



Implementing an IBM b-type SAN with 8 Gbps Directors and Switches

Learn about the latest additions to the IBM b-type portfolio

Refresh and enhance your skills and awareness

Increase your SAN knowledge



Jon Tate
Uwe Dubberke
Michael Engelbrecht
Shanmuganthan Kumaravel
Jose Rodriguez Ruibal

ibm.com/redbooks

Redbooks



International Technical Support Organization

**Implementing an IBM b-type SAN
with 8 Gbps Directors and Switches**

March 2011

Note: Before using this information and the product it supports, read the information in “Notices” on page xvii.

Eleventh Edition (March 2011)

This edition applies to Data Center Fabric Manager v10.1.4 and Fabric Operating System v6.4.x.

© Copyright International Business Machines Corporation 2000-2011. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	xvii
Trademarks	xviii
Preface	xix
The team who wrote this book	xix
Now you can become a published author, too!	xxii
Comments welcome	xxii
Stay connected to IBM Redbooks	xxiii
Summary of changes	xxv
March 2011, Eleventh Edition	xxv
Chapter 1. Product introduction	1
1.1 Overview of the product	2
1.1.1 Hardware features	2
1.1.2 Hardware naming convention: IBM and Brocade	2
1.2 Product descriptions	3
1.2.1 IBM System Storage SAN24B-4 switch	3
1.2.2 IBM System Storage SAN40B-4 switch	4
1.2.3 IBM System Storage SAN80-B4 switch	6
1.2.4 IBM Converged Switch B32	7
1.2.5 SAN32B-E4(2498-E32)	8
1.2.6 IBM System Storage SAN06B-R (2498-R06)	9
1.2.7 IBM System Storage SAN384B Director	9
1.2.8 IBM System Storage SAN768B Director	11
1.2.9 Brocade 4 Gbps SAN Switch Modules	14
1.2.10 Brocade 8 Gbps SAN Switch Modules	15
1.2.11 Fabric Operating System	16
1.2.12 Management tools	16
1.2.13 Licensing	17
1.2.14 Security	17
1.2.15 Virtual Fabrics	18
1.2.16 Support	19
Chapter 2. Data Center Fabric	21
2.1 Methodologies of current SAN design	22
2.1.1 Availability considerations	25
2.1.2 Benefits of a consolidated SAN design	25
2.2 IBM/Brocade Data Center Fabric value enhanced	25

2.3 IBM/Brocade backbone models	27
2.4 Scalability at the core	30
2.4.1 Scalability advantages	31
2.4.2 Moving a backbone-class switch from core to edge	33
2.5 Blade options.	36
2.6 Advance Feature and Licensing	36
2.6.1 Advanced Feature summary	37
2.6.2 Fabric Operating Systems.	37
Chapter 3. Hardware features.	39
3.1 The hardware	40
3.1.1 Entry level, midrange, and director models.	40
3.1.2 Switch and director model types	40
3.2 Generic features	41
3.2.1 Auto-sensing speed negotiation	41
3.2.2 Zoning	41
3.2.3 Frame filtering	42
3.2.4 Routing	42
3.2.5 Service functions.	43
3.2.6 Port Fencing	43
3.2.7 ISL Trunking	43
3.2.8 Diagnostics	45
3.3 Products and features	46
3.3.1 2499-384.	46
3.3.2 2499-192.	73
3.3.3 2109-M48	78
3.3.4 b-type top of rack switches	85
Chapter 4. Fabric Operating System	91
4.1 Fabric Operating System overview	92
4.2 Fabric Operating System v6.2.0 features	92
4.2.1 New features.	92
4.2.2 Feature descriptions	94
4.3 Fabric OS v6.3.1 and v6.4.0 updates	98
4.3.1 Changes in Fabric OS version 6.3.1	99
4.3.2 Changes in Fabric OS version 6.4.0	101
4.4 Firmware upgrade considerations.	109
4.4.1 Licensing changes	109
4.4.2 Fabric scalability	110
4.4.3 FICON support	112
4.5 Additional important notes and guidelines.	113
Chapter 5. Management tools	117
5.1 Web Tools	118

5.2 Fabric Watch	118
5.3 SNMP	121
5.4 Data Center Fabric Manager	121
5.4.1 Time-saving tools	123
5.4.2 Exceeding service level agreements	124
Chapter 6. Implementation	125
6.1 Implementation	126
6.1.1 Initial setup	126
6.1.2 The command-line interface initial setup	131
6.1.3 SAN768B, SAN384B, and SAN256B configuration procedure	136
6.1.4 Connecting to the switch	139
6.1.5 Setting the switch name	140
6.1.6 The Port Identifier format	142
6.1.7 Setting the date	144
6.1.8 Firmware update	145
6.1.9 SAN256B optional modem setup	147
6.2 SAN32B-3 implementation using EZSwitchSetup	152
6.2.1 Implementing EZSwitchSetup	153
6.2.2 Using Switch Manager to manage a switch	166
6.2.3 Basic troubleshooting with EZSwitchSetup	167
Chapter 7. License administration	169
7.1 Licensed features	170
7.1.1 Ports on Demand	170
7.1.2 Full Fabric	171
7.1.3 8 Gbps	171
7.1.4 Inter-Chassis Link (ICL)	172
7.1.5 Adaptive Networking	172
7.1.6 Frame Based ISL Trunking	172
7.1.7 Fabric Watch	173
7.1.8 Advanced Performance Monitoring	173
7.1.9 Extended Fabrics	174
7.1.10 ISL Trunking	174
7.1.11 Integrated Routing	174
7.1.12 High Performance Extension over FCIP/FC	174
7.1.13 FICON Management Server	175
7.2 Using Web Tools to administer licenses	175
7.2.1 Using Web Tools to check licensed ports	175
7.2.2 Installed licenses	177
7.3 Tips on solving licensing issues	179
Chapter 8. Web Tools	181
8.1 Web Tools walk-through	182

8.1.1	Web Tools, the EGM license, and DCFM	182
8.1.2	System requirements	187
8.1.3	Java installation on the workstation	189
8.1.4	Java plug-in configuration	190
8.1.5	Value line licenses	193
8.1.6	Opening Web Tools	193
8.1.7	Requirements for the examples in this chapter	200
8.1.8	Overview of the Web Tools user interface	201
8.2	Web Tools buttons	204
8.2.1	Status button	205
8.2.2	Temp button	210
8.2.3	Power button	212
8.2.4	Fan button	212
8.2.5	HA button	213
8.2.6	Beacon button	215
8.2.7	Switch Status Policy button	216
8.2.8	Legend button	216
8.3	Name Server task	217
8.4	Zone Admin task	220
8.5	Admin Domain task	220
8.5.1	Requirements for Admin Domains	224
8.5.2	Creating an Admin Domain	225
8.6	Port Admin task	231
8.6.1	Renaming a port	237
8.6.2	Editing the configuration	237
8.6.3	Enabling and disabling a port	241
8.6.4	Persistent enable and persistent disable options for a port	241
8.6.5	Enabling or disabling trunking for a specific port	242
8.6.6	Enabling or disabling N_Port ID virtualization (NPIV)	242
8.6.7	Port swap	242
8.6.8	F_Port Trunking	242
8.6.9	Re-authenticating	245
8.6.10	F_Port BB credit	245
8.6.11	QoS Enable/Disable	245
8.6.12	Port beaconing	245
8.6.13	WWN to N_Port mapping	246
8.6.14	Port Administration window on the SAN256B and SAN768B	248
8.6.15	Port Administration for the FCOE switch	249
8.6.16	Port Administration for the IBM System Storage SAN06B-R	249
8.7	Switch Admin task	250
8.7.1	Switch Administration window layout	251
8.7.2	Switch tab	252
8.7.3	Network tab	255

8.7.4	Firmware Download tab	259
8.7.5	SNMP tab	263
8.7.6	License tab	268
8.7.7	User tab	276
8.7.8	Configure tab	281
8.7.9	Routing tab	286
8.7.10	Extended Fabric tab	286
8.7.11	AAA Service tab	287
8.7.12	Trace tab	290
8.7.13	Security Policies tab	291
8.7.14	FICON CUP tab	296
8.7.15	Trunking tab	296
8.8	Telnet/SSH Client task	296
8.9	Fabric Watch task	297
8.9.1	Alarm Notification tab	299
8.9.2	Threshold Configuration tab	300
8.9.3	Configuration Report tab	308
8.9.4	Memory and CPU Usage monitor with Fabric Watch	308
8.9.5	Modifying settings for switches with one power supply	310
8.9.6	Email Configuration	314
8.10	IBM SAN ICL connectivity	314
8.10.1	Before you begin	315
8.10.2	ICL cabling	317
Chapter 9. IBM System Storage Data Center Fabric Manager		323
9.1	DCFM products	324
9.1.1	DCFM Professional	324
9.1.2	DCFM Enterprise Edition	325
9.1.3	Enhanced Group Management	326
9.1.4	DCFM Enterprise scalability	326
9.1.5	DCFM operating system support	327
9.2	DCFM installation	327
9.2.1	Installation of DCFM Enterprise Edition on Windows platform	328
9.2.2	DCFM server and client	336
9.3	DCFM GUI orientation	337
9.3.1	Front panel	337
9.3.2	Main toolbar	339
9.3.3	Product list	340
9.3.4	Connectivity Map	344
9.3.5	Master Log	349
9.3.6	Performance Legend	350
9.3.7	Minimap	351
9.3.8	Status bar	352

9.3.9 Fabric tracking	353
9.3.10 WWN display	355
9.3.11 Object naming	356
9.4 DCFM Fabric Discovery	356
9.4.1 Seed switch	356
9.4.2 Setting up the discovery	358
9.4.3 DCFM Discovery Verification	362
9.5 DCFM reports	363
9.5.1 Fabric Summary Report and Port Report	364
9.5.2 Generating performance reports	367
9.5.3 Generating zoning reports	367
9.6 Event logs	369
9.7 Performance management	373
9.7.1 Performance measures	375
9.7.2 Collecting performance data	375
9.7.3 Real time performance data	376
9.7.4 Historical performance data	378
9.7.5 Performance thresholds	381
9.7.6 Connection utilization	384
9.8 Encryption configuration	386
9.9 User management	388
9.10 DCFM Server Management Console	390
9.10.1 Changing server port numbers	392
9.10.2 Restoring the database	393
9.10.3 Configuring authentication	395
9.10.4 Capturing technical support information	396
9.10.5 Gathering switch information for support	398
9.10.6 Viewing technical support information	400
9.10.7 HMC upgrade	401
Chapter 10. Host Connectivity Manager	403
10.1 HCM features	404
10.1.1 Software features	404
10.1.2 Tree node pop-up menus	405
10.2 Getting started with HCM software	406
10.2.1 HCM software launch	407
10.2.2 Command line utility	408
10.2.3 HCM configuration data	408
10.2.4 Remembering the password	409
10.2.5 Skipping login	409
10.2.6 Changing an HCM application password	409
10.2.7 Changing an HCM agent password	410
10.2.8 Resetting a password or restoring a factory default password	411

10.2.9	Backing up data after an uninstall	412
10.2.10	Backing up HCM data using HCM	412
10.2.11	Restoring HCM data using HCM	413
10.2.12	HCM main window	414
10.2.13	HCM product icons	415
10.2.14	Discovery	416
10.2.15	Setting up out-of-band discovery for an adapter	416
10.2.16	Logging off HCM	418
10.3	Host configuration	418
10.3.1	Host security authentication	418
10.3.2	Configuring security authentication using the GUI	418
10.3.3	Configuring security authentication using the CLI	420
10.3.4	Buffer credits	421
10.3.5	Basic port configuration	422
10.3.6	Opening the Basic Port Configuration dialog box	422
10.3.7	Port logging level	423
10.3.8	Port speed	426
10.3.9	Frame data field size	428
10.3.10	Persistent binding	428
10.3.11	QoS (HBA only)	430
10.3.12	Path Time Out	433
10.3.13	Target rate limiting	435
10.3.14	Boot over SAN	437
10.3.15	Configuring Boot over SAN	438
10.3.16	Boot code image upload	439
10.3.17	Updating the boot code using the GUI	439
10.3.18	Virtual port configuration	440
10.3.19	Creating a virtual port	440
10.3.20	Deleting a virtual port	442
10.3.21	HCM logging levels	443
10.3.22	Advanced port configuration	444
10.3.23	Opening the Advanced Port Configuration dialog box	444
10.3.24	NPIV	446
10.3.25	Name configuration	446
10.3.26	Exporting the properties for a WWN	451
10.3.27	Importing the properties for a WWN	452
10.3.28	Importing properties in EFCM format	453
10.3.29	VLAN configuration	455
10.3.30	Adding a VLAN	456
10.3.31	Editing a VLAN	457
10.3.32	Removing a VLAN	458
10.4	Monitoring	459
10.4.1	Performance monitoring	459

10.4.2	Polling frequency rate	460
10.4.3	Resetting statistics	462
10.4.4	Master Log	463
10.4.5	Filtering event log entries	464
10.4.6	Application log	465
10.4.7	Syslog support	466
10.4.8	Opening the Syslog Server Configuration dialog box	466
10.4.9	Removing a host server	467
Chapter 11.	Virtual Fabrics	469
11.1	IBM/Brocade Virtual Fabric	470
11.1.1	Virtual Fabrics introduction	470
11.1.2	Logical switches and logical fabrics	470
11.2	What Virtual Fabrics are	472
11.2.1	Logical switch	473
11.2.2	Logical fabric	474
11.2.3	ISL sharing	474
11.2.4	Administrative Domains	474
11.2.5	User accounts	475
11.3	Configuring Virtual Fabrics	475
11.3.1	Changing the context to a different logical switch	476
11.3.2	Enabling Virtual Fabrics	476
11.3.3	Disabling Virtual Fabrics	479
11.3.4	Logical switch management	481
11.3.5	Modifying the base switch	482
11.3.6	Creating a logical switch	483
11.3.7	Deleting a logical switch	488
11.3.8	Displaying the logical switch configuration	489
11.3.9	Changing the fabric ID of a logical switch	490
11.3.10	Changing a logical switch to a base switch	491
11.3.11	Configuring a logical switch for XISL use	492
11.3.12	Creating a logical fabric using XISLs	494
11.4	A real life example of Virtual Fabrics	495
11.4.1	The scenario	496
11.4.2	Enabling Virtual Fabric on the switches	499
11.4.3	Creating logical switches	502
11.4.4	Assigning ports to the newly created switch	504
11.4.5	Creating the base switch	507
11.4.6	Creating a user to manage the Virtual Fabric	509
Chapter 12.	Basic zoning	513
12.1	Zoning in general	514
12.1.1	Mixed fabrics	514

12.1.2 Zone configurations	515
12.2 Zoning using DCFM	516
12.2.1 Administrative Domains	519
12.2.2 Implementing Administrative Domains	520
12.3 Implementing zoning	521
12.3.1 Managing zoning	521
12.3.2 Creating an alias	523
12.3.3 Creating a zone	528
12.3.4 Creating a zone configuration	533
12.3.5 Enabling zone configurations	534
12.3.6 Adding a zone to a existing zone configuration	537
12.3.7 Analyzing a zone configuration	539
12.4 Basic zoning using Web Tools	543
12.4.1 To start zoning with Web Tools	544
12.4.2 Creating an alias	547
12.4.3 Creating a zone	550
12.4.4 Using Web Tools to create a zone	551
12.4.5 Creating a zone configuration	553
12.4.6 Enabling zone configurations	555
12.4.7 Analyzing a zone configuration	557
12.4.8 Zoning and E_Ports	559
12.4.9 Broadcast zone	561
12.5 Backing up a zone configuration	561
12.5.1 Backing up a zone configuration to an FTP server	562
12.5.2 Backing up a zone configuration to a Brocade USB device	566
12.5.3 Downloading a zone configuration from a USB device	572
12.6 Zoning using CLI	576
12.6.1 Using CLI to create a zone	577
12.6.2 Using CLI to create a zone configuration	579
12.6.3 Backing up a zone configuration using the CLI	581
12.6.4 Backing up a zone configuration using a USB drive	582
12.6.5 Downloading a zone configuration from an FTP server	584
12.6.6 Downloading a zone configuration from a USB device	585
Chapter 13. Multiple switches and fabrics	587
13.1 Multiple switch environments	588
13.1.1 Gateway links	588
13.1.2 Buffer credit recovery	589
13.1.3 ISL Trunking	590
13.1.4 Connecting switches over distance	595
13.1.5 Routing policies	601
13.2 Merging fabrics	608
13.2.1 Duplicate domain IDs	610

13.2.2	Zoning configuration conflicts	611
13.2.3	Merging fabrics example	612
13.2.4	Merging with a configuration cleared switch	619
13.2.5	Operating parameter conflicts	620
13.2.6	InteropMode	621
Chapter 14.	Security	623
14.1	User accounts overview	624
14.1.1	User authentication	624
14.1.2	Role-Based Access Control	625
14.1.3	Local database user accounts	626
14.2	Account management	627
14.2.1	Displaying account information	627
14.2.2	Creating an account	628
14.2.3	Modifying User and Account settings	631
14.3	Security protocols	635
14.3.1	Security protocol support	635
14.3.2	Secure file copy	637
14.4	Simple Network Management Protocol	638
14.4.1	SNMP and Virtual Fabrics	640
14.4.2	Security level	641
14.4.3	snmpConfig command	641
14.5	Secure Sockets Layer protocol	645
14.5.1	Browser and Java support	645
14.5.2	SSL configuration overview	646
14.5.3	Certificate authorities	647
14.6	Secure Shell protocol	653
14.6.1	SSH public key authentication	654
14.6.2	Configuring SSH authentication	655
14.7	Telnet protocol	658
14.7.1	Blocking Telnet	658
14.7.2	Unblocking Telnet	659
14.7.3	Listener applications	659
14.8	Ports and applications used by switches	660
14.8.1	Access defaults	660
14.8.2	Port configuration	661
14.9	Security policies	662
14.9.1	ACL policies overview	662
14.9.2	ACL policy management	662
14.9.3	FCS policies	663
14.9.4	Overview of FCS policy management	665
14.9.5	Creating an FCS policy	665
14.9.6	Modifying the order of FCS switches	666

14.9.7	FCS policy distribution	667
14.9.8	DCC policies	669
14.9.9	DCC policy restrictions	670
14.9.10	Creating a DCC policy	670
14.9.11	Creating a device policy	671
14.9.12	Deleting a device policy	672
14.9.13	Activating policy changes	673
14.9.14	SCC policies	673
14.9.15	Creating an SCC policy	674
14.9.16	Authentication policy for fabric elements	674
14.9.17	E_Port authentication	676
14.9.18	AUTH policy restrictions	679
14.9.19	Viewing current authentication parameter settings for a switch	681
14.9.20	Setting authentication protocol used by the switch to DH-CHAP	681
14.9.21	Re-authenticating E_Ports	682
14.9.22	Secret key pairs	683
14.9.23	Viewing a list of secret key pairs in the current switch database	683
14.9.24	Setting a secret key pair	684
14.9.25	Distributing the local ACL policies	685
14.9.26	IP Filter policy	686
14.9.27	Creating an IP Filter policy	687
14.9.28	Cloning an IP Filter policy	687
14.9.29	Saving an IP Filter policy	688
14.9.30	Activating an IP Filter policy	688
14.9.31	Deleting an IP Filter policy	688
14.9.32	IP Filter policy rules	689
14.9.33	IP Filter policy enforcement	691
14.9.34	Adding a rule to an IP Filter policy	692
14.9.35	Deleting a rule in an IP Filter policy	693
14.9.36	Aborting a transaction associated with IP Filter	693
14.9.37	IP Filter policy distributions	693
14.9.38	IP Filter policy restrictions	694
Chapter 15.	Adaptive Networking	695
15.1	Traffic Management	696
15.1.1	Committed rate considerations on FCIP	696
15.1.2	Adaptive Rate Limiting considerations	697
15.1.3	Trunking across multiple FCIP circuits	698
15.1.4	Supported packet loss and delay	698
15.1.5	Scalability considerations	699
15.2	Ingress rate limiting	699
15.2.1	Ingress Rate limiting with the CLI	701
15.2.2	Ingress Rate Limiting with Web Tools	703

15.3	Traffic Isolation	705
15.3.1	TI zone failover	707
15.3.2	FSPF routing rules and traffic isolation	709
15.3.3	TI zone misconfiguration example	711
15.3.4	Supported configurations	712
15.3.5	Virtual Fabric configuration	712
15.3.6	TI zones using CLI	713
15.3.7	Other zoning CLI commands	719
15.3.8	TI zones with DCFM	721
15.4	QoS: SID/BID traffic prioritization	724
15.4.1	QoS zones	726
15.4.2	QoS E_Ports	728
15.4.3	Supported configurations and limitations	729
15.4.4	QoS with CLI	730
15.4.5	Web Tools and QoS Zones	734
15.4.6	DCFM and QoS zones	736
Chapter 16.	Performance monitoring	739
16.1	Performance monitoring with Web Tools	740
16.2	Basic Performance Monitoring	744
16.2.1	Basic Performance Monitoring with Web Tools	745
16.2.2	Throughput examples	745
16.3	Advanced Performance Monitoring	753
16.3.1	Virtual Fabrics considerations	753
16.3.2	Performance Monitors	754
16.3.3	Displaying Performance Monitors with the CLI	755
16.3.4	SID/DID Performance Monitor	755
16.3.5	End-to-end monitoring with DCFM	761
16.3.6	Filter-based performance monitoring	762
16.3.7	ISL performance monitoring	768
16.3.8	Top Talker monitors	769
16.3.9	Top Talkers monitors in port mode	771
16.3.10	Top Talkers monitors in fabric mode	774
16.3.11	Top Talkers monitoring considerations	776
16.3.12	Trunk monitoring	777
16.3.13	Saving and restoring the monitoring configuration	777
16.4	SCSI commands with Web Tools	778
16.4.1	SCSI versus IP traffic	780
16.4.2	ALPA error	781
16.5	Bottleneck detection	781
16.5.1	Latency bottleneck	782
16.5.2	Congestion bottleneck	782

Chapter 17. Health and troubleshooting	785
17.1 SAN Health	786
17.1.1 New features of SAN Health	786
17.1.2 Implementing SAN Health	787
17.2 Error logs	806
17.2.1 Capturing a trace dump	806
17.2.2 The supportsave command	807
17.2.3 DCFM support information	812
17.3 General troubleshooting	812
17.3.1 Troubleshooting device connectivity	813
17.3.2 Trace route	816
17.4 Port Fencing	819
17.4.1 Port Fencing using DCFM	820
17.4.2 Port Fencing using CLI	828
17.4.3 Enabling Port Fencing for E_Port class link loss	834
17.4.4 Testing the configuration	834
17.4.5 Basic troubleshooting commands	835
Related publications	841
IBM Redbooks publications	841
Other resources	841
Referenced websites	842
Help from IBM	843
Index	845

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:


This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
BladeCenter®
DS4000®
DS8000®
eServer™

FICON®
IBM®
Redbooks®
Redbooks (logo) ®
System Storage®

System x®
System z®
TotalStorage®

The following terms are trademarks of other companies:

Snapshot, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

“Do everything that is necessary and absolutely nothing that is not.”

This IBM® Redbooks® publication, written at a Data Center Fabric Manager v10.1.4 and Fabric Operating System v6.4.x level, consolidates critical information while also covering procedures and tasks that you are likely to encounter on a daily basis when implementing an IBM b-type SAN.

The products that we describe in this book have more functionality than we can possibly cover in a single book. A storage area network (SAN) is a powerful infrastructure for consolidation, distance solutions, and data sharing. The quality applications that the IBM SAN portfolio provides can help you take full advantage of the benefits of SAN.

In this book, we cover the latest additions to the IBM b-type SAN family and show how you can implement them in an open systems environment. In particular, we focus on the Fibre Channel Protocol (FCP) environment. We address the key concepts that these products bring to the market and, in each case, we provide an overview of the functions that are essential to building a robust SAN environment.

It is our intent to show *how* to implement the functions and features of the IBM b-type portfolio and, to get the best from this book, you must be familiar with SANs, basic SAN tasks, and the terminology associated with SANs. If not, we advise that you read the following IBM Redbooks publications before you start this one:

- ▶ *Introduction to Storage Area Networks*, SG24-5470
- ▶ *IBM System Storage/Brocade Multiprotocol Routing: An Introduction and Implementation*, SG24-7544

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Jon Tate is a Project Manager for IBM System Storage® SAN Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2 support for IBM storage products. Jon has 24 years of experience in

storage software and management, services, and support, and is both an IBM Certified IT Specialist and an IBM SAN Certified Specialist. He is also the UK Chairman of the Storage Networking Industry Association.

Uwe Dubberke is an IBM Certified Specialist for High End Disk Solutions, working as a field specialist (RDS) for DASD and SAN products in IBM Germany. Since starting in 1990 at IBM he has been responsible for various high-end customers as an Account CE. He has also worked as an SE. Since 1999 he has been a virtual member of the EMEA Central Region Hardware Support Center in Mainz, and since 2005 he has also been a virtual member of the SAN Support Group, also in Mainz. He holds a degree in Electrical Engineering with a specialization in communications engineering from the University of Applied Sciences of Gelsenkirchen (Germany). Uwe has co-authored other Redbooks publications about the DS8000 and SSD.

Michael Engelbrecht is a Senior SSR in IBM Global Technical Services, MTS. He has worked with IBM for 29 years and for the last nine years he has worked for the Hardware Field Support team for SSA Sub Sahara Africa. Before that he was a Networking Specialist with many years of networking experience and a large range of networking equipment, specializing in ATM and Frame relay. He is currently a member of the VFE team for CEE and MEA on all RMSS products, as well as regional support for all SAN products for Sub Sahara Africa.

Shanmuganthan Kumaravel is an IBM Technical Services Specialist for the ITD-SSO MR Storage team of IBM India. He supports SAN, and disk products of both IBM and Hewlett Packard since August 2008. Prior to this he worked for HP product support providing remote support on HP SAN storage products, servers and operating systems including HP UNIX and Linux. Shan is a Brocade Certified SAN Designer (BCSD), Brocade Certified Fabric Professional (BCFP), and an HP Certified System Engineer (HPCSE).

Jose Rodriguez Ruibal is the Technical Sales Leader for the IBM System x® Networking team, based in Montpellier, France, and covering the southwest Europe region. He has more than 12 years of experience in IT, and has worked for IBM for more than eight years. His experience includes serving as Benchmark Manager in the IBM PSSC Benchmark Center in Montpellier, working as an IT Architect for Nokia while living in Finland for three years, and IT Architect and Team Leader for the IBM STG OEM and Next Generation Networks teams in EMEA. Prior to joining IBM, he worked for Red Hat and other consulting firms. He holds an MSC and a BSC in Computer Engineering and Computer Systems from Nebrija University, Madrid. His areas of expertise include Business Development, Strategic OEM Alliances and long-term IT projects in the Telecom, Media and Defense industries, high-level IT architecture and complex solutions design, Linux® and all x86 hardware. Jose has co-authored other Redbooks publications on Linux solutions, on IBM x86 servers and Performance Tuning for x86 servers.

Thanks to the following people for their contributions to this project:

Sangam Racherla

Lori Bideaux

International Technical Support Organization, San Jose Center

Doris Konieczny

IBM Storage Systems Group

Khalid Ansari

Jure Arzensek

George DeBiasi

Brian Cartwright

Gareth Edwards

Kerry Edwards

Sven Eichelbaum

Michael Engelbrecht

Steve Garraway

Joe Hew

Cameron Hildebran

Uwe Hofmann

Thomas Jahn

Kamalakkannan Jayaraman

Mark Kornakiewicz

Jin Su Kim

Carsten Larsen

Andy McManus

Dariusz Myszka

Jeannie Vangness

Sangam Racherla

Pauli Ramo

Simon Richardson

Glen Routley

Chris Seiwert

Marcus Thordal

Eric Wong

The authors of previous versions of this book

A special mention must go to Brocade for their unparalleled support of this residency in terms of equipment and support in many areas throughout. Namely:

Jim Baldyga

Yong Choi

Silviano Gaona

Jason Russo

Brian Steffler
Marcus Thordal
Steven Tong
Mansi Botadra
Brocade Communications Systems

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- Send your comments in an email to:

redbooks@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes

for SG24-6116-10

for Implementing an IBM b-type SAN with 8 Gbps Directors and Switches

as created or updated on March 30, 2011.

March 2011, Eleventh Edition

This revision reflects the addition, deletion, or modification of new and changed information described here.

New information

- ▶ DCFM
- ▶ Virtual Fabrics
- ▶ Host Connectivity Manager
- ▶ Adaptive Networking
- ▶ Fabric OS v6.4.

Changed information

- ▶ Screen captures updated to reflect latest software available at time of writing



Product introduction

In this chapter we describe the IBM System Storage and TotalStorage® SAN b-type family, including the hardware naming conventions (IBM versus Brocade), as well as the various components involved.

1.1 Overview of the product

In this section, we describe the IBM System Storage and TotalStorage SAN b-type family. For the most up-to-date information, see the following website:

<http://www-03.ibm.com/systems/storage/san/index.html>

1.1.1 Hardware features

The b-type fabric directors and switches provide a flexible, intelligent platform for networking storage. With models ranging from entry-level 8-port fabric switches all the way up to 768-port backbones, this family addresses the requirements of both small departments and global enterprises. The 1, 2, 4, 8, and 10 Gbps solutions are available to support high-performance requirements. Express models that are pre-configured with Small Form-factor Pluggable (SFP) optical transceivers are available for several of the switches within the b-type family.

1.1.2 Hardware naming convention: IBM and Brocade

Table 1-1 lists the b-type family products, along with their equivalent Brocade names. Note that the table references the switches using their standard IBM names as well as the IBM type and model throughout this text.

Table 1-1 IBM System Storage SAN b-type family products

IBM name	IBM machine type and model	Brocade name
IBM System Storage SAN24B-4	2498-B24	Brocade 300
IBM System Storage SAN40B-4	2498-B40	Brocade 5100
IBM System Storage SAN80B-4	2049-B80	Brocade 5300
IBM Converged Switch B32	3758-L32	Brocade 8000
IBM System Storage SAN06B-R	2498-R06	Brocade 7800
IBM System Storage SAN384B	2499-192	Brocade DCX-4S
IBM Encryption Switch SAN32B-E4	2498-E32	Brocade BES
IBM System Storage SAN768B	2499-384	Brocade DCX
Brocade 4 Gbps SAN Switch Modules for IBM eServer™ BladeCenter®	32R1813 32R1812	Brocade 4020

IBM name	IBM machine type and model	Brocade name
Brocade 8 Gbps SAN Switch Modules for IBM eServer BladeCenter	42C1828 44X1920 44X1921	Brocade 5470

The b-type family includes the following switches, which will only be referenced in this book. These products include:

- ▶ The entry level SAN06B-R multiprotocol routers as well as the SAN768B and SAN384B router blade. These components are discussed in depth in *IBM System Storage b-type Multiprotocol Routing: An Introduction and Implementation*, SG24-7544-03, available at this website:
<http://www.redbooks.ibm.com/abstracts/sg247544.html?Open>
- ▶ The Converged Switch B32. These components are discussed in depth in *IBM Converged Switch B32*, SG24-7935-00, available at this website:
<http://www.redbooks.ibm.com/redpieces/abstracts/sg247935.html?Open>
- ▶ The Encryption switch SAN32B-E4 as well as the SAN768B and SAN384B encryption blade. These components are discussed in depth in *Implementing the IBM System Storage SAN32B-E4 Encryption Switch*, SG24-7922, available at this website:
<http://www.redbooks.ibm.com/abstracts/sg247922.html?Open>

1.2 Product descriptions

In this section, we provide descriptions of the various product components.

1.2.1 IBM System Storage SAN24B-4 switch

The SAN24B-4 is a high performance scalable switch that provides 8, 16, or 24 fabric ports enabled. With auto-sensing link speeds at 1, 2, 4, and 8 Gbps and a flexible design to configure this switch as a fabric switch or an Access Gateway, it is suitable for small to mid-sized businesses.

Figure 1-1 shows the SAN24B-4 fabric switch.



Figure 1-1 SAN24-B fabric switch

The SAN24B-4 fabric switch requires Fabric OS v6.1.0 or later. The switch offers easy to use Web Tools, 8 Gb FC, Long Distance support, Advanced Zoning, Full-Fabric support, Fiber Watch, Advanced Performance Monitoring, Enhanced Group Management, and ISL Trunking. The base switch also offers eight default ports and Ports on Demand (POD) licenses are available in 8-port increments. With flexible architecture based on GoldenEye2 ASIC, the switch supports F, FL, E, and M Ports at 8 Gbps. The switch also has USB port support for firmware download, configuration upload and download, and supportsave.

It also supports NPIV and Access Gateway which is included in the base FOS.

The SAN24B-4 has a 1U form factor and is a single FRU with no field replaceable parts. The switch has one power supply and three integrated fans.

Important:

- ▶ Access Gateway mode is supported only in 24-port configurations, and *only* 2 GB Brocade branded USB drives are supported on the USB port.
- ▶ The 4 and 8 Gbps link speeds are supported *only* with Brocade branded SFPs.

1.2.2 IBM System Storage SAN40B-4 switch

The SAN40B-4 is a high performance enterprise fabric switch with 40 ports at 8 Gbps link speeds. This switch supports features such as Full-Fabric Support, Fabric Watch, Long Distance support, Advanced Performance Monitoring, Integrated Routing, FICON® CUP, and ISL Trunking. The switch requires Fabric OS v6.1, and port hardware is based on the Condor2 ASIC. One ASIC can support all 40 ports at 1, 2, 4, and 8 Gbps link speeds, and they can be configured as F, FL, E, M, and EX Ports.

It also supports NPIV and Access Gateway which is included in the base FOS.

Figure 1-2 shows the SAN40B-4 fabric switch.



Figure 1-2 SAN40B-4 fabric switch

The SAN40B-4 fabric switch requires Fabric OS v6.1.0 or later. The base model of the switch has 24 ports enabled, and the POD licenses are available in 8 port increments. Integrated Routing is a licensed feature which is supported on every port of the switch and requires a POD license for all 40 ports. The ports on the switch are grouped in 8-port groups matching the trunk group, and with ISL Trunking speeds of up to 64 Gbps can be achieved per trunk.

Dynamic Path selection can be used for optimizing the performance and load balancing, and the switch can be managed using Web Tools. The built-in USB port can be used for firmware download, configuration upload and download, and supportsave, and the switch supports non-disruptive firmware downloads.

New features in Fabric OS v6.2.0 make the switch Virtual Fabric capable. A single physical chassis can be subdivided into two or more logical switches creating a logical fabric with other switches.

Integrated Routing is a licensed feature that is supported on every port of the switch and requires the POD license for all 40 Ports.

Two hot-swappable, redundant 125W power supply and fan assemblies are included with the switch and these are field replaceable units (FRU). Each FRU has an ON/OFF switch AC plug and a power supply/fan status LED, and the switch has a 1U form factor.

Important:

- ▶ The USB port supports *only* 2 GB Brocade branded USB drives.
- ▶ The 4 and 8 Gbps link speeds are supported *only* with Brocade branded SFPs.

1.2.3 IBM System Storage SAN80-B4 switch

The SAN80-B4 is an 80-port, 8 Gbps enterprise fabric switch with 2U form factor. This switch supports features such Full-Fabric Support, Fabric Watch, Long Distance support, Advanced Performance Monitoring, Integrated Routing, FICON CUP and ISL Trunking.

It also supports NPIV and Access Gateway which is included in the base FOS.

Figure 1-3 shows the SAN80-B4 fabric switch.



Figure 1-3 SAN80-B4 fabric switch

The SAN80B-4 fabric switch requires Fabric OS v6.1.0 or later. Port hardware is based on the GoldenEye2 ASIC. Each ASIC can support 32 ports at 1, 2, 4, and 8 Gbps link speeds. The base model of the switch comes with 48 ports enabled, and the POD licenses are available in 16-port increments.

New features in Fabric OS v6.2.0 make the switch Virtual Fabric capable. A single physical chassis can be subdivided into two or more logical switches creating a logical fabric with other switches.

Integrated Routing is a licensed feature that is supported on every port of the switch and requires the POD license for all 80 Ports.

The ports on the switch are grouped in 8-port groups matching the trunk group, and with ISL Trunking speeds of up to 64 Gbps can be achieved per trunk. Dynamic Path selection can be used for optimizing the performance and load balancing, and the switch can be managed using Web Tools.

The built-in USB port can be used for firmware download, configuration upload and download, and supports save, and the switch supports non-disruptive firmware downloads.

