp
```

3. Enable interoperability with certain legacy Cisco switches.

```
switch(config-mstp)# cisco-interoperability enable
```

## Disabling Cisco interoperability (MSTP)

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.

```
switch(config)# protocol spanning-tree mstp
```

3. Disable interoperability with certain legacy Cisco switches.

```
switch(config-mstp)# cisco-interoperability disable
```



## Mapping a VLAN to an MSTP instance

In MSTP mode, use the **instance** command to map a VLAN to an MTSP. You can group a set of VLANs to an instance. This command can be used only after the VLAN is created. VLAN instance mapping is removed from the configuration if the underlying VLANs are deleted.

To map a VLAN to an MSTP instance, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.  

```
switch(config)# protocol spanning-tree mstp
```
2. Enter the **protocol** command to enable MSTP.  

```
switch(config-mstp)# instance 5 vlan 300
```
3. Map a VLAN to an MSTP instance.  

```
switch(config-mstp)# end
```
4. Return to privileged EXEC mode.  

```
switch# copy running-config startup-config
```

## Specifying the maximum number of hops for a BPDU (MSTP)

In MSTP mode, use this command to configure the maximum number of hops for a BPDU in an MSTP region. Specifying the maximum hops for a BPDU prevents the messages from looping indefinitely on the interface. When you change the number of hops, it affects all spanning tree instances. The range is 1 through 40. The default is 20 hops.

To configure the maximum number of hops for a BPDU in an MSTP region, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.  

```
switch(config)# protocol spanning-tree mstp
```
2. Enter the **protocol** command to enable MSTP.  

```
switch(config-mstp)# max-hops hop_count
```
3. Enter the **max-hops** command to configure the maximum number of hops for a BPDU in an MSTP region.  

```
switch(config-mstp)# end
```
4. Return to privileged EXEC mode.  

```
switch# copy running-config startup-config
```

## Specifying a name for an MSTP region

In MSTP mode, use this command to assign a name to an MSTP region. The region name has a maximum length of 32 characters and is case-sensitive.

To assign a name to an MSTP region, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.  

```
switch(config)# protocol spanning-tree mstp
```
3. Enter the **region** command to assign a name to an MSTP region.  

```
switch(config-mstp)# region sydney
```
4. Return to privileged EXEC mode.  

```
switch(config-mstp)# end
```
5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.  

```
switch# copy running-config startup-config
```

## Specifying a revision number for MSTP configuration

In MSTP mode, use this command to specify a revision number for an MSTP configuration. The range is 0 through 255. The default is 0.

To specify a revision number for an MSTP configuration, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.  

```
switch(config)# protocol spanning-tree mstp
```
3. Enter the **revision** command to specify a revision number for an MSTP configuration.  

```
switch(config-mstp)# revision 17
```
4. Return to privileged EXEC mode.  

```
switch(config-mstp)# end
```
5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.  

```
switch# copy running-config startup-config
```

## Clearing spanning tree counters

In privileged EXEC mode, use this command to clear spanning tree counters on all interfaces or on the specified interface.

To clear spanning tree counters, perform the following steps from privileged EXEC mode.

1. Use the **clear** command to restart the protocol migration process on all the interfaces.  

```
switch# clear spanning-tree counter
```
2. Use the **clear** command to restart the protocol migration process associated with a specific port-channel or DCB port interface.  

```
switch# clear spanning-tree counter interface tengigabitethernet 0/1
```

## Clearing spanning tree-detected protocols

In privileged EXEC mode, restart the protocol migration process (force the renegotiation with neighboring switches) on all interfaces or on the specified interface.

To restart the protocol migration process, perform the following tasks from privileged EXEC mode.

1. Use the **clear** command to clear all spanning tree counters on all interfaces:

```
switch# clear spanning-tree detected-protocols
```

2. Use the **clear** command to clear the spanning tree counters associated with a specific port-channel or DCB port interface:

```
switch# clear spanning-tree detected-protocols interface tengigabitethernet
0/1
```

## Displaying STP-related information

Enter the **show spanning-tree brief** command in privileged EXEC mode to display all STP, RSTP, MSTP, PVST+, or R-PVST+-related information.

---

### NOTE

The **show spanning-tree brief** command output shows the port state as *ERR*, not *root\_inc*, when root guard is in effect.

---

# Configuring STP, RSTP, or MSTP on DCB interface ports

This section details the commands for enabling and configuring STP, RSTP, or MSTP on individual 10-gigabit Ethernet DCB interface ports.

## Enabling automatic edge detection

From the DCB interface, use this command to automatically identify the edge port. The port can become an edge port if no BPDU is received. By default, automatic edge detection is disabled.

To enable automatic edge detection on the DCB interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(conf-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to enable automatic edge detection on the DCB interface.

```
switch(conf-if-te-0/1)# spanning-tree autoedge
```

## Configuring the path cost

From the DCB interface, use this command to configure the path cost for spanning tree calculations. The lower the path cost means there is a greater chance of the interface becoming the root port. The range is 1 through 200000000. The default path cost is 2000 for a 10-gigabit Ethernet interface.

Additionally, you may specify the spanning tree cost for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

On the Brocade VDX 8770, the VLAN value can be 1 through 4086. VLAN IDs 4087 through 4094 are internally-reserved VLAN IDs. On all other Brocade VDX switches, the VLAN value can be 1 through 3962, as 3963 through 4094 are reserved.

To configure the path cost for spanning tree calculations on the DCB interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.  
  
The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:  

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```
3. Enter the **no shutdown** command to enable the DCB interface.  

```
switch(conf-if-te-0/1)# no shutdown
```
4. Enter the **spanning-tree** command to configure the path cost for spanning tree calculations on the DCB interface.  

```
switch(conf-if-te-0/1)# spanning-tree cost 10000
```
5. Optional: Enter the **spanning-tree** command to configure the path cost for spanning tree calculations on the DCB interface.  

```
switch(conf-if-te-0/1)# spanning-tree cost 10000 vlan 10
```
6. Return to privileged EXEC mode.  

```
switch(conf-if-te-0/1)# end
```
7. Enter the **copy** command to save the *running-config* file to the *startup-config* file.  

```
switch# copy running-config startup-config
```

## Enabling a port (interface) as an edge port

From the DCB interface, use this command to enable the port as an edge port to allow the port to quickly transition to the forwarding state. To configure a port as an edge port, follow these guidelines:

- A port can become an edge port if no BPDU is received.
- When an edge port receives a BPDU, it becomes a normal spanning tree port and is no longer an edge port.

- Because ports that are directly connected to end stations cannot create bridging loops in the network, edge ports transition directly to the forwarding state and skip the listening and learning states.
- This command is only for RSTP and MSTP. Use the **spanning-tree portfast** command for STP (see “[Enabling port fast \(STP\)](#)” on page 367).

To enable the DCB interface as an edge port, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(config-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to enable the DCB interface as an edge port.

```
switch(config-if-te-0/1)# spanning-tree edgeport
```

## Enabling guard root

From the DCB interface, use this command to enable the guard root feature on the switch. This feature provides a way to enforce the root bridge placement in the network. With guard root enabled on an interface, the switch is able to restrict which interface is allowed to be the spanning tree root port or the path to the root for the switch. The root port provides the best path from the switch to the root switch. By default, guard root is disabled.

Guard root protects the root bridge from malicious attacks and unintentional misconfigurations where a bridge device that is not intended to be the root bridge becomes the root bridge. This causes severe bottlenecks in the data path. Guard root ensures that the port on which it is enabled is a designated port. If the guard-root-enabled port receives a superior BPDU, it goes to a discarding state.

Additionally, you may enable guard root for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

On the Brocade VDX 8770, the VLAN value can be 1 through 4086. VLAN IDs 4087 through 4094 are internally reserved VLAN IDs. On all other Brocade VDX switches, the VLAN value can be 1 through 3962, as 3963 through 4094 are reserved.

To enable guard root on a DCB interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(conf-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to enable guard root on a DCB interface.

```
switch(conf-if-te-0/1)# spanning-tree guard root
```

5. Enter the **spanning-tree** command to enable guard root for a specific VLAN.

```
switch(conf-if-te-0/1)# spanning-tree guard root vlan 10
```

### Specifying the STP hello time

From the DCB interface, use this command to set the time interval between BPDUs sent by the root switch. Changing the hello time affects all spanning tree instances.

The **max-age** setting must be greater than the **hello-time** setting (see [“Specifying the bridge maximum aging time”](#) on page 357). The range is 1 through 10 seconds. The default is 2 seconds.

To specify the MSTP hello time on a DCB interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(conf-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to specify the hello time on a DCB interface.

```
switch(conf-if-te-0/1)# spanning-tree hello-time 5
```

5. Return to privileged EXEC mode.

```
switch(conf-if-te-0/1)# end
```

6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

### Specifying restrictions for an MSTP instance

From the DCB interface, use this command to specify restrictions on the interface for an MSTP instance.

To specify restrictions for an MSTP instance on a DCB interface, perform the following steps.

1. Enter the **configure terminal** command to access global configuration mode from privileged EXEC mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(conf-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to specify the restrictions for an MSTP instance on a DCB interface.

```
switch(conf-if-te-0/1)# spanning-tree instance 5 restricted-tcn
```

5. Return to privileged EXEC mode.

```
switch(conf-if-te-0/1)# end
```

6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Specifying a link type

From the DCB interface, use this command to specify a link type. Specifying the **point-to-point** keyword enables rapid spanning tree transitions to the forwarding state. Specifying the **shared** keyword disables spanning tree rapid transitions. The default setting is point-to-point.

To specify a link type on a DCB interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(conf-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to specify the link type on the DCB interface.

```
switch(conf-if-te-0/1)# spanning-tree link-type shared
```

## Enabling port fast (STP)

From the DCB interface, use this command to enable port fast on an interface to allow the interface to transition quickly to the forwarding state. Port fast immediately puts the interface into the forwarding state without having to wait for the standard forward time.

---

### NOTE

If you enable the **portfast bpduguard** option on an interface and the interface receives a BPDU, the software disables the interface and puts the interface in the ERR\_DISABLE state.

---

**CAUTION**

Enabling portfast on ports can cause temporary bridging loops, in both trunking and nontrunking mode.

Use the **spanning-tree edgeport** command for MSTP, RSTP, and R-PVST+ (see [“Enabling a port \(interface\) as an edge port”](#) on page 364).

To enable port fast on the DCB interface for STP, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(config-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to enable port fast on the DCB interface.

```
switch(config-if-te-0/1)# spanning-tree portfast
```

## Specifying the port priority

From the DCB interface, use this command to specify the port priority. The range is 0 through 240 in increments of 16. The default is 128.

Additionally, you may specify the spanning tree priority for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

On the Brocade VDX 8770, the VLAN value can be 1 through 4086. VLAN IDs 4087 through 4094 are internally reserved VLAN IDs. On all other Brocade VDX switches, the VLAN value can be 1 through 3962, as 3963 through 4094 are reserved.

To specify the port priority on the DCB interface, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

**NOTE**

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.



```
switch(conf-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to specify the port priority on the DCB interface.

```
switch(conf-if-te-0/1)# spanning-tree priority 32
```

5. Optional: Enter the **spanning-tree** command to specify the port priority for a specific VLAN.

```
switch(conf-if-te-0/1)# spanning-tree priority 32 vlan 10
```

## Restricting the port from becoming a root port

From the DCB interface, use this procedure to restrict a port from becoming a root port. The default is to allow the DCB interface to become a root port. This procedure affects MSTP only.

To restrict the DCB interface from becoming a root port, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(conf-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to restrict the DCB interface from becoming a root port.

```
switch(conf-if-te-0/1)# spanning-tree restricted-role
```

## Restricting the topology change notification

From the DCB interface, use this command to restrict the topology change notification BPDUs sent on the interface. By default, the restriction is disabled. This procedure affects MSTP only.

To restrict the topology change notification BPDUs sent on the DCB interface, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.

```
switch(conf-if-te-0/1)# no shutdown
```

4. Enter the **spanning-tree** command to restrict the topology change notification BPDUs sent on the DCB interface.

```
switch(conf-if-te-0/1)# spanning-tree restricted-tcn
```

### Enabling spanning tree

From the DCB interface, use this command to enable spanning tree on the DCB interface. By default, spanning tree is disabled.

To enable spanning tree on the DCB interface, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.  
  
The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.  
  
switch(conf-if-te-0/1)# **no shutdown**
4. Enter the **spanning-tree** command to enable spanning tree on the DCB interface.

```
switch(conf-if-te-0/1)# no spanning-tree shutdown
```

### Disabling spanning tree

From the DCB interface, use this command to disable spanning tree on the DCB interface. By default, spanning tree is disabled.

To disable spanning tree on the DCB interface, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.  
  
The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.  
  
switch(conf-if-te-0/1)# **no shutdown**
4. Enter the **spanning-tree** command to enable spanning tree on the DCB interface.

```
switch(conf-if-te-0/1)# spanning-tree shutdown
```

## Spanning Tree Protocol and VCS mode

Network OS v4.0.0 and later supports any version of STP to run in VCS mode and function correctly between interconnecting VCS nodes, or between VCS and other vendor's switches. This feature is called Distributed Spanning Tree Protocol (DiST).

The purpose of DiST is as follows:

- To support VCS to VCS connectivity and automatic loop detection and prevention.
- To assist deployment plans for replacing the legacy xSTP enabled switches in your network.
- To eliminate the need for standalone mode in order to be able to run xSTP.

DiST supports any of the following flavors of xSTP:

- IEEE STP
- IEEE RSTP
- IEEE MSTP
- Cisco PVST
- Cisco PVRST

DiST treats one VCS as one virtual xSTP bridge from an external view. The xSTP protocol is running between VCS and standalone nodes. Each VCS has a unique RBridge ID and Priority. DiST can be enabled on VCS edge ports connecting to other VCS nodes or standalone nodes. The Port Ids used for xSTP are dynamically assigned and unique within the VCS.

It is important to note that in fabric cluster mode, it is assumed that the global xSTP configuration is the same on all the member nodes of the VCS. Mismatched global configurations of xSTP across different nodes in VCS are not supported. For example, if one of the nodes in a VCS is configured for RSTP and another one is configured for STP or MSTP, this scenario is unsupported and the fabric will not form.

Each RBridge runs the spanning tree instance in a distributed manner. Each spanning tree instance considers all the edge ports and the best information from the remote RBridges to arrive at the spanning tree topology. Each RBridge updates all the other members about its best information for a given spanning tree instance.

Each RBridge maintains a table of best information from the other RBridges in the cluster. This table is identical across all the RBridges in the cluster. This information is used to derive the port roles for the local edge ports. The shared information of the whole cluster is considered in the spanning tree calculations for port roles and states of local edge ports. Thus, all the remote RBridges' edge port information could affect the port role selection and port state transitions for the local edge ports. This ensures that each RBridge considers the port roles and states of all the other RBridges to arrive at a final spanning tree topology.

In the event of a change of the "best" information on any member RBridge, that RBridge would update its own next best information to the other RBridges. Some of the scenarios in which this could happen are the following:

- Operational status change of port associated with the "best" information
- Reception of superior information by another edge port on the RBridge
- Reception of superior or inferior information by the "best" port on the RBridge
- Nonreception of BPDUs on the best port for a given period of time

The xSTP update information is received by all member nodes of the cluster. Each node updates its internal database with the received information. If this results in a best-information change, the update is applied on to the logical port for the node. This triggers the xSTP state machine for all local ports.

Configuring DiST

The CLI commands to configure xSTP in the VCS mode are exactly the same as the CLI commands in standalone mode. By default, all the ports are spanning tree disabled. Server ports must be configured as an xSTP edge port.

**NOTE**  
DiST is supported on the VCS edge ports only. DiST cannot be enabled on ISL ports participating in the TRILL-based fabric within VCS. DiST does not update the port state of ISL ports.

xSTP can be enabled on the VCS by means of the **protocol spanning-tree** command. An interface begins participating in the spanning tree once it is configured by the **spanning-tree enable** command. Refer to [“Spanning Tree Protocol configuration and management”](#) on page 355.

[Table 62](#) describes the behavior of interface based on global and interface level configuration. By default, spanning tree is disabled at the global and interface configuration levels.

TABLE 62 Interface behavior by global and interface configuration

| Global config | Interface config | Interface type where STP BPDU is received | Action                     |
|---------------|------------------|-------------------------------------------|----------------------------|
| Disable       | N/A              | Layer 2 (switchport, FCoE)                | Flood to all Layer 2 ports |
| Disable       | N/A              | Layer 3                                   | Drop                       |
| Enable        | Disable          | Layer 2 (switchport, FCoE)                | Drop                       |
| Enable        | N/A              | Layer 3                                   | Drop                       |
| Enable        | Enable           | (switchport, FCoE)                        | Trap                       |

# Configuring UDLD

## In this chapter

- UDLD..... 373

## UDLD

UniDirectional Link Detection (UDLD) protocol is a nonstandard Layer 2 protocol that detects when a physical link becomes unidirectional by means of the exchange of UDLD protocol data units (PDUs). A unidirectional loop can lead to the creation of a loop in a network, which the Spanning Tree Protocol (STP) could inadvertently allow to occur.

Note the following requirements for UDLD:

- NOS 4.0 or later.
- UDLD runs only on physical ports assigned to a port channel.
- UDLD is supported on directly connected switches only.
- UDLD works in standalone mode or VCS mode. In VCS mode, UDLD applies only to edge ports.
- UDLD can interoperate with Brocade IP products.

## How UDLD works

Figure 37 shows a simple four-switch network in which two paths connect to each switch. STP blocks traffic on as many ports as necessary so that only one operational path exists from the STP root bridge to all nodes in the network.




FIGURE 37 Four-switch example for UDLD

In [Figure 37](#), STP detects that the port on switch D that is connected to switch C should be put into a blocked state. Therefore, no data traffic gets transmitted or received on this port. Data traffic remains blocked as long as switch D receives bridge protocol data units (BPDUs) from both switches C and B.

If the link between switch C and switch D becomes unidirectional (for reasons such as hardware failure or incorrect cabling) in the direction from D to C, switch D ages out the status that it was receiving BPDUs from switch C. This eventually causes STP to put the port in a forwarding state, thus allowing all data traffic. This creates a loop for all BUM traffic that enters the network. BUM traffic can go from switch B to switch D to switch C to switch A, and then back to switch B.

To prevent this loop from forming, UDLD can be used to detect that the link between switch C and switch D has become unidirectional.

The UDLD protocol is disabled by default. To use the UDLD protocol, you must first enable the protocol globally and then enable UDLD on each desired individual physical port. For a configuration example, see [“Configuring UDLD”](#) on page 374.

UDLD determines that a link has become unidirectional if the UDLD daemon stops receiving UDLD PDUs from the other end of the link. The UDLD daemon then blocks the physical link. The physical link remains up but the line protocol goes down. During this time, the link continues to transmit and receive UDLD PDUs.

---

#### NOTE

In a VCS environment, the UDLD protocol is applicable only to the edge ports in the VCS. A configuration command to enable the UDLD protocol on a logical port or a non-edge port will be rejected.

---

## Configuring UDLD

Follow the steps below to configure UDLD on your switch.

1. Enter global configuration mode by entering the **configure** command from the desired switch:

```
switch# configure
```

2. To enable the UDLD protocol, as well as to enter protocol UDLD configuration mode, enter the **protocol udld** command.

```
switch(config)# protocol udld
```

3. (Optional) You can change the interval at which UDLD PDUs are transmitted from edge ports. The default interval, in counts of one hundred milliseconds, is 5 (500 milliseconds). To change the interval to 2,000 milliseconds, enter the **hello 20** command:

```
switch(config-udld)# hello 20
```

4. (Optional) You can change the timeout multiplier value to affect the UDLD PDU timeout interval. The UDLD timeout interval is the product of the hello time interval at the other end of the link and the timeout multiplier value. To change the timeout multiplier from the default of 5 to the value 8, run the **multiplier 8** command:

```
switch(config-udld)# multiplier 8
```

5. Enter interface subconfiguration mode for the edge port on which you want to enable UDLD:

```
switch(config-udld)# end
switch# configure
switch(config)# interface te 5/0/1
```

```
switch(config-int-te-5/0/1)# udld enable
```

6. Repeat the preceding step for each edge port on which you wish to enable UDLD.

---

**NOTE**

When the UDLD protocol is enabled on one end of a link, the timeout period might elapse before the UDLD protocol is enabled on the other end of the link. In this case, the link becomes temporarily blocked. When the UDLD protocol is enabled at the other end of the link and a UDLD PDU is received, UDLD automatically unblocks the link.

---

## Other UDLD commands

Other UDLD commands include:

- **clear udld statistics**—Clears either all unidirectional link detection (UDLD) protocol statistics or clears the statistics on a specified port.
- **show udld**—Shows global UDLD information.
- **show udld interface**—Shows UDLD information for one or all ports.
- **show udld statistics**—Shows either all UDLD statistics or shows the statistics on a specified port.

For more information about how to use UDLD commands, refer to the *Network OS Command Reference*.





# Configuring Link Aggregation

---

## In this chapter

- [Link aggregation overview](#) ..... 377
- [Virtual LAG overview](#) ..... 379
- [LACP configuration guidelines and restrictions](#) ..... 384
- [Default LACP configuration](#) ..... 384
- [LACP configuration and management](#) ..... 384
- [LACP troubleshooting tips](#) ..... 386

## Link aggregation overview

Link aggregation allows you to bundle multiple physical Ethernet links to form a single logical trunk providing enhanced performance and redundancy. The aggregated trunk is referred to as a Link Aggregation Group (LAG). The LAG is viewed as a single link by connected devices, the Spanning Tree Protocol, IEEE 802.1Q VLANs, and so on. When one physical link in the LAG fails, the other links stay up and there is no disruption to traffic.

To configure links to form a LAG, the physical links must be the same speed and all links must go to the same neighboring device. Link aggregation can be done by manually configuring the LAG or by dynamically configuring the LAG using the IEEE 802.3ad Link Aggregation Control Protocol (LACP).

When queuing traffic from multiple input sources to the same output port, all input sources are given the same weight, regardless of whether the input source is a single physical link or a trunk with multiple member links.

---

### NOTE

The LAG or LAG interface is also referred to as a port-channel.

---

The benefits of link aggregation are summarized as follows:

- Increased bandwidth (The logical bandwidth can be dynamically changed as the demand changes.)
- Increased availability
- Load sharing
- Rapid configuration and reconfiguration

The Brocade VDX family of switches supports the following trunk types:

- Static, standards-based LAG
- Dynamic, standards-based LAG using LACP
- Static, Brocade-proprietary LAG
- Dynamic, Brocade-proprietary LAG using proprietary enhancements to LACP

## Link Aggregation Group configuration

You can configure a maximum of 24 LAGs with up to 16 links per standard LAG, or four links per Brocade-proprietary LAG. Each LAG is associated with an aggregator. The aggregator manages the Ethernet frame collection and distribution functions.

On each port, link aggregation control does the following:

- Maintains configuration information to control port aggregation.
- Exchanges configuration information with other devices to form LAGs.
- Attaches ports to and detaches ports from the aggregator when they join or leave a LAG.
- Enables or disables an aggregator's frame collection and distribution functions.

Each link in the Brocade VDX hardware can be associated with a LAG; a link cannot be associated with more than one LAG. The process of adding and removing links to and from a LAG is controlled statically, dynamically, or through LACP.

Each LAG consists of the following components:

- A MAC address that is different from the MAC addresses of the LAG's individual member links.
- An interface index for each link to identify the link to neighboring devices.
- An administrative key for each link. Only links having the same administrative key value can be aggregated into a LAG. On each link configured to use LACP, LACP automatically configures an administrative key value equal to the port-channel identification number.

## Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standards-based protocol that allows two partner systems to dynamically negotiate attributes of physical links between them to form logical trunks. LACP determines whether a link can be aggregated into a LAG. If a link can be aggregated into a LAG, LACP puts the link into the LAG. All links in a LAG inherit the same administrative characteristics. LACP operates in two modes:

- Passive mode—LACP responds to Link Aggregation Control Protocol Data Units (LACPDU) initiated by its partner system but does not initiate the LACPDU exchange.
- Active mode—LACP initiates the LACPDU exchange regardless of whether the partner system sends LACPDU.

## Dynamic link aggregation

Dynamic link aggregation uses LACP to negotiate which links can be added and removed from a LAG. Typically, two partner systems sharing multiple physical Ethernet links can aggregate a number of those physical links using LACP. LACP creates a LAG on both partner systems and identifies the LAG by the LAG ID. All links with the same administrative key and all links that are connected to the same partner switch become members of the LAG. LACP continuously exchanges LACPDU to monitor the health of each member link.

## Static link aggregation

In static link aggregation, links are added into a LAG without exchanging LACPDU between the partner systems. The distribution and collection of frames on static links is determined by the operational status and administrative state of the link.

## Brocade-proprietary aggregation

Brocade-proprietary aggregation is similar to standards-based link aggregation but differs in how the traffic is distributed. It also has additional rules that member links must meet before they are aggregated:

- The most important rule requires that there is not a significant difference in the length of the fiber between the member links, and that all member links are part of the same port-group. For example, the Brocade VDX 6720-24 has two port groups; te0/1 to te0/12 and te0/13 to te0/24. All of the member links should be members of either group 1 to 12 or group 13 to 24, but not both.
- A maximum of four Brocade LAGs can be created per port-group.

## LAG distribution process

The LAG aggregator is associated with the collection and distribution of Ethernet frames. The collection and distribution process is required to guarantee the following:

- Inserting and capturing control PDUs.
- Restricting the traffic of a given conversation to a specific link.
- Load balancing between individual links.
- Handling dynamic changes in LAG membership.

## Virtual LAG overview

Configuring a virtual LAG (vLAG) is similar to configuring a LAG. Once the Brocade VCS Fabric detects that the LAG configuration spans multiple switches, the LAG automatically becomes a vLAG.

LACP on the Brocade VCS Fabric emulates a single logical switch by sending the same LACP system ID and sending the same admin and operational key.

Note these features of vLAG:

- Only ports with the same speed are aggregated.
- Brocade proprietary LAGs are not available for vLAGs.
- LACP automatically negotiates and forms the vLAG.
- A port-channel interface is created on all the vLAG members.
- The Brocade VCS Fabric relies on you to consistently configure all nodes in the vLAG.
- Similar to static LAGs, vLAGs are not able to detect configuration errors.
- A zero port vLAG is allowed.
- IGMP snooping fits into the primary link of a vLAG to carry multicast traffic.
- Interface statistics are collected and shown per vLAG member switch. The statistics are not aggregated across switches participating in a vLAG.
- In order to provide link and node level redundancy, the Brocade VCS Fabric supports static vLAGs.

A Brocade VCS Fabric vLAG functions with servers that do not implement LACP because it supports static vLAGs as well.

## Configuring the vLAG

Network OS v4.0.0 supports the speed option to set the “Allowed Speed” of the port-channel to either 1 Gbps or 10 Gbps. The default is 10 Gbps. If the port-channel is 1 Gbps, then the speed needs to be configured prior to enabling the port-channel. Otherwise, the physical links are throttled down because of a speed mismatch. Refer to the *Network OS Command Reference* for information on the **speed** command.

### Notes:

- FCoE and DCB capabilities are not supported by vLAG. FCoE traffic is treated similarly to normal LAN data traffic.
- Static vLAGs are not supported on internal ports.
- Perform this procedure on all member nodes of a vLAG.

To configure the vLAG, perform the following steps in global configuration mode.

1. Configure a LAG between two switches within the Brocade VCS Fabric.

See “[Link Aggregation Group configuration](#)” on page 378 for more information. Once the Brocade VCS Fabric detects that the LAG configuration spans multiple switches, the LAG automatically becomes a vLAG.

2. Configure each VLAG to treat FCoE MAC addresses as being multi-homed hosts, similar to LAN traffic.

The default configuration is to treat FCoE traffic as non-VLAG traffic. This command must be performed on every switch in the vLAG.

```
switch(config)# interface port-channel 10
```

3. Use the **end** command to return to privileged EXEC mode.

```
switch(config-Port-channel-10)# end
switch#
```

4. Use the **showport-channel detail** command to verify the port-channel details.

```
switch# show port-channel detail
LACP Aggregator: Po 27
Aggregator type: Standard
Ignore-split is disabled
Actor System ID - 0x8000,00-05-33-6f-18-18
Admin Key: 0027 - Oper Key 0027
Receive link count: 4 - Transmit link count: 4
Individual: 0 - Ready: 1
Partner System ID - 0x8000,00-05-1e-cd-6e-9f
Partner Oper Key 0027
Member ports on rbridge-id 231:
Link: Te 231/0/22 (0xE718160201) sync: 1 *
Link: Te 231/0/23 (0xE718170202) sync: 1
Link: Te 231/0/36 (0xE718240305) sync: 1
Link: Te 231/0/37 (0xE718250306) sync: 1
```

5. Use the **showport port-channel** command to verify the port-channel interface details.

```
switch# show port port-channel tengigabitethernet 1/0/21
LACP link info: te0/21 -0x18150014
Actor System ID: 0x8000,01-e0-52-00-01-00
Actor System ID Mapped Id: 0
Partner System ID: 0x0001,01-80-c2-00-00-01
```

```
Actor priority: 0x8000 (32768)
Admin key: 0x000a (10) Operkey: 0x0000 (0)
Receive machine state : Current
Periodic Transmission machine state : Slow periodic
Muxmachine state : Collecting/Distr
Admin state: ACT:1 TIM:0 AGG:1 SYN:0 COL:0 DIS:0 DEF:1 EXP:0
Operstate: ACT:1 TIM:0 AGG:1 SYN:1 COL:1 DIS:1 DEF:0 EXP:0
Partner operstate: ACT:1 TIM:0 AGG:1 SYN:1 COL:1 DIS:1 DEF:0 EXP:0
Partner oper port: 100
```

## Configuring the vLAG ignore split

The **vlag ignore-split** command is for LACP-based vLAGs. The scope of this configuration is per port-channel. In scenarios where the vLAG spans more than one node, it minimizes the extent of packet loss in the event of one of the nodes in the vLAG going down.


In a case where connectivity between nodes is lost because of a fabric split (as opposed to one of members going down), there will be duplication of multicast/broadcast packets.

Brocade recommends that you build redundancy in the fabric so that individual links aren't single points of failure.

[Figure 38](#) displays a dual vLAG configuration with three legs of RB2, RB3, and RB4. If RB2, RB3, or RB4 reboots while Host-1 is communicating to Host-2 or Host3, a momentary traffic disruption may occur.

**NOTE**

With ignore-split active, a vLAG node reboot can result in a more than one second loss while interoperating with a Linux server/nic-team/CNA, because of premature egress of traffic from the server.



**FIGURE 38** vLAG configuration of the ignore split

To reduce vLAG failover down time, you must set the ignore split on all of the legs in the vLAG (RB2, RB3 and RB4 in this case).

**NOTE**

By default, vLAG ignore split is already activated in VCS.

To configure the vLAG ignore split, perform the following steps from global configuration mode.

1. Log in to RB2, the first leg of the vLAG 1.
2. Access the port channel for the first leg.  

```
switch(config)# interface port-channel 1
```
3. Activate vLAG ignore split.  

```
switch(config-Port-channel-1)# vlag ignore-split
```
4. Log in to RB3, the second leg of vLAG 1.
5. Access the port channel for the second leg.  

```
switch(config)# interface port-channel 2
```
6. Activate vLAG ignore split.

```
switch(config-Port-channel-2)# vlag ignore-split
```

7. Access the port channel for the third leg.

```
switch(config)# interface port-channel 3
```

8. Activate vLAG ignore split.

```
switch(config-Port-channel-3)# vlag ignore-split
```

## Configuring load balancing on a remote RBridge

This feature allows you to configure the load balancing feature on a remote RBridge, which is not a member of the vLAG (also known as a non-local RBridge), to forward traffic to a vLAG. To distribute the traffic among the possible paths towards the VLAG, you can configure the vlag-load-balancing flavor on RB2. Available flavors are listed in [Table 63](#).

**TABLE 63** Load balancing flavors

| Flavor                  | Definition                                                                         |
|-------------------------|------------------------------------------------------------------------------------|
| dst-mac-vid             | Destination MAC address and VID based load balancing.                              |
| src-mac-vid             | Source MAC address and VID based load balancing.                                   |
| src-dst-mac-vid         | Source and Destination MAC address and VID based load balancing.                   |
| src-dst-ip              | Source and Destination IP address based load balancing.                            |
| src-dst-ip-mac-vid      | Source and Destination IP and MAC address and VID based load balancing.            |
| src-dst-ip-port         | Source and Destination IP and TCP/UDP port based load balancing.                   |
| src-dst-ip-mac-vid-port | Source and Destination IP, MAC address, VID and TCP/UDP port based load balancing. |

Additionally, an RBridge can be set to a different flavor for different vLAGs present in the cluster. This feature is available for each RBridge and each VLAG, so different load-balancing flavors can be set for traffic directed towards different VLAGs. The **show running-config rbridge-id** command displays the configuration information.

The following example sets the flavor to “destination MAC address and VID based load balancing.”

### Example

```
Switch(config)# rbridge-id 2
switch(config-rbridge-id-2)# fabric port-channel 20 load-balance dst-mac-vid
switch(config-rbridge-id-2)# end
switch# show running-config rbridge-id 2
rbridge-id 2
 interface-nodespecific ns-vlan 10
 interface-nodespecific ns-ethernet 100
 fabric vlag 10 load-balance src-dst-mac-vid
 fabric vlag 20 load-balance dst-mac-vid
 no protocol vrrp
switch# show fabric port-channel load-balance 10
Fabric Vlag Load-Balance Information

Rbridge-Id : 2
Vlag : 10
Load-Balance Flavor : Source and Destination MAC address and VID based load
balancing
```

```
switch# show fabric port-channel all
Fabric Vlag Load-Balance Information

Rbridge-Id : 2
Vlag : 10
```

## LACP configuration guidelines and restrictions

This section applies to standards-based and Brocade-proprietary LAG configurations, except where specifically noted otherwise.

Follow these LACP configuration guidelines and restrictions when configuring LACP:

- All ports on the Brocade VDX hardware can operate only in full-duplex mode.
- Brocade-proprietary LAGs only—All LAG member links must be part of the same port-group.
- Switchport interfaces—Interfaces configured as “switchport” interfaces cannot be aggregated into a LAG. However, a LAG can be configured as a switchport.

## Default LACP configuration

Table 64 lists the default LACP configuration.

| TABLE 64 Default LACP configuration |                                            |
|-------------------------------------|--------------------------------------------|
| Parameter                           | Default setting                            |
| System priority                     | 32768                                      |
| Port priority                       | 32768                                      |
| Timeout                             | Long (standard LAG) or short (Brocade LAG) |

## LACP configuration and management

---

**NOTE**  
Enter the **copy running-config startup-config** command to save your configuration.

---

### Enabling LACP on a DCB interface

To add additional interfaces to an existing LAG, repeat this procedure using the same LAG group number for the new interfaces.

To enable LACP on a DCB interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710. The prompt for these ports is in the format:

```
switch(config-if-gi-22/0/1)#
```



```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.
4. Enter the **channel-group** command to configure the LACP for the DCB interface.

```
switch(config-if)# channel-group 4 mode active type standard
```

## Configuring the LACP system priority

You configure an LACP system priority on each switch running LACP. LACP uses the system priority with the switch MAC address to form the system ID and also during negotiation with other switches.

The system priority value must be a number in the range of 1 through 65535. The higher the number, the lower the priority. The default priority is 32768.

To configure the global LACP system priority, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Specify the LACP system priority.

```
switch(config)# lacp system-priority 25000
```

## Configuring the LACP timeout period on a DCB interface

The LACP timeout period indicates how long LACP waits before timing out the neighboring device. The **short** timeout period is 3 seconds and the **long** timeout period is 90 seconds. The default is **long**.

To configure the LACP timeout period on a DCB interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the DCB interface.
4. Specify the LACP timeout period for the DCB interface.

```
switch(config-if-te-0/1)# lacp timeout short
```

## Clearing LACP counter statistics on a LAG

To clear LACP counter statistics, enter the **clear** command to clear the LACP counter statistics for the specified LAG group number.

**Example of clearing LACP counters for a specific LAG.**

```
switch# clear lacp 42 counters
```

## Clearing LACP counter statistics on all LAG groups

To clear LACP counter statistics, enter the **clear** command to clear the LACP counter statistics for all LAG groups.

**Example of clearing LACP counters.**

```
switch# clear lacp counters
```

## Displaying LACP information

Use the **show** command to display LACP statistics and configuration information. See the *Network OS Command Reference* for information.

# LACP troubleshooting tips

To troubleshoot problems with your LACP configuration, use the following troubleshooting tips.

If a standard IEEE 802.3ad-based dynamic trunk is configured on a link and the link is not able to join the LAG:

- Make sure that both ends of the link are configured as **standard** for the trunk type.
- Make sure that both ends of the link *are not* configured for **passive** mode. They must be configured as **active/active**, **active/passive**, or **passive/active**.
- Make sure that the port-channel interface is in the administrative “up” state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.
- Make sure the speed parameter is configured to 1000 if the port-channel is using the gigabit interface.
- Make sure that the links that are part of the LAG are connected to the same neighboring switch.
- Make sure that the system ID of the switches connected by the link is unique. This can be verified by entering the **show lacp sys-id** command on both switches.
- You can verify the system ID of the switches in the Brocade VCS Fabric cluster with the **show lacp sys-id** command.
- Make sure that LACPDUs are being received and transmitted on both ends of the link and that there are no error PDUs. This can be verified by entering the **show lacp counters number** command and looking at the receive mode (rx) and transmit mode (tx) statistics. The statistics should be incrementing and should not be at zero or a fixed value. If the PDU rx count is not incrementing, check the interface for possible CRC errors by entering the **show interface link-name** command on the neighboring switch. If the PDU tx count is not incrementing, check the operational status of the link by entering the **show interface link-name** command and verifying that the interface status is “up.”

If a Brocade-based dynamic trunk is configured on a link and the link is not able to join the LAG:

- Make sure that both ends of the link are configured as **Brocade** for trunk type.
- Make sure that both ends of the link *are not* configured for **passive** mode. They must be configured as **active/active**, **active/passive**, or **passive/active**.
- Make sure that the port-channel interface is in the administrative “up” state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.
- Make sure that the links that are part of the LAG are connected to the same neighboring switch.
- Make sure that the system ID of the switches connected by the link is unique. This can be verified by entering the **show lacp sys-id** command on both switches.
- Make sure that LACPDUs are being received and transmitted on both ends of the link and there are no error PDUs. This can be verified by entering the **show lacp counters number** command and looking at the rx and tx statistics. The statistics should be incrementing and should not be at zero or a fixed value. If the PDU rx count is not incrementing, check the interface for possible CRC errors by entering the **show interface link-name** command on the neighboring switch.
- Make sure that the fiber length of the link has a deskew value of 7 microseconds. If it does not, the link will not be able to join the LAG and the following RASLOG message is generated:  
`Deskew calculation failed for link <link-name>.`

When a link has this problem, the **show port-channel** command displays the following message:

```
Mux machine state : Deskew not OK.
```

If a Brocade-based static trunk is configured on a link and the link is not able to join the LAG:

- Make sure that both ends of the link are configured as **Brocade** for trunk type and verify that the mode is “on.”
- Make sure that the port-channel interface is in the administrative “up” state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.

If a standards-based static trunk is configured on a link and the link is not able to join the LAG:

- Make sure that both ends of the link are configured as **standard** for trunk type and verify that the mode is “on.”
- Make sure that the port-channel interface is in the administrative “up” state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.



# Configuring LLDP

---

## In this chapter

|                                                       |     |
|-------------------------------------------------------|-----|
| • LLDP overview .....                                 | 389 |
| • Layer 2 topology mapping.....                       | 390 |
| • DCBX overview.....                                  | 391 |
| • LLDP configuration guidelines and restrictions..... | 393 |
| • Default LLDP configuration.....                     | 393 |
| • LLDP configuration and management.....              | 393 |

## LLDP overview

The IEEE 802.1AB Link Layer Discovery Protocol (LLDP) enhances the ability of network management tools to discover and maintain accurate network topologies and simplify LAN troubleshooting in multi-vendor environments. To efficiently and effectively operate the various devices in a LAN you must ensure the correct and valid configuration of the protocols and applications that are enabled on these devices. With Layer 2 networks expanding dramatically, it is difficult for a network administrator to statically monitor and configure each device in the network.

Using LLDP, network devices such as routers and switches advertise information about themselves to other network devices and store the information they discover. Details such as device configuration, device capabilities, and device identification are advertised. LLDP defines the following:

- A common set of advertisement messages.
- A protocol for transmitting the advertisements.
- A method for storing the information contained in received advertisements.

---

### NOTE

LLDP runs over the data-link layer which allows two devices running different network layer protocols to learn about each other.

---

LLDP information is transmitted periodically and stored for a finite period. Every time a device receives an LLDP advertisement frame, it stores the information and initializes a timer. If the timer reaches the time to live (TTL) value, the LLDP device deletes the stored information ensuring that only valid and current LLDP information is stored in network devices and is available to network management systems.

## Layer 2 topology mapping

The LLDP protocol lets network management systems accurately discover and model Layer 2 network topologies. As LLDP devices transmit and receive advertisements, the devices store information they discover about their neighbors. Advertisement data such as a neighbor's management address, device type, and port identification is useful in determining what neighboring devices are in the network.

---

**NOTE**

The Brocade LLDP implementation supports up to two neighbors.

---

The higher level management tools, such as the Brocade Network Advisor, can query the LLDP information to draw Layer 2 physical topologies. The management tools can continue to query a neighboring device through the device's management address provided in the LLDP information exchange. As this process is repeated, the complete Layer 2 topology is mapped.

In LLDP the link discovery is achieved through the exchange of link-level information between two link partners. The link-level information is refreshed periodically to reflect any dynamic changes in link-level parameters. The basic format for exchanging information in LLDP is in the form of a type, length, value (TLV) field.

LLDP keeps a database for both local and remote configurations. The LLDP standard currently supports three categories of TLVs. Brocade's LLDP implementation adds a proprietary Brocade extension TLV set. The four TLV sets are described as follows:

- Basic management TLV set. This set provides information to map the Layer 2 topology and includes the following TLVs:
  - Chassis ID TLV—Provides the ID for the switch or router where the port resides. This is a mandatory TLV.
  - Port description TLV—Provides a description of the port in an alphanumeric format. If the LAN device supports RFC-2863, the port description TLV value equals the "ifDescr" object. This is a mandatory TLV.
  - System name TLV—Provides the system-assigned name in an alphanumeric format. If the LAN device supports RFC-3418, the system name TLV value equals the "sysName" object. This is an optional TLV.
  - System description TLV—Provides a description of the network entity in an alphanumeric format. This includes system name, hardware version, operating system, and supported networking software. If the LAN device supports RFC-3418, the value equals the "sysDescr" object. This is an optional TLV.
  - System capabilities TLV—Indicates the primary functions of the device and whether these functions are enabled in the device. The capabilities are indicated by two octets. The first octet indicates Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station, respectively. The second octet is reserved. This is an optional TLV.
  - Management address TLV—Indicates the addresses of the local switch. Remote switches can use this address to obtain information related to the local switch. This is an optional TLV.

- IEEE 802.1 organizational TLV set. This set provides information to detect mismatched settings between local and remote devices. A trap or event can be reported once a mismatch is detected. This is an optional TLV. This set includes the following TLVs:
  - Port VLANID TLV—Indicates the port VLAN ID (PVID) that is associated with an untagged or priority tagged data frame received on the VLAN port.
  - PPVLAN ID TLV—Indicates the port- and protocol-based VLAN ID (PPVID) that is associated with an untagged or priority tagged data frame received on the VLAN port. The TLV supports a “flags” field that indicates whether the port is capable of supporting port- and protocol-based VLANs (PPVLANs) and whether one or more PPVLANs are enabled. The number of PPVLAN ID TLVs in a Link Layer Discovery Protocol Data Unit (LLDPDU) corresponds to the number of the PPVLANs enabled on the port.
  - VLAN name TLV—Indicates the assigned name of any VLAN on the device. If the LAN device supports RFC-2674, the value equals the “dot1QVLANStaticName” object. The number of VLAN name TLVs in an LLDPDU corresponds to the number of VLANs enabled on the port.
  - Protocol identity TLV—Indicates the set of protocols that are accessible at the device's port. The protocol identity field in the TLV contains a number of octets after the Layer 2 address that can enable the receiving device to recognize the protocol. For example, a device that wishes to advertise the spanning tree protocol includes at least eight octets: 802.3 length (two octets), LLC addresses (two octets), 802.3 control (one octet), protocol ID (two octets), and the protocol version (one octet).
- IEEE 802.3 organizational TLV set. This is an optional TLV set. This set includes the following TLVs:
  - MAC/PHY configuration/status TLV—Indicates duplex and bit rate capabilities and the current duplex and bit rate settings of the local interface. It also indicates whether the current settings were configured through auto-negotiation or through manual configuration.
  - Power through media dependent interface (MDI) TLV—Indicates the power capabilities of the LAN device.
  - Link aggregation TLV—Indicates whether the link (associated with the port on which the LLDPDU is transmitted) can be aggregated. It also indicates whether the link is currently aggregated and provides the aggregated port identifier if the link is aggregated.
  - Maximum Ethernet frame size TLV—Indicates the maximum frame size capability of the device's MAC and PHY implementation.

## DCBX overview

Storage traffic requires a lossless communication which is provided by DCB. The Data Center Bridging (DCB) Capability Exchange Protocol (DCBX) is used to exchange DCB-related parameters with neighbors to achieve more efficient scheduling and a priority-based flow control for link traffic.

DCBX uses LLDP to exchange parameters between two link peers; DCBX is built on the LLDP infrastructure for the exchange of information. DCBX-exchanged parameters are packaged into organizationally specific TLVs. The DCBX protocol requires an acknowledgement from the other side of the link, therefore LLDP is turned on in both transmit and receive directions. DCBX requires version number checking for both control TLVs and feature TLVs.

DCBX interacts with other protocols and features as follows:

- LLDP—LLDP is run in parallel with other Layer 2 protocols such as RSTP and LACP. DCBX is built on the LLDP infrastructure to communicate capabilities supported between link partners. The DCBX protocol and feature TLVs are treated as a superset of the LLDP standard.
- QoS management—DCBX capabilities exchanged with a link partner are passed down to the QoS management entity to set up the Brocade VDX hardware to control the scheduling and priority-based flow control in the hardware.

The DCBX QoS standard is subdivided into two features sets:

- “Enhanced Transmission Selection”
- “Priority Flow Control”

## Enhanced Transmission Selection

In a converged network, different traffic types affect the network bandwidth differently. The purpose of Enhanced Transmission Selection (ETS) is to allocate bandwidth based on the different priority settings of the converged traffic. For example, Inter-process communications (IPC) traffic can use as much bandwidth as needed and there is no bandwidth check; LAN and SAN traffic share the remaining bandwidth. Table 65 displays three traffic groups: IPC, LAN, and SAN. ETS allocates the bandwidth based on traffic type and also assigns a priority to the three traffic types as follows: Priority 7 traffic is mapped to priority group 0 which does not get a bandwidth check, priority 2 and priority 3 are mapped to priority group 1, priorities 6, 5, 4, 1 and 0 are mapped to priority group 2.

The priority settings shown in Table 65 are translated to priority groups in the Brocade VDX hardware.

**TABLE 65** ETS priority grouping of IPC, LAN, and SAN traffic

| Priority | Priority group | Bandwidth check |
|----------|----------------|-----------------|
| 7        | 0              | No              |
| 6        | 2              | Yes             |
| 5        | 2              | Yes             |
| 4        | 2              | Yes             |
| 3        | 1              | Yes             |
| 2        | 1              | Yes             |
| 1        | 2              | Yes             |
| 0        | 2              | Yes             |

## Priority Flow Control

With Priority Flow Control (PFC), it is important to provide lossless frame delivery for certain traffic classes while maintaining existing LAN behavior for other traffic classes on the converged link. This differs from the traditional 802.3 PAUSE type of flow control where the pause affects all traffic on an interface.

PFC is defined by a one-byte bitmap. Each bit position stands for a user priority. If a bit is set, the flow control is enabled in both directions (Rx and Tx).



**NOTE**

The Brocade VDX 6740 supports only two PFCs.

## LLDP configuration guidelines and restrictions

Follow these LLDP configuration guidelines and restrictions when configuring LLDP:

- Brocade's implementation of LLDP supports Brocade-specific TLV exchange in addition to the standard LLDP information.
- Mandatory TLVs are always advertised.
- The exchange of LLDP link-level parameters is transparent to the other Layer 2 protocols. The LLDP link-level parameters are reported by LLDP to other interested protocols.

**NOTE**

DCBX configuration simply involves configuring DCBX-related TLVs to be advertised. Detailed information is provided in [“LLDP configuration and management”](#) on page 393.

## Default LLDP configuration

[Table 66](#) lists the default LLDP configuration.

**TABLE 66** Default LLDP configuration

| Parameter                                         | Default setting |
|---------------------------------------------------|-----------------|
| LLDP global state                                 | Enabled         |
| LLDP receive                                      | Enabled         |
| LLDP transmit                                     | Enabled         |
| Transmission frequency of LLDP updates            | 30 seconds      |
| Hold time for receiving devices before discarding | 120 seconds     |
| DCBX-related TLVs to be advertised                | dcbx-tlv        |

## LLDP configuration and management

**NOTE**

Enter the **copy running-config startup-config** command to save your configuration changes.

### Enabling LLDP globally

The **protocol lldp** command enables LLDP globally on all interfaces unless it has been specifically disabled on an interface, or the global LLDP disable command has been executed. LLDP is globally enabled by default.

To enable LLDP globally, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

## Disabling and resetting LLDP globally

The **no protocol lldp** command returns all configuration settings made using the protocol lldp commands to their default settings. LLDP is globally enabled by default.

The **disable** command disables LLDP globally.

To reset LLDP globally, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Reset LLDP globally.

```
switch(config)# no protocol lldp
```

To disable LLDP globally perform the following step from global configuration mode.

1. Enter the **protocol lldp** command to enter protocol configuration mode.

```
switch(config)# protocol lldp
```

2. Enter the disable command to disable LLDP globally.

```
switch(conf-lldp)# disable
```

## Configuring LLDP global command options

After entering the **protocol lldp** command from global configuration mode, you are in LLDP configuration mode, which is designated with the switch(conf-lldp)# prompt. Using the keywords in this mode, you can set nondefault parameter values that apply globally to all interfaces.

### *Specifying a system name for the Brocade VDX hardware*

The global system name for LLDP is useful for differentiating between switches. By default, the “host-name” from the chassis/entity MIB is used. By specifying a descriptive system name, you will find it easier to configure the switch for LLDP.

To specify a global system name for the Brocade VDX hardware, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Specify an LLDP system name for the DCB switch.

```
switch(conf-lldp)# system-name Brocade_Alpha
Brocade_Alpha(conf-lldp)#
```

### *Specifying an LLDP system description for the Brocade VDX hardware*

---

**NOTE**

Brocade recommends you use the operating system version for the description or use the description from the chassis/entity MIB. Do not use special characters, such as #!\$, as part of the system name and description.

---

To specify an LLDP system description for the Brocade VDX hardware, perform the following steps from privileged EXEC mode. The system description is seen by neighboring switches.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Specify a system description for the Brocade VDX hardware.

```
switch(config-lldp)# system-description IT_1.6.2_LLDP_01
```

### *Specifying a user description for LLDP*

To specify a user description for LLDP, perform the following steps from privileged EXEC mode. This description is for network administrative purposes and is not seen by neighboring switches.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Specify a user description for LLDP.

```
switch(config-lldp)# description Brocade-LLDP-installed-july-25
```

### *Enabling and disabling the receiving and transmitting of LLDP frames*

By default both transmit and receive for LLDP frames is enabled. To enable or disable the receiving (rx) and transmitting (tx) of LLDP frames, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **mode** command to do the following:

- Enable only receiving of LLDP frames:  
switch(config-lldp)# mode rx
- Enable only transmitting of LLDP frames:  
switch(config-lldp)# mode tx
- Enable both transmit and receive modes.  
switch(config-lldp)# no mode

### *Configuring the transmit frequency of LLDP frames*

To configure the transmit frequency of LLDP frames, perform the following steps from privileged EXEC mode. The default is 30 seconds.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Configure the transmit frequency of LLDP frames.

```
switch(config-lldp)# hello 45
```

### *Configuring the hold time for receiving devices*

To configure the hold time for receiving devices, perform the following steps from privileged EXEC mode. This configures the number of consecutive LLDP hello packets that can be missed before removing the neighbor information. The default is 4.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Configure the hold time for receiving devices.

```
switch(conf-lldp)# multiplier 6
```

### *Advertising the optional LLDP TLVs*

To advertise the optional LLDP TLVs, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Advertise the optional LLDP TLVs.

```
switch(conf-lldp)# advertise optional-tlv management-address port-description
system-capabilities system-name system-description
```

### *Configuring the advertisement of LLDP DCBX-related TLVs*

By default, For a switch in standalone mode only “dcbx-tlv” is advertised

For a switch in Brocade VCS Fabric mode the following TLVs are advertised by default:

- dcbx-tlv
- dcbx-fcoe-app-tlv
- dcbx-fcoe-logical-link-tlv

To configure the LLDP DCBX-related TLVs to be advertised, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Advertise the LLDP DCBX-related TLVs by using these commands:

- switch(conf-lldp)# advertise dcbx-fcoe-app-tlv
- switch(conf-lldp)# advertise dcbx-fcoe-logical-link-tlv
- switch(conf-lldp)# advertise dcbx-tlv
- switch(conf-lldp)# advertise dot1-tlv
- switch(conf-lldp)# advertise dot3-tlv

### *Configuring iSCSI priority*

The iSCSI priority setting is used to configure the priority to be advertised in the DCBx iSCSI TLV.

The iSCSI TLV is used only to advertise the iSCSI traffic configuration parameters to the attached CEE enabled servers and targets. No verification or enforcement of the usage of the advertised parameters by the iSCSI server or target is done by the switch.

To configure the iSCSI priority, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Configure the iSCSI priority.

```
switch(conf-lldp)# iscsi-priority 4
```

---

#### **NOTE**

The default iscsi-priority is 4 and does not display unless you change the iscsi-priority to a different value.

---

4. Advertise the TLV.

```
switch (conf-lldp)# advertise dcbx-iscsi-app-tlv
```

### *Configuring LLDP profiles*

You can configure up to 64 profiles on a switch. Using the **no profile name** command deletes the entire profile.

To configure LLDP profiles, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)# protocol lldp
```

3. Configure the profile name.

```
switch(conf-lldp)# profile UK_LLDP_IT
```

4. Specify a description for the profile.

```
switch(conf-lldp-profile-UK_LLDP_IT)#description standard_profile_by_Jane
```

5. Enable the transmitting and receiving of LLDP frames.

```
switch(conf-lldp-profile-UK_LLDP_IT)# no mode
```

6. Configure the transmission frequency of LLDP updates.

```
switch(conf-lldp-profile-UK_LLDP_IT)# hello 10
```

7. Configure the hold time for receiving devices.

```
switch(conf-lldp-profile-UK_LLDP_IT)# multiplier 2
```

8. Advertise the optional LLDP TLVs.

```
switch(conf-lldp)# advertise optional-tlv management-address port-description
system-capabilities system-name system-description
```

9. Advertise the LLDP DCBX-related TLVs.

```
switch(conf-lldp-profile-UK_LLDP_IT)# advertise dot1-tlv
switch(conf-lldp-profile-UK_LLDP_IT)# advertise dot3-tlv
switch(conf-lldp-profile-UK_LLDP_IT)# advertise advertise dcbx-tlv
switch(conf-lldp-profile-UK_LLDP_IT)# advertise dcbx-fcoe-logical-link-tlv
switch(conf-lldp-profile-UK_LLDP_IT)# advertise dcbx-fcoe-app-tlv
switch(conf-lldp-profile-UK_LLDP_IT)# advertise dcbx-iscsi-app-tlv
```

---

#### NOTE

Brocade recommends against advertising dot1.tlv and dot3.tlv LLDPs if your network contains CNAs from non-Brocade vendors,. This configuration may cause functionality problems.

---

10. Return to privileged EXEC mode.

```
switch(conf-lldp-profile-UK_LLDP_IT)# end
```

11. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(conf-lldp-profile-UK_LLDP_IT)# end
switch# copy running-config startup-config
```

## Configuring the iSCSI profile

You can configure an iSCSI profile to be applied to individual interfaces. However, the priority bit must be set manually for each interface. Using the **no profile name** command deletes the entire profile.

To configure iSCSI profiles, perform the following steps from privileged EXEC mode.

1. Configure the CEE map, if it has not already been created.

For information on the **cee-map** command structure, refer to the *Network OS Command Reference*.

```
switch(config)# cee-map default
switch(conf-ceemap)# priority-group-table 1 weight 50 pfc
switch(conf-ceemap)# priority-group-table 2 weight 30 pfc on
switch(conf-ceemap)# priority-group-table 3 weight 20 pfc on
switch(conf-ceemap)# priority-table 1 1 1 1 2 3 1 1
```

The **priority-table** command syntax is as follows:

```
priority-table PGID0 PGID1 PGID2 PGID3 PGID4 PGID5 PGID6 PGID7
```

- PGID0 = Set the Priority Group ID for all packets with CoS = 0.  
PGID value range is 0..7 for DWRR Priority Group and 15.0..15.7 for Strict Priority Group.
- PGID1 = Set the Priority Group ID for all packets with CoS = 1.  
PGID value range is 0..7 for DWRR Priority Group and 15.0..15.7 for Strict Priority Group.
- PGID2 = Set the Priority Group ID for all packets with CoS = 2.  
PGID value range is 0..7 for DWRR Priority Group and 15.0..15.7 for Strict Priority Group.
- PGID3 = Set the Priority Group ID for all packets with CoS = 3.  
PGID value range is 0..7 for DWRR Priority Group and 15.0..15.7 for Strict Priority Group.
- PGID4 = Set the Priority Group ID for all packets with CoS = 4.  
PGID value range is 0..7 for DWRR Priority Group and 15.0..15.7 for Strict Priority Group.

- PGID5 = Set the Priority Group ID for all packets with CoS = 5.  
PGID value range is 0..7 for DWRR Priority Group and 15.0..15.7 for Strict Priority Group.
- PGID6 = Set the Priority Group ID for all packets with CoS = 6.  
PGID value range is 0..7 for DWRR Priority Group and 15.0..15.7 for Strict Priority Group.
- PGID7 = Set the Priority Group ID for all packets with CoS = 7.  
PGID value range is 0..7 for DWRR Priority Group and 15.0..15.7 for Strict Priority Group.(Reserved for Fabric Priority)

Priority-Table in CEE map configuration requires that PGID 15.0 is dedicated for CoS7. Because of this restriction, make sure that PGID15.0 is configured *only* as the last parameter for Priority-Table configuration.

An explanation of syntax “priority-table 1 2 2 2 2 2 2 15.0” is as follows:

This shows the definition of a CEE Map with Priority to Priority Group mapping of CoS=1, CoS=2, CoS=3, CoS=4, CoS=5, and CoS=6 to a DWRR Priority Group ID of 2, and CoS=0 to a Priority Group ID of 1, and CoS=7 to a Strict Priority Group.

This is one way to provision the CEE Priority to Priority Group Table, which maps each of the eight ingress CoS into a Priority Group.

In VCS mode, traffic classes are either all strict priorities (802.1Q default) or a combination of strict and DWRR traffic classes.

2. Enter LLDP configuration mode.

```
switch(conf-ceemap)# protocol lldp
```

3. Create an LLDP profile for iSCSI.

```
switch(conf-lldp)# profile iscsi_config
```

4. Advertise the iSCSI TLV.

```
switch(conf-lldp-profile-iscsi_config)# advertise dcbx-iscsi-app-tlv
```

5. Enter configuration mode for the specific interface.

The **gigabitethernet** ports on the Brocade VDX 6710 do not allow the **cee default** command. This port type does not support PFC or iSCSI App TLV.

```
switch (conf-lldp-profile-iscsi_config)# interface te 0/1
```

6. Apply the CEE provisioning map to the interface.

```
switch(conf-if-te-0/1)# cee default
```

7. Apply the LLDP profile you created for iSCSI.

```
switch(conf-if-te-0/1)# lldp profile iscsi_config
```

8. Set the iSCSI priority bits for the interface.

```
switch(conf-if-te-0/1)# lldp iscsi-priority 4
```

9. Repeat steps 5 through 8 for additional interfaces.



## Configuring LLDP interface-level command options

Only one LLDP profile can be assigned to an interface. If you do not use the **lldp profile** option at the interface level, the global configuration is used on the interface. If there are no global configuration values defined, the global default values are used.

To configure LLDP interface-level command options, perform the following steps from privileged EXEC mode.

1. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 0/10
```

2. Apply an LLDP profile to the interface.

```
switch(config-if-te-0/10)# lldp profile network_standard
```

3. Return to privileged EXEC mode.

```
switch(config-if-te-0/10)# end
```

4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Clearing LLDP-related information

To clear LLDP-related information, perform the following steps from privileged EXEC mode.

1. Use the **clear** command to clear LLDP neighbor information.

```
switch# clear lldp neighbors interface tengigabitethernet 0/1
```

2. Use the **clear** command to clear LLDP statistics.

```
switch# clear lldp statistics interface tengigabitethernet 0/1
```

## Displaying LLDP-related information

To display LLDP-related information, perform the following steps from privileged EXEC mode.

1. Use the **show lldp** command to display LLDP general information.

```
switch# show lldp
```

2. Use the **show lldp** command to display LLDP interface-related information.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch# show lldp interface tengigabitethernet 0/1
```

3. Use the **show lldp** command to display LLDP neighbor-related information.

```
switch# show lldp neighbors interface tengigabitethernet 0/1 detail
```



# Configuring ACLs

---

## In this chapter

|                                                       |     |
|-------------------------------------------------------|-----|
| • ACL overview .....                                  | 403 |
| • Default ACL configuration .....                     | 404 |
| • ACL configuration guidelines and restrictions ..... | 404 |
| • ACL configuration and management .....              | 405 |
| • IP ACL .....                                        | 408 |

## ACL overview

---

### NOTE

In the Brocade Network OS v4.0.0 release, both Ingress Layer 2 MAC access control lists (ACLs) and Layer 3 IP ACLs are supported. With the introduction of Network OS v4.0.0 extended IP ACLs support IPv6 to access the switch or cluster from the management plane; however, management plane ACLs do not support the use of remarks.

---

ACLs filter traffic for the Brocade VDX hardware platforms and permit or deny frames on ingress interfaces that have the ACLs applied to them. You can apply ACLs on the three kinds of Layer 2 interfaces that Brocade Network OS v4.0.0 supports: physical 1-, 10-, and 4-gigabit Ethernet, VLAN, and port-channel (both static and dynamic LAGs), and Layer 3 IP virtual interfaces.

Each ACL is a unique collection of “permit” and “deny” statements (rules) that apply to frames. When a frame is received on an interface, the switch compares the fields in the frame against any ACLs applied to the interface to verify that the frame has the required permissions to be forwarded. The switch compares the frame sequentially against each rule in the ACL, and either forwards the frame or drops the frame.

The switch examines ACLs associated with options configured on a given interface. As frames enter the switch on an interface, ACLs associated with all inbound options configured on that interface are examined.

The primary benefits of ACLs are as follows:

- Provide a measure of security.
- Save network resources by reducing traffic.
- Block unwanted traffic or users.
- Reduce the chance of denial of service (DOS) attacks.

There are two types of ACLs:

- Standard ACLs—Permit and deny traffic according to the source MAC address in the incoming frame. Use standard MAC ACLs if you only need to filter traffic based on source addresses.
- Extended ACLs—Permit and deny traffic according to the source and destination MAC addresses in the incoming frame, as well as EtherType.

MAC ACLs are supported on the following interface types:

- Physical interfaces
- Logical interfaces (LAGs)
- VLANs

IP ACLs are supported on the following interface types:

- Logical interfaces (LAGs)
- VLANs

## Default ACL configuration

When none of the policies is enforced on the switch, these default ACL rules are effective in Network OS:

- seq 0 permit tcp any any eq 22
- seq 1 permit tcp any any eq 23
- seq 2 permit tcp any any eq 897
- seq 3 permit tcp any any eq 898
- seq 4 permit tcp any any eq 111
- seq 5 permit tcp any any eq 80
- seq 6 permit tcp any any eq 443
- seq 7 permit udp any any eq 161
- seq 8 permit udp any any eq 111
- seq 9 permit tcp any any eq 123
- seq 10 permit tcp any any range 600 65535
- seq 11 permit udp any any range 600 65535

## ACL configuration guidelines and restrictions

Follow these ACL configuration guidelines and restrictions when configuring ACLs:

- The order of the rules in an ACL is critical. The first rule that matches the traffic stops further processing of the frames.
- Standard ACLs and extended ACLs cannot have the same name.

- Applying a permit or deny UDP ACL to the management interface enacts an implicit deny for TCP; however, ping will succeed.
- Applying a permit or deny ACL for a specific UDP port enacts an implicit deny for all other UDP ports.
- Applying a permit or deny ACL for a specific TCP port enacts an implicit deny for all other TCP ports.
- There is a default “permit” rule added at the end of the rules list of a Layer 2 (L2) ACL. This implicit rule permits all L2 streams that do not match any of the configured rules in the sequence list associated with the ACL.
- The default action of “permit any” is inserted at the end of a bounded L2 ACL. This default rule is not exposed to the user.
- There is a default “deny” rule added at the end of the rule list of a Layer 3 (L3) ACL.
- The default action of “deny any” is inserted at the end of a bounded L3 ACL. This default rule is not exposed to the user.
- Applying a hard-drop ACL in place of a permit or deny ACL enables packets to be dropped and overrides the control packet trap entries, but does not override the permit entry that occurs before the rule in the ACL.
- You cannot delete an ACL if it is applied to an interface.

## ACL configuration and management

Invoke the **copy running-config startup-config** command to save your configuration changes.

### Creating a standard MAC ACL and adding rules

---

#### NOTE

You can use the **resequence** command to change the sequence numbers assigned to the rules in a MAC ACL. For detailed information, see [“Reordering the sequence numbers in a MAC ACL”](#) on page 408.

---

A MAC ACL does not take effect until it is applied to a Layer 2 interface. Refer to [“Applying a MAC ACL to a DCB interface”](#) on page 406 and [“Applying a MAC ACL to a VLAN interface”](#) on page 407.

To create a standard MAC ACL and add rules, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Create a standard MAC ACL and enter ACL configuration mode.

In this example, the name of the standard MAC ACL is “test\_01.”

