DATA CENTER

Brocade Switch Cookbook

October 2013

BROCADE

CONTENTS

Brocade Switch Cookbook 1

Contents 2

Chapter 1: Brocade Switch Introduction 10

- Overview of Brocade 6505 10
- Overview of Brocade 6510 12
- Platform Components of 6510 13
- Port side view of Brocade 6510 14
- Nonport side view of Brocade 6510 15
- Overview of Brocade 6520 16
- Platform components of 6520 17
- Port side of the Brocade 6520 19
- Non-port side of the Brocade 6520 19
- Overview of Brocade 5100 20
- Port side view of Brocade 5100 20
- Overview of Brocade 300 22
- Port side view of Brocade 300 22
- Overview of Brocade 5300 23
- Port side view of Brocade 5300 24
- Access Gateway support 25

Chapter 2: Basic Configuration 27

- Assigning IP address 27
  - Creating serial connection 27
  - Using DHCP (Dynamic Host Configuration Protocol) to assign IP address 27
  - Static IP address 28
  - IPv6 Auto configuration 29
  - View IP configuration 29
- Domain IDs 29
  - Viewing your Domain ID 29
  - Setting your Domain ID 30
- Ports 30
  - Enable a port 30
  - Disable a port 31
  - Default Port Name 31
  - Ports on Demand license 32
  - Setting Port Speed 33
  - Setting Port name 33
  - Swapping port area IDs 34
- Customizing the switch name 34
- Checking Switch Status 34
- Chassis names 34
- Customizing chassis names 35
- 36
- Switch activation and deactivation 36
  - Disabling a switch 36

Enabling a switch 36
Switch shutdown 36
Powering off a Brocade switch 37
Setting the date and time 37
Synchronizing the local time with an external source 37

.... 38
Webtools 38
Configuring IP and netmask information 38
Switch configuration 39
Enabling and disabling a switch 40
Changing the switch name 40
Changing the switch domain ID 40
Viewing and printing a switch report 40
Switch restart 41
Performing a fast boot 41
Performing a reboot 41
Configuring fabric settings 41
Assigning a name to a port 42
Enabling and disabling a port 43
Persistent enabling and disabling ports 43
Enabling and disabling NPIV ports 43
Configuring BB credits on an F_Port 44
Chapter 3: Account Management 45
Overview 45
User Roles 45
Local database user accounts 45
Creating account: 45
Displaying Account Information: 46
Deleting account: 46
Modifying account: 46
Changing password for current login account 47
Changing password for different login account 47
Local account database distribution 48
Distributing the local user database 48
Accepting distribution of user databases on the local switch 48
Rejecting distributed user databases on the local switch 48
Password policies 48
Enabling the admin lockout policy 49
Unlocking an account 49
Disabling the admin lockout policy 49
Authentication servers on the switch 49
Adding a RADIUS or LDAP server to the switch configuration 49
Enabling and disabling a RADIUS or LDAP server 50
Deleting a RADIUS or LDAP server from the configuration 50

Changing a RADIUS or LDAP server configuration 50
Changing the order in which RADIUS or LDAP servers are contacted for service 50
Displaying the current RADIUS configuration 50
Configuring local authentication as backup 51
Chapter 4: Setting the Protocols 52
Secure Copy 52
Setting up SCP for configUploads and downloads 52
Secure Shell protocol 52
Allowed-user 52
Configuring SSH authentication 52
Deleting keys on the switch 54
Telnet protocol 54
Blocking Telnet 54
Unblocking Telnet 55
Chapter 5: Configuration file and Firmware management 56
Configuration file backup 56
Uploading a configuration file in interactive mode 56
Configuration file restoration 56
Installing firmware 57
Firmware download from network 57
Firmware download from a USB device 58
.... 59
Webtools 59
Creating a configuration backup file 59
Restoring a configuration 61
Uploading and downloading from USB storage 63
Performing a firmware download 63
Chapter 6: Licensing 66
Licensing overview 66
8G License 66
For more on Time based licenses see Fabric OS Administrator’s Guide. 68
Viewing installed licenses 68
Adding a licensed feature 68
Removing a licensed feature 69
Ports on Demand 69
Activating Ports on Demand 70
Displaying the port license assignments 70
.... 70
Web tools 70
Licensed feature management 70
Activating a license on a switch 71
Removing a license from a switch 71
Universal time based licensing 72

Chapter 7: Virtual Fabrics 73
Overview 73
Enabling Virtual Fabric mode 73
Logical Switches 74
Creating Logical Switches 74
Assigning/Removing ports to logical switches 76
Displaying logical switch configuration 77
Changing Fabric ID of switch 77
Setting /Removing IP address for fabric 78
Logical Fabric and ISL sharing 78
Configuring the switch to use XISL 78
Deleting a logical switch 79
Disable Virtual Fabrics 79
Web Tools 80
NPIV overview 83
Configuring NPIV 84
Viewing NPIV port configuration information 85
Chapter 8: Zoning 87
Overview 87
Zoning Configurations 87
Zone aliases 88
Creating an aliases 88
Adding members to alias 88
Removing members from alias 89
Deleting an aliases 89
Viewing aliases in defined configuration 89
Creating a Zone 90
Adding members to zone 90
Removing members from zone 90
Deleting a zone 91
Viewing zoning in defined configuration 91
Default zoning mode 91
Zoning database size 92
Zoning Configurations 92
Creating zoning configurations 92
Adding zones to zoning configuration 92
Removing zones from zoning configuration 93
Enable a zoning configuration 93
Disabling a zoning configuration 94
Deleting a zone configuration 94
Clearing changes to configuration 94
Viewing all zone configuration information 95
Viewing selected zone configuration 95
Viewing configuration in effective zone database 95
Clearing all zone configurations 96
Zone object maintenance 96

Copying a zone object 96
Deleting a zone object....................97
Renaming a zone object 98
Zoning configuration management....................98

....................99
Zone merging 99
For more on zone merging please refer to Fabric OS Administrator’s Guide. 99
Web Tools....................99
Zoning configurations 99
Opening the Zone Administration window 99
Setting the default zoning mode 99
Zoning management 100
Refreshing fabric information 101
Saving local zoning changes.................... 101
Select a zoning view.................... 101
Creating and populating zone aliases 102
Adding and removing members of a zone alias 102
Renaming zone aliases 102
Deleting zone aliases 103
Creating and populating zones.................... 103
Adding and removing members of a zone 103
Renaming zones.................... 104
Cloning zones 104
Deleting zones.................... 104
Creating zone configurations.................... 105
Adding or removing zone configuration members.................... 105
Renaming zone configurations.................... 105
Cloning zone configurations 105
Deleting zone configurations 106
Enabling zone configurations 106
Disabling zone configurations 106
Displaying enabled zone configurations 107
Adding a WWN to multiple aliases and zones 107
Removing a WWN from multiple aliases and zones.................... 107
Replacing a WWN in Multiple Aliases and Zones 108
Searching for zone members.................... 108
Clearing the Zoning Database 108
Zone configuration analysis.................... 109
Best practices for zoning 109
Dynamic Fabric Provisioning: Fabric Assigned World Wide Name 109
Chapter 9: Routing and Trunking....................110
Inter-Switch Links (ISLs) 110
For more on In-flight Encryption and Compression please refer to Fabric OS Adminstrator’s Guide. 112
Forward Error Correction (FEC) on Brocade 6510’s Condor3 ASIC 112
This enables Condor 3 to recover bit errors in a 10 and 16 Gbps data stream (frames and primitives)....................112
Condor3 can correct up to 11 error bits in every 2112-bit transmission enhancing reliability of transmission..112
This feature is enabled by default on Condor 3 ports. 112

For more FEC please refer to Fabric OS Adminstrator's Guide. 112
Routing policies......112
Displaying the current routing policy...... 112
Setting the routing policy 112
Setting up the AP route policy...... 113
Route selection 113
Dynamic Load Sharing 113
Setting DLS 113
Trunking overview 114
Basic trunk group configuration......115
Re-initializing ports for trunking...... 115
Enabling Trunking on a port...... 115
Enabling Trunking on a switch 115
Displaying trunking information 115
F_Port trunking 116
Enabling F_Port trunking...... 116
Disabling F_Port trunking...... 116
FC-FC Routing......116
Setting up the FC-FC routing service 117
Verifying the setup for FC-FC routing...... 117
Assigning backbone fabric IDs 118
Creating an FCIP tunnel 118
Inter-fabric link configuration 119
Configuring an IFL for both edge and backbone connections 119
Setting router port cost for an EX_Port 122
Configuring EX_Port frame trunking...... 122
LSAN zone configuration 123
Controlling device communication with the LSAN...... 123

......125
Web Tools......125
Disabling or enabling ISL trunking 125
Viewing trunk group information 126
F_Port trunk groups...... 127
Creating and maintaining F_Port trunk groups...... 128
FC-FC routing management 129
Opening the FC Routing module...... 130
Configuring an EX_Port 130
Editing the configuration of an EX_Port 131
Viewing LSAN zones 131
Viewing LSAN Devices...... 132
Configuring the backbone fabric ID...... 132
Enabling Access Gateway mode...... 132
Disabling Access Gateway mode...... 133
Viewing the Access Gateway settings 133
Port configuration...... 134
Creating port groups 134
Editing or Viewing port groups 136

Deleting port groups....... 137
Access Gateway policy modification....... 138
Path Failover and Failback policies....... 138
Modifying Path Failover and Failback policies....... 138
Enabling the Automatic Port Configuration policy 138
Chapter 10: Brocade 8000.......140
Brocade 8000 overview 140
FCoE and Layer 2 capabilities.......140
A number of FCoE enhancements have been made with FOS 7.0.0 and higher. For more on these please refer to Fabric OS Adminstrator's Guide........141
Fibre Channel capabilities.......141
Port side of Brocade 8000.......141
CEE Command Line Interface 142
Accessing the CEE CLI through the console or Telnet....... 142
Saving your configuration changes 143
Saving configuration changes with the write command 143
CEE Commands 144
Using the do command as a shortcut 145
CEE Integrations 145
Integrating a Brocade 8000 switch on a SAN....... 146
CEE and LAN integration....... 146
CEE map attributes 146
Configuring DCBX 147
Configuring Spanning Tree Protocol....... 148
Configuring VLAN Membership....... 149
Configuring the CEE Interfaces....... 150
Server connections to the Brocade 8000 switch 151
Fibre Channel configuration for the CNA 151
Ethernet configuration for the CNA 151
Minimum CEE configuration to allow FCoE traffic flow.......151
Fabric OS Native and Access Gateway modes.......153
Enabling Access Gateway mode....... 153
Disabling Access Gateway mode....... 154
Web Tools.......154
FCOE configuration tasks....... 154
Quality of Service (QoS) configuration....... 155
Adding a DCB map 155
Adding a traffic class map 157
LLDP-DCBX configuration....... 158
Configuring global LLDP characteristics 158
Adding an LLDP profile....... 160
Configuring DCB interfaces....... 163
Configuring a link aggregation group (LAG) 164
Configuring VLANs....... 167
Configuring FCoE login groups....... 168
Displaying FCoE Port Information....... 170
Displaying LAG information....... 173

Displaying VLAN information 173
Displaying FCoE login groups 174
Displaying QoS information 174
Displaying LLDP-DCBX information 175
Displaying DCB interface statistics 176
Enabling and disabling a DCB interface 179
Enabling and disabling a LAG 180
Enabling and disabling LLDP 180
Enabling and disabling QoS priority-based flow control 180
Enabling and disabling FCoE ports 180

Chapter 11: Brocade 7800 181

Overview of Brocade 7800 Extension Switch 181

Feature comparison - base 7800 and with the Upgrade License 182

Available licenses 183

Port Side of Brocade 7800 184

Installing SFPs and cabling the Brocade 7800 185

7800 Switch license option 186

Fibre Channel routing services configuration 186

FCIP 186

VE_Ports and FCIP tunnels on the 7800 switch 187
Configuration preparation 187

Configuration steps 187

Setting VE_ports to persistently disabled state 188
Configuring VEX_ports 188
Configuring the media type for GbE ports 0 and 1 188
Configuring a GbE or XGE port IP address 188
Configuring an IP route 189
Validating IP connectivity 189
Creating an FCIP tunnel 190
Creating additional FCIP circuits 190
Verifying the FCIP tunnel configuration 191
Enabling persistently disabled ports 191
Modifying an FCIP tunnel 191
Modifying an FCIP circuit 191
Deleting an IP interface 191
Deleting an IP route 191
Deleting an FCIP tunnel 192
Deleting an FCIP circuit 192

Chapter 12: Getting technical help 192

SupportSave 193

Few Notes on FOS v7.0.0 Updates .... Error! Bookmark not defined.

CHAPTER 1: BROCADE SWITCH INTRODUCTION

Overview of Brocade 6505

The Brocade 6505 is a 24-port auto-sensing 2, 4, 8, or 16 Gbps Fibre Channel (FC) switch that delivers the latest Brocade single-chip architecture for Fibre Channel Storage Area Networks (SANs). The Brocade 6505 is a small-to-midsize business-class switch that is designed to handle smaller-scale SAN requirements. The Brocade 6505 provides up to 24 ports in a single height (1U) switch that enables the creation of very dense fabrics in a relatively small space. The Brocade 6505 offers Ports on Demand (POD) licensing as well. “Base” models of the switch contain 12 ports, and an additional 12-port POD license can be purchased. The base model also offers a single power supply and fan module with a second module available as an upgrade for redundancy.

The Brocade 6505 supplies Reliability, Availability, and Serviceability (RAS) performance and the scalability requirements of an enterprise switch along with interoperability and ease-of-use advantages.

- Up to 24 auto-sensing ports of high-performance 16-Gbps technology in a single domain.
- Ports on Demand scaling from 12 to 24 ports.
- 2, 4, 8, and 16 Gbps auto-sensing Fibre Channel switch and router ports.
    - 2, 4, and 8 Gbps performance is enabled by 8 Gbps SFP+ transceivers.
    - 4, 8, and 16 Gbps performance is enabled by 16 Gbps SFP+ transceivers.
- Universal ports self-configure as E, F, or M ports. EX_Ports can be activated on a per-port basis with the optional Integrated Routing license. D-port functionality is also available for diagnostics.
- Airflow is set for port side exhaust.
- Inter-Switch Link (ISL) Trunking, which allows up to eight ports (at 2, 4, 8, or 16 Gbps speeds) between a pair of switches combined to form a single, logical ISL with a speed of up to 128 Gbps (256 Gbps full duplex) for optimal bandwidth utilization and load balancing. The base model permits one eight-port trunk plus one four-port trunk.
- Dynamic Path Selection (DPS), which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- Brocade-branded SFP+ optical transceivers that support any combination of Short Wavelength (SWL), Long Wavelength (LWL), and Extended Long Wavelength (ELWL) optical media among the switch ports.
- Extended distance support enables native Fibre Channel extension up to 7,500 km at 2 Gbps.
- Support for unicast traffic type.
- Brocade Fabric OS, which delivers distributed intelligence throughout the network and enables a wide range of value-added applications including Brocade Advanced Web Tools, Brocade Enhanced Group Management, and Brocade Zoning.
- Licensable fabric services include:
    - Adaptive Networking with QoS
    - Brocade Extended Fabrics
    - Brocade Fabric Watch
    - ISL Trunking
    - Advanced Performance Monitoring (APM)
    - Server Application Optimization (SAO)
- Support for Access Gateway configuration where server ports connected to the fabric core will be virtualized.
- Hardware zoning is accomplished at the port level of the switch and by World Wide Name (WWN). Hardware zoning permits or denies delivery of frames to any destination port address.
- Extensive diagnostics and system-monitoring capabilities for enhanced high Reliability, Availability, and Serviceability (RAS).
- The Brocade EZSwitchSetup wizard that makes SAN configuration a three-step point-and-click task.

- Real-time power monitoring enables users to monitor real-time power usage of the fabric at a switch level.
- Port-to-port latency minimized to 800 nanoseconds through the use of cut-through frame routing at 16 Gbps.

Platform Components of 6505

- A system motherboard that features a PowerPC 440EPx Reduced Instruction Set Computer (RISC) CPU running at 667 MHz, with integrated peripherals.
  Brocade 6505 Hardware Reference Manual 3
  53-1002449-01
- An RJ45 10/100 BaseT Ethernet system management port, in conjunction with Brocade EZSwitchSetup, that supports switch IP address discovery and configuration, eliminating the need to attach a serial cable to configure the switch IP address and greatly increasing the ease of use.
- One RS-232 serial port with an RJ45 connector for initial switch setup (if not using EZSwitchSetup) and factory default restoration.
- A USB 2.0 port that provides storage for firmware updates, output of the supportSave command, and storage for configuration uploads and downloads.
- One power supply and fan assembly in the base model. There are two fans per assembly. A second assembly is available for redundancy and hot-swap capability.
- One LED (green/amber) per FC port to indicate status.
- One LED (green) for system power.
- One LED (green/amber) for system status.
- Two Ethernet port LEDs (integrated with RJ45) for speed and port activity. (A green LED for port speed and an amber LED for port activity.)
- SEEPROM for switch identification.
- Voltage monitor.
- Fan monitor.
- Temperature monitor.
- Real-time clock (RTC) with battery.

Port side of the Brocade 6505

The port side of the Brocade 6505 includes the system status LED, the console port, the Ethernet port and accompanying LEDs, the USB port, and the Fibre Channel ports and corresponding port status LEDs.

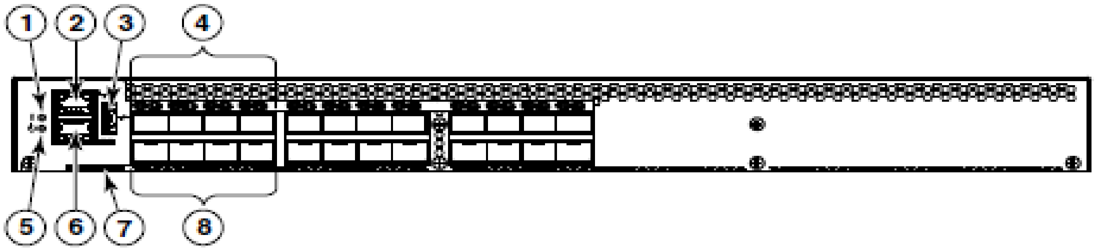

1	System status LED	5	System power LED
2	Management Ethernet port with LEDs	6	Serial console port
3	USB port	7	Switch ID pull-out tab
4	FC ports 0-3 (all LEDs above)	8	FC ports 4-7

Figure 1: Port side of the Brocade 6505

Nonport side of the Brocade 6505

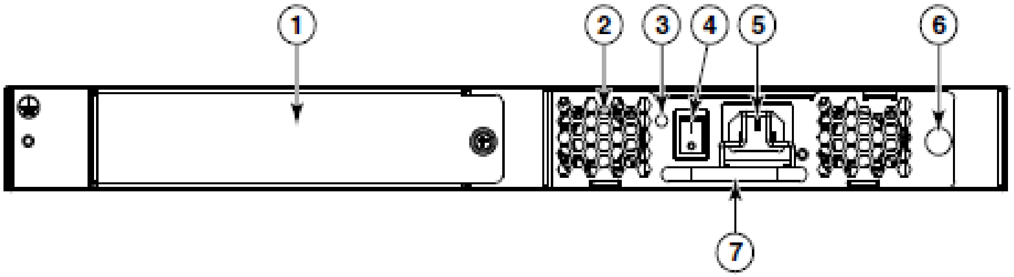

1	Filler panel	5	Power plug receptacle (with plug retainer)
2	Power supply and fan assembly #1	6	Captive screw
3	Power supply and fan assembly LED	7	Handle
4	On/off switch		

Figure 2: Nonport side of the Brocade 6505

Overview of Brocade 6510

The Brocade 6510 is a 48-port auto-sensing 2, 4, 8, or 16 Gbps as well as 10 Gbps Fibre Channel (FC) switch that delivers the latest Brocade single-chip architecture for Fibre Channel Storage Area Networks (SANs). The Brocade 6510 is an enterprise-class switch that is designed to handle the large-scale SAN requirements of an enterprise, and can also be used to address the SAN requirements of a small to medium-sized workgroup.
It provides 48 ports in a single (1U) height switch that enables the creation of very dense fabrics in a relatively small space.

The Brocade 6510 supplies Reliability, Availability, and Serviceability (RAS) performance and scalability requirements of an enterprise switch along with interoperability and ease-of-use advantages.

- Up to 48 auto-sensing ports of high-performance 16 Gbps technology in a single domain.
- Ports on Demand scaling from 24 to 36 or 48 ports.
- 2, 4, 8, and 16 Gbps auto-sensing Fibre Channel switch and router ports.
    - 2, 4, and 8 Gbps performance is enabled by 8 Gbps SFP+ transceivers.
    - 4, 8, and 16 Gbps performance is enabled by 16 Gbps SFP+ transceivers.
- 10 Gbps manual set capability on FC ports (requires the optional 10 Gigabit FCIP/Fibre Channel license).
    - 10 Gbps performance is enabled by 10 Gbps SFP+ transceivers.
    - Ports can be configured for 10 Gbps for metro connectivity (on the first eight ports only).
- Universal ports self-configure as E, F, M, or D ports. EX_Ports can be activated on a per port basis with the optional Integrated Routing license.
    - Brocade Diagnostic Port (D-Port) feature provides physical media diagnostic, troubleshooting and verification services.
- In-flight data compression and encryption on up to two ports provides efficient link utilization and security.
- Options for port side exhaust (default) or nonport side exhaust airflow for cooling.
- Virtual Fabric support to improve isolation between different VFs.
- Fibre Channel Routing (FCR) service, available with the optional Integrated Routing license provides improved scalability and fault isolation.
- FICON, FICON Cascading, and FICON Control Unit Port ready.
- Inter-Switch Link (ISL) Trunking (licensable), which allows up to eight ports (at 2, 4, 8, or 16 Gbps speeds) between a pair of switches combined to form a single, logical ISL with a speed of up to 128 Gbps (256 Gbps full duplex) for optimal bandwidth utilization and load balancing.
- Dynamic Path Selection (DPS), which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- Brocade-branded SFP+ optical transceivers that support any combination of Short Wavelength (SWL) and Long Wavelength (LWL) or Extended Long Wavelength (ELWL) optical media among the switch ports.
- Extended distance support enables native Fibre Channel extension up to 7,500 km at 2 Gbps.
- Support for unicast, multicast (255 groups), and broadcast data traffic types.
- Brocade Fabric OS, which delivers distributed intelligence throughout the network and enables a wide range of value-added applications including Brocade Advanced Web Tools and Brocade Zoning. Optional Fabric Services include: Adaptive Networking with QoS, Brocade Extended Fabrics, Brocade Enhanced Group Management, Brocade Fabric Watch, ISL Trunking, and End-to-End Performance Monitoring (APM).
- Support for Access Gateway configuration where server ports connected to the fabric core will be virtualized.
- Hardware zoning is accomplished at the port level of the switch and by World Wide Name (WWN). Hardware zoning permits or denies delivery of frames to any destination port address.
- Extensive diagnostics and system-monitoring capabilities for enhanced high Reliability, Availability, and Serviceability (RAS).
- 10G Fibre Channel integration on the same port provides for DWDM metro connectivity on the same switch (can be done on first eight ports only).
- The Brocade EZSwitchSetup wizard that makes SAN configuration a three-step point-and-click task.
- Real time power monitoring enables users to monitor real time power usage of the fabric at a switch level.
- Port-to-port latency minimized to 800 nanoseconds through the use of cut-through frame routing at 16 Gbps.

Platform Components of 6510

- A system motherboard that features a PowerPC 440EPx Reduced Instruction Set Computer (RISC) CPU running at 667 MHz, with integrated peripherals, and that provides high performance with low power consumption.
- An RJ45 10/100 Base T Ethernet system management port, in conjunction with EZSwitchSetup, that supports switch IP address discovery and configuration, eliminating the need to attach a serial cable to configure the switch IP address and greatly increasing the ease of use.

- One RS-232 serial port with RJ45 connector for initial switch setup (if not using EZSwitch Setup) and factory default restoration (the integral LEDs remain unlit at all times).
- A USB port that provides storage for firmware updates, output of the supportSave command and storage for configuration uploads and downloads.
- Two hot-swappable, redundant power supply and fan FRUs. There are two fans per FRU.
- Rack-mount design (1U form factor) in a 19-inch EIA rack.
- One LED (green/amber) per FC port to indicate status.
- One LED (green) for system power.
- One LED (green/amber) for system status.
- Two Ethernet LEDs (integrated with RJ45) for speed and activity status.
- SEEPROM for switch identification.
- Voltage monitoring.
- Fan monitoring including flow direction.
- Temperature monitoring.
- Real-time clock (RTC) with battery.
- The Brocade EZSwitchSetup wizard that makes SAN configuration a three-step point-and-click task.

Port side view of Brocade 6510

The port side of the Brocade 6510 includes the system status LED, console port, Ethernet port and LEDs, USB port, and Fibre Channel ports and the corresponding port status LEDs.

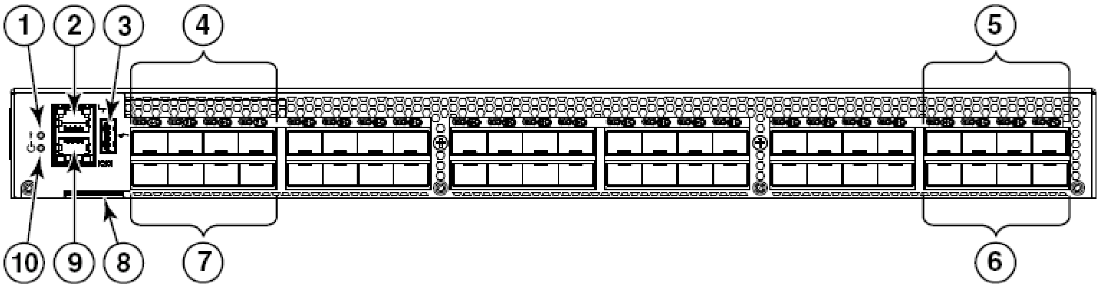

1	System status LED	6	FC ports 44-47
2	Management Ethernet port with LEDs	7	FC ports 4-7
3	USB port	8	Switch ID pull-out tab
4	FC ports 0-3(all LEDs above)	9	Serial console port
5	FC ports 40-43	10	System power LED

Figure 1: Port side of the Brocade 6510.

NOTE:

1. The two LEDs on the serial console port are non-functional. Brocade 6510 has dual power supplies with integrated fans. It is 1U with reversible airflow option part numbers: Port side to non-port side and non-port side to port side.
2. The power supply FRU is hot-swappable with N+1 redundancy. The unit is auto-ranging to accommodate input voltages of 85 V to 264 V ~5 A to 2.5 A.

Nonport side view of Brocade 6510

Figure 4 shows the nonport side of the Brocade 6510, which contains the power supply (including the AC power receptacle and AC power switch) and fan assemblies.

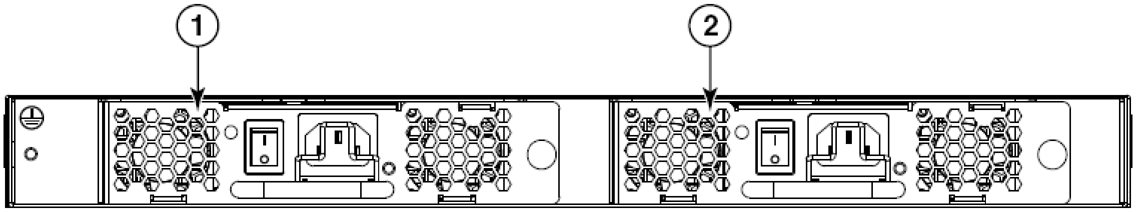

1 Power supply/fan assembly #2 2 Power supply/fan assembly #1

Figure 2: Nonport side of the Brocade 6510

Overview of Brocade 6520

The Brocade 6520 is a 96-port auto-sensing 2, 4, 8, or 16 Gbps, as well as 10 Gbps, Fibre Channel (FC) switch that delivers the latest Brocade multi-chip architecture for Fibre Channel Storage Area Networks (SANs). The Brocade 6520 is an enterprise-class switch that is designed to handle the large-scale SAN requirements of an enterprise, and can also be used to address the SAN requirements of a small to medium-sized workgroup.

The Brocade 6520 provides up to 96 ports in a double height (2U) switch that enables the creation of very dense fabrics in a relatively small space.

The Brocade 6520 offers Ports on Demand (POD) licensing as well. “Base” models of the switch contain 48 ports, and up to two additional 24-port POD licenses can be purchased to fill all 96 ports.

Port activation works through a process called Dynamic Ports on Demand (DPOD). With DPOD, ports are licensed as they come online. For instance, if you have a base model with 48 port licenses, the first 48 ports to come online, regardless of their numbering, are licensed. Once all the licenses have been assigned, you can manually move those licenses from one port to another if you choose.

The first eight ports can be configured to run at 10 Gbps with the appropriate licensing.

The Brocade 6520 supplies Reliability, Availability, and Serviceability (RAS) performance and scalability requirements of an enterprise switch along with interoperability and ease-of-use advantages.

The Brocade 6520 is only 24 inches deep and has airflow direction options. You can order either port side exhaust (the default configuration) or non-port side exhaust airflow to accommodate specific installations.

The Brocade 6520 offers the following features and capabilities:

- Up to 96 auto-sensing ports of high-performance 16 Gbps technology in a single domain.
- Ports on Demand scaling from 48 to 72 or 96 ports.
- Port licensing via DPOD
- 2, 4, 8, and 16 Gbps auto-sensing Fibre Channel switch and router ports.
  - 2, 4, and 8 Gbps performance is enabled by 8 Gbps SFP+ transceivers.
  - 4, 8, and 16 Gbps performance is enabled by 16 Gbps SFP+ transceivers.
- 10 Gbps manual set capability on FC ports (requires the optional 10 Gigabit FCIP/FibreChannel license) on the first eight ports only.
  - Ports can be configured for 10 Gbps for metro connectivity.
  - 10 Gbps performance is enabled by 10 Gbps Fibre Channel SFP+ transceivers.
- FC ports will self-configure as E_ports and F_ports. EX_ports can be activated on a per-port basis with the optional Integrated Routing license.
  - Mirror ports (M_ports) and diagnostic ports (D_ports) must be manually configured.
  - The Brocade Diagnostic Port (D_port) feature provides physical media diagnostic, troubleshooting, and verification services.
- In-flight data compression and encryption on up to 16 ports (up to 8 ports at 16 Gbps) provides efficient link utilization and security.
- Options for port side exhaust (default) or non-port side exhaust airflow for cooling.
- Virtual Fabric (VF) support to improve isolation between different VFs.
- Fibre Channel Routing (FCR) service, available with the optional Integrated Routing license, provides improved scalability and fault isolation.
- Inter-Switch Link (ISL) Trunking (licensable), which allows up to eight ports (at 2, 4, 8, or 16 Gbps speeds) between a pair of switches combined to form a single, logical ISL with a speed of up to 128 Gbps (256 Gbps full duplex) for optimal bandwidth utilization and load balancing. There is no limit to how many trunk groups can be configured.
- Dynamic Path Selection (DPS), which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- Brocade-branded SFP+ optical transceivers that support any combination of Short Wavelength (SWL) and

Long Wavelength (LWL) or Extended Long Wavelength (ELWL) optical media among the switch ports.

- Extended distance support enables native Fibre Channel extension up to 7,500 km at 2 Gbps.
- Support for unicast data traffic types.
- Brocade Fabric OS, which delivers distributed intelligence throughout the network and enables a wide range of value-added applications including Brocade Advanced Web Tools and Brocade Zoning. Optional Fabric Services include Adaptive Networking with QoS, Brocade Extended Fabrics, Brocade Enhanced Group Management, Brocade Fabric Watch, ISL Trunking, and End-to-End Advanced Performance Monitoring (APM).
- Hardware zoning is accomplished at the port level of the switch and by World Wide Name (WWN). Hardware zoning permits or denies delivery of frames to any destination port address.
- Extensive diagnostics and system-monitoring capabilities for enhanced high Reliability, Availability, and Serviceability (RAS).
- 10Gbps Fibre Channel integration on the same port provides for DWDM metro connectivity on the same switch (can be done on first eight ports only with appropriate licensing).
- The Brocade EZSwitchSetup wizard that makes SAN configuration a three-step point-and-click task.
- Real-time power monitoring enables users to monitor real-time power usage of the fabric at a switch level.

- Local port latency minimized to 700 nanoseconds (ns) through the use of cut-through frame routing at 16 Gbps.
- Switch latency of 2100 ns (L2 latency without forward error correction)

Platform components of 6520

- A system motherboard that features a primary CPU running at 1.20 GHz, with integrated peripherals.
- One 2 GB DDR2 running at 400 MHz.
- Boot memory of 8 MB.
- One 2 GB compact flash card.
- Up to 96 16 Gbps Fibre Channel ports.
- An RJ45 10/100/1000 BaseT Ethernet system management port (RJ45 connector), in conjunction with EZSwitchSetup, that supports switch IP address discovery and configuration, eliminating the need to attach a serial cable to configure the switch IP address.
- One RS-232 console (serial) port with an RJ45 connector for initial switch setup (if not using
- EZSwitchSetup) and factory default restoration.
- One USB 2.0 port that provides storage for firmware updates, output of the supportSave command, and storage for configuration uploads and downloads.
- Two hot-swappable, 80+ Platinum certified, redundant power supplies.
- Three hot-swappable fan FRUs.
- One LED (green/amber) per FC port to indicate status.
- One LED (green) for system power.

- One LED (green/amber) for system status.
- Two Ethernet LEDs (integrated with RJ45) for speed and activity status.
- Two LEDs per power supply: one green for AC line in status and one green/amber for DC power out.
- One LED (green/amber) per fan.
- SEEPROM for switch identification.
- Voltage monitoring.
- Fan monitoring including flow direction.
- Temperature monitoring.
- Real-time clock (RTC) with battery.

Port side of the Brocade 6520

The port side of the Brocade 6520 includes the system status LED, console port, Ethernet port and LEDs, USB port, and Fibre Channel ports and the corresponding port status LEDs. Figure 1 shows the port side of the Brocade 6520.

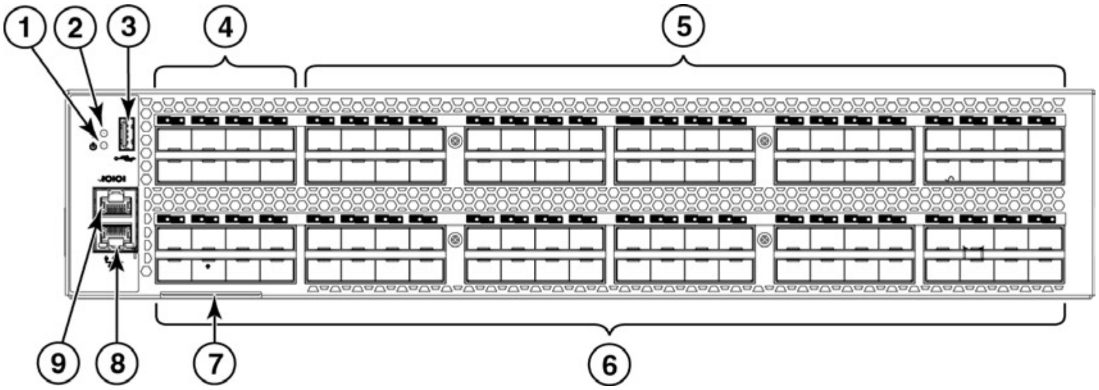

1 System power LED

2 System status LED

3 USB port

4 FC ports 0-7 (all LEDs above)

5 FC ports 8-47

6 FC ports 48-95

7 Switch ID pull-out tab

8 Management Ethernet port with LEDs

9 Serial console port

FIGURE 1 **Port side view of the Brocade 6520**

Non-port side of the Brocade 6520

Figure 2 shows the non-port side of the Brocade 6520, which contains the power supplies (including the AC power receptacle) and fans.

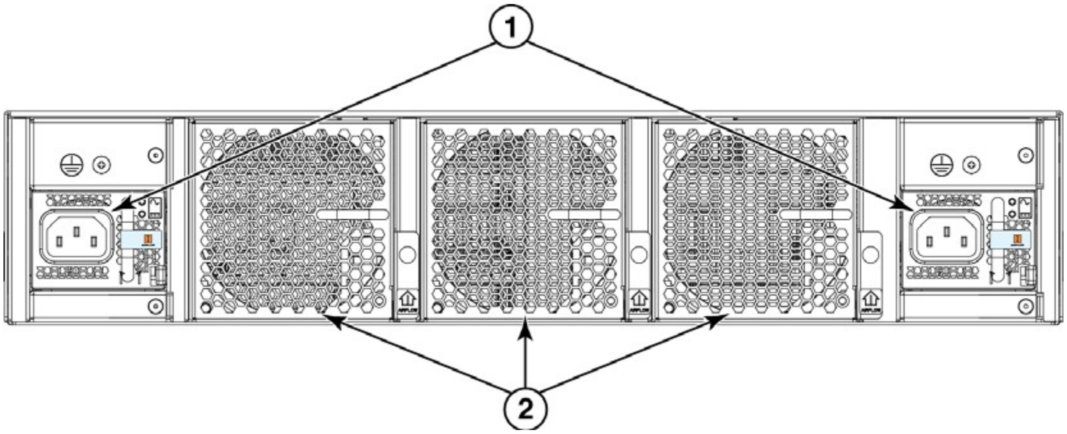

1 Power supplies with integral fans

2 Fans

FIGURE 2 **Non-port side of the Brocade 6520**

Overview of Brocade 5100

The Brocade 5100 is an Enterprise class 1U, 40-port Fibre Channel 1, 2, 4 or 8 Gbps Fibre Channel switch that offers the next generation Brocade, single-chip architecture for Storage Area Networks (SANs). The Brocade 5100 is designed to function in large-scale enterprise SANs and can also fit the requirements of small to medium-sized work groups.
Because the Brocade 5100 has a slim 1U height and a high port count, you can use the Brocade 5100 to create very dense fabrics in a relatively small space. With its flexible Ports On Demand (POD) capability, the Brocade 5100 provides excellent overall value as the foundation of a SAN with the ability to grow with an organization's SAN needs.

- Up to 40 ports of high-performance 8 Gbps technology and POD scaling from 24 to 32 or 40 ports.
- Support for 1, 2, 4, and 8 Gbps auto-sensing Fibre Channel switch and router ports.
- FICON®, FICON Cascading and FICON Control Unit Port ready.
- Two hot-swappable, redundant integrated power supply and fan FRUs.
- Universal ports that self-configure as E, F, M, or FL ports. Ex_Ports are activated on a per port basis with the optional Integrated Routing license.
- Fibre Channel Routing (FCR) service that provides improved scalability and fault isolation (through the optional Integrated Routing license).
- An RJ45 Ethernet management port that in conjunction with EZSwitchSetup, supports switch IP address discovery and configuration, eliminating the need to attach a serial cable to configure the switch IP address and greatly increasing the ease of use.
- USB port that provides storage for firmware updates, output of the supportSave command and storage for configuration uploads and downloads
- Inter-Switch Link (ISL) Trunking (licensable), which allows up to eight ports (at 1, 2, 4, or 8 Gbps speeds) between a pair of switches combined to form a single, logical ISL with a speed of up to 128 Gbps full duplex for optimal bandwidth utilization and load balancing.
- Dynamic Path Selection (DPS), which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- Industry-leading extended distance support, which enables native Fibre Channel extension greater than 590 km.
- Expanded security for up to 16,000 hardware zones. Hardware zoning is accomplished at the port level of the switch or by World Wide Name (WWN). Hardware zoning permits or denies delivery of frames to any destination port address.
- Unicast, multicast (255 groups), and broadcast data traffic type, are support.
- Brocade Small Form-Factor Pluggable (SFP) or SFP+ optical transceivers support any combination of Short Wavelength (SWL), Long Wavelength (LWL) or Extended Long Wavelength (ELWL) optical media among the switch ports.
- Brocade Fabric Operating System (Fabric OS), which delivers distributed intelligence throughout the network and enables a wide range of value-added applications including Brocade Advanced Web Tools and Brocade Zoning. Optional Fabric Services include: Adaptive Networking with QoS, Brocade Extended Fabrics, Brocade Enhanced Group Management, Brocade Fabric Watch, ISL Trunking, Integrated Routing, and End-to-End Performance Monitoring (APM).
- Port-to-port latency minimized to 700 nanoseconds through the use of cut-through frame routing at 8 Gbps.

Port side view of Brocade 5100

The port side of the Brocade 5100 includes the system status LED, console port, Ethernet port and LEDs, USB port, and Fibre Channel ports and the corresponding port status LEDs. Figure 3 shows the port side of the Brocade 5100.

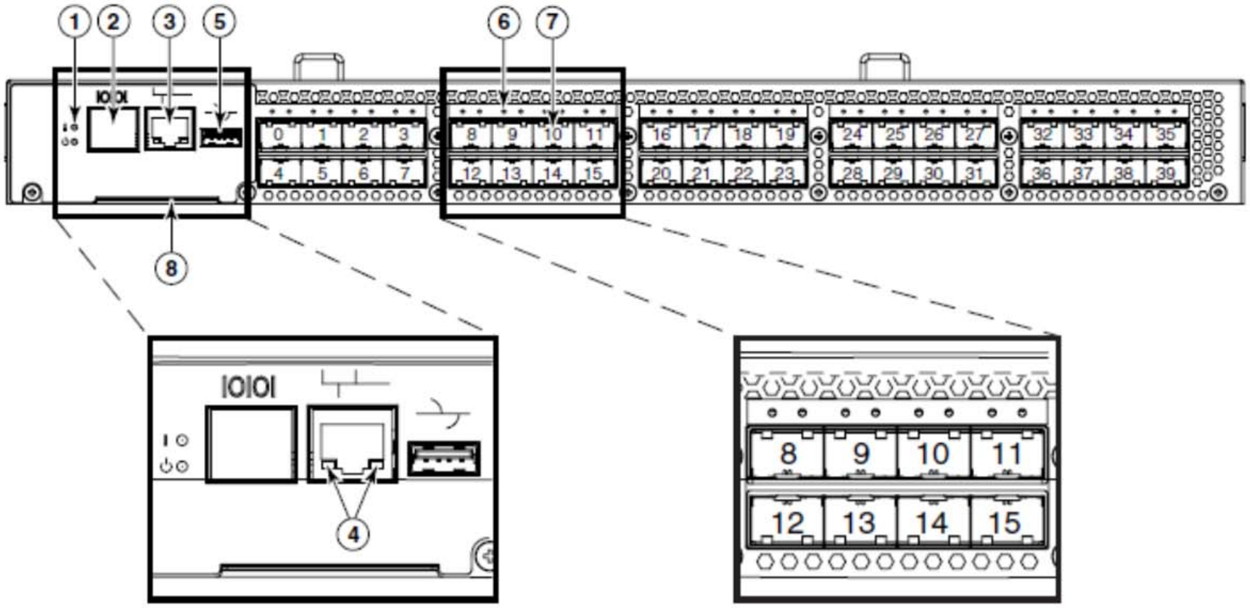

1 System status (top) and power (bottom) LEDs

2 System RS232 console port (RJ-45)

3 System Ethernet port (RJ-45)

4 Ethernet port LEDs (green/amber)

5 USB port

6 Fibre Channel port status LED

7 Fibre Channel ports

8 Switch ID pull-out tab

Figure 3: Port-side view of the Brocade 5100

The Fibre Channel ports on the Brocade 5100 are numbered from left to right, in eight-port groups from 0 to 39 as illustrated in Figure 4.

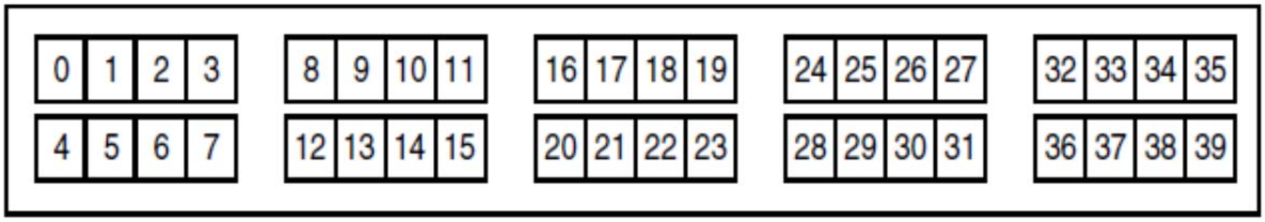

0	1	2	3	8	9	10	11	16	17	18	19	24	25	26	27	32	33	34	35
4	5	6	7	12	13	14	15	20	21	22	23	28	29	30	31	36	37	38	39

Figure 4: Port Numbering on the Brocade 5100

Overview of Brocade 300

The Brocade 300 is a cost-effective and highly-scalable 1, 2, 4, or 8 Gbps switch, designed for small to mid-sized businesses. Like all Brocade switches, the Brocade 300 runs on the Brocade Fabric Operating System (Fabric OS) and is compatible with other Brocade switches, which enables seamless connectivity into heterogeneous SAN environments.
The Brocade 300 is a dual purpose device that you can use either as a full-functioned switch or as an N_Port ID Virtualization (NPIV) access gateway. When functioning as an access gateway, the Brocade 300 provides a single platform for all SAN connectivity.

- 1U chassis that can be installed as a standalone unit or mounted in a standard Electronic Industries Association (EIA) 48.26 cm (19 inches) cabinet.
- One built-in fixed power unit (not a FRU).
- Three built-in fans (there are no fan FRUs) that allows a single fan failure and permits the switch to continue to function properly.
- On-demand scaling of 8 to 24 8 Gbps ports.
- ASIC technology supporting 1, 2, 4 and 8 Gbps auto-sensing Fibre Channel ports.
- A flexible design that enables the Brocade 300 to function as either a full-functioned switch or an NPIV access gateway.
- RJ45 Ethernet management port that in conjunction with EZSwitchSetup supports switch IP address discovery and configuration.
- USB port that provides storage for firmware updates, output of the supportSave command and storage for configuration uploads and downloads.
- Inter-Switch-Link Trunking (licensable) which enables up to eight ports (at 1, 2, 4, or 8 Gbps speeds) between a pair of switches to be combined to form a single, logical ISL switch with a speed of up to 64 Gbps (128 Gbps full duplex) for optimal bandwidth utilization and load balancing.
- Dynamic Path Selection (DPS) which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.

Port side view of Brocade 300

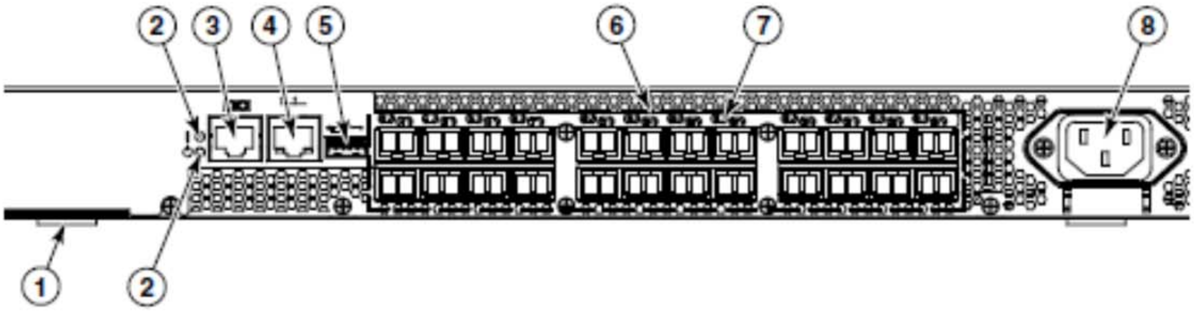

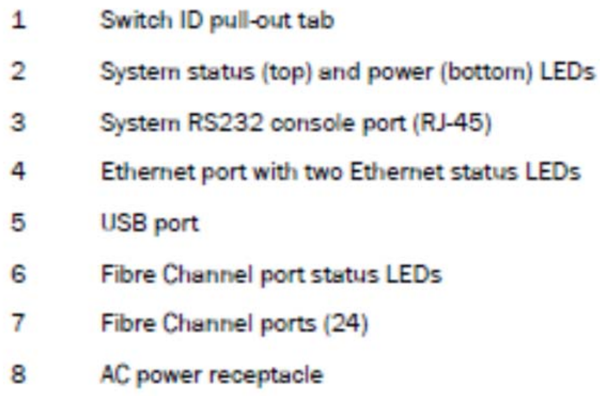

1	Switch ID pull-out tab
2	System status (top) and power (bottom) LEDs
3	System RS232 console port (RJ-45)
4	Ethernet port with two Ethernet status LEDs
5	USB port
6	Fibre Channel port status LEDs
7	Fibre Channel ports (24)
8	AC power receptacle

Figure 5: Port-side view of the Brocade 300

Overview of Brocade 5300

The Brocade 5300 is an 80-port 1, 2, 4, or 8 Gbps Fibre Channel switch that delivers Brocade sixth generation ASIC technology and architecture for Fibre Channel Storage Area Networks (SANs). The Brocade 5300 is designed for the needs of enterprise environments that require a high-port footprint for port aggregation. With its high auto-sensing port count and ports-on-demand flexibility, the Brocade 5300 is an ideal solution as a fan-out switch from a director core, or as the core switch in a fabric. The Brocade 5300 satisfies demanding Reliability, Availability, and Serviceability (RAS), performance and scalability requirements of an enterprise switch while delivering interoperability and ease-of-use advantages found only in the Brocade product family. The Brocade 5300 is the latest enterprise offering from the Brocade family of entry-to-enterprise products, and offers the following features and capabilities:

- Up to 80 auto-sensing ports of high-performance 8 Gbps technology in a single domain.
- Ports On Demand scaling from 48 to 64 or 80 ports.
- Full 1:1 subscription on all 80 ports at 8 Gbps.
- 1, 2, 4 and 8 Gbps auto-sensing Fibre Channel switch and router ports.
- FICON and FICON Control Unit Port ready.
- Fibre Channel Routing (FCR) service, which provides improved scalability and fault isolation, along with multi-vendor interoperability through the optional Integrated Routing license.
- Two hot-swappable, redundant power supply FRUs.
- Three hot-swappable fan FRUs in an N+1 configuration to provide hardware-redundant cooling.
- Universal ports that self-configure as E, F, M or FL ports. Ex_Ports are activated on a per port basis with the optional Integrated Routing license.
- An RJ45 Ethernet management port, in conjunction with EZSwitchSetup that supports switch IP address discovery and configuration, eliminating the need to attach a serial cable to configure the switch IP address and greatly increasing the ease of use.
- USB port that provides storage for firmware updates, output of the supportSave command and storage for configuration uploads and downloads.
- Inter-Switch Link (ISL) Trunking (licensable), which allows up to eight ports (at 1, 2, 4, or 8 Gbps speeds) between a pair of switches combined to form a single, logical ISL with a speed of up to 64 Gbps (128 Gbps full duplex) for optimal bandwidth utilization and load balancing.
- Dynamic Path Selection (DPS), which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- Industry-leading extended distance support, which enables native Fibre Channel extension greater than 590 km (366 miles).

- Brocade Small Form-Factor Pluggable (SFP) or SFP+ optical transceivers that support any combination of Short Wavelength (SWL), Long Wavelength (LWL), or Extended Long Wavelength (ELWL) optical media among the switch ports.
- Unicast, multicast (255 groups), and broadcast data traffic type support.
- Brocade Fabric Operating System (FOS), which delivers distributed intelligence throughout the network and enables a wide range of value-added applications including Brocade Advanced
- Web Tools and Brocade Zoning. Optional Fabric Services include: Adaptive Networking with QoS, Brocade Extended Fabrics, Brocade Enhanced Group Management, Brocade Fabric Watch, ISL Trunking, Integrated Routing, and End-to-End Performance Monitoring (APM).
- Port-to-port latency minimized to 2100 nanoseconds through the use of cut-through frame routing at 8 Gbps.

Port side view of Brocade 5300

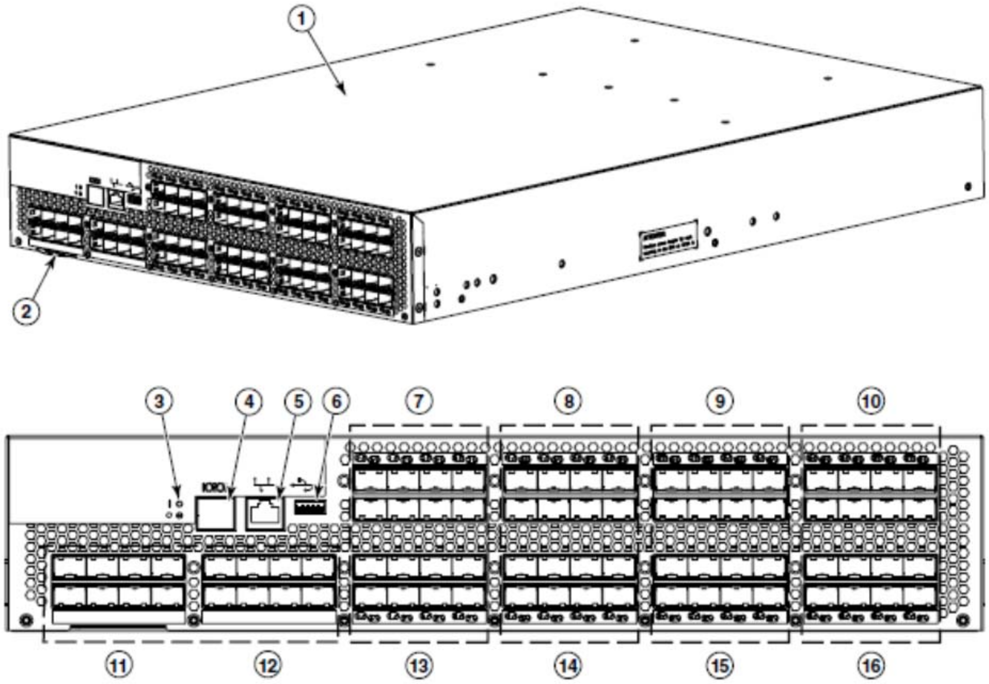

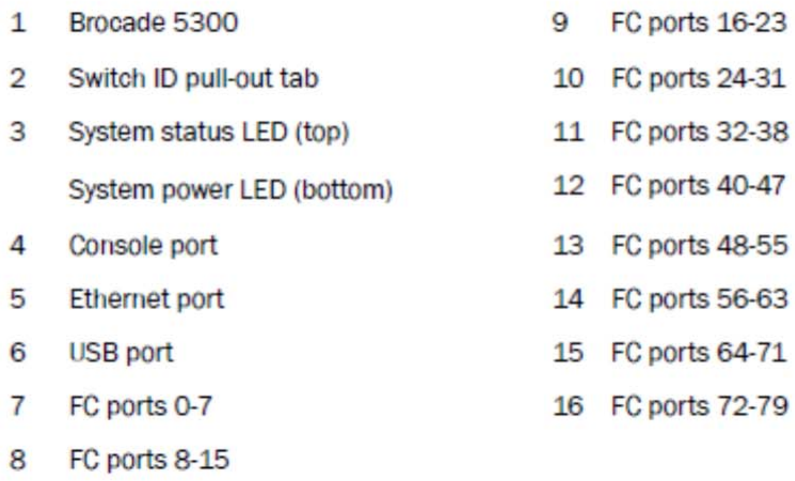

1	Brocade 5300	9	FC ports 16-23
2	Switch ID pull-out tab	10	FC ports 24-31
3	System status LED (top)	11	FC ports 32-38
	System power LED (bottom)	12	FC ports 40-47
4	Console port	13	FC ports 48-55
5	Ethernet port	14	FC ports 56-63
6	USB port	15	FC ports 64-71
7	FC ports 0-7	16	FC ports 72-79
8	FC ports 8-15		

Figure 6: Port-side view of the Brocade 5300

Access Gateway support

Brocade 5100, Brocade 300 and Brocade 6510 have Access gateway support. Brocade 6520 and 5300 does not support access gateway mode.

1. Before disabling a 6510 switch to enable Access Gateway mode, save the current configuration file using the configupload command in case you might need this configuration again.
2. At the terminal application prompt, type **SwitchDisable**, and press Enter to disable switch mode.
3. If you are converting an Brocade 300 or Brocade 5100 currently configured as a switch to Access Gateway mode, type **configUpload Save** and press Enter to save the current configuration.
4. To enable Access Gateway mode, type **ag - -modeEnable**, and press Enter. The switch automatically reboots and comes back online in Access Gateway mode.
5. Enter the **ag - -modeShow** command to ensure that the switch is in Access Gateway mode.
6. Enter **ag - -mapshow** to display the F_port to N_port mapping. The F_ports connect to servers, and the N_ports connect to Fabrics.

NOTE:

- After you enable AG mode, some fabric information is erased, such as the zone and security databases.
- Enabling AG mode is disruptive because the switch is disabled and rebooted.
- Ensure that no zoning or Admin Domain (AD) transaction buffers are active. If any transaction buffer is active, enabling Access Gateway mode will fail with the error, “Failed to clear Zoning/Admin Domain configuration.”

Brocade 5100

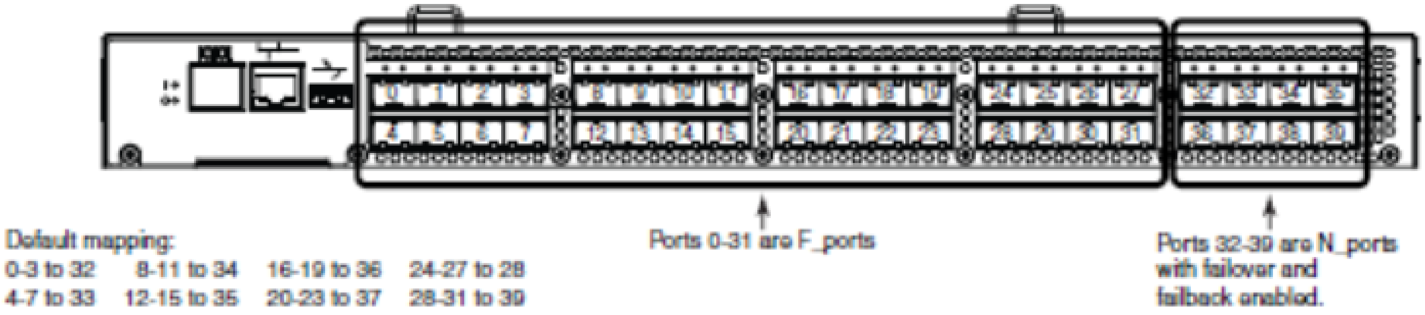

Brocade 300

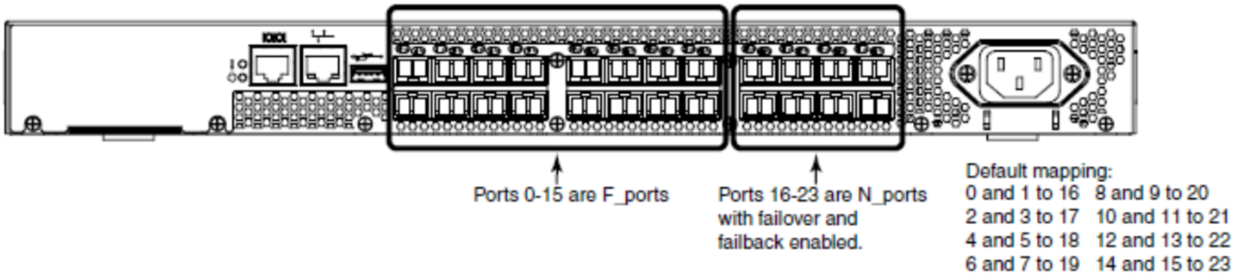

Figure 7: Port-side view of the Brocade 5300

Access Gateway (AG) enhancements with FOS 7.0.1

For AG a number of enhancements have been made to F_Port Static Mapping , Brocade Network Advisor and a new feature Advanced Performance Monitor (APM) has been added.

Advanced Performance Monitor support

Advanced Performance Monitoring (APM) is a licensed feature that allows you to monitor traffic on a specific port It supports End to End monitors on F-ports for Frame Monitors on F and N ports and requires an APM license.

For more on these enhancements with FOS 7.0.1 please refer to Fabric OS Adminstrator’s Guide.

Chapter 2: Basic Configuration

Assigning IP address

We need to create a console connection to the switch to assign an IP address to the Ethernet interface. This IP address can be used later to access the remotely (Telnet, SSH, FTP etc) or to perform management activities.

Creating serial connection

Connect the serial cable to the RJ-45 serial port (shown in Figure 1 as number 2) on the switch and to an RS-232 serial port on the workstation. If the serial port on the workstation is RJ-45 instead of RS-232, remove the adapter on the end of the serial cable and insert the exposed RJ-45 connector into the RJ-45 serial port on the workstation.

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

- In a Windows environment:

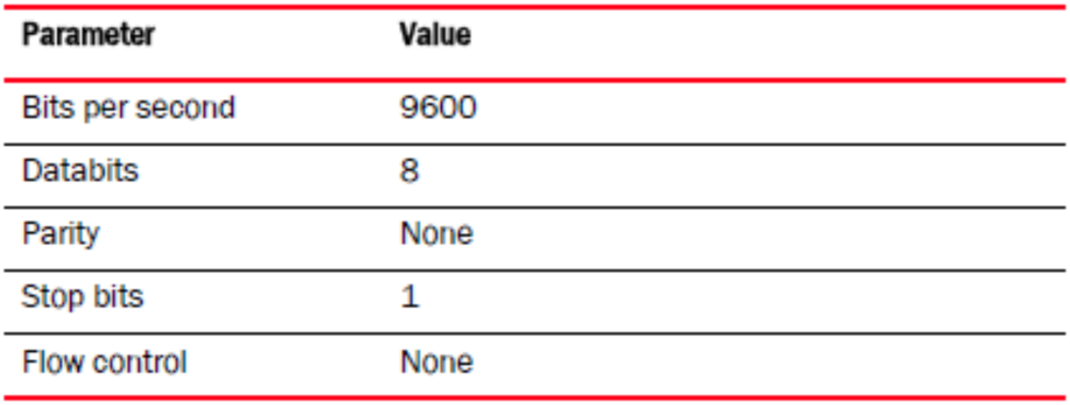

Parameter	Value
Bits per second	9600
Databits	8
Parity	None
Stop bits	1
Flow control	None

- In a UNIX environment, enter the following string at the prompt:
  tip /dev/ttyb -9600

  If ttyb is already in use, use ttya instead and enter the following string at the prompt:
  tip /dev/ttya -9600

Using DHCP (Dynamic Host Configuration Protocol) to assign IP address

On Brocade 300, 5100,5300, 6505, 6510 and 6520 DHCP is enabled by default. If you have the DHCP server on the same IP subnet as the switch then you can use DHCP to assign IP address to your switch. If not, you have to assign static IP address.

Enabling DHCP

Connect the DHCP-enabled switch to the network, power on the switch, and the switch automatically obtains the Ethernet IP address, Ethernet subnet mask, 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. Enabling DHCP after the Ethernet information has been configured releases the current Ethernet network interface settings, including Ethernet IP address, Ethernet subnet mask, and gateway IP address. The Fibre Channel IP address and subnet mask are static and are not affected by DHCP; for instructions on setting the FC IP address, see "Static Ethernet addresses" below.

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

3. If already set up, skip the Ethernet IP address, Ethernet subnet mask, Fibre Channel IP address, and Fibre Channel subnet mask prompts by pressing Enter.
4. Enable DHCP by entering on.

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

Disabling DHCP

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

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the ipAddrSet command.
3. Enter the network information in dotted-decimal notation for the Ethernet IPv4 address or in semicolon-separated notation for IPv6.
   If a static Ethernet address is not available when you disable DHCP, enter 0.0.0.0 at the Ethernet IP address prompt.
4. Skip the Fibre Channel prompts by pressing Enter.
5. When you are prompted for DHCP[On], disable it by entering off.

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

Static IP address

Once the console connection is up you can configure a static IP address as follows:

1. Login to the switch using the default password, which is *password*.
2. Use the **ipaddrset** command to assign an IP address

```
switch:admin> ipaddrset
Ethernet IP Address [192.168.10.10]: 192.168.10.12
Ethernet Subnetmask [255.255.255.0]: 255.255.255.0
Gateway IP Address [192.168.10.1]: 192.168.10.1
DHCP [Off]:off
switch:admin>
```

The values inside the square braces are the default or previously assigned values. Just press enter at the prompts if wish to keep the values in the braces.

If you are going to use an IPv6 address, enter the network information in semicolon-separated notation as prompted.

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

IP address is being changed...Done.

IPv6 Auto configuration

Here are the steps to enable/disable IPv6 auto configuration.

1. Enter the **ipAddrSet -ipv6 -auto** command to enable IPv6 auto configuration for all managed entities on the target platform.
2. Enter the **ipAddrSet -ipv6 -noauto** command to disable IPv6 auto configuration for all managed entities on the target platform.

View IP configuration

To view the IP configuration of the switch use **ipaddrShow** command

```
switch:admin> ipaddrshow
SWITCH
Ethernet IP Address: 192.168.10.12
Ethernet Subnetmask: 255.255.255.0
Gateway IP Address: 192.168.10.1
DHCP: Off
switch:admin>
```

Domain IDs

Domain IDs are set dynamically on Brocade switches. The default value is 1. You can change the domain ID if you want to control the ID number or resolve conflict while merging fabrics. Conflicts can be automatically resolved if one of the two switch's domain ID is not set persistently.

ATTENTION
Do not use domain ID 0. The use of this domain ID can cause the switch to reboot continuously.
Avoid changing the domain ID on the FCS switch in secure mode. To minimize down time, change the domain IDs on the other switches in the fabric.

Below are the steps to view and set the Domain Ids.

Viewing your Domain ID

1. Log in using account with admin privileges.
2. Issue the **fabricShow** command

```
DS_5100:admin> fabricshow
Switch ID   Worldwide Name          Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 10:00:00:05:1e:02:0e:de 10.246.54.240   0.0.0.0         "DS_200B"
  2: fffc02 10:00:00:05:1e:02:93:75 10.246.54.241   0.0.0.0         "DS_5100"
  4: fffc04 10:00:00:05:1e:44:b6:00 10.246.54.79    10.10.10.10    >"ED_DCX_B"

The Fabric has 3 switches
```

The switch with the arrow (>) next to its name is the principal switch. Below is the description of the output.

- **Switch ID:** The switch’s domain_ID and embedded port D_ID. The numbers are broken down as follows:
  Example **64: fffc40**
  64 is the switch domain_ID
  fffc40 is the hexidecimal format of the embedded port D_ID.

- **Worldwide Name:** The switch’s WWN.

- **Enet IP Addr:** The switch’s Ethernet IP address for IPv4- and IPv6-configured switches. For IPv6 switches, only the static IP address displays.

- **FC IP Addr:** The switch’s Fibre Channel IP address.

- **Name:** The switch’s symbolic or user-created name in quotes. An arrow (>) indicates the principal switch.

Setting your Domain ID

Here are the steps to configure Domain ID manually

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

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

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

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

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

Ports

By default, all licensed ports are enabled. You can disable and re-enable them as necessary. Ports that you activate with the “Ports on Demand” license must be enabled explicitly, as described in “Ports on Demand”. If ports are persistently disabled and you use the **portEnable** command to enable a disabled port, the port will revert to being disabled after a power cycle or a switch reboot. To ensure the port remains enabled, use the **portCfgPersistentEnable** command as instructed below.

Enable a port

1. Log in with account that has admin privileges
2. Issue the **portEnable** *portnumber* command.

   ```
   switch:admin> portenable 10
   ```

3. Issue the **portCfgPersistentEnable** *portnumber* command to enable a port that has been persistently disabled.

   ```
   switch:admin> portcfgpersistentenable 10
   ```

Disable a port

1. Log in with account that has admin privileges
2. Issue the **portDisable** *portnumber* command.

```
switch:admin> portdisable 10
```

3. Issue the **portCfgPersistentDisable** *portnumber* command to persistently disable a port.

```
switch:admin> portcfgpersistentdisable 10
```

Default Port Name

FOS v7.0.1 assigns a port name to all FC ports on a switch by default

Class of Switch	Default PortName Format	Example
Switch	port<port_no>	port5
Director	slot<slot_no> port<port_no>	slot1 port5

NOTES:

Port Name Behavior:

  - Pre FOS7.0: portname is NULL if not configured
  - FOS7.0/+: default portname assigned if not configured

- The port name assignments do not reside in the configuration
- Configured port name always overrides default port name
- Issuing portcfgdefault causes any configured portname to reset to defaults
- DCFM, CLI and any interface that queries for portname will be provided with the default port name if not configured

For example:

```
B6510:FID128:admin> portshow 0

portIndex:   0

portName: port0

portHealth: HEALTHY

<truncated output>
```

Ports on Demand license

To enable additional ports, you must install Ports On Demand (POD) licenses. To install a POD license, you can either use the supplied license key or generate a license key. Typically the switch is shipped with a paper pack that specifies the transaction key to use with the Software License Keys link. Use this transaction key on my.brocade.com Web site and follow the instructions to generate the key.