The switch has two hot-swappable, redundant 300 W power supplies and three hot-swappable fan assemblies. Both the power supplies and the fan assemblies are field replaceable units, and they have a status LED on them.

Important:

- ▶ The USB port supports *only* 2 GB Brocade branded USB drives.
- ▶ The 4 and 8 Gbps link speeds are supported *only* with Brocade branded SFPs.

1.2.4 IBM Converged Switch B32

The IBM Converged Switch B32 is a top-of-rack Fibre Channel over Ethernet (FCoE) switch in a compact 1U form factor. It features eight 8 Gbps Fibre Channel ports along with 24 Converged Enhanced Ethernet (CEE) ports with 10 Gigabit Ethernet capabilities. The CEE ports are capable of transporting both storage and LAN traffic eliminating the need for separate SAN and LAN adapters and cables.

The IBM Converged Switch B32 connects to servers through Converged Network Adapters (CNA). The consolidated SAN and LAN server ports and corresponding cables simplify configuration and cabling in server cabinets to reduce acquisition costs.

Figure 1-4 shows the IBM Converged Switch B32.



Figure 1-4 IBM Converged Switch B32

The IBM Converged Switch B32 requires Fabric Operating System v6.1.2_cee or later. The IBM Converged Switch B32 is designed to support Fibre Channel over Ethernet (FCoE), Fibre Channel, Converged Enhanced Ethernet (CEE), and traditional Ethernet protocol connectivity for servers and storage.

FCoE is a new protocol that can expand Fibre Channel into the Ethernet environment, and it helps to combine and leverage the advantages of two technologies, Fibre Channel protocol and Ethernet. The IBM Converged Switch B32 offer the following capabilities:

- ▶ A 32-port multiprotocol switch for server I/O consolidation

- ▶ Enterprise-class availability for business continuance
- ▶ Improved return of investment and investment protection
- ▶ Fabric security for mission-critical information
- ▶ The Converged Switch B32 components are discussed in depth in *IBM Converged Switch B32*, SG24-7935-00, available at this website:
<http://www.redbooks.ibm.com/redpieces/abstracts/sg247935.html?open>

1.2.5 SAN32B-E4(2498-E32)

The IBM System Storage SAN32B-E4 Encryption Switch is a high performance 32 port auto-sensing 8 Gbps Fibre Channel switch with data encryption, decryption, and compression features.

This is a SAN fabric solution that has the capability of encrypting data-at-rest for heterogeneous disk LUNs, tape drives, and virtual tape libraries. The encrypting of the data is done using Advanced Encryption Standard (AES) 256-bit algorithms. The encryption and decryption engines provide in-line encryption services with up to 96 Gbps throughput for disk I/O (mix of ciphertext and clear text traffic) and up to 48 Gbps throughput for tape I/O (mix of ciphertext and clear text traffic).

The SAN32B-E4 shown in Figure 1-5 is a 2U form factor for standard 19-inch rack mount.




Figure 1-5 SAN32B-E4

Management is provided by using the same integrated management tools as the rest of the IBM System Storage b-type family. This approach allows simple installation, configuration, and everyday administration.

The SAN32B-E4 offers 32 8 Gbps high-speed FC ports. In addition, it has hot swap redundant power supplies and fans to provide for high availability. The FC ports support universal (F/FL/E/EX/M) port configurations with full duplex, auto-sensing of 1, 2, 4, and 8 Gbps port speeds. This can be set to fixed port speed; speed matching between 1, 2, 4, and 8 Gbps ports. A number of

hot-pluggable 8 Gbps SFPs (small form-factor pluggable) are available for installation into the FC ports on the switch. These are 8 Gbps SW, 10 Km LW and 25 km ELW. Only Brocade branded SFPs will operate in this switch.

The Encryption switch SAN32B-E4 components are discussed in depth in *Implementing the IBM System Storage SAN32B-E4 Encryption Switch*, SG24-7922, available at this website:

<http://www.redbooks.ibm.com/abstracts/sg247922.html?Open>

1.2.6 IBM System Storage SAN06B-R (2498-R06)

IBM System Storage SAN06B-R (2498-R06) is a rack based product with 8 Gbps FC routing, switching capabilities along with the Fibre Channel Over IP (FCIP) hardware feature. This has sixteen 8 Gbps FC ports and 6 GbE Ethernet ports. Figure 1-6 shows the IBM System Storage SAN06B-R (2498-R06).



Figure 1-6 IBM System Storage SAN06B-R

This model is available in 2 different configurations, as a Base Fabric Switch configuration which comes with 4 active 8 Gbps FC ports and 2 active 1 GbE ports. It can be upgraded with an additional 12 8 Gbps FC ports and 4 1GbE ports. Figure 1-6 shows the base ports and the upgrade license enabled ports.

1.2.7 IBM System Storage SAN384B Director

The IBM System Storage SAN384B fabric backbone is a core and edge network switching platform used to interconnect storage devices, hosts, and servers in Storage Area Networks (SANs). The SAN384B is a sixth generation platform, designed to facilitate server, SAN, and data center consolidation while helping to reduce infrastructure and administrative costs for IT environments.

Following the introduction of the industry-leading 384-port 8 Gbps SAN768B fabric backbone in 2008, the new SAN384B scales to 192 ports at up to full 8 Gbps speed through its four modular blade slots, and leverages the same breakthrough technology as the larger model to deliver industry-leading performance, scalability, and energy efficiency. The IBM b-type fabric backbone models were designed to address the data growth and application demands of evolving enterprise data centers.

Figure 1-7 shows the IBM System Storage SAN384B fabric backbone.



Figure 1-7 IBM System Storage SAN384B, no door

The IBM System Storage SAN384B requires Fabric OS v6.2.0. New features in Fabric OS v6.2.0 make the switch Virtual Fabric capable. A single physical chassis can be subdivided into two or more logical switches creating a logical fabric with other switches.

The SAN384B base model includes the following components enclosed in an 8U chassis with doors:

- ▶ Two vertical cable management combs
- ▶ Two control processor blades
- ▶ Two core blades
- ▶ Dual power supplies
- ▶ Two blower fans
- ▶ One exhaust duct kit
- ▶ Ship group
- ▶ Fabric Operating System (v6.2 or later)

Be aware of the following considerations:

- ▶ Fibre Channel (FC) switch blades are not included in the base.
- ▶ Fiber optic transceivers are not included in the base.
- ▶ Customers will need to order the 64-port, 48-port, 32-port, or 16-port 8 Gbps FC Switch blades and populate them with transceivers. The 64-port FC switch blade requires the new mini small form-factor pluggable (mSFP), which is of a reduced width to allow for 60 ports. The mSFP uses the standard SC connector type.

Integrated Routing is a licensed feature that is supported on every FC port of the switch.

Important:

- ▶ Only 2 GB Brocade branded USB drives are supported for use on the USB port.
- ▶ The SAN384B supports all features and functions as indicated and requires Fabric OS v6.2 or later. Blades that use the Condor2 ASIC *must* use Brocade branded SFPs.

You can find more information about the IBM System Storage Fabric Backbones at the following website:

<http://www-03.ibm.com/systems/storage/san/b-type/san384b/>

1.2.8 IBM System Storage SAN768B Director

The IBM System Storage SAN768B Fabric Backbone is a core switching platform that is used to interconnect storage devices, hosts, and servers in a storage area network (SAN). The SAN768B is a fifth generation platform, designed to meet the growing connectivity, virtualization, and the cost efficiency needs of enterprise data centers.

Figure 1-8 shows a view from the front, with no door.



Figure 1-8 Front, side angle, no door of SAN768B

The SAN768B is designed for the following tasks:

- ▶ Deliver breakthrough performance with 8 Gbps Fibre Channel connectivity.
- ▶ Provide long-term scalability to facilitate server and storage expansion by offering the highest density footprint with up to 768 Fibre Channel ports using two SAN768Bs that are connected with Inter-Chassis Links (ICL).
- ▶ Take advantage of proven reliability and new technology to deliver enterprise-class reliability, availability, and serviceability.
- ▶ Support multiprotocol infrastructure for both Fibre Channel and IP traffic.
- ▶ Improve energy efficiency by combining high bandwidth with low power consumption.

The SAN768B is designed to address key customer requirements while helping to protect investments already made by the deployment of existing SANs, servers, storage hardware, and advanced functions, as follows:

- ▶ It offers forward and backward compatibility with IBM System Storage SAN director, switch, and router models, 1, 2, 4, and 8 Gbps autosensing capability, as well as advanced fabric services and management tools.
- ▶ It enables interoperability between IBM System Storage and TotalStorage b-type and m-type SAN switches and directors.

Preferred for larger midrange to enterprise level SAN applications, the SAN768B Fabric Backbone integrates a new generation of hardware, including a minimum of two control processors, two core blades, four power supplies, and three fans in a 14U rack height, with the following advanced functions:

- ▶ Full Fabric operation and universal port operation on all ports (F_Port, E_Port, FL_Port, M_Port, EX_Port, and N_Port support on selected blades)
- ▶ ISL and ISL Trunking, Advanced Zoning, and FICON CUP
- ▶ Intelligent management and monitoring with Web Tools, Fabric Watch, and Performance Monitor
- ▶ USB port support for firmware download, configuration upload and download, and supportsave

Important: Only 2 GB Brocade branded USB drives are supported for use on the USB port.

- ▶ Adaptive Networking Services with Quality of Service (QoS) is a licensed feature with Fabric Operating System v6.0 and later (Fabric OS) and uses network intelligence to anticipate congestion and to make adjustments in the fabric dynamically so that application traffic continues to flow

Important: The SAN768B supports all features and functions as indicated and requires Fabric OS v6.0 or later. Blades that use Condor2 ASIC *must* use Brocade branded SFPs.

New features in Fabric OS v6.2.0 make the switch Virtual Fabric Capable. A single physical chassis can be subdivided into two or more logical switches creating a logical fabric with other switches.

Integrated Routing is a licensed feature that is supported on every FC port of the switch.

You can find more information about the 768B at the following website:

<http://www-03.ibm.com/systems/storage/san/b-type/san768b/>

1.2.9 Brocade 4 Gbps SAN Switch Modules

Figure 1-9 shows the Brocade 4 Gbps SAN Switch Module.



Figure 1-9 Brocade 4 Gbps SAN Switch Module

The Brocade Module supports 1, 2, and 4 Gbps. You can choose between the 10-port (part number 32R1813) or the 20-port module (part number 32R1812). The modules provide the ability to implement non-disruptive software upgrades. It includes Web Tools and Advanced Zoning, with optional features including Fabric Watch, Advanced ISL Trunking, Extended Fabric Activation, Advanced Security Activation, and Advanced Performance Monitoring.

Access Gateway is a standard feature that is available to all Brocade 4 Gbps SAN Switch Modules with Fabric OS v5.2.1 or later.

For more information about Access Gateway and its implementation, review *Implementing the Brocade Access Gateway for IBM BladeCenter*, REDP-4343, which is available at the following website:

<http://www.redbooks.ibm.com/abstracts/redp4343.html?open>

1.2.10 Brocade 8 Gbps SAN Switch Modules

Figure 1-9 shows the Brocade 8 Gbps SAN Switch Module.




Figure 1-10 Brocade 8 Gbps SAN Switch Module

The Brocade Module supports 2,4, and 8 Gbps. You can choose between the 10-port (part number 44X1921), the 20-port module (part number 44X1920) or the 20-port enterprise model (part number 42C1828). The modules provide the ability to implement non-disruptive software upgrades. It includes Web Tools and Advanced Zoning, with optional features including Fabric Watch, Advanced ISL Trunking, Extended Fabric Activation, Advanced Security Activation, and Advanced Performance Monitoring.

The Enterprise 20-port model (42C1828) is a fully enabled switch with a complete set of licensed features that maximizes performance, ensures availability and simplifies management for the most demanding applications and expanding virtualization environments including Fabric Watch which monitors and creates alerts based on the health of switch and fabric elements.

For a complete list of all the Brocade 8 Gbps SAN switch module features and capabilities, see the following website:

<http://www-03.ibm.com/systems/bladecenter/hardware/openfabric/fibrechannel.html>

For more information about Access Gateway and its implementation, review *Implementing the Brocade Access Gateway for IBM BladeCenter*, REDP-4343, which is available at the following website:

<http://www.redbooks.ibm.com/abstracts/redp4343.html?Open>

1.2.11 Fabric Operating System

Fabric Operating System (Fabric OS) provides enterprise-class, ultra-high availability, reliability, and security capabilities for a wide range of SAN environments. Fabric OS runs on the b-type family of Fibre Channel directors and switches, and provides transparent interoperability between 1, 2, 4, 8, and 10 Gbps devices as well as the reliable, high-performance data transport that is critical for scalable SAN fabrics interconnecting thousands of servers and storage devices. Fabric OS v6.x is common across all current members of the IBM System Storage and TotalStorage SAN b-type family.

Brocade Fabric OS v6.4.1 is supported for the following hardware platforms:

- ▶ IBM System Storage SAN24B-4
- ▶ IBM System Storage SAN40B-4
- ▶ IBM System Storage SAN80B-4
- ▶ IBM Converged Switch B32
- ▶ IBM System Storage SAN06B-R
- ▶ IBM Encryption Switch SAN32B-E4
- ▶ IBM System Storage SAN384B
- ▶ IBM System Storage SAN768B
- ▶ FC 8 Gbps, 16-port switch blade
- ▶ FC 8 Gbps, 32-port switch blade
- ▶ FC 8 Gbps, 48-port switch blade. FICON supported in FOS v6.2+
- ▶ FC 8 Gbps, 64-port switch blade
- ▶ FC 4 Gbps Routing blade - 16x 4 Gbps FC ports, 2x GigE FCIP ports. Two max per 384. FC7887 required for optional FCIP
- ▶ FC 10 Gbps, 6-port switch blade. Requires FC2510 or FC2520
- ▶ FCoE 10GbE blade - 24 x 10GbE CEE/FCoE ports. Two max per 384. Can not add any other intelligent blades.
- ▶ FC 8 Gbps Extension blade - 12x 8 Gbps FC ports, 10x GigE FCIP ports. Optional 2x 10 GigE ports with FC7892. Two max per 384.
- ▶ FC Encryption Blade
- ▶ Brocade 4 Gbps SAN Switch Modules for IBM eServer BladeCenter
- ▶ Brocade 8Gb SAN Switch Module for IBM BladeCenter

1.2.12 Management tools

To ensure open fabric management, Fabric OS provides standard management interfaces, a full range of management tools, and an API that enables the development of third-party SAN management applications.

The following tools simplify SAN fabric management by centralizing control and enabling automation of repetitive administrative tasks:

- ▶ **Web Tools:** A built-in Web-based application that provides administration and management functions on a per switch basis
- ▶ **Data Center Fabric Manager (DCFM):** A client/server-based external application that centralizes management of IBM/Brocade multiprotocol fabrics within and across data centers, including support for FCoE and CEE
- ▶ **Fabric Watch:** A Fabric OS built-in tool that allows the monitoring of key switch elements: power supplies, fans, temperature, error counters, and so on
- ▶ **SNMP:** A feature that enables storage administrators to manage storage network performance, find and solve storage network problems, and plan for storage network growth

1.2.13 Licensing

Within the b-type family, licensing is performed at both a hardware and software level. The “pay-as-you-grow” flexibility with Ports On Demand allows scalability in 4-port, 8-port, or 16-port increments on the switch platforms.

Features such as ISL Trunking, Advanced Performance Monitoring (APM), Fabric Watch, and Extended Fabrics are software licensed and available across all platforms. CUP for FICON, Adaptive Networking, and Integrated Routing are additional licenses available on certain platforms.

Important: Secure Fabric OS features are now included in Fabric OS v6.0 and later, and Secure Fabric OS is no longer available as a licensed feature.

1.2.14 Security

Security within a SAN varies and can include external security, restricting physical access to directors and switches; software-based security where the use of zoning restricts which hosts and storage can communicate, and hardware-based security where the use of frame filtering monitors each frame and enforces its path through a SAN fabric.

Authentication, Authorization, and Accounting (AAA) services are available through the local switch user and password database or external RADIUS server. In Fabric OS v6.0 and later, the AAA services now support an external Lightweight Directory Access Protocol (LDAP) server running Microsoft® Active Directory Services, and the Fabric OS includes Active Directory/LDAP client.

Fabric OS v6.0 and later also provides support for Federal Information Processing Standards (FIPS).

Standards: FIPS Standards Publication 140-2 was issued by the National Institute of Standards and Technology (NIST) and includes both software and hardware component requirements to handle sensitive, unclassified data. The FIPS standards are used by the departments and agencies of the United States federal government.

Secure Fabric OS

All the features of Secure Fabric OS are migrated to Fabric OS v6.0 and later, and Secure Fabric OS is no longer available as a licensed feature.

Role Based Access Control

Fabric OS v6.4.1 uses Role-Based Access Control (RBAC) to control access to all Fabric OS operations. Role-Based Access Control (RBAC) defines the capabilities that a user account has based on the role that the account is assigned. For each role, there is a set of predefined permissions on the jobs and tasks that can be performed on a fabric and its associated fabric elements.

For the Command Line Interface (CLI), you can display a list of all command help topics for a given login level. For example, if you are logged in as user and enter the **help** command, a list of all user-level commands that can be executed is displayed. The same rule applies to the admin, securityAdmin, and the switchAdmin roles.

1.2.15 Virtual Fabrics

With the release of Brocade Fabric OS (FOS) 6.2 and higher, customers have a new option, an ANSI standard based implementation of Virtual Fabrics. The Virtual Fabrics feature adds two new capabilities; Logical Switches and Logical Fabrics. Both are available in base FOS firmware.

With Virtual Fabrics, customers can partition a physical switch into multiple Logical Switches. Each Logical Switch belongs to a Logical Fabric, which has independent data paths, fabric configuration (zoning, Quality of Service (QoS), fabric mode, and so on) and management. With or without Virtual Fabrics, customers benefit from advanced Fabric OS (FOS) features, designed to deliver scalability, performance, and High Availability (HA), with simple management.

The Virtual Fabrics feature is available on 8 Gbps platforms that are Virtual Fabric-capable (VF-capable), including these:

- ▶ IBM System Storage SAN768B Backbone
- ▶ IBM System Storage SAN384B Backbone
- ▶ IBM System Storage SAN80B Switch
- ▶ IBM System Storage SAN40B Switch

For investment protection, products that are not VF-capable, such as the Brocade 48000 Director, earlier 2 Gbps and 4 Gbps FOS and m-series platforms running M-Enterprise OS (M-EOS) software can seamlessly connect to Logical Switches in VF-capable products without any reconfiguration.

To simplify management, customers use the newest management platform, Brocade Data Center Fabric Manager (DCFM). After being created, Logical Switches and Logical Fabrics are managed in exactly the same way as their physical counterparts. Standard FOS Command-Line Interface (CLI) commands can be used to perform configuration and management functions for Virtual Fabrics or to script them.

The Virtual Fabrics feature is described in detail in Chapter 11, “Virtual Fabrics” on page 469.

1.2.16 Support

The IBM Resource Library website provides support for IBM users and is available at the following location:

<http://www.ibm.com/systems/storage/san/b-type/library.html>



Data Center Fabric

The data center is being transformed from a static physical infrastructure of dedicated servers and storage that hosts fixed applications, to a dynamic virtual infrastructure where applications run on virtual servers. In this chapter, we discuss the role of the Data Center Fabric in the evolution of the data center.

2.1 Methodologies of current SAN design

As SAN designs have evolved from 1 Gbps through 4 Gbps and now 8 Gbps, introducing next generation technology (such as the IBM System Storage 768B and 384B) has followed a pattern. To understand this pattern, it is useful to review some of the common methodologies of current SAN design.

Today, most SAN designs use a variant of what is known as a *core-to-edge* network design. In this design, the network elements, typically switches, are designated as either core or edge switches.

In Figure 2-1 we show a core-to-edge SAN design.




Figure 2-1 Core-to-edge SAN design with director type switches (SAN256B)

As we can see, the edge switches connect node devices such as servers and the core switches connect the storage devices and the edge switches. For this reason, core switches are sometimes called *backbone* switches. The core-to-edge, two-layer design has the following advantages:

- ▶ **Storage:** Because it is usually the most expensive component of the SAN, storage needs to be attached to the most powerful component of the FC network. The goal is to have fully utilized storage ports. The total bandwidth going through the inter-switch links (ISL) to the core switch depends on the bandwidth of the core switch to storage port connection.
- ▶ **Server:** Gathering servers connected to the edge switch allows many servers to be funneled into the core of the network.
- ▶ **Separated SAN design:** The use of a core switch for edge duty is appropriate for nodes that have a *one-to-many* relationship in the network, such as storage elements.
- ▶ **Scalability:** Rapid growth of the network infrastructure can be resolved by scaling by adding more switches, at the core and at the edge.
- ▶ **High availability:** Firstly, the two fabric design as described in Figure 2-1 on page 22 helps to address the requirement for high availability in the case of a single switch failure in the fabric, or a total fabric failure. Secondly, the core switches have built in high availability components.
- ▶ **Reliability:** Switches that are functioning as the core of the SAN network connecting storage are more reliable than the edge switches which connects node devices as servers. They use more sophisticated error detection and correction mechanisms to ensure reliability of data, for example:
 - Bus monitoring and control of blades and other field-replaceable units (FRUs).
 - Dual control processors that enable hot, non-disruptive fast firmware upgrades.

Because most SAN communications occur between servers and storage elements, the SAN design must provide the proper bandwidth to funnel communications for all server connections through to the storage elements.

The IBM technology of the new core SAN768B/SAN384B switches can enable the next *enterprise-class SAN design* by providing higher speeds and a denser fabric core, allowing a denser concentration of both server and storage connections to the fabric.

Similar to the design decisions that SAN architects made as Fibre Channel speeds moved from 2 Gbps to 4 Gbps, the transition from 4 Gbps to 8 Gbps will allow the most current FOS SAN infrastructure to use the IBM SAN768B/SAN384B at the core of the fabric and propagate directors, such as the SAN256B, towards the edge, as shown in Figure 2-2.

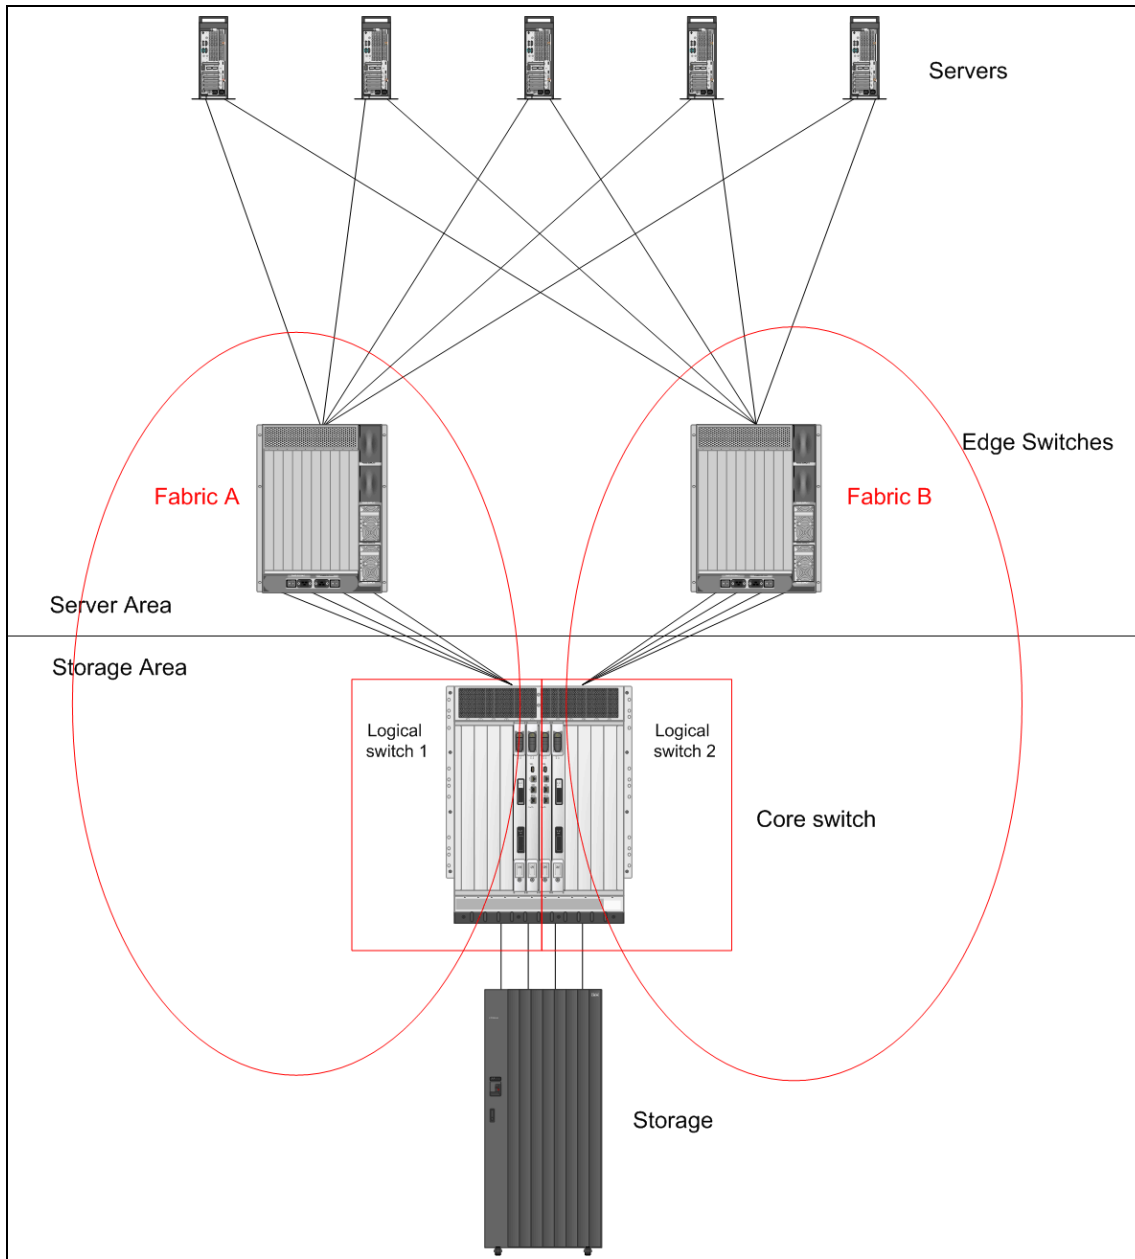


Figure 2-2 Core-to-edge design with IBM System Storage SAN768B/SAN384B as a core and SAN 256B as edge switches

2.1.1 Availability considerations

For high availability reasons, we keep the two fabric design. The IBM/Brocade backbone switch has been divided into two logical switches. Virtual fabrics and logical switches are described in Chapter 11, “Virtual Fabrics” on page 469.

2.1.2 Benefits of a consolidated SAN design

This type of consolidated SAN design has the following benefits:

- ▶ Ease of management by reducing the number of SAN switch elements in the fabric
- ▶ Use of fewer ISLs, as well as fewer optics for ISLs; 8 Gbps ISL provide improved performance
- ▶ Improved availability characteristics
- ▶ More connectivity with lower power and cooling profiles than the equivalent number of edge switches
- ▶ Improved data center space efficiencies
- ▶ Greater bandwidth on the core/storage edge of the SAN
- ▶ Reducing overhead and complexity
- ▶ Improving resource utilization
- ▶ Increasing performance
- ▶ Deploying long-term solutions

2.2 IBM/Brocade Data Center Fabric value enhanced

IBM/Brocade introduced the new class of 8 Gbps modular switching platform, which is designed for the next generation enterprise data centers. It supports open system and mainframe environments.

The building blocks for the IBM/Brocade Data Center Fabric are the IBM System Storage SAN768B and IBM System Storage SAN384B, and are called *backbones*.

IBM/Brocade backbones, a new class of fabric infrastructure, delivers the high-performance, non-disruptive scalability and continuous availability necessary for converged data center fabrics. The IBM/Brocade solution accelerates the transformation of today's physical data center into tomorrow's virtual data center.

The IBM/Brocade backbone delivers this performance at excellent value on an expandable technology platform designed to seamlessly add tomorrow's advanced technologies.

Here are the main advantages of the IBM/Brocade Backbone solution:

- ▶ Un-congested 8 Gbps throughput (up to 4 times more powerful than other offerings):
 - Offers greater server, storage, network, and data center consolidation
 - Broadens server virtualization with more virtual machines on fewer physical servers
 - Reduces equipment, facility, and overhead costs
 - Scales efficiently and non-disruptively to meet storage growth and application demands without increasing complexity
- ▶ A 10X energy efficiency advantage:
 - Provides lowest power consumption, cooling (BTU/hour), and carbon emissions
 - Frees limited power and cooling resources for more servers (physical/virtual) and storage arrays
 - Lowers energy costs and helps achieve “green” initiatives
- ▶ High-speed, highly secure fabric based encryption:
 - Prevents exposure to sensitive information
- ▶ Native b/m-series connectivity
- ▶ Adaptive Networking:
 - Optimizes fabric behavior and ensures ample bandwidth for critical applications
- ▶ Logical partitioning:
 - Provides ability to logically separate fabrics
- ▶ Future-built for FCoE/CEE:
 - Reduces and simplifies server network connections as new protocols emerge

2.3 IBM/Brocade backbone models

There are two IBM/Brocade backbone models:

- ▶ IBM System Storage SAN384B
- ▶ IBM System Storage SAN768B

Their features are described in Chapter 1, “Product introduction” on page 1.

In this and the following chapters, we introduce details of the backbone family of switches.

The differentiation between *director type* and *backbone* switches is as follows:

- ▶ Director:
 - Designed primarily to provide Layer 2 Fibre Channel/FICON connectivity
 - At the core of midsize enterprise data center SANs
 - At the edge of large data center fabrics
 - Requires high bandwidth (moderate 8 Gbps utilization), RAS, and energy efficiency
 - Has moderate virtual server demands (hundreds of virtual machines; less than 20 virtual machines per host; Tier 2 applications)
- ▶ Backbones:
 - Require the highest bandwidth (extensive 8 Gbps utilization)
 - Have extensive virtual server demands (several hundred to thousands of virtual machines; generally 15 or more virtual machines per host; Tier 1 applications in addition to Tier 2 applications)
 - Want integrated FC routing as opposed to using a special blade (and blade slot)
 - Have security/compliance requirements warranting encryption of data rest on a broad scale (as opposed to using/adding separate encryption switches)
 - Want the ability to logically partition a SAN fabric and manage by application, business group, customer, or traffic type
 - Want the flexibility to deploy direct CEE/FCoE network connections in the future as opposed to using a separate CEE/FCoE “top-of-rack” switch

In Figure 2-3 we show the smaller model of the family: the IBM System Storage SAN384B, which can be used as a mid-size enterprise fabric core or large enterprise edge/application engine.




Figure 2-3 IBM System Storage SAN384B

The following main features are included:

- ▶ Up to 256 FC/FICON ports
- ▶ An 8-slot horizontal card cage:
 - Two Active/Passive Control Processor (CP) Blades (same as the larger model)
 - Two Active/Active Core Switching (CR) Blades:
 - DCX can operate with one CR blade at half slot bandwidth.
 - Each CR blade has two ICL ports.

In Figure 2-4 we show the larger model of the backbone family.




Figure 2-4 IBM System Storage SAN768B

The following main features are included:

- ▶ A 12-slot vertical card cage:
 - Two Active/Passive Control Processor (CP) blades
 - Two Active/Active Core Switching (CR) blades:
 - DCX can operate with one CR blade at half slot bandwidth.
 - Each CR blade has two special Inter Chassis Link (ICL) ports.
- ▶ Up to 512 ports at full 8 Gbps speed
- ▶ 3.072 Tbps Chassis Bandwidth:
 - 256 Gbps bandwidth per slot

The backbone switches have separate (and redundant) control processor and core switching blades. If a control processor blade fails, the second control processor blade (in standby mode) takes over without any degradation in performance.

Consideration: However, the core switching blades are active/active: If a core switching blade fails, all ports will fully operate but with half the aggregate chassis and slot bandwidth.

2.4 Scalability at the core

If one backbone switch as the core is not sufficient to meet the needs of the fabric, there is a solution: Inter Chassis Links (ICLs), which are ISL connections between two CR (core) blades on the switch.

Inter-Chassis Links (ICLs): ICLs harness unused ports to connect the switching backplane of one SAN768B chassis directly with the switching backplane of another SAN768B or SAN384B chassis. This additional connection means that it does not consume usable ports. ICL connections operate as hardware trunked ISLs.

The connection means is through copper cables between each of the core switching blades on the SAN768B or SAN384B chassis. The copper cables are supplied by IBM and are 2 meters in length. Because of the short connectivity distance, chassis connected with ICLs will reside in the same cabinet or in adjacent cabinets. ICL cables can be connected in any fashion from one core blade in one chassis to the other core blade in the other chassis.

The best way to connect is using the same connector or blade in each chassis for simplicity. ICLs are an optional licensed feature of the SAN768B. Feature number 7870 provides two cables and feature numbers 7882 or 7885 provides a license. These features must be ordered for each of the chassis using ICL connections.

Use of ICLs does not collapse the switches domains into a single domain.

Both the SAN768B or SAN384B switches have the special Fibre Channel ICL ports to connect two or three backbone chassis, enabling the SAN768B and SAN384B switch to scale to the following capacity:

- ▶ Dual-chassis: Up to 1024 ports on a SAN768B or 512 ports on a SAN384B.
- ▶ Three chassis: Up to 1536 ports on a SAN768B or 768 ports on a SAN384B.

ICL: FOS V6.3 or higher is required for a three way ICL configuration.

2.4.1 Scalability advantages

The advantages of these connections are as follows:

- ▶ Port count is doubled at the core of the fabric without use of any user ports.
- ▶ Speed is locked at 8 Gbps and cannot be changed.

Attention: Configuration change is not allowed on ICL ports.

- ▶ No SFPs or FC ports are required.

For the IBM System Storage SAN768B:

- ▶ Each ICL port consists of sixteen 8Gbps FC connections.
- ▶ Four ICLs deliver 512 Gbps of bandwidth.

For the IBM System Storage SAN384B:

- ▶ Each ICL port consists of eight 8 Gbps FC connections.
- ▶ Four ICLs deliver 256 Gbps of bandwidth.

For both backbone models, it appears that each ICL is managed as:

- ▶ One (SAN384B) 8-port ISL trunk
- ▶ Two (SAN768B) 8-port ISL trunks

These models have the following features:

- ▶ Frame-based trunking is enabled across each ICL.
- ▶ DPS distributes exchanges across all frame trunks.
- ▶ If an ICL fails, traffic automatically flows over the remaining ICLs.

With the usage of ICL we can preserve E_Ports in the chassis for any other server/storage/switch connection:

- ▶ 64 x 8 Gbps E_Ports/per chassis for SAN768B (sixteen 8 Gbps per ICL * 4 ICLs)
- ▶ 32 x 8Gbps E_Ports/per chassis for SAN384B (eight 8 Gbps per ICL * 4 ICLs)

In Figure 2-5 there are four copper pin ICL ports per chassis which are used to connect two backbones with special 2m ICL cables.




Figure 2-5 ICL ports

Bear in mind that proper cable connections optimize load distribution across ICLs. Each ICL cable has a color coded connector (silver text/black background or black text/silver background). The connection can be done by matching ICL cable connector color with ICL port color on core blades. The multiple supported combinations of two chassis connection are outlined in SAN768B and SAN384B hardware guides (IBMSAN768B/DCX ICL Cable Replacement Procedure or IBM384B/DCX-4S ICL Cable Replacement Procedure).

Figure 2-6 shows a supported connection of two SAN384B switches.

Because the blades in the SAN384B are installed horizontally, the familiar top-to-bottom orientation for other products becomes a left-to-right orientation. Connect the cables from the left (top) connectors of the CR4S-8 blades in the first chassis to the right (bottom) connectors of the CR4S-8 blades in the second chassis. Similarly, connect the cables from the right (bottom) connectors of the CR4S-8 blades in the first chassis to the left (top) connectors of the CR4S-8 blades in the second chassis.

Cables: It is also acceptable to attach the cables from slot 3 on one chassis to slot 3 on the second chassis (or slot 6 to slot 6) as long as the left-to-right (top-to-bottom) rule is followed.




Figure 2-6 ICL cable connection between two SAN384B backbone switches

2.4.2 Moving a backbone-class switch from core to edge

As we mentioned previously, the smaller backbone switch can act as a mid-size enterprise fabric core, or a large enterprise edge/application switch. So, it can be moved to the edge switch area and connected to the core switch.