```
switch(config)# mac access-list standard test_01
switch(conf-macl-std)#
```

3. Enter the **deny** command to create a rule in the MAC ACL to drop traffic with the source MAC address.

```
switch(conf-macl-std)# deny 0022.3333.4444 count
```

4. Enter the **permit** command to create a rule in the MAC ACL to permit traffic with the source MAC address.

```
switch(conf-macl-std)# permit 0022.5555.3333 count
```

5. Use the **seq** command to create MAC ACL rules in a specific sequence.

```
switch(conf-macl-std)# seq 100 deny 0011.2222.3333 count
switch(conf-macl-std)# seq 1000 permit 0022.1111.2222 count
```

6. Return to privileged EXEC mode.

```
switch(conf-macl-std)# end
```

7. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Creating an extended MAC ACL and adding rules

---

### NOTE

You can use the **resequence** command to change the sequence numbers assigned to the rules in a MAC ACL. For detailed information, see [“Reordering the sequence numbers in a MAC ACL”](#) on page 408.

---

The MAC ACL name length is limited to 64 characters. A MAC ACL does not take effect until it is applied to a Layer 2 interface. Refer to [“Applying a MAC ACL to a DCB interface”](#) on page 406 and [“Applying a MAC ACL to a VLAN interface”](#) on page 407.

To create an extended MAC ACL and add rules, run the following steps in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

2. Create an extended MAC ACL and enter ACL configuration mode.

```
switch(config)# mac access-list extended test_02
```

3. Create a rule in the MAC ACL to **permit** traffic with the source MAC address and the destination MAC address.

```
switch(conf-macl-ext)# permit 0022.3333.4444 0022.3333.5555
```

4. Use the **seq** command to insert the rule anywhere in the MAC ACL.

```
switch(conf-macl-ext)# seq 5 permit 0022.3333.4444 0022.3333.5555
```

5. Return to privileged EXEC mode.

```
switch(conf-macl-ext)# end
```

6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Applying a MAC ACL to a DCB interface

Ensure that the ACL you want to apply exists and is configured to filter traffic in the manner you need for this DCB interface. An ACL does not take effect until it is expressly applied to an interface using the **access-group** command. Frames can be filtered as they enter an interface (ingress direction).

---

### NOTE

The DCB interface must be configured as a Layer 2 switch port before an ACL can be applied as an access-group to the interface.

---

To apply a MAC ACL to a DCB interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the DCB interface type and slot/port number.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, VDX 6740 and 6740T, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following example format:

```
switch(config)# interface tengigabitethernet 0/1
switch(config-if-gi-22/0/1)#
```

3. Enter the **switchport** command to configure the interface as a Layer 2 switch port.
4. Enter the **mac-access-group** command to specify the MAC ACL that is to be applied to the Layer 2 DCB interface in the ingress direction.

```
switch(conf-if-te-0/1)# mac access-group test_02 in
```

## Applying a MAC ACL to a VLAN interface

Ensure that the ACL that you want to apply exists and is configured to filter traffic in the manner that you need for this VLAN interface. An ACL does not take effect until it is expressly applied to an interface using the **access-group** command. Frames can be filtered as they enter an interface (ingress direction).

To apply a MAC ACL to a VLAN interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to apply the MAC ACL to the VLAN interface.
3. Enter the **mac-access-group** command to specify the MAC ACL that is to be applied to the VLAN interface in the ingress direction.

```
switch(config)# interface vlan 50
```

```
switch(config-Vlan-50)# mac access-group test_02 in
```

## Modifying MAC ACL rules

You cannot modify the existing rules of a MAC ACL. However, you can remove the rule and then recreate it with the desired changes.

If you need to add more rules between existing rules than the current sequence numbering allows, you can use the **resequence** command to reassign sequence numbers. For detailed information, see [“Reordering the sequence numbers in a MAC ACL”](#) on page 408.

Use a sequence number to specify the rule you wish to modify. Without a sequence number, a new rule is added to the end of the list, and existing rules are unchanged.

---

### NOTE

Using the **permit** and **deny** keywords, you can create many different rules. The examples in this section provide the basic knowledge needed to modify MAC ACLs.

---



---

### NOTE

This example assumes that “test\_02” contains an existing rule number 100 with the “deny any any” options.

---

To modify a MAC ACL, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **mac** command to specify the ACL called test\_02 for modification.

```
switch(config)# mac access-list extended test_02
```

3. Enter the **no seq** command to delete the existing rule 100.

```
switch (conf-macl-ext)# no seq 100
```

-or-

Enter the **seq** command to recreate rule number 100 by recreating it with new parameters.

```
switch(conf-macl-ext)# seq 100 permit any any
```

## Removing a MAC ACL

A MAC ACL cannot be removed from the system unless the access-group applying the MAC ACL to a DCB or a VLAN interface is first removed.

To remove a MAC ACL, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **mac** command to specify and delete the ACL that you want to remove. In this example, the extended MAC ACL name is “test\_02.”

```
switch(config)# no mac access-list extended test_02
```

## Reordering the sequence numbers in a MAC ACL

You can reorder the sequence numbers assigned to rules in a MAC ACL. Reordering the sequence numbers is useful when you need to insert rules into an ACL and there are not enough available sequence numbers. The default initial sequence number is 10 and the default increment is 10 for both standard and extended MAC ACLs.

The first rule receives the number specified by the starting-sequence number that you specify. Each subsequent rule receives a number larger than the preceding rule. The difference in numbers is determined by the increment number that you specify. The starting-sequence number and the increment number must range from 1 through 65535.

For example, in the task listed below the **resequence** command assigns a sequence number of 50 to the rule named test\_02. The second rule has a sequence number of 55, and the third rule has a sequence number of 60. The example is using IPv4.

```
switch# resequence ip access-list mac test_02 50 5
```

## IP ACL

The IP ACLs control access to the switch. The policies do not control the egress and outbound management traffic initiated from the switch. The IP ACLs support both IPv4 and IPv6 simultaneously.



An IP ACL is a set of rules that are applied to the interface as a packet filtering firewall. Each rule defines whether traffic of a combination of source and destination IP address, protocol, or port, is to be denied or permitted.

Each ACL must have a unique name, but there is no limit to the number of ACLs to be defined. An ACL can contain rules for only one version of IP (either IPv4 or IPv6). Only one ACL by the version of IP can be active on the interface at a time. In other words, one ACL for IPv4 addresses and one ACL for IPv6 address on the interface for packet filtering can be active at the same time.

For filtering the traffic, each rule of the ACL applied to the interface is checked in the ascending order of their sequence numbers. A maximum of 2048 rules can be added to an access list. When the ACL is applied to an interface, only the 256 lowest-numbered rules are applied. If an ACL does not contain any rules and is applied to the interface, it becomes “no-op” and all ingress traffic is denied through the interface. For Layer 2 ACL, if there are no rules applied to the interface then the action is permitted through that interface. But in Layer 3 ACL or IP ACL, it is denied.

After an IP ACL rule is created, it is not possible to modify any of its options.

The default configuration of the switch consists of two ACLs; one IPv4 ACL and one IPv6 ACL is applied to the interface.

There are two types of IP access lists:

- Standard: Contains rules for only the source IP address. The rules are applicable to all ports of that source IP address.
- Extended: Contains rules for a combination of IP protocol, source IP address, destination IP address, source port, and destination port.

## IP ACL parameters

Table 67 lists the parameters and their definitions for IP ACLs.

### NOTE

For Network OS 3.0 and later, on the Brocade VDX 67xx series, the only supported parameter for Extended IP ACL rules is the **eq** parameter.

**TABLE 67 IP ACL parameters**

| ACL / Rule type      | IP ACL parameter | IP ACL parameter definition                                                                                                                                                                                                                              |
|----------------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Standard IP ACL      | name             | The name of the standard IP ACL. The name must begin with a-z, A-Z, or 0-9. Underscores and hyphens are also accepted except as the first character. The ACL name must be unique among all ACL types (L2/L3) and cannot contain more than 63 characters. |
| Standard IP ACL rule | seq              | The sequence number of the rule. The number must be from 0 through 4294967290. A rule without a sequence number is allocated one. The allocated sequence can be changed by the user using the <b>resequence</b> command.                                 |
|                      | permit/deny      | Specifies whether to permit or deny traffic for the combination specified in the rule.                                                                                                                                                                   |
|                      | any/host         | The IP address of the host from which ingress traffic must be filtered.                                                                                                                                                                                  |

TABLE 67 IP ACL parameters (Continued)

| ACL / Rule type      | IP ACL parameter              | IP ACL parameter definition                                                                                                                                                                                                                                              |
|----------------------|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Extended IP ACL      | name                          | The name of the extended IP Access Control List. The name must begin with a-z, A-Z, or 0-9. Underscores and hyphens are also accepted except as the first character. The ACL name must be unique among all ACL types (L2/L3) and cannot contain more than 63 characters. |
|                      | seq                           | The sequence number of the rule. The number must be from 0 through 65535. A rule without a sequence number is allocated one. The allocated sequence can be changed by the user using the <b>resequence</b> command.                                                      |
| Extended IP ACL Rule | permit/deny                   | Specifies whether to permit or deny traffic for the combination specified in the rule.                                                                                                                                                                                   |
|                      | protocol                      | Indicates the type of IP packet to be filtered.                                                                                                                                                                                                                          |
|                      | any/host                      | The IP address of the host from which inbound traffic must be filtered.                                                                                                                                                                                                  |
|                      | any                           | The IP address of the host to which egress or control of outbound traffic must be blocked. Because the egress and outbound traffic is blocked, the destination address is always “any” (which also covers the Virtual IP address of a host).                             |
|                      | port-number                   | Indicates the source or destination port for which the filter is applicable. This is applicable for both UDP and TCP. The number is from 0 through 65535.                                                                                                                |
|                      | range                         | If there is more than one destination port that must be filtered through the ACL rule, use the range parameter to specify the starting port and end port.                                                                                                                |
|                      | eq                            | If there is only one destination port that must be filtered through the ACL rule, use the eq parameter.                                                                                                                                                                  |
|                      | dscp value                    | Compares the specified value against the DSCP value of the received packet. The range of valid values is from 0 through 63.                                                                                                                                              |
|                      | ack, fin, rst, sync, urg, psh | Any combination of the TCP flags may be specified.                                                                                                                                                                                                                       |
|                      | log                           | Packets matching the filter is sent to the CPU and a corresponding log entry is generated. The optional log parameter enables the logging mechanism. This option is only available with permit and deny.                                                                 |
|                      | hard drop                     | Overrides the trap behavior for control frames and data frames such as echo request (ping).                                                                                                                                                                              |

## Creating a standard IP ACL

To create a standard IP ACL, perform the following steps in global configuration mode.

1. Use the **ip access-list standard** command to enter the configuration mode.

```
switch(config)# ip access-list standard stdACL3
```

2. Use the **seq** command to enter the rules for the ACL. You can enter multiple rules.

```
switch(config-ip-std)# seq 5 permit host 10.20.33.4
switch(config-ip-std)# seq 15 deny any
```

3. Use the **exit** command to return to global configuration mode. Your changes are automatically saved.

```
switch(config-ip-std)# exit
switch(config)#
```

## Creating an extended IP ACL

To create an extended IP ACL, perform the following steps in global configuration mode.

1. Use the **ip access-list extended** command to enter the configuration mode.

```
switch(config)# ip access-list extended extdACL5
```

2. Use the **seq** command to enter the rules for the ACL. You can enter multiple rules.

```
switch(config-ip-ext)# seq 5 deny tcp host 10.24.26.145 any eq 23
switch(config-ip-ext)# seq 7 deny tcp any any eq 80
switch(config-ip-ext)# seq 10 deny udp any any range 10 25
switch(config-ip-ext)# seq 15 permit tcp any any
```

3. Use the **exit** command to return to global configuration mode. Your changes are automatically saved.

```
switch(config-ip-ext)# exit
switch(config)#
```

## Applying an IP ACL to a management interface

To apply the IP ACLs, perform the following steps in global configuration mode.

1. Use the **interface** command to enter the configuration mode for the management interface.

```
switch(config)# interface Management 3/1
```

2. Use the **ip access-group** command to apply the IPv4 standard ACL.

```
switch(config-Management-3/1)# ip access-group stdACL3 in
```

3. Use the **ip access-group** command to apply the IPv6 standard ACL.

```
switch(config-Management-3/1)# ipv6 access-group stdV6ACL1 in
```

4. Use the **ip access-group** command to apply the IPv4 extended ACL.

```
switch(config-Management-3/1)# ip access-group extdACL5 in
```

5. Use the **exit** command to return to global configuration mode. Your changes are automatically saved.

```
switch(config-Management-3/1)# exit
switch(config)#
```

---

**NOTE**

Applying a permit or deny UDP ACL to the management interface enacts an implicit deny for TCP and vice versa.

---

## Binding an ACL in standalone mode or fabric cluster mode

In standalone or fabric cluster mode, an ACL can be applied to any node present in the cluster by specifying its RBridge ID. One ACL per IPv4 and one ACL per IPv6 can be applied to the management interface. Applying a new ACL replaces the ACL that was previously applied. The **no** command form removes an ACL from the interface. Removing the active ACL results in default behavior of “permit any.”

You can bind an IP ACL in the ingress direction for the management interface, and you are not required to create an ACL before binding it to the management interface.

On a management interface, the default action of “permit any” is inserted at the end of an ACL that has been bound.

To bind an ACL to a management interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
interface Management rbridge-id/port [no] {ip | ipv6} access-group NAME {in}
```

---

**NOTE**

Before downgrading firmware, you must unbind any ACLs on the management interface, or the downgrade will be blocked.

---

## Displaying the IP ACL configuration

To display the IP ACL configuration, use the **show running-config ip access-list** command in privileged EXEC mode.

```
switch# show running-config ip access-list
ip access-list standard stdACL3
 seq 5 permit host 10.20.33.4
 seq 7 permit any
!
ip access-list extended extdACL5
 seq 5 deny tcp host 10.24.26.145 any eq 23
 seq 7 deny tcp any any eq 80
 seq 10 deny udp any any range 10 25
 seq 15 permit tcpp any any
```

# Configuring QoS

---

## In this chapter

|                                                       |     |
|-------------------------------------------------------|-----|
| • Standalone QoS.....                                 | 413 |
| • Rewriting .....                                     | 414 |
| • Queueing .....                                      | 414 |
| • Congestion control.....                             | 428 |
| • Multicast rate limiting.....                        | 435 |
| • BUM storm control.....                              | 436 |
| • Scheduling .....                                    | 437 |
| • Data Center Bridging map configuration .....        | 440 |
| • Brocade VCS Fabric QoS .....                        | 444 |
| • Restrictions for Layer 3 features in VCS mode ..... | 445 |
| • Port-based Policer .....                            | 445 |

## Standalone QoS

Standalone Quality of Service (QoS) provides you with the capability to control how the traffic is moved from switch to switch. In a network that has different types of traffic with different needs (specified by Class of Service, or CoS), the goal of QoS is to provide each traffic type with a virtual pipe. FCoE uses traffic class mapping, scheduling, and flow control to provide quality of service.

Traffic running through the switches can be classified as either multicast traffic or unicast traffic. Multicast traffic has a single source but multiple destinations. Unicast traffic has a single source with a single destination. With all this traffic going through inbound and outbound ports, QoS can be set based on egress port and priority level of the CoS.

QoS can also be set on interfaces where the end-station knows how to mark traffic with QoS and it lies with the same trusted interfaces. An untrusted interface occurs when the end-station is untrusted and is at the administrative boundaries.

The principal QoS features are as follows:

- **Rewriting**—Rewriting or marking a frame allows for overriding header fields such as the priority and VLAN ID.
- **Queueing**—Queueing provides temporary storage for frames while waiting for transmission. Queues are selected based on ingress ports, egress ports, and configured user priority level.
- **Congestion control**—When queues begin filling up and all buffering is exhausted, frames are dropped. This has a detrimental effect on application throughput. Congestion control techniques are used to reduce the risk of queue overruns without adversely affecting network throughput. Congestion control features include IEEE 802.3x Ethernet Pause, Tail Drop, Ethernet Priority Flow Control (PFC), and Random Early Discard (RED).

- **Multicast rate limiting**—Many multicast applications cannot be adapted for congestion control techniques and the replication of frames by switching devices can exacerbate this problem. Multicast rate limiting controls frame replication to minimize the impact of multicast traffic. This feature is called BUM Storm Control on Brocade VDX 8770-4, VDX 8770-8, and later platforms.
- A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. **BUM storm control** allows you to limit the amount of broadcast, unknown unicast, and multicast (BUM) traffic admitted to the system to prevent disruptions on Layer 2 physical ports. All traffic received at a physical port in excess of a configured maximum rate for BUM traffic will be discarded. You can also specify whether to shut down an interface if the maximum rate has been exceeded within a 5-second sampling period and receive a LOG indication for the disabled interface. This feature is only supported on VDX 8770-4, VDX 8770-8, and later platforms.
- **Data Center Bridging**—DCB describes an enhanced Ethernet that will enable convergence of various applications in data centers (LAN, SAN, and IPC) onto a single interconnect technology.

## Rewriting

Rewriting a frame header field is typically performed by an edge device. Rewriting occurs on frames as they enter or exit a network because the neighboring device is untrusted, unable to mark the frame, or is using a different QoS mapping.

The frame rewriting rules set the Ethernet CoS and VLAN ID fields. Egress Ethernet CoS rewriting is based on the user-priority mapping derived for each frame as described later in the queueing section.

## Queueing

Queue selection begins by mapping an incoming frame to a configured user priority, then each user-priority mapping is assigned to one of the switch's eight unicast traffic class queues or one of the eight multicast traffic class queues.

### User-priority mapping

There are several ways an incoming frame can be mapped into a user-priority. If the neighboring devices are untrusted or unable to properly set QoS, then the interface is considered untrusted. All traffic must be user-priority mapped using explicit policies for the interface to be trusted; if it is not mapped in this way, the IEEE 802.1Q default-priority mapping is used. If an interface is trusted to have QoS set then the CoS header field can be interpreted.

In standalone mode:

- All incoming priority 7 tagged packets are counted in queue 7 (TC7).
- Untagged control frames are counted in queue 7 (TC7).

---

#### NOTE

The user priority mapping described in this section applies to both unicast and multicast traffic.

---

### *Default user-priority mappings for untrusted interfaces*

When Layer 2 QoS trust is set to *untrusted* then the default is to map all Layer 2 switched traffic to the port default user priority value of 0 (best effort), unless configured to a different value.

Table 68 presents the Layer 2 QoS *untrusted* user-priority generation table.

**TABLE 68** Default priority value of untrusted interfaces

| Incoming CoS | User Priority                    |
|--------------|----------------------------------|
| 0            | port <user priority> (default 0) |
| 1            | port <user priority> (default 0) |
| 2            | port <user priority> (default 0) |
| 3            | port <user priority> (default 0) |
| 4            | port <user priority> (default 0) |
| 5            | port <user priority> (default 0) |
| 6            | port <user priority> (default 0) |
| 7            | port <user priority> (default 0) |

#### **NOTE**

Nontagged Ethernet frames are interpreted as incoming CoS value of 0 (zero).

You can override the default user-priority mapping by applying explicit user-priority mappings.

When neighboring devices are trusted and able to properly set QoS, then Layer 2 QoS trust can be set to CoS and the IEEE 802.1Q default-priority mapping is applied.

Table 69 presents the Layer 2 CoS user-priority generation table conforming to 802.1Q default mapping. You can override this default user-priority table per port if you want to change (mutate) the CoS value.

**TABLE 69** IEEE 802.1Q default priority mapping

| Incoming CoS | User Priority |
|--------------|---------------|
| 0            | 0             |
| 1            | 1             |
| 2            | 2             |
| 3            | 3             |
| 4            | 4             |
| 5            | 5             |
| 6            | 6             |
| 7            | 7             |

### Configuring the QoS trust mode

The QoS trust mode controls user priority mapping of incoming traffic. The Class of Service (CoS) mode sets the user priority based on the incoming CoS value. If the incoming packet is not priority tagged, then fallback is to the Interface Default CoS value.

---

#### NOTE

When a CEE map is applied on an interface, the **qos trust** command is not allowed. The CEE map always puts the interface in the CoS trust mode.

---

To configure the QoS trust mode, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 2/1/2
```

3. Set the interface mode to CoS trust.

Standalone mode:

```
switch(conf-if-te-2/1/2)# qos trust cos
```

VCS mode:

```
switch(conf-if-te-2/1/2)# qos dscp-cos test
```

```
switch(conf-if-te-2/1/2)# qos dscp-traffic-class test
```

---

#### NOTE

To deactivate CoS trust from an interface, enter **no qos trust cos**.

---

4. Return to privileged EXEC mode.

```
switch(conf-if-te-0/2)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

### Verifying CoS trust

To verify applied CoS trust, you can enter the following command from global configuration mode, where *tengigabitethernet 0/2* is the interface name.

```
switch(config)# do show qos interface tengigabitethernet 0/2
```

### Configuring user-priority mappings

To configure user-priority mappings, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```



2. Specify the Ethernet interface.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 1/2/2
```

3. Configure the interface to priority 3.

```
switch(conf-if-te-1/2/2)# qos cos 3
```

4. Return to privileged EXEC mode.

```
switch(conf-if-te-1/2/2)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

### *Creating a CoS-to-CoS mutation QoS map*

To create a CoS-to-CoS mutation, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Create the CoS-to-CoS mutation QoS map. In this example “test” is the map name.

```
switch(config)# qos map cos-mutation test 0 1 2 3 4 5 6 7
```

3. Enter the **do copy** command to save the *running-config* file to the *startup-config* file.

```
switch(config)# do copy running-config startup-config
```

### *Applying a CoS-to-CoS mutation QoS map to an interface*

To apply a CoS-to-CoS mutation QoS map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8.

```
switch(config)# interface tengigabitethernet 2/1/2
```

3. Activate or apply changes made to the CoS-to-CoS mutation QoS map. In this example, “test” is the map name.

```
switch(conf-if-te-2/1/2)# qos cos-mutation test
```

---

#### **NOTE**

To deactivate the mutation map from an interface, enter **no qos cos-mutation** *name*.

---

- 4. Specify the trust mode for incoming traffic.

Use this command to specify the interface ingress QoS trust mode, which controls user-priority mapping of incoming traffic. The untrusted mode overrides all incoming priority markings with the Interface Default CoS. The CoS mode sets the user priority based on the incoming CoS value, if the incoming packet is not priority tagged, then fallback is to the Interface Default CoS value.

```
switch(config-if-te-2/1/2)# qos trust cos
```

- 5. Return to privileged EXEC mode.

```
switch(config-if-te-2/1/2)# end
```

- 6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

*Verifying CoS-to-CoS mutation QoS mapping*

To verify applied QoS maps, you can use one or both of the following options from global configuration mode.

- Verify the CoS mutation mapping for a specific map by using the **do show qos maps cos-mutation** command and the map name.

```
switch(config)# do show qos maps cos-mutation test
```

- Verify all QoS mapping by using the **do show qos maps** command with just the **cos-mutation** parameter only.

```
switch(config)# do show qos maps cos-mutation
```

- Verify the interface QoS mapping by using the **do show qos interface** command.

```
switch(config)# do show qos interface te 2/1/2
```

*Configuring the DSCP trust mode*

Like QoS trust mode, the Differentiated Services Code Point (DSCP) trust mode controls the user-priority mapping of incoming traffic. The user priority is based on the incoming DSCP value. When this feature is not enabled, DSCP values in the packet are ignored.

When DSCP trust is enabled, [Table 70](#) shows default mapping of DSCP values to user priority.

**TABLE 70** Default DSCP priority mapping

| DSCP Values | User Priority |
|-------------|---------------|
| 0-7         | 0             |
| 8-15        | 1             |
| 16-23       | 2             |
| 24-31       | 3             |
| 32-39       | 4             |
| 40-47       | 5             |
| 48-55       | 6             |
| 56-63       | 7             |

**NOTE**

Note the restrictions for using this feature in VCS mode under [“Restrictions for Layer 3 features in VCS mode”](#) on page 445.

To configure DSCP trust mode, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8.

```
switch(config)# interface tengigabitethernet 10/0/2
```

3. Set the interface mode to QoS DSCP trust.

```
switch(config-if-te-10/0/2)# qos trust dscp
```

**NOTE**

To deactivate the DSCP trust mode from an interface, enter **no qos trust dscp**.

4. Return to privileged EXEC mode.

```
switch(config-if-te-10/0/2)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

***Verifying DSCP trust***

To verify applied DSCP trust, you can enter the following command from global configuration mode, where *tengigabitethernet 10/0/2* is the interface name.

```
switch(config)# do show qos interface tengigabitethernet 10/0/2
```

***Creating a DSCP mutation map*****NOTE**

This feature is only supported on Brocade VDX 8770-4, VDX 8770-8, and later models.

To create a DSCP mutation and remap the incoming DSCP value of the ingress packet to egress DSCP values, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Create the DSCP mutation map by specifying a map name. The following command uses “test” as the map name and places the system in DSCP mutation mode so that you can map to traffic classes.

```
switch(config)# qos map dscp-mutation test
```

3. Once the system is in DSCP mutation mode for the configured map (in this case *dscp-mutation-test*), you can map ingress DSCP values to egress DSCP values by using the **mark** command as in the following examples:

```
switch(dscp-mutation-test)# mark 1,3,5,7 to 9
switch(dscp-mutation-test)# mark 11,13,15,17 to 19
switch(dscp-mutation-test)# mark 12,14,16,18 to 20
switch(dscp-mutation-test)# mark 2,4,6,8 to 10
```

This sets the following:

- DSCP values 1, 3, 5, and 7 are set to output as DSCP number 9.
- DSCP values 11, 13, 15, and 17 are set to output as DSCP number 19.
- DSCP values 12, 14, 16, and 18 are set to output as DSCP number 20
- DSCP values 2, 4, 6, and 8 are set to output as DSCP number 10.

4. Enter the **do copy** command to save the *running-config* file to the *startup-config* file.

```
switch(config)# do copy running-config startup-config
```

### *Applying a DSCP mutation map to an interface*

To apply a configured DSCP mutation map to an interface, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

```
switch(config)# interface tengigabitethernet 3/1/2
```

3. Activate or apply changes made to the DSCP mutation map to the interface. In this example “test” is the map name.

```
switch(conf-if-te-3/1/2)# qos dscp-mutation test
```

---

#### **NOTE**

To deactivate a map from an interface, enter **no qos dscp-mutation name**.

---

4. Specify the DSCP trust mode for incoming traffic.

Standalone mode:

```
switch(conf-if-te-2/1/2)# qos trust dscp
```

VCS mode:

```
switch(conf-if-te-2/1/2)# qos dscp-cos test
switch(conf-if-te-2/1/2)# qos dscp-traffic-class test
```

5. Return to privileged EXEC mode.

```
switch(conf-if-te-3/1/2)# end
```

6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

### *Verifying DSCP mutation mapping*

To verify applied DSCP maps, you can use one or both of the following options from global configuration mode.

- Verify DSCP mapping for a specific map using the **do show qos maps dscp-mutation** command and the map name.

```
switch(config)# do show qos maps dscp-mutation test
```

- Verify all DSCP mapping by using the **do show qos maps** command with the **dscp-mutation** operand only.

```
switch(config)# do show qos maps dscp-mutation
```

- Verify DSCP mutation mapping for an interface by using the **show qos interface** command and specifying the interface:

```
switch(config)# do show qos interface te 3/1/2
```

### *Creating a DSCP-to-CoS mutation map*

You can use the incoming DSCP value of ingress packets to remap the outgoing 802.1P CoS priority values by configuring a DSCP-to-COS mutation map on the ingress interface. Use the following steps.

---

#### NOTE

Note the restrictions for using this feature in VCS mode under [“Restrictions for Layer 3 features in VCS mode”](#) on page 445.

---

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Create the DSCP-to-CoS map by specifying a map name. The following command uses “test” as the map name and places the system in dscp-cos map mode so that you can map DSCP values to CoS values.

```
switch(configure)# qos map dscp-cos test
```

3. Once the system is in dscp-cos map mode for the configured map (in this case dscp-cos-test), you can map incoming DSCP values to outgoing CoS priority values by using the **mark** command as in the following examples:

```
switch(dscp-cos-test)# mark 1,3,5,7 to 3
switch(dscp-cos-test)# mark 11,13,15,17 to 5
switch(dscp-cos-test)# mark 12,14,16,18 to 6
switch(dscp-cos-test)# mark 2,4,6,8 to 7
```

This sets the following:

- DSCP values 1, 3, 5, and 7 are set to output as CoS priority 3.
- DSCP values 11, 13, 15, and 17 are set to output as CoS priority 5
- DSCP values 12, 14, 16, and 18 are set to output as CoS priority 6
- DSCP values 2, 4, 6, and 8 are set to output as CoS priority 7.

4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(config)# copy running-config startup-config
```

### *Applying a DSCP-to-CoS map to an interface*

To apply a DSCP-to-CoS mutation map to an interface, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8.

```
switch(config)# interface tengigabitethernet 1/1/2
```

3. Activate or apply changes made to the DSCP-to-CoS mutation map. In this example, “test” is the map name.

```
switch(config-if-te-1/1/2)# qos dscp-cos test
```

---

#### NOTE

To deactivate a map from an interface, enter **no qos dscp-cos name**.

---

4. Specify the DSCP trust mode for incoming traffic.

Use this command to specify the interface ingress DSCP trust mode, which controls user priority mapping of incoming traffic. The untrusted mode will not classify packets based on DSCP. The DSCP trust mode classifies packets based on the incoming DSCP value. If the incoming packet is priority tagged, fallback is to classify packets based on the CoS value.

```
switch(config-if-te-1/1/2)# qos trust dscp
```

5. Return to privileged EXEC mode.

```
switch(config-if-te-1/1/2)# end
```

6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

### Verifying a DSCP-to-CoS mutation map

To verify applied DSCP-to-CoS maps, you can use one or both of the following options from global configuration mode.

- Verify DSCP mapping for a specific map using the **do show qos maps dscp-cos** command and the map name.

```
switch(config)# do show qos maps dscp-cos test
```

- Verify all DSCP mapping by using the **do show qos maps** command with the **dscp-cos** operand only.

```
switch(config)# do show qos maps dscp-cos
```

- Verify DSCP-to-CoS mutation mapping for an interface by using the **show qos interface** command and specifying the interface:

```
switch(config)# do show qos interface te 1/1/2
```

### Traffic class mapping

The Brocade switch supports eight unicast traffic classes to provide isolation and to control servicing for different priorities of application data. Traffic classes are numbered from 0 through 7, with higher values designating higher priorities.

The traffic class mapping stage provides some flexibility in queue selection:

- The mapping may be many-to-one, such as mapping one-byte user priority (256 values) to eight traffic classes.
- There may be a nonlinear ordering between the user priorities and traffic classes.

*Unicast traffic*

Table 71 presents the Layer 2 default traffic class mapping supported for a CoS-based user priority to conform to 802.1Q default mapping.

**TABLE 71** Default user priority for unicast traffic class mapping

| User priority | Traffic class |
|---------------|---------------|
| 0             | 1             |
| 1             | 0             |
| 2             | 2             |
| 3             | 3             |
| 4             | 4             |
| 5             | 5             |
| 6             | 6             |
| 7             | 7             |

You are allowed to override these default traffic class mappings per port. Once the traffic class mapping has been resolved, it is applied consistently across any queueing incurred on the ingress and the egress ports.

*Multicast traffic*

The Brocade switch supports eight multicast traffic classes for isolation and to control servicing for different priorities of application data. Traffic classes are numbered from 0 through 7, with higher values designating higher priorities. The traffic class mapping stage provides some flexibility in queue selection.

Table 72 presents the Layer 2 default traffic class mapping supported for a CoS-based user priority to conform to 802.1Q default mapping.

**TABLE 72** Default user priority for multicast traffic class mapping

| User Priority | Traffic class |
|---------------|---------------|
| 0             | 1             |
| 1             | 0             |
| 2             | 2             |
| 3             | 3             |
| 4             | 4             |
| 5             | 5             |
| 6             | 6             |
| 7             | 7             |

Once the traffic class mapping has been resolved for inbound traffic, it is applied consistently across all queueing incurred on the ingress and egress ports.

You can configure an interface with either a CoS-to-traffic class-map or a DSCP-to-traffic class-map.



## Mapping CoS-to-Traffic-Class

To map a CoS-to-Traffic-Class, perform the following steps from privileged EXEC mode.

---

### NOTE

Creating a CoS-to-Traffic-class-map is available only in standalone mode.

---

1. Enter global configuration mode.  

```
switch# configure terminal
```
2. Create the CoS-Traffic-Class mapping by specifying a name and the mapping.  

```
switch(config)# qos map cos-traffic-class test 1 0 2 3 4 5 6 7
```
3. Return to privileged EXEC mode.  

```
switch(config)# end
```
4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.  

```
switch# copy running-config startup-config
```

## Applying CoS-to-Traffic-Class mapping to an interface

To activate a CoS-to-traffic class mapping, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.  

```
switch# configure terminal
```
2. Specify the Ethernet interface.  

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8.

```
switch(config)# interface tengigabitethernet 12/2/2
```
3. Activate the CoS-to-Traffic-Class mapping. In this case, “test” is the map name.  

```
switch(conf-if-te-12/2/2)# qos cos-traffic-class test
```

---

### NOTE

To deactivate the mutation map from an interface, enter **no qos cos-traffic-class** *name*.

---

4. Return to privileged EXEC mode.  

```
switch(conf-if-te-12/2/2)# end
```
5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.  

```
switch# copy running-config startup-config
```

## Verifying CoS-to-Traffic-Class mapping

To verify a CoS-to-Traffic-Class mapping, you can use one or both of the following options from global configuration mode.

- Verify CoS-Traffic-Class mapping for a specific map by using the **do show qos maps cos-traffic-class** command and specifying a map name.

```
switch(config)# do show qos map cos-traffic-class test
```

- Verify all COS-to-Traffic-Class mapping by using the **do show qos maps** command with the **cos-traffic-class** operand only.

```
switch(config)# do show qos maps cos-traffic-class
```

- Verify CoS-to-Traffic-Class mapping for an interface by using the **do show qos interface** command and specifying the interface:

```
switch(config)# do show qos interface te 12/2/2
```

### *Mapping DSCP-to-Traffic-Class*

Ingress DSCP values can be used to classify traffic for the ingress interface into a specific traffic class by means of a DSCP-to-Traffic class-map. To map a DSCP-to-Traffic-Class, perform the following steps from privileged EXEC mode.

---

#### NOTE

Note the restrictions for using this feature in VCS mode under [“Restrictions for Layer 3 features in VCS mode”](#) on page 445.

---

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Create the DSCP-Traffic-Class mapping by specifying a map name. The following command uses “test” as the map name and places the system in dscp-traffic-class mode so that you can configure mapping for the map that you created.

```
switch(config)# qos map dscp-traffic-class test
```

3. Once the system is in dscp-traffic-class mode for the configured map (in this case dscp-traffic-class-test), you can map DSCP values to traffic classes by using the **mark** command as in the following examples:

```
switch(dscp-traffic-class-test)# mark 1,3,5,7 to 3
switch(dscp-traffic-class-test)# mark 11,13,15,17 to 5
switch(dscp-traffic-class-test)# mark 12,14,16,18 to 6
switch(dscp-traffic-class-test)# mark 2,4,6,8 to 7
```

This sets the following:

- DSCP values 1, 3, 5, and 7 are mapped to traffic class 3.
- DSCP values 11, 13, 15, and 17 are mapped to traffic class 5.
- DSCP values 12, 14, 16, and 18 are mapped to traffic class 6.
- DSCP values 2, 4, 6, and 8 are mapped to traffic class 7.

4. Return to privileged EXEC mode:

```
switch(dscp-traffic-class-test)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

### *Applying DSCP-to-Traffic-Class mapping to an interface*

To activate a DSCP-to-Traffic Class mapping, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8.

```
switch(config)# interface tengigabitethernet 1/1/2
```

3. Activate the DSCP-to-Traffic-Class mapping. In this case, “test” is the map name.

```
switch(conf-if-te-1/1/2)# qos dscp-traffic-class test
```

---

#### **NOTE**

To deactivate a DSCP-to-Traffic-Class map from an interface, enter **no qos dscp-traffic-class name**.

---

4. Return to privileged EXEC mode.

```
switch(conf-if-te-1/1/2)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

### *Verifying DSCP-to-Traffic-Class mapping*

To verify a DSCP-to-Traffic-Class mapping, you can use one or both of the following options from global configuration mode.

- Verify the DSCP-Traffic-Class mapping for specific map by using the **do show qos maps dscp-traffic-class** command and specifying a map name.

```
switch(config)# do show qos maps dscp-traffic-class test
```

- Verify all DSCP-Traffic-Class mapping with the just the **do show qos maps** command with the *dscp-traffic-class parameter only*.

```
switch(config)# do show qos maps dscp-traffic-class
```

- Verify DSCP-to-Traffic-Class mapping for an interface by using the **show qos interface** command and specifying the interface:

```
switch(config)# show qos interface te 1/1/2
```

# Congestion control

Queues can begin filling up for a number of reasons, such as over subscription of a link or backpressure from a downstream device. Sustained, large queue buildups generally indicate congestion in the network and can affect application performance through increased queueing delays and frame loss.

Congestion control covers features that define how the system responds when congestion occurs or active measures taken to prevent the network from entering a congested state.

---

**NOTE**  
You cannot configure CoS thresholds and multicast tail drop on Brocade VDX 8770-4 and VDX 8770-8 platforms. Random Early Discard (RED) is supported only on VDX 8770-4 and VDX 8770-8 platforms.

---

## Tail drop

Tail drop queueing is the most basic form of congestion control. Frames are queued in FIFO order and queue buildup can continue until all buffer memory is exhausted. This is the default behavior when no additional QoS has been configured.

The basic tail drop algorithm does not have any knowledge of multiple priorities and per traffic class drop thresholds can be associated with a queue to address this. When the queue depth breaches a threshold, then any frame arriving with the associated priority value will be dropped. [Figure 39](#) illustrates how you can utilize this feature to ensure that lower-priority traffic cannot totally consume the full buffer memory.




FIGURE 39 Queue depth

Thresholds can also be used to bound the maximum queueing delay for each traffic class. Additionally, if the sum of the thresholds for a port is set below 100 percent of the buffer memory, then you can also ensure that a single port does not monopolize the entire shared memory pool allocated to the port. The tail drop algorithm can be extended to support per-priority drop thresholds. When the ingress port CoS queue depth breaches a threshold, then any frame arriving with the associated priority value will be dropped.

### *Changing the multicast tail drop threshold*

To change the tail drop threshold, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.  

```
switch# configure terminal
```
2. Change the Tail Drop threshold for each multicast traffic class. In this example, a threshold of 1000 packets is used.  

```
switch(config)# qos rcv-queue multicast threshold 1000 1000 1000 1000 1000
1000 1000 1000
```
3. Return to privileged EXEC mode.  

```
switch(config)# end
```
4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.  

```
switch# copy running-config startup-config
```

## Configuring CoS thresholds

Every port has associated with it a total of 9 CoS thresholds, one for the port tail-drop threshold and the other eight are thresholds for per priority. To give a fair allocation of buffers for the traffic from all priorities, the port buffers are allocated among different priorities. That is achieved through per-priority tail-drop thresholds. The port tail-drop threshold represents the amount of buffers given to the port, and the per-priority tail-drop threshold (called the CoS tail-drop threshold from here on) represents the buffers allocated to each CoS.

Whenever the buffers allocated to a priority are fully exhausted, all the traffic coming in on that priority is dropped. In the absence of per-priority tail-drop thresholds (and with only port tail-drop thresholds), the buffers would be consumed on a first-come, first-served basis, resulting in an unfair share of buffers among all the priorities. If you know which priority traffic is most seen, then providing a sufficient number of buffers for those priorities results in fewer packet drops for those priorities.

Therefore, instead of using the standard priority values, you can assign anywhere from 0% through 100% priority to any threshold, with the sum value of all eight priorities to not exceed 100%. For example, using the priorities 5 5 5 5 50 20 2 8 sums up to 100%, as shown in the following example:

```
switch(conf-if-te-0/1)# qos rcv-queue cos-threshold 5 5 5 5 50 20 2 8
switch(conf-if-te-0/1)# do show qos in te 0/1
Interface TenGigabitEthernet 0/1
CoS-to-Traffic Class map 'default'
 In-CoS: 0 1 2 3 4 5 6 7

Out-CoS/TrafficClass: 0/1 1/0 2/2 3/3 4/4 5/5 6/6 7/7
Per-Traffic Class Tail Drop Threshold (bytes)
 TC: 0 1 2 3 4 5 6 7

Threshold: 10180 10180 10180 10180 10180 40723 4072 16289
```

The tail drop thresholds are not allowed to exceed 100%, but can be below 100%. For example, if the tail drop thresholds entered sum, to less than 100%, then the buffer allocation is made on the basis of what has been configured.

## Random Early Discard

---

### NOTE

This feature is only supported on Brocade VDX 8770-4, VDX 8770-8, and later models.

---

Traditionally, Random Early Discard (RED) is used for TCP traffic streams, which are generally more aggressive, as well as reactive, to network drops. If RED is not configured, queues build up at the switch and become full, resulting in tail drop. Tail drop situations can cause head-of-line blocking issues at the switch, which is not desirable. By configuring RED, you set a probability for dropping packets before traffic in the queue reaches a specific threshold. This allows congestion to ease more gradually, avoids retransmit synchronization, resolves “bursty” TCP connections during congestion conditions, and controls packet latency.

Configure RED using the following parameters:

- RED profile identification (0–384)
- Minimum threshold of a queue (0–100%)
- Maximum threshold of a queue (0–100%)
- Drop probability (0-100%)

The ASIC driver maps the configured minimum and maximum percentages to the actual queue size in bytes, depending on the bandwidth of the port (buffers are allocated to a port according to port speed). When buffers in the queue build up to the set minimum threshold, packets being enqueued are randomly dropped. The drop probability parameter defines the randomness of the drops. When the queues exceed the minimum threshold, packets are dropped according to the configured drop probability value. When the queue buffers exceed the set maximum threshold, packets are dropped with 100% probability. The higher the probability set, the more likely packets will be dropped when the minimum percentage is reached.

You can also map a specific CoS priority value (0 through 7) to a specific RED profile.

## Configuring RED profiles

To configure an egress RED profile, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Configure a RED profile. For the profile ID, 10 is used in this case. The *min-threshold*, *max-threshold*, and *drop-probability* values are percentages.

```
switch(config)# qos red-profile 10 min-threshold 10 max-threshold 80
drop-probability 80
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Map CoS priority to RED profile on interface

To map a CoS priority value for a port to the RED profile created under “Configuring RED profiles” on page 431, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

```
switch(config)# interface tengigabitethernet 1/2/2
```

3. Map the profile to use a CoS priority for a port. In the following example, CoS priority 3 is mapped to RED profile ID 10.

```
switch(conf-if-te-1/2/2)# qos random-detect cos 3 red-profile-id 10
```

---

### NOTE

To deactivate the map from an interface, enter **no qos random-detect cos value**.

---

4. Return to privileged EXEC mode.

```
switch(conf-if-te-1/2/2)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Verifying RED profiles

Verify a configured RED profiles by using the **show qos red profiles** command.

```
switch# show qos red profiles
```

Examine the applied RED profiles for an interface by using the **show qos interface interface-name** command. Note that besides the RED profile, this will display all QoS configurations applied to the interface, such as DSCP trust, DSCP-to-DSCP map, CoS-Traffic Class map, and others.

```
switch# show qos red statistics interface te 1/2/2
```

### *Considerations for RED*

Consider the following when configuring RED.

- Up to four RED profiles can be applied to each port group. On the 48 x 10G line card, the port groups consist of ports 1-8, 9-16, 17-24, 25-32, 33-40, and 41-48. On the 12 x 40G line card, the port groups consist of ports 1-2, 3-4, 5-6, 7-8, 9-10, and 11-12.
- Trunk ports cannot share RED profiles with any other ports because the bandwidth for a trunk port changes according to the number of active links in the trunk.
- When queue thresholds in a RED profile are configured by percentage, the switch maps this to a total number of bytes as buffers allocated to a port depend on the port speed.
- A total of 384 RED profiles are supported per chassis.

### **Link aggregation considerations**

Consider the following when using RED profiles for link aggregation (LAG) interfaces:

- RED profiles can be enabled on LAG interfaces. However, the profile is configured on the individual member interfaces of the LAG.
- Because LAG members may belong to different port groups, one of the port groups may not have enough resources available to support a new RED configuration for the member interface. In this case, an error log will indicate that the RED application failed on the specific member interface. When a new member is added to the port-channel, the same error may occur if the new member belongs to a port group with all resources used. To apply the RED profile on the failed member interface, you must remove the RED configuration on other all interfaces in the port group so that resources are available and remove or add the member interface to the LAG.

## **Ethernet Pause**

Ethernet Pause is an IEEE 802.3 standard flow-control mechanism for back pressuring a neighboring device. Pause messages are sent by utilizing the optional MAC control sublayer. A PAUSE frame contains a 2-byte pause number, which states the length of the pause in units of 512bits. When a device receives a PAUSE frame, it must stop sending any data on the interface for the specified length of time, once it completes the transmission of any frame in progress. You can use this feature to reduce Ethernet frame losses by using a standardized mechanism. However, the pause mechanism does not have the ability to selectively back-pressure data sources multiple hops away, or to exert any control per VLAN or per priority, so it is disruptive to all traffic on the link.

Ethernet Pause includes the following features:

- All configuration parameters can be specified independently per interface.
- Pause On/Off can be specified independently for TX and RX directions. No support is provided for disabling autonegotiation.
- Pause generation is based on input (receive) queueing. Queue levels are tracked per input port. When the instantaneous queue depth crosses the high-water mark, then a PAUSE frame is generated. If any additional frames are received and the queue length is still above the low-water mark, then additional PAUSE frames are generated. Once the queue length drops below the low-water mark, then the generation of PAUSE frames ceases.
- A PAUSE frame that is received and processed halts transmission of the output queues associated with the port for the duration specified in the PAUSE frame.



### 1-Gbps pause negotiation

When a 1-Gbps local port is already online, and the **qos flowcontrol** command is issued, the pause settings take effect immediately on that local port. However, when the link is toggled, pause is renegotiated. The local port will advertise the most recent **qos flowcontrol** settings. After autonegotiation completes, the local port pause settings may change, depending on the outcome of the pause negotiation, per 802.3 Clause 28B, as shown in [Table 73](#).

**TABLE 73** Pause negotiation results

| Advertised LOCAL cfg | Advertised REMOTE cfg | Negotiated result                                         |
|----------------------|-----------------------|-----------------------------------------------------------|
| Rx=off Tx=on         | Rx=on Tx=on           | asymmetrical: LOCAL Tx=on --> pause --> REMOTE Rx=on      |
| Rx=on Tx=on          | Rx=off Tx=on          | asymmetrical: LOCAL Rx=on <-- pause <-- REMOTE Tx=on      |
| Rx=on Tx=n/a         | Rx=on Tx=n/a          | symmetrical: LOCAL Tx/Rx=on <-- pause --> REMOTE Tx/Rx=on |
| Rx=n/a Tx=n/a        | Rx=off Tx=off         | disable pause both sides                                  |

### Enabling Ethernet Pause

This task configures FlowControl, in addition to enabling the Ethernet pause frames. Brocade recommends that you also configure the flow control parameters on the connecting device, and not leave the options set to “auto”.

#### NOTE

The Ethernet Pause option is available only in standalone mode.

To enable Ethernet Pause, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8.

```
switch(config)# interface tengigabitethernet 3/0/2
```

3. Enable Ethernet Pause on the interface for both TX and RX traffic.

```
switch(conf-if-te-3/0/2)# qos flowcontrol tx on rx on
```

#### NOTE

To deactivate Ethernet pause on an interface, enter **no qos flowcontrol**.

4. Return to privileged EXEC mode.

```
switch(conf-if-te-3/0/2)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

6. Verify the Ethernet Pause with the **show qos flowcontrol** command.

```
switch# show qos flowcontrol interface all 3/0/2
```

## Ethernet Priority Flow Control

Ethernet Priority Flow Control (PFC) is a basic extension of Ethernet Pause. The Pause MAC control message is extended with eight 2-byte pause numbers and a bitmask to indicate which values are valid. Each pause number is interpreted identically to the base Pause protocol; however, each number is applied to the corresponding Ethernet priority/class level. For example, the Pause number 0 applies to priority zero, Pause number 1 applies to priority one, and so on. This addresses one shortcoming of the Ethernet Pause mechanism, which is disruptive to all traffic on the link. However, it still suffers from the other Ethernet Pause limitations.

---

### NOTE

The Brocade VDX 6740 supports only two PFCs.

---

Ethernet Priority Flow Control includes the following features:

- Everything operates exactly as in Ethernet Pause described above, except there are eight high-water and low-water thresholds for each input port. This means queue levels are tracked per input port plus priority.
- Pause On/Off can be specified independently for TX and RX directions per priority.
- Pause time programmed into Ethernet MAC is a single value covering all priorities.
- Both ends of a link must be configured identically for Ethernet Pause or Ethernet Priority Flow Control because they are incompatible.

### *Enabling an Ethernet PFC*

To enable Ethernet PFC, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 1/1/2
```

3. Enable an Ethernet PFC on the interface.

```
switch(conf-if-te-1/1/2)# qos flowcontrol pfc 3 tx on rx on
```

---

### NOTE

To disable Ethernet PFC from an interface, enter **no qos flowcontrol pfc cos value**.

---

4. Return to privileged EXEC mode.

```
switch(conf-if-te-1/1/2)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Multicast rate limiting

Multicast rate limiting provides a mechanism to control multicast frame replication and cap the effect of multicast traffic.

Multicast rate limit is applied to the output of each multicast receive queue. Rate limits apply equally to ingress receive queueing (first-level expansion) and egress receive queueing (second-level expansion), because the same physical receive queues are utilized. You can set policies to limit the maximum multicast frame rate differently for each traffic class level and cap the total multicast egress rate out of the system.

Multicast rate limiting includes the following features:

- All configuration parameters are applied globally. Multicast rate limits are applied to multicast receive queues as frame replications are placed into the multicast expansion queues. The same physical queues are used for both ingress receive queues and egress receive queues, so rate limits are applied to both ingress and egress queueing.
- The rate limit value represents the maximum multicast expansion rate in packets per second (pps).

---

### NOTE

Multicast rate limiting is not supported on Brocade VDX 8770-4 and 8770-8 platforms. For these products, refer to [“BUM storm control”](#) on page 436.

---

## Creating a receive queue multicast rate-limit

To create the receive queue multicast rate-limit, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Create a lower maximum multicast frame expansion rate. In this example, the rate is set to 10000 pps.

```
switch(config)# qos rcv-queue multicast rate-limit 10000
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## BUM storm control

A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. Broadcast, unicast, and unknown multicast (BUM) storm control can prevent disruptions on Layer 2 physical ports.

BUM storm control allows you to limit the amount of broadcast, unknown unicast, and multicast ingress traffic on a specified interface or on the entire system. All traffic received in excess of the configured rate gets discarded. You also have the option to specify whether to shut down an interface if the maximum defined rate is exceeded within a five-second sampling period. When a port is shut down, you receive a log message. You must then manually re-enable the interface by using the **no shut** command.

### Considerations

- BUM storm control must be configured on one of the following physical interfaces:
  - 10-gigabit Ethernet
  - 1-gigabit Ethernet
  - 40-gigabit Ethernet
- BUM storm control and input service-policy are mutually exclusive features. Only one can be enabled at a time on a given interface.
- BUM storm control replaces the multicast rate-limit feature for Brocade VDX 8770-4 and 8770-8, and later platforms. This command is not supported on Brocade VDX 6xxx modules, such as the VDX 6710, 6720, and 6730.

### Configuring BUM storm control

To configure storm control on 10-gigabit Ethernet interface 101/0/2, with the *broadcast* traffic type and limit-rate of 1000000 bps, perform the following steps:

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface for the traffic you want to control. In the following example, interface 101/0/2 is in *rbridge-id/slot/port* format:

```
switch(config)# interface tengigabitethernet 101/0/2
```

3. Issue the **storm-control ingress** command to set a traffic limit for broadcast traffic on the interface:

```
switch(conf-if-te-101/0/2)# storm-control ingress broadcast 1000000
```

---

#### NOTE

To deactivate storm control from an interface, enter **no storm-control ingress {broadcast | unknown-unicast | multicast} {limit-bps | limit-percent} rate [{monitor | shutdown}]**.

---

4. Verify the storm control verification with the **show storm-control** command.

```
switch(conf-if-te-101/0/2)# do show storm-control
```

| Interface | Type            | rate (Mbps) | conformed   | violated    | total       |
|-----------|-----------------|-------------|-------------|-------------|-------------|
| Tel02/4/1 | broadcast       | 100,000     | 12500000000 | 12500000000 | 25000000000 |
| Tel02/4/1 | unknown-unicast | 100,000     | 12500000000 | 12500000000 | 25000000000 |

|           |                 |         |              |              |              |
|-----------|-----------------|---------|--------------|--------------|--------------|
| Tel02/4/1 | multicast       | 100,000 | 125000000000 | 125000000000 | 250000000000 |
| Tel02/4/2 | broadcast       | 100,000 | 125000000000 | 125000000000 | 250000000000 |
| Tel02/4/3 | broadcast       | 100,000 | 125000000000 | 125000000000 | 250000000000 |
| Tel02/4/4 | unknown-unicast | 100,000 | 125000000000 | 125000000000 | 250000000000 |

## Scheduling

Scheduling arbitrates among multiple queues waiting to transmit a frame. The Brocade switch supports both Strict Priority (SP) and Deficit Weighted Round Robin (DWRR) scheduling algorithms. Also supported is the flexible selection of the number of traffic classes using SP-to-DWRR. When there are multiple queues for the same traffic class, then scheduling takes these equal-priority queues into consideration.

### Strict priority scheduling

Strict priority scheduling is used to facilitate support for latency-sensitive traffic. A strict priority scheduler drains all frames queued in the highest-priority queue before continuing on to service lower-priority traffic classes. A danger with this type of service is that a queue can potentially starve out lower-priority traffic classes.

Figure 40 describes the frame scheduling order for an SP scheduler servicing two SP queues. The higher-numbered queue, SP2, has a higher priority.




FIGURE 40 Strict priority schedule — two queues

### Deficit weighted round robin scheduling

Weighted Round Robin (WRR) scheduling is used to facilitate controlled sharing of the network bandwidth. WRR assigns a weight to each queue; that value is then used to determine the amount of bandwidth allocated to the queue. The round robin aspect of the scheduling allows each queue to be serviced in a set order, sending a limited amount of data before moving onto the next queue and cycling back to the highest-priority queue after the lowest-priority queue is serviced.

Figure 41 describes the frame scheduling order for a WRR scheduler servicing two WRR queues. The higher-numbered queue is considered higher priority (WRR2), and the weights indicate the network bandwidth should be allocated in a 2:1 ratio between the two queues. In Figure 41 WRR2 receives 66 percent of bandwidth and WRR1 receives 33 percent. The WRR scheduler tracks the extra bandwidth used and subtracts it from the bandwidth allocation for the next cycle through the queues. In this way, the bandwidth utilization statistically matches the queue weights over longer time periods.




FIGURE 41 WRR schedule – two queues

Deficit Weighted Round Robin (DWRR) is an improved version of WRR. DWRR remembers the excess used when a queue goes over its bandwidth allocation and reduces the queue’s bandwidth allocation in the subsequent rounds. This way the actual bandwidth usage is closer to the defined level when compared to WRR.

Traffic class scheduling policy

The traffic classes are numbered from 0 to 7; higher-numbered traffic classes are considered higher priority. The Brocade switch provides full flexibility in controlling the number of SP-to-WRR queues. The number of SP queues is specified as SP1 through 8, then the highest-priority traffic classes are configured for SP service and the remaining eight are WRR serviced. Table 74 describes the set of scheduling configurations supported.

When you configure the QoS queue to use strict priority 4 (SP4), then traffic class 7 will use SP4, traffic class 6 will use SP3, and so on down the list. You use the strict priority mappings to control how the different traffic classes will be routed in the queue.

TABLE 74 Supported scheduling configurations

| Traffic Class | SP0  | SP1  | SP2  | SP3  | SP4  | SP5  | SP6  | SP8 |
|---------------|------|------|------|------|------|------|------|-----|
| 7             | WRR8 | SP1  | SP2  | SP3  | SP4  | SP5  | SP6  | SP8 |
| 6             | WRR7 | WRR7 | SP1  | SP2  | SP3  | SP4  | SP5  | SP7 |
| 5             | WRR6 | WRR6 | WRR6 | SP1  | SP2  | SP3  | SP4  | SP6 |
| 4             | WRR5 | WRR5 | WRR5 | WRR5 | SP1  | SP2  | SP3  | SP5 |
| 3             | WRR4 | WRR4 | WRR4 | WRR4 | WRR4 | SP1  | SP2  | SP4 |
| 2             | WRR3 | WRR3 | WRR3 | WRR3 | WRR3 | WRR3 | SP1  | SP3 |
| 1             | WRR2 | WRR2 | WRR2 | WRR2 | WRR2 | WRR2 | WRR2 | SP2 |
| 0             | WRR1 | WRR1 | WRR1 | WRR1 | WRR1 | WRR1 | WRR1 | SP1 |

Figure 42 shows that extending the frame scheduler to a hybrid SP+WRR system is fairly straightforward. All SP queues are considered strictly higher priority than WRR so they are serviced first. Once all SP queues are drained, then the normal WRR scheduling behavior is applied to the non-empty WRR queues.

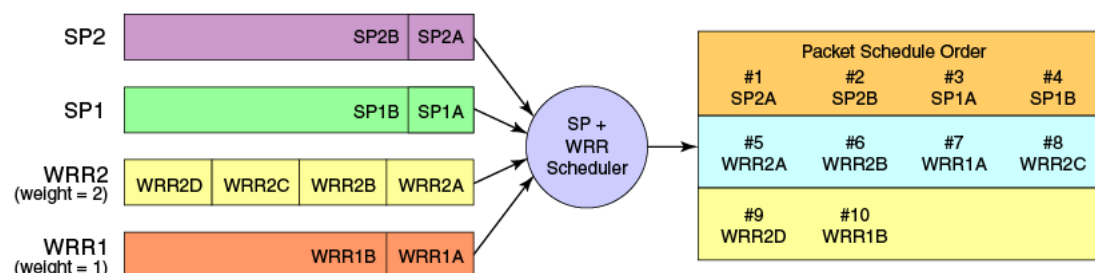


FIGURE 42 Strict priority and Weighted Round Robin scheduler

### *Scheduling the QoS queue*

To specify the schedule used, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the schedule to use and the traffic class to bandwidth mapping.

```
switch(config)# qos queue multicast scheduler dwrr 10 20 20 10 10 10 10 10
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

## Multicast queue scheduling

The multicast traffic classes are numbered from 0 to 7; higher numbered traffic classes are considered higher priority. A fixed mapping from multicast traffic class to equivalent unicast traffic class is applied to select the queue scheduling behavior. Table 75 presents the multicast traffic class with the equivalence mapping applied.

Once the multicast traffic class equivalence mapping has been applied, then scheduling and any scheduler configuration are inherited from the equivalent unicast traffic class. See Table 74 for details on exact mapping equivalencies.

TABLE 75 Multicast traffic class equivalence mapping

| Multicast traffic class | Equivalent unicast traffic class |
|-------------------------|----------------------------------|
| 0                       | 0                                |
| 1                       | 1                                |
| 2                       | 2                                |
| 3                       | 3                                |
| 4                       | 4                                |

TABLE 75 Multicast traffic class equivalence mapping (Continued)

| Multicast traffic class | Equivalent unicast traffic class |
|-------------------------|----------------------------------|
| 5                       | 5                                |
| 6                       | 6                                |
| 7                       | 7                                |

Unicast ingress and egress queueing utilizes a hybrid scheduler that simultaneously supports SP+WRR service and multiple physical queues with the same service level. Multicast adds additional multicast expansion queues. Because multicast traffic classes are equivalent to unicast service levels, they are treated exactly as their equivalent unicast service policies.

*Scheduling the QoS multicast queue*

To schedule the QoS multicast queue, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.  
`switch# configure terminal`
2. Specify the schedule to use and the traffic class to bandwidth mapping.  
`switch(config)# qos queue multicast scheduler dwrr 10 20 20 10 10 10 10 10`
3. Return to privileged EXEC mode.  
`switch(config)# end`
4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.  
`switch# copy running-config startup-config`

# Data Center Bridging map configuration

The DCB QoS covers frame classification, priority and traffic class (queue) mapping, congestion control, and scheduling. Under the DCB Provisioning model all of these features are configured on the basis of two configuration tables, Priority Group Table and Priority Table.

DCB Priority Group Table defines each Priority Group ID (PGID) and its scheduling policy (Strict Priority versus DWRR, DWRR weight, relative priority), and partially defines the congestion control (PFC) configuration. There are 16 rows in the DCB Priority Group Table. [Table 76](#) presents the default DCB Priority Group Table configuration.



**NOTE**

Only a single CoS can be mapped to a PFC-enabled priority queue. The switch automatically maps the CoS number to the same TC number when PFC is enabled. The PGID can be anything from 0 through 7. If your configuration violates this restriction an error message displays and the Priority Group Table is set back to the default values.

When the DCB map is applied, and the interface is connected to the CNA, only one Strict Priority PGID (PGID 15.0 through PGID 15.7) is allowed.

**TABLE 76** Default DCB Priority Group Table configuration

| PGID | Bandwidth% | PFC |
|------|------------|-----|
| 15.0 | —          | N   |
| 15.1 | —          | N   |
| 15.2 | —          | N   |
| 15.3 | —          | N   |
| 15.4 | —          | N   |
| 15.5 | —          | N   |
| 15.6 | —          | N   |
| 15.7 | —          | N   |
| 0    | 0          | N   |
| 1    | 0          | N   |
| 2    | 0          | N   |
| 3    | 0          | N   |
| 4    | 0          | N   |
| 5    | 0          | N   |
| 6    | 0          | N   |
| 7    | 0          | N   |

Strict Priority versus DWRR is derived directly from the PGID value. All PGIDs with prefix 15 receive Strict Priority scheduling policy, and all PGIDs in the range 0 through 7 receive DWRR scheduling policy. Relative priority between Priority Group is exactly the ordering of entries listed in the table, with PGID 15.0 being highest priority and PGID 7 being lowest priority. Congestion control configuration is partially specified by toggling the PFC column On or Off. This provides only partial configuration of congestion control because the set of priorities mapped to the Priority Group is not known, which leads into the DCB Priority Table.

The DCB Priority Table defines each CoS mapping to Priority Group, and completes PFC configuration. There are eight rows in the DCB Priority Table as shown in [Table 77](#).

**TABLE 77** Default DCB priority table

| CoS | PGID |
|-----|------|
| 0   | 15.6 |
| 1   | 15.7 |
| 2   | 15.5 |

TABLE 77 Default DCB priority table (Continued)

| CoS | PGID |
|-----|------|
| 3   | 15.4 |
| 4   | 15.3 |
| 5   | 15.2 |
| 6   | 15.1 |
| 7   | 15.0 |

### Creating a DCB map

To create a DCB map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.  

```
switch# configure terminal
```
2. Select the DCB map by using the **cee-map** command.  
The only map name allowed is “default.”  

```
switch(config)# cee-map default
```
3. Return to privileged EXEC mode.  

```
switch(config)# exit
```
4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.  

```
switch# copy running-config startup-config
```

### Defining a priority group table

To define a priority group table map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.  

```
switch# configure terminal
```
2. Specify the name of the DCB map to define by using the **cee-map** command.

---

**NOTE**  
The only map name allowed is “default.”

---

- ```
switch(config)# cee-map default
```
3. Define the DCB map for PGID 0.

```
switch(config-cee-map-default)# priority-group-table 0 weight 50 pfc on
```
 4. Define the DCB map for PGID 1.

```
switch(config-cee-map-default)# priority-group-table 1 weight 50 pfc off
```
 5. Return to privileged EXEC mode.

```
switch(config-cee-map-default)# end
```
 6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Defining a priority-table map

To define a priority-table map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the name of the DCB map to define by using the **cee-map** command.

```
switch(config)# cee-map default
```

3. Define the map.

```
switch(config-cee-map)# priority-table 1 1 1 0 1 1 1 15.0
```

NOTE

For information about priority-table definitions, see the “**cee-map (configuration)**” command in the *NOS Command Reference, 3.0*.

4. Return to privileged EXEC mode.

```
switch(config-cee-map)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Applying a DCB provisioning map to an interface

To apply a DCB provisioning map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface. In this example, 101/0/2 is used.

```
switch(config)# interface tengigabitethernet 101/0/2
```

3. Apply the DCB map on the interface.

```
switch(conf-if-te-101/0/2)# cee default
```

NOTE

To deactivate the map on the interface, enter **no cee**.

4. Return to privileged EXEC mode.

```
switch(conf-if-te-101/0/2)# end
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Verifying the DCB maps

To verify the CoS DCB map, use the **show cee maps default** command from privileged EXEC mode.