1. Login to my.brocade.com
2. Scroll down to Licensing tools and click on 'Enter the Software Portal'

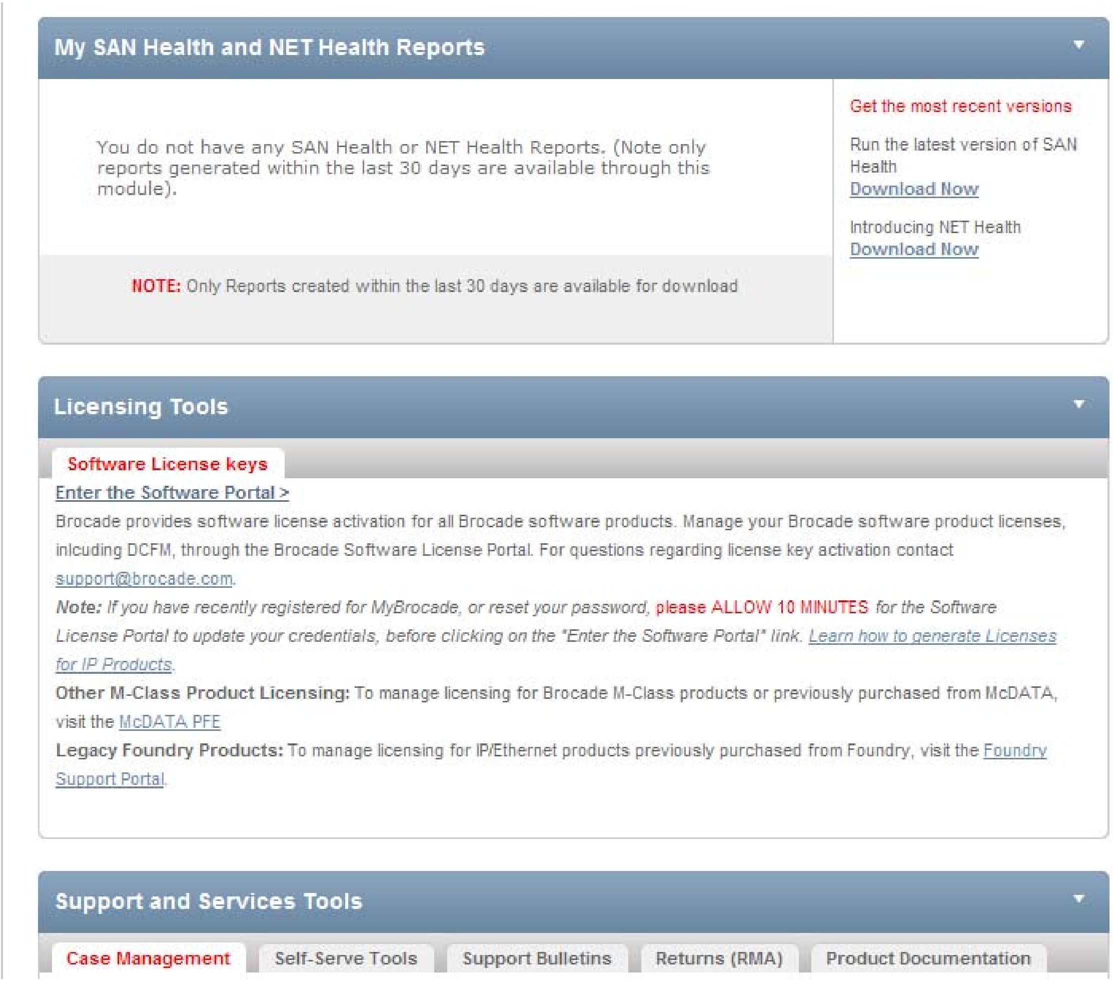

My SAN Health and NET Health Reports

You do not have any SAN Health or NET Health Reports. (Note only reports generated within the last 30 days are available through this module).

NOTE: Only Reports created within the last 30 days are available for download

Get the most recent versions

Run the latest version of SAN Health
Download Now

Introducing NET Health
Download Now

Licensing Tools

Software License keys

Enter the Software Portal >

Brocade provides software license activation for all Brocade software products. Manage your Brocade software product licenses, inlcuding DCFM, through the Brocade Software License Portal. For questions regarding license key activation contact support@brocade.com.

Note: *If you have recently registered for MyBrocade, or reset your password,* please ALLOW 10 MINUTES *for the Software License Portal to update your credentials, before clicking on the "Enter the Software Portal" link. Learn how to generate Licenses for IP Products.*

Other M-Class Product Licensing: To manage licensing for Brocade M-Class products or previously purchased from McDATA, visit the McDATA PFE

Legacy Foundry Products: To manage licensing for IP/Ethernet products previously purchased from Foundry, visit the Foundry Support Portal.

Support and Services Tools

You can also use this site to generate other license keys for your switch. After you have installed the license keys using the **licenseAdd** command, you must enable the ports. You can do so without disrupting switch operation by using the **portEnable** command on each port individually. Alternatively, you can disable and re-enable the switch to activate all ports simultaneously.

```
switch:admin> licenseadd DXXtN3LmRSMWCSW3XmfSBPfrWKLZ3HMTN73rP9GANJMA
adding license-key [DXXtN3LmRSMWCSW3XmfSBPfrWKLZ3HMTN73rP9GANJMA]
```

Use the **licenseShow** command to view all the licenses.

Setting Port Speed

1. Log in with account that has admin privileges
2. Issue the **portCfgSpeed** *portnumber* <speed> command.

The following example sets the speed for port 3 to 8 Gbps:

```
ecp:admin> portcfgspeed 3 8
done.
```

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

```
ecp:admin> portcfgspeed 3 0
done.
```

3. Issue the **switchCfgSpeed** <speed> command to set all ports to same speed setting.

The following example sets the speed for all ports on the switch to 8 Gbps:

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

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

```
switch:admin> switchcfgspeed 0
Committing configuration...done.
```

Following things can be entered for speed settings:

```
Speed_Level:  0  - Auto Negotiate (Hardware)
              1  - 1Gbps
              2  - 2Gbps
              4  - 4Gbps
              8  - 8Gbps
              ax - Auto Negotiate (Hardware) + retries
              s  - Auto Negotiate (Software)
```

Setting Port name

To set a name for the port we use the **portName** command. The name of the port is shown in **portShow** output. It is not to be confused with the port World-Wide Name (pWWN).

1. Log in with account that has admin privileges
2. Use command **portname** *portnumber* **-n** *"desired name"*

```
switch:admin> portname 1 -n "To DCX"

switch:admin> portshow 1
portIndex:   1
```

```
portName: To DCX
portHealth: No Fabric Watch License
(output truncated)
```

Swapping port area IDs

If a device that uses port binding is connected to a port that fails, you can use port swapping to make another physical port use the same PID as the failed port. The device can then be plugged into the new port without the need to reboot the device.
Use the following procedure to swap the port area IDs of two physical switch ports. In order to swap port area IDs, the port swap feature must be enabled, and both switch ports must be disabled. The swapped area IDs for the two ports remain persistent across reboots, power cycles, and failovers.

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

```
ecp:admin>portdisable 1
```

4. Enter the **portSwap** command.

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

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

Customizing the switch name

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

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

Record the new switch name for future reference.

Checking Switch Status

1. Log in with account that has admin privileges
2. Use the **switchShow** command to check the status of the all ports
3. Use **switchStatusShow** command to check the status of switch

Chassis names

Brocade recommends that you customize the chassis name for each platform. Some system logs identify devices by platform names; if you assign meaningful platform names, logs are more useful. All chassis names have a limit of 15 characters, except for the Brocade 300, 5100, 5300, and VA-40FC switches, and the 5410, 5424, 5450, and 5480 embedded switches, which allow 31 characters. Chassis names must begin with a letter, and can contain letters, numbers, or the underscore character.

Customizing chassis names

1. Connect to the switch and log in as admin.
2. Enter the chassisName command.

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

3. Record the new chassis name for future reference.

Fabric Name

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

- Each name must be unique for each logical switch within a chassis; duplicate fabric names are not allowed.
- A fabric name can be from 1 through 128 alphanumeric characters.
- All switches in a logical fabric must be running Fabric OS v7.0.0. Switches running earlier versions of the firmware can co-exist in the fabric, but do not show the fabric name details.
- You must have admin permissions to configure the fabric name.

Configuring the fabric name

To set and display the fabric name, use the command fabricname as shown in the following example:

```
switch:user> fabricname –set myfabric@1
```

Using the fabricname –set command without a fabric name takes the existing fabric name and synchronizes it across the entire fabric. An error message displays if no name is configured.
To set a fabric name that includes spaces, use the command fabricname as shown in the following example:

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

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

```
switch:user> fabricname –set 'red fabric $$'
```

To clear the fabric name, use the fabricname –clear command.

Upgrade and downgrade considerations

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

Config file upload and download considerations

A new key, “fabric name” is added to store the user configuration. You can only configure fabric names using config download when the switch is offline

Fabric Name Conflicts

- In an existing fabric, a fabric name change will **not** result in segmenting ISLs, rather the latest name will be propagated to other switches in the same fabric
- When fabrics are merged, conflicting fabric names are not merged, the fabric name that was configured before the merge is retained on each respective switch
- In case of a conflict, the user must choose the correct name and set the fabric name manually
  - Re-running the command from any switch in the fabric will propagate it fabric wide
- Two fabrics are merged with different fabric names configured

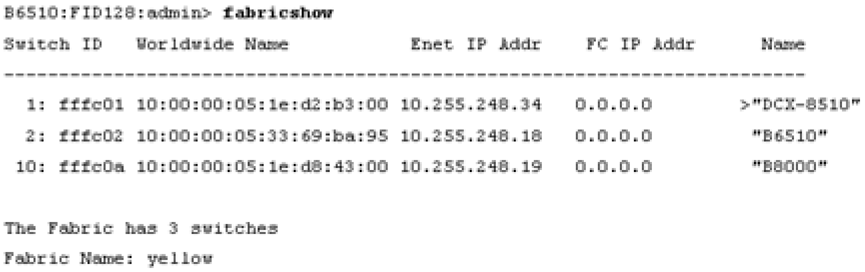

```
B6510:FID128:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 10:00:00:05:1e:d2:b3:00 10.255.248.34   0.0.0.0        >"DCX-8510"
  2: fffc02 10:00:00:05:33:69:ba:95 10.255.248.18   0.0.0.0         "B6510"
 10: fffc0a 10:00:00:05:1e:d8:43:00 10.255.248.19   0.0.0.0         "B8000"

The Fabric has 3 switches
Fabric Name: yellow
```

Switch activation and deactivation

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

Disabling a switch

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

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

Enabling a switch

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

All Fibre Channel ports that passed POST are enabled. If the switch has interswitch links (ISLs) to a fabric, it joins the fabric.

Switch shutdown

To avoid corrupting your file system, Brocade recommends that you perform graceful shutdowns of Brocade enterprise-class platforms.

Warm reboot refers to shutting down the appliance per the instructions below, also known as a *graceful shutdown*. *Cold boot* refers to shutting down the appliance by suddenly shutting down power and then turning it back on, also known as a *hard boot*.

Powering off a Brocade switch

The following procedure describes how to gracefully shut down a switch.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **sysShutdown** command.
3. At the prompt, enter y.

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

4. Wait until the following message displays:

```
Broadcast message from root (ttyS0) Wed Jan 25 16:12:09 2006...
The system is going down for system halt NOW !!
INIT: Switching to runlevel: 0
INIT: Sending processes the TERM signal
Unmounting all filesystems.
The system is halted
flushing ide devices: hda
Power down.
```

5. Power off the switch.

Setting the date and time

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

   date "mmddHHMMyy"
   The values represent the following:
   - mm is the month; valid values are 01 through 12.
   - dd is the date; valid values are 01 through 31.
   - HH is the hour; valid values are 00 through 23.
   - MM is minutes; valid values are 00 through 59.
   - yy is the year, valid values are 00-37 and 70-99 (year values from 70-99 are interpreted as 1970-1999, year values from 00-37 are interpreted as 2000-2037).

Example of showing and setting the date

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

Synchronizing the local time with an external source

The tsClockServer command accepts multiple server addresses in IPv4, IPv6, or DNS name formats. When multiple NTP server addresses are passed, tsClockServer sets the first obtainable address as the active NTP server. The rest are stored as backup servers that can take over if the active NTP server fails. The principal or primary FCS switch synchronizes its time with the NTP server every 64 seconds.

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

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

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

Example of setting the NTP server

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

Example of displaying the NTP server

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

Example of setting up more than one NTP server using a DNS name

```
switch:admin> tsclockserver "10.1.2.4;10.1.2.5;ntp.localdomain.net"
Updating Clock Server configuration...done.
Updated with the NTP servers
Changes to the clock server value on the principal or primary FCS switch are
propagated to all switches in the fabric.
```

VC Level Credit Loss Detection and Automatic Recovery on 16G Platforms

Buffer credit loss can lead to performance degradation, hence early detection and recovery is essential. In FOS 7.0.1 Buffer credit loss is detected and recovered automatically at Virtual Channel level on Condor3 FC ISLs (both sides of the link must be Condor 3).

For more on this enhancement see Fabric OS Administrator's Guide.

Webtools

Configuring IP and netmask information

1. Click the Network tab.

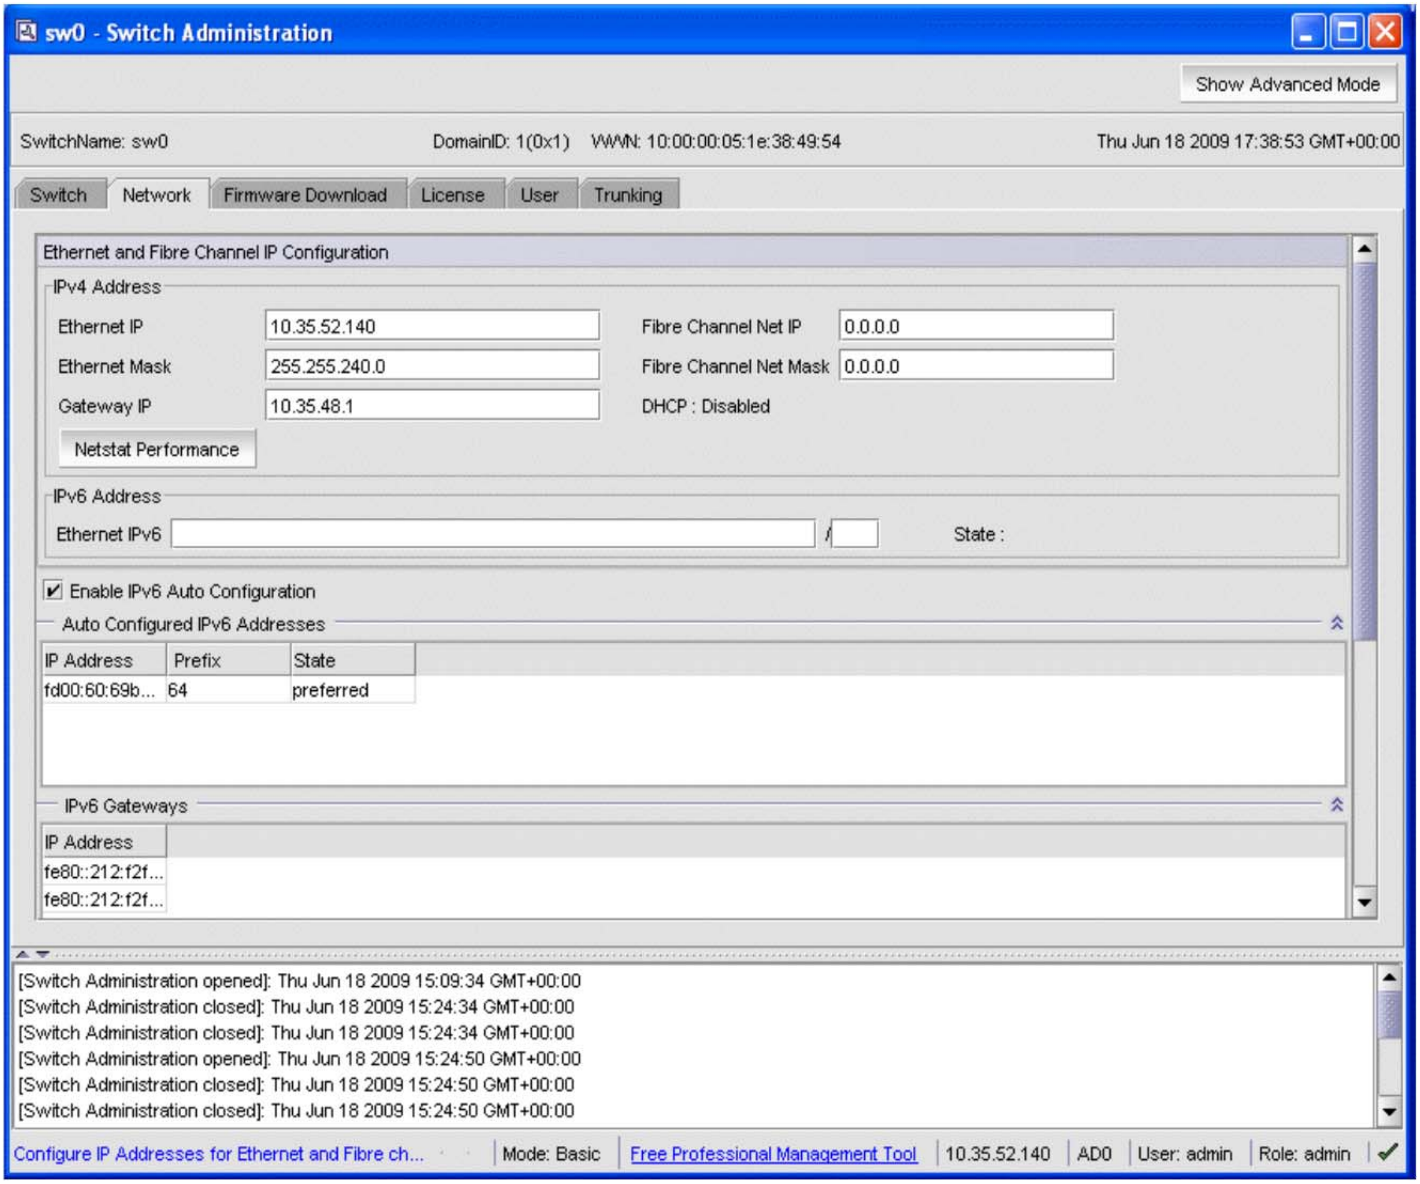

SwitchName: sw0 DomainID: 1(0x1) WWN: 10:00:00:05:1e:38:49:54 Thu Jun 18 2009 17:38:53 GMT+00:00

Ethernet and Fibre Channel IP Configuration

IPv4 Address

Ethernet IP 10.35.52.140

Ethernet Mask 255.255.240.0

Gateway IP 10.35.48.1

Fibre Channel Net IP 0.0.0.0

Fibre Channel Net Mask 0.0.0.0

DHCP : Disabled

IPv6 Address

Ethernet IPv6 / State :

☑ Enable IPv6 Auto Configuration

Auto Configured IPv6 Addresses

IP Address	Prefix	State
fd00:60:69b...	64	preferred

IPv6 Gateways

IP Address
fe80::212:f2f...
fe80::212:f2f...

```
[Switch Administration opened]: Thu Jun 18 2009 15:09:34 GMT+00:00
[Switch Administration closed]: Thu Jun 18 2009 15:24:34 GMT+00:00
[Switch Administration closed]: Thu Jun 18 2009 15:24:34 GMT+00:00
[Switch Administration opened]: Thu Jun 18 2009 15:24:50 GMT+00:00
[Switch Administration closed]: Thu Jun 18 2009 15:24:50 GMT+00:00
[Switch Administration closed]: Thu Jun 18 2009 15:24:50 GMT+00:00
```

Configure IP Addresses for Ethernet and Fibre ch... | Mode: Basic | Free Professional Management Tool | 10.35.52.140 | AD0 | User: admin | Role: admin | ✓

2. In the appropriate IP address section, enter the IP address you want to use for the IP interface.
3. Use the IPv4 Address section or the IPv6 Address section to specify IP addresses.
4. In the IPv4 Address section:
   a. In the Ethernet IP field, enter the Ethernet IP address.
   b. In the Ethernet Mask field, enter the Ethernet Mask address.
   c. In the GateWay IP address field, enter the gateway IP address.
5. In the IPv6 Address section, in the Ethernet IPv6 field, enter the Ethernet IP address.
6. You can also enable automatic configuration of IPv6 addresses by selecting Enable IPV6 Auto Configuration. The automatically generated IPv6 addresses are displayed under Auto Configured IPV6 Addresses. Eight auto-configured addresses are created per switch, and up to
   24 for a 48000, DCX, or DCX-4S chassis (eight per chassis, and eight per each installed CP).

Switch configuration

Use the Switch tab of the Switch Administration window to perform basic switch configuration.

Enabling and disabling a switch

You can identify whether a switch is enabled or disabled in the Switch Administration window by looking at the lower-right corner. If you rest the cursor over the icon, the system displays text that indicates the status of the switch. The steps are as given below.

1. Open the Switch Administration window.
2. Click the Switch tab.
3. In the Switch Status section, click Enable to enable the switch or Disable to disable the switch.
4. Click Apply.

The system displays a confirmation window that asks if you want to save the changes to the switch. You must click Yes to save the changes.

Changing the switch name

Switches can be identified by IP address, domain ID, World Wide Name (WWN), or switch names. Names must begin with an alphabetic character, but otherwise can consist of alphanumeric, hyphen, and underscore characters. The maximum number of characters is 30, unless FICON mode is enabled. When FICON mode is enabled, the maximum number of characters is 24.

NOTE
Some system messages identify a switch service by the chassis name. If you assign meaningful chassis names and switch names, system logs are easier to use.

1. Open the Switch Administration window.
2. Click the Switch tab.
3. Enter a new name in the Name field and click Apply.

Changing the switch domain ID

Although domain IDs are assigned dynamically when a switch is enabled, you can request a specific ID to resolve a domain ID conflict when you merge fabrics. Follow the steps below.

1. Open the Switch Administration window.
2. Disable the switch.
3. Click the Switch tab.
4. Enter a new domain ID in the Domain ID field. The domain ID range depends on the switch interop mode:
   - For IM0, the range is between 1 and 239
   - For IM2, it depends on the selected offset value
   - For IM3, it depends on the selected offset value

5. Click Apply.
6. Enable the switch.

Viewing and printing a switch report

The switch report includes the following information:

- A list of switches in the fabric
- Switch configuration parameters
- A list of ISLs and ports
- Name Server information
- Zoning information
- SFP serial ID information

Perform the following steps to view or print a report.

1. Open the Switch Administration window.
2. Click the Switch tab.

3. Click View Report.
4. In the new window that displays the report, view or print the report using your browser.

Switch restart

When you restart the switch, the restart takes effect immediately. Ensure that there is no traffic or other management on the switch, because traffic is interrupted during the restart; however, frames are not dropped. Be sure to save your changes before the restart, because any changes not saved are lost.

Performing a fast boot

A fast boot reduces boot time significantly by bypassing the power-on self test (POST).

1. Open the Switch Administration window.
2. Click Fastboot.
3. On the Fastboot Confirmation window, click Yes to continue.
4. Click Apply.

Performing a reboot

Use the following procedure to reboot the CP and execute the normal power-on booting sequence.

1. Open the Switch Administration window.
2. Click Reboot.
3. On the Reboot Confirmation window, click Yes to continue.
4. Click Apply.

Configuring fabric settings

Perform the following steps to configure the fabric settings.

1. Open the Switch Administration window.
2. Disable the switch.
3. Click the Configure tab.
4. Click the Fabric subtab.

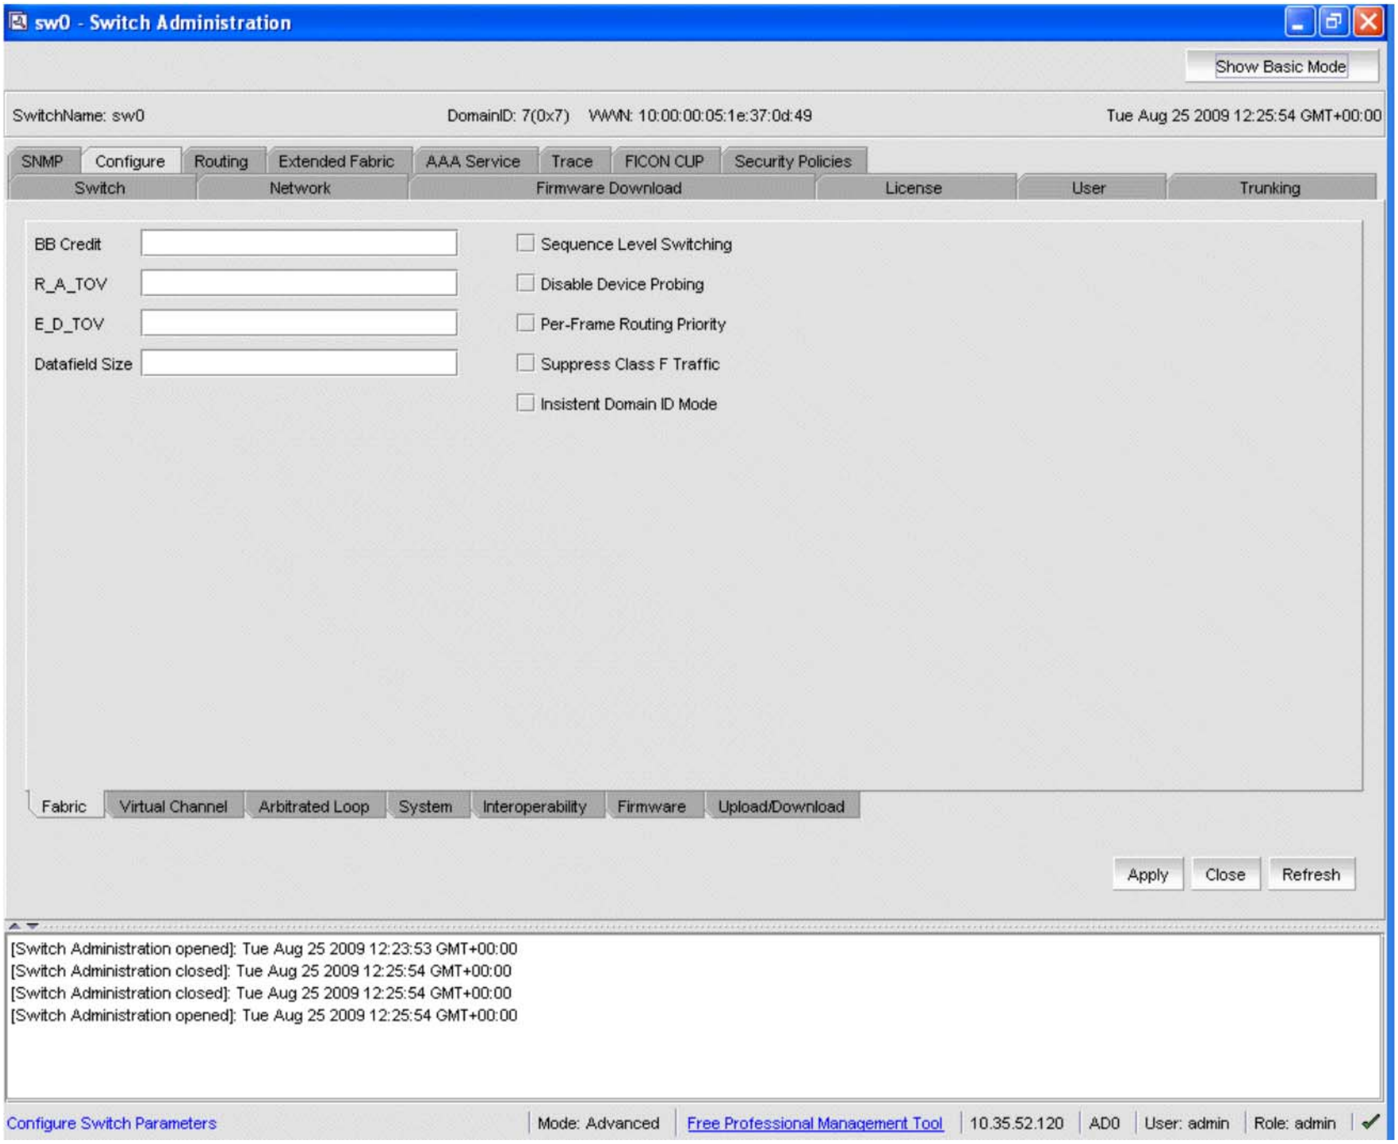

SwitchName: sw0 DomainID: 7(0x7) WWN: 10:00:00:05:1e:37:0d:49 Tue Aug 25 2009 12:25:54 GMT+00:00

BB Credit

R_A_TOV

E_D_TOV

Datafield Size

□ Sequence Level Switching

□ Disable Device Probing

□ Per-Frame Routing Priority

□ Suppress Class F Traffic

□ Insistent Domain ID Mode

```
[Switch Administration opened]: Tue Aug 25 2009 12:23:53 GMT+00:00
[Switch Administration closed]: Tue Aug 25 2009 12:25:54 GMT+00:00
[Switch Administration closed]: Tue Aug 25 2009 12:25:54 GMT+00:00
[Switch Administration opened]: Tue Aug 25 2009 12:25:54 GMT+00:00
```

Mode: Advanced | 10.35.52.120 | AD0 | User: admin | Role: admin ✓

5. Make the fabric parameter configuration changes.
6. Click Apply.
7. Enable the switch.

Assigning a name to a port

Port names are optional. You can assign a name to an FC or FCIP port to make port grouping easier.

You can rename FC and FCIP ports too. You cannot rename GbE ports. The Port Name column in the Ports tab displays the port name, if one exists.

Port names can be from 1 through 32 alphanumeric characters, unless Ficon Management Server (FMS) mode is enabled; if FMS mode is enabled, port names should be limited from 1 through 24 alphanumeric characters. The comma (,), semicolon (;), and “at” symbol (@) are not allowed.

NOTE

Although it is not required, it is recommended that port names be unique.

1. Click a port in the Switch View to open the Port Administration window.

2. Click the FC Ports tab.
3. From the tree on the left, click the switch that contains the port you want to rename.
4. From the table, select the port you want to rename

5. Click Rename.
6. Type a name for the port and click Rename.

To delete the existing port name, leave the field blank and click Rename.

Enabling and disabling a port

Use the following procedure to enable or disable a port.

1. Click a port in the Switch View to open the Port Administration window.
2. Click the FC Ports or GigE Ports tab.
3. From the tree on the left, click the switch or slot that contains the port you want to enable or disable.
4. From the table, select one or more ports.

Use Shift+click and Ctrl+click to select multiple ports. You can select multiple ports from the table. You cannot select multiple ports from the tree.

5. Click Enable or Disable.

If the button is gray (unavailable), the port is already in the enabled or disabled state. For example, if the Enable button is unavailable, the port is already enabled.

If you select multiple ports in both enabled and disabled states, both buttons are active. When you click either button, the action is applied to all selected ports. Click Yes in the confirmation window.

Persistent enabling and disabling ports

Use the following procedure to enable or disable an FC port so that it remains enabled or disabled across switch restarts.

NOTE
Ports cannot be persistently enabled or disabled when FMS is enabled.

1. Click a port in the Switch View to open the Port Administration window.
2. Click the FC Ports or GigE Ports tab.
3. From the tree on the left, click the switch or slot that contains the port.
4. From the table, select one or more ports.
   Use Shift-click and Ctrl-click to select multiple ports. You can select multiple ports from the table. You cannot select multiple ports from the tree.
5. Click Persistent Enable or Persistent Disable.

NOTE
Persistent Enable or Disable is not supported in FMS mode.

If the button is gray (unavailable), the port is already in that state or FMS mode is enabled on the switch. For example, if the Persistent Enable button is unavailable, the port or ports are already persistently enabled over restarts.
If you select multiple ports in both enabled and disabled states, both buttons are active. When you click either button, the action is applied to all selected ports.

6. Click Yes in the confirmation window.

Enabling and disabling NPIV ports

The NPIV license must be installed on a switch before NPIV functionality can be enabled on any port. For detailed information about understanding and configuring NPIV ports, refer to the *Fabric OS Administrator's Guide*. With Web Tools, you can only enable or disable the NPIV functionality on a port. Perform the following procedure to enable or disable NPIV ports.

NOTE
NPIV feature cannot be disabled when Access gateway mode is enabled.

1. Click a port in the Switch View to open the Port Administration window.
2. Click the FC Ports tab.
3. From the tree on the left, select the logical port you want to enable or disable.
4. Click Enable NPIV or Disable NPIV.

Configuring BB credits on an F_Port

From 6.3.0 you can configure the BB credits value on an F_Port. Follow the steps given below.

1. Click a port in the Switch View to open the Port Administration window.
2. Click the FC Ports tab.
3. Click Show Advanced Mode.

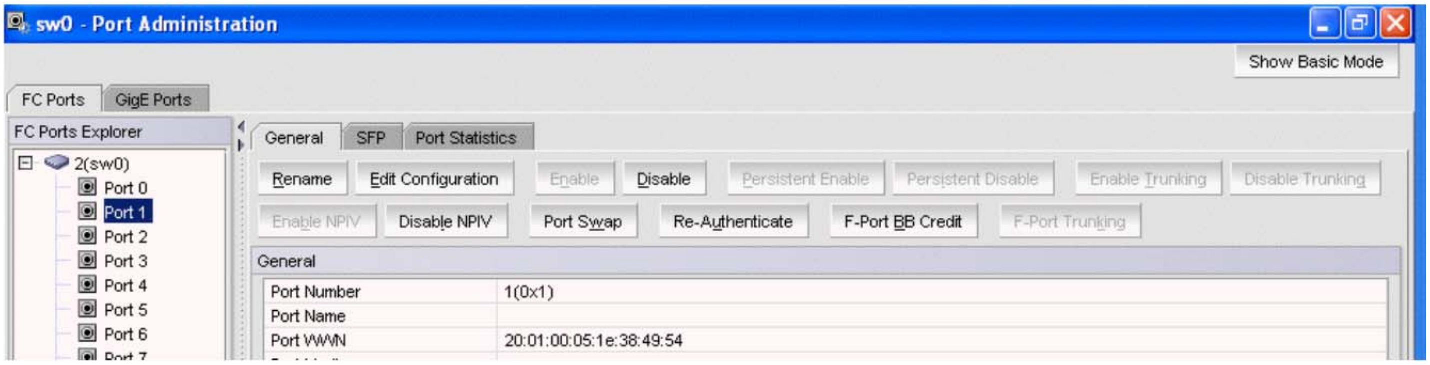

Port Number	1(0x1)
Port Name	
Port WWN	20:01:00:05:1e:38:49:54

4. Click F-Port BB Credit.

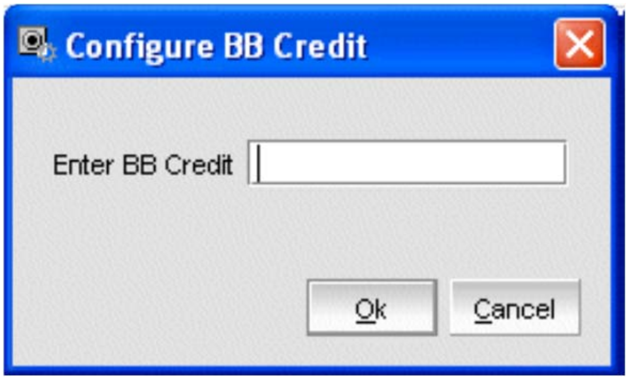

5. Enter the BB credit value in the Enter BB Credit field. The default value is 8.

NOTE
You cannot modify the default BB credit value for VE and ICL ports.

6. Click Ok.
 The value is displayed in the table of the Port Administration window. If no value is configured the F-Port BB Credit column displays the default value.

Chapter 3: Account Management

Overview

In addition to the default accounts—root, factory, admin, and user—Fabric OS supports up to 252 additional user-defined accounts in each logical switch (domain). These accounts expand your ability to track account access and audit administrative activities.

Fabric OS provides three options for authenticating users—remote RADIUS services, remote LDAP service, and the local switch user database. All options allow users to be centrally managed using the following methods:

- **Remote RADIUS server**: Users are managed in a remote RADIUS server. All switches in the fabric can be configured to authenticate against the centralized remote database.
- **Remote LDAP server**: Users are managed in a remote LDAP server. All switches in the fabric can be configured to authenticate against the centralized remote database.
- **Local user database**: Users are managed using the local user database.

User Roles

- **Admin**: All administrative commands excluding chassis-specific commands.
- **BasicSwitchAdmin**: Mostly monitoring with limited switch (local) commands.
- **FabricAdmin**: All switch and fabric commands, excludes user management and Admin Domains commands.
- **Operator**: Routine switch maintenance commands.
- **SecurityAdmin**: All switch security and user management functions.
- **SwitchAdmin**: Most switch (local) commands, excludes security, user management, and zoning commands.
- **User**: Nonadministrative use, such as monitoring system activity.
- **ZoneAdmin**: Zone management commands only.

Local database user accounts

Creating account:

1. Login to switch using an account with administrator privileges
2. Use the **userConfig - - add** command.

```
switch:admin> userconfig –add Larry -r admin
Setting initial password for Larry
Enter new password:
Re-type new password:
Account Larry has been successfully added.
```

The usage for command userConfig - - add is as follows:

userconfig ***- -****add username -r role [-h AD_ID] [-a AD_ID_list] [-d description] [-p password] [-x]*

Displaying Account Information:

1. Login to switch using an account with administrator privileges
2. Use **userconfig –show** [<username> | -a | -r <role name>]:

```
switch:admin> userconfig --show -a
Account name: admin
Description: Administrator
Enabled: Yes
Password Last Change Date: Mon Aug 23 2010 (UTC)
Password Expiration Date: Not Applicable (UTC)
Locked: No
Role: admin
AD membership: 0-255
Home AD: 0

Account name: Larry
Description:
Enabled: Yes
Password Last Change Date: Tue Sep  7 2010 (UTC)
Password Expiration Date: Not Applicable (UTC)
Locked: No
Role: admin
AD membership: 0
Home AD: 0
```

Deleting account:

1. Login to switch using an account with administrator privileges
2. Use the **userConfig - - delete** command.

```
switch:admin> userconfig --delete Larry
About to delete account Larry
ARE YOU SURE  (yes, y, no, n): [no] Y

Broadcast message from root (pts/0) Tue Sep  7 18:45:50 2010...

Security Policy, Password or Account Attribute Change: Larry will be logged out
Account Larry has been successfully deleted.
switch:admin>
```

Modifying account:

1. Login using account with administrator privelages.
2. Use **userconfig –change** <username> [-r <rolename (admin | user | switchadmin | zoneadmin | fabricadmin | basicswitchadmin | operator | securityadmin [[-h <AD_ID>] [-a <AD_ID list>]][-d <description>] [-e yes | no] [-x] [-u]:

In the example below, we look at account Larry which has admin privileges. We change his privileges to securityadmin.

```
DS_4900B:admin> userconfig --show -a
Account name: Larry
Description:
Enabled: Yes
Password Last Change Date: Tue Sep  7 2010 (UTC)
Password Expiration Date: Not Applicable (UTC)
Locked: No
Role: admin
AD membership: 0
Home AD: 0


DS_4900B:admin> userconfig --change Larry -r securityadmin

Broadcast message from root (pts/0) Tue Sep  7 19:30:16 2010...

Security Policy, Password or Account Attribute Change: Larry will be logged out
DS_4900B:admin>
```

Changing password for current login account

User can change password for his account as follows:

1. Log into your account
2. Enter **passwd** command and follow the prompts

```
DS_4900B:admin> passwd
Changing password for admin
Enter old password:
Enter new password:
Re-type new password:
passwd: all authentication tokens updated successfully
Saving password to stable storage.
Password saved to stable storage successfully.
DS_4900B:admin>
```

Changing password for different login account

1. Log into the switch as user with admin privelages
2. Enter **passwd** <username> command and follow the prompts

```
DS_4900B:admin> passwd Larry
Changing password for Larry
Enter old password:
Enter new password:
```

```
Re-type new password:
passwd: all authentication tokens updated successfully
Saving password to stable storage.
Password saved to stable storage successfully.
```

Local account database distribution

Fabric OS allows you to distribute the user database and passwords to other switches in the fabric. When the switch accepts a distributed user database, it replaces the local user database with the user database it receives.
By default, switches accept the user databases and passwords distributed from other switches.
The 'Locked' status of a user account is not distributed as part of local user database distribution.
When distributing the user database, the database may be rejected by a switch for one of the following reasons:

- One of the target switches does not support local account database distribution.
- One of the target switch's user database is protected.
- One of the remote switches has logical switches defined.

Distributing the local user database

When distributing the local user database, all user-defined accounts residing in the receiving switches are logged out of any active sessions.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **distribute -p PWD -d** command.

Accepting distribution of user databases on the local switch

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **fddCfg –localaccept PWD** command.

Rejecting distributed user databases on the local switch

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **fddCfg –localreject PWD** command.

Password policies

You can use the **passwdCfg –set** command to modify following parameters

- Password strength
- Password history
- Password expiration
- Account lockout

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

Example of a password strength policy
The following example shows a password strength policy that requires passwords to contain at least 3 uppercase characters, 4 lowercase characters and 2 numeric digits; the minimum length of the password is 9 characters.

```
passwdcfg –set -uppercase 3 -lowercase 4 -digits 2 -minlength 9
```

To display the current password configuration parameters:

```
switch:admin> passwdcfg –showall
passwdcfg.minlength: 8
passwdcfg.lowercase: 0
passwdcfg.uppercase: 0
passwdcfg.digits: 0
passwdcfg.punctuation: 0
passwdcfg.history: 1
passwdcfg.minpasswordage: 0
passwdcfg.maxpasswordage: 0
passwdcfg.warning: 0
passwdcfg.lockoutthreshold: 0
passwdcfg.lockoutduration: 30
passwdcfg.status: 0
```

Enabling the admin lockout policy

1. Log in to the switch using an account that is an Admin role or securityAdmin role.
2. Enter the **passwdCfg –enableadminlockout** command.

Unlocking an account

1. Log in to the switch using an account that is an Admin role or securityAdmin role.
2. Enter the **userConfig –change** ***account_name*** **-u** command specifying the name of the user account that is locked out.

Disabling the admin lockout policy

1. Log in to the switch using an account that is an Admin role or securityAdmin role.
2. Enter the **passwdCfg –disableadminlockout** command.

Authentication servers on the switch

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

Adding a RADIUS or LDAP server to the switch configuration

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

At least one RADIUS or LDAP server must be configured before you can enable the RADIUS or LDAP service. If no RADIUS or LDAP configuration exists, turning on the RADIUS authentication mode triggers an error message. When the command succeeds, the event log indicates that the configuration is enabled or disabled.

Enabling and disabling a RADIUS or LDAP server

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **aaaConfig –authspec** command to enable RADIUS or LDAP using the local database.

You must specify the type of server as either RADIUS or LDAP, but not both. Local is used for local authentication if the user authentication fails on the RADIUS or LDAP server.

Example of enabling RADIUS

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

Deleting a RADIUS or LDAP server from the configuration

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

When the command succeeds, the event log indicates that the server is removed.

Changing a RADIUS or LDAP server configuration

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

Changing the order in which RADIUS or LDAP servers are contacted for service

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the aaaConfig –move command.

When the command succeeds, the event log indicates that a server configuration is changed.

Displaying the current RADIUS configuration

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

If a configuration exists, its parameters are displayed. If RADIUS or LDAP service is not configured, only the parameter heading line is displayed. Parameters include:

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

Configuring local authentication as backup

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

Example of enabling local authentication, enter the following command for RADIUS

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

Example for LDAP

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

When local authentication is enabled and the RADIUS or LDAP servers fail to respond, you can login to the default switch accounts (admin and user) or any user-defined account. You must know the passwords of these accounts. When the command succeeds, the event log indicates that local database authentication is disabled or enabled.

Chapter 4: Setting the Protocols

Secure Copy

The secure copy protocol (SCP) runs on port 22. It encrypts data during transfer, thereby avoiding packet sniffers that attempt to extract useful information during data transfer. SCP relies on SSH to provide authentication and security

Setting up SCP for configUploads and downloads

1. Log in to the switch as admin.
2. Type the **configure** command.
3. Type y or yes at the *cfgload attributes* prompt.
4. Type y or yes at the *Enforce secure configUpload/Download* prompt.

```
switch:admin> configure
Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
Configure...
System services (yes, y, no, n): [no] n
ssl attributes (yes, y, no, n): [no] n
http attributes (yes, y, no, n): [no] n
snmp attributes (yes, y, no, n): [no] n
rpcd attributes (yes, y, no, n): [no] n
cfgload attributes (yes, y, no, n): [no] y
Enforce secure config Upload/Download (yes, y, no, n): [no] y
Enforce signature validation for firmware (yes, y, no, n): [no]
```

Secure Shell protocol

To ensure security, Fabric OS supports secure shell (SSH) encrypted sessions. SSH encrypts all messages, including the client transmission of the password during login. The SSH package contains a daemon (sshd), which runs on the switch. The daemon supports a wide variety of encryption algorithms, such as Blowfish-Cipher block chaining (CBC) and Advanced Encryption Standard (AES).

Commands that require a secure login channel must originate from an SSH session. If you start an SSH session, and then use the login command to start a nested SSH session, commands that require a secure channel will be rejected

Allowed-user

The default admin user must set up the allowed-user with the admin role. By default, the admin is the configured allowed-user. While creating the key pair, the configured allowed-user can choose a passphrase with which the private key is encrypted. Then the passphrase must always be entered when authenticating to the switch. The allowed-user must have an admin role that can perform OpenSSH public key authentication, import and export keys, generate a key pair for an outgoing connection, and delete public and private keys. After the allowed-user is changed, all the public keys related to the old allowed-user are lost.

Configuring SSH authentication

Incoming authentication is used when the remote host needs to authenticate to the switch. Outgoing authentication is used when the switch needs to authenticate to a server or remote host, more commonly used for the configUpload command. Both password and public key authentication can coexist on the switch.
After the allowed-user is configured, the remaining setup steps must be completed by the allowed-user.

1. Log in to the switch as the default admin.
2. Change the allowed-user's role to admin, if applicable.

```
switch:admin> userconfig –change username -r admin
```

Where *username* is the name of the user you want to perform SSH public key authentication, import, export, and delete keys.

3. Set up the allowed-user by typing the following command:

```
switch:admin> sshutil allowuser username
```

Where *username* is the name of the user you want to perform SSH public key authentication, import, export, and delete keys.

4. Generate a key pair for host-to-switch (incoming) authentication by logging in to your host as admin, verifying that SSH v2 is installed and working (refer to your host's documentation as necessary) by typing the following command:

```
switch:admin> ssh-keygen -t dsa
```

If you need to generate a key pair for outgoing authentication, skip steps 4 and 5 and proceed to step 6.

Example of RSA/DSA key pair generation

```
alloweduser@mymachine: ssh-keygen -t dsa
Generating public/private dsa key pair.
Enter file in which to save the key (/users/alloweduser/.ssh/id_dsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /users/alloweduser/.ssh/id_dsa.
Your public key has been saved in /users/alloweduser/.ssh/id_dsa.pub.
The key fingerprint is:
32:9f:ae:b6:7f:7e:56:e4:b5:7a:21:f0:95:42:5c:d1 alloweduser@mymachine
```

5. Import the public key to the switch by logging in to the switch as the allowed-user and entering the **sshUtil importpubkey** command to import the key.

Example of adding the public key to the switch

```
switch:alloweduser> sshutil importpubkey
Enter IP address:192.168.38.244
Enter remote directory:~auser/.ssh
Enter public key name(must have .pub suffix):id_dsa.pub
Enter login name:auser
Password:
Public key is imported successfully.
```

6. Generate a key pair for switch-to-host (outgoing) authentication by logging in to the switch as the allowed user and entering the **sshUtil genkey** command.
   You may enter a passphrase for additional security.

Example of generating a key pair on the switch

```
switch:alloweduser> sshutil genkey
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
```

```
Key pair generated successfully.
```

7. Export the public key to the host by logging in to the switch as the allowed-user and entering the **sshUtil exportpubkey** command to export the key.

Example of exporting a public key from the switch

```
switch:kghanta> sshutil exportpubkey
Enter IP address:192.168.38.244
Enter remote directory:~auser/.ssh
Enter login name:auser
Password:
public key out_going.pub is exported successfully.
```

8. Append the public key to a remote host by logging in to the remote host, locating the directory where authorized keys are stored, and appending the public key to the file.
   You may need to refer to the host's documentation to locate where the authorized keys are stored.
9. Test the setup by using a command that uses SCP and authentication, such as **firmwareDownload** or **configUpload**.

Deleting keys on the switch

1. Log in to the switch as the allowed-user.
2. Use the **sshUtil delprivkey** command to delete the private key.
   or
   Use the **sshUtil delpubkeys** command to delete all public keys.

Telnet protocol

Telnet is enabled by default. To prevent passing clear text passwords over the network when connecting to the switch, you can block the Telnet protocol using an IP Filter policy.

ATTENTION
Before blocking Telnet, make sure you have an alternate method of establishing a connection with the switch.

Blocking Telnet

If you create a new policy using commands with just one rule, all the missing rules have an implicit deny and you lose all IP access to the switch, including Telnet, SSH, and management ports.

1. Connect to the switch and log in as admin.
2. Clone the default policy by typing the **ipFilter –clone** command.

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

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

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

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

```
switch:admin> ipfilter –show
```

5. Add a rule to the policy, by typing the **ipFilter –addrule** command.

```
switch:admin> ipfilter –addrule BlockTelnet -rule 1 -sip any -dp 23 –proto tcp -act deny
```

6. Save the new ipfilter policy by typing the **ipfilter –save** command.
7. Verify the new policy is correct by typing the **ipFilter –show** command.
8. Activate the new ipfilter policy by typing the **ipfilter –activate** command.

```
switch:admin> ipfilter –activate BlockTelnet
```

9. Verify the new policy is active (the default_ipv4 policy should be displayed as *defined*).

```
switch:admin> ipfilter –show
Name: BlockTelnet, Type: ipv4, State: defined
Rule Source IP Protocol Dest Port Action
1 any tcp 23 deny
2 any tcp 22 permit
3 any tcp 22 permit
4 any tcp 897 permit
5 any tcp 898 permit
6 any tcp 111 permit
7 any tcp 80 permit
8 any tcp 443 permit
9 any udp 161 permit
10 any udp 111 permit
11 any udp 123 permit
12 any tcp 600 - 1023 permit
13 any udp 600 - 1023 permit
Name: default_ipv4, Type: ipv4, State: defined
Rule Source IP Protocol Dest Port Action
1 any tcp 22 permit
2 any tcp 23 permit
3 any tcp 897 permit
4 any tcp 898 permit
5 any tcp 111 permit
6 any tcp 80 permit
7 any tcp 443 permit
8 any udp 161 permit
9 any udp 111 permit
10 any udp 123 permit
11 any tcp 600 - 1023 permit
12 any udp 600 - 1023 permit
```

Unblocking Telnet

1. Connect to the switch through a serial port or SSH and log in as admin.
2. Type in the **ipfilter –delete** command.
3. To permanently delete the policy, type the **ipfilter –save** command.

Chapter 5: Configuration file and Firmware management

Configuration file backup

In case the configuration is lost or unintentional changes are made, keep a backup copy of the configuration file. You should keep individual backup files for all switches in the fabric and avoid copying configurations from one switch to another. The **configUpload** command, by default, only uploads the switch context configuration for the logical switch context in which the command is executed.

In non-Virtual Fabric mode, you must use the **configUpload -all** command to include both the switch and the chassis information. In Virtual Fabric mode, the **configUpload -all** command can be selected to upload all logical switches and the chassis configuration. Only administrators with the chassis role permission are allowed to upload other FIDs or the chassis configuration.

The following information is *not saved* in a backup:

- dnsConfig information
- Passwords

Before beginning, verify that you can reach the FTP server from the switch. Using a Telnet connection, save a backup copy of the configuration file from a logical switch to a host computer.

Uploading a configuration file in interactive mode

1. Verify that the FTP or SCP service is running on the host computer.
2. Connect to the switch and log in as admin.
3. Enter the **configUpload** command. The command becomes interactive and you are prompted for the required information.
4. Store a soft copy of the switch configuration information in a safe place for future reference.

```
switch:admin> configupload
Protocol (scp, ftp, local) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: UserFoo
Path/Filename [<home dir>/config.txt]: switchConfig.txt
Section (all|chassis|FID# [all]): chassis
Password: <hidden>
configUpload complete
```

Configuration file restoration

1. Verify that the FTP service is running on the server where the backup configuration file is located.
2. Connect to the switch and log in using an account assigned to the admin role, and if necessary with the chassis-role permission.
3. If there are any changed parameters in the configuration file that do not belong to SNMP, Fabric Watch, or ACL, disable the switch by entering the **switchDisable** command.
4. Enter the **configDownload** command.
   The command becomes interactive and you are prompted for the required information.
5. At the “Do you want to continue [y/n]” prompt, enter y.
6. Wait for the configuration to be restored.
7. If you disabled the switch, enter the **switchEnable** command when the process is finished.

```
switch:admin> configdownload
Protocol (scp, ftp, local) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
```

```
User Name [user]: UserFoo
Path/Filename [<home dir>/config.txt]:
Section (all|chassis|FID# [all]): all
*** CAUTION ***
This command is used to download a backed-up configuration
for a specific switch. If using a file from a different
switch, this file's configuration settings will override
any current switch settings. Downloading a configuration
file, which was uploaded from a different type of switch,
may cause this switch to fail. A switch reboot might be
required for some parameter changes to take effect.
configDownload operation may take several minutes
to complete for large files.
Do you want to continue [y/n]: y
Password: <hidden>
configDownload complete.
```

NOTE:
Use configUpload and configDownload command with –vf option to manage config files for logical switches. You must perform the configDownload command on the switch after restoring the Virtual Fabric configuration to fully restore your switch or chassis configuration.

Installing firmware

Firmware download from network

1. Take the following appropriate action based on what service you are using:
   - If you are using FTP or SCP, verify that the FTP or SSH server is running on the host server and that you have a valid user ID and password on that server.
   - If your platform supports a USB memory device, verify that it is connected and running.
2. Obtain the firmware file from the Brocade Web site at *http://www.brocade.com* and store the file on the FTP or SSH server or the USB memory device.
3. Unpack the compressed files preserving directory structures.
   The firmware is in the form of RPM packages with names defined in a *.plist* file. The *.plist* file contains specific firmware information and the names of packages of the firmware to be downloaded.
4. Connect to the switch and log in as admin.
5. Issue the **firmwareShow** command to check the current firmware version on connected switches. Upgrade their firmware if necessary before proceeding with upgrading this switch.
6. Enter the **firmwareDownload** command and respond to the prompts.
7. At the “Do you want to continue [y/n]” prompt, enter y.
8. After the HA reboot, connect to the switch and log in again as admin.
9. If you want snapshots of the upgrade progress, use a separate session and enter the **firmwareDownloadStatus** command to monitor the firmware download.
10. After the firmware commit is completed, which takes several minutes, enter the **firmwareShow** command to display the firmware level of both partitions.

Example of an interactive firmware download

```
switch:admin> firmwareDownload
Server Name or IP Address: 10.1.2.3
User Name: userfoo
File Name: /userfoo/firmware/v6.4.0
Network Protocol(1-auto-select, 2-FTP, 3-SCP) [1]: 2
```

```
Password: <hidden>
Checking system settings for firmwareDownload...
Trying address-->AF_INET IP: 10.1.2.3, flags : 2
System settings check passed.
You can run firmwaredownloadstatus to get the status
of this command.
```

Firmware download from a USB device

The Brocade 300, 5100, 5300, 6505, 6510 and 6520 support a firmware download from Brocade branded USB device attached to the switch. Before the USB device can be accessed by the **firmwareDownload** command, it must be enabled and mounted as a file system. The firmware images to be downloaded must be stored under the relative path from ***/usb/usbstorage/brocade/firmware*** or use the absolute path in the USB file system. Multiple images can be stored under this directory.

There is a *firmwarekey* directory where the public key signed firmware is stored.

When the firmwareDownload command line option, **-U** (upper case), is specified, the firmwareDownload command downloads the specified firmware image from the USB device. When specifying a path to a firmware image in the USB device, you can only specify the relative path to */firmware* or the absolute path.

Enabling USB

1. Log in to the switch using an account assigned to the admin role.
2. Enter the **usbStorage -e** command.

Viewing the USB file system

1. Log in to the switch using an account assigned to the admin role.
2. Enter the **usbStorage -l** command.

```
Brcd:admin> usbstorage –l
firmware\ 381MB 2010 Mar 28 15:33
v6.4.0\ 381MB 2010 Mar 28 10:39
config\ 0B 2010 Mar 28 15:33
support\ 0B 2010 Mar 28 15:33
firmwarekey\ 0B 2010 Mar 28 15:33
Available space on usbstorage 79%
```

Downloading from USB using the relative path

1. Log in to the switch as admin.
2. Enter the **firmwareDownload -U** command.

```
ecp:admin>firmwaredownload –U v6.4.0
```

Downloading from USB using the absolute path

1. Log in to the switch as admin.
2. Enter the **firmwareDownload** command with the -U operand.

```
ecp:admin>firmwaredownload –U /usb/usbstorage/brocade/firmware/v6.4.0
```

Webtools

Creating a configuration backup file

Keep a backup copy of the configuration file in case the configuration is lost or unintentional changes are made. You should keep individual backup files for all switches in the fabric. You should avoid copying configurations from one switch to another.

3. Open the Switch Administration window.
4. Select Show Advanced Mode.
5. Select the Configure tab.

The Configure screen is displayed.

6. Select the Upload/Download tab.

The Upload/Download configuration screen is displayed. By default, Config Upload is chosen under Function, and Network is chosen as the source of the configuration file.

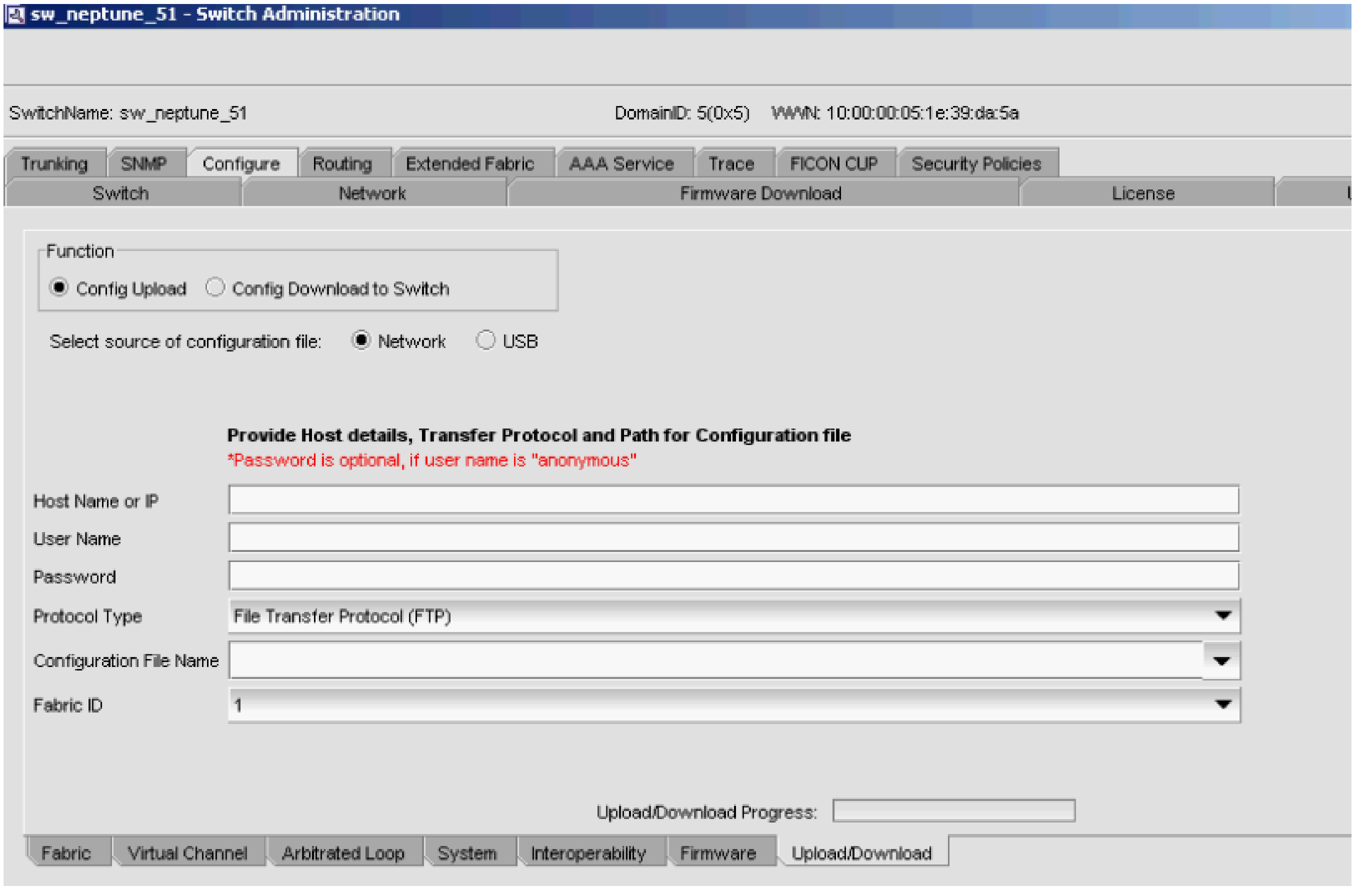

SwitchName: sw_neptune_51 DomainID: 5(0x5) WWN: 10:00:00:05:1e:39:da:5a

Function

● Config Upload ○ Config Download to Switch

Select source of configuration file: ● Network ○ USB

Provide Host details, Transfer Protocol and Path for Configuration file

*Password is optional, if user name is "anonymous"

Host Name or IP

User Name

Password

Protocol Type: File Transfer Protocol (FTP)

Configuration File Name

Fabric ID: 1

Upload/Download Progress:

7. If you upload from a network, type the host name or IP address in the Host Name or IP field, the user ID and password required for access to the host in the User Name and Password fields, and choose the Protocol Type used for the upload. The default is FTP. If you choose
   "Secure Copy Protocol (SCP)," you cannot specify "anonymous" in the User Name field.

If you choose USB as the configuration file source, the network parameters are not needed and are not displayed. You can skip to step 6.

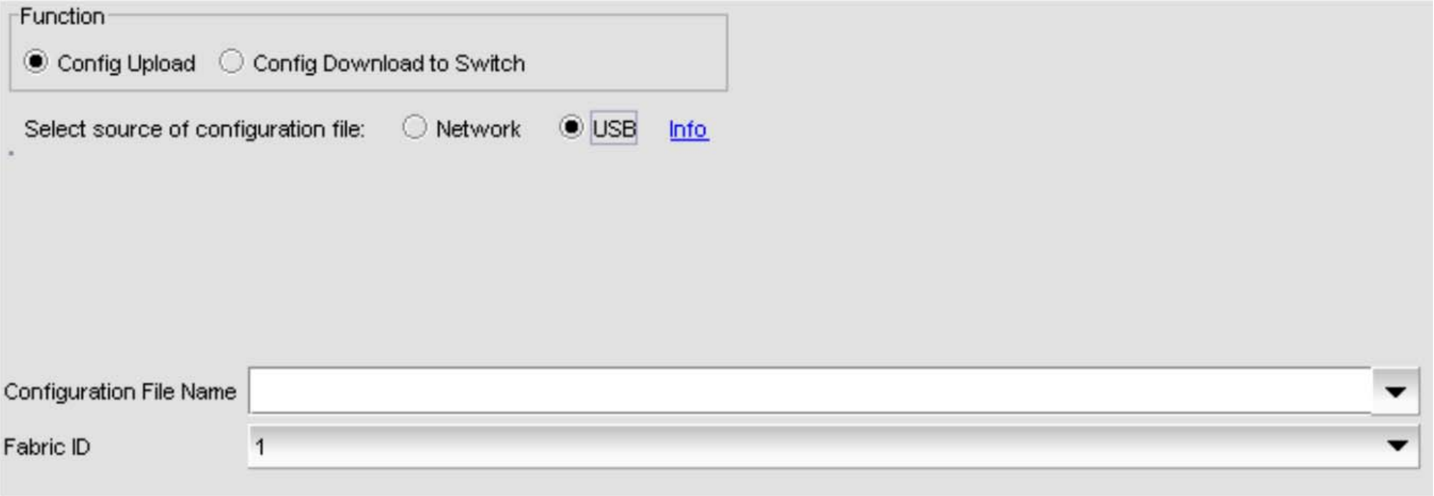

Function

● Config Upload ○ Config Download to Switch

Select source of configuration file: ○ Network ● USB Info

Configuration File Name

Fabric ID 1

An info link is enabled when USB is chosen as the source of the configuration file. If you click on info, the following information message is displayed

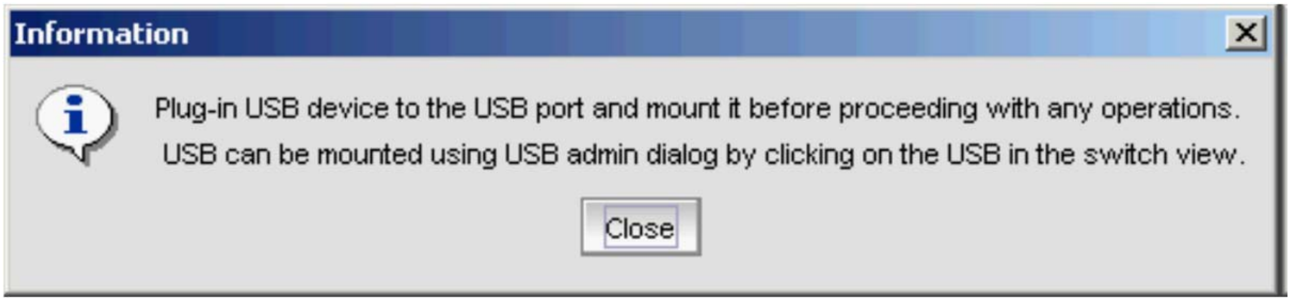

Information

Plug-in USB device to the USB port and mount it before proceeding with any operations. USB can be mounted using USB admin dialog by clicking on the USB in the switch view.

8. Type the configuration file with a fully-qualified path, or select the configuration file name in the Configuration File Name field.
9. Use the Fabric ID selector to select the fabric ID of the logical switch from which the configuration file is to uploaded. The selector will show all the virtual fabric IDs that have been defined, the default of 128 for the physical switch, chassis level configuration, and all chassis and switches.

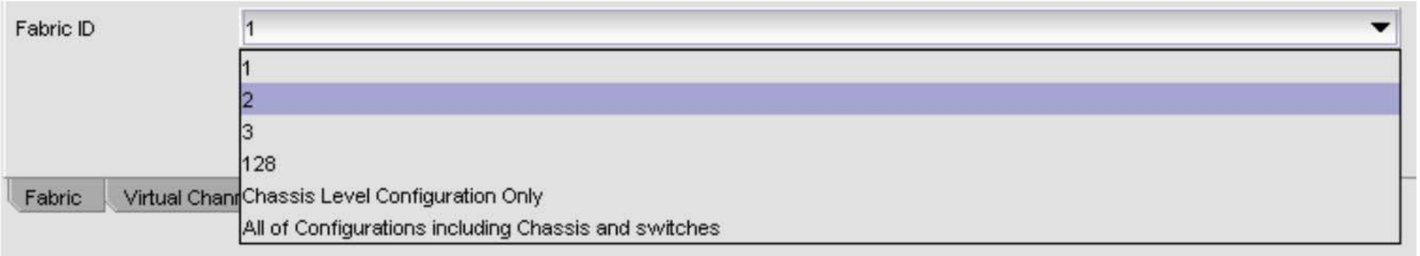

Fabric ID 1

1
2
3
128
Chassis Level Configuration Only
All of Configurations including Chassis and switches

NOTE
If you are using a USB device, it must be connected and mounted before you upload or download.

10. Click Apply.
    You can monitor the progress by watching the Upload/Download Progress bar.

Restoring a configuration

Restoring a configuration involves overwriting the configuration on the switch by downloading a previously saved backup configuration file. Perform this procedure during a planned down time.
Make sure that the configuration file you are downloading is compatible with your switch model.
Configuration files from other model switches might cause your switch to fail.

1. Open the Switch Administration window.
2. Select Show Advanced Mode.
3. Select the Configure tab.

The Configure screen is displayed.

4. Select the Upload/Download tab.

The Upload/Download configuration screen is displayed ().
By default, Config Upload is chose under Function, and Network is chosen as the source of the configuration file.