In 2.2, “IBM/Brocade Data Center Fabric value enhanced” on page 25 we discussed moving the director class switch SAN256B to the edge. By moving this switch to the edge, and connecting it to the core we are forced to use the ports from the core and edge switch as well.

In this case we have not preserved the E_Ports for any other use. On the other hand, the same 8 Gbps, 10 Gbps, blades used in the SAN256B director are also utilized in the IBM/Brocade backbone switches. Backbones, however, deliver four times the slot bandwidth of the SAN256B (256 Gbps as opposed to 64 Gbps).

Figure 2-7 shows the ICL connection between core and edge switches, which includes the following characteristics:

- ▶ 256 Gbps of aggregate ICL bandwidth:
 - Based on ICL bandwidth of smaller backbone switch
- ▶ Moves the backbone-class from the core to the edge with:
 - Full 8 Gb
 - Integrated Routing
 - Virtual Fabrics
 - QoS
 - 2X CS and CP blades

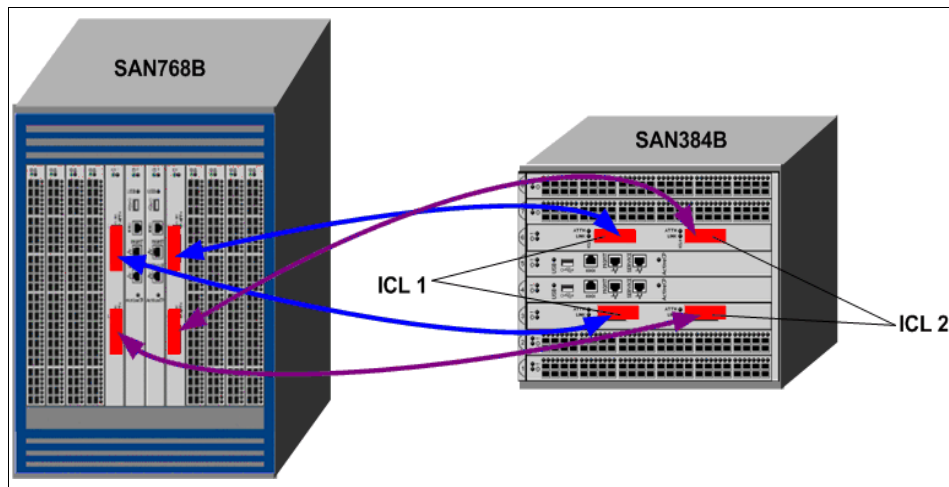


Figure 2-7 ICL cable connections between 768B and 384B

As of FOS v6.3 and higher, the options to connect SAN384B and SAN768B switches have increased because there is now the possibility of connecting a three way ISL. All possible options can be found in the *IBM System Storage SAN768B Installation, Service, and User Guide*, GA32-0574-04.

Figure 2-8 shows an example of a three way ISL between three SAN768B switches.




Figure 2-8 Three way ISL

2.5 Blade options

For both switches in the family, we have the following blade options:

- ▶ Condor2 ASIC Port blades: 8 Gbps universal FC/FICON ports (E, F, FL, M, EX) up to a maximum of eight:
 - FC8-16: 16 x 8 Gbps FC ports
 - FC8-32: 32 x 8 Gbps FC ports
 - FC8-48: 48 x 8 Gbps
 - FC8-64: 64 x 8 Gbps
- ▶ The following other blades are supported:
 - FC10-6; up to eight
 - 6-port 10 Gbps blade for dark fiber connections
 - Encryption blade
 - 16 x 8 Gbps front-end user ports
 - 2 GbE ports for out-of-band cluster HA interconnect
 - Data cryptographic and data compression capabilities
 - 1:1 subscription at 16 ports
 - Auto-sensing Link Speeds at 1, 2, 4, and 8 Gbps
 - FR4-18i SAN Extension blade:
 - 16 x 4 Gbps FC ports, 2 GigE ports
 - FCIP and 4 Gbps FC Routing (alternative to integrated 8 Gbps FC Routing)
 - Up to 8 blades if FCIP only; up to 2 blades if FCIP and/or FCR

2.6 Advance Feature and Licensing

The Advanced Feature and License information is described in Table 2-1.

Table 2-1 Advanced Features for SAN76bB/ SAN384B

Feature	License
Integrated Routing	IR
Virtual Fabric	None (included in FOS 6.2+)
Top Talkers	Advanced Performance Monitoring
Ingress Rate Limiting	Adaptive Networking

Feature	License
Traffic Isolation Zones	None (included in FOS 6.0+)
Fabric QoS	Adaptive Networking
Native Interop with M-Series	None
Future 10GigE FCoE/CEE	Per future blade

2.6.1 Advanced Feature summary

Here we provide an Advanced Feature summary:

- ▶ Integrated Routing:
 - Up to 128 EX_ports
- ▶ Virtual Fabric (requires FOS v6.2)
- ▶ Logical Switch (Logical Fabric); maximum of eight per chassis:
 - Units that ship with FOS v6.2 are VF enabled by default
- ▶ Native interoperability with m-series switches (Interopmode 2, Interopmode 3):
 - SAN768B Model: FC and FICON
 - SAN384B Model: FC only (no FICON cascading with M-Series)
- ▶ 10 GigE FCoE/CEE future-built architecture:
 - Both models will utilize the same FCoE/CEE blade in the future

Licenses: Integrated Routing and ICL licenses are different for the SAN768B from the SAN384B. Otherwise, all other licenses are the same for both models.

2.6.2 Fabric Operating Systems

The following list shows which Fabric Operating Systems (FOS) are supported by the SAN768B and the SAN384B:

- ▶ SAN768B model requires FOS v6.0+
 - Virtual Fabric requires FOS v6.2+
- ▶ SAN384B model requires FOS v6.2+



Hardware features

In this chapter, we present the hardware features of the IBM System Storage and TotalStorage b-type family of products. We first discuss the generic features, then describe the features that are specific to the various models.

3.1 The hardware

The IBM System Storage and TotalStorage b-type family of products provide a range of entry and midrange switches and enterprise class directors.

3.1.1 Entry level, midrange, and director models

The entry level, midrange, and director models provide 1, 2, 4, and 8 Gbps port-to-port non-blocking throughput. With auto-sensing feature, they are capable of connecting to older 1 Gbps host servers, storage, and switches. Hub-based Fibre Channel Arbitrated Loop (FC-AL) solutions reduce performance as devices are added by sharing the bandwidth, but an IBM System Storage and TotalStorage SAN Fabric throughput continues to increase as additional ports are interconnected.

All of these models are fully interoperable with the previous IBM System Storage and TotalStorage SAN switches, and can be added to existing fabrics, enabling transition from existing Fibre Channel storage networks to the faster technology.

3.1.2 Switch and director model types

Table 3-1 lists the current and historic switch and director model types with speed and port capabilities, the current supported version of Fabric Operating System (Fabric OS), and the type of Application Specific Integrated Circuit (ASIC).

Table 3-1 Director and switch models

Switch type	# Ports	Port speed	Fabric OS version	ASIC type
2499-192	16 to 256	1, 2, 4 and 8 Gbps	6.4.1+	Condor2
2499-384	16 to 768	1, 2, 4, and 8 Gbps	6.4.1+	Condor2
2498-B80	80	1, 2, 4 and 8 Gbps	6.4.1+	GoldenEye2
2498-B80	48, 64, and 80	1, 2, 4, and 8 Gbps	6.4.1+	GoldenEye2
2498-B40	24, 32, and 40	1, 2, 4, and 8 Gbps	6.4.1+	Condor2
2498-E32	32	1, 2, 4, and 8 Gbps	6.4.1+	Condor2
2498-B24	8, 16, or 24	1, 2, 4, and 8 Gbps	6.4.1+	GoldenEye2
2498-R06	16 6	1, 2, 4, and 8 Gbps 6 x 1 GbE ports	6.4.1+	GoldenEye2

Switch type	# Ports	Port speed	Fabric OS version	ASIC type
3758-L32	8 24	1, 2, 4 and 8 Gbps 10Gig Converged Enhanced Ethernet (CEE)	6.4.1+	Condor2 Anvil
2005-B5K	16, 24, or 32	1, 2, and 4 Gbps	6.4.1+	Condor
2005-B64	32, 48, or 64	1, 2, and 4 Gbps	6.4.1+	Condor
2005-B32	16, 24, or 32	1, 2, and 4 Gbps	6.4.1+	Condor
2005-B16	8, 12, or 16	1, 2, and 4 Gbps	6.4.1+	GoldenEye
2005-H16	16	1 and 2 Gbps	5.3.x	Bloom II
2005-H08	8	1 and 2 Gbps	5.3.x	Bloom II
2109-M48	16 to 384	1, 2, 4 and 8 Gbps	6.4.1+	Condor
2109-M14	32 to 128	1 and 2 Gbps	5.3.x	Bloom II

3.2 Generic features

In this section, we describe the standard features that are available on all of the b-type family.

3.2.1 Auto-sensing speed negotiation

The IBM System Storage and TotalStorage SAN uses internal Application Specific Integrated Circuits that support link operation at either 8, 4, 2, or 1 Gbps. As a device is connected to a port, the link speed is negotiated to the highest speed that is supported by the device. This speed selection is auto-negotiated by the ASIC driver on a per-port basis. If multiple devices are connected to a port (for example, on an FL_Port), the driver auto-negotiates for the highest common speed and sets the transmitter and receiver accordingly. This *auto-sensing negotiation* allows easy configuration.

3.2.2 Zoning

You can use *zoning* to arrange fabric-connected devices into logical groups (zones) dynamically across the physical topology of the fabric.

Zones can include selected storage and hosts within a fabric restricting information access to only the member devices in the defined zone. Although zone members can access only other members in their zones, individual devices can be members of more than one zone. This approach enables the secure sharing of storage resources, a primary benefit of storage networks.

We discuss zoning further in Chapter 12, “Basic zoning” on page 513.

3.2.3 Frame filtering

Frame filtering enables another level of storage area network (SAN) fabric monitoring and management. Zoning is a fabric management service that you can use to create logical subsets of devices within a SAN and enable partitioning of resources for management and access control purposes. Frame filtering enables the switch to provide zoning functions with finer granularity. Frame filtering can be used to set up port level zoning, worldwide name zoning, device level zoning, protocol level zoning, and LUN level zoning. After the filter is set up, the complicated function of zoning and filtering can be achieved at wire speed. Frame filtering is also used with performance monitoring, allowing you to monitor either “End to End” traffic flow or device-based I/O requirements.

3.2.4 Routing

The switch or director’s control processor maintains two routing tables, one for *unicast* and one for *multicast*. The unicast routing tables are constructed during fabric initialization. The multicast tables are initially empty, except for broadcast addresses. When the tables have been constructed, they are loaded into each ASIC.

The unicast tables change if ports or links come online or go offline, or if some other topology changes occur. These updates are triggered by a Registered State Change Notification (RSCN). When new paths become available, the control processor can change the routing tables in order to share the traffic load.

The multicast tables change as ports register with the alias server to create, join, or leave a multicast group. Each time a table changes, it must be reloaded into the ASICs.

3.2.5 Service functions

The ASIC interrupts the embedded processor when a frame arrives that has an error (for example, incorrect source ID), when a frame times out, or when a frame arrives for a destination that is not in its routing tables. When a frame arrives for a destination that is not in its routing tables, the frame might be addressed to an illegal destination ID, or it might be addressed to one of the *service functions* that are provided by the embedded processor such as SNMP, name server, or alias server.

3.2.6 Port Fencing

Port Fencing is a feature that is added to the switches starting with Fabric OS v6.1. Using this feature, ports that operate outside the bounds of normal operation can be disabled. Port Fencing requires a Fabric Watch License and can be configured based on the Fabric Watch event monitoring. When the port is disabled, user intervention is required for the port to be enabled again. Port Fencing has to be configured using the command-line interface (CLI) and is not supported in Web Tools.

3.2.7 ISL Trunking

The current IBM System Storage and TotalStorage SAN b-type switches have an optional feature called *ISL Trunking*. ISL Trunking is ideal for optimizing performance and simplifying the management of a multi-switch SAN fabric.

When two to four or eight adjacent ISLs in the same trunking group, depending on switch models, are used to connect two switches, the switches automatically group the ISLs into a single logical ISL, or *trunk*. The throughput of the resulting trunk is the sum of the throughputs of the participating links.

ISL Trunking is designed to significantly reduce traffic congestion. As shown in Figure 3-1, four 4 Gbps ISLs are combined into a single logical ISL with a total bandwidth of 16 Gbps. The trunk can support any number of connections, although we only show five connections in our example. Be aware that prior to implementing the trunking, the four parallel ISLs result in a throughput of 10 Gb due to the fact that two of the connections are sharing the same ISL. Following the implementation of trunking, this throughput increases to 14 Gb, that is, full throughput.

To balance the load across all of the ISLs in the trunk, each incoming frame is sent across the first available physical ISL in the trunk. As a result, transient workload peaks for one system or application are much less likely to impact the performance of other devices of the SAN fabric.




Figure 3-1 SAN b-type ISL Trunking

Because the full bandwidth of each physical link is available with ISL Trunking, no bandwidth is wasted by inefficient load sharing. As a result, the entire fabric is used more efficiently. Fabric OS and management software, such as Fabric Watch, also view the group of physical ISLs as a single logical ISL. A failure of a single ISL in a trunk causes only a reduction of the available bandwidth and not a failure of the complete route. Therefore, no re-calculation of the routes at that time is required. Bandwidth is restored automatically when the ISL is repaired.

Trunk master: If an older 2 Gbps switch is involved in either end of a trunk, one of the links forming the trunk is chosen as the trunk master. If that trunk master link fails, the trunk must select a new master, causing a slight disruption to traffic. Trunks between the new 4 Gbps switches do not have this restriction.

ISL Trunking helps to simplify fabric design, lower provisioning time, enhance switch-to-switch performance, simplify management, and improve the reliability of the SAN fabrics. In-order delivery is still guaranteed by the switch ASICs.

Table 3-2 lists the maximum number of ISLs that a single trunk supports, as well as the maximum trunk speed for different IBM System Storage and TotalStorage b-type switch models. If you have to form an ISL trunk between two different switch models, the lower of the maximum values for both number of ports supported and port speed apply.

Table 3-2 Maximum trunk capacity

Device type	Ports/trunk	Port speed	Trunk speed
SAN24B-4	8	8 Gbps	64 Gbps
SAN40B-4	8	8 Gbps	64 Gbps
SAN80B-4	8	8 Gbps	64 Gbps
B32	8	8 Gbps	64 Gbps
SAN256B	8	8 Gbps	64 Gbps
SAN384B	8	8 Gbps	64 Gbps
SAN768B	8	8 Gbps	64 Gbps

3.2.8 Diagnostics

The switch supports a set of power-on self tests (POSTs), as well as tests that can be invoked using a CLI. These diagnostics are used during the manufacturing process as well as for fault isolation of the product in customer installations. The POST and diagnostic commands concentrate on the Fibre Channel ports and verify the functionality of the switch. Post diagnostics are written to run in the Fabric OS environment. However, as the Fabric OS does not run without a working SDRAM, a SDRAM/boot EEPROM test is run as part of the pre-Fabric OS startup code to verify that the basic processor connected memories are functioning properly.

Loop-back paths for frame traffic are provided in the hardware for diagnostic purposes. A loop-back path within the ASIC, at the final stages of the Fibre Channel interface, can be used to verify that the internal Fibre Channel port logic is functioning properly, as well as paths between the interface and the central memory.

Additionally, the Serial Link macro within the ASIC includes a serial data loop-back function that can be enabled through a register in the corresponding ASIC.

Diagnostics are provided to allow traffic to be circulated between two switch ports that are connected with an external cable. This allows the diagnostics to verify the integrity of the final stage of the SERDES interface, as well as the SFP module.

3.3 Products and features

In this section, we provide an overview of the products and features in the IBM/Brocade portfolio. For the latest information, see this website:

<http://www-03.ibm.com/systems/storage/san/index.html>

Both the Condor2 and GoldenEye2 ASICs now support 8 Gbps port throughput capability throughout the current product range from the 8-port B24 switch to the SAN768B Fabric Backbone. Additional functionality of these ASICs provides larger trunking capabilities and integrated SERDES.

We discuss the support for these new features in the following sections.

3.3.1 2499-384

The 2499-384, also known as the IBM System Storage SAN768B Fabric Backbone, is designed for larger mid-range to enterprise-level SAN applications and is a core switching platform used to interconnect storage devices, hosts, and servers in SANs. The SAN768B is designed to meet the growing connectivity, virtualization, and cost-efficiency needs of enterprise data centers. As the core of the Data Center Fabric (DCF) architecture, the SAN768B Fabric Backbone is highly robust and can support both open systems and mainframe environments. With breakthrough performance, scalability, and energy efficiency, the SAN768B is designed to meet a wide range of technology challenges for evolving enterprise data centers and provides long-term investment protection.

To support growing server and storage environments and to facilitate virtualization, the SAN768B delivers over four times the bandwidth of industry-leading SAN directors and perform at speeds of 1, 2, 4, 8, or 10 Gbps. Using 4 Gbps or 8 Gbps Fibre Channel SFPs, a single chassis provides up to 384 Fibre Channel ports. Four Inter-Chassis Link (ICL) ports are used to connect two or three SAN768B chassis, and preserve the equivalent of up to 128 8 Gbps E_Ports for server and storage connectivity.

As a member of the IBM System Storage and TotalStorage SAN b-type family products, the SAN768B is designed to participate in fabrics containing other b-type and m-type devices manufactured by Brocade. This versatile hardware can serve as a new top tier (or backbone) in a complex fabric and provide connections to other b-type and m-type directors, switches and routers.

Figure 3-2 shows the front view of the SAN 768B Fabric Backbone.




Figure 3-2 SAN768B Fabric Backbone

With the chassis size identical to the SAN256B, the SAN768B houses four 2000-watt power supplies and three 220 mm blower fans. Both the blowers and the power supplies plug directly into the backplane and are both individual FRUs. The power supplies are auto-sensing from 110 to 240 V single phase and frequency range 47 to 63 Hz. With a DC power consumption of 1515 watts when configured as a fully loaded system, the SAN768B is extremely energy efficient.

Power consumption: The 1515 watts power consumption is for a fully loaded system configured with two CP8 blades, two CR8 blades, eight FC8-48 blades with 384 SWL SFPs, and three blowers.

Table 3-3 shows the power consumption of each blade.

Table 3-3 Blade power consumption

Blade	Power Consumption (Watts)
CP8(CP Blade)	38
CR8(Core Blade)	95
FC8-16(16 Port Blade)	50
FC8-32(32 Port Blade)	80
FC8-48(48 Port Blade)	115
FC8-64(64 Port Blade)	126
FR4-18i(FCIP Blade)	140
Encryption blade	235
FA4-18(Application Blade)	168
FC10-6(10 Gbps Blade)	120

The SAN768B also has two world wide name (WWN) cards per chassis that are located between the power supplies and are covered with a plate. The WWN cards have one EEPROM on each card to store the FRU S/N, runtime hours, OEM specific information and event/error logs on each. The data stored on the WWN cards is CRC checked when the data is written.

Figure 3-4 shows the WWN card status LEDs.




Figure 3-3 WWN card status LEDs

Inter-Chassis Link

The Inter-Chassis Link (ICL) allows three SAN768Bs or SAN384Bs to be connected together without sacrificing user ports (for more details, see 2.4, “Scalability at the core” on page 30). Each ICL is the equivalent of 16 x 8 Gbps of bandwidth. A total of four ICL connections can be made between two SAN768Bs, providing a total bandwidth of 64 x 8 Gbps and creating a dual core fabric with 768 user ports available.

Figure 3-4 shows an example of two SAN768Bs with the ICL connected and three SAN768Bs connected with ICL. An ICL kit that includes two ICL licenses, and four ICL cables, are required to establish an ICL.

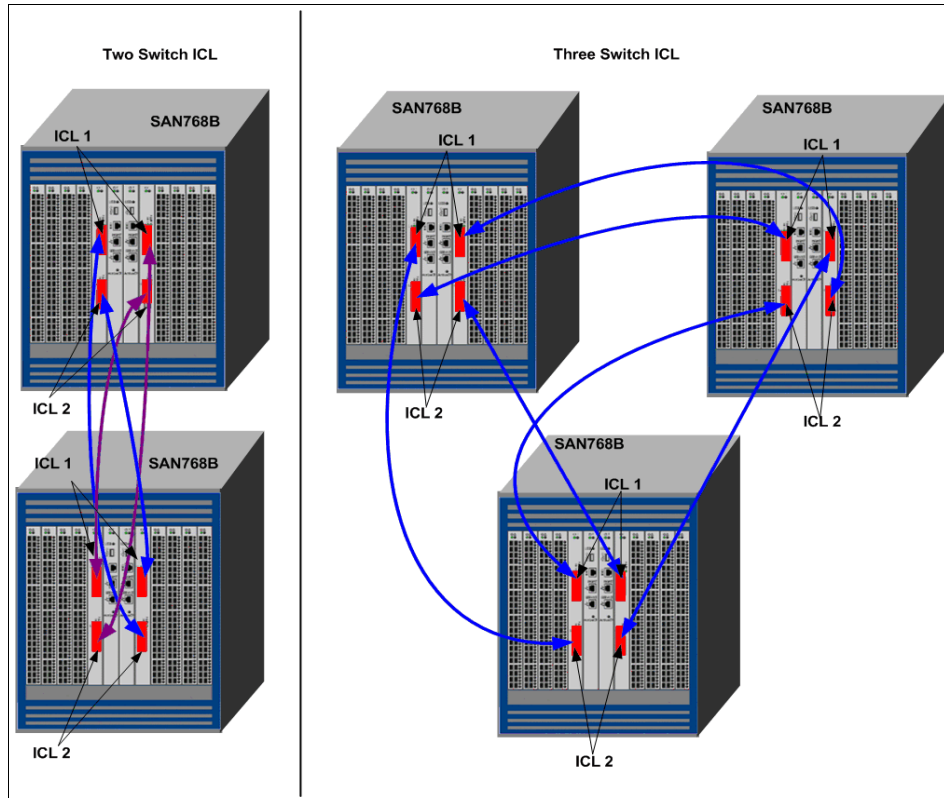


Figure 3-4 SAN768B ICL connection

ICLs: Use of ICLs does not collapse two switch domains into a single domain.

Figure 3-5 shows the ICL cable connector.




Figure 3-5 SAN768B: ICL Cable Connector

Key features of the SAN768B include these:

- ▶ Operating on the Condor2 ASIC, the SAN768B supports full-line-rate performance with 8 Gbps ports.
- ▶ The high density blade architecture provides up to 512 8 Gbps Fibre Channel ports per single chassis or up to 1024 8 Gbps Fibre Channel ports in a dual chassis configuration and up to 1536 8 Gbps ports on a three chassis configuration.
- ▶ The SAN768B is designed to provide high performance with 4.6 Tbps per chassis (512 ports × 8 Gbps data rate + 512 Gbps ICL bandwidth).
- ▶ Multiprotocol capabilities and fabric interoperability features include support for Fibre Channel, FICON, FCIP, and IPFC. It also supports 10 Gbps Ethernet, Converged Enhanced Ethernet (CEE), Data at Rest Encryption, and Fibre Channel over Ethernet (FCoE) for future enhancements.
- ▶ The SAN768B supports fabric services and applications such as Adaptive Networking services, including Quality of Service (QoS), Ingress Rate Limiting, Traffic Isolation, N_Port ID Virtualization (NPIV), and Top Talkers.
- ▶ Special application blades designed for fabric based storage virtualization, continuous data protection, and replication and online data migration are supported on the SAN768B.
- ▶ High availability design of the SAN768B includes features such as the passive backplane, separate and redundant control processor and core switching blades, and redundant WWN cards. The high availability design also features four hot-pluggable redundant power supplies and three blower fans.
- ▶ Energy efficient design of the SAN768B makes it possible to use less than one-half watt per Gbps.

Blades on the SAN768B

The SAN786B architecture utilizes a wide variety of blades for increasing port density. The SAN768B is a 12-slot chassis consisting of 8-port blades, 2 control processor (CP8) blades, and 2 core (CR8) blades, and can provide 512 Gbps bi-directional bandwidth per slot.

The core frame routing functionality is handled by the CR8 blade, and the entire unit is capable of handling up to 16 000 hard zones. Port blades are available in 16, 32, and 48 port configurations and operate on the Condor2 ASIC. The 4 Gbps and 8 Gbps SFPs used on the blades that operate on the Condor2 ASIC must be Brocade branded.

Figure 3-6 shows the port side layout of SAN768B and slot numbers where the blades can be inserted. Slots 1 through 4 show 48 port blades; slots 5 and 8 have the core blades 0 and 1, respectively; slots 6 and 7 have the control processor blades 0 and 1, respectively; slots 9 and 10 show 32 port blades, and slots 11 and 12 show 16 port blades.




Figure 3-6 SAN768B: Port Side View

Control processor blade overview

The SAN768B has two control processor (CP8) blades in slots 6 and 7. Figure 3-7 shows the blade.

Each CP8 blade has a USB port, an RS-232 Console Port, two IP network ports, and dual processors. One of the processors and the service IP port are for future use. The USB port only supports Brocade branded USB drives and can be used for firmware download, supportsave, configuration upload, and configuration download.




Figure 3-7 CP8 blade

Table 3-4 describes the various LEDs located on the CP8 blade and the messages that can be interpreted from the status of these LEDs.

Table 3-4 LEDs on CP8 blade

LED	Color	Location	Description
Power	Green	Front Panel	On = CP power is operational. Off = CP power is failed.
Attention	Amber	Front Panel	On = If on for more than 5 seconds, board is defective or faulty.
Active	Blue	Front Panel	On = This CP is the active CP. Off = This CP is either booting, negotiating to be active or the standby CP.
Ethernet Link	Green	Front Panel RJ45 Top	On = Ethernet Port MAC link has been established at 100/1000 Mbps. Off = No Link or 10 Mbps.
Ethernet Activity	Green	Front Panel RJ45 Bottom	On Blinking = TX and RX Frames activity present. Off = No TX and RX activity.
USB Port	Green	External	On = USB port is enabled. Off = USB port is disabled.

Core blade

The CR8 blade provides the core routing of frames either from blade to blade or from SAN768B to SAN768B (or SAN384B) through an ICL cable. Each CR8 blade has four Condor2 ASICs and two ICL ports.

Figure 3-8 shows the CR8 blade with two ICL ports.




Figure 3-8 CR8 blade

The CR8 blade has two LEDs on the top, and their status can be interpreted as shown in Table 3-5.

Table 3-5 LEDs on CR8 Blade

LED	Color	Description
Power	Green	On = CP power is operational. Off = CP power has failed.
Attention	Amber	On = If on for more than 5 seconds, board is defective or faulty. Off = Normal operation.

The two ICL connectors have two LEDs each, and their status can be interpreted as shown in Table 3-6.

Table 3-6 LEDs on ICL connectors

Green	Amber	Condition
Off	Off	Cable is not present OR local end is not ready OR far end is not ready.
Off	On	N/A
On	Off	Cable is present AND local end is ready AND far end is ready.
On	On (Blinking)	Cable is present AND local end is ready AND far end is ready AND attention is required.

Condor2 ASIC port blades

The Condor2 ASIC is the next generation ASIC used on the port and the core blades, and provides all the functions that the Condor ASIC does but has more ports, higher speeds, and more buffer space. The SAN768B supports a variety of Condor2 ASIC port blades for increased port density.

FCOE10-24

This blade provides CEE/FCoE connectivity for server I/O consolidation (24 x 10 GbE CEE ports; up to two blades per chassis), shown in Figure 3-9.




Figure 3-9 FCOE10-24

Encryption Engine

The Encryption Blade provides plug-in encryption of data on disk or tape, supporting industry-standard AES-256 and DataFort-compatible encryption mode (16 8 Gbps Fibre Channel ports; up to four blades per chassis and requires DCFM management), shown in Figure 3-10.




Figure 3-10 Encryption Engine

FX8-24

This blade provides SAN extension over IP networks (12 8 Gbps Fibre Channel ports with license options providing up to 10 1 GbE ports, and up to two 10 GbE ports per blade; up to four blades), shown in Figure 3-11.




Figure 3-11 FX8-24 blade

FC 4 Gbps Routing blade

The Application Blade provides 16x 4 Gbps Fibre Channel ports and two Gigabit Ethernet ports per blade (up to four blades), shown in Figure 3-12.




Figure 3-12 FS4-16i

FC10-6

This blade provides six 10 Gbps Fibre Channel ports (up to eight blades)

References: These blades are described in detail in the following books:

- ▶ *IBM Converged Switch B32*, SG24-7935-00, available at this website:
<http://www.redbooks.ibm.com/redpieces/abstracts/sg247935.html?open>
- ▶ *Implementing the IBM System Storage SAN32B-E4 Encryption Switch*, SG24-7922, available at this website:
<http://www.redbooks.ibm.com/abstracts/sg247922.html?open>
- ▶ *IBM System Storage b-type Multiprotocol Routing: An Introduction and Implementation*, SG24-7544-03, available at this website:
<http://www.redbooks.ibm.com/abstracts/sg247544.html?open>

FC8-16

The FC8-16 is a 16 port blade that can operate on one Condor2 ASIC. Operating at speeds of 1, 2, 4, and 8 Gbps, this blade supports F/FL/E ports and provides a 1:1 subscription at all speeds.

Figure 3-13 shows the FC8-16 blade.




Figure 3-13 FC8-16 blade

Figure 3-14 shows the SAN768B architecture for an FC8-16 blade.

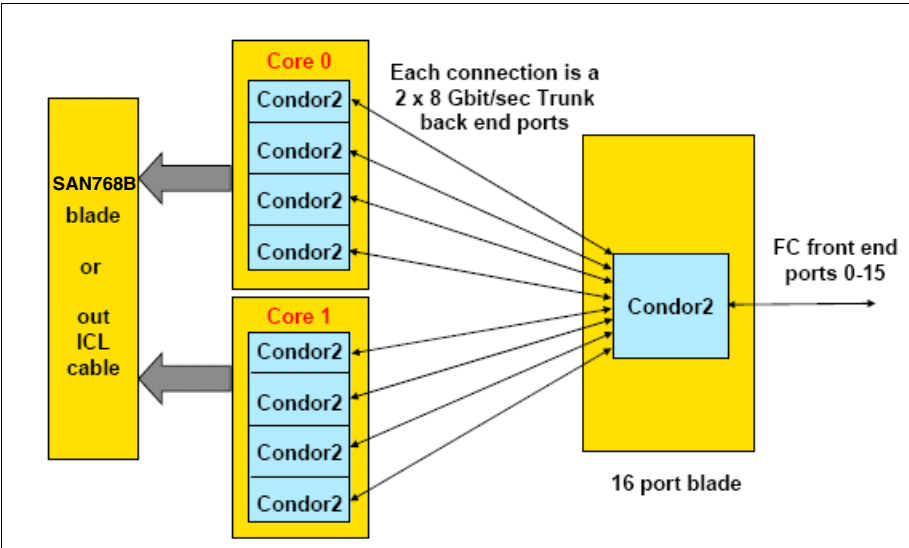


Figure 3-14 SAN768B architecture for a FC8-16 blade

Figure 3-15 shows the port area numbers when all the slots in the SAN768B are configured with FC8-16 port blades.




Figure 3-15 Port area numbers: FC8-16 blade

FC8-32

The FC8-32 is a 32-port blade that can operate on two Condor2 ASICs. Operating at speeds of 1, 2, 4, and 8 Gbps, this blade supports F/FL/E ports and provides a 1:1 subscription at all speeds. The FC8-32 is shown in Figure 3-16.




Figure 3-16 FC8-32 blade

The SAN768B architecture for an FC8-32 blade is shown in Figure 3-17.

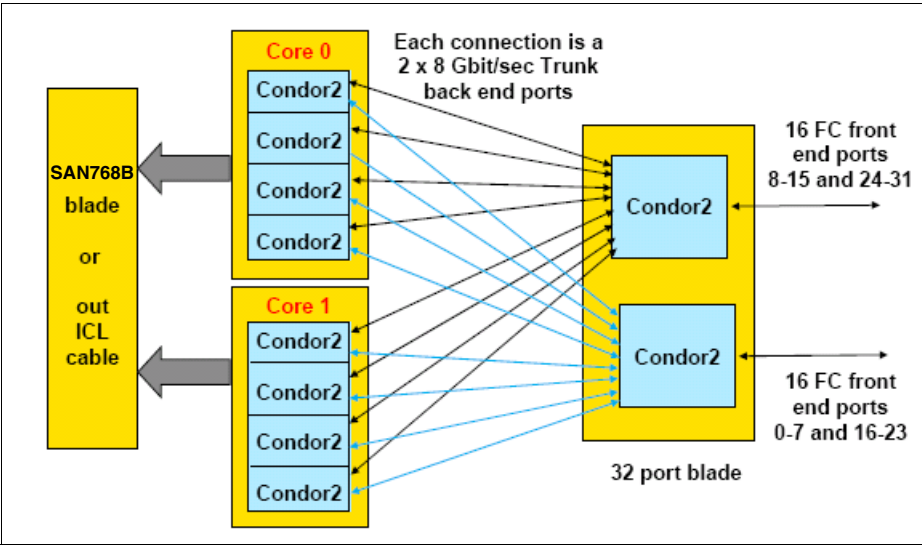


Figure 3-17 SAN768B architecture for a FC8-32 blade

Figure 3-18 shows the port area numbers for the FC8-32 blade. Ports 0-7 and 16-23 are on one ASIC, and ports 8-15 and 24-31 are on the other ASIC.

15	143	31	159	47	175	63	191	79	207	95	223	111	239	127	255
14	142	30	158	46	174	62	190	78	206	94	222	110	238	126	254
13	141	29	157	45	173	61	189	77	205	93	221	109	237	125	253
12	140	28	156	44	172	60	188	76	204	92	220	108	236	124	252
11	139	27	155	43	171	59	187	75	203	91	219	107	235	123	251
10	138	26	154	42	170	58	186	74	202	90	218	106	234	122	250
9	137	25	153	41	169	57	185	73	201	89	217	105	233	121	249
8	136	24	152	40	168	56	184	72	200	88	216	104	232	120	248
7	135	23	151	39	167	55	183	71	199	87	215	103	231	119	247
6	134	22	150	38	166	54	182	70	198	86	214	102	230	118	246
5	133	21	149	37	165	53	181	69	197	85	213	101	229	117	245
4	132	20	148	36	164	52	180	68	196	84	212	100	228	116	244
3	131	19	147	35	163	51	179	67	195	83	211	99	227	115	243
2	130	18	146	34	162	50	178	66	194	82	210	98	226	114	242
1	129	17	145	33	161	49	177	65	193	81	209	97	225	113	241
0	128	16	144	32	160	48	176	64	192	80	208	96	224	112	240
1 2 3 4 Card Slots 9 10 11 12															

Figure 3-18 Port area number: FC8-32 blade

FC8-48

The FC8-48 is a 48 port blade that can operate on two Condor2 ASICs. Operating at speeds of 1, 2, 4, and 8 Gbps, this blade supports F/E ports and provides a 24:16 subscription at 8 Gbps speeds. The FC8-48 is shown in Figure 3-19.




Figure 3-19 FC8-48 blade

Figure 3-20 shows the SAN768B architecture for an FC8-48 blade.

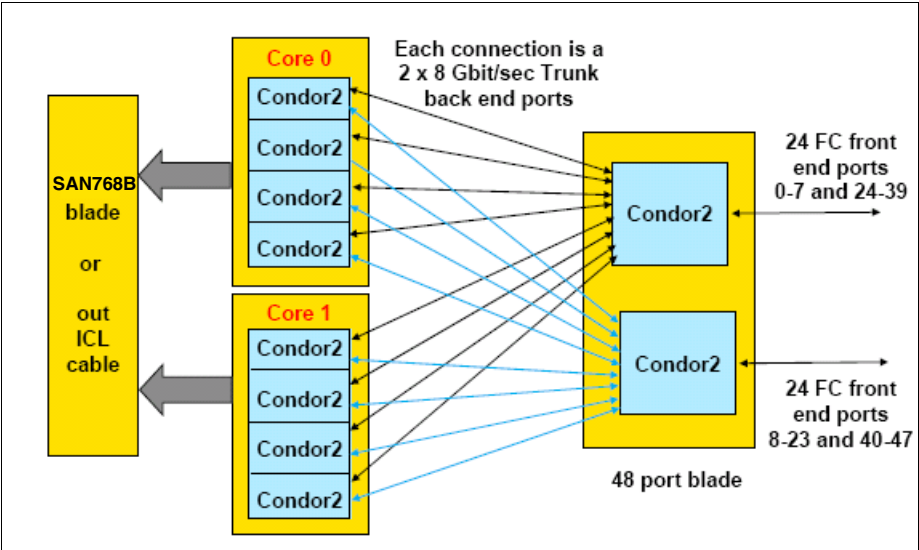


Figure 3-20 SAN768B architecture for a FC8-48 blade

The port area numbers for the FC8-48 blade are shown in Figure 3-21. Each ASIC supports 24 ports. Ports 0-7 and 24-39 are on one ASIC, and ports 8-23 and 40-47 are on the other ASIC.