```
switch# show cee maps default
```

Brocade VCS Fabric QoS

Brocade VCS Fabric QoS requires very little user configuration. The only options to modify are the fabric priority and the lossless priority.

Brocade VCS Fabric reserves a mapping priority and fabric priority of seven (7). Any traffic that enters the Brocade VCS Fabric cluster from upstream that is using the reserved priority value is automatically remapped to a lower priority.

Changing the mapping or fabric priority is not required. By default the values are set to zero (0) for both of the remapped priorities.

In Brocade VCS Fabric mode:

- All incoming priority 7 tagged packets are redefined to the default or user-defined value.
- Untagged control frames are counted in queue 7 (TC7).

All switches in the Brocade VCS Fabric cluster must have matching remapping priority values and the same priority-group-table values.

Configuring Brocade VCS Fabric QoS

To configure the remapping priorities for the Brocade VCS Fabric, perform the following steps from global configuration mode.

1. Use the **cee-map** command to enter CEE map configuration mode.

```
switch(config)# cee-map default
```

2. Use the **remap lossless priority** command to set the lossless priority for Brocade VCS Fabric QoS.

The default lossless remap priority is set to 0.

```
switch(config-cee-map-default)# remap lossless-priority priority 2
```

3. Use the **remap fabric priority** command to set the fabric priority for Brocade VCS Fabric QoS.

The default FCoE remap fabric priority is set to 0.

```
switch(fabric-cee-map-default)# remap fabric-priority priority 2
```

4. Use the **exit** command to return to global configuration mode.

```
switch(config-cee-map)# exit
```

5. Specify the incoming Ethernet interface.

```
switch(config)# interface tengigabitethernet 22/0/1
```

6. Apply the CEE Provisioning map to the interface.

```
switch(conf-if-te-22/0/1)# cee default
```

Restrictions for Layer 3 features in VCS mode

When the switch is in VCS mode, the lossless priority for carrying FCoE traffic and the fabric priority for carrying fabric traffic must be isolated from any Layer 3 QoS markings and classification. Therefore, specific restrictions apply to some Layer 3 DSCP QoS features when the switch is working in VCS mode:

The following are restrictions for using applicable Layer 3 DSCP-Traffic-Class map, DSCP-CoS map, and DSCP Trust features in VCS mode. Note that DSCP mutation maps and the egress RED feature are not affected in VCS mode.

- DSCP trust is disabled in VCS mode as it is for CoS trust.
- There are no default DSCP maps in VCS mode. Default maps occur when DSCP trust is enabled in standalone mode.
- A nondefault DSCP-Traffic-Class map has the following restrictions:
 - A DSCP value cannot be classified to Traffic Class 7.
 - A DSCP value cannot be classified to a queue that carries lossless traffic (by default Traffic Class 3).
- A nondefault DSCP-CoS map has the following restrictions:
 - A DSCP value cannot be marked to CoS 7.
 - A DSCP value cannot be marked to lossless priority (by default CoS 3).
- Lossless priorities are identified through the CEE map.
- To enable DSCP based marking or classification, a nondefault DSCP-Traffic-Class map and a DSCP-CoS map have to be applied on the interface.
- To apply a DSCP-Traffic-Class or DSCP-CoS map to an interface, the CoS and Traffic Class values have to be remarked for lossless priorities. For example, when DSCP-Traffic-Class map “abcd” is created, it will have the default contents. When this map is applied to an interface, an error will display that the fabric and lossless priorities are used in the map and it cannot be applied on the interface.
- When a valid DSCP-Traffic-Class map and DSCP-CoS map are applied on the interface, then DSCP trust is enabled with the configured maps.

Port-based Policer

The port-based Policer feature controls the amount of bandwidth consumed by an individual flow or aggregate of flows by limiting the inbound and outbound traffic rate on an individual port according to criteria defined by the user. The Policer provides rate control by prioritizing or dropping ingress and egress packets classified according to a two-rate, three-color marking scheme defined by RFC 4115. This feature is supported only on Brocade VDX 8770-4, VDX 8770-8, and later models.

The Policer supports the following features.

- A color-based priority mapping scheme for limiting traffic rate:
 - One-rate, two-color policing with “conform” color options. “Violate” color traffic will be dropped.
 - Two-rate, three-color policing with “conform” and “exceed” color options. “Violate” color traffic will be dropped.
- A policing option that allows packet headers to be modified for IP precedence.

- Policing options that allow packet headers to be modified for Class of Service (CoS).
- Policing options that allow packet headers to be modified for Differentiated Services Code Point (DSCP).
- Policing options that allow packets to be assigned to a traffic class (0–7).

Color-based priority

The following is the color-based priority mapping scheme for limiting traffic rate:

- Traffic flagged to the green or “conform” color priority conforms to the committed information rate (CIR) as defined by the *cir-rate* variable for the policy-map (refer to [“Policing parameters”](#) on page 453). This rate can be anything from 40000 to 400000000000 bps.
- Traffic flagged as yellow or “exceed” exceeds the CIR, but conforms to the Excess Information Rate (EIR) defined by the *eir-rate* variable for the policy-map (refer to [“Policing parameters”](#) on page 453). This rate can be set from 0 through 400000000000 bps.
- Traffic flagged as red or “violate” are not compared to CIR or EIR and will be dropped.

Using policing parameters, you can define metering rates, such as CIR and EIR, and actions for traffic flagged as conforming or exceeding the rates. As a simple example, traffic within the “conform” rate may be sent at a certain CoS priority, traffic flagged at the “exceed” rate may be sent at a lower priority, and traffic that violates the set rates can be dropped (default and only option).

Configuring Policer functions

To configure port-based Policer functions, perform the following steps while in switch global configuration mode and using Policer CLI commands and parameters:

1. Configure a class map to classify traffic according to traffic properties that you will configure with the policing parameters while adding the class map to a policy-map. Refer to [“Configuring a class map”](#) on page 446.
2. Configure a police priority-map to add color-based priority mapping. Refer to [“Configuring a police priority-map”](#) on page 447. This is an optional step. If you do not define priority mapping for a color, the map defaults to priorities 0, 1, 2, 3, 4, 5, 6, and 7 (in other words, nothing is modified).
3. Configure a policy-map to associate QoS and policing parameters to traffic belonging to specific classification maps. Each policy-map can contain only one classification map. Refer to [“Configuring the policy-map”](#) on page 448.
4. Bind the policy-map to a specific interface using the **service-policy** command. Refer to [“Binding the policy-map to an interface”](#) on page 452.

Configuring a class map

The classification map or “class map” classifies traffic based on match criteria that you configure using the **class-map** command. If traffic matches this criteria, it belongs to the class. Currently, the only match criteria is “match any.” With the “match any” criterion, traffic with any MAC address, IP address, VLAN ID, IP precedence, Access Control List (ACL) security, or other identification belongs to the class.

When you add the class map to a policy-map, the traffic belonging to the class is subject to actions of the QoS and Policer parameters configured for the class-map in the policy-map. For more information on these parameters, refer to [“Policing parameters”](#) on page 453.

To configure a class map, use the following steps:

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Create an access-list (either MAC or IP) to define the traffic

```
switch(config)# mac access-list standard ACL1
switch(conf-macl-std)# permit host 0000.00aa.aa00
switch(conf-macl-std)# exit
```

3. Create a class map by providing a class map name. This enables class-map configuration mode.

```
switch(config)# class-map class-1
```

The name for the class map (in this case default) can be a character string up to 64 characters.

NOTE

The “default” class-map name is reserved and intended to match everything. It is always created and cannot be removed.

4. Provide match criteria for the class.

```
switch(config-classmap)# match access-list mac acl ACL1
```

5. Exit the class-map configuration mode.

```
switch(config-classmap)# exit
```

6. Return to privileged EXEC mode.

```
switch(config)# end
```

7. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

NOTE

Enter the **no map class-map** *name* command while in global configuration mode to remove the classification map

Configuring a police priority-map

Add color-based priority CoS mapping by configuring a police priority-map. A police priority-map remaps frame class of service CoS values (802.1p priority bits in VLAN tag) to conform or exceed color values when rates conform to or exceed limits set in a classification map.

The police priority-map will re-mark CoS values according to color-based green (conform), yellow (exceed), and red (violate) priorities. Creating a police priority-map is optional. If you do not define priority mapping for a color, the map defaults to priorities of 0, 1, 2, 3, 4, 5, 6, and 7 (in other words, nothing is modified). You can configure a maximum of 32 priority-maps (one reserved as a default), but only one map can be associated with a Policer.

NOTE

You can set a priority-map when creating a policy-map by using appropriate Policer attributes.

To configure a priority-map, use the following steps:

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Create a priority-map by providing a priority-map name. This enables police priority-map configuration mode.

```
switch(config)# police-priority-map pmap1
```

The name for the priority-map (in this case pmap1) can be a character string up to 64 characters.

3. Create color-based priority mapping. The following example sets the CoS for traffic that conforms to the CIR set in the policy-map.

```
switch(config-policemap)# conform 0 1 1 2 1 2 2 1 1
```

The following example sets the CoS for traffic that exceeds the CIR setting, but conforms to the EIR set in the policy-map.

```
switch(config-policemap)# exceed 3 3 3 3 4 5 6 7
```

4. Exit the police priority-map configuration mode.

```
switch (config-policemap)# exit
```

5. Return to privileged EXEC mode.

```
switch(config)# end
```

6. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

To delete color-based CoS mapping, use the no operand as in the following examples:

- To delete the conform color, use the following example:

```
switch(config-policemap)# no conform
```
- To delete the exceed color, use the following example:

```
switch(config-policemap)# no exceed
```
- To delete an entire police priority-map, use the following example:

```
switch(config)# no police-priority-map name
```

Configuring the policy-map

Configure a rate-limit policy-map to associate QoS and policing parameters with traffic belonging to a specific classification map. A policy-map can only contain one classification map. You can apply one policy-map per interface per traffic direction (inbound and outbound) by using the **service-policy** command.

To configure a policy-map, add a classification map, and configure QoS and policing parameters for the classification map, use the following steps:

1. Enter the global configuration mode.

```
switch# configure terminal
```

2. Configure a policy-map by providing a policy-map name. This enables policy-map configuration mode.

```
switch(config)# policy-map policymap1
```

The name for the policy-map (in this case policymap1) can be a character string up to 64 characters.

To delete a policy-map, use the **no** operand as in the following example.

```
switch(config)# no policy-map policymap1
```

3. Configure a class map in the policy-map by providing the class map name. This enables policy class map configuration mode. Note that the class map name in the following example matches the name provided when you create the class map by using the **class-map** command (refer to [“Configuring a class map”](#) on page 446).

```
switch(config-policymap)# class default
```

4. Set QoS and policing parameters for the class map as shown in the following example. For information on all of the optional parameters for this command, refer to the *Network OS Command Reference*.

```
(config-policymap-class)# police cir 40000 cbs 5000 eir 40000 ebs 3000  

set-priority pmap1 conform-set-dscp 61 conform-set-tc 7 exceed-  

63 exceed- set-tc 3
```

The CIR parameter is mandatory for a class map. All other parameters are optional. Note that the parameter for set-priority (pmap1) includes the name for the created priority-map (refer to [“Configuring a police priority-map”](#) on page 447). For details on setting QoS and policing parameters, refer to [“Policing parameters”](#) on page 453.

To delete the mandatory CIR parameter, you must delete all Policer parameters while in the policy-map class configuration mode as in the following example:

```
switch(config-policymap-class)# no police
```

To delete any optional parameter, use the **no** operand while in the policy-map class police configuration mode. The following example removes the EBS setting.

```
switch(config-policymap-class-police)# no ebs
```

5. Exit the policy class map configuration mode.

```
switch(config-policymap-class)# exit
```

6. Exit the policy-map configuration mode.

```
switch(config-policymap)# exit
```

7. Return to privileged EXEC mode.

```
switch(config)# end
```

8. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Configuring parameters for a class map (policy-class-map Policer mode)

You can configure QoS and policing parameters when configuring a class map in the policy-class-map configuration mode as shown in the preceding procedure, or you can create or modify parameters for a class map in the policy-class-map Policer attributes mode, as shown in the following example. Using the policy-class-map Policer mode is a convenience for adding or modifying single or multiple attributes for a policy.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Configure a policy-map by providing a policy-map name. This enables policy-map configuration mode.

```
switch(config)# policy-map policymap2
```

The name for the priority-map (in this case policymap2) can be a character string up to 64 characters.

To delete a policy-map, use the *no* operand as in the following example.

```
switch(config)# no policy-map policymap2
```

3. Configure a class map in the policy-map by providing the class map name. This enables policy-class-map configuration mode.

```
switch(config-policymap)# class default
```

NOTE

To configure a class map in the policy-map you must create the class map first using the **class-map** command while in global configuration mode. Refer to [“Configuring a class map”](#) on page 446.

4. Provide a policing parameter for the class map. This enables policy-class-map Policer configuration mode.

```
switch(config-policymap-class)# police cir 4000000
```

5. Enter another parameter as applicable.

```
(config-policymap-class-police)# cbs 50000
```

6. Enter additional parameters as applicable.

```
switch(config-policymap-class-police)# eir 800000 ebs 400000 conform-set-tc 3  
exceed-set-prec 4
```

7. Exit the policy-class-map Policer configuration mode.

```
switch(config-policymap-class-police)# exit
```

8. Exit the policy-class-map configuration mode.

```
switch (config-policymap-class)# exit
```

9. Exit the policy-map configuration mode.

```
switch(config-policymap)# exit
```

10. Return to privileged EXEC mode.

```
switch(config)# end
```

11. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Attach the mutation to the class

User can specify the mutation-map to be used on the port. This can lead to a possible contradiction if there are other user-defined classes used in the same policy-map that has a set cos action configured. In this case-defined cos takes priority over the mutation map.

Perform the following task in global configuration mode.

1. Select the policy map.

```
switch(config)# policy-map p1
```

2. Select the class.

```
switch(config-policymap)# class class-default
```

3. Specify the mutation map to be used on the port. Different kinds of mutations can be used depending on the command. For complete information, refer to *Network OS Command Reference*. The available commands are:

- cos-mutation
- cos-traffic-class
- dscp-mutation
- dscp-cos
- dscp-traffic-class

```
switch(config-policyclass)# cos-mutation plsmap
```

Attach the port-shaper to the class

You can specify the shaping rate per port attached to the policy map. You can use this command to smooth out the traffic egressing an interface. This command is allowed only for the egress direction. This command is mutually exclusive of the **scheduler** and **police** commands.

Perform the following task in global configuration mode.

1. Select the policy map.

```
switch(config)# policy-map p1
```

2. Select the class.

```
switch(config-policymap)# class class-default
```

3. Specify the shaping rate for the port.

```
switch(config-policyclass)# port-shape 3000
```

Attach the scheduler to the class

User can specify the scheduling attributes along with per TC shape rate. There are total of eight queues on an interface. The number of DWRR queues present depends on the SP_COUNT value. For example, if the SP_COUNT is 2, then there are two strict priority queues and six DWRR queues. This command is allowed only for the egress direction and is mutually exclusive of the **scheduler** and **police** commands.

Perform the following task in global configuration mode.

1. Select the policy map.

```
switch(config)# policy-map p1
```

2. Select the class.

```
switch(config-policymap)# class class-default
```

3. Specify the scheduling attributes. For complete information, refer to *Network OS Command Reference*.

```
switch(config-policyclass)# scheduler 3 31000 32000 33000 dwrr 20 20 20 10 10
```

Attach the priority mapping table to the CEE map [define MQC]

Priority-mapping-table is MQC way to bring in partial CEE Map content in an MQC class. MQC does not allow ingress and egress feature to be present in a same policy-map. By definition, they are two different entities and should be provisioned through two separate policy-maps. However, CEE map provisions ingress and egress features in a same provisioning command. Due to this conflict, only the following features are inherited from a CEE map.

- A. Priority-Group Table.
- B. Priority-Mapping Table.
- C. PFC Configuration.
- D. Lossless Priority Remapping.
- E. Fabric Priority Remapping.

In Brocade switches, the CEE map scheduler configuration is global. Unless an egress scheduling policy is applied on an interface, the default scheduler is present.

Perform the following task in global configuration mode.

1. Select the policy map.

```
switch(config)# policy-map p1
```

2. Select the class.

```
switch(config-policymap)# class class-default
```

3. Attach policy map to the CEE map.

```
switch(config-policyclass)# priority-mapping-table import cee default
```

Binding the policy-map to an interface

Use the **service-policy** command to associate a policy-map to an interface to apply policing parameters.

1. Enable the global configuration mode.

```
switch# configure terminal
```

2. Specify the Ethernet interface, as in the following 10-gigabit Ethernet example

```
(config)# interface te 1/1/2
```

3. Bind a policy-map to egress traffic on the interface. The following associates binds policymap1 to outbound traffic on the interface.

```
switch(config-if-te-1/1/2)# service-policy out policymap1
```

You can unbind the policy-map by using the **no** operand.

```
switch(config-if-te-1/1/2)# no service-policy out
```

4. Bind a policy-map to inbound traffic on the interface. The following associates binds policymap1 to inbound traffic on the interface.

```
switch(config-if-te-1/1/2)# service-policy in policymap1
```

You can unbind the policy-map by using the **no** operand.

```
switch(config-if-te-1/1/2)# no service-policy in
```

5. Return to privileged EXEC mode.

```
switch(conf-if-te-1/1/2)# end
```

6. Copy the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Policer binding rules

Consider the following rules when binding a policy-map to an interface:

- You can bind the same policy-map to multiple interfaces but only one policy per interface per direction is allowed.
- You cannot bind policy-maps to an interface if a class map is not associated with the policy-map.
- If policy-map is bound to an interface and the policy-map does not have mandatory Policer attributes, then traffic coming on that interface will be treated as conformed traffic. Packets on that interface will be marked as green and color based actions such as DSCP, CoS mapping and so on will not be applied.

Policing parameters

Policing parameters provide values for CIR, CBS, EIR, and EBS, for classifying traffic by a specific class for color-based priority mapping. They also specify specific actions to perform on traffic with a color-class priority, such as having packet DSCP priority, traffic class (internal queue assignment), or traffic class (internal queue assignment) set to specific values.

CIR and CBS

The Committed Information Rate (CIR) is the maximum number of bits that a port can receive or send during one-second over an interface. For CIR, there are two parameters that define the available traffic: CIR and the Committed Burst Size (CBS). The CIR represents a portion of the interface's total bandwidth expressed in bits per second (bps). It cannot be larger than the interface's total bandwidth. CBS controls the bursty nature of the traffic. Traffic that does not use the configured CIR accumulates credits until the credits reach the configured CBS. These credits can be used when the rate temporarily exceeds the configured CIR. When credits are not available, the traffic is either dropped or subject to the policy set for the Excess Information Rate (EIR). The traffic limited by the CIR can have its priority, traffic class, and DSCP values changed.

CIR is mandatory policing parameter for configuring a class map.

- **cir** *cir-rate*

The **cir** parameter defines the value of the CIR as the rate provided in the *cir-rate* variable. Acceptable values are in multiples of 40000 in the range 40000–40000000000 bps.

- **cbs** *cbs-size*

The **cbs** parameter defines the value of the CBS as the rate provided in the *cbs-size* variable. Acceptable values are 1250–5000000000 bytes in increments of 1 byte.

EIR and EBS

The Excess Information Rate (EIR) provides an option for traffic that has exceeded the CIR. For EIR, there are two parameters that define the available traffic: the EIR and the Excess Burst Size (EBS). The EIR and EBS operate exactly like the CIR and CBS, except that they act only upon traffic that has been passed to the EIR because it could not be accommodated by the CIR. Like the CIR, the EIR provides an initial bandwidth allocation to accommodate inbound and outbound traffic. Like the CBS, the bandwidth available for burst traffic from the EBS is subject to the amount of bandwidth that is accumulated during periods when traffic allocated by the EIR policy is not used. When inbound or outbound traffic exceeds the bandwidth available (accumulated credits or tokens), it is be dropped. The traffic rate limited by the EIR can have its priority, traffic class, and DSCP values changed.

EIR and EBS parameters are optional policing parameters. If not set, they are considered disabled.

- **eir** *eir-rate*

The **eir** parameter defines the value of the EIR as the rate provided in the *eir-rate* variable. Acceptable values are in multiples of 40000 in the range 0–40000000000 bps.

- **ebs** *ebs-size*

The **ebs** parameter defines the value of the EBS as the rate provided in the *ebs-size* variable. Acceptable values are 1250–5000000000 bytes in increments of 1 byte.

Parameters that apply actions to conform and exceed traffic

Following are policing parameters that apply actions to conform or exceed color traffic:

- **conform-set-dscp** *dscp-num*

The **conform-set-dscp** parameter specifies that traffic with bandwidth requirements within the rate configured for CIR will have its packet DSCP priority set to the value specified in the *dscp-num* variable. Acceptable values for *dscp-num* are 0–63.

- **conform-set-prec** *prec-num*

The **conform-set-prec** parameter specifies that traffic with bandwidth requirements within the rate configured for CIR will have its packet IP precedence value (first 3 bits of DSCP) set to the value in the *prec-num* variable. Acceptable values for *prec-num* are 0–7.

- **conform-set-tc** *trafficclass*

The **conform-set-tc** parameter specifies that traffic with bandwidth requirements within the rate configured for CIR will have its traffic class (internal queue assignment) set to the value in the *trafficclass* variable. Acceptable values for *trafficclass* are 0–7.

- **exceed-set-dscp** *dscp-num*

The **exceed-set-dscp** parameter specifies that traffic with bandwidth requirements that exceeds the rate configured for CIR and sent to the EIR bucket will have its packet DSCP priority set to the value in the *dscp-num* variable. Acceptable values for *dscp-num* are 0–63.

- **exceed-set-prec** *prec-num*

The **exceed-set-prec** parameter specifies that traffic with bandwidth requirements that exceed the rate configured for CIR and sent to the EIR bucket will have its packet IP precedence set to the value in the *prec-num* variable. Acceptable values for *prec-num* are 0–7.

- **exceed-set-tc** *trafficclass*

The **exceed-set-tc** parameter specifies that traffic with bandwidth requirements that exceed the rate configured for CIR and is in the limit of what is configured for EIR will have its traffic class (internal queue assignment) set to the value in the *trafficclass* variable. Acceptable values for *trafficclass* are 0–7.

- **set-priority** *priority-mapname*

The **set-priority** parameter specifies the mapping used for setting QoS priority (802.1p priority) in the packet. The *priority-mapname* name variable should be same as configured for the priority-map (police-priority-map), which will have a set priority and color type (conform or exceed).

Displaying policing settings and policy-maps

Use the following commands to display policies configured in policy, class, and priority-maps.

Policy-maps

In the following example, the **show policymap** command is used to display Policer policies and parameters set for the 10-gigabit Ethernet interface 4/1 inbound traffic.

```
switch(conf-if-te-5/1/33)# do show policy-map interface tengigabitethernet 5/1/33
Ingress Direction :
  Policy-Map pmap1
    Class default
      Police cir 43454
      Stats:
        Operational cir:39944 cbs:6518 eir:0 ebs:0
        Conform Byte:0 Exceed Byte:0 Violate Byte:0

Egress Direction :
  Policy-Map pmap1
    Class default
      Police cir 43454
      Stats:
        Operational cir:39944 cbs:6518 eir:0 ebs:0
        Conform Byte:0 Exceed Byte:0 Violate Byte:0
```

Entering **show policymap** without identifying an interface and specify inbound or outbound traffic displays policy-maps bound on all switch interfaces.

```
switch(conf-if-te-5/1/33)# do show policy-map
Number of policy maps : 1
Policy-Map pmap1
  Bound To: te 5/1/33(in), te 5/1/33(out)
sw0(conf-if-te-5/1/33)#

switch(conf-if-te-5/1/33)# do show policy-map detail pmap1
Policy-Map pmap1
  Class default
    Police cir 43454
  Bound To: te 5/1/33(in), te 5/1/33(out)
```

The following example displays the running configured policy-map by means of the **show running-config policy-map** command.

```
switch(conf-if-te-5/1/33)# do show running-config policy-map
```



```

policy-map pmap1
  class default
    police cir 43454

switch(conf-if-te-5/1/33)# do sh ru int te 5/1/33
interface TenGigabitEthernet 5/1/33
  service-policy in pmap1
  service-policy out pmap1
  fabric isl enable
  fabric trunk enable
  no shutdown

```

Class maps

The following example displays the running configured class map name and configured match attribute by means of the **show running-config class-map** command.

```

switch(config-classmap)# do show running-config class-map
class-map cee
!
class-map class_map1
  match access-group stdacl1
!
class-map default

```

Police-priority-maps

The following example displays the running configured police priority-map name and mapping of CoS values for conform and exceed color priorities by means of the **show running-config police-priority-map** command.

```

switch# show running-config police-priority-map
police-priority-map prio_map1
  conform 3 3 3 5 6 1 1 1
  exceed 2 2 2 1 1 1 1 2

```

Considerations and limitations

Consider the following when configuring the port-based Policer feature.

Best practices

Follow these best practices when configuring the port-based Policer feature:

- Avoid mapping lossy priority to lossless priority in conform and exceed CoS maps.
- Configure rate (CIR or EIR) and burst size (CBS or EBS) based on interface speed.
- Set conform and exceed token count (Tc) to the same values to avoid any reordering issues

Configuration rules and considerations

The following are rules for configuring maps and using policing parameters for the Policer feature:

- A policy-map, class map, priority-map name must be unique among all maps of that type.
- A Policer name must begin with a-z, or A-Z. You can use underscore, hyphen, and numeric values 0-9 except as the first character.

- You cannot delete a policy-map, class map, or priority-map if it is active on the interface.
- You cannot delete a class map from a policy-map when the policy-map is active on the interface.
- Configure *CIR* and *EIR* in multiples of 40000 bps.
- Percentage as a rate limit is not supported,
- Policer actions are applicable only to data traffic. Control traffic, FCoE, and internal VLAN traffic is not subjected to policing.
- The egress Policer can overwrite ingress Policer results such as CoS mapping and DSCP mapping.
- If a policy-map is applied to an interface and no Policer attributes are present in that policy-map, then ingress and egress packets on that interface is marked as green (conforming).
- If the configured cbs value is less than 2*MTU value, then 2*MTU is programmed as the CBS in the hardware. For example, if you configure CBS at 4000 bytes and the MTU on an interface is 3000 bytes, when a policy-map is applied on this interface, the CBS programmed in the hardware is 2*MTU (6000 bytes).
- If CBS and EBS values are not configured, then these values are derived from CIR and EIR values, respectively. Burst size calculation is as follows:

$$\text{Burst size (cbs or ebs)} = 1.2 * \text{information rate (CIR/EIR)} / 8$$

- If you do not configure EIR and EBS, then the single-rate, two-color scheme is applied (packets are marked as either green or red).
- You must configure rate limit threshold values on an interface based on interface speed. No validation is performed for user-configured values against interface speed.

Limitations

- The incremental step size for CIR or EIR is set to 40000 bps.
- The Policer operates in color-blind mode. In other words, color is evaluated at ingress and egress Policers independently. This may result in packets that are marked as yellow in the inbound Policer to be evaluated as green at the outbound Policer, depending on Policer settings.
- Because inbound queue scheduling is performed before outbound policing, setting traffic class (set-conform-tc or set-exceed-tc) based on policing results does not effect packet forwarding at the outbound side.
- Packets drops caused by any action other than ACLs are included in Policer counters.
- Layer 3 control packets are policed at the outbound side.
- Policing is enabled on lossless priorities at the outbound side.

Considerations for vLAGs

Because a virtual link aggregation group (vLAG) spans multiple switches, it is not possible to associate flows on each LAG member port to a common Policer. Instead, apply the same policy-map on individual member ports so that traffic flow on member ports is controlled by a Policer configured on that member port. The total rate-limit threshold value on a vLAG consists of the cumulative values of rate-limit thresholds on all member ports.

Policer behavior for control packets

Port-based Policer behavior for Layer 2 and Layer 3 control packets is shown in table [Table 78](#).

TABLE 78 Policer behavior for L2 and L3 control packets

Protocol	Ingress Policer	Egress Policer
LLDP	Enabled if protocol is not enabled and disabled if protocol is enabled.	Disabled
LACP	Enabled if protocol is not enabled and disabled if protocol is enabled.	Disabled
STP	Enabled if protocol is not enabled and disabled if protocol is enabled.	Disabled
DOT1X	Enabled if protocol is not enabled and disabled if protocol is enabled.	Disabled
PIM	Disabled	Enabled
OSPF	Disabled	Enabled
IGMP	Disabled	Enabled
VRRP/VRRP-E	Disabled	Enabled

When a Layer 2 control protocol is not enabled on an interface, packets are dropped during ingress and are subjected to ingress policing. Layer 3 control packets, irrespective of whether they are protocol-enabled or not, will not be subject to ingress and egress policing.

Lossless traffic

The following are considerations for lossless traffic:

- Policing is applicable only for lossy traffic. Lossless traffic should not get policed. For port-based policing, apply a policy-map to an interface even if PFC is configured on that interface. The CoS value (priority) on which PFC is applied is excluded from being policed.
- Remapped priority values should not include lossless priorities. Do not remap lossy traffic priorities to lossless traffic priorities and vice-versa.
- Policer attributes **conform-set-tc** and **exceed-set-tc** should not be set to a lossless traffic class.

Configuring 802.1x Port Authentication

In this chapter

- [802.1x protocol overview](#) 461
- [802.1x authentication configuration tasks](#) 462
- [Interface-specific administrative tasks for 802.1x](#) 462

802.1x protocol overview

The 802.1x protocol defines a port-based authentication algorithm involving network data communication between client-based supplicant software, an authentication database on a server, and the authenticator device. In this situation the authenticator device is the Brocade VDX hardware.

As the authenticator, the Brocade VDX hardware prevents unauthorized network access. Upon detection of the new supplicant, the Brocade VDX hardware enables the port and marks it “unauthorized.” In this state, only 802.1x traffic is allowed. All other traffic, such as DHCP and HTTP, is blocked. The Brocade VDX hardware transmits an Extensible Authentication Protocol (EAP) Request to the supplicant, which responds with the EAP Response packet. The Brocade VDX hardware then forwards the EAP Response packet to the RADIUS authentication server. If the credentials are validated by the RADIUS server database, the supplicant may access the protected network resources.

NOTE

802.1x port authentication is not supported by LAG (Link Aggregation Group) or interfaces that participate in a LAG.

NOTE

The EAP-MD5, EAP-TLS, EAP-TTLS and PEAP-v0 protocols are supported by the RADIUS server and are transparent to the authenticator switch.

When the supplicant logs off, it sends an EAP Logoff message to the Brocade VDX hardware, which then sets the port back to the “unauthorized” state.

802.1x configuration guidelines and restrictions

When configuring 802.1x, be aware of this 802.1x configuration guideline and restriction: If you globally disable 802.1x, then all interface ports with 802.1x authentication enabled automatically switch to force-authorized port-control mode.

802.1x authentication configuration tasks

The tasks in this section describe the common 802.1x operations that you will need to perform. For a complete description of all the available 802.1x CLI commands for the Brocade VDX hardware, see the *Network OS Command Reference*.

Configuring authentication

The **radius-server** command attempts to connect to the first RADIUS server. If the RADIUS server is not reachable, the next RADIUS server is contacted. However, if the RADIUS server is contacted and the authentication fails, the authentication process does not check for the next server in the sequence.

Perform the following steps to configure authentication.

1. Enter global configuration mode.

```
switch# configure terminal
```
2. Use the **radius-server** command to add the RADIUS to the switch as the authentication server. This command can be repeated for additional servers. However, this command moves the new RADIUS server to the top of the access list.

```
switch(config)# radius-server host 10.0.0.5
```
3. Enable 802.1x authentication globally

```
switch(config)# dot1x enable
```
4. Use the **interface** command to select the interface port to modify.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 1/12
```
5. Use the **dot1x authentication** command to enable 802.1x authentication.

```
switch(conf-if-te-1/12)# dot1x authentication
```
6. Return to privileged EXEC mode.

```
switch(conf-if-te-1/12)# end
```
7. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Interface-specific administrative tasks for 802.1x

It is essential to configure the 802.1x port authentication protocol globally on the Brocade VDX hardware, and then enable 802.1x and make customized changes for each interface port. Since 802.1x was enabled and configured in “[802.1x authentication configuration tasks](#)”, use the administrative tasks in this section to make any necessary customizations to specific interface port settings.

802.1x readiness check

The 802.1x readiness check monitors 802.1x activity on all the switch ports and displays information about the devices connected to the ports that support 802.1x. You can use this feature to determine if the devices connected to the switch ports are 802.1x-capable.

The 802.1x readiness check is allowed on all ports that can be configured for 802.1x. The readiness check is not available on a port that is configured by the **dot1x force-unauthorized** command.

When you configure the **dot1x test eapol-capable** command on an 802.1x-enabled port, and the link comes up, the port queries the connected client about its 802.1x capability. When the client responds with a notification packet, it is 802.1x-capable. A RASLog message is generated if the client responds within the timeout period. If the client does not respond to the query, the client is not 802.1x-capable, and a syslog message is generated saying the client is not EAPOL-capable.

Follow these guidelines to enable the readiness check on the switch:

- The readiness check is typically used before 802.1x is enabled on the switch.
- 802.1x authentication cannot be initiated while the 802.1x readiness test is in progress.
- The 802.1x readiness test cannot be initiated while 802.1x authentication is active.
- 802.1x readiness can be checked on a per-interface basis. Readiness check for all interfaces at once is not supported.
- The 802.1x test timeout is shown in **show dot1x** command.
- The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

This example shows how to enable a readiness check on a switch to query a port. It also shows the response received from the queried port verifying that the device connected to it is 802.1x-capable:

```
switch# dot1x test eapol-capable interface gigabitethernet 0/13
DOT1X_PORT_EAPOL_CAPABLE:DOT1X: MAC 00-01-02-4b-f1-a3 on gigabitethernet0/13 is
EAPOL capable.
```

Configuring 802.1x on specific interface ports

To configure 802.1x port authentication on a specific interface port, perform the following steps from privileged EXEC mode. Repeat this task for each interface port you wish to modify.

1. Enter the **configure terminal** command to access global configuration mode.
2. Use the **interface** command to select the interface port to modify.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 1/12
```

3. Use the **dot1x authentication** command to enable 802.1x authentication.

```
switch(conf-if-te-1/12)# dot1x authentication
```

4. Return to privileged EXEC mode.

```
switch(conf-if-te-1/12)# end
```

5. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Configuring 802.1x timeouts on specific interface ports

NOTE

While you are free to modify the timeouts, Brocade recommends that you leave timeouts set to their default values.

To configure 802.1x timeout attributes on a specific interface port, perform the following steps from privileged EXEC mode. Repeat this task for each interface port you wish to modify.

1. Enter the **configure terminal** command to access global configuration mode.
2. Use the **interface** command to select the interface port to modify.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 1/12
```

3. Configure the timeout interval.

```
switch(conf-if-te-1/12)# dot1x timeout supp-timeout 40
```

4. Return to privileged EXEC mode.

```
switch(conf-if-te-1/12)# end
```

5. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Configuring 802.1x re-authentication on specific interface ports

To configure 802.1x port re-authentication on a specific interface port, perform the following steps from privileged EXEC mode. Repeat this task for each interface port you wish to modify.

1. Enter the **configure terminal** command to access global configuration mode.
2. Use the **interface** command to select the interface port to modify.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
switch(config)# interface tengigabitethernet 1/12
```

3. Enable 802.1x authentication for the interface port.

```
switch(conf-if-te-1/12)# dot1x authentication
```

4. Configure reauthentication for the interface port.

```
switch(conf-if-te-1/12)# dot1x reauthentication
switch(conf-if-te-1/12)# dot1x timeout re-authperiod 4000
```

5. Return to privileged EXEC mode.


```
switch(conf-if-te-1/12)# end
```

6. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Configuring 802.1x port-control on specific interface ports

To configure 802.1x port-control on a specific interface port, perform the following steps from privileged EXEC mode. Repeat this task for each interface port you wish to modify.

1. Use the **configure terminal** command to enter global configuration mode.
2. Use the **interface** command to select the interface port to modify.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 1/12
```

3. Enable 802.1x authentication for the interface port.
4. Change the port authentication mode to auto, force-authorized or force-unauthorized.

```
switch(conf-if-te-1/12)# dot1x authentication
```

```
switch(conf-if-te-1/12)# dot1x port-control
auto/force-authorized/force-unauthorized
```

5. Return to privileged EXEC mode.
6. Save the *running-config* file to the *startup-config* file.

```
switch(conf-if-te-1/12)# end
```

```
switch# copy running-config startup-config
```

Re-authenticating specific interface ports

To re-authenticate supplicant connected to a specific interface port, perform the following steps from privileged EXEC mode. Repeat this task for each interface port you wish to reauthenticate.

1. Use the **configure terminal** command to enter global configuration mode.
2. Use the **interface** command to select the interface port to modify.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 1/12
```

3. Start re-authentication on a port where dot1x is already enabled.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8.

```
switch(conf-if-te-1/12)# dot1x reauthenticate
```

4. Return to privileged EXEC mode.

```
switch(conf-if-te-1/12)# end
```

5. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Disabling 802.1x on specific interface ports

To disable 802.1x authentication on a specific interface port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Use the **interface** command to select the interface port to modify.

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 1/12
```

3. Use the **no dot1x port-control** command to disable 802.1x Authentication.

```
switch(conf-if-te-1/12)# no dot1x authentication
```

4. Return to privileged EXEC mode.

```
switch(conf-if-te-1/12)# end
```

5. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Disabling 802.1x globally

To disable 802.1x authentication globally, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Use the **no dot1x enable** command to disable 802.1x Authentication.

```
switch(config)# no dot1x enable
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Save the *running-config* file to the *startup-config* file.

```
switch# copy running-config startup-config
```

Checking 802.1x configurations

To check 802.1x configurations, perform the following steps from privileged EXEC mode.

1. To view all dot1x configuration information, use the **show dot1x** command with the **all** operand.

```
switch# show dot1x all
```

2. To check 802.1x configurations for specific interface ports, use the **interface** command to select the interface port to modify.

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)#interface tengigabitethernet 1/12
```

3. To check 802.1x authentication statistics on specific interface ports, use the **show dot1x** command with the **statistics interface** operand.

```
switch# show dot1x statistics interface tengigabitethernet 1/12
```

4. To check all diagnostics information of the authenticator associated with a specific interface port, use the **show dot1x** command with the **diagnostics interface** operand.

```
switch# show dot1x diagnostics interface tengigabitethernet 1/12
```

5. To check all statistical information of the established session, use the **show dot1x** command with the **session-info interface** operand.

```
switch# show dot1x session-info interface tengigabitethernet 1/12
```

30 Interface-specific administrative tasks for 802.1x

Configuring sFlow

In this chapter

- [sFlow protocol overview](#) 469
- [Configuring the sFlow protocol globally](#) 470
- [Configuring sFlow for interfaces](#) 471
- [Hardware support matrix for sFlow](#) 473

sFlow protocol overview

The sFlow protocol is an industry-standard technology for monitoring high-speed switched networks. The sFlow standard consists of an sFlow agent that resides anywhere within the path of the packet and an sFlow collector that resides on a central server. This release is compliant with sFlow Version 5.

The sFlow agent combines the flow samples and interface counters into sFlow datagrams and forwards them to the sFlow Collector at regular intervals. The datagrams consist of information on, but not limited to, packet header, ingress and egress interfaces, sampling parameters, and interface counters. Packet sampling is typically performed by the ASIC. The sFlow collector analyzes the sFlow datagrams received from different devices and produces a network-wide view of traffic flows. You can configure up to five collectors, using both IPv4 and IPv6 addresses.

The sFlow datagram provides information about the sFlow version, its originating agent's IP address, a sequence number, one sample, and protocol information.

The sFlow agent uses two forms of operation:

- Time-based sampling of interface counters
- Statistical sampling of switched packets

Note the following limitations:

- If both port-based sampling and flow-based sampling are enabled on an interface, samples are based on port rate only.
- Flow-based sFlow is not supported on port channels, FCoE ports, or ISL ports.

Interface flow samples

A flow sample is based on random packets being forwarded to the sFlow collector at defined numeric intervals, either for the entire Brocade switch or for a single port interface. For example, every 4,096th packet is forwarded to the sFlow collector for analysis and storage.

The sampling rate is adaptive, and the sFlow agent is free to schedule the sampling to maximize internal efficiency.

NOTE

This type of random sampling provides estimated flow rates, but not perfect accuracy.

Packet counter samples

A polling interval defines how often the sFlow octet and packet counter for a specific interface are sent to the sFlow collector, but the sFlow agent is free to schedule the polling in order to maximize internal efficiency.

Configuring the sFlow protocol globally

Brocade recommends that you globally configure sFlow on the Brocade switch first, and then enable sFlow on specific interface ports and make custom alterations, because sFlow parameters at the interface level can differ from those at the global level. For details, refer to [“Configuring sFlow for interfaces”](#) on page 471.

Enabling sFlow globally does not enable it on all interface ports. sFlow must be explicitly enabled on all the required interface ports. Refer to [“Enabling and customizing sFlow on specific interfaces”](#) on page 471.

NOTE

On the Brocade VDX 8770, Switched Port Analyzer (SPAN), and sFlow can be enabled at the same time. However, on the Brocade VDX 6720, SPAN and sFlow cannot be enabled at the same time.

For complete information on the sFlow CLI commands for the Brocade switch, refer to the *Network OS Command Reference*.

To configure sFlow globally, perform the following steps in global configuration mode.

1. Globally enable the sFlow protocol.

```
switch(config)# sflow enable
```

2. Designate the IP address (up to five addresses) for the sFlow collector server. Optionally, you can designate the port number.

NOTE

Both IPv4 and IPv6 addresses are supported. However, each address must be entered individually by means of a separate **sflow collector** command.

```
switch(config)# sflow collector 10.10.138.176 6343
switch(config)# sflow collector fd00::1900:4545:3:200:f8ff:fe21:67cf port
6343switch(config)# sflow collector fd00::200:f8ff:fe21:67cf
```

3. Set the sFlow polling interval (in seconds).

```
switch(config)# sflow polling-interval 35
```

4. Set the sFlow sample-rate.

```
switch(config)# sflow sample-rate 4096
```

5. Return to privileged EXEC mode.

```
switch(config)# end
```

6. Confirm the sFlow configuration status by using the **show sflow** or **show sflow all** commands.

```
switch# show sflow
sFlow services are:                enabled
Global default sampling rate:      32768 pkts
Global default counter polling interval: 20 secs
Collector server address           Number of samples sent
-----
fd00::1900:4545:3:200:f8ff:fe21:67cf      0
fd00::200:f8ff:fe21:67cf                 0
10.10.138.176                             0

switch# show sflow all
sFlow services are:                enabled
Global default sampling rate:      32768 pkts
Global default counter polling interval: 20 secs
Collector server address           Number of samples sent
-----
fd00::1900:4545:3:200:f8ff:fe21:67cf : 6343    0
fd00:fd00::200:f8ff:fe21:67cf :              0
10.10.138.176: 6343                        0
```

7. Clear any existing sFlow statistics to ensure accurate readings.

```
switch# clear sflow statistics
```

Configuring sFlow for interfaces

After the global sFlow configuration, sFlow must be explicitly enabled on all the required interface ports.

NOTE

When sFlow is enabled on an interface port, it inherits the sampling rate and polling interval from the global sFlow configuration.

This section presents the following tasks:

- [Enabling and customizing sFlow on specific interfaces](#)
- [Configuring an sFlow policy map and binding it to an interface](#)
- [Configuring flow-based sFlow](#)
- [Disabling sFlow on specific interfaces](#)

Enabling and customizing sFlow on specific interfaces

NOTE

On the Brocade VDX 8770, SPAN and sFlow can be enabled at the same time. However, on the Brocade VDX 6710, VDX 6720, or VDX 6730, SPAN and sFlow cannot be enabled at the same time.

Perform the following steps in privileged EXEC mode to enable and customize sFlow on an interface. This task assumes that sFlow has already been enabled at the global level; refer to [“Configuring the sFlow protocol globally”](#) on page 470.

1. Enter the **interface** command to specify the DCB interface type, the RBridge ID, and the slot/port number.

NOTE

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following example format:

```
switch(config-if-gi-22/0/1)#
```

```
switch(config)# interface tengigabitethernet 0/16
```

2. Configure the sFlow polling interval.

```
switch(config-if-te-0/16)# sflow polling interval 35
```

3. Use the **sflow enable** command to enable sFlow on the interface.

```
switch(config-if-te-0/16)# sflow enable
```

4. Set the sFlow sample-rate.

```
switch(config-if-te-0/16)# sflow sample-rate 8192
```

5. Confirm the sFlow configuration status on a specific interface.

```
switch# show sflow interface tengigabitethernet 22/0/16
sFlow info for interface TenGigabitEthernet 22/0/16
-----
Configured sampling rate:      100 pkts
Actual sampling rate:         100 pkts
Counter polling interval:     100 secs
Samples received from hardware: 32
Port backoff-threshold :      272
Counter samples collected :    147
```

Configuring an sFlow policy map and binding it to an interface

Perform the following steps, beginning in global configuration mode.

1. Create standard MAC access control list (ACL).

```
switch# mac access-list standard acl1
switch(config-macl-std)# permit any
```

2. Create a class map and attach the ACL to the class map.

```
switch(config-macl-std)# class-map class1
switch(config-classmap)# match access-group acl1
```

3. Create a policy map and attach the class map to the policy map.

```
switch(config-classmap)# policy-map policy1
switch(config-policymap)# class class1
```

4. Add an sFlow profile name by using the **map** command.

This example assigns the profile name “policy1.”

```
switch(config-policymap-class)# map sflow policy1
```

5. Bind the policy map to an interface.

```
switch(config-if-te-1/8/1)# service-policy in policy1
```


Disabling sFlow on specific interfaces

NOTE

Disabling sFlow on the interface port does not completely shut down the network communication on the interface port.

To disable sFlow on a specific interface, perform the following steps in interface configuration mode.

1. Disable the sFlow interface.

```
switch(config-if)# no sflow enable
```

2. Return to privileged EXEC mode.

```
switch(config-if)# end
```

3. Confirm the sFlow configuration status on the specific interface.

NOTE

The **gigabitethernet rbridge-id/slot/port** operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following example format:

```
switch(config-if-gi-22/0/1)#
```

```
switch# show sflow interface tengigabitethernet 0/12
```

Hardware support matrix for sFlow

Table 79 describes which sFlow features are supported on Brocade hardware.

TABLE 79 sFlow feature support

Feature	Brocade VDX 8770	Brocade VDX 67xx
sFlow global configurations for enabling sFlow, polling interval, collector, and sample rate	All are supported	All are supported
sFlow data source interface	Supports 1-Gbps, 10-Gbps, and 40-Gbps interfaces	Supports on 10-Gbps interfaces only
sFlow data source: Front port trunks and VLANs	Not supported	Not supported
sFlow scanning for inbound, outbound, or both directions on a port	Supports inbound only	Supports inbound only
sFlow counter polling support on per-port, per-VLAN, or per-trunk	Supports only per-port counter polling	Supports only per-port counter polling
All standard if_counters and Ethernet counters	Supported	Supported
Multiple collector configuration	A maximum of five collectors can be configured.	A maximum of five collectors can be configured.
Extended Gateway, Extended router, and NAT/MPLS/URL header formats	Not supported	Not supported
Subagent-ID	Filled with slot number of the interface	Filled with a zero (0)

TABLE 79 sFlow feature support (Continued)

Feature	Brocade VDX 8770	Brocade VDX 67xx
Agent IP address	Preference 1: Chassis IP Preference 2: Management CP IP	Management IP
Maximum packets per second	272 pkts/sec/ASIC VDX 8770-4: 6528 pkts/sec VDX 8770-8: 13056 pkts/sec	96 pkts/sec/ASIC Each ASIC supports eight front-user ports on the 48x10G and 48x1G line cards, and supports two front-user ports on the 12x48G Line card.
Sample rate calculation	Dropped packets (such as errors and ACL dropped packets) are not counted for the calculations used for sample generation	Dropped packets (such as errors and ACL dropped packets) are counted for the calculations used for sample generation
Maximum sFlow raw packet header size	228 bytes The hardware truncates the packet.	128 bytes The software truncates the packet.
SPAN and sFlow configurations	SPAN and sFlow can be enabled at the same time	SPAN and sFlow cannot be enabled at the same time

Flow-based sFlow

Note the following considerations and limitations for flow-based sFlow:

- A maximum of 16 profiles is allowed on a Brocade VDX 8770, and 8 on a Brocade VDX 6740.
- The purpose of flow-based sFlow is not to drop or trap packets on the basis of ACLs, but rather just to match traffic. Packets are still sampled and allowed, and sFlow samples are generated for any rule.
- Port-based sFlow takes precedence over flow-based sFlow. Both cannot operate simultaneously.
- On a Brocade VDX 8770, if a packet is classified as a Layer 3 IPv4 packet, with a match on destination address and IP type, and even if a packet's destination or source address matches a Layer 2 ACL, flow-based sFlow samples are not generated. This is not the case with a Brocade VDX 6740.
- Because all samples of different rates are sent to the collector from a single port, only the lowest sampling rate is used.

On a Brocade VDX 6740, if the traffic is Layer 2 then samples with a MAC ACL sample rate are collected. If the traffic is Layer 3, then samples with a Layer 3 ACL sample rate are collected.

Configuring flow-based sFlow

Flow-based sFlow is used to analyze a specific type of traffic (flow based on access control lists, or ACLs). This involves configuring an sFlow policy map and binding it to an interface.

NOTE

The "deny ACL" rule is not supported for flow-based sflow. Only the permit action is supported.

Perform the following steps, beginning in global configuration mode.

1. Create an sFlow profile. Be sure to specify that sampling rate as a power of 2.

```
switch(config)# sflow-profile profile1 sampling-rate 256
```

2. Create a standard MAC ACL.

```
switch# mac access-list standard acl1
switch(config-macl-std)# permit any
```

3. Create a class map and attach the ACL to the class map.

```
switch(config-macl-std)# class-map class1
switch(config-classmap)# match access-group acl1
```

4. Create a policy map and attach the class map to the policy map.

```
switch(config-classmap)# policy-map policy1
switch(config-policymap)# class class1
```

5. Add an sFlow profile name by using the **map** command.

This example assigns the profile name "policy1."

```
switch(config-policymap-class)# map sflow profile1
```

6. Switch to interface configuration mode.

```
switch(config-policymap-class)# exit
switch(config)# interface ten 1/8/1
switch(config-if-te-1/8/1)#
```

7. Bind the policy map to an interface.

```
switch(config-if-te-1/8/1)# service-policy in policy1
```

Disabling flow-based sFlow on specific interfaces

NOTE

Disabling sFlow on the interface port does not completely shut down the network communication on the interface port.

To disable sFlow on a specific interface, perform the following steps in interface configuration mode.

1. Disable the sFlow interface.

```
switch(config-if)# no sflow enable
```

2. Return to privileged EXEC mode.

```
switch(config-if)# end
```

3. Switch to interface configuration mode.

```
switch(config-policymap-class)# exit
switch(config)# interface ten 1/8/1
switch(config-if-te-1/8/1)#
```

4. Disable flow-based sFlow by removing the policy map.

```
switch(config-if-te-1/8/1)# no service-policy in
```

5. Confirm the sFlow configuration status on the specific interface.

NOTE

The **gigabitethernet** *rbridge-id/slot/port* operand is used only for the Brocade VDX 6710, Brocade VDX 8770-4, and Brocade VDX 8770-8. The prompt for these ports is in the following example format:

```
switch(config-if-gi-22/0/1)#
```

```
switch# show sflow interface tengigabitethernet 0/12
```

Configuring Switched Port Analyzer

In this chapter

• Switched Port Analyzer protocol overview	477
• Configuring ingress SPAN	478
• Configuring egress SPAN	479
• Configuring bidirectional SPAN	479
• Deleting a SPAN connection from a session	480
• Deleting a SPAN session	480
• SPAN in management cluster	481
• Using RSPAN	482

Switched Port Analyzer protocol overview

Switched Port Analyzer is used on a network switch to send a copy of network packets seen on one switch port to a network monitoring connection on another switch port. If you are interested in listening or snooping on traffic that passes through a particular port, Switched Port Analyzer (SPAN) artificially copies the packets to a port connected to your analyzer. Usually, this traffic is limited to incoming or outgoing packets, but Network OS v4.0.0 allows bidirectional traffic monitoring on the source port.

SPAN guidelines and limitations

Note the following guidelines for and limitations of SPAN connections:

- For the Brocade VDX 6720-24:
 - The mirror port can be any port in the switch.
 - Only one port per switch can be configured as a destination port for ingress mirroring.
 - Only one port per switch can be configured as a destination port for egress mirroring.
- For the Brocade VDX 6720-60:
 - The mirror port should be in the same port-group as the source port.
 - Only one port per port-group can be configured as a destination port for ingress mirroring.
 - Only one port per port-group can be configured as a destination port for egress mirroring.
- The mirror port should not be configured to carry normal traffic.
- A port cannot be mirrored to multiple locations in the same direction.
- A port cannot be made a destination port for bidirectional mirroring if a different port supported by that ASIC is already configured as destination port for any type of mirroring.

- If a port is configured as a destination port for bidirectional mirroring, no other port supported by that ASIC can be made a destination port for any type of mirroring.
- The destination mirror port can handle from 1 to 40Gbps (line rate) worth of mirror traffic, depending on the capability of the destination port. If multiple ports, or both flows on the same port, are mirrored to the same destination mirror port, then only the destination port's capacity worth of mirror traffic is mirrored and the remaining traffic is ignored.
- If the source port receives burst traffic and the destination mirror port cannot handle all the bursts, some of the burst traffic is not mirrored.
- Mirroring of ISL ports is supported, but the destination port should reside on the same RBridge.
- Mirroring of LAG or Port-Channel interfaces is not supported, but LAG members can be mirrored.
- TRILL ports cannot be designated as a destination port.
- TRILL ports can be a source port, but mirroring is restricted to the port local to the source node ports.
- Inter-ASIC port mirroring is not allowed on the Brocade VDX 6720-60.
- Ethernet Pause frames are not mirrored.
- Mirroring of trunk port is not supported, though the ASIC supports the mirroring of a trunk. To mirror a trunk, you must individually enable mirroring on all member ports.
- The multicast and broadcast statistics are correctly updated on TX ports for mirrored traffic.
- All commands except for **shutdown** and **no shutdown** are blocked on a destination mirror port.
- The interface counters are cleared when a port is successfully designated as a destination mirror port.
- The **show interface** command hides the Receive Statistics and Rate Info (Input) information for a destination mirror port.
- The MTU of a port should be set to the default value of 2500 bytes before it is made a destination mirror port. When the port is successfully designated as the destination mirror, the MTU of that port is automatically set to the maximum value of 9216 bytes. When the port becomes a nondestination mirror, the MTU is restored to the default value.
- Port mirroring is supported on any physical front-end user-configurable port. The source port can be part of a LAG, VLAG, VLAN, or any other user configuration
- A maximum of 512 mirror sessions are supported in management cluster and fabric cluster modes, but 24 sessions in standalone mode.

Configuring ingress SPAN

To configure SPAN for incoming packets only, do the following.

1. Open a monitor session and assign a session number

```
switch(config)# monitor session 1
```
2. Configure the source port and the destination port, with the **rx** parameter for received packets.

The destination port is always an external port. The source and destination ports must be in the same port group for the Brocade VDX 6720-60.

```
switch(config-session-1)# source tengigabitethernet 1/0/15 destination
tengigabitethernet 1/0/18 direction rx
```

NOTE

If the following error is displayed, use the interface **no lldp** command to disable LLDP on the destination port before preceding:

```
% Error: Destination port cannot be in L2/L3/Qos/ACL/802.1x/LAG
member/Lldp/Port-profile/non-default-MTU
```

3. *Optional:* Use the **description** command to add a label to the monitor session.

```
switch(config-session-1)# description Hello World!
```

4. Repeat [step 1](#) and [step 2](#) as needed for additional ports.

A monitor session can have only one source port. For additional ports you must create additional monitor sessions as needed for additional port mirroring sessions.

Configuring egress SPAN

To configure SPAN for outgoing packets only, do the following.

1. Open a monitor session and assign a session number

```
switch(config)# monitor session 1
```

2. Configure the source port and the destination port, with the **tx** parameter for transmitted packets.

The destination port is always an external port. The source and destination ports must be in the same port group for the Brocade VDX 6720-60.

```
switch(config-session-1)# source tengigabitethernet 1/0/15 destination
tengigabitethernet 1/0/18 direction tx
```

NOTE

If the following error is displayed, use the interface **no lldp** command to disable LLDP on the destination port before preceding:

```
% Error: Destination port cannot be in L2/L3/Qos/ACL/802.1x/LAG
member/Lldp/Port-profile/non-default-MTU
```

3. *Optional:* Use the **description** command to add a label to the monitor session.

```
switch(config-session-1)# description Hello World!
```

4. Repeat [step 1](#) and [step 2](#) as needed for additional ports.

A monitor session can have only one source port. For additional ports you must create additional monitor sessions as needed for additional port mirroring sessions.

Configuring bidirectional SPAN

To configure SPAN for packets traveling in both directions, do the following.

32 Deleting a SPAN connection from a session

1. Open a monitor session and assign a session number

```
switch(config)# monitor session 1
```

2. Configure the source port and the destination port, with the **both** parameter for all packets.

The destination port is always an external port. The source and destination ports must be in the same port group for the Brocade VDX 6720-60.

```
switch(config-session-1)# source tengigabitethernet 1/0/15 destination  
tengigabitethernet 1/0/18 direction both
```

NOTE

If the following error is displayed, use the interface **no lldp** command to disable LLDP on the destination port before preceding:

```
% Error: Destination port cannot be in L2/L3/Qos/ACL/802.1x/LAG  
member/Lldp/Port-profile/non-default-MTU
```

3. *Optional:* Use the **description** command to add a label to the monitor session.

```
switch(config-session-1)# description Hello World!
```

4. Repeat [step 1](#) and [step 2](#) as needed for additional ports.

A monitor session can have only one source port. For additional ports you must create additional monitor sessions as needed for additional port mirroring sessions.

Deleting a SPAN connection from a session

To remove a single connection from a SPAN session, do the following.

1. Display the existing configuration of the monitor session.

```
switch# show monitor session 1
```

2. Open an existing monitor session.

```
switch(config)# monitor session 1
```

3. Use the **no** option to delete a particular port connection.

```
switch(config-session-1)# no source tengigabitethernet 1/0/15 destination  
tengigabitethernet 1/0/18 direction both
```

4. Display the monitor session again to confirm the deletion of the connection.

```
switch# show monitor session 1
```

Deleting a SPAN session

To remove a SPAN session, do the following.

1. Display the existing configuration of the monitor session.

```
switch# show monitor session 1
```

2. Delete the existing monitor session by using the **no** option.

```
switch(config)# no monitor session 1
```


3. Return to Privileged EXEC mode with the **exit** command.
4. Display the monitor session again to confirm the deletion of the connection.

```
switch# show monitor session 1
```

SPAN in management cluster

SPAN in management cluster supports mirroring of a source port to a destination port lying on a different switch in the management cluster. SPAN in management cluster is configured in the same manner, with the exception of the **source** command.

The source command controls the source and destination switches in the management cluster by the interface designation. The source and destination port can be anywhere in the management cluster. In this example, the source is set as the third switch in the management cluster by the **3/0/15** operand. However the destination is set to the fifth switch in the management cluster by the **5/0/18** operand.

```
switch(config-session-1)# source tengigabitethernet 3/0/15 destination
tengigabitethernet 5/0/18 direction tx
```

This configuration rule applies to **ingress**, **egress**, and **both** directions of SPAN. Otherwise, configure SPAN as you would in standalone mode. Refer to [“Switched Port Analyzer protocol overview”](#) on page 477.

The **show monitor** and **show span path** commands display the source and destination switches, as shown in this example:

```
switch# show span path session 1
Session                               :1
Path                                  :Te 1/0/10 -> Te 3/0/15 (ISL-exit port) -> Te
5/0/18

switch# show monitor
Session                               :1
Description                           :Test monitor session
State                                 :Enabled
Source interface                       :Te 3/0/15 (Up)
Destination interface                  :Te 5/0/18 (Up)
Direction                             :Rx
```

In addition to the standard SPAN limitations (refer to “[SPAN guidelines and limitations](#)” on page 477), note the following guidelines and limitations for SPAN in management cluster:

- The Brocade VDX 6720 is not supported as SPAN source node in management cluster but it can act as a destination node.
- The Brocade VDX 6740, VDX 6740T, and VDX 6730 are supported as both source and destination nodes for SPAN in management cluster.
- SPAN in management cluster supports up to 512 sessions.
- For SPAN with the destination port residing on a remote node, the **show span path session session-number** command shows the path taken by the mirrored packets in the cluster.

Using RSPAN

Remote SPAN, or RSPAN, extends SPAN by enabling remote monitoring of multiple switches across your network. The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches. The SPAN traffic from the sources is copied onto the RSPAN VLAN and then forwarded over trunk ports that are carrying the RSPAN VLAN to any RSPAN destination sessions monitoring the RSPAN VLAN.

NOTE

RSPAN is supported only on Brocade VDX 8770 and VDX 6740 platforms.

RSPAN consists of an RSPAN source interface and an RSPAN VLAN. The configured source port is mirrored to the RSPAN VLAN and the ports that are members of this VLAN receive the mirrored traffic.

All participating switches must be trunk-connected at Layer 2, and the remote VLAN must be configured on all the switches participating in the RSPAN session.

Note the following configurations and restrictions for RSPAN:

Basic considerations:

- All participating switches must be connected by Layer 2 trunks.
- ISL mirroring is not supported on RSPAN.
- The source and destination ports cannot both be TRILL (ISL) ports.
- RSPAN supports multi-hop.
- RSPAN can support both the fabric cluster and management cluster modes. However, using RSPAN in management cluster mode uses unnecessary ISL bandwidth because it floods the traffic on the ISL as well as the trunk port.
- If the source port is not Layer 2 and untagged traffic is mirrored, it will be dropped for RSPAN because untagged and unclassified traffic is dropped on an ISL trunk.
- On the Brocade VDX 6740, if source port is in unknown mode, that is neither Layer 2 nor Layer 3, the packets are dropped and are not mirrored.
- Ethernet Pause frames are not mirrored.

VLAN considerations:

- Before you configure an RSPAN session, you must create the RSPAN VLAN.
- A native VLAN cannot be made the RSPAN VLAN.

- The VLAN used for RSPAN should not be used for other purposes; furthermore, if the VLAN has ports as its members, it cannot be made an RSPAN VLAN. Only when the session is deconfigured, and the VLAN is deleted as an RSPAN VLAN, should the VLAN number be used for another purpose.
- You can configure any VLAN as an RSPAN VLAN as long as all participating network devices support the configuration of RSPAN VLANs and you use the same RSPAN VLAN for each RSPAN session in all participating network devices.
- You must configure the RSPAN VLANs on all source, intermediate, and destination network devices.
- Do not configure any ports in an RSPAN VLAN except the ports selected to carry RSPAN traffic. However, all configurations are allowed on the RSPAN destination port.
- The vlan-id of the packets marked for RSPAN will change to the RSPAN vlan-id.
- Access ports can be added to an RSPAN VLAN as destination ports.
- MAC address learning is disabled in the RSPAN VLAN.

Mirroring across RSPAN

Network OS v4.0.0 uses ISL to mirror packets across RBridges to reach the destination. All the SPAN standalone commands function for RBridges with the following exceptions:

- The source port cannot be a Brocade VDX6720-60 or VDX6720-24 port.
- If there is congestion at the ingress span queue resulting from bandwidth-related backpressure at the ISL, the SPAN mirrored packets will be dropped and traffic will be lost.
- FCoE mirroring is not supported.
- TRILL ports cannot be destination ports.
- A TRILL port can be a source port, but its mirroring is restricted to the local node only.
- For the traffic flowing across the switch, if the source port is in unknown mode (the node is in Layer 2 or Layer 3), then the untagged packets are dropped.
- Ethernet Pause frames are not mirrored across RBridges.

Configuring RSPAN

The principal difference between configuring SPAN and RSPAN is that RSPAN requires a remote VLAN to be created first, by means of the **rspan-vlan** command. This example demonstrates the configuration of a bidirectional RSPAN.

1. Create a remote VLAN on the destination interface.

```
switch(config)# interface vlan 100
```

2. Execute the **rspan-vlan** command to make the VLAN remote.

```
switch(config-vlan-100)# rspan-vlan
```

3. Exit the VLAN configuration mode.

```
switch(config-vlan-100)#end
```

4. Open a monitor session and assign a session number

```
switch(config)# monitor session 1
```

5. Configure the source port and the destination port, with the **both** parameter for bidirectional port mirroring.

By modifying the direction parameter, you can control whether this is an ingress, egress, or a bidirectional SPAN.

In the case of RSPAN, the destination is the VLAN, instead of a destination interface.

```
switch(config-session-1)# source tengigabitethernet 1/0/15 destination
rspan-vlan 100 direction both
```

6. *Optional:* Use the **description** command to add a label to the monitor session.

```
switch(config-session-1)# description Hello World!
```

7. Use the **switchport** command to add a port to the RSPANVLAN to access the mirrored packets.

```
switch(config-session-1)# exit
switch (config)# interface ten 1/0/15
switch(conf-if-te-1/0/15)# switchport access rspan-vlan 100
```

8. Display the results of the configuration.

```
switch(conf-if-te-1/0/14)# do show vlan rspan-vlan
```

Configuring SFP breakout mode

In this chapter

- [SFP breakout overview](#) 485
- [Configuring breakout mode \(static configuration\)](#) 487
- [Breakout configuration examples](#) 489
- [Breakout limitations](#) 492
- [High availability considerations](#) 492

SFP breakout overview

SFP breakout is a new port configuration parameter. Breakout interfaces are those interfaces created on the breakout SFP. The number of interfaces created is dependent on the SFP type. For example, when a Quad SFP (QSFP) is not in breakout mode, only one 40-Gbps interface exists; however, when that QSFP has breakout mode enabled, four 10-Gbps interfaces are created. These interfaces, no matter whether breakout mode is enabled or disabled, are administered and operate exactly the same as any other interface created on a regular SFP with no breakout capability. As a result, existing DCE module operations are not affected.