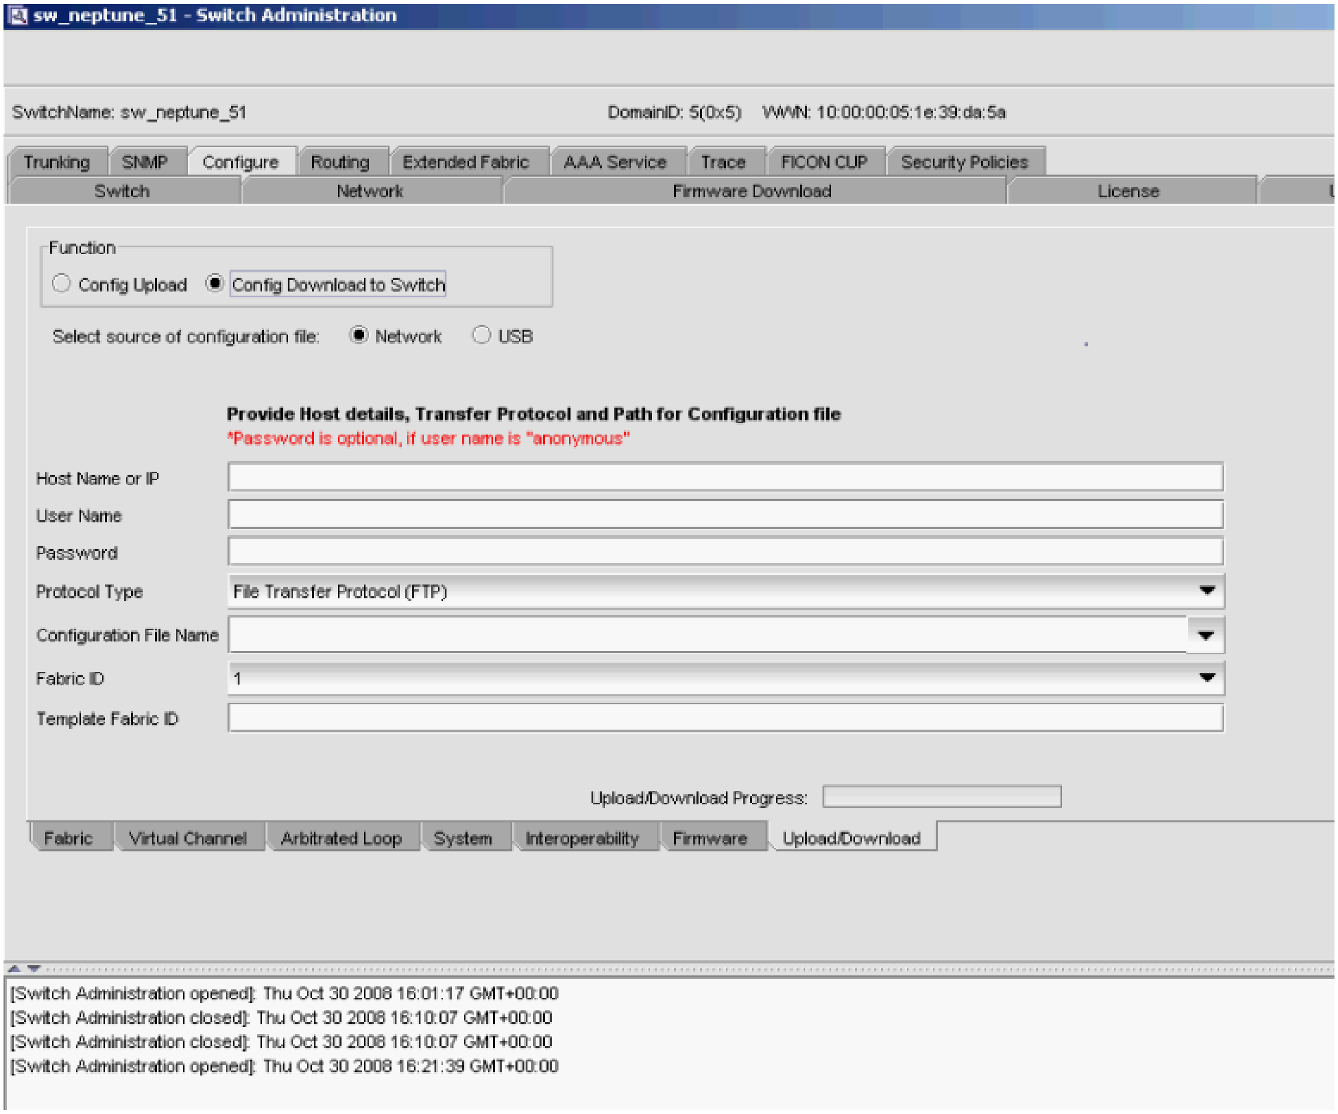

sw_neptune_51 - Switch Administration

SwitchName: sw_neptune_51 DomainID: 5(0x5) WWN: 10:00:00:05:1e:39:da:5a

Function: ○ Config Upload ● Config Download to Switch

Select source of configuration file: ● Network ○ USB

Provide Host details, Transfer Protocol and Path for Configuration file
*Password is optional, if user name is "anonymous"

Host Name or IP	
User Name	
Password	
Protocol Type	File Transfer Protocol (FTP)
Configuration File Name	
Fabric ID	1
Template Fabric ID	

Upload/Download Progress:

```
[Switch Administration opened]: Thu Oct 30 2008 16:01:17 GMT+00:00
[Switch Administration closed]: Thu Oct 30 2008 16:10:07 GMT+00:00
[Switch Administration closed]: Thu Oct 30 2008 16:10:07 GMT+00:00
[Switch Administration opened]: Thu Oct 30 2008 16:21:39 GMT+00:00
```

5. Under Function, select Config Download to Switch.

6. If you download from a network, type the host name or IP address in the Host Name or IP field, the user ID and password required for access to the host in the User Name and Password fields, and choose the Protocol Type used for the upload. The default is FTP. If you choose
   "Secure Copy Protocol (SCP)," you cannot specify "anonymous" in the User Name field.
7. If you choose USB as the configuration file source, the network parameters are not needed and are not displayed, and you can skip to step 6.

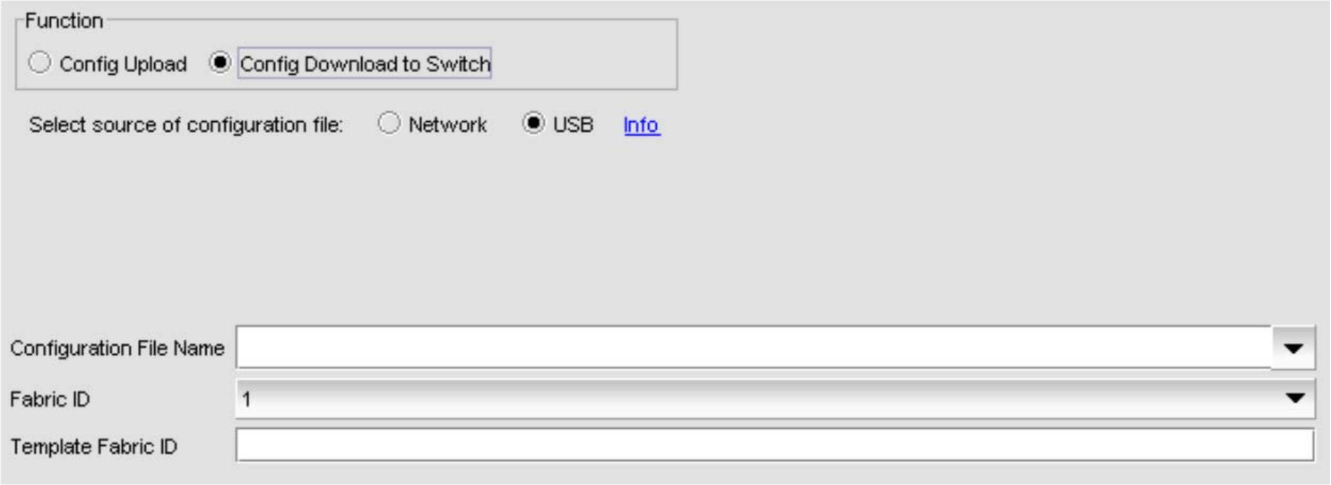

Function

○ Config Upload ● Config Download to Switch

Select source of configuration file: ○ Network ● USB Info

Configuration File Name

Fabric ID 1

Template Fabric ID

An info link is enabled when USB is chosen as the source of the configuration file. If you click info, the following information message is displayed.

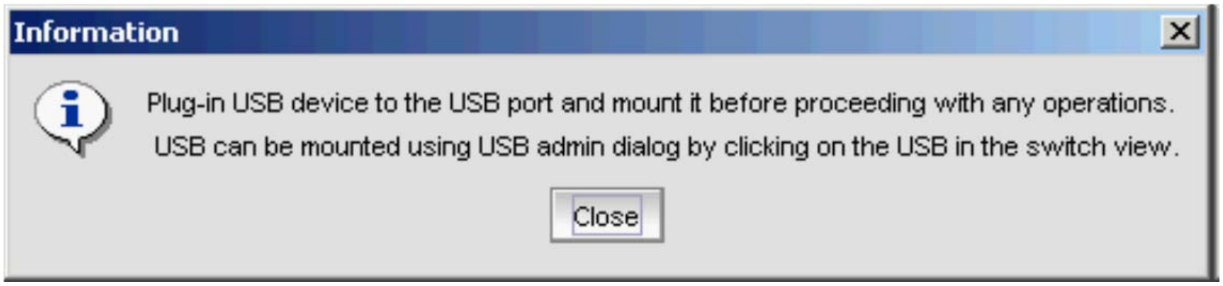

Information

Plug-in USB device to the USB port and mount it before proceeding with any operations. USB can be mounted using USB admin dialog by clicking on the USB in the switch view.

8. Type the configuration file with a fully-qualified path, or select the configuration file in the Configuration File Name field.
9. Use the Fabric ID selector to select the fabric ID of the logical switch to which the configuration file is to downloaded. The selector will show all the virtual fabric IDs that have been defined, the default of 128 for the physical switch, chassis level configuration, and all chassis and switches.

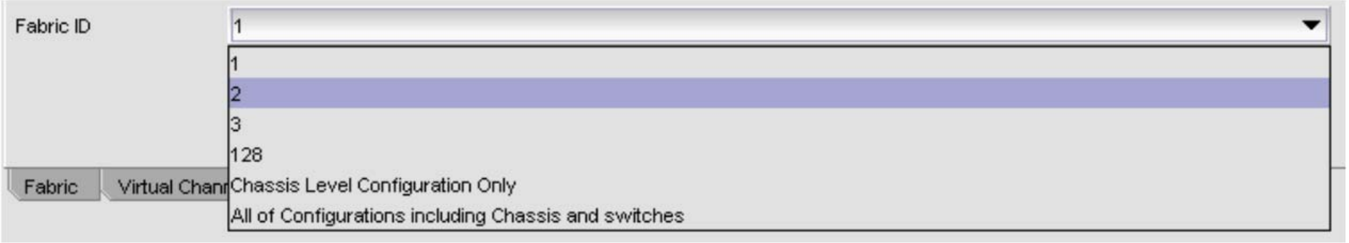

Fabric ID 1

1
2
3
128
Chassis Level Configuration Only
All of Configurations including Chassis and switches
Fabric
Virtual Chan

10. Type the fabric ID of the logical switch in Template Fabric ID.

NOTE
If you are using a USB device, it must be connected and mounted before you upload or download.

11. Click Apply.

    You can monitor the progress by watching the Upload/Download Progress bar.

Uploading and downloading from USB storage

If you choose to upload or download from a USB device, you must left-click the USB port to launch the USB Port Management wizard.

1. Select Mount USB Device, and select Yes at the confirmation prompt.
2. Right click on a configuration file to access Export, Copy, and Search options

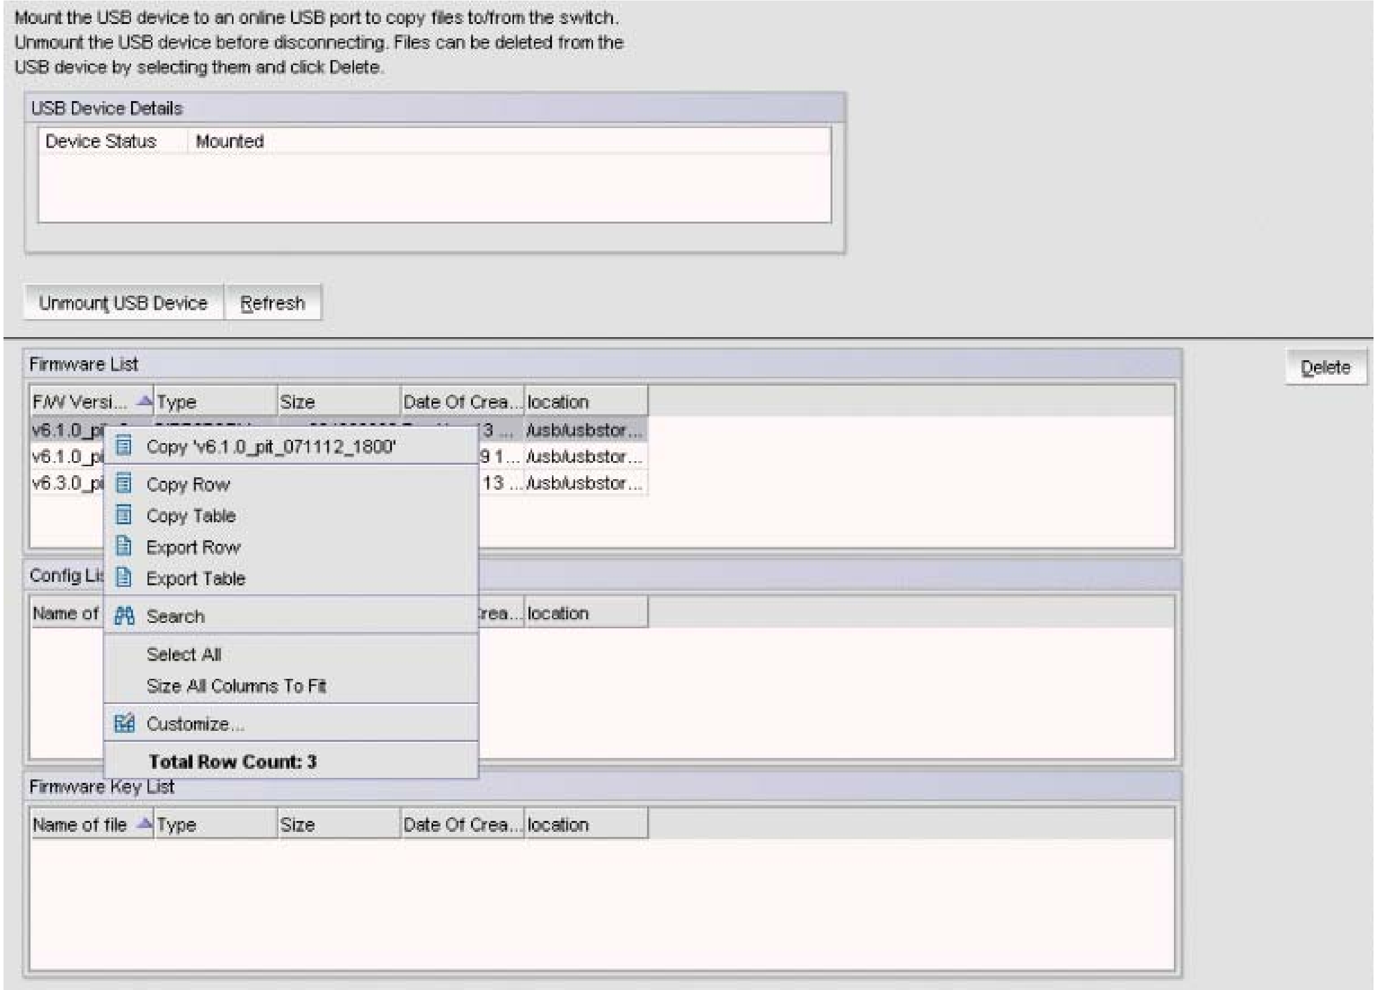

Mount the USB device to an online USB port to copy files to/from the switch. Unmount the USB device before disconnecting. Files can be deleted from the USB device by selecting them and click Delete.

USB Device Details

Device Status	Mounted

Firmware List

F/W Versi...	Type	Size	Date Of Crea...	location
v6.1.0_pi			3 ...	/usb/usbstor...
v6.1.0_pi			9 1...	/usb/usbstor...
v6.3.0_pi			13 ...	/usb/usbstor...

Config Lis

Firmware Key List

Name of file	Type	Size	Date Of Crea...	location

3. Click Copy to upload and Export to download.

Performing a firmware download

During a firmware download, the switch restarts and the browser temporarily loses connection with the switch. When the connection is restored, the version of the software running in the browser is different from the new software version that was installed and activated on the switch. You mus close all of the Web Tools windows and log in again to avoid a firmware version mismatch. Note that for chassis-based switches, you might get popup messages that imply the loss of connection is temporary and will soon be resolved. You must still close all windows and re-log in.

When you request a firmware download, the system first checks the file size being downloaded. If the compact flash does not have enough space, Web Tools displays a message and the download does not occur. If this happens, contact your switch support supplier.

NOTE
You can perform a firmware download only when the current Admin Domain owns the switch.

1. Open the Switch Administration window as described on.
2. Click the Firmware Download tab.

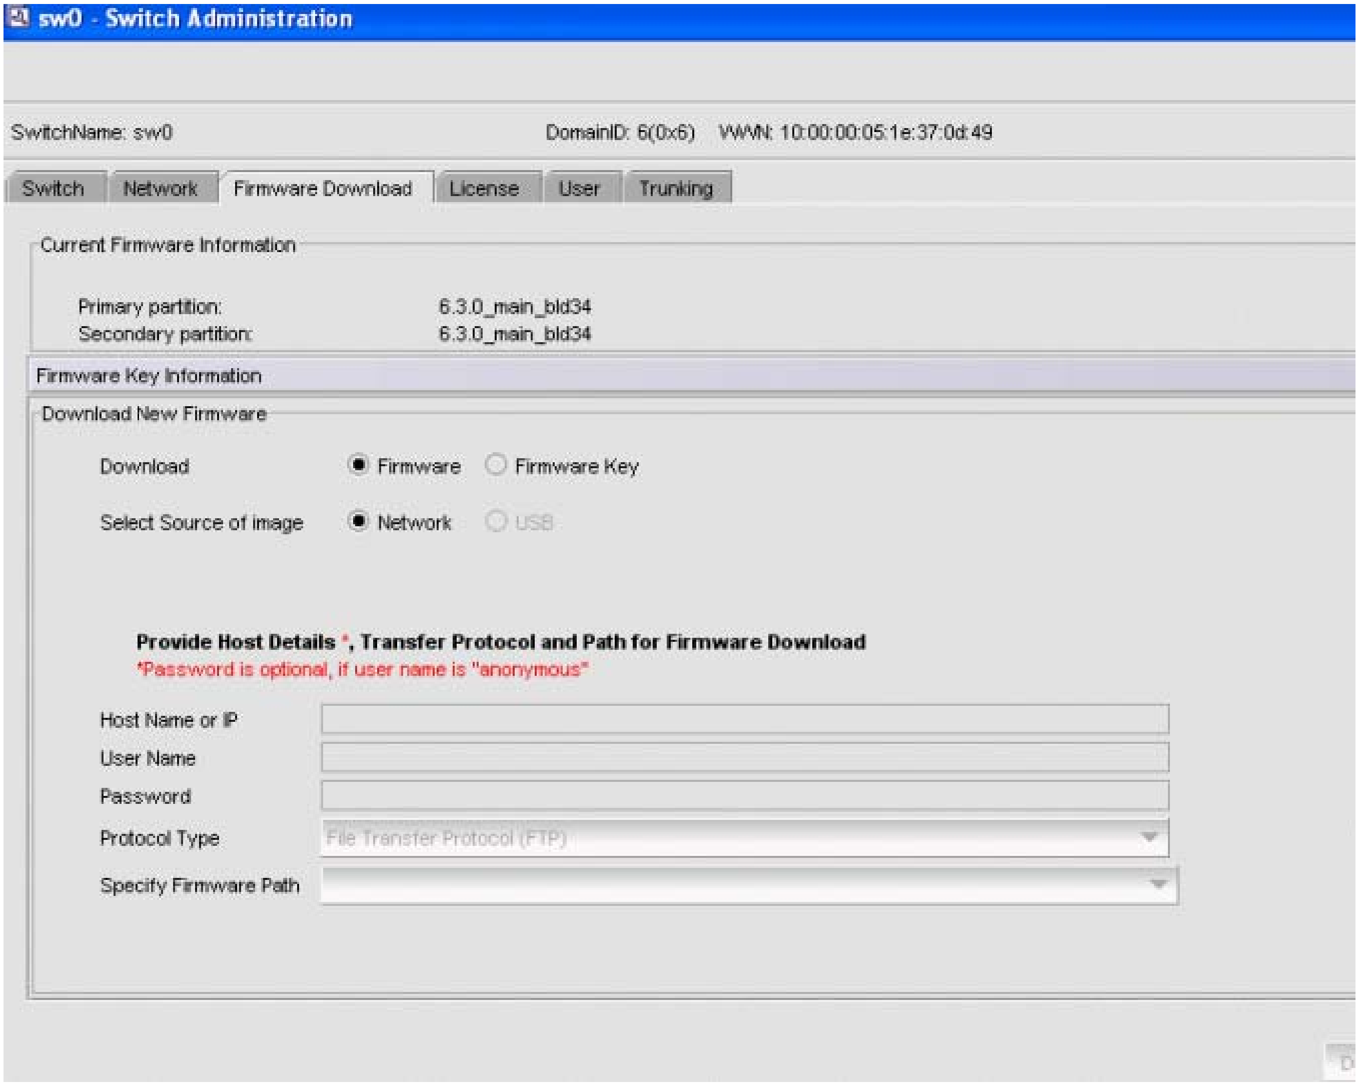

sw0 - Switch Administration

SwitchName: sw0 DomainID: 6(0x6) WWN: 10:00:00:05:1e:37:0d:49

Current Firmware Information

Primary partition: 6.3.0_main_bld34
Secondary partition: 6.3.0_main_bld34

Download ● Firmware ○ Firmware Key

Select Source of image ● Network ○ USB

Provide Host Details *, Transfer Protocol and Path for Firmware Download
*Password is optional, if user name is "anonymous"

Protocol Type File Transfer Protocol (FTP)

3. Choose whether you are downloading the firmware or the firmware key.
4. Choose whether the download source is located on the network or a USB device.

   When you select the USB button, you can specify only a firmware path or directory name. No other fields on the tab are available. The USB button is available if the USB is present on the switch.
5. Type the host name or IP address, user name, password, and fully-qualified path to the file *release.plist*.

   You can enter the IP address in either IPv4 or IPv6 format.

   The path name should follow the structure below:

//<directory>/<fos_version_directory>/release.plist where the *<directory>* is the path up to the entry point of *<fos_version_directory>* and

<fos_version_directory> is where the unzipped version of Fabric OS is located. For example:

//directory_1/my_directory/v6.3.0/release.plist

6. Select the protocol type in the Protocol Type field.

   If you choose "Secure Copy Protocol (SCP)," you cannot specify "anonymous" in the User field.

7. Click Apply. The firmware download begins. You can monitor the progress by looking at the Firmware Download progress bar.

   About halfway through the download process, after the firmware key is downloaded to the switch, connection to the switch is lost and Web Tools invalidates the current session. (Web Tools invalidates all windows because upfront login is always enabled and cannot be disabled.

8. Close all Web Tools windows and log in again.

9. If the firmware download is in progress when you log in, you can continue to monitor its progress.

Chapter 6: Licensing

Licensing overview

Feature licenses may be part of the licensed paperpack supplied with your switch software; if not, you can purchase licenses separately from your switch vendor, who will provide you with transaction keys to unlock the features. License keys are provided on a per-product and per-feature basis. Each switch within a fabric needs its own licensing.

NOTE
To preserve licenses on your switch, perform a configUpload prior to upgrading or downgrading your Fabric OS.

If you downgrade your Fabric OS to an earlier version, some licenses associated with specific features of Fabric OS may not work.
Licences can be associated with a feature version. If a feature has a version-based license, that license is valid only for a particular version of the feature. If you want a newer version of the feature, you must purchase a new license. If a license is not version-based, then it is valid for all versions of the feature.

8G License

ATTENTION
This license is installed by default and you should not remove it.

The 8 Gbps licensing applies to the Brocade 300, 5100,5300 switches. This license does NOT apply to the Brocade 6505, 6510 and 6520
The following list describes the basic rules of using, adding, or removing 8G licenses.

- Without an 8G license, even if there is an 8 Gbps SFP plugged into a port in an applicable platform, the port would be enabled to run at a maximum speed of 4 Gbps.
- To obtain an 8G license, only the License ID from the switch is required. When you add the 8G license, you must enter either the **portDisable** and **portEnable** commands on each individual port on the switch, or the **switchDisable** and **switchEnable** commands on the switch, to enable the 8 Gbps functionality.
- When you remove the 8G license, the ports which are online and already running at 8 Gbps are not disturbed until the port goes offline or the switch is rebooted. The switch ports return to their pre-licensed state maximum speed of 4 Gbps

10G licensing

The 10 Gbps FCIP/Fibre Channel license (10G license) enables the following features:

- 10 Gbps access on the 16 Gbps FC ports on the Brocade 6510 and 6520 switches, this feature is new since the Fabric OS v7.0.0 and v7.1.0 (on 6520 switch) release. When this license is applied to the Brocade 6510 switch, it is applied to the whole chassis.

Add the 10G license to the chassis using the LicenseAdd command, as for any license.

After applying a 10G license to the Brocade 6510 and 6520 chassis or to a 16 Gbps FC blade, you must also configure the port octet (portCfgOctetSpeedCombo command) with the correct port octet speed group and configure each port to operate at 10 Gbps (portCfgSpeed command). It is necessary to configure the port octet because only certain combinations of port speeds are allowed within the port octet. No license is required for the octet group. If the speed configuration operation succeeds and a 10G-capable SFP is inserted in the port connector, the port will allow operation at 10Gbps when the link becomes active at that speed.

NOTE

10 Gbps FC capability is restricted to the ports in the first port octet group on each blade or chassis to which the license is applied. So for Brocade 6510 first 8 ports can be configured to operate at 10 Gbps.
Before removing a 10 Gbps license from an entire platform (licenseRemove command) or from a specific blade (licenseSlotCfg –remove command), you must first deconfigure all affected FC ports to no longer operate at 10Gbps.

NOTE

An FC port that is operating at 10G FC speed on a 16G FC blade or 16G FC switch does not need an Extended Fabrics license to be used for FC long distance connectivity.
FC ports licensed and configured to operate at 10 Gbps on a Brocade 6510 switch or 16 Gbps FC port blade cannot interoperate with 10 Gbps ports on an FC10-6 port blade or with 10 Gbps FC ports on the M6140 platform. The new FC ports use different protocols and physical connections

Enabling 10 Gbps operation on an FC port

To enable 10 Gbps operation on an FC port on a Brocade 6510 switch follow these steps:

1. Connect to the switch and log in using an account with admin permissions, or an account with OM permissions for the license and switchportconfiguration classes of RBAC commands.
2. Use the licenseAdd command to add the 10G license.
3. *Bladed platforms only:* Use the licenseShow command to check the results of automatic license assignment. If the results are not what you intended, use the licenseSlotCfg command to reassign the license to the desired blades.
4. Use the licenseShow command to verify the license.
5. Use the portCfgOctetSpeedCombo command to set the combination speed for the first port octet to a setting that supports 10 Gbps operations. Valid settings for 10 Gbps operations include:
    - 2—autonegotiated or fixed port speeds of 10 Gbps, 8 Gbps,4 Gbps, and 2 Gbps
    - 3—autonegotiated or fixed port speeds of 16 Gbps and 10 Gbps
6. Use the portCfgSpeed command to set the port speed on each port you want to operate at 10 Gbps.

Example of assigning a 10G license on an FC port blade and enabling 10 Gbps operation on a port

This example assigns a license to the Brocade 6510 switch and enables 10 Gbps operation on port 2.

```
6510-switch:admin> licenseadd aTFPNFXGLmABANMGtT4LfSBJSDLWTYD3EFrr4WGAEMBA
6510-switch:admin> licenseshow
aTFPNFXGLmABANMGtT4LfSBJSDLWTYD3EFrr4WGAEMBA
10 Gigabit FCIP/Fibre Channel (FTR_10G) license
Capacity 1
 Consumed 1

6510-switch:admin> portcfgoctetspeedcombo 2
6510-switch:admin> portcfgspeed 2 10
6510-switch:admin>
```

Time-based licenses

A Time-based license applies a try-before-you-buy approach to certain features so that you can experience the feature and its capabilities prior to buying the license. Once you have installed the license, you are given a time limit to use the feature. The following lists the types of licenses that have this time-based trial feature:

- 10 Gigabit FCIP/Fibre Channel license
- Advanced Extension
- Advanced FICON Acceleration license

- Adaptive Networking
- Advanced Performance Monitoring
- Fabric
- Fabric Watch
- Extended Fabric
- High Performance Extension over FCIP/FC
- Integrated Routing
- Trunking

For more on Time based licenses see Fabric OS Administrator's Guide.

Viewing installed licenses

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

Adding a licensed feature

To enable a feature, go to the feature's appropriate section in this manual. Enabling a feature on a switch may be a separate task from adding the license.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Activate the license using the **licenseAdd** command.
3. Verify the license was added by entering the **licenseShow** command. The licensed features currently installed on the switch are listed. If the feature is not listed, enter the **licenseAdd** command again.

Some features may require additional configuration, or you may need to disable and re-enable the switch to make them operational; see the feature documentation for details.

```
switch:admin> licenseshow
aAYtMJg7tmMZrTZ9JTWBC4SXWLJMY3QfBJYHG:
Fabric license
Remote Switch license
Remote Fabric license
Extended Fabric license
Entry Fabric license
Fabric Watch license
Performance Monitor license
Trunking license
4 Domain Fabric license
FICON_CUP license
High-Performance Extension over FCIP/FC license
Full Ports on Demand license - additional 16 port upgrade license
2 Domain Fabric license
Integrated Routing license
Storage Application Services license
FICON Tape license
FICON XRC license
Adaptive Networking license
Inter Chassis Link license
Enhanced Group Management license
8 Gig FC license
DataFort Compatibility license
Server Application Optimization license
```

Removing a licensed feature

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **licenseShow** command to display the active licenses.
3. Remove the license key using the **licenseRemove** command.
   The license key is case-sensitive and must be entered exactly as given. The quotation marks are optional. After removing a license key, the licensed feature is disabled when the switch is rebooted or when a switch disable and enable is performed.
4. Enter the licenseShow command to verify the license is disabled.

```
switch:admin> licenseshow
bQebzbRdScRfc0iK:
Entry Fabric license
Fabric Watch license
switch:admin> licenseremove "bQebzbRdScRfc0iK"
removing license key "bQebzbRdScRfc0iK"
```

After a reboot (or switchDisable and switchEnable), only the remaining licenses appear:

```
switch:admin> licenseshow
SybbzQQ9edTzcc0X:
Fabric license
```

If there are no license keys, licenseShow displays “No licenses.”

Ports on Demand

The Brocade models in the following list can be purchased with the number of licensed ports indicated. As your needs increase, you can activate unlicensed ports up to a particular maximum by purchasing and installing the optional Ports on Demand licensed product:

Brocade 300– Can be purchased with eight ports and no E_Port, eight ports with full fabric access, or 16 ports with full fabric access. A maximum of 16 ports is allowed; eight-port systems can be upgraded in four-port increments. An E_Port license upgrade is also available for purchase.

Brocade 5100– Can be purchased with 24, 32, or 40 licensed ports. A maximum of 40 ports is allowed.

Brocade 5300– Can be purchased with 48, 64, or 80 licensed ports. A maximum of 80 ports is allowed.

Brocade 6505—Can be purchased with 12 or 24 licensed ports. A maximum of 24 ports is allowed.

Brocade 6510— Can be purchased with a maximum of 48 licensed ports. Configurations can be 24, 36, or 48 licensed ports.

Brocade 6520 – Can be purchased with a maximum of 96 licensed ports. Configurations can be 48, 72 or 96 licensed ports.

ATTENTION

Licenses are not interchangeable between units. For example, if you bought a POD license for a Brocade 300, you cannot use that license on a Brocade 5100. The licenses are based on the switches' License Identifiers and are not interchangeable.

Ports on Demand is ready to be unlocked in the switch firmware. Its license key may be part of the licensed paperpack supplied with switch software, or you can purchase the license key separately from your switch vendor. You may need to generate a license key from a transaction key supplied with your purchase. If so, launch an Internet browser and go to the Brocade Web site at *http://www.brocade.com*. Click Products > Software Products > Software License Keys and follow the instructions to generate the key.

Each Ports on Demand license activates the next group of ports in numerical order in either four-port or eight-port increments, depending on the model. Before installing a license key, you must insert transceivers in the ports to be activated. Remember to insert the transceivers in the lowest group of inactive port numbers first. For example, if only 16 ports are currently active and you are installing one Ports on Demand license key, make sure to insert the transceivers in ports 16 through 23. If you later install a second license key, insert the transceivers in ports 24 through 31. For details on inserting transceivers, see the switch’s *Hardware Reference Manual*

Activating Ports on Demand

1. Connect to the switch and log in using an account assigned to the admin role.
2. Verify the current states of the ports, using the **portShow** command.
   In the portShow output, the Licensed field indicates whether the port is licensed.
3. Install the Brocade Ports on Demand license.
4. Use the **portEnable** command to enable the ports.
   Alternatively, you can disable and re-enable the switch to activate ports.
5. Use the **portShow** command to check the newly activated ports.

Displaying the port license assignments

When you display the available licenses, you can also view the current port assignment of those licenses and the POD method state of dynamic or static.

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

Example of manually assigned POD licenses.

```
switch:admin> licenseport –show
24 ports are available in this switch
Full POD license is installed
Static POD method is in use
24 port assignments are provisioned for use in this switch:
12 port assignments are provisioned by the base switch license
12 port assignments are provisioned by a full POD license
24 ports are assigned to installed licenses:
12 ports are assigned to the base switch license

12 ports are assigned to the full POD license
Ports assigned to the base switch license:
1, 2, 3, 4, 5, 6, 7, 8, 17, 18, 19, 20
Ports assigned to the full POD license:
0, 9, 10, 11, 12, 13, 14, 15, 16, 21, 22, 23
```

Web tools

Licensed feature management

The licensed features currently installed on the switch are listed in the License tab of the Switch Administration window. If the feature is listed, such as the EGM license, it is installed and immediately available. When you enable some licenses, such as ISL Trunking, you might need to change the state of the port to enable the feature on the link. For time-based licenses, the expiry date is included.

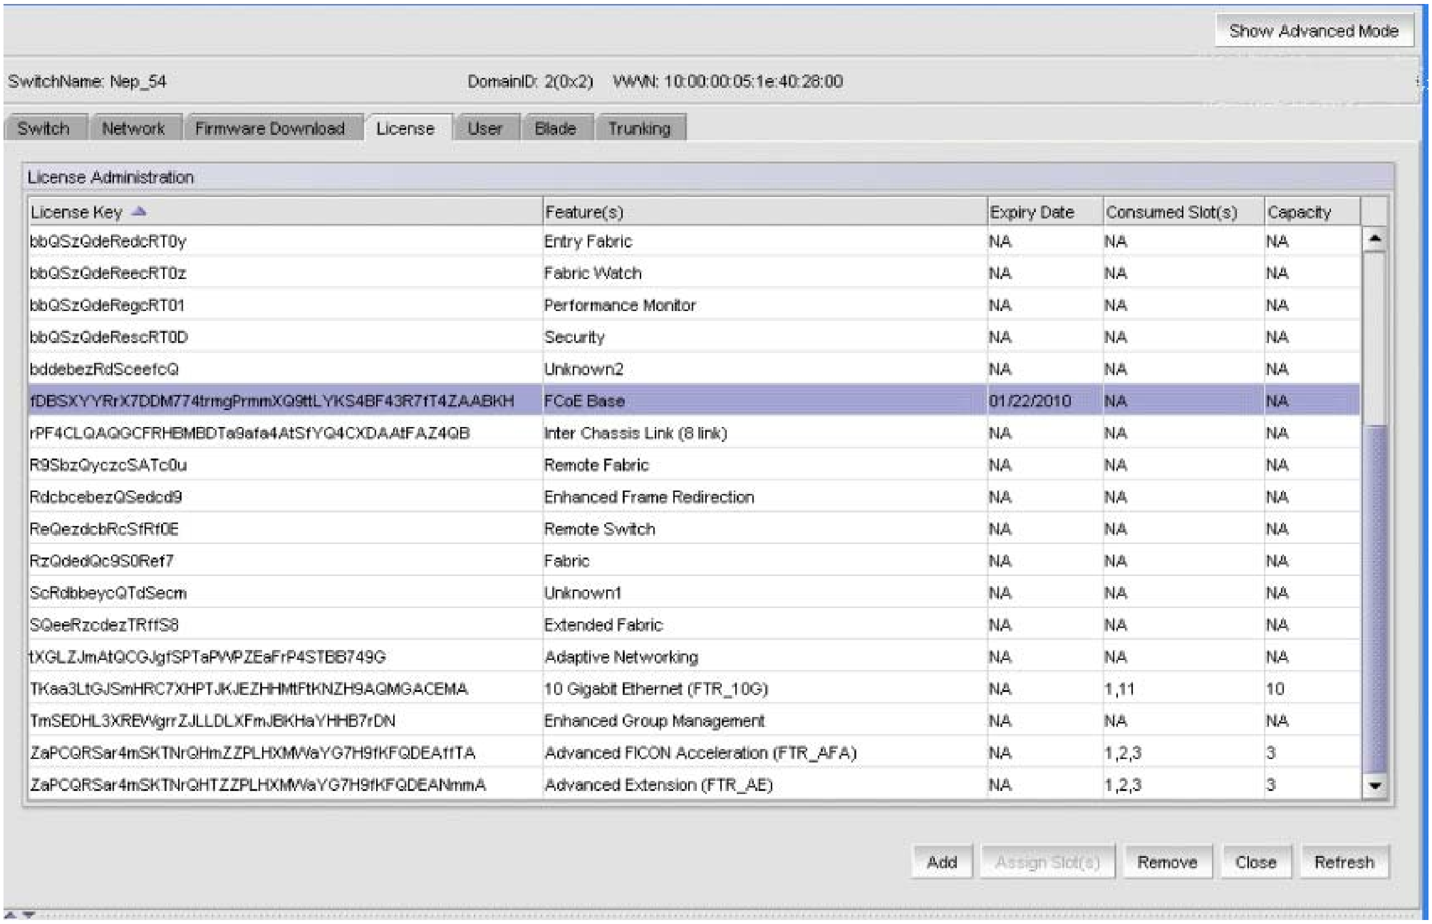

SwitchName: Nep_54 DomainID: 2(0x2) WWN: 10:00:00:05:1e:40:28:00

License Key	Feature(s)	Expiry Date	Consumed Slot(s)	Capacity
bbQSzQdeRedcRT0y	Entry Fabric	NA	NA	NA
bbQSzQdeReecRT0z	Fabric Watch	NA	NA	NA
bbQSzQdeRegcRT01	Performance Monitor	NA	NA	NA
bbQSzQdeRescRT0D	Security	NA	NA	NA
bddebezRdSceefcQ	Unknown2	NA	NA	NA
fDBSXYYRrX7DDM774trmgPrmmXQ9ttLYKS4BF43R7fT4ZAABKH	FCoE Base	01/22/2010	NA	NA
rPF4CLQAQGCFRHBMBDTa9afa4AtSfYQ4CXDAAtFAZ4QB	Inter Chassis Link (8 link)	NA	NA	NA
R9SbzQyczcSATc0u	Remote Fabric	NA	NA	NA
RdcbcebezQSedcd9	Enhanced Frame Redirection	NA	NA	NA
ReQezdcbRcSfRf0E	Remote Switch	NA	NA	NA
RzQdedQc9S0Ref7	Fabric	NA	NA	NA
ScRdbbeycQTdSecm	Unknown1	NA	NA	NA
SQeeRzcdezTRffS8	Extended Fabric	NA	NA	NA
tXGLZJmAtQCGJgfSPTaPWPZEaFrP4STBB749G	Adaptive Networking	NA	NA	NA
TKaa3LtGJSmHRC7XHPTJKJEZHHMtFtKNZH9AQMGACEMA	10 Gigabit Ethernet (FTR_10G)	NA	1,11	10
TmSEDHL3XREWgrrZJLLDLXFmJBKHaYHHB7rDN	Enhanced Group Management	NA	NA	NA
ZaPCQRSar4mSKTNrQHmZZPLHXMWaYG7H9fKFQDEAffTA	Advanced FICON Acceleration (FTR_AFA)	NA	1,2,3	3
ZaPCQRSar4mSKTNrQHTZZPLHXMWaYG7H9fKFQDEANmmA	Advanced Extension (FTR_AE)	NA	1,2,3	3

Right-click a license key to export data, copy data, or search the table.

Activating a license on a switch

Before you can unlock a licensed feature, you must obtain a license key. You can either use the license key provided in the paperpack document supplied with switch software or refer to the *Fabric OS Administrator's Guide* for instructions on how to obtain a license key at the Brocade Website (*www.brocade.com*).

Use the following procedure to activate a license.

1. Open the Switch Administration window.
2. Click the License tab and click Add.

   The Add License dialog box displays.
3. Paste or type a license key in the field.
4. Click Add License.
5. Click Refresh to display the new licenses in the License tab.

   Some licenses, such as the Trunking or the 7500E and 7800 upgrade license, do not take effect until the switch is restarted.

Removing a license from a switch

Use the following procedure to remove a license from a switch in the Switch Administration window.

ATTENTION

Use care when removing licenses. If you remove a license for a feature, that feature will no longer work.

1. Open the Switch Administration window.
2. Click the License tab.
3. Click the license you want to remove.
4. Click Remove.

Universal time based licensing

After v6.3.0, Web Tools supports universal time based licensing. Each universal key is for a single feature, and can be used on any product that supports the feature, for a defined trial period. At the end of the trial period, the feature gets disabled. You can extend the universal key license. For time-based licenses, the Expiry Date is displayed in the License Administration table.

.

The following features are supported for universal time based license:

The following features are supported for universal time-based license:

- Fabric
- Extended Fabric
- Fabric Watch
- Performance Monitor
- Trunking
- High-Performance Extension over FCIP/FC
- Advanced Extension
- Advanced FICON Acceleration
- FICON Management Server (CUP)
- Enhanced Group Management (EGM)
- 10GbE
- Integrated Routing
- Adaptive Networking
- Server Application Optimization

Chapter 7: Virtual Fabrics

Overview

Virtual fabric is an architecture to virtualize hardware boundaries. It is a suite consisting of following features:

- Logical Switch
- Logical Fabric
- Device Sharing

Note:

- Virtual Fabrics is just a name of the feature. You can create a fabric called logical fabric using this feature.
- Virtual Fabrics and Admin Domains are mutually exclusive and are not supported at the same time on the switch

The following platforms are Virtual Fabrics-capable in the pizzabox form factor:

- Brocade 5100
- Brocade 5300
- Brocade 6510
- Brocade 6520
- Brocade 7800

Enabling Virtual Fabric mode

Virtual Fabric mode is not supported on Brocade 300 and 6505 Switches.

A fabric is said to be in Virtual Fabrics mode (VF mode) when the Virtual Fabrics feature is enabled. Before you can use the Virtual Fabrics features, such as logical switch and logical fabric, you must enable VF mode. VF mode is disabled by default on switches that you upgrade to Fabric OS 6.2.0 or later. VF mode is enabled by default on a new chassis.

Please note that if Virtual Fabric mode is enabled on a switch that is part of a multi-switch fabric, the fabric remains intact and the configuration and zoning information is retained on the switches. This is because when VF mode is enabled, a single logical default switch is created and contains all the ports on the switch.

Steps to enable VF mode:

1. Log in using an account having admin privileges.
2. Issue **fosconfig - -show** to check if VF mode is enabled.

```
switch:admin> fosconfig --show
FC Routing service:            disabled
iSCSI service:                 Service not supported on this Platform
iSNS client service:           Service not supported on this Platform
Virtual Fabric:                disabled
Ethernet Switch Service:       enabled
```

3. Delete all Admin Domains prior to enabling the VF mode.
4. Issue **fosconfig - - enable vf** to enable VF mode

```
switch:admin> fosconfig --enable vf
WARNING:  This is a disruptive operation that requires a reboot to take effect.
All EX ports will be disabled upon reboot.
Would you like to continue [Y/N]: Y
VF has been enabled.  Your system is being rebooted.
```

Enabling Virtual Fabrics creates a single logical switch in the physical chassis. This logical switch is called the *default logical switch*, and it initially contains all of the ports in the physical chassis. It has a fabric ID of 128. In this example, the switch has 10 ports, labeled P0 through P9.

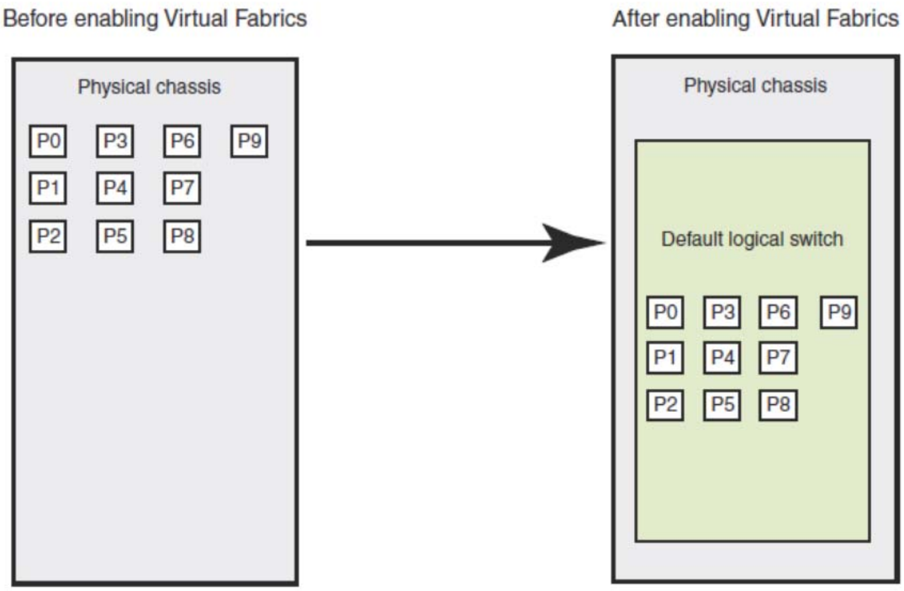

Logical Switches

Creating Logical Switches

You can create more logical switches within your physical switch (up to 8 logical switches possible).

Please note that when creating a new logical switch that does not contain any ISL ports, this switch will not reflect any of the zoning or configuration information of the default switch. It will exist as a single switch fabric with its own zoning and configuration information.

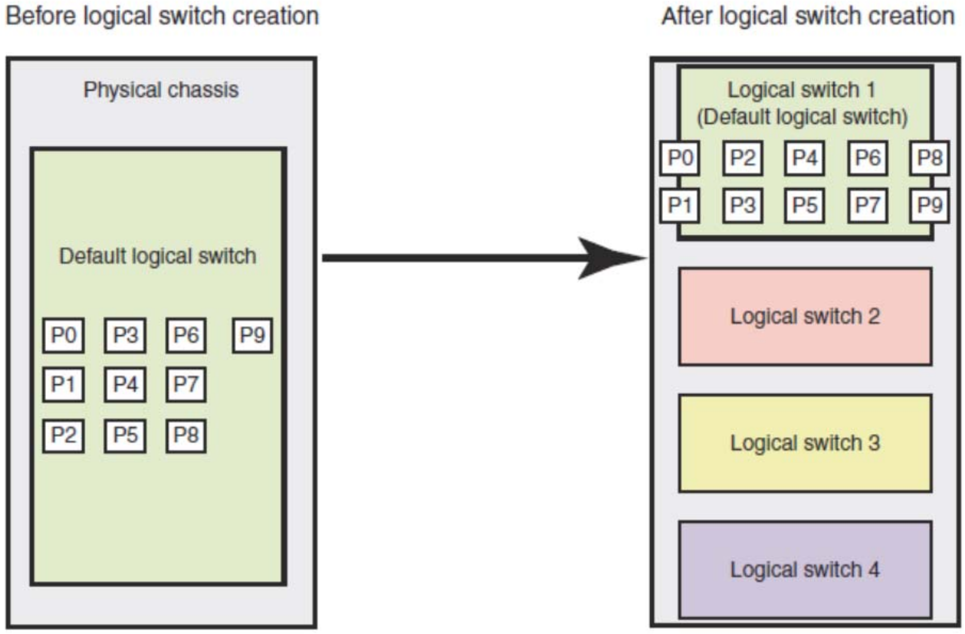

Steps to create logical switches:

1. Enter the following command to create a logical switch:
   lscfg –create *fabricID* [**-base**] [**-force**]
   where *fabricID* is the fabric ID that is to be associated with the logical switch.
   Specify the **-base** option if the logical switch is to be a base switch.
   Specify the **-force** option to execute the command without any user prompts or confirmation.

2. Set the context to the new logical switch.
   setcontext *fabricID*
   where *fabricID* is the fabric ID of the logical switch you just created.

3. Disable the logical switch.
   Switchdisable

4. Configure the switch attributes, including assigning a unique domain ID.
   Configure

5. Enable the logical switch:
   Switchenable

```
sw0:FID128:admin> lscfg –create 4
About to create switch with fid=4. Please wait...
Logical Switch with FID (4) has been successfully created.
Logical Switch has been created with default configurations.
Please configure the Logical Switch with appropriate switch
and protocol settings before activating the Logical Switch.

sw0:FID128:admin> setcontext 4
Please change passwords for switch default accounts now.
Use Control-C to exit or press 'Enter' key to proceed.

Password was not changed. Will prompt again at next login
until password is changed.
```

```
switch_4:FID4:admin> switchdisable
switch_4:FID4:admin> configure

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

Domain: (1..239) [1] 14
WWN Based persistent PID (yes, y, no, n): [no]
...
(output truncated)
WARNING: The domain ID will be changed. The port level zoning may be affected
switch_4:FID4:admin> switchenable
```

Assigning/Removing ports to logical switches

When you create a logical switch, it has no ports assigned to it. You add ports to a logical switch by moving the ports from one logical switch to another. When you move a port from one logical switch to another, the port is automatically disabled. Any performance monitors that were installed on the port are deleted. If monitors are required in the new logical switch, you must manually reinstall them on the port after the move.

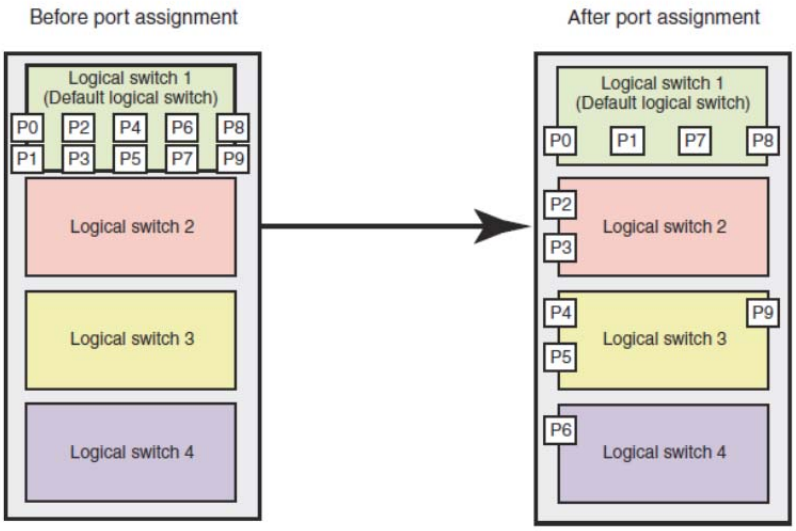

Steps to add ports to logical switch:

1. Connect to the physical chassis and log in using an account assigned to the admin role.
2. Enter the following command to move ports from one logical switch to another:
   lscfg –config *fabricID* **-port** *port* [**-force**]
3. The ports are automatically disabled, then removed from their current logical switch and assigned to the logical switch specified by *fabricID*.
4. Specify the **-force** option to execute the command without any user prompts or confirmation.

```
sw0:FID128:admin> lscfg –config 5 -port 1-3
```

```
This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Making this configuration change. Please wait...
Configuration change successful.
Please enable your ports/switch when you are ready to continue.
```

In the above example we assigned ports 1, 2 and 3 to logical switch with fabric ID 5.

Displaying logical switch configuration

1. Connect to the physical chassis and log in using an account assigned to the admin role.
2. Enter the command **lscfg –show** to display a list of all logical switches and the ports assigned to them

```
sw0:FID128:admin> lscfg –show
Created switches: 128(ds) 4 5
Port 0 1 2 3 4 5 6 7 8 9
-------------------------------------------------------------
FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |
Port 10 11 12 13 14 15 16 17 18 19
-------------------------------------------------------------
FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 5 | 5 |
Port 20 21 22 23 24 25 26 27 28 29
-------------------------------------------------------------
FID 5 | 128 | 4 | 4 | 128 | 128 | 128 | 128 | 128 | 128 |
Port 30 31 32 33 34 35 36 37 38 39
-------------------------------------------------------------
FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |
```

Changing Fabric ID of switch

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the following command to change the fabric ID of a logical switch:
   lscfg –change *fabricID* **-newfid** *newFID* [**-force**]
   Specify the -force option to execute the command without any user prompts or confirmation.
3. Enable the logical switch

```
sw0:FID128:admin> lscfg –change 5 -newfid 7
Changing of a switch fid requires that the switch be disabled.
Would you like to continue [y/n]?: y
Disabling switch...
All active login sessions for FID 5 have been terminated.
Checking and logging message: fid = 5.
Please enable your switch.
sw0:FID128:admin> fosexec –fid 7 -c "switchenable"
```

If you are in the context of the logical switch whose fabric ID you want to change, you are automatically logged out when the fabric ID changes. To avoid being logged out, make sure you are in the context of a different logical switch from the one whose fabric ID you are changing.

Setting /Removing IP address for fabric

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ipAddrSet -ls** command.
3. Enter the network information in dotted-decimal notation for the Ethernet IPv4 address with a CIDR prefix.

Example of setting an IP address for a logical switch in a Virtual Fabric with an FID of 123 in non-interactive mode with the CIDR prefix:

```
switch:admin> ipaddrset -ls 123 –add 11.1.2.4/24
```

Enter the ipAddrSet -ls FID - -delete command.

```
switch:admin> ipaddrset -ls 123 –delete
```

Logical Fabric and ISL sharing

When you divide a chassis into logical switches, you can designate one of the switches to be a base switch. A *base switch* is a special logical switch that is used for interconnecting the physical chassis. A base switch can be connected to other base switches through a special ISL, called a *shared ISL* or *extended ISL* (XISL). An extended ISL is an ISL that connects base switches. The XISL is used to share traffic among different logical fabrics. Fabric formation across an XISL is based on the FIDs of the logical switches.

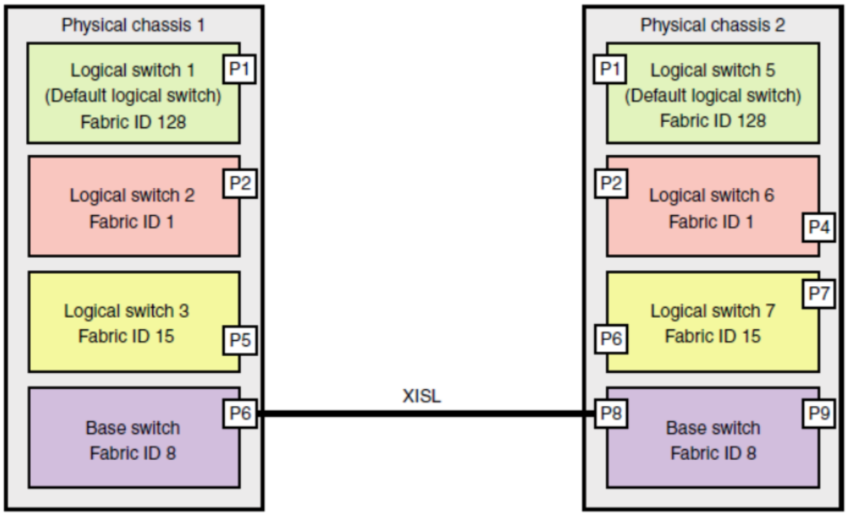

Configuring the switch to use XISL

When you create a logical switch, by default it is configured to use XISLs. Use the following procedure to allow or disallow the logical switch to use XISLs in the base fabric.

1. Connect to the physical chassis and log in using an account assigned to the admin role.
2. Set the context to the logical switch you want to manage, if you are not already in that context.
   setcontext *fabricID*
   where *fabricID* is the fabric ID of the logical switch you want to switch to and manage.
3. Enter the switchShow command and check the value of the Allow XISL Use parameter.
4. Disable the logical switch.

switchdisable

5. Enter the following command:
   configure
   Enter y after the Fabric Parameters prompt:
   Fabric parameters (yes, y, no, n): [no] **y**
   Enter y at the Allow XISL Use prompt to allow XISL use; enter n at the prompt to disallow XISL use:
   Allow XISL Use (yes, y, no, n): **y**
6. Respond to the remaining prompts or press Ctrl-d to accept the other settings and exit.
7. Enable the logical switch.
   Switchenable

Deleting a logical switch

Before deleting a logical switch you must remove all the ports from the switch. You cannot delete the default switch. If you are in the context of the switch you are deleting, you will be automatically logged out when issue the command to delete the switch. To avoid getting logged out make sure you are in context of a different switch while issuing the command.

Here are the steps to delete the logical switch

1. Connect to the physical chassis and log in using an account assigned to the admin role.
2. Remove all ports from the logical switch as described in the section **Assigning/Removing ports to logical switches.**
3. Enter the following command to delete the logical switch:
   lscfg –delete *fabricID* [-force]
   where *fabricID* is the fabric ID of the logical switch to be deleted.
   Specify the -force option to execute the command without any user prompts or confirmation.

```
switch_4:FID4:admin> lscfg –delete 7
All active login sessions for FID 7 have been terminated.
Switch successfully deleted.
```

Disable Virtual Fabrics

Here are the steps to disable virtual fabrics:

1. Connect to the physical chassis and log in using an account assigned to the admin role with the chassis-role permission.
2. Enter the following command to check whether VF mode is disabled:
   fosconfig –show
3. Delete all of the non-default logical switches, as described in the above section “Deleting a logical switch”.
4. Enter the following command to disable VF mode:
   fosconfig –disable vf
5. Enter **y** at the prompt

```
switchA:FID128:admin> fosconfig –show

FC Routing service: disabled
iSCSI service: Service not supported on this Platform
iSNS client service: Service not supported on this Platform
Virtual Fabric: enabled

switch:admin> fosconfig –disable vf
WARNING: This is a disruptive operation that requires a reboot to take
```

```
effect.
Would you like to continue [Y/N] y
```

Enabling Virtual Fabrics in an active muti switch fabric:

Although enabling virtual fabrics on a Brocade Switch requires a reboot but there is no impact on existing zoning configuration. After enabling virtual fabrics existing switch zoning configuration gets imported on the Default Logical Switch FID 128. There is no impact on ISLs as all ISLs links would be under Default Switch.

Here are the sample command outputs from the switch that we tested in our lab:

Switch ESNSVT_6510 with Virtual Fabric Disabled

```
ESNSVT_6510:root> fosconfig --show
FC Routing service:             disabled
iSCSI service:                  Service not supported on this Platform
iSNS client service:            Service not supported on this Platform
Virtual Fabric:                 disabled
Ethernet Switch Service:        Service not supported on this Platform

ESNSVT_6510:root> fabricshow
Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 10:00:00:05:33:6e:4e:96 192.168.50.1   0.0.0.0         "ESNSVT_6510"
  2: fffc02 10:00:00:05:33:a9:f0:ae 192.168.50.5   0.0.0.0         "ESNSVT_6505"
  6: fffc06 10:00:00:05:1e:52:fc:00 192.168.50.10   0.0.0.0        >"ESNSVT_DCX4S"
 97: fffc61 10:00:00:05:1e:43:68:00 192.168.50.15   0.0.0.0         "ESNSVT_DCX"

ESNSVT_6510:root> cfgactvshow
Effective configuration:
 cfg:   ZonesetA
 zone:  ZoneA   50:06:04:8a:d5:f0:ea:9e
                50:06:04:8a:d5:f0:ea:ae
                10:00:00:00:c9:3c:f7:c4

ESNSVT_6510:root> cfgshow
Defined configuration:
 cfg:   ZonesetA
                ZoneA
 zone:  ZoneA   50:06:04:8A:D5:F0:EA:9E; 50:06:04:8A:D5:F0:EA:AE;
                10:00:00:00:C9:3C:F7:C4
 alias: TestHost
                1,1
 alias: TestStorage
                1,3

Effective configuration:
 cfg:   ZonesetA
 zone:  ZoneA   50:06:04:8a:d5:f0:ea:9e
                50:06:04:8a:d5:f0:ea:ae
                10:00:00:00:c9:3c:f7:c4
```

```
ESNSVT_6510:root> switchshow
switchName:     ESNSVT_6510
switchType:     109.1
switchState:    Online
switchMode:     Native
switchRole:     Subordinate
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:00:05:33:6e:4e:96
zoning:         ON (ZonesetA)
switchBeacon:   OFF
FC Router:      OFF
FC Router BB Fabric ID: 128
Address Mode:   0

Index Port Address Media Speed       State   Proto
==================================================
   0   0   010000   id    N16      Online      FC  E-Port  10:00:00:05:1e:43:68:00
"ESNSVT_DCX" (Trunk master)
   1   1   010100   --    N16      No_Module   FC
   2   2   010200   --    N16      No_Module   FC
   3   3   010300   --    N16      No_Module   FC
   4   4   010400   id    N16      Online      FC  E-Port  10:00:00:05:1e:43:68:00
"ESNSVT_DCX" (upstream)(Trunk master)
   5   5   010500   --    N16      No_Module   FC
   6   6   010600   --    N16      No_Module   FC
```

```
ESNSVT_6510:root> fosconfig --enable vf

WARNING:  This is a disruptive operation that requires a reboot to take effect.

All EX ports will be disabled upon reboot.

Would you like to continue [Y/N]: y
```

Switch ESNSVT_6510 with Virtual Fabric Enabled

```
ESNSVT_6510:FID128:root> fosconfig --show
FC Routing service:             disabled
iSCSI service:                  Service not supported on this Platform
iSNS client service:            Service not supported on this Platform
Virtual Fabric:                 enabled
Ethernet Switch Service:        Service not supported on this Platform
ESNSVT_6510:FID128:root>
```

```
ESNSVT_6510:FID128:root> fabricshow
Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 10:00:00:05:33:6e:4e:96 192.168.50.1   0.0.0.0         "ESNSVT_6510"
  2: fffc02 10:00:00:05:33:a9:f0:ae 192.168.50.5   0.0.0.0         "ESNSVT_6505"
  6: fffc06 10:00:00:05:1e:52:fc:00 192.168.50.10   0.0.0.0        >"ESNSVT_DCX4S"
 97: fffc61 10:00:00:05:1e:43:68:00 192.168.50.15   0.0.0.0         "ESNSVT_DCX"
```

```
ESNSVT_6510:FID128:root> cfgactvshow

Effective configuration:
 cfg:   ZonesetA
 zone:  ZoneA   50:06:04:8a:d5:f0:ea:9e
                50:06:04:8a:d5:f0:ea:ae
                10:00:00:00:c9:3c:f7:c4
ESNSVT_6510:FID128:root> cfgshow
Defined configuration:
 cfg:   ZonesetA
                ZoneA
 zone:  ZoneA   50:06:04:8A:D5:F0:EA:9E; 50:06:04:8A:D5:F0:EA:AE;
                10:00:00:00:C9:3C:F7:C4
 alias: TestHost
                1,1
 alias: TestStorage
                1,3

Effective configuration:
 cfg:   ZonesetA
 zone:  ZoneA   50:06:04:8a:d5:f0:ea:9e
                50:06:04:8a:d5:f0:ea:ae
                10:00:00:00:c9:3c:f7:c4

ESNSVT_6510:FID128:root> switchshow
switchName:     ESNSVT_6510
switchType:     109.1
switchState:    Online
switchMode:     Native
switchRole:     Subordinate
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:00:05:33:6e:4e:96
zoning:         ON (ZonesetA)
switchBeacon:   OFF
FC Router:      OFF
Allow XISL Use: OFF
LS Attributes:  [FID: 128, Base Switch: No, Default Switch: Yes, Address Mode 0]

Index Port Address Media Speed       State   Proto
==================================================
   0   0   010000   id    N16      Online      FC  E-Port  10:00:00:05:1e:43:68:00
"ESNSVT_DCX" (Trunk master)
   1   1   010100   --    N16      No_Module   FC
   2   2   010200   --    N16      No_Module   FC
   3   3   010300   --    N16      No_Module   FC
   4   4   010400   id    N16      Online      FC  E-Port  10:00:00:05:1e:43:68:00
"ESNSVT_DCX" (upstream)(Trunk master)
   5   5   010500   --    N16      No_Module   FC
   6   6   010600   --    N16      No_Module   FC
```

Web Tools

The following platforms are Virtual Fabrics-capable:

- Brocade DCX and DCX-4S
- Brocade 5300
- Brocade 5100
- Brocade 6510
- Brocade 6520
- Brocade DCX 8510-4
- Brocade DCX 8510-8

Virtual Fabrics cannot be configured or managed from Web Tools. Configuration and management is done from either the Brocade Network Advisor, or the Fabric OS command line interface. For information about configuring and managing Virtual Fabrics, refer to the *Brocade Network Advisor User Manual* if you are using Brocade Network Advisor, or *Fabric OS Administrator's Guide* if you are using the Fabric OS command line interface.

You can use Web Tools to view Virtual Fabrics and logical switch configurations.

NPIV overview

N_Port ID Virtualization (NPIV) enables a single Fibre Channel protocol port to appear as multiple, distinct ports, providing separate port identification within the fabric for each operating system image behind the port (as if each operating system image had its own unique physical port). NPIV assigns a different virtual port ID to each Fibre Channel protocol device. NPIV is designed to enable you to allocate virtual addresses without affecting your existing hardware implementation. The virtual port has the same properties as an N_Port, and is therefore capable of registering with all services of the fabric.

Each NPIV device has a unique device PID, Port WWN, and Node WWN, and should act the same as all other physical devices in the fabric; in other words, multiple virtual devices emulated by NPIV appear no different than regular devices connected to a non-NPIV port. The same zoning rules apply to NPIV devices as non-NPIV devices. Zones can be defined by *domain,port* notation, by WWN zoning, or both. To perform zoning to the granularity of the virtual N_Port IDs, you must use WWN-based zoning.

If you are using *domain,port* zoning for an NPIV port, and all the virtual PIDs associated with the port are included in the zone, then a port login (PLOGI) to a non-existent virtual PID is not blocked by the switch; rather, it is delivered to the device attached to the NPIV port. In cases where the device is not capable of handling such unexpected PLOGIs, you should use WWN-based zoning.

The following example shows the number of NPIV devices in the output of the switchShow command. The number of NPIV devices is equal to the sum of the base port plus the number of NPIV public devices. The base port is the N_Port listed in the switchShow output. Based on the formula, index 010000 shows only 1 NPIV device and index 010300 shows 222 NPIV devices.

```
switch:admin> switchshow
switchName: 5100
switchType: 71.2
switchState: Online
switchMode: Access Gateway Mode
switchWwn: 10:00:00:05:1e:41:49:3d
switchBeacon: OFF
Index Port Address Media Speed State Proto
=================================================
0 0 010000 id N4 Online FC F-Port 20:0c:00:05:1e:05:de:e4 0xa06601
1 1 010100 id N4 Online FC F-Port 1 N Port + 4 NPIV public
2 2 010200 id N4 Online FC F-Port 1 N Port + 119 NPIV public
3 3 010300 id N4 Online FC F-Port 1 N Port + 221 NPIV public
```

Configuring NPIV

The NPIV feature is enabled by default. You can set the number of virtual N_Port_IDs per port to a value between 1 and 255 per port. The default setting is 126. To specify the number of virtual N_Port_IDs per port on a switch, use the portCfgNPIVport command to enable or disable the feature. Once the feature is enabled on the port, you can specify the number of logins per port. If the feature has been disabled, then the NPIV port configuration will not work.

The addressing mode can limit the maximum number of NPIV logins to 127 or 63 depending on the mode. The portCfgNPIVPort command can set the maximum number of NPIV login limit to anything from 1 to 255, regardless of the addressing mode. Whichever of these two (addressing mode or the value configured through the portCfgNPIVPort) is lower will be the maximum number that can be logged in.

CAUTION

The portDisable command disables the port and stops all traffic flowing to and from the port.Perform this command during a scheduled maintenance.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the portDisable command.
3. Enter the portCfgNPIVPort –setloginlimit command with the port number and the number of logins per port.
4. Press Enter.
5. Enter the portEnable command to enable the port.

```
switch:admin> portcfgnpivport –setloginlimit 1 176
NPIV Limit Set to 176 for Port 1
switch:admin> portcfgshow 1
Area Number: 1
Speed Level: AUTO(HW)
Fill Word: 1(Arbff-Arbff)
AL_PA Offset 13: OFF
Trunk Port ON
Long Distance OFF

VC Link Init OFF
Locked L_Port OFF
Locked G_Port OFF
Disabled E_Port OFF
Locked E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable OFF
LOS TOV enable OFF
NPIV capability ON
QOS E_Port OFF
Port Auto Disable: OFF
Rate Limit OFF
EX Port OFF
Mirror Port OFF
Credit Recovery ON
F_Port Buffers OFF
NPIV PP Limit: 176
CSCTL mode: OFF
Enabling and disabling NPIV
On the Brocade 300, 4100, 4900, 5000, 5100, 5300, and 8000 switches, the Brocade 5410,
5424, 5450, 5460, 5470, and 5480 embedded switches, the Brocade 48000 director, the
Brocade DCX and DCX-4S enterprise-class platforms, and the FA4-18 blade, NPIV is enabled for
every port.
```

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

2. To enable or disable NPIV on a port, enter the **portCfgNPIVPort** command with either the --enable or –disable option. The following example shows NPIV being enabled on port 10 of a Brocade 5100:

```
switch:admin> portCfgNPIVPort –enable 10
```

Viewing NPIV port configuration information

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **portCfgShow** command to view the switch ports information.