135	271	161	287	167	303	183	319	199	335	215	351	231	367	247	383
134	270	160	286	166	302	182	318	198	334	214	350	230	366	246	382
133	269	149	285	165	301	181	317	197	333	213	349	229	365	245	381
132	268	148	284	164	300	180	316	196	332	212	348	228	364	244	380
131	267	147	283	163	299	179	315	195	331	211	347	227	363	243	379
130	266	146	282	162	298	178	314	194	330	210	346	226	362	242	378
129	265	145	281	161	297	177	313	193	329	209	345	225	361	241	377
128	264	144	280	160	296	176	312	192	328	208	344	224	360	240	376
16	263	31	279	47	295	63	311	79	327	95	343	111	359	127	375
14	262	30	278	46	294	62	310	78	326	94	342	110	358	126	374
13	261	29	277	45	293	61	309	77	325	93	341	109	357	125	373
12	260	28	276	44	292	60	308	76	324	92	340	108	356	124	372
11	259	27	275	43	291	59	307	75	323	91	339	107	355	123	371
10	258	26	274	42	290	58	306	74	322	90	338	106	354	122	370
9	257	25	273	41	289	57	305	73	321	89	337	105	353	121	369
8	256	24	272	40	288	56	304	72	320	88	336	104	352	120	368
7	143	23	159	39	175	55	191	71	207	87	223	103	239	119	255
6	142	22	158	38	174	54	190	70	206	86	222	102	238	118	254
5	141	21	157	37	173	53	189	69	205	85	221	101	237	117	253
4	140	20	156	36	172	52	188	68	204	84	220	100	236	116	252
3	139	19	155	35	171	51	187	67	203	83	219	99	235	115	251
2	138	18	154	34	170	50	186	66	202	82	218	98	234	114	250
1	137	17	153	33	169	49	185	65	201	81	217	97	233	113	249
0	136	16	152	32	168	48	184	64	200	80	216	96	232	112	248
1 2 3 4 Card Slots 9 10 11 12															

Figure 3-21 Port area number: FC8-48 blade

FC8-64

The FC8-64 is a 64 port blade that can operate on four Condor2 ASICs. Operating at speeds of 2, 4, and 8 Gbps, this blade supports F/E ports and provides a 2:1 oversubscription at 8 Gbps speeds and 1:1 at 4 Gbps. FL ports are not supported on this blade. The FC8-48 is shown in Figure 3-22.

Attention: The 8-Gb 64-port blade (FC3864) cannot be installed in the same chassis as a FCOE10-24 blade (FC3880).




Figure 3-22

Figure 3-23 on page 67 shows the SAN768B architecture for an FC8-64 blade.

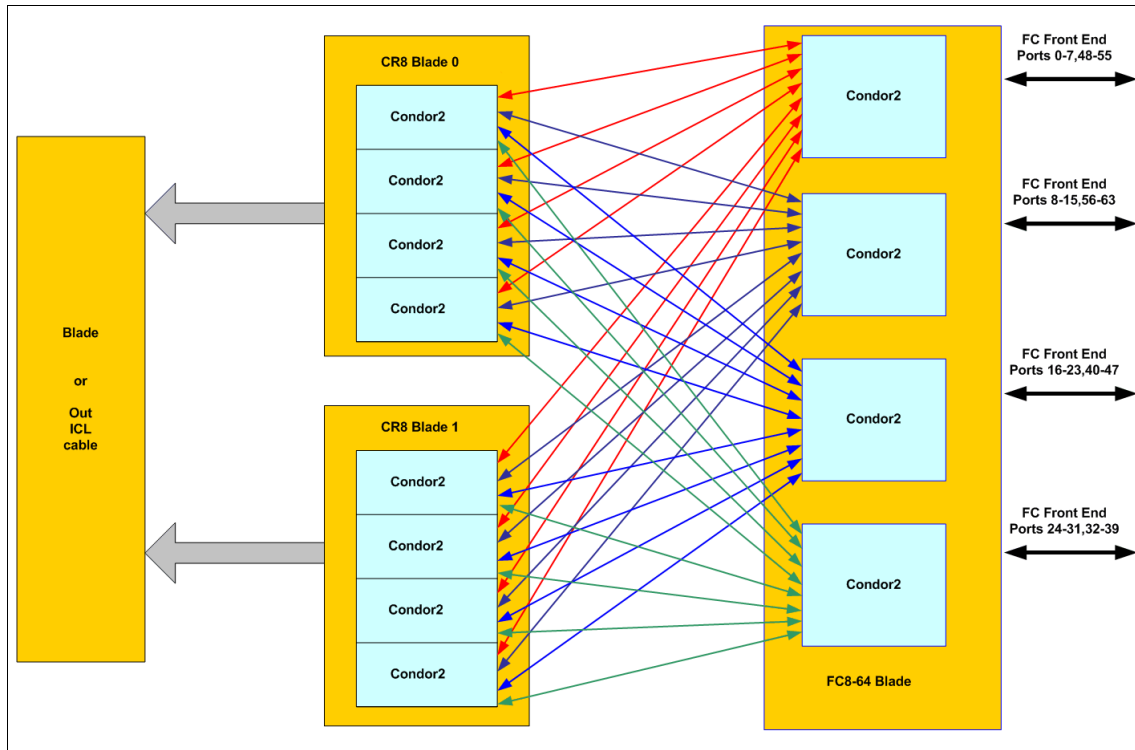


Figure 3-23 FC8-64 architecture

As shown there are four Condor2 ASICs on the FC8-64 blade. The port layout and port groups that share an ASIC are shown in Figure 3-24 on page 68.

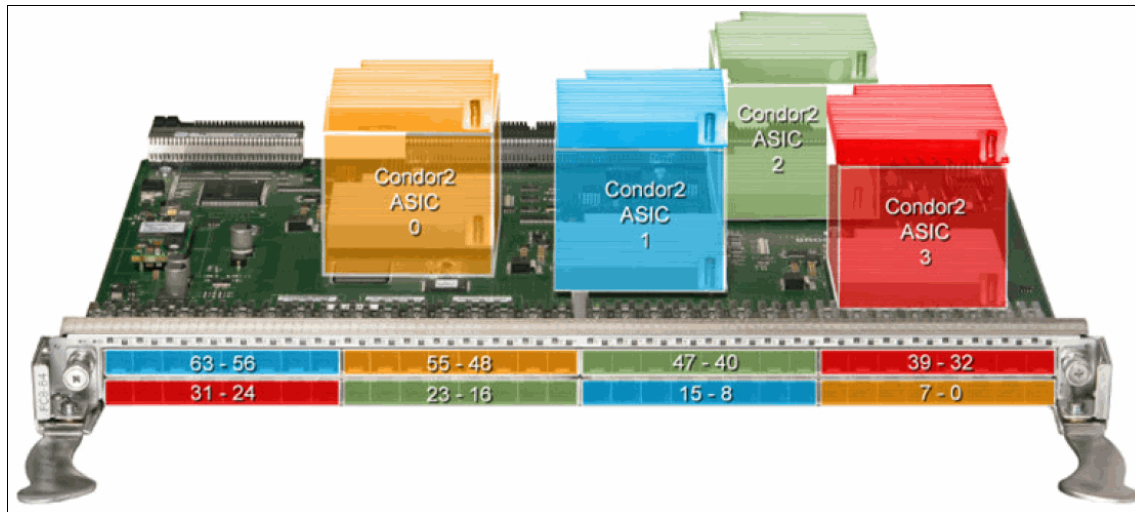


Figure 3-24 Port assignments for FC8-64 blade.

The ports on the FC8-64 are assigned port numbers as shown in Figure 3-25 on page 69.

143	783	159	799	175	815	191	831	207	847	223	863	239	879	255	895
142	782	158	798	174	814	190	830	206	846	222	862	238	878	254	894
141	781	157	797	173	813	189	829	205	845	221	861	237	877	253	893
140	780	156	796	172	812	188	829	204	844	220	860	236	876	252	892
139	779	155	795	171	811	187	827	203	843	219	859	235	875	251	891
138	778	154	794	170	810	186	826	202	842	218	858	234	874	250	890
137	777	153	793	169	809	185	825	201	841	217	857	233	873	249	889
136	776	152	792	168	808	184	824	200	840	216	856	232	872	248	888
135	775	151	791	167	807	183	823	199	839	215	855	231	871	247	887
134	774	150	790	166	806	182	822	198	838	214	854	230	870	246	886
133	773	149	789	165	805	181	821	197	837	213	853	229	869	245	885
132	772	148	788	164	804	180	820	196	836	212	852	228	868	244	884
131	771	147	787	163	803	179	819	195	835	211	851	227	867	243	883
130	770	146	786	162	802	178	818	194	834	210	850	226	866	242	882
129	769	145	785	161	801	177	817	193	833	209	849	225	865	241	881
128	768	144	784	160	800	176	816	192	832	208	848	224	864	240	880
15	271	31	287	47	303	63	319	79	335	95	351	111	367	127	383
14	270	30	286	46	302	62	318	78	334	94	350	110	366	126	382
13	269	29	285	45	301	61	317	77	333	93	349	109	365	125	381
12	268	28	284	44	300	60	316	76	332	92	348	108	364	124	380
11	267	27	283	43	299	59	315	75	331	91	347	107	363	123	379
10	266	26	282	42	298	58	314	74	330	90	346	106	362	122	378
9	265	25	281	41	297	57	313	73	329	89	345	105	361	121	377
8	264	24	280	40	296	56	312	72	328	88	344	104	360	120	376
7	263	23	279	39	295	55	311	71	327	87	343	103	359	119	375
6	262	22	278	38	294	54	310	70	326	86	342	102	358	118	374
5	261	21	277	37	293	53	309	69	325	85	341	101	357	117	373
4	260	20	276	36	292	52	308	68	324	84	340	100	356	116	372
3	259	19	275	35	291	51	307	67	323	83	339	99	355	115	371
2	258	18	274	34	290	50	306	66	322	82	338	98	354	114	370
1	257	17	273	33	289	49	305	65	321	81	337	97	353	113	379
0	256	16	272	32	288	48	304	64	320	80	336	96	352	112	368
Slot 1		Slot 2		Slot 3		Slot 4		Slot 9		Slot 10		Slot 11		Slot 12	

Figure 3-25 FC8-64 port numbering layout

Firmware upgrade requirements for FC8-64

Disable the ports in SAN768B or SAN384B Logical Switches that use 10 bit addressing mode that have 8 bit areas in the range 0x70-0x8F before upgrading to FOS v6.4.+. Otherwise the firmware upgrade will fail with an error message.

This step is necessary even if there is no plan to use FC8-64 blades after performing a firmware upgrade to FOS v6.4.0. However, if areas 0x70-0x8F are not in use this step is not necessary. Use the **portAddress** CLI command to find out the areas in use within a Logical Switch.

If there is a downgrade to a FOS below 6.4 then the FC8-64 blade must be removed or the process will fail.

mSFP Transceivers

The FC8-64 requires a mini Small Form-Factor Pluggable (mSFP). This is a smaller format SFP which has the same internal technology as a standard SFP.

The mSFP has a reduced width and gap between optics to cater for the increased number of ports on a single blade.

- ▶ 2mm less in width compared to a regular SFP
- ▶ Tx and Rx spacing decreased from 6.25mm to 5.25mm
- ▶ Cables with new Tx/Rx spacing connectors

The mSFP has a Pull/Push tab for easy insertion and removal as shown in Figure 3-26.




Figure 3-26 mSFP and SFP

Attention: An mSFP will fit into a standard SFP slot, but this is not a supported configuration, and the port will be faulted and taken offline.

Restriction: Only SWL optics are available for mSFP, and V6.1+ OS is required.

Port area numbers

You can also use Figure 3-21 as a reference for the port area numbers, as follows:

- Ports 0 through 15 on all blades are mapped to area numbers 0 through 127.
- Ports 16 through 31 on all the blades are mapped to area numbers 128 through 255.
- Ports 32 through 47 on all the blades are mapped to area numbers 256 through 383.
- Ports 48 through 61 on the FC8-64 blades are mapped to area numbers 768 through 895.

Figure 3-27 explains the port area numbers that are assigned, which depends on the slot in which the blade is installed.

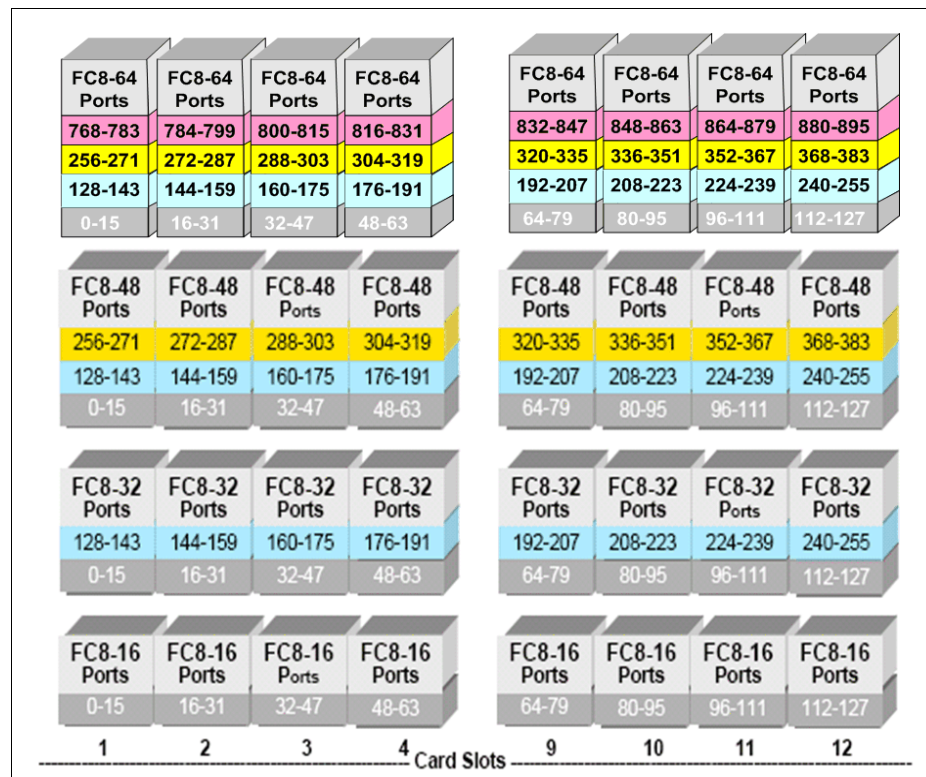


Figure 3-27 SAN768B port area numbers with mixed blades

We further explain the port area numbers for the SAN768B using the following blade configuration:

- ▶ Slot 1 - 32 port blade
- ▶ Slot 2 - 16 port blade
- ▶ Slot 3 - 64 port blade
- ▶ Slot 4 - empty
- ▶ Slot 9 - 48 port blade
- ▶ Slot 10 - 32 port blade
- ▶ Slot 11 - empty
- ▶ Slot 12 - 64 port blade

Figure 3-28 shows the blades installed in the SAN768B, and Figure 3-29 shows the actual port area numbers that are assigned for this setup.

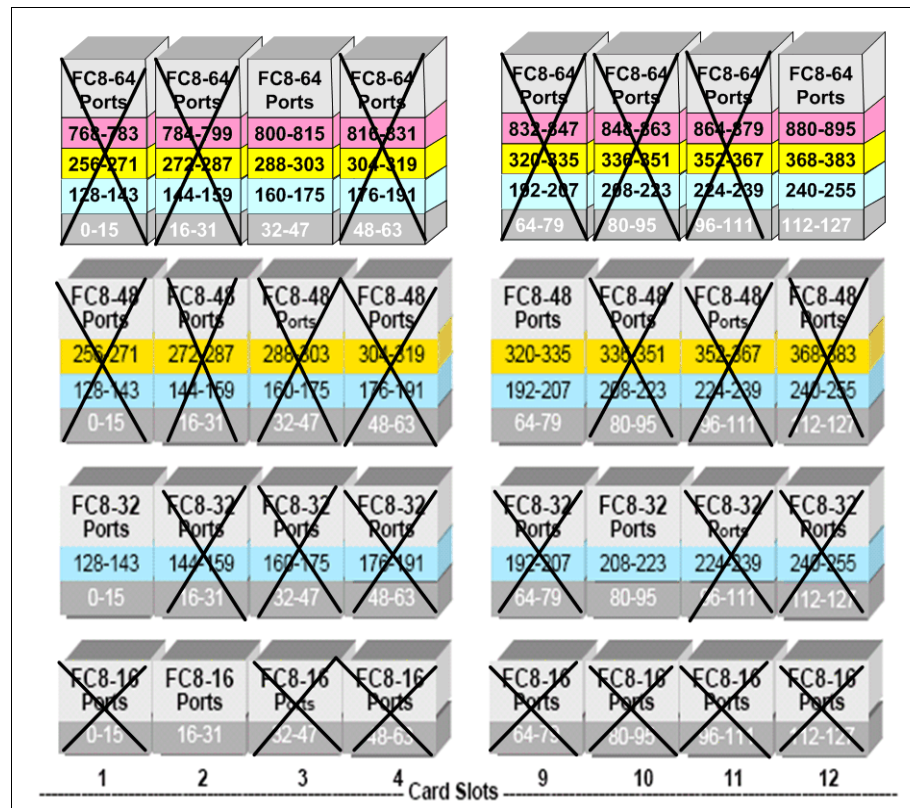


Figure 3-28 SAN768B with different blades installed




Figure 3-29 SAN768B: Port area numbers for different blades

3.3.2 2499-192

The 2499-192, also known as the IBM System Storage SAN384B, is a fabric backbone product line extension offering 192x 8 Gbps ports (half the port count offered by the 2499-384, also known as the IBM System Storage SAN768B fabric backbone) in a horizontal chassis.

The SAN384B offers flexible deployment and investment protection in both new and existing storage networks. It can be deployed as a lower cost core backbone solution in midsize enterprise network environments that do not require the throughput and port density of the larger SAN768B fabric backbone. Large enterprise customers can also implement the SAN384B at the network edge to provide complete, scalable, and cost effective backbone-class capabilities throughout their data centers. The SAN384B can also connect natively to IBM b-type and m-type network environments without disruption.

Figure 3-30 shows the front view of the SAN384B.



Figure 3-30 SAN384B front view

The mechanical chassis of the SAN384B is a horizontal chassis enclosure with 8 slots accessed from the front side and 2 blowers, 2 power supplies and dual WWN cards accessed from the non-port side.

Figure 3-31 shows the back view of the SAN384B.



Figure 3-31 SAN384B back view

The SAN384B is shipped with two 2000-watt power supplies and two 220 mm blower fans. Both the blowers and the power supplies plug directly into the backplane and are both individual FRUs. The power supplies are auto-sensing from 110 to 240 V single phase and frequency range 47 to 63 Hz, with a DC power consumption of 753 watts when configured as a fully loaded system.

For blade power consumption, see Table 3-3 on page 48.

Power consumption: The power consumption of the SAN384B is 753 watts when configured as a fully loaded system. This calculation is based on a system configuration with two CP8 control processors, two CR4S-8 core blades, four FC8-48 port blades with 192 SWL SFPs, and two blowers.

Key features of the SAN384B include these:

- ▶ 8 Gbps Fibre Channel interfaces supporting 1 Gbps, 2 Gbps, 4 Gbps, and 8 Gbps auto-sensing Fibre Channel ports, which self-configure as E, F, or FL ports.
- ▶ Four high-performance port blades (64 port, 48 port, 32 port, and 16 port).
- ▶ Dual redundant control processors and core switching blades designed to provide enhanced high availability and enable non-disruptive software upgrades.
- ▶ Up to 768 x 8 Gbps user ports in connected in a three core or 512 x 8 Gbps user ports connected in a dual core with the use of Inter Chassis Links (ICLs), and up to 256 ports in a single domain.
- ▶ The multiprotocol design of the SAN384B supports blades for Fibre Channel over IP (FCIP) and routing that are the cornerstone of highly intelligent fabrics.
- ▶ Inter-Switch Link (ISL) trunking allows up to eight ports between a pair of switches to be combined to form a single, logical ISL with an aggregate speed of up to 64 Gbps for optimal bandwidth utilization.
- ▶ Dynamic Path Selection can evenly balance up to eight equal cost paths, including trunks and ISLs, at speeds from 1 Gbps to 8 Gbps. With this function, two SAN384B chassis can have up to 256 Gbps of hardware load balanced paths.
- ▶ Frame filtering resources have been increased to provide 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).

Important: Only 2 GB Brocade branded USB drives are supported for use on the USB port.

Blades on the SAN384B

The SAN384B supports all the blades supported by the SAN768B: the FC8-16, FC8-32, FC8-48, and FC8-64 8 Gb port blades.

For more information, see “Blades on the SAN768B” on page 51.

Important:

- ▶ The SAN384B supports all features and functions as indicated and requires Fabric OS v6.4.1+.
- ▶ Blades that use Condor2 ASIC *must* use Brocade branded SFPs.

Figure 3-32 shows the CP, CR, and port blades.

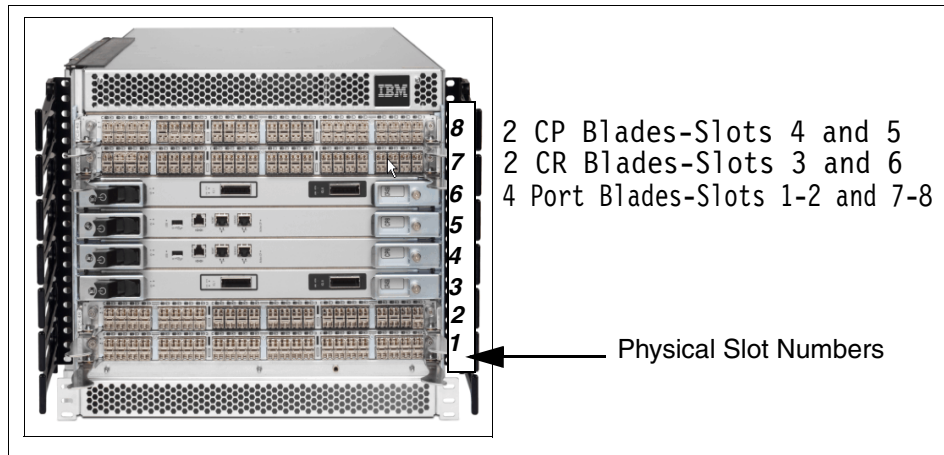


Figure 3-32 Blades

Port assignment

Port numbering for the FC blades is as follows:

- ▶ FC8-16 port blade: Ports are numbered from 0 through 15 from right to left.
- ▶ FC8-32 port blade: Ports are numbered from 0 through 15 from right to left on the lower row of ports and 16 through 31 from right to left on the upper row of ports.
- ▶ FC8-48 port blade: Ports are numbered from 0 through 23 from right to left on the lower row of ports and 24 through 47 from right to left on the upper row of ports.
- ▶ FC8-64 port blade: Ports are numbered from 0 through 31 from right to left on the lower row of ports and 32 through 63 from right to left on the upper row of ports. Trunking groups are permitted with up to eight ports per group. Trunking groups are as follows: 0-7, 8-15, 16-23, 24-31, 32-39, 40-47, 48-55, and 56-63.

Figure 3-33 shows the port assignment on a SAN384B with 64 port layout. The FC port numbering is the same for all blades types in the different slots. For example, an FC8-16 blade in slot 7 will start with port number 123 and end on port number 143.

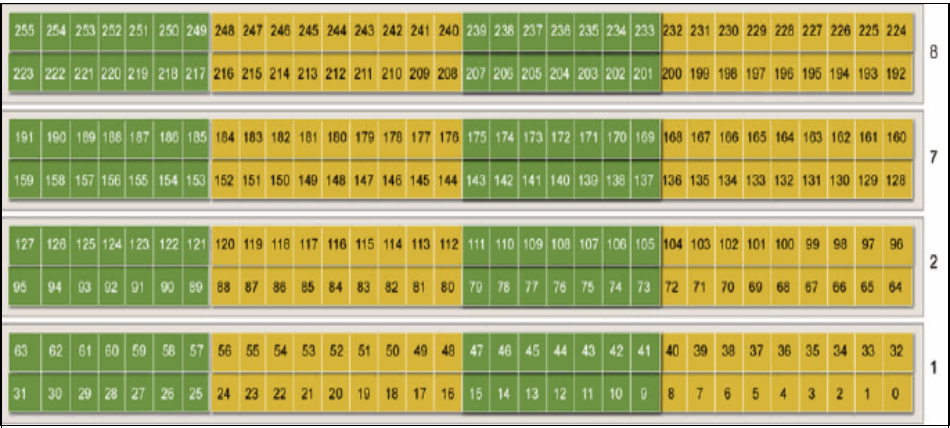


Figure 3-33 FC8-64 in a SABN384B

Inter-Chassis Link

The Inter-Chassis Link (ICL) allows up to three SAN384Bs (or SAN768Bs) to be connected together without sacrificing user ports. This is described in 2.4, “Scalability at the core” on page 30.

3.3.3 2109-M48

The IBM System Storage SAN256B (2109-M48) director is a single domain 384-port machine capable of running its ports at 1, 2, 4, 8, or 10 Gbps. The M48 includes support for FICON, FICON/Fibre Channel intermixing, FICON CUP, and FICON cascading, enabling it to address the demands for integrated System z® and open system server enterprise SANs. The chassis includes two control processor blades and, with improved port density, enables up to 384 ports in 14U space. Other standard software features include Web Tools, Zoning, Fabric Watch, Trunking, and Advanced Performance Monitoring. Optional software products include Extended Fabric Activation and FICON with CUP Activation.

The SAN256B is a 14U chassis with 10 slots for various blades. The CP4 control processor blades, which run on the Condor ASIC, are installed in slot 5 and slot 6 of the chassis, and the remainder are used for port blades.

The chassis has two power supplies installed as a standard and has room for four power supplies. Each of these power supplies can provide 1000 W of power. Two power supplies are required for normal operation. The SAN256B has three blower assemblies installed as a standard, and the blower side of the SAN256B also has the WWN card installed.

The WWN card retains important information about the chassis and switch identity data, chassis serial number, IP address assigned to each CP card slot, switch configuration, and FRU history logs.

Figure 3-34 shows the front view of the SAN256B SAN Director.



Figure 3-34 SAN256B SAN Director

Figure 3-35 shows the non-port side view.




Figure 3-35 SAN256B non-port side view

The SAN256B architecture utilizes a wide variety of blades for increasing port density. (We describe more about the FC4-16, FC4-32, and FC4-48 blades in the sections that follow.)

Important: Starting with Fabric OS v6.1, SAN256B can now support all the 8 Gbps Condor2 ASIC-based port blades. 8 Gbps speeds can be achieved only for local switching on the same 8 Gbps blade. We describe the FC8-16, FC8-32, and FC8-48 blades further in 3.3.1, “2499-384” on page 46.

The FC4-16 is a 16-port Fibre Channel blade that can support 1, 2, and 4 Gbps as shown in Figure 3-36.




Figure 3-36 FC4-16: 16 port blade

On the 16-port blade, all ports have 64 Gbps (128 Gbps full duplex) of possible external input, and the same internal bandwidth available. In other words, the blade has a 1:1 subscription ratio. It is useful for extremely high-performance servers, supercomputing environments, high-performance shared storage subsystems, and SANs with unpredictable traffic patterns.

The FC4-32 is a 32-port Fibre Channel blade that can support 1, 2, and 4 Gbps as shown in Figure 3-37.




Figure 3-37 FC4-32: 32 port blade

The 32-port blade is designed with a 16:8 subscription ratio at 4 Gbps for non-local traffic, and a 1:1 ratio at 2 Gbps for any traffic pattern. If some or all of the attached servers and storage devices run at 2 Gbps, or if I/O profiles are “bursty,” the 32-port blade typically provides the same performance as the 16-port blade.

The FC4-48 is a 48-port Fibre Channel blade that can support 1, 2, and 4 Gbps as shown in Figure 3-38.

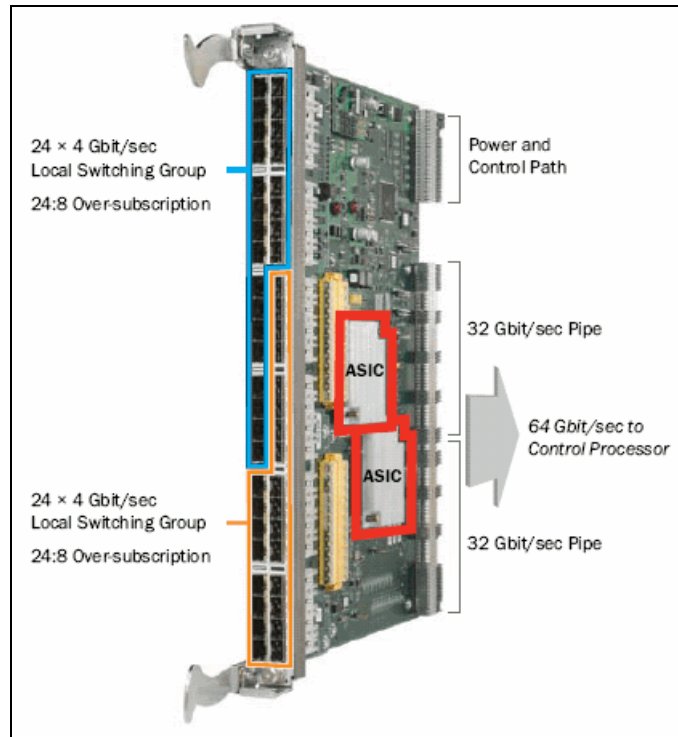


Figure 3-38 FC4-48: 48 port blade

At 24:8, the 48-port blade has a higher backplane over-subscription ratio but also has larger port groups. The backplane connectivity of this blade is identical to the 32-port blade. The only difference is that, rather than just 16 ports per ASIC, the 48-port blade exposes 24 outward-facing ports (96 Gbps or 192 Gbps full duplex of local switching per ASIC).

SAN256B applications

This blade is especially useful for high-density SAN deployments, in the following situations:

- ▶ Large numbers of servers have to be connected to the director.
- ▶ Some or all hosts are running below line rate much of the time.
- ▶ Potential localization of most traffic flows is achievable.

The SAN256B also supports the FC10-6 blade for 10 Gbps ISLs between two directors. This blade is designed to extend the value of the Fibre Channel infrastructure to include 10 Gbps FC and DWDM connected SANs. The blade provides the same high availability features that exist in the Brocade 48000 director today, satisfying the stringent requirements expected of an enterprise-class director. There are two local switch groups on the blade. Ports 0-2 and Ports 3-5 are locally switched with no data transferred over the backplane.

Support: ISL Trunking is not supported with the FC10-6 blade.

The control processor (CP4) cards are new by design, including faster processor units, and make use of two 32-port Condor ASICs as the switching core.

The 16, 32, and 48-port cards make use of cut-through routing, ensuring that frames destined for ports on the same card never leave the ASIC. This integrated feature called local switching provides significant performance benefits.

SAN256B numbering scheme

The SAN256B (2109-M48) uses a numbering scheme that progresses from left to right and bottom to top in numerical order. The reference location is from the cable side to chassis:

- ▶ Blade assemblies are numbered from 1 to 10, from left to right.
- ▶ Power supplies are numbered from 1 to 4, from bottom to top.
- ▶ Fans are numbered from 1 to 3, from left to right.
- ▶ The physical ports of the 16-port card are numbered 0 to 15, from bottom to top.
- ▶ The physical ports of the 32-port card are numbered 0 to 15 on the left column and 16 to 31 on the right column, from bottom to top.
- ▶ The physical ports of the 48 port blade are numbered 0 to 23 from bottom to top on the left set of ports and 24 to 47 from bottom to top on the right set of ports.

Figure 3-39 shows the logical decimal port numbering for the SAN256B with 48-port cards.

135	271	151	287	167	303	183	319	199	335	215	351	231	367	247	383
134	270	150	286	166	302	182	318	198	334	214	350	230	366	246	382
133	269	149	285	165	301	181	317	197	333	213	349	229	365	245	381
132	268	148	284	164	300	180	316	196	332	212	348	228	364	244	380
131	267	147	283	163	299	179	315	195	331	211	347	227	363	243	379
130	266	146	282	162	298	178	314	194	330	210	346	226	362	242	378
129	265	145	281	161	297	177	313	193	329	209	345	225	361	241	377
128	264	144	280	160	296	176	312	192	328	208	344	224	360	240	376
127	263	143	279	159	295	175	311	191	327	207	343	223	359	239	375
126	262	142	278	158	294	174	310	190	326	206	342	222	358	238	374
125	261	141	277	157	293	173	309	189	325	205	341	221	357	237	373
124	260	140	276	156	292	172	308	188	324	204	340	220	356	236	372
123	259	139	275	155	291	171	307	187	323	203	339	219	355	235	371
122	258	138	274	154	290	170	306	186	322	202	338	218	354	234	370
121	257	137	273	153	289	169	305	185	321	201	337	217	353	233	369
120	256	136	272	152	288	168	304	184	320	200	336	216	352	232	368
119	255	135	271	151	287	167	303	183	319	199	335	231	367	247	383
118	254	134	270	150	286	166	302	182	318	198	334	230	366	246	382
117	253	133	269	149	285	165	301	181	317	197	333	229	365	245	381
116	252	132	268	148	284	164	300	180	316	196	332	228	364	244	380
115	251	131	267	147	283	163	299	179	315	195	331	227	363	243	379
114	250	130	266	146	282	162	298	178	314	194	330	226	362	242	378
113	249	129	265	145	281	161	297	177	313	193	329	225	361	241	377
112	248	128	264	144	280	160	296	176	312	192	328	224	360	240	376
111	247	127	263	143	279	159	295	175	311	191	327	223	359	239	375
110	246	126	262	142	278	158	294	174	310	190	326	222	358	238	374
109	245	125	261	141	277	157	293	173	309	189	325	221	357	237	373
108	244	124	260	140	276	156	292	172	308	188	324	220	356	236	372
107	243	123	259	139	275	155	291	171	307	187	323	219	355	235	371
106	242	122	258	138	274	154	290	170	306	186	322	218	354	234	370
105	241	121	257	137	273	153	289	169	305	185	321	217	353	233	369
104	240	120	256	136	272	152	288	168	304	184	320	216	352	232	368
103	239	119	255	135	271	151	287	167	303	183	319	231	367	247	383
102	238	118	254	134	270	150	286	166	302	182	318	230	366	246	382
101	237	117	253	133	269	149	285	165	301	181	317	229	365	245	381
100	236	116	252	132	268	148	284	164	300	180	316	228	364	244	380
99	235	115	251	131	267	147	283	163	299	179	315	227	363	243	379
98	234	114	250	130	266	146	282	162	298	178	314	226	362	242	378
97	233	113	249	129	265	145	281	161	297	177	313	225	361	241	377
96	232	112	248	128	264	144	280	160	296	176	312	224	360	240	376
95	231	111	247	127	263	143	279	159	295	175	311	223	359	239	375
94	230	110	246	126	262	142	278	158	294	174	310	222	358	238	374
93	229	109	245	125	261	141	277	157	293	173	309	221	357	237	373
92	228	108	244	124	260	140	276	156	292	172	308	220	356	236	372
91	227	107	243	123	259	139	275	155	291	171	307	219	355	235	371
90	226	106	242	122	258	138	274	154	290	170	306	218	354	234	370
89	225	105	241	121	257	137	273	153	289	169	305	217	353	233	369
88	224	104	240	120	256	136	272	152	288	168	304	216	352	232	368
87	223	103	239	119	255	135	271	151	287	167	303	231	367	247	383
86	222	102	238	118	254	134	270	150	286	166	302	230	366	246	382
85	221	101	237	117	253	133	269	149	285	165	301	229	365	245	381
84	220	100	236	116	252	132	268	148	284	164	300	228	364	244	380
83	219	99	235	115	251	131	267	147	283	163	299	227	363	243	379
82	218	98	234	114	250	130	266	146	282	162	298	226	362	242	378
81	217	97	233	113	249	129	265	145	281	161	297	225	361	241	377
80	216	96	232	112	248	128	264	144	280	160	296	224	360	240	376
79	215	95	231	111	247	127	263	143	279	159	295	223	359	239	375
78	214	94	230	110	246	126	262	142	278	158	294	222	358	238	374
77	213	93	229	109	245	125	261	141	277	157	293	221	357	237	373
76	212	92	228	108	244	124	260	140	276	156	292	220	356	236	372
75	211	91	227	107	243	123	259	139	275	155	291	219	355	235	371
74	210	90	226	106	242	122	258	138	274	154	290	218	354	234	370
73	209	89	225	105	241	121	257	137	273	153	289	217	353	233	369
72	208	88	224	104	240	120	256	136	272	152	288	216	352	232	368
71	207	87	223	103	239	119	255	135	271	151	287	231	367	247	383
70	206	86	222	102	238	118	254	134	270	150	286	230	366	246	382
69	205	85	221	101	237	117	253	133	269	149	285	229	365	245	381
68	204	84	220	100	236	116	252	132	268	148	284	228	364	244	380
67	203	83	219	99	235	115	251	131	267	147	283	227	363	243	379
66	202	82	218	98	234	114	250	130	266	146	282	226	362	242	378
65	201	81	217	97	233	113	249	129	265	145	281	225	361	241	377
64	200	80	216	96	232	112	248	128	264	144	280	224	360	240	376

Figure 3-39 IBM System Storage SAN256B director 384-port numbering scheme

3.3.4 b-type top of rack switches

Here we describe the current B-Type rack mounted FC switches. The following advanced function switches are also available:

- ▶ 3758-L32 converged switch, full details, and functionality, described in *IBM Converged Switch B32*, SG24-7935-00, available at this website:
<http://www.redbooks.ibm.com/redpieces/abstracts/sg247935.html?Open>
- ▶ 2005-R04 and 2498-R06 multi protocol routers and extension switches, full details, and functionality, which can be found in *IBM System Storage b-type Multiprotocol Routing: An Introduction and Implementation*, SG24-7544-03, available at this website:
<http://www.redbooks.ibm.com/abstracts/sg247544.html?Open>
- ▶ 2498-E32 encryption switch, full details and functionality, which can be found in *Implementing the IBM System Storage SAN32B-E4 Encryption Switch*, SG24-7922, available at this website:
<http://www.redbooks.ibm.com/abstracts/sg247922.html?Open>

2498-B24

The 2498-B24 (also known as the SAN24B-4) is designed for medium-sized SAN environments. It can be used to create a wide range of high-performance SAN solutions, from simple, single-switch configurations to larger, multi-switch configurations. As an entry-level eight-port storage consolidation solution, it can support up to seven servers with a single path to either disk or tape. The Ports on Demand feature is designed to enable a base switch to grow to 16 and 24 ports to support more servers and more storage devices without taking the switch offline.