While the DCE module has no dependency on the breakout interfaces operation, the interface creation and deletion is also similar to that for regular interfaces, so that modules do not require special handling as a result of an SFP mode change. Currently, only static configuration is supported, whereby the reincarnation of a breakout interface resembles the sequence that an interface executes during a slotpoweroff-no linecard-slotpoweron operation. If the existing DCE module supports this line-card removal sequence, the changed SFP mode operational sequence is implicitly supported. The slight difference is that the SFP mode change initiates a “no linecard” equivalent at the port level, rather than attempting to remove all configurations for all interfaces on the line card.

A breakout interface basically supports all operations or configurations that a regular interface supports (with few exceptions, which are noted in [“Breakout limitations”](#) on page 492). As such, it has the following properties:

- Has its own admin and operational state.
- Has its own ASIC resources interface statistics.
- Supports any configuration applicable to any regular SFP interface.
- Can be a port-channel or vLAG member.
- Can be static or dynamic depending on the targeted platform or line card.

The default state for an SFP is “no breakout.”

Table 80 lists the current platforms that support breakout mode.

TABLE 80 Breakout supported platforms

Platform	Line card configuration	QSFP ports
VDX 6740 VDX 6740T	48 10G plus 4 40G	4
VDX 8770-4 VDX 8770-8	12x 40G	12

The SFP connector-id identifies a physical front-end SFP and has the same meaning as the port-id used in the interface name. The connector-id infers which interfaces are created or deleted as a result of the breakout mode change. All interfaces at this connector must be disabled before the command is accepted.

For example, consider line card 2 port 1 on a node with rbridge-id 3 that has a Quad SFP (QSFP) in breakout disabled mode. If breakout is then set to enabled, the existing interface Fo 3/2/1 is deleted and four new interfaces—Te 3/2/1:1, Te 3/2/1:2, Te 3/2/1:3 and Te 3/2/1:4—are created. When the interface is deleted, any configuration on this interface is deleted as well. The new interfaces created have the default port configuration. If breakout is then set to disabled, the 4 Te interfaces are deleted and a single Fo interface is created.

The condition in which the command is executed and becomes effective is both line-card-dependent and platform-dependent.

- On chassis systems, if a line card supports only static configuration, then this command is supported only when the line card is in the power off state, and the mode change is effective only after the switch is powered on.
- On the Brocade VDX 6740 and 6740T, this command is supported only when the interface is in the shutdown state, and the mode change is effective only after the switch reboots.

The interface created on an SFP in breakout-enabled mode uses a new interface naming nomenclature. A numeric suffix is added to the existing interface name, which is separated by a colon. This nomenclature identifies that a port is in breakout mode.

Table 81 shows an example of an SFP in breakout mode and its respective interface names.

TABLE 81 SFP breakout values

SFP # (rbridge/slot/port)	SFP type	Interface name	
		Breakout disabled	Breakout enabled
3/2/1	QSFP (4 x10G)	Fo 3/2/1	Te 3/2/1:1
			Te 3/2/1:2
			Te 3/2/1:3
			Te 3/2/1:4

Configuring breakout mode (static configuration)

For a chassis system

To configure breakout mode on a blade in a chassis, complete the following procedure.

1. Slotpoweroff the linecard.
 - a. Applications will remove all interfaces.
 - b. Operational CLIs will not show any interfaces on this linecard.
 - c. Interface and interface configurations still exist in the DCM database. The **show running interface** command will show interface and interface configurations.
2. Apply the SFP breakout command to the target port(s).
 - a. Existing interface(s) under the old mode and the associated configurations are removed from the DCM database. If the target is a QSFP, the disabled breakout deletes four Te interfaces and the enabled breakout deletes one Fo interface. Configuration of a nontarget port is not affected.
 - b. SFP breakout mode is updated in the DCM database.
3. Copy the running-config to the startup-config [(or standalone (and fabric cluster modes)).
 - a. If you are in standalone mode or fabric cluster mode, you must save the running-config to the startup-config. The SFP configuration must be saved in the startup database, which is used after the line card is rebooted.
 - b. If you are in logical chassis cluster mode, you do not need to copy the running-config to the startup-config.
4. Slotpoweron the linecard.
 - a. The platform creates interfaces corresponding to the SFP breakout mode of each port. For QSFP, a single Fo interface is created in disable mode and four Te interfaces are created in enable mode.
 - b. The SFP interfaces under the new mode come up in default configurations, just as if the system were booted up for the first time. Unaffected interfaces retain their original configurations before the slotpoweroff was applied.

The following example shows the enable breakout mode on port 11.

```
switch# show ip int bri
```

Interface	IP-Address	Status	Protocol
=====	=====	=====	=====
FortyGigabitEthernet 1/2/1	unassigned	up	up
FortyGigabitEthernet 1/2/2	unassigned	up	up
FortyGigabitEthernet 1/2/3	unassigned	up	up
FortyGigabitEthernet 1/2/4	unassigned	up	up
FortyGigabitEthernet 1/2/5	unassigned	up	up
FortyGigabitEthernet 1/2/6	unassigned	up	up
FortyGigabitEthernet 1/2/7	unassigned	up	up
FortyGigabitEthernet 1/2/8	unassigned	up	up
FortyGigabitEthernet 1/2/9	unassigned	up	up
FortyGigabitEthernet 1/2/10	unassigned	up	up
FortyGigabitEthernet 1/2/11	unassigned	up	up

33 Configuring breakout mode (static configuration)

```
FortyGigabitEthernet 1/2/12      unassigned      up      up

#Power off line card
switch# power-off linecard 2
switch# show ip int bri

Interface                        IP-Address      Status      Protocol
=====
switch# config t
switch(config)# hardware
switch(config-hardware)#
switch(config-hardware)# connector 1/2/11
switch(config-connector-1/2/11)# sfp breakout

# Interfaces under new mode are not created at this time..
switch# show ip int bri
Interface                        IP-Address      Status      Protocol
=====

#In SA/FC mode, save running-config to start-up config"
switch# copy running-config startup-config

#Power on line card
switch# power-on linecard 2
switch# show ip int bri

Interface                        IP-Address      Status      Protocol
=====
FortyGigabitEthernet 1/2/1      unassigned      up      up
FortyGigabitEthernet 1/2/2      unassigned      up      up
FortyGigabitEthernet 1/2/3      unassigned      up      up
FortyGigabitEthernet 1/2/4      unassigned      up      up
FortyGigabitEthernet 1/2/5      unassigned      up      up
FortyGigabitEthernet 1/2/6      unassigned      up      up
FortyGigabitEthernet 1/2/7      unassigned      up      up
FortyGigabitEthernet 1/2/8      unassigned      up      up
FortyGigabitEthernet 1/2/9      unassigned      up      up
FortyGigabitEthernet 1/2/10     unassigned      up      up
FortyGigabitEthernet 1/2/12     unassigned      up      up
TenGigabitEthernet 1/2/11:1     unassigned      up      down
TenGigabitEthernet 1/2/11:2     unassigned      up      down
TenGigabitEthernet 1/2/11:3     unassigned      up      down
TenGigabitEthernet 1/2/11:4     unassigned      up      down
```

For a standalone switch

Static configuration supports certain line cards whereby the SFP mode change must be done in the line-card power off state to facilitate the re-allocation of line-card-wide ASIC resources among the new set of interfaces.

For the Brocade VDX 6740 and 6740T, the SFP interfaces must be in the shutdown state before the configuration is allowed.

Perform the following configuration procedure for each SFP.

1. Shut down all interfaces that exist on this SFP.
 - a. To *disable* breakout, four Te interfaces are shut down on the Quad SFP (QSFP).
 - b. To *enable* breakout, a single Fo interface is shut down on the QSFP.
2. Apply the SFP breakout command to the target port(s).
 - a. Existing interface(s) under the old mode, along with the associated configurations, are removed from the DCM database. If the target is a QSFP, disabling breakout mode results in deleting four Te interfaces and enabling breakout results in deleting a single Fo interface. The configuration of nontargeted interfaces is not affected.
 - b. SFP breakout mode is updated in the DCM database.
3. After all SFPs are configured, copy the running-config to the startup-config. (This applies to standalone and fabric cluster modes.)
 - a. If you are in standalone mode or fabric cluster mode, you must save the running-config to the startup-config. The SFP configuration must be saved in the startup database, which is used after you reboot the switch.
 - b. If you are in logical chassis cluster mode, you do not need to copy the running-config to the startup-config.
4. Reboot the switch.
 - a. The Brocade VDX 6740 or 6740T creates interfaces corresponding to the SFP breakout mode of each port. For a QSFP, a single Fo interface is created in disable mode and four Te interfaces are created in enable mode.
 - b. The SFP interfaces under the new mode come up in default configurations as if the system were booting up for the first time. Unaffected interfaces retain their original configurations before the switch reboot is applied.

Breakout configuration examples

The following configuration examples show how to set a 40G QSFP into breakout mode, reserve a 40G QSFP port while in breakout mode, and release a 40G QSFP port while in breakout mode.

Setting a 40G QSFP port into breakout mode

The following example shows you how to set a 40G QSFP port into breakout mode and then manually reserving and releasing the 40G port DPOD reservation.

```
switch# config
Entering configuration mode terminal
switch(config)# interface FortyGigabitEthernet 48/0/49
switch(config-if-fo-48/0/49)# shut
switch(config-if-fo-48/0/49)# exit
switch(config)# hardware
switch(config-hardware)# connector 48/0/49
switch(config-connector-48/0/49)# sfp breakout
%Warning: Sfp Breakout is a disruptive command.
Please save the running-config to startup-config and a power-cycle for the
changes to take place.
```

```
switch(config-connector-48/0/49)# do copy running-config startup-config
This operation will modify your startup configuration. Do you want to
continue? [y/n]:y
switch(config-connector-48/0/49)# do reload
```

Warning: Unsaved configuration will be lost. Please run `copy running-config startup-config` to save the current configuration if not done already.

```
Are you sure you want to reload the switch? [y/n]:y
The system is going down for reload NOW !!
```

```
switch# show ip int br
```

Interface	IP-Address	Vrf	Status	Protocol
FortyGigabitEthernet 48/0/50	unassigned	default-vrf	up	down
FortyGigabitEthernet 48/0/51	unassigned	default-vrf	up	down
FortyGigabitEthernet 48/0/52	unassigned	default-vrf	up	down
TenGigabitEthernet 48/0/1	unassigned	default-vrf	up	up ISL)
TenGigabitEthernet 48/0/2	unassigned	default-vrf	up	down
TenGigabitEthernet 48/0/3	unassigned	default-vrf	up	down
TenGigabitEthernet 48/0/47	unassigned	default-vrf	up	down
TenGigabitEthernet 48/0/48	unassigned	default-vrf	up	down
TenGigabitEthernet 48/0/49:1	unassigned	default-vrf	up	down(ISL)
TenGigabitEthernet 48/0/49:2	unassigned	default-vrf	up	down(ISL)
TenGigabitEthernet 48/0/49:3	unassigned	default-vrf	up	down(ISL)
TenGigabitEthernet 48/0/49:4	unassigned	default-vrf	up	down(ISL)
Vlan 1	unassigned administratively		down	down
Vlan 4093	unassigned		up	down
Vlan 4095	unassigned administratively		down	down

```
switch#
```

Reserving a 40G QSFP port while in breakout mode

The following example shows you how to reserve a 40G QSFP port while in breakout mode

```
switch# config
Entering configuration mode terminal
switch(config)# dpod 48/0/
Possible completions:
  1  2  3  4  5  6  7  8  9  10 11 12 13 14 15 16 17 18 19
 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38
 39 40 41 42 43 44 45 46 47 48 49 50 51 52
switch(config)# dpod 48/0/49:1 reserve
Invalid InterfaceId.
switch(config)# dpod 48/0/49 reserve
switch(config-dpod-48/0/49)# do show dpod
rbridge-id: 48
  48 10G ports are available in this switch
   4 40G ports are available in this switch
  10G Port Upgrade license is installed
  40G Port Upgrade license is installed
  Dynamic POD method is in use

48 10G port assignments are provisioned for use in this switch:
   24 10G port assignments are provisioned by the base switch license
   24 10G port assignments are provisioned by the 10G Port Upgrade license
   2 10G ports are assigned to installed licenses:
```

```

    2 10G ports are assigned to the base switch license
    0 10G ports are assigned to the 10G Port Upgrade license
10G ports assigned to the base switch license:
    48/0/1, 48/0/31
10G ports assigned to the 10G Port Upgrade license:
    None
10G ports not assigned to a license:
    48/0/2, 48/0/3, 48/0/4, 48/0/5, 48/0/6, 48/0/7, 48/0/8, 48/0/9, 48/0/10,
48/0/11
    48/0/12, 48/0/13, 48/0/14, 48/0/15, 48/0/16, 48/0/17, 48/0/18, 48/0/19,
48/0/20, 48/0/21
    48/0/22, 48/0/23, 48/0/24, 48/0/25, 48/0/26, 48/0/27, 48/0/28, 48/0/29,
48/0/30, 48/0/32
    48/0/33, 48/0/34, 48/0/35, 48/0/36, 48/0/37, 48/0/38, 48/0/39, 48/0/40,
48/0/41, 48/0/42
    48/0/43, 48/0/44, 48/0/45, 48/0/46, 48/0/47, 48/0/48
46 license reservations are still available for use by unassigned ports

    4 40G port assignments are provisioned for use in this switch:
        0 40G port assignments are provisioned by the base switch license
        4 40G port assignments are provisioned by the 40G Port Upgrade license
    2 40G ports are assigned to installed licenses:
        0 40G ports are assigned to the base switch license
        2 40G ports are assigned to the 40G Port Upgrade license
40G ports assigned to the base switch license:
    None
40G ports assigned to the 40G Port Upgrade license:
    48/0/49, 48/0/50
40G ports not assigned to a license:
    48/0/51, 48/0/52
2 license reservations are still available for use by unassigned ports

```

Releasing a 40G QSFP port while in breakout mode

The following example shows you how to release a 40G QSFP port while in breakout mode

```

switch(config-dpod-48/0/49)# dpod 48/0/49 release
Port should be Offline to change POD assignment.
switch(config-dpod-48/0/49)# exit
switch(config)# interface TenGigabitEthernet 48/0/49:1
switch(conf-if-te-48/0/49:1)# shut
switch(conf-if-te-48/0/49:1)# interface TenGigabitEthernet 48/0/49:2
switch(conf-if-te-48/0/49:2)# shut
switch(conf-if-te-48/0/49:2)# interface TenGigabitEthernet 48/0/49:3
switch(conf-if-te-48/0/49:3)# shut
switch(conf-if-te-48/0/49:3)# interface TenGigabitEthernet 48/0/49:4
switch(conf-if-te-48/0/49:4)# shut
switch(conf-if-te-48/0/49:4)# exit
switch(config)# dpod 48/0/49 release
switch(config-dpod-48/0/49)#

```

Breakout limitations

In most circumstances, breakout interfaces behave the same as nonbreakout (normal) interfaces with regard to port attributes and states. Each breakout interface maintains its administrative state, operational state, and statistics. The exception is at the physical layer, whereby the hardware platform does not have per-breakout interface information.

- SFP media

In breakout mode, there is only SFP and no per-breakout media information. The **show media** command displays the same media information for all breakout interfaces.

NOTE

The TX Power Field in the **show media** command is not supported by the 40G optics.

- LED

For VDX 6740 and 6740T platforms, the LED state for a breakout interface is deterministic. For all other supported platforms, the LED state for a breakout interface *is not* deterministic.

In addition to the hardware limitations, note the following:

- Fabric Inter-Switch Links (ISLs) are not supported in breakout mode.

The following configuration and operational behavior is expected on the breakout interfaces.

- Breakout interface default ISL admin state is disabled. This is equivalent to invoking the **no fabric isl enable** command. Therefore, it will never come up as a Ve port. Invoking the **show ip interface brief** command always shows the protocol state as “down.”
- The command **fabric isl enable** is not supported on a breakout-enabled interface; however, it is still supported on the interface in nonbreakout mode (for example, the Fo interface on the QSFP).

High availability considerations

The active control processor (CP) needs to sync the SFP breakout state to the standby CP. After a failover, you should reconfirm the existence of interfaces, as well as note any possible configuration differences between what is expected for breakout mode and what is actually present.

Because changing the SFP mode requires associated interfaces to be disabled first, there is no data-traffic disruption during failover.

Network OS Layer 3 Routing Features

This section describes Layer 3 routing features of Network OS, and includes the following chapters:

- In-band Management 495
- IP Route Policy 505
- IP Route Management. 509
- Configuring OSPF..... 525
- Configuring VRRP..... 541
- Configuring VRF 555
- Configuring BGP..... 559
- Configuring SFP breakout mode..... 485
- Configuring Remote Monitoring 211
- Configuring IGMP..... 587

In-band Management

In this chapter

- [In-band management overview 495](#)
- [Configuring a standalone in-band management interface. 497](#)
- [Configuring an in-band management interface using OSPF 498](#)
- [Base configuration for a standalone in-band management. 499](#)

In-band management overview

In-band management on the Brocade VDX switches allows you to manage devices through Layer 3-enabled front-end Ethernet ports. An in-band management interface is relatively easy to configure and the most cost-effective management solution, because management traffic and data traffic use the same physical port (a design principle referred to as “fate-sharing”). Therefore, no special infrastructure is required to support management traffic. The downside is that any problem in the data network can potentially cause loss of connectivity, and thus loss of management function, to the managed devices. Therefore, it is highly recommended that you configure a dedicated serial connection for any device in your network as an out-of-band fallback solution in the event in-band management becomes unavailable.

In-band management facilitates management tasks such as downloading firmware, SNMP polling, SNMP traps, troubleshooting, and configuration when an out-of-band management interface is not available. [Table 82](#) lists some of the applications you can use with in-band management. The application listing is not meant to be exhaustive.

TABLE 82 Supported applications for in-band management

Application	Description
FWDL	Download firmware from an external server to a remote device using FTP or SCP.
SCP	Transfer files by using the Secure Copy Protocol.
SSH	Connect to a device through the Secure Shell application.
SNMP	Manage devices through the Simple Network Management Protocol.
Telnet	Connect to a device by using Telnet.

Prerequisites

The management station must be able to acquire an IP address and the routes to the management network. You can configure the management station to use a static IP address or to acquire an IP address dynamically or through protocols such as DHCP. A default gateway can be used to forward all the packets from the management station to the management network. Refer to [“Configuring the Ethernet management interface”](#) in [Chapter 3, “Basic Switch Management”](#)

In addition, you must configure IP routes and subnets. The front-end Ethernet port that you configure for management access acts as a router with IP forwarding implemented to allow communication with the target device. If the management station and the managed devices are in separate subnets, it is necessary to configure IP routes throughout the network to allow the communication to take place. You can configure the management interface to use either dynamic routing protocols, such as Open Shortest Path First (OSPF), or static routing.

- To configure the in-band management interface to use static routing, refer to [“Configuring static routes”](#) in [Chapter 36, “IP Route Management”](#).
- To configure the in-band management interface to use dynamic routing, refer to [Chapter 38, “Configuring OSPF”](#).

On switches running Network OS v 3.0.0 and later, in-band management is supported in VCS-enabled mode to manage devices through a Layer 2 or Layer 3 network. In standalone mode, a management station may be directly connected to another node in standalone mode. On switches running firmware prior to Network OS v3.0.0, in-band management is supported only in standalone mode.

NOTE

Standalone mode is not supported on the Brocade VDX 8770 switches.

In-band management does not require any special configuration commands. Because management traffic rides over the existing IP routing infrastructure, the commands needed to configure an in-band management interface are the same you would use to configure IP interfaces supported by static or dynamic routing protocols to provide connectivity to target devices.

Supported interfaces

In-band management is supported on the interfaces shown in [Table 83](#). Refer to the **interface** command documentation in the *Network OS Command Reference* for more information on the configuration options available for each of these interfaces.

TABLE 83 Ports configurable for in-band management

Interface	Addressing	Description
Management (Ma)	rbridge-id/slot	Management interface
GigabitEthernet (Gi)	rbridge-id/slot/port	1-GbE physical interface
TenGigabitEthernet (Te)	rbridge-id/slot/port	10-GbE physical interface
FortyGigabitEthernet (Fo)	rbridge-id/slot/port	40-GbE physical interface
Port-channel (Po)	interface-id (IP or Po in standalone mode only)	Port Channel interface
Virtual Ethernet (Ve)	interface-id (corresponding VLAN ID)	Virtual Ethernet Interface

NOTE

A virtual Ethernet (Ve) interface is a logical port associated with a Layer 3 Virtual LAN (VLAN) configured on a Layer 3 switch. You can configure routing parameters on the virtual interface to enable the Layer 3 switch to route protocol traffic from one Layer 3 VLAN to the other, without using an external router. A corresponding VLAN must be configured before you can configure the Ve interface.

Configuring a standalone in-band management interface

Figure 43 shows the configuration of an in-band management interface in standalone mode. In this example, the management station IP address and Ethernet port interface IP addresses for Switch-A and Switch-B are all in the same subnet, and therefore, no routing protocols are needed for the management station to connect to Switch-B through Switch-A. The management station (a server or workstation) connects to the physically attached switch-A, and switch-A can connect to the physically connected Switch-B.

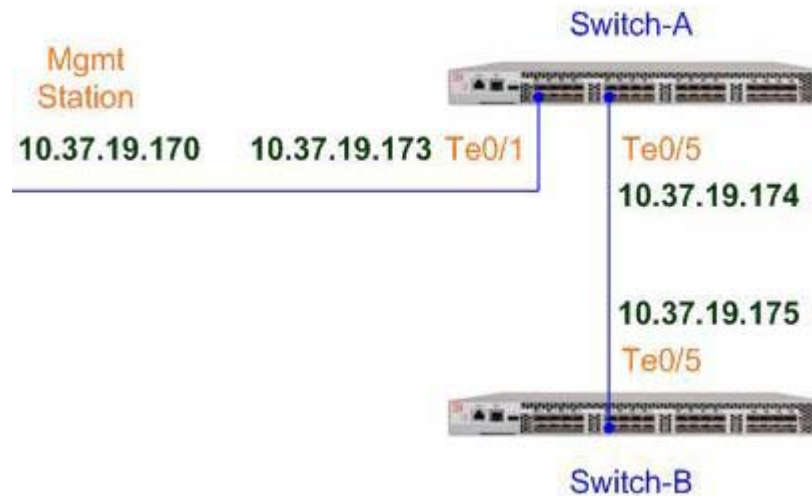


FIGURE 43 A management station with a networked device in standalone mode.

The configuration shown in Figure 43 supports the following operations:

- Connecting from the management station to Switch-A through an SSH or Telnet session.
- Transferring files between the management station and Switch-A using secure copy (SCP) or FTP.
- Transferring files between Switch-A and Switch-B using secure copy (SCP).
- Using any of the applications in Table 82 between Switch-A and Switch-B.

Provisioning in standalone mode

The following procedure configures the in-band management interface shown in Figure 43.

1. Connect to the switch through the serial console or the management interface if available.
2. Issue the **configure terminal** command to enter global configuration mode.
3. Enter the **interface** command followed by the interface type you want to configure.

For a standalone in-band management interface, only a physical user port (1 GbE, 10 GbE, or 40 GbE) needs to be configured with IP addresses. There is no need to configure either a VLAN or a Ve interface.

4. Enter the **ip address** *IPv4_address/prefix_length* command to set the IPv4 address for the interface.

NOTE

You must configure a primary IP address only. Secondary IP addresses are not supported.

5. Enter the **ip mtu** command to set the interface IP Maximum Transmission Unit (MTU) in bytes.
6. Enter the **arp-ageing-timeout** command to configure the interface timeout parameter (in minutes) for the Address Resolution Protocol (ARP).

The default timeout value is 4 hours.

7. Clear the ARP cache by using the **do clear-arp-cache** command with the **no-refresh** option to delete unused ARP entries.
8. Configure a proxy ARP per interface by using the **ip proxy-arp** command.
9. Display the configuration by using the **show ip interface** command.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface TenGigabitEthernet 1/0/1
switch(config-if-te-1/0/1)# no shutdown
switch(config-if-te-1/0/1)# ip address 1.1.1.1/24
switch(config-if-te-1/0/1)# ip mtu 1200
switch(config-if-te-1/0/1)# arp-ageing-timeout 300
switch(config-if-te-1/0/1)# do clear-arp-cache no-refresh
switch(config-if-te-1/0/1)# ip proxy-arp
switch(config-if-te-1/0/1)# exit

switch# show ip interface TenGigabitEthernet 1/0/1
TenGigabitEthernet 10/1 is up protocol is up
Primary Internet Address is 1.1.1.1/24 broadcast is 1.1.1.255
IP MTU is 1200
Proxy Arp is Enabled
ICMP unreachable are always sent
ICMP mask replies are never sent
IP fast switching is enabled
```

Configuring an in-band management interface using OSPF

Figure 44 shows the configuration of an in-band management interface connected to a VCS fabric.

- In this scenario, a Brocade VDX 6720 (RB1 local), is connected to a modular Brocade VDX 8770 switch (RB2) in a VCS fabric through fabric ports.
- Another Brocade VDX 6720 switch (C1) is connected to a VCS fabric member (RB1) through the front panel port in RB1.
- OSPF is configured between RB1 and RB2 for Layer 3 forwarding, and to export and exchange routes.
- A management station (MS) is connected to C1 through a Telnet or SSH session.
- In a standalone configuration (“[Base configuration for a standalone in-band management](#)”) C1 and RB1 are in standalone mode. The management station connects to RB1 through C1 using physical front-panel port connections.

- In VCS mode (“[Base configuration in VCS fabric cluster mode](#)” on page 500), the management station connects to RB2 through C1 using OSPF and dynamic route distribution in a VCS fabric.




FIGURE 44 In-band management in a VCS fabric with dynamic routes (OSPF)

Base configuration for a standalone in-band management

The following configuration establishes an in-band management connection from the management station to RB1 through C1. For the purpose of this example, C1 and RB1 operate in standalone mode (VCS mode is disabled).

1. Configure the front-end Ethernet port on RB1 through a serial connection.
 - a. Connect to RB1 using the serial console.
 - b. Issue the **configure terminal** command to enter global configuration mode.
 - c. Configure a front-end Ethernet port by using the **interface vlan vlan_id** command.
 - d. Issue the **rbridge-id rbridge-id** command to enter RBridge ID configuration mode.
 - e. Enter the **interface ve ve_id** command to configure the virtual Ethernet interface (Ve).
The Ve ID must correspond to the existing VLAN ID.
 - f. Enter the IP address of the interface.
 - g. Enter the **no shutdown** command to enable the interface.
 - h. Enter the **do show vcs** command to verify that RB1 is in standalone mode and not a part of a VCS fabric.

```
RB1# configure terminal
Entering configuration mode terminal.
RB1(config)# interface vlan 2
RB1(config)# rbridge-id 1
```

34 Base configuration for a standalone in-band management

```
RB1(config-rbridge-id-1)# interface ve 2
RB1(config-Ve-2)# ip address 2.2.2.17/24
RB1(config-Ve-2)# no shutdown
RB1(config--Ve2)# exit
RB1(config)# do show vcs
```

```
state      : Disabled
```

2. C1 is a management station and automatically Telnets into node RB1.
3. Verify the in-band management connection between the management station to C1, and between C1 and RB1 (standalone test).

- a. Connect to C1 through the management interface by using an SSH session.
- b. On C1, establish a Telnet connection from C1 to RB1.

```
C1# telnet 2.2.2.17/24
Trying 2.2.2.17...
Connected to 2.2.17.24.
Escape character is '^]'.

```

You are now logged in to RB1 (note the prompt change).

- c. Verify the Telnet notification on RB1.

```
RB1# 1970/01/01-02:16:27, [SEC-1203], 13406, M1, INFO, RB1, Login
information: Login successful via TELNET/SSH/RSN. IP Addr: 2.2.17.24

```

- d. Through the Telnet in-band management connection from C1, verify that RB1 is in standalone mode.

```
RB1# show vcs
```

```
state      : Disabled
```

You can now perform in-band management functions on RB1 through the management interface SSH connection to C1, such as downloading firmware or managing SNMP.

Base configuration in VCS fabric cluster mode

The following configuration establishes an in-band management connection from a management station (MS) to C1, and from C1 to RB1 and RB2. RB1 is the local switch, RB2 is the remote switch. Both switches operate in VCS-enabled mode.

1. Set up an OSPF network (area 0) on RB1 to enable dynamic routing by using the following procedure.

- a. Connect to RB1 by using a serial connection.
- b. Issue the **configure terminal** command to enter global configuration mode.
- c. Enter the **interface vlan** command to configure a VLAN on RB1.

```
RB1# configure terminal
Entering configuration mode terminal.
RB1(config)# interface vlan 100

```

- d. Issue the **rbridge-id rbridge-id** command to enter RBridge ID configuration mode.
- e. Configure OSPF by using the **router ospf** command followed by the **area** command. Enter **0** to configure OSPF area 0.
- f. Enter the **exit** command to return to RBridge ID configuration mode.

```
RB1(config-Vlan-100)# rbridge-id 17
RB1(config-rbridge-id-17)# router ospf
RB1(config-ospf-router)# area 0
RB1(config-ospf-router)# exit
RB1(config-rbridge-id-17)#
```

- g. Configure a virtual Ethernet interface for the VLAN and assign the IP address for RB1.
- h. Set OSPF area 0 with the **ip ospf area** command.
- i. Enter the **no shutdown** command to enable the interface.

```
RB1(config-rbridge-id-17)# interface ve 100
RB1(config-Ve-100)# ip address 100.100.100.17/24
RB1(config-Ve-100)# ip ospf area 0
RB1(config-Ve-100)# no shutdown
```

2. Set up an OSPF network (area 0) on RB2 to enable dynamic routing. The configuration steps mirror the configuration on RB1.

NOTE

If you are configuring this in a logical chassis cluster mode, you do not configure the VLAN again on RB2 because RB1 (the principal node) would distribute the configuration to all nodes in the logical chassis cluster.

- a. Connect to RB2 by using a serial connection.
- b. Issue the **configure terminal** command to enter global configuration mode.
- c. Enter the **interface vlan** command to configure a VLAN on RB2.
- d. Issue the **rbridge-id rbridge-id** command to enter RBridge ID configuration mode.
- e. Configure OSPF by using the **router ospf** command followed by the **area** command. Enter **0** to configure OSPF area 0.
- f. Enter the **exit** command to return to RBridge ID configuration mode.

```
RB2# configure terminal
Entering configuration mode terminal.
RB2(config)# interface vlan 100
RB2(config-rbridge-id-18)# rbridge-id 18
RB2(config-rbridge-id-18)# router ospf
RB2(config-ospf-router)# area 0
RB2(config-ospf-router)# exit
RB2(config-rbridge-id-18)#
```

- g. Configure the virtual Ethernet interface for the VLAN and assign the IP address for RB2.
- h. Set OSPF area 0 with the **ip ospf area** command.
- i. Enter the **no shutdown** command to enable the interface.

```
RB2(config-rbridge-id-18)# interface ve 100
RB2(config-Ve-100)# ip address 100.100.100.18/24
RB2(config-Ve-100)# ip ospf area 0
RB2(config-Ve-100)# no shutdown
```

3. Verify the OSPF configuration in the VCS fabric.

- a. On RB2, verify adjacency between RB1 and RB2. RB2 is displayed as a neighbor.

```
RB2(config-Ve-100)# do show ip ospf neighbor
```

```
Port Address Pri State Neigh Address Neigh ID Ev Opt Cnt
Ve 100 100.100.100.18 1 FULL/BDR 100.100.100.17 100.100.100.17 6 2 0
```

- b. On RB2, verify the IP routes between RB1 and RB2.

```
RB2(config-Ve-100)# do show ip route
```

```
Total number of IP routes: 2
Type Codes-B:BGP D:Connected I:ISIS O:OSPF R:RIP S:Static;
Cost-Dist/Metric
BGP Codes - i:iBGP e:eBGP
ISIS Codes - L1:Level-1 L2:Level-2
OSPF Codes - i:Inter Area 1:External Type 1 2:External Type 2 s:Sham Link
Destination Gateway Port Cost Type Uptime
1 2.2.2.0/24 100.100.100.17 Ve 100 110/10 O2 0m14s
2 100.100.100.0/24 DIRECT Ve 100 0/0 D 5m2s
```

- c. On RB1, **distribute** the routes between RB1 and RB2 as shown in the example.

```
RB1# configure terminal
```

```
Entering configuration mode terminal.
```

```
RB1(config)# rbridge-id 17
```

```
RB1(config-rbridge-id-17)# router ospf
```

```
RB1(conf-ospf-router)# redistribute connected
```

```
RB1(conf-ospf-router)# exit
```

```
RB1(config)#
```

- d. Verify the VCS fabric configuration on RB1.

```
RB1(config)# do show vcs
```

```
Config Mode : Local-Only
```

```
VCS ID : 2
```

```
Total Number of Nodes : 2
```

```
Rbridge-Id WWN Management IP Status HostName
```

```
-----
17 10:00:00:05:33:77:31:9C* 10.24.73.80 Online RB1
18 >10:00:00:05:33:77:23:6C 10.24.73.85 Online RB2
```

- e. Verify the VCS fabric configuration on RB2.

```
RB2# show vcs
```

```
Config Mode : Local-Only
```

```
VCS ID : 2
```

```
Total Number of Nodes : 2
```

```
Rbridge-Id WWN Management IP Status HostName
```

```
-----
17 10:00:00:05:33:77:31:9C* 10.24.73.80 Online RB1
18 >10:00:00:05:33:77:23:6C 10.24.73.85 Online RB2
```

4. Verify the in-band management connection to RB1 and RB2.

- a. Connect to C1 from the management station using an SSH connection.

- b. On C1, verify connectivity to RB1 (local test) by issuing the **ping** command to the front-end Ethernet interface (Ve port) IP address for RB1.

```
C1# ping 2.2.2.17
```

```
PING 2.2.2.17 (2.2.2.17): 56 data bytes
```

```
64 bytes from 2.2.2.17: icmp_seq=0 ttl=64 time=8.206 ms
```

```
64 bytes from 2.2.2.17: icmp_seq=1 ttl=64 time=7.126 ms
--- 2.2.2.17 ping statistics ---
```

- c. From C1, verify connectivity to RB2 (remote test) by issuing the **ping** command to the front-end Ethernet interface (VE port) IP address for RB2.

```
C1# ping 100.100.100.18
PING 100.100.100.18 (100.100.100.18): 56 data bytes
64 bytes from 100.100.100.18: icmp_seq=0 ttl=63 time=21.239 ms
64 bytes from 100.100.100.18: icmp_seq=1 ttl=63 time=19.889 ms
--- 100.100.100.18 ping statistics ---
```

- d. From C1, establish a Telnet connection to RB1 by issuing the **telnet** command to the front-end Ethernet interface (Ve port) IP address for RB1.

This test verifies the in-band management interface between C1 and the local RBridge, RB1. If the test succeeds, you can perform management functions on RB1 through C1.

```
C1# telnet 2.2.2.17
Trying 2.2.2.17...
Connected to 2.2.2.17
Escape character is '^['.
```

- e. Verify the Telnet RASLog notification on RB1 through C1.

The RASLog message displays on the console.

```
RB1# 2012/05/10-17:31:09, [SEC-1203], 312942, M1, INFO, C1, Login
information: Login successful via TELNET/SSH/RSH. IP Addr: 2.2.2.1
```

- f. Verify the VCS fabric configuration on RB1 through C1.

The output is identical to the output generated in step 3d. The difference is that you are now connected to RB1 through the in-band management interface.

```
RB1# show vcs
Config Mode   : Local-Only
VCS ID       : 2
Total Number of Nodes   : 2
Rbridge-Id   WWN          Management IP Status HostName
-----
17          10:00:00:05:33:77:31:9C* 10.24.73.80 Online RB1
18          >10:00:00:05:33:77:23:6C 10.24.73.85 Online RB2
```

- g. From C1, establish a Telnet connection to RB2 by issuing the **telnet** command to the front-end Ethernet interface (VE port) IP address for RB2.

This test verifies the in-band management interface between C1 and the remote RBridge, RB2. If the test succeeds, you can perform management functions on RB2 through C1.

```
C1# telnet 100.100.100.18
Trying 100.100.100.18...
Connected to 100.100.100.18.
Escape character is '^['.
```

- h. Verify the Telnet RASLog notification on RB2 through C1.

The RASLog message displays on the console.

```
RB2# 2012/05/10-17:31:09, [SEC-1203], 312942, M1, INFO, C1, Login
information: Login successful via TELNET/SSH/RSH. IP Addr: 2.2.2.1
```

- i. Verify the VCS fabric configuration on RB2 through C1.

34 Base configuration for a standalone in-band management

The output is identical to the output generated in step 3d. The difference is that you are now connected to RB2 through the in-band management interface.

```
RB2# show vcs
Config Mode   : Local-Only
VCS ID        : 2
Total Number of Nodes   : 2
Rbridge-Id  WWN                Management IP Status HostName
-----
17          10:00:00:05:33:77:31:9C* 10.24.73.80 Online  RB1
18          >10:00:00:05:33:77:23:6C 10.24.73.85 Online  RB2
```

If all verification steps produce the desired results, the configuration is successful, and you can use the management interface on C1 to perform management functions on both the local (RB1) and the remote (RB2) switch.

For more information on IP address configuration in the VCS fabric, refer to [“VCS Virtual IP address configuration”](#) on page 131.

IP Route Policy

In this chapter

- About IP route policy 505
- Configuring IP route policy 506

About IP route policy

IP route policy controls how routes or IP subnets are transported from one subsystem to another subsystem. The IP route policy may perform “permit” or “deny” actions so that matched routes may be allowed or denied to the target subsystem accordingly. Additionally, IP route policy may also be used for modify the characteristics of a matched route and IP subnet pair.

There are two types of IP route policies supported; prefix-list and route-map.

IP prefix-list

An IP prefix-list is identified by its name. Each IP prefix-list may consist of one or more instances. The following is an example of IP prefix-list,

```
switch# ip prefix-list test 1 deny 1.2.0.0/16 ge 17 le 30
switch# ip prefix-list test 2 permit 1.1.0.0/16
```

A matching condition of prefix-list instance contains two portions 1) IP subnet prefix and 2) optionally prefix (mask) length, where *ge* (greater or equal) is the lower limit of the mask length, and *le* (less or equal) is the upper limit of the mask length. If no *ge* and/or *le* is given in an instance, the exact match of subnet prefix length is needed.

In the example above, a route is considered match for instance 1 if this route is inside subnet 1.2.0.0/16 AND whose mask length is between 17 and 30. That is, route 1.2.1.0/24 matches, but route 1.2.1.1/32 does not due to mask length.

Similar to route-map, when finding match, each prefix-list instance is looked at in order specified by its instance ID. The look-up terminates at the first match. A route that does not find match in prefix list is denied.

At present, prefix-list is not used by itself. The IP prefix-list can be used as part of route-map match clauses. In this context, “permit” stands of matching this pattern, and “deny” stands for not matching this route pattern.

Route-map

A route-map is identified by its name. Each route-map may consist of one or more instances. Each route-map instances may contain zero or more matching clauses, and zero or more set clauses.

At present, a route-map instance is largest configuration granularity. That is, end-user is required to add AND delete route-maps via its instance. For example, when removing a route-map, an end-user is required to remove this route-map by all its instances. A route-map instance may contain more than one match conditions. The overall matching condition of the instance is true only if all matching conditions are met. The following is an example of route-map:

```
switch# Route-map test deny 1
      Match interface te 0/1
switch# Route-map test permit 2
      Match ip next-hop prefix-list pre-test
      Set tag 5000
```

In the example above, route-map test comprises of two instances; instance 1 denies entry for any routes whose next-hop interface is te 0/1 and instance 2 allows entry for routes whose next-hop match the IP subnets specified in the prefix-list pre-test (not shown). Additionally, each matched route has its tag set to 5000.

NOTE

The maximum number of OSPF networks that can be advertised and processed in a single area in a router is limited to 600.

A route-map instance does not need to contain a matching condition. It implies that the matching condition for this instance is true.

A route-map instance may contain more than one set clause. All set clauses are applied to the match routes when applicable.

When a route-map is applied, each instance is looked at in the order specified by the instance ID. If there is a match, the instance's action are applied, and its set clauses are applied if the action is permitted. The search terminates at the first match. A route that does not find match in a route-map is denied.

Configuring IP route policy

Similar to ACLs, route-map and IP prefix need to be applied for their specified policy to take effect. The following example applies a route-map to the redistribution of static routes to into an OSPF domain. For complete information on these commands, refer to the *Network OS Command Reference*.

To set an IP route policy, perform the following steps in Privileged EXEC mode.

1. Enter the **router ospf** command to enable the OSPF protocol.

```
switch# router ospf
      redistribution static route-map test
      area 0
```

2. Enter the **ip route** command to create the prefix for a static route.

```
switch# ip route 11.11.11.0/24 2.2.2.1
```

3. Enter the **ip route** command to create the next hop in the static route. Repeat as needed.

```
switch# ip route 11.11.11.0/24 2.2.2.2
```

4. Enter the **route-map** command to create the route map instance.

```
switch# route-map test permit 1  
      match ip address prefix-list pretest
```

5. Enter the **ip prefix-list** command to configure the IP prefix instance.

```
switch# ip prefix-list pretest 2 permit 1.1.1.0/24
```

In the example above, when the **route-map test permit 1** command executes, only the static route **1.1.1.0/24** is exported into the OSPF domain because there are no matching rules in the ip prefix-list pretest for route **11.11.11.0/24**. The default action of prefix list is deny (no match), therefore the route **11.11.11.0/24** is not exported into OSPF domain.

You can configure the router to explicitly permit or deny specific IP addresses. The router permits all IP addresses by default. If you want permit to remain the default behavior, define individual filters to deny specific IP addresses. If you want to change the default behavior to deny, define individual filters to permit specific IP addresses. Once you define a filter, the default action for addresses that do not match a filter is “deny”. To change the default action to “permit”, configure the last filter as “permit any any”.

IP Route Management

In this chapter

- Overview of IP route management 509
- How IP route management determines best route 509
- Configuring static routes 510
- Other routing commands 511

Overview of IP route management

IP route management is the term used to refer to software that manages routes and next hops from different sources in a routing table, from which the Brocade device selects the best routes for forwarding IP packets. This route management software gets activated automatically at system bootup and does not require preconfiguration.

IP route management runs on all platforms configured for Layer 3, and provides the following:

- Maintains routes submitted by other protocols.
- Supports route redistribution.
- Supports router identification.
- Selects and synchronizes routes to the forwarding information base (FIB).
- Synchronizes the Layer 3 interface to the FIB.
- Supports the following Layer 3 interfaces: virtual ethernet (Ve), router port, loopback, and management.

NOTE

IP route management supports both IPv4 and IPv6 routes.

How IP route management determines best route

The sources of routes that are added into IP route management are:

- Dynamic routes from routing protocols. Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) are both supported.
- Static configured routes: You can add routes directly to the route table. When you add a route to the IP route table, you are creating a static IP route.
- Directly connected routes from interface configuration: When you add an IP interface, the Brocade device automatically creates a route for the network.

Administrative distance can be configured for route types other than connected routes. IP route management prefers routes with lower administrative distances.

Configuring static routes

You can add a static route to IP route management by using the **ip route** commands in RBridge ID configuration mode. With these commands, you can specify either the next-hop gateway or the egress interface for the route.

Specifying the next-hop gateway

To configure a static route to 10.95.7.0, using 10.95.6.157 as the next-hop gateway, use the **ip route** command in RBridge ID configuration mode, as shown in this example:

```
switch (config)# rbridge-id 30
switch (config-rbridge-id-30)# ip route 10.95.7.0/24 10.95.6.157
```

Specifying the egress interface

To configure a static IP route with an IPv4 address on a 10-gigabit Ethernet port, enter an **ip route** command such as the following.

```
switch (config)# rbridge-id 30
switch (config-rbridge-id-30)# ip route 192.128.2.0/24 te 101/4/1
```

The command configures a static IP route for destination network 192.128.2.0/24. Because an Ethernet port is specified instead of a gateway IP address as the next hop, the Brocade device forwards traffic for the 192.128.2.0/24 network to the 10-gigabit Ethernet port 101/4/1.

This example is the same command using IPv6.

```
switch (config)# rbridge-id 30
switch (config-rbridge-id-30)# ipv6 route fe80::21b:edff:fe0b:3c00/64 te 101/4/1
```

Configuring the default route

A default route is configured with an all-zero prefix/netmask (for example, 0.0.0.0/0). The default route is an example of a special static route with a destination prefix of zero. All traffic that does not have other matching routes is forwarded using the default route.

Once the maximum number of routes are installed in the IP route table and if you delete some of those routes, the **clear ip route all** command needs to be executed for the routes to be refreshed, so that previously uninstalled routes can be re-installed up to the maximum limit.

The same gateway for the default route is updated on the management interface, on both MMs.

The first configured management port default route entry becomes the gateway on the management interface. Any other default routes configured later do not change the management interface gateway value. When multiple such entries exist, and you remove them one by one, the last entry is updated as the gateway.

To configure a default route with a next hop of 10.95.6.157, enter the following **ip route** command.

```
switch(config)# rbridge-id 30
switch(config-rbridge-id-30)# ip route 0.0.0.0/0 10.95.6.157
```

Other routing commands

Refer to *Network OS Command Reference* for more information about all IP routing-related commands. For example:

- The **ip route** command offers an option that allow you to specify a tag value of a route for route filtering with a route map. The command also offers an option for specifying a cost metric.
- The **ip load-sharing** command can be used to balance IP traffic across up to eight equal paths.
- The **ip route next-hop ospf** command allows a Brocade device to use routes learned from OSPF to resolve a configured static route.
- The **ip route next-hop bgp** command allows a Brocade device to use routes learned from BGP to resolve a configured static route.
- The **ip route next-hop-recursion** command allows a Brocade device to resolve a route by using as many as 10 recursive-level lookups of other routes.

Configuring PIM

In this chapter

• PIM overview	513
• PIM Sparse Mode	514
• Topologies	515
• PIM Sparse device types	517
• Prerequisites	518
• Standards Conformity	518
• Limitations	518
• Supportability	519
• Configuration Example	520

PIM overview

The Protocol Independent Multicast (PIM) protocol is a family of IP multicast protocols. PIM does not rely on any particular routing protocol for creating its network topology state. Instead, PIM uses routing information supplied by other traditional routing protocols such as the Routing Information Protocol, Open Shortest Path First, Border Gateway Protocol, and Multicast Source Discovery Protocol.

PIM messages are sent encapsulated in an IP packet with the IP protocol field set to 103. Depending on the type of message, the packet is either sent to the PIM All-Router-Multicast address (224.0.0.13) or sent as unicast to a specific host.

As with IP multicast, the main use case of PIM is for the source to be able to send the same information to multiple receivers by using a single stream of traffic. This helps minimize the processing load on the source as it needs to maintain only one session irrespective of the number of actual receivers. It also minimizes the load on the IP network since the packets are sent only on links which lead to an interested receiver.

Several types of PIM exist, but in this release Brocade supports only PIM Sparse Mode (PIM-SM). PIM-SM explicitly builds unidirectional shared trees rooted at a rendezvous point (RP) per group, and optionally creates shortest-path trees per source.

Important Notes

- PIM can be enabled on the Brocade VDX 6740, VDX 6740T, and the VDX 8770 only.
- PIM-SM can be used in VCS mode only.

PIM Sparse Mode

PIM-SM (Sparse Mode) is the most commonly deployed flavor of PIM. PIM-SM is most effective in large networks sparsely populated with hosts interested in multicast traffic. It is assumed that most hosts within a network are not interested in all multicast data streams.

PIM Sparse devices are organized into domains. A PIM Sparse domain is a contiguous set of devices that all implement PIM and are configured to operate within a common boundary.

PIM-SM creates unidirectional shared trees which are rooted at a common node in the network called the rendezvous point (RP). The RP acts as the messenger between the source and the interested hosts or routers.

There are various ways of identifying an RP within a network. It can either be statically configured per PIM router or configured using Bootstrap Router (BSR). Within a network, the RP should always be upstream compared to the destination hosts.

Once the RP has been identified, each interested host and/or router sends join messages to the RP for the group that they are interested in. To reduce incoming join messages to a RP, the local network selects one of its upstream routers as the designated router (DR). All hosts below a DR send IGMP join messages to the DR. The DR sends only one join message to the RP on behalf of all its interested hosts.

PIM-SM also provides the option of creating a source-based tree rooted at a router adjacent to the tree. This provides the destination hosts with an option of switching from the shared tree to the source-based tree if this is a shorter path between the source and the destination.

Topologies

This section shows diagrams of two supported PIM topologies.

Figure 45 shows the components for a single-VCS PIM topology.




FIGURE 45 Single VCS deployment

The following requirements apply to the single-VCS deployment depicted in Figure 45:

- Top of rack switches can be Brocade VDX 6710, VDX 6720, VDX 6730, VDX 6740, VDX 6740T or VDX 8770 models. However, top of rack switches are typically only Layer 2-capable when used in this context as part of a PIM environment, and PIM can be enabled on the Brocade VDX 8770 and VDX 6740 only.
- Top of rack switches must have IGMP-snooping enabled.
- Aggregation-layer switches must be Brocade VDX 8770 or VDX 6740 models only.

- Aggregation-layer switches can be PIM-enabled.
- L3 (VRRP-E and OSPF) can be configured on all interfaces with L3 connectivity to the data-center core.
- IGMP snooping must be enabled on the aggregation-layer switches.
- PIM DR-priority is configured on ve interfaces of all PIM-capable aggregation routers to optimize load-sharing abilities within the aggregation.

Figure 46 shows the components for a two-tier VCS PIM topology.




FIGURE 46 Two-tier VCS deployment

The following requirements apply to the two-tier-VCS deployment depicted in Figure 46:

- Top of rack switches can be Brocade VDX 6710, VDX 6720, VDX 6730, or VDX 8770 models. However, Top of rack switches are typically only L2-capable when used in this context as part of a PIM environment, and PIM can be enabled on the Brocade VDX 8770 or VDX 6740 only.
- Top of rack VCS are typically only L2 capable.
- Top of rack switches must have IGMP-snooping enabled.
- Aggregation-layer VCS must be VDX 8770 or VDX 6740 models only.
- Aggregation-layer switches can be PIM-enabled.
- L3 (VRRP-E and OSPF) can be configured on all interfaces with L3 connectivity to the data-center core.
- IGMP snooping must be enabled on the aggregation-layer switches.
- PIM can be enabled on all Brocade VDX 8770 or VDX 6740 switches where VRRP-E is enabled.
- PIM DR-priority is configured on ve interfaces of all PIM-capable aggregation routers to optimize load-sharing abilities within the aggregation.

PIM Sparse device types

Devices that are configured with PIM Sparse interfaces also can be configured to fill one or more of the following roles:

- PIM multicast border router (PMBR)—A PIM device that has interfaces within the PIM domain and other interface outside the PIM domain. PBMRs connect the PIM domain to the Internet.
- Bootstrap router (BSR)—A router that distributes rendezvous point (RP) information to the other PIM Sparse devices within the domain. Each PIM Sparse domain has one active BSR. For redundancy, you can configure ports on multiple devices as candidate BSRs. The PIM Sparse protocol uses an election process to select one of the candidate BSRs as the BSR for the domain. The BSR with the highest BSR priority (a user-configurable parameter) is elected. If the priorities result in a tie, then the candidate BSR interface with the highest IP address is elected.

The BSR must be configured as part of the L3 core network.

- Rendezvous point (RP)—The meeting point for PIM Sparse sources and receivers. A PIM Sparse domain can have multiple RPs, but each PIM Sparse multicast group address can have only one active RP. PIM Sparse devices learn the addresses of RPs and the groups for which they are responsible from messages that the BSR sends to each of the PIM Sparse devices.

The RP must be configured as part of the L3 core network.

NOTE

Brocade recommends that you configure the same ports as candidate BSRs and RPs.

- PIM designated router (DR)—Once the RP has been identified, each interested host and/or router sends join messages to the RP for the group that they are interested in. The local network selects one of its upstream routers as the designated router (DR). All hosts below a DR send IGMP join messages to the DR. The DR in turn sends only one join message to the RP on

behalf of all its interested hosts. The RP receives the first few packets of the multicast stream, encapsulated in the PIM register message, from the source hosts. These messages are sent as a unicast to the RP. The RP de-encapsulates these packets and forwards them to the respective DRs.

NOTE

DR election is based first on the router with the highest configured DR priority for an interface (if DR priority has been configured), and based next on the router with the highest IP address. To configure DR priority, use the **ip pim dr-priority** command. For more information about this command, refer to *Fabric OS Command Reference, Supporting Fabric OS v7.1.0* (or later).

Prerequisites

- PIM requires the following to function properly:
- The system should support receiving and transmitting unicast as well as multicast packets.
 - A Routing Information Base (RIB) must be accessible for obtaining routing information.
 - An IPC mechanism must be available.
 - A timer mechanism must be available.
 - An IGMP module should be available for correct operation of PIM when working as a DR.

Standards Conformity

Table 84 lists the level of Brocade conformity for various PIM-related RFCs.

TABLE 84 RFCs supported

Standard	Level (Y/N/Partial)	Notes
RFC 4601	Y	PIM-SM Protocol Specification
RFC 3973	N	PIM-DM Protocol Specification
RFC 5059	Partial	BSR mechanism for PIM supported
RFC 5060	Partial	PIM MIB supported
RFC 5240	N	BSR MIB
RFC 4610	N	Anycast-RP
RFC 3618	N	MSDP

Limitations

In this release, PIM can be enabled on the Brocade VX 6740, VDX 6740T, and VDX 8770 only. Also, only PIM Sparse Mode (PIM-SM) is supported at this time.

Static RP is not supported within the VCS cluster.

All PIM-enabled aggregation layer devices should have a direct Layer 3 connection to RP.

The following PIM features are not supported in this release:

- Non-stop routing (NSR)
- IP version 6
- VRF
- Prefix list
- Configuring the switch as the BSR candidate. However, the switch will be able to receive and process the BSR messages from other routers.
- Configuring the switch as the RP candidate.

Supportability

This release of Fabric OS includes the following PIM support:

- 32 virtual interfaces. The virtual interfaces can be either L3 VLAN or router ports
- 32 output interfaces
- 4,000 L3 multicast group IDs
- 2,000 (S,G) forwarding entries
- 256 (*, G) forwarding entries
- A learning rate of 32 routes per second.

Configuration Example

This section shows you an example PIM Sparse deployment and configuration, based on [Figure 47](#).




FIGURE 47 Example deployment using single VCS

Notes

- VLAGs must belong to PIM-enabled VLANs. For more information, refer to the “Configuring Link Aggregation” chapter. Set up your VLAGs before performing any PIM-specific configuration.
- Make sure the rendezvous point (RP) is configured. This should be a third-party box for dynamic RP functionality, or either the Brocade VDX 8770 or the VDX 6740 for static RP functionality.
- Make sure the bootstrap router (BSR), if applicable to your setup, is configured. The BSR can be any third-party box that supports PIM, BSR and rendezvous point (RP) functionality. If you are using a Brocade MLX switch as the bootstrap router, refer to the *Brocade MLX Series and NetIron Family Configuration Guide* for more information.

NOTE

If you are statically configuring the RP per PIM router, use the **router pim** and **rp-addr** commands, as described in the *Network OS Command Reference*.

- In the example shown in [Figure 47](#):
 - M1 and M2 must be Brocade VDX 8770 or VDX 6740 switches.
 - M1 is the designated router (DR) for virtual LAN 10 (labeled v10) and virtual LAN 30 (labeled v30).
 - M2 is the designated router (DR) for virtual LAN 20 (labeled v20) and virtual LAN 40 (labeled v40).
 - The switches labeled C-24 and C-60 can be any combination of Brocade VDX 6710, VDX 6720, VDX 6730, VDX 6740, or VDX 8770 models. These switches are pure L2 devices and need IGMP snooping enabled only.

The following steps show you how to configure PIM Sparse for the scenario depicted in [Figure 47](#). These steps show you where to enable IGMP snooping, where to create IP addresses for Vx interfaces, and where to enable PIM Sparse:

1. Enable IGMP snooping on each access-level switch by performing the following steps on each of these switches:
 - a. From the switch console, in privileged EXEC mode, enter global configuration mode by issuing the **configure** command:


```
switch# configure
```
 - b. Enter VLAN interface configuration mode for the first VLAN:


```
switch (config)# int vlan 10
```
 - c. Enable IGMP snooping by issuing the following command:


```
switch(config-Vlan-10)# ip igmp snooping enable
```
 - d. Exit interface configuration mode:


```
switch(config-Vlan-10)# exit
```
 - e. Enter VLAN interface configuration mode for the second VLAN:


```
switch (config)# int vlan 20
```
 - f. Enable IGMP snooping by issuing the following command:


```
switch(config-Vlan-20)# ip igmp snooping enable
```
 - g. Exit interface configuration mode:


```
switch(config-Vlan-20)# exit
```
 - h. Enter VLAN interface configuration mode for the third VLAN:


```
switch (config)# int vlan 30
```
 - i. Enable IGMP snooping by issuing the following command:


```
switch(config-Vlan-30)# ip igmp snooping enable
```
 - j. Exit interface configuration mode:


```
switch(config-Vlan-30)# exit
```
 - k. Enter VLAN interface configuration mode for the fourth VLAN:


```
switch (config)# int vlan 40
```
 - l. Enable IGMP snooping by issuing the following command:


```
switch(config-Vlan-40)# ip igmp snooping enable
```

- m. Exit interface configuration mode:


```
switch(config-Vlan-40)# exit
```
2. Do the following on switch M1 in [Figure 47](#):
 - a. From the switch console, in privileged EXEC mode, enter global configuration mode by issuing the **configure** command:


```
switch# configure
```
 - b. Enter VLAN interface configuration mode for the first VLAN:


```
switch (config)# int vlan 10
```
 - c. Enable IGMP snooping by issuing the following command:


```
switch(config-Vlan-10)# ip igmp snooping enable
```
 - d. Exit interface configuration mode:


```
switch(config-Vlan-10)# exit
```
 - e. Enter Rbridge subconfiguration mode for the Rbridge associated with this switch (for this example, it is assumed that an Rbridge ID of 1 has already been configured for this switch):


```
switch (config)# rbridge-id 1
```
 - f. Issue the **router pim** command to enable PIM for this switch:


```
switch(config-rbridge-id-17)# router pim
```
 - g. To add a static RP configuration, add the rendezvous point (RP) address for the router PIM.