The following example shows whether a port is configured for NPIV:

```
switch:admin> portcfgshow
Ports of Slot 0 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
-----------------+--+--+--+--+--+---+--+--+--+--+---+--+--+--+--+---+--+--+--
Speed AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN
Trunk Port ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON
Long Distance .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
VC Link Init .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Locked L_Port .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Locked G_Port .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Disabled E_Port .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
ISL R_RDY Mode .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
RSCN Suppressed .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Persistent Disable.. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
NPIV capability ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON
```

3. Use the switchShow and portShow commands to view NPIV information for a given port. If a port is an F_Port, and you enter the switchShow command, then the port WWN of the N_Port is returned. For an NPIV F_Port, there are multiple N_Ports, each with a different port WWN. The switchShow command output indicates whether or not a port is an NPIV F_Port, and identifies the number of virtual N_Ports behind it. Following is sample output from the switchShow command:

```
switch:admin> switchshow
switchName:switch
switchType:66.1
switchState:Online
switchMode:Native
switchRole:Principal
switchDomain:1
switchId:fffc01
switchWwn:10:00:00:05:1e:82:3c:2a
zoning:OFF
switchBeacon:OFF
FC Router:OFF
FC Router BB Fabric ID:128
Area Port Media Speed State Proto
=====================================
0 0 id N1 Online F-Port 1 Nport + 1 NPIV devices.
1 1 id N4 No_Light
2 2 id N4 Online F-Port 20:0e:00:05:1e:0a:16:59
3 3 id N4 No_Light
4 4 id N4 No_Light
...
```

<output truncated>

4. Use the portShow command to view the NPIV attributes and all the N_Port (physical and virtual) port WWNs that are listed under *portWwn of device(s) connected*. Following is sample output for the portShow command:

```
switch:admin> portshow 2
portName: 02
portHealth: HEALTHY
Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x24b03 PRESENT ACTIVE F_PORT G_PORT NPIV LOGICAL_ONLINE LOGIN
NOELP LED ACCEPT
portType: 10.0

portState: 1Online
portPhys: 6In_Sync
portScn: 32F_Port
port generation number: 148
portId: 630200
portIfId: 43020005
portWwn: 20:02:00:05:1e:35:37:40
portWwn of device(s) connected:
c0:50:76:ff:fb:00:16:fc
c0:50:76:ff:fb:00:16:f8
...
<output truncated>
...
c0:50:76:ff:fb:00:16:80
50:05:07:64:01:a0:73:b8
Distance: normal
portSpeed: N2Gbps
Interrupts: 0 Link_failure: 16 Frjt: 0
Unknown: 0 Loss_of_sync: 422 Fbsy: 0
Lli: 294803 Loss_of_sig: 808
Proc_rqrd: 0 Protocol_err: 0
Timed_out: 0 Invalid_word: 0
Rx_flushed: 0 Invalid_crc: 0
Tx_unavail: 0 Delim_err: 0
Free_buffer: 0 Address_err: 1458
Overrun: 0 Lr_in: 15
Suspended: 0 Lr_out: 17
Parity_err: 0 Ols_in: 16
2_parity_err: 0 Ols_out: 15
CMI_bus_err: 0
```

Viewing virtual PID login information

Use the portLoginShow command to display the login information for the virtual PIDs of a port.

Following is sample output from the portLoginShow command:

```
switch:admin> portloginshow 2
Type PID World Wide Name credit df_sz cos
======================================================
fe 630240 c0:50:76:ff:fb:00:16:fc 101 2048 c scr=3
fe 63023f c0:50:76:ff:fb:00:16:f8 101 2048 c scr=3
fe 63023e c0:50:76:ff:fb:00:17:ec 101 2048 c scr=3
```

```
...
<output truncated>
...
ff 630202 c0:50:76:ff:fb:00:17:70 192 2048 c d_id=FFFFFC
ff 630201 c0:50:76:ff:fb:00:16:80 192 2048 c d_id=FFFFFC
```

Chapter 8: Zoning

Overview

Zoning enables you to partition your storage area network (SAN) into logical groups of devices that can access each other. A device can communicate only with other devices connected to the fabric within its specified zone. Devices can belong to more than one zone. When using a mixed fabric—that is, a fabric containing two or more switches running different release levels of fabric operating systems—you should use the switch with the highest Fabric OS level to perform zoning tasks.

You can establish a zone by identifying zone objects using one or more of the following *zoning schemes*:

- Domain,index (D,I)
  All members are specified by *domain ID*, *port number*, or *domain, index number* pair or aliases.
- World Wide Name (WWN)
  All members are specified only by World Wide Name (WWNs) or aliases of WWNs. They can be node or port versions of the WWN.
- Mixed zoning
  A zone containing members specified by a combination of *domain,port* or *domain,index* or aliases, and WWNs or aliases of WWNs.

In any scheme, you can identify zone objects using aliases.

Zoning Configurations

A zone configuration is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is in effect, all zones that are members of that configuration are in effect. Several zone configurations can reside on a switch at once, and you can quickly alternate between them. However, only one zone configuration can be enabled at a time.

The different types of zone configurations are:

- Defined Configuration
  The complete set of all zone objects defined in the fabric.

- Effective Configuration
  A single zone configuration that is currently in effect. The effective configuration is built when you enable a specified zone configuration.

- Saved Configuration
  A copy of the defined configuration plus the name of the effective configuration, which is saved in flash memory. (You can also provide a backup of the zoning configuration and restore the zoning configuration.) There might be differences between the saved configuration and the defined configuration if you have modified any of the zone definitions and have not saved the configuration.

- Disabled Configuration
  The effective configuration is removed from flash memory.

When you disable the effective configuration, the Advanced Zoning feature is disabled on the fabric, and all devices within the fabric can communicate with all other devices (unless you previously set up a default zone). This does not mean that the zoning database is deleted, however, only that there is no configuration active in the fabric. On power-up, the switch automatically reloads the saved configuration. If a configuration was active when it was saved, the same configuration is reinstated on the local switch.

Zone aliases

Creating an aliases

A zone alias is a logical group of ports or WWNs. You can simplify the process of creating zones by first specifying aliases, which eliminates the need for long lists of individual zone member names.

Here are the steps:

1. Log in using a command with admin privileges.
2. Enter **aliCreate** using the following syntax
   alicreate "*aliasname*", "*member*[; member...]"
3. Enter **cfgSave** command to save the changes to defined configuration.

```
switch:admin> alicreate "array1", "2,32; 2,33; 2,34; 4,4"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Adding members to alias

1. Log in using a command with admin privileges.
2. Enter **aliAdd** using the following syntax
   aliadd "*aliasname*", "*member*[; member...]"
3. Enter **cfgSave** command to save the changes to defined configuration.

```
switch:admin> aliadd "array1", "1,2"
switch:admin> aliadd "loop1", "5,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Removing members from alias

1. Log in using a command with admin privileges.
2. Enter **aliAdd** using the following syntax
   aliadd "*aliasname*", "*member*[; member...]"
3. Enter **cfgSave** command to save the changes to defined configuration.

```
switch:admin> aliremove "array1", "1,2"
switch:admin> aliremove "loop1", "5,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Deleting an aliases

1. Log in using a command with admin privileges.
2. Enter **aliDelete** using the following syntax
   alidelete "*aliasname*"
3. Enter **cfgSave** command to save the changes to defined configuration.

```
switch:admin> alidelete "array1"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Viewing aliases in defined configuration

1. Connect to the switch and log in as admin.
2. Enter the **aliShow** command, using the following syntax
   alishow "*pattern*"[, *mode*]

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

The following example shows all zone aliases beginning with “arr”.

```
switch:admin> alishow "arr*"
alias: array1 21:00:00:20:37:0c:76:8c
alias: array2 21:00:00:20:37:0c:66:23
```

Creating a Zone

1. Connect to the switch and log in as admin.
2. Enter the **zoneCreate** command, using the following syntax:
   zonecreate "*zonename*", "*member*[; member...]"
3. To create a broadcast zone, use the reserved name “broadcast”.
4. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> zonecreate "greenzone", "2,32; 2,33; 2,34; 4,4"
switch:admin> zonecreate "bluezone", "21:00:00:20:37:0c:66:23; 4,3"
switch:admin> zonecreate "broadcast", "1,2; 2,33; 2,34"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Adding members to zone

1. Connect to the switch and log in as admin.
2. Enter the **zoneAdd** command, using the following syntax:
   zoneadd "*zonename*", "*member*[; member...]"
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> zoneadd "greenzone", "1,2"
switch:admin> zoneadd "bluezone", "21:00:00:20:37:0c:72:51"
switch:admin> zoneadd "broadcast", "1,3"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This action will only save the changes on the
Defined configuration.
Any changes made on the Effective configuration will not take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Removing members from zone

1. Connect to the switch and log in as admin.
2. Enter the **zoneRemove** command, using the following syntax:
   zoneremove "zonename", "member[; member...]"
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> zoneremove "greenzone", "1,2"
switch:admin> zoneremove "bluezone", "21:00:00:20:37:0c:72:51"
switch:admin> zoneremove "broadcast", "2,34"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Deleting a zone

1. Connect to the switch and log in as admin.
2. Enter the **zoneDelete** command, using the following syntax:
   zonedelete "*zonename*"
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> zonedelete "broadcast”
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Viewing zoning in defined configuration

1. Connect to the switch and log in as admin.
2. Enter the **zoneShow** command, using the following syntax:
   zoneshow[–sort] ["pattern"] [, mode]

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

The following example shows all zones beginning with A, B, or C, in ascending order:

```
switch:admin> zoneshow –sort "[A-C]*"
zone: Blue_zone 1,1; array1; 1,2; array2
zone: Bobs_zone 4,5; 4,6; 4,7; 4,8; 4,9
```

Default zoning mode

The default zoning mode controls device access if zoning is not implemented or if there is no effective zone configuration. The default zoning mode has two options:

- All Access—All devices within the fabric can communicate with all other devices.
- No Access—Devices in the fabric cannot access any other device in the fabric.

The default zone mode applies to the entire fabric, regardless of switch model.
The default setting is All Access.

1. Connect to the switch and log in as admin.
2. Enter the **cfgActvShow** command to view the current zone configuration.
3. Enter the **defZone** command with one of the following options:
   defzone –noaccess
   defzone –allaccess
4. Enter either the **cfgSave**, **cfgEnable**, or **cfgDisable** command to commit the change and distribute it to the fabric. The change will not be committed and distributed across the fabric if you do not enter one of these commands

```
switch:admin> defzone –noaccess
You are about to set the Default Zone access mode to No Access
Do you want to set the Default Zone access mode to No Access ? (yes, y, no, n):
```

```
[no] y
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
Updating flash ...
```

Zoning database size

To view the size of zoning database use **cfgSize** command.

```
switch:admin> cfgsize

Zone DB max size - 1045274 bytes
Available Zone DB size - 1030583 bytes

committed - 13679
transaction - 0
```

Zoning Configurations

Creating zoning configurations

You can store a number of zones in a zoning configuration database. When enabling a new zone configuration, ensure that the size of the defined configuration does not exceed the maximum configuration size supported by all switches in the fabric. This is particularly important if you downgrade to a Fabric OS version that supports a smaller zone database than the current Fabric OS. In this scenario, the zone database in the current Fabric OS would have to be changed to the smaller zone database before the downgrade. You can use the **cfgSize** command to check both the maximum available size and the currently saved size on all switches. The **cfgSize** command reports the maximum available size on the current switch only. It cannot determine the maximum available size on other switches in the fabric. The minimum zoning database size is 4 bytes, even if the zoning database is empty.

1. Connect to the switch and log in as admin.
2. Enter the **cfgCreate** command, using the following syntax:
   cfgcreate "*cfgname*", "*member*[; member...]"
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> cfgcreate "NEW_cfg", "purplezone; bluezone; greenzone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Adding zones to zoning configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgAdd** command, using the following syntax:
   cfgadd "*cfgname*", "*member*[; member...]"
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> cfgadd "newcfg", "bluezone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Removing zones from zoning configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgRemove** command, using the following syntax:
   cfgadd "*cfgname*", "*member*[; member...]"
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> cfgremove "newcfg", "bluezone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Enable a zoning configuration

The following procedure ends and commits the current zoning transaction buffer to nonvolatile memory. If a transaction is open on a different switch in the fabric when this procedure is run, the transaction on the other switch is automatically aborted. A message displays on the other switches to indicate that the transaction was aborted.

1. Connect to the switch and log in as admin.
2. Enter the **cfgenable** command, using the following syntax:
   cfgenable "*cfgname"*
3. Enter **y** at the prompt.

```
switch:admin> cfgenable "USA_cfg"

You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the current configuration selected. If the
update includes changes to one or more traffic isolation zones, the update may result in localized
disruption to traffic on ports associated with the traffic isolation zone changes.

Do you want to enable 'USA_cfg' configuration (yes, y, no, n): [no] y
```

```
zone config "USA_cfg" is in effect

Updating flash ...
```

Disabling a zoning configuration

When you disable the current zone configuration, the fabric returns to non-zoning mode. All devices can then access each other or not, depending on the default zone access mode setting.

1. Connect to the switch and log in as admin.
2. Enter the **cfgdisable** command, using the following syntax:
   cfgdisable
3. Enter **y** at the prompt.

```
switch:admin> cfgdisable

You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the current configuration selected. If the
update includes changes to one or more traffic isolation zones, the update may result in localized
disruption to traffic on ports associated with the traffic isolation zone changes.

Do you want to disable zoning configuration (yes, y, no, n): [no] y
```

Deleting a zone configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgDelete** command, using the following syntax:
   cfgdelete "*cfgname*"
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> cfgdelete "testcfg"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This action will only save the changes on the
Defined configuration. Any changes made on the Effective configuration will not take effect until it is
re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Clearing changes to configuration

1. Enter the **cfgTransAbort** command.

When this command is executed, all changes since the last save operation (performed with the cfgSave, cfgEnable, or cfgDisable command) are cleared.

In the following example, assume that the removal of a member from zone1 was done in error:

```
switch:admin> zoneremove "zone1","3,5"
```

```
switch:admin> cfgtransabort
```

Viewing all zone configuration information

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command with no operands.

```
switch:admin> cfgshow

Defined configuration:
cfg: USA1 Blue_zone
cfg: USA_cfg Purple_zone; Blue_zone
zone: Blue_zone
1,1; array1; 1,2; array2
zone: Purple_zone
1,0; loop1
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
Effective configuration:
cfg: USA_cfg
zone: Blue_zone
1,1
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2
21:00:00:20:37:0c:76:22
21:00:00:20:37:0c:76:28
zone: Purple_zone
1,0
21:00:00:20:37:0c:76:85
21:00:00:20:37:0c:71:df
```

Viewing selected zone configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command and specify a pattern.
   cfgshow "*pattern*"[, *mode*]

The following example displays all zone configurations that start with “Test”:

```
switch:admin> cfgshow "Test*"
cfg: Test1 Blue_zone
cfg: Test_cfg Purple_zone; Blue_zone
```

Viewing configuration in effective zone database

1. Connect to the switch and log in as admin.

2. Enter the **cfgActvShow** command.

```
switch:admin> cfgactvshow
Effective configuration:
cfg: NEW_cfg
zone: Blue_zone
1,1
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2
21:00:00:20:37:0c:76:22
21:00:00:20:37:0c:76:28
zone: Purple_zone
1,0
21:00:00:20:37:0c:76:85
21:00:00:20:37:0c:71:df
```

Clearing all zone configurations

1. Connect to the switch and log in as admin.
2. Enter the **cfgClear** command to clear all zone information in the transaction buffer.

ATTENTION
Be careful using the cfgClear command because it deletes the defined configuration.

```
switch:admin> cfgclear
The Clear All action will clear all Aliases, Zones, FA Zones and configurations in the Defined
configuration.
cfgSave may be run to close the transaction or cfgTransAbort may be run to cancel the transaction.
Do you really want to clear all configurations? (yes, y, no, n): [no]
```

3. Enter one of the following commands, depending on whether an effective zoning configuration exists:

- If no effective zoning configuration exists, enter the **cfgSave** command.
- If an effective zoning configuration exists, enter the **cfgDisable** command to disable and clear the zone configuration in nonvolatile memory for all switches in the fabric.

Zone object maintenance

The following procedures describe how to copy, delete, and rename zone objects. Depending on the operation, a zone object can be a zone member, a zone alias, a zone, or a zone configuration.

Copying a zone object

When you copy a zone object, the resulting object has the same name as the original. The zone object can be a zone configuration, a zone alias, or a zone.

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command to view the zone configuration objects you want to copy.
   cfgshow "*pattern*"[, *mode*]

For example, to display all zone configuration objects that start with “Test”:

```
switch:admin> cfgshow "Test*"
cfg: Test1 Blue_zone
cfg: Test_cfg Purple_zone; Blue_zone
```

3. Enter the **zone –copy** command, specifying the zone objects you want to copy, along with the new object name. Note that zone configuration names are case-sensitive; blank spaces are ignored and it works in any Admin Domain other than AD255.

```
switch:admin> zone –copy Test1 US_Test1
```

4. Enter the **cfgShow** command to verify the new zone object is present.

```
switch:admin> cfgshow "Test*"
cfg: Test1 Blue_zone
cfg: Test_cfg Purple_zone; Blue_zone
switch:admin> cfgShow "US_Test1"
cfg: US_Test1
Blue_zone
```

5. If you want the change preserved when the switch reboots, enter the **cfgSave** command to save it to nonvolatile (flash) memory.
6. Enter the **cfgEnable** command for the appropriate zone configuration to make the change effective

Deleting a zone object

The following procedure removes all references to a zone object and then deletes the zone object. The zone object can be a zone member, a zone alias, or a zone.

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command to view the zone configuration objects you want to delete.

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Purple_zone; White_zone; Blue_zone
zone: Blue_zone
1,1; array1; 1,2; array2
zone: Purple_zone
1,0; loop1
zone: White_zone
1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
Effective configuration:
cfg: USA_cfg
zone: Blue_zone
1,1
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2
21:00:00:20:37:0c:76:22
21:00:00:20:37:0c:76:28
```

```
zone: Purple_zone
1,0
21:00:00:20:37:0c:76:85
21:00:00:20:37:0c:71:df
```

3. Enter **the zone –expunge** command to delete the zone object. Zone configuration names are case-sensitive; blank spaces are ignored and it works in any Admin Domain other than AD255.

```
switch:admin> zone –expunge "White_zone"
You are about to expunge one configuration or member. This action could result in removing many
zoning configurations recursively.
[Removing the last member of a configuration removes the configuration.]
Do you want to expunge the member? (yes, y, no, n): [no] yes
```

4. Enter **yes** at the prompt.
5. Enter the **cfgShow** command to verify the deleted zone object is no longer present.
6. If you want the change preserved when the switch reboots, enter the **cfgSave** command to save it to nonvolatile (flash) memory.
7. Enter the **cfgEnable** command for the appropriate zone configuration to make the change effective.

Renaming a zone object

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command to view the zone configuration objects you want to rename.

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Purple_zone; White_zone; Blue_zone
zone: Blue_zone
1,1; array1; 1,2; array2
zone: Purple_zone
1,0; loop1
zone: White_zone
1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Enter the **zoneObjectRename** command to rename zone configuration objects. Note that zone configuration names are case-sensitive; blank spaces are ignored and it works in any Admin Domain other than AD255.

```
switch:admin> zoneObjectRename "White_zone", "Purple_zone"
```

4. Enter the **cfgShow** command to verify the renamed zone object is present.
5. If you want the change preserved when the switch reboots, enter the **cfgSave** command to save it to nonvolatile (flash) memory.
6. Enter the **cfgEnable** command for the appropriate zone configuration to make the change effective.

Zoning configuration management

You can add, delete, or remove individual elements in an existing zone configuration to create a appropriate configuration for your SAN environment. After the changes have been made, save the configuration to ensure the configuration is permanently saved in the switch and that the configuration is replicated throughout the fabric. The switch configuration file can also be uploaded to the host for archiving and it can be downloaded from the host to a switch in the fabric using configUpload and configDownload commands respectively. Refer the Fabric *OS Command Reference* for additional information on uploading and downloading the configuration file.

Zone merging

When a new switch is added to the fabric, it automatically takes on the zone configuration information from the fabric. If you are adding a switch that is already configured for zoning, clear the zone configuration on that switch before connecting it to the zoned fabric. Adding a new fabric that has no zone configuration information to an existing fabric is very similar to adding a new switch. All switches in the new fabric inherit the zoning configuration data. If a zone configuration is in effect, then the same configuration becomes the effective configuration for the new switches.

For more on zone merging please refer to Fabric OS Administrator's Guide.

Web Tools

Zoning configurations

The Zone Administration window is where all of the zoning tasks are performed.
When performing zoning tasks for switches in a mixed fabric—that is, a fabric containing two or more switches running different fabric operating systems—you should use the switch with the highest Fabric OS level.

Opening the Zone Administration window

You cannot open the Zone Administration window from AD255 (physical fabric).

1. Select a switch from the Fabric Tree.
2. Click Zone Admin in the Manage section of the Tasks menu.

The Zone Administration window opens.

Setting the default zoning mode

The default zoning mode has two options:

- All Access—All devices within the fabric can communicate with all other devices.
- No Access—Devices in the fabric cannot access any other device in the fabric.

Web Tools supports default zoning on switches running firmware v5.1.0 or later. Default zoning on legacy switches (switches running firmware versions prior to v 5.1.0) are not supported. Legacy switches can use default zoning; however, they cannot manipulate the default zone or default configuration.

NOTE
To use Admin Domains, you must set the default zoning mode to No Access prior to setting up the Admin Domains. To use the Admin Domain feature, the EGM license must be enabled on the switch; otherwise access to this feature is denied. You cannot change the default zoning mode to All Access
if user-specified Admin Domains are present in the fabric.

1. Open the Zone Administration window.
2. Click Zoning Actions > Set Default Mode, and then select the access mode.

Zoning management

You can monitor and manage basic and traffic isolation zoning through the Web Tools Zone Administration. The information in the Zone Administration window is collected from the selected switch.

If the FCS policy is activated in the fabric, zoning can be administered only from the primary FCS switch. If the selected switch has an Advanced Zoning license installed, but is not the primary FCS switch, the Zone Admin option is displayed, but not activated.
You must be logged into the switch using a user name with one of the following roles associated with it to make changes to the zoning: zoneAdmin, admin, or fabricAdmin. All other roles allow only a view or read-only access. Most of the zoning operations are disabled in read-only mode.
A snapshot is taken of all the zoning configurations at the time you launch the Zone Administration window; this information *is not updated automatically* by Web Tools.

When you log in to a virtual switch, or select a virtual switch using the drop down list under Fabric
Tree section in the Switch Explorer window, only the ports that are associated with the Virtual Fabric ID you selected are displayed in the member selection list. You can use the Add Other button to add ports of other switches in the fabric.

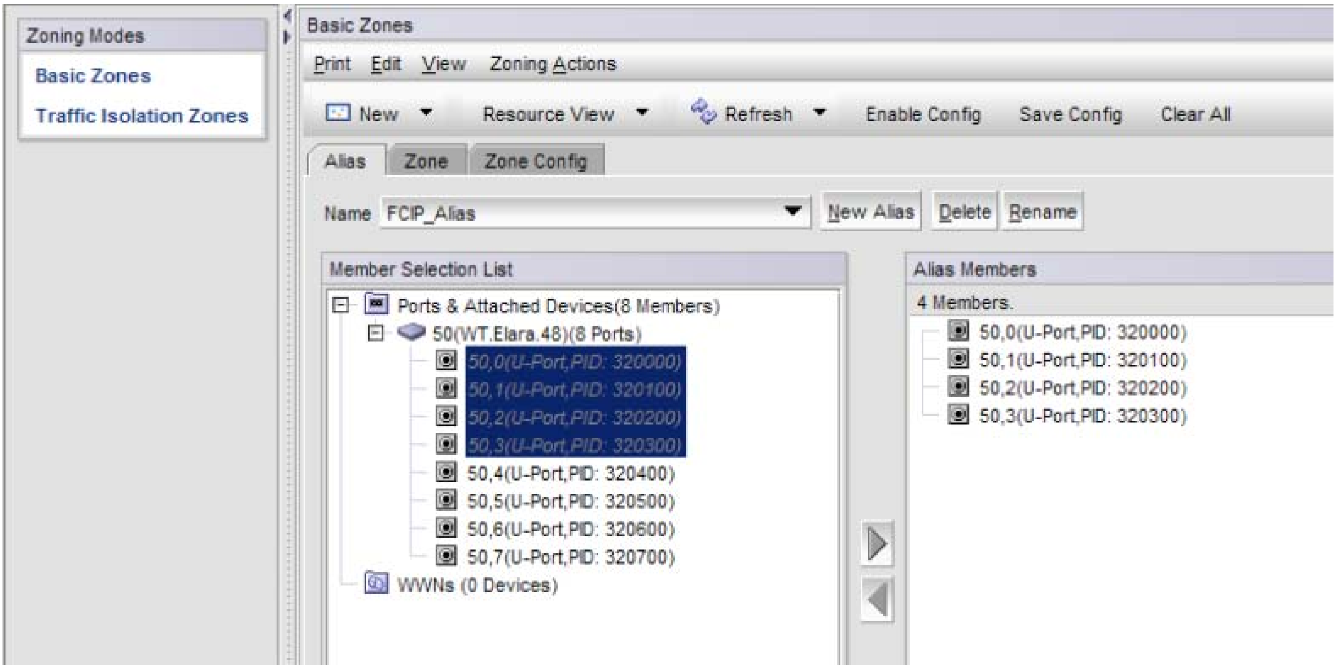

Note the following:

- “Saving” means updating the zoning database on the switch with the local changes from the Web Tools buffer.
- “Refreshing” means copying the current state of the zoning database on the switch to the Web Tools buffer, overwriting its current contents.

In the Zone Administration window, all WWNs also display vendor names.

NOTE
The Member Selection List only lists the ports of the current switch and the devices of all the switches in the fabric. Slot and port information of other switches are not displayed in the tree.
You can click the Alias tab to display which aliases the port or device is a member of. Also, you can right-click the device nodes and click View Device Detail to display detailed information about the selected device.

The Member Selection List panel displays only physical FC ports. To verify whether you have any unzoned devices, you must use Brocade Network Advisor to analyze zone configurations.

The Member Selection List displays virtual initiators if the chassis has an FC4-16IP blade in it; they are shown under a separate folder icon called Virtual Initiators. If the chassis has the Brocade
7500E Extended Switch license installed, the tree displays only two FC ports; otherwise all logical ports display if you have the 7500 without the extended license installed.
Admin Domain considerations: The Member Selection List panel displays a filtered list of ports that are as follows:

- Direct port members are zoneable and are displayed in the tree.
- Indirect port members to which owned devices are attached are displayed in the tree, but cannot be added to a zone or alias.
- Direct device members are zoneable and are displayed in the tree.
- Indirect device members (devices that are currently attached to owned ports) are also zoneable and displayed in the tree. But if such a device is later moved to a non-owned port it will no longer be displayed or zoneable.
- Switches and blades are displayed only if they contain owned ports or devices, regardless of switch ownership, such as the FS8-18 Encryption blade.
- Ports that are indirect members only because the switch is owned are not displayed.
- When no user- defined Admin Domains are present on the switch, AD0 shows the port count. If there are user-defined Admin Domains, AD0 does not show port count and the user-defined AD shows port count.

Refreshing fabric information

This function refreshes the display of *fabric elements only* (switches, ports, and devices). It does not affect any zoning element changes or update zone information in the Zone Administration window. You can refresh the fabric element information displayed at any time.

1. In the Zone Administration window, click View > Refresh From Live Fabric.

This refreshes the status for the fabric, including switches, ports, and devices.

Saving local zoning changes

All information displayed and all changes made in the Zone Administration window are buffered until you save the changes. In that case any other user looking at the zone information for the switch will not see the changes you have made until you save them.
Saving the changes propagates any changes made in the Zone Administration window (buffered changes) to the zoning database on the switch. If another user has a zoning operation in progress at the time that you attempt to save changes, a warning is displayed that indicates that another zoning transaction is in progress on the fabric. You can select to abort the other transaction and override it with yours.
If the zoning database size exceeds the maximum allowed, you cannot save the changes. The zoning database summary displays the maximum zoning database size.
This action updates the entire contents of the Zone Administration window, not just the selected zone, alias, or configuration. You can save your changes at any time during the Zone Administration session.

1. Make the zoning changes in the Zone Administration window.
2. Click Zoning Actions > Save Config.

Select a zoning view

You can choose how zoning elements are displayed in the Zone Administration window. The zoning view you select determines how members are displayed in the Member Selection List panel. The views filter the fabric and device information displayed in the Member Selection
List for the selected view, making it easier for you to create and modify zones, especially when creating "hard zones."
Depending on the method you use to zone, certain tabs might or might not be available in the Zone Administration window.
There are two views of defining members for zoning:

- Fabric View—Displays the physical hierarchy of the fabric, a list of the attached and imported physical devices (by WWN), and a list of the FC Virtual Initiators on switches that support iSCSI.
  In the Fabric View, you can select ports for port-based zoning or devices for WWN-based zoning.
- Devices Only—Displays a list of the attached and imported physical devices by WWN. You cannot select ports for port-based or mixed zoning schemes, nor can you select virtual initiators for iSCSI FC Zone creation.

Use the following procedure to define the way you want to view the fabric resource.

1. Launch the Zone Administration window.
2. Click View > Choose Fabric Resources View.
3. Choose the way you want to view the fabric resource and click OK.

Creating and populating zone aliases

An alias is a logical group of port index numbers and WWNs. Specifying groups of ports or devices as an alias makes zone configuration easier, by enabling you to configure zones using an alias rather than inputting a long string of individual members. You can specify members of an alias using the following methods:

- Identifying members by switch domain and port index number pair, for example, 2, 20.
- Identifying members by device node and device port WWNs.

Use the following procedure to create a zone alias.

1. Open the Zone Administration window.
2. Select a format to display zoning members in Member Selection List
3. Click the Alias tab and click New Alias.
   The Create New Alias dialog box displays.
4. On Create New Alias, type a name for the new alias and click OK.
   The new alias is displayed in the Name list.
5. Expand the Member Selection List to view the nested elements.
   The choices available in the Member Selection List depend on the selection in the View menu.
6. Click elements in the Member Selection List that you want to include in the alias.
   The Add Member button becomes active.
7. Click Add Member to add alias members.
   Selected members move to the Alias Members window.
8. *Optional*: Repeat steps 6 and 7 to add more elements to the alias.
9. *Optional*: Click Add Other to include a WWN or port that is not currently a part of the fabric.
10. Click Actions > Save Config to save the configuration changes.

Adding and removing members of a zone alias

Use the following procedure to add or remove zone alias members.

1. Open the Zone Administration window.
2. Click the Alias tab.
3. Select the alias you want to modify from the Name list.
4. Select an element in the Member Selection List that you want to add to the alias, or select an element in the Alias Members list that you want to remove.
5. Click Add Member to add the selected alias member, or click Remove Member to remove the selected alias member.
   The alias is modified in the Zone Admin buffer. At this point you can either save your changes or save and enable your changes.
6. Click Zoning Actions > Save Config to save the configuration changes.

Renaming zone aliases

The new alias name cannot exceed 64 characters and can contain alphabetic, numeric, and underscore characters. Use the following procedure to change the name of a zone alias.

1. Open the Zone Administration window.
2. Click the Alias tab and select the alias you want to rename from the Name list.
3. Click Rename.

The Rename an Alias dialog box displays.

4. Type a new alias name and click OK.

   The alias is renamed in the Zone Admin buffer. At this point you can either save your changes or save and enable your changes.

5. Click Zoning Actions > Save Config to save the configuration changes.

Deleting zone aliases

You can remove a zone alias from the Zone Admin buffer. When a zone alias is deleted, it is no longer a member of the zones of which it was once a member.

NOTE
If you delete the only member zone alias, an error message is issued when you attempt to save the configuration.

1. Open the Zone Administration window.
2. Click the Alias tab.
3. Select the alias you want to delete from the Name list.
4. Click Delete.

   The Confirm Deleting Alias dialog box opens.

5. Click Yes.

   The selected alias is deleted from the Zone Admin buffer. At this point you can either save your changes or save and enable your changes.

6. Click Zoning Action > Save Config to save the configuration changes.

Creating and populating zones

A zone is a region within the fabric where specified switches and devices can communicate. A device can communicate only with other devices connected to the fabric within its specified zone.
Use the following procedure to create a zone.

1. Open the Zone Administration window.

   Select a format to display zoning members in the Member Selection List.

2. Click the Zone tab.
3. Click New Zone.

   The Create New Zone dialog box displays.

4. On Create New Zone, enter a name for the new zone, and click OK.

   LSAN zones and QoS zones have specific naming requirements:

   The new zone displays in the Name list.

5. Expand the Member Selection List to view the nested elements.

   The choices available in the list depend on the selection made in the View menu.

6. Select an element in the Member Selection List that you want to include in your zone. Note that

   LSAN zones should contain only port WWN members.

   The Add Member button becomes active.

7. Click Add Member to add the zone member.

   The selected member is moved to the Zone Members window.

8. *Optional*: Repeat steps 7 and 8 to add more elements to your zone.
9. *Optional*: Click Add Other to include a WWN or port that is not currently a part of the fabric.

   At this point you can either save your changes or save and enable your changes.

10. Click Zoning Actions > Save Config to save the configuration changes.

Adding and removing members of a zone

Use the following procedure to add or remove zone members.

1. Open the Zone Administration window.
2. Click the Zone tab.
3. Select the zone you want to modify from the Name list.

The zone members for the selected zone are listed in the Zone Members list.

4. Highlight an element in the Member Selection List that you want to include in your zone, or highlight an element in the Zone Members list that you want to delete.
5. Click Add Member to add a zone member, or click Remove Member to remove a zone member.

   The zone is modified in the Zone Admin buffer. At this point you can either save your changes or save and enable your changes.
6. Click Zoning Actions > Save Config to save the configuration changes.

Renaming zones

Use the following procedure to change the name of a zone.

1. Open the Zone Administration window.
2. Click the Zone tab.
3. Elect the zone you want to rename from the Name list.
4. Click Rename.
5. On Rename a Zone, type a new zone name and click OK.

   The zone is renamed in the Zone Admin buffer. At this point you can either save your changes or save and enable your changes.
6. Click Zoning Actions > Save Config to save the configuration changes.

Cloning zones

NOTE
To perform clone operations for zoning, the EGM license must be installed on the switch; otherwise, access to this feature is denied and an error message displays.

The EGM license is required only for 8 Gbps platforms, such as the Brocade DCX and DCX-4S enterprise-class platforms, the Encryption Switch, the 300, 5300, and 5100 switches. For non-8 Gbps platforms, all functionalities are available without EGM license.
Use the following procedure to clone a zone configuration.

1. Open the Zone Administration window.
2. Click the Zone tab.
3. Select the zone you want to clone from the Name list.
4. Click Clone
5. On Clone an Existing Zone, enter a name for the copied zone.
6. Click OK.

   The selected zone is copied from the Zone Admin buffer.
7. Click Zoning Actions > Save Config to save the configuration changes.

Since no changes were made to the effective configuration, you do not need to enable the configuration.

Deleting zones

Use the following procedure to delete a zone.

1. Open the Zone Administration window.
2. Click the Zone tab.
3. Select the zone you want to delete from the Name menu and click Delete.
4. On the confirmation dialog box, click Yes.
5. The selected zone is deleted from the Zone Admin buffer. At this point you can either save your changes or save and enable your changes.
6. Click Zoning Actions > Save Config to save the configuration changes.

Creating zone configurations

Use the following procedure to create a zone configuration. After creating a zone configuration, you must explicitly enable it for it to take effect.

NOTE
Any changes made to the currently enabled configuration does not appear until you re-enable the configuration.

1. Open the Zone Administration window.
2. Select a format to display zoning members in the Member Selection List
3. Click the Zone Config tab and click New Zone Config.
4. On Create New Config, type a name for the new configuration and click OK.
   The new configuration displays in the Name list.
5. Expand the Member Selection List to view the nested elements.
   The choices available in the list depend on the selection made in the View menu.
6. Select an element in the Member Selection List that you want to include in your configuration.
   The Add Member button becomes active.
7. Click Add Member to add configuration members.
   Selected members are moved to the Config Members Window.
8. Repeat steps 6 and 7 to add more elements to your configuration.
9. Click Zoning Actions > Save Config to save the configuration changes.

Adding or removing zone configuration members

Use the following procedure to add or remove members of a zone configuration.

NOTE
You can make changes to a configuration that is currently enabled; however, changes do not appear until you re-enable the configuration.

1. Open the Zone Administration window.
2. Click the Zone Config tab.
3. Select the configuration you want to modify from the Name list.
4. Click an element in the Member Selection List that you want to include in your configuration or
   click an element in the Config Members that you want to delete.
5. Click Add Member to add a configuration member or Remove Member to remove a
   configuration member.
6. Click Zoning Actions > Save Config to save the configuration changes.

Renaming zone configurations

The new name cannot exceed 64 characters and can contain alphabetic, numeric, and underscore characters. Use the following procedure to change the name of a zone configuration.

NOTE
You cannot rename the currently enabled configuration.

1. Open the Zone Administration window.
2. Click the Zone Config tab.
3. Select the configuration you want to rename from the Name list and click Rename.
4. On Rename a Config, type a new configuration name and click OK.
   The configuration is renamed in the configuration database.
5. Click Zoning Actions > Save Config to save the configuration changes.

Cloning zone configurations

You must use Web Tools with the EGM license to perform cloning operations for zone configurations; otherwise, access to this feature is denied and an error message displays.

Use the following procedure to clone a zone configuration.

1. Open the Zone Administration window.
2. Click the Zone Config tab.
3. Select the zone configuration you want to clone from the Name list.
4. Click Clone.
5. On Copy An Existing Zone Config, enter a name for the copied zone and click OK.

   The selected zone is copied from the Zone Admin buffer.
6. Click Zoning Actions > Save Config to save the configuration changes.

   No changes were made to the effective configuration. You do not need to enable the configuration.

Deleting zone configurations

Use the following procedure to delete a zone configuration.

NOTE
You cannot delete a enabled configuration.

1. Open the Zone Administration window.
2. Click the Zone Config tab.
3. Select the configuration you want to delete from the Name list and click Delete.
4. On the confirmation dialog box, click Yes.

   The selected configuration is deleted from the configuration database.
5. Click Zoning Actions > Save Config to save the configuration changes.

Enabling zone configurations

Several zone configurations can reside on a switch at the same time, 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.

When you enable a zone configuration from Web Tools, the entire zoning database is automatically saved, and then the selected zone configuration is enabled.

If the zoning database size exceeds the maximum allowed, you cannot enable the zone configuration. The zoning database summary displays the maximum zoning database size.

1. Open the Zone Administration window.
2. Click Zoning Actions > Enable Config.
3. On Enable Config, select the configuration to be enabled from the menu.
4. Click OK to save and enable the selected configuration.

Disabling zone configurations

When you disable the active configuration, the Advanced Zoning feature is disabled on the fabric, and all devices within the fabric can communicate with all other devices. This does not mean that the zoning database is deleted, however, only that there is no configuration active on the fabric.

When you disable a zone configuration from Web Tools, keep in mind that the entire zoning database is automatically saved, and then the selected zone configuration is disabled.

NOTE
When you disable the active configuration, Advanced Zoning is disabled on the fabric, and according to the default zone set, devices within the fabric can or cannot communicate with other devices.

1. Open the Zone Administration window.
2. Click Zoning Actions > Disable Zoning.

   The Disable Config warning message displays.
3. Click Yes to save and disable the current configuration.

Displaying enabled zone configurations

The enabled zone configuration screen displays the actual content of the single zone configuration that is currently enabled on the fabric, whether it matches the configuration that was enabled when the current Zone Administration session was launched or last refreshed. The zones are displayed, and their contents (ports, WWNs) are displayed next to them. Aliases are not displayed in the enabled zone configuration. If there is no active zone configuration enabled on the switch, a message is displayed to that effect.

NOTE
The enabled configuration is listed in the lower-right corner of the Zone Administration window.

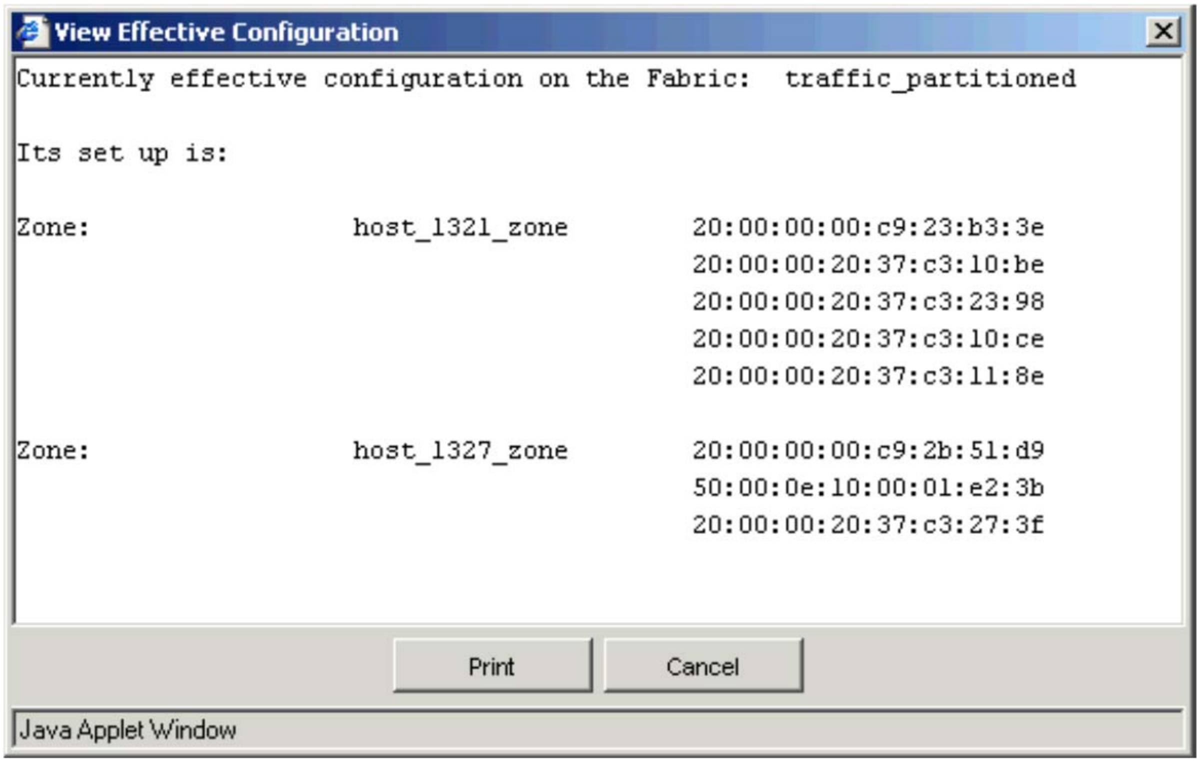

```
Currently effective configuration on the Fabric:  traffic_partitioned

Its set up is:

Zone:                host_1321_zone          20:00:00:00:c9:23:b3:3e
                                             20:00:00:20:37:c3:10:be
                                             20:00:00:20:37:c3:23:98
                                             20:00:00:20:37:c3:10:ce
                                             20:00:00:20:37:c3:11:8e

Zone:                host_1327_zone          20:00:00:00:c9:2b:51:d9
                                             50:00:0e:10:00:01:e2:3b
                                             20:00:00:20:37:c3:27:3f
```

Adding a WWN to multiple aliases and zones

This procedure enables you to configure a WWN as a member in a zone configuration prior to adding that device to the fabric. Specifically, it is useful if you want to add a WWN to all or most zoning entities. The added WWN does not need to currently exist in the fabric.

1. Open the Zone Administration window.
2. Click Edit > Add WWN.
   The Add WWN dialog box opens.
3. Type a WWN value in the WWN field and click OK.
   The Add WWN dialog box displays all the zoning elements that will include the new WWNs. All of the elements are selected by default.
4. Click items in the list to select or unselect, and click Add to add the new WWN to all the selected zoning elements.
   The WWN is added to the Zone Admin buffer and can be used as a member.

Removing a WWN from multiple aliases and zones

Use this procedure if you want to remove a WWN from all or most zoning entities.

1. Open the Zone Administration window.
2. Click Edit > Delete WWN.

The Delete WWN dialog box opens.

3. Type a WWN value in the WWN field and click OK.

   The Delete WWN dialog box displays all the zoning elements that include the WWN.

4. Click items in the list to select or unselect, and click Delete to delete the WWN from all the selected zoning elements.

   The WWN is deleted from the selected items in the Zone Admin buffer

Replacing a WWN in Multiple Aliases and Zones

This procedure enables you to replace a WWN throughout the Zone Admin buffer. This is helpful when exchanging devices in your fabric and helps you to maintain your current configuration.

1. Launch the Zone Administration window.
2. Click Edit > Replace WWN.

   The Replace WWN dialog box opens.

3. Type the WWN to be replaced in the Replace field.
4. Type the new WWN in the By field and click OK.

   The Replace WWN dialog box is displayed. It lists all the zoning elements that include the WWN.

5. Click an item in the list to select or unselect, and click Replace to replace the WWN in all the selected zoning elements.

   The former WWN is replaced in the Zone Admin buffer by the new WWN, including within any alias or zone in which the old WWN was a member.

Searching for zone members

You can search zone member selection lists for specified strings of text. If you know some identifying information about a possible member of a zoning entity, you can select the tab and view for that entity and then search through its member selection list using the Search for Zone Member option. If the target entity is an alias or zone, then the search domain includes elements like switch names and domain numbers, port names and “domain, port” addresses, device WWNs and manufacturer names, and also any aliases that might already have been defined. If the target entity is a configuration, then zones are also included, along with the elements they contain.

The search starts from the top of the list, and when the target element is found, it is also selected in the Member Selection List so it can be added or its parent or children can be found. By default, the Member Selection List is searched from beginning to end one time. If you select the wraparound option, the search continues to loop from the beginning to the end of the Member Selection List.

1. Open the Zone Administration window.
2. Click Edit > Search Member.
3. Type the zone member name in the Member Name field.

   Optional: Narrow the search by selecting one or more of the check boxes, such as Match Case.

4. Click Next to begin the zone member search

Clearing the Zoning Database

Use the following procedure to disable the active zoning configuration, if one exists, and delete the entire zoning database. You must disable any active configuration before you can delete the zoning database.

ATTENTION

This action not only disables zoning on the fabric, but also deletes the entire zoning database. This results in all devices being able to communicate with each other.

1. Open the Zone Administration window.
2. Click Actions > Clear All.

   The Disable Config warning opens.

3. Click Yes to do *all* of the following:

   - Disable the current configuration.
   - Clear the entire contents of the current Web Tools Zone Admin buffer.

- Delete the entire persistent contents of the fabric zoning database.

The wizard allows you to define one and only one name for each device port (WWN). Devices with one or more aliases are considered already named and are not displayed.

Zone configuration analysis

You must use Brocade Network Advisor to analyze the following zone configurations:

- Add unzoned devices
- Remove offline or inaccessible devices
- Replace offline devices
- Define device alias

Best practices for zoning

The following are recommendations for using zoning:

- Always zone using the highest Fabric OS-level switch.
  Switches with lower Fabric OS versions do not have the capability to view all the functionality that a newer Fabric OS provides as functionality is backwards compatible but not forwards compatible.
- Zone using the core switch versus an edge switch.
- Zone using a director over a switch.
  A director has more resources to handle zoning changes and implementations.
- Zone on the switch you connect to when bringing up Web Tools (the proxy switch).

Dynamic Fabric Provisioning: Fabric Assigned World Wide Name

In order to simplify and accelerate server deployment and improve operational efficiency, FOS v7.0.1 provides Fabric Assigned WWN or FA-PWWN capability. This feature allows users to create a virtual WWN for a server instead of using the server's physical port WWN (PWWN) to create zoning and LUN mapping/masking. When a FA-PWWN capable server is attached to the SAN, this feature allows the fabric to assign this virtual WWN to that server. This feature requires servers to be using Brocade HBAs/Adapters. Please consult Brocade HBA/Adapter driver documentation and Release Notes to confirm minimum requirements for this feature. For Brocade Network Advisor support, please consult Brocade Network Advisor documentation and Release Notes.

Chapter 9: Routing and Trunking

Routing Traffic

In the following section we will see routing related configurations.

Inter-Switch Links (ISLs)

Inter-Switch Link (ISL) Trunking optimizes network performance by forming trunking groups that can distribute traffic between switches across a shared bandwidth.. When connecting two switches together, you need to verify that the following parameters are different:

- Domain ID
- Switch name
- Chassis name

You must also verify the following fabric parameters are identical on each switch for a fabric to merge:

- R_A_TOV
- E_D_TOV
- Data field size
- Sequence level switching
- Disable device probing
- Suppress class F traffic
- Per-frame route priority
- BB credit
- PID format

This information can be found by issuing **configure** command

```
DS_4900B:admin> configure

Configure...

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

```
Domain: (1..239) [2]

R_A_TOV: (4000..120000) [10000]

E_D_TOV: (1000..5000) [2000]

WAN_TOV: (0..30000) [0]

MAX_HOPS: (7..19) [7]

Data field size: (256..2112) [2112]

Sequence Level Switching: (0..1) [0]

Disable Device Probing: (0..1) [0]

Suppress Class F Traffic: (0..1) [0]

Per-frame Route Priority: (0..1) [0]

Long Distance Fabric: (0..1) [0]

BB credit: (1..27) [16]

Disable FID Check (yes, y, no, n): [no]
```

There are non-fabric parameters that must match as well, such as zoning. Some fabric services, such as Management Server must match. If it is enabled in the fabric, then the switch you are introducing into the fabric must also have it enabled. If you experience a segmented fabric, refer to the *Fabric OS Troubleshooting and Diagnostics Guide* to fix the problem.

In-flight Encryption and Compression over ISLs

The in-flight encryption and compression feature of Fabric OS allows frames to be encrypted or compressed at the egress point of an ISL between two Brocade switches, and then to be decrypted or decompressed at the ingress point of the ISL. This feature uses port-based encryption and compression. It is supported on 16 Gbps E_Ports, only and the devices at either end of the ISL must run Fabric OS 7.0.0 or later software.
Encryption and compression can be enabled at the same time for an ISL, or you can enable either encryption or compression selectively.

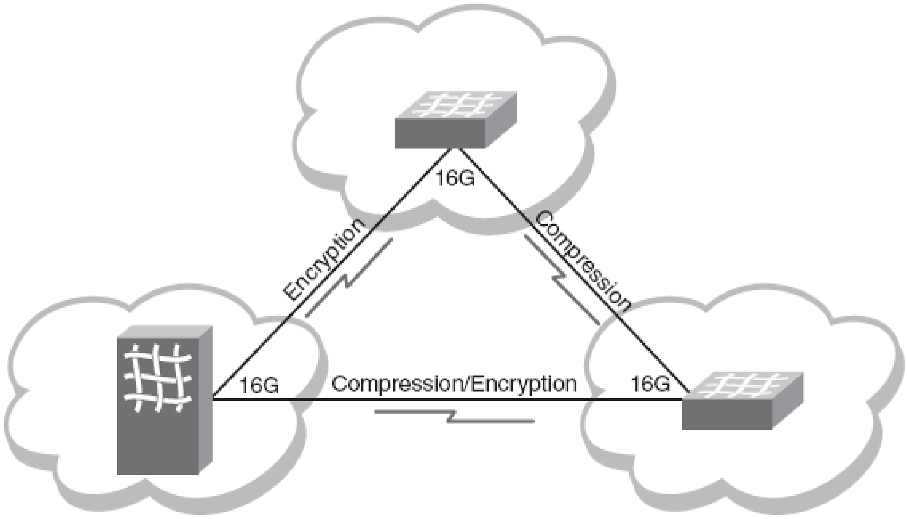

Figure: Encryption and Compression on 16 Gbps ISLs

For more on In-flight Encryption and Compression please refer to Fabric OS Adminstrator's Guide.

Forward Error Correction (FEC) on Brocade 6510's Condor3 ASIC

This enables Condor 3 to recover bit errors in a 10 and 16 Gbps data stream (frames and primitives)
Condor3 can correct up to 11 error bits in every 2112-bit transmission enhancing reliability of transmission
This feature is enabled by default on Condor 3 ports.

For more FEC please refer to Fabric OS Adminstrator's Guide.

Routing policies

By default, all routing protocols place their routes into a routing table. You can control the routes that a protocol places into each table and the routes from that table that the protocol advertises by defining one or more routing policies and then applying them to the specific routing protocol.
The routing policy is responsible for selecting a route based on one of two user-selected routing policies:

- Port-based routing
- Exchange-based routing

Displaying the current routing policy

1. Connect to the switch and log in as admin.
2. Enter the **aptPolicy** command with no parameters.

The current policy is displayed, followed by the supported policies for the switch.
Example of the output from the **aptPolicy** command.

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

```
switch:admin> aptpolicy
Current Policy: 3 1(ap)
3 0(ap): Default Policy
1: Port Based Routing Policy
3: Exchange Based Routing Policy
        0: AP Shared Link Policy
        1: AP Dedicated Link Policy
```

Setting the routing policy

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command to disable the switch.

3. Take the appropriate following action based on the route policy you choose to implement:
    - If Exchange-based policy is required, enter the **aptPolicy 3** command.
    - If Port-based policy is required, enter the **aptPolicy 1** command.

Setting up the AP route policy

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

Route selection

We can only select dynamic route selection on Brocade 300, 5100 and 5300.

Dynamic Load Sharing

The exchange-based routing policy depends on the Fabric OS Dynamic Load Sharing feature (DLS) for dynamic routing path selection. When using the exchange-based routing policy, DLS is enabled by default and cannot be disabled. In other words, you cannot enable or disable DLS when the exchange-based routing policy is in effect. When the port-based policy is in force, you can enable DLS to optimize routing. When DLS is enabled, it shares traffic among multiple equivalent paths between switches. DLS recomputes load sharing when any of the following occurs:

- a switch boots up
- an E_Port goes offline and online
- an EX_Port goes offline
- a device goes offline

Setting DLS

1. Connect to the switch and log in as admin.
2. Enter the **dlsShow** command to view the current DLS setting.

One of the following messages appears:

- "DLS is set" indicates that dynamic load sharing is turned on.
- "DLS is not set" indicates that dynamic load sharing is turned off.
- "DLS is set with Lossless enabled." DLS is enabled with the Lossless feature. Load sharing is recomputed with every change in the fabric, and existing routes can be moved to maintain optimal balance. In Lossless mode, no framers are lost during this operation.
- "DLS is set by default with current routing policy. DLS is set with Lossless enabled."
  Indicates that the current routing policy (exchange-based) requires DLS to be enabled by default. In addition, the lossless option is enabled. Frame loss is prevented during a load sharing re-computation. If you get this message, you cannot perform step 3, so you are done with this procedure.

3. Enter the **dlsSet** command to enable DLS or enter the **dlsReset** command to disable it.

Example of setting and resetting DLS.

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

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

Trunking overview

The trunking feature optimizes the use of bandwidth by allowing a group of inter-switch links (ISLs) to merge into a single logical link. Trunking is automatically implemented for any eligible ISLs after you install the Brocade ISL Trunking license. The license must be installed on each switch that participates in trunking.

Brocade's trunking feature supports the following trunking configurations:

- ISL trunking configurations are only applicable to E_Ports.
- F_Port trunking configurations are only applicable to two separate Fabric OS switches where all the ports on each switch reside in the same quad and are running at the same speed.
- EX_Port frame trunking configurations are between an FC router and the edge fabric.
- F_Port Masterless trunking configurations are on edge switches running in Access Gateway mode where the trunk ports are F_Ports, which are connected as N_Ports

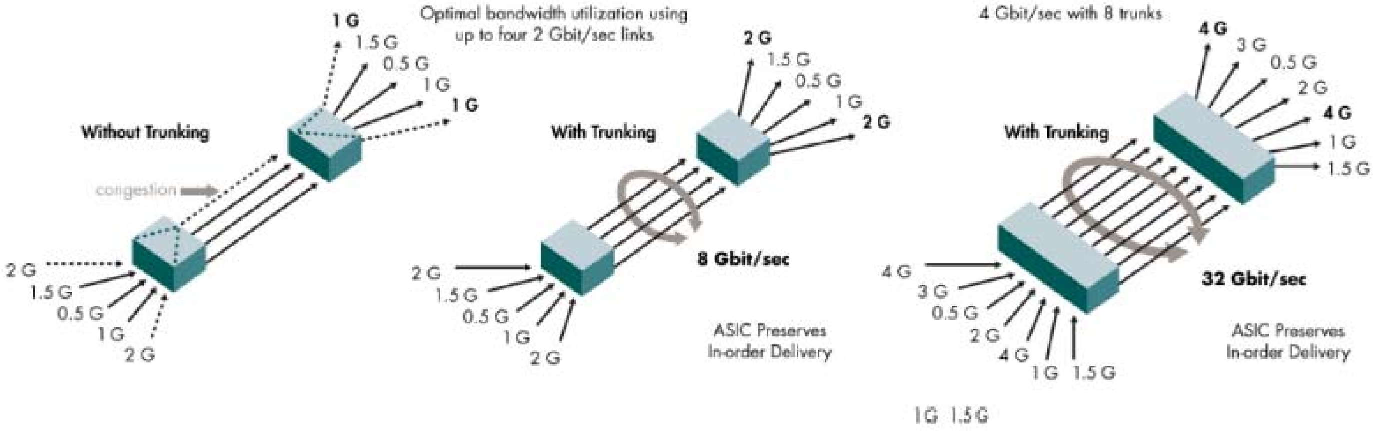

Following is the criteria for managing trunking connections:

- You can have up to eight ports in one trunk group to create high performance 32 Gbps ISL trunks between switches and up to 64 Gbps if there are eight ISLs with 8 Gbps each if 8 Gbps is supported.
- There must be a direct connection between participating switches.
- In Fabric OS v6.1.0 and later, you can configure EX_Ports to use frame-based trunking just like regular E_Ports. The EX_Port restrictions are the same as E_Ports. An E_Port or EX_Port trunk can be up to eight ports wide. All the ports must be adjacent to each other using the clearly marked groups on the front of the product
- The switch must be set to **interopMode 0** for Brocade Native mode, which supports all stand-alone Brocade switches, but provides no interoperability support.
- The port ISL R_RDY mode must be disabled (using the **portCfgIslMode** command).

Basic trunk group configuration

Re-initializing ports for trunking is required after you install the ISL Trunking license. You must re-initialize the ports being used for ISLs so that they recognize that trunking is enabled. This procedure needs to be performed only one time. To re-initialize the ports, you can either disable and then re-enable the switch, or disable and then re-enable the affected ports.

You can enable or disable Trunking for a single port or for an entire switch. When you issue the **portCfgTrunkPort** or **switchCfgTrunk** command to update the trunking configuration, the ports to which the configuration applies are disabled and re-enabled with the new trunk configuration. As a result, traffic through those ports can be disrupted.

Re-initializing ports for trunking

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **islShow** command to determine which ports are used for ISLs.
3. Enter the **portDisable** command for each ISL port.
4. Enter the **portEnable** command for each port that you disabled in step 3.

Enabling Trunking on a port

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the portCfgTrunkPort command to enable trunking. In the following example, trunking is being enabled on slot 1, port 3.

```
switch:admin> portcfgtrunkport 1/3 1
```

Enabling Trunking on a switch

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **switchCfgTrunk** command.
   Mode 1 enables and mode 0 disables ISL Trunking for all ports on the switch.

```
switch:admin> switchcfgtrunk 1
Committing configuration...done.
```

Displaying trunking information

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **trunkShow** command.
   This example shows trunking groups 1, 2, and 3; ports 4, 13, and 14 are masters.

```
switch:admin> trunkshow
1:        6-> 4 10:00:00:60:69:51:43:04 99 deskew 15 MASTER

2:        15-> 13 10:00:00:60:69:51:43:04 99 deskew 16 MASTER
          12-> 12 10:00:00:60:69:51:43:04 99 deskew 15
         14-> 14 10:00:00:60:69:51:43:04 99 deskew 17
         13-> 15 10:00:00:60:69:51:43:04 99 deskew 16

3:       24-> 14 10:00:00:60:69:51:42:dd 2 deskew 15 MASTER
```

F_Port trunking

F_Port trunking is enabled between two separate Fabric OS switches that support trunking and where all the ports on each switch reside in the same quad and are running the same speed. Trunk groups form when you connect two or more cables on one Fabric OS switch to another Fabric OS switch with ports in the same port group or quad. A port group or a quad is a set of sequential ports, for example ports 0-3 in the figure shown below. The Brocade 300, 5100, 5300 platforms support a trunk group with up to eight ports. The trunking groups are based on the user port number, with contiguous eight ports as one group, such as 0-7, 8-15, 16-23 and up to the number of ports on the switch.

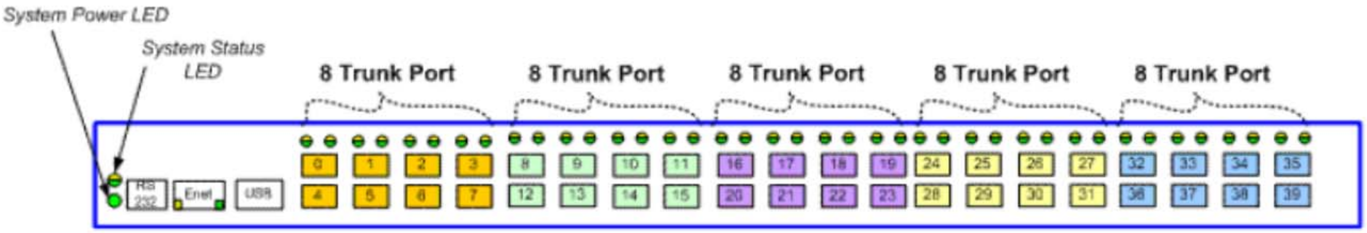

Enabling F_Port trunking

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **portDisable** command to disable the ports that are to be assigned to the trunk area.
3. Enter the **portTrunkArea –enable** command to create the trunk area.

```
switch:admin> portdisable 0-2
switch:admin> porttrunkarea --enable 0-2 -index 2
Trunk index 2 enabled for ports 0, 1, and 2.
```

Disabling F_Port trunking

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the portDisable command to disable the ports that are to be removed from the trunk area.
3. Enter the **portTrunkArea –disable** command to remove ports from the trunk area.

FC-FC Routing

The FC-FC routing service provides Fibre Channel routing (FCR) between two or more fabrics without merging those fabrics. A Fibre Channel router (*FC router*) is a switch running the FC-FC routing service. The FC-FC routing service can be simultaneously used as an FC router and as a SAN extension over wide area networks (WANs) using FCIP. FCR supports backbone-to-edge routing, allowing devices in the backbone to communicate with devices on the edge fabric.

Integrated Routing

Integrated Routing is a licensed feature that allows 8-Gbps FC ports to be configured as EX_Ports (or VEX_Ports) supporting Fibre Channel routing. This license eliminates the need to add a Brocade 7500 for FC-FC routing purposes. Using 8-Gbps ports for Fibre Channel routing provides double the bandwidth for each FCR connection (when connected to another 8-Gbps-capable port). It is supported on Brocade 5100 and 5300 but not on Brocade 300.

Setting up the FC-FC routing service

To set up the FC-FC Routing Service, perform the following tasks in the order listed:

- Verify that you have the proper setup for FC-FC routing.
- Assign backbone fabric IDs.
- Configure FCIP tunnels if you are connecting Fibre Channel SANs over IP-based networks
- Configure IFLs for edge and backbone fabric connection.
- Modify port cost for EX_Ports, if you want to change from the default settings.
- Configure trunking on EX_Ports that are connected to the same edge fabric
- Configure LSAN zones to enable communication between devices in different fabrics.

Verifying the setup for FC-FC routing

Before configuring a fabric to connect to another fabric, you must perform the following verification checks on the FC router.

1. Log in to the switch or director as admin and enter the version command. Verify that Fabric OS v6.4.0 is installed on the FC router as shown in the following example.

```
switch:admin> version
Kernel: 2.6.14.2
Fabric OS: v6.4.0
Made on: Fri Jan 22 01:15:34 2010
Flash: Mon Jan 25 20:53:48 2010
BootProm: 1.0.9
```

2. Enter the **interopMode** command and verify that Fabric OS switch interoperability with switches from other manufacturers is disabled.

```
switch:admin> interopmode

InteropMode: Off
usage: InteropMode [0|2|3 [-z McDataDefaultZone] [-s McDataSafeZone]]
0: to turn interopMode off
2: to turn McDATA Fabric mode on
Valid McDataDefaultZone: 0 (disabled), 1 (enabled)
Valid McDataSafeZone: 0 (disabled), 1 (enabled)
3: to turn McDATA Open Fabric mode on
```

If InteropMode is on, FC routing is not supported. To turn off interoperability mode, disable the switch and enter the interopMode 0 command

3. Verify that the Fabric Wide Consistency Policy is not in 'strict' mode by issuing the **fddCfg –showall** command. When it is in strict mode, ACL cannot support Fibre Channel routing in the fabric.

```
switch:admin> fddcfg –showall
Local Switch Configuration for all Databases:-
DATABASE - Accept/Reject
-------------------------------------
SCC - accept
DCC - accept
PWD - accept
Fabric-Wide Consistency Policy :- "SCC:S;DCC"
```

If the Fabric Wide Consistency Policy has the letter “S” in it in the edge fabric or the backbone fabric, do not connect the edge fabric to the FC router. The letter “S” (shown in the preceding sample output) indicates the policy is strict. The fabric-wide policy must be tolerant before you can connect fabrics to the FC router.

Assigning backbone fabric IDs

1. Log in to the switch or director.
2. Enter the **switchDisable** command if EX_Ports are online.
3. Enter the **fosConfig –disable fcr** command to disable the FC-FC Routing Service.
   The default state for the FCR is disabled.
4. Enter the **fcrConfigure** command. At the prompt, enter the fabric ID, or press Enter to keep the current fabric ID, which is displayed in brackets.
   Verify the backbone fabric ID is different from that set for edge fabrics.
   Multiple FC routers attached to the same backbone fabric must have the same backbone fabric ID.
5. Enter the **fosConfig –enable fcr** command.
6. Enter the **switchEnable** command.

```
switch:admin> switchdisable
switch:admin> fosconfig –disable fcr
FC Router service is disabled
switch:admin> fcrconfigure
FC Router parameter set. <cr> to skip a parameter
Please make sure new Backbone Fabric ID does not conflict with any configured
EX-Port's Fabric ID Backbone fabric ID: (1-128)[128]
switch:admin> fosconfig –enable fcr
FC Router service is enabled
switch:admin> switchenable
```

Creating an FCIP tunnel

As you plan the tunnel configurations, be aware that uncommitted rate tunnels use a minimum of 1000 Kbps, up to a maximum of available uncommitted bandwidth on the GbE port. The total bandwidth available on a GbE port is 1 Gbps. You can configure tunnels as bidirectional entities with different commit rates in both directions.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Create an FCIP tunnel using the portCfg fciptunnel command. The command syntax is as follows.

```
portCfg fciptunnel [slot/]ge0|ge1 create tunnel_id remote_ip_addr local_ip_addr comm_rate
[-c] [-s] [-f] [-t] [-M] [-n remote_wwn] [-k timeout] [-r retransmissions] [-m time] [-q control_dscp] [-Q data_dscp] [-v vlan_id] [-p
control_L2CoS] [-P data_L2CoS] [-ike ike_number] [-ipsec ipsec_number] [-key preshared_key] [-d FCIP_tunnel_description]
[-bstr 0|1 TCP Byte Streaming]
```

Example of creating an FCIP tunnel

The following example creates one end of a tunnel over ge0 between remote IP address 192.168.10.1 and local IP address 192.168.20.1 with a tunnel id of 0, over VLAN 100, with a layer 2 class of service of 3 for control traffic, and a layer 2 class of service of 7 for data traffic.

```
portcfg fciptunnel 8/ge0 create 2 192.168.10.1 192.168.20.1 0 -v 100 -p 3 -P 7
```

Inter-fabric link configuration

Before configuring an IFL, be aware that you cannot configure both IFLs (EX_Ports, VEX_Ports) and ISLs (E_Ports) from a backbone fabric to the same edge fabric.

Configuring an inter-fabric link involves disabling ports and cabling them to other fabrics, configuring those ports for their intended use, and then enabling the ports.

To configure an 8-Gbps IFL, both the EX_Port and the connecting E_Port must be 8-Gbps ports.

Configuring an IFL for both edge and backbone connections

1. On the FC router, disable the port that you are configuring as an EX_Port (the one connected to the Fabric OS switch) by issuing the portDisable command.
   switch:admin> **portdisable 7**
   You can verify that port 7 has been disabled by issuing the portShow command for the port.
2. Configure each port that connects to an edge fabric as an EX_Port or VEX_Port. Note the following:
   - portCfgVEXPort works only on VE_Ports.
   - portCfgEXPort (only on the FC ports on the FC router) commands work only on ports that are capable of FC-FC routing.

   Use the portCfgEXPort or portCfgVEXPort command to:
- Enable or disable EX_Port or VEX_Port mode.
- Set the fabric ID (avoid using fabric IDs 1 and 128, which are the default IDs for backbone connections).