Figure 3-40 shows the SAN24B-4 fabric switch.




Figure 3-40 SAN24B-4

The switch requires Fabric OS v6.1+, and port hardware is based on the GoldenEye2 ASIC. One ASIC can support all 24 ports at 1, 2, 4, and 8 Gbps link speeds and they can be configured as F, FL, E, M and EX Ports. This switch provides a 1:1 subscription on all 24 ports.

The switch supports non-blocking architecture designed to provide up to 284 Gbps aggregate throughput with a 24-port configuration.

Standard configuration includes 8-port activation and capability to attach to hosts and storage device. Capability to attach to other SAN devices is standard on 2498-B24 and optional on 249824E.

Frame-based trunking with up to eight 8 Gbit/sec ports per ISL trunk with optional license; up to 64 Gbit/sec per ISL trunk (8 ports × 8 Gbit/sec [data rate])
Exchange-based load balancing across ISLs with DPS included in Fabric OS.

The switch also has a built-in USB port that can be used for firmware download, configuration upload and download, and supports save.

Figure 3-41 shows the SAN24B-4 fabric switch port layout and trunk groups.




Figure 3-41 SAN24B-4 port layout

2498-B40

The 2498-B40 (also known as the SAN40B-4) is a high performance enterprise fabric switch with 40 ports at 8 Gbps link speeds. This switch supports features such as Web Tools, Advanced Zoning, Full-Fabric support, Fabric Watch, and Enhanced Group Management standard. Optional features include ISL Trunking, Extended Fabrics, Advanced Performance Monitoring, Adaptive Networking, FICON CUP, and Integrated Routing.

The switch requires Fabric OS v6.1+, and port hardware is based on the Condor2 ASIC. One ASIC can support all 40 ports at 1, 2, 4 and 8 Gbps link speeds and they can be configured as F, FL, E, M and EX Ports. This switch provides a 1:1 subscription on all 40 ports.

Figure 3-42 shows the SAN40B-4 fabric switch.



Figure 3-42 SAN40B-4 fabric switch

The base model of the switch has 24 ports enabled, and the POD licenses are available in 8-port increments. The ports on the switch are grouped in 8-port groups matching the trunk group, and ISL Trunking speeds of up to 64 Gbps can be achieved per trunk. Integrated Routing is a licensed feature that is supported on every port of the switch and requires the POD license for all 40 ports.

The switch also has a built-in USB port that can be used for firmware download, configuration upload and download, and supportsave.

Figure 3-43 shows the port numbering scheme on the SAN40B-4.




Figure 3-43 SAN40B-4 ports

Two hot-swappable, redundant 125 W power supply and fan assemblies are included with the switch. These are field replaceable units (FRU). Each FRU has an ON/OFF switch AC plug and a power supply and fan status LED, and the switch has a 1U form factor.

Important: The USB port supports *only* a 2 GB Brocade branded USB drive. The 4 Gbps and 8 Gbps link speeds are supported *only* with Brocade branded SFPs.

2498-B80

The 2498-B80 (also know as SAN80B-4) is an 80-port, 8 Gbps enterprise fabric switch with 2U form factor. This switch supports features such as Web Tools, Advanced Zoning, Full-Fabric support, Fabric Watch, and Enhanced Group Management standard. Optional features include ISL Trunking, Extended Fabrics, Advanced Performance Monitoring, Adaptive Networking, FICON CUP, and Integrated Routing.

Figure 3-44 shows the SAN80-B4 fabric switch.



Figure 3-44 SAN80-B4 fabric switch

The switch requires Fabric OS v6.1 and port hardware is based on the GoldenEye2 ASIC. Each ASIC can support 32 ports at 1, 2, 4, and 8 Gbps link speeds and the switch has 9 ASICs. Ports can be configured as F, FL, E, M, and EX Ports.

The base model of the switch comes with 48 ports enabled and offers a 1:1 subscription at all 80 ports that can be activated with the POD licenses that are available in 16 port increments. Integrated Routing is a licensed feature that is supported on every port of the switch and requires the POD license for all 80 ports.

The ports on the switch are grouped in 8-port groups matching the trunk group. Figure 3-45 shows the port numbering scheme.

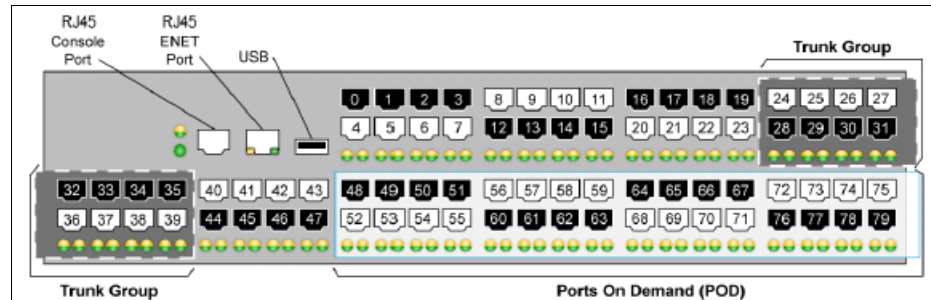


Figure 3-45 SAN80B-4 ports

ISL trunk speeds of up to 64 Gbps can be achieved per trunk. Dynamic Path Selection can be used for optimizing the performance and load balancing, and the switch can be managed using Web Tools. The built-in USB port can be used for firmware download, configuration upload and download, and supports save. The switch supports non-disruptive firmware downloads.

Important: The USB port supports *only* 2 GB Brocade branded USB drives. The 4 Gbps and 8 Gbps link speeds are supported *only* with Brocade branded SFPs.

The switch has two hot-swappable, redundant 300 W power supplies and three hot-swappable fan assemblies. Both the power supplies and the fan assemblies are FRUs, and they have a status LED. With a nominal power consumption of 260 W, this switch is extremely energy efficient.

Figure 3-46 shows the non-port side view of the switch and illustrates how the power supplies and the fans are identified in the Fabric OS.




Figure 3-46 SAN80B-4 non-port side view



Fabric Operating System

In this chapter, we provide an overview of the Fabric Operating System (FOS) software for the switches, including the firmware that initializes and manages the switch hardware and the diagnostics. We also introduce Fabric OS v6.4.0 and explain its new features and the changes introduced.

4.1 Fabric Operating System overview

The Fabric OS manages the operation of the switch and delivers the same, and compatible, functionality to all the different models of switches and directors. The switch firmware is designed to make the switches easy to install and use while retaining the flexibility that is required to accommodate user requirements.

The Fabric OS includes all the basic switch and fabric support software as well as optionally licensed software that you enable using license keys. It is composed of two major software components:

- ▶ Firmware that initializes and manages the switch hardware
- ▶ Diagnostics

Fabric OS v5.x and v6.x are Linux-based operating systems, while Fabric OS v3.x and prior were based on the VxWorks operating system.

In this chapter we also introduce the changes in versions 6.3.1 and 6.4.0, to help you understand the differences between versions, and understand the upgrades from one version to the next one.

4.2 Fabric Operating System v6.2.0 features

Releases: In this section you can find information about Fabric OS v6.2.0, which is here for legacy and compatibility reasons. This information is mainly for your reference, but it is important to understand the main changes between releases. We strongly advise you to update to the newer version of the Fabric OS to ensure that you have all the latest features available and bugs fixed.

4.2.1 New features

In addition to the support for the new hardware platform that was introduced with this release, many new features were introduced in Fabric OS v6.2.0:

- ▶ Virtual Fabrics:
 - Full VF feature support on SAN768B, SAN384B, SAN80B-4, SAN40B-4
 - Single physical chassis can be subdivided into two or more logical switches creating a logical fabric with other switches
 - Per-port assignment of ports to logical switches
 - Shared ISLs provide connectivity for multiple logical fabrics

- ▶ FCR and FCIP enhancements:
 - FCIP (SCSI) Read Tape Pipelining
 - Enhancements to SoTCP
 - LSAN Tagging
 - Support for Pathinfo over MetaSAN
 - Use FSPF cost in FCR backbone fabric to find shortest path to edge fabric
 - In-band management link over FCIP connections for the IBM SAN18B-R
 - TCP Byte Streaming for FCIP connections used with WAN optimization hardware
 - Improved FCIP statistics support, including TCP connection history, high water mark information and connection snapshot capability
- ▶ Support temporary licenses for Adaptive Networking, Integrated Routing, and Fabric Watch.
- ▶ Security enhancements:
 - IPv6 Auto-configuration
 - IPsec with IPv6 (for management port)
 - Configurable switch-wide policy requiring authentication of all HBAs
 - RADIUS enhancements allowing password expiration and source IP address information
 - LDAP enhancement allowing for alternate domain UPN
 - IPv6 certified for JITC Approved Product List
- ▶ FICON enhancements:
 - RNID support for CUP
 - New FC addressing modes for support with Virtual Fabrics
 - Support for FC8-48 blade with VF-enabled SAN768B/SAN384B for FICON environments
- ▶ Access gateway enhancements:
 - AG mode supported on IBM SAN40B-4
- ▶ Encryption enhancements:
 - Data encryption support in Virtual Fabrics environments
 - Support for tape encryption and compression.
 - Support for up to four FS8-18 Encryption blades in a single SAN768B or SAN384B chassis
- ▶ Brocade HBA feature support:
 - Beacon adjacent switch port from HCM
 - Fabric based boot LUN discovery
 - QoS nameserver support allowing query for QoS zone information
 - Support for FC Ping

- ▶ Miscellaneous features:
 - FC ping support between switches (ping switch WWN)
 - Provide path information through the CLI
 - Frame Redirection support in interopmode 3 (McDATA Open Fabric Mode)
 - Support for M-EOSn 239 Domain ID mode through FCR in interopmode 3
 - System-wide RASLOG
 - Port Auto-Disable support
 - Ethernet Port Bonding for management ports
 - New CLI command to configure F_Port receive buffer credits

4.2.2 Feature descriptions

In the topics that follow we describe some of the features introduced with Fabric OS v6.2.0 in greater depth.

Virtual Fabrics

In this section we describe the Virtual Fabrics capability:

- ▶ Virtual Fabrics (VF) is a new capability supported on the IBM SAN768B, SAN384B, SAN80B-4, and SAN40B-4 switches, and newer models. After being enabled, VF allows the user to divide a single physical chassis or switch into multiple “logical switches” by assigning individual ports to a logical switch. Each of these logical switches is managed as a completely independent layer 2 Fibre Channel switch, and can be deployed in independent fabrics known as “logical fabrics.”
- ▶ VF also allows the user to create a special logical switch known as the “base switch,” used for connectivity to other base switches and also as a backbone fabric for Fibre Channel Routing. Individual logical fabrics can utilize this shared base fabric for connectivity to other switches, providing efficient use of resources by sharing common ISL and ICL connections among multiple logical fabrics.
- ▶ The Virtual Fabrics feature is part of the base Fabric OS and does not require a license. Virtual Fabrics is fully compatible with legacy IBM b-series products as well as m-series switches and directors.

FCR and FCIP enhancements

In this section we describe various FCR and FCIP enhancements:

- ▶ FCIP (SCSI) Read Tape Pipelining: Anticipates host read operations and buffers data to reduce latency from the FCIP WAN.
- ▶ LSN tagging: Provides special behavior for designated LSN zones using either a new “speed” tag or “enforce” tag. The enforce tag allows individual FCRs to be configured to only import devices in specific LSN zones,

increasing scalability. The speed tag allows designated targets to remain imported, allowing sensitive hosts to discover targets much faster. This is useful when performing boot over LSAN operations.

- ▶ SoTCP enhancements: Improved network congestion management in Slow Start Mode helps to prevent host I/O from timing out.
- ▶ Support for Pathinfo over MetaSAN: The **Pathinfo** command has been enhanced to provide path information across routed fabrics, especially useful when troubleshooting connectivity problems across routed fabrics.
- ▶ Use of FSPF cost in FCR backbone fabric: Uses the most efficient route to reach a destination fabric, preventing the use of an FCIP link when an ISL is available.
- ▶ In-band management link over FCIP connections for the IBM SAN18B-R: Allows a management station to communicate with a remote SAN18B-R through the GE ports. This allows a single management station located on the WAN side of the SAB18B-R to communicate with the management interface on the CP for management tasks such as firmware downloads, SNMP polling, SNMP traps, troubleshooting, and configuration.

Enhanced native connectivity with McDATA products

In this section we describe enhanced native connectivity with McDATA:

- ▶ Frame Redirection: Fabric OS v6.2.0 and M-EOS v9.9 now support Frame Redirection in McDATA Open Fabric Mode (interopmode 3) fabrics. Frame Redirection zones must be created and activated from FOS platforms.
- ▶ FCR Support for M-EOSn 239 DID mode: FOS platforms now support EX_Port connections to McDATA Open Fabric Mode Mi10ks using the 239 DID setting.

Security enhancements

In this section we describe several security enhancements:

- ▶ IPv6 auto-configuration: Configurable stateless IPv6 auto-configuration support.
- ▶ IPSec with IPv6: Supports greater security for management ports by providing configurable security policies for IPv4/6 addresses.
- ▶ Switch-wide policy requiring HBA authentication: New configurable switch-wide setting requires the FC-SP bit to be set in FLOGI. If bit is not set, the FLOGI is rejected and the port will be disabled.
- ▶ RADIUS enhancements: New warning for RADIUS login allows users to configure how many days in advance they need to be notified of password expiration.

- ▶ LDAP enhancements: Added ability to provide an alternate UPN (userPrincipalName) to domain authentication.

Encryption enhancements

Additional support for application based tape encryption and compression, or tape encryption on its own, has been included with FOS 6.2.0

Contact your application vendor regarding their implementation of tape encryption.

Optionally licensed software

The Fabric OS v6.2.0 release includes all basic switch and fabric support software, as well as the following optionally licensed software that is enabled using license keys:

- ▶ Ports on Demand: Allows customers to instantly scale the fabric by provisioning additional ports through license key upgrade (applies to select models of switches).
- ▶ Extended Fabrics: Provides greater than 10km of switched fabric connectivity at full bandwidth over long distances (depending on platform this can be up to 3000km). Contact IBM or your vendor for the latest supported distances.
- ▶ ISL Trunking: Provides the ability to aggregate multiple physical links into one logical link for enhanced network performance and fault tolerance. Also includes Access Gateway ISL Trunking on those products that support Access Gateway deployment.
- ▶ Advanced Performance Monitoring: Enables performance monitoring of networked storage resources. This license includes the TopTalkers feature.

Top Talkers: The Top Talkers feature was introduced in Fabric OS v6.0.0 and is part of the optional Advanced Performance Monitoring license. This feature provides real-time information about the top n bandwidth consuming flows that pass through a specific point in the network. You can enable Top Talkers on individual F_Ports to provide information about top consumers of bandwidth for all E_Port connections on a switch.

- ▶ High Performance Extension over FCIP/FC: Formerly known as FC-IP Services, (available for the FR4-18i blade, IBM SAN18B-R, and SAN04B-R): This license key also includes the FC-Fastwrite feature and IPsec capabilities.
- ▶ Accelerator for FICON: This license enables unique FICON emulation support for IBM's Global Mirror application as well as Tape Pipelining for all FICON tape and virtual tape systems to significantly improve XRC and tape backup/recovery performance over virtually unlimited distance for SAN18B-R, upgraded SAN04B-R and FR4-18i.

- ▶ **Fabric Watch:** Monitors mission-critical switch operations. Fabric Watch now includes new Port Fencing capabilities.
- ▶ **FICON Management Server:** Also known as “CUP” (Control Unit Port), enables host-control of switches in Mainframe environments.
- ▶ **ICLs, or Inter Chassis Links:** Provide dedicated high-bandwidth links between two IBM SAN768B or SAN384B chassis, without consuming valuable front-end 8Gbps ports. Each SAN768B/SAN384B must have the ICL license installed in order to enable the ICL connections. (Available on the SAN768B/SAN384B only).
- ▶ **Enhanced Group Management:** This license, available only on the SAN768B, SAN384B, and other 8 Gbps platforms, enables full management of the device in a datacenter fabric with deeper element management functionality and greater management task aggregation throughout the environment. This license is used in conjunction with IBM's Data Center Fabric Manager (DCFM) application software.
- ▶ **Adaptive Networking:** Adaptive Networking provides a rich framework of capability allowing a user to ensure high priority connections obtain the network resources necessary for optimum performance, even in congested environments. The QoS SID/DID Prioritization and Ingress Rate Limiting features are the first components of this license, and are fully available on all 8 Gbps platforms.

Ingress: Ingress Rate Limiting was introduced with Fabric OS v6.0.0, and this feature allows the Application Specific Integrated Circuit (ASIC) to delay the return of BB_Credits to the external device. By doing so, a user can limit the throughput on the ingress side of a port, thereby removing potential congestion scenarios within a fabric caused by heavy bandwidth consumption by low priority applications. Ingress rate limiting is only supported on F/FL ports, and is only available on 8 Gbps capable ports.

QoS feature was also introduced with Fabric OS v6.0.0 and is available on all 8 Gbps capable ports on the 8 Gbps platforms. When congestion is detected, QoS allocates the largest portion of available bandwidth to high priority traffic and the smallest amount to low priority traffic. SID/DID flow pairs not explicitly set as having high or low priority automatically default to medium priority.

- ▶ **Integrated Routing:** This license allows ports in a SAN768B, SAN384B, SAN80B-4 or SAN40B-4 to be configured as EX_ports supporting Fibre Channel Routing. This eliminates the need to add an FR4-18i blade or use a SAN Router for FCR purposes, and also provides double the bandwidth for each FCR connection when connected to another 8 Gbps capable port.

- ▶ SAN04B-R Upgrade (For the IBM SAN04B-R only): This license allows customers to upgrade a 4- port (2 FC ports and 2 GE ports) SAN04B-R base to a full 18-port (16 FC ports and 2 GE ports) SAN18B-R configuration and feature capability. The upgraded SAN04B-R includes the complete High Performance Extension license feature set.
- ▶ Server Application Optimization: This new license introduced with FOS v6.2, when deployed with Brocade Server Adapters, optimizes overall application performance for physical servers and virtual machines by extending virtual channels to the server infrastructure. Application specific traffic flows can be configured, prioritized, and optimized throughout the entire data center infrastructure.

Temporary license support

The following licenses are available for 45-day temporary use, with a maximum of two temporary license per feature and per switch (90 days maximum):

- ▶ Fabric (E_port) license
- ▶ Extended Fabric license
- ▶ Trunking license
- ▶ High Performance Extension license
- ▶ Advanced Performance Monitoring license
- ▶ Adaptive Networking license (introduced in FOS v6.2)
- ▶ Fabric Watch license (introduced in FOS v6.2)
- ▶ Integrated Routing license (introduced in FOS v6.2)

4.3 Fabric OS v6.3.1 and v6.4.0 updates

Since the last edition of this book, two versions of Fabric OS have been released, and in this section we discuss the changes in versions 6.3.1 and 6.4.0. In terms of the CLI, the implementation is the same, but here we introduce the changes, and make sure that all the new features or differences from previous versions are documented. Additionally, we show screen captures of the changed DCFM interface.

Clarification of basic concepts: Before getting into details, let us review some important terms related to FCIP that you need to be familiar with, and that will help in the understanding of the new features in the two last releases:

- ▶ Circuit: A communication that is established between a source IP address to destination IP address.
- ▶ Tunnel: A collection of one or more circuits between two switches. Note that in the case of two or more circuits in the tunnel, the tunnel is trunked.

- ▶ VE Port: Virtual_E port that is behind one or more physical ports on each side of the tunnel.

In Figure 4-1 we can see a graphical representation of these terms.

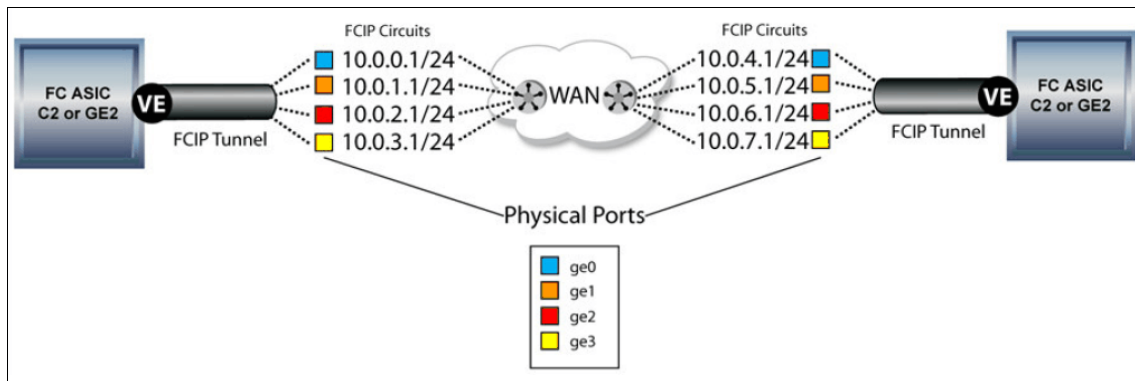


Figure 4-1 FCIP terms explanation

We now list the main changes in the two latest releases of the firmware that concern FCIP implementation. These changes do not affect the commands and the way to proceed or execute one FCIP configuration, but might have an impact on how to plan your infrastructure and how to deploy it. These changes can also explain why certain features will work after you update the firmware, but you cannot modify them afterwards, because the Fabric OS requires compliance with certain values. If this is the case in the latest release, we explain it.

4.3.1 Changes in Fabric OS version 6.3.1

Fabric OS 6.3.1 provides the following general enhancements:

- ▶ Virtual Fabrics with the 24 ports 8 Gbps blade's VE_Ports as XISLs; the VE_Ports on a blade can be assigned to a base switch for XISL use
- ▶ 10 GbE to 1 GbE tunnel and circuit connectivity configuration on the 24 ports blade, which allows connection between 10 GbE ports and 1 GbE ports on another Ethernet switch of the family
- ▶ Access Gateway support on the IBM Converged switch B32
- ▶ Added support for FCoE Initialization Protocol (FIP) v1.0 (equivalent to the FC-BB-5 version 2.0 standards, dated June 4 2009)
- ▶ VLAN tagging/802.1p
- ▶ IPSec (only for the IBM System Storage SAN06B-R)
- ▶ Port Beaconsing
- ▶ DCFM v10.4 support for all the foregoing features

In the rest of this section we explain some of these changes, and describe what has been introduced.

IPSec on the IBM SAN06B-R

Internet Protocol Security (IPSec) uses cryptographic security to ensure private, secure communications over IP networks. The IPSec protocol is used to secure the FCIP tunnel over a public IP WAN by using encryption.

IPSec feature capabilities

With the latest release of the firmware, all the b-type products will be able to handle IPSec. This was already possible in release 6.3.1 with the IBM System Storage SAN06B-R, and now it is supported in the FC routing blade.

There is no longer any need for an external IPSec appliance.

IPSec feature attributes

The FC routing switches will use the following IPSec attributes:

- ▶ The SAN06B-R and the FC routing blade will use the internal FPGA for hardware encryption operations.
- ▶ IPSec is configured using a pre-configured policy:
 - Internet Key Exchange (IKE) v2 for key negotiation
 - Encapsulating Security Payload (ESP) transport mode for IPSec (only the payload is encrypted)
 - Advanced Encryption Standard (AES) with 256 bit keys of encryption; AES-CGM-ESP is the mode AES uses
 - Security Association (SA) lifetime is roughly 2 GB of data sent through the SA
- ▶ IPSec will *only* support IPv4 for the Fabric OS release 6.4.0, and IPv6 is not supported at this time.

VLAN Tagging (IEEE 802.1p) / DSCP

This feature allows for differentiation of Fabric OS QoS levels across FCIP links. The L2CoS and Differentiated Services Code Point (DSCP) are used to assign QoS across the IP WAN. L2CoS refers to the Class of Service field defined by the IEEE 802.1p s

The L2COS and DSCP are configured on a per circuit basic. The values are configured per Fabric OS QoS priority basis within a circuit, and each fabric QoS can even be defined to have its own L2CoS bits and DSCP value.

Port beaconing

The port beaconing feature is enabled using the command **portbeacon --enable [slot/]port**. The normal LED output is suppressed, and the LEDs will flash amber and green in a 2.5 second pattern.

The beaconing will remain enabled until disabled using the command **portbeacon --disable [slot/]port**, and this is not persistent across reboots.

4.3.2 Changes in Fabric OS version 6.4.0

In this latest version (at the time of writing) of Fabric OS, there are more changes:

- ▶ IPv6 (IPv6 support for IPsec is not supported)
- ▶ IPsec (for FCoE 10GbE blade, already supported on SAN06B-R)
- ▶ Virtual EX_Port (for FCoE 10GbE blade, and already supported on SAN06B-R)
- ▶ Software compression for maximum compression ratio (for FCoE 10GbE blade, and already supported on SAN06B-R)
- ▶ DSCP
- ▶ Scalability increase (4 blades per chassis for FCoE 10GbE blade)
- ▶ Lossless DLS on FC ports
- ▶ TPerf enhancements
- ▶ Supports up to 100 ms with 0.1% packet loss for FCIP tunnels with one or both ends on 10 GbE
- ▶ DCFM v10.4 support for all the foregoing features

DCFM screen captures of the previously listed enhancements

In the following set of screen captures, you will see the enhancements in the DCFM software. This section is intended to show you how it looks in the DCFM interface in order to have a clear overview of the changes.

As you can see in Figure 4-2, DCFM now supports the configuration of FCIP tunnel advanced settings. You can reach this menu when creating or editing a tunnel, under the “Advanced Settings” button.

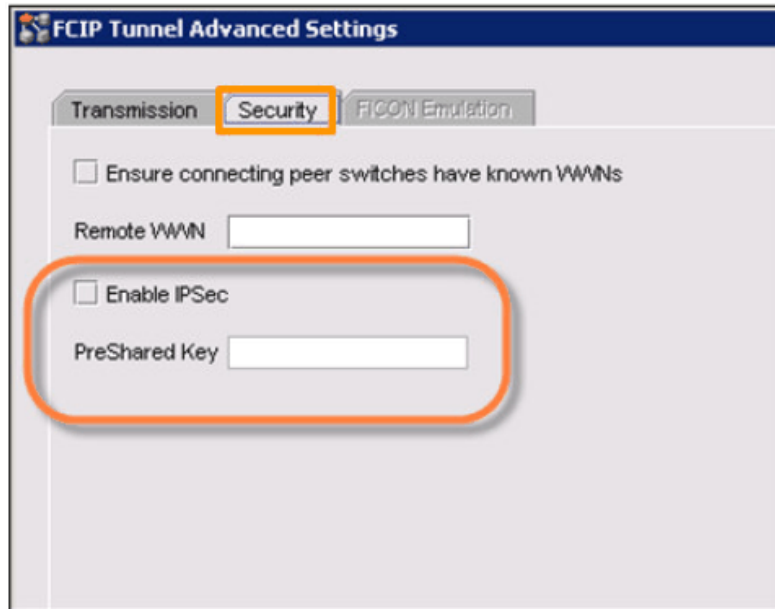


Figure 4-2 FCIP tunnel configuration on DCFM

In Figure 4-3, we can see also how DCFM is now able to help us configure the VLAN tagging options for FCIP, and DCFM will be able to discover both fabrics for the VLAN.

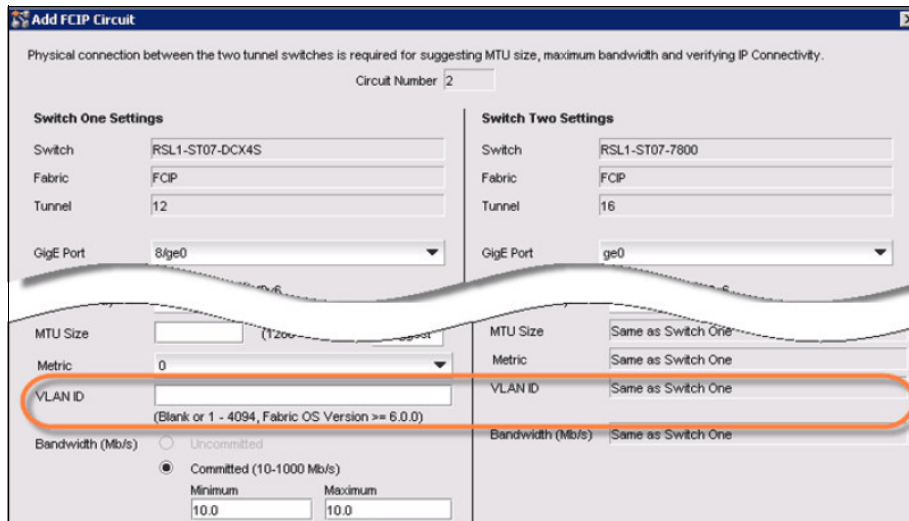


Figure 4-3 VLAN tagging configuration on DCFM

Figure 4-4 shows how to configure DSCP and L2CoS using DCFM, in the “FCIP circuit Advanced Settings” page for the circuit.

The image shows a screenshot of the "FCIP Circuit Advanced Settings" dialog box. The "Transmission" tab is selected. The following settings are visible:

- ☒ Selective Ack
- Keep Alive Time Out (ms): 10000 (500-7200000)
- Min. Retransmission Time (ms): 100
- Max. Retransmits: 8 (1-8)

The following settings are highlighted with an orange box:

L2Cos	F-Class	DSCP	F-Class	Value	Range
L2Cos F-Class	0	DSCP F-Class	0	0	(0-63)
L2Cos Low	0	DSCP Low	0	0	(0-63)
L2Cos Medium	0	DSCP Medium	0	0	(0-63)
L2Cos High	0	DSCP High	0	0	(0-63)

At the bottom right, there are buttons for OK, Cancel, and Help.

Figure 4-4 DSCP and L2CoS options

Figure 4-5 illustrates the configuration panel for the VEX_Port options, using DCFM.

Add FCIP Tunnel

Configure the settings for the tunnel to be added on the selected switch and a second switch if selected.
At least one circuit (IP interface) should be set up for adding a tunnel.

Switch One Settings

Switch: RSL1-ST07-DCX4S
Fabric: FCIP
Tunnel: 8/12 (Requires Circuit Configuration)
Description:
Port Type: ☐ VE Port ☒ VEX Port
Fabric ID:
Interop Mode: Brocade

Switch Two Settings

Switch:
Fabric:
Tunnel:
Description:
Port Type: ☒ VE Port ☐ VEX Port
Fabric ID:
Interop Mode: Brocade

FICON settings unavailable when VLAN or VEX Port is specified

Figure 4-5 VEX_Port configuration panel

IPv6 support will be added in the configuration panel, when available, as shown in Figure 4-6.

Add FCIP Circuit

Physical connection between the two tunnel switches is required for suggesting MTU size, maximum bandwidth and verifying IP Connectivity.
Circuit Number: 2

Switch One Settings

Switch: RSL1-ST07-DCX4S
Fabric: FCIP
Tunnel: 12
GigE Port: 8/ge0
IP Address Type: ☐ IPv4 ☒ IPv6
IP Address:
Prefix Length:
Default route will get created using the above IP address.
☐ Create Non-Default Route

Switch Two Settings

Switch: RSL1-ST07-7800
Fabric: FCIP
Tunnel: 16
GigE Port: ge0
IP Address Type: ☒ IPv4 ☐ IPv6
IP Address:
Subnet Mask:
Default route will get created using the above IP address.
☐ Create Non-Default Route

Figure 4-6 IPv6 support

There is a new “Delete” button in the FCIP tunnel configuration panel that will let us delete a tunnel from DCFM. This is shown in Figure 4-7.

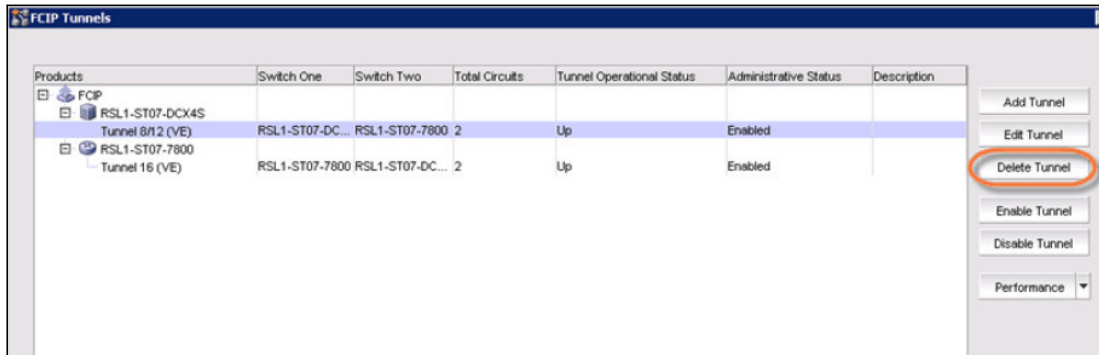


Figure 4-7 Delete button for FCIP tunnel on DCFM

There is a new “metric” field in DCFM that can be configured when adding a new FCIP circuit from the GUI. This is shown in Figure 4-8.