```
switch(config-rbridge-id-17)# rp-address 22.22.22.22
```
 - h. Enter interface subconfiguration mode for the ve interface associated with VLAN 10:


```
switch(config-rbridge-id-1)# int ve 10
```
 - i. Enter the **no shut** command to activate the ve interface and bring the ports online:


```
switch (config-ve-10)# no shut
```
 - j. Assign a unique IP address for the interface:


```
switch (config-ve-10)# ip addr 10.1.1.11/24
```
 - k. Enable PIM Sparse for this interface:


```
switch (config-ve-10)# ip pim-sparse
```
 - l. Enter the **end** command:


```
switch (config-ve-10)# end
```
 - m. Repeat [step b](#) through [step l](#) for each of the other VLANs in the example in [Figure 47](#).
3. Do the following on M2 in [Figure 47](#):
 - a. From the switch console, in privileged EXEC mode, enter global configuration mode by issuing the **configure** command:


```
switch# configure
```
 - b. Enter VLAN interface configuration mode for the first VLAN:


```
switch (config)# int vlan 10
```
 - c. Enable IGMP snooping by issuing the following command:

```
switch(config-Vlan-10)# ip igmp snooping enable
```

- d. Exit interface configuration mode:

```
switch(config-Vlan-10)# exit
```

- e. Enter Rbridge ID configuration mode for the Rbridge associated with this switch (for this example, it is assumed that an Rbridge ID of 2 has already been configured for this switch):

```
switch (config)# rbridge-id 2
```

- f. Issue the **router pim** command to enable PIM for this switch:

```
switch(config-rbridge-id-2)# router pim
```

- g. Enter interface subconfiguration mode for the ve interface associated with VLAN 10:

```
switch(config-rbridge-id-2)# int ve 10
```

- h. Enter the **no shut** command to activate the ve interface and bring the ports online:

```
switch (config-ve-10)# no shut
```

- i. Assign a unique IP address for the interface:

```
switch (config-ve-10)# ip addr 10.1.1.12/24
```

- j. Enable PIM Sparse for this interface:

```
switch (config-ve-10)# ip pim-sparse
```

- k. Enter the **end** command:

```
switch (config-ve-10)# end
```

- l. Repeat [step b](#) through [step k](#) for each of the other VLANs in the example in [Figure 47](#).

NOTE

For more information about PIM, refer to the *Network OS Command Reference*. Global PIM CLIs are found in Rbridge ID configuration mode, while the interface-level PIM CLIs are found under their respective interface subconfiguration modes.

Configuring OSPF

In this chapter

- Overview of OSPF 525
- OSPF over VRF 527
- OSPF in a VCS environment 529
- Using designated routers 531
- Performing basic OSPF configuration 532
- Changing other settings 539
- Disabling OSPF on the router 539

Overview of OSPF

Open Shortest Path First (OSPF) is a link-state routing protocol that uses link-state advertisements (LSAs) to update neighboring routers about a router's interfaces. Each router maintains an identical area-topology database to determine the shortest path to any neighboring router.

- OSPF must be configured in a Virtual Cluster Switching (VCS) environment.
- The following platforms support OSPF:
 - Brocade VDX 6710-54
 - Brocade VDX 6720
 - Brocade VDX 6730
 - Brocade VDX 6740
 - Brocade VDX 6740T
 - VDX 8770-4
 - VDX 8770-8
- OSPF can be configured on either a point-to-point or broadcast network.
- OSPF can be enabled on the following interfaces: gigabitethernet, tengigabitethernet, fortygigabitethernet, loopback, and ve.
- On enabling OSPF over a loopback interface, the network is advertised as a stub network in the router LSA for the attached area. OSPF control packets, such as *hellos*, are not transmitted on loopback interfaces and adjacencies will not form.

OSPF is built upon a hierarchy of network components. The highest level of the hierarchy is the *Autonomous System (AS)*. An autonomous system is defined as a number of networks, all of which share the same routing and administration characteristics.

An AS can be divided into multiple *areas* as shown in [Figure 48](#) on page 527. Each area represents a collection of contiguous networks and hosts. Areas limit the amount of advertisements sent (called *flooding*) within the network. An area is represented in OSPF by either an IP address or a number.

The *backbone area* (also known as area 0 or area 0.0.0.0) forms the core of an OSPF network. All other areas are connected to it, and inter-area routing happens via routers connected to the backbone area and to their own associated areas. The backbone area is the logical and physical structure for the OSPF domain and is attached to all nonzero areas in the OSPF domain.

The backbone area is responsible for distributing routing information between non-backbone areas. The backbone must be contiguous, but it does not need to be physically contiguous; backbone connectivity can be established and maintained through the configuration of *virtual links*.

You can further consolidate routes at an area boundary by defining an *area range*. The area range allows you to assign an aggregate value to a range of IP addresses. This aggregate value becomes the address that is advertised instead all of the individual addresses it represents being advertised. You can assign up to 32 ranges in an OSPF area.

An OSPF router can be a member of multiple areas. Routers with membership in multiple areas are known as *Area Border Routers (ABRs)*. Each ABR maintains a separate topological database for each area the router is in. Each topological database contains all LSA databases for each router within a given area. The routers within the same area have identical topological databases. An ABR is responsible for forwarding routing information or changes among its border areas.

An *Autonomous System Boundary Router (ASBR)* is a router that is running multiple protocols and serves as a gateway to routers outside the OSPF domain and those operating with different protocols. The ASBR is able to import and translate different protocol routes into OSPF through a process known as *redistribution*. (For more information about redistribution, see the **redistribute** command in *Network OS Command Reference*.)

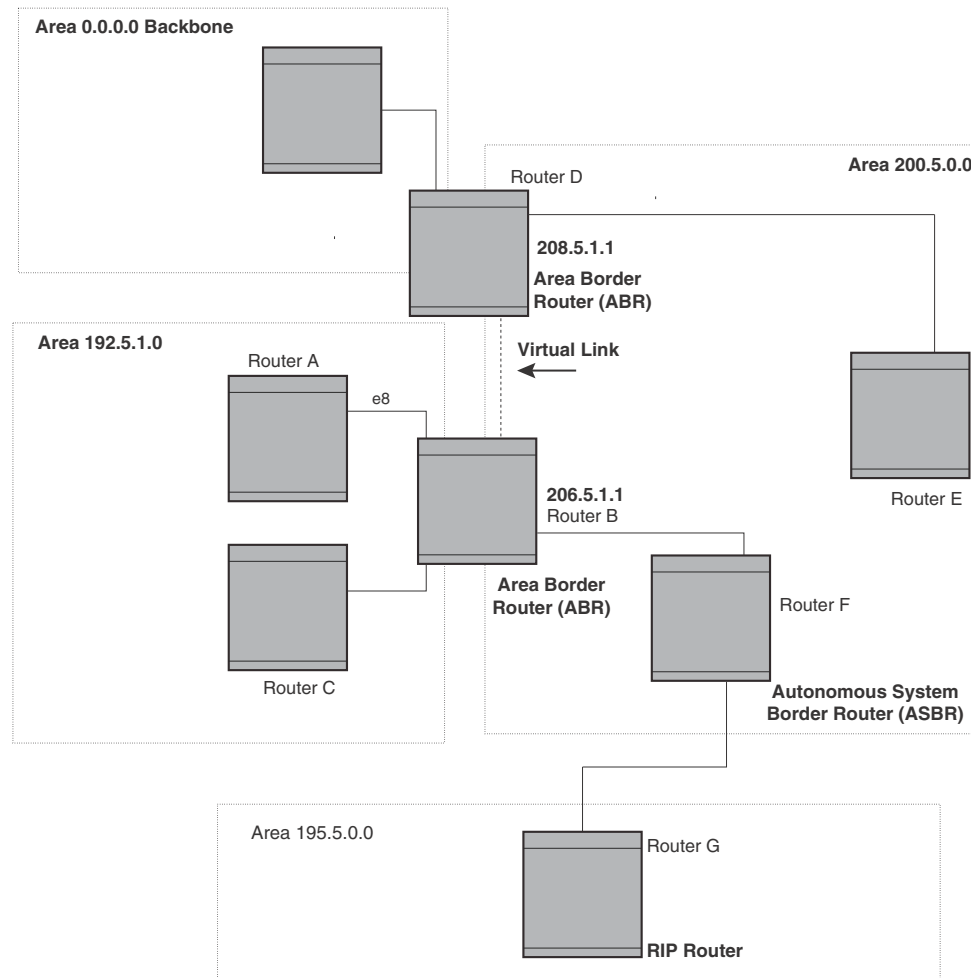


FIGURE 48 OSPF operating in a network

OSPF over VRF

With NOS 4.0 and later, OSPF can run over multiple Virtual Forwarding and Routing (VRF) mechanisms. OSPF maintains multiple instances of the routing protocol to exchange route information among various VRFs. A multi-VRF-capable router maps an input interface to a unique VRF, based on user configuration. These input interfaces can be physical or SVIs. By default, all input interfaces are attached to the default VRF. All OSPF commands supported in NOS4.0 and later are available over default and non-default OSPF instances.

Some VRF considerations include:

- To enable OSPF on a default VRF and to enter OSPF VRF router configuration mode, run the **router ospf** command in RBridge ID configuration mode, as shown in the following example:

```
switch# configure
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# router ospf
switch(config-router-ospf-vrf-default-vrf)#
```

- To enable OSPF on a non-default VRF and to enter OSPF VRF router configuration mode, run the **router ospf vrf name** command in RBridge ID configuration mode, as shown in the following example:

```
switch# configure
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# router ospf vrf vrfname
switch(config-router-ospf-vrf-vrfname)#
```

- OSPF **show** commands include the optional **vrf name** keyword to display data from non-default OSPF instances.
- Router-ID calculation for an OSPF instance includes IP addresses that are attached to the same VRF. The same subnets can coexist on multiple VRFs.

NOTE

For more information about OSPF over VRF, see [Chapter 40, “Configuring VRF”](#).

OSPF in a VCS environment

Figure 49 shows one way in which OSPF can be used in a VCS fabric cluster environment. Routers RB1 and RB2, as well as the MLX switches, are configured with OSPF. Switches RB3, RB4, and RB5 are Layer 2 switches.

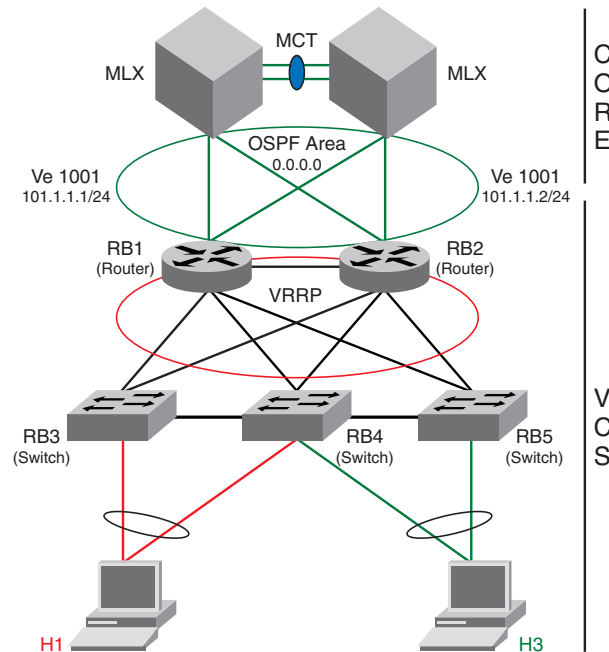


FIGURE 49 OSPF example in a VCS environment

1. On Router RB1, do the following:
 - a. Enter the **conf t** command to enter terminal configuration mode.
 - b. Enter the **interface vlan** command followed by the VLAN number to create a VLAN for the router.
 - c. Enter the **exit** command to exit interface configuration mode.
 - d. Enter the **rbridge-id** command followed by the Rbridge ID to enter RBridge configuration mode.
 - e. Enter the **router ospf** command to enable the OSPF routing protocol and to enter OSPF VRF router configuration mode.
 - f. Enter the **area** operand followed by the area ID to create this OSPF area on this router.
 - g. Enter the **exit** command to exit OSPF VRF router configuration mode.
 - h. Enter the **interface ve** command followed by the VLAN number to enter interface configuration mode.
 - i. Enter the **ip address** operand followed by the IP address/subnet of the interface.
 - j. Enter the **ip ospf area** operand followed by the area ID to assign the interface to this area.

- k. Enter the **no shutdown** command:

```
RB1# conf t
RB1(config)# interface vlan 1001
RB1(config-Vlan-1001)# exit
RB1(config)# rbridge-id 1
RB1(config-rbridge-id-1)# router ospf
RB1(config-router-ospf-vrf-default-vrf)# area 0.0.0.0
RB1(config-router-ospf-vrf-default-vrf)# exit
RB1(config-rbridge-id-1)# interface ve 1001
RB1(config-Ve-1001)# ip address 101.1.1.1/24
RB1(config-Ve-1001)# ip ospf area 0.0.0.0
RB1(config-Ve-1001)# no shutdown
```

2. On Router RB2, do the following:

- a. Enter the **conf t** command to enter terminal configuration mode.
- b. Enter the **interface vlan** command followed by the VLAN number to create a VLAN for the router.
- c. Enter the **exit** command to exit interface configuration mode.
- d. Enter the **rbridge-id** command followed by the RBridge ID to enter RBridge configuration mode.
- e. Enter the **router ospf** command to enable the OSPF routing protocol and to enter OSPF VRF router configuration mode.
- f. Enter the **area** operand followed by the area ID to create this OSPF area on this router.
- g. Enter the **exit** command to exit OSPF VRF router configuration mode.
- h. Enter the **interface ve** command followed by the VLAN number to enter interface configuration mode.
- i. Enter the **ip address** operand followed by the IP address/subnet of the interface.
- j. Enter the **ip ospf area** operand followed by the area ID to assign the interface to this area.
- k. Enter the **no shutdown** command:

```
RB2# conf t
RB2(config)# interface vlan 1001
RB2(config-Vlan-1001)# exit
RB2(config)# rbridge-id 2
RB2(config-rbridge-id-2)# router ospf
RB2(config-router-ospf-vrf-default-vrf)# area 0.0.0.0
RB2(config-router-ospf-vrf-default-vrf)# exit
RB2(config-rbridge-id-2)# interface ve 1001
RB2(config-Ve-1001)# ip address 101.1.1.2/24
RB2(config-Ve-1001)# ip ospf area 0.0.0.0
RB2(config-Ve-1001)# no shutdown
```

3. Assign VLAN 1001 to a vLAG.

Using designated routers

In an OSPF broadcast network, OSPF elects one router to serve as the designated router (DR) and another router on the segment to act as the backup designated router (BDR). This minimizes the amount of repetitive information that is forwarded on the network. OSPF forwards all messages to the designated router.

On broadcast networks such as LAN links, all routers on the LAN other than the DR and BDR form full adjacencies with the DR and BDR and pass LSAs only to them. The DR forwards updates received from one neighbor on the LAN to all other neighbors on that same LAN. One of the main functions of a DR is to ensure that all the routers on the same LAN have identical LSDBs. Therefore, on broadcast networks, an LSDB is synchronized between a DROther (a router that is not a DR or a BDR) and its DR and BDR.

NOTE

In an OSPF point-to-point network, where a direct Layer 3 connection exists between a single pair of OSPF routers, there is no need for designated or backup designated routers.

In a network with no designated router and no backup designated router, the neighboring router with the highest priority is elected as the DR, and the router with the next highest priority is elected as the BDR, as shown in Figure 50. Priority is a configurable option at the interface level; refer to the **ip ospf priority** command in *Network OS Command Reference*.




FIGURE 50 Designated and backup router election

If the DR goes off line, the BDR automatically becomes the DR. The router with the next highest priority becomes the new BDR.

If two neighbors share the same priority, the router with the highest router ID is designated as the DR. The router with the next highest router ID is designated as the BDR. The DR and BDRs are recalculated after the OSPF protocol is disabled and re-enabled by means of the **[no] router ospf** command.

NOTE

By default, the Brocade device's router ID is the IP address configured on the lowest numbered loopback interface. If the device does not have a loopback interface, the default router ID is the lowest numbered IP address configured on the device.

When multiple routers on the same network are declaring themselves DRs, then both the priority and router ID are used to select the designated router and backup designated routers.

The DR and BDR election process is performed when one of the following events occurs:

- An interface is in a waiting state and the wait time expires.
- An interface is in a waiting state and receives a hello packet that addresses the BDR.
- A change in the neighbor state occurs, such as the following:
 - A neighbor state transitions from ATTEMPT state to a higher state.
 - Communication to a neighbor is lost.
 - A neighbor declares itself to be the DR or BDR for the first time.

Performing basic OSPF configuration

To begin using OSPF on the router, perform these steps:

1. Follow the rules in the [“Configuration rules”](#) on page 532.
2. Run the **router ospf** command in RBridge ID configuration mode to enable OSPF on the router. This is shown in [“OSPF over VRF”](#) on page 527.
3. Assign the areas to which the router will be attached. Refer to [“Assigning OSPF areas”](#) on page 532.
4. Assign individual interfaces to the OSPF areas. Refer to [“Assigning interfaces to an area”](#) on page 537.
5. Assign a virtual link to any ABR that does not have a direct link to the OSPF backbone area. Refer to [“Assigning virtual links”](#) on page 537.
6. Refer to [“Changing other settings”](#) on page 539.

Configuration rules

Note the following:

- If a router is to operate as an ASBR, you must enable the ASBR capability at the system level.
- Redistribution must be enabled on routers configured to operate as ASBRs.
- All router ports must be assigned to one of the defined areas on an OSPF router. When a port is assigned to an area, all corresponding subnets on that port are automatically included in the assignment.

Assigning OSPF areas

Once OSPF is enabled on the system, you can assign areas. Assign an IP address or number as the *area ID* for each area. The area ID is representative of all IP addresses (subnets) on a router port. Each port on a router can support one area.

An area can be *normal*, a *stub*, or a *not-so-stubby area* (NSSA):

- **Normal** – OSPF routers within a normal area can send and receive external link state advertisements (LSAs).
- **Stub** – OSPF routers within a stub area cannot send or receive external LSAs. In addition, OSPF routers in a stub area must use a default route to the area’s Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) to send traffic out of the area.

- **NSSA** – The ASBR of an NSSA can import external route information into the area.
 - ASBRs redistribute (import) external routes into the NSSA as type 7 LSAs. Type 7 External LSAs are a special type of LSA generated only by ASBRs within an NSSA, and are flooded to all the routers within only that NSSA.
 - ABRs translate Type 7 LSAs into type-5 External LSAs, which can then be flooded throughout the AS. You can configure summary-addresses on the ABR of an NSSA so that the ABR converts multiple Type 7 external LSAs received from the NSSA into a single Type 5 external LSA.

When an NSSA contains more than one ABR, OSPF elects one of the ABRs to perform the LSA translation for NSSA. OSPF elects the ABR with the highest router ID. If the elected ABR becomes unavailable, OSPF automatically elects the ABR with the next highest router ID to take over translation of LSAs for the NSSA. The election process for NSSA ABRs is automatic.

Example

To set up the backbone area shown in [Figure 48](#) on page 527:

1. In privileged EXEC mode on Router A, issue the **configure** command to enter global configuration mode.
2. Enter the **interface vlan** command followed by the VLAN number to create a VLAN.
3. Enter the **rbridge-id** command followed by the RBridge ID to enter RBridge configuration mode.
4. Enter the **interface ve** command followed by the VLAN number to enter interface configuration mode.
5. Enter the **ip address** operand followed by the IP address/subnet for the interface.
6. Issue the **ip ospf area** operand followed by the area ID to assign the interface to this area.

```
Router A# configure
Router A(config) # interface vlan 1001
Router A(config-Vlan-1001) # rbridge 10
Router A(config-rbridge-id-10) # interface ve 1001
Router A(config-Ve-1001) # ip address 101.1.1.1/24
Router A(config-Ve-1001) # ip ospf area 0.0.0.0
```

Supported link state advertisements

The link state advertisements (LSAs) supported for each area type are as follows:

- Backbone (area 0) supports LSAs 1, 2, 3, 4, 5, and 7.
- Nonbackbone, not stub area supports LSAs 1, 2, 3, 4, and 5.
- Stub area supports LSAs 1, 2, and 3.
- Totally stubby area (TSA) supports LSAs 1 and 2. Also supports a single LSA 3 per ABR, advertising a default route.
- No so stubby area (NSSA) supports LSAs 1, 2, 3, and 7.

Assigning a totally stubby area

By default, the device sends summary LSAs (LSA Type 3) into stub areas. You can further reduce the number of link state advertisements (LSA) sent into a stub area by configuring the device to stop sending summary LSAs (Type 3 LSAs) into the area. This is called assigning a *totally stubby area* (TSA). You can disable the summary LSAs when you are configuring the stub area or later after you have configured the area.

This feature disables origination of summary LSAs, but the device still accepts summary LSAs from OSPF neighbors and floods them to other neighbors.

When you enter a command to disable the summary LSAs, the change takes effect immediately. If you apply the option to a previously configured area, the device flushes all the summary LSAs it has generated (as an ABR) from the area.

NOTE

This feature applies only when the device is configured as an Area Border Router (ABR) for the area. To completely prevent summary LSAs from being sent to the area, disable the summary LSAs on each OSPF router that is an ABR for the area.

To disable summary LSAs for a stub area, enter a command such as the following.

```
switch(config-router-ospf-vrf-default-vrf)# area 40 stub 99 no-summary
```

Assigning a not-so-stubby area (NSSA)

The OSPF not-so-stubby area (NSSA) feature enables you to configure OSPF areas that provide the benefits of stub areas, but that also are capable of importing external route information. OSPF does not flood external routes from other areas into an NSSA, but does translate and flood route information from the NSSA into other areas such as the backbone.

NSSAs are especially useful when you want to summarize Type 5 External LSAs (external routes) before forwarding them into an OSPF area. The OSPF specification prohibits summarization of Type 5 LSAs and requires OSPF to flood Type 5 LSAs throughout a routing domain. When you configure an NSSA, you can specify a summary-address for aggregating the external routes that the NSSA's ABR exports into other areas.

Figure 51 shows an example of an OSPF network containing an NSSA.

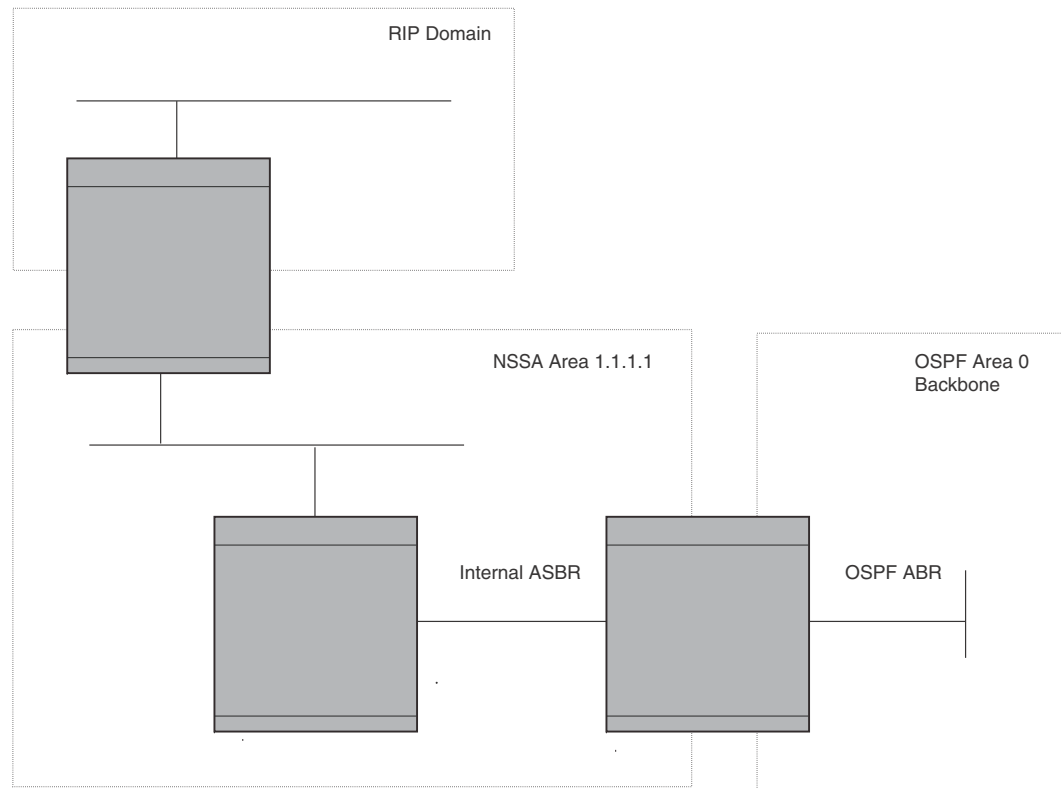


FIGURE 51 OSPF network containing an NSSA

This example shows two routing domains, a RIP domain and an OSPF domain. The ASBR inside the NSSA imports external routes from RIP into the NSSA as Type 7 LSAs, which the ASBR floods throughout the NSSA.

The ABR translates the Type 7 LSAs into Type 5 LSAs. If a summary-address is configured for the NSSA, the ABR also summarizes the LSAs into an aggregate LSA before flooding the Type 5 LSAs into the backbone.

Because the NSSA is partially stubby the ABR does not flood external LSAs from the backbone into the NSSA. To provide access to the rest of the Autonomous System (AS), the ABR generates a default Type 7 LSA into the NSSA.

Configuring an NSSA

To configure OSPF area 1.1.1.1 as an NSSA:

1. In privileged EXEC mode, issue the **configure** command to enter global configuration mode.
2. Enter the **rbridge-id** command followed by the RBridge ID to enter RBridge configuration mode.

3. Enter the **router ospf** command to enable OSPF on the router.
4. Enter the **area** operand followed by the area ID, then enter the **nssa** operand followed by the nssa ID.

```
switch# configure
switch(config)# rbridge-id 101
switch(config-rbridge-id-101)# router ospf
Switch(config-router-ospf-vrf-default-vrf)# area 1.1.1.1 nssa 1
```

Configuring a summary-address for the NSSA

If you want the ABR that connects the NSSA to other areas to summarize the routes in the NSSA before translating them into Type 5 LSAs and flooding them into the other areas, configure a summary-address. The ABR creates an aggregate value based on the summary-address. The aggregate value becomes the address that the ABR advertises instead of advertising the individual addresses represented by the aggregate. You can configure up to 32 ranges in an OSPF area.

To configure a summary-address in NSSA 1.1.1.1 (This example assumes that you have already configured NSSA 1.1.1.1.):

1. In privileged EXEC mode, issue the **configure** command to enter global configuration mode.
2. Enter the **rbridge-id** command followed by the RBridge ID to enter RBridge configuration mode.
3. Enter the **router ospf** command to enable OSPF on the router and to enter router OSPF configuration mode.
4. Enter the **area** operand followed by the area ID, then enter the **nssa** operand followed by the NSSA ID.
5. Enter the **summary-address** command followed by the IP address and mask for the summary route.

```
switch# configure
switch(config)# rbridge-id 101
switch(config-rbridge-id-101)# router ospf
switch(config-router-ospf-vrf-default-vrf)# area 1.1.1.1 nssa 10
switch(config-router-ospf-vrf-default-vrf)# summary-address 209.157.1.0
255.255.255.0
```

Assigning an area range (optional)

You can assign a *range* for an area, but it is not required. Ranges allow a specific IP address and mask to represent a range of IP addresses within an area, so that only that reference range address is advertised to the network, instead of all the addresses within that range. Each area can have up to 32 range addresses.

Example

To define an area range for subnets on 0.0.0.10 and 0.0.0.20:

1. In privileged EXEC mode, issue the **configure** command to enter global configuration mode.
2. Issue the **rbridge-id** command followed by the RBridge ID to enter RBridge configuration mode.
3. Issue the **router ospf** command to enable OSPF on the router.
4. Issue the **area** operand followed by the area ID, then enter the range, and repeat as necessary.

```
switch# configure
switch(config)# rbridge-id 101
```



```
switch(config-rbridge-id-101)# router ospf
switch(config-router-ospf-vrf-default-vrf)# area 0.0.0.10 range 192.45.0.0
255.255.0.0
switch(config-router-ospf-vrf-default-vrf)# area 0.0.0.20 range 192.45.0.0
255.255.0.0
```

Assigning interfaces to an area

Once you define OSPF areas, you can assign interfaces to the areas. All router ports must be assigned to one of the defined areas on an OSPF router. When a port is assigned to an area, all corresponding subnets on that port are automatically included in the assignment.

For example, to assign a tenGigabitEthernet interface 101/0/1 to a router area whose IP address is 192.5.0.0:

1. In privileged EXEC mode, issue the **configure** command to enter global configuration mode.
2. Issue the **rbridge-id** command followed by the RBridge ID to enter Rbridge sub-configuration mode.
3. Issue the **interface** command followed by the interface ID to enter interface configuration mode.
4. Issue the **ip ospf area** command followed by the IP address of the area.

```
switch# configure
switch(config)# rbridge-id 101
switch(config-rbridge-id-101)# int te 101/0/1
switch(config-if-te-101/0/1)# ip ospf area 192.5.0.0
```

If you want to set an interface to passive mode, use the **ip ospf passive** command. If you want to block flooding of outbound LSAs on specific OSPF interfaces, use the **ip ospf database-filter all out** command. (Refer to the *Network OS Command Reference* for details.)

Assigning virtual links

All ABRs (area border routers) must have either a direct or indirect link to the OSPF backbone area (0.0.0.0 or 0). If an ABR does not have a physical link to the area backbone, the ABR can configure a *virtual link* to another router within the same area, which has a physical connection to the area backbone.

The path for a virtual link is through an area shared by the neighbor ABR (router with a physical backbone connection), and the ABR requires a logical connection to the backbone.

Two parameters fields must be defined for all virtual links—transit area ID and neighbor router:

- The *transit area ID* represents the shared area of the two ABRs and serves as the connection point between the two routers. This number should match the area ID value.
- The *neighbor router* field is the router ID (IP address) of the router that is physically connected to the backbone, when assigned from the router interface requiring a logical connection. When assigning the parameters from the router with the physical connection, be aware that the router ID is the IP address of the router requiring a logical connection to the backbone.

NOTE

By default, the Brocade device's router ID is the IP address configured on the lowest numbered loopback interface. If the device does not have a loopback interface, the default router ID is the lowest numbered IP address configured on the device. When you establish an area virtual link, you must configure it on both of the routers (both ends of the virtual link).

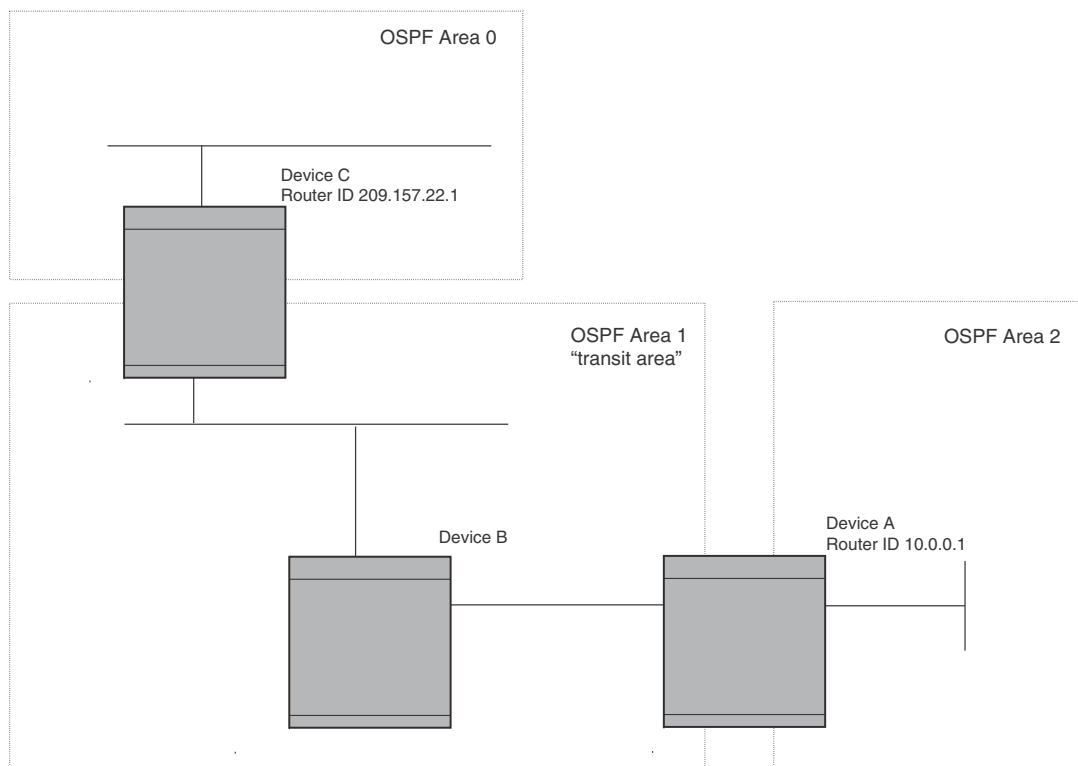


FIGURE 52 Defining OSPF virtual links within a network

Figure 52 shows an OSPF area border router, Device A, that is cut off from the backbone area (area 0). To provide backbone access to Device A, you can add a virtual link between Device A and Device C using area 1 as a transit area. To configure the virtual link, you define the link on the router that is at each end of the link. No configuration for the virtual link is required on the routers in the transit area.

To define the virtual link on Device A:

1. In privileged EXEC mode, issue the **configure** command to enter global configuration mode.
2. Enter the **rbridge-id** command followed by the RBridge ID to enter RBridge configuration mode.
3. Enter the **router ospf** command to enable OSPF on the router.
4. Enter the **area** operand followed by the area ID, and repeat as necessary.
5. Enter the **area** operand followed by the area address in decimal or dotted-decimal format, then enter the **virtual-link** operand followed by ID of the OSPF router at the remote end of the virtual link.

```
Device A# configure
Device A(config)# rbridge-id 101
Device A(config-rbridge-id-101)# router ospf
```

```
Device A(config-router-ospf-vrf-default-vrf)# area 2
Device A(config-router-ospf-vrf-default-vrf)# area 1
Device A(config-router-ospf-vrf-default-vrf)# area 1 virtual-link 209.157.22.1
```

To configure the virtual link on Device C:

1. In privileged EXEC mode, issue the **configure** command to enter global configuration mode.
2. Enter the **rbridge-id** command followed by the RBridge ID to enter RBridge configuration mode.
3. Enter the **router ospf** command to enable OSPF on the router.
4. Enter the **area** operand followed by the area ID, and repeat as necessary.
5. Enter the **area** operand followed by the area address in decimal or dotted-decimal format, then enter the **virtual-link** operand followed by ID of the OSPF router at the remote end of the virtual link

```
Device C# configure
Device C(config)# rbridge-id 101
Device C(config-rbridge-id-101)# router ospf
Device C(config-router-ospf-vrf-default-vrf)# area 0
Device C(config-router-ospf-vrf-default-vrf)# area 1
Device C(config-router-ospf-vrf-default-vrf)# area 1 virtual-link 10.0.0.1
```

Changing other settings

Refer to the *Network OS Command Reference* for other commands you can use to change default OSPF settings. Some commonly configured items include the following:

- Changing reference bandwidth to change interface costs by using the **auto-cost reference-bandwidth** command.
- Defining redistribution filters for the Autonomous System Boundary Router (ASBR) by using the **redistribute** command.

Disabling OSPF on the router

To disable OSPF on the router, use the **no router ospf** command:

1. In privileged EXEC mode, issue the **configure** command to enter global configuration mode.
2. Enter the **rbridge-id** command followed by the RBridge ID to enter RBridge configuration mode.
3. Issue the **no router ospf** command to disable OSPF on the router.

```
switch# configure
switch(config)# rbridge-id 101
switch(config-rbridge-id-101)# no router ospf
```

If you disable OSPF, the device removes all the configuration information for the disabled protocol from the running configuration. Moreover, when you save the configuration to the startup configuration file after disabling one of these protocols, all the configuration information for the disabled protocol is removed from the startup configuration file.

38 Disabling OSPF on the router

If you are testing an OSPF configuration and are likely to disable and re-enable the protocol, you might want to make a backup copy of the startup configuration file containing the protocol's configuration information. This way, if you remove the configuration information by saving the configuration after disabling the protocol, you can restore the configuration by copying the backup copy of the startup configuration file into the flash memory.

If the management default route information is available in the Chassis ID (CID) card, the OSPF default route is overwritten by the management default route when the switch reboots. In order to prevent this, remove the management default route after the switch reboots. The OSPF default route is automatically re-instated. Refer to [Chapter 43, "Using the Chassis ID \(CID\) Recovery Tool"](#).

Configuring VRRP

In this chapter

- Overview of virtual routers 541
- Guidelines 543
- VRRP/VRRP-E packet behavior 544
- VRRP basic configuration example 544
- Enabling preemption 546
- Using track ports and track priority with VRRP and VRRP-E 547
- Using short-path forwarding (VRRP-E only) 547
- Multigroup configuration for VRRP/VRRP-E 550

Overview of virtual routers

A virtual router is a collection of physical routers that can use the Virtual Router Redundancy Protocol (VRRP) to provide redundancy to routers within a LAN. Two or more VRRP-configured routers can create a virtual router.

VRRP eliminates a single point of failure in a static, default-route environment by dynamically assigning virtual IP routers to participating hosts. The interfaces of all routers in a virtual router must belong to the same IP subnet. There is no restriction against reusing a virtual router ID (VRID) with a different address mapping on different LANs.

Figure 53 shows a basic VRRP setup to illustrate some basic VRRP concepts. Router 1 and Router 2 are two physical routers that can be configured to compose one virtual router. This virtual router provides redundant network access for Host1. If Router 1 were to fail, Router 2 would provide the default gateway out of the subnet.

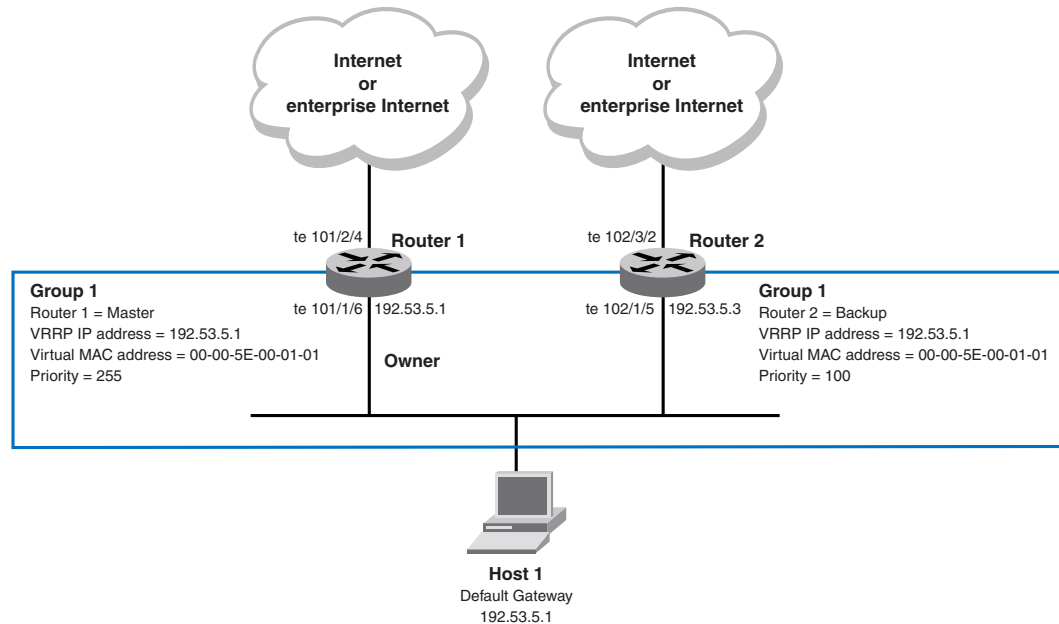


FIGURE 53 Basic VRRP setup

The virtual router shown in Figure 53 is identified as Group 1. A physical router forwards packets for the virtual router. This physical router is called the *master* router.

The following are some common VRRP-related terms and concepts:

- Virtual router—A collection of physical routers that can use either the VRRP or VRRP Extended (VRRP-E) protocol to provide redundancy to routers within a LAN.

NOTE

Most of the information in this chapter applies to both VRRP and VRRP-E, and, therefore, the term “VRRP” is often used to mean either VRRP or VRRP-E. Where there are differences between the two protocols, these differences are explicitly described.

- Virtual router group—A group of physical routers that are assigned to the same virtual router.
- Virtual router address—The address you are backing up:
 - For VRRP: The virtual router IP address must belong to the same subnet as a real IP address configured on the VRRP interface, and can be the same as a real IP address configured on the VRRP interface. The virtual router whose virtual IP address is the same as a real IP address is the IP address *owner* and the default *master*.
 - For VRRP-E: The virtual router IP address must belong to the same subnet as a real IP address configured on the VRRP-E interface, but cannot be the same as a real IP address configured on the VRRP-E interface.

- **Owner**—This term applies only to the VRRP protocol, not to VRRP-E. The owner is the physical router whose real interface IP address is the IP address that you assign to the virtual router. The owner responds to packets addressed to any of the IP addresses in the corresponding virtual router. The owner, by default, is the master (see “Master” below) and has the highest priority (255).
- **Master**—The physical router that responds to packets addressed to any of the IP addresses in the corresponding virtual router. For VRRP, if the physical router whose real interface IP address is the IP address of the virtual router, then this physical router is always the master. For VRRP-E, the router with the highest priority becomes the master. The **priority** command is used to set priority for a physical router.
- **Backup**—Routers that belong to a virtual router but are not the master. Then, if the master becomes unavailable, the backup router with the highest priority (a configurable value) becomes the new master. By default, routers are given a priority of 100.

NOTE

VRRP operation is independent of the Open Shortest Path First (OSPF) protocol or the Border Gateway Protocol (BGP).and is unaffected when enabled on interfaces running those protocols.

Guidelines

Virtual routers must be configured in a Virtual Cluster Switching (VCS) environment.

- The following platforms support VRRP and VRRP-E:
 - Brocade VDX 6710-54
 - Brocade VDX 6720
 - Brocade VDX 6730
 - Brocade VDX 6740
 - Brocade VDX 6740T
 - Brocade VDX 8770-4
 - Brocade VDX 8770-8
- Brocade supports two VRRP protocols:
 - **Standard VRRP**—The standard router redundancy protocol, VRRP v2 supports the IPv4 environment. Also, the Brocade version of standard VRRP is compliant with RFC 3768.
 - **VRRP-E (Extended)**—A Brocade proprietary protocol similar to standard VRRP that is not standard compliant and cannot interoperate with VRRP.
- **Supported ports:**
 - For VRRP—**fortygigabitethernet**, **tengigabitethernet**, **gigabitethernet**, and **ve**.
 - For VRRP-E—**ve** ports only.
- Only IPv4 support is provided. IPv6 and VRRPv3 are not supported.

- Maximum number of supported configured VRRP and VRRP-E instances (an instance is a session configured on a router):
 - Brocade VDX 8770: 1,024
 - Brocade VDX 6740 and 6740T: 255
 - Brocade VDX 6710, VDX 6720, and VDX 6730: 64.
- Maximum number of virtual IP addresses per virtual router session is 16 for VRRP and one for VRRP-E.

VRRP/VRRP-E packet behavior

There are some differences in how VRRP and VRRP-E handle ARP and VRRP control packets.

Gratuitous ARP

VRRP: Sent only once when the VRRP router becomes the master.

VRRP-E: Sent every two seconds by the virtual router master because VRRP-E control packets do not use the virtual MAC address.

The source MAC address of the gratuitous ARP sent by the master is the virtual MAC address.

When a router (either master or backup) sends an ARP request or a reply packet, the MAC address of the sender is the MAC address of the router interface. One exception is if the owner sends an ARP request or a reply packet, in which case the MAC address of the sender is the virtual MAC address.

Only the master answers an ARP request for the virtual router IP address. Any backup router that receives this request forwards the request to the master.

VRRP control packets

VRRP: VRRP control packets are IP protocol type 112 (reserved for VRRP), and are sent to VRRP multicast address 224.0.0.18.

VRRP-E: control packets are UDP packets destined to port 8888, and are sent to the all-router multicast address 224.0.0.2.

Source MAC in VRRP control packets

VRRP: The virtual MAC address is the source.

VRRP-E: The physical MAC address is the source.

VRRP basic configuration example

You can implement the IPv4 VRRP configuration shown in [Figure 53](#) on page 542 by entering just a few commands. This section contains information for configuring each router shown in [Figure 53](#) on page 542. This is for a VCS fabric cluster mode environment.

Configuring Router 1 as master for VRRP

1. From the Router 1 switch (host name sw1 for this example) console, in privileged EXEC mode, enter configuration mode by issuing the **configure** command.

```
sw1# configure
```

2. Enter the **rbridge-id** command, using the RBridge ID (which has an asterisk next to it when you run a **do show vcs** command).

```
sw1(config)# rbridge-id 101
```

3. Globally enable both the VRRP and VRRP-E protocols.

```
sw1(config-rbridge-id-101)# protocol vrrp
```

4. Configure the tengigabitethernet interface link for Router 1.

```
sw1(config-rbridge-id-101)# int te 101/1/6
```

5. Configure the IP address of the interface.

```
sw1(conf-if-te-101/1/6)# ip address 192.53.5.1/24
```

6. Assign Router 1 to a group called Group 1.

```
sw1(conf-if-te-101/1/6)# vrrp-group 1
```

NOTE

You can assign a group number in the range of 1 through 255.

7. Assign a virtual router IP address.

```
sw1(config-vrrp-group-1)# virtual-ip 192.53.5.1
```

NOTE

For VRRP, the physical router whose IP address is the same as the virtual router group IP address becomes the owner and master. However, for VRRP-E, you use the **priority** command to assign the highest priority to the router you want as master.

Configuring Router 2 as backup for VRRP

1. From the Router 2 switch (host name sw2 for this example) console, in privileged EXEC mode, enter configuration mode by issuing the **configure** command.

```
sw2# configure
```

2. Enter the **rbridge-id** command, using the RBridge ID (which has an asterisk next to it when you run a **do show vcs** command).

```
sw2(config)# rbridge-id 102
```

3. Globally enable both the VRRP and VRRP-E protocols.

```
sw2(config-rbridge-id-102)# protocol vrrp
```

4. Configure the tengigabitethernet interface for Router 2.

```
sw2(config-rbridge-id-102)# int te 102/1/5
```

5. Configure the IP address of interface:

```
sw2(conf-if-te-102/1/5)# ip address 192.53.5.3/24
```

NOTE

This router will become the backup router to Router 1.

6. Assign Router 2 to the same VRRP group as Router 1.

```
sw2(config-if-te-102/1/5)# vrrp-group 1
```

7. To assign Group 1 a virtual IP address, use the same virtual IP addresses you used for Router 1.

```
sw2(config-vrrp-group-1)# virtual-ip 192.53.5.1
```

VRRP-E differences for basic configuration

If you were to configure the two routers shown in [Figure 53](#) on page 542, you would need to consider the following items specific to VRRP-E:

- On the Brocade VDX 8770, 6710, 6720, and 6730 platforms, the **protocol vrrp** command enables both VRRP and VRRP-E.
- On the Brocade VDX 6740 and 6740T platforms, the **protocol vrrp** command enables only VRRP; the **protocol vrrp-extended** command enables only VRRP-E. VRRP and VRRP-E cannot be simultaneously enabled on the VDX 6740 or 6740T.
- The *group* command for VRRP-E is **vrrp-extended-group group-id**.
- VRRP-E virtual routers can be configured on Ve interfaces only.

Enabling preemption

You can allow a backup router that is acting as the master to be preempted by another backup router with a higher priority value.

Default: Preemption is enabled for VRRP, disabled for VRRP-E.

NOTE

If preemption is disabled for VRRP, the owner router is not affected because the owner router always preempts the active master.

To enable preemption for a virtual router, run the **preempt-mode** command in virtual-router-group configuration mode, as shown in the following example:

```
switch(config-vrrp-group-5)# preempt-mode
```

Using track ports and track priority with VRRP and VRRP-E

A track port allows you to monitor the state of the interfaces on the other end of a the route path. A track port also allows the virtual router to lower its priority if the exit path interface goes down, allowing another virtual router in the same VRRP (or VRRP-E) group to take over.

Rules

- Track priorities must be lower than VRRP/VRRP-E priorities.
- The dynamic change of router priority can trigger mastership switchover if preemption is enabled. However, if the router is an owner (applicable only for VRRP), the mastership switchover will not occur.
- Maximum number of interfaces that can be tracked for a virtual router is 16.
- Port tracking is allowed for physical interfaces and port-channels.

Track priority example

Using [Figure 55](#) on page 550 as an example, you can configure interface ve 10 on Router 1 to track interface 101/2/4. Then, if 101/2/4 goes down, interface ve 10 can respond by lowering the Router 1 VRRP priority by the track-port priority value. The backup routers detect this change and negotiate to become the new master, thus providing a master with an uninterrupted path out of the network.

Perform the following:

1. Enter interface configuration mode and run the following command:

```
switch(config)# int ve 10
```

2. Run the following command to enter group configuration mode.

```
switch(config-ve-10)# vrrp-group 1
```

3. Run the following command to set the track port and priority:

```
switch(config-vrrp-group-1)# track te 101/2/4 priority 60
```

Using short-path forwarding (VRRP-E only)

VRRP-E is enhanced with the VRRP-E extension for Server Virtualization feature so that Brocade devices attempt to bypass the VRRP-E master router and directly forward packets to their destination through interfaces on the backup router. This is called *short-path forwarding*. A backup router participates in a VRRP-E session only when short-path-forwarding is enabled.

VRRP-E active-active load-balancing is achieved with the ingress RBridge, by hashing either the L2-7 header information (Brocade VDX 8770) or the destination MAC address (Brocade VDX 67xx) to determine the path. All nodes in the VCS are aware of all VRRP-E sessions and the participating RBridges in each session.

If short-path forwarding is enabled, traffic travels through the short-path forwarding path (dashed line in [Figure 54](#)) to reach the client.

In Figure 54, the virtual servers are dynamically moved between Host Server 1 and Host Server 2.




FIGURE 54 Short path forwarding

Enabling short-path forwarding

Under the VRRP-E group-configuration level, there is an option to enable short-path-forwarding.

For example, follow these steps:

1. Enter global configuration mode.

```
switch# configure
```

2. In global configuration mode, entering the **int vlan** command.

```
switch(config)# int vlan 10
```

3. Exit the interface configuration mode by entering the **end** command.

```
switch(config-Vlan-10)# end
```

4. Enter global configuration mode.

```
switch# configure
```

5. Enter RBridge ID configuration mode.

```
switch (config)# rbridge-id 101
```

6. In RBridge ID configuration mode, enter the **int ve** command.

```
switch(config-rbridge-id-101)# int ve 10
```

7. In interface configuration mode, enter the **vrrp-extended-group** command.

```
switch(config-Ve-10)# vrrp-extended-group 100
```

8. In group configuration mode, enter the **short-path-forwarding** command.

```
switch(config-vrrp-extended-group-100)# short-path-forwarding
```

Packet routing with short-path forwarding (VRRP-E only)

If you enable short-path forwarding for VRRP-E, all packets sent by the local subnet of the virtual IP address are routed to the WAN instead of being switched to the master router.

Multigroup configuration for VRRP/VRRP-E

Figure 55 depicts a commonly employed virtual router setup. This setup introduces redundancy by configuring two virtual router groups — the first group has Router 1 as the master and Router 2 as the backup, and the second group has Router 2 as the master and Router 1 as the backup. This type of configuration is sometimes called *Multigroup VRRP*.

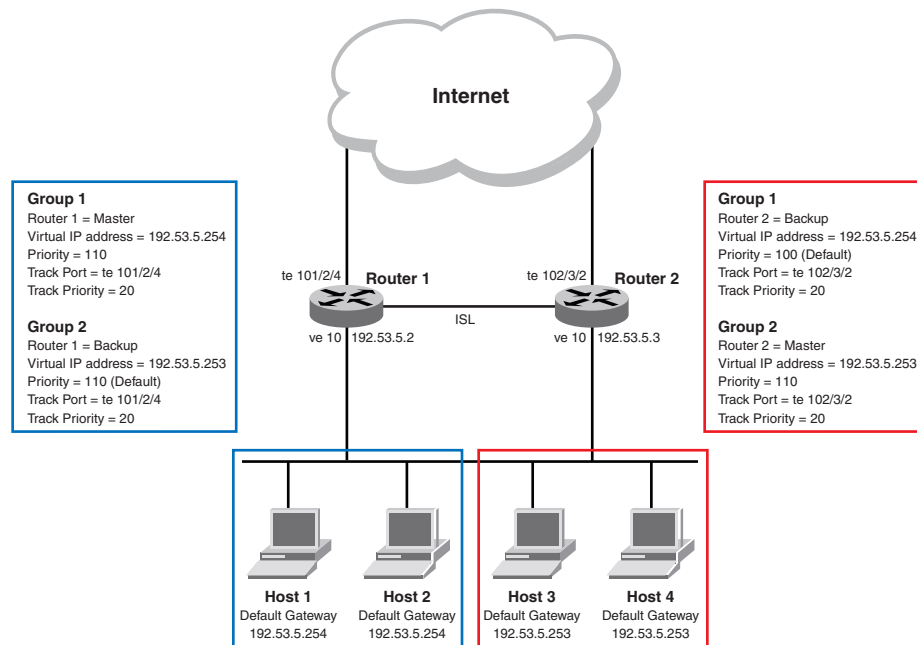


FIGURE 55 Router 1 and 2 configured for dual redundant network access for the host

In this example, Router 1 and Router 2 use VRRP-E to load share as well as provide redundancy to the hosts. The load sharing is accomplished by creating two VRRP-E groups, each with its own virtual IP addresses. Half of the clients point to Group 1's virtual IP address as their default gateway, and the other half point to Group 2's virtual IP address as their default gateway. This enables some of the outbound Internet traffic to go through Router 1 and the rest to go through Router 2.

Router 1 is the master for Group 1 (master priority = 110) and Router 2 is the backup for Group 1 (backup priority = 100). Router 1 and Router 2 both track the uplinks to the Internet. If an uplink failure occurs on Router 1, its backup priority is decremented by 20 (track-port priority = 20) to 90, so that all traffic destined to the Internet is sent through Router 2 instead.

Similarly, Router 2 is the master for Group 2 (master priority = 110) and Router 1 is the backup for Group 2 (backup priority = 100). Router 1 and Router 2 are both tracking the uplinks to the Internet. If an uplink failure occurs on Router 2, its backup priority is decremented by 20 (track-port priority = 20) to 90, so that all traffic destined to the Internet is sent through Router 1 instead.

Configuring a multi-group virtual router cluster

To implement the configuration shown in Figure 55 on page 550, configure one VRRP-E router to act as a master in the first virtual router group and a backup in the second virtual group. Then, configure the second VRRP-E router to act as a backup in the first virtual group and master in the second virtual group.

NOTE

This example is for VRRP-E. There are minor syntax differences for VRRP, which you can determine by consulting *Network OS Command Reference*, 3.0. This example is for a VCS fabric cluster mode environment.

Configuring Router 1 as master for first virtual router group

Make sure that VCS is enabled and then perform these steps:

1. Enter the **rbridge-id** command, using the RBridge ID (which has an asterisk next to it when you run a **do show vcs** command).

```
swl01(config)# rbridge-id 101
```

2. Configure the VRRP-E protocol globally.

```
swl01(config-rbridge-id-101)# protocol vrrp
```

3. Configure the Ve interface link for Router 1.

```
swl01(config-rbridge-id-101)# int ve 10
```

NOTE

You first would need to create **int vlan 10** in global configuration mode.

4. Configure the IP address of the Ve link for Router 1.

```
swl01(conf-Ve-10)# ip address 192.53.5.2/24
```

5. To assign Router 1 to a VRRP-E group called Group 1, enter the command:

```
swl01(conf-Ve-10)# vrrp-extended-group 1
```

6. Configure the tengigabitethernet port 101/2/4 as the tracking port for the interface ve 15, with a track priority of 20.

```
swl01(config-vrrp-extended-group-1)# track te 101/2/4 priority 20
```

7. Configure an IP address for the virtual router.

```
swl01(config-vrrp-extended-group-1)# virtual-ip 192.53.5.254
```

NOTE

(For VRRP-E only) The address you enter with the **virtual-ip** command cannot be the same as a real IP address configured on the interface.

8. To configure Router 1 as the master, set the priority to a value higher than the default (which is 100).

```
swl01(config-vrrp-group-1)# priority 110
```

Configuring Router 1 as backup for second virtual router group

1. Enter the **rbridge-id** command, using the RBridge ID (which has an asterisk next to it when you run a **do show vcs** command).

```
swl01(config)# rbridge-id 101
```

2. Configure the Ve interface link for Router 1.

```
sw101(config-rbridge-id-101)# int ve 15
```

3. Configure the IP address of the Ve link for router 1.

```
sw101(config-ve-15)# ip address 192.54.6.2/24
```

4. Assign Router 1 to a group called Group 2.

```
sw101(config-ve-15)# vrrp-extended-group 2
```

5. Configure the tengigabitethernet port 101/2/4 as the tracking port for the interface ve 10, with a track priority of 20.

```
sw101(config-vrrp-extended-group-2)# track te 101/2/4 priority 20
```

6. Configure an IP address for the virtual router.

```
sw101(config-vrrp-extended-group-2)# virtual-ip 192.53.6.253
```

NOTE

(For VRRP-E only) The address you enter with the **virtual-ip** command cannot be the same as a real IP address configured on the interface.

Configuring Router 2 as backup for first virtual router group

Make sure that VCS is enabled and then perform these steps:

1. Enter the **rbridge-id** command, using the RBridge ID (which has an asterisk next to it when you run a **do show vcs** command).

```
sw102(config)# rbridge-id 102
```

2. Configure the VRRP protocol globally.

```
sw102(config-rbridge-id-102)# protocol vrrp
```

3. Configure the Ve interface link for Router 2.

```
sw102(config-rbridge-id-102)# int ve 10
```

4. Configure the IP address of the Ve link for Router 2.

```
sw102(conf-ve-10)# ip address 192.53.5.3/24
```

5. Assign Router 2 to the group called Group 1.

```
sw102(conf-ve-10)# vrrp-extended-group 1
```

6. Configure the tengigabitethernet port 102/3/2 as the tracking port for interface ve 10, with a track priority of 20.

```
sw102(config-vrrp-extended-group-1)# track te 102/3/2 priority 20
```

7. Configure an IP address for the virtual router.

```
sw102(config-vrrp-extended-group-1)# virtual-ip 192.53.5.254
```

NOTE

(For VRRP-E only) The address you enter with the **virtual-ip** command cannot be the same as a real IP address configured on the interface.

Configuring Router 2 as master for second virtual router group

1. Enter the **rbridge-id** command, using the RBridge ID (which has an asterisk next to it when you run a **do show vcs** command).

```
sw102(config)# rbridge-id 102
```

2. Configure the Ve interface link for Router 2.

```
sw102(config-rbridge-id-102)# int ve 15
```

3. Configure the IP address of the Ve link for Router 2.

```
sw102(conf-Ve-15)# ip address 192.54.6.3/24
```

4. Assign Router 2 to the group called Group 2.

```
sw102(conf-Ve-15)# vrrp-extended-group 2
```

5. Configure the tengigabitethernet port 102/2/4 as the tracking port for interface ve 15, with a track priority of 20.

```
sw102(config-vrrp-extended-group-2)# track te 102/2/4 priority 20
```

6. To configure Router 2 as the master, set the priority to a value higher than the default (which is 100):.

```
sw102(config-vrrp-extended-group-2)# priority 110
```

7. Configure an IP address for the virtual router.

```
sw101(config-vrrp-extended-group-2)# virtual-ip 192.53.6.253
```

NOTE

(For VRRP-E only) The address you enter with the **virtual-ip** command cannot be the same as a real IP address configured on the interface.

39 Multigroup configuration for VRRP/VRRP-E

Configuring VRF

In this chapter

- [Overview of VRF 555](#)
- [Enabling VRRP for VRF 557](#)
- [OSPF VRF-Lite for customer-edge routers 558](#)

Overview of VRF

VRFs (Virtual Routing and Forwarding) is a technology that controls information flow within a network by isolating the traffic by partitioning the network into different logical VRF domains. Every VRF-capable router supports one routing table for each VRF instance. Each VRF-capable router can function as a group of multiple virtual routers on the same physical router. VRF, in conjunction with virtual private network (VPN) solutions, guarantees privacy of information and isolation of traffic within its logical VRF domain.

Typical full-blown implementations of VRFs are designed to support BGP/MPLS VPNs, whereas VRF-lite implementations typically are much simpler with moderate scalability (as compared to BGP/MPLS VPNs). These two flavors share a lot in common but differ in the interconnect schemes, routing protocols used over the interconnect, and also in the VRF classification mechanisms. Brocade Network OS v4.0.0 supports the VRF-lite implementation. All references to VRF in this document implicitly indicate VRF-lite. [Figure 56](#) shows a typical single VCS comprising Customer Edge 1, Provider Edge, and Customer Edge 2 routers.




FIGURE 56 VRF configuration diagram

ORANGE (v11) and GREEN (v12) are the two VPNs supporting two different customer sites. Both of them have overlapping IP subnets; 100.100.100.0/24 and 100.100.200.0/24.

Configuring VRF

VRF is supported on the Brocade VDX 8770 and VDX 6740, supporting up to 32 VRFs.

This configuration gives an example of a typical VRF-lite use case and is not meant to give an ideal configuration.

The examples in this section are based on the network diagram in [Figure 56](#).

(See the *Network OS Command Reference* for detailed information on VRF commands.)

1. Configure the edge routers.

Example of enabling routing and configuring VRF on the ORANGE edge router. Repeat this example for the GREEN edge router.

```
switch(config)# rbridge-id 1
switch(config-rbridge-id-1)# vrf orange
switch(config-vrf-orange)# rd 1:1
switch(config-vrf-orange)# address-family ipv4 max-route 3600
```

2. Enable OSPF protocol in VRF configuration mode.

All OSPF commands that are present under OSPF router configuration mode are applicable to the new OSPF VRF router configuration mode for a nondefault VRF, the same as is done for a default-VRF.

```
switch(vrf-ipv4)# router ospf vrf orange
switch(config-router-ospf-vrf-orange)# area 10
switch(vrf-ipv4)# exit
switch(config-vrf-orange)# exit
```

3. Bind the interface to a VRF.

NOTE

Once VRF is enabled on an interface, all Layer 3 configurations on the interface are removed, and you will need to configure them again.

```
switch(config)# rbridge-id 1
switch(config)# interface ve 128
switch(config-ve-128)# vrf forwarding orange
switch(config-vrf-orange)# address-family ipv4 max-route 3600
```

4. Configure the static routes.

```
switch(vrf-ipv4)# ip route 28.31.1.0/30 122.30.2.1
```

5. Configure static ARP for the interface.

```
switch(vrf-ipv4)# arp 2.2.2.3 0000.0011.0022 int ve 128
```

6. Confirm the VRF configuration with the **show vrf** command (using **do** in this configuration mode).

```
switch(vrf-ipv4)# do show vrf
Total number of VRFs configured: 1
Status Codes - A:active, D:pending deletion, I:inactive
Name          Default RD IFL ID vrf|v4|v6 Interfaces
orange        1:1          131072 A | A| I  Ve 128
```

Enabling VRRP for VRF

To enable VRRP/VRRP-E for a VRF region, an interface should be assigned to a VRF first and then followed by VRRP/VRRPE configuration. The VRRP protocol cannot be enabled or disabled on a specific VRF. It is enabled or disabled globally on the switch under rbridge-id mode.

Perform the following task in global configuration mode.

1. Set the switch into rbridge-id mode.

```
switch(config)# rbridge-id 1
```

2. Set the protocol to VRRP.

```
switch(config)# protocol vrrp
```

3. Select the interface

```
switch(config)# interface ve 128
```

4. Select the IP route.

```
switch(config)# ip address 172.128.20.10/24
```

5. Enable VRRP/VRRPE protocol for the interface.

```
switch(config)# vrrp-extended 10
```

6. Set the virtual IP address.