The following example configures the EX_Port (or VEX_Port) and assigns a Fabric ID of 30 to port 7.

```
switch:admin> portcfgexport 7 -a 1 -f 30
switch:admin> portcfgexport 7
Port 7/10 info
Admin: enabled
State: NOT OK
Pid format: Not Applicable
Operate mode: Brocade Native
Edge Fabric ID: 30
Preferred Domain ID: 160
Front WWN: 50:06:06:9e:20:38:6e:1e
Fabric Parameters: Auto Negotiate
R_A_TOV: Not Applicable
E_D_TOV: Not Applicable
Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

3. Enter the **portEnable** command to enable the ports that you disabled in step 1.
   switch:admin> **portenable 7**
   Physically attach ISLs from the Fibre Channel router to the edge fabric.
4. Enter the portCfgShow command to view ports that are persistently disabled.

```
switch:admin> portcfgshow 7
Area Number: 74
Speed Level: AUTO
Trunk Port OFF
Long Distance OFF
VC Link Init OFF
Locked L_Port OFF
Locked G_Port OFF
```

```
Disabled E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable OFF
NPIV capability ON
EX Port ON
Mirror Port ON
FC Fastwrite ON
```

5. After identifying such ports, enter the **portCfgPersistentEnable** command to enable the port,and then the **portCfgShow** command to verify the port is enabled.

```
switch:admin> portcfgpersistentenable 7
switch:admin> portcfgshow 7
Area Number: 74
Speed Level: AUTO
Trunk Port OFF
Long Distance OFF
VC Link Init OFF
Locked L_Port OFF
Locked G_Port OFF
Disabled E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable OFF
NPIV capability ON
EX Port ON
Mirror Port ON
FC Fastwrite ON
```

Enter either the **portCfgEXPort** or **portShow** command to verify that each port is configured correctly:

```
switch:admin> portcfgexport 7
Port 7 info
Admin: enabled
State: NOT OK
Pid format: Not Applicable
Operate mode: Brocade Native
Edge Fabric ID: 30
Preferred Domain ID: 160
Front WWN: 50:06:06:9e:20:38:6e:1e
Fabric Parameters: Auto Negotiate
R_A_TOV: Not Applicable
E_D_TOV: Not Applicable
Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A

switch:admin_06> portshow 7
portName:
portHealth: OFFLINE
Authentication: None
EX_Port Mode: Enabled
Fabric ID: 30
```

```
Front Phantom: state = Not OK Pref Dom ID: 160
Fabric params: R_A_TOV: 0 E_D_TOV: 0 PID fmt: auto
Authentication Type: None
Hash Algorithm: N/A
DH Group: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
portDisableReason: None
portCFlags: 0x1
portFlags: 0x1 PRESENT U_PORT EX_PORT
portType: 10.0
portState: 2 Offline
portPhys: 2 No_Module
portScn: 0
port generation number: 0
portId: 014a00
portIfId: 4372080f
portWwn: 20:4a:00:60:69:e2:03:86
portWwn of device(s) connected:
Distance: normal
portSpeed: N4Gbps
LE domain: 0
FC Fastwrite: ON
Interrupts: 0 Link_failure: 0 Frjt : 0
Unknown: 0 Loss_of_sync: 0 Fbsy : 0
Lli: 0 Loss_of_sig: 2
Proc_rqrd: 0 Protocol_err: 0
Timed_out: 0 Invalid_word: 0
Rx_flushed: 0 Invalid_crc: 0
Tx_unavail: 0 Delim_err: 0
Free_buffer: 0 Address_err: 0
Overrun: 0 Lr_in: 0
Suspended: 0 Lr_out: 0
Parity_err: 0 Ols_in: 0
2_parity_err: 0 Ols_out: 0
CMI_bus_err: 0
Port part of other ADs: No
```

6. Enter the **switchShow** command to verify the EX_Port (or VEX_Port), edge fabric ID, and name of the edge fabric switch (containing the E_Port or VE_Port) are correct.
7. Enter the **fcrFabricShow** command to view any edge fabric's switch names and ensure links are working as expected

```
switch:admin> fcrfabricshow
FCR WWN: 10:00:00:05:1e:13:59:00, Dom ID: 2, Info: 10.32.156.52
1080::8:800:200C:1234/64,
"fcr_7500"
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)
------------------------------------------------------------------------
7 10 10:00:00:05:1e:34:11:e5 10.32.156.33 "7500" 1080::8:8FF:FE0C:417A/64
4 116 10:00:00:05:1e:37:00:44 10.32.156.34 "7500"
FCR WWN: 10:00:00:05:1e:12:e0:00, Dom ID: 100, Info:10.32.156.50
1080::8:60F:FE0C:456A/64
"fcr_7500"
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)
----------------------------------------------------------------------------------
4 95 10:00:00:05:1e:37:00:45 10.32.156.31 "7500"
FCR WWN: 10:00:00:05:1e:12:e0:00, Dom ID: 100, Info: 10.32.156.50,
"fcr_Brocade 7500"
```

```
EX_Port FID Neighbor Switch Info (WWN, enet IP, name)
------------------------------------------------------------------------
4 95 10:00:00:05:1e:37:00:45 10.32.156.31 "Brocade 7500"
5 95 10:00:00:05:1e:37:00:45 10.32.156.31 "Brocade 7500"
6 95 10:00:00:05:1e:37:00:45 10.32.156.31 "Brocade 7500"
```

Setting router port cost for an EX_Port

The router port cost value for an EX_Port is set automatically when the EX_Port is created. However, you can modify the cost for that port. You can configure the EX_ or VEX_Port with values of either 1000 or 10,000. If you want to differentiate between two EX_Port links with different speeds, you can assign 1000 to one link and 10,000 to the other link.

1. Enter the **portDisable** command to disable any port on which you want to set the router port cost.

```
switch:admin> portdisable 7
```

2. Enable EX_Port or VEX_Port mode with the **portCfgEXPort** or **portCfgVEXPort** command.

```
switch:admin> portcfgexport 7 -a 1
```

3. Enter the **fcrRouterPortCost** command to display the router port cost for each EX_Port.

```
switch:admin> fcrrouterportcost

Port Cost
-------------------------
7          1000
```

You can also use the **fcrRouteShow** command to display the router port cost.

4. Enter the **fcrRouterPortCost** command with a port and slot number, to display the router port cost for a single EX_Port.

```
switch:admin> fcrrouterportcost 7
Port Cost
-------------------------
7          1000
```

5. Enter the appropriate form of the **fcrRouterPortCost** command based on the task you want to perform:
   - To set the router port cost for a single EX_Port, enter the command with a port and a specific cost:

```
switch:admin> fcrrouterportcost 7 10000
```

   - To set the cost of the EX_Port back to the default, enter a cost value of 0:

```
switch:admin> fcrrouterportcost 7 0
```

6. Enter the **portEnable** command to enable the ports that you disabled in step 1.

```
switch:admin> portenable 7
```

Configuring EX_Port frame trunking

With EX_Port frame trunking, you can use the same CLI commands as you do for E_Port trunking. Administration control for EX_Port trunking is available through root, admin, and switch admin access. The procedures for administering EX_Port frame trunking are the same as for E_Port trunking. You initialize trunking on ports with portCfgTrunkPort or switchCfgTrunk, and monitor traffic with the portPerfShow command.

LSAN zone configuration

An 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. You can define and manage LSANs using Brocade Advanced Zoning.

LSAN zones and fabric-to-fabric communications

Zoning is enforced by all involved fabrics; any communication from one fabric to another must be allowed by the zoning setup on both fabrics. If the SANs are under separate administrative control, then separate administrators maintain access control.

Controlling device communication with the LSAN

The following procedure illustrates how LSANs control which devices can communicate with each other. The procedure shows the creation of two LSANs (called *lsan_zone_fabric75* and *lsan_zone_fabric2*), which involve the following devices and connections:

- Switch1 and the host in fabric75.
- Switch2, Target A, and Target B in fabric2.
- Switch1 is connected to the FC router using an EX_Port or VEX_Port.
- Switch2 is connected to the FC router using another EX_Port or VEX_Port.
- Host has WWN 10:00:00:00:c9:2b:c9:0c (connected to switch1).
- Target A has WWN 50:05:07:61:00:5b:62:ed (connected to switch2).
- Target B has WWN 50:05:07:61:00:49:20:b4 (connected to switch2).

1. Log in as admin and connect to switch1.
2. Enter the **nsShow** command to list the WWN of the host (10:00:00:00:c9:2b:c9:0c).

NOTE
The nsShow output displays both the port WWN and node WWN; the port WWN must be used for LSANs.

```
switch:admin> nsshow
{
Type Pid COS PortName NodeName
TTL(sec)
N 060f00; 2,3; 10:00:00:00:c9:2b:c9:0c; 20:00:00:00:c9:2b:c9:0c; na
FC4s: FCP
NodeSymb: [35] "Emulex LP9002 FV3.91A3 DV5-5.20A6 "
Fabric Port Name: 20:0f:00:05:1e:37:00:44
Permanent Port Name: 10:00:00:00:c9:2b:c9:0c
The Local Name Server has 1 entry }
```

3. Enter the **zoneCreate** command to create the LSAN *lsan_zone_fabric75*, which includes the host.

```
switch:admin> zonecreate "lsan_zone_fabric75", "10:00:00:00:c9:2b:c9:0c"
```

4. Enter the **zoneAdd** command to add Target A to the LSAN.

```
FID75Domain5:admin> zoneadd "lsan_zone_fabric75", "50:05:07:61:00:5b:62:ed"
```

5. Enter the **cfgAdd** or **cfgCreate** and **cfgEnable** commands to add and enable the LSAN configuration.

```
switch:admin> cfgadd "zone_cfg", "lsan_zone_fabric75"
```

```
switch:admin> cfgenable "zone_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.

Do you want to enable 'zone_cfg' configuration (yes, y, no, n): [no] y
zone config "zone_cfg" is in effect
Updating flash ...
```

6. Log in as admin to fabric2.
7. Enter the **nsShow** command to list Target A (50:05:07:61:00:5b:62:ed) and Target B (50:05:07:61:00:49:20:b4).

```
switch:admin> nsshow
{
Type Pid COS PortName NodeName TTL(sec)
NL 0508e8; 3; 50:05:07:61:00:5b:62:ed; 50:05:07:61:00:1b:62:ed; na
FC4s: FCP [IBM DNEF-309170 F90F]
Fabric Port Name: 20:08:00:05:1e:34:11:e5
Permanent Port Name: 50:05:07:61:00:5b:62:ed
NL 0508ef; 3; 50:05:07:61:00:49:20:b4; 50:05:07:61:00:09:20:b4; na
FC4s: FCP [IBM DNEF-309170 F90F]
Fabric Port Name: 20:08:00:05:1e:34:11:e5
Permanent Port Name: 50:05:07:61:00:49:20:b4
The Local Name Server has 2 entries }
```

8. Enter the **zoneCreate** command to create the LSAN *lsan_zone_fabric2*, which includes the host (10:00:00:00:c9:2b:6a:2c), Target A, and Target B.

```
switch:admin> zonecreate "lsan_zone_fabric2",
"10:00:00:00:c9:2b:c9:0c;50:05:07:61:00:5b:62:ed;50:05:07:61:00:49:20:b4"
```

9. Enter the **cfgShow** command to verify that the zones are correct.

```
switch:admin> cfgshow
Defined configuration:
zone: lsan_zone_fabric2
10:00:00:00:c9:2b:c9:0c; 50:05:07:61:00:5b:62:ed;
50:05:07:61:00:49:20:b4
Effective configuration:
no configuration in effect
```

10. Enter the **cfgAdd** and **cfgEnable** commands to create and enable the LSAN configuration.

```
switch:admin> cfgadd "zone_cfg", "lsan_zone_fabric2"
switch:admin> cfgenable "zone_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'zone_cfg' configuration (yes, y, no, n): [no] y
zone config "zone_cfg" is in effect
Updating flash ...
```

11. Log in as an admin and connect to the FC router.
12. Enter the following commands to display information about the LSANs.
    lsanZoneShow -s shows the LSAN.

```
switch:admin> lsanzoneshow -s
Fabric ID: 2 Zone Name: lsan_zone_fabric2
10:00:00:00:c9:2b:c9:0c Imported
50:05:07:61:00:5b:62:ed EXIST
50:05:07:61:00:49:20:b4 EXIST
Fabric ID: 75 Zone Name: lsan_zone_fabric75
10:00:00:00:c9:2b:c9:0c EXIST
50:05:07:61:00:5b:62:ed Imported
```

fcrPhyDevShow shows the physical devices in the LSAN.

```
switch:admin> fcrphydevshow
Device WWN Physical
Exists PID
in Fabric
-----------------------------------------
75 10:00:00:00:c9:2b:c9:0c c70000
2 50:05:07:61:00:5b:62:ed 0100ef
2 50:05:07:61:00:5b:62:ed 0100e8
Total devices displayed: 3
```

- fcrProxyDevShow shows the proxy devices in the LSAN.

```
switch:admin> fcrproxydevshow
Proxy WWN Proxy Device Physical State
Created PID Exists PID
in Fabric in Fabric
----------------------------------------------------------------------------
75 50:05:07:61:00:5b:62:ed 01f001 2 0100e8 Imported
2 10:00:00:00:c9:2b:c9:0c 02f000 75 c70000 Imported
Total devices displayed: 2
```

On the FC router, the host and Target A are imported, because both are defined by *lsan_zone_fabric2* and *lsan_zone_fabric75*. However, target B is defined by *lsan_zone_fabric75* and is not imported because *lsan_zone_fabric2* does not allow it.
When a PLOGI, PDISC, or ADISC arrives at the FC router, the SID and DID of the frame are checked. If they are LSAN-zoned at both SID and DID edge fabrics, the frame is forwarded to the DID. If they are not zoned, only the PLOGI is dropped; for the remaining frames zoning enforcement takes place in the edge fabrics.

VCS/VDX6730 to FC SAN Connectivity

This feature enables connectivity between hosts (using FCoE) connected to VCS/VDX platforms and FC storage connected to FC SAN via FCR. An E-port on a VDX6730 platform running NOS v2.1.1 is connected to an EX_port on an FCR running FOS v7.0.1 to enable this functionality.
Note:
Integrated Routing license is not required to share devices between VDX/VCS Ethernet fabric and FC SAN fabric.

It is recommended to use 5300, DCX/DCX-4S, DCX8510-8, DCX8510-4 for FCR functionality for higher scalability.

A new FCR EX_port mode 5 is used to connect VCS/VDX6730 to FCR

Web Tools

Disabling or enabling ISL trunking

The trunking feature requires using Web Tools with the EGM license. If you attempt to use this feature without the EGM license, the following error message displays.

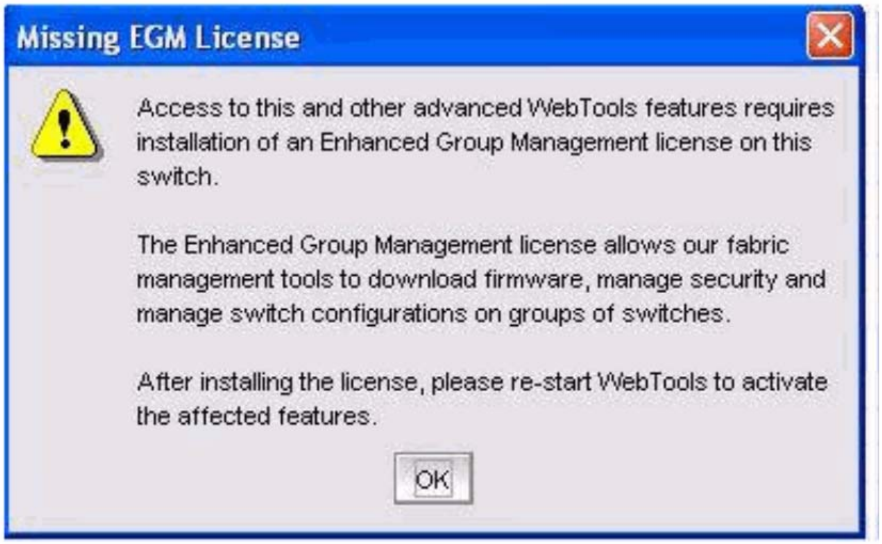

Missing EGM License

Access to this and other advanced WebTools features requires installation of an Enhanced Group Management license on this switch.

The Enhanced Group Management license allows our fabric management tools to download firmware, manage security and manage switch configurations on groups of switches.

After installing the license, please re-start WebTools to activate the affected features.

When the trunking license is activated, trunks are automatically established on eligible ISLs and trunking capability is enabled by default on all ports. Use the following procedure to disable trunking on a port or to re-enable trunking if it has been disabled.
Trunking is not supported on logical ports or GbE ports.

1. Click a port in the Switch View to open the Port Admin window.
2. Click the FC Ports tab.
   Trunking mode does not apply to GbE ports.
3. From the tree on the left, click the switch name or slot name.
4. From the table, select the port that you want to trunk.
   You can select multiple ports from the table. You cannot select multiple ports from the tree.
   Trunking mode does not apply to logical ports.
5. Click the Show Advanced Mode of Ports Admin.
   If the button is unavailable, the port is already in that state.
   Click Yes in the confirmation window

Viewing trunk group information

Use the Trunking tab on the Switch Administration window to view trunk group information

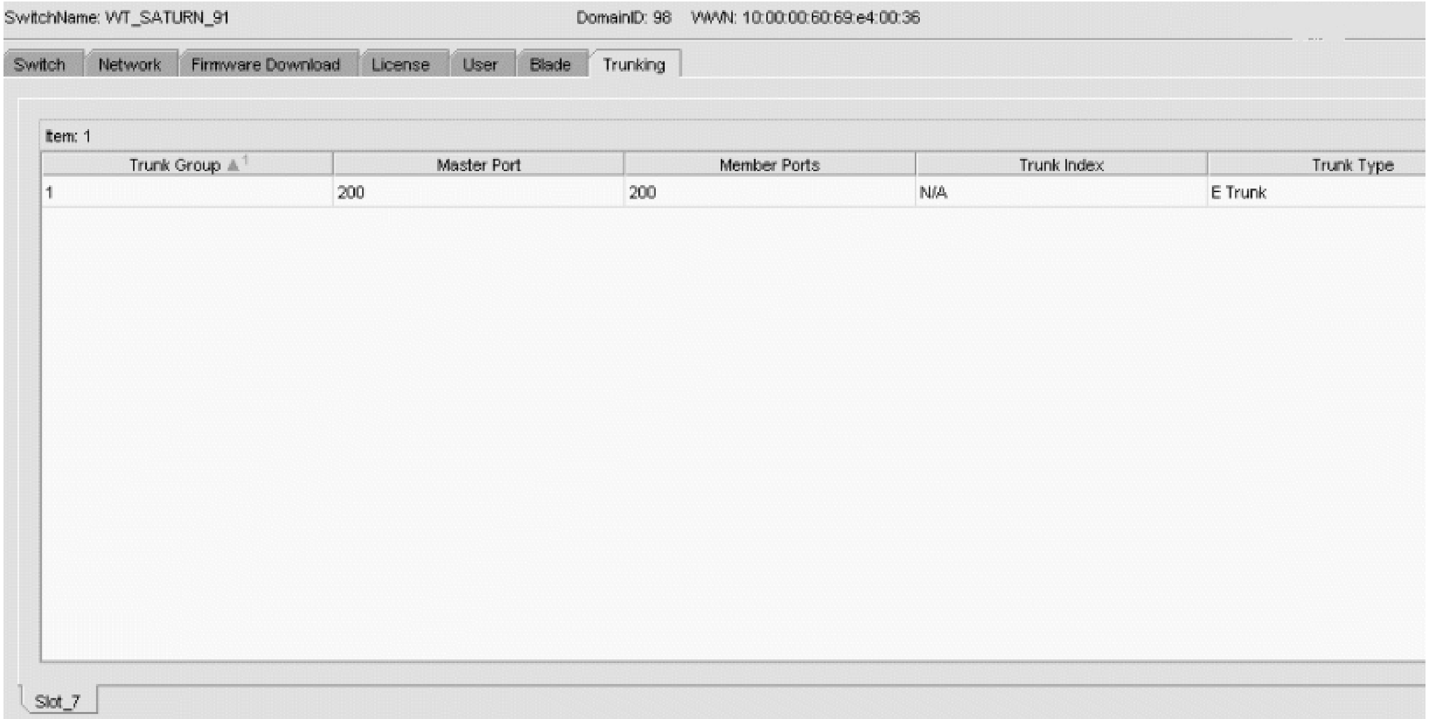

SwitchName: WT_SATURN_91 DomainID: 98 WWN: 10:00:00:60:69:e4:00:36

Item: 1

Trunk Group ▲1	Master Port	Member Ports	Trunk Index	Trunk Type
1	200	200	N/A	E Trunk

The following trunking attributes can be displayed from the Port Admin view by selecting Show Advanced Mode:

- Trunk port state, either master or slave.
- Trunk master port (does not apply to F_Port trunking).
- Trunk index (applies only to F_Port trunking).

F_Port trunk groups

F_Port trunking provides extra bandwidth and robust connectivity for hosts and targets connected by switches in Access Gateway mode. There are five general criteria for establishing F_Port trunking:

- The F_Port trunking feature requires installing the EGM license; otherwise if you attempt to use this feature in Web Tools without the license, the following error message displays.

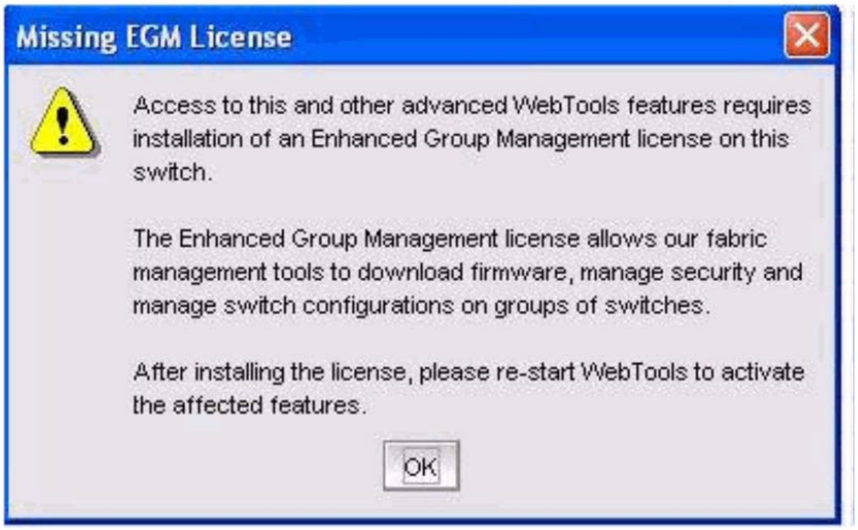

Missing EGM License

Access to this and other advanced WebTools features requires installation of an Enhanced Group Management license on this switch.

The Enhanced Group Management license allows our fabric management tools to download firmware, manage security and manage switch configurations on groups of switches.

After installing the license, please re-start WebTools to activate the affected features.

NOTE
The EGM license is required only for 8 Gbps platforms, such as the Brocade DCX enterprise-class platform, the Encryption Switch, the 300, 5300, and 5100 switches. For non-8
Gbps platforms, all functionalities are available without EGM license.

- Trunking must be enabled on the ports.
- The trunking license must be enabled on the switch in Access Gateway mode.
- The ports should not be configured for long distance connections.
- The ports should not be port-swapped.

When you create an F_Port trunk, you create a logical entity called a trunk index (TI), which represents the physical ports. The TI represents all ports in the trunk. If a master port fails, and a slave port takes over, the TI stays the same.

NOTE
If F_Port trunking is configured, a firmware downgrade is not allowed.

Creating and maintaining F_Port trunk groups

User this procedure to create an F_Port trunk group, and to add or remove member ports.

1. Select Port Admin.
2. Click Show Advanced Mode.
3. Select any port from the port group in which you want to create the trunk group.
4. Select F_Port Trunking.

The F_Port Trunking dialog box displays.

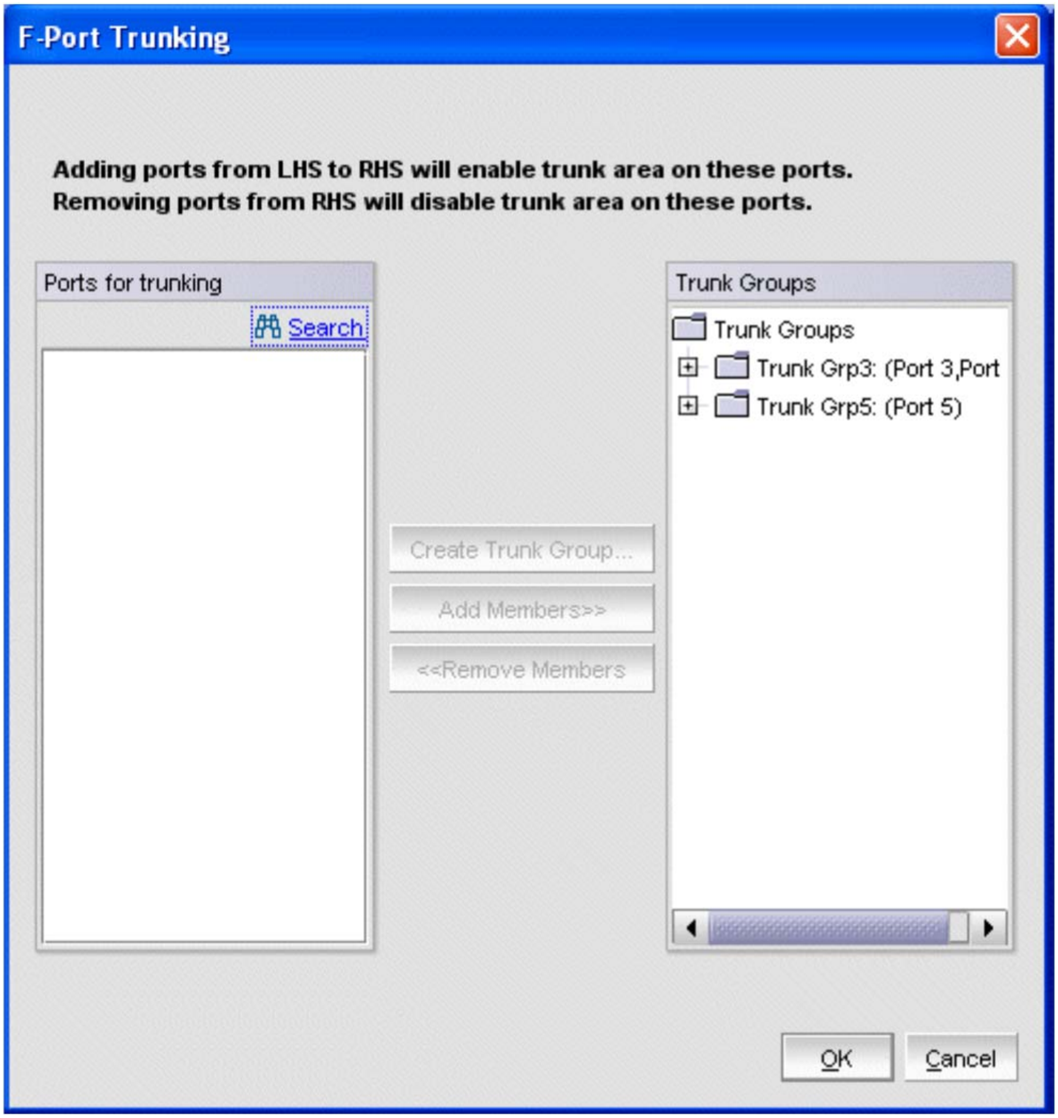

F-Port Trunking

Adding ports from LHS to RHS will enable trunk area on these ports.
Removing ports from RHS will disable trunk area on these ports.

5. Select one or more ports in the Ports for trunking pane.
   A dialog box displays, asking you to select a trunk index.
6. Select the trunk index from the drop-down box populated with the index for all the ports.
   A trunk group is created, identified by the trunk index, and containing the port you selected.
7. Select the trunk group you just created.
   Add Members becomes active.
8. Additional ports can be added by selecting a port from Ports for trunking table and then clicking Add Members.
   To remove a port from the trunk group, select the port from Trunk Groups table and then click Remove Members.
9. Click OK when you are finished.

FC-FC routing management

You can perform Fibre Channel Routing operations using Web Tools, Web Tools with the EGM license, and Integrated Routing license. You can manage FC-FC Routing through the FC Routing module. The FC Routing module has tabbed panes that display EX_Ports, LSAN fabrics, LSAN zones, LSAN devices, and general FCR information.

The FC Routing module provides a dynamic display. Any changes in the FCR configuration on the switch are automatically updated in the FC Routing module within 30 to 90 seconds, depending on the network traffic.
The switch must be FC Router-capable. The only things you need to configure on the FC Router are the EX_Ports and the backbone fabric ID. You configure LSAN zones on the fabrics from where devices need to be shared. You can configure LSAN zones on the backbone fabric to allow edge fabrics to share devices in the backbone fabric.
You can log in with any role and launch the FC routing module. To modify the data, you must log in as switchadmin, fabricadmin, basicswitchadmin, or operator. If you log in as user, zoneadmin, or securityadmin, you can only view the data.
If the FC-FC Routing service is disabled, the LSAN zones, LSAN fabric, and devices tabs will continue to show the existing entries, but shows the entries related to the *backbone fabric* only.
EX_Port configurations must be removed to disable FC-FC Routing service.

Opening the FC Routing module

The FCR button in the Switch View launches the FC Routing module.

NOTE
When the Virtual Fabrics capability is enabled on the switch, Fabric ID cannot be set using the Set Fabric ID button.

Use the following procedure to open the FC Routing module.

1. Select a switch from the Fabric Tree.
   The selected switch displays in the Switch View.
2. Click FCR in the Manage section of the Tasks menu.
   The FC Routing module displays. If FC-FC Routing is disabled, a message to that effect displays on all the tabs in the module.

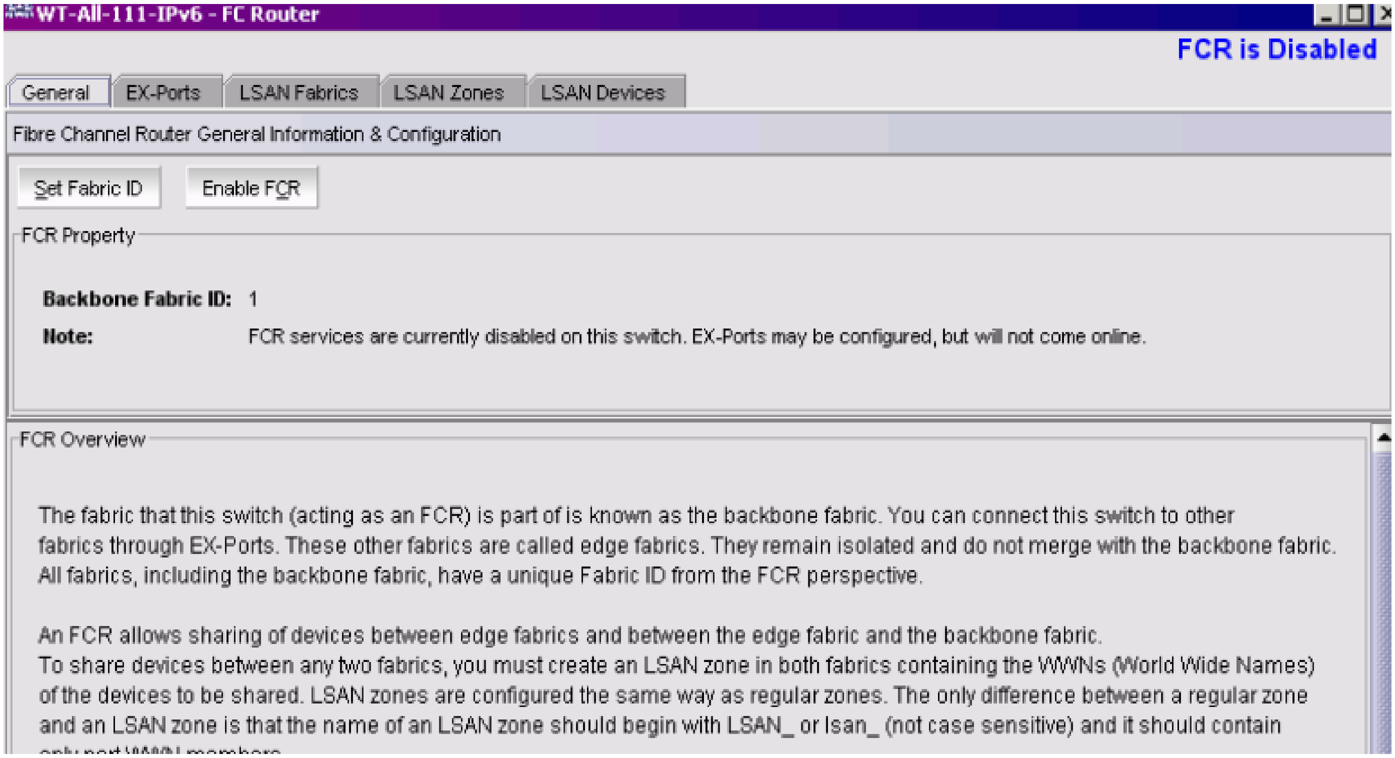

FCR is Disabled

Fibre Channel Router General Information & Configuration

FCR Property

Backbone Fabric ID: 1
Note: FCR services are currently disabled on this switch. EX-Ports may be configured, but will not come online.

FCR Overview

The fabric that this switch (acting as an FCR) is part of is known as the backbone fabric. You can connect this switch to other fabrics through EX-Ports. These other fabrics are called edge fabrics. They remain isolated and do not merge with the backbone fabric. All fabrics, including the backbone fabric, have a unique Fabric ID from the FCR perspective.

An FCR allows sharing of devices between edge fabrics and between the edge fabric and the backbone fabric.
To share devices between any two fabrics, you must create an LSAN zone in both fabrics containing the WWNs (World Wide Names) of the devices to be shared. LSAN zones are configured the same way as regular zones. The only difference between a regular zone and an LSAN zone is that the name of an LSAN zone should begin with LSAN_ or lsan_ (not case sensitive) and it should contain only port WWN members.

Configuring an EX_Port

Use the following procedure to configure an EX_Port.

1. Select Tasks > Manage > FCR.
2. Click the EX_Ports tab.
3. Click New in the task bar to configure one or more EX_Ports.

   NOTE
   For Brocade 7800 extension switch and FX8-24 blade, New button is enabled only if Integrated Routing license is present.
   This opens the port configuration wizard, which guides you through the port configuration process.

4. Follow the instructions in the wizard to configure the EX_Port.
   You must specify the Fabric ID and, if configuring an FC port, the speed and long distance mode. You can choose any unique fabric ID as long as it is consistent for all EX_Ports that connect to the same edge fabric.

Editing the configuration of an EX_Port

Use the following procedure to edit the configuration of an EX_Port.

1. Select Tasks > Manage > FCR.
   Click the EX_Ports tab.
2. Select a port to configure, by clicking in the row.
3. Click Edit Configuration in the task bar.

This opens the port configuration wizard, which guides you through the port configuration process.
The current configuration values are displayed in the wizard steps.
If you choose to configure a disabled port, the wizard provides the Enable Port after configuration check box. If you select this check box, the disabled port is automatically enabled after configuration. If you leave this box cleared, the port remains in the same state after configuration.

Configuring FCR router port cost

In FCR, EX_Ports can be assigned router port cost. The cost of the link is a positive number. The router port path or tunnel path is chosen based on the minimum cost per connection. If multiple paths exist with the same minimum cost, there will be load sharing over these paths. If multiple paths exist where one path costs lower than the others, then the lowest cost path is used.
Every link has a default cost. For an EX_Port 1 Gbps, 2 Gbps, 4 Gbps, and 8 Gbps links, the default cost is 1000. For a VEX_Port, the default cost is 10000. If the cost is set to 0, the default cost will be used for that link.

4. Open the Switch View window.
5. Click FCR in Manage section of the Tasks menu.
6. Click the Ex_Ports tab.
7. Click the Router Port Cost button.

Viewing LSAN zones

The LSAN Zones tab displays all the LSAN zones, in both a tabular and tree form. If FC-FC Routing is disabled, the table and the tree node in this tab display only the LSAN zones present in the backbone fabric.
For more detailed information about a specific LSAN zone, click a zone name in the table and then click the View Details button in the task bar. You can also click the zone name in the tree on the left side of the window.

The LSAN matrix is mapping of LSAN Zones with the edge fabric they are going to communicate with. When an LSAN matrix is created in the backbone fabric, only the LSAN zones mapped in the edge fabrics are displayed in the LSAN Zones tab.

Viewing LSAN Devices

The LSAN Devices tab displays information about the physical and proxy devices and displays these devices in a tree on the left side of the window. (If FC-FC Routing is disabled, the tables and tree nodes in this tab are empty.

Click the LSAN Devices element in the tree to display a count of all the physical and proxy LSAN devices. Note that this count is for all of the LSAN fabrics.

Click the Physical Devices or Proxy Devices element in the tree to see a detailed list of the physical or proxy devices. Click the device name in the tree for more detailed information about a specific device.

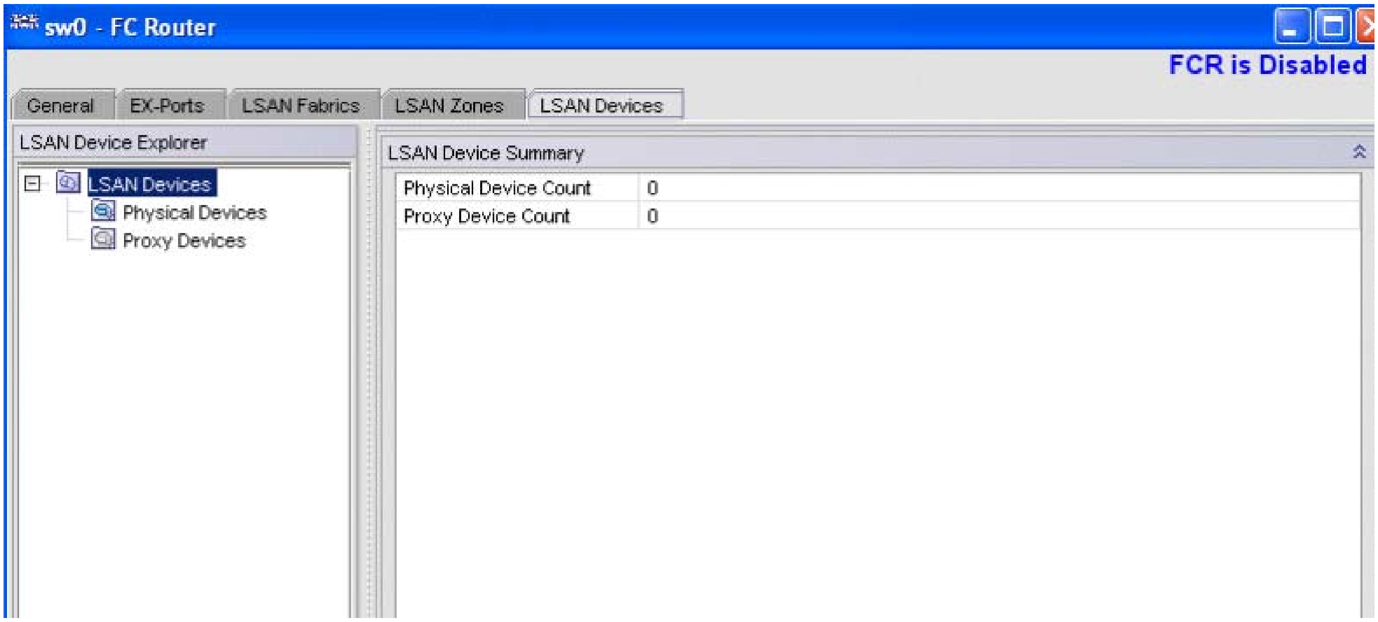

Physical Device Count	0
Proxy Device Count	0

Configuring the backbone fabric ID

To configure the backbone fabric ID, you must disable the switch; however, all the Ex_Ports must be removed before invoking this operation. After the fabric ID is changed, you can configure these ports again.

The fabric ID for a backbone fabric must be different from the fabric IDs of all other edge fabrics; otherwise, a fabric ID conflict error can occur.

Make sure that all switches in the backbone fabric have the same fabric ID.

1. Open the Switch View window.
2. Click FCR in the Manage section of the Tasks menu.
3. Click the EX-Ports tab.
4. Remove the entire EX_ports configuration and disable the switch.
5. Click the General tab.
6. Click Set Fabric ID in the task bar.

   The Configure Backbone Fabric ID window displays.
7. Select a fabric ID from the drop-down menu.

   The fabric ID is a number from 1 through 128. Web Tools warns you if you select a fabric ID that is already in use.
8. Click OK.
9. Enable the switch and manually enable FC-FC Routing Service.

Enabling Access Gateway mode

When you enable Access Gateway mode some fabric information, such as the zone and security databases, is erased. To recover this information, save the switch configuration before enabling
Access Gateway mode.

To save the switch configuration using Web Tools, click Switch Admin in the Manage section under

Tasks, and then click the Configure > Upload/Download subtab and upload the configuration file.

NOTE
You cannot enable Access Gateway mode if Management Server is enabled. To disable Management Server, enter the MsplmgmtDeactivate command.

1. Select a switch.
2. Click Switch Admin in the Manage section under Tasks.

   The Switch Administration window opens.
3. Click Disable in the Switch Status section.

   You can enable Access Gateway mode only after the switch is disabled.
4. Click Enable in the Access Gateway Mode section.
5. Click Apply.
6. Click Yes to restart the switch in Access Gateway mode

Disabling Access Gateway mode

Use the following procedure to disable Access Gateway mode.

1. Select a switch.
2. Click Switch Admin in the Manage section under Tasks.

   The Switch Administration window opens.
3. Click Disable in the Switch Status section.

   You can disable Access Gateway mode only after the switch is disabled.
4. Click Disable in the Access Gateway Mode section.
5. Click Apply.
6. Click Yes to restart the device in native switch mode.

Viewing the Access Gateway settings

You can view the effective Access Gateway settings for the selected switch. The view can be customized.

1. Click Access Gateway Devices in the Monitor section under Tasks.

   The Access Gateway Device Display window opens

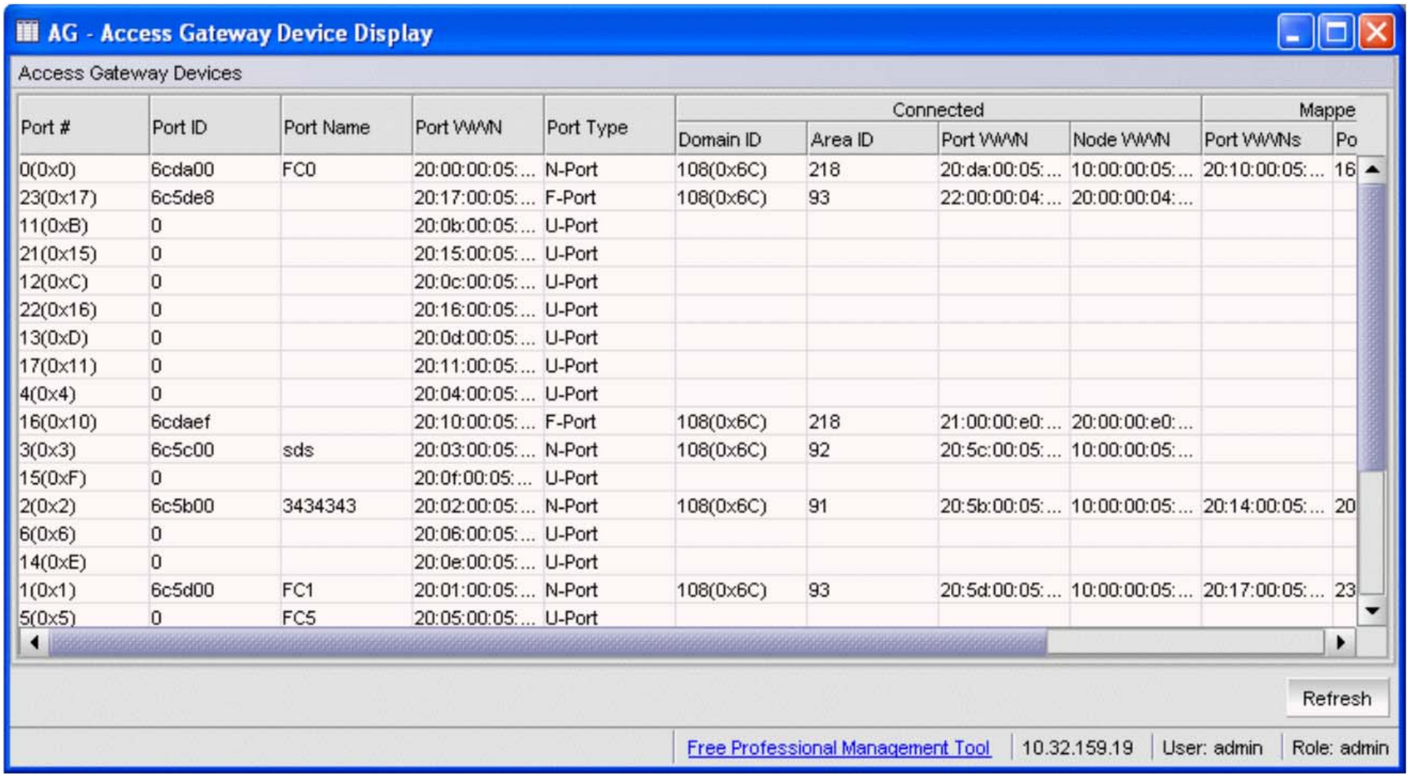

Access Gateway Devices

Port #	Port ID	Port Name	Port WWN	Port Type	Connected				Mappe	
					Domain ID	Area ID	Port WWN	Node WWN	Port WWNs	Po
0(0x0)	6cda00	FC0	20:00:00:05:...	N-Port	108(0x6C)	218	20:da:00:05:...	10:00:00:05:...	20:10:00:05:...	16
23(0x17)	6c5de8		20:17:00:05:...	F-Port	108(0x6C)	93	22:00:00:04:...	20:00:00:04:...		
11(0xB)	0		20:0b:00:05:...	U-Port						
21(0x15)	0		20:15:00:05:...	U-Port						
12(0xC)	0		20:0c:00:05:...	U-Port						
22(0x16)	0		20:16:00:05:...	U-Port						
13(0xD)	0		20:0d:00:05:...	U-Port						
17(0x11)	0		20:11:00:05:...	U-Port						
4(0x4)	0		20:04:00:05:...	U-Port						
16(0x10)	6cdaef		20:10:00:05:...	F-Port	108(0x6C)	218	21:00:00:e0:...	20:00:00:e0:...		
3(0x3)	6c5c00	sds	20:03:00:05:...	N-Port	108(0x6C)	92	20:5c:00:05:...	10:00:00:05:...		
15(0xF)	0		20:0f:00:05:...	U-Port						
2(0x2)	6c5b00	3434343	20:02:00:05:...	N-Port	108(0x6C)	91	20:5b:00:05:...	10:00:00:05:...	20:14:00:05:...	20
6(0x6)	0		20:06:00:05:...	U-Port						
14(0xE)	0		20:0e:00:05:...	U-Port						
1(0x1)	6c5d00	FC1	20:01:00:05:...	N-Port	108(0x6C)	93	20:5d:00:05:...	10:00:00:05:...	20:17:00:05:...	23
5(0x5)	0	FC5	20:05:00:05:...	U-Port						

10.32.159.19 | User: admin | Role: admin

Port configuration

You can configure the port types (N_Port, F_Port) on each individual port on an Access Gateway enabled switch. When you configure ports, you can specify a global configuration policy using the Port Configuration Policy button. By default, Advanced is selected and sets the initial defaults for port types, groups, and the F_Port-to-N_Port mappings. When the policy is Automatic, the port type assignments and mappings are configured automatically based on device and switch connections and internal load-balancing and grouping; user controls are disabled.

When you configure ports, perform the tasks in the following order:

- Configure N_Ports, if necessary.
  Use the Edit Configuration button to configure a port.
- Configure N_Port groups.
- Configure F_Port-to-N_Port mappings.

You can set up primary and secondary mappings. The secondary mapping is the N_Port to which an F_Port is mapped when the primary N_Port mapping goes offline.

Creating port groups

You can group a number of N_Ports (and its mapped F_Ports) together to connect to multiple independent fabrics or to create performance optimized ports. To group a number of ports, you must create a new port group and assign desired N_Ports to it. The N_Port grouping option is enabled by default, and all N_Ports are members of a default port group 0 (pg0). Access Gateway prevents failover of F_Ports across N_Port groups.

NOTE
If you want to distribute F_Ports among groups, you can leave all ports in the default port group 0, or you can disable N_Port grouping.

Use the following procedure to create port groups.

1. Click a port in the Switch View to open the Port Administration window.

2. Make sure that you have selected Advanced from the Port Configuration Policy drop-down list.
3. Click Configure N_Port Groups.

NOTE
Configure N_Port Groups will be disabled if you select Automatic from the Port Configuration Policy drop-down list.

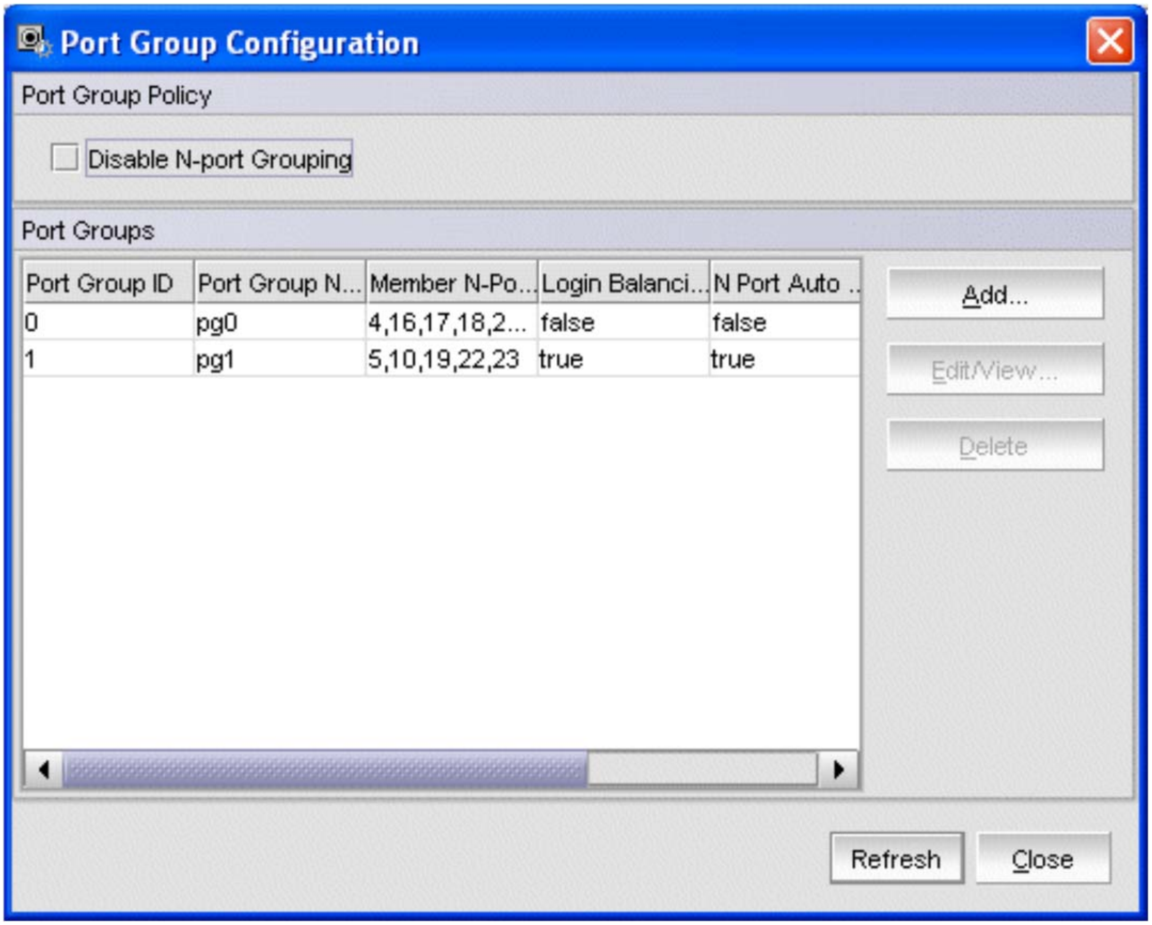

☐ Disable N-port Grouping

Port Group ID	Port Group N...	Member N-Po...	Login Balanci...	N Port Auto ..
0	pg0	4,16,17,18,2...	false	false
1	pg1	5,10,19,22,23	true	true

4. On Port Group Configuration, click Add.

The Add Port Group window displays.

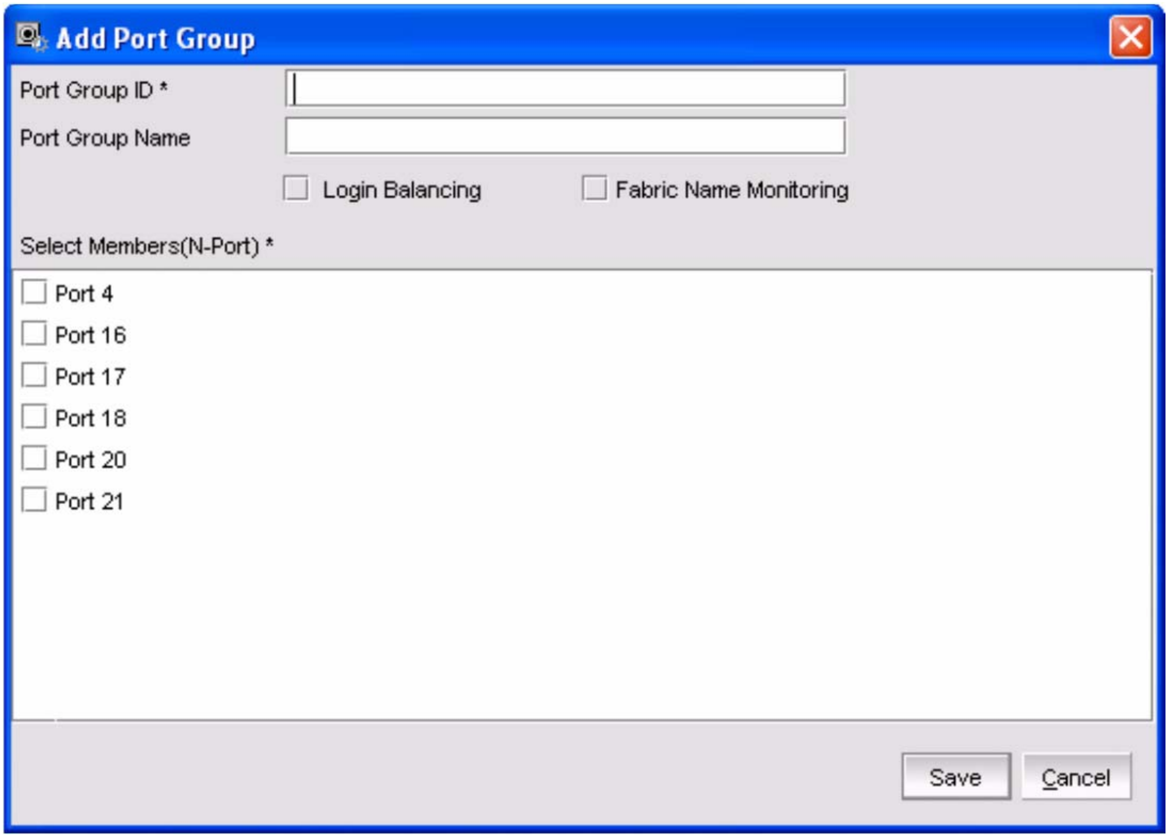

5. Enter the id for the new port group in the Port Group ID* field.
6. Enter the name for the new port group in the Port Group Name field.
7. Select the Login Balancing check box to enable login balance for the port group.
8. Select the Fabric Name Monitoring check box to manually configure the managed fabric name monitoring.
9. Under the Select Members (N-Port)* section, select the required ports you want to group.
10. Click Save.

Editing or Viewing port groups

Use the following procedure to edit port groups.

1. Click a port in the Switch View to open the Port Administration window.
2. Click Configure N_Port Groups.
3. On Port Group Configuration dialog box, select the group that you want to edit and then click Edit/View.

   The Edit/View Port Group window displays.

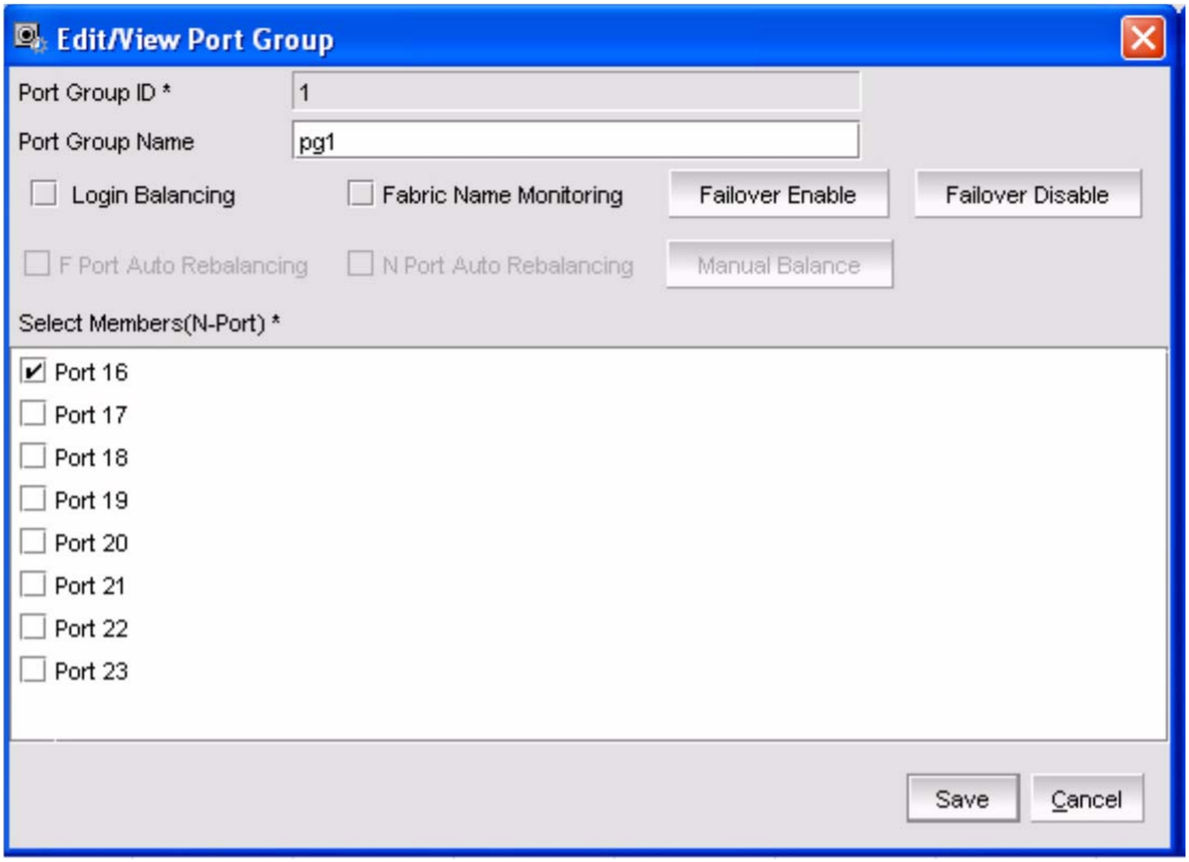

Edit/View Port Group

Port Group ID * 1

Port Group Name pg1

□ Login Balancing □ Fabric Name Monitoring

□ F Port Auto Rebalancing □ N Port Auto Rebalancing

Select Members(N-Port) *

☑ Port 16
□ Port 17
□ Port 18
□ Port 19
□ Port 20
□ Port 21
□ Port 22
□ Port 23

4. Edit the name of the port group in the Port Group Name field.
5. Select the Login Balancing check box and the Fabric Name Monitoring check box if you want to enable these features. Clear the check boxes to disable these features.

   On selecting Login Balancing check box, F Port Auto Rebalancing and N-Port Auto Rebalancing check boxes and Manual Balancing button gets enabled.
6. Click Failover Enable.

   A confirmation dialog box displays. Click Yes to enable failover to all the ports in the port group or click No if you do not want to enable failover.
7. Click Failover Disable.

   A confirmation dialog box displays. Click Yes to disable failover to all the ports in the port group or click No if you do want to disable failover.
8. Under the Select Members(N-Port)* section, select the required ports you want to group and clear the check boxes for the ports you want to remove from the port group.
9. Click Save.
10. Click Close on the Port Group Configuration dialog box.

Deleting port groups

Use the following procedure to delete port groups.

NOTE
You cannot delete the default port group 0.

1. Click a port in the Switch View to open the Port Administration window.
2. Click Configure N_Port Groups.
3. On Port Group Configuration dialog box, select the group that you want to delete and then click

Delete.

A confirmation dialog box displays.

4. Click Yes to confirm the action.
5. Click Close.

Access Gateway policy modification

Although you can control a number of policies on switches in Access Gateway mode, Web Tools only provides the ability to enable and disable the policies. For more information on these policies please refer to *Access Gateway Administrator's Guide*.

Path Failover and Failback policies

The Path Failover and Failback policies determine the behavior of the F_Port if the primary mapped N_Port they are mapped to goes offline or is disabled. The Path Failover and failback policies are attributes of the N_Port. By default, the Path Failover and Failback policies are enabled for all N_Ports.

Modifying Path Failover and Failback policies

Use the following procedure to modify Path Failover and Failback policies.

1. Click a port in the Switch View to open the Port Administration window.
2. Select the N_Port for which you want to modify the policy.
3. Click Edit Configuration.

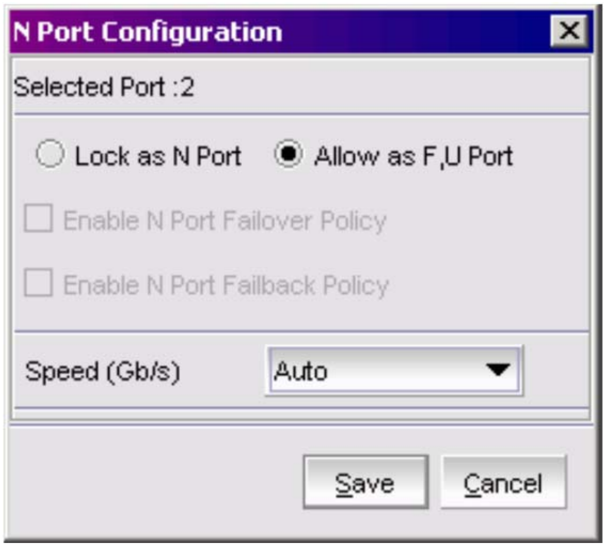

N Port Configuration

Selected Port :2

○ Lock as N Port ● Allow as F,U Port

□ Enable N Port Failover Policy

□ Enable N Port Failback Policy

Speed (Gb/s) Auto

4. Select the appropriate check box to modify the policy.
5. Click Save.

Enabling the Automatic Port Configuration policy

The Automatic Port Configuration (APC) policy is a global configuration policy for a switch in Access Gateway mode. By default, this policy is disabled. If you created an N_Port grouping and switching over to the automatic mode, those port groups will be lost. After you enable the APC policy, you cannot define custom port type configurations, port mappings, Path Failover, and Failback settings.

Use the following procedure to enable auto rebalancing from the Switch Administration window.

1. Click a port in the Switch View to open the Port Administration window.
2. Select Automatic from the Port Configuration Policy drop-down list.

NOTE
When Port Configuration Policy is set to Advanced, you can enable the auto rebalancing options from the Configure N_Port Groups dialog box through the Port Administration window.

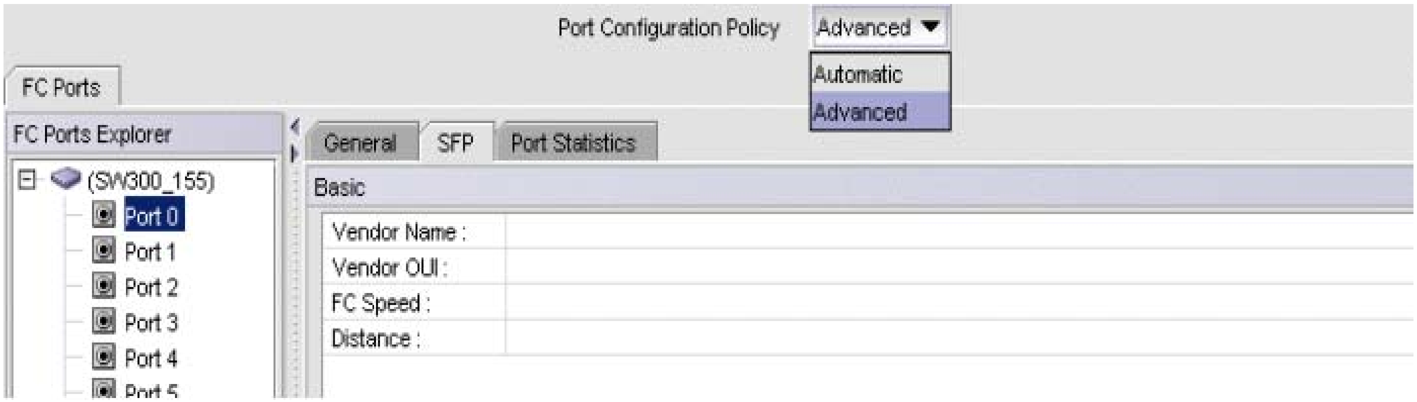

Port Configuration Policy: Advanced

3. Click Yes in the confirmation window.
4. In the Switch Explorer window select Switch Admin.

   The Switch Administration window displays.

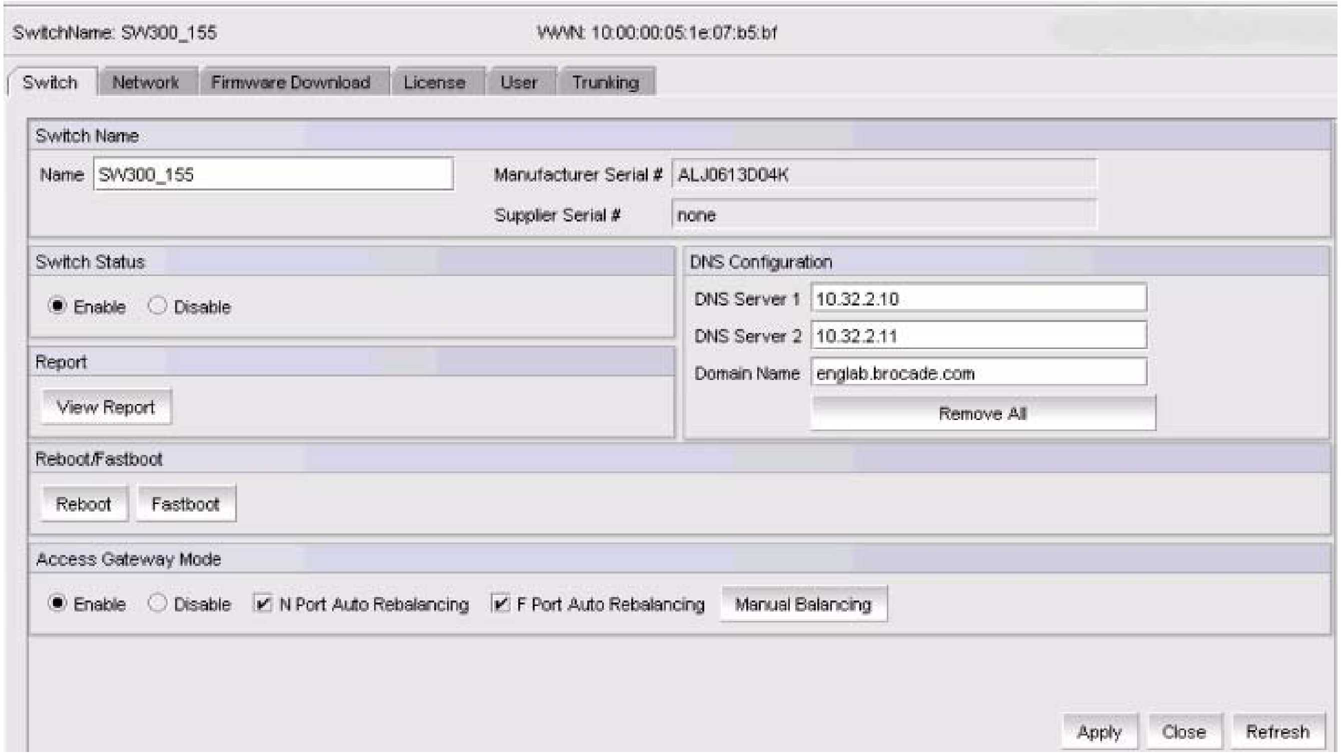

SwitchName: SW300_155 WWN: 10:00:00:05:1e:07:b5:bf

Switch Name
Name: SW300_155
Manufacturer Serial #: ALJ0613D04K
Supplier Serial #: none

Switch Status: ● Enable ○ Disable

DNS Configuration
DNS Server 1: 10.32.2.10
DNS Server 2: 10.32.2.11
Domain Name: englab.brocade.com

Access Gateway Mode: ● Enable ○ Disable ☑ N Port Auto Rebalancing ☑ F Port Auto Rebalancing

5. Click Refresh.
6. Under the Access Gateway Mode section, do the following:
   - Select the N Port Auto Rebalancing check box to enable N_Port rebalancing.
   - Select F Port Auto Rebalancing check box to enable F_Port rebalancing.

   - Click Manual Balancing and a confirmation dialog box displays. Click Yes to change F Port - N Port Mapping or click No to cancel the changes. Click Apply to apply the changes.

Chapter 10: Brocade 8000

Brocade 8000 overview

The Brocade 8000 is a 24-port 10 GbE line-rate, low latency lossless Converged Enhanced Ethernet (CEE) and an 8-port auto-sensing 1, 2, 4, or 8 Gbps Fibre Channel (FC) switch that delivers the latest Brocade ASIC technology and architecture for Fibre Channel Storage Area Networks (SANs). The fully-configured Brocade 8000 enables the Fibre Channel over Ethernet (FCoE) protocol and is a high performance 8 Gbps Fibre Channel switch designed for the needs of enterprise environments that require a high-port footprint for port aggregation and desire the simplified management environment that comes with reducing the total number of domains to manage.

The Brocade 8000 is also available as a CEE-only base model switch with no FC capability. The CEE-only switch can later be upgraded to full FCoE configuration with the purchase of an additional license. This version of the switch is available only with Fabric OS v6.3.0 and above.
The Brocade 8000 supplies Reliability, Availability, and Serviceability (RAS) performance and scalability requirements of an enterprise switch along with the interoperability and ease-of-use advantages found only in the Brocade product family.

The Brocade 8000 can also be configured in Access Gateway mode that lets you configure your Enterprise fabric to handle additional N_Ports instead of domains. By reducing the number of domain IDs and ports you simplify configuration and management in a large fabric.
Switch modules in AG mode are logically transparent to the host and the fabric. You can increase the number of hosts that have access to the fabric without increasing the number of switch module

The Brocade 8000 offers the following features and capabilities:

- A system motherboard that features a Freescale MPC8548 Reduced Instruction Set Computer (RISC) CPU running at 1.3 GHz with integrated peripherals, and that provides high performance with low power consumption.
- An RJ45 Ethernet management port
- A USB port that provides storage for firmware updates, output of the supportSave command and storage for configuration uploads and downloads.
- Two hot-swappable, redundant power supply FRUs.
- Three hot-swappable fan assembly FRUs in an N+1 configuration to provide hardware-redundant cooling.
- Extensive diagnostics and system-monitoring capabilities for enhanced high Reliability, Availability, and Serviceability (RAS).
- FCoE to FC latency of 1670 nanoseconds.

FCoE and Layer 2 capabilities

The Brocade 8000 has the following capabilities for the Ethernet functions:

- 24 ports 10 GbE CEE.
- Low latency, lossless, deterministic interconnect required for FCoE.
- FCoE support along with Fabric Provided MAC Address (FPMA) discovery. FOS delivers these features and also enables support for Priority-based Flow Control (802.1Qbb). Data Center Bridging eXchange (DCBX) - Capabilities Exchange and Enhanced Transmission Selection (802.1Qaz) delivers the lossless and deterministic FCoE requirement.
- Enables hardware-assisted MAC learning and aging.
- Support for 32K MAC addresses and 4K Vlans.
- Support for Layer 2 protocols STP/MSTP/RSTP (802.1q) and Link Aggregation (802.1ad).
- Brocade 10G SFP+ (SR and LR) and Brocade-branded Twinax copper cables.
- CEE port to CEE port latency of 570 nanoseconds (same ASIC) and 1050 nanoseconds (different ASIC).

A number of FCoE enhancements have been made with FOS 7.0.0 and higher. For more on these please refer to Fabric OS Adminstrator's Guide.

Fibre Channel capabilities

The fully-configured Brocade 8000 offers the following features for the Fibre Channel functions (license required):

- Up to 8 auto-sensing ports of high-performance 8 Gbps technology in a single domain.
- Full 1:1 subscription on 8 Gbps ports.
- 1, 2, 4, and 8 Gbps auto-sensing Fibre Channel switch and router ports.
- Universal ports self-configure as E, F, M, or FL ports.
- Inter-Switch Link (ISL) Trunking (licensable), which allows up to eight ports (at 1, 2, 4, or 8 Gbps speeds) between a pair of switches combined to form a single, logical ISL with a speed of up to 64 Gbps (128 Gbps full duplex) for optimal bandwidth utilization and load balancing.
- Dynamic Path Selection (DPS), which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- Brocade-branded SFP optical transceivers that support any combination of Short Wavelength (SWL) and Long Wavelength (LWL) optical media among the switch ports.
- Support for unicast, multicast (255 groups), and broadcast data traffic types.
- Brocade FOS, which delivers distributed intelligence throughout the network and enables a wide range of value-added applications including Brocade Advanced Web Tools and Brocade Zoning. Optional Fabric Services include: Adaptive Networking with QoS, Brocade Extended Fabrics, Brocade Enhanced Group Management, Brocade Fabric Watch, ISL Trunking, and End-to-End Performance Monitoring (APM).
- Port-to-port latency minimized to 700 nanoseconds through the use of cut-through frame routing at 8 Gbps.

Port side of Brocade 8000

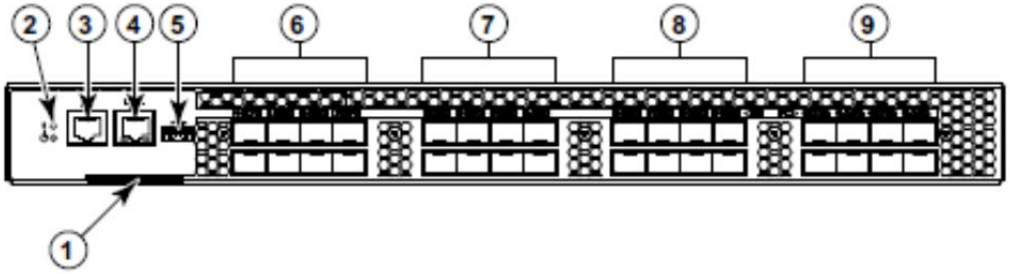

1	Switch ID pull-out tab	5	USB port
2	System status LED (top)	6	GbE ports 0-7
	System power LED (bottom)	7	GbE ports 8-15
3	Serial console port	8	GbE ports 16-23
4	Management Ethernet port	9	FC ports 0-7

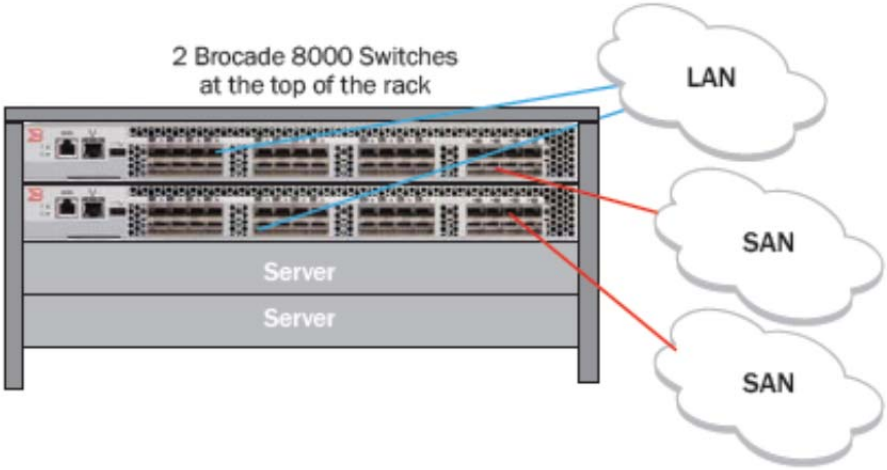

The Brocade 8000 runs traditional Fabric OS (FOS) software and can be managed using the same tools traditionally used for SAN management. Using the FOS Command Line Interface (CLI), administrators have access to all commands and utilities common to other Brocade switches. In addition, FOS software on the Brocade 8000 enables Brocade Web Tools to support the following features for configuring and managing a Converged Ethernet Network:

- CEE interface display and configuration
- FCoE trunk display and configuration
- CEE configuration including link aggregation (LACP), Virtual LANs (VLANs), Quality of Service (QoS), and LLDP (Link Layer Discovery Protocol)/ DCBX protocol (Data Center Bridging eXchange)
- FCoE login groups

CEE Command Line Interface

The Brocade 8000 introduces a new CLI designed to support the management of CEE and L2 Ethernet switching functionality. The CEE CLI uses an industry-standard hierarchical shell familiar to Ethernet/IP networking administrators.

All conventional port-related Fabric OS CLI commands are only applicable to Fibre Channel. These commands have no knowledge of the Ethernet ports. The CEE features and CEE ports can only be configured through the CEE CLI interface which is accessed by entering the cmsh command from the Fabric OS shell.

The system starts up with the default Fabric OS configuration and the CEE startup configuration.

After logging in you are in the Fabric OS shell. Some Fabric OS commands are available in the CEE shell. Enter the fos ? command at the CEE CLI. Privileged EXEC mode command prompt to view the available Fabric OS commands. The traditional Fabric OS command help found in the Fabric OS shell is not available through the CEE shell.

Accessing the CEE CLI through the console or Telnet

The procedure to access the CEE 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:admin> cmsh
switch#
```