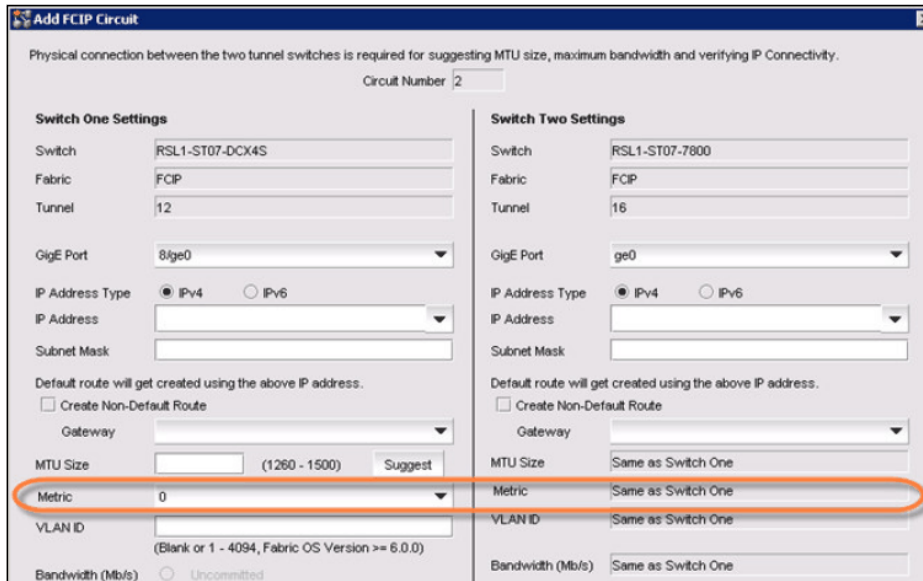


Figure 4-8 Metric field when adding an FCIP circuit on DCFM

Changes on the Traffic Isolation Zones with Fabric OS v6.4.0

With the latest 6.4.0 firmware of Fabric OS, there have been changes to Traffic Isolation (TI) Zones. These changes are listed here:

- ▶ Device ports can now be in multiple Traffic Isolation Zones at the same time.
- ▶ Devices in a *failover disabled* TI zone are able to communicate with local devices that are not part of the same TI zone.

- Prior to Fabric OS v6.4.0, devices would need to be in a failover enabled TI zone in order to communicate with local devices.
- There is an enhancement that ensures that the Domain Controller connectivity between the switches will never affect the TI zones.

Support: These new features are supported in Condor2 and GoldenEye2 ASICs. They are not supported on the 8 Gbps blades for DCX.

In the following sections we will describe the changes made in the latest release at the time of writing.

Overlapping TI Zones with failover disabled

Even with failover disabled, now it is possible to have devices in more than one zone. By enabling this feature, a link failure will not affect the alternative path.

The best way to illustrate this new change is with an example. As we can see in Figure 4-9, in the diagram we have a Channel Device that is a member of both the green and red zones with failover disabled. We also have two Control Units (CU), A and B, that are members of the two zones.




Figure 4-9 Link failure on TI Zone with failover disabled

This example illustrates how devices can now be members of multiple TI zones. The Channel is a member of both the green and the red TI zones, which have failover disabled. CU A is a member of the red and blue zones, and CU B is a member of the green and blue zones.

In the event of an ISL failure, which in this example would be in the ISL between the Channel and CU A, and the ISL goes offline, the communication with CU B will be maintained. However, the traffic between the Channel and CU A will be halted because failover is disabled.

Local device communication

The devices that are members of a TI zone can now communicate with local devices that are not members of the failover disabled TI zone. In versions prior to 6.4, this communication was blocked.

In Figure 4-10, we can see that now it is possible that a host can communicate with a local device even in a TI zone. In this example the host needs access to the tape library using a dedicated, failover disabled TI zone (in blue), and also to the local storage. The green zone in the figure represents this new capability to access local devices, even when the host is connected to a failover disabled zone.




Figure 4-10 Example of local device communication

Domain controller only routes

Suppose that we have three Domain Controller links between three switches and they are installed in serial in such a way that switch 2 is between the other two switches. In our case the links between the two first switches are failover enabled, and the link between the two others is not. In the case of a link failure between one of the two links in the communication of the first two switches, we would have a situation.

Prior to version 6.4, in the case of a link failure between more than two switches, if the main ISL between two switches was down, the first switch was not able to connect to the third switch through the second switch. Now, even in the case of a link failure between the first two controllers, the communication will reach the third controller.

Switch-to-switch management: Domain Controller refers to switch-to-switch communication coming from the embedded port PID 0xFFFFxx (where xx is the domain ID). This traffic is for switch-to-switch management. The traffic information will be updated on all the elements of the fabric that are connected to the management port.

Lossless DLS overview

With Fabric OS v6.4.0, there are improvements with respect to Lossless DLS. The Lossless Dynamic Load Sharing (DLS) allows for rebalancing of port paths without causing I/O failures. The Lossless mode means that no frame will be lost during rebalancing and only takes effect if DLS is enabled. If the end device also requires the order of frames to be maintained during the rebalancing operations, then In Order Delivery (IOD) can be enabled.

The combination of Lossless DLS and IOD is supported only in specific topologies, such as in a FICON environment.

Lossless DLS works by:

- ▶ Pausing ingress traffic (by not returning credits) with no impact on frames already in flight
- ▶ If IOD is enabled, giving sufficient time for frames already received on the original path to be transmitted
- ▶ Rebalancing the traffic to use all available paths
- ▶ Resuming traffic on the new path

Fabric OS v6.4.0 now allows Lossless DLS to be used with exchange-based routing. The IOD option can now be enabled or disabled, when in previous versions of Fabric OS, it was only supported on port-based routing.

Dynamic Path Sharing: Lossless DLS with exchange-based routing is also known as Dynamic Path Sharing (DPS).

This enhancement has some limitations:

- ▶ Lossless DLS cannot use XISLs in FICON environments.
- ▶ The entire data path in the switch fabric must be Condor2/GoldenEye2 ASICs.
- ▶ It is not supported on switches or blades that do not comply with the previous statements.

4.4 Firmware upgrade considerations

In this section we cover the considerations to upgrade the firmware to version 6.4.0 or higher.

In the IBM SAN768B with Virtual Fabrics enabled, the addressing mode defaults to 10-bit addressing. The 10-bit addressing works by borrowing the top two bits from the ALPA field of a 24-bit Fibre Channel address. 10-bit addressing is required in order to support the 512 user FC ports plus ICL ports in a single chassis. As a result of the additional ports provided by the 8 ports blade, the area numbers between 0x70 and 0x0F are no longer unique.

The following considerations apply when upgrading the firmware:

- ▶ Disable the ports on the IBM SAN768B logical switches that use 10-bit addressing mode and that have 8-bit areas in the range 0x70-0x8F.
 - Otherwise the firmware upgrade will fail with an error message.
 - Additional areas 0x70-0x8F are needed to support the 8 port blades. This is necessary even if users do not plan to use those blades.
 - If the areas are not in use, this step is not necessary.
- ▶ An 8-bit area is assigned to a port in the following scenario:
 - A port is connected to a switch in Access Gateway mode.
 - A port is connected to a Brocade HBA.
 - A user has explicitly bound an 8-bit area to a port (as in the case of FICON)
- ▶ Use the command **portaddress --show** to find areas in use.

This situation only applies to the IBM SAN768B when Virtual Fabrics are in use, not to the IBM SAN384B, which has no shared areas and therefore would not be a problem in this case.

4.4.1 Licensing changes

In Fabric OS v6.4.0, the slot-based licenses are automatically assigned to an appropriate slot when they are installed. The licenses can be unassigned and reassigned to different slots as needed using the **licenseslotconfig** command.

The slot-based licensing was first introduced with Fabric OS version 6.3.0. When a license is added, it is applied to the lowest slot number containing a blade that uses that license. With the old model, if the license needed to be moved to another slot, it had to be removed from the assigned slot and then added to the desired slot.

4.4.2 Fabric scalability

Scalability limits for Fabric OS v6.2.0 are essentially the same as those limits supported by Fabric OS v6.1.0. Fabrics of up to 6000 virtual or physical connections (WWNs logged into a single fabric) and 56 domains (domain support is the same as on previous Fabric OS releases) can be supported on the SAN768B, SAN384B and the SAN80B-4. Other products running Fabric OS v6.1.0 retain the same fabric limits as Fabric OS v5.3.x for non-routed fabrics (Layer 2 only, 56 domains and 2560-ports).

When operating in Native Connectivity modes (InteropMode 2 or 3), different scalability limits are supported. For both InteropMode 2 and 3, fabrics of up to 2048 virtual or physical connections (WWNs logged into a single fabric) and 31 domains are supported. This is an increase from Fabric OS v6.0.0 for InteropMode 3, which only supported up to 800 connections and 15 domains in a fabric.

The Virtual Fabrics capabilities supported in FOS v6.2.0 or later introduce additional factors to consider when assessing scalability. Specifically, when looking at the limits that an individual chassis or switch can support, it is no longer just a factor of the size of the Layer 2 fabric or the number of devices imported from edge fabrics. Virtual Fabrics allows a single physical chassis to participate in up to 8 separate Layer 2 fabrics, not including additional impacts from imported devices from FC Routing.

To account for this, each physical switch has limits which are supported for the aggregate environment. This means that if a single physical switch has three individual Logical Switches, each participating in an independent Logical Fabric, the total number of domains and participating host/storage devices in all three Logical Fabrics must be counted and compared against the physical switch limits.

The individual Logical Fabric limits are the same as those noted for a traditional Layer 2 fabric.

Table 4-1 lists supported limits specific to Virtual Fabrics enabled environments.

Table 4-1 Supported limits

	SAN40B/ SAN80B/ SAN384B and SAN768B
Max # of Logical Switches per Chassis/switch (including default and base switch)	3/4/8
Total # of participating devices (all Logical Fabrics) per chassis	6000
Total # of fabrics (Logical Switches and FCR-connected edge fabrics) per chassis	32/32/48
Total # of base switches creating a single base fabric	12
Total # of Logical Fabrics utilizing a single base fabric	48

Supported FCR scalability limits have increased in a few select areas and some new limits are included to reflect the new Integrated Routing support. Table 4-2 lists the Supported Routing scalability limits.

Table 4-2 Fibre Channel Routing scalability limits

Maximum number of edge fabrics per metaSAN	48
Maximum number of edge fabrics per chassis	16 (SAN18B-R and FR4-18i in SAN256B, SAN768B or SAN384B) 32 (SAN40B-4 and SAN80B-4) 48 (SAN768B and SAN384B)
Maximum number of local switches per edge fabric	26
Maximum number of WWNs per edge fabric	1500
Maximum number of local switches per edge fabric (M-EOS fabric) ^a	16
Maximum number of WWNs per edge fabric (M-EOS fabric) ^a	1500
Maximum number of imported devices per fabric (M-EOS fabric) ^a	1000
Maximum number of L2 switches per backbone fabric	12
Maximum number of FCRs per backbone fabric	12
Maximum number of WWNs per backbone fabric	1024
Maximum number of imported devices per fabric	1000
Maximum number of devices per metaSAN	10000

Maximum number of LSAN zones per metaSAN	5000 ^b
Maximum number of devices per LSAN zone	64
Maximum number of hops between edge switches	19
Ex_Ports per FCR (SAN256B/SAN768B with FR4-18i)	32
Ex_Ports per FCR (SAN768B with Integrated Routing)	64
Ex_Ports per chassis with Integrated Routing (SAN768B and SAN384B/SAN80B-4/ SAN40B-4)	128/80/40

- a. M-EOS fabrics must be running M-EOS 9.6.2 firmware or later.
- b. All BB FCRs with Fabric OS v6.0.0 and above. For M-EOS edge fabrics prior to v9.6, the limit is 1024 zones. For M-EOS edge fabrics operating with v9.6.x or later, the limit is 2048 zones.

Other important considerations include these:

- ▶ IPFC over FCR is only supported for edge to edge.
- ▶ FC Fast Write is only supported for edge to edge.
- ▶ The backbone cannot run in InteropMode 2 (McDATA Native Interop) or mode 3 (Open mode). It must be in Fabric OS native mode.
- ▶ All limits apply to Integrated Routing as well as FCR on SAN18B-R/FR4-18i unless otherwise noted.

4.4.3 FICON support

Fabric OS v6.1.0b provides FICON CUP support in Fabric OS/MEOS mixed fabrics operating in InteropMode 2. This support is available in fabrics with SAN768B, SAN40B-4, SAN80B-4, and SAN256M. Fabric OS v6.1.0b also adds support for configuring the MIHPTO (Missing Interrupt Handler Primary Time-out) value.

Fabric OS v6.1.0b includes enhanced CUP statistics counters comparable to those supported in M-EOS.

The FC4-48 and FC8-48 Fibre Channel port blades are *not* supported to connect to System z environments using FICON channels or using FCP zLinux on System z. To attach the SAN256B or SAN768B to the System z environment, use an FC4-16, FC4-32, FC8-16, or FC8-32 Fibre Channel port blade.

The SAN384B is not supported for FICON Cascading in interopmode 2 or 3 for use in mixed fabrics with M-EOS platforms.

4.5 Additional important notes and guidelines

Here we provide other important notes and guidelines:

- ▶ Virtual Fabrics:
 - On Virtual Fabrics capable platforms, the Virtual Fabrics feature must be enabled after upgrading to FOS v6.2.0 in order to utilize the related capabilities, including Logical Switches and Logical Fabrics. On units that ship with FOS v6.2.0 or later installed, the Virtual Fabrics feature is enabled by default on capable platforms.
 - When creating Logical Fabrics that include switches that are not Virtual Fabrics capable, it is possible to have two Logical Switches with different FIDs in the same fabric. Use extra caution to verify that the FIDs match for all switches in the same Logical Fabric.
 - The **aptpolicy** can be configured per logical switch. The Admin Guide indicates that it is a chassis level setting.
 - In order to support non-disruptive Hot Code Load on a Brocade 5100 with VF enabled, the total zoning DB size for the entire chassis must not exceed 1 MB.
 - A switch with Virtual Fabrics enabled cannot use Port Mirroring or participate in a fabric that is using IP Filter or Password Database distribution or Administrative Domains. The Virtual Fabrics feature must be disabled prior to deploying in a fabric using these features.
- ▶ Licensing Behavior:
 - When operating a switch with Fabric OS v6.2, some licenses might display as “Unknown.”
 - This is due to changes in licensing requirements for some features that no longer require a license key that might still be installed on a switch.
- ▶ Frame Redirection
 - In v6.2.0, Frame Redirection zoning is not allowed with Default Zoning (“all access” in IM0 and default zone in IM2).
 - This was allowed in prior releases. There is no SW enforcement to block the upgrade.

- ▶ Adaptive Networking/Flow-Based QoS Prioritization:
 - When using QoS in a fabric with 4G ports or switches, FOS v6.0 or later must be installed on all products in order to pass QoS info. E_Ports from the DCX to other switches must come up *after* 6.0 is running on those switches.
 - Flow based QoS is *not* supported on FC8 blades in the Brocade 48000.
 - Any products that are not capable of operating with FOS v6.0 *cannot* exist in a fabric with Flow based QoS. Major problems will occur if previous generation 2G products exist in the fabric.
- ▶ FCR switches:
 - All FCR switches need to be running FOS v6.2.0.
 - This is done in order to support M-EOS 239 Domain Mode on the i10K.
- ▶ FCR Backbone Fabric ID change:
 - With FC8 blades, the switch must be disabled to change the backbone fabric ID.
 - With routing and dual backbone fabrics, the backbone fabric ID must be changed to keep the IDs unique.
- ▶ Traffic Isolation over FCR:
 - All switches and Fibre Channel Routers both in edge and backbone fabrics must be running FOS v6.1.0 or later in order to support this feature.
 - In order for Traffic Isolation over FCR to function properly, the associated TI zones in each fabric (edge and backbone) need to have failover *enabled*.
- ▶ Integrated Routing:
 - To allow Hot Code Load on a Brocade 5100 when using Integrated Routing, the edge switch connected to the 5100 must be running Fabric OS v6.1 or later code.
 - Integrated Routing EX_Ports are only supported in the base switch on a switch with VF enabled.
 - Integrated Routing and TopTalkers are not concurrently supported in FOS v6.2. To use Integrated Routing, be sure to first disable TopTalkers prior to configuring EX_Ports.
 - Ports 16-47 on the FC8-48 blade cannot be used as Ex_Ports. Use only ports 0-15 for FCR on the 48-port blade.

- ▶ FCS Automatic Distribution:
 - When using the FCS Automatic Distribution feature in Fabric OS v6.0 or later, all switches in the fabric must be running Fabric OS v6.0 or later. If any switches are running Fabric OS v5.x or earlier, only manual distribution can be used.
 - Fabric OS v6.0 or later will allow only FCS automatic distribution when in strict mode, requiring only switches with Fabric OS v6.0 or later.
- ▶ Access Gateway:
 - When in Access Gateway mode, the Automatic Port Configuration policy might not work when attached to M-EOS switches. M-EOS ports must be set to G_Port to prevent problems with port type discovery.
 - Ports 16-47 on the FC8-48 blade cannot be used for Access Gateway F_Port Trunking connections.
- ▶ 10 Gbps Interoperability:
 - 10 Gbps interoperability between FC10-6 and McDATA blades is not supported.
 - However, the FC10-6 blade is supported in a chassis running in InteropMode 2 or 3 (FC10-6 to FC10-6 connections only).
 - An FC10-6 blade will not synchronize with a McDATA 10 Gbps blade, but this will not impact the system negatively.
- ▶ Traffic Isolation over FCR:
 - All switches and Fibre Channel Routers both in edge and backbone fabrics must be running Fabric OS v6.1.0 in order to support this feature.
 - It is essential to have “fail-over” policy *enabled* in all edge fabrics that are part of the traffic isolation zones, in order for the proper functioning of Traffic Isolation over FCR.
- ▶ FICON:
 - For the DCX, FICON CUP is not allowed with a 48-port blade in the Default Logical Switch.
 - All ports on a 48 port blade must be assigned to a user-defined Logical Switch to use them in a FICON CUP enabled switch.
- ▶ FICON CUP Cascading:
 - All switches must be running Fabric OS v6.1.0b in order to support this feature.

- ▶ Port Fencing:
 - The default settings for port fencing have very low thresholds and can fence ports that experience a small number of errors. It is best to increase these threshold values for use in production environments. Different platforms might require different threshold settings for optimum behavior. Port Fencing is only available with the optional Fabric Watch license.
 - When the port fencing feature is enabled for ITW or CRC errors, the first set of errors detected on an active link that meet the custom high threshold level set by the user (or the default threshold level) is always ignored to account for expected link transition errors. The port is only disabled upon detection of a second set of errors, i.e. the next time the user-set threshold level (or default threshold level) is reached. This prevents a port from being disabled due to normal link transition behaviors.
 - When using the Port Fencing feature, you must first run the **fwalarmsfilterset** command. This command enables the port and allows you to receive Port Fencing messages.
 - Port Fencing can be inadvertently disabled from Web Tools. This happens when you do the following operations:
 - i. Open the Fabric Watch configuration window.
 - ii. Check the **SNMP Trap** check box in the *Above* row.
 - This change in WebTools disables Port Fencing. If this happens, you must re-enable the Port Fencing bit from the command line interface. See 8.9.2, “Threshold Configuration tab” on page 300.
- ▶ Port Mirroring:
 - Proper behavior of Port Mirroring functionality requires that the entire frame path must contain either only 4 Gbps ASICs or 8 Gbps ASICs. If a frame path contains a mix of 4 Gbps and 8 Gbps ASICs, then this functionality will not work as intended.
 - On the SAN80B-4, the port mirroring feature has a limitation where all port mirror resources must stay within the same ASIC port group. The resources are the configure mirror port, Source Device, and Destination Device or ISL, if the Destination Device is located on another switch. The ASIC port groups are 0-15, 16-31, 32-47, 48-63, and 64-79. The routes will be broken if the port mirror resources are spread across multiple port groups.
 - Port Mirroring is not supported on a switch with the Virtual Fabrics feature enabled.



Management tools

In this chapter, we briefly discuss the various built-in and external management tools available to IBM/Brocade SAN-switch users. Examples of these include Fabric OS v6.4 Web Tools and the latest management software, Data Center Fabric Manager (DCFM).

5.1 Web Tools

Web Tools is an easy-to-use, graphical user interface (GUI) that you can use to monitor and manage single or small fabrics, switches, and ports from a standard workstation. You perform tasks using a Java™-capable Web browser from a workstation anywhere on the same network. Web Tools provide the administrative control point for Advanced Fabric Services, including Advanced Zoning, ISL Trunking, Advanced Performance Monitoring, and Fabric Watch.

Web Tools also provide an interface to Telnet commands to perform special switch functions and diagnostics that are available only through the Telnet interface. For some switch models, Web Tools provide a simplified interface, EZSwitchSetup, that allows less-experienced users to perform basic management tasks.

Highlights of the Web Tools features include these:

- ▶ Simplified management with a single interface for viewing all switches in the fabric and their current status, including Registered State Change Notification (RSCN) filtering
- ▶ Centralized administration and configuration tasks for the entire fabric, specific switches, or even individual ports
- ▶ Increased flexibility by performing administrative and configuration tasks from any remote location through a Web browser and Internet connection
- ▶ Ability to view real-time performance data for monitoring and tuning
- ▶ Ability to take advantage of an easy-to-use shortcut panel for commonly performed administrative functions

Beginning with Fabric Operating System v6.1 (Fabric OS), the Web Tools license is no longer required. Web Tools is enabled automatically on all the devices that are running Fabric OS v6.1 or later.

We discuss the current functionality that is available in Web Tools in Chapter 8, “Web Tools” on page 181.

5.2 Fabric Watch

Fabric Watch is a storage area network (SAN) health monitor for switches and was an optional licensed feature that is supported on Fabric OS v2.2 or higher, and it comes with no cost in the current versions of the Enterprise switches. Fabric Watch enables each switch to monitor its SAN fabric constantly for potential faults and to provide an alert automatically to any problems long before they become costly failures.

Fabric Watch tracks a variety of SAN fabric elements, events, and counters. It monitors fabric-wide events, ports, SFPs, environmental parameters, and enables early fault detection and isolation as well as performance measurement. Fabric Watch is easy to configure and can be used to select custom fabric elements and alert thresholds or choose from a selection of preconfigured settings. Fabric Watch can also be integrated easily with enterprise systems management solutions.

By implementing Fabric Watch, you can improve SAN availability rapidly and improve performance without installing new software or system administration tools.

For a growing number of organizations, SAN fabrics are a mission-critical part of their systems architecture. These fabrics can include hundreds of elements, such as hosts, storage devices, switches, and interswitch links (ISLs). A flexible solution like Fabric Watch can optimize SAN value by tracking a wide spectrum of fabric events.

For example, Fabric Watch monitors the following situations:

- ▶ Fabric resources, including fabric reconfiguration, zoning changes, and new logins
- ▶ Switch environmental functions such as temperature, power supply, and fan status, along with security violations
- ▶ Port state transitions, errors, and traffic information for multiple port classes as well as operational values for supported models of “smart” GBICs/SFPs
- ▶ Performance information for AL_PA and end-to-end metrics

Fabric Watch allows you to define how often to measure each switch and fabric element, and to specify notification thresholds. Whenever fabric elements exceed these thresholds, Fabric Watch automatically provides notification using several methods, including email messages, SNMP traps, and log entries.

Fabric Watch provides the following types of automatic notifications:

- ▶ A continuous alarm provides a warning message whenever a threshold is breached; it continues to send alerts until the condition is corrected. For example, if a switch exceeds its temperature threshold, Fabric Watch activates an alarm at every measurement interval until the temperature returns to an acceptable level.
- ▶ A triggered alarm generates the first warning when a threshold condition is reached and a second alarm when the threshold condition is cleared.

Fabric Watch provides event notifications in several different formats to ensure that event details are accessible from all platforms and operating systems. In response to an event, Fabric Watch can record event data as any (or all) of the following types:

- ▶ Simple Network Management Protocol (SNMP) trap: The SNMP performs an operation called a *trap* that notifies a management station (a workstation that runs network management applications using SNMP protocol) when events occur. The software has to be configured to receive trap information from the network device and the SNMP agent on the switch has to be configured to send the trap to the management station. An SNMP trap forwards the following information to an SNMP management station:
 - Name of the element whose counter registered an event
 - Class, area, and index number of the threshold that the counter crossed
 - Event type
 - Value of the counter that exceeded the threshold
 - State of the element that triggered the alarm
 - Source of the trap

The trap stores event information but does not actively send alerts. Port changes do not generate SNMP traps. Support for SNMP makes Fabric Watch readily compatible with both network and enterprise management solutions.

- ▶ Switch Event log: Following an event, Fabric Watch adds an entry to the internal event log of an individual switch and can store up to 1024 error messages. The error log stores the event information but does not actively send alerts.
- ▶ Port Log Lock: The port log locks to retain detailed information about an event, preventing the information from being overwritten as the log becomes full. This alarm stores event information but does not actively send alerts, which is done automatically when some thresholds are exceeded and an alert is triggered.
- ▶ RAPI Trap: This Fabric Watch alarm actively alerts you to events by forwarding all event information to a designated proxy switch. The host API configures the proxy switch automatically based on firmware version. The switch forwards the information to a server and alerts the SAN manager to event activity.
- ▶ Email notification: Following an event, Fabric Watch sends email alerts to a specified email address with information about a switch event. An email alert can send information about any error from any element, area, and class. The email specifies the threshold and describes the event, much like an error message.

Fabric Watch is designed for rapid deployment: Simply enabling Fabric Watch permits immediate fabric monitoring, and is also designed for rapid custom configuration. You can create and modify configuration files easily using a text editor and then distribute configurations to all the switches in the SAN through the Fabric OS configuration management utility. Fabric Watch also comes with preconfigured profiles for rapid implementation.

5.3 SNMP

Simple Network Management Protocol (SNMP) is an industry-standard method of monitoring and managing network devices. This protocol promotes interoperability, because SNMP-capable systems must adhere to a common set of framework and language rules.

Understanding the components of SNMP makes it possible to use third-party tools to view, browse, and manipulate switch variables (MIBs) remotely as well as to set up an enterprise-level management process. Every switch and director supports SNMP.

5.4 Data Center Fabric Manager

Data Center Fabric Manager (DCFM) Enterprise is a comprehensive network management application that enables end-to-end management of data center fabrics. Data Center Fabric Manager Enterprise manages and secures the flow of data across multiple fabrics—empowering organizations to achieve their goals related to Service Level Agreements (SLA), security, and compliance while containing their operating expenses.

To account for the enormous growth in data moving within and across data centers, Data Center Fabric Manager Enterprise provides unprecedented scalability and performance that helps maximize data availability. In addition, it features easy-to-use administration tools that streamline or automate repetitive tasks so organizations can achieve unprecedented levels of productivity and efficiency.

As a key component of the IBM/Brocade Data Center Fabric (DCF) architecture, Data Center Fabric Manager Enterprise is designed for unified management of data center fabrics—from storage ports all the way to the Host Bus Adapters (HBA), both physical and virtual. It configures and manages the IBM/Brocade Backbone product family along with IBM/Brocade directors, routers, and switches. Moreover, it supports encryption capabilities for data-at-rest and HBA products, as well as new storage networking technologies such as Virtual Fabrics and emerging protocols such as FCoE and CEE.

DCFM: Organizations that need to manage smaller SAN environments can utilize DCFM Professional, an application included with IBM/Brocade switches that contains a subset of DCFM Enterprise features.

DCFM provides the following benefits:

- ▶ Centralizes management of IBM/Brocade multiprotocol fabrics within and across data centers, including support for FCoE and CEE as they become available
- ▶ Supports data center virtualization, including visualizing and setting QoS levels for applications running on virtual machines and support for Virtual Fabrics
- ▶ Reduces expenses and maximizes productivity by automating tasks and providing easy-to-use, wizard-driven operations
- ▶ Helps organizations meet SLAs through industry-leading monitoring, troubleshooting, diagnostics, and event notification capabilities
- ▶ Secures data flow from applications to storage across data center fabrics by managing user access controls and ensuring consistent security settings
- ▶ Delivers real-time and historical performance monitoring to enable proactive problem diagnosis, maximize resource utilization, and facilitate capacity planning
- ▶ Integrates seamlessly with leading third-party automation solutions to provide a holistic approach to data center management

Figure 5-1 shows a sample DCFM main window.

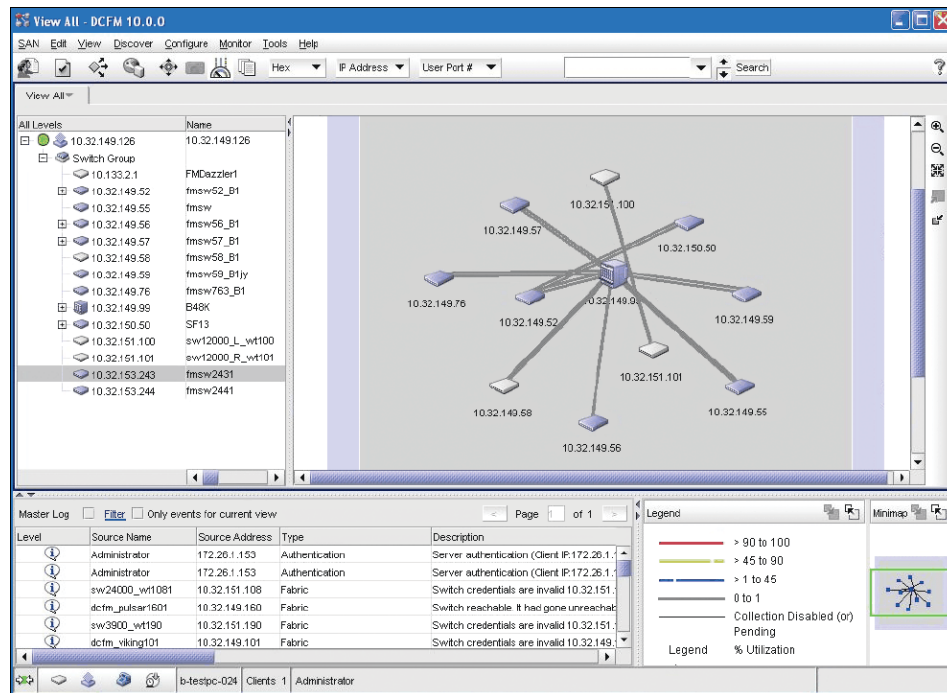


Figure 5-1 DCFM main window

5.4.1 Time-saving tools

As the demand for IT administrator resources continues to increase, Data Center Fabric Manager Enterprise provides time-saving tools that automate repetitive tasks through an intuitive wizard-driven approach. As a result, Data Center Fabric Manager Enterprise offers these capabilities:

- ▶ Streamlines firmware downloads, switch configuration backups/restores, and the collection of supportSave diagnostics data across groups of switches
- ▶ Enables organizations to edit zone information and preview the impact of proposed changes without affecting the production fabric
- ▶ Provides a configurable fabric snapshot/compare capability to track changes to fabric objects and membership
- ▶ Generates reports for status changes, port usage, zoning activation, and more—either on-demand, automatically, on a scheduled basis, or in response to a fabric event
- ▶ Configures and manages FICON and cascaded FICON environments while providing powerful analysis and diagnostic facilities

- ▶ Configures, monitors, and optimizes Fibre Channel over IP (FCIP) tunnels across WAN links
- ▶ Provides comprehensive support for Fibre Channel Routing, including configuration, zoning, visualization, analysis, and troubleshooting tools

5.4.2 Exceeding service level agreements

In order to gain a competitive advantage, organizations must meet or exceed established SLAs with their internal and external stakeholders. These SLAs depend on an efficient data center fabric that provides critical connectivity from servers (and their applications) to corresponding storage.

Data Center Fabric Manager Enterprise seamlessly manages multiple data center fabrics to help optimize performance and maximize the availability of data and networked resources. It does so by utilizing industry-leading capabilities for monitoring, troubleshooting, diagnostics, and proactive event notification.

Data Center Fabric Manager Enterprise also includes these functions:

- ▶ Proactively monitors critical fabric health information at varying levels of detail
- ▶ Configures Quality of Service (QoS) priorities for applications running on both physical and virtual servers to optimize performance for high-priority applications
- ▶ Collects and displays performance statistics in real-time and historical views for proactive problem determination
- ▶ Enables organizations to capture, back up, and compare switch configuration profiles through advanced replication capabilities
- ▶ Performs automatic data collection and triggers Call Home notification for easier fault isolation, diagnosis, and remote support
- ▶ Provides proactive alerts with real-time logging, diagnostic, and fault isolation capabilities to resolve issues before they impact SLAs

Upgrades: Previous versions of IBM/Brocade management software Fabric Manager (FM) or Enterprise Fabric Connectivity Manager (EFCM) can be upgraded to DCFM. An automatic migration facility will migrate most of the data and configurations from EFCM or Fabric Manager to DCFM as long as the minimum versions levels are met (EFCM 9.6 or higher or Fabric Manager 5.4 or higher).

DCFM Installation, configuration, and management are covered in depth in Chapter 9, “IBM System Storage Data Center Fabric Manager” on page 323.



Implementation

In this chapter, we discuss the initial setup to implement the switches. We then describe the EZSwitchSetup, a starter kit that greatly simplifies the setup and implementation of storage area network (SAN) switches.

6.1 Implementation

In the topics that follow, we show how to implement the switches initially.

6.1.1 Initial setup

Prior to configuring the IBM System Storage SAN switch, it must be physically mounted and connected to the appropriate electrical outlets. The amount of planning and preparation that is required for the installation is dependent upon the SAN product that is installed. See the Brocade hardware reference guide for the model that you plan to install, because this guide highlights the key aspects for consideration. You must arrange for your IBM service representative to install the chassis or rack physically in the location that you have planned.

After the switch is installed and turned on, it requires some initial configuration parameters to be set. All of the b-type switches require the same initial setup. The fundamental steps have not changed from the earlier switch models.

Sequence to turn on a switch

Timing: These steps take a minimum of three minutes to complete.

When you turn on or restart the switch, the following sequence of steps occurs:

1. Early power-on self-test (POST) diagnostics run. POST runs before Fabric Operating System (Fabric OS) starts.
2. The Fabric OS initializes.
3. The hardware initializes. The switch resets, the internal addresses are assigned, the Ethernet port initializes, the serial port initializes, and the front panel initializes.
4. A full POST runs.
5. The links initialize. Receiver and transmitter negotiation runs to bring the connected ports online.
6. During the Fabric Login (FLOGI), link parameters exchange to determine whether any ports are connected to other switches. If so, FLOGI negotiates which switch becomes the principal switch.
7. Domain addresses are assigned. After the principal switch is identified, port addresses are assigned. Each switch tries to keep the same domain ID that it used previously. Previous IDs are stored in the configuration Flash memory.
8. The routing table is constructed. After the addresses are assigned, the unicast routing tables are constructed.

9. Normal Nx_Port operation is enabled.

The chart in Figure 6-1 describes the initialization sequence of a port when a device is connected to it.

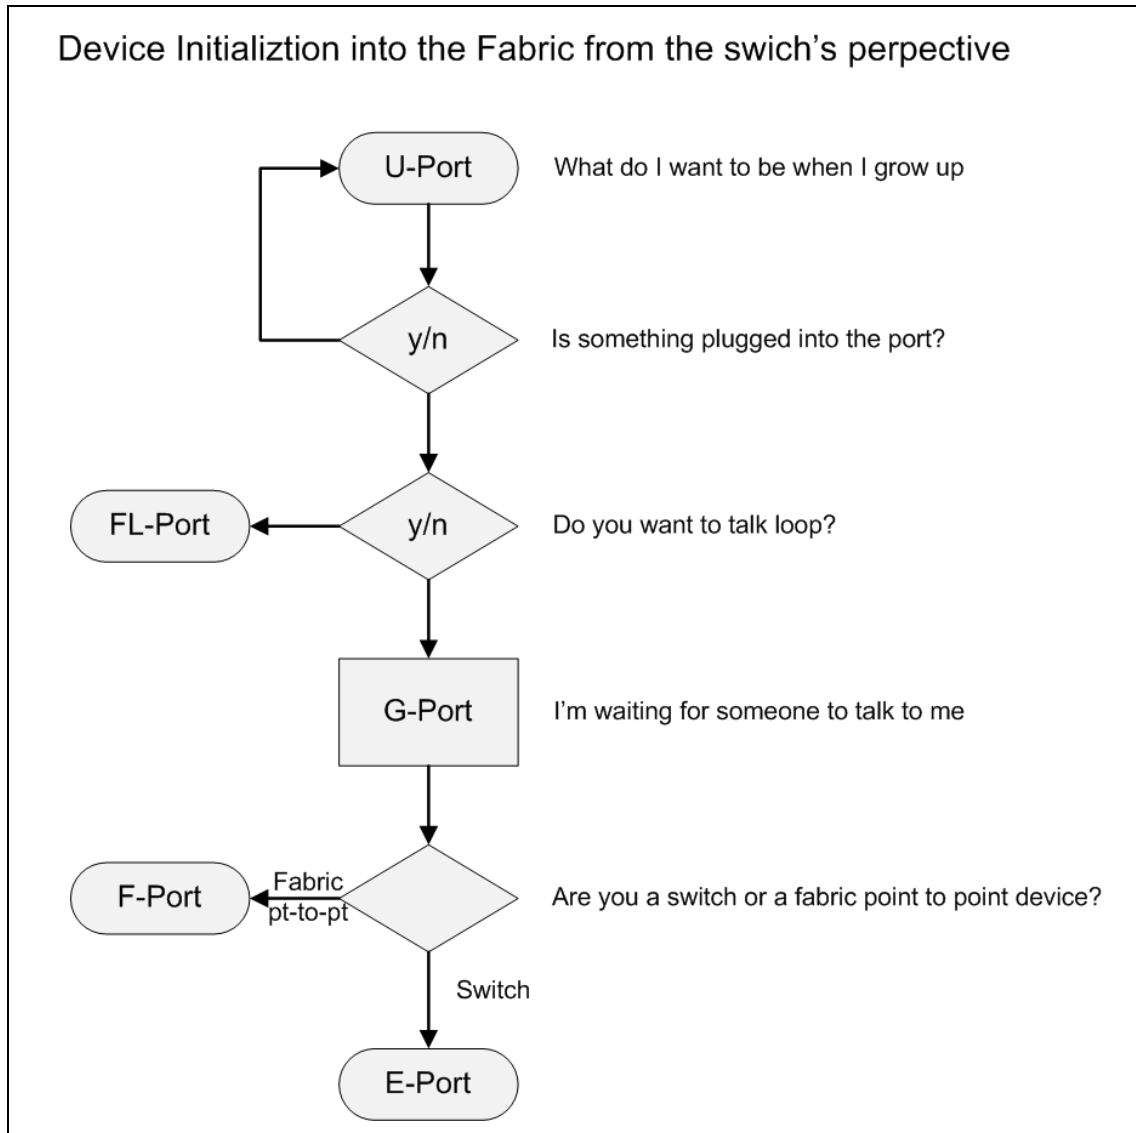


Figure 6-1 Flowchart showing device initialization

Basic setup functions

After you install the switch or director into a rack, and it has passed successfully through the POST tests, you need to perform some basic setup functions.