```
switch(config)# virtual-ip 172.128.20.1
```

OSPF VRF-Lite for customer-edge routers

When a type 3, 5, or 7 LSA is sent from a provider edge (PE) router to a customer edge (CE) router, the DN (down) bit in the LSA options field must be set. This prevents any type 3, 5, or 7 LSA messages sent from the CE router to the PE router from being distributed any farther. The PE router ignores messages with the DN bit set and does not add these routes to the VRF routing table.

When you enable VRF-Lite on the CE router, the DN setting is ignored, allowing the CE router to add these routes to the VRF routing table.

To enable VRF-Lite:

```
switch(config)# router ospf vrf 1  
switch(config-ospf-router-vrf-1)# vrf-lite-capability
```

Use the **no vrf-lite-capability** command to disable VRF-Lite. This applies to the VRF instance only. It does not apply to the default VRF.

Configuring BGP

In this chapter

- [BGP overview](#) 559
- [BGP configuration examples](#) 580

BGP overview

Border Gateway Protocol (BGP) is an exterior gateway protocol that performs inter-autonomous system (AS) or inter-domain routing. It peers to other BGP-speaking systems over TCP to exchange network reachability and routing information. BGP primarily performs two types of routing: inter-AS routing, and intra-AS routing. BGP peers belonging to different autonomous systems use the inter-AS routing, referred as Exterior BGP (EBGP). On the other hand, within an AS BGP can be used to maintain a consistent view of network topology, to provide optimal routing, or to scale the network.

BGP is a path vector protocol and implements this scheme on large scales by treating each AS as a single point on the path to any given destination. For each route (destination), BGP maintains the AS path and uses this to detect and prevent loops between autonomous systems.

The Open Shortest Path First (OSPF) protocol (supported since NOS v3.0.0) provides dynamic routing within the VCS and internal domain. However, even though OSPF suffices for most of the routing needs within the VCS, an exterior gateway protocol such as BGP is needed for inter-domain routing outside the VCS domain.

Support for BGP on NOS platforms is for BGP4 (compliant with RFC 1771 and 4271), and provides the following:

- Connectivity from the VCS to a core/external network or cloud
- A foundation to support virtual routing and forwarding (VRF) for multi-tenancy and remote-VCS access and route distribution across VRFs
- A foundation to support VRF scaling (OSPF does not scale well with lots of VRFs)
- A foundation to support OSPF Interior Gateway Protocol (IGP) scaling needs in future

NOTE

An L3 License is required to enable BGP routing.

Limitations and considerations

Note the following limitations and considerations with the initial release of BGP support on NOS platforms:

- IPv6 is not supported.
- VRF functionality is not supported.

- There is no backward compatibility. In case of a downgrade, BGP configurations are lost.
- There is no support for graceful restart or nonstop routing. Following a failover or switchover, BGP routes and active neighbors are lost. However, the configuration is restored with a restart.
- RASLogs are generated when a BGP session begins.
- RASTRACE logs are available with module ID “261 BGP.”

Deployment scenarios

BGP is typically used in a VCS fabric at the aggregation layer and in connecting to the core. Routing bridges at the aggregation layer can either connect directly to the core, or connect through an MLX. The topologies below illustrate connectivity with and without an MLX. The details of these topologies are discussed in subsequent sections.

Figure 57 illustrates connectivity to the core through an MLX. The R Bridges use OSPF and IBGP to communicate with each other, connecting to the MLX through IBGP. The MLX connects in turn to the core through EBGP.




FIGURE 57 Connectivity to the core through an MLX

Figure 58 illustrates the previous topology but without an MLX.




FIGURE 58 Connectivity to the core without an MLX

Figure 59 illustrates the role of BGP in communicating through multiple VCS clusters and autonomous systems.




FIGURE 59 BGP with multiple VCS clusters and autonomous systems

Figure 60 illustrates a BGP topology that incorporates a route-reflector server and route-reflector clients.




FIGURE 60 BGP route-reflector server and clients

BGP Peering

Unlike OSPF or other IGP protocols, BGP does not have neighbor detection capability. BGP neighbors (or peers) must be configured manually. A router configured to run BGP is called a BGP “speaker.” A BGP speaker connects to another speaker (either in the same or a different AS) by using a TCP connection to port 179 (the well-known BGP port), to exchange the routing information. The TCP connection is maintained throughout the peering session. While the connection between BGP peers is alive, two peers communicate by means of the following types of messages:

- OPEN
- UPDATE
- KEEPALIVE
- NOTIFICATION
- ROUTE REFRESH

BGP peering can be internal or external, depending on whether the two BGP peers belong to the same AS or different ASs. A BGP session between peers within a single AS is referred to as an Interior BGP (IBGP) session; a session between peers belonging to different ASs is referred to as an Exterior BGP (EBGP) session.

In order to establish a TCP connection between two IBGP peers, the IP reachability should be established either by means of the underlying IGP protocol (OSPF) or by means of static routes. When routes are advertised within IBGP peers, the following primary actions are taken in contrast to EBGP peering:

- Routes learned from an IBGP peer are not usually advertised to other IBGP peers, in order to prevent loops within an AS.
- Path attributes are not usually changed, in order to maintain the best path selection at other nodes within an AS.
- The AS path and next hop are not normally changed.

BGP Message Types

All BGP messages use a common packet header, with the following byte lengths:

Marker	Length	Type	Data
16	2	1	variable

NOTE

All values in the following tables are in bytes.

Type can be OPEN, UPDATE, NOTIFICATION, or KEEPALIVE, as described below.

OPEN message

After establishing TCP connection, BGP peers exchange OPEN message to identify each other.

Version	Autonomous System	Hold-Time	BGP Identifier	Optional Parameter Len	Optional Parameters
1	2 or 4	2	4	1	4

Version

Only BGP4 version 4 is supported.

Autonomous System

Both 2-byte and 4-byte AS numbers are supported.

KEEPALIVE and HOLDDTIME

A BGP **timer** command specifies both **keep-alive** and **hold-time** operands that manage the intervals for BGP KEEPALIVE and HOLDDTIME messages. The first operand sets the number of seconds the device waits for UPDATE/KEEPALIVE message before closing the TCP connection. The second operand sets the number of seconds that BGP maintains a session with a neighbor without receiving messages. When two neighbors have different hold-time values, the lowest value is used. A hold-time value of 0 means “always consider neighbor to be active.”

BGP Identifier

Indicates the router (or device) ID of the sender. When router-id is not configured, device-id is taken from the loopback interface. Otherwise, the lowest IP address in the system is used.

Parameter List

Optional list of additional parameters used in peer negotiation.

UPDATE message

The UPDATE message is used to advertise new routes, withdraw previously advertised routes, or both.

Withdrawn Routes Length	Withdrawn Routes	Total Path Attributes Len	Path Attributes	NLRI
2	variable	2	variable	variable

Withdrawn Routes Length

Indicates the length of next (withdrawn routes) field. It can be 0.

Withdrawn Routes

Contains list of routes (or IP-prefix/Length) to indicate routes being withdrawn.

Total Path Attribute Len

Indicates length of next (path attributes) field. It can be 0.

Path Attributes

Indicates characteristics of the advertised path. Possible attributes: Origin, AS Path, Next Hop, MED (Multi-Exit Discriminator), Local Preference, Atomic Aggregate, Aggregator.

NLRI

Network Layer Reachability Information— the set of destinations whose addresses are represented by one prefix. This field contains a list of IP address prefixes for the advertised routes.

NOTIFICATION message

In case of an error that causes the TCP connection to close, the closing peer sends a notification message to indicate the type of error.

Error Code	Error Subcode	Error Data
1	1	variable

Error Code

Indicates the type of error, which can be one of following:

- Message header error
- Open message error
- Update message error
- Hold timer expired
- Finite state-machine error
- Cease (voluntarily)

Error Subcode

Provides specific information about the error reported.

Error Data

Contains data based on error code and subcode.

KEEPALIVE message

Because BGP does not regularly exchanges route updates to maintain a session, KEEPALIVE messages are sent to keep the session alive. A KEEPALIVE message contains just the BGP header without data field. Default KEEPALIVE time is 60 seconds and is configurable.

REFRESH message

A REFRESH message is sent to a neighbor requesting that the neighbor resend the route updates. This is useful when the inbound policy has been changed.

BGP Attributes

BGP attributes are passed in UPDATE messages to describe the characteristics of a BGP path by the advertising router. At a high level, there are only two types of attributes: well-known and optional. All of the well-known attributes as described in RFC 4271 are supported in NOS v4.0.0.

Best-Path Algorithm

The BGP decision process is applied to the routes contained in the Routing Information Base, Incoming (RIB-In), which contains routes learned from inbound update messages. The output of the decision process is the set of routes that will be advertised to BGP speakers in local or remote autonomous systems and are stored in the Adjacency RIB, Outgoing (RIB-Out).

When multiple paths for the same route prefix are known to a BGP4 device, the device uses the following algorithm to weigh the paths and determine the optimal path for the route. (The optimal path depends on various parameters, which can be modified.)

1. Verify that the next hop can be resolved by means of Interior Gateway Protocol (IGP).
2. Use the path with the largest weight.
3. Prefer the path with the higher local preference.
4. Prefer the route that was self-originated locally.

5. Prefer the path with the shortest AS-path. An AS-SET counts as 1. A confederation path length, if present, is not counted as part of the path length. (An **as-path ignore** command disables the comparison of the AS path lengths of otherwise equal paths.)
6. Prefer the path with the lowest origin type. From low to high, route origin types are valued as follows:
 - IGP is lowest.
 - EGP is higher than IGP, but lower than INCOMPLETE.
 - INCOMPLETE is highest.
7. Prefer the path with the lowest MED.

The device compares the MEDs of two otherwise equivalent paths if and only if the routes were learned from the same neighboring AS. This behavior is called deterministic MED. Deterministic MED is always enabled and cannot be disabled.

To ensure that the MEDs are always compared, regardless of the AS information in the paths, the **always-compare-med** command can be used. This option is disabled by default.

The **med-missing-as-worst** command can be used to make the device regard a BGP4 route with a missing MED attribute as the least-favorable path when the MEDs of the route paths are compared.

MED comparison is not performed for internal routes that originate within the local AS or confederation, unless the **compare-med-empty-aspath** command is configured.

8. Prefer paths in the following order:
 - Routes received through EBGp from a BGP4 neighbor outside of the confederation
 - Routes received through EBGp from a BGP4 device within the confederation or routes received through IBGP.
9. If all the comparisons above are equal, prefer the route with the lowest IGP metric to the BGP4 next hop. This is the closest internal path inside the AS to reach the destination.
10. If the internal paths also are the same and BGP4 load sharing is enabled, load-share among the paths. Otherwise go to Step 11.

NOTE

For EBGp routes, load sharing applies only when the paths are from neighbors within the same remote AS. EBGp paths from neighbors in different ASs are not compared, unless multipath multi-as is enabled.

11. If **compare-routerid** is enabled, prefer the path that comes from the BGP4 device with the lowest device ID. If a path contains originator ID attributes, then the originator ID is substituted for the router ID in the decision.
12. Prefer the path with the minimum cluster-list length.
13. Prefer the route that comes from the lowest BGP4 neighbor address.

Configuration fundamentals

This section provides an overview of the configuration considerations and basic steps required to enable BGP functionality.

Similar to other Layer 3 protocols in NOS, BGP is supported only in the VCS mode of operation. Each RBridge in a VCS fabric running BGP acts as an individual BGP device. BGP can form IBGP peering with other RBridges in the same VCS fabric running BGP.

Configuring BGP

To enable BGP on an RBridge, enter Bridge ID configuration mode and issue the **router bgp** command:

```
switch(config-rbridge-id-12)# router bgp  
switch(config-bgp-router)#
```

There are two CLI modes for BGP:

- Global BGP
- BGP Address-Family IPv4 Unicast

After issuing the **router bgp** command, you first enter into BGP configuration mode, where an address-family-specific configuration can be applied. In order to apply an IPv4 address-family-specific configuration, issue the **address-family ipv4 unicast** command, described later.

To create a route map, enter the **route-map** command from RBridge ID configuration mode. Then declare a route-map name by using a **permit** or **deny** statement and an instance number.

To remove the entire BGP configuration, issue the **no router bgp** command.

Device ID

BGP automatically calculates the device identifier it uses to specify the origin in routes it advertises. If a router-id configuration is already present in the system, then device-id is used as the router-id. Otherwise, BGP first checks for a loopback interface, and the IP address configured on that interface is chosen as the device-id. However, if a loopback interface is not configured, the device-id is chosen from lowest-numbered IP interface address configured on the device. Once device-id is chosen, the device identifier is not calculated unless the IP address configured above is deleted.

Local AS number

The local AS number (ASN) identifies the AS in which the BGP device resides. It can be configured from BGP global mode:

```
switch(config-bgp-router)# local-as 6500  
switch(config-bgp-router)#
```

Use well-known private ASNs in the range from 64512 through 65535 if the AS number of the organization is not known.

IPv4 unicast address family

Currently only the IPv4 unicast address family is supported. This configuration is applied in the IPv4 address-family unicast submode of BGP:

```
switch(config-bgp-router)# address-family ipv4 unicast
switch(config-bgp-ipv4u)#
```

The following configurations are allowed under BGP IPv4 address-family unicast mode:

- Network (including static networks)
- Route aggregation
- Route redistribution
- Route reflection
- Dampening
- Default route origination
- Multipathing (including maximum paths)
- Address-family-specific neighbor configuration
- Explicit specification of networks to advertise

The following illustrates CLI options in address-family mode:

```
switch(config-bgp-router)# address-family ipv4 unicast
switch(config-bgp-ipv4u)# ?
Possible completions:
  aggregate-address      Configure BGP aggregate entries
  always-propagate       Allow readvertisement of best BGP routes not
                          in IP Forwarding table
  bgp-redistribute-internal Allow redistribution of iBGP routes into IGPs
  client-to-client-reflection Configure client to client route reflection
  dampening              Enable route-flap dampening
  default-information-originate Originate Default Information
  default-metric          Set metric of redistributed routes
  do                     Run an operational-mode command
  exit                   Exit from current mode
  help                   Provide help information
  maximum-paths           Forward packets over multiple paths
  multipath              Enable multipath for ibgp or ebgp neighbors
                          only
  neighbor               Specify a neighbor router
  network                Specify a network to announce via BGP
  next-hop-enable-default Enable default route for BGP next-hop lookup
  next-hop-recursion      Perform next-hop recursive lookup for BGP
                          route
  no                     Negate a command or set its defaults
  pwd                    Display current mode path
  redistribute            Redistribute information from another
                          routing protocol
  rib-route-limit         Limit BGP rib count in routing table
  static-network          Special network that do not depends on IGP
                          and always treat as best route in BGP
  table-map              Map external entry attributes into routing
                          table
  top                    Exit to top level and optionally run command
  update-time            Configure igp route update interval
```


BGP global mode

Configurations that are not specific to address-family configuration are available in the BGP global configuration mode:

```
switch(config-bgp-router)# ?
Possible completions:
  address-family          Enter Address Family command mode
  always-compare-med      Allow comparing MED from different neighbors
  as-path-ignore          Ignore AS_PATH length for best route selection
  capability              Set capability
  cluster-id              Configure Route-Reflector Cluster-ID
  compare-med-empty-aspath Allow comparing MED from different neighbors even with empty
                           as-path attribute
  compare-routerid        Compare router-id for identical BGP paths
  default-local-preference Configure default local preference value
  distance                Define an administrative distance
  enforce-first-as        Enforce the first AS for EBGp routes
  fast-external-fallover  Reset session if link to EBGp peer goes down
  install-igp-cost        Install igp cost to nexthop instead of MED value as BGP route
                           cost
  local-as                Configure local AS number
  log-dampening-debug     Log dampening debug messages
  maxas-limit             Impose limit on number of ASes in AS-PATH attribute
  med-missing-as-worst    Consider routes missing MED attribute as least desirable
  neighbor                Specify a neighbor router
  timers                  Adjust routing timers
```

Neighbor configuration

For each neighbor a device is going to peer with, there must be a neighbor configuration that specifies an IP address (which must be the primary IP address of interface connection to get established) and an AS number of the neighbor. For each neighbor, you can specify a set of attributes. However, in case a set of neighbors share same set of attributes, then it is advisable to use a peer-group. The peer-group configuration is described in the next subsection.

The following illustrates the configuration of a neighbor's IP address and AS number:

```
switch(config-bgp-router)# neighbor 10.231.64.10 remote-as 6500
switch(config-bgp-router)#
```

Note that the neighbor configuration appears in both the global and address-family modes of BGP. The neighbor parameters/attributes that are common to all of the address families appear in the BGP global mode; the parameters/attributes that are specific to an address family appear within the BGP address-family submode. Even though only the IPv4 unicast address family is supported currently, the options of the **neighbor** command are divided, to support future address families such as IPv6.

The following running-config excerpt illustrates the neighbor configurations that are allowed in BGP global mode and IPv4 address-family submode:

```
router bgp
  local-as 6500
  neighbor 10.231.64.10 advertisement-interval 60
  neighbor as-override
  neighbor 10.231.64.10 capability as4-enable
  neighbor 10.231.64.10 description "Example Neighbor Configuration"
  neighbor 10.231.64.10 ebgp-multihop 2
  neighbor 10.231.64.10 enforce-first-as
```

```
neighbor 10.231.64.10 local-as 64900
neighbor 10.231.64.10 maxas-limit in disable
neighbor 10.231.64.10 next-hop-self always
neighbor 10.231.64.10 password default
neighbor 10.231.64.10 remote-as 1200
neighbor 10.231.64.10 remove-private-as
neighbor 10.231.64.10 shutdown generate-rib-out
neighbor 10.231.64.10 soft-reconfiguration inbound
neighbor 10.231.64.10 timers keep-alive 120 hold-time 240
neighbor 10.231.64.10 update-source loopback lo0
address-family ipv4 unicast
neighbor 10.231.64.10 default-originate route-map test-map
neighbor 10.231.64.10 filter-list [ 1 ] in
neighbor 10.231.64.10 maximum-prefix 15000 teardown
neighbor 10.231.64.10 prefix-list test-prefix in
neighbor 10.231.64.10 route-map in test-map
neighbor 10.231.64.10 send-community both
neighbor 10.231.64.10 unsuppress-map test-map-2
neighbor 10.231.64.10 weight 10
```

As mentioned above, a set of configurations can be specified for each neighbor, with support for the following:

- Advertisement interval
- Default route origination
- Enforcing of first AS in AS-path list as AS of originator
- AS path filter list
- Enforcing of local ASN of neighbor
- Enabling/disabling of 4-byte ASN capability at the BGP global level
- Maximum AS path length
- Ignoring of AS path lengths of otherwise equal paths
- Maximum routes learned from neighbor
- Enforcing of nexthop as self in routes advertised
- MD5 password authentication
- Prefix list for route filtering
- Remote AS
- Removing of private ASN while advertising routes
- Route map filtering
- Sending community attributes
- Shutting down of neighbor without removing the configuration
- Applying policy changes without resetting neighbor
- Keepalive and hold time
- Specifying of routes not to be suppressed in route aggregation
- Specifying of source IP to be used in TCP connection to neighbor
- Adding of weight to each route received from neighbor

Peer groups

Neighbors having the same attributes and parameters can be grouped together by means of the **peer-group** command. You must first create a peer-group, after which you can associate neighbor IP addresses with the peer-group. All of the attributes that are allowed on a neighbor are allowed on a peer-group as well.

Configurations for both creating a peer-group and associating neighbors to the peer-group are available in BGP global mode. As an example, you can create a peer-group named “external-group” as follows:

```
switch(config-bgp-router)# neighbor external-group peer-group
```

Subsequently, you can associate neighbors to “external-group,” and configure attributes on the peer-group as illustrated below:

```
switch(config-bgp-router)# neighbor 172.29.233.2 peer-group external-group
switch(config-bgp-router)# neighbor 10.120.121.2 peer-group external-group
switch(config-bgp-router)# neighbor external-group remote-as 1720
```

An attribute value configured explicitly for a neighbor takes precedence over the attribute value configured on peer-group. In case neither the peer-group nor the individual neighbor has the attribute configured, the default value for the attribute is used.

Four-byte AS numbers

Four-byte autonomous system numbers (ASNs) can be optionally configured on a device, peer-group, or neighbor. If this is enabled, the device announces and negotiates “AS4” capability with its neighbors. You can configure AS4 capability to be enabled or disabled at the BGP global level:

```
switch(config-bgp-router)# capability as4-enable
```

You can do the same at the neighbor or peer-group level:

```
switch(config-bgp-router)# neighbor 172.29.233.2 capability as4-enable
```

You can configure AS4 capability to be enabled for a neighbor while still keeping AS4 numbers disabled at the global level, or vice-versa. The neighbor AS4 capability configuration takes precedence. If AS4 capability is not configured on the neighbor, then the peer-group configuration takes effect. The global configuration is used if AS4 capability is configured neither at the neighbor nor at the peer-group level. If a device having a 4-byte ASN tries to connect to a device that does not have AS4 support, peering will not be established.

Route redistribution

The redistribution of static, connected, and OSPF routes into BGP is supported. Similarly, routes learnt through BGP can also be redistributed into OSPF. An optional route-map can be specified, and this map will be consulted before routes are added to BGP. Management routes are not redistributed.

You configure redistribution under IPv4 address-family mode:

```
switch(config-bgp-router)# address-family ipv4 unicast
switch(config-bgp-ipv4u)# redistribute ?
```

Possible completions:

connected	Connected
ospf	Open Shortest Path First (OSPF)
static	Static routes

While redistributing routes learned by OSPF, you can specify the type of routes to be redistributed. You can choose to redistribute internal, external type1, or external type2 routes.

The **redistribute ospf** command redistributes internal OSPF routes in a way that is equivalent to the effect of the **redistribute ospf match internal** command, as illustrated in this running-config excerpt:

```
router bgp
  local-as 6500
  address-family ipv4 unicast
    redistribute ospf match internal
    redistribute ospf match external1
    redistribute ospf match external2
```

Advertised networks

As described in the previous section, you can advertise routes into BGP by redistributing static, connected, or OSPF routes. However, you can explicitly specify routes to be advertised by BGP by using the **network** command in IPv4 address-family submode:

```
switch(config-bgp-ipv4u)# network 10.40.25.0/24
```

Before BGP can advertise this route, the routing table must have this route already installed.

Another use of the **network** command is to specify a route to be local. In case the same route is received by means of EBGp, the local IGP route will be preferred. The **backdoor** parameter changes the administrative distance of the route to this network from the EBGp administrative distance (20 by default) to the local BGP4 weight (200 by default), tagging the route as a backdoor route. Use this parameter when you want the device to prefer IGP routes such as RIP or OSPF routes over the EBGp route for the network.

```
switch(config-bgp-ipv4u)# network 10.40.25.0/24 backdoor
```

The **neighbor weight** command specifies a weight that the device adds to routes that are received from the specified BGP neighbor. (BGP4 prefers larger weights over smaller weights.)

Static networks

Before advertising any route, BGP checks for its existence in the routing table. If you want BGP to advertise a stable route that does not depend on its existence in the routing table, then use the **static-network** command to advertise that network:

```
switch(config-bgp-ipv4u)# static-network 10.40.25.0/24
```

When the configured route is lost, BGP installs the “null0” route in the routing table. Later, when the route is resolved, the null0 route is removed. You can override the administrative local distance of 200 by specifying the distance value in the command:

```
switch(config-bgp-ipv4u)# static-network 10.40.25.0/24 distance 300
```

Route reflection

A BGP device can act as a route-reflector client or as a route reflector. You can configure a BGP peer as a route-reflector client from the device that is going to reflect the routes and act as the route reflector:

```
switch(config-bgp-ipv4u)# neighbor 10.61.233.2 route-reflector-client
```

When there is more than one route-reflector, they should all belong to the same cluster. By default, the value for **cluster-id** is used as the device ID. However, the device ID can be changed:

```
switch(config-bgp-router)# cluster-id ipv4-address 10.30.13.4  
switch(config-bgp-router)# cluster-id 2300
```

The route-reflector server reflects the routes as follows:

- Routes from the client are reflected to client as well as to nonclient peers.
- Routes from nonclient peers are reflected only to client peers.

In case route-reflector clients are connected in a full IBGP mesh, you may wish to disable client-to-client reflection on the route reflector:

```
switch(config-bgp-ipv4u)# no client-to-client-reflection
```

A BGP device advertises only those routes that are preferred ones and are installed into the Routing Table Manager (RTM). When a route could not be installed into the RTM because the routing table was full, the route reflector may not reflect that route. In case the route reflector is not placed directly in the forwarding path, you can configure the route reflector to reflect routes even though those routes are not in the RTM:

```
switch(config-bgp-ipv4u)# always-propagate
```

Route flap dampening

Unstable routes can trigger a lot of route state changes. You can configure dampening in IPv4 address-family mode to avoid this churn by penalizing the unstable routes:

```
switch(config-bgp-ipv4u)# dampening
```

The above command uses default values for the dampening parameters described below:

half-life

Number of minutes after which penalty for a route becomes half of its value. The route penalty allows routes that have remained stable for a period despite earlier instability to become eligible for reuse after the interval configured by this parameter. The decay rate of the penalty is proportional to the value of the penalty. After the half-life expires, the penalty decays to half its value. A dampened route that is no longer unstable can eventually again become eligible for use. You can configure the half-life to be from 1 through 45 minutes. The default is 15 minutes.

reuse

Minimum penalty below which routes becomes reusable again. You can set the reuse threshold to a value from 1 through 20000. The default is 750 (0.75, or three-fourths, of the penalty assessed for one flap).

suppress

Maximum penalty above which route is suppressed by the device. You can set the suppression threshold to a value from 1 through 20000. The default is 2000 (more than two flaps).

max-suppress-time

Maximum number of minutes a route can be suppressed by the device, regardless of how unstable the route is. You can set maximum suppress time to a value from 1 through 255 minutes. The default is 40 minutes.

NOTE

A dampening value for half-life can also be adjusted through a route map, by means of the **set dampening** option in the **route-map** command.

Default route origination

While redistributing routes from OSPF, BGP does not advertise default route 0.0.0.0/0 even if it exists. You can enable default route origination by using the **default-information-originate** command under IPv4 address-family mode:

```
switch(config-bgp-ipv4u)# default-information-originate
```

In order to advertise a default route without OSPF redistribution, use the **network** command.

Multipath load sharing

Unlike IGP, BGP does not perform multipath load sharing by default. Therefore, the maximum number of paths across which BGP can balance the traffic is set to 1 by default. You can change this value by using the **maximum-paths** command in IPv4 address-family mode:

```
switch(config-bgp-ipv4u)# maximum-paths 4
```

You can specify the maximum path value explicitly for IBGP or EBGP:

```
switch(config-bgp-ipv4u)# maximum-paths ebgp 2
```

In case the maximum-path value must be picked up from the **ip load-sharing** configuration on the router, use the following:

```
sw0(config-bgp-ipv4u)# maximum-paths use-load-sharing
```

You can enable multipathing for both IBGP or EBGP. By default, the AS numbers of the two paths should match for them to be considered for multipathing. However, you can remove this restriction by specifying a **multi-as** option, as illustrated in the following running-config excerpt:

```
router bgp
  local-as 6500
  address-family ipv4 unicast
    multipath ebgp
    multipath ibgp
    multipath multi-as
```

Configuring the default route as a valid next-hop

BGP does not use the default route installed in the device to resolve the BGP next-hop. You can change this in IPv4 address-family mode to enable BGP to use default route for next-hop resolution:

```
sw0(config-bgp-ipv4u)# next-hop-enable-default
```

Next-hop recursion

Next-hop recursion is disabled by default in BGP, but you can enable it. When next-hop recursion is disabled, only one route lookup is performed to obtain the next-hop IP address. If the lookup result does not succeed or the result points to another BGP path, then route destination is considered unreachable. Enable it as follows:

```
sw0(config-bgp-ipv4u)# next-hop-recursion
```

When next-hop recursion is enabled, if the first lookup for the destination IP address results in an IBGP path that originated in the same AS, the device performs lookup for the IP address of the next-hop gateway. This goes on until the final lookup results in an IGP route. Otherwise, the route is declared unreachable.

Route filtering

The following route filters are supported:

- AS-path filter
- Community filter
- Prefix list
- Route map
- Table map

NOTE

Support for access lists in route filtering is not available, and has been replaced by prefix-list filtering. BGP does not use community and extended-community filters directly. Rather, it uses them indirectly through route-map filtering by means of the **route-map** command.

Timers

You can change keepalive and hold-time values from their default values of 60 and 180 seconds, respectively:

```
sw0(config-bgp-router)# timers keep-alive 10 hold-time 60
```

A hold-time value of 0 means that the device will wait indefinitely for messages from a neighbor without tearing down the session.

Once the IGP routes are changed, BGP routing tables are affected after 5 seconds by default. You can change this value by using the **update-time** command:

```
sw0(config-bgp-ipv4u)# update-time 0
```

An **update-time** value of 0 will trigger BGP route calculation immediately after the IGP routes are changed.

Using route maps

A route map is a named set of match conditions and parameter settings that the device can use to modify route attributes and to control redistribution of the routes into other protocols. A route map consists of a sequence of instances, the equivalent of rows in a table. The device evaluates a route according to route map instances in ascending numerical order. The route is first compared against instance 1, then against instance 2, and so on. When a match is found, the device stops evaluating the route.

Route maps can contain **match** clauses and **set** statements. Each route map contains a **permit** or **deny** statement for routes that match the **match** clauses:

- If the route map contains a **permit** statement, a route that matches a match statement is permitted; otherwise, the route is denied.
- If the route map contains a **deny** statement, a route that matches a match statement is denied.
- If a route does not match any **match** statements in the route map, then the route is denied. This is the default action. To change the default action, configure the last **match** statement in the last instance of the route map to **permit any any**.
- If there is no **match** statement, the software considers the route to be a match.
- For route maps that contain address filters, AS-path filters, or community filters, if the action specified by a filter conflicts with the action specified by the route map, the route map action takes precedence over the filter action.

If the route map contains **set** statements, routes that are permitted by the route map **match** statements are modified according to the **set** statements.

Match statements compare the route against one or more of the following:

- The route BGP4 MED (metric)
- The IP address of the next hop device
- The route tag
- For OSPF routes only, the route type (internal, external type 1, or external type 2)
- An AS-path access control list (ACL)
- A community ACL
- An IP prefix list

For routes that match all of the **match** statements, the route map **set** statements can perform one or more of the following modifications to the route attributes:

- Prepend AS numbers to the front of the route AS-path. By adding AS numbers to the AS-path, you can cause the route to be less preferred when compared to other routes based on the length of the AS-path.
- Add a user-defined tag or an automatically calculated tag to the route.
- Set the community attributes.
- Set the local preference.
- Set the MED (metric).
- Set the IP address of the next-hop device.
- Set the origin to IGP or INCOMPLETE.

- Set the weight.

When you configure parameters for redistributing routes into BGP4, one of the optional parameters is a route map. If you specify a route map as one of the redistribution parameters, the device matches the route against the match statements in the route map. If a match is found and if the route map contains **set** statements, the device sets the attributes in the route according to the set statements.

To create a route map, you define instances of the map by a sequence number.

The route-map *name* is a string of characters that defines the map instance. Map names can be up to 80 characters long. The following is the complete syntax of the **route-map** command:

```
[no] route-map name [permit | deny] instance_number ]
[continue sequence number?? ]
match [as-path ACL_name |
      community ACLname |
      interface [fortygigabitethernet rbridge-id/slot/port gigabitethernet rbridge-id/slot/port |
                loopback port tengigabitethernet rbridge-id/slot/port ve port] |
      ip [address prefix-list name] | next-hop prefix-list name | route-source prefix-list name] |
      metric number |
      protocol bgp [external | internal | static-network] |
      route-type [internal | type-1 | type-2] |
      tag number |
set [as-path [prepend | tag] |
    automatic-tag |
    comm-list |
    community community_number | additive | local-as | no-advertise | no-export | none] |
    dampening number |
    distance number ]
    ip next-hop [A.B.C.D | peer-address] |
    local-preference number] |
    metric [add number | assign | none | sub] |
    metric-type [external | internal | type-1 | type-2] |
    origin [igp | incomplete] |
    route-type [internal | type-1 | type-2] |
    tag number |
    weight number]
```

Operands for this command are defined below.

The **permit** | **deny** options specify the action the device will take if a route matches a match statement:

- If you specify **deny**, the device does not advertise or learn the route.
- If you specify **permit**, the device applies the match and set clauses associated with this route map instance.

The *instance-number* parameter specifies the instance of the route map you are defining. The following illustrates a creation of a route-map instance 10, which is done in RBridge ID mode. Note the change in the command prompt.

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# routemap myroutemap1 permit 10
switch(config-route-map-myroutemap1/permit/10)#
```

To delete a route map, enter a command such as the following in RBridge ID configuration mode. When you delete a route map, all the permit and deny entries in the route map are deleted.

```
switch(config-rbridge-id-5)# no route-map myroutemap1
```

This command deletes a route map named myroutemap1. All entries in the route map are deleted.

To delete a specific instance of a route map without deleting the rest of the route map, enter a command such as the following.

```
switch(config-rbridge-id-5)# no route-map myroutemap1 permit 10
```

This command deletes the specified instance from the route map but leaves the other instances of the route map intact.

Specifying the match conditions

Use the **match** command to define the match conditions for instance 10 of the route map myroutemap1.

```
switch(config-route-map-myroutemap1/permit/10)# match ?
```

Operands for the route-map **match** statement are as follows:

community *num*

Specifies a community ACL. (The ACL must already be configured.)

ip address | **next-hop** *acl-num* | **prefix-list** *string*

Specifies an ACL or IP prefix list. Use this parameter to match based on the destination network or next-hop gateway. To configure an IP ACL for use with this command, use the **ip access-list** command. To configure an IP prefix list, use the **ip prefix-list** command.

ip route-source *acl* | **prefix** *name*

Matches on the source of a route (the IP address of the neighbor from which the device learned the route).

metric *num*

Compares the route MED (metric) to the specified value.

next-hop *address-filter-list*

Compares the IP address of the route next-hop to the specified IP address filters. The filters must already be configured.

route-type [**internal** | **type-1** | **type-2**]

Applies only to OSPF routes.

- **internal** sets an internal route type.
- **type-1** sets an OSPF external route type 1.
- **type-2** sets an OSPF external route type 2.

tag *tag-value*

Compares the route tag to the specified tag value.

protocol bgp static-network

Matches on BGP4 static-network routes.

protocol bgp external

Matches on EBGP (external) routes.

protocol bgp internal

Matches on IBGP (internal) routes.

Setting parameters in the routes

Use the following command to define a **set** statement that prepends an AS number to the AS path on each route that matches the corresponding match statement.

```
switch(config-routemap-myroutemap1/permit10)# set as-path prepend 7701000
```

Operands for the route-map **set** statement are as follows:

as-path prepend *num,num,...*

Adds the specified AS numbers to the front of the AS-path list for the route. Values range from 1 through 65535 for two-byte ASNs, and from 1 through 4294967295 if AS4s have been enabled.

automatic-tag

Calculates and sets an automatic tag value for the route. (This parameter applies only to routes redistributed into OSPF.)

comm-list

Deletes a community from the community attributes field for a BGP4 route.

community

Sets the community attribute for the route to the number or well-known type specified.

dampening [*half-life*]

Sets route dampening parameters for the route; *half-life* specifies the number of minutes after which the route penalty becomes half its value.

ip next hop *ip-addr*

Sets the next-hop IP address for a route that matches a **match** statement in the route map.

ip next-hop peer-address

Sets the BGP4 next hop for a route to the neighbor address.

local-preference *num*

Sets the local preference for the route. Values range from 0 through 4294967295.

metric [+ | -] *num* | **none**

Sets the Multi-Exit Discriminator (MED) value for the route. Values range from 0 through 4294967295. The default is 0.

- **set metric** *num* sets the metric for the route to the specified number.
- **set metric +** *num* increases the route metric by the specified number.
- **set metric -** *num* decreases route metric by the specified number.
- **set metric none** removes the metric from the route (removes the MED attribute from the BGP4 route).

metric-type *type-1* | *type-2*

Changes the metric type of a route redistributed into OSPF.

metric-type internal

Sets the route MED to the same value as the IGP metric of the BGP4 next-hop route, for advertising a BGP4 route to an EBGP neighbor.

next-hop ip-addr

Sets the IP address of the next-hop device.

origin igp incomplete

Sets the route's origin to IGP or INCOMPLETE.

tag

Keyword that sets the tag to be an AS-path attribute. (This parameter applies only to routes redistributed into OSPF.)

weight num

Sets the weight for the route. Values range from 0 through 65535.

BGP configuration examples

This section expands upon the preceding overview and provides additional BGP management examples:

- [Adjusting defaults to improve routing performance](#)
- [Using route maps with match and set statements](#)
- [Using route-map continue statements](#)
- [Clearing configurations](#)

Adjusting defaults to improve routing performance

The following examples illustrate a variety of options for enabling and fine-tuning route flap dampening.

To enable default dampening as an address-family function:

```
switch(config)# rbridge-id 10
switch(config-rbridge-id-10)# router bgp
switch(config-bgp-router)# address-family ipv4 unicast
switch(config-bgp-ipv4u)# dampening
```

To change the all dampening values as an address-family function:

```
switch(config)# rbridge-id 10
switch(config-rbridge-id-10)# router bgp
switch(config-bgp-router)# address-family ipv4 unicast
switch(config-bgp-ipv4u)# dampening 20 200 2500 40
```

NOTE

To change any of the parameters, you must specify all the parameters in the command in the following order: half-life, reuse, suppress, mas-suppress-time. To leave any parameters unchanged, enter their default values.

For more details about the use of route maps, including more flap-dampening options, see the following.

Using route maps with match and set statements

The following examples illustrate using route maps with match and set statements.

Matching on an AS-path ACL

To configure a route map that matches on AS-path ACL 1:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# route-map myaclroutemap1 permit 10
switch(config-route-map-myaclroutemap1/permit/10)# match as-path 1
```

Matching on a community ACL

To configure a route map that matches on community ACL 1:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# ip community-list standard 1 permit 123:2
switch(config-rbridge-id-5)# route-map mycommroutemap1 permit 10
switch(config-route-map-mycommroutemap1/permit/10)# match community 1
```

*Matching on a destination network***NOTE**

You can use the results of an IP ACL or an IP prefix list as the match condition.

To configure a route map that matches on a destination network:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# route-map mynetroutemap1 permit 10
switch(config-route-map-mynetroutemap1/permit/10)# match ip address prefix-list 1
```

Matching on a next-hop device

To configure a route map that matches on a next-hop device:

```
switch(config)# rbridge-id-5
switch(config-rbridge-id-5)# route-map myhoproutemap1 permit 10
switch(config-route-map-myhoproutemap1/permit/10)# match ip next-hop
prefix-list 1
```

Matching on a route source

To configure a route map that matches on a route source:

```
switch(config)# rbridge-id 5
switch(config)# access-list 10 permit 192.168.6.0 0.0.0.255
switch(config-rbridge-id-5)# route-map mysourceroutemap1 permit 1
switch(config-route-map-mysourceroutemap1/permit/10)# match ip route-source
prefix-list 10
```

Matching on routes containing a specific set of communities

To configure a route map that matches on a set of communities:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# ip community-list standard std_1 permit 12:34
no-export
switch(config-rbridge-id-5)# route-map mycommroutemap2 permit 1
switch(config-routemap-mycommroutemap2/permit/1)# match community std_1
exact-match
```

NOTE

The first command configures a community ACL that contains community number 12:34 and community name “no-export.” The remaining commands configure a route map that matches the community attributes field in BGP4 routes against the set of communities in the ACL. A route matches the route map only if the route contains all the communities in the ACL and no other communities.

To configure an additional community-based route map for comparison with the first:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# ip community-list standard std_2 permit 23:45 56:78
switch(config-rbridge-id-5)# route-map mycommroutemap3 permit 1
switch(config-routemap-mycommroutemap3/permit/1)# match community std_1 std_2
exact-match
```

NOTE

These commands configure an additional community ACL, `std_2`, that contains community numbers 23:45 and 57:68. Route map `mycommroutemap3` compares each BGP4 route against the sets of communities in ACLs `std_1` and `std_2`. A BGP4 route that contains *either but not both* sets of communities matches the route map. For example, a route containing communities 23:45 and 57:68 matches. However, a route containing communities 23:45, 57:68 and 12:34, or communities 23:45, 57:68, 12:34, and “no-export” does not match. To match, the route communities must be the same as those in exactly one of the community ACLs used by the **match community** statement.

Matching on a BGP4 static network

To configure a route map that matches on a static network:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# route-map mystaticroutemap3 permit 1
switch(config-routemap-mystaticroutemap3/permit/1)# match protocol bgp
static-network
switch(config-routemap-mystaticroutemap3/permit/1)# set local-preference 150
switch(config-routemap-mystaticroutemap3/permit/1)# set community no-export
switch(config-routemap-mystaticroutemap3/permit/1)# exit
switch(config)# router bgp
switch(config-bgp)# neighbor 192.168.6.0 route-map out mystaticroutemap3
```

Matching on an interface

To configure a route map that matches on an interface:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# route-map myintroutemap1 permit 99
switch(config-rbridge-id-5)# match interface ten 5/1/1 ten 5/3/2
```

Setting a BGP4 route MED to equal the next-hop route IGP metric

To set a route's Multi-Exit Discriminator (MED) to the same value as the IGP metric of the BGP4 next-hop route, when advertising the route to a neighbor:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# route-map mymedroutemap1 permit 1
switch(config-routemap-mymedroutemap/permit/1)# match ip address 1
switch(config-routemap-mymedroutemap/permit/1)# set metric-type internal
```

Setting the next-hop of a BGP4 route

To set the next-hop address of a BGP4 route to a neighbor address:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# route-map mhoproutemap1 permit 1
switch(config-routemap-myhoproutemap/permit/1)# match ip address 1
switch(config-routemap-myhoproutemap/permit/1)# set ip next-hop peer-address
```

Deleting a community from a BGP4 route

To delete a community from a BGP4 route's community attributes field:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# ip community-list standard std_3 permit 12:99 12:86
```

```
switch(config-rbridge-id-5)# route-map mcommroutemap1 permit 1
switch(config-routemap-mycommroutemap/permit/1)# match ip address 1
switch(config-routemap-mycommroutemap/permit/1)# set comm-list std_3 delete
```

NOTE

The first command configures a community ACL containing community numbers 12:99 and 12:86. The remaining commands configure a route map that matches on routes whose destination network is specified in ACL 1, and deletes communities 12:99 and 12:86 from those routes. The route does not need to contain all the specified communities in order for them to be deleted. For example, if a route contains communities 12:86, 33:44, and 66:77, community 12:86 is deleted.

Using route-map continue statements

To configure **continue** statements in a route map:

```
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# route-map mcontroutemap1 permit 1
switch(config-routemap-mycontroutemap/permit/1)# match metric 10
switch(config-routemap-mycontroutemap/permit/1)# set weight 10
switch(config-routemap-mycontroutemap/permit/1)# match metric 10
switch(config-routemap-mycontroutemap/permit/1)# continue 2
switch(config-routemap-mycontroutemap/permit/1)# route-map mcontroutemap1
    permit 2
switch(config-routemap-mycontroutemap/permit/2)# match tag 10
switch(config-routemap-mycontroutemap/permit/2)# set weight 20
```

NOTE

This configures the route map to continue to evaluate and execute **match** statements after a successful match occurs. The **continue** statement proceeds to the route map with the specified sequence number. If no sequence number is specified, the statement proceeds to the route map with the next sequence number (an “implied” continue).

Using a route map to configure dampening

To apply the dampening half-life established in a route map:

```
switch(config)# rbridge-id 10
switch(config-rbridge-id-10)# route-map myroutemap permit 1
switch(config-route-map-myroutemap/permit/1)# set dampening 20
```

NOTE

To change any of the parameters, you must specify all the parameters with the command. To leave any parameters unchanged, enter their default values.

Clearing configurations

To refresh all BGP4 neighbor routes:

```
switch# clear ip bgp neighbor all
```

To refresh a route to a specific neighbor:

```
switch# clear ip bgp neighbor 10.11.12.13
```


To clear BGP4 routes:

```
switch# clear ip bgp routes 10.0.0.0/16
```

To clear the BGP4 message counters:

```
switch# clear ip bgp traffic
```

To unsuppress all suppressed BGP4 routes:

```
switch# clear ip bgp dampening
```

To clear the dampening statistics for a BGP4 route:

```
switch# clear ip bgp flap-statistics 10.0.0.0/16
```

41 BGP configuration examples

Configuring IGMP

In this chapter

• IGMP overview	587
• Configuring IGMP snooping	588
• Configuring IGMP snooping querier	589
• Monitoring IGMP snooping	590
• IGMP scalability	590

IGMP overview

The forwarding of multicast control packets and data through a Layer 2 switch configured with VLANs is most easily achieved by the Layer 2 forwarding of received multicast packets on all the member ports of the VLAN interfaces. However, this simple approach is not bandwidth efficient, because only a subset of member ports may be connected to devices interested in receiving those multicast packets. In a worst-case scenario, the data would get forwarded to all port members of a VLAN with a large number of member ports (for example, all 24 ports), even if only a single VLAN member is interested in receiving the data. Such scenarios can lead to loss of throughput for a switch that gets hit by a high rate of multicast data traffic.

Internet Group Management Protocol (IGMP) snooping is a mechanism by which a Layer 2 switch can effectively address this issue of inefficient multicast forwarding to VLAN port members. Snooping involves “learning” forwarding states for multicast data traffic on VLAN port members from the IGMP control (join/leave) packets received on them. The Layer 2 switch also provides for a way to configure forwarding states statically through the CLI.

Multicast routing

Multicast routers use IGMP to learn which groups have members on each of their attached physical networks. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership.

NOTE

“Multicast group memberships” means that at least one member of a multicast group on a given attached network is available.

There are two ways that hosts join multicast routing groups:

- By sending an unsolicited IGMP join request.
- By sending an IGMP join request as a response to a general query from a multicast router.

In response to the request, the switch creates an entry in its Layer 2 forwarding table for that VLAN. When other hosts send join requests for the same multicast, the switch adds them to the existing table entry. Only one entry is created per VLAN in the Layer 2 forwarding table for each multicast group.

vLAG and LAG primary port

The current data center Ethernet (DCE) implementation of vLAGs and LAGs uses the concept of a so-called primary port. One of the member ports of the vLAG and LAG is selected to be the primary port, and all multicast traffic egressing from the LAG or vLAG is sent on the primary port. Thus, normal hash-based forwarding is not performed for multicast traffic, whether it is control traffic or data. Now, consider the case where RBridge R1 receives an IGMP join request for group G1 on Po10, shown in Figure 61. This causes Po10 to be added to the list of IGMP receivers for group G1. Now, assume that the primary port of the vLAG is the link connecting R4 and S1. Therefore, any multicast traffic received by the cluster for group G1 egresses on vLAG Po10 from R4 and not from R1, even though the original join was received on R1.

If the primary port for the vLAG changes, such as if the link between R4 and S1 in Figure 61 went down, then multicast traffic would egress out of the new primary port on the vLAG. In the above case, the new primary port would be the link connecting R1 and S1.




FIGURE 61 IGMP snooping in Brocade VCS Fabric mode

Configuring IGMP snooping

By default, IGMP snooping is globally disabled on all VLAN interfaces. Refer to the *CEE Command Reference* for complete information about the commands in this section.

Use the following procedure to configure IGMP on a Data Center Bridging (DCB)/Fibre Channel over Ethernet (FCoE) switch.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **ip igmp snooping enable** command to enable IGMP for all interfaces.

This command ensures that IGMP snooping is active on all interfaces.

```
switch(config)# ip igmp snooping enable
```

3. Enter the **interface** command to select the VLAN interface number.

```
switch(config)# interface vlan 10
```

4. Activate the default IGMP querier functionality for the VLAN.

```
switch(config-vlan-10)# ip igmp snooping querier enable
```

5. *Optional:* Activate the IGMP querier functionality with additional features.

NOTE

The IGMP snooping configuration must be the same for all switches in the same VFC Fabric cluster. For example, if you configure **ip igmp snooping enable** on VLAN 2 on a VDX switch, you must configure that same command on VLAN 2 on all VDX switches in the cluster.

Configuring IGMP snooping querier

If your multicast traffic is not routed because Protocol-Independent Multicast (PIM) and IGMP are not configured, use the IGMP snooping querier in a VLAN.

IGMP snooping querier sends out IGMP queries to trigger IGMP responses from switches that are to receive IP multicast traffic. IGMP snooping listens for these responses to map the appropriate forwarding addresses.

NOTE

An IGMP snooping querier cannot be configured on the same interface as a multicast router (mrouter) interface.

Refer to the *CEE Command Reference* for complete information about the commands in this section.

Use the following procedure to configure the IGMP snooping querier.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to select the VLAN interface number.

```
switch(config)# interface vlan 25
```

3. Activate IGMP querier functionality for the VLAN.

The valid range is 1 to 18000 seconds. The default is 125 seconds.

```
switch(config-vlan-25)# ip igmp query-interval 125
```

4. Set the last member query interval.

The valid range is 1000 to 25500 milliseconds. The default is 1000 milliseconds.

```
switch(config-vlan-25)# ip igmp last-member-query-interval 1000
```

5. Set the Max Response Time (MRT).

The valid range is 1 to 25 seconds. The default is 10 seconds.

```
switch(config-vlan-25)# ip igmp query-max-response-time 10
```

6. Activate the IGMP querier functionality for the VLAN.

```
switch(config-vlan-25)# ip igmp snooping querier enable
```

Monitoring IGMP snooping

Monitoring the performance of your IGMP traffic allows you to diagnose any potential issues on your switch. This helps you utilize bandwidth more efficiently by setting the switch to forward IP multicast traffic only to connected hosts that request multicast traffic.

Refer to the *CEE Command Reference* for complete information about the commands in this section.

Use the following procedure to monitor IGMP snooping on a DCB/FCoE switch.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **show ip igmp groups** command to display all information on IGMP multicast groups for the switch.

Use this command to display the IGMP database, including configured entries for either all groups on all interfaces, all groups on specific interfaces, or specific groups on specific interfaces.

```
switch# show ip igmp groups
```

3. Use the **show ip igmp statistics** command to display the IGMP statistics for a VLAN or interface.

```
switch# show ip igmp snooping statistics interface vlan 10
```

4. Use the **show ip igmp mrouter** to display mrouter port-related information for all VLANs, or a specific VLAN.

```
switch# show ip igmp snooping mrouter
```

- or -

```
switch# show ip igmp snooping mrouter interface vlan 10
```

5. When you have reviewed the IGMP statistics for the switch, refer to [“Configuring IGMP snooping”](#) on page 588 or [“Configuring IGMP snooping querier”](#) on page 589 to make any needed corrections.

NOTE

Refer to the *CEE Command Reference* for additional information on IGMP CLI commands.

IGMP scalability

This section describes the scalability limits of IGMP snooping feature for Network OS v4.0.0 in various modes of switch operation, and explains the various metrics involved in describing the scalability limits.

Note the following IGMP metrics:

- Maximum number of IGMP groups supported—This metric is based on the available hardware resources, such as multicast group ID (MGID), configuration replay, and Ethernet Name Server (eNS) distribution bandwidth.

- Maximum number of VLANs supported with IGMP snooping configuration—This metric is limited by the general-query packet-generation capacity of IGMP software processes running on the switch, as well as by eNS distribution bandwidth.
- Maximum IGMP packet-processing rate per switch—The scalability number described by this metric suggests the upper limit on the number of packets that can be processed by IGMP software processes running on the switch. If the packets are incoming from multiple ports/VLANs, the same processing bandwidth is shared.
- Maximum IGMP packet-processing rate per Brocade VCS Fabric cluster—This metric specifies the upper limit on the maximum rate of IGMP packets incoming to a logical Brocade VCS Fabric switch. It is limited by the eNS distribution bandwidth and the number of nodes in the Brocade VCS Fabric cluster.

Standalone mode

In standalone mode, the VDX switch functions as an isolated device, possibly connected as a top-of-rack (TOR) switch. [Table 85](#) details the metric levels.

TABLE 85 Standalone mode metrics

Metric	Limit	Comments
Maximum number of IGMP groups supported	2000	Join requests are sent on four ports of the same switch.
Maximum number of VLANs supported with IGMP configuration	128	
Maximum IGMP packet-processing rate per switch	512 packets/sec	

Brocade VCS Fabric cluster mode

When supporting a flat Layer 2 network in a data center, VDX switches can be connected in any order to form a cluster. The number of nodes involved in a cluster ranges from four nodes to 24 nodes.

- [Table 86](#) and [Table 87](#) detail metric levels for four-node and 24-node clusters, respectively.
- [Table 87](#) details metric levels for topologies that use the Brocade VDX 8770-4 and VDX 8770-8.
- [Table 89](#) details metric levels for IP multicast.

TABLE 86 Four-node cluster metrics

Metric	Limit	Comments
Maximum number of IGMP groups supported	2000	Join requests are sent on four ports of the same switch.
Maximum number of VLANs supported with IGMP configuration	128	
Maximum IGMP packet-processing rate per switch	512 packets/sec	
Maximum IGMP packet-processing rate per Brocade VCS Fabric cluster	512 packets/sec	

TABLE 87 Twenty-four-node cluster metrics

Metric	Limit	Comments
Maximum number of IGMP groups supported	2000	Join requests are sent on four ports of the same switch.
Maximum number of VLANs supported with IGMP configuration	128	
Maximum IGMP packet-processing rate per switch	512 packets/sec	
Maximum IGMP packet processing rate per Brocade VCS Fabric cluster	512 packets/sec	

TABLE 88 VDX 8770-4 and VDX 8770-8 cluster metrics

Metric	Limit	Comments
Maximum number of IGMP groups supported	8000	Join requests are sent on four ports of the same switch.
Maximum number of VLANs supported with IGMP configuration	256	
Maximum IGMP packet processing rate per switch	512 packet/second	
Maximum IGMP packet processing rate per Brocade VCS Fabric cluster	512 packet/second	

TABLE 89 IP multicast metrics

Metric	Limit	Comments
Number of Layer 3 forwarding entries	256	
Number of IGMP snooping forwarding entries	8000	
Number of multicast flows	10000	
PIM interfaces supported	32	
IGMP interfaces supported	32	
IGMP snooping interfaces supported	256	
Learning rate for PIM-SM	32 flows/second	
Learning rate for IGMP snooping	512 groups/second	

Troubleshooting

This section describes troubleshooting information, and includes the following chapters:

- [Using the Chassis ID \(CID\) Recovery Tool 595](#)
- [Troubleshooting 599](#)

Using the Chassis ID (CID) Recovery Tool

In this chapter

- Chassis ID card usage 595
- Automatic auditing and verification of CID card data 596
- Running the CID recovery tool 596

Chassis ID card usage

Each Brocade VDX 8770-4 and VDX 8770-8 contains two chassis ID cards (CIDs) called *CD1* and *CD2*. Most data on each card is identical, and CID2 is used only as a backup if CID1 encounters an issue.

The data contained on the CID card is essential for correct operation of the switch and is accessed most frequently during system startup.

Each CID contains two serial electronically erasable programmable read-only memory (SEEPR0M) devices:

- Critical SEEPR0M. This SEEPR0M is read-only.
- Noncritical SEEPR0M. This SEEPR0M can be written to by the software.

Critical SEEPR0M data

The critical SEEPR0M contains [the following](#):

- A header with the CID part number, serial number, and other data about the CID card. If this data is corrupted or cannot be accessed, the card is identified as faulty in RASLogs:
 - [EM-1003], M1 | FFDC, CRITICAL,..., CID 2 has unknown hardware identifier: FRU faulted.
 - [FW-1432], M1, WARNING, sw0, Switch status change contributing factor Cid-Card: 1 bad.
- A chassis part number and serial number. Cluster configuration management uses the serial number to uniquely identify the chassis in the fabric.
- An eight-byte number that represents both the license ID and World Wide Name (WWN) base value for the chassis. The license ID is used to validate installed licenses. Licenses are invalid if the license ID is not available. The WWN is used to identify the switch in a fabric.

Non-Critical SEEPROM data

The non-critical SEEPROM contains the following data sets.

- The FRU history table, which contains logs of insertions and removals of FRUs into and from the chassis. The content of this table is not audited or verified.
- The IP data table, which contains management module and chassis management IP addresses/masks, the IP default gateway, and the chassis name.
- A power-off list, which controls the order in which blades are automatically powered off if an impending power loss is detected.
- A set of Data Center Ethernet (DCE) data containing chassis MAC addresses, without which the switch will not function.

Automatic auditing and verification of CID card data

The contents of both CID cards are verified on a periodic basis and whenever an event indicates that an issue may exist.

Under normal circumstances, the CID card audit is run for about one hour after a system startup or restart, then repeated every 24 hours. If no errors occur, no action is taken.

If CID card errors occur, if mismatches between data sets on the two CID cards are detected, or if a card is inserted, RASLogs are shown on the console:

- [EM-1020],... M1, ERROR... A problem was found on one or both CID cards (x), please run the *cidrecov* tool to get more information and recovery options.
- [EM-1021],... M1, INFO,... A CID card has been inserted, a CID verification audit will be run to detect any mismatches or other problems.
- [EM-1022],... M1, WARNING,... A CID card access problem has been encountered, please run the *cidrecov* tool to get more information and recovery options.

Running the CID recovery tool

You should run the CID recovery tool when instructed by RASLogs, and you can also run the tool if you suspect an issue with one or both of the CID cards. To run the CID recovery tool, enter the **cidrecov** command in privileged EXEC mode on the NOS command line:

```
sw0# cidrecov
```

Data corruption or mismatches

If *cidrecov* detects any CID 1 or CID 2 non-critical SEEPROM corruption or mismatches, the tool displays related data and the following data-recovery options as applicable for each data-set error:

- Exit. Select this option if you do not want to change any data values.
- Recover with default values. Select this option if you want to reset all data in the data set to the factory defaults. For IP data, dummy IP addresses and masks are written. DCE and chassis-configuration data are based on the chassis type.

A system restart repopulates IP addresses and chassis names that appear in the startup configuration file. If you want to manually change the IP data, you can use the **ip-address**, **chassis virtual-ip** and **chassis-name** commands. For more information, refer to the *Network OS Command Reference*.

- “Recover BAD from GOOD.” This option is offered only if one CID card contains good data and the other card contains corrupt data. If you select this option, *cidrecov* copies the good data onto the affected card.
- “Recover CID 2 from CID 1” and “Recover CID 1 from CID 2.” These options are offered only if the data on both CID cards is good but there is a mismatch. You can select which card to use to overwrite data on the other card.

Example

The following is an example of running the *cidrecov* tool, receiving errors that can be fixed, and selecting the “Recover BAD from GOOD” option (note that the example below contains only some of the actual output):

```
sw0# cidrecov

CID 1 Non-Critical Seeprom is Inaccessible or Corrupted.

      CID Non-Critical Seeprom Problem Details

CID 1 Non-Critical Seeprom IP address Control Data Checksum Bad !!!!

      CID Recovery Options

0. Exit

1. Recover with default values

2. Recover BAD from GOOD

Enter Selection > 2
Copy IP Data table...
  Copy 384 bytes from CID 2 to CID 1, num blks 1 resid 128
  Read block 1 from CID 2 succeeded
  Write block 1 to CID 1 succeeded
  Read last block from CID 2 succeeded
  Write last block to CID 1 succeeded
  copy successful

Copy succeeded for all data types attempted
IP Address CID Recovery completed.
```

CID card failure

If the critical SEEPROM of a CID card contains *any* errors, or if the non-critical SEEPROM cannot be read, then recovery is not possible, and the following message is displayed:

```
Recovery is not possible. Please contact Brocade Technical Support for
replacement of the inaccessible CID(s).
```

43 Running the CID recovery tool

Troubleshooting

In this chapter

- Troubleshooting overview 599
- Gathering troubleshooting information 599
- Understanding troubleshooting hotspots 601
- Troubleshooting procedures 610
- Troubleshooting and diagnostic tools 643

Troubleshooting overview

This chapter provides tips and procedures for troubleshooting issues that may occur while operating a Brocade switch running Network OS. It also introduces some of the common troubleshooting tools.

Gathering troubleshooting information

The following information is helpful for incident investigation and resolution when you contact your switch-support provider:

- A network diagram and topology information
- A record of the steps and events leading to the incident
- Lists of applications, management agents, and scripts running at the time of the incident
- Supportsave files
- Output from the **show media** command if the issue is related to SFP transceivers
- Outputs from any commands run while attempting to troubleshoot the problem yourself
- Any network traces captured using Wireshark software or other network analyzer.
- Terminal Access Controller Access-Control System (TACACS) server version if the issue is related to TACACS.

Capturing supportsave data

The **copy support** command not only runs diagnostic commands, but also gathers core dumps, trace files, and other relevant data. In the same action, the command also copies all this information to a remote host. Once on the remote host, your switch provider can proceed to analyze the problem. Meanwhile, your switch can be returned to production with minimal downtime.

To capture supportsave data, complete the following steps:

1. Log in to the switch.
2. In privileged EXEC mode, enter the **copy support** command to capture the supportSave data.

The **copy support** command has options to copy the supportSave files to a remote server using FTP or SCP, or you can save to a local USB device. You can use the command in a single command line, or in interactive mode.

The following example uses the single command line mode to copy the supportSave files to a remote host using FTP.

```
switch# copy support ftp host 10.38.33.131 user admin directory 108
Password: *****
```

The following example uses the interactive form of the command and FTP:

```
switch# copy support-interactive
Server Name or IP Address: 10.38.33.131
Protocol (ftp, scp): ftp
User: admin
Password: *****
Directory: /home/admin/support
VCS support [y/n]? (y): y
```

An approach to troubleshooting

This section outlines a methodology for troubleshooting issues. It introduces steps that you might consider using, depending on the issue in question.

1. Check whether the switch has all the required licenses:
 - License requirements include the POD license, VCS Fabric license, and FCoE license.
 - License types include POD1, POD2, VCS Fabric (multi-node license for more than two nodes), and FCoE.
 - No VCS Fabric license is needed for a one-node or two-node VCS Fabric cluster.
 - The FCoE license needs VCS Fabric mode enabled to be installed.
 - After adding or modifying an FCoE or POD license, always reboot the switch to activate the license.
2. Verify the topology and switch configuration as conveyed by the switch
3. Enter the **copy support** command.
4. Run other relevant show commands (for example, **show logging raslog**) to look for clues or triggers of the reported failure.
5. Check the utilization of various resources.
 - a. Enter the **show process cpu** command to determine CPU use.
 - b. Enter the **show process me** command to determine memory use.
 - c. Enter the **show mac-address-table count** command to determine the number of MAC addresses used.
 - d. Enter the **show fabric route topology** command to determine the number of routes.
 - e. Enter the **show fabric all** command to determine the number of VCS Fabric nodes.

- f. Enter the **show media** command to investigate any optics issues.
6. Conduct data-path fabric continuity tests:
 - a. Issue pings from and to the end-stations or devices.
 - b. Check the counters in the output of the **show interface** command to see if packets are coming in or are dropped as errors.
 - c. Verify that optics used are Brocade-certified. Enter the **show media interface** command and verify that the Vendor name field shows “Brocade.” Check also that the Tx and Rx Power fields are not zero.
 - d. Verify that the MAC address table learns the MAC addresses.
 - e. If the switch is part of a VCS Fabric cluster, verify that the MAC address tables are synchronized properly across all Brocade VDX switches in the cluster.
 - f. Check whether LLDP reports neighbors.
 - g. Check the Ethernet Name Server (ENS) functionality by ensuring that the MAC address table reports MAC addresses learned from other VCS Fabric switches.
 - h. Use the **l2tracroute** command for validating the data-path fabric continuity. This command helps identify where the packets are being dropped within the fabric.

The command prompts for some basic and allows you to choose to enter some extended parameters. Currently supported basic parameters include:

- Source Address (SA) and Destination Address (DA) of dynamically learned MAC addresses
- VLAN
- Edge routing bridge (RBridge) ID

Currently supported extended parameters include:

- Protocol type (IP)
- Source and destination IP addresses
- IP protocol type (recommend TCP)
- Source and destination port numbers

The purpose of IP parameters is to provide a way to make the traceroute packet traverse a specific ECMP link.



CAUTION

The following step affects configuration and should be used with care.

7. To track certain flows within the fabric, use permit ACLs and monitor the hit increments.

Understanding troubleshooting hotspots

This section provides relevant background information and best practices guidance related to features of Network OS where problems have been reported. With this guidance, you should be able to avoid many potential problems.

Licensing

When a licenced feature does not work, one likely cause is that the license has not been installed correctly. Follow the guidelines and procedures in Chapter 7, “Administering Licenses” to ensure your features are licensed properly and those licenses installed correctly.

For license recovery procedures, refer to [“License not properly installed”](#) on page 623.

STP interoperability with Brocade MLX or other switches

- To use the Spanning Tree Protocol (STP) in a network with Brocade MLX switches, or switches from other vendors such as Juniper or Cisco, you may have to configure the interface to send BPDUs to the shared spanning tree MAC address 0100.0ccc.cccd. Without this setting, the RPVST/PVST root bridge is not recognized on VLANs other than VLAN 1.

To interoperate with MLX switches or other vendors’ switches, enter the following command in interface configuration mode:

```
switch(conf-if-te-0/1)# spanning-tree bpdu-mac 0100.0ccc.cccd
```

- If a Brocade IP switch has a VLAN is configured with tagged ports and Rapid Spanning Protocol (RSTP) is enabled under the VLAN (PVST), then BPDUs from the tagged ports received by the Brocade VDX switch will be dropped if pvst-mode is not configured under the ports that are in the VLAN and connected to the Brocade VDX switches.

The following example shows a configuration on a Brocade IP switch with tagged ports and RSTP enabled under the VLAN:

```
vlan 2
tagged ethe 1/24 ethe 2/1 to 2/2
router-interface ve 2
rstp priority 100
```

If the conditions are met, then all the ports should have pvst-mode configured so that tagged BPDUs pass through the Brocade VDX switch. If pvst-mode is not enabled, enable it as follows:

```
Brocade(config)# interface ethernet 2/1
Brocade(config-if-2/1)# pvst-mode
```

Load balancing distribution

Understanding issues related to load balancing requires some basic knowledge of the criteria used by load balancing algorithms. [Table 90](#) provides details for each feature that provides load balancing.

TABLE 90 Load balancing algorithms

Feature	Algorithm
ECMP IP	<p>Paths are selected on the basis of a hash derived from the following parameters:</p> <ul style="list-style-type: none"> • Source MAC address • Destination MAC address • VID • IP protocol • Source IP address • Destination IP address • Layer 4 source port • Layer 4 destination port <p>You can configure the hashing fields using the fabric-ecm load-balance and fabric-ecmp load-balance-hash-swap commands.</p> <p>For related recovery procedures, refer to “ECMP not load balancing as expected” on page 616.</p>
ECMP FCoE	<p>Paths are selected on the basis of a hash derived form the following parameters:</p> <ul style="list-style-type: none"> • Input Port ID • Source MAC address • Destination MAC address • VID • FID • SID • DID • OXID
LACP	Provides adaptive load balancing based on up to seven criteria (7-tuple), depending upon what fields are available in the frame.
Brocade trunk	Provides equal packet load balancing (round-robin) among member links.

Static assignment of the routing bridge ID

Duplicate routing bridge (RBridge) IDs are a common source of error when a switch is added to an Ethernet fabric. Before adding a switch to an Ethernet fabric, you must assign it a unique RBridge ID. If the new switch is to be added to an existing VCS Fabric cluster, it must be assigned the same VCS ID as other switches in the cluster. Once the switch is added, the principal routing bridge performs the negotiation in the control plane to include the new switch and rebuild the fabric. The data plane remains unaffected.

Procedures for recovering from duplicate routing IDs are provided in [“RBridge ID is duplicated”](#) on page 633.

FSPF route change

When the Fabric Shortest Path First (FSPF) algorithms select a new route, a temporary disruption of traffic can occur. This behavior is normal as the old path is first deleted and then the new path is programmed. Such path changes can occur when FSPF calculates a new shortest route, or when the current path is down.

VCS Fabric mode and standalone mode

Some key differences exist between standalone mode and VCS Fabric mode that you should be aware of when troubleshooting your system:

- Interfaces are disabled by default in standalone mode but enabled by default in VCS Fabric mode. Thus, when you apply the default configuration, you should take this into account.
- The interface can be configured as a Layer 2 switch port in standalone and VCS Fabric modes.
- Switching between VCS Fabric mode and standalone mode and then reverting back to the original mode results in losing the configuration and booting up using the default configuration.
- A port-profile port is allowed only for Layer 2 ports.
- The out-of-band management through the management port allows default gateways to be configured.

vLAG

You should be aware of the following aspects of the vLAG feature before troubleshooting vLAG problems:

- Multicast (BUM) traffic in vLAG
- Edge-port feature requirements
- Failover

Multicast traffic in vLAG

Flooding traffic always goes through a primary link of the vLAG. You should consider this restriction when provisioning bandwidth for most traffic. This link is marked with an asterisk(*) in the output of the **show port-channel** command.

```
switch# show port-channel 38
LACP Aggregator: Po 38
Aggregator type: Standard
Admin Key: 0038 - Oper Key 0038
Partner System ID - 0x8000,01-e0-52-00-20-00
Partner Oper Key 0038
Member ports:
Link: Te 0/13 (0x180D0102) sync: 1
Link: Te 0/14 (0x180E0103) sync: 1 *
```

Edge-port feature requirements for vLAG

LACP can be configured on edge ports only with either Brocade or Standard types. If Brocade is chosen, so that Link Reset (LR) primitives are exchanged properly, make sure that the edge peering device is a Brocade Converged Network Adapter (CNA), a standalone Brocade VDX switch, or a Brocade 8000 switch.

Failover and vLAG

For the fast failover convergence requirements, Brocade recommends using the **vlag ignore-split** command, which enables subsecond failover times. This command is added automatically to all port-channel configurations when a Brocade VDX switch is upgraded to Network OS v2.1.x or when a new port-channel is added while running under Network OS v2.1.x.

When planning to deploy this feature in production, use care to prevent a “split-brain” scenario, in which vLAG members detach from each other. Brocade recommends having more than one interswitch link (ISL) between the vLAG member switches and physically routing them through separate conduits and cable trays. Secondly, Brocade strongly recommends using topologies that are certified by Brocade.

NOTE

Brocade does not recommend using vLAG failover in a network with Cisco or Juniper switches that are connected using copper. Brocade has observed greater than one-second failover times in networks with this hardware.

The following topics discuss the split-brain scenario and how to mitigate it.

Understanding “split-brain”

A split-brain can occur when the end-hosts or edge switches are connected to two separate cluster switches by way of a vLAG (using LACP). The end-devices see those two cluster switches as one switch because they have the same system ID advertised in LACP.

Under rare conditions, when all the ISLs between the two cluster switches are broken and both the cluster switches continue to advertise the same system ID to their LACP partner, a “segmented fabric” or “split-brain” condition exists, where the end-host or edge switch might not detect this segmentation and could continue to treat both the vLAG switches as one switch.

ATTENTION

This condition can cause packet duplication or unexpected packet loss.

Network OS protects traffic during split-brain conditions

By default, Network OS has a capability to recover gracefully from the split-brain scenario. When all the ISLs between the VDX cluster switches go down, the switch with the lower RBridge ID uses LACP to inform the edge-switch partner that it has segmented out of the port-channel. It does this by changing its advertised system ID. When the edge switch learns a different system ID on one of its members, it removes this member from that port-channel, and continues to function with only one vLAG member—the switch with the higher RBridge ID. The other vLAG member switch still has the link up, but remains segmented out of the original port-channel (sync: 0). This capability prevents duplication of packets or potential packet drops resulting from a split-brain scenario.

When a member switch is reloaded

Reloading the switch with the lower RBridge ID has no impact.

When the switch with the higher RBridge ID is reloaded, the other vLAG member sees all of its ISLs as down. Though this is *not* a real split-brain scenario, the switch with the lower RBridge ID may not be able to differentiate, and thus would inform the partner about a changed system ID. The partner edge switch would see two events:

- The system ID on one link changes.
- The other interface goes down.

In such a case, LACP will renegotiate and reform the port-channel, which could flap the port-channel, impacting traffic momentarily. The same effect could occur when the switch boots up and joins the fabric again.

Thus, if the switch with the higher RBridge ID is reloaded, the potential impact could be a port-channel flap that can momentarily disrupt traffic. Note that this effect does not occur when the switch with the lower RBridge ID is reloaded.

Avoiding traffic disruption during switch reload

Network OS switches offer flexibility to the user by providing a special vLAG **ignore-split** option that you can configure for the logical port-channel. This option should be configured on both vLAG member ports.

Configuring this option prevents the switch with the lower RBridge ID from changing its system ID, so both switches will continue to advertise the same system ID. This action prevents the partner edge switch from detecting a change when one of the member switches is reloaded and the traffic is handled gracefully.

Using the vLAG ignore-split option

To use the vLAG **ignore-split** option, redundancy should be built around ISLs to prevent a situation in which all ISLs are broken at the same time. Brocade recommends using multiple ISLs, and routing those ISLs through different physical paths or conduits to eliminate the possibility of accidental damage to all links at the same time.

Principal routing bridge availability

If a new principal routing bridge is introduced into a working VCS Fabric cluster, or if the principal routing bridge is lost and a new switch must be elected, the fabric is rebuilt from the control-plane viewpoint, whereas the data plane continues to forward traffic without disruption. The primary responsibilities of the principal routing bridge in a VCS Fabric are:

- RBridge ID allocation
- Ownership of virtual management IP address
- Keeping the configuration database synchronized

Brocade trunks

Brocade trunks is the only aggregation method that works using ISLs.

Brocade ISL trunks are formed automatically with other switches using Line Reset (LR) primitives signaling with the peer switch.

All ISL ports connected to the same neighbor Brocade switch attempt to form a trunk. For a successful trunk formation, all ports on the local switch must be part of the same port group and must be configured at the same speed. The number of ports allowed per trunk group is six. The trunk is turned on by default.

Table 91 shows allocation of port numbers to port groups for Brocade VDX switches.

TABLE 91 Port groups

Network OS switch	Port groups
Brocade VDX 6720-24 and Brocade VDX 6730-32 (switches with 24 Ethernet ports)	te0/1 through te0/12
	te0/13 through te0/24
Brocade VDX 6720-60 and Brocade VDX 6730-76	te0/1 through te0/10
	te0/11 through te0/20
	te0/21 through te0/30
	te0/31 through te0/40
	te0/41 through te0/50
	te0/51 through te0/60
Brocade VDX 6710	te0/1 through te0/6
	gi0/1 through gi0/14
	gi0/15 through gi0/27
	gi0.28 through gi0/48

NOTE

Brocade trunks are not supported over 1-Gbps links.

To utilize the advantages of Brocade trunking between VDX switches, Brocade recommends having at least a two-member trunk and multiple ECMP paths. Brocade also recommends routing the cables in a trunk through separate conduits to ensure connectivity in case a conduit is accidentally cut.

NIC teaming with vLAG

NIC teaming permits link aggregation between server and switch. It can be one of two types: active/passive model or active/active model. For the active/passive model, you may not need to configure a LAG on the switch side, as unique MAC addresses will be seen on only one link.