To return to the Fabric OS CLI, enter the following command.

```
switch#exit
switch:admin>
```

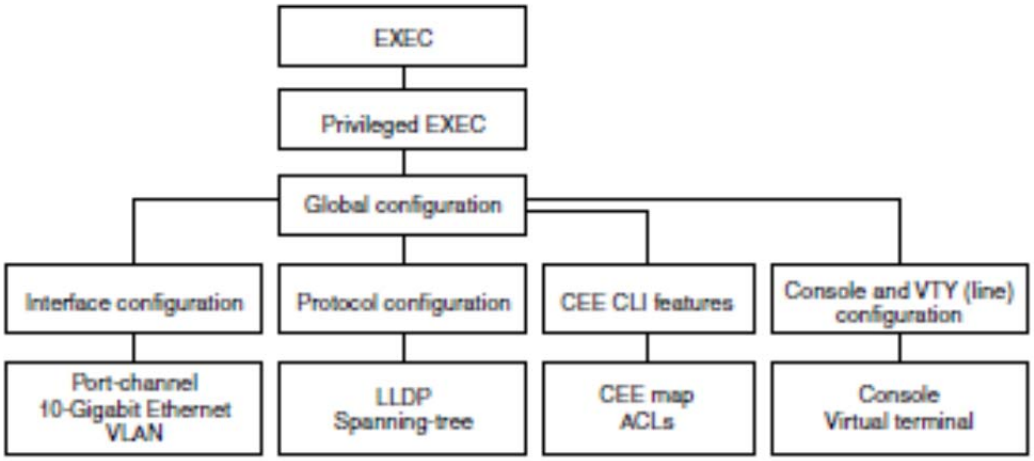

NOTE
The CEE configuration is not affected by configUpload and configDownload commands entered in the Fabric OS shell.

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, you must use either the write memory command or the copy command to commit the *running-config* file to the *startup--config* file.
Saving configuration changes with the copy command
Perform this task from Privileged EXEC mode.

Enter the copy command to save the *running-config* file to the *startup-config* file.

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

Saving configuration changes with the write command

Perform this task from Privileged EXEC mode.

Enter the write memory command to save the *running-config* file to the *startup-config* file.

```
switch# write memory
Overwrite the startup config file (y/n): y
Building configuration...
```

CEE Commands

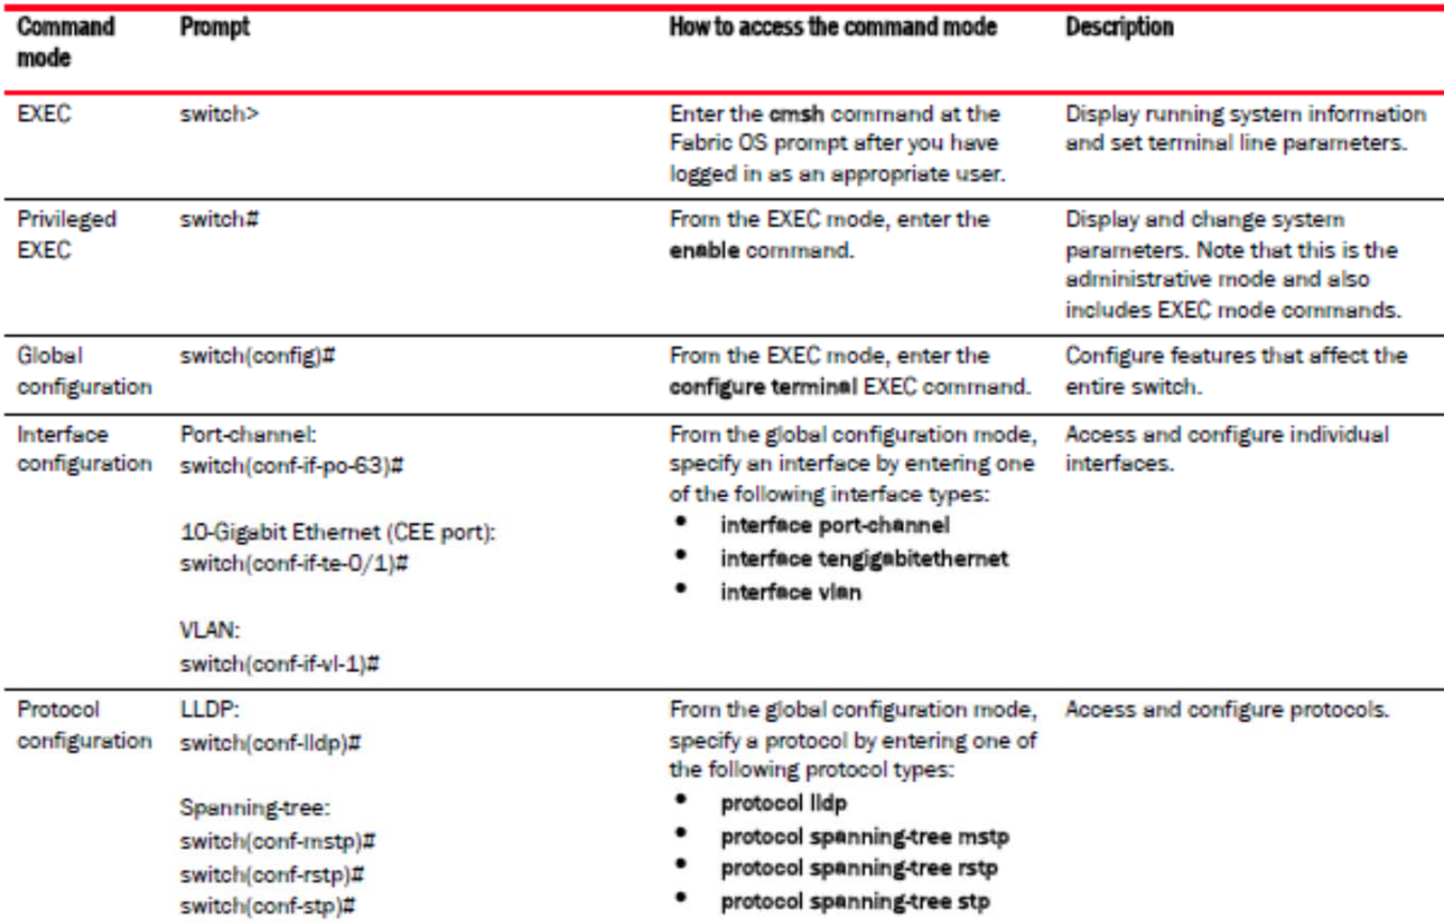

Command mode	Prompt	How to access the command mode	Description
EXEC	switch>	Enter the **cmsh** command at the Fabric OS prompt after you have logged in as an appropriate user.	Display running system information and set terminal line parameters.
Privileged EXEC	switch#	From the EXEC mode, enter the **enable** command.	Display and change system parameters. Note that this is the administrative mode and also includes EXEC mode commands.
Global configuration	switch(config)#	From the EXEC mode, enter the **configure terminal** EXEC command.	Configure features that affect the entire switch.
Interface configuration	Port-channel: switch(conf-if-po-63)# 10-Gigabit Ethernet (CEE port): switch(conf-if-te-0/1)# VLAN: switch(conf-if-vl-1)#	From the global configuration mode, specify an interface by entering one of the following interface types: • **interface port-channel** • **interface tengigabitethernet** • **interface vlan**	Access and configure individual interfaces.
Protocol configuration	LLDP: switch(conf-lldp)# Spanning-tree: switch(conf-mstp)# switch(conf-rstp)# switch(conf-stp)#	From the global configuration mode, specify a protocol by entering one of the following protocol types: • **protocol lldp** • **protocol spanning-tree mstp** • **protocol spanning-tree rstp** • **protocol spanning-tree stp**	Access and configure protocols.

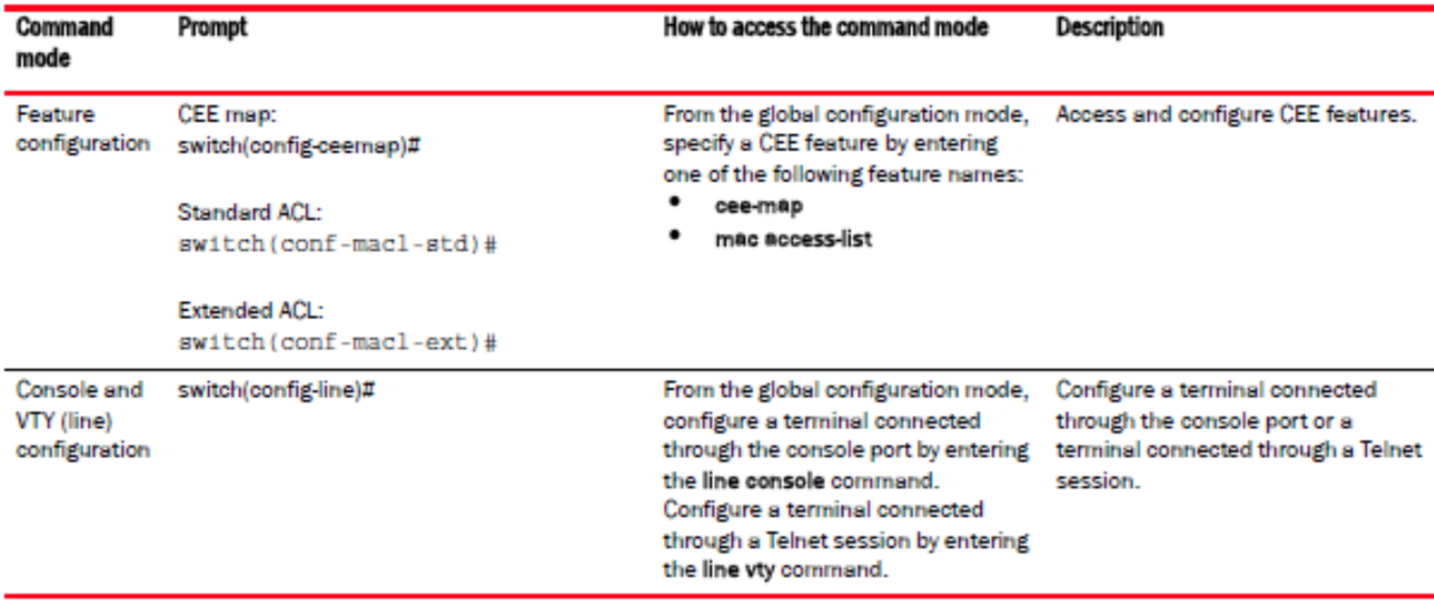

Command mode	Prompt	How to access the command mode	Description

Feature configuration	CEE map:
	switch(config-ceemap)#
	Standard ACL:
	`switch(conf-macl-std)#`
	Extended ACL:
	`switch(conf-macl-ext)#`
		From the global configuration mode, specify a CEE feature by entering one of the following feature names:
		• **cee-map**
		• **mac access-list**
			Access and configure CEE features.

Console and VTY (line) configuration	switch(config-line)#	From the global configuration mode, configure a terminal connected through the console port by entering the **line console** command. Configure a terminal connected through a Telnet session by entering the **line vty** command.	Configure a terminal connected through the console port or a terminal connected through a Telnet session.

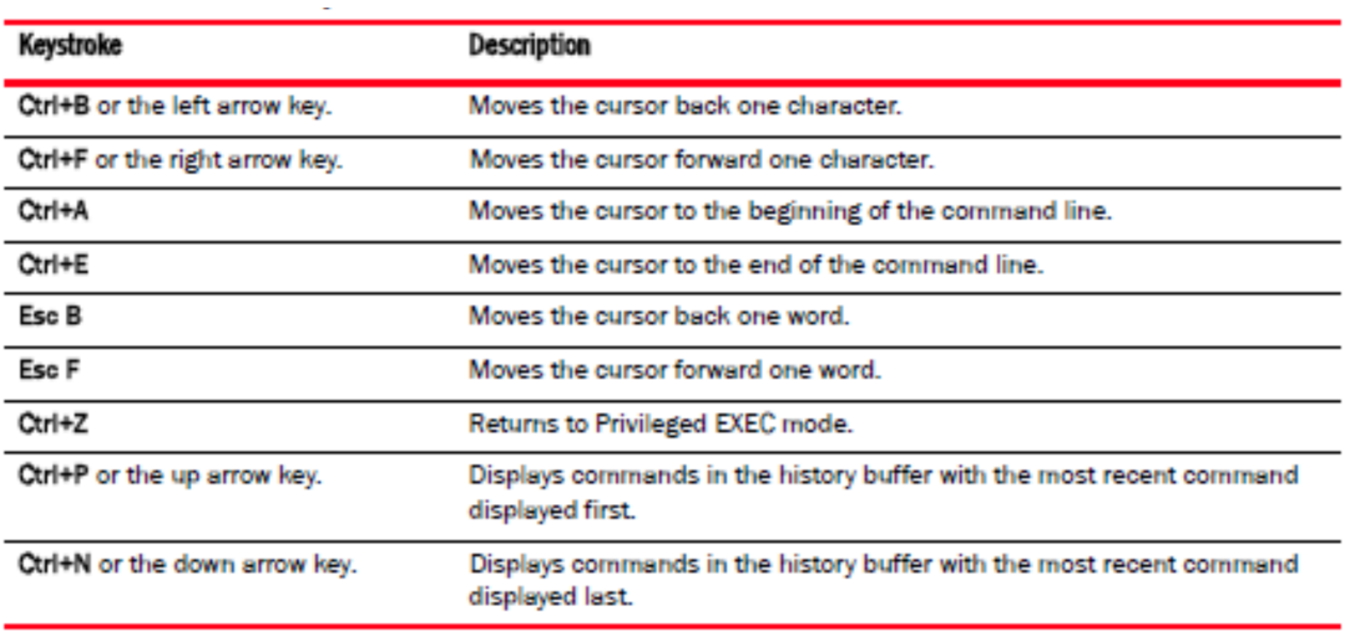

Keystroke	Description
Ctrl+B or the left arrow key.	Moves the cursor back one character.
Ctrl+F or the right arrow key.	Moves the cursor forward one character.
Ctrl+A	Moves the cursor to the beginning of the command line.
Ctrl+E	Moves the cursor to the end of the command line.
Esc B	Moves the cursor 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.

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 the EXEC or Privileged EXEC mode.

For example, if you are configuring an 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.

However, by using the do command with the dir command you can ignore the need to change configuration modes, as shown in the example below.

```
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
switch(conf-lldp)#
```

CEE Integrations

FC SANs are typically deployed in a core-edge topology with servers connecting to edge switches in the fabric. Since the Brocade 8000 FC switching module operates with the same features and functionality of a regular FC switch, this topology is preserved when the Brocade 8000 switch is introduced into the fabric. The Brocade 8000 switch can be treated as just another edge switch connecting to the core FC infrastructure. The only difference is that servers are directly attached using a CNA supporting the FCoE protocol instead of an HBA supporting the FC protocol.

Connecting the Brocade 8000 switch to an existing FC SAN follows the same process as adding a new FC edge switch into a SAN. Most SAN environments include redundant fabrics (A and B). A typical installation involves connecting a Brocade 8000 switch to Fabric A, verifying stability, and then installing a second Brocade 8000 switch into Fabric B.

FCoE devices log in to one of the six FCoE ports on the Brocade 8000 switch. The FCoE ports provide FC services to FCoE initiators and enable bridging between FCoE initiators and FC targets. FCoE ports differ from regular FC ports in that they are not directly associated with an external physical port on the switch. Instead, each FCoE port supports up to four logical traffic paths.

Brocade's implementation of FCoE on the Brocade 8000 switch provides integral NPIV support so that multiple FCoE initiators can log in to a single FCoE interface.
When a CNA logs into the fabric, it is assigned a new MAC address using a function called Fabric Provided MAC Address (FPMA). This address is used for all FCoE communication. The first three bytes of the MAC address are provided by the FC-MAP and the last three bytes are determined by the FCID. The VF_Port or FC entity that the CNA logs in to determines the FCID.

Integrating a Brocade 8000 switch on a SAN

Perform the following process to install a new Brocade 8000 switch.

1. On the Brocade 8000 switch, verify that the Zone database is empty and change the domain ID to a unique number. If there are any non-default fabric configuration changes in the existing fabric, ensure that these are also configured on the new switch. For details, see the "Administering Advanced Zoning" and "Performing Basic Configuration Tasks-Domain IDs" sections of the *Fabric OS Administrator's Guide*.
2. Power off the Brocade 8000 switch and connect the Inter-Switch Link (ISL) cables to the core FC switch or director.

NOTE
Connecting a new Brocade 8000 switch to the fabric while it is powered off ensures that reconfiguration will not occur.

3. Power on the Brocade 8000 switch and verify that the ISLs are online and the fabric is merged.
4. Check to make sure the existing Zone database files for the fabric were copied over to the Brocade 8000 switch.
5. Use the FOS CLI command **nsShow** to display any FCoE or FC devices connected to the switch. Any CNAs should be able to log in to the fabric and can be zoned using standard management tools, including the FOS CLI or Web Tools.
6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.
7. Repeat this procedure for the second Brocade 8000 switch attached to Fabric B.

CEE and LAN integration

Because Brocade FCoE hardware is IEEE 802.1Q compliant, it easily integrates into the existing LAN infrastructure in a variety of data center network topologies. In a typical installation, the Brocade 8000 switch acts as an access layer switch connecting to a distribution or core layer switch in the LAN.

The following steps are the basic process for integrating the Brocade FCoE hardware on a LAN.

1. Create a **CEE map** for the Brocade FCoE hardware to define the traffic types on your LAN.
2. Define your present **DCBX setup** for TLV.
3. Configure the Brocade FCoE hardware for your present type of **STP**.
4. Assign the Brocade FCoE hardware to the correct **VLAN membership** and **VLAN group.**
5. Assign the **CEE interfaces** on the Brocade FCoE hardware to the correct VLAN groups.
6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

CEE map attributes

The following information is needed for CEE configuration:

- The types of traffic flowing through an interface, FCoE, TCP/IP, and so on.
- The minimum bandwidth required for each traffic type.
- Which traffic type needs lossless behavior.

Brocade uses CEE Maps to simplify the configuration of QoS and flow control. Users assign different priorities to different traffic types and enable lossless connectivity. A CEE map configures two features: Enhanced Transmission Selection (ETS) and Priority Flow Control (PFC).

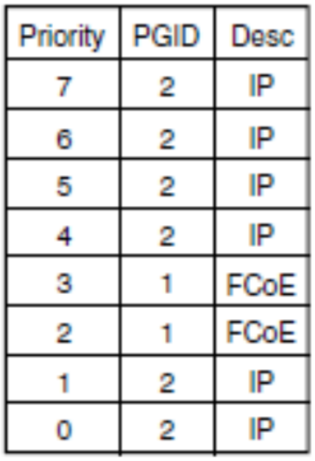

Priority	PGID	Desc
7	2	IP
6	2	IP
5	2	IP
4	2	IP
3	1	FCoE
2	1	FCoE
1	2	IP
0	2	IP

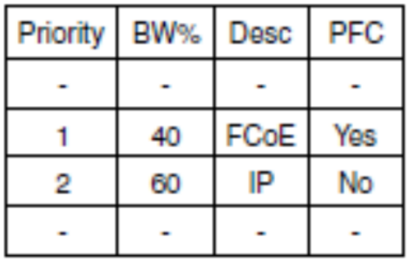

Priority	BW%	Desc	PFC
-	-	-	-
1	40	FCoE	Yes
2	60	IP	No
-	-	-	-

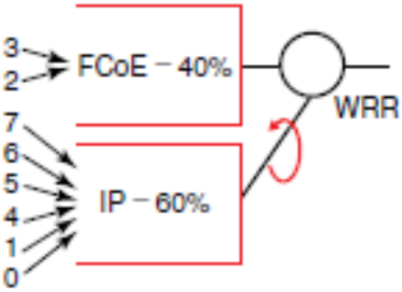

For the given example, a CEE Map named “srvgroup” is created using the following syntax.
Perform the following steps in global configuration mode.

1. Define the name of the CEE map

Example of setting the CEE map name as “srvgroup”.

```
switch(config)#cee-map srvgroup
```

2. Specify the traffic requirements for each PGID using priority-group-table

Example of setting two traffic requirements.

```
switch(config)#priority-group-table 1 weight 40 pfc
switch(config)#priority-group-table 2 weight 60
```

3. The priority-table is then used to specify which priorities are mapped to which PGID. The priorities are defined from lowest to highest.

Example of setting the priority mappings.

```
switch(config)#priority-table 2 2 1 1 2 2 2 2
```

Enter the copy command to save the *running-config* file to the *startup-config* file.

```
switch(config)#end
switch#copy running-config startup-config
```

Configuring DCBX

DCBX (Data Center Bridging eXchange Protocol) runs on CEE links and is an extension of the Link Layer Discovery Protocol (LLDP). The primary goal of DCBX is to allow the discovery of CEE-capable hosts and switches and allow CEE-specific parameters—such as those for ETS and PFC—to be sent before the link is shared. DCBX parameters use a type-length-value (TLV) format. By default, DCBX is turned on, but there are two TLVs that must be enabled to support FCoE on a CEE link:

- dcbx-fcoe-app-tlv – IEEE Data Center Bridging eXchange FCoE Application TLV.

- dcbx-fcoe-logical-link-tlv - IEEE Data Center Bridging eXchange FCoE Logical Link TLV. The presence of this TLV declares that the FCoE part of the converged link is UP.

To configure the TLVs for DCBX, perform the following steps in global configuration mode.

1. Set the protocol type to LLDP.

```
switch(config)#protocol lldp
```

2. Activate the protocol.

```
switch(conf-lldp)#no disable
```

3. Activate the TLV formats using the advertise command in Protocol LLDP Configuration Mode.

```
switch(conf-lldp)#advertise dcbx-fcoe-app-tlv
switch(conf-lldp)#advertise dcbx-fcoe-logical-link-tlv
```

4. Enter the copy command to save the *running-config* file to the *startup-config* file.

```
switch(conf-lldp)#exit
switch(config)#end
switch#copy running-config startup-config
```

Configuring Spanning Tree Protocol

Spanning Tree Protocol is a mechanism to detect and avoid loops in Ethernet networks by establishing a fixed path between all the switches in a LAN. The Brocade FCoE hardware supports three spanning tree variations: Standard Spanning Tree (STP), Rapid Spanning Tree (RSTP), and Multiple Instance Spanning Tree (MSTP).
It is best practice that an access layer switch, such as the Brocade 8000 switch, does not become the root switch. Changing the bridge or STP priority helps to ensure that this does not occur. The example below performed from the CEE CLI configures the Brocade 8000 switch for RSTP and sets the bridge priority to the highest value ensuring it will not become the root switch in an existing LAN.

To configure RSTP, perform the following steps in global configuration mode.

1. Configure the Brocade 8000 switch for RSTP.

```
switch(config)#protocol spanning-tree rstp
```

2. Set the bridge priority to the highest value so it does not become the root switch in an existing LAN.

```
switch(conf-rstp)#bridge-priority 61440
```

3. Enter the copy command to save the *running-config* file to the *startup-config* file.

```
switch(conf-rstp)#exit
switch(config)#end
switch#copy running-config startup-config
```

Configuring VLAN Membership

IEEE 802.1q Virtual LANs (VLANs) provide the capability to overlay the physical network with multiple virtual networks. VLANs allow network traffic isolation into separate virtual networks reducing the size of administrative and broadcast domains.

A VLAN contains end stations that have a common set of requirements which can be in independent physical locations. 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 subnets and all the end stations in a particular IP subnet belong to the same VLAN.

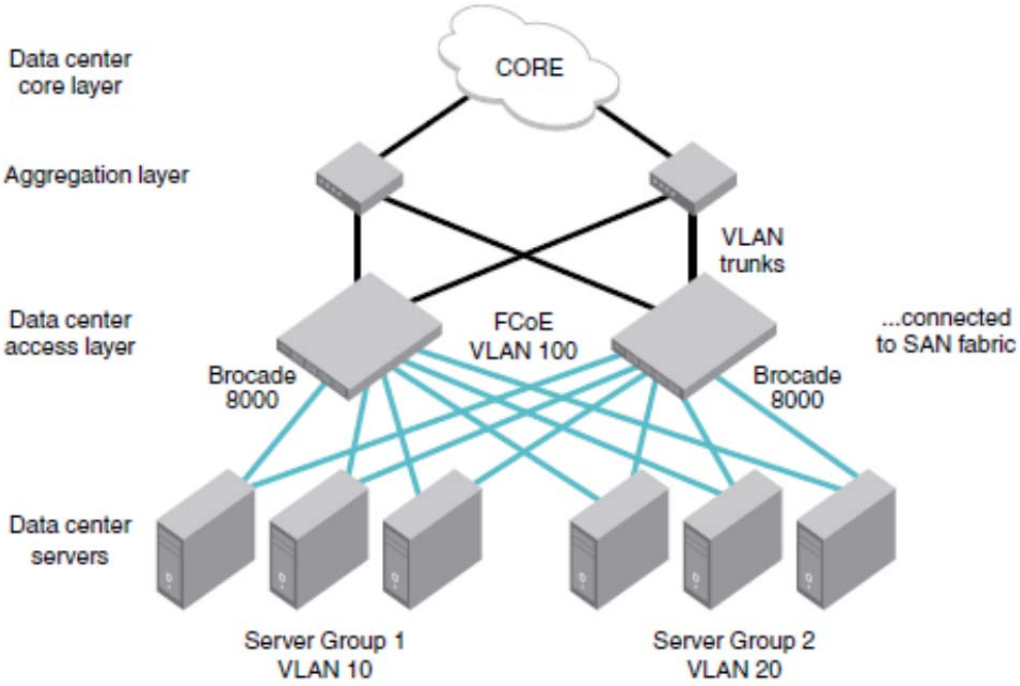

In the sample network shown above, there are three VLANs: VLAN 100, VLAN 10, and VLAN 20. VLAN 10 and 20 are used to isolate the L2 traffic from the two server groups. These VLANs carry IP traffic from the servers to the data center LAN. Any routing between these VLANs is performed at the distribution layer of the network. VLAN 100 is a special VLAN used for FCoE traffic between the servers and storage connected to the Fibre Channel fabric and must be configured as an FCoE Forwarder (FCF). Only FCF-capable VLANs can carry FCoE traffic.

In addition to creating a special VLAN for FCoE traffic, VLAN classifiers are applied to incoming EtherTypes for FCoE Initiation Protocol (FIP) and FCoE. VLAN classifiers are rules used to dynamically classify Ethernet frames on an untagged interface to VLANs.

To configure VLAN membership, perform the following steps in global configuration mode.

1. Create the VLAN interfaces on the Brocade FCoE hardware using the CEE CLI.
   Example of creating two VLAN interfaces and assigning each one to a server group.

```
switch(config)#interface vlan 10
switch-cmsh(conf-if-vl-10)#description server group 1
switch(config)#interface vlan 20
switch-cmsh(conf-if-vl-20)#description server group 2
switch(config)#interface vlan 100
switch-cmsh(conf-if-vl-100)#description FCoE VLAN
```

```
switch-cmsh(conf-if-vl-100)#fcf forward
```

2. Create VLAN rules and a VLAN classifier group for these two EtherTypes.
   Example of creating VLAN rules and classifier groups.

```
switch(config)#vlan classifier rule 1 proto fip encap ethv2
switch(config)#vlan classifier rule 2 proto fcoe encap ethv2
switch(config)#vlan classifier group 1 add rule 1
switch(config)#vlan classifier group 1 add rule 2
```

3. Apply the VLAN classifier group to any CEE interface.
4. Enter the copy command to save the *running-config* file to the *startup-config* file.

```
switch(config)#end
switch#copy running-config startup-config
```

Configuring the CEE Interfaces

Traffic from downstream CEE interfaces can be assigned to a VLAN using several methods:

- The VLAN tag contained in the incoming frame
- The VLAN classifiers
- The Port-VLAN ID (PVID)

Because the Ethernet uplink ports from the Brocade FCoE hardware to the distribution layer switches will carry traffic for multiple VLANs, they are configured as 802.1q trunk ports.

The downstream CEE ports connected to the server CNAs are configured as access ports with a PVID of either 10 or 20. The VLAN classifier group created for the FIP and FCoE EtherTypes must be applied to the interfaces in order to place FCoE traffic on the correct VLAN. The CEE map is also applied to the interface.

To configure the CEE interfaces, perform the following steps in global configuration mode.

1. Assign VLANs to the uplink Ethernet port.
   You must repeat this step for all uplink interfaces.
   Example of assigning VLAN 10 and VLAN 20 to the uplink Ethernet port.

```
switch(config)#interface TenGigabitEthernet 0/1
switch(conf-if-te-0/1)#switchport
switch(conf-if-te-0/1)#switchport mode trunk
switch(conf-if-te-0/1)#switchport trunk allowed vlan add 10
switch(conf-if-te-0/1)#switchport trunk allowed vlan add 20
switch(conf-if-te-0/1)#no shutdown
```

2. Apply the VLAN classifier group to the interfaces.
   Example of applying a VLAN classifier group 1 to the interfaces.

```
switch(config)#interface TenGigabitEthernet 0/10
switch(conf-if-te-0/1)#switchport
switch(conf-if-te-0/1)#switchport mode access
switch(conf-if-te-0/1)#switchport access vlan 10
switch(conf-if-te-0/1)#vlan classifier activate group 1 vlan 100
switch(conf-if-te-0/1)#no shutdown
```

3. Apply the CEE map to the interfaces.
   Example of setting the map name to srvgroup.

```
switch(conf-if-te-0/1)#cee srvgroup
```

4. Enter the copy command to save the *running-config* file to the *startup-config* file.

```
switch(conf-if-te-0/1)#exit
switch(config)#end
switch#copy running-config startup-config
```

Server connections to the Brocade 8000 switch

Converged Network Adapters (CNAs) support FCoE and Ethernet LAN communication over the same cable from the server to a CEE switch, such as the Brocade 8000 switch as shown in figure. The CNA is presented to the host operating system as both an Ethernet NIC and a Fibre Channel HBA so that network configuration and server management practices do not change.

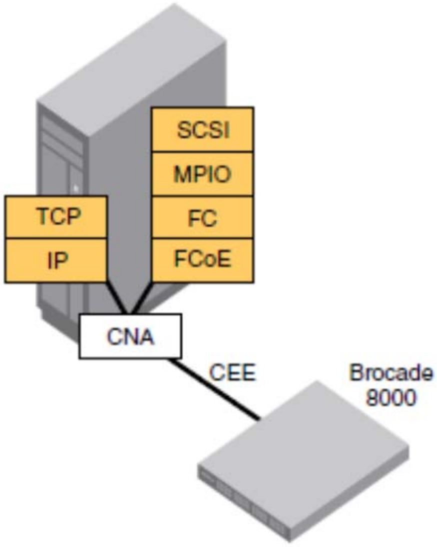

The CNA supports CEE features required to support lossless connectivity and QoS of different traffic types. Although modification of parameters is possible with some CNAs, most adapters are set up in a “Willing” mode, meaning that they automatically accept CEE configurations for QoS and PFC from the connected switch using the DCBX protocol.

Fibre Channel configuration for the CNA

The CNA discovers storage on the FC SAN and presents LUNs to the operating system in the same manner as an HBA. The same multipathing software needed for high availability in a traditional SAN can be used in a converged network.

Ethernet configuration for the CNA

Most CNAs support some type of Network Teaming or Link Aggregation protocol to allow the use of multiple ports in parallel, to improve performance or create redundancy for higher availability. For highest availability it is always recommended that you install two CNAs into a server and connect each to a different Brocade 8000 switch.

Minimum CEE configuration to allow FCoE traffic flow

The following process shows the minimum configuration steps required to run FCoE on the Brocade 8000 switch. Treat the sample code for each step as a single CLI batch file.

To set the minimum CEE configuration, perform the following steps in global configuration mode.

1. Configure the CEE interface as a Layer 2 switch port.
   Example of configuring the switch port as a 10-Gigabit Ethernet interface.

```
switch(config)#interface tengigabitethernet 0/0
switch(config-if)#switchport
switch(config-if)#no shutdown
switch(config-if)#exit
switch(config)#end
```

2. Create an FCoE VLAN and add an interface to it.
   Example of creating a FCoE VLAN and adding a single interface.

```
switch(config)#vlan classifier rule 1 proto fcoe encap ethv2
switch(config)#vlan classifier rule 2 proto fip encap ethv2
switch(config)#vlan classifier group 1 add rule 1
switch(config)#vlan classifier group 1 add rule 2
switch(config)#interface vlan 1002
switch(conf-if-vl-1002 )#fcf forward
switch(conf-if-vl-1002 )#interface tengigabitethernet 0/0
switch(config-if-te-0/0)#switchport
switch(config-if-te-0/0)#switchport mode converged
switch(config-if-te-0/0)#switchport mode converged allowed vlan add 1002
switch(config-if-te-0/0)#vlan classifier activate group 1 vlan 1002
switch(config-if-te-0/0)#cee default
switch(config-if-te-0/0)#no shutdown
switch(config-if-te-0/0)#exit
```

3. Create a CEE Map to carry LAN and SAN traffic and apply it to an interface.
   Example of creating a CEE map for 10-Gigabit Ethernet interface.

```
switch(config)#cee-map default
switch(conf-cee-map)#priority-group-table 1 weight 40 pfc
switch(conf-cee-map)#priority-group-table 2 weight 60
switch(conf-cee-map)#priority-table 2 2 2 1 2 2 2 2
switch(conf-cee-map)#interface tengigabitethernet 0/2
switch(conf-if-te-0/2)#cee default
switch(conf-if-te-0/2)#exit
```

4. Configure LLDP for FCoE.
   Example of configuring LLDP for 10-Gigabit Ethernet interface.

```
switch(config)#protocol lldp
switch(conf-lldp)#advertise dcbx-fcoe-app-tlv
switch(conf-lldp)#advertise dcbx-fcoe-logical-link-tlv
```

5. Enter the copy command to save the *running-config* file to the *startup-config* file.

```
switch(conf-lldp)#exit
switch(config)#end
switch#copy running-config startup-config
```

For detailed configurations refer the *Converged Enhanced Ethernet Administrators Guide*.

Fabric OS Native and Access Gateway modes

The Brocade 8000 can function in either Fabric OS Native mode or Brocade Access Gateway mode. The switch is shipped in Fabric OS Native mode by default.

- You can enable Access Gateway mode using Fabric OS commands or Web Tools.
- When you enable Access Gateway, you can use the default F_Port-to-N_Port mappings or change this mapping using command line interface (CLI) or WebTools, after you configure an IP address using instructions under Switch IP Address in Chapter 2.
- Access Gateway simplifies SAN deployment by using NPIV. NPIV provides Fibre Channel switch functions that improve switch scalability, manageability, and interoperability.

NOTE
Access Gateway cannot be connected directly into an array; it requires a fabric to support NPIV.

- Fabric OS features available to the Brocade 8000 depend on whether the switch is configured in Access Gateway or Fabric OS Native mode.
- In Fabric OS Native mode, the switch provides up to eight external Fibre Channel ports. These universal and self-configuring ports are capable of becoming one of the following port types:
  - F_Port (fabric enabled)
  - FL_Port (fabric loop enabled)
  - E_Port (expansion port)
  - M_Port (mirror port)
- In Access Gateway mode, the switch also provides up to eight external Fibre Channel ports. However, these ports are configured as N_Ports, and you cannot reconfigure these as any other port type.
- The Brocade 8000 provides up to 24 internal CEE ports, divided into six groups or trunks. In Access Gateway mode, these CEE ports are configured as F_Ports, and you cannot reconfigure these as any other port type. Each CEE port group is mapped to one of the eight Fibre Channel ports (N_Ports).
- CEE ports are divided into the following groups:
  - 8, 9, 10, 11
  - 12, 13, 14, 15
  - 16, 17, 18, 19
  - 20, 21, 22, 23
  - 24, 25, 26, 27
  - 28, 29, 30, 31
- Although you can change the default F_Port to N_Port mapping (CEE port group to Fibre Channel port mapping), consider these points for the CEE port groups:
  - All four CEE ports in the port group are mapped to the same Fibre Channel N_Port.
  - You cannot map individual CEE ports within the same port group to different Fibre Channel ports.
  - Any Access Gateway operation that involves moving F_Ports will move all FCoE ports in the group.
  - All four CEE ports in a group will failover or failback to one Fibre Channel N port.
  - All four CEE ports are considered as a unit for rebalancing.

Enabling Access Gateway mode

Note the following when enabling Access Gateway mode:

- After you enable AG mode, some fabric information is erased, such as the zone and security databases.
- Enabling AG mode is disruptive because the switch is disabled and rebooted.
- Ensure that no zoning or Admin Domain (AD) transaction buffers are active. If any transaction buffer is active, enabling Access Gateway mode will fail with the error, “Failed to clear Zoning/Admin Domain configuration.”

Use the following steps to enable Access Gateway mode using Fabric OS commands.

1. Before disabling a switch to enable Access Gateway mode, save the current configuration file using the **configupload** command in case you might need this configuration again.

2. Enter the **switchshow** command to verify the switch mode.
   - **"Access Gateway Mode"** displays for switchMode if the switch is in Access Gateway mode.
   - **"Native"** displays for switchMode if the switch is in Fabric OS Native mode.
3. Enter **switchDisable** to disable the switch. Access Gateway mode can only be enabled or disabled when the switch is in a disabled state.
4. Enter **ag –modeEnable** to enable Access Gateway mode.
5. Enter the **ag –modeshow** command to verify that AG mode is enabled.

```
switch:admin> ag –modeshow
Access Gateway mode is enabled.
```

Disabling Access Gateway mode

When you disable Access Gateway mode, the switch automatically reboots and comes back online using the fabric switch configuration. The Access Gateway parameters, such as F_Port-to-N_Port mapping, Failover, and Failback are automatically removed. When the switch reboots, it starts in Fabric OS Native mode.

Use the following steps to disable Access Gateway mode using Fabric OS commands.

1. Enter the **switchshow** command to verify the switch mode.
   "Access Gateway Mode" displays if the switch is in Access Gateway mode.
   Interopmode "0" or "Native" displays if the switch is in Fabric OS Native mode.
2. Enter **switchDisable** to disable the switch. Access Gateway mode can only be disabled or enabled when the switch is in a disabled state.
3. Enter **ag –modeDisable** to disable Access Gateway mode.
4. Enter the **ag –modeshow** command to verify that AG mode is disabled.

```
switch:admin> ag –modeshow
Access Gateway mode is NOT enabled
```

Web Tools

FCOE configuration tasks

There are several tasks related to FCOE configuration. The following lists the high level tasks in a suggested order:

- Quality of Service (QoS) configuration (optional) - If you intend to implement a specific QoS scheme to prioritize data traffic, it is recommended that you finish your QoS configuration before you begin port configuration. QoS values are referenced when you configure ports.
- LLDP-DCBX configuration (optional) - If you intend to implement DCBX, it is recommended that you finish LLDP-DCBX configuration before you configure ports. LLDP-DCBX values are referenced when you configure ports.
- DCB interface configuration (mandatory).
- Link Aggregation Groups (LAG) configuration - Ports must be configured before they can be placed into a LAG. The parameters applied to the LAG will reflect on each port which is member of a LAG.
- VLAN configuration (optional) - Port and LAG names are referenced in VLAN configuration, and must be defined before you can successfully complete a VLAN configuration.
- Login group configuration (optional) - Login group configuration is not dependent on any of the above configurations. It can be done as a separate task.

Quality of Service (QoS) configuration

As a general concept, QoS is a mechanism for classifying and scheduling data traffic based on priority settings. QoS can be used to control traffic congestion, allocate bandwidth, and carry data traffic with different characteristics over a common interface.

The following two configuration options are available:

- You can create a DCB map. A DCB map defines priority and priority group tables that support Enhanced Transmission Selection (ETS). ETS allows allocation of bandwidth to different traffic classes. DCB maps also allow you to enable Priority Flow Control (PFC).
- You can create a traffic class map. A traffic class map can be used to map a specific class of traffic to a specific Class of Service (CoS).

Adding a DCB map

A DCB map defines priority and priority group tables that support Enhanced Transmission Selection (ETS). ETS allows bandwidth to be allocated based on priority settings through an exchange of priority group tables.

1. Select the DCB tab on the Switch Administration panel.
2. Select the QoS tab.
3. Select the DCB Map tab.

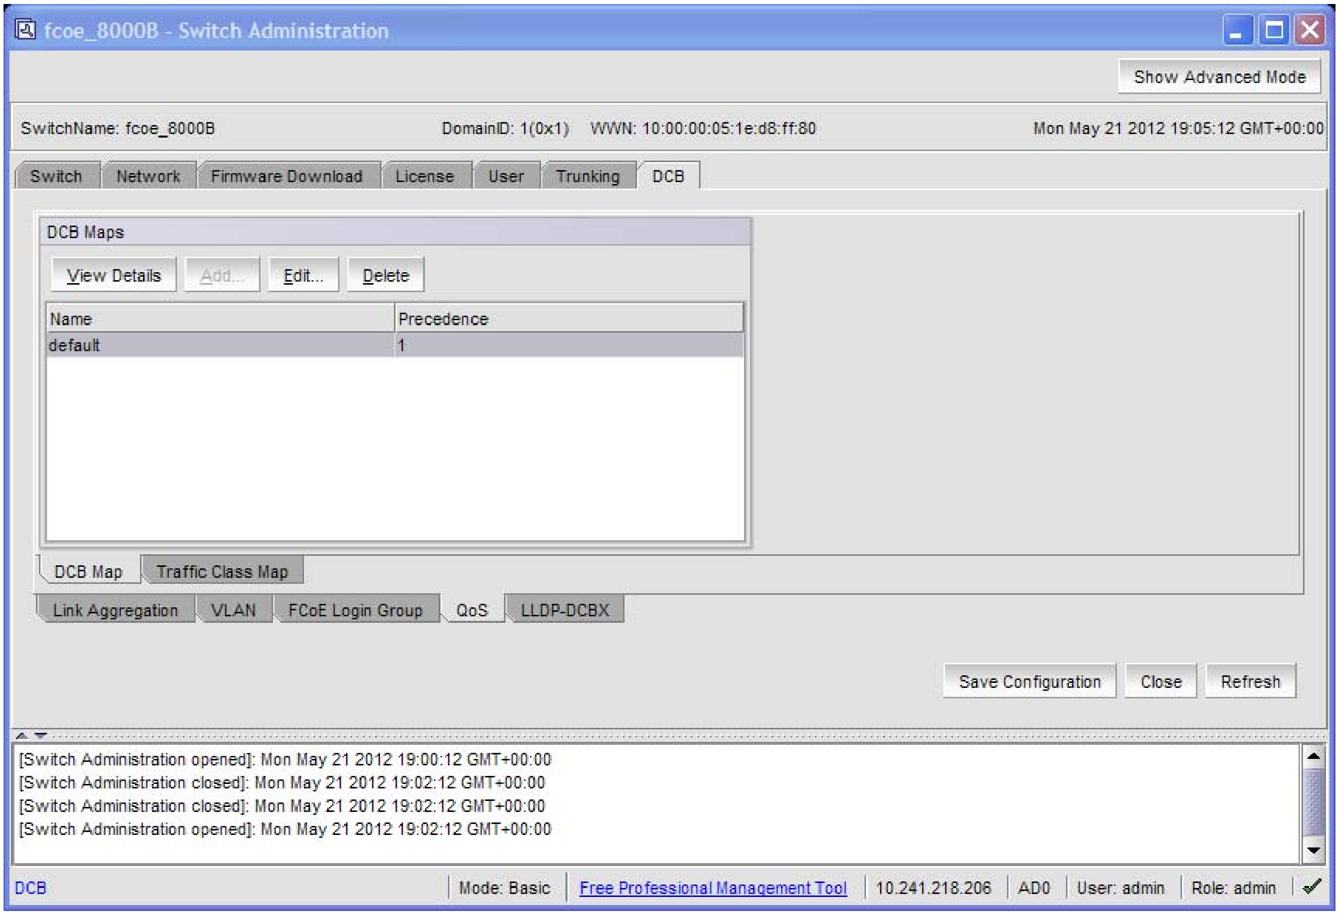

SwitchName: fcoe_8000B DomainID: 1(0x1) WWN: 10:00:00:05:1e:d8:ff:80 Mon May 21 2012 19:05:12 GMT+00:00

Name	Precedence
default	1

```
[Switch Administration opened]: Mon May 21 2012 19:00:12 GMT+00:00
[Switch Administration closed]: Mon May 21 2012 19:02:12 GMT+00:00
[Switch Administration closed]: Mon May 21 2012 19:02:12 GMT+00:00
[Switch Administration opened]: Mon May 21 2012 19:02:12 GMT+00:00
```

4. Select Add.

The DCB Map Configuration dialog box is displayed .

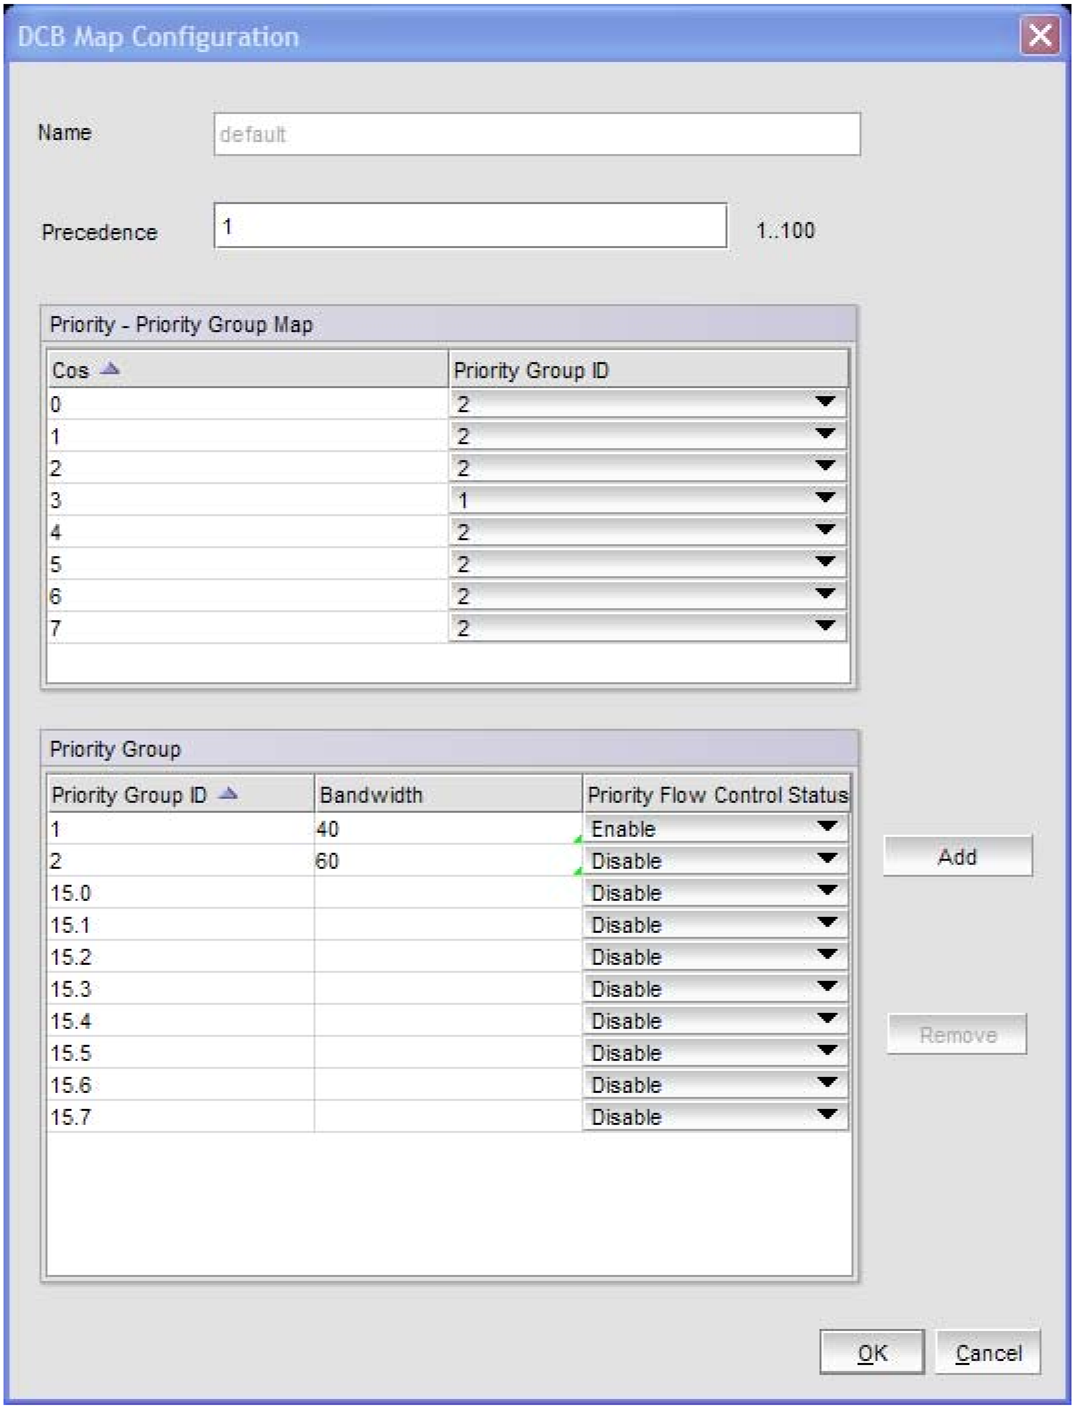

DCB Map Configuration

Name: default

Precedence: 1 (1..100)

Priority - Priority Group Map

Cos	Priority Group ID
0	2
1	2
2	2
3	1
4	2
5	2
6	2
7	2

Priority Group

Priority Group ID	Bandwidth	Priority Flow Control Status
1	40	Enable
2	60	Disable
15.0		Disable
15.1		Disable
15.2		Disable
15.3		Disable
15.4		Disable
15.5		Disable
15.6		Disable
15.7		Disable

5. Type a name for the DCB map in the Name field.
6. Type a precedence value in the Precedence field. The value is specified as a number. The allowable range is 1 to 100. The default is 1.

   The precedence value controls QoS scheduling policies. If different DCB maps have conflicting policies, the scheduler gives precedence to the DCB map with the highest precedence value (the highest number).

   When the DCB Map Configuration dialog box is displayed, the default values shown in the

Priority Group Map match the IEEE 802.1Q recommendation for systems supporting eight traffic classes. The Priority Group Map shows the Layer 2 Cos values mapped to Priority Group ID (PGID). PGID values are in the form <policy>.<priority>. A policy value of 15 indicates Priority values run from 0 (highest priority) to 7 (lowest priority). Note that this is contrary to the Cos values, which run from 7 (highest priority) to 0 (lowest priority).

7. Create a new priority group by clicking Add next to the Priority Group table.
8. Edit the Bandwidth entry to indicate the desired percentage of total bandwidth.
9. Change the Priority Flow Control Status to Enabled to enable PFC for the entry.
10. Click OK.

Adding a traffic class map

CoS priorities can be mapped to traffic classes using a traffic class map.

1. Select the DCB tab on the Switch Administration panel.
2. Select the QoS tab.
3. Select the Traffic Class Map tab.
4. Select Add.

The Traffic Class Map Configuration dialog box is displayed. This dialog box has the same structure as the Priority Group Map in the DCB Configuration dialog box.The default CoS-to-traffic class structure is based on IEEE 802.1Q recommendations, as in the default Priority Group Map shown.

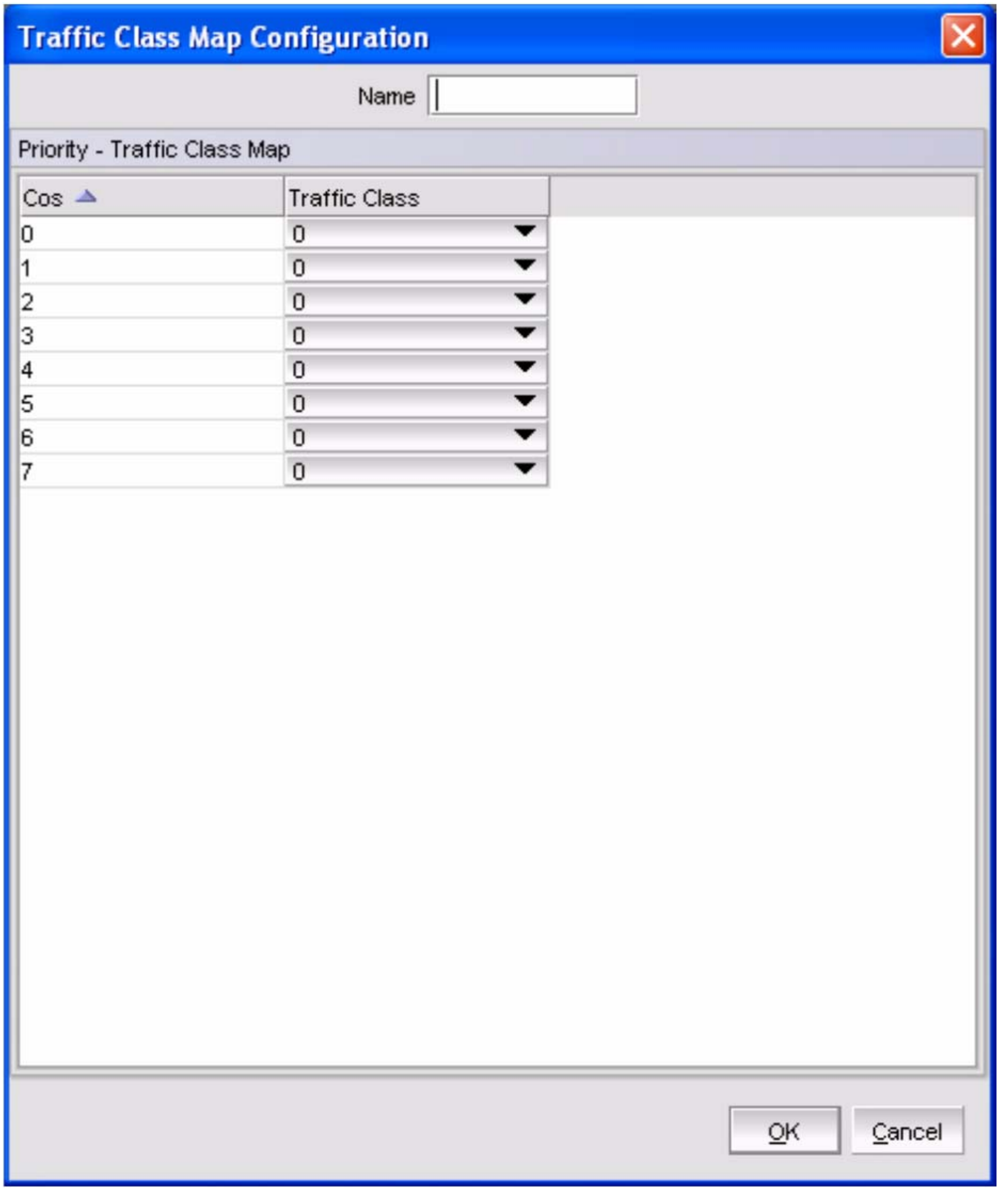

Cos	Traffic Class
0	0
1	0
2	0
3	0
4	0
5	0
6	0
7	0

5. Type a name for the traffic class map in the Name field.
6. Select the Traffic Class that you want to assign to the Cos priority.
7. Click OK.

LLDP-DCBX configuration

Link Layer Discovery Protocol (LLDP) is a IEEE standard for collecting and distributing device information. Data Center Bridging Exchange (DCBX) extends LLDP by providing a protocol for discovering, initializing, and managing DCB-compliant devices.

There are two configuration procedures:

- Configuring global LLDP characteristics.
- Configuring an LLDP profile.

Configuring global LLDP characteristics

Configuring at the global level enables you to apply changes to every port.

1. Select the DCB tab on the Switch Administration panel.

2. Select the LLDP-DCBX tab.
3. Select the Global tab.

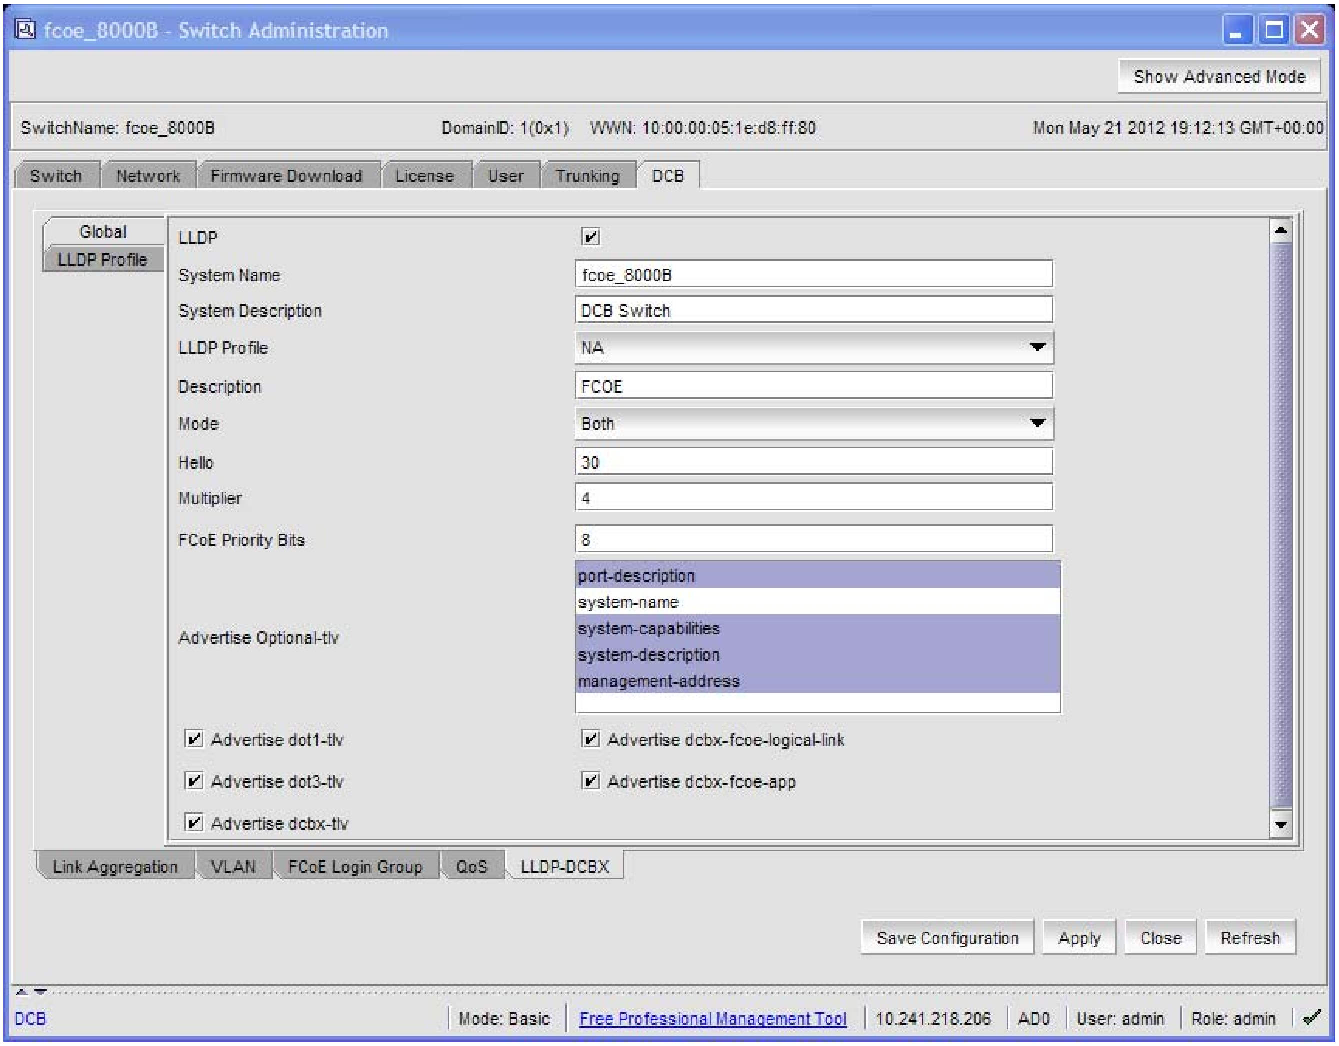

SwitchName: fcoe_8000B DomainID: 1(0x1) WWN: 10:00:00:05:1e:d8:ff:80 Mon May 21 2012 19:12:13 GMT+00:00

LLDP ☑

System Name: fcoe_8000B

System Description: DCB Switch

LLDP Profile: NA

Description: FCOE

Mode: Both

Hello: 30

Multiplier: 4

FCoE Priority Bits: 8

Advertise Optional-tlv: port-description, system-name, system-capabilities, system-description, management-address

☑ Advertise dot1-tlv ☑ Advertise dcbx-fcoe-logical-link

☑ Advertise dot3-tlv ☑ Advertise dcbx-fcoe-app

☑ Advertise dcbx-tlv

4. Select the LLDP check box to enable LLDP globally. You can clear the check box to disable LLDP.
5. Type a name for the configuration in the System Name field. Optionally, add a description in the System Description field.
6. Choose the Mode.