By connecting to the switch using a terminal emulator, you can see the switch POST tests as they progress.

Example 6-1 shows the startup of a SAN80B switch.

Example 6-1 SAN80B startup

The system is coming up, please wait...

U-Boot 1.1.3 (Apr 13 2009 - 21:30:54)

CPU: 8548_E, Version: 2.0, (0x80390020)

Core: E500, Version: 2.0, (0x80210020)

Clock Configuration:

CPU:1199 MHz, CCB: 399 MHz,

DDR: 199 MHz, LBC: 49 MHz

L1: D-cache 32 kB enabled

I-cache 32 kB enabled

Board: Thor

CPU Board Revision 255.198 (0xffc6)

PCI2: disabled

I2C: ready

DRAM: Initializing DDRSDRAM

memsize = 400

DDR: 1024 MB

Now running in RAM - U-Boot at: 3fb8e000

trap_init : 0x0

system inventory subsystem initialized

FLASH: 4 MB

L2 cache 512KB: enabled

ATA interface setup: io_base=0xf8f00000, port=0x1f0, ctl=0x3f6

PCI:

Skip our host bridge

00 11 1657 0011 0280 1a

00 12 1657 0011 0280 1a

00 13 1657 0011 0280 1a

01 01 1657 0011 0280 1a

01 02 1657 0011 0280 1a

01 03 1657 0011 0280 1a

01 04 1657 0011 0280 32

00 14 12d8 8150 0604 3b


```

02 01 1657 0011 0280 1a
02 02 1657 0011 0280 1a
02 03 1131 1561 0c03 1a
02 03 1131 1562 0c03 1a
00 15 12d8 8150 0604 1a
In:  serial
Out: serial
Err: serial
Net:
ENET0: PHY is Broadcom BCM5241 10/100 BaseT PHY (143bc31)
ENET1: PHY is not applicable
ENET2: PHY is not applicable
ENET3: PHY is not applicable

Checking system RAM - press any key to stop test

Checking memory address: 00100000

System RAM test using Default POST RAM Test succeeded.

set_bootstatus: BS_LOAD_OS, platform_idx = 6
Hit ESC to stop autoboot: 0
Map file at LBA sector 0x9dbff
## Booting image at 00400000 ...
   Image Name:   Linux-2.6.14.2
   Image Type:   PowerPC Linux Multi-File Image (gzip compressed)
   Data Size:    2890659 Bytes = 2.8 MB
   Load Address: 00000000
   Entry Point:  00000000
   Contents:
     Image 0:    1814143 Bytes = 1.7 MB
     Image 1:    1076503 Bytes = 1 MB
   Uncompressing Multi-File Image ... ## Current stack ends at
0x3FB6CBC0 => set upper limit to 0x00800000
## initrd at 0x005BAECC ... 0x006C1BE2 (len=1076503=0x106D17)
   Loading Ramdisk to 1fef9000, end 1ffffd17 ... OK
   initrd_start = 1fef9000, initrd_end = 1ffffd17
## Transferring control to Linux (at address 00000000) ...
tlbcam_index=11
mpc85xx_setup: Doing Pcie bridge setup
SILKWORM_HWSEM: This BD 64 is not supported
PCI: Cannot allocate resource region 2 of PCI bridge 1
PCI: Cannot allocate resource region 2 of PCI bridge 2
Installing Linux 2.6 Kernel
Attempting to find a root file system on hda1...

```

```
INIT: version 2.78 booting
Bypassing firmware validation.
INIT: Entering runlevel: 3
uptime: 429388564; sup_qid: 0
```

Fabric OS (IBM_SAN80B_217)

```
IBM_SAN80B_217 console login:
2009/07/17-22:51:22, [HAM-1004], 565, CHASSIS, INFO, SAN80B_182,
Processor rebooted - Reboot
SNMP Research EMANATE/Lite Agent Version 16.2.0.9
Copyright 1989-2006 SNMP Research, Inc.
ki_state_set: redundant ACTIVATE for instance 1
2009/07/17-22:51:33, [FCR-1069], 566, FID 128, INFO, IBM_SAN80B_217,
The FC Routing service is enabled.
2009/07/17-22:51:33, [FCR-1068], 567, FID 128, INFO, IBM_SAN80B_217,
The FC Routing service is disabled.
All service instances on Active
POST1: Started running Fri Jul 17 22:51:53 GMT 2009
POST1: Test #1 - Running turboramtest
POST1: Script PASSED with exit status of 0 Fri Jul 17 22:52:03 GMT 2009
took (0:0:10)
POST2: Started running Fri Jul 17 22:52:03 GMT 2009
POST2: Test #1 - Running portloopbacktest (SERDES)
POST2: Test #2 - Running portloopbacktest (BI LINKS FE_BI->CORE_BI)
POST2: Running diagshow
POST2: Script PASSED with exit status of 0 Fri Jul 17 22:52:40 GMT 2009
took (0:0:37)
2009/07/17-22:52:41, [BL-1000], 568, CHASSIS, INFO, SAN80B_182,
Initializing ports...
2009/07/17-22:52:51, [BL-1021], 569, CHASSIS, INFO, SAN80B_182, Retry
1, internal port retry initialization, ports: 3 188 .
2009/07/17-22:52:52, [BL-1001], 570, CHASSIS, INFO, SAN80B_182, Port
initialization completed.
2009/07/17-22:53:08, [SNMP-1008], 571, FID 128, INFO, IBM_SAN80B_217,
The last device change happened at : Fri Jul 17 22:53:06 2009
```

Fabric OS (IBM_SAN80B_217)

IBM_SAN80B_217 console login:

To get to the console login prompt, you must press the Enter key. It is useful to be aware of the standard boot up sequence for your switch so that, if a problem occurs, it is easy to distinguish between standard and abnormal behavior.

6.1.2 The command-line interface initial setup

To manage a switch, director, or a SAN768B backbone from a remote workstation on a network, you have to set the IP address, subnet mask, and gateway address for the Ethernet management interface on the switch. On SAN768B, SAN384B, and SAN256B, set the networking parameters for the following three entities:

- ▶ Logical switch
- ▶ CP 0
- ▶ CP 1

You can modify these settings using the **ipAddrSet** command. (We show the steps to modify the settings in 6.1.3, “SAN768B, SAN384B, and SAN256B configuration procedure” on page 136.)

The default IP address and subnet mask for the SAN24B-4, SAN40B-4, and SAN80B-4 switches are 10.77.77.77 and 255.255.255.0.

The default IP addresses and subnet masks for a SAN768B, SAN384B, or SAN256B are as follows:

- ▶ CP 0: 10.77.77.75 and 255.255.255.0 (CP 0 is the CP Card in slot 6)
- ▶ CP 1: 10.77.77.74 and 255.255.255.0 (CP 1 is the CP Card in slot 7)

There is no default IP address for logical switch SW 0 on SAN768B, SAN384B, or SAN256B.

For successful implementation, set the following parameters as well:

- ▶ Domain ID: For switches to be connected together within a fabric, each switch must have a unique domain ID. The default domain ID for a switch is 1. If two switches are connected through an ISL and if they both have the same Domain ID, they become segmented. You can modify domain IDs using the **configure** command. We show an example of how to use this command in “Connecting to the switch” on page 139.
- ▶ Switch names: Set a switch name to identify different switches within a site. This name is very helpful in easily identifying a switch to which you are connected. Using the **switchName** command, you can assign your own switch names, which can be up to 15 characters long, must begin with an alpha character, and can include alpha, numeric, and underscore characters.

Configuration examples

We describe the steps to configure these settings in the sections that follow, using two examples:

- ▶ SAN32B-3
- ▶ SAN768B

To configure these settings will take approximately 15 minutes. The following items are required:

- ▶ SAN switch installed physically and connected to a power source.
- ▶ A workstation that has a terminal emulator application. In our examples, we used **putty.exe** for both serial and telnet connections.
- ▶ The serial cable that is provided with the switch for connecting the switch to the workstation. If your workstation does not have a 9-pin serial port, you might require an adapter. We used a USB serial adapter to connect.
- ▶ Unused IP address or addresses plus gateway IP address and subnet mask. SAN256B, SAN384B and SAN768B requires three IP addresses (SW 0, CP 0, and CP 1), while SAN24B-4, SAN40B-4, and SAN80B-4 need one IP address.
- ▶ Ethernet cable for connecting the switch to the workstation or to a network that contains the workstation.
- ▶ SWL or LWL SFPs and fiber optic cables as required.

Attention: Do not connect the switch to your LAN until the IP settings are configured properly and they do not conflict with any other devices in your network.

It is important to leave at least 3.28 ft (1 m) of slack for each port cable. This extra length provides room to remove and replace the switch, allows for inadvertent movement of the rack, and helps prevent the cables from being bent to less than the minimum bend radius.

Use hook-and-loop straps to secure and to organize fiber optic cables. Do not use tie wraps on fiber optic cables, because these wraps are easily overtightened and can damage the optic fibers.

Setting the IP address using the serial port

In this section, we describe the steps necessary to set the IP address using the serial port on an IBM SAN32B-3. The procedure is the same for all IBM/Brocade switches except for the SAN256B, SAN384B and SAN768B. (We show the steps for these two products in 6.1.3, “SAN768B, SAN384B, and SAN256B configuration procedure” on page 136.)

Follow these steps to set the IP address:

1. Remove the shipping plug from the serial port and insert the serial cable that is provided with the switch.
2. Connect the other end of the serial cable to an RS-232 serial port on the workstation. If you do not have a male DB-9 serial port connector on your workstation, you can use a USB serial adapter.

Tip: The serial cable shipped with the switch is a straight-through cable, not a cross-over cable. Label the cable as such to minimize confusion at a later date.

3. Verify that the switch is on and initialization has completed by confirming that the system and power status LEDs are both on and green.
4. Disable any serial communication programs running on the workstation, such as PDA synchronization.
5. Open a terminal emulator application (such as HyperTerminal or **putty.exe** on a Windows® workstation or TERM in a UNIX® environment), and configure as follows:
 - a. In Microsoft Windows environment, adjust the following parameters and values if necessary:
 - Bits per second: 9600
 - Databits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: None

Figure 6-2 shows the PuTTY serial connection configuration options.

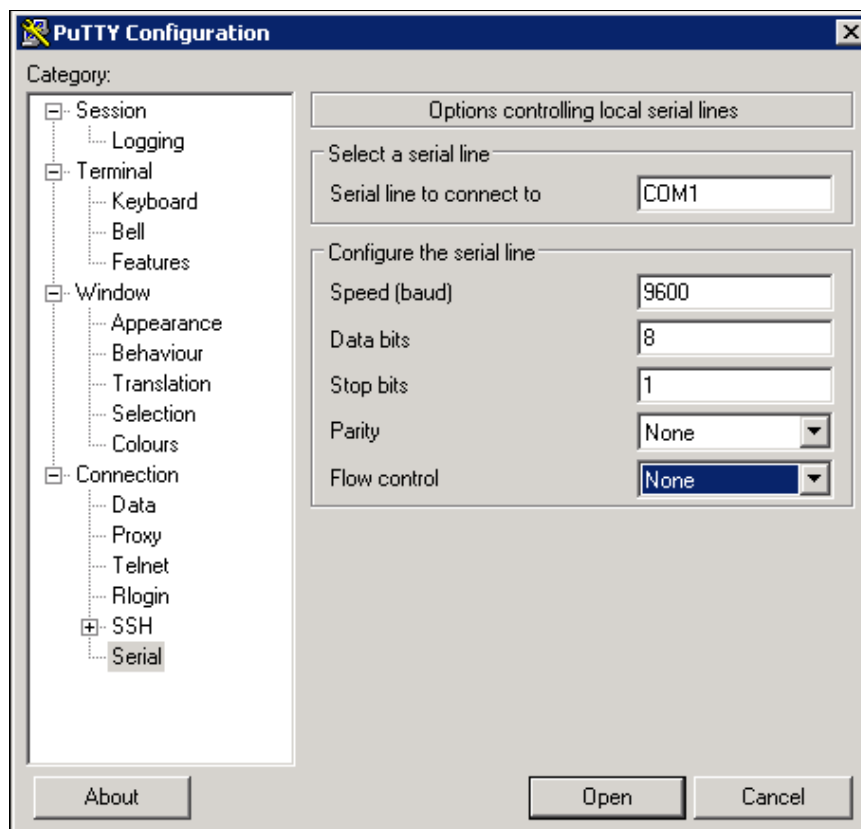


Figure 6-2 PuTTY connection options

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

tip /dev/ttyb -9600

From the terminal emulator application, log on to the switch through the serial connection. The default administrative logon is **admin** and the default password is **password**. If you have just turned on the switch, you might have to press Enter to display the login prompt following the message: Port Initialization Completed.

When logging into a new switch, you are requested to change the password. To skip this request, press Ctrl+C. You are prompted to change the password again at your next login. If you choose to change the password at this stage, you are prompted to change the password for each of the generic user accounts (root, factory, admin, and user). When all the passwords are changed, they are saved to stable storage. Change the password prior to connecting the switch to your network.

6. Enter the **ipAddrSet** command at the prompt.

Then, enter the appropriate values at the corresponding prompts, as shown in Example 6-2.

Example 6-2 Entering network settings with ipAddrSet command

```
IBM_SAN80B_4_217:admin> ipAddrSet
Ethernet IP Address [10.64.210.217]: Enter new IP address
Ethernet Subnetmask [255.255.240.0]: Enter new subnet mask
Gateway IP Address [10.64.208.1]: Enter new gateway ip address
DHCP [Off]:
```

7. Verify that the address was set correctly by entering the **ipAddrShow** command.

Example 6-3 displays the values that you entered in the previous step.

Example 6-3 ipAddrShow command output

```
IBM_SAN80B_4_217:admin> ipAddrShow

SWITCH
Ethernet IP Address: 10.64.210.217
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.64.208.1
DHCP: Off
IBM_SAN80B_4_217:admin>
```

8. After verifying that the IP address is correct, remove the serial cable, and replace the shipping plug in the serial port.

Serial Port: The serial port is intended only for use during the initial setting of the IP address and for service purposes. Do not use the serial port for day-to-day management and monitoring operations.

9. Record the IP address for future reference.

After the IP address is set, you can connect the switch to the managing workstation by Ethernet cable (this can be a direct cross-over connection or through a network) by following these steps:

1. Remove the shipping cover from the Ethernet port.
2. Insert one end of an Ethernet cable in the Ethernet port.
3. Connect the other end of the Ethernet cable to the workstation or to an Ethernet network that contains the workstation.

Important: The switch can now be accessed remotely, using Telnet or Web Tools. As a result, it is important to ensure that the switch is not modified simultaneously from any other connections during the remaining steps.

6.1.3 SAN768B, SAN384B, and SAN256B configuration procedure

The initial configuration of a SAN768B, SAN384B, or SAN256B requires a serial connection. (Note that the following examples were carried out on a SAN768B; however, the procedure is identical for a SAN384B and a SAN256B.)

Follow these steps to establish a serial connection and log in to the director:

1. Make sure that the SAN768B is turned on and that POST is complete by verifying that all power LED indicators on the port blades and CP blades display a steady green light.
2. Use the serial cable that is provided with the SAN768B to connect the serial console port on the active CP blade to a workstation.

Attention: The active CP blade is indicated by an illuminated blue LED. The LED on the standby CP blade must be off. The serial port is intended primarily for use during the initial setting of the IP addresses and for service purposes.

3. Access the SAN768B using a terminal emulator application (such as HyperTerminal or **putty.exe** on Windows or TERM in a UNIX environment).
4. Disable any serial communication programs running on the workstation (such as synchronization programs).
5. Open the terminal emulator application and configure as follows:
 - Bits per second: 9600
 - Databits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: None

For most UNIX systems, enter the following string at the prompt:

tip /dev/ttyb -9600

When the terminal emulator application stops reporting information, press Enter. You receive the following login prompt:

Fabric OS (IBM_SAN768B)
IBM_SAN768 console login:

6. Log in to the SAN768B as **admin**. The default password is **password**.

Passwords: At the initial login, you are prompted to enter new admin and user passwords.

7. Change the passwords. Passwords can be 8 to 40 characters long. They must begin with an alphabetic character and can include numeric characters, the period (.), and the underscore (_). Passwords are case-sensitive, and they are not displayed when you enter them on the command line. To skip modifying the password, press *Ctrl+C*, as shown in Example 6-4.

Example 6-4 Chassis console login

Fabric OS (IBM_SAN768B)

IBM_SAN768B console login: **admin**
Password:

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.
IBM_SAN768B:admin>
.

-
8. View the active CP LED on the CP blades in slots 6 and 7 or enter the **haShow** command to verify which CP blade is active. Example 6-5 shows the output for the **haShow** command. You can modify the configuration only through a login session to the active CP blade.

Example 6-5 hashow command

IBM_SAN768:admin> **hashow**
Local CP (Slot 7, CP1): Active, Cold Recovered
Remote CP (Slot 6, CP0): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
IBM_SAN768:admin>

To configure the IP addresses for the logical switch and both CP blades (from the active CP blade), follow these steps:

1. Log in to the active CP as admin using the serial cable connection.
2. Set up the logical switch IP address and subnet mask by entering the **ipAddrSet -sw 0** command at the command prompt. Enter the requested information at the prompts, which are shown in Example 6-6.

Example 6-6 Setting the logical switch IP address

```
IBM_SAN768:admin> ipAddrSet -chassis
Ethernet IP Address [10.64.210.210]: Enter the IP address
Ethernet Subnetmask [255.255.240.0]: Enter the subnet mask
IBM_SAN768:admin>
```

Reserved addresses: The addresses 10.0.0.0 through 10.0.0.255 are reserved and used internally by the Brocade DCX. External IP addresses must not use this range.

3. Set up the CP0 blade IP address by entering the **ipAddrSet -cp 0** command at the prompt. This is the CP blade in slot 6. Enter the requested information at the prompts, as shown in Example 6-7.

Example 6-7 Setting the CP 0 IP address

```
IBM_SAN768_210:admin> ipAddrSet -cp 0
Host Name [IBM_SAN768_210_cp0]: Enter the hostname
Ethernet IP Address [10.64.210.211]: Enter the new IP address
Ethernet Subnetmask [255.255.240.0]: Enter the new IP subnet mask
Gateway IP Address [10.64.208.1]: Enter the gateway IP address
IBM_SAN768_210:admin>
```

4. Set up the CP1 blade IP address by entering the **ipAddrSet -cp 1** command at the prompt, as shown in Example 6-8. This is the CP blade in slot 6.

Example 6-8 Setting the CP 1 IP address

```
IBM_SAN768_210:admin> ipAddrSet -cp 1
Host Name [IBM_SAN768_210_cp1]: Enter the hostname
Ethernet IP Address [10.64.210.212]: Enter the new IP address
Ethernet Subnetmask [255.255.240.0]: Enter the new IP subnet mask
Gateway IP Address [10.64.208.1]: Enter the gateway IP address
IBM_SAN768_210:admin>
```

After entering all the IP addresses, you can use the **ipAddrShow** command to verify the settings. Example 6-9 shows the output of this command on our SAN768B.

Example 6-9 The ipAddrShow command output

```
IBM_SAN768_210:admin> ipAddrShow

CHASSIS
Ethernet IP Address: 10.64.210.210
```

Ethernet Subnetmask: 255.255.240.0

CP0

Ethernet IP Address: 10.64.210.211

Ethernet Subnetmask: 255.255.240.0

Host Name: IBM_SAN768_210_cp0

Gateway IP Address: 10.64.208.1

CP1

Ethernet IP Address: 10.64.210.212

Ethernet Subnetmask: 255.255.240.0

Host Name: IBM_SAN768_210_cp1

Gateway IP Address: 10.64.208.1

Backplane IP address of CP0 : 10.0.0.5

Backplane IP address of CP1 : 10.0.0.6

IPv6 Autoconfiguration Enabled: Yes

Local IPv6 Addresses:

IPv6 Gateways:

IBM_SAN768_210:admin>

Reboot not needed: Although the SAN768B hardware reference manual suggests that a reboot is required when changing the IP address, this reboot is not necessary. You can change the IP address online without rebooting the director.

The terminal serial port can be used to monitor error messages through a serial connection. It is not intended for use as a command interface during normal operations. If this port is not going to be in ongoing use, remove the serial cable and protect the port from dust by replacing the shipping cap. This completes the initial configuration.

6.1.4 Connecting to the switch

After using a serial connection to configure the IP addresses for the SAN768B, you have to connect both the active and the standby CP blade to the local area network (LAN). Connect the CP blades to a private network/VLAN to provide additional security to your SAN as well as to protect it from network broadcast storms or other problems.

By establishing an Ethernet connection, you can complete the configuration using either the serial session or a Telnet session or through the graphical management interfaces (Web Tools and Data Center Fabric Manager). However, you must ensure that the SAN768B configuration is not modified from other connections at the same time.

To establish an Ethernet connection, follow these steps:

1. Remove the shipping plug from the Ethernet port on the active CP blade.
2. Insert one end of an Ethernet cable into the Ethernet port.
3. Connect the other end to an Ethernet 10/100 Base-T LAN.

The SAN768B can now be accessed by remote connection using any of the available management tools, such as Telnet, Web Tools, or Data Center Fabric Manager.

4. Repeat steps 1 through 3 for the standby CP blade.
5. To complete any additional configuration procedures through a Telnet session, log in to the SAN768B using Telnet with the admin login. The default password is **password**.

Important: When managing the SAN768B backbone, use the **-chassis** IP address for management GUI and Telnet access. Unless you are carrying out activities to a specific CP, this address prevents unpredictable results.

6.1.5 Setting the switch name

The switch name of the SAN768B can be up to 15 characters long, can include alpha, numeric, and underscore characters, and must begin with an alpha character. Setting meaningful names for your switches simplifies the management of your SAN. Ideally, you need to define an appropriate naming convention and use this naming convention to provide standardized names for your switches.

Customizing the name

To customize the name, follow these steps:

1. Enter the **switchName** command with the new name in quotes (see Example 6-10). The change will be committed but the prompt will not change until the telnet session is reconnected.

Example 6-10 Changing the SAN768B name

```
switch:admin> switchName "IBM_SAN768B_210"  
Committing configuration...  
Done.  
switch:admin>
```

2. Record the new name for future reference.

Setting the Domain ID

Each switch in the fabric must have a unique Domain ID. The Domain ID can be set using the **configure** command. You can also allow the Domain ID to be set automatically. The default Domain ID for the SAN768B is 1.

To set the Domain ID, follow these steps:

1. Enter the **fabricShow** command to determine the current Domain IDs available.
2. Enter the **switchDisable** command to disable the SAN768B.
3. Enter the **configure** command. Enter **y** at the Fabric parameters prompt:
Fabric parameters (yes, y, no, n): [no] **y**
4. Then, enter a unique Domain ID:
Domain: (1..239) [1] **3**
5. Complete the remaining prompts or press Ctrl+D to accept the other settings and to exit.
6. Enter the **switchEnable** command to re-enable the SAN768B.
7. Add SFPs and fiber optic cables to the ports as required.

Cables: The ports and cables that are used in trunking groups must meet specific requirements.

8. Remove the shipping plug from the ports to be used.

9. Position the SFP so that the key (the tab near the cable-end of the SFP) is on top, and insert the SFP into the port until it is firmly seated and the latching mechanism makes a clicking sound. For specific instructions, see the SFP manufacturer's documentation.

Attention: The SFP module is keyed so that it can only be inserted correctly into the port. If the module does not slide in easily, make sure it is not upside down.

10. Connect the fiber optic cables to the SFPs as appropriate to the fabric topology by positioning each cable so that the key (the ridge on one side of the cable connector) is aligned with the slot in the SFP, then inserting the cable into the SFP until it is firmly seated and the latching mechanism makes a clicking sound.

Attention: The cable is keyed so that it can only be inserted correctly into the SFP. If the cable does not slide in easily, try turning it over.

11. Verify the correct operation of the switch.
12. Enter the following command at the Telnet prompt to verify the switch and port status:

switchShow

Backups: This command provides information about the status of the switch and the ports. Always back up the configuration after any initial configuration changes and then perform backups periodically thereafter. This ensures that a complete configuration is available if ever required for uploading to a replacement switch. Switch configuration is backed up by issuing a **configUpload** to the FTP server.

6.1.6 The Port Identifier format

The Port Identifier (PID) format is a fabric wide parameter that must be set to the same value on all switches in the fabric. SAN devices use PID for routing and zoning services. Historically, the IBM/Brocade SAN products support the following PID format types:

- ▶ VC encoded
- ▶ Native (PID format 0)
- ▶ Core (PID format 1)
- ▶ Extended edge (PID format 2)

Remember: Changing the switch PID format might require a reboot of UNIX servers that bind by port ID.

If the switch PID format is set to a value other than 1 on existing switches, you can change it by following these steps:

1. Disable the switch with the **switchDisable** command:
`switchDisable`
2. Then, run the **configure** command:
`configure`
3. Enter **y** when prompted to set Fabric parameters:
Fabric parameters (yes, y, no, n): [no] **y**
4. Press Enter to use default parameters for settings until you are prompted for the switch PID format setting. Set the parameter to **1**:
Core Switch PID Format: (0..1) [0] **1**
Switch PID Format: (1..2) [1] **1**
5. Continue to press Enter to skip other settings. You receive the following message:
Committing configuration...done.
6. Enable the switch using the **switchenable** command.
7. Fastboot the switch using the **fastboot** command.

6.1.7 Setting the date

Now is also a good opportunity to set the date and time in the switch. Although a switch with the incorrect date and time will function properly, it is best to make these values realistic, because they are used for time stamping during logging of events. We suggest that you set these parameters prior to any further operations, because you might find this information very helpful if you have to troubleshoot at a later date.

Set the day and time using the **date MMDDhhmmYY** command, where:

MM	Month
DD	Day
hh	hour
mm	minutes
YY	Year

See Example 6-12 showing the use of this command.

Example 6-12 Setting the date and time

```
IBM_SAN384B_213:admin> date 0714164509  
Tue Jul 14 16:45:00 UTC 2009
```

Time: Use an NTP server to ensure that all switches in your environment are on the same time.

The steps for the installation are now complete, although it is best to upgrade to the latest level of firmware that is available before making the switch available for use.

6.1.8 Firmware update

When the switch is delivered, it might not have the latest firmware installed. It is a good practice to update the firmware to the latest version (or to the version in line with other switches in your SAN) before putting it in production. In this section, we show an example of firmware update using a Telnet session. Other possible ways to update the code are as follows:

- ▶ Web Tools:

We explain how to do a SAN switch firmware update with Web Tools in 8.7.4, “Firmware Download tab” on page 259.

- ▶ Data Center Fabric Manager

Set the Telnet session timeout value to **0**. This effectively disables the timeout (so that your session will not time out during the firmware upgrade procedure). Use the following command:

timeout 0

Because the new timeout value takes effect with the next login, you now need to log out and log back in.

Normally, the next steps are to configure the switch upload (with the **configUpload** command) and to save the support information (with the **supportSave** command). However, because you are performing the initial setup and have not yet configured the switch, you do not perform these steps now.

Therefore, continue with the actual firmware update, using the **firmwareDownload** command. Example 6-13 uses the FTP server with IP address 10.64.210.103 and logs in as **IBM** with password **password**. The Fabric OS v6.2.0e files are stored on this FTP server.

Example 6-13 Firmware update with firmwareDownload command

```
SAN32B_3_146:admin> firmwaredownload -p ftp  
10.64.210.103,ibm,/,password  
Server IP: 10.64.210.103, Protocol IPv4  
Checking system settings for firmwaredownload...  
System settings check passed.
```

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

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

```
Do you want to continue [Y]: y  
Firmware is being downloaded to the switch. This step may take up to 30  
minutes.  
Preparing for firmwaredownload...  
Start to install packages...  
dir #####  
ldconfig #####  
glibc #####  
glibc-linuxthreads #####  
bash #####  
readline #####  
terminfo #####  
termcap #####  
...  
Lines deleted for clarity  
...  
tz #####  
mtracer-tool #####  
sysstat #####  
ipv6 #####  
awk #####  
ipsec #####  
kernel-module-ipsec #####  
Writing kernel image into flash.
```

```
.....  
Finished writing kernel image.  
Removing unneeded files, please wait ...  
Finished removing unneeded files.
```

All packages have been downloaded successfully.
Firmware has been downloaded to the secondary partition of the switch.
HA Rebooting ...

To verify that the switch firmware was updated properly, you can use the **firmwareShow** and **version** commands. Example 6-14 shows the output of these two commands. The switch now runs Fabric OS v6.1.0.

Example 6-14 Verifying the switch firmware version

```
SAN32B_3_146:admin> firmwareshow  
Appl      Primary/Secondary Versions  
-----  
FOS       v6.2.0e  
          v6.2.0e  
SAN32B_3_146:admin> version  
Kernel:    2.6.14.2  
Fabric OS: v6.2.0e  
Made on:   Tue Apr 14 21:26:47 2009  
Flash:     Wed Jul 15 17:02:37 2009  
BootProm:  4.6.6
```

6.1.9 SAN256B optional modem setup

Each CP blade in the SAN256B contains a modem serial port for connection to a Hayes-compatible modem. The modem serial ports are wired as standard DTE ports and have the same commands, log in capabilities, and operational behavior as the terminal serial ports. However, asynchronous informational messages and other unsolicited text are not sent to the modem ports. No additional software is required to use modems with the director.

Modems: The director detects modems only during power-on, reboot, or a CP blade failover sequence. Set up the modems before powering on the director. For increased security, any active modem sessions are disconnected automatically if the modem cable is disconnected. For optimal security, disconnect the modem cable when it is not in use.

You can ensure high availability of the modem connection by connecting a separate modem to each CP blade and then connecting both modems to a shared telephone line. This connection ensures an available telephone connection to the active CP blade even if a failover occurs; however, it is necessary to log in again after a failover. When both CP blades are connected to a shared telephone line, callers are dialed in to the active CP blade automatically, which answers on the first ring. If the active CP blade cannot answer for any reason, the standby CP blade answers on the seventh ring and allows login to proceed.

Connection: If a modem connection is set up, connect a modem to each CP blade, as shown in Figure 6-3.




Figure 6-3 Optional modem line and data connections

Connecting modems

Attention: Set up the modems before turning on the director and connect it to the fabric.

The following items are required to set up two modems to work with the director:

- ▶ Two Hayes-compatible modems, such as the Zoom/Modem V.92 EXT Model 3049
- ▶ Two standard modem cables, DB25 (male) to DB9 (female)
- ▶ One RJ-11 “Y” adapter for standard Telco wiring or equivalent circuitry (three total connections)
- ▶ One analog telephone line

Attention: Turn off the director before connecting cables to the modem ports.

Complete the following steps to connect the modems to the director:

1. Set up the two modem units and corresponding power connections, but do not turn on the modems until all cables are attached.
2. Connect the modem cables to the modems and to the director RS-232 modem ports.
3. Connect the telephone line inputs on the modems to the RJ-11 Y connector. This effectively places both modems on a single telephone line.
4. Optionally connect a telephone handset to one of the phone connections on the modems.
5. Connect the “Y” adapter to an appropriate analog telephone line and document the dial-in number for later use.
6. Turn on the modems and verify that the Modem Ready indicator illuminates on both units.
7. Turn on the director, to allow the director to recognize the modems.

When the modems are connected, you can use a Telco system to dial in to the modems and verify that they answer and communicate as expected. If a dial-out modem facility is not available, you can use a terminal emulation program on a computer workstation (or mobile computer) that has an attached modem.

This procedure is only required if a dial-out modem facility is not already available for testing the director modem connections.

Perform the following steps to set up the optional remote modem:

1. Connect the remote modem to the workstation, as shown in Figure 6-4.
2. Disable any serial communication programs running on the workstation (such as a synchronization program for a PDA).




Figure 6-4 Remote modem setup

3. Launch the terminal emulator application and configure it as shown in Table 6-1.

Table 6-1 Configuration parameters

Parameter	Value
Port Speed	115200 ^a
Data Protocol	Standard EC
Compression	Enabled
Flow Control	Hardware
Databits	8
Parity	None
Stop Bits	1
Modulation	Standard

a. Port usually defaults to the highest speed that the modem supports but might negotiate at a slower speed.

4. Follow the instructions from the modem manufacturer to set up and verify modem operation.

Verifying the modem connection

This section provides information about how to verify that the modems are connected correctly.

Equipment: This procedure requires either a Telco system to dial in to the modems or a terminal emulation program on a mobile computer or workstation that has an attached modem.

Complete the following steps to verify the modem connection:

1. Verify that both modem cables are connected firmly.
2. Turn on the modems, if they are not already on.

Boot: The modems must be turned on and operational before the director is turned on, to allow the director to detect the modems during boot.

3. Verify that both modems indicate that they are ready by illuminating their Clear to Send (CS), Terminal Ready (TR), and Modem Ready (MR) indicators. If this illumination does not occur, ensure that the modems are connected to a power source and are turned on. Check all modem cable connections.
4. Verify that POST is complete on the director (a minimum of 3 minutes).
5. Dial in to the telephone number that is assigned to the director, using a Telco system to dial in to the modems.
6. Observe the modem lamps. The Ring indicator will flash briefly as the telephone rings. If the Ring indicator does not flash on both units, check the incoming telephone lines to the modems again.
7. Verify that after one ring, the modem that is associated with the active CP blade (usually in slot 5), illuminates the Off Hook (OH) indicator on the modem and a login prompt is presented to the remote client.
8. Log in to the switch from the remote client as admin. The default password is **password**.

Cable: If the OH indicator illuminates on the standby CP blade modem, recheck the modem cable connection to the active CP blade.

9. Log out of the modem session.

10. Remove the Telco connector from the active CP blade modem, leaving the Telco line from the standby CP blade connected to the “Y” connector. See Figure 6-3 on page 148.

Cable: The modem session is disconnected automatically if the modem cable is detached while a session is active.

11. Dial in to the telephone number that is assigned to the director.
12. Observe the modem lamps. The Ring indicator will flash only on the modem that is connected to the standby CP blade.
13. Verify that after seven rings, the OH indicator on the standby CP blade modem is illuminated. A login prompt is presented to the remote client, and a message confirms that the standby CP blade is being logged in to. You can log in or disconnect the session, as desired.
14. Reconnect the Telco connector to the active CP blade modem. The director modems are ready for use.

6.2 SAN32B-3 implementation using EZSwitchSetup

The EZSwitchSetup starter kit greatly simplifies the setup and implementation of supported switches. The kit ships with the switch and includes a serial cable and a CD that contains the setup software. It makes the switch setup as simple as a “click-and-go” solution. It runs only in a single switch fabric.