For the active/active model, the same MAC address may appear on both the links terminating on a switch (or pair of switches). In such a case, you must configure a LAG on the switch side.

Selecting the MTU

Always set the switch MTU to the maximum host MTU plus 100 bytes. This method is recommended because the definition of MTU sometimes varies among different vendors. If the switch MTU is set to the same as the connected host MTU, packets could be dropped.

Avoiding oversubscription

Under certain congestion conditions, you may observe incrementing packet drops representing “tail-drops” in the output of the **show qos rcv-queue interface tengigabitethernet** command, as shown underlined in the following example:

```
switch# show qos rcv-queue interface tengigabitethernet 5/0/1
Interface TenGigabitEthernet TenGigabitEthernet 5/0/1
In-use 0 bytes, Total buffer 144144 bytes
0 packets dropped
      In-use   Max
CoS   Bytes   Bytes
-----
  0    0      18018
  1    0      18018
  2    0      18018
  3    0      18018
  4    0      18018
  5    0      18018
  6    0      18018
  7    0      18018
```

In such conditions, you must first identify the bottleneck, and then take action to mitigate the congestion.

Identifying the congestion bottleneck

To identify the bottleneck in the Brocade VDX network, enter the **show interface** command at various locations, and identify interfaces with incrementing TX and RX discards. Depending upon the TX or RX discards, the congestion could be anywhere downstream.

Mitigating the congestion

Try the following actions to mitigate congestion:

- Increase bottleneck bandwidth.
 - Add more links to the LAG and ECMP paths.
 - Use higher-speed interfaces.
- Implement flow control on the bottleneck and on neighboring devices.
- Implement QoS congestion management schemes.
 - Classify, mark, and prioritize critical traffic.
 - Modify scheduling schemes. Consider and compare the effects of using strict priority or deficit weighted round-robin (DWRR) scheduling schemes.

For the flow control solution, enable flow control either on the ports receiving the traffic from end-devices (servers or personal computers) and the connected end-device itself, or enable flow control on the port-channel as shown in the following example.

```
switch(conf-if-te-1/0/24)# interface port-channel 100
switch(config-Port-channel-100)# qos flowcontrol tx on rx on
```

Once flow control is enabled, enter the **show qos rcv-queue interface tengigabitethernet** command again and check the output. It should no longer be reporting packet drops. If the packet drops continue or the ingress rate is considerably lower than expected, contact your switch support provider for further investigation.

We recommend enabling asymmetric flow control with Brocade VDX switches. For any two adjacent devices, one device should have Rx ON and Tx OFF, while the other device should have Rx OFF and Tx ON.

Refer to [“Congestion control and queuing”](#) on page 319 for further details about congestion control.

ACL limits issues

If you keep within the supported limits of ACL usage as shown in [Table 92](#), you are unlikely to run into system limits issues. ACLs should instantiate quickly and correctly.

TABLE 92 ACL limits per switch in VCS mode

Feature	Limit
Number of standard or extended ACLs created but not applied	512
Number of Layer 3 standard or extended ACLs created but not applied	512
Number of rules per standard or extended ACL	2048
Maximum number of Layer 2 or Layer 3 standard or extended ACL rules	100k
Number of physical interfaces on which an ACL is applied concurrently	48 (60 in standalone mode)
Number of VLAN interfaces on which ACL is applied concurrently	100
Number of ACL counters	252
Number of TCAM table entries	1000
Number of ACL rules	6000
Number of applied, co-existing standard and extended ACLs	50

In addition, up 30,720 MAC addresses are supported.

As you approach or exceed combinations of these limits, it is possible you might encounter slow instantiation of ACL rules, process exceptions, or ACL failure due to MAC learning issues.

Delays of several minutes can occur in the instantiation of ACL rules and counters if the number of ACLs or VLANs is excessive. The L2SYS process message queue can become full, or CPU context switching and process scheduling can increase to the point that ACL instantiation proceeds slowly. Periodic monitoring with the **show statistics access-list mac** command will show not more than 252 ACL rules with a nonzero and incrementing frame count for rules that are correctly instantiated and have hardware counters allocated.

Process exceptions can sometimes occur with the L2SYSD process when combinations of ACL limits are approached or exceeded.

Constant MAC learning and flushing can occur when ASIC table limitations are exceeded. Layer 2 frame switching can fail if the number of MAC address table entries is exceeded.

Troubleshooting procedures

This section describes some potential problems you may encounter and suggestions on how to investigate or resolve each issue. If these steps do not lead to resolution of the problem, prepare a case for your switch provider, as described in [“Getting technical help”](#) on page xl.

- [“AMPP is not working”](#) on page 610
- [“Continuous panic reboots”](#) on page 614
- [“Corrupted CID card”](#) on page 614
- [“CPU use is unexpectedly high”](#) on page 615
- [“ECMP not load balancing as expected”](#) on page 616
- [“ENS functionality check”](#) on page 616
- [“FCoE devices unable to log in”](#) on page 617
- [“General debugging for traffic forwarding”](#) on page 619
- [“ISL does not come up on some ports”](#) on page 620
- [“License not properly installed”](#) on page 623
- [“Packets dropped in hardware”](#) on page 624
- [“Password recovery for the Brocade VDX 8770”](#) on page 630
- [“Ping failure”](#) on page 633
- [“QoS configuration causes tail drops”](#) on page 633
- [“QoS is not marking or treating packets correctly”](#) on page 633
- [“RBridge ID is duplicated”](#) on page 633
- [“SNMP MIBs report incorrect values”](#) on page 634
- [“SNMP traps are missing”](#) on page 634
- [“Telnet operation into the switch fails”](#) on page 634
- [“Trunk member not used”](#) on page 635
- [“Upgrade failure”](#) on page 637
- [“VCS Fabric cannot be formed”](#) on page 637
- [“vLAG cannot be formed”](#) on page 638
- [“Zoning conflict resolution”](#) on page 640
- [“Zone does not form correctly”](#) on page 640

AMPP is not working

Configuring Brocade Automatic Migration of Port Profiles (AMPP) is complex. It works in standalone mode and VCS Fabric mode. For details on configuring AMPP, refer to [Chapter 21, “Configuring AMPP”](#).

Problems encountered while using AMPP are usually the result of configuration errors in the port-profile itself, errors in the associated virtual machine (VM) configuration, or compatibility problems between the host adapters and AMPP. Specifically, AMPP problems can be caused by the following conditions:

- A port-profile configuration does not exist on the target switch or does not contain a basic switchport and VLAN configuration.
Refer to [“Verifying the port-profile configuration”](#) on page 611.
- The VM MAC address does not appear in the MAC address table.
Refer to [“Verifying the VM MAC address”](#) on page 612.
- The port-profile is not activated or is not associated with the correct MAC address.
Refer to [“Verifying the port-profile state”](#) on page 612.
- The VM kernel MAC addresses are not associated correctly with the port-profile on the respective switches.
Refer to [“Verifying the VM kernel MAC addresses”](#) on page 612.
- The VM and its associated hosts do not share a common storage device.
Refer to [“Verifying a shared storage device”](#) on page 613.
- The port-profile was learned on a nonprofiled VLAN.
Refer to [“Verifying the status of a learned profiled MAC address”](#) on page 613.
- A conflicting port-profile is applied to the same interface.
Refer to [“Verifying that port profiles do not conflict”](#) on page 613.
- The Ethernet Name Server is not functioning correctly.
Refer to [“Verifying the Ethernet Name Server”](#) on page 613.
- An ESX host has an incompatible network adapter or driver installed.
Refer to [“Verifying an ESX host”](#) on page 613.

Verifying the port-profile configuration

A valid port-profile must exist on the target switch. It must contain a basic switchport and VLAN configuration.

1. In the privileged EXEC mode, enter the **show running-config port-profile** command to verify that the port-profile configuration exists on the target switch, and that it contains a basic switchport and VLAN configuration.

```
switch# show running-config port-profile
port-profile default
  vlan-profile
  switchport
  switchport mode trunk
  switchport trunk allowed vlan all
  switchport trunk native-vlan 1
  !
!
port-profile pp1
  vlan-profile
  !
!
port-profile pp2
  vlan-profile
  !
!
```

2. If the port-profile configuration does not exist or is missing the required switchport or VLAN configuration, create the port-profile as described in [“Configuring AMPP port-profiles”](#) on page 162.

Verifying the VM MAC address

For the correct functioning of AMPP, the MAC address for the VM and its associated hosts must appear in the MAC address table.

1. Enter the **show mac-address-table** command to verify that the VM MAC addresses appear in the switch MAC address table.

```
switch# show mac-address-table
VlanId  Mac-address      Type    State    Ports
1       0000.0010.0001    Static  Inactive Te 4/0/3
1       0000.0010.0002    Static  Inactive Te 4/0/3
Total MAC addresses : 2
```

2. If a VM MAC address is not present, contact your switch support provider for further investigation and provide this data.

Verifying the port-profile state

For the correct functioning of AMPP, the port-profile must be active and must be associated with the correct MAC address.

1. Enter the **show port-profile status** command to verify that the port-profile is activated and is associated with the correct MAC address.

```
switch# show port-profile status
Port-Profile  PPID    Activated  Associated MAC  Interface
pp1           1       No        None           None
pp2           2       No        None           None
```

2. Correct any misconfigurations as follows:

- If the port-profile is not activated, enter the **port-profile profile-name activate** command to activate it.
- If the port-profile is not associated with a MAC address, enter the **port-profile port-profile-name static** command to perform the association.

```
switch(config)# port-profile PP3 static 0050.5600.10030
```

- If the port-profile is associated with the wrong MAC address, enter the **no port-profile port-profile-name static** command to break the association with the incorrect MAC address, and then reassociate the port with the correct MAC address.

```
switch(config)# no port-profile PP3 static 0050.5600.10020
```

```
switch(config)# port-profile PP3 static 0050.5600.10030
```

Refer to [“Configuring a new port-profile”](#) on page 304, for details about activating a port-profile and associating a port-profile with a MAC address.

Verifying the VM kernel MAC addresses

Confirm that the virtual machine (VM) kernel MAC addresses are also associated with the port-profile on the respective switches. If not, perform the association as described in [“Verifying the port-profile configuration”](#) on page 611.

Verifying a shared storage device

Confirm that the VM and its associated hosts are sharing a storage device. If not, then reconfigure the VM and hosts to share a storage device.

Verifying the status of a learned profiled MAC address

For correct functioning of AMPP, the MAC address must be learned from a valid source—a profiled VLAN. This procedure determines whether a MAC address was learned from a valid source.

Enter the **show mac-address-table port-profile** command to check the status on learned profiled MAC addresses.

```
switch# show mac-address-table port-profile
Legend: Untagged(U), Tagged (T), Not Forwardable(NF) and Conflict(C)
VlanId    Mac-address      Type      State      Port-Profile      Ports
1         0050.5679.5351   Dynamic   Active     Profiled(U)       Te 111/0/10
1         0050.567b.7030   Dynamic   Active     Profiled(U)       Te 111/0/12
1         005a.8402.0000   Dynamic   Active     Profiled(T)       Te 111/0/24
1         005a.8402.0001   Dynamic   Active     Profiled(NF)      Te 111/0/24
1         005a.8402.0002   Dynamic   Active     Not Profiled      Te 111/0/24
1         005a.8402.0003   Dynamic   Active     Not Profiled      Te 111/0/24
1         005a.8402.0004   Dynamic   Active     Not Profiled      Te 111/0/24
(output truncated)
Total MAC addresses : 17
```

Check for and investigate MAC addresses identified in the output as “Not Profiled.”

Verifying that port profiles do not conflict

1. Enter the **show port-profile name pp1_name name pp2_name validate** command to validate whether multiple port-profiles applied on an interface can co-exist without conflict.

```
switch# show port-profile name pp1 name pp2 validate
Port-Profile      Port-Profile      Conflicts
-----
pp1               pp2
vlan-profile      vlan-profile      No
qos-profile       qos-profile       No
security-profile  security-profile  No
```

2. If a conflict exists, reconfigure one of the port-profiles to avoid the conflict.

Refer to [Chapter 21, “Configuring AMPP,”](#) for information about the rules for co-existence.

Verifying the Ethernet Name Server

AMPP requires each VCS Fabric switch in the cluster have the same view of the MAC address table. Any differences in the view indicate a failure of the Ethernet Name Server (ENS). Refer to [“ENS functionality check”](#) on page 616 for details.

Verifying an ESX host

Verify that each ESX host has the correct Converged Network Adapter (CNA) installed with appropriate drivers, and does not use the Cisco Nexus 1000V software switch, as that switch might send out specially crafted packets.

Continuous panic reboots

If your switch is having continuous panic reboots, perform the following procedure.

1. Bring the switch to a stable state with the original binary or swap the other partition with a good image.
 - a. In bootenv, execute **setenv OSLoadOptions 2, saveenv** and **reset**.
 - b. Immediately after reboot before panic, try to login and execute **do chkconfig fabos off**.

This boots the switch in single-user mode, with just Linux OS and not loading the Network OS modules. Make the needed changes to the filesystem, such as changing the partition by using **bootenv** or replacing the image with original good binaries.

In single-user mode the network is not accessible, so use the command **ifconfig eth0** for any network support. Furthermore, the command **init 3** boots the switch after reboot for case a) above. Otherwise, reset the switch by using **chkconfig fabos on; reboot**.

2. If the switch comes to a halt after five repeated reboots, try cleaning up the DCMD database with the following command:

```
rm -rf /etc/fabos/Dcmd/*.cfg; rm -rf /etc/fabos/Dcmd/WaveDatabase; rm -rf
/etc/fabos/Ccmd/*.cfg; rm -rf /etc/fabos/Ccmd/WaveDatabase; reboot
"/sbin/reboot -f
```

3. Download a good nightly build to return the switch to a stable state.

Corrupted CID card

In the case of a corrupted CID card, perform the following steps.

1. Link the **wwncardshow** command to see the extent of the damage. (This does not have to be done for single boards.)

```
switch# ln -s /fabos/cliexec/em /fabos/bin/wwncardshow
```

2. Display the wwncardshow data.

```
switch# wwncardshow ipdata
packet count is 2
++ Wwn Card IP Data ++
Type Num Field Address Mask Cfg/Zone
-----
CP 0 Eth IP: 255.255.255.255 255.255.255.255
CP 1 Eth IP: 255.255.255.255 255.255.255.255
Chassis GW IP: 255.255.255.255
      LicID: 10:00:00:ff:ff:ff:ff:ff enet cfg
      Name: VDX 6710-54 Gen# : -1/0
Sw 0 Eth IP: 10.17.10.84 255.255.240.0
      FC IP: 0.0.0.0 0.0.0.0
      GW IP: 10.17.0.1
      WWN: 10:00:00: 05:33:14:b2:70
      Name: swd77 Gen# : 0/0
Sw 1 Eth IP:
      FC IP:
      GW IP:
      WWN: 10:00:00:05:33:14:b2:71
      Name: Gen# : 0/0
```

Items that are FFs, 255s, or zeros are unacceptable. Only the first two groups count, and the items that must be correct are the following:

- The CP Eth IP entries. They need valid data only if that CP/MM is present.
 - The chassis LicID entry.
 - The Sw 0 Eth IP entry.
 - The Sw 0 GW IP entry.
 - The Sw 0 WWN entry.
3. To correct the CP Eth IP entries, run **ipaddrset -cp x**, where x is 0 for MM1 and 1 for MM2, and put in correct data at the prompts. Then un **ipaddrset -chassis** and enter the correct data as needed.

Sometimes, if the entries have enough 255/0xFFs in them, **ipaddrset** does not update the values properly, in which case you have to use **test_sysmod** to clear a couple of entries.

4. To correct Sw 0 WWN, run **wwn -d626 xx:xx:xx:xx:xx:xx:xx** with the correct wwn value. The system must be rebooted for the change to take effect (at the prompt or manually).
5. To correct chassis LicID, you need the test_sysmod tool. Mount a filesystem (if necessary get eth0 up manually with ifconfig, or set the gateway first).

```
switch# run test_sysmod
test_sysmod
```

6. At the first menu, enter **11** for WWN testing, then **2** for copy WWN to LID, and then enter **1** to confirm. Perform a **Ctrl-C** to exit.
7. The system must be rebooted for the change to take effect. Exit test_sysmod with **Ctrl-C**.

If you have lost both the WWN and license ID, then you must perform [step 4](#) first. If you do not know the value, it is available in the MAC address in the boot environment variables (for pizza boxes only).

This value can be entered in the **wwn** command by inserting 10:00: before the MAC value).

8. Finally, if you can't correct the IP addresses, there is one more option in test_sysmod that can help. At the main menu, enter **11** for WWN testing and then **1** for clear WWN IP data entry, then **0**, **1**, **2**, or **3** for entries that had a lot of FFs. If you clear all of the entries that are corrupted with FFs, you should be able to run **ipaddrset** to restore the real addresses.
9. Reboot the switch in order for the change to take effect and make the **ipaddrset** command available.

Verifying SEEPROM data

1. To verify the SEEPROM, copy the test_symod file to /fabos/bin as test_sysmod, and select option **10** for i2c and option **27** to Verify FRU Seeprom. The test begins automatically.
2. Use the offset of 0x6a4c, as that is where the IP table starts (size 256), but any offset (and size less than or equal to 256) will access that device.

CPU use is unexpectedly high

Unexpectedly high CPU use is usually the result of a process consuming a large percentage of available CPU cycles. It can prevent access to the switch by Telnet or make an ISL nonfunctional.

If you suspect high CPU use, complete the following steps.

1. In privileged EXEC mode, enter the **show process cpu** command to determine which process is causing the high CPU reading.
2. Shut down the corresponding interface or delete the configuration suspected of causing the high CPU use.

ECMP not load balancing as expected

Equal cost multipath (ECMP) routing increases throughput by balancing traffic across multiple routes that tie for best cost. If you suspect that traffic is not being balanced as expected, complete the following steps.

1. In privileged EXEC mode, enter the **show fabric route topology** command to see whether ECMP routes are expected.

```
switch# show fabric route topology
```

```
Total Path Count: 1
```

Src RB-ID	Dst RB-ID	Out Index	Out Interface	Nbr Hops	Nbr Cost	Index	Interface	BW	Trunk
66	1	124	Fi 66/0/4	1	500	129	Fi 1/-1/-1	32G	Yes

If the output shows multiple equal-cost paths between the source and destination switches, then ECMP load balancing is expected.

2. Check the interface utilization to verify whether it matches with the expected number of flows.
3. Enter the **I2traceroute** command to investigate whether Layer 2, Layer 3, and Layer 4 flows hash to separate ECMP links.

To avoid disruption of operation inherent in ECMP, the correctly functioning Brocade routing strategy routes a specific flow along one deterministic route. Additional flows take available equal-cost routes. This step verifies whether this flow hashing strategy is functioning correctly.

For details about using the **I2traceroute** command, refer to [“Layer 2 traceroute”](#) on page 644.

ENS functionality check

The Ethernet Name Server (ENS) is working correctly when the content of MAC address tables is the same among switches in the same VCS Fabric cluster. Perform the following checks to ensure that ENS is working correctly:

- Check the that fabric membership information is what you expect. Refer to [“Verifying the fabric”](#) on page 617.
- Ensure that MAC addresses are not moving among ports. Refer to [“Checking for MAC address movement among ports”](#) on page 617.
- Ensure that no edge port has an external loopback. Refer to [“Verifying edge ports have no external loopback”](#) on page 617.

Verifying the fabric

Enter the **show fabric all** command and ensure that information about all switches in the VCS Fabric cluster is displayed.

```
switch# show fabric all
```

```
VCS Id: 1
```

```
Config Mode: Local-Only
```

Rbridge-id	WWN	IP Address	Name
1	50:00:51:E4:44:40:0E:04	0.0.0.0	"fcr_fd_1"
2	50:00:51:E4:44:50:0F:09	0.0.0.0	"fcr_xd_2_128"
60	10:00:00:05:33:5F:EA:A4	10.24.81.65	"switch"
66	10:00:00:05:33:67:26:78	10.24.81.66	>"switch"

```
The Fabric has 4 Rbridge(s)
```

Checking for MAC address movement among ports

MAC address movement from port to port occurs when the same source address is detected on multiple ports. This condition is sometimes known as “MAC address flapping.”

To check for MAC address flapping, enter the **show mac-address-table** command multiple times and check the output.

Verifying edge ports have no external loopback

Physically check for extended loopback.

FCoE devices unable to log in

The inability to log in from a device connected through FCoE is usually because either the port or LLDP has been incorrectly configured. Potential reasons include:

- The default profile map has not been applied correctly. Refer to [“Verifying the default profile map”](#) on page 618.
- Required TLVs have not been advertized under LLDP. Refer to [“Verifying TLVs”](#) on page 618.

CNAs not logging into the switch

If CNAs are not logging into the switch, perform the following procedure.

1. Check that the physical port is provisioned for FCOE.

```
switch# show fcoe interface ethernet | include "1/0/5"
TenGigaBitEthernet 1/0/5      default
```

2. If the physical port is not provisioned, provision the interface for FCOE.
3. If the CNA is still not logging in, check that the logical FCOE interface is online by using the **show running-config interface fcoe** command, as in the following example:

```
switch# show running-config interface fcoe
interface Fcoe 1/11/1
no shutdown
```

```

!
interface Fcoe 1/11/2
no shutdown
!
interface Fcoe 1/11/3
no shutdown
!
interface Fcoe 1/11/4
no shutdown
!
interface Fcoe 1/11/5
no shutdown
!
interface Fcoe 1/11/6
no shutdown
!
interface Fcoe 1/11/7
no shutdown

```

4. Remove the FCOE provisioning and reprovision the physical interface.
5. If that does not work, execute the **shut**, and then the **no shut** command on the FCOE logical interface.
6. If it still fails, collect the **supportsave** information and contact support. Refer also to [“Gathering troubleshooting information”](#) on page 599, which provides information about Network OS supportSave files.

Verifying the default profile map

1. In privileged EXEC mode, enter the **show running-config interface tengigabitethernet** command to determine whether the default profile map has been applied to the interface.

```

switch# show running-config interface tengigabitethernet 5/0/1
interface TenGigabitEthernet 5/0/1
 fcoeport default
 shutdown

```

2. If the default profile map has not been applied to the interface, or the initiator and target do not share the same VLAN ID, in interface configuration mode, enter the **fcoeport default** command to apply it.

```

switch(conf-if-te-0/1)# fcoeport default

```

This command not only applies the default profile map, but also associates the initiator and target with the same VLAN ID.

Verifying TLVs

The following TLVs—dcbx-fcoe-app-tlv, dcbx-fcoe-logical-link-tlv, and dcbx-tlv— must be advertized under LLDP or FCoE devices will not be able to log in.

1. In the privileged EXEC mode, enter the **show running-config protocol lldp** command to verify that the required TLVs are advertized.

```

switch# show running-config protocol lldp
protocol lldp
 advertise dcbx-fcoe-app-tlv
 advertise dcbx-fcoe-logical-link-tlv
 advertise dcbx-tlv

```

2. If any of the required TLVs is missing, in protocol configuration mode, enter the corresponding **advertise** command.

```
switch# configure terminal
switch(config)# protocol lldp
switch(config-lldp)# advertise dcbx-fcoe-app-tlv
switch(config-lldp)# advertise dcbx-fcoe-logical-link-tlv
switch(config-lldp)# advertise dcbx-tlv
```

General debugging for traffic forwarding

If the traffic is not being forwarded, perform the following steps:

1. Check for db packet capture. Below are the commands to enable and view a capture

```
db 8/0/1 rte enable capture all
db 8/0/1 rte start capture
db 8/0/1 rte show capture
```

After the **start capture** command, the system sends a stream and performs **show capture**. This displays most of the capture information:

- a. It shows all the fields resolved—whether it is trap, drop, or fwd.
- b. It shows the packet itself.
- c. It shows the Routing Engine (RTE) Layer 2 history, as in the result of the Layer 2 table hit or miss.
- d. It shows the RTE Layer 3 history, as in the result of the Layer 3 table hit or miss. If Layer 2 table has a success and Layer 3 table failed, then check for routing issues.

For example, in the entry for trap (Ping to box), the Routed fields should display Ipv4Rtd and the entries hit. For a TRAP hit, it should display trape:1 (Bit set to 1 to indicate packet is trapped).

- e. Packet Capture displays the last four packets, Make sure those are the fwd packets (for example, check for SA DA MAC, pkttyp:0806).

```
db 8/0/1 rte enable capture all
db 8/0/1 rte start capture
db 8/0/1 rte show capture
```

- f. If the FWD and packets are not being forwarded, then it is an ASIC problem. If it is a DROP, proceed to [step 2](#).

2. If result is DROP:

- a. Execute the **show ip route** command. If the route is not present, then it is an RTM issue.
- b. Execute the **show arp command**. If ARP is not resolved for the corresponding next hop, then it is an ARP issue. If it is VLAN along with ARP, MAC should be resolved. If MAC is not resolved then it is an L2SYS issue.
- c. If [step b](#) passes, enter the **debug show ip lpm** command to display the routes in hardware, and verify that the corresponding destination ARP address is present. If it is not present, then it is an L3FWD issue. Collect the information from **debug show ip lpm**, attach the file `/tmp/fib_wlv_ioctl`, along with supportsave data and contact support. Specify that it failed in this step.

Refer also to [“Gathering troubleshooting information”](#) on page 599, which provides information about Network OS supportSave files.

- d. If [step c](#) passes and traffic is still dropping, then it is an ASIC issue.

NOTE

For additional information about packet capture, see [“Using the packet capture utility”](#) on page 652.

ISL does not come up on some ports

The failure of an interswitch link (ISL) between two switches in a VCS Fabric cluster can occur for various reasons:

- The ISL configuration is disabled. Refer to [“Verifying the status of ISLs”](#) on page 620.
- The ISL is segmented. Refer to [“Verifying the status of ISLs”](#) on page 620.
- VCS Fabric mode is not enabled on one of the switches. Refer to [“Verifying VCS Fabric configuration and RBridge ID”](#) on page 621.
- Different VCS IDs on each of the switches. Refer to [“Verifying VCS Fabric configuration and RBridge ID”](#) on page 621.
- LLDP is not reporting its neighbors. Refer to [“Verifying LLDP”](#) on page 622.
- An overloaded CPU fails to generate keepalive packets. Refer to [“Checking for CPU overload”](#) on page 623.

Verifying the status of ISLs

If any port looks suspicious, begin by checking the status of ISLs.

1. On the switches at each end of the broken link, in privileged EXEC mode, enter the **show fabric isl** command to view the status of ISL connections.

```
switch1# show fabric isl
Rbridge-id: 2    #ISLs: 2
Src      Src      Nbr      Nbr
Index   Interface  Index   Interface  Nbr-WWN      BW    Trunk  Nbr-Name
-----
1        Te 2/0/1    1        Te 3/0/1    10:00:00:05:1E:CD:7A:7A  10G   Yes   "switch1"
2        Te 2/0/2    ?        Te ?/?/?    ????:??:??:??:??:??:? (segmented - incompatible)
26       Te 2/0/26   56       Te 25/0/56  10:00:00:05:33:40:2F:C9  60G   Yes   "Edget12r31_25"
34       Te 2/0/34   58       Te 26/0/58  10:00:00:05:33:41:1E:B7  40G   Yes   "Edget12r32_26"
```

Ports on which the ISL link is broken appear with the text “(segmented - incompatible).” Ports for which the ISL configuration is disabled do not appear in the output.

2. Enter the **show fabric islports** command to gather more information about the status of suspect ports.

```
switch# show fabric islports
Name:      switch
Type:      107.4
State:     Online
Role:      Fabric Subordinate
VCS Id:    10
Config Mode: Local-Only
Rbridge-id: 11
WWN:       10:00:00:05:33:6d:7f:77
FCF MAC:   00:05:33:6d:7f:77

Index   Interface  State  Operational State
```

```

=====
 1      Te 11/0/1   Up      ISL 10:00:00:05:33:00:77:80 "switch" (upstream)(Trunk Primary)
 2      Te 11/0/2   Down
 3      Te 11/0/3   Down
 4      Te 11/0/4   Up      ISL (Trunk port, Primary is Te 11/0/1)
 5      Te 11/0/5   Down
 6      Te 11/0/6   Down
 7      Te 11/0/7   Down
 8      Te 11/0/8   Down
 9      Te 11/0/9   Down
10      Te 11/0/10  Down
11      Te 11/0/11  Up      ISL 10:00:00:05:1e:00:50:00 "switch" (Trunk Primary)
121     Fi 11/0/1   Up      LS ISL 50:00:53:37:b6:93:5e:02 "fcr_fd_160" (downstream)
(Trunk Primary)
122     Fi 11/0/2   Up      LS ISL (Trunk port, Primary is Fi 11/0/1)
123     Fi 11/0/3   Down
124     Fi 11/0/4   Down
125     Fi 11/0/5   Down
126     Fi 11/0/6   Down
127     Fi 11/0/7   Down

```

3. If the port state is “Down,” enable the port with the **no shutdown** command.

```

switch# configure terminal
Entering configuration mode terminal
switch(config)# interface tengigabitethernet 178/0/9
switch(config-if-te-178/0/9)# no shutdown

```

4. If the port state is “Up,” but the ISL is segmented, examine the Operational State string for further clues to the reason for the segmentation.

Refer to the *Network OS Command Reference* for details about the **show fabric islports** command and help in interpreting the Operational State string for a segmented ISL.

Verifying VCS Fabric configuration and RBridge ID

For the ISL to function correctly, the following criteria must be true:

- Both switches must have VCS Fabric mode enabled.
- Both switches must have the same VCS ID.
- Each switch must have a unique RBridge ID.

To check the criteria, complete the following steps.

1. Enter the **show vcs** command on each switch.
2. Depending on the output, proceed as follows:
 - If the VCS Fabric mode is not enabled on either switch, enter the **vcs enable** command to enable it.

```

switch1# show vcs
Config Mode   : Local-Only
VCS ID        : 1
Total Number of Nodes      : 1
Rbridge-Id    WWN           Management IP   Status   HostName
-----
66            >10:00:00:05:33:67:26:78*  10.24.81.66   Online   switch1

```

```
switch2# show vcs
state      : Disabled
switch2# vcs vcsid 1 enable
```

- If the **show vcs** command indicates that the VCS ID is not the same on each switch, enter the **vcs vcsid** command to correct the VCS ID on the switch that is in error.

```
switch1# show vcs
Config Mode : Local-Only
VCS ID      : 1
Total Number of Nodes      : 1
Rbridge-Id  WWN              Management IP  Status  HostName
-----
66          >10:00:00:05:33:67:26:78* 10.24.81.66      Online  switch1
```

```
switch2# show vcs
Config Mode : Local-Only
VCS ID      : 2
Total Number of Nodes      : 1
Rbridge-Id  WWN              Management IP  Status  HostName
-----
66          >10:00:00:05:33:67:26:78* 10.24.81.77      Online  switch1
```

```
switch2# vcs vcsid 1
```

- If both switches have the same RBridge ID, enter the **vcs rbridge-id** command to change the RBridge ID to a unique value.

```
switch1# show vcs
Config Mode : Local-Only
VCS ID      : 1
Total Number of Nodes      : 1
Rbridge-Id  WWN              Management IP  Status  HostName
-----
66          >10:00:00:05:33:67:26:78* 10.24.81.66      Online  switch1
```

```
switch2# show vcs
Config Mode : Local-Only
VCS ID      : 1
Total Number of Nodes      : 1
Rbridge-Id  WWN              Management IP  Status  HostName
-----
66          >10:00:00:05:33:67:26:78* 10.24.81.77      Online  switch1
```

```
switch2# vcs rbridge-id 77
```

Verifying LLDP

When ISLs are functioning correctly, the **show lldp neighbors** command reports on each neighbor switch in the VCS Fabric cluster.

1. Enter the **show lldp neighbors** command to verify that LLDP reports on all of its neighbors.

```
switch1# show lldp neighbors
Local Intf  Dead Interval  Remaining Life  Remote Intf  Chassis ID  Tx  Rx
Te 66/0/55  120              106            port1        0005.1e78.f004 20300 19914
Te 66/0/60  120              108            port0        0005.1e55.16c8 20300 19911
```

2. If neighbors are missing, perform further debugging or contact your switch support provider.

Checking for CPU overload

An abnormally high CPU load can cause an ISL to malfunction. Use the **show process cpu** command as described in [“CPU use is unexpectedly high”](#) on page 615 to troubleshoot an overloaded CPU.

License not properly installed

If a licensed feature is not functioning, a probable reason is that the license for that feature has not been installed correctly. Either the license was not installed, or it was installed and a required system reboot was not performed.

If, on a Brocade VDX 6720-24 or VDX 6730-32 switch, only eight Ethernet ports are working, it is probable that no DPOD license is installed. Similarly, if on a Brocade VDX 6720-60 or Brocade VDX 6730-76, only 40 Ethernet ports are working, it is probable that no DPOD license is installed.

If, on a Brocade VDX 6720-60 or Brocade VDX 6730-76, 50 Ethernet ports are working but the remaining 10 are not, it is likely that you have the DPOD1 license installed, but not the DPOD2 license.

If you are unable to add a third switch to a VCS Fabric cluster, it is likely that the VCS Fabric license is not installed.

If you are unable to connect an FCoE device or unable to use Fibre Channel ports on a Brocade VDX 6730 switch, it is likely that the FCoE license is not installed.

For detailed licensing information, refer to [Chapter 7, “Administering Licenses”](#).

If you suspect a license is not properly installed, complete the following steps.

1. In privileged EXEC mode, enter the **show license** command to see which licenses are currently installed.

```
switch# show license

rbridge-id: 66
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
  FCoE Base license
  Feature name:FCOE_BASE
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
  VCS Fabric license
  Feature name:VCS_FABRIC
```

2. If the FCoE or DPOD license appears in the **show license** command output, but the feature does not work for the expected ports, the probable cause is that the affected ports were not re-enabled after installing the license.

NOTE

After adding an FCoE or DPOD license, you must disable and re-enable all affected ports. The VCS Fabric license does not require re-enabling.

You can disable and then enable each affected port, or you can enter the **chassis disable** command followed by the **chassis enable** command to re-enable the entire chassis.

```
switch# chassis disable
switch# chassis enable
```

3. If the license does not appear in the **show license** command output, then it was not installed. In privileged EXEC mode, enter the **license add licstr** command to install the license. For FCoE and DPOD licenses, you must also disable and enable the switch or port.

```
switch# license add licstr "*B
slSETgzTgeVGUDeQR4WIfRx7mmXODdSwENoRGENAmX3Ca3uHeZgXK0b,jzxyzfzKLrMsPN8ClSxvD
QRRT8VyuULyyKTO0ryU6qm4sljjiSAeV,COedzCx1v6ycQgnYMeSVp#"

```

```
License Added [*B
slSETgzTgeVGUDeQR4WIfRx7mmXODdSwENoRGENAmX3Ca3uHeZgXK0b,jzxyzfzKLrMsPN8ClSxvD
QRRT8VyuULyyKTO0ryU6qm4sljjiSAeV,COedzCx1v6ycQgnYMeSVp# ]

```

For license change to take effect, please disable/enable port or switch...

```
switch# chassis disable
switch# chassis enable

```

Packets dropped in hardware

This section discusses how to troubleshoot problems in which loss of packets occurs in all traffic, on specific traffic flows, in specific types of traffic, consistently, or intermittently. Dropped packets could occur for many reasons, including the following:

- High latency in an end device. Refer to [“Packets dropped because of high-latency end device”](#) on page 624.
- Broken data path. Refer to [“Verifying the data path”](#) on page 627.
- Noise on an optical line caused by too many CRC errors, packet errors, or NIC interoperability errors. Refer to [“Checking for noise on an optical line”](#) on page 629.

Packets dropped because of high-latency end device

Packets can sometimes be dropped because of buffer overrun within the fabric caused by end devices taking longer to respond than expected. For example, an overloaded disk array can cause such latency, as can a host that does not process data as quickly as expected. Devices that stop receiving data for an extended period of time can cause excessive latency.

The ultimate solution to these problems is to fix the end device itself. However, some adjustments to the switch and fabric configuration can help to reduce the problem.

To detect and relieve congestion and dropped packets resulting from latency in end devices, complete the following steps:

1. Enter the **show lldp neighbors detail** command to check under “DCBX TLVs” that the end device is DCB-ready and confirm that the end device is also advertising its DCB capabilities.

```
switch# show lldp neighbors detail
Neighbors for Interface Te 66/0/55

MANDATORY TLVs
=====
Local Interface: Te 66/0/55 (Local Interface MAC: 0005.3367.26d3)
Remote Interface: port1 (Remote Interface MAC: 0005.1e78.f004)
Dead Interval: 120 secs
Remaining Life : 104 secs
Chassis ID: 0005.1e78.f004
LLDP PDU Transmitted: 2412 Received: 2372

```



```

OPTIONAL TLVs
=====

DCBX TLVs
=====
Version : CEE
DCBX Ctrl OperVersion: 0 MaxVersion: 0 SeqNo: 1 AckNo: 4
DCBX ETS OperVersion: 0 MaxVersion: 0 Enabled: 1 Willing: 1 Error: 0
Enhanced Transmission Selection (ETS)
  Priority-Group ID Map:
    Priority : 0 1 2 3 4 5 6 7
    Group ID : 2 2 2 1 2 2 2 15
  Group ID Bandwidth Map:
    Group ID : 0 1 2 3 4 5 6 7
    Percentage: 0 40 60 0 0 0 0 0
  Number of Traffic Classes supported: 8
DCBX PFC OperVersion: 0 MaxVersion: 0 Enabled: 1 Willing: 1 Error: 0
Priority-based Flow Control (PFC)
  Enabled Priorities: 3
  Number of Traffic Class PFC supported: 8
FCoE App OperVersion: 0 MaxVersion: 0 Enabled: 1 Willing: 1 Error: 0
FCoE Application Protocol
  User Priorities: 3

```

2. Enter the **show qos flowcontrol interface** command to check for pause frames.

```

switch# show qos flowcontrol interface tengigabitethernet 66/0/55
Interface TenGigabitEthernet 66/0/55
Mode PFC
DCBX enabled for PFC negotiation
TX 4926331124 frames

```

CoS	TX Admin	TX Oper	RX Admin	RX Oper	Output Paused 512 BitTimes
0	Off	Off	Off	Off	0
1	Off	Off	Off	Off	0
2	Off	Off	Off	Off	0
3	On	On	On	On	0
4	Off	Off	Off	Off	0
5	Off	Off	Off	Off	0
6	Off	Off	Off	Off	0
7	Off	Off	Off	Off	0

3. Enter the **show qos queue interface** command to check the CoS statistics.

```

switch# show qos queue interface tengigabitethernet 66/0/60
Interface TenGigabitEthernet 66/0/60

```

CoS	RX Packets	RX Bytes	TC	TX Packets	TX Bytes
0	1600	354184	0	0	0
1	0	0	1	7962	636960
2	0	0	2	0	0
3	8508	544832	3	18	6048
4	0	0	4	0	0
5	0	0	5	0	0
6	0	0	6	0	0
7	0	0	7	2123	282360
untag	2082	216528			

4. Enter the **show qos rcv-queue interface** command to check for indicators of congestion, including dropped packets, buffer consumption, and real-time queue statistics.

```
switch# show qos rcv-queue interface tengigabitethernet 66/0/55
Interface TenGigabitEthernet TenGigabitEthernet 66/0/55
In-use 27216 bytes, Total buffer 144144 bytes
0 packets dropped
```

TC	In-use Bytes	Max Bytes
0	0	252
1	0	252
2	0	252
3	27216	75284
4	0	252
5	0	252
6	0	57456
7	0	9576

5. Enter the **show qos interface** command to check the QoS configuration.

```
switch# show qos interface tengigabitethernet 66/0/55
Interface TenGigabitEthernet 66/0/55
Provisioning mode cee
Priority Tag disable
CEE Map default
FCoE CoS: 3
FCoE Provisioned
Default CoS 0
Interface trust cos
      In-CoS:  0  1  2  3  4  5  6  7
-----
      Out-CoS/TrafficClass: 0/6 1/6 2/6 3/3 4/6 5/6 6/6 0/7
Per-Traffic Class Tail Drop Threshold (bytes)
      TC:    0  1  2  3  4  5  6  7
-----
      Threshold:  252  252  252 75284  252  252 57456  9576
Flow control mode PFC
      CoS3 TX on, RX on
Multicast Packet Expansion Rate Limit 3000000 pkt/s, max burst 4096 pkts
Multicast Packet Expansion Tail Drop Threshold (packets)
      TrafficClass:  0  1  2  3  4  5  6  7
-----
      Threshold:  64  64  64  64  64  64  64  64
Traffic Class Scheduler configured for 1 Strict Priority queues
      TrafficClass:  0  1  2  3  4  5  6  7
-----
      DWRRWeight:  0  0  0 40  0  0 60 ---
Multicast Packet Expansion Traffic Class Scheduler
      TrafficClass:  0  1  2  3  4  5  6  7
-----
      DWRRWeight:  12 13 12 13 12 13 12 13
```

6. Reconfigure QoS. Refer to [Chapter 29, “Configuring QoS”](#).

Verifying the data path

This procedure checks whether fabric continuity might be the reason for dropped packets.

1. Enter the **ping** command to test for a complete path to the end device

```
switch# ping dest-address 10.24.81.2
PING 10.24.81.2 (10.24.81.2): 56 octets data
64 octets from 10.24.81.2: icmp_seq=0 ttl=128 time=9.4 ms
64 octets from 10.24.81.2: icmp_seq=1 ttl=128 time=0.3 ms
64 octets from 10.24.81.2: icmp_seq=2 ttl=128 time=0.3 ms
64 octets from 10.24.81.2: icmp_seq=3 ttl=128 time=0.3 ms
64 octets from 10.24.81.2: icmp_seq=4 ttl=128 time=0.3 ms

--- 10.24.81.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0.3/2.1/9.4 ms
```

2. Enter the **show interface** command to see whether packets are coming in or are dropped as errors. Specifically, examine the output fields shown underlined in the following example.

```
switch# show interface tengigabitethernet 66/0/60
TenGigabitEthernet 66/0/60 is up, line protocol is up (connected)
Hardware is Ethernet, address is 0005.3367.26d8
  Current address is 0005.3367.26d8
Pluggable media present
Interface index (ifindex) is 283874428169
MTU 2500 bytes
LineSpeed Actual    : 10000 Mbit
LineSpeed Configured : Auto, Duplex: Full
Flowcontrol rx: off, tx: off
Last clearing of show interface counters: 22:07:59
Queueing strategy: fifo
Receive Statistics:
  15254 packets, 1395269 bytes
  Unicasts: 10641, Multicasts: 2637, Broadcasts: 1976
  64-byte pkts: 10874, Over 64-byte pkts: 3294, Over 127-byte pkts: 117
  Over 255-byte pkts: 969, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
  Over 1518-byte pkts(Jumbo): 0
  Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
  Errors: 0, Discards: 0
Transmit Statistics:
  12633 packets, 1155963 bytes
  Unicasts: 18, Multicasts: 12615, Broadcasts: 0
  Underruns: 0
  Errors: 0, Discards: 0
Rate info (interval 299 seconds):
  Input 0.000128 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 1d00h40m
```

3. Enter the **show media interface** command to check that the optics used are Brocade-certified. Check the Vendor Name field shown underlined in the following example.

Check also the TX Power and RX Power fields to ensure they are not zero.

```
switch# show media interface tengigabitethernet 66/0/60
Interface   TenGigabitEthernet 66/0/60
Identifier   3   SFP
Connector    7   LC
Transceiver  00000000000000010 10_GB/s
```

```

Name      id
Encoding   6
Baud Rate 103 (units 100 megabaud)
Length 9u  0 (units km)
Length 9u   0 (units 100 meters)
Length 50u  8 (units 10 meters)
Length 62.5u 3 (units 10 meters)
Length Cu   0 (units 1 meter)
Vendor Name BROCADE
Vendor OUI  00:05:1e
Vendor PN   57-0000075-01
Vendor Rev  A
Wavelength 850 (units nm)
Options     001a
BR Max      0
BR Min      0
Serial No   AAA209282044472
Date Code   090709
Temperature 35 Centigrade
Voltage     3356.4 (mVolts)
Current     5.564 (mAmps)
TX Power    568.9 (uWatts)
RX Power    549.9 (uWatts)

```

If the Vendor Name field shows anything other than BROCADE, replace the optics with Brocade-certified optics.

4. Enter the **show mac-address-table** command to verify that the MAC address table learns new values.

The new MAC address should appear here.

```

switch# show mac-address-table
VlanId  Mac-address  Type  State  Ports
1002    0efc.0042.7300  FPMA  Active Te 66/0/55
1002    0efc.0042.7302  FPMA  Active Te 66/0/55
1002    0efc.0042.7800  FPMA  Active Te 66/0/60
Total MAC addresses : 3

```

5. Enter the **show lldp neighbors** command to verify that LLDP reports all neighbors.

```

switch# show lldp neighbors
Local Intf  Dead Interval  Remaining Life  Remote Intf  Chassis ID  Tx  Rx
Te 66/0/55  120           101            port1       0005.1e78.f004 3000 2948
Te 66/0/60  120           117            port0       0005.1e55.16c8 2999 2945

```

If the output does not show all neighbors, contact your switch support provider.

6. Enter the **show mac-address-table** command to verify the Ethernet Name Service functionality and to see whether MAC addresses learned from other VCS Fabric switches are present.

Enter this command on other switches in the fabric to ensure that those switches can see this MAC address.

```

switch# show mac-address-table
VlanId  Mac-address  Type  State  Ports
1002    0efc.0042.7300  FPMA  Active Te 66/0/55
1002    0efc.0042.7302  FPMA  Active Te 66/0/55
1002    0efc.0042.7800  FPMA  Active Te 66/0/60
Total MAC addresses : 3

```

7. Enter the **l2tracert** command to validate the data-path fabric continuity.
 - Enter dynamically learned source MAC address and destination MAC address for the data path.
 - Among the extended commands, use IP, SIP, DIP, TCP, Src Port, and Dest Port commands.
 - Enter the IP command parameters to ensure that the traceroute packet traverses a specific ECMP link.

For details on using the **l2tracert** command, refer to “[Layer 2 traceroute](#)” on page 644.

Checking for noise on an optical line

Excessive noise on an optical line can result in dropped packets because of excessive CRC errors, NIC interoperability errors, or other conditions.

1. Enter the **show interface** command and check the output for CRC errors or TX discards; examine the fields shown underlined in the following example.

```
switch# show interface tengigabitethernet 66/0/55
TenGigabitEthernet 66/0/55 is up, line protocol is up (connected)
Hardware is Ethernet, address is 0005.3367.26d3
  Current address is 0005.3367.26d3
Pluggable media present
Interface index (ifindex) is 283874100484
MTU 2500 bytes
LineSpeed Actual    : 10000 Mbit
LineSpeed Configured : Auto, Duplex: Full
Flowcontrol rx: off, tx: off
Last clearing of show interface counters: 21:51:35
Queueing strategy: fifo
Receive Statistics:
  15433457505 packets, 32164575799774 bytes
  Unicasts: 15433454934, Multicasts: 2571, Broadcasts: 0
  64-byte pkts: 11357, Over 64-byte pkts: 242664576, Over 127-byte pkts: 0
  Over 255-byte pkts: 0, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
  Over 1518-byte pkts(Jumbo): 15190781568
  Runt: 0, Jabbers: 0, CRC: 0, Overruns: 0
  Errors: 0, Discards: 0
Transmit Statistics:
  21456965161 packets, 32549136821934 bytes
  Unicasts: 15313174675, Multicasts: 6143790486, Broadcasts: 0
  Underruns: 0
  Errors: 0, Discards: 0
Rate info (interval 299 seconds):
  Input 3345.136864 Mbits/sec, 200572 packets/sec, 33.45% of line-rate
  Output 3386.493904 Mbits/sec, 281345 packets/sec, 33.86% of line-rate
Time since last interface status change: 1d00h24m
```

2. If errors are reported in the previous step, check the SFP transceiver and cable on the local switch and on the peer switch at the other end of the cable.
 - a. Enter the **show media interface** command on each switch and check the Vendor Name field to check that the optics are Brocade-certified.
Replace any non-Brocade SFP transceiver.
 - b. Try replacing the SFP transceiver.
 - c. Try replacing the cable.

Password recovery for the Brocade VDX 8770

Use this procedure if you do not have access to the root account.

NOTE

If the **ESC** key is not effective during reboot, turn the power off and back on, and then try again. If the **ESC** key is still not effective, check the serial console cable. If the cable is connected correctly, then the unit must be returned for service or repair.

To reset the root password to its factory default value on a Brocade VDX 8770-4 or 8770-8 switch, set a password for the admin account, and then restore nondefault user accounts following these steps:

Use this procedure if you do not have access to the root account.

To reset the root password to its factory default value on a Brocade VDX 8770-4 or 8770-8 switch, set a password for the admin account, and then restore nondefault user accounts by following these steps:

1. Connect to the serial console port of the switch.
2. Manually reboot the switch.
3. When prompted to stop test or stop AutoBoot, press **ESC**. The Boot PROM menu is displayed with the following options:
 - Start system.
Used to reboot the system.
 - Recover password.
Used to generate a character string for your support provider to recover the Boot PROM password. Use this feature only when directed by technical support personnel.
 - Enter command shell.
Used to enter the command shell to reset all passwords on the system.

```
Checking system RAM - press any key to stop test
Checking memory address: 00100000
System RAM test terminated by keyboard
set_bootstatus: BS_LOAD_OS, platform_idx = 6
Hit ESC to stop autoboot: 0
1) Start system.
2) Recover password.
3) Enter command shell.
Option?
```
4. Type **3** at the prompt to open the command shell:
5. Type the Boot PROM password, if prompted, and press **Enter**.


```
Network OS Password Recovery Notes 9
53-1002130-02
Root password recovery using the Boot PROM interface for Brocade 8770
platforms 1
The Boot PROM has a password only if one was defined.
```

NOTE

If you are prompted to enter a new Boot PROM password, make sure it is at least 8 characters long. Do not select this option unless specifically instructed by support personnel.

6. Append "S" to the boot arguments so that the switch boots into the single user mode. Enter the

following command:

```
=> setenv bootargs "root=/dev/sda1 rootfstype=ext4 quiet S"
```

7. Enter the **printEnv** command to verify the change:

```
=> printenv
AutoLoad=yes
LoadIdentifiers=Fabric Operating System;Fabric Operating System
OSLoadOptions=quiet
OSRootPartition=sda2;sda1
SkipWatchdog=yes
autoreset_mac=true
baudrate=9600
bootargs=root=/dev/sda1 rootfstype=ext4 quiet S
bootcmd=execute_internal_bootcmd
(output truncated)
```

8. Save the changes:

```
=> saveenv
Saving Environment to Flash.....Done
```

9. Enter the **reset** command to bring up the device in the single-user mode:

```
=> reset
BootROM version: 1.0.48
Copyright (C) 2011 Brocade Communication.
CPU0: P4080E, Version: 2.0, (0x82080020)
(output truncated)
```

10. Enter the **mount** command with the following parameters to remount the root partition as read/write capable:

```
sh-2.04# mount -vo remount,rw,noatime /
/dev/root on / type ext4 (rw,noatime)
```

11. Mount the secondary partition. Examine the output of the **printEnv** command in step 7 to check which partition the root points to in the boot arguments (bootargs=root setting). If the root partition is sda2, then use sda1 in this command. If the root partition is sda1, use sda2.

```
sh-2.04# mount /dev/sda2 /mnt
```

12. Enter the **passwddefault** command to reset the root password to its factory default value as follows:

```
sh-2.04# /sbin/passwddefault
```

NOTE

For the Brocade Network OS, the **passwddefault** command restores the passwords of factory default accounts to their default values and removes nondefault user accounts that are present. Error messages seen during the execution of **passwddefault** (applicable to Network OS v3.0.0) should be ignored.

In a dual management module (MM) chassis, the **passwddefault** command should be entered on the standby MM for password recovery.

13. Reset the boot arguments by removing "S".

```
sh-2.04# bootenv bootargs "root=/dev/sda1 rootfstype=ext4 quiet"
```

14. Reboot the switch using the **partman -r** command.

```
sh-2.04# partman -r
```

15. Log on to the switch over a serial interface or Telnet connection and access any of the factory default accounts (root/admin/user) by entering the default password.

16. Start the Network OS command line:

```
switch:root> noscli
WARNING: The default password of 'admin' and 'user' accounts have not been
changed.
Welcome to the Brocade Network Operating System Software
admin connected from 127.0.0.1 using console on switch
```

17. Enter global configuration mode:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)#
```

18. Use the following syntax of the **username** command to reset passwords for the admin or user accounts, or for any other nondefault users:

username account-name password new-password

The following example resets the admin password to the default value of “password”:

```
switch(config)# username admin password password
```

19. To restore the nondefault user accounts, perform the following steps.

- a. Copy the running-config to a file:

```
switch# copy running-config flash://running-config.cfg
2012/07/09-11:51:21, [DCM-1108], 4930, M2, INFO, VDX8770-4, Running
configuration file has been uploaded successfully to the remote location.
```

- b. Copy the default-config to the startup-config to reset the startup-config:

```
switch# copy default-config startup-config
```

- c. Reboot the switch:

```
switch# reload
Warning: Unsaved configuration will be lost. Please run `copy
running-config startup-config` to save the
current configuration if not done already.
Are you sure you want to reload the switch? [y/n]:y
The system is going down for reload NOW !!
```

- d. Copy the file saved in step 19a to the running-config:

```
switch# copy flash://running-config.cfg running-config
Loading.
2012/07/09-12:08:13, [DCM-1105], 5456, M2, INFO, VDX8770-4, Copy of the
downloaded config file to the current running-config has completed
successfully on this node.
```

- e. Copy the running-config to the startup-config:

```
switch# copy running-config startup-config
The password recovery procedure is now complete. YYou can now use normal
password management procedures from the admin account.
```


Ping failure

If pings do not successfully traverse the switch, try the following operations.

1. Trace the packet flow and check whether ARP or ICMP packets are getting dropped.
2. Trace which direction is failing by using interface statistics.
3. Locate the device that is dropping the packets.
4. Look for any error counters incrementing on that device.
5. Check the MAC address table to determine whether the MAC addresses are learned on the correct port or port-channel.

QoS configuration causes tail drops

Tail-drop queueing is the most basic form of congestion control. Normal operation is first-in, first-out (FIFO) until all buffers are exhausted. After that, new frames are dropped. You can reduce the impact of such drops by configuring thresholds for each COS priority through the **qos rcv-queue multicast threshold** command. Refer to [Chapter 29, “Configuring QoS”](#).

QoS is not marking or treating packets correctly

Use the Switched Port Analyzer (SPAN) feature to mirror the ingress and egress ports to check that QoS is marking and treating packets correctly. Refer to [Chapter 32, “Configuring Switched Port Analyzer,”](#) for details.

RBridge ID is duplicated

Switches with the same RBridge ID cannot coexist in the same VCS Fabric cluster. Any attempt to add a switch with the same RBridge ID as an existing cluster switch will fail. The ISL between the two switches will not be formed; it will be segmented.

1. On the new switch, enter the **show vcs** command to determine the RBridge ID.

```
switch2# show vcs
Config Mode   : Local-Only
VCS ID        : 1
Total Number of Nodes      : 1
Rbridge-Id     WWN                Management IP  Status  HostName
-----
22             >10:00:00:05:33:13:B3:5A*  10.24.84.41   Online
```

2. On any switch in the functioning VCS Fabric cluster, enter the **show vcs** command to see the RBridge IDs of all the switches in the cluster.

```
switch1# show vcs
Config Mode   : Local-Only
VCS ID        : 1
Total Number of Nodes      : 2
Rbridge-Id     WWN                Management IP  Status  HostName
-----
60             10:00:00:05:33:5F:EA:A4   10.24.81.65   Online  switch1
66             >10:00:00:05:33:67:26:78*  10.24.81.66   Online  switch 2
```

3. If the new switch has the same RBridge ID as any switch in the existing cluster, on the new switch, in privileged EXEC mode, enter the **vcs rbridge-id** command to change the RBridge ID to a unique value.

```
switch2# vcs rbridge-id 77
```

SNMP MIBs report incorrect values

If SNMP MIBs report incorrect values, complete the following steps.

1. Ensure you are using a supported MIB browser.
2. Ensure that the issue is seen consistently.
3. Ensure that the SNMP configuration is correct.
4. If the MIB browser is supported, the SNMP configuration is correct, and you still see the issue consistently, contact your switch support provider.

SNMP traps are missing

If SNMP traps are missing, complete the following steps.

1. Ensure that the correct SNMP configuration is enabled. Refer to [Chapter 8, “SNMP,”](#) for details.
2. Ensure that the SNMP host is reachable.
3. If the problem still persists, contact your switch support provider.

As a workaround, set a trap configuration for syslog messages.

Telnet operation into the switch fails

Assuming a correct IP address and correct login credentials, failure to access the switch using Telnet could be for one of the following reasons:

- The management port is down. Refer to [“Verifying the status of the management port”](#) on page 634 for details.
- Access to the management interface is denied by an ACL. Refer to [“Checking for a deny ACL”](#) on page 635 for details.
- The switch CPU is overloaded. Refer to [“Checking for overloaded CPU”](#) on page 635 for details.

Verifying the status of the management port

1. On the system console, enter the **show system** command to check the status of the management port, shown underlined in the following example.

```
switch# show system
Stack MAC                : 00:05:33:67:26:78

-- UNIT 0 --
Unit Name                 : switch
Switch Status             : Online
Hardware Rev              : 107.4
TengigabitEthernet Port(s) : 60
Up Time                   : up 1 day, 2:52
Current Time              : 23:40:50 GMT
```

```

NOS Version           :
Jumbo Capable         : yes
Burned In MAC         : 00:05:33:67:26:78
Management IP         : 10.24.81.66
Management Port Status : UP

```

```

-- Power Supplies --
PS1 is faulty
PS2 is OK

```

```

-- Fan Status --
Fan 1 is Ok
Fan 2 is Ok
Fan 3 is Ok

```

2. If the status of the management port is DOWN, enter the **interface management** command to configure the management port correctly. Refer to [“Configuring the Ethernet management interface”](#) on page 49.
3. If the problem persists, contact your switch support provider.

Checking for a deny ACL

On the system console, enter the **show running-config ip access-list** command and check the output to determine whether an ACL is denying access to the management port.

Checking for overloaded CPU

An overloaded switch CPU can prevent Telnet access. Refer to [“CPU use is unexpectedly high”](#) on page 615.

Trunk member not used

If you suspect that one or more members of a trunk are not being used, complete the following steps.

1. Enter the **show running-config interface** command to determine which interfaces have trunking enabled.

```

switch# show running-config interface
interface Management 66/0
  no ip address dhcp
  ip address 10.24.81.66/20
  ip route 0.0.0.0/0 10.24.80.1
  ipv6 address ""
  no ipv6 address autoconfig
!
interface TenGigabitEthernet 66/0/1
  fabric isl enable
  fabric trunk enable
  no shutdown
!
interface TenGigabitEthernet 66/0/2
  fabric isl enable
  fabric trunk enable
  no shutdown
!

```

```

interface TenGigabitEthernet 66/0/3
  fabric isl enable
  fabric trunk enable
  no shutdown
!
(output truncated)

```

2. Verify the status of the ISL port and link,
 - a. Enter the **show fabric isl** command to verify whether the ISL is up.
 - b. Enter the **show fabric islports** command to examine the status of each port.

Refer to “[Verifying the status of ISLs](#)” on page 620 for details and corrective action.

3. Enter the **show interface** command for each trunk link and examine the rate information to check for an equal distribution of traffic on the interfaces in the trunk. The rate information is shown underlined in the following example.

```

switch# show interface tengigabitethernet 66/0/12
TenGigabitEthernet 66/0/12 is up, line protocol is down (link protocol down)
Hardware is Ethernet, address is 0005.3367.26a8
  Current address is 0005.3367.26a8
Pluggable media not present
Interface index (ifindex) is 283871281409
MTU 2500 bytes
LineSpeed Actual   : Nil
LineSpeed Configured : Auto, Duplex: Full
Flowcontrol rx: off, tx: off
Last clearing of show interface counters: 1d00h42m
Queueing strategy: fifo
Receive Statistics:
  0 packets, 0 bytes
  Unicasts: 0, Multicasts: 0, Broadcasts: 0
  64-byte pkts: 0, Over 64-byte pkts: 0, Over 127-byte pkts: 0
  Over 255-byte pkts: 0, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
  Over 1518-byte pkts(Jumbo): 0
  Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
  Errors: 0, Discards: 0
Transmit Statistics:
  0 packets, 0 bytes
  Unicasts: 0, Multicasts: 0, Broadcasts: 0
  Underruns: 0
  Errors: 0, Discards: 0
Rate info (interval 299 seconds):
  Input 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 1d03h16m

```

4. Having found a trunk member that carries no traffic while the other trunk members are busy, from the same **show interface** command output, check the interface status, configuration, and error statistics.
 - If the interface is disabled, enable it with the **no shutdown** command.
 - If misconfiguration is apparent, refer to [Chapter 9, “Fabric,”](#) for information on how to configure fabric trunks.
 - If you see significant errors in the error statistics counters, depending on the error, check the SFP transceiver and cable on the local switch and on the peer switch at the other end of the cable.

- a. Enter the **show media interface** command on each switch and check the Vendor Name field to check that the optics are Brocade-certified.
- b. Replace any non-Brocade SFP transceiver.
- c. Try replacing the SFP transceiver.
- d. Try replacing the cable.

Upgrade failure

If a failure occurs during firmware upgrade, complete the following steps.

1. Revert to the previous firmware version.
2. Contact your switch support provider to evaluate whether retrying the upgrade is appropriate.

VCS Fabric cannot be formed

A VCS Fabric can fail to form for several reasons:

- The required licenses are not active. Refer to [“Verifying VCS Fabric licenses”](#) on page 637.
- The VCS Fabric configuration is incorrect. The following configuration issues will prevent the VCS Fabric from forming:
 - VCS Fabric mode has not been enabled.
 - The VCS ID on the constituent switches is not the same.
 - Multiple switches have the same RBridge ID.
 - ISL ports that connect the switches are not up.

Refer to [“Verifying the VCS Fabric configuration”](#) on page 638.

Verifying VCS Fabric licenses

If the VCS Fabric cluster has just one or two switches, no VCS Fabric license is required. For more than two switches to exist in a VCS Fabric cluster, you must have the VCS Fabric license installed.

1. Enter the **show license** command to check whether the required VCS Fabric license is installed.

```
switch# show license
rbridge-id: 66
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
FCoE Base license
Feature name:FCOE_BASE
```

2. If the VCS Fabric license is not listed in the output of the **show license** command, enter the **license add licstr** command to enable the license.

```
switch# license add licstr "*B
r84pNRtHKdRZuJmwAUT63GORXIpBhBZK0ckRq6Bvv13Strvw1:fUjANF
av5W:gWx3hH2:9RsMv3BHfeCRFM2gj9NlkrdIiBPBOa4xfSD2jf,Xx1RwksliX8fH6gpx7,73t#"

Adding license [*B r84pNRtHKdRZuJmwAUT63GORXIpBhBZK0ckRq6Bvv13Strvw1:fUjANF
av5W:gWx3hH2:9RsMv3BHfeCRFM2gSLj9NlkrdIiBPBOa4xfSD2jf,Xx1RwksliX8fH6gpx7,73t#]
```

3. Enter the **show license** command again to verify that the license has been added.

NOTE

It is not necessary to reboot the switch to enable the VCS Fabric license.

```
switch# show license
Rbridge-Id: 66
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
    FCoE Base license
    Feature name:FCOE_BASE
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
    VCS Fabric license
    Feature name:VCS
```

Refer to [Chapter 7, “Administering Licenses,”](#) for more information about license management.

Verifying the VCS Fabric configuration

To verify the VCS Fabric configuration, complete the following steps.

1. Enter the **show vcs** command on each switch to verify that the VCS Fabric mode is enabled, the VCS ID on each switch is the same, and the RBridge ID on each switch is different.
2. Enter the **show fabric isl** command to verify whether the ISL is up.
3. Enter the **show fabric islports** command to examine the status of each port.

Refer to [“ISL does not come up on some ports”](#) on page 620 for details and corrective action.

vLAG cannot be formed

A vLAG trunk can fail to form for several reasons:

- The link between the VCS Fabric switches does not exist. Refer to [“Verifying the link between the VCS Fabric switches”](#) on page 638.
- A bad connection causes abnormal reception or transmission of LACPDUs. Refer to [“Verifying LACPDUs”](#) on page 639.
- Port-channel numbers are not the same on the VCS Fabric switches. Refer to [“Verifying the vLAG configuration”](#) on page 639.
- The peer switches are not configured in the same LACP mode (static or dynamic). Refer to [“Verifying the LACP mode of each switch”](#) on page 639.
- A 1-Gbps port-channel has been upgraded to Network OS 2.1.x. Refer to [“Explicitly setting the speed for a 1-Gbps port-channel”](#) on page 640.

Verifying the link between the VCS Fabric switches

The link between switches could be broken for various reasons:

- A port is not activated.
- The ISL is segmented.
- The VCS Fabric is not properly formed.
- The CPU is overload.

Refer to [“ISL does not come up on some ports”](#) on page 620 for details on detecting and correcting the problem.

Verifying LACPDUs

LACPDUs should be transmitted and received on both ends of the vLAG. This procedure verifies whether that is happening, and also checks for PDU errors.

1. On both switches, enter the **show lacp counter** command to verify that LACPDUs are transmitted and received, and there are no error PDUs.

```
switch# show lacp counter 10
% Traffic statistics
Port          LACPDU          Marker          Pckt err
              Sent    Recv    Sent    Recv    Sent  Recv
% Aggregator  Po 10 1000000
Te 0/1      65      0      0      0      0      0
Te 0/2      64      0      0      0      0      0
Te 0/3      64      0      0      0      0      0
Te 0/4       0      0      0      0      0      0
```

In this case, LACPDUs are being transmitted by the switch, but none are being received.

2. If the output shows that LACPDUs are not being transmitted and received correctly, or packet errors are showing, contact your switch support provider.

Verifying the vLAG configuration

The port-channel number must be the same across all vLAG member switches, or the vLAG will not form.

1. On each vLAG member switch, in privileged EXEC mode, enter the **show port-channel summary** command.