   For Mode, the choices are Tx (transmit), Rx, (receive) or Both. The default is Both.
7. In the Hello field, enter a time value in seconds. The Hello value sets the interval between hello bridge protocol data units sent by the root switch configuration messages. The range is 4 to 180 seconds. The default is 30 seconds.
8. In the Multiplier field, set the number of consecutive misses allowed before LLDP considers the interface to be down. The range is 1 to 10. The default is 4. The multiplier is related to the Hello time interval. Using the defaults, you wait four times (the multiplier value) at 30 second intervals (the hello value) before giving up on the interface.
9. In the FCOE Priority Bits field, type a value that indicates the desired user priority. Each bit represents a user priority associated with FCoE traffic. The range is 0-255. The default is 8.

10. Choose the parameters you want to exchange. Note that the term TLV indicates packaging of parameters into a Brocade-specific Type/Length/Value (TLV):

    - Advertise Optional-tlv - Advertises the following optional TLVs:

    - system-description - Describes switch or blade characteristics.
    - port-description - Describes the configured port.
    - system-name - Specifies the system name.
    - system-capabilities - Describes the system capabilities.
    - management-address - The IP address of the management port on the 8000 switch.

    - Advertise dot1-tlv - Select this check box to advertise to any attached device to send IEEE 802.1 LLDP type, length, and values.

    - Advertise dot3-tlv - Select this check box to advertise to any attached device to send IEEE 802.3 LLDP type, length, and values.

    - Advertise dcbx-tlv - Select this check box to advertise to any attached device the respective LLDP type, length, and values.

    - Advertise dcbx-fcoe-logical-link - Select this check box to advertise to any attached device to send DCBX protocol over LLDP to negotiate the logical link type, length, and values.

    - Advertise dcbx-fcoe-app - Select this check box to advertise application type, length, and values to ensure interoperability of traffic over DCBX protocol running over LLDP.

11. Click Apply.
12. Click Save Configuration.

Adding an LLDP profile

The LLDP profile determines LLDP settings per port.

1. Select the DCB tab on the Switch Administration panel.
2. Select the LLDP-DCBX tab.
3. Select the LLDP Profile tab.

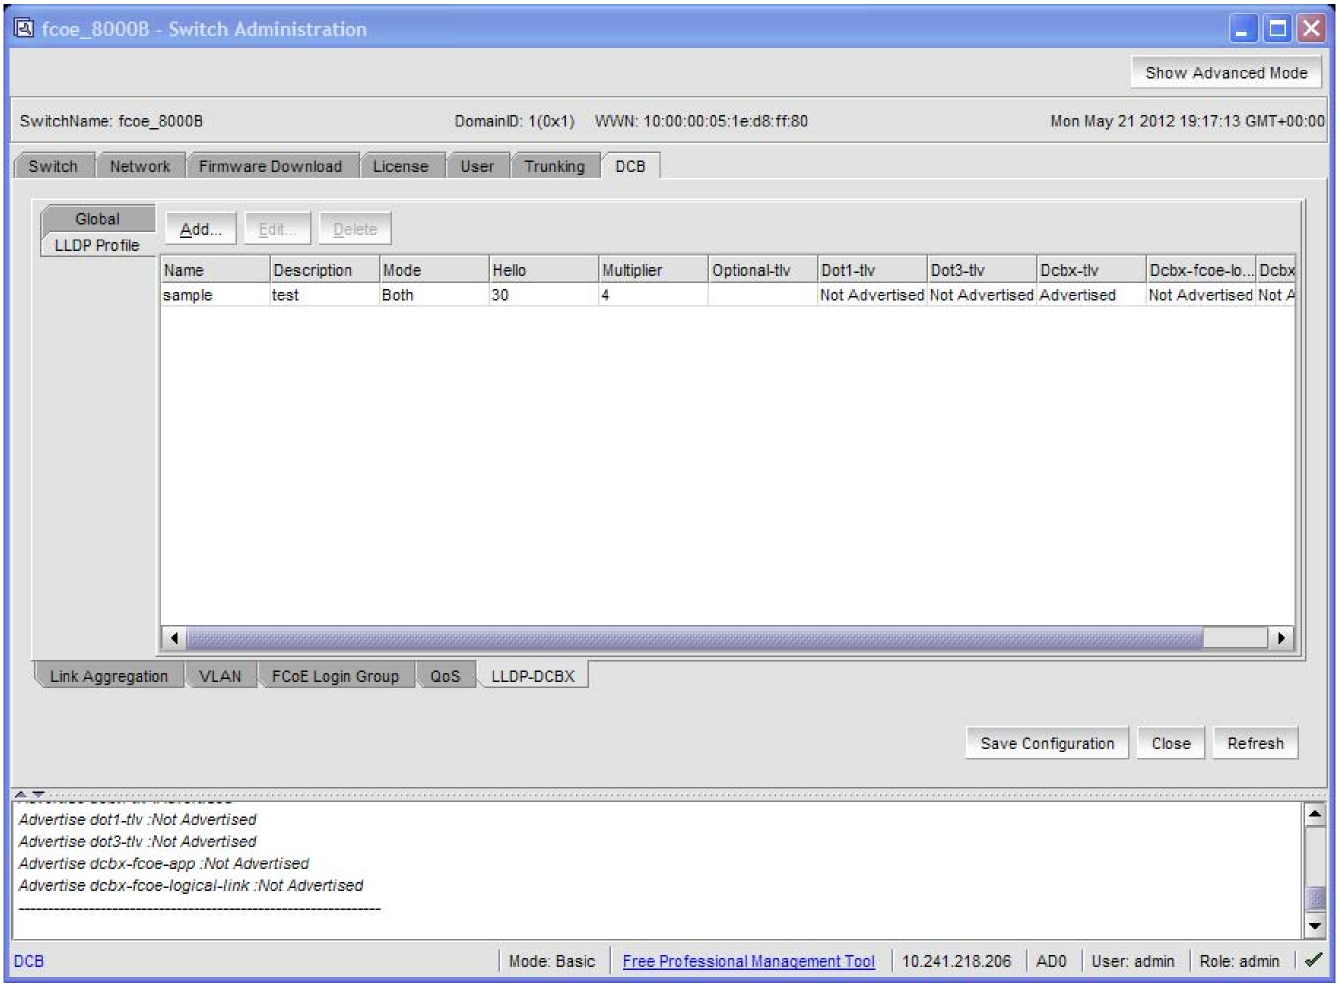

SwitchName: fcoe_8000B DomainID: 1(0x1) WWN: 10:00:00:05:1e:d8:ff:80 Mon May 21 2012 19:17:13 GMT+00:00

Name	Description	Mode	Hello	Multiplier	Optional-tlv	Dot1-tlv	Dot3-tlv	Dcbx-tlv	Dcbx-fcoe-lo...	Dcbx
sample	test	Both	30	4		Not Advertised	Not Advertised	Advertised	Not Advertised	Not A

```
Advertise dot1-tlv :Not Advertised
Advertise dot3-tlv :Not Advertised
Advertise dcbx-fcoe-app :Not Advertised
Advertise dcbx-fcoe-logical-link :Not Advertised
------------------------------------------------------------
```

Mode: Basic | 10.241.218.206 | AD0 | User: admin | Role: admin ✓

4. Click Add.

    The LLDP Configuration dialog box is displayed .

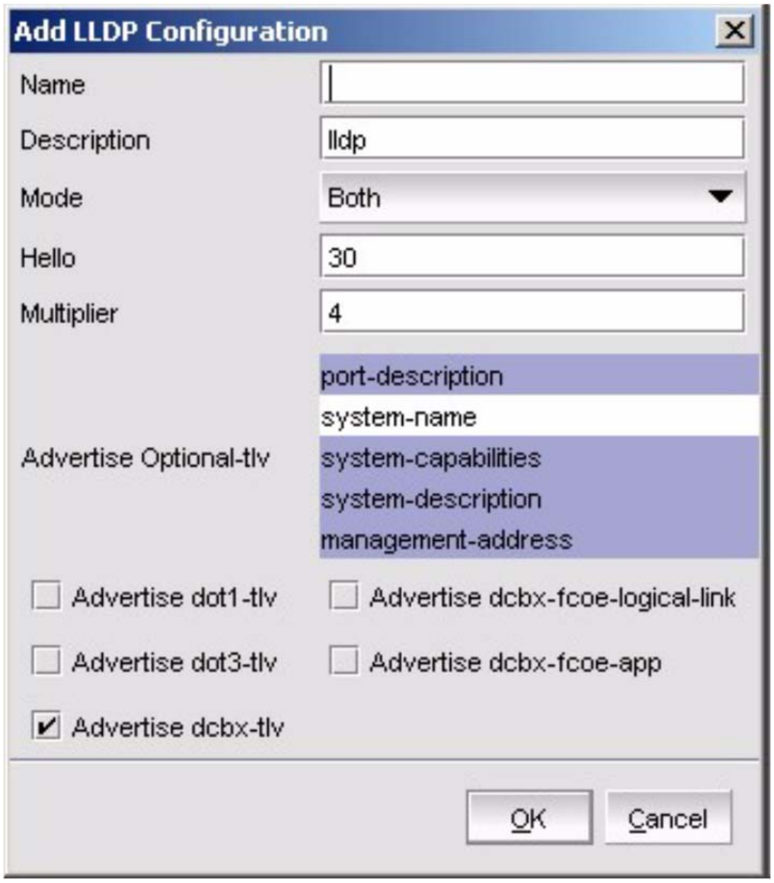

Name

Description lldp

Mode Both

Hello 30

Multiplier 4

Advertise Optional-tlv
port-description
system-name
system-capabilities
system-description
management-address

□ Advertise dot1-tlv
□ Advertise dcbx-fcoe-logical-link
□ Advertise dot3-tlv
□ Advertise dcbx-fcoe-app
☑ Advertise dcbx-tlv

5. Type a name for the configuration in the Name field.
6. Optionally, add a description in the Description field.
7. Choose the Mode.

   For Mode, the choices are Tx (transmit), Rx, (receive) or Both. The default is Both.
8. In the Hello field, enter a time value in seconds. The Hello value sets the interval between hello bridge protocol data units sent by the root switch configuration messages. The range is 4 to

   180 seconds. The default is the global configuration range.
9. In the Multiplier field, set the number of consecutive misses allowed before LLDP considers the interface to be down. The range is 1 to 10. The default is the global configuration range.

   The multiplier is related to the Hello time interval. Using the defaults, you wait four times (the multiplier value) at 30 second intervals (the hello value) before giving up on the interface.
10. Choose the parameters you want to exchange. Note that the term TLV indicates packaging of parameters into a Brocade-specific Type/Length/Value (TLV).

    - Advertise Optional-tlv - Advertises the following optional TLVs:
      - system-description - Describes switch or blade characteristics.
      - port-description - Describes the configured port.
      - system-name - Specifies the system name.
      - system capabilities - Describes the system capabilities.
      - management-address - The IP address of the management port on the 8000 switch.
    - Advertise dot1-tlv - Advertises to any attached device to send IEEE 802.1 LLDP type, length, and values.
    - Advertise dot3-tlv - Advertises to any attached device to send IEEE 802.3 LLDP type, length, and values.
    - Advertise dcbx-tlv - Advertises to any attached device the respective LLDP type, length, and values.

- Advertise dcbx-fcoe-logical-link - Advertises to any attached device to send DCBX protocol over LLDP to negotiate the logical link type, length, and values.

- Advertise dcbx-fcoe-app - Advertises application type, length, and values to ensure interoperability of traffic over DCBX protocol running over LLDP.

11. Click Save Configuration.

Configuring DCB interfaces

DCB interfaces are configured from the Port Administration panel.

1. Select the DCB Interfaces tab on the Port Administration panel.
2. Select the port you want to configure under the DCB Interface Explorer.
3. Select the General tab. Normally, this tab is pre-selected.

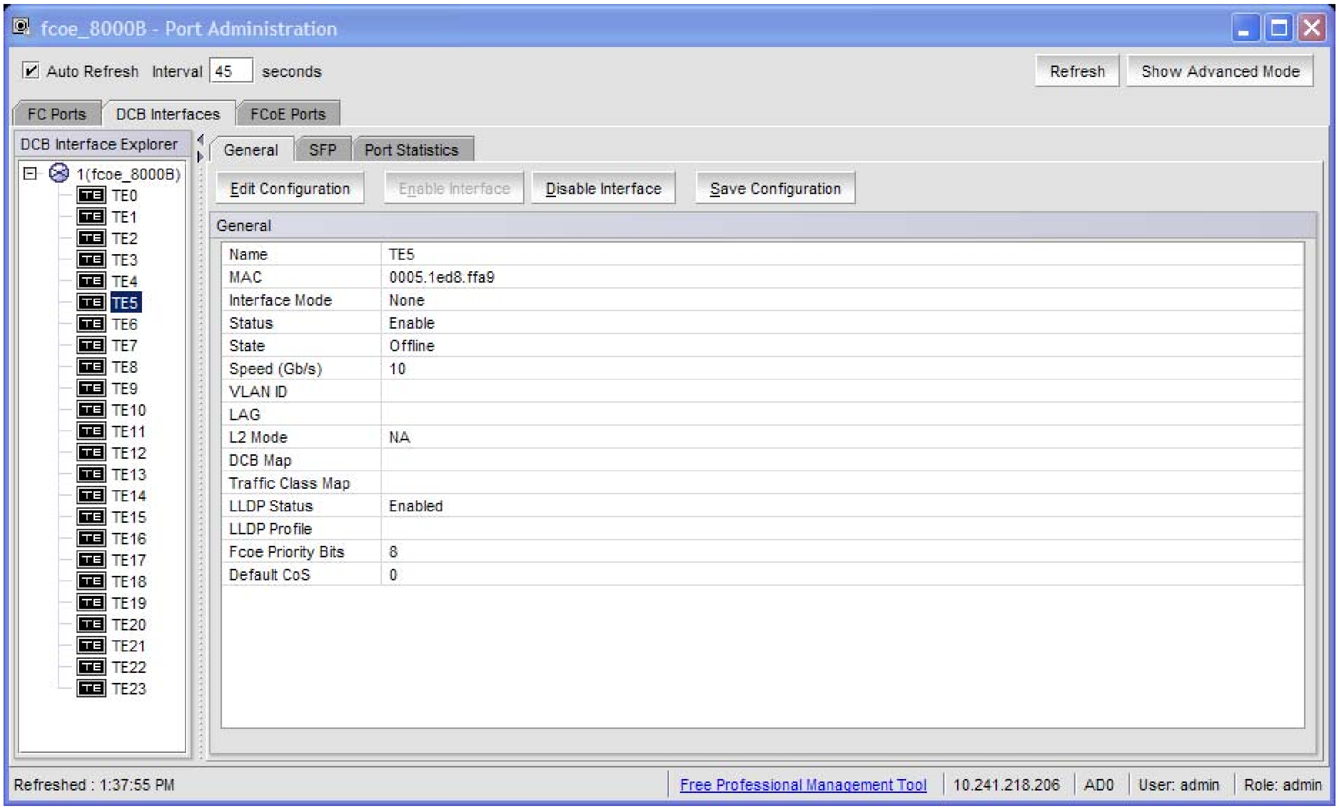

Name	TE5
MAC	0005.1ed8.ffa9
Interface Mode	None
Status	Enable
State	Offline
Speed (Gb/s)	10
VLAN ID	
LAG	
L2 Mode	NA
DCB Map	
Traffic Class Map	
LLDP Status	Enabled
LLDP Profile	
Fcoe Priority Bits	8
Default CoS	0

4. Select Edit Configuration.

   The DCB Edit Configuration dialog box is displayed.

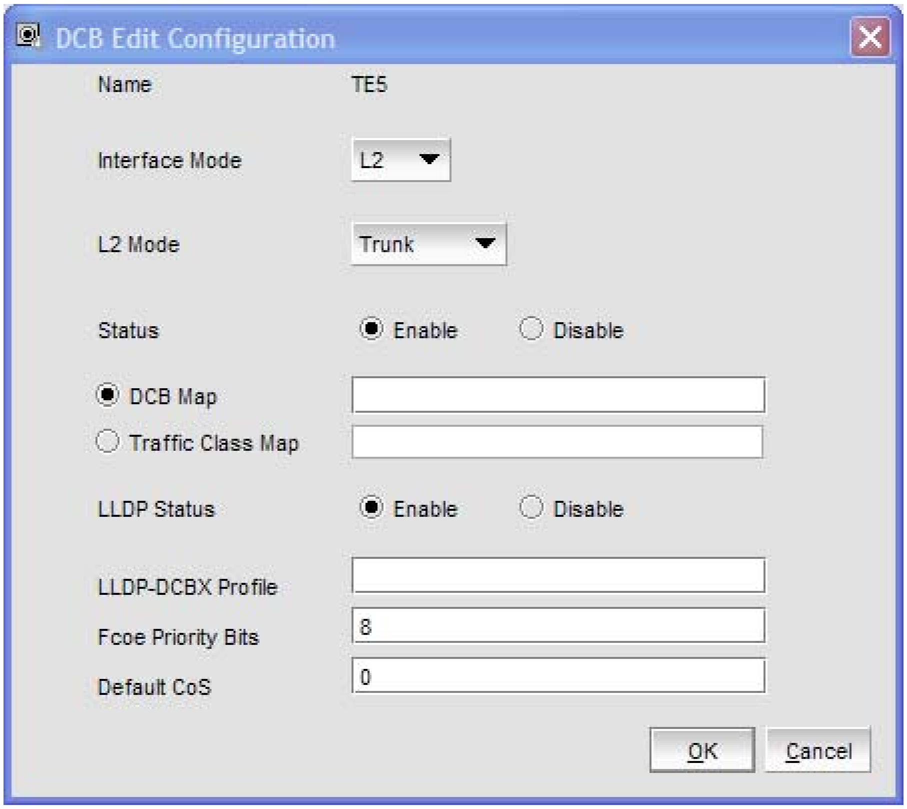

Name TE5

Interface Mode L2

L2 Mode Trunk

Status ● Enable ○ Disable

● DCB Map

○ Traffic Class Map

LLDP Status ● Enable ○ Disable

LLDP-DCBX Profile

Fcoe Priority Bits 8

Default CoS 0

5. Select the Interface Mode.

   The options are None and L2. The default is None.

   If you intend to use this port in a Link Aggregation Group (LAG), choose None. L2 mode will be applied when you configure the LAG.

6. Select the L2 Mode. The choices are Access, Trunk, and Converged. The default is Access.

   The L2 mode setting determines operation within a VLAN:

   - Access mode allows only one VLAN association, and all frames are untagged.

   - Trunk mode allows more that one VLAN association, and tagged frames are allowed.

   - Converged mode interface can be Native (untagged or access) in one VLAN and it could be non-native (trunk or tagged) type in another VLAN.

7. If you are using a DCB map or Traffic Class Map to apply QoS traffic priority, select the appropriate button, and enter the name of the map you want to use.
8. Enter the profile name in the LLDP-DCBX Profile field for using a specific profile for the interface.
9. In the FCOE Priority Bits field, type a value that indicates the desired user priority. Each bit represents a user priority that is associated with FCoE traffic. The range is 0-255. The default is 8.
10. Assign a default class of service in the Default CoS field. The default CoS range is 0-7. The default is 0.
11. Click OK.
12. Click Enable for Status and LLDP Status. This can be done at a later time.

Configuring a link aggregation group (LAG)

FCoE ports can be grouped to create a LAG. The LAG is treated as a single interface.

1. Select the DCB Interfaces tab on the Switch Administration panel.

   Select the Link Aggregation tab .

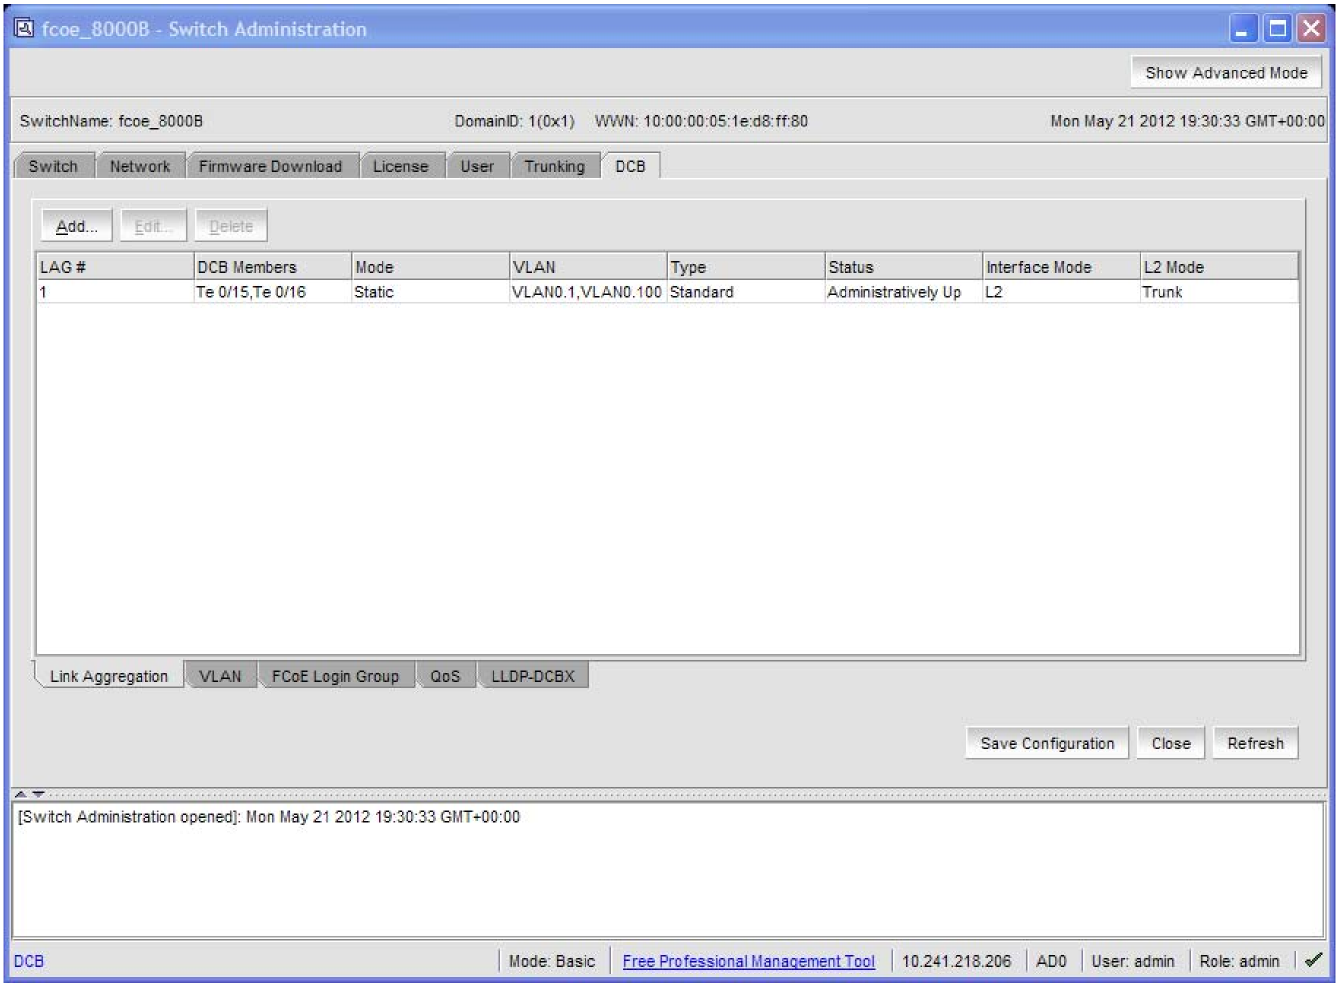

SwitchName: fcoe_8000B DomainID: 1(0x1) WWN: 10:00:00:05:1e:d8:ff:80 Mon May 21 2012 19:30:33 GMT+00:00

LAG #	DCB Members	Mode	VLAN	Type	Status	Interface Mode	L2 Mode
1	Te 0/15,Te 0/16	Static	VLAN0.1,VLAN0.100	Standard	Administratively Up	L2	Trunk

```
[Switch Administration opened]: Mon May 21 2012 19:30:33 GMT+00:00
```

2. Click Add.

   The Add LAG Configuration dialog box is displayed . Note that only ports that you defined with an Interface Mode of None can be a LAG Member.

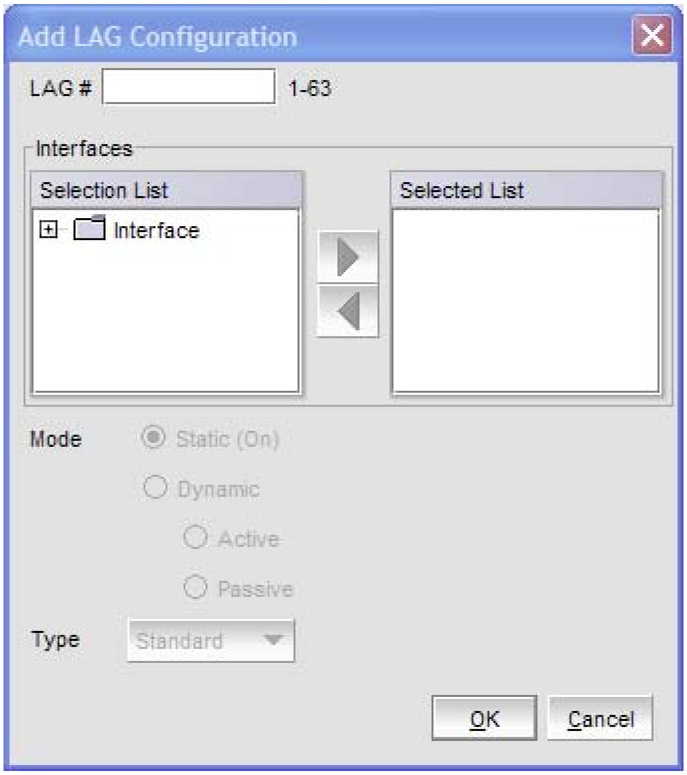

Add LAG Configuration

LAG # 1-63

Interfaces

Mode ● Static (On) ○ Dynamic ○ Active ○ Passive

Type Standard

3. Select the Mode.

   The choices are Static and Dynamic.

   Static mode does not use Link Aggregation Control Protocol (LACP) to negotiate and manage link aggregation. Link participation in the LAG is determined by the link's operational status and administrative state.

   Dynamic mode uses LACP. LACP allows partner systems to examine the attributes of the links that connect them and dynamically form a LAG. When you choose Dynamic mode, the Active and Passive options are enabled:

   - If you choose Active, your switch will initiate an exchange of LACP data units.

   - If you choose Passive, your switch will wait to receive LACP data units from its partner system and then respond.

   Passive is the default behavior.
4. Select the Type.

   Type refers to the type of trunking used by the LAG. The choices are Standard and Brocade.
5. Select the Interface Mode.

   The options are None and L2. The default is None.
6. Select the L2 Mode.

   The L2 mode setting determines operation within a VLAN:

   - Access mode allows only one VLAN association, and all frames are untagged.

   - Trunk mode allows more than one VLAN association, and allows tagged frames.
7. Select the operational Status.

   The choices are Administratively Up and Administratively Down.
8. Click OK.

Configuring VLANs

The Virtual LAN (VLAN) capability allows multiple virtual LANs within a single physical LAN infrastructure. The physical interface must be configured as L2 prior to configuring a VLAN, either as an individual interface, or as a LAG. Before you start the VLAN configuration procedure, you need to know which interfaces or LAGs you want to associate with each VLAN.

1. Select the DCB tab on the Switch Administration panel.
2. Select the VLAN tab.

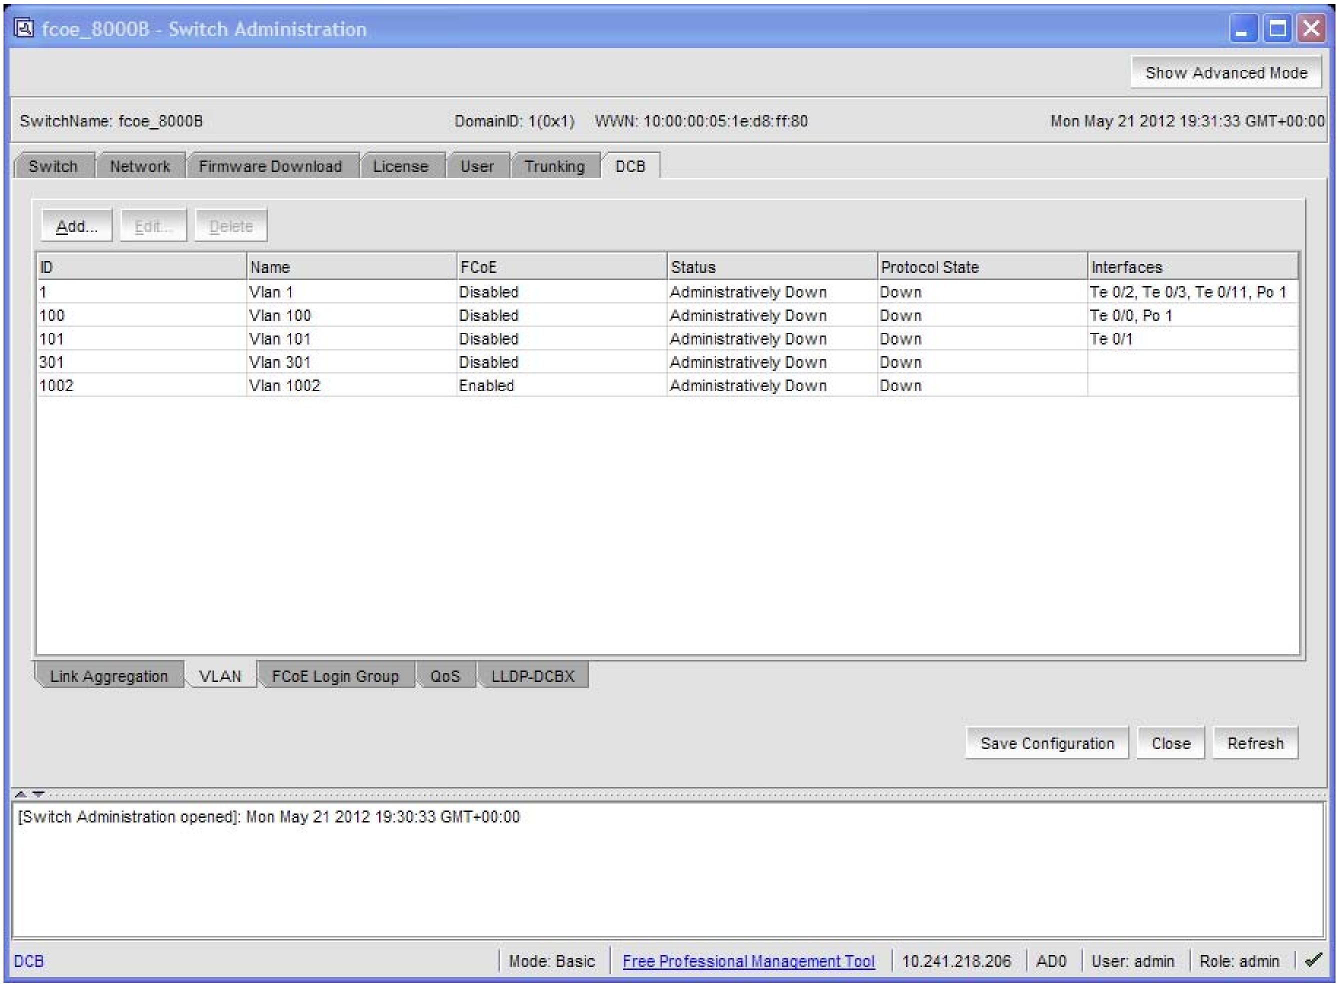

SwitchName: fcoe_8000B DomainID: 1(0x1) WWN: 10:00:00:05:1e:d8:ff:80 Mon May 21 2012 19:31:33 GMT+00:00

ID	Name	FCoE	Status	Protocol State	Interfaces
1	Vlan 1	Disabled	Administratively Down	Down	Te 0/2, Te 0/3, Te 0/11, Po 1
100	Vlan 100	Disabled	Administratively Down	Down	Te 0/0, Po 1
101	Vlan 101	Disabled	Administratively Down	Down	Te 0/1
301	Vlan 301	Disabled	Administratively Down	Down	
1002	Vlan 1002	Enabled	Administratively Down	Down	

[Switch Administration opened]: Mon May 21 2012 19:30:33 GMT+00:00

3. Click Add.

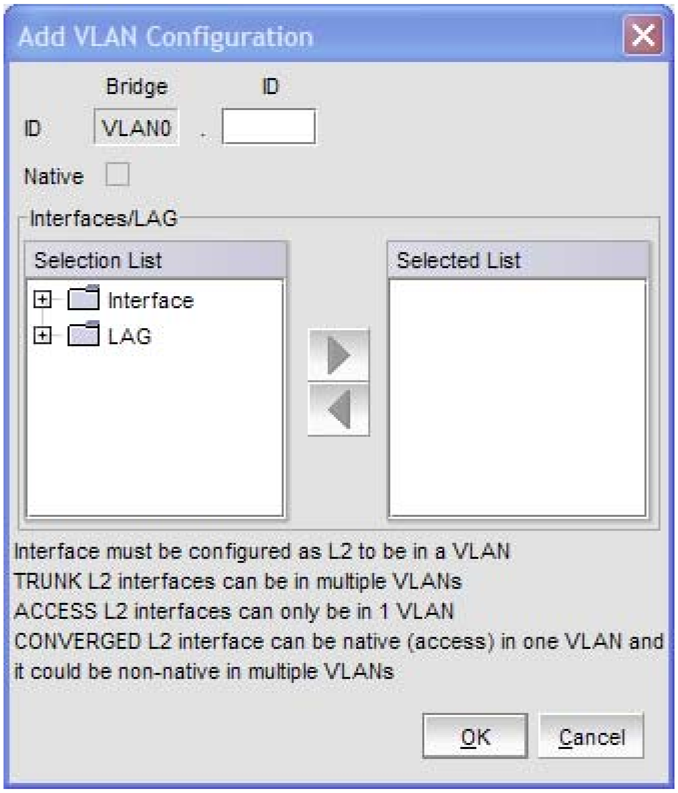

Interface must be configured as L2 to be in a VLAN
TRUNK L2 interfaces can be in multiple VLANs
ACCESS L2 interfaces can only be in 1 VLAN
CONVERGED L2 interface can be native (access) in one VLAN and
it could be non-native in multiple VLANs

4. Specify a VLAN ID. The format is VLAN<bridge number><ID>. In this Fabric OS release, no bridge instances are supported, so the bridge number is always 0, and the value under Bridge is statically defined as VLAN0. The <ID> is an integer from 1 to 3583, which must be typed in the ID field.
5. Select the Native check box.
6. Under the Selection List, click the plus sign next to the Interface and LAG folders, and select individual interfaces and LAGs you want to associate with the VLAN ID.
7. Click Add to move the interfaces or LAGs to the Selected List.

   Note the reminder that interfaces must be configured as L2, and that the interfaces or LAGs must be in Trunk mode to be associated with multiple VLANs, Access mode interfaces can be associated with only one VLAN, and the Converged mode interface can be Native in one VLAN and it could be non-native type in more than one VLAN.
8. Click OK.
9. Repeat the procedure for additional VLANs.
10. To edit VLAN, select the detail from the table in the VLAN tab and click Edit. The

    FCoE check box is selected by default. Click OK to enable FCoE. Clear the check box to disable FCoE.

Configuring FCoE login groups

FCoE login groups control which FCoE switches are allowed to log in to a fabric.

1. Select the DCB tab on the Switch Administration panel.
2. Select the FCoE Login tab).

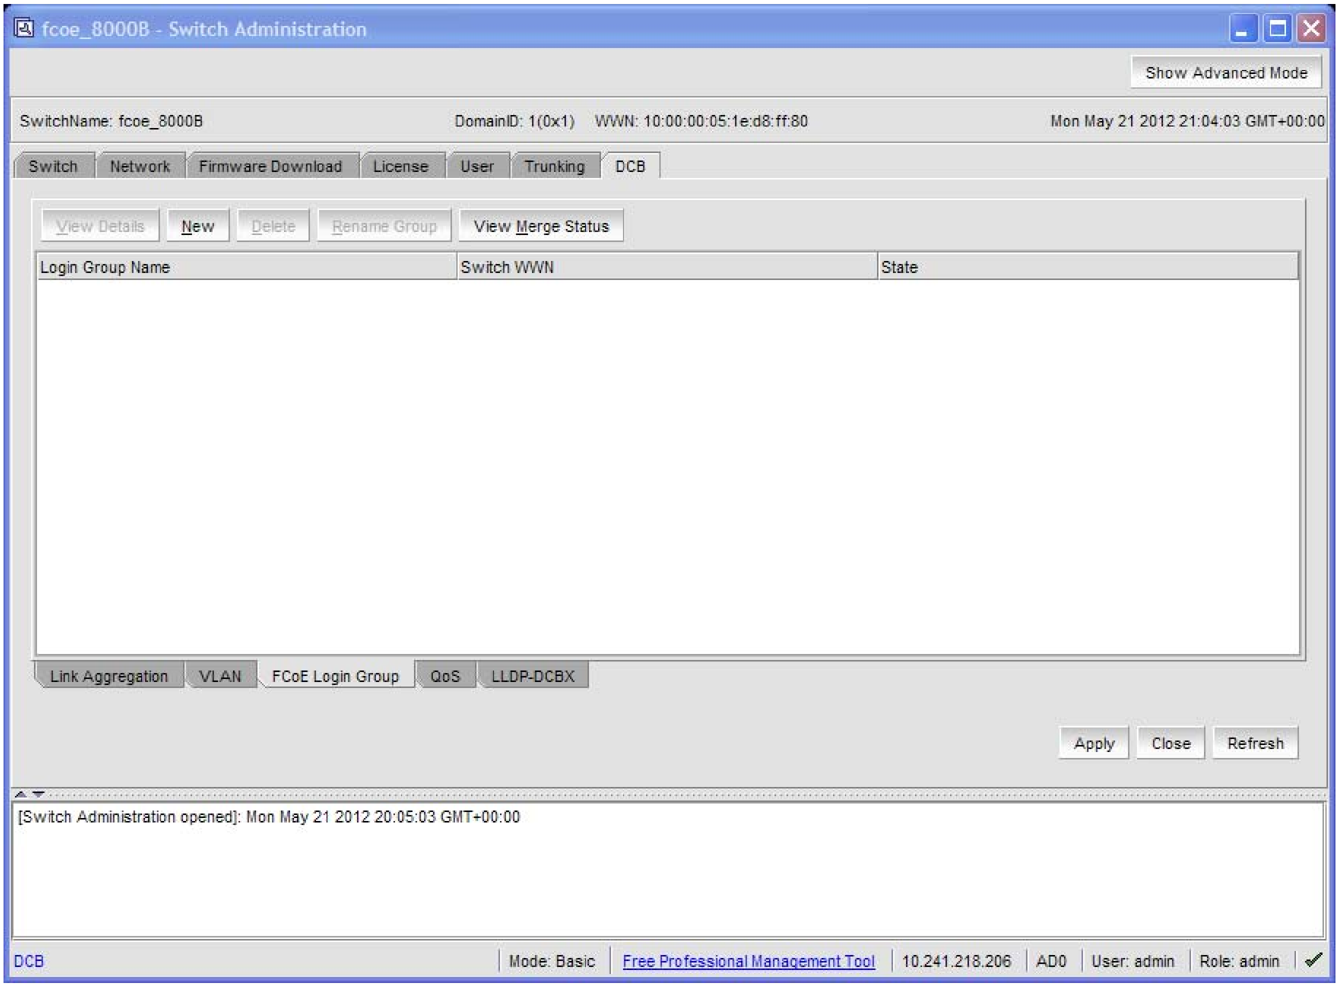

SwitchName: fcoe_8000B DomainID: 1(0x1) WWN: 10:00:00:05:1e:d8:ff:80 Mon May 21 2012 21:04:03 GMT+00:00

Login Group Name	Switch WWN	State

[Switch Administration opened]: Mon May 21 2012 20:05:03 GMT+00:00

Mode: Basic | 10.241.218.206 | AD0 | User: admin | Role: admin

3. Click New.

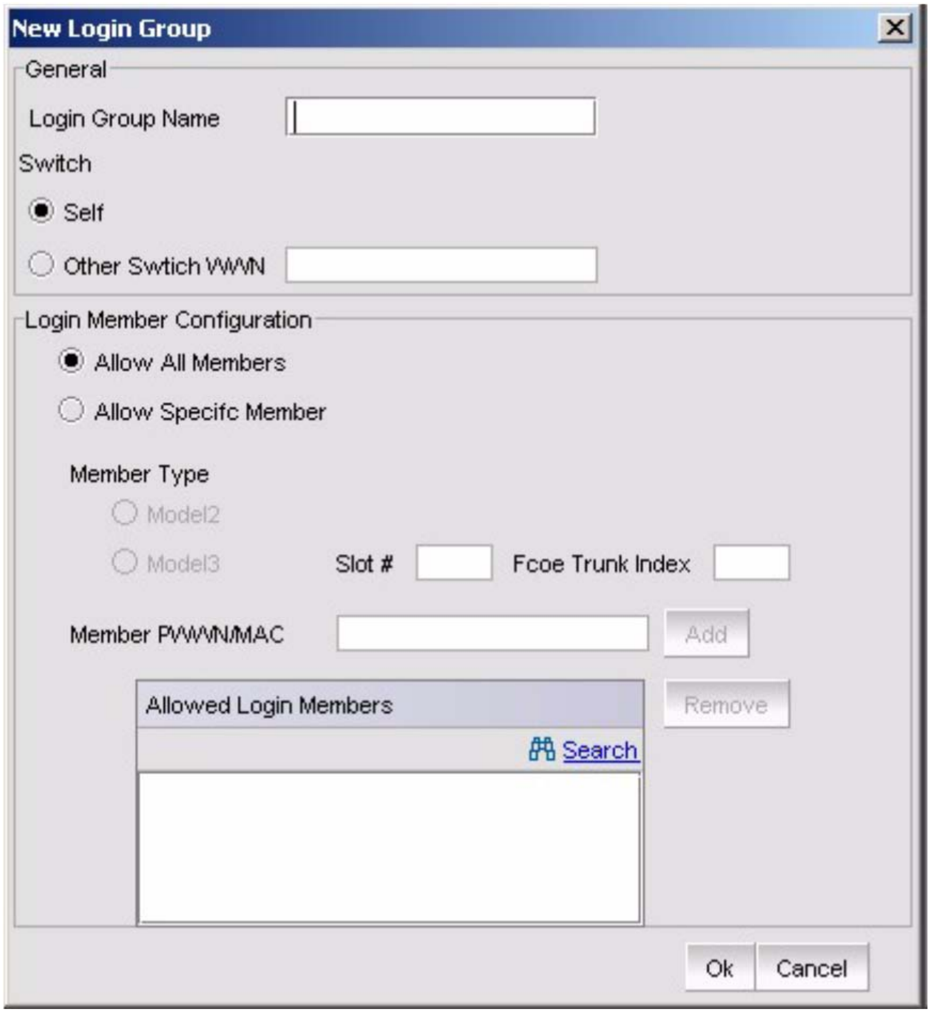

New Login Group

General

Login Group Name

Switch

● Self

○ Other Swtich WWN

Login Member Configuration

● Allow All Members

○ Allow Specifc Member

Member Type

○ Model2

○ Model3 Slot # Fcoe Trunk Index

Member PWWN/MAC

4. Type a name for the login group in the Login Group Name field.
5. Select the switch WWN, The choices are Self, which is the WWN of the switch you are logged into, or Other Switch WWN. If you choose Other Switch WWN, you must type the WWN of that switch in the provided field.
6. Under Login Member Configuration, click either Allow All Members, or Allow Specific Member.
    - If you choose Allow All Members, all devices attached to FCoE ports are allowed to log in to the switch.
    - If you choose Allow Specific Member, you can control which devices can log in, using

   Member Type, Member PWWN/MAC, and the Add and Remove buttons, as described below.
   a. Select Model2 as Member Type for an 8000 switch and proceed to step c.
   b. Select Model3 as Member Type and enter values in the Slot # and Fcoe Trunk Index fields.
   c. Type the port WWN in hexadecimal format in the Member PWWN/MAC field, and click
   Add.
   The WWN is displayed under Allowed Login Members. If you decide a member should not be on the list, highlight the entry and click Remove.
7. Click OK

Displaying FCoE Port Information

There are 24 internal FCoE Ports that bridge FC and Ethernet traffic. You can view FCoE port information from the Port Administration panel.

1. Select the FCoE Ports tab on the Port Administration panel.
   The initial view shows a summary of all FCoE ports on the switch.

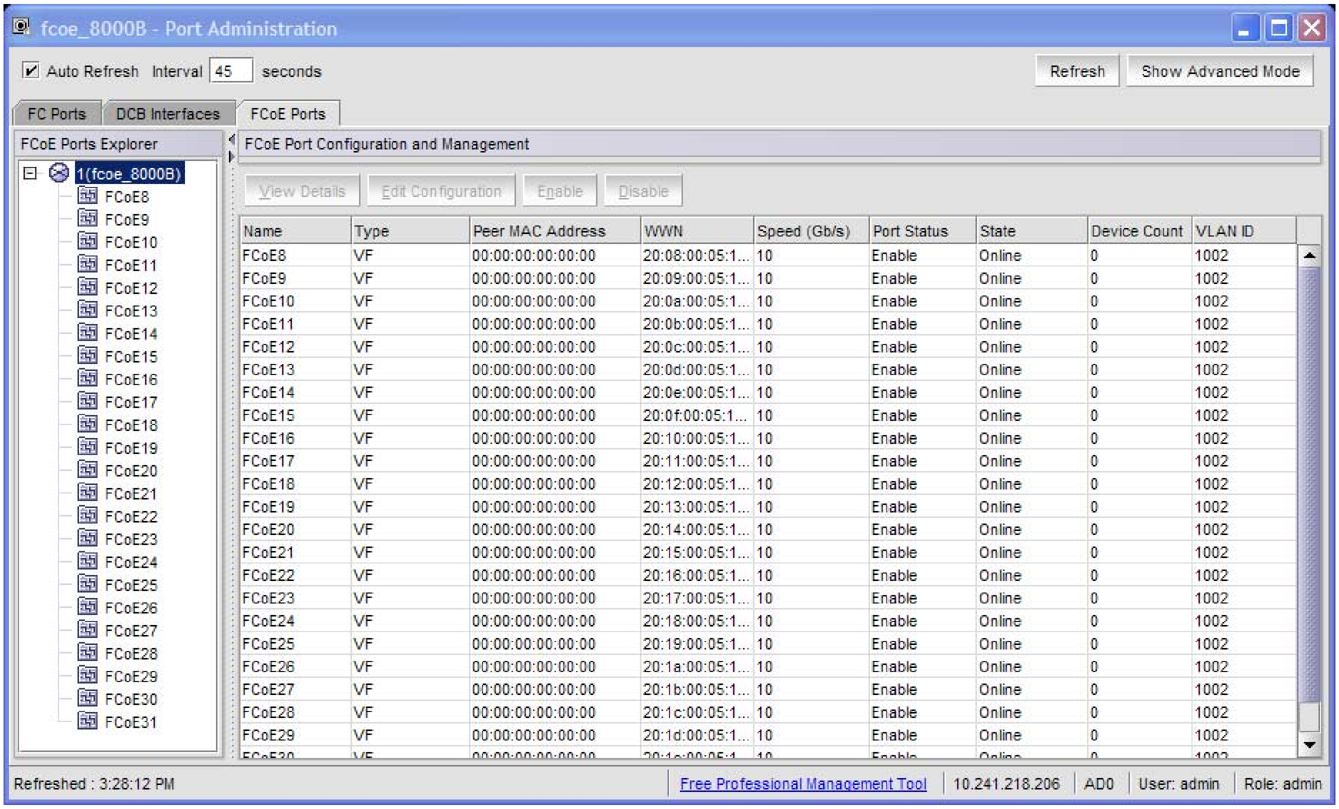

Name	Type	Peer MAC Address	WWN	Speed (Gb/s)	Port Status	State	Device Count	VLAN ID
FCoE8	VF	00:00:00:00:00:00	20:08:00:05:1...	10	Enable	Online	0	1002
FCoE9	VF	00:00:00:00:00:00	20:09:00:05:1...	10	Enable	Online	0	1002
FCoE10	VF	00:00:00:00:00:00	20:0a:00:05:1...	10	Enable	Online	0	1002
FCoE11	VF	00:00:00:00:00:00	20:0b:00:05:1...	10	Enable	Online	0	1002
FCoE12	VF	00:00:00:00:00:00	20:0c:00:05:1...	10	Enable	Online	0	1002
FCoE13	VF	00:00:00:00:00:00	20:0d:00:05:1...	10	Enable	Online	0	1002
FCoE14	VF	00:00:00:00:00:00	20:0e:00:05:1...	10	Enable	Online	0	1002
FCoE15	VF	00:00:00:00:00:00	20:0f:00:05:1...	10	Enable	Online	0	1002
FCoE16	VF	00:00:00:00:00:00	20:10:00:05:1...	10	Enable	Online	0	1002
FCoE17	VF	00:00:00:00:00:00	20:11:00:05:1...	10	Enable	Online	0	1002
FCoE18	VF	00:00:00:00:00:00	20:12:00:05:1...	10	Enable	Online	0	1002
FCoE19	VF	00:00:00:00:00:00	20:13:00:05:1...	10	Enable	Online	0	1002
FCoE20	VF	00:00:00:00:00:00	20:14:00:05:1...	10	Enable	Online	0	1002
FCoE21	VF	00:00:00:00:00:00	20:15:00:05:1...	10	Enable	Online	0	1002
FCoE22	VF	00:00:00:00:00:00	20:16:00:05:1...	10	Enable	Online	0	1002
FCoE23	VF	00:00:00:00:00:00	20:17:00:05:1...	10	Enable	Online	0	1002
FCoE24	VF	00:00:00:00:00:00	20:18:00:05:1...	10	Enable	Online	0	1002
FCoE25	VF	00:00:00:00:00:00	20:19:00:05:1...	10	Enable	Online	0	1002
FCoE26	VF	00:00:00:00:00:00	20:1a:00:05:1...	10	Enable	Online	0	1002
FCoE27	VF	00:00:00:00:00:00	20:1b:00:05:1...	10	Enable	Online	0	1002
FCoE28	VF	00:00:00:00:00:00	20:1c:00:05:1...	10	Enable	Online	0	1002
FCoE29	VF	00:00:00:00:00:00	20:1d:00:05:1...	10	Enable	Online	0	1002
FCoE30	VF	00:00:00:00:00:00	20:1e:00:05:1...	10	Enable	Online	0	1002

Refreshed : 3:28:12 PM | Free Professional Management Tool | 10.241.218.206 | AD0 | User: admin | Role: admin

2. To view information for a specific port, select the trunk in the FCOE Ports Explorer or select the port in the FCoE Port Configuration and Management table and click View Details.

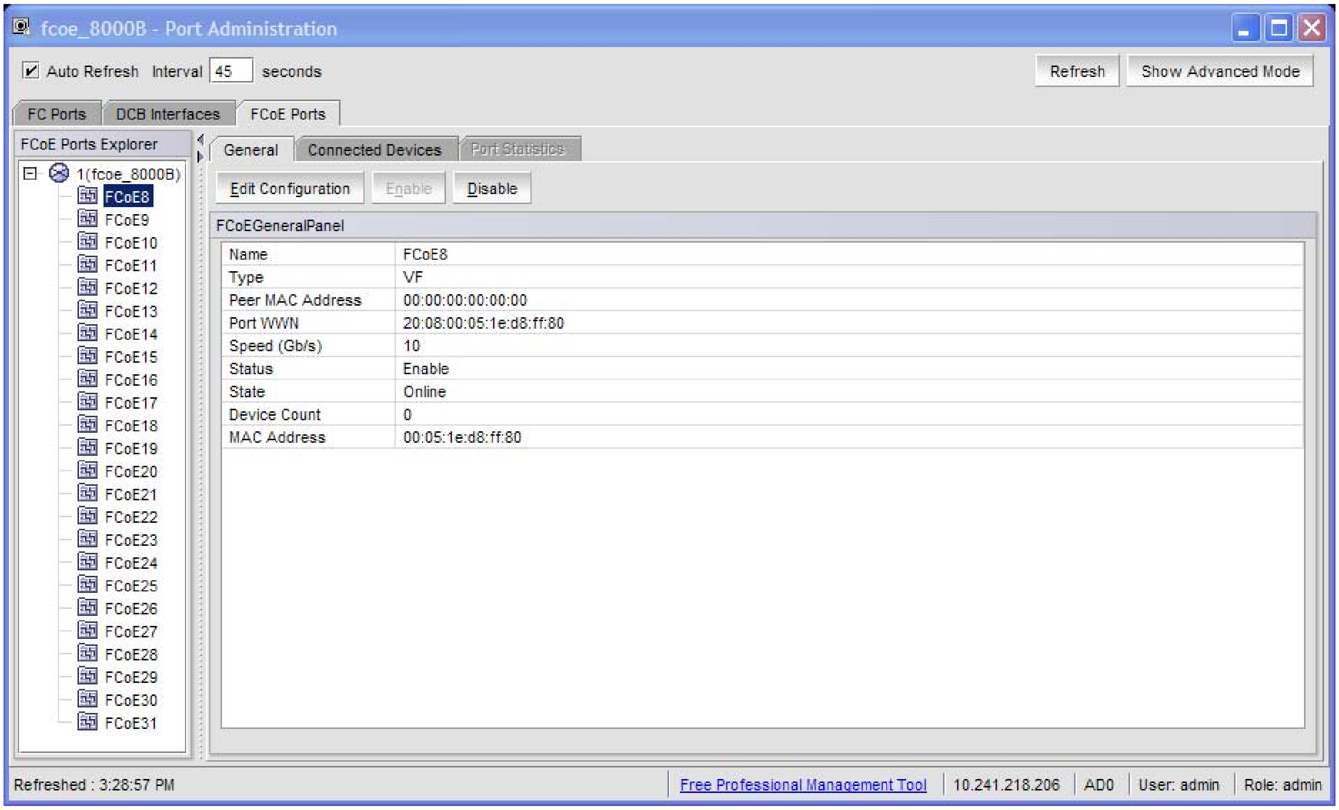

☑ Auto Refresh Interval 45 seconds

Name	FCoE8
Type	VF
Peer MAC Address	00:00:00:00:00:00
Port WWN	20:08:00:05:1e:d8:ff:80
Speed (Gb/s)	10
Status	Enable
State	Online
Device Count	0
MAC Address	00:05:1e:d8:ff:80

Refreshed : 3:28:57 PM | 10.241.218.206 | AD0 | User: admin | Role: admin

The Connected Devices tab shows information about devices connected to the switch. Six columns of information are displayed:

- Device WWN shows the WWN of the connected device.
- Device MAC shows the MAC address of the connected device.
- Connected Peer Type shows the port type on the connected device.
- Is Directly Connected indicates whether or not the device is directly connected to the trunk.
- FCoE Port MAC shows the FCoE port MAC address.
- Switch Port shows the switch port WWN.

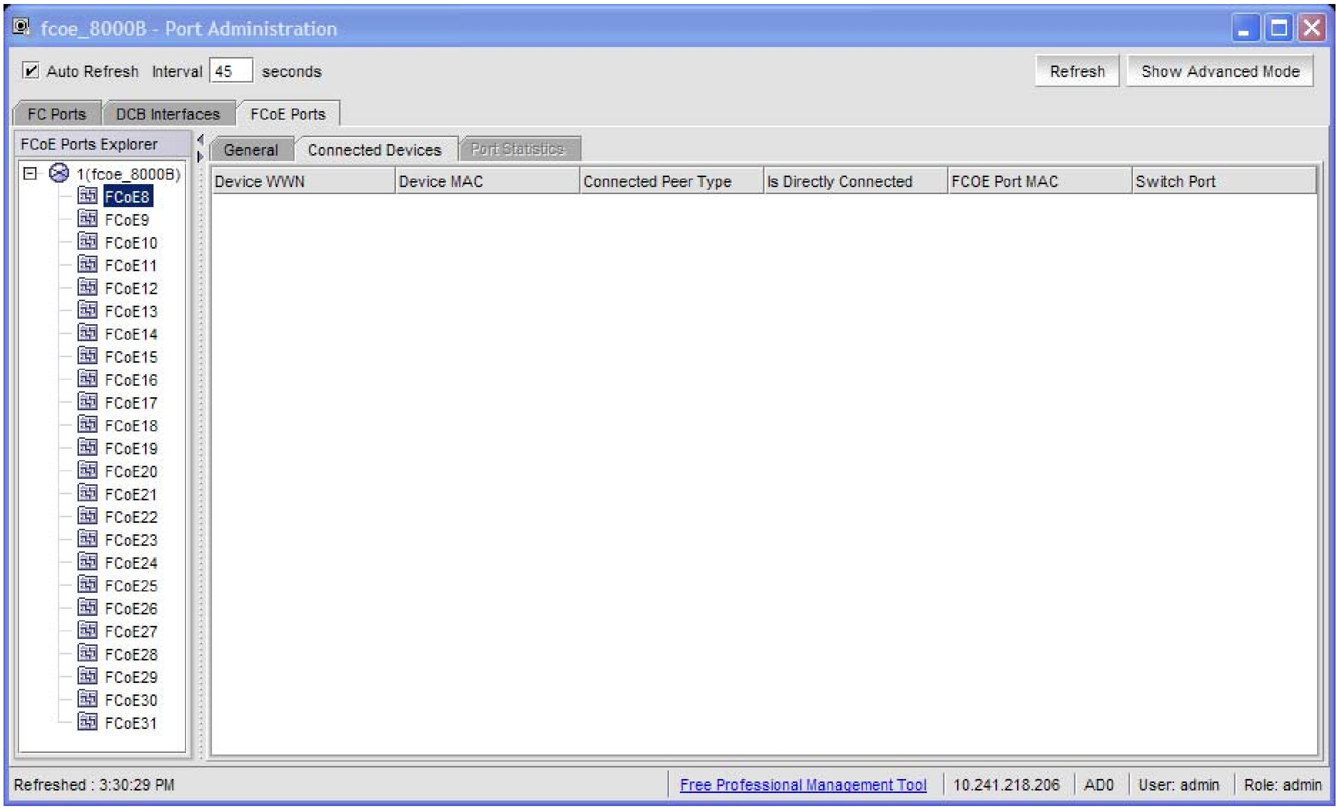

fcoe_8000B - Port Administration

☑ Auto Refresh Interval 45 seconds

Device WWN	Device MAC	Connected Peer Type	Is Directly Connected	FCOE Port MAC	Switch Port

Refreshed : 3:30:29 PM | Free Professional Management Tool | 10.241.218.206 | AD0 | User: admin | Role: admin

Displaying LAG information

Use the following procedure to display LAG information.

1. Select the DCB tab on the Switch Administration panel.
2. Select the Link Aggregation tab

LAG information is displayed .

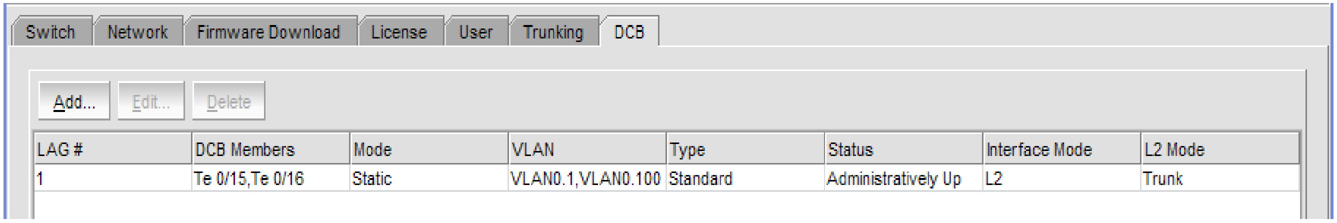

LAG #	DCB Members	Mode	VLAN	Type	Status	Interface Mode	L2 Mode
1	Te 0/15,Te 0/16	Static	VLAN0.1,VLAN0.100	Standard	Administratively Up	L2	Trunk

Displaying VLAN information

Use the following procedure to display VLAN information.

1. Select the DCB tab on the Switch Administration panel.
2. Select the VLAN tab.

VLAN information is displayed .

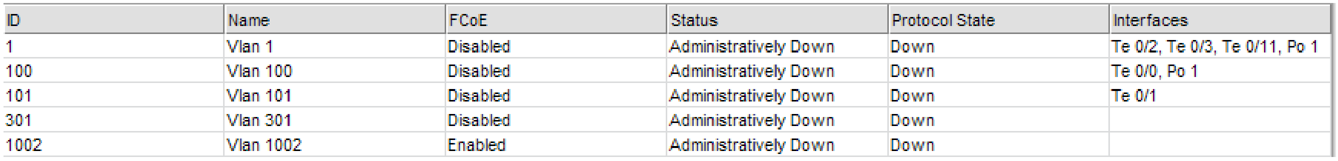

ID	Name	FCoE	Status	Protocol State	Interfaces
1	Vlan 1	Disabled	Administratively Down	Down	Te 0/2, Te 0/3, Te 0/11, Po 1
100	Vlan 100	Disabled	Administratively Down	Down	Te 0/0, Po 1
101	Vlan 101	Disabled	Administratively Down	Down	Te 0/1
301	Vlan 301	Disabled	Administratively Down	Down	
1002	Vlan 1002	Enabled	Administratively Down	Down	

Displaying FCoE login groups

Use the following procedure to display FCoE login group information.

1. Select the DCB tab on the Switch Administration panel.
2. Select the FCoE Login tab.

FCOE login group information is displayed .

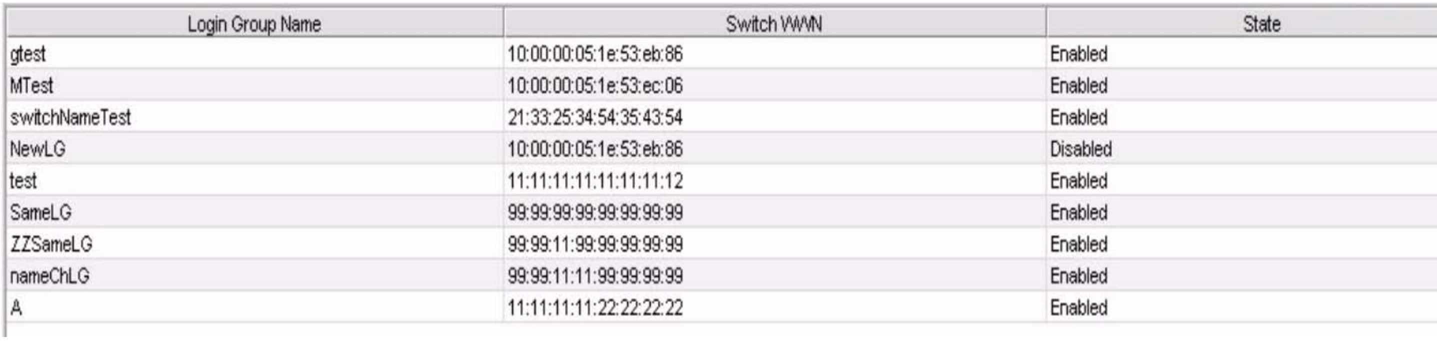

Login Group Name	Switch WWN	State
gtest	10:00:00:05:1e:53:eb:86	Enabled
MTest	10:00:00:05:1e:53:ec:06	Enabled
switchNameTest	21:33:25:34:54:35:43:54	Enabled
NewLG	10:00:00:05:1e:53:eb:86	Disabled
test	11:11:11:11:11:11:11:12	Enabled
SameLG	99:99:99:99:99:99:99:99	Enabled
ZZSameLG	99:99:11:99:99:99:99:99	Enabled
nameChLG	99:99:11:11:99:99:99:99	Enabled
A	11:11:11:11:22:22:22:22	Enabled

Displaying QoS information

Use the following procedure to display QoS information.

1. Select the DCB tab on the Switch Administration panel.
2. Select the QoS tab.

From the QoS tab, you can select the DCB Map tab to display DCB Map information or select the Traffic Class Map tab to display Traffic Class Maps information.

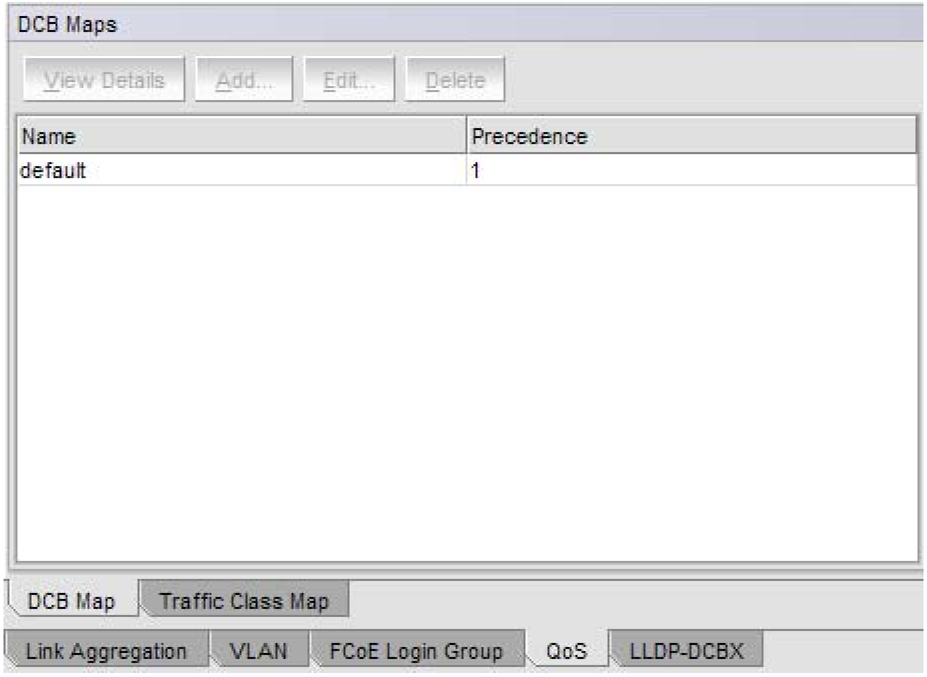

Name	Precedence
default	1

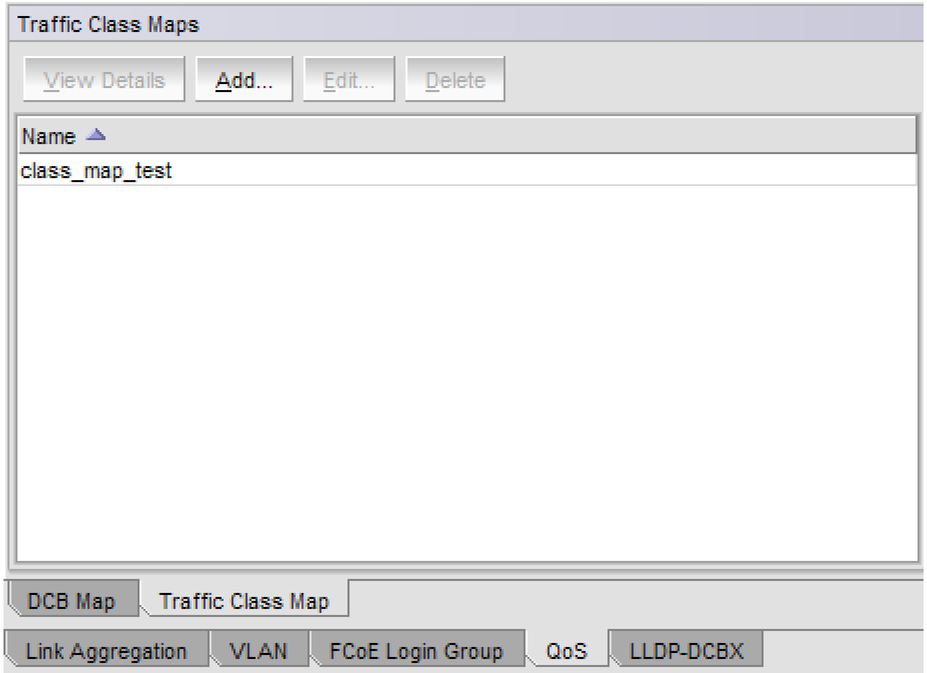

Name ▲
class_map_test

Displaying LLDP-DCBX information

Use the following procedure to display LLDP-DCBX information.

1. Select the DCB tab on the Switch Administration panel.
2. Select the LLDP-DCBX tab.

- To display global settings, select the Global tab.