If you follow the standard switch configuration practice, you implement a new switch by connecting a serial cable, setting up a tool such as Hyperterm to communicate, and implementing the `ipaddrset` command to configure the IP address. Then, you can then connect to the network using an Ethernet cable, using a Web browser to access Web Tools, or alternatively using Telnet to enter CLI mode and to configure the switch further. From here, you can set up zoning, assuming that all devices are connected and also that switch status monitoring uses Web Tools, SNMP, or an external application.

EZSwitchSetup greatly simplifies this process by walking you through all the steps automatically using a GUI-based interface.

In this section, we explain how to use EZSwitchSetup to configure a SAN32B-3 switch.

6.2.1 Implementing EZSwitchSetup

Before you begin, you need to obtain an IP address, subnet mask, and default gateway address for the switch. Then, follow these steps:

1. Using either a Windows system that is located physically close to the switch or a mobile computer, insert the CD, which starts the EZSwitchSetup program automatically, as shown in Figure 6-5. Click **OK** to start the installation.




Figure 6-5 EZSwitchSetup - startup panel

InstallAnywhere guides you through the simple installation. The EZSwitchSetup program runs automatically when the installation is complete.

2. Following the instructions, as shown in Figure 6-6 and Figure 6-7, connect the power cord to the switch, the Ethernet cable between the mobile computer and switch, and the serial cable from your mobile computer to the switch. Wait for the switch to turn on fully before continuing. It can take up to 3 minutes for the switch to be in a ready state with both the System Status and Power LEDs green.




Figure 6-6 EZSwitchSetup - Introduction




Figure 6-7 EZSwitchSetup - Connect Cables

3. Click **Next**, and the software starts a discovery by entering the switch WWN as shown in Figure 6-8. When the switch is found, you can move to the next panel by clicking **Next** again.

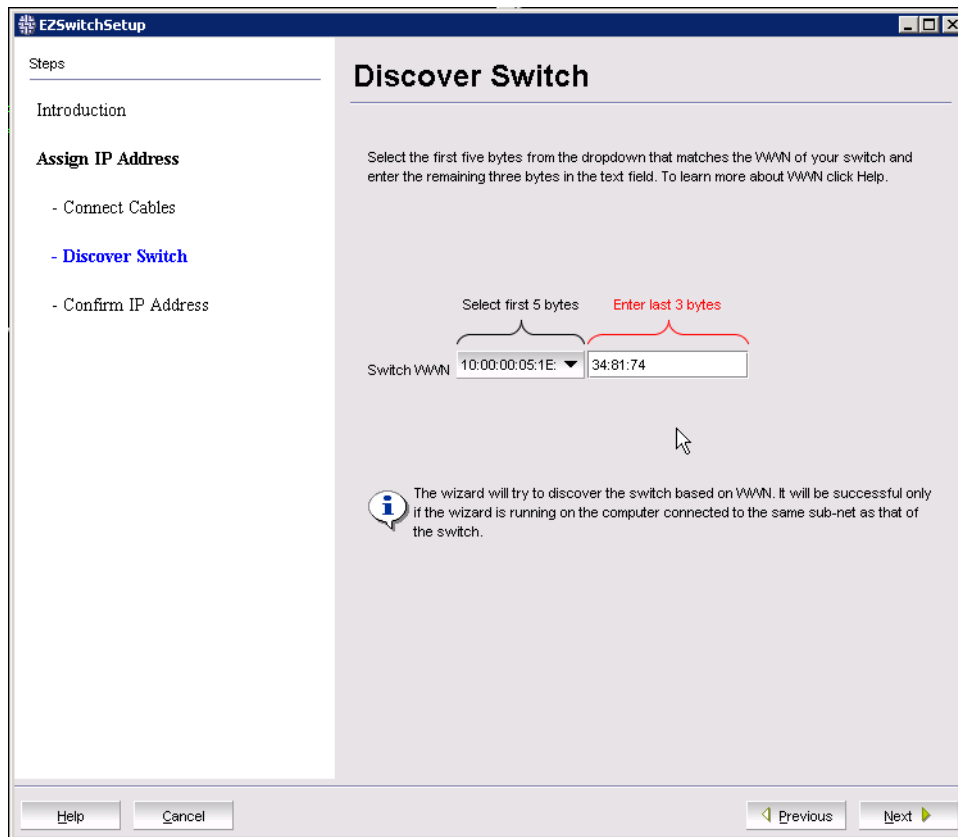


Figure 6-8 EZSwitchSetup - Discover Switch

4. EZSwitchSetup logs in to the switch using the admin ID. If, for whatever reason, the default password has changed, EZSwitchSetup prompts you for the new password.

5. At the next panel (shown in Figure 6-9), to accept the IP address, select the **No** radio button and click **Next**.

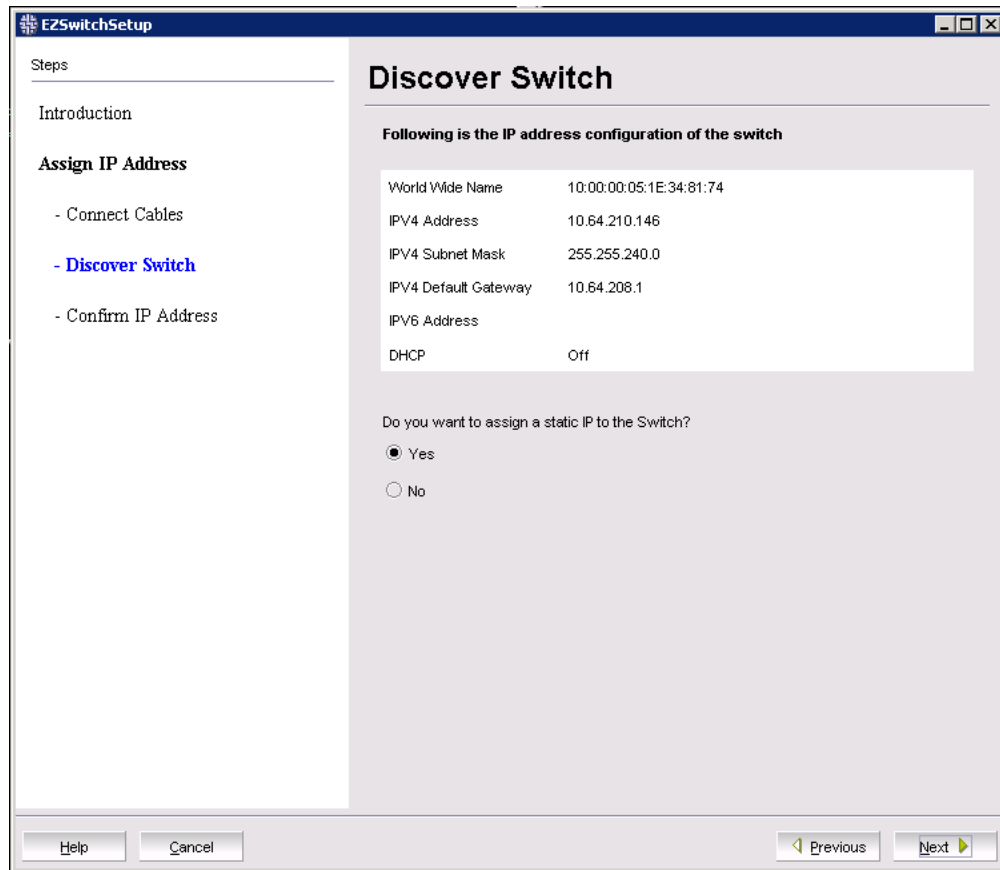


Figure 6-9 EZSwitchSetup - Set Switch IP Address

6. Now, the switch IP settings are complete, see Figure 6-10. Click **Continue** to spawn the second part of configuration, which is a Java plug-in for your existing browser. If you have a firewall enabled, you might have to permit access to the Internet in order to continue.

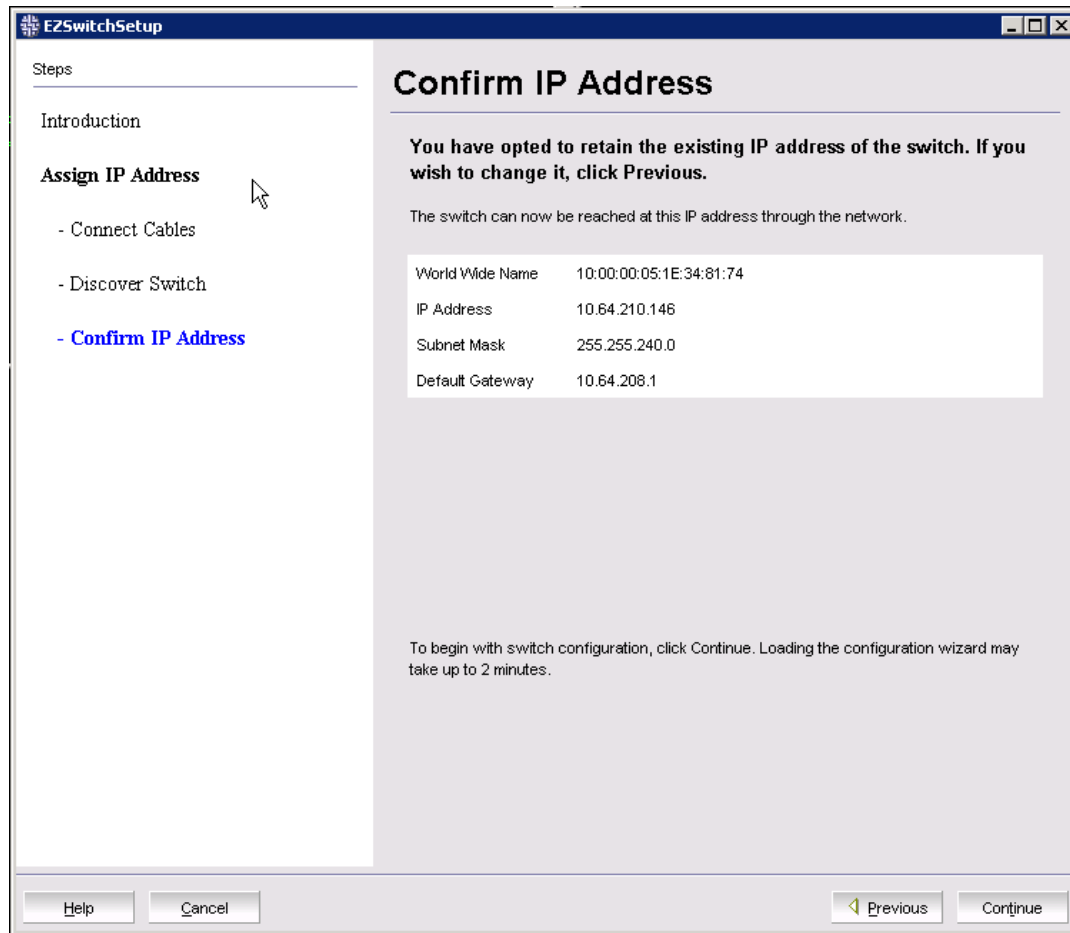


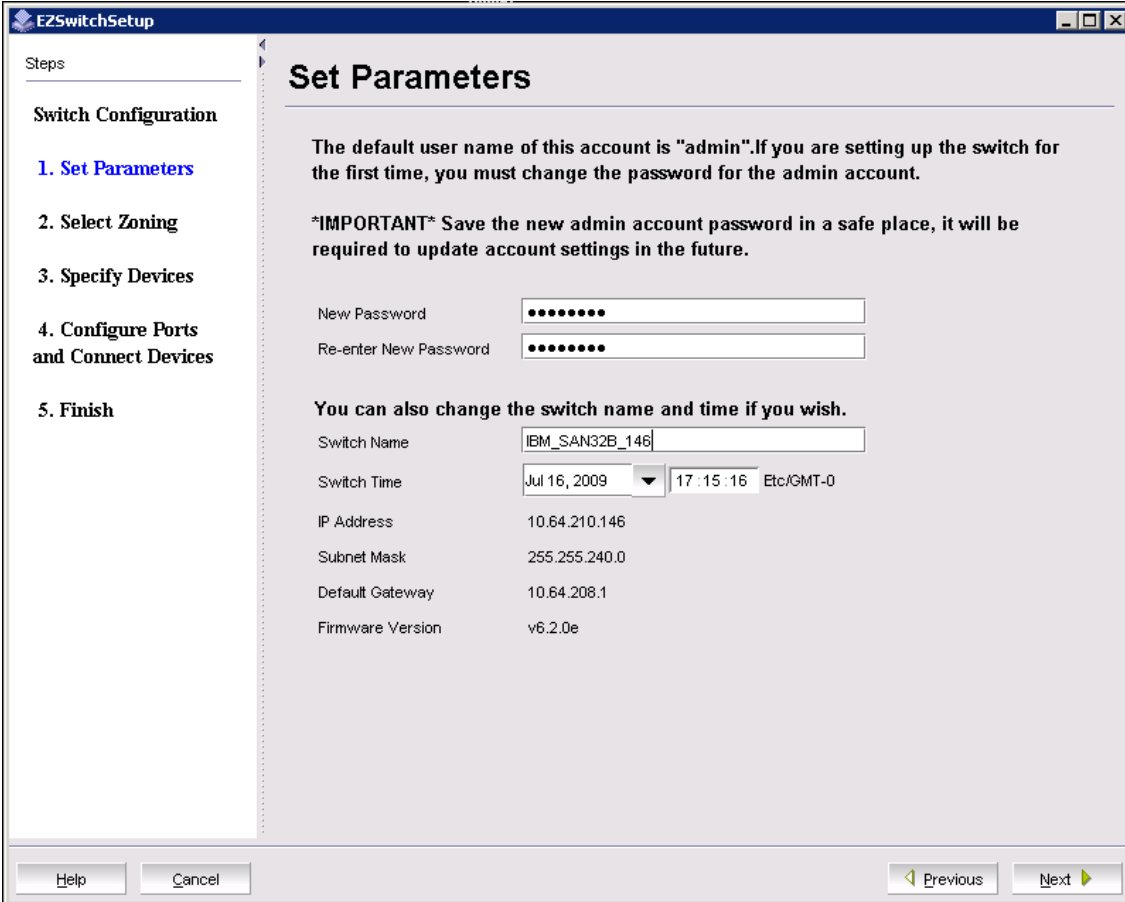
Figure 6-10 EZSwitchSetup - Confirm IP Address

7. This part of the configuration is controlled by the switch setup wizard. First, the Welcome window displays (Figure 6-11). Click **Next**.



Figure 6-11 EZSwitchSetup - Welcome window

8. Now, set the switch parameters, as shown in Figure 6-12. In this step, you can change the Admin password, the switch name, and the date and time.



The image shows a screenshot of the EZSwitchSetup application window, specifically the 'Set Parameters' step. The window has a title bar with the text 'EZSwitchSetup'. On the left side, there is a 'Steps' pane with a list of steps: '1. Set Parameters' (highlighted in blue), '2. Select Zoning', '3. Specify Devices', '4. Configure Ports and Connect Devices', and '5. Finish'. The main area of the window is titled 'Set Parameters'. It contains the following text: 'The default user name of this account is "admin". If you are setting up the switch for the first time, you must change the password for the admin account.' followed by '*IMPORTANT* Save the new admin account password in a safe place, it will be required to update account settings in the future.' Below this text are two password input fields: 'New Password' and 'Re-enter New Password', both showing masked characters (dots). Further down, there is a section titled 'You can also change the switch name and time if you wish.' followed by several configuration fields: 'Switch Name' (text box with 'IBM_SAN32B_146'), 'Switch Time' (date and time picker showing 'Jul 16, 2009' and '17:15:16 Etc/GMT-0'), 'IP Address' (text box with '10.64.210.146'), 'Subnet Mask' (text box with '255.255.240.0'), 'Default Gateway' (text box with '10.64.208.1'), and 'Firmware Version' (text box with 'v6.2.0e'). At the bottom of the window, there are buttons for 'Help', 'Cancel', 'Previous', and 'Next'.

Figure 6-12 EZSwitchSetup - Set Parameters

You can access EZSwitchSetup again later by entering **switchIP/EZsetup.html** as the address field in a Web browser or, alternatively, by selecting the setup option from the switch manager.

9. After you set up these values, click **Next** to open the zoning configuration panel, as shown in Figure 6-13.

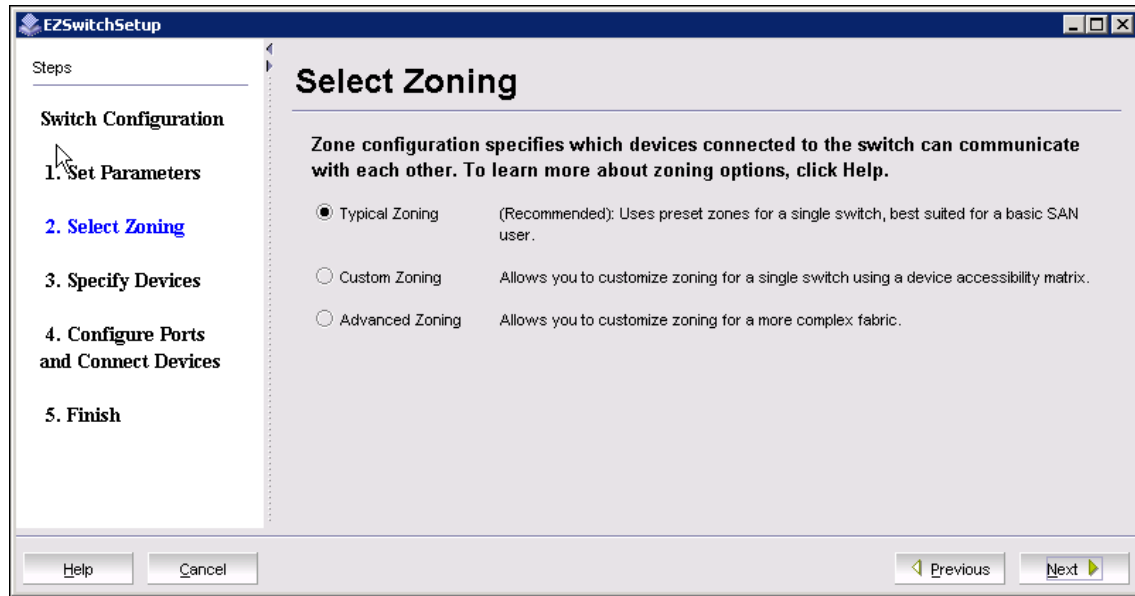


Figure 6-13 EZSwitchSetup - Select Zoning

Here you are presented with three options:

- Typical Zoning
- Custom Zoning
- Advanced Zoning

In our example, we select **Typical Zoning** and then click **Next**.

10. Figure 6-14 shows the next step. In typical zoning scenario, we simply need to specify the number of switch ports for HBA and for storage connections.

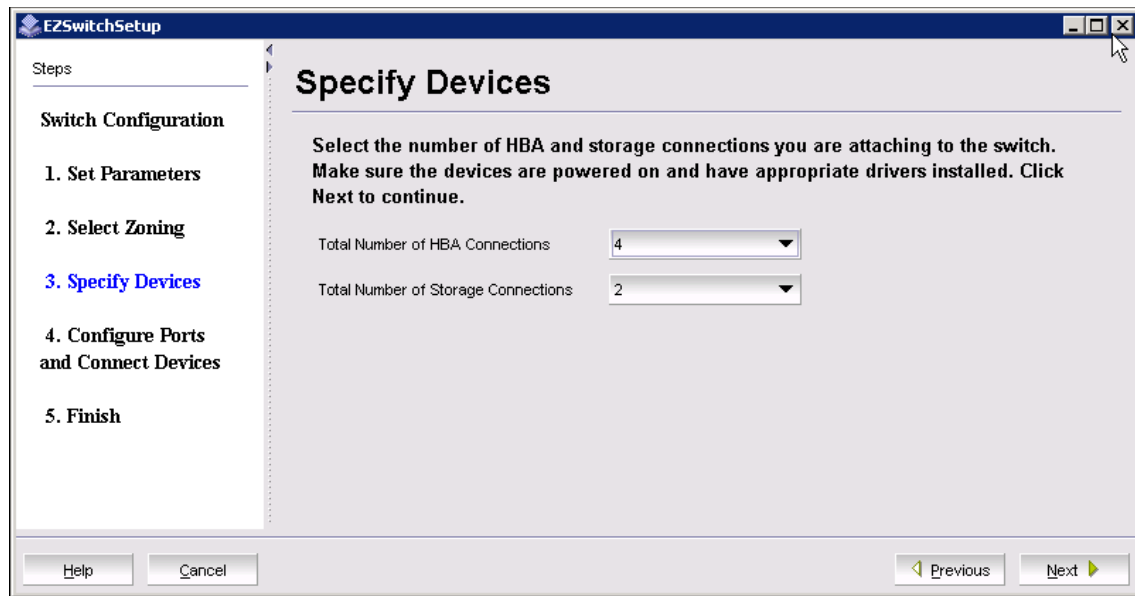


Figure 6-14 EZSwitchSetup - Specify Devices

11. EZSwitchSetup wizard then displays the device connection window (see Figure 6-15), which suggests the ports that you need to use for the requested connections. Here, you need to connect the hosts and storage physically as suggested by the software.

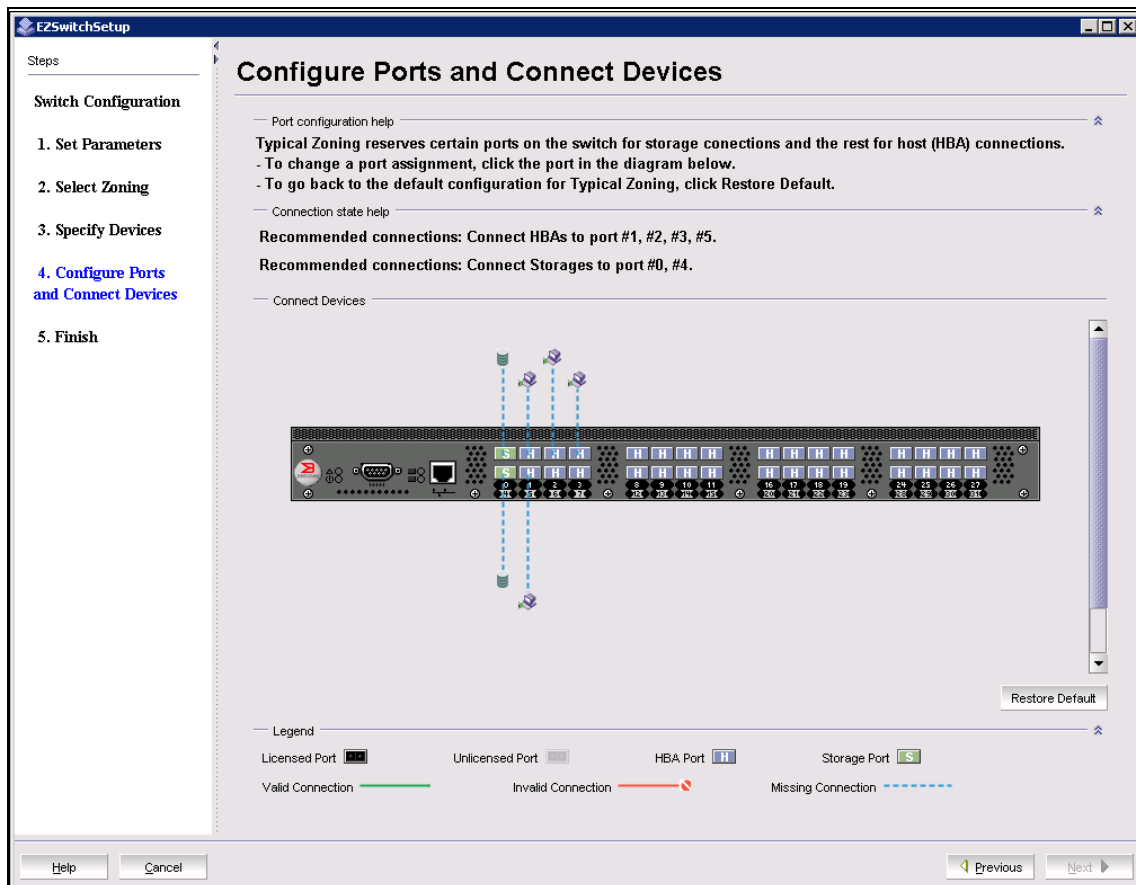


Figure 6-15 EZSwitchSetup - Configure Ports and Connect Devices

Because we have no devices connected to the switch yet, Figure 6-16 shows this situation. Valid connections would be shown with a solid green line. Invalid connections would be highlighted with the color red.

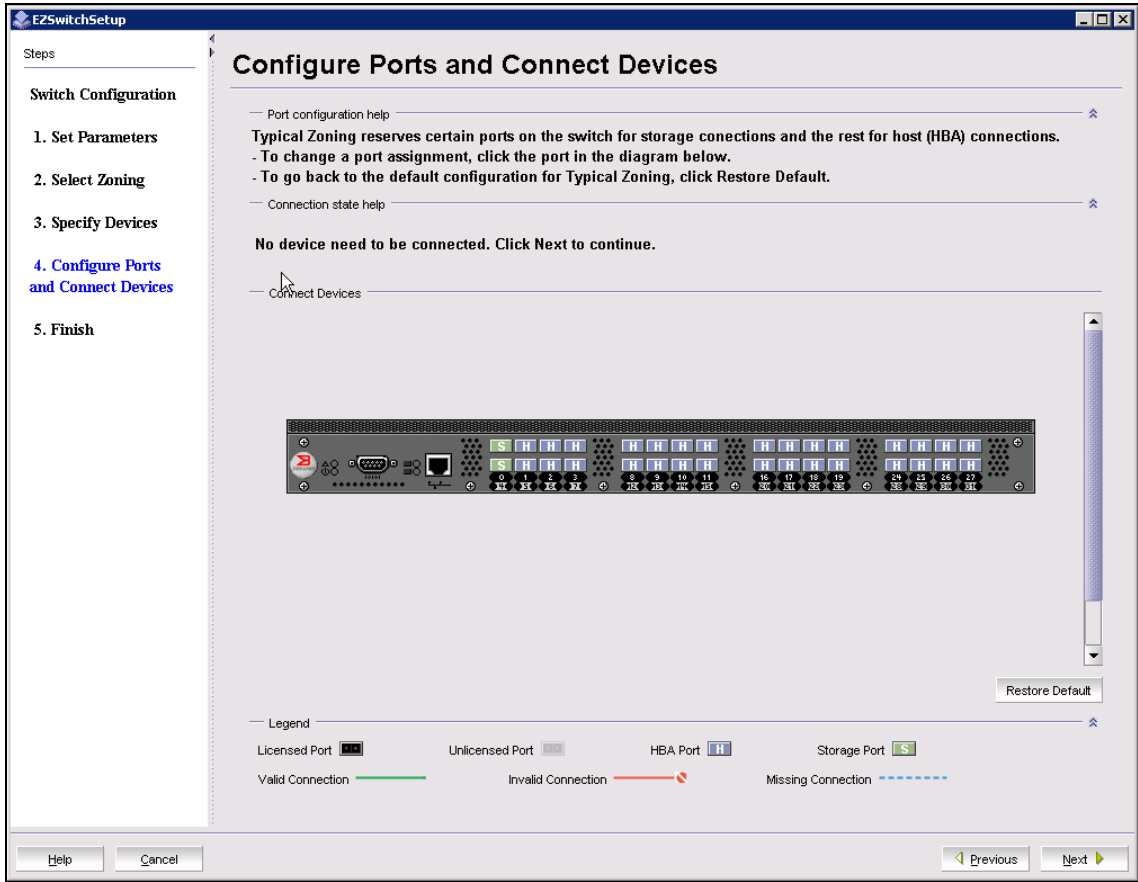


Figure 6-16 EZSwitchSetup - showing that all devices are connected

12. The final setup panel (shown in Figure 6-17) displays a summary of the switch configuration. At this point, the process is complete.




Figure 6-17 EZSwitchSetup - Finish

6.2.2 Using Switch Manager to manage a switch

If you have set up a switch with EZSwitchSetup and if it is not connected to other switches, then you use the Switch Manager utility to manage it (instead of Web Tools). Figure 6-18 shows an example of the Switch Manager user interface.

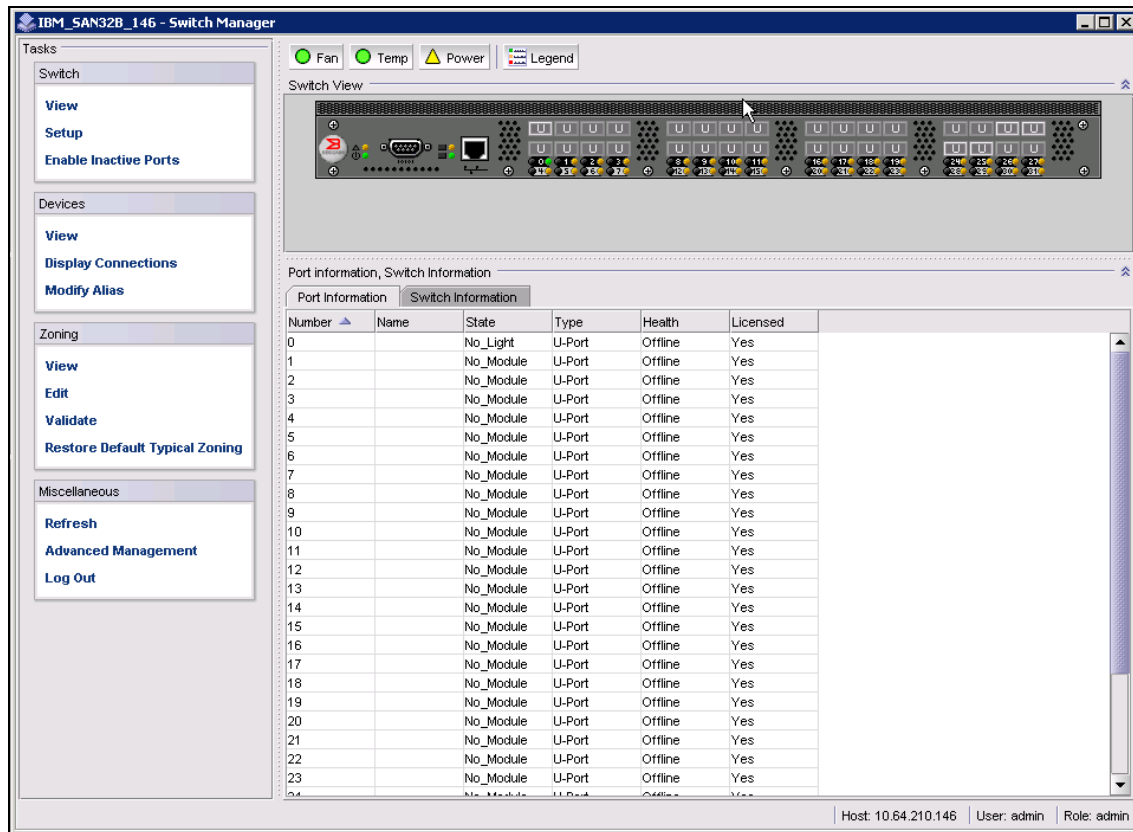


Figure 6-18 Switch Manager user interface

In comparison with Web Tools, the Switch Manager is a simple utility that is designed to manage stand-alone SAN switches. You launch the Switch Manager the same way that you launch Web Tools by specifying the switch IP address in the Web browser address field.

6.2.3 Basic troubleshooting with EZSwitchSetup

If reinstalling or upgrading EZSwitchSetup fails, you need to uninstall the previous version first, and then reinstall.

If EZSwitchSetup encounters a launch problem, check whether there is already a copy of EZSwitchSetup running on another user's system. Only one copy of the program is allowed to run at any given time.

If during the EZSwitchSetup process, you encounter an operation failure, check the serial and Ethernet connection and fix it if necessary, and then launch EZSwitchSetup again.

EZSwitchSetup does not fully recognize storage that is presented to the SAN in initiator and target mode. This issue can occur, for example, if a DS4000® with remote mirroring enabled is connected. Usually, a simple SAN will not involve this type of configuration, and as such it is unlikely that you will experience this issue. However, if it occurs, you can circumvent the issue by adding only the hosts in the initial setup. You can then add the storage when you have proceeded past the Switch Setup Complete window.

After you add the switch to a fabric, you can no longer access the EZSwitchSetup wizard, as shown in the error message in Figure 6-19.




Figure 6-19 Web Tools EZ error message



License administration

In this chapter, we discuss the licensed features that are available on IBM/Brocade SAN products. We also describe how to use Web Tools and CLI for license administration (as explained in 8.7.6, “License tab” on page 268).

7.1 Licensed features

IBM/Brocade SAN products normally ship with the basic set of features enabled. Numerous advanced features are available as add-ons and can be enabled by applying appropriate license keys. This ability to add-on features allows “pay-as-you-grow” flexibility so that customers with basic requirements do not have to pay for features that they do not need, while customers who need additional functionality can purchase the exact set of features they require.

Fabric OS v6.1.0 introduces some changes to the licensing scheme. Web Tools and zoning are no longer licensed features. They are now part of the basic Fabric OS, and there are a number of new licenses, such as Integrated Routing, Inter-Chassis Link (ICL), and Adaptive Networking and for the IBM Converged Switch B32 Frame-Based ISL Trunking.

7.1.1 Ports on Demand

Storage area network (SAN) switches typically ship with lower than the maximum number of ports enabled. You can enable additional Ports on Demand in increments of 8 or 16 ports.

Additional Ports On Demand

For example, the SAN24B-4 ships with eight ports activated (ports 0 through 7). You can enable additional eight ports (8 through 15) by purchasing the Ports on Demand (PoD) license. You can enable another eight ports (16 through 23) with an additional PoD license.

The SAN40B-4 is delivered with 24 enabled ports. The first increment of eight ports brings the number of enabled ports to 32, and the second increment allows you to use the total number of 40 ports.

Finally, the SAN80B-4 ships with 48 activated ports. You can increase the number of ports to 64 ports with the first PoD license and to a total of 80 ports with a second PoD license.

On the previous generation of IBM/Brocade SAN products, a simple rule of thumb was that switches with the PoD capability are shipped with 50% of the ports configured and that PoD upgrades are available in 25% increments. For example:

- ▶ SAN16B (16 physical ports) ships with eight ports activated initially. A PoD license is available in increments of four ports.
- ▶ SAN32B (32 physical ports) ships with 16 ports activated and, with a PoD license, can be upgraded in 8-port blocks.
- ▶ SAN64B (64 physical ports) ships with 32 ports activated and, with a PoD license, can be upgraded in 16-port blocks.

Enabling these licenses using Web Tools or the CLI (using **licenseAdd** and **portEnable**) are both non-disruptive. If you remove a PoD license by mistake, the affected ports continue to operate until the switch is disabled or rebooted.

Dynamic Ports On Demand

You can use Dynamic Ports on Demand on Brocade SAN switch modules for the BladeCenter. If some Blade servers do not have the Fibre Channel expansion card installed, then the corresponding port on the SAN switch module is never needed. So, you only need to enable a subset of all the switch module ports. With Dynamic Ports on Demand, you can use a particular subset of ports, according to which Blade servers actually require SAN connectivity. You can create the subset dynamically (ports are assigned as they come online, until the number of licensed ports is reached) or statically (you specify a PoD ports subset that is used for ports enablement).

7.1.2 Full Fabric

Base switches act as single-switch fabrics without E_Port capability. Some older switches had to upgrade with the Full Fabric license to get this function. When this license is applied, the switch can connect to other switches through E_Ports to join a multi-domain fabric. With newer products this function is already included in the base switch.

7.1.3 8 Gbps

This license is a pre-installed license. You can see the license using **licenseshow** on the 8 Gbps platforms, and it must never be removed. This license is for informational purposes only.

7.1.4 Inter-Chassis Link (ICL)

Two SAN768B chassis, two SAN384B (or one SAN768B and one SAN384B) can be interconnected using four ICL cables.

The ICL license enables ICL ports and is, therefore, a mandatory requirement before any ICL connections can be made. You must enable this license on both SAN768B or SAN384B chassis.

7.1.5 Adaptive Networking

Use the Adaptive Networking feature to introduce SAN traffic control in congested environments. The feature is available on 8 Gbps capable switches and provides the following features:

- ▶ Quality of Service (QoS) SID/DID prioritization:

This technique utilizes virtual channels (VC) across an ISL connection. The VCs are assigned high, medium, or low priority. High priority VCs get 60% of bandwidth, medium priority VCs get 30%, and low priority