```
switch# show port-channel summary
Static Aggregator: Po 15
Aggregator type: Standard
Member ports:
Te 0/6
Te 0/7
Te 0/14
Te 0/15
...
switch2# show port-channel summary
switch2#
```

2. If the port-channel does not appear on both switches, on the switch where it does not appear, in global configuration mode, enter the **interface port-channel** command to create the port-channel.

```
switch2(config)# interface port-channel 15
```

Refer to [Chapter 26, “Configuring Link Aggregation,”](#) for details.

Verifying the LACP mode of each switch

A vLAG must be configured either statically on both ends of the vLAG, or dynamically on both ends of the vLAG. Refer to [Chapter 26, “Configuring Link Aggregation,”](#) for details.

Explicitly setting the speed for a 1-Gbps port-channel

To set the port speed to 1 Gbps, complete the following steps.

1. In interface configuration mode, shut down the port-channel.

```
switch(config-Port-channel-2)# shutdown
```

2. Set the port-channel speed to 1 Gbps.

```
switch(config-Port-channel-2)# speed 1000
```

3. Re-enable all port members in the port-channel.

```
switch(config-Port-channel-2)# no shutdown
```

Zoning conflict resolution

When merging two fabrics, multiple zoning CLI sessions can be launched on the same switch, or on different switches. This section describes these situations and how they are automatically resolved.

Dual-CLI sessions from the same switch: If you start a zone transaction from CLI-Session1 and then try to perform a zone modification from CLI-Session2, the CLI-Session2 zone transaction is not allowed due to CLI-Session2 not being the owner of the open transaction. If CLI-Session1 logs out, this ends the open transaction and aborts any current zone modifications. CLI-Session2 is then able to perform zone modifications. Therefore, the zone transaction locking mechanism works on a single switch from the CLI perspective and there is no dangling transaction.

Dual-CLI sessions from different switches: If you start a CLI zone transaction on Switch1 and started another CLI zone transaction on Switch2, when committing the zone transaction from Switch1, the open zone transaction from Switch2 is aborted by Switch1. The following message is posted on Switch2 at the time of zone commit from Switch1:

```
2012/03/09-21:45:26, [ZONE-1027], 3285, FID 128, INFO, switch, Zoning transaction
aborted Zone Config update Received
```

Zone does not form correctly

Some problems you might encounter when configuring zones include potential Fibre Channel router issues. For a more detailed discussion of possible Fibre Channel issues, refer to the *Fabric OS Troubleshooting and Diagnostics Guide*.

Some of the following problems may contribute to the zone not forming correctly:

- A Brocade VDX switch gets isolated when an RBridge ID matches the front domain ID or translate domain ID in a mixed network. Refer to [“Recovering an isolated switch in a mixed FCoE fabric”](#) on page 641.
- A “FID over-subscribed” message occurs during attempts to connect a backbone fabric to an edge fabric. Refer to [“Recovering from FID oversubscription”](#) on page 641.
- A “FID conflict” message occurs during attempts to connect a backbone fabric to an edge fabric. Refer to [“Recovering from Fabric ID conflict”](#) on page 642.
- Interfabric link (IFL) traffic does not flow over the intended link. Refer to [“Rebalancing traffic over multiple IFLs”](#) on page 642.
- Zone merge was expected to be blocked following reboot, but was not blocked. Refer to [“Blocking zone merge after reboot”](#) on page 642.

- Stale translate domains exist in an edge fabric. Refer to [“Removing stale translate domains”](#) on page 643.

Recovering an isolated switch in a mixed FCoE fabric

In an FCoE fabric that spans Network OS switches and Fabric OS switches, a Network OS switch with an RBridge ID that matches a front phantom domain ID or translate phantom domain ID of a connecting Fibre Channel router can become isolated.

FCoE connectivity across the Fibre Channel link between VCS Fabric clusters and Fibre Channel routers uses domain IDs to identify switches. Within a VCS Fabric cluster, a domain ID is the same as an RBridge ID. When you connect to a Fibre Channel router, the Fibre Channel router service in the Fibre Channel fabric emulates virtual phantom Fibre Channel domains in the FCoE fabric. Each Fibre Channel router-enabled switch emulates a single front phantom domain and each FC fabric is represented by a translate phantom domain.

To recover an isolated Network OS switch, complete the following steps.

1. Disable all FC routers that connect to the VCS Fabric cluster.
2. Reboot the isolated Network OS switch.
3. Re-enable all disabled FC routers.

To prevent switch isolation, follow these steps on each FC router that attaches to a VCS Fabric cluster.

1. Enter the **portCfgExPort -d** Fabric OS command to set a unique front phantom domain ID.
2. Enter the **forXlateConfig importedFID exportedFID preferredDomainID** command to set a unique translate phantom domain ID.

Refer to the *Fabric OS Command Reference* for details about the **portCfgExPort** and **forXlateConfig** commands.

Recovering from FID oversubscription

A “FID over-subscribed” message occurs when different Fibre Channel backbones attempt to connect to the same edge fabric using different Fabric IDs (FIDs). When you assign a FID to the edge fabric (**portCfgExPort -f** command), you must use the same FID as any other Fibre Channel backbone that connects to the edge fabric.

To resolve this problem, complete the following steps.

1. On the Fibre Channel router on the backbone with the errant FID configured, disable the EX_Port.
2. Enter the **portCfgExPort -f** command to configure the EX_Port with the same FID as the EX_Port on the other Fibre Channel router that connects to the same edge fabric.
3. Re-enable the EX_Port.

Refer to the “Configuring an IFL for both edge and backbone connections” section in the *Fabric OS Administrator’s Guide* for details.

Recovering from Fabric ID conflict

The “FID conflict” message occurs when a backbone fabric connects to two or more edge fabrics that have the same Fabric ID (FID). Every edge fabric that a Fibre Channel router connects to must have a Fabric ID configured for the EX_Port that is unique on that Fibre Channel backbone. This error is most likely to occur when an edge fabric temporarily splits, causing it to appear as two edge fabrics with the same Fabric ID. This symptom might occur during VCS Fabric or Fibre Channel fabric upgrade, or as a result of a Brocade VDX or Fibre Channel switch reboot or crash.

Problem resolution depends on the cause of the problem. If the error is due to a temporary split, the problem will go away when the fabrics merge again.

If the problem is not due to a temporary fabric split, the most likely cause is misconfiguration. In this case, enter the **portCfgExPort -f** command to reconfigure one of the EX_Ports with a unique fabric ID.

Rebalancing traffic over multiple IFLs

If traffic across multiple interfabric links (IFLs) between a Fibre Channel router and an edge fabric is not balanced as you intended, it may be because the Fibre Channel router cannot determine an FSFP path from the Fibre Channel backbone to the target in the edge fabric. It uses all paths.

To direct the traffic the way you intend, on the FC router, use the **fcrRouterPortCost** command to configure a cost for each IFL. Traffic will flow across the lowest-cost IFL.

1. Connect to the FC router and log in using an account with admin permissions.
2. Disable the EX_Port.
3. Enter the **fcrRouterPortCost** command to configure the link cost. Set the cost to 1000 if you want the link to carry traffic during normal operation.

If you want the link to not carry traffic under normal operation, set the cost to 10000 and set the cost of at least one other link to 1000. The default value is 1000, which you get when you enter a value of 0.

4. Re-enable the port.

For details about the **fcrRouterPortCost** command, refer to the *Fabric OS Command Reference*.

Blocking zone merge after reboot

To be sure of blocking zone merge following a switch reboot, enter the **no fabric isl enable** command to disable the ISL between neighboring Brocade VDX switches.



CAUTION

Brocade recommends that you do NOT use the shutdown command. If you use the shutdown command, then following switch reboot, the zone merge could happen before the shutdown command is replayed by the running configuration.

To block zone merge following reboot, follow these steps on each ISL port.

1. In global configuration mode, enter the **interface tengigabitethernet** (or **interface gigabitethernet**) command to enter interface configuration mode.
2. Enter the **no fabric isl enable** command.

Removing stale translate domains

A translate domain becomes stale when the edge fabric it represents becomes unreachable. By default, the stale translate domain is not deleted until the local edge fabric is rebuilt.

To delete a stale translate domain and avoid the disruption caused by rebuilding the local edge fabric, complete the following steps.

1. Connect to the FC router and log in using an account with admin permissions.
2. On the FC router, enter the **fcrXlateConfig –show stalexd** command to list any stale translate domains.
3. Enter the **fcrXlateConfig –delete stalexd** command to delete the stale translate domain.

Refer to the *Fabric OS Command Reference* for details about the **fcrXlateConfig** command.

Troubleshooting and diagnostic tools

This section describes the various troubleshooting and diagnostic tools available with Brocade Network OS and provides some guidelines for their use:

- [“Layer 2 traceroute”](#) on page 644
- [“show commands”](#) on page 648
- [“Debug commands”](#) on page 649
- [“SPAN port and traffic mirroring”](#) on page 650
- [“Hardware diagnostics”](#) on page 651
- [“Viewing routing information”](#) on page 652
- [“Using the packet capture utility”](#) on page 652

Refer also to [“Gathering troubleshooting information”](#) on page 599, which provides information about Network OS supportSave files.

Layer 2 traceroute

TRILL OAM provides the **I2traceroute** command to verify the fabric path continuity. When the **I2traceroute** command is used with extended options, it provides granular control over the Layer 2 path that a Layer 2 traceroute packet takes.

Layer 2 traceroute packets

To use the Layer 2 traceroute tool, you need to understand the structure of the Layer 2 traceroute packet when observed on the wire, when it is a request frame, and when it is a response frame.

Figure 62 shows what a normal Layer 2 packet looks like when traversing through an Ethernet fabric, without Layer 2 traceroute applied.

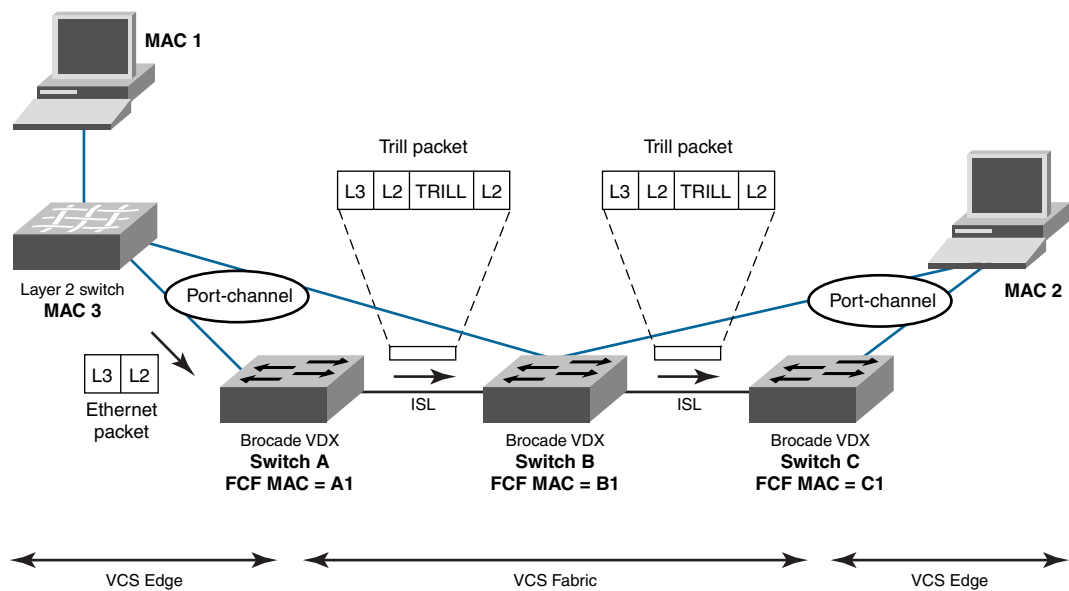


FIGURE 62 Normal Layer 2 packet traversing a VCS fabric

In Figure 62, an Ethernet packet arrives from MAC 1 at the VCS fabric edge. TRILL header information is added while the packet passes through the VCS fabric. The TRILL information is removed on leaving the VCS fabric, and a regular Ethernet packet arrives at MAC 2. Table 93 shows the Layer 2 packet header details.

TABLE 93 Packet header details—Layer 2 packet traverses VCS fabric

Ethernet packet	TRILL packet—first hop	TRILL packet—second hop
L2 DA = MAC 2	Outer L2 DA = B1	Outer L2 DA = C1
L2 SA = MAC 1	Outer L2 SA = A1	Outer L2 SA = B1
	Outer 802.1q tag	Outer 802.1q tag
	Outer etype = TRILL	Outer etype = TRILL
	TRILL destination RBridge ID = C	TRILL destination RBridge ID = C
	TRILL source RBridge ID = A	TRILL source RBridge ID = A
	TRILL flags	TRILL flags
	Inner L2 DA = MAC 2	Inner L2 DA = MAC 2
	Inner L2 SA = MAC 1	Inner L2 SA = MAC 1
	Inner 802.1q tag	Inner 802.1q tag
	Inner etype = 0x800	Inner etype = 0x800

When viewing packets while using the **l2tracert** command, you will see TRILL OAM header information added to the packets as they traverse the VCS Fabric. Starting the trace on Switch A, TRILL OAM first verifies path continuity with its immediate neighbor, in this case Switch B. It does this as shown in Figure 63, by sending a Layer 2 traceroute request packet with the time-to-live (TTL) TRILL attribute set to 1. Switch B replies with reachability information regarding the next hop.

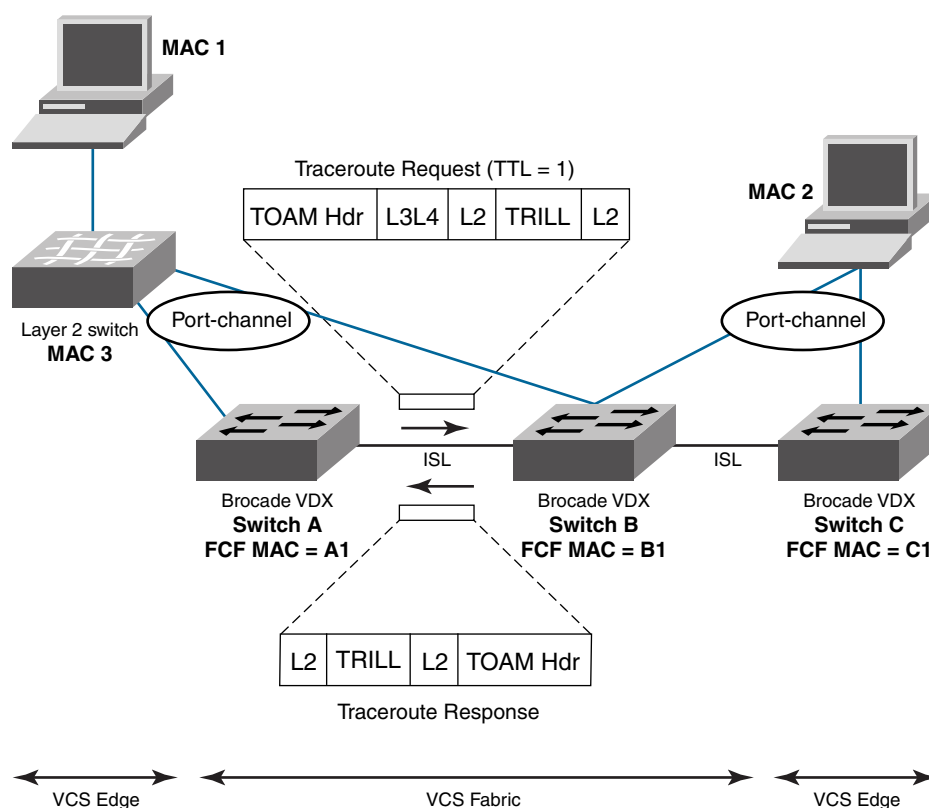


FIGURE 63 Verifying path continuity with immediate neighbor

Table 94 shows the packet header information for the request and response. The added TRILL OAM information is shown in bold typeface.

TABLE 94 Packet header details with Layer 2 traceroute—first hop

Traceroute request packet header	Traceroute reply packet header
Outer L2 DA = B1	Outer L2 DA = B1
Outer L2 SA = A1	Outer L2 SA = A1
Outer 802.1q tag	Outer 802.1q tag
Outer etype = TRILL	Outer etype = TRILL
TRILL destination RBridge ID = C	TRILL destination RBridge ID = A
TRILL source RBridge ID = A	TRILL source RBridge ID = B
TRILL flags: TTL = 1	TRILL flags: TTL = MAX (63)
Inner L2 DA = MAC 2	Inner L2 DA = A1
Inner L2 SA = MAC 1	Inner L2 SA = B1
Inner 802.1q tag	Inner 802.1q tag
Inner etype = 0x800	Inner etype = TRILL OAM
TOAM Opcode = 5 (request)	TOAM Opcode = 4 (reply)
	C reachable

Having successfully exchanged packets with the immediate neighbor (Switch B) and established the reachability of Switch C, the Layer 2 traceroute feature issues another request with TTL set to 2. Switch B decrements the TTL count and forwards the packet to Switch C, which returns a response to Switch A. Refer to Figure 64.

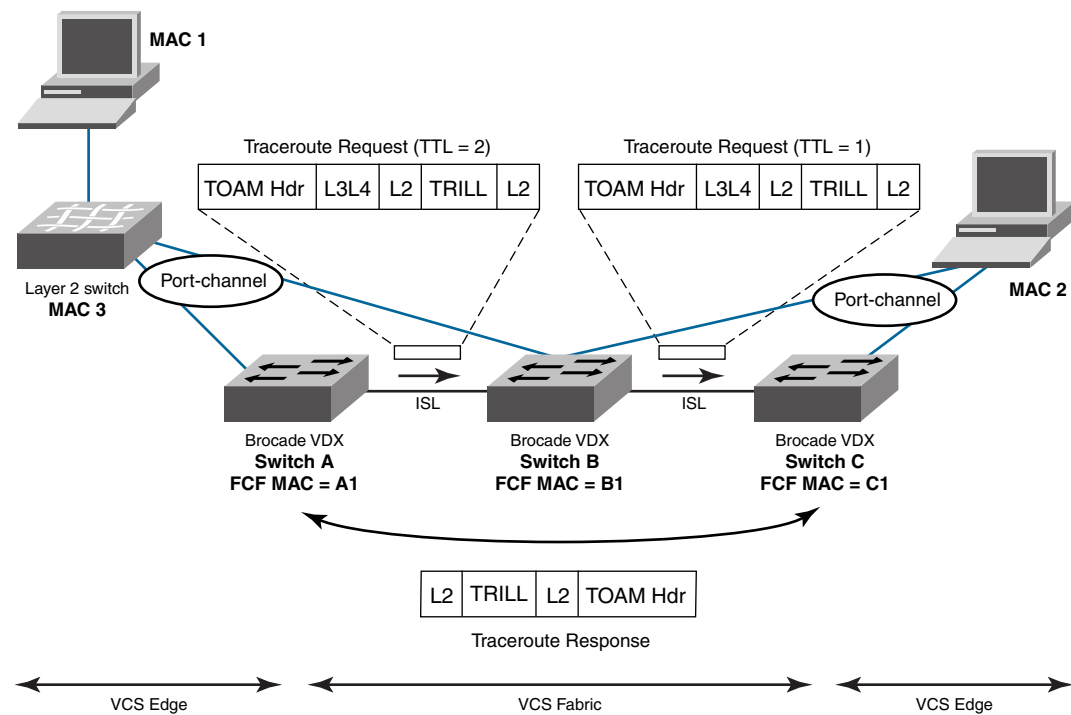


FIGURE 64 Verifying path continuity—second hop

Table 95 shows the packet header information for the request and response packets. Information specific to the Layer 2 traceroute feature is shown in bold typeface.

TABLE 95 Packet header details with Layer 2 traceroute—second hop

Traceroute request—first hop (TTL = 2)	Traceroute request—second hop (TTL = 1)	Traceroute reply
Outer L2 DA = B1	Outer L2 DA = C1	Outer L2 DA = B1->A1
Outer L2 SA = A1	Outer L2 SA = B1	Outer L2 SA = C1->B1
Outer 802.1q tag	Outer 802.1q tag	Outer 802.1q tag
Outer etype = TRILL	Outer etype = TRILL	Outer etype = TRILL
TRILL destination RBridge ID = C	TRILL destination RBridge ID = C	TRILL destination RBridge ID = A
TRILL source RBridge ID = A	TRILL source RBridge ID = A	TRILL source RBridge ID = C
TRILL flags: TTL = 2	TRILL flags: TTL = 1	TRILL flags: TTL = MAX (63)
Inner L2 DA = MAC 2	Inner L2 DA = MAC 2	Inner L2 DA = A1
Inner L2 SA = MAC 1	Inner L2 SA = MAC 1	Inner L2 SA = B1
Inner 802.1q tag	Inner 802.1q tag	Inner 802.1q tag
Inner etype = 0x800	Inner etype = 0x800	Inner etype = TRILL OAM
TOAM Opcode = 5 (request)	TOAM Opcode = 5 (request)	TOAM Opcode = 4 (reply)

Tracing a route with the l2traceroute command

In the following example, the **l2traceroute** command verifies the path between port 3/0/1 (source MAC address 0050.5685.0003) and port 2/0/9 (destination MAC address 0024.3878.3720).

1. enter the **show mac-address-table** command to display all known MAC addresses in the network.

```
switch# show mac-address-table
VlanId  Mac-address      Type   State   Ports
-----
100     0024.3878.e720    Dynamic Active Po 11
100     0050.5685.0001    Dynamic Active Po 1
101     0000.0000.0003    Dynamic Active Po 1
101     0024.3878.e720    Dynamic Active Po 11
101     0050.5685.0003    Dynamic Active Po 1
Total MAC addresses : 5
```

From the output, choose the source and destination MAC address:

- Source MAC address: 0050.5685.0003
- Destination MAC address: 0024.3878.e720

2. Enter the **l2traceroute** command.

```
switch2# l2traceroute
Source mac address      : 0050.5685.0003
Destination mac address : 0024.3878.e720
Vlan [1-3962]           : 101
Edge rbridge-id [1-239] : 3
Extended commands [Y/N]? : y
Protocol Type [IP]       : IP
Source IP address        : 101.101.101.10
Destination IP address   : 101.101.101.101
IP Protocol Type [TCP/UDP] : TCP
Source port number [0-65535] : 3000
Dest port number [0-65535] : 22
Rbridge  Ingress              Egress              Rtt(usec)
-----
3        Te 3/0/1(std-lag, Po 1) Te 3/0/20(isl)       0
2        Te 2/0/20(isl)          Te 2/0/9(std-lag, Po 11) 34041
```

Be advised of the following points:

- The MAC addresses used should be present in the MAC address-table (dynamic or static).
- The **l2tracert** command can be used in VCS Fabric mode only.
- Make use of IP parameters to influence path selection.

show commands

Table 96 lists some **show** commands that are often used for troubleshooting. Refer to the *Network OS Command Reference* for details of all **show** commands.

TABLE 96 show commands used for troubleshooting

Command group	Commands	Specific fields or purpose
System commands	show system	
	show license	
	show running-config	
	show startup-config	
	show logging raslog	
	show version	
	show chassis	
	show environment	
	show vlan brief	
	show mac-address-table	
	show process cpu	
Interface commands	show process memory	
	show firmwaredownloadstatus	
	show interface	
	show media	
	show ip int brief	
	show qos flowcontrol interface	Check pause-frames
	show qos queue interface	Check the CoS statistics
	show qos rcv-queue interface	Check packet drops, buffer consumption, real-time queue statistics
	show qos int	Check the QoS configuration on an interface
Diagnostic commands	show diags status	
	show diags post results detailed	
	show diag burninerrshow	
	show diag burninstatus	

TABLE 96 show commands used for troubleshooting (Continued)

Command group	Commands	Specific fields or purpose
Feature commands	show port-channel detail show lacp counter show port-profile status show fcoe login show fcoe interface brief show fcoe internal fcf-mac-address show lldp neighbors detail show lldp statistics show qos interface all show uddi statistics	
VCS Fabric commands	show vcs show fabric trunk all show fabric all show fabric isl show fabric islports show fabric route linkinfo show fabric route multicast show fabric route neighbor-state show fabric route pathinfo show fabric route topology show name-server detail all	

Debug commands

You can perform the following operations related to debugging features:

- To enable debugging on a feature, use the **debug** command.
`debug feature required-keywords`
- To check whether debugging is enabled on a feature, use the **show debug** command.
`show debug feature`
- To disable debugging, use the **no debug** command.
`no debug feature required-keywords`

Use caution when debugging in real time on a production switch, because real-time debugging is CPU-intensive. Brocade recommends checking the debug output on a lab switch first, and then if the output looks acceptable, enable it on the production switch to get more data. In addition, to reduce CPU load, Brocade recommends using keywords such as **events** and **summary** that limit the extent of debugging rather than more comprehensive options such as **detail** and **all**.

Debugging operations are used mainly for debugging control plane protocols such as LACP and LLDP. For example, to view received LLDP packets on the console, use the following command.

```
switch# debug lldp packets all rx
```

If the switch is accessed through Telnet, enable logging using a terminal monitor.

The following are the most often used debug commands:

- debug lldp packets** *interface* [rx | tx | both]

- `debug lacp pdu [rx | tx]`
- `debug spanning-tree bpdu [rx | tx]` —Standalone mode only
- `debug dot1x packet` —Standalone mode only

SPAN port and traffic mirroring

In certain instances, you may need to examine packets in transit across links to understand the traffic pattern on a specific port. In such situations, Switched Port Analyzer (SPAN) can be configured to copy the traffic (with the desired direction) on the specific Ethernet port to a mirror port where a sniffing device is connected. You can then analyze the packets captured by the sniffing device.

```
switch(config)# monitor session 1
switch(conf-mon-sess-1)# source tengigabitethernet 1/0/10 destination
tengigabitethernet 1/0/15 direction both

switch# show monitor 1
Session :1
Description :Test SPAN Session
State :Enabled
Source interface : 1/0/10 (Up)
Destination interface : 1/0/15 (Up)
Direction :Both
```

The source and destination ports must belong to the same ASIC. The Brocade VDX 6720-24 and Brocade VDX 6730-32 switches have just one ASIC, so source and destination can be any 10-GbE port. Other Brocade VDX switches have multiple ASICs; [Table 97](#) shows the mapping of ports to these ASICs.

TABLE 97 ASICs and ports

Network OS switch	ASIC	Port numbers
Brocade VDX 6720-60 and Brocade VDX 6730-76	0	te0/1 through te0/10
	1	te0/11 through te0/20
	2	te0/21 through te0/30
	3	te0/31 through te0/40
	4	te0/41 through te0/50
	5	te0/51 through te0/60
Brocade VDX 6710	0	te0/1 through te0/6 and gi0/1 through gi0/14
	1	gi0/15 through gi0/27
	2	gi0/28 through gi0/48

The destination port cannot be an ISL, Layer 2, Layer 3, QoS, ACL, 802.1x, LAG member, LLDP, or port-profile port. The source port cannot be an ISL port. In VCS Fabric mode, only edge ports are eligible for mirroring.

Hardware diagnostics

The following diagnostic types currently exist:

- Power-on self-test (POST)
- Offline diagnostics

Online diagnostics are not currently supported on Brocade VDX switches.

POST diagnostics

POST is run on bootup and the results are stored. Use the **show diag post results** command to view the stored results.

To enable POST, enter the **diag post rbridge-id rbridge-id enable** command.

Offline diagnostics

Before proceeding, note the following Caution:



CAUTION

Offline diagnostics—otherwise known as system verification tests—are disruptive tests that check the individual hardware components thoroughly and report the findings. You must disable the chassis before running these tests. Do not run production traffic during this time.

Enter the **diag systemverification** command to run the entire set of offline diagnostics. This command can take up to two hours to finish, so Brocade recommends the less disruptive **diag systemverification short** command, which typically takes 10 to 15 minutes. Alternatively, you can run subsets of the offline commands that check various parts of the hardware. [Table 98](#) shows the complete list of supported offline commands.

TABLE 98 Offline diagnostic commands

Offline diagnostic command	Purpose
diag burninerrclear	Clears the errors that are stored in the nonvolatile storage during the burn-in process.
diag clearerror	Clears the diagnostics failure status.
diag portledtest	Runs various action modes on the port LEDs and validates the functionality.
diag portloopbacktest	Sends frames between various ASICs on the switch and validates the ASIC functionality.
diag setcycle	Configures all the parameters required for the system verification test.
diag systemverification	Runs a combination of various hardware diagnostic tests.
diag turboramtest	Performs a turbo static RAM (SRAM) test of the ASIC chips.

Table 99 lists the show commands that provide output from offline diagnostics.

TABLE 99 Offline diagnostic show commands

Show offline diagnostic command	Purpose
show diag burninerrshow	Displays the errors that are stored in the nonvolatile storage during burn-in.
show diag burninstatus	Displays the diagnostics burn-in status.
show diag setcycle	Displays the current values used in system verification.
show diag status	Displays the currently running diagnostics tests.

For details of the commands listed in Table 98 and Table 99, refer to the *Network OS Command Reference*.

Viewing routing information

The **show fabric route pathinfo** command displays routing and statistical information from a source port index on the local switch to a destination port index on another switch in the same VCS Fabric cluster, a different VCS Fabric cluster, a connected Fabric OS backbone fabric, or Fabric OS edge fabric. This routing information describes the full path that a data stream travels between these ports, including all intermediate switches.

The routing and statistics information are provided by every switch along the path, based on the current routing table information and statistics calculated continuously in real time. Each switch represents one hop.

Use the **show fabric route pathinfo** command to display routing information from a source port on the local switch to a destination port on another switch. The command output describes the exact data path between these ports, including all intermediate switches.

To use the **show fabric route pathinfo** command across remote fabrics, you must specify both the VCS ID (or Fabric ID) and the RBridge ID (or domain ID) of the remote switch. When obtaining path information across remote fabrics, the destination switch must be identified by its RBridge ID or domain ID. Identifying the switch by name or WWN is not accepted.

For details about the **show fabric route pathinfo** command, refer to the *Network OS Command Reference*.

Using the packet capture utility

When a packet is received at a switch's source port, it is routed through the switch to the destination port. If a problem occurs, SPAN or sFlow are commonly used. However, SPAN requires a network analyzer, and sFlow requires a collector. The packet capture utility, executed by means of the **capture packet interface** command, makes it possible to capture packets destined toward the CPU, as well as transit packets if a trap is enabled by means of ACL logging. This command can provide significant help in debugging, especially for Layer 2 TRILL and Layer 3 packets.

There are two ways to view the results of packet capture:

- By using the **show capture packet interface** command on the target switch.
- By viewing the results in an automatically generated file.

For command details, refer to the *Network OS Command Reference*.

Captured packets are stored in a circular buffer, and they are also written to an automatically generated “pktcapture.pcap” file, which can store up to 1500 K of data in flash memory (the equivalent of approximately 10k packets, each having an average size of 100 bytes). Once this file is full, it is saved at *_old.pcap and data are written to a new pktcapture.pcap file. These files can be exported and viewed through a packet analyzer such as Wireshark.

NOTE

Up to 100 packets per interface can be captured. Once the buffer is filled, the oldest packets are replaced with the most recent.

Note the following limitations:

- Support is provided only on physical interfaces (1-, 10-, and 40-gigabit Ethernet), not on logical interfaces. To see packets on logical interfaces, first enable the capture on the corresponding physical interfaces.
- In the initial release, support for capturing transit traffic requires ACL logging.
- Packets that are dropped in the ASIC cannot be captured.



CAUTION

Capturing packets over multiple sessions and over long durations can affect system performance.

TACACS+ Accounting Exceptions

In this chapter

- [Command accounting limitations.](#) 655

Command accounting limitations

TACACS+ command accounting is subject to the following limitations:

- The TACACS+ command accounting logs only the base command name. For example, if the command executed is **secpolicy defined-policy SCC_POLICY**, only the **secpolicy** command is logged in the TACACS+ server.
- the **no radius-server** command is logged as the **radius-server** command.
- A few commands are not accounted for. Refer to [Table 100](#) for a listing of unsupported operational commands. Refer to [Table 101](#) for a listing of unsupported configuration commands.

TABLE 100 Unsupported commands in privileged EXEC mode

Command name	Command Description
cipherset	Configures FIPS-compliant secure ciphers for LDAP and SSH.
clear	Clears the specified parameter.
clear arp	Clears Address Resolution Protocol (ARP) configuration data.
clear counters	Clears statistics from the switch.
clear dot1x	Clears IEEE 802.1X Port-Based Access Control configuration data.
clear fcoe	Clears FCoE configuration data.
clear ip	Clears Internet Protocol (IP) configuration data.
clear lacp	Clears Link Aggregation Control Protocol (LACP) configuration data.
clear lldp	Clears Link Layer Discovery Protocol (LLDP).configuration data.
clear mac-address-table	Clears the MAC address table.
clear mcagt	Clears the MCAGT agent.
clear policy-map-counters	Clears the policy map counters.
clear sflow	Clears sFlow configuration data.
clear spanning-tree	Clears Spanning Tree Protocol (STP) configuration data.
clear vrrp	Clears Virtual Router Redundancy Protocol (VRRP) configuration data.
configure	Configures access mode.
copy	Copies data.

TABLE 100 Unsupported commands in privileged EXEC mode (Continued)

Command name	Command Description
debug	Sets debugging options.
delete	Delete a specified file.
dir	Displays a directory listing.
dot1x	Executes IEEE 802.1X Port-Based Access Control options.
exit	Exits to the top level and optionally runs a command.
fips	Executes FIPS-related operations.
help	Provides help information.
history	Configures the size of the history log.
logout	Terminates the current login session.
mac-rebalance	Rebalances MAC on a port channel
ping	Executes the ping command.
quit	Terminates the current session.
rename	Renames a file.
reload	Reboots the system.
resequence	Re-orders a list.
send	Sends a message to terminal of one or all users.
terminal	Configures terminal properties.
show arp	Displays the Address Resolution Protocol (ARP) configuration.
show bpdudrop	Displays the Bridge Protocol Data Unit (BPDU) drop configuration.
show cee maps	Displays CEE maps.
show cipher	Displays ciphers for LDAP and SSH.
show cli	Displays CLI session parameters.
show clock	Displays the date and time settings.
show debug arp	Displays ARP packet debugging information.
show diag	Displays diagnostic information.
show dot1x	Displays IEEE 802.1X Port-Based Access Control configuration data.
show edge-loop-detection globals	Displays system-wide Edge-Loop-Detection status information.
show fcoe login	Displays the FCoE CNA Login information.
show file	Displays the contents of a file.
show history	Displays command history.
show interface	Displays interface status and configuration.
show ip	Displays Internet Protocol (IP) information.
show lacp counter	Displays Link Aggregation Control Protocol (LACP) counters.
show lldp	Displays Link Layer Discovery Protocol (LLDP) configuration data
show monitor	Displays interface status and configuration.

TABLE 100 Unsupported commands in privileged EXEC mode (Continued)

Command name	Command Description
show netconf-state	Displays NETCONF statistics.
show ntp	Displays the active NTP server.
show parser dump	Displays a parser dump.
show policy-map	Displays the configured rate-limiting policy maps.
show port	Displays port parameters.
show port-channel	Displays the port-channel configuration.
show port-profile	Displays the port profile configuration
show qos	Display the Quality of Service (QoS) configuration.
show running-config	Displays the running configuration.
show sflow	Displays the sFlow configuration.
show spanning-tree	Displays the Spanning Tree Protocol configuration.
show ssm	Displays the switch services subsystem.
show startup-db	Displays the startup configuration.
show storm-control	Displays storm control configuration.
show statistics	Displays accounting information.
show system	Displays runtime system information.
show rmon	Displays the Remote Monitoring Protocol (RMON) configuration.
show vcs	Displays VCS information.
show vlan	Displays the VLAN configuration
show mac-address-table	Displays the MAC address table.
show startup-config	Displays the contents of the startup-configuration file.
show zoning	Displays zoning information.
traceroute	Executes the traceroute command.

TABLE 101 Unsupported commands in global configuration mode

Command name	Command Description
abort	Aborts the current configuration session.
diag	Manages diagnostic commands.
do	Executes an operational command while in global configuration mode.
end	Terminates the current configuration session.
exit	Exits from the current mode.
help	Provides help information.
pwv	Displays the current mode path.
service	Performs password encryption services.
top	Exits to the top level and optionally runs a command.
no vlan	Disables VLAN configuration.

A Command accounting limitations

Supported time zones and regions

Time zones and regions supported by the Network Time Protocol are listed in the following tables.

- **Africa**—[Table 102](#) on page 659
- **America**—[Table 103](#) on page 660
- **Antarctica**—[Table 104](#) on page 661
- **Arctic**—[Table 105](#) on page 661
- **Asia**—[Table 106](#) on page 661
- **Atlantic**—[Table 107](#) on page 662
- **Australia**—[Table 108](#) on page 662
- **Europe**—[Table 109](#) on page 662
- **Indian**—[Table 110](#) on page 663
- **Pacific**—[Table 111](#) on page 663

Africa

TABLE 102 Region/city time zones in Africa

Africa/Luanda	Africa/Banjul	Africa/Mogadishu
Africa/Ouagadougou	Africa/Conakry	Africa/Sao_Tome
Africa/Bujumbura	Africa/Malabo	Africa/Mbabane
Africa/Porto-Novo	Africa/Bissau	Africa/Ndjamena
Africa/Gaborone	Africa/Nairobi	Africa/Lome
Africa/Kinshasa	Africa/Monrovia	Africa/Tunis
Africa/Lubumbashi	Africa/Maseru	Africa/Dar_es_Salaam
Africa/Bangui	Africa/Tripoli	Africa/Kampala
Africa/Brazzaville	Africa/Casablanca	Africa/Johannesburg
Africa/Abidjan	Africa/Bamako	Africa/Lusaka
Africa/Douala	Africa/Nouakchott	Africa/Harare
Africa/Djibouti	Africa/Blantyre	
Africa/Algiers	Africa/Maputo	
Africa/Cairo	Africa/Windhoek	
Africa/El_Aaiun	Africa/Niamey	
Africa/Asmara	Africa/Lagos	
Africa/Ceuta	Africa/Kigali	
Africa/Addis_Ababa	Africa/Khartoum	
Africa/Libreville	Africa/Freetown	
Africa/Accra	Africa/Dakar	

America

TABLE 103 Region/city time zones in America

America/Antigua	America/Guatemala	America/Edmonton
America/Anguilla	America/Guyana	America/Cambridge_Bay
America/Curacao	America/Tegucigalpa	America/Yellowknife
America/Argentina/Buenos_Aires	America/Port-au-Prince	America/Inuvik
America/Argentina/Cordoba	America/Guadeloupe	America/Dawson_Creek
America/Argentina/San_Luis	America/Jamaica	America/Vancouver
America/Argentina/Jujuy	America/St_Kitts	America/Whitehorse
America/Argentina/Tucuman	America/Cayman	America/Thunder_Bay
America/Argentina/Catamarca	America/St_Lucia	America/Iqaluit
America/Argentina/La_Rioja	America/Marigot	America/Pangnirtung
America/Argentina/San_Juan	America/Adak	America/Resolute
America/Argentina/Mendoza	America/Martinique	America/Rankin_Inlet
America/Argentina/Rio_Gallegos	America/Montserrat	America/Winnipeg
America/Argentina/Ushuaia	America/Mexico_City	America/Rainy_River
America/Aruba	America/Cancun	America/Regina
America/Barbados	America/Merida	America/Montevideo
America/St_Barthelemy	America/Monterrey	America/St_Vincent
America/La_Paz	America/Mazatlan	America/Caracas
America/Noronha	America/Chihuahua	America/Tortola
America/Belem	America/Hermosillo	America/St_Thomas
America/Fortaleza	America/Tijuana	America/New_York
America/Recife	America/Managua	America/Detroit
America/Araguaina	America/Panama	America/Kentucky/Monticello
America/Maceio	America/Lima	America/Indiana/Indianapolis
America/Bahia	America/Miquelon	America/Indiana/Vincennes
America/Sao_Paulo	America/Puerto_Rico	America/Indiana/Knox
America/Campo_Grande	America/Asuncion	America/Indiana/Winamac
America/Cuiaba	America/Paramaribo	America/Indiana/Marengo
America/Santarem	America/El_Salvador	America/Indiana/Vevay
America/Porto_Velho	America/Grand_Turk	America/Chicago
America/Boa_Vista	America/Swift_Current	America/Indiana/Tell_City
America/Manaus	America/Dawson	America/Indiana/Petersburg
America/Eirunepe	America/Santiago	America/Menominee
America/Rio_Branco	America/Bogota	America/North_Dakota/Center
America/Nassau	America/Costa_Rica	America/North_Dakota/New_Salem
America/Belize	America/Havana	America/Denver
America/St_Johns	America/Dominica	America/Boise
America/Halifax	America/Santo_Domingo	America/Shiprock
America/Glace_Bay	America/Guayaquil	America/Phoenix
America/Moncton	America/Grenada	America/Los_Angeles
America/Goose_Bay	America/Cayenne	America/Anchorage
America/Blanc-Sablon	America/Godthab	America/Juneau
America/Montreal	America/Danmarkshavn	America/Yakutat
America/Toronto	America/Scoresbysund	America/Nome
America/Nipigon	America/Thule	America/Port_of_Spain

Antarctica

TABLE 104 Region/city time zones in Antarctica

Antarctica/McMurdo	Antarctica/Mawson	Antarctica/Vostok
Antarctica/South_Pole	Antarctica/Davis	Antarctica/Dumont_dUrville
Antarctica/Rothera	Antarctica/Casey	Antarctica/Syowa

Arctic

TABLE 105 Region/city time zone in Arctic

Arctic/Longyearbyen

Asia

TABLE 106 Region/city time zones in Asia

Asia/Dubai	Asia/Tokyo	Asia/Gaza
Asia/Kabul	Asia/Bishkek	Asia/Qatar
Asia/Yerevan	Asia/Phnom_Penh	Asia/Yekaterinburg
Asia/Baku	Asia/Pyongyang	Asia/Omsk
Asia/Dhaka	Asia/Seoul	Asia/Novosibirsk
Asia/Bahrain	Asia/Kuwait	Asia/Krasnoyarsk
Asia/Brunei	Asia/Almaty	Asia/Irkutsk
Asia/Thimphu	Asia/Qyzylorda	Asia/Yakutsk
Asia/Shanghai	Asia/Aqtobe	Asia/Vladivostok
Asia/Harbin	Asia/Aqtau	Asia/Sakhalin
Asia/Chongqing	Asia/Oral	Asia/Magadan
Asia/Urumqi	Asia/Vientiane	Asia/Kamchatka
Asia/Kashgar	Asia/Beirut	Asia/Anadyr
Asia/Nicosia	Asia/Colombo	Asia/Riyadh
Asia/Tbilisi	Asia/Rangoon	Asia/Singapore
Asia/Hong_Kong	Asia/Ulaanbaatar	Asia/Damascus
Asia/Jakarta	Asia/Hovd	Asia/Bangkok
Asia/Pontianak	Asia/Choibalsan	Asia/Dushanbe
Asia/Makassar	Asia/Macau	Asia/Dili
Asia/Jayapura	Asia/Kuala_Lumpur	Asia/Ashgabat
Asia/Jerusalem	Asia/Kuching	Asia/Taipei
Asia/Kolkata	Asia/Katmandu	Asia/Samarkand
Asia/Baghdad	Asia/Muscat	Asia/Tashkent
Asia/Tehran	Asia/Manila	Asia/Ho_Chi_Minh
Asia/Amman	Asia/Karachi	Asia/Aden

Atlantic

TABLE 107 Region/city time zones in Atlantic

Atlantic/Bermuda	Atlantic/Faroe	Atlantic/Azores
Atlantic/Cape_Verde	Atlantic/South_Georgia	Atlantic/St_Helena
Atlantic/Canary	Atlantic/Reykjavik	
Atlantic/Stanley	Atlantic/Madeira	

Australia

TABLE 108 Region/city time zones in Australia

Australia/Lord_Howe	Australia/Sydney	Australia/Darwin
Australia/Hobart	Australia/Brisbane	Australia/Perth
Australia/Currie	Australia/Lindeman	Australia/Eucla
Australia/Melbourne	Australia/Adelaide	

Europe

TABLE 109 Region/city time zones in Europe

Europe/Andorra	Europe/Gibraltar	Europe/Warsaw
Europe/Tirane	Europe/Athens	Europe/Lisbon
Europe/Vienna	Europe/Zagreb	Europe/Bucharest
Europe/Mariehamn	Europe/Budapest	Europe/Belgrade
Europe/Sarajevo	Europe/Dublin	Europe/Kaliningrad
Europe/Brussels	Europe/Isle_of_Man	Europe/Moscow
Europe/Sofia	Europe/Rome	Europe/Volgograd
Europe/Minsk	Europe/Jersey	Europe/Samara
Europe/Zurich	Europe/Vaduz	Europe/Stockholm
Europe/Prague	Europe/Vilnius	Europe/Ljubljana
Europe/Berlin	Europe/Luxembourg	Europe/Bratislava
Europe/Copenhagen	Europe/Riga	Europe/San_Marino
Europe/Tallinn	Europe/Monaco	Europe/Istanbul
Europe/Madrid	Europe/Chisinau	Europe/Kiev
Europe/Helsinki	Europe/Podgorica	Europe/Uzhgorod
Europe/Paris	Europe/Skopje	Europe/Zaporozhye
Europe/London	Europe/Malta	Europe/Simferopol
Europe/Guernsey	Europe/Amsterdam	Europe/Vatican
Europe/Oslo		

Indian

TABLE 110 Region/city time zones in India

Indian/Cocos	Indian/Antananarivo	Indian/Mahe
Indian/Christmas	Indian/Mauritius	Indian/Kerguelen
Indian/Chagos	Indian/Maldives	Indian/Mayotte
Indian/Comoro	Indian/Reunion	

Pacific

TABLE 111 Region/city time zones in Pacific

Pacific/Pago_Pago	Pacific/Kwajalein	Pacific/Palau
Pacific/Rarotonga	Pacific/Saipan	Pacific/Guadacanal
Pacific/Easter	Pacific/Noumea	Pacific/Fakaofu
Pacific/Galapagos	Pacific/Norfolk	Pacific/Tongatapu
Pacific/Fiji	Pacific/Nauru	Pacific/Funafuti
Pacific/Truk	Pacific/Niue	Pacific/Johnston
Pacific/Ponape	Pacific/Auckland	Pacific/Midway
Pacific/Kosrae	Pacific/Chatham	Pacific/Wake
Pacific/Guam	Pacific/Tahiti	Pacific/Honolulu
Pacific/Tarawa	Pacific/Marquesas	Pacific/Efate
Pacific/Enderbury	Pacific/Gambier	Pacific/Wallis
Pacific/Kiritimati	Pacific/Port_Moresby	Pacific/Apia
Pacific/Majuro	Pacific/Pitcairn	

Index

Symbols

? (CLI help), 23

Numerics

802.1x

LAG, 461

overview, 461

timeouts, 464

A

AAA service requests, 235

access

remote access policies, 243

Access Control Lists

See ACL

access interface, configuring, 336

access mode, 331, 336

ACL

configuration guidelines and restrictions, 404

configuration procedures

applying a MAC ACL to a CEE interface, 406

applying a MAC ACL to a VLAN interface, 407

creating extended MAC ACL and adding rules,
406

creating standard MAC ACL and adding rules, 405

important notes, 405

modifying a MAC ACL, 407

removing a MAC ACL, 408

reordering the sequence numbers, 408

default configuration, 404

extended ACL, defined, 404

overview, 321, 403

standard ACL, defined, 404

ACLs

FIPS, 272

activating

POD, 105

Active Directory

LDAP, 260

adding

alias members, 155

admin login, 18

alias

adding members, 155

creating, 155

deleting, 157

removing members, 156

AMPP

access-group, 307

ACL, 307

flow control, 306

port-profile, 302

port-profile states, 303

priority, 306

QoS profile, 306

security profile, 307

VLAN profile, 305

authentication

configuring, 235

authentication server, 461

authenticator, 461

Auto Migrating Port Profile (AMPP), 299

B

Base Port Set, 104

basic management TLV sets, 390

BGP

- advertised networks, 572
- attributes, 565
- best-path algorithm, 565
- clearing configurations, 584
- configuration examples, 580
- configuration fundamentals, 567
- default route origination, 574
- deployment scenarios, 560
- device ID, 567
- four-byte AS numbers, 571
- global mode, 569
- IPv4 unicast address family, 568
- limitations, 559
- local AS number, 567
- message types, 563
- multipath load sharing, 574
- neighbor configuration, 569
- next-hop, 575
- peer groups, 571
- peering, 562
- route filtering, 575
- route flap dampening, 573
- route maps, 576
- route redistribution, 571
- route reflection, 573
- static networks, 572
- timers, 575

Bootstrap router (BSR), 517

Border Gateway Protocol (BGP), 559

bridge

- forwarding delay, 357
- hello time, 359
- maximum aging time, 357
- priority, 356

Brocade

- proprietary aggregation, 379

Brocade MIB files, 114

BUM

- configuration, 436
- considerations, 436

BUM storm control, 436

C

CA certificate, 261

CEE interface

- applying a MAC ACL, 406
- configuring for STP, RSTP, MSTP, 363
- configuring the hello time for MSTP, 366
- disable or enable STP on the interface, 370
- enabling and disabling, 334
- enabling as an edge port for RSTP, MSTP, 364
- enabling guard root for STP, RSTP, MSTP, 365
- enabling LACP, 384
- enabling port fast, 367
- path cost, 364
- restricting the port from becoming a root port for STP, RSTP, MSTP, 369
- restricting the topology change notification for STP, RSTP, MSTP, 369
- spanning-tree defaults, 354
- specifying a link type, 367
- specifying restrictions for an MSTP instance, 366
- specifying the port priority for STP, RSTP, MSTP, 368

cerutil, 261

chassis ID (CID)

- card usage, 595
- critical SEEPROM data, 595
- non-critical SEEPROM data, 595

chassis ID(CID) recovery tool, 596

chassis name, 31

- setting the, 32

Cisco interoperability, disabling for MSTP, 360

Cisco interoperability, enabling for MSTP, 360

classifier groups, VLAN, 339

classifier rules, VLAN, 338

CLI, 18

CLI, CEE

- command completion, 24
- command modes, 19
- command syntax, 23
- configuration guidelines and restrictions, 17
- displaying commands, 23
- help, 23
- keyboard shortcuts, 22
- output modifiers, 24
- RBAC permissions, 18

color-based priority mapping, 446

command completion, CEE CLI, 24

command line interface, 18, 27

command modes, CEE, 19

command output modifiers, 24

command syntax, 23

- configuration, 71
 - backup, 76
 - default, 73
 - display, 74
 - in Brocade VCS Fabric mode, 80
 - restore, 76
 - restore default, 77
 - running, 73
 - save, 74
 - startup, 73, 74
- configuration file, 71
 - display settings, 71
 - save to a host, 71
- configuration management
 - saving changes, 17
- configure
 - FCoE VLAN, 334
- configuring
 - authentication, 235
 - LINUX RADIUS server, 241
 - SNMP, 111–??
- congestion control
 - QoS, 428
 - queuing, 319
- connection
 - serial, 27
- console, 27
- contacting your switch support provider, 63
- converged mode, 331
- creating
 - alias, 155

D

- Data Center Bridging (DCB) Capability Exchange Protocol
 - See DCBX
- DCB map, configuring, 442
- DCB maps, verifying, 444
- DCB provisioning map, applying, 443
- DCBX
 - Enhanced Transmission Selection, 392
 - overview
 - Priority Flow Control, 392
- default account, 29
- deleting
 - alias, 157
- dictionary.brocade, 242
- displaying
 - configuration settings, 71

- Double POD License, 104
- dynamic link aggregation, 378
- Dynamic POD license, 96, 100, 102
- Dynamic Ports on Demand, 104

E

- EAP, 461
- edge detection, 363
- edge port, enabling a CEE interface as an edge port for
 - RSTP, MSTP, 364
- Enhanced Transmission Selection
 - See ETS
- ERR, 363
- error disable timeout, 358
- error disable timeout interval, 358
- Ethernet port, 29
- Ethernet, forwarding, 317
- ETS
 - overview
 - priority grouping of IPC, LAN, and SAN traffic, 392

F

- FCoE
 - Layer 2 Ethernet overview, 317
 - overview, 313
 - queuing, 325
 - terminology, 314
 - VLAN forwarding, 318
- FCoE initialization protocol
 - See FIP
- Fibre Channel Association, xxxix
- file
 - delete, 71
 - rename, 71
- file management, 71
- filtering VLAN ingress, 331
- FIP
 - FIP discovery, 322
 - login, 323
 - logout, 323
 - name server, 324
 - registered state change notification (RSCN), 324

FIPS

- certificates, deleting, 258, 259
- certificates, installing, 257, 258
- certificates, verifying, 258, 259
- conditional tests, 267
- conditional tests, enabling, 271
- KATs, enabling, 271
- selftests, 271
- zeroization, 266, 271, 275

FIPS-compliant state

- ACLs for, 272
- enabling, 269
- LDAP, configuring for, 275

firmware, 81

- decompressing, 82
- download, 81
- download from USB, 87
- remote download, 86
- show version, 84
- upgrades, 82

firmware download, 90

- USB device, 87

flash

- deleting a file, 72
- listing contents, 72
- renaming a file, 72

flash memory, 71

flow control, 322

frame classification, incoming, 319

G

guard root, enabling on a CEE interface, 365

H

HA failover, 229

hello time, 366

hops, configuring for MSTP, 361

host name, 31

- setting the, 31

I

IAS

- remote access policies, 243

IEEE 802.1 organizational TLV set, 391

IEEE 802.3 organizational TLV set, 391

IGMP

- interface, 588
- interval, 589
- mrrouter, 588
- MRT, 589
- querier, 589
- query-interval, 588
- tcn, 588
- timer, 588
- vlan, 588

incoming frame classification, 319

ingress VLAN filtering, 331

instance

- MSTP, mapping a VLAN to, 361
- specifying restrictions for an MSTP instance, 366

interface, 325, 326, 334

IP route management

- configuring default route, 510
- configuring static routes, 510
- how best routes are determined, 509
- overview, 509

iSCSI priority, 398

K

keyboard shortcuts, CEE CLI, 22

L

LACP

- configuration guidelines and restrictions, 384
- configuration procedures
 - clearing counters, 386
 - configuring system priority, 385
 - configuring timeout period, 385
 - displaying LACP information, 386
 - enabling on a CEE interface, 384
 - important notes, 384
- default LACP configuration, 384
- overview
- troubleshooting tips, 386

LAGs

- 802.1x, 461
- distribution process, 379
- overview

Layer 2

- ACL
- Ethernet overview, 317

- Layer 2 Ethernet, 317
- Layer 2 forwarding, 317
- Layer 3 feature restrictions, 445
- LDAP, 261
 - Active Directory, 260
 - attributes, 256
 - cerutil, 261
 - maprole, 261
 - role, 260
- ldapca, 261
- license, 95
 - Dynamic Ports on Demand, 96
 - expiration, 98
 - individual, 98
 - permanent, 97
 - removing, 102
 - reservation, 108
 - temporary, 97
 - time-based, 97
 - universal, 98
 - upgrade, 99
 - VCS Fabric license, 96
- license ID, 97
- licenses
 - remove feature, 102
- Licensing, 95
- link aggregation
 - Brocade-proprietary, 379
 - dynamic, 378
 - LACP, 378
 - LAG distribution process, 379
 - LAGs, 378
 - overview, 253, 377
 - static, 378
- Link Aggregation Control Protocol
 - See LACP
- link aggregation group
 - See LAGs
- Link Layer Discovery Protocol
 - See LLDP
- link type, specifying, 367
- Linux, configuring RADIUS on, 241

- LLDP
 - configuration guidelines and restrictions, 393
 - configuration procedures
 - clearing LLDP-related information, 401
 - disabling LLDP globally, 394
 - displaying LLDP-related information, 401
 - enabling LLDP globally, 393
 - global command options, 394
 - important notes, 393
 - interface-level command options, 401
 - DCBX overview
 - default configuration, 393
 - Layer 2 topology mapping, 390
 - overview, 389
 - TLV sets, 390
- loading Brocade MIBs, 114
- logical chassis cluster
 - adding a node, 44
 - characteristics, 35
 - configuration, 36
 - creating, 37
 - description, 35
 - mode conversions, 43, 44
 - mode transitions, 39, 41
 - principal node setting, 38, 43
 - removing a node, 44
 - replacing a node, 45
- login
 - FIP, 323
- logout
 - FIP, 323
- LSAN zone, 147
 - example of, 176
 - naming, 175
- LSAN zoning
 - configuration of, 175

M

- MAC addresses
 - configuration guidelines and restrictions, 333
- management port, 27
- map, 325, 326
- maprole, 261
- MIB loading order, 116
- MSTP
 - configuration procedures
 - PVST configuration, 355
 - default configuration, 354
 - overview, 351

- MTU, configuring, 334
- multicast rate limiting, QoS, 435
- Multiple Spanning Tree Protocol
 - See MSTP

N

- name server, 324
- network
 - flow control, 322
 - trunking, 321

O

- Open Shortest Path First (OSPF), 525
- OSPF
 - area border routers (ABRs), 526
 - area ranges, 536
 - autonomous system boundary routers (ASBRs), 526
 - backbone area, 526
 - configuration, 532
 - designated routers, 531
 - Link State Advertisements, 533
 - not-so-stubby areas, 535
 - supported platforms, 525
 - totally stubby areas, 533
 - VCS environment, 529
 - virtual links, 537
- output modifiers, CEE CLI, 24
- overview
 - ACL, 403
 - link aggregation, 253, 377
 - MSTP, 351
 - PVST, 353
 - RSTP, 349
 - STP, 345

P

- packet capture utility, 652
- paperpack, 95
- password, 29
- path cost
 - CEE interface, configuring for STP, RSTP, MSTP, 364
 - port channel, 358
- PEAP, 461

- PIM
 - configuration, 520
 - IGMP snooping, 521
- PIM designated router (DR), 517
- PIM multicast border router (PMBR), 517
- PIM-SM (Sparse Mode), 514
- POD
 - activating, 105
- port
 - activating POD, 105
- port assignment
 - release, 108
- port assignments
 - display, 107
 - override, 107
 - reserve, 107
- port configuration for STP, RSTP, MSTP, 363
- port fast, enabling on a CEE interface, 367
- port priority, specifying on a CEE interface, 368
- port-profile, 303
- Ports on Demand
 - Base port set, 104
 - Double POD set, 104
 - port assignments, 105
 - Single POD set, 104
- Priority Flow Control (PFC), 392
- priority group table, mapping, 442
- priority mapping, QoS, 414
- priority-table, mapping, 443
- Protocol Independent Multicast (PIM), 513
- PVST
 - default configuration, 354
 - overview, 353

Q

QoS

configuration procedures

- activating CoS-to-Traffic-Class mapping, 425
- activating DSCP-to-Traffic-Class mapping, 427
- applying a CEE provisioning map, 443
- applying a CoS-to-CoS mutation QoS map, 417
- applying a DCB provisioning map, 443
- applying a DSCP mutation map, 420
- applying a DSCP-to-COS mutation map, 421
- changing the multicast Tail Drop threshold, 429
- configuring Brocade VCS fabric QoS, 444
- configuring CoS thresholds, 429
- configuring QoS trust mode, 416
- configuring RED profiles, 431
- configuring the DSCP trust mode, 418
- configuring user priority mappings, 416
- creating a CEE map, 442
- creating a CoS-to-CoS mutation QoS map, 417
- creating a DCB map, 442
- creating a DSCP mutation map, 419
- creating a DSCP-to-COS mutation map, 421
- creating a receive queue multicast rate-limit, 435
- defining a priority group table, 442
- defining a priority-table map, 443
- enabling an Ethernet PFC, 434
- enabling Ethernet pause, 433
- enabling RED profile to use CoS priority, 431
- mapping a priority group table, 442
- mapping a priority-table, 443
- mapping CoS-to-Traffic-Class, 425
- mapping DSCP-to-Traffic-Class, 426
- scheduling the QoS multicast queue, 440
- scheduling the QoS queue, 439
- verifying a DSCP-to-COS mutation map, 422
- verifying CEE maps, 444
- verifying CoS trust, 416
- verifying CoS-to-CoS mutation QoS mapping, 418
- verifying CoS-to-Traffic-Class mapping, 425
- verifying DCB maps, 444
- verifying DSCP mutation mapping, 420
- verifying DSCP trust, 419
- verifying DSCP-to-Traffic-Class mapping, 427,

431

- congestion control, 428
- data center bridging map configuration overview, 440
- multicast rate limiting, 435
- overview, 299, 413, 469, 477
- port-based policer
 - binding the policy map to an interface, 452
 - configuring a class map, 446
 - configuring a police priority map, 447
 - configuring parameters for a class map, 450
 - configuring the policy map, 448
 - considerations and limitations, 457
 - displaying policing settings and policy maps, 456
 - overview, 445
 - policer binding rules, 453
 - policing parameters, 453
- queuing
 - traffic class mapping, 422
 - user-priority mapping, 414
- queuing overview, 414
- rewriting frame header field, 414
- scheduling, 437

Quality of Service

See QoS

querier

- interval, 589
- MRT, 589
- VLAN, 589

queuing

- congestion control, 319
- FCoE, 325
- QoS, 414

R

RADIUS, 461

RADIUS server

- LINUX configuration, 241

Rapid Spanning Tree Protocol

See RSTP

RBAC permissions

region name, specifying for MSTP, 362

registered state notification protocol (RSCN), 324

remote access policies, 243

remove feature, 102

removing

- alias members, 156
- licensed feature, 102

Rendezvous point (RP), 517

revision number, specifying for MSTP, 362

RJ-45, 29

Role-Based Action Control

See RBAC

root port, CEE interface, restricting for Spanning Tree, 369

root_inc, 363

RSTP

configuration guidelines and restrictions

MSTP configuration guidelines and restrictions,
348

configuration procedures, 355

default configuration, 354

overview, 349

S

saving configuration, 17

scheduling, QoS, 437

secure shell, 27

security

IAS remote access policies, 243

serial connection, 27

Simple Network Management Protocol. See: SNMP.

Single POD License, 104

SNMP

configuring, 111-??

SNMP v3, 121

Spanning Tree Protocol

See STP

spanning-tree defaults, 354

SSH, 27

static link aggregation, 378

STP

configuration guidelines and restrictions, 348

configuration procedures, 355

default configuration, 354

overview, 345

VDX switches, 345

supplicant, 461

switch, 31

attributes, 31

certificates, deleting for FIPS, 258, 259

certificates, installing for FIPS, 257, 258

certificates, verifying for FIPS, 258, 259

port configuration, 336

switchType, 33

system priority, configuring for LACP, 385

T

Telnet, 27

telnet, 18

time-based licenses, 97

timeout period, configuring for LACP, 385

TLV sets

basic management TLV, 390

IEEE 802.1 organizational TLV set, 391

IEEE 802.3 organizational TLV set, 391

topology change notification, CEE interface, restricting for
Spanning Tree, 369

topology mapping, LLDP, 390

traffic class mapping, QoS, 422

transmit hold count, 360

troubleshooting tips, LACP, 386

trunk interface, configuring, 336

trunk mode, 331, 336

trunking, 321

U

understanding MIBs, 113

understanding SNMP bASICs, 112

UniDirectional Link Detection (UDLD)

commands, 375

configuring, 374

example, 373, 374

requirements, 373

USB, 76

USB device, 87

user-priority mapping, QoS, 414

V

VCS restrictions for DSCP features, 445

Virtual LANs

See VLAN

Virtual Router Redundancy Protocol (VRRP), 541

VLAN

applying a MAC ACL, 407

configuration guidelines and restrictions, 333

configuration procedures

configuring a CEE interface as a Layer 2 switch
port, 336

configuring a CEE interface as an access or trunk

- interface, 336
 - configuring the MTU on an interface, 334
 - displaying VLAN information, 339
 - enabling and disabling a CEE interface, 334
 - important notes, 334
 - VLAN classifier groups, 339
 - VLAN classifier rules, 338
- default configuration, 334
- FDB
 - overview, 333
- forwarding, 318
- important management notes, 334
- ingress VLAN filtering, 331
- overview, 331
- tagging, 318
- VRRP, 542
 - backup, 543
 - control packets, 544
 - gratuitous ARP, 544
 - master, 543
 - multigroup configuration, 550
 - owner, 542
 - track ports, 547
 - track priority, 547
 - virtual router address, 542
 - virtual router group, 542
- VRRP-E, 542
 - packet routing with short-path forwarding, 549
 - short-path forwarding, 547

- zone configuration
 - creating, 162
 - deleting, 165
 - deleting all, 166
 - disabling, 165
 - enabling, 164
 - zone, adding, 163
 - zone, removing from, 163
- zoning
 - database size for, 154
 - default mode, setting, 153
 - default modes, All Access, 153
 - default modes, No Access, 153
 - default modes, overview, 153
 - defined configuration, 149
 - defined configuration, viewing, 160
 - downgrade, 151
 - enabled configuration, 149
 - enabled configuration, viewing, 161
 - example of, 168
 - LSAN configuration, 175
 - LSAN, example of, 176
 - LSAN, overview, 147
 - management of, 151
 - objects for, 148, 149
 - overview, 145
 - supported firmware, 150
 - transaction abort, 166
 - transaction commit, 152

Z

zeroization, 266, 271, 275

zone

- alias, adding members, 155
- alias, deleting, 157
- alias, removing members, 156
- aliases, creating and managing, 154
- configuration, definition, 149
- creating, 157
- deleting, 159
- member, adding, 155, 156, 157, 158
- member, deleting, 156, 157, 159
- merging, 169
- resources, optimizing, 145
- splitting a fabric, 171
- WWN, adding, 155, 156, 157, 158
- WWN, deleting, 156, 157, 159
- zone configuration, adding to, 163