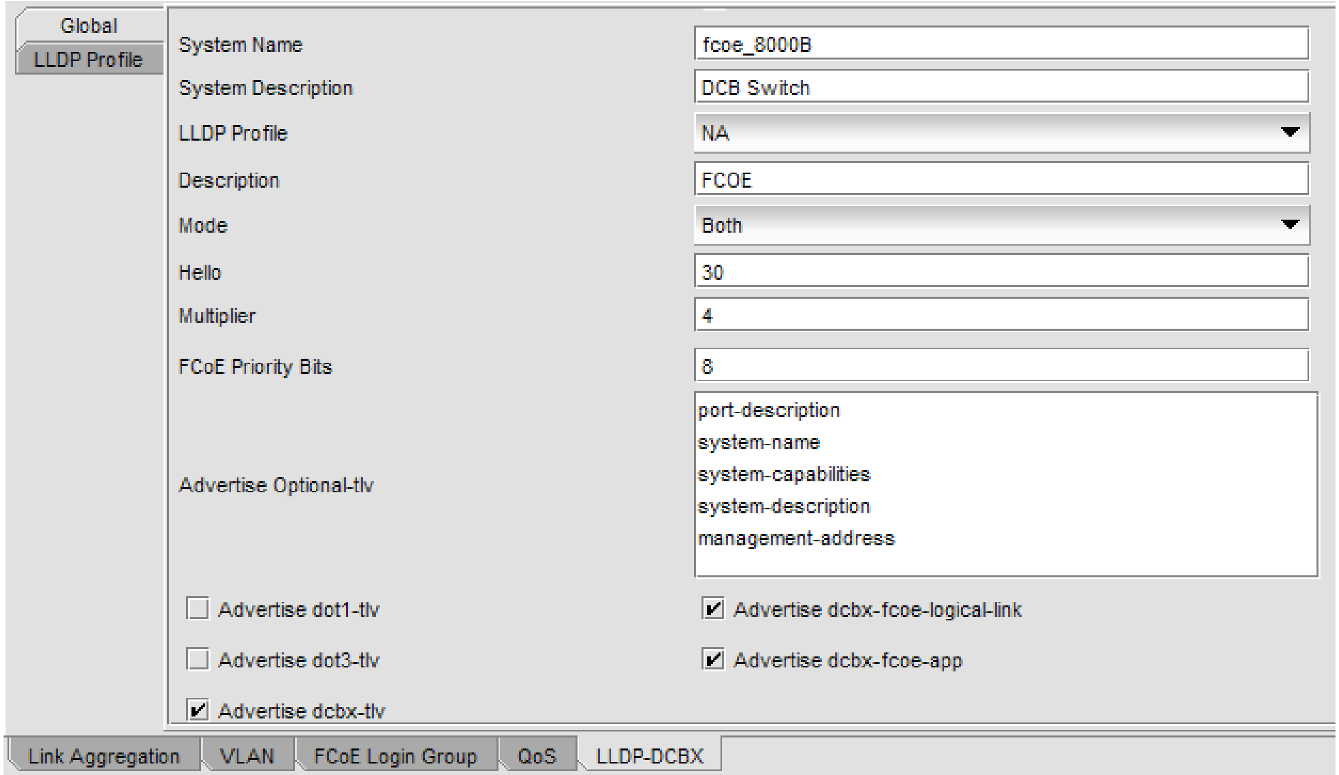

System Name	fcoe_8000B
System Description	DCB Switch
LLDP Profile	NA
Description	FCOE
Mode	Both
Hello	30
Multiplier	4
FCoE Priority Bits	8
Advertise Optional-tlv	port-description system-name system-capabilities system-description management-address

☐ Advertise dot1-tlv
☑ Advertise dcbx-fcoe-logical-link
☐ Advertise dot3-tlv
☑ Advertise dcbx-fcoe-app
☑ Advertise dcbx-tlv

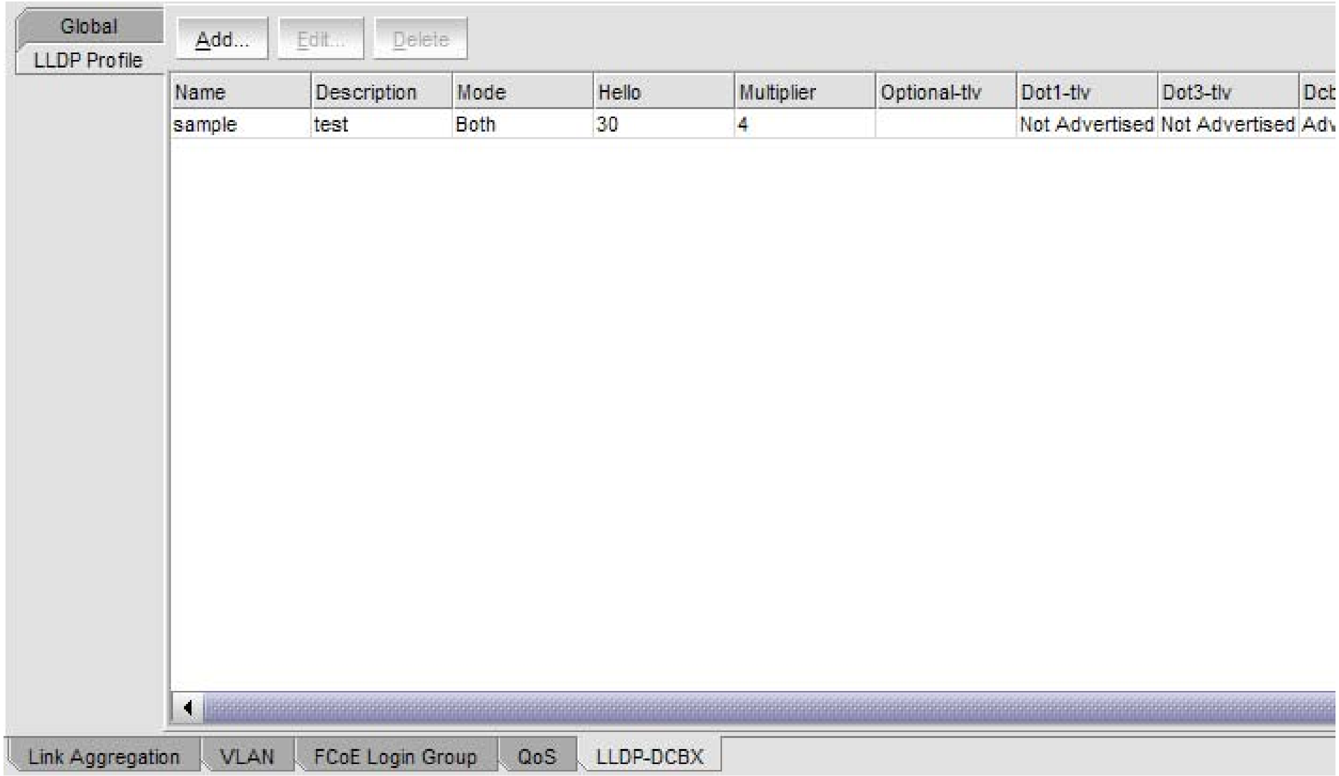

Name	Description	Mode	Hello	Multiplier	Optional-tlv	Dot1-tlv	Dot3-tlv	Dcb
sample	test	Both	30	4		Not Advertised	Not Advertised	Adv

Displaying DCB interface statistics

The DCB interface Port Statistics tab shows basic and advanced statistics, and allows you to change statistics collection parameters. Use the following procedure to display DCB interface statistics.

1. Select the DCB Interfaces tab on the Port Administration panel.
2. Under the DCB Interface Explorer, select a port.
3. Select the Port Statistics tab.

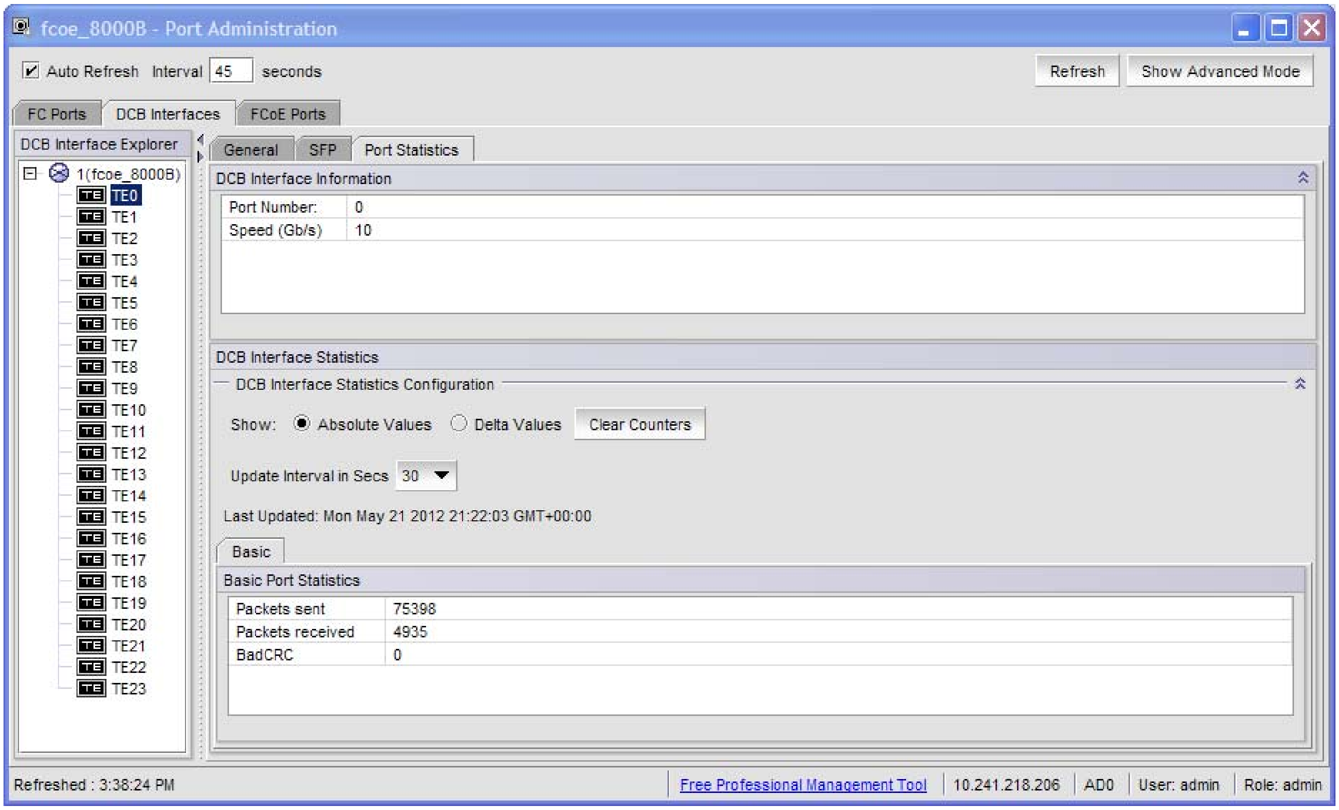

☑ Auto Refresh Interval 45 seconds

Port Number:	0
Speed (Gb/s)	10

Show: ● Absolute Values ○ Delta Values

Update Interval in Secs 30

Last Updated: Mon May 21 2012 21:22:03 GMT+00:00

Packets sent	75398
Packets received	4935
BadCRC	0

Refreshed : 3:38:24 PM | 10.241.218.206 | AD0 | User: admin | Role: admin

The DCB Interface Statistics Configuration section allows you to do the following:

- Toggle between showing Absolute Values or Delta Values (values that have changed since the last data collection).
- Use the Clear Counters button to clear the counters in port statistics.
- Change the retrieval interval.

To view additional information, select Show Advanced Mode. An Advanced tab and an Error Detail tab are added next to Basic Mode.

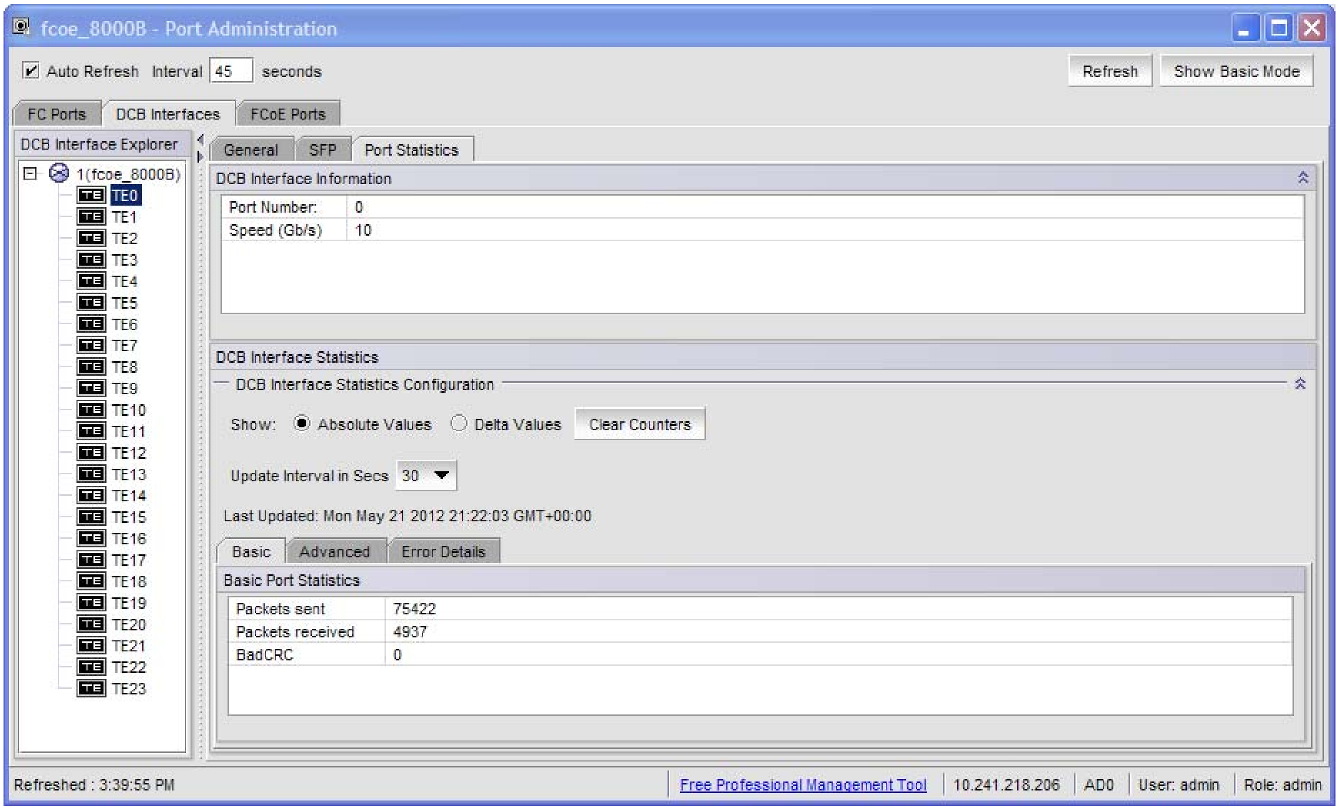

☑ Auto Refresh Interval 45 seconds

DCB Interface Information

Port Number:	0
Speed (Gb/s)	10

DCB Interface Statistics

DCB Interface Statistics Configuration

Show: ● Absolute Values ○ Delta Values

Update Interval in Secs 30

Last Updated: Mon May 21 2012 21:22:03 GMT+00:00

Basic Port Statistics

Packets sent	75422
Packets received	4937
BadCRC	0

Refreshed : 3:39:55 PM | 10.241.218.206 | AD0 | User: admin | Role: admin

The Advanced tab shows DCB transmission statistics

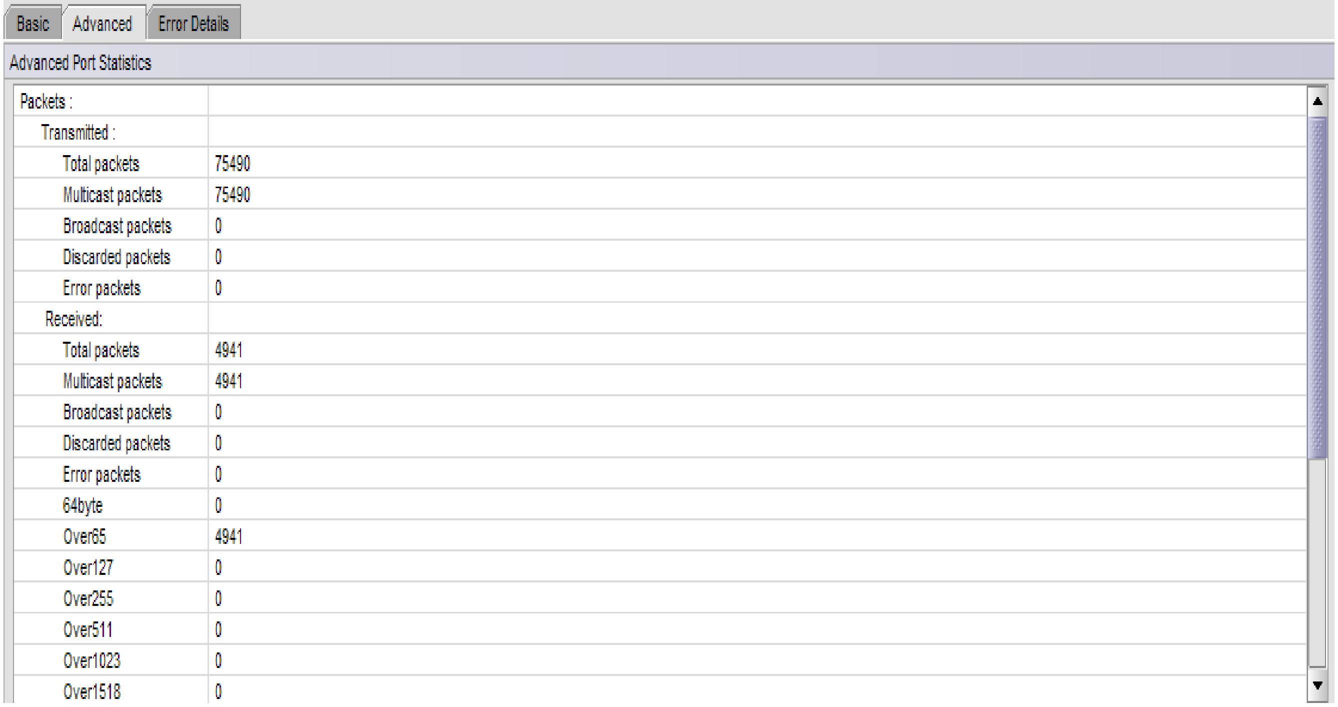

Advanced Port Statistics

Packets :	
Transmitted :	
Total packets	75490
Multicast packets	75490
Broadcast packets	0
Discarded packets	0
Error packets	0
Received:	
Total packets	4941
Multicast packets	4941
Broadcast packets	0
Discarded packets	0
Error packets	0
64byte	0
Over65	4941
Over127	0
Over255	0
Over511	0
Over1023	0
Over1518	0

The Error Details tab shows transmission error statistics

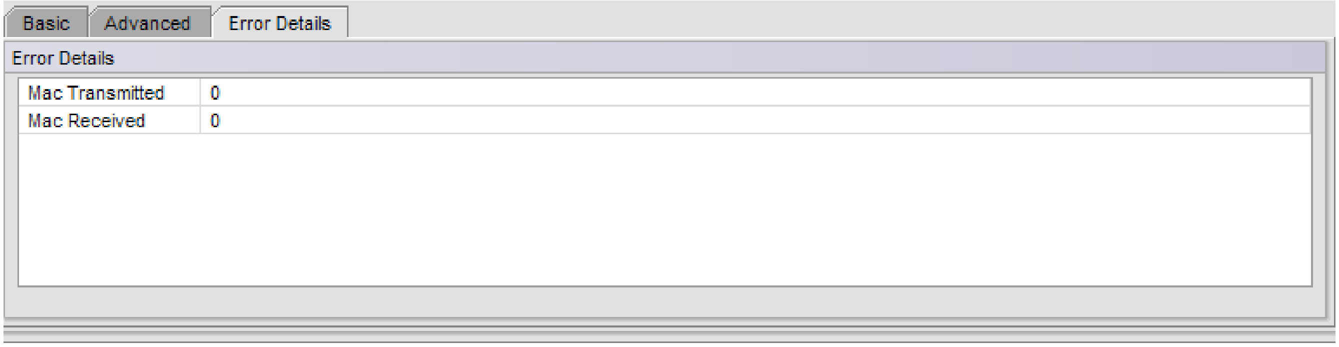

Error Details

Mac Transmitted	0
Mac Received	0

Enabling and disabling a DCB interface

DCB interfaces can be enabled and disabled from a right-click menu on the Switch View, or from the Port Administration panel.
To enable or disable a DCB interface from the Switch View, perform the following steps.

1. Right-click the port to display the right-click menu.

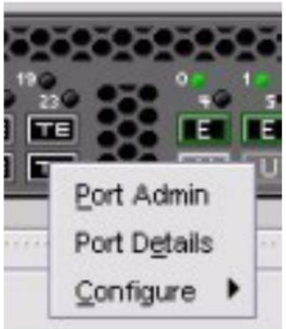

2. Select Configure to display the Enable and Disable options.

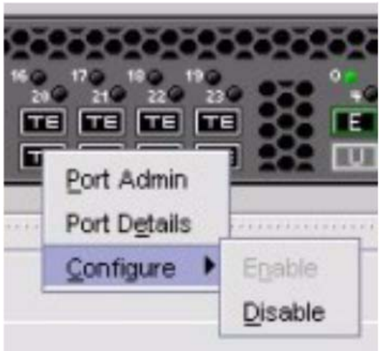

To enable or disable a DCB interface from the Port Administration panel, do the following:

3. Select the DCB Interfaces tab on the Port Administration panel.
4. Under the DCB Interface Explorer, select the port you want to enable or disable.
5. Select the General tab. This tab is normally pre-selected. You can follow either of the following options to enable or disable the interface:
   - Click Enable Interface or Disable Interface to enable or disable the interface, as desired.
   - Click Edit Configuration to open the DCB Edit Configuration dialog box. Select Enable or
6. Disable for Status to enable or disable the interface.

Enabling and disabling a LAG

To enable or disable a LAG, perform the following steps.

1. From the Switch Administration panel, select the DCB tab.
2. Select the Link Aggregation tab.
3. Click Add.
   The LAG Configuration dialog box is displayed.
4. Change the Status to Administratively Up or Administratively Down.

Enabling and disabling LLDP

To enable or disable LLDP on a DCB interface, perform the following steps.

1. Select the DCB Interfaces tab on the Port Administration panel.
2. Under the DCB Interface Explorer, select the port.
3. Select the General tab.
4. Select Edit Configuration.
   The DCB Edit Configuration dialog box is displayed.
5. For LLDP Status, select Enable or Disable.

Enabling and disabling QoS priority-based flow control

Priority-based flow control (PFC) can be used to control network congestion. PFC can be used to selectively pause lower priority traffic classes to ensure that high priority and delay-sensitive traffic are not affected by network congestion. For example, if a large storage transfer is monopolizing the network and causing congestion, PFC can be used to pause the storage transfer so other traffic may use the network.
To enable or disable PFC, perform the following steps.

1. Select the DCB tab on the Switch Administration panel.
2. Select the QoS tab.
3. Select the DCB Maps tab.
4. Under Priority Group, enable or disable Priority Flow Control Status per each Priority Group ID.

Enabling and disabling FCoE ports

You can enable and disable FCoE Ports individually from the Port Administration panel.

1. Select the FCoE Ports tab on the Port Administration panel.
2. Select the port you want to enable or disable under the FCOE Ports Explorer, or from the list.
3. Click Enable or Disable to change the current status of the port.

You can also enable or disable by selecting Edit Configuration, and selecting Enable or Disable on the FCoE Edit Configuration dialog box

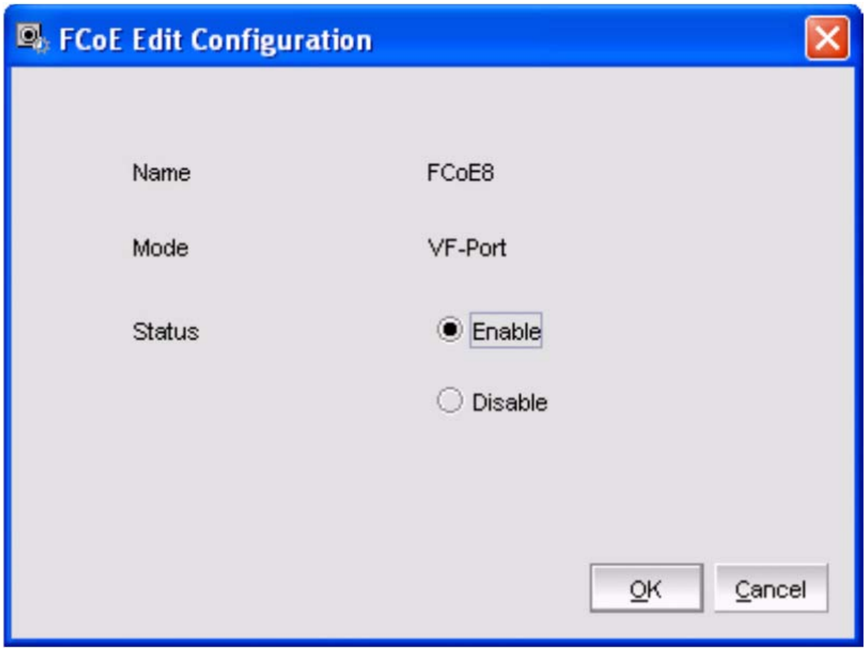

Name FCoE8

Mode VF-Port

Status ● Enable ○ Disable

Chapter 11: Brocade 7800

Overview of Brocade 7800 Extension Switch

The Brocade 7800 Extension Switch is intended as a platform for Fibre Channel over IP (FCIP). This enables transmission of Fibre Channel data over long distances via IP networks by wrapping Fibre Channel frames in IP packets. Each end of the FCIP communication path must be a compatible FCIP device, either the Brocade 7800 or the FX8-24 blade in a DCX-family chassis.

A minimum level of Brocade Fabric Operating System (FOS) 6.3 is required to use the Brocade 7800. Refer to the *Fabric OS Administrator's Guide* for information on configuring these features. The base model of the switch is shipped with six Fibre Channel SFP ports and two physical Gigabit Ethernet (GbE) ports active. It includes FOS 6.3 and is compatible with the entire Brocade switch family. It can operate independently or in a fabric containing multiple Extension Switches.

Fully licensed Brocade 7800 provides the following functionality features:

- FCIP capability
- Up to 8 FCIP tunnels.
- Each FCIP tunnel is represented and managed as a virtual Fibre Channel E_Port (VE_Port).
- Fibre Channel Routing Services functionality can be used over the FCIP link.
- Fabrics connected through FCIP merge if the ports are configured as VE_Ports, and do not merge if one end of the connection is configured as a VEx_Port. If VE_Ports are used in a Fibre Channel Routing Services backbone fabric configuration, then the backbone fabric merges but the Ex_Port attached to edge fabrics do not merge.
- FCIP Trunking with load balancing and network-based failure recovery
- Adaptive Rate Limiting
- Configurable maximum and minimum committed bandwidth per FCIP tunnel
- Minimum rate is guaranteed rate
- FC frame compression before FCIP encapsulation
- Fibre Channel Routing
- SO-TCP with reorder resistance
- FastWrite over FCIP (not over FC)
- Open Systems Tape Pipelining over FCIP
- XRC acceleration and FICON tape pipelining over FCIP
- FICON CUP

- FCIP QoS
- TCP performance graphing in Web Tools

The Brocade 7800 provides the following hardware features:

- Up to 16 Fibre Channel SFP ports supporting Fibre Channel Routing Services with link speeds up to 1, 2, 4, or 8 Gbps
- Up to six 1 GbE ports supporting the FCIP and Fibre Channel Routing Services features with transmit link speeds up to 1-Gbps on each port:
- Two ports (ge0 and ge1) can be configured for use with either copper or optical cables.
- Two PPC440EPx Processors running @ 667 MHz.
- One GoldenEye2 switch ASIC for 1/2/4/8 Gbps FC switching.
- One Cavium CN 5740 running with eight MIPS cores @ 750 MHz for data path processing
- One Blaster FPGA for FC compression, offloads like chksum generation/checks, etc.
- One 10/100/1000 Base-T Ethernet port for management interface.
- One RJ45 terminal port.
- One USB port that provides storage for firmware updates, output of the supportSave command and storage for configuration uploads and downloads.
- Two redundant, hot-swappable combined power supply/fan assembly FRUs.
- Five internal temperature sensors.

Feature comparison - base 7800 and with the Upgrade License

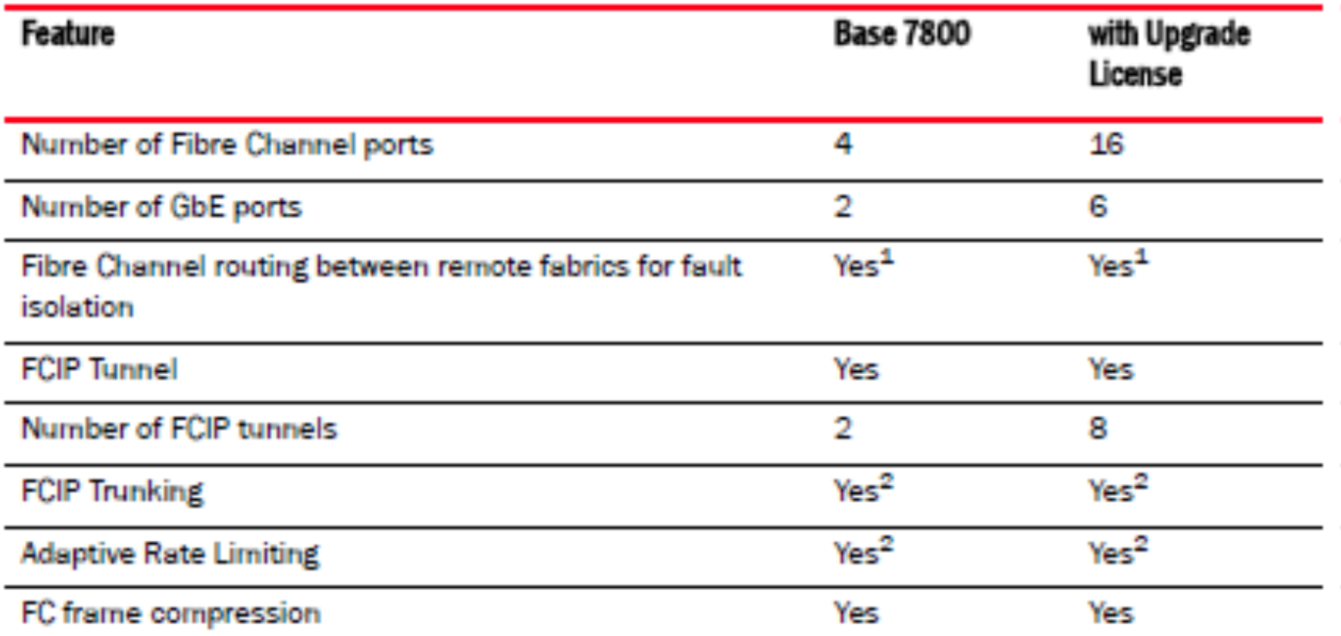

Feature	Base 7800	with Upgrade License
Number of Fibre Channel ports	4	16
Number of GbE ports	2	6
Fibre Channel routing between remote fabrics for fault isolation	Yes[1]	Yes[1]
FCIP Tunnel	Yes	Yes
Number of FCIP tunnels	2	8
FCIP Trunking	Yes[2]	Yes[2]
Adaptive Rate Limiting	Yes[2]	Yes[2]
FC frame compression	Yes	Yes

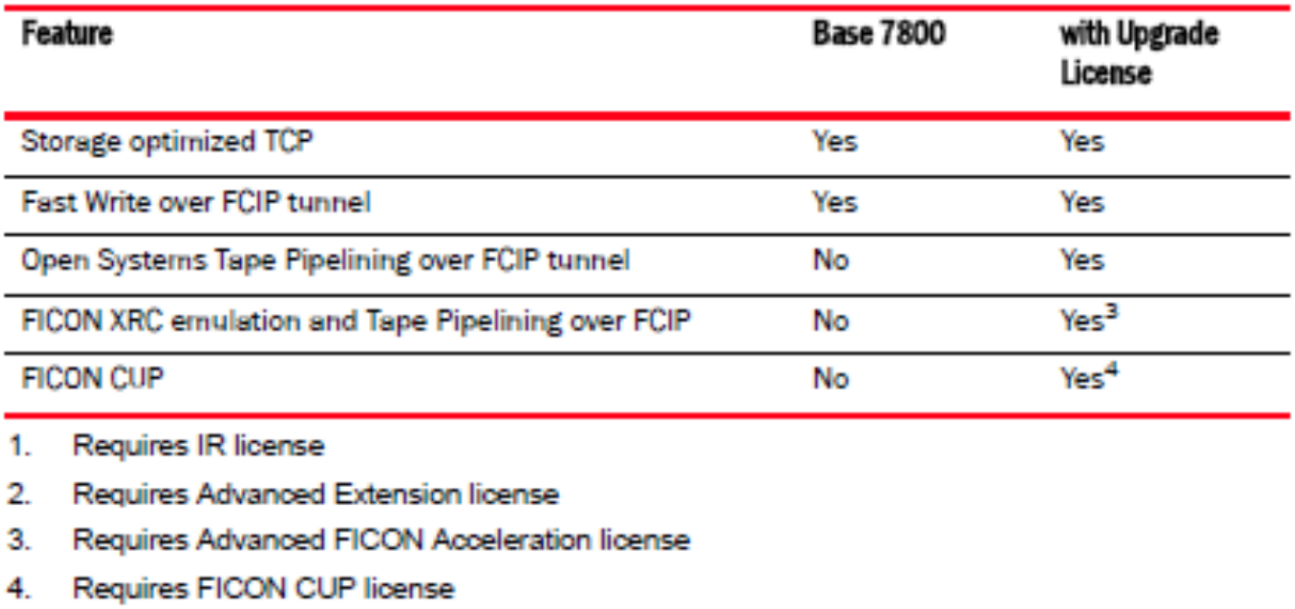

Feature	Base 7800	with Upgrade License
Storage optimized TCP	Yes	Yes
Fast Write over FCIP tunnel	Yes	Yes
Open Systems Tape Pipelining over FCIP tunnel	No	Yes
FICON XRC emulation and Tape Pipelining over FCIP	No	Yes[3]
FICON CUP	No	Yes[4]

1. Requires IR license
2. Requires Advanced Extension license
3. Requires Advanced FICON Acceleration license
4. Requires FICON CUP license

- Before the installation of the Upgrade License, ports beyond the basic four FC and two GbE are shown as *Disabled* with the switchShow command.
- On the base 7800, the two GbE ports (ge0 and ge1) can be configured for use with either copper or optical cables (physically separate ports provided).
- FC frame compression is not the same as IP compression and is disabled by default. It can be enabled using the portCfg command.
- FCIP tunnel bandwidth has a minimum rate of 1544 Kbps (T1 rate). Configuration requests of lower rates will be rejected.
- FCIP Trunking is available which will “virtualize” two or more TCP connections (circuits) as part of a single FCIP tunnel. Up to four circuits can be configured for a single FCIP tunnel.
- Multiple FCIP tunnels can share the same GbE port. At the same time, VE_ and VEx_Ports are not associated with a single physical GbE port.

Available licenses

The following features are available with the purchase of a specific license key for the Brocade 7800.

- Advanced Extension
- Integrated Routing (IR)
- Advanced Acceleration for FICON
- FICON CUP
- Extended Fabric
- Adapative Networking
- Server Application Optimization
- ISL Trunking
- Fabric Watch
- Advanced Performance Monitoring

A number of FCIP feature enhancements have been made with Fabric OS v7.0.0 and higher. For example:

- Ability to create circuits with bandwidth larger than 1 Gbps

- Removal of subnet restrictions
- Lossless failover of 10 Gbps ports
- Support for 10 Gbps Adaptive Rate Limiting (ARL)
- Updates to Virtual Fabric support
- Compression enhancements
- QoS enhancements
- Inband management
- Updated latency support

For more on these enhancements please refer to Fabric OS Administrator Guide.

Port Side of Brocade 7800

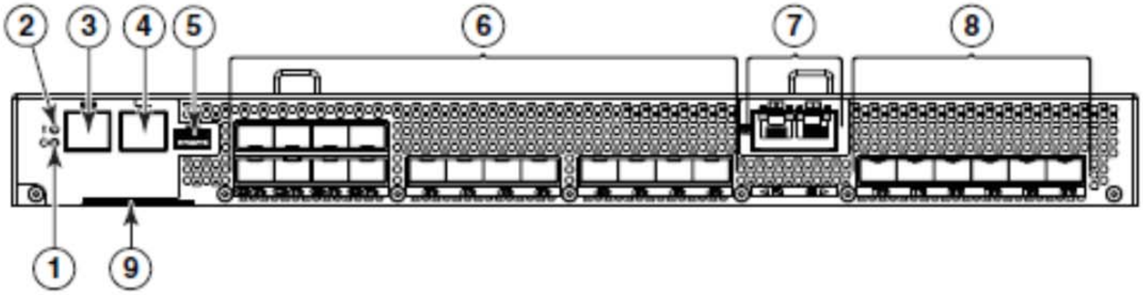

1 System Power LED

2 System Status LED

3 Console Port (RJ45)

4 Ethernet Management Port

5 USB Port

6 Fibre Channel Ports (16)

7 GbE ports - copper RJ45(2)

8 GbE ports - optical SFP (6)

9 Serial number pull-out tab

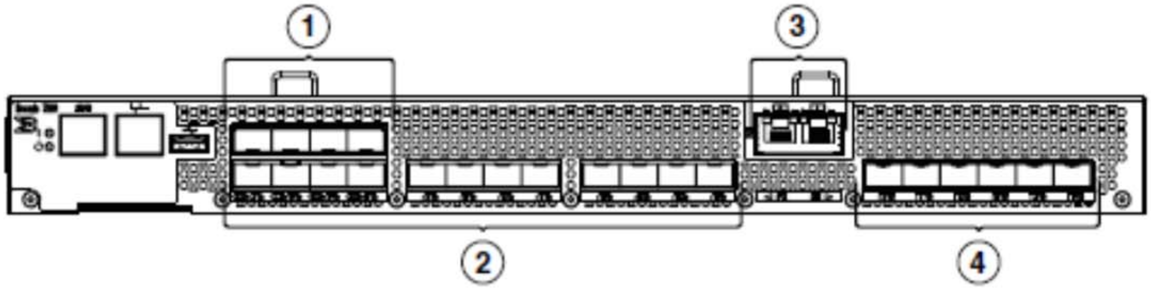

1 Fibre Channel Ports 0 through 3

2 Fibre Channel Ports 4 through 15

3 GbE ports ge0-ge1 (copper only)

4 GbE ports ge0 through ge5 (SFP)

You can have two trunking groups on a fully icensed Brocade 7800. Groups 1 would consist of FC ports 0-7 and group 2 would be ports 8-15.

The GbE ports can only be used once you have configured FCIP and enabled the VE_Ports.

Installing SFPs and cabling the Brocade 7800

Perform the following steps to install SFPs and cable the switch.

1. Install the SFP transceivers in the Fibre Channel ports, as required. The ports selected for use in trunking groups must meet specific requirements.
2. If you have chosen to use the optical ports for ge0 and ge1, install those SFPs. If you have licensed the additional GbE ports, install the SFP transceivers in GbE ports ge2

For instance, to select the optical option for port ge0, use the following command.

```
switch:admin> portcfggemediatype ge0 optical
```

3. Connect the cables to the transceivers.
4. Check the LEDs to verify that all components are functional.
5. Verify the correct operation of the Brocade 7800 by entering the **switchShow** command from the workstation.

```
sw7800:admin> switchshow
switchName: sw7800
switchType: 83.3
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 220
switchId: fffcdc
switchWwn: 10:00:00:05:1e:55:a2:00
zoning: ON (DEFAULT_CFG_LSAN)
switchBeacon: OFF
FC Router: ON
FC Router BB Fabric ID: 1
Index        Port         Address      Media        Speed        State    Proto
==============================================
0            0            dc0000       id           N8           No_Light FC
1            1            dc0100        id          N8           No_Light FC
2            2            dc0200       id           N8            No_Light FC
.
.
.
(output truncated)
21           21            dc1500       --           –            Offline VE
22           22           dc1600       --           --           Offline VE
23           23            dc1700       --           --           Offline VE
             ge0                        id           1G            No_Light FCIP
             ge1                         id          1G            No_Light FCIP
             ge2                        id           1G           No_Light FCIP
             ge3                        id           1G           No_Light FCIP
             ge4                        id            1G          No_Light FCIP
             ge5                        id           1G           No_Light FCIP
sw7800:admin>
```

7800 Switch license option

Some of the capabilities of the 7800 switch require feature licenses. These include the following:

- The 7800 upgrade license to enable full hardware capabilities, full FCIP tunnel capabilities, support of advanced capabilities like open systems tape pipelining (OSTP), FICON CUP support, and separately licensed advanced FICON acceleration capabilities.
- The Advanced Extension License to enable FCIP trunking and Adaptive Rate Limiting (ARL).
- The Advanced FICON acceleration license to enable accelerated tape read/write and accelerated data mirroring over distance in FICON environments.
- The IR is required for FCR. The IR license is required to configure VEX_ports.

NOTE

FCR is not supported on a 7800 switch that has been partitioned for virtual fabrics.

Fibre Channel routing services configuration

The ports on the Brocade 7800 are initially set to **persistently disabled.**
If you want to enable the FC ports as a standard E_Port or F_port use the **portcfgpersistentenable** command to enable the ports.
If you are using the FC ports as EX_Ports you must configure the Fibre Channel Routing services feature prior to enabling the ports.

FCIP

Fibre Channel over IP (FCIP) enables you to use existing IP wide area network (WAN) infrastructure to connect Fibre Channel SANs. FCIP supports applications such as remote data replication (RDR), centralized SAN backup, and data migration over very long distances that are impractical or very costly using native Fibre Channel connections. FCIP tunnels are used to pass Fibre Channel I/O through an IP network. FCIP tunnels are built on a physical connection between two peer switches or blades. Fibre Channel frames enter FCIP through virtual E_ports (VE_ports or VEX_ports) and are encapsulated and passed to TCP layer connections. The TCP connections insure in-order delivery of FC frames and lossless transmission. The Fibre Channel fabric and all Fibre Channel targets and initiators are unaware of the presence of the IP network.

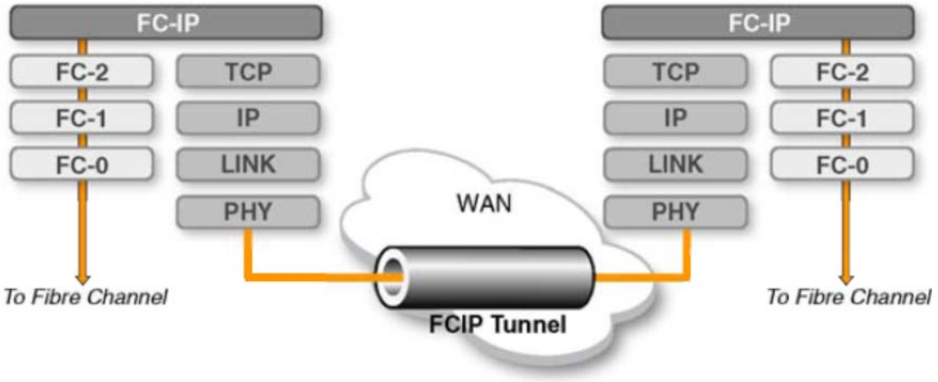

Because FCIP uses TCP connections over an existing wide area network, consult with the WAN carrier and IP network administrator to be sure that the network hardware and software equipment operating in the data path can properly support the TCP connections. When consulting, keep the following in mind:

- Routers and firewalls that are in the data path must be configured to pass FCIP traffic (TCP port 3225) and IPSec traffic, if IPsec is used (UDP port 500).
- To enable recovery from a WAN failure or outage, be sure that diverse, redundant network paths are available across the WAN.
- Be sure the underlying WAN infrastructure is capable of supporting the redundancy and performance expected in your implementation.

VE_Ports and FCIP tunnels on the 7800 switch

A 7800 switch can support eight VE_Ports. VE_Ports are numbered from 16 to 23. Each FCIP tunnel is identified with a VE_port number. Up to eight FCIP tunnels may be created. The 7800 switch supports VEX_ports to avoid the need to merge fabrics.

Configuration preparation

Before you begin to configure FCIP, do the following:

- Determine the amount of bandwidth that will be required for the RDR, FICON or tape application to be deployed.
- The WAN link has been provisioned and tested for integrity.
- Cabling within the data center has been completed.
- Equipment has been physically installed and powered on.
- Make sure you have admin access to all switches and blades you need to configure.
- For the 7800 switch, determine if copper or optical ports will be used for GbE ports 0 and 1.
- For the FX8-24 blade, determine which of the three possible GbE port operating modes will be used.
- Obtain IP addresses for each GbE port you intend to use, plus the netmask and MTU size.

- Determine the gateway IP address and netmask as needed for each route across the WAN. You may also assign a metric to each route to prioritize their use based on expected performance.
- Determine if there is any reason to turn off selective acknowledgement (SACK). Because SACK improves performance for most installations, it is turned on by default.
- Determine the VE_port numbers you want to use. The VE_port numbers serve as tunnel IDs.
- Determine source and destination IP addresses for circuit 0, and the minimum and maximum committed rates for circuit 0. These values are set by the **portCfg fciptunnel create** command.
- Determine how many additional FCIP circuits you want to create. You will need the source and destination IP addresses for the circuit, and the minimum and maximum committed rates for the circuit. You will need to know if you intend to assign metrics to circuits to implement standby circuits. For all circuits except circuit 0, these values are set by the **portCfg fcipcircuit create** command.

Configuration steps

The following is a list of the major steps in configuring FCIP on the 7800 switch or FX8-24 blade:

- Persistently disable VE_ports.
- If required, configure VEX_ports.
- For the 7800 switch, set the media type for GbE ports 0 and 1.
- Assign IP addresses to the GbE ports.
- Create one or more IP routes using the **portCfg iproute** command.
- Test the IP connection using the **portCmd –ping** command.
- Create FCIP tunnels and FCIP circuits, and enable or disable features.
- Persistently enable the VE_ports.

Setting VE_ports to persistently disabled state

VE_Ports used on an FCIP tunnel must be persistently disabled before you can configure FCIP tunnels. You must change their state from persistently enabled to persistently disabled. Once the
FCIP tunnels have been fully configured on both ends of the tunnel, you can persistently enable the ports.

1. Enter the **portCfgShow** command to view ports that are persistently disabled.
2. Enter the **portCfgPersistentDisable** command to disable any VE_ports that you will use in the FCIP tunnel configuration.

Configuring VEX_ports

If you are going to use a VEX_port in your tunnel configuration, use the **portCfgVEXPort** command to configure the port as a VEX_port. VEX_Ports can be used to avoid merging fabrics over distance in FCIP implementations. If the fabric is already connected, disable the GbE ports and do not enable them until *after you have configured the VEX_Port.* This prevents unintentional merging of the two fabrics.

The following example configures a VEX_port, enables admin, and specifies fabric ID 2 and preferred domain ID 220:

```
switch:admin> portcfgvexport 18 -a 1 -f 2 -d 220
```

Configuring the media type for GbE ports 0 and 1

Two media types are supported for GbE ports 0 and 1 on the 7800 switch; copper and optical. The media type must be set for GbE ports 0 and 1 using the **portcfggemediatype** command. The command options are as follows: ge0|ge1 ge0 for port 0 or ge1 for port 1. copper|optical The media type.

The following example configures port 1 (ge1) in optical mode.

```
switch:admin> portcfggemediatype ge1 optical
```

When you enter this command without specifying <media_type>, the current media type for the specified GbE port is displayed as in the following example.

```
switch:admin> portcfggemediatype ge1
Port ge1 is configured in optical mode
```

Configuring a GbE or XGE port IP address

You must configure an IP address, netmask, and an MTU size for each GbE port that you intend to use. This is done using the **portCfg ipif** create command. The following examples create the addressing needed for the basic sample configuration in figure

The following command creates an IP interface for port ge0 on the FX8-24 blade in slot 8 of the
Brocade DCX-4S.

```
switch:admin> portcfg ipif 8/ge0 create 192.168.1.24 255.255.255.0 1500
```

The following command creates an IP interface for port ge0 on the Brocade 7800 switch.

```
switch:admin> portcfg ipif ge0 create 192.168.1.78 255.255.255.0 1500
```

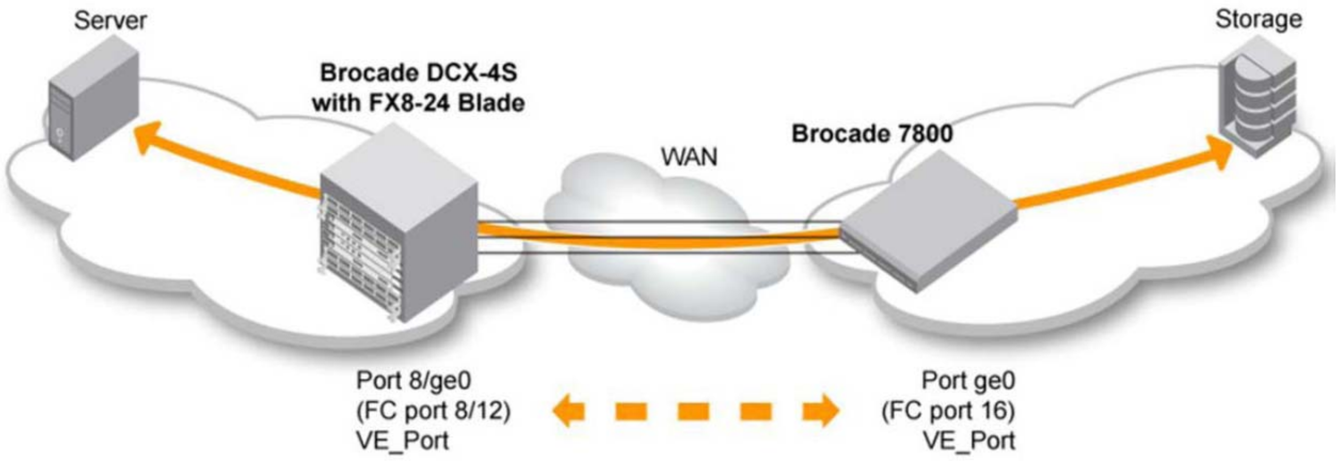

Configuring an IP route

Routing is based on the destination IP address presented by an FCIP circuit. If the destination address is not on the same subnet as the GbE port IP address, you need to configure an IP route with an IP gateway as the destination, using the portCfg iproute create command. Up to 32 IP routes may be defined for each GbE port. Figure 11 adds an IP route for the basic sample configuration.

The following command creates an IP route to destination network 192.168.11.0 for port ge0 on the FX8-24 blade in slot 8 of the Brocade DCX-4S. The route is through local gateway 192.168.1.1.

```
switch:admin> portcfg iproute 8/ge0 create 192.168.11.0 255.255.255.0 192.168.1.1
```

The following command creates an IP route to destination network 192.168.1.0 for port ge0 on the Brocade 7800 switch. The route is through local gateway 192.168.11.1. The metric for the route is 0. The metric should be the same on both ends.

```
switch:admin> portcfg iproute ge0 create 192.168.1.0 255.255.255.0 192.168.11.1
```

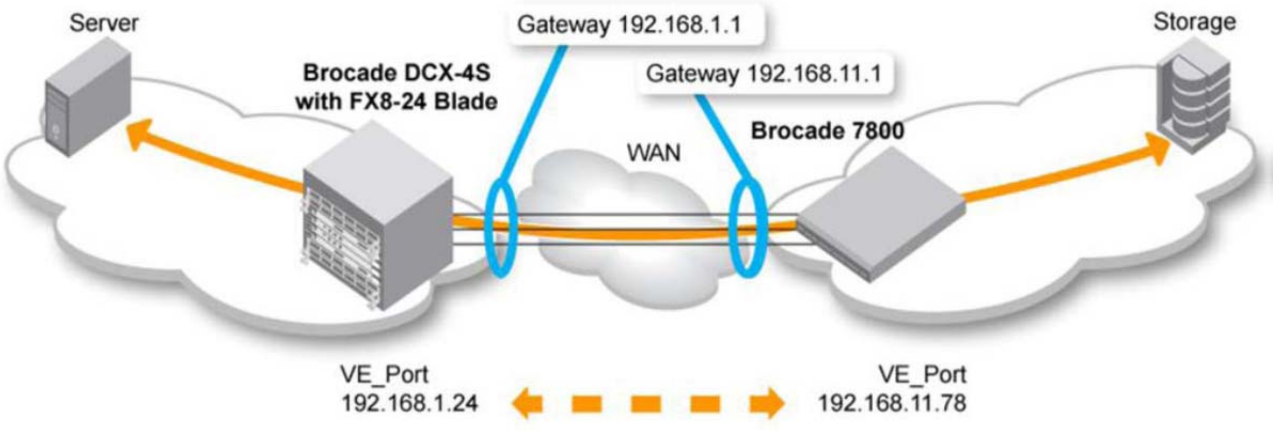

Validating IP connectivity

After you have established the IP interfaces and an IP route, you can issue a portcmd - -ping command to verify connectivity.

The following example tests the connectivity between the FX8-24 blade and 7800 switch in the basic sample configuration from the 7800 switch. The -s option specifies the source address, and the -d option specifies the destination address.

```
switch:admin> portcmd --ping ge0 -s 192.168.11.78 -d 192.168.1.24
```

Creating an FCIP tunnel

FCIP tunnels are created using the **portCfg fciptunnel** create command.
The following command creates the FX8-24 end of the tunnel. VE_port 12 is specified. Circuit parameters are included to create circuit 0. The 7800 switch destination address is specified first, followed by the FX8-24 source address. ARL minimum and maximum committed rates are specified for circuit 0.

```
switch:admin> portcfg fciptunnel 8/12 create 192.168.11.78 192.168.1.24 -b 15500 -B 1000000
```

The following command creates the 7800 end of the tunnel. VE_port 16 is specified. Circuit parameters are included to create circuit 0 on the 7800. The circuit parameters must match up correctly with the circuit parameters on the FX8-24 end of the circuit. The FX8-24 destination address is specified first, followed by the 7800 source address. Matching ARL minimum and maximum committed rates must be specified on both ends of circuit 0.

```
switch:admin> portcfg fciptunnel 16 create 192.168.1.24 192.168.11.78 -b 15500 -B 1000000
```

You can create a tunnel with no circuit parameters. This may be useful in staging a configuration without committing specific circuit parameters.

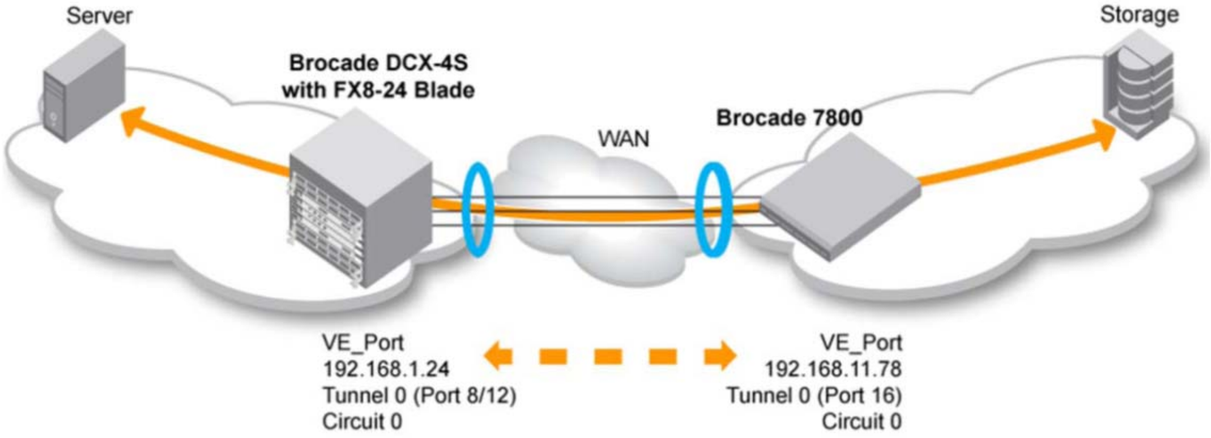

Creating additional FCIP circuits

If the Advanced Extension license is enabled, additional FCIP circuits can be created and added to an FCIP tunnel using the **portCfg fcipcircuit create** command. The following examples adds a circuit to the tunnel in the basic sample configuration.
The following command creates circuit 1 on the FX8-24 end of the tunnel.

```
switch:admin> portcfg fcipcircuit 8/12 create 1 192.168.11.79 192.168.1.25 –b 15500 -B 1000000
```

The following command creates circuit 1 on the 7800 end of the tunnel.

```
switch:admin> portcfg fcipcircuit 16 create 1 192.168.1.25 192.168.11.79 -b 15500 -B 1000000
```

Verifying the FCIP tunnel configuration

After you have created local and remote FCIP configurations, verify that the FCIP tunnel and circuit parameters are correct using the portshow fciptunnel command.

Enabling persistently disabled ports

Ports must be disabled while they are being configured. Before an FCIP tunnel can be used, the associated ports must be persistently enabled.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **portCfgShow** command to view ports that are persistently disabled.
3. After identifying the ports, enter the **portCfgPersistentEnable** command to enable the ports.
4. Enter the **portCfgShow** command to verify the port is persistently enabled.

Modifying an FCIP tunnel

FCIP tunnel characteristics and options can be modified as needed, using the **portCfg fcipTunnel** command with the modify option. The command syntax is as follows:

```
portCfg fciptunnel ve_port modify <options>
Where:
ve_port Each tunnel is assigned to a specific VE_port. The VE_port number serves as
the tunnel ID. The range is 16 through 23.
```

CAUTION
Using the modify option disrupts traffic on the specified FCIP tunnel for a brief period of time.

Modifying an FCIP circuit

FCIP circuit characteristics and options can be modified as needed, using the **portCfg fcipcircuit** command with the modify option. The command syntax is as follows:

```
portCfg fcipcircuit ve_port modify circuit_id <options>
Where:
ve_port Each FCIP tunnel is assigned to a specific VE_port. The VE_port number
serves as the tunnel ID. Specify the VE_Port of the tunnel that contains the
FCIP circuit you want to modify.
circuit_id The numeric ID assigned when the circuit was created.
```

Deleting an IP interface

You can delete an IP interface using the **portcfg ipif** command with the delete option. The command syntax is as follows:

```
portcfg ipif ge<n> delete ipaddr
```

Deleting an IP route

You can delete an IP route to a gateway destination IP address using the **portcfg iproute** with the delete option. The command syntax is as follows:

```
portcfg iproute ge<n> delete dest_IPv4_addr netmask
```

Deleting an FCIP tunnel

When you delete an FCIP tunnel, you also delete all associated FCIP circuits. Use the portCfg **fciptunnel** command with the delete option to delete FCIP tunnels. The command syntax is as follows:

```
portcfg fciptunnel ve_port delete
```

CAUTION
The fciptunnel delete command does not prompt you to verify your deletion. Be sure you want to delete the tunnel before you press Enter.

Deleting an FCIP circuit

You can delete individual FCIP circuits using the portCfg fcipcircuit command with the delete option. The command syntax is as follows:

```
portcfg fcipcircuit ve_port delete circuit_id
```

CHAPTER 12: GETTING TECHNICAL HELP

Perform the following steps before contacting your support contact:

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

    The serial number label is located as follows:

    - *Brocade 300, 4100, 4900, 5100, 5300, 7500, 7800, 8000, VA-40FC, and Brocade Encrypytion Switch*—On the switch ID pull-out tab located inside the chassis on the port side on the left
    - *Brocade 5000*—On the switch ID pull-out tab located on the bottom of the port side of the Switch
    - *Brocade 7600*—On the bottom of the chassis
    - *Brocade 48000*—Inside the chassis next to the power supply bays
    - *Brocade DCX*—On the bottom right on the port side of the chassis
    - *Brocade DCX-4S*—On the bottom right on the port side of the chassis, directly above the cable management comb
3. World Wide Name (WWN)
    - Use the **licenseIdShow** command to display the chassis' WWN.

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

SupportSave

Use this command to collect RASLOG, TRACE, supportShow, core file, FFDC data and other support information to a remote FTP location. On platforms that support USB, the information can also be stored on an attached USB device. On a dual-CP system, information is saved for the local and the remote CP. SupportShow information is available on Active and Standby CPs. To reduce the chance of missing the correct trace dump, supportSave retrieves old (the dump created prior to the current one) and new (the dump triggered by the command) trace dumps.

The files generated by this command are compressed before being sent off the switch. The core files and panic dumps remain on the switch after the command is run. The FFDC data are removed after the command has finished.

If there are blade processor (BP) blades installed on the switch, a support file (a.tar.gz file) is generated from each slot. This command accepts IPv4 and IPv6 addresses. If the configured IP address is in IPv6 format, the RAS auto file transfer and event notification to syslog will not work in the case where the Fabric OS version is downgraded. It is required to reconfigure auto file transfer and syslog with IPv4 IP addresses.
In a Virtual Fabric environment, supportSave saves all chassis-based information and iterates through the defined switch-based information for all logical switches. Chassis permissions are required to execute this command.
System-wide supportSave is supported on platforms running Fabric OS v6.2.0 or later. The command collects support data from the Active CP (and its Co-CPU), the standby CP (and its Co-CPU), and all AP blades

```
switch:admin> supportsave
This command will collect RASLOG, TRACE, supportShow, core file, FFDC data
and other support information and then transfer them to a FTP/SCP server
or a USB device. This operation can take several minutes.
NOTE: supportSave will transfer existing trace dump file first, then
automatically generate and transfer latest one. There will be two trace dump
files transfered after this command.
OK to proceed? (yes, y, no, n): [no] y
Host IP or Host Name: 192.168.126.115
User Name: admin
Password:
Protocol (ftp or scp): ftp
Remote Directory: /temp/support

Saving support information for chassis:HL_5100_66, module:RAS...
Saving support information for chassis:HL_5100_66, module:TRACE_OLD...
Saving support information for chassis:HL_5100_66, module:TRACE_NEW...
Saving support information for chassis:HL_5100_66, module:FABRIC...
Saving support information for chassis:HL_5100_66, module:CORE_FFDC...
Saving support information for chassis:HL_5100_66, module:DIAG...
Saving support information for chassis:HL_5100_66, module:RTE...
Saving support information for chassis:HL_5100_66, module:ISCSID_DBG...
Saving support information for chassis:HL_5100_66, module:AGDUMP...
Saving support information for chassis:HL_5100_66, module:SSHOW_PLOG...
Saving support information for chassis:HL_5100_66, module:SSHOW_OS...
Saving support information for chassis:HL_5100_66, module:SSHOW_EX...
Saving support information for chassis:HL_5100_66, module:SSHOW_FABRIC...
Saving support information for chassis:HL_5100_66, module:SSHOW_SERVICE...
Saving support information for chassis:HL_5100_66, module:SSHOW_SEC...
Saving support information for chassis:HL_5100_66, module:SSHOW_NET...
```

```
......(output truncated)
```

To collect support information on a Brocade 5100 and save it to an attached USB device timeout values are doubled):

```
switch:admin> supportsave -U -d -t 2 mysupportsave

This command will collect RASLOG, TRACE, supportShow, core file, FFDC data
and other support information and then transfer them to a FTP/SCP server
or a USB device. This operation can take several minutes.

NOTE: supportSave will transfer existing trace dump file first, then automatically generate and transfer latest one. There will be two trace dump files transferred after this command.

OK to proceed? (yes, y, no, n): [no] y
Saving support information for chassis:ras095_chassis, module:RAS...
Saving support information for chassis:ras095_chassis, module:TRACE_OLD...
Saving support information for chassis:ras095_chassis, module:TRACE_NEW...
Saving support information for chassis:ras095_chassis, module:FABRIC...
Saving support information for chassis:ras095_chassis, module:CORE_FFDC...
No core or FFDC data files found!
Saving support information for chassis:ras095_chassis, module:DIAG..
Saving support information for chassis:ras095_chassis, module:RTE...
Saving support information for chassis:ras095_chassis, module:ISCSID_DBG...
Saving support information for chassis:ras095_chassis, module:AGDUMP...
Saving support information for chassis:ras095_chassis, module:SSHOW_PLOG...
(output truncated)
```

To run supportSave without confirmation on a Brocade DCX with AP blades included using supportFTP parameters (only Active CP output is shown):

```
switch:admin> supportsave -n -c
Saving support information for chassis:ras020_chassis, module:RAS...........
Saving support information for chassis:ras020_chassis, module:TRACE_OLD...
Saving support information for chassis:ras020_chassis, module:TRACE_NEW...
Saving support information for chassis:ras020_chassis, module:FABRIC.......
Saving support information for chassis:ras020_chassis, module:CORE_FFDC...
Saving support information for chassis:ras020_chassis, slot:4...
slot 4 support file transfer done.
Saving support information for chassis:ras020_chassis, slot:12...
slot 12 support file transfer done.
Saving support information for chassis:ras020_chassis, module:DIAG.....
Saving support information for chassis:ras020_chassis, module:RTE...
Saving support information for chassis:ras020_chassis, module:ISCSID_DBG...
Saving support information for chassis:ras020_chassis, module:AGDUMP...

Saving support information for chassis:ras020_chassis, module:SSHOW_PLOG.....
Saving support information for chassis:ras020_chassis,
module:SSHOW_OS.................................
Saving support information for chassis:ras020_chassis, module:SSHOW_EX.....
Saving support information for chassis:ras020_chassis,
module:SSHOW_FABRIC........
(output truncated)
```

Notes on FOS v7.1.0

Feature Enhancements

16G Platform Specific Enhancements

- ### D_Port Enhancements
  - D_Port (including auto config support) from Brocade 16G Adapter to 16G Switch
  - D_Port support on Access Gateway
  - D_Port on optical ICLs (no Electrical/Optical Loopback support)
  - Other D_Port extensions – users can specify number of frames, frame size, test duration, etc.
- ### Other 16G Platform RAS Enhancements
  - FEC, Credit recovery from Brocade 16G HBA to 16G Switch (Requires HBA driver 3.2)
- ### In-flight Encryption and Compression Enhancements
  - Support more ports for enc/comp at reduced speeds
- ### Fabric Services Enhancements
  - Name Server/Zoning CLI enhancements
  - FDMI enhancements
  - Performance Optimization on long distance ISLs
  - ICL bandwidth balancing on DCX8510
- ### FCR Enhancements
  - Remove support for Interop mode 2 and Interop mode 3
  - In-flight encryption/compression on EX_Ports
  - Pathinfo over FCR
  - Credit recovery on EX_Ports
- ### Additional RAS Enhancements
  - Bottleneck detection, Backend link monitoring, Edge hold time, Credit recovery RASlogs enhancements
  - RASlog Management, Audit log for CLI, CLI history enhancements
  - SFP monitoring, Pathinfo enhancements
- ### **Access Gateway Enhancements**
  - D_Port support (16G only)
  - Credit recovery enhancements (both 8G and 16G)
  - FEC support on F-ports and N-ports (16G only)
- ### **FOS Security and User Management Enhancements**
  - Support for TACACS+ in FOS
  - LDAP support enhancements
  - Support for Open LDAP
  - LDAP RFEs
- ### **FCIP Enhancements**
  - Virtual Fabrics support on 7800

  - IPsec on XGE0 of FX8-24
  - Enhanced stats monitoring and RASlog enhancements
  - Riverbed and Silverpeak WAN Optimizer interoperability

- **FICON Enhancements**
  - FICON support for XISL use
  - FICON CUP support in 7800 Logical Switches
  - NO XISL support in 7800
  - Lossless XISL
  - Various diagnostics enhancements

- **Encryption Platform (BES/FS8-18) Enhancements**
  - Data at rest encryption support for thin-provisioned LUNs
  - KMIP client support
  - Encryption Engine Performance measurement related enhancements
  - Various RFEs

Scalability, Interoperability, Firmware Upgrade

- **Scalability – No change in scalability limits compared to FOS v7.0.x**

- **Interoperability Consideration**
  - Platform running FOS v7.1 does not support EX port configuration in Interop mode 2 or Interop mode 3
  - Device sharing between a switch running FOS v7.1 and McDATA fabrics is allowed via FCR running latest supported FOS v7.0.x firmware

- **Firmware Upgrade Consideration**
  - Non disruptive upgrade to FOS v7.1 from FOS v7.0.x
  - Disruptive upgrade to FOS v7.1 from FOS v6.4.x

Licensing Changes

- **Both the Adaptive Networking and SAO licenses become part of the base FOS v7.1 firmware**
  - No impact to any platform that already has the Adaptive Networking and/or SAO licenses installed on a switch that is upgraded to FOS v7.1
  - Platforms with either of these licenses installed will continue to reflect the license(s) as installed when upgraded to FOS v7.1

REFERENCES

Brocade Fabric OS Adminstrator's Guide v7.1.0

Brocade 6505 Hardware Reference Manual

Brocade 6505 Data Sheet

Brocade 6510 Hardware Reference Manual

Brocade 6510 Data Sheet

Brocade 6520 Hardware Reference Manual

Brocade 6520 Data Sheet

Brocade 8000 Hardware Reference Manual

Brocade 8000 Data Sheet

Brocade 5300 Hardware Reference Manual

Brocade 5300 Data Sheet

Brocade 300 Hardware Reference Manual

Brocade 300 Data Sheet

Brocade 7800 Hardware Reference Manual

Brocade 7800 Data Sheet

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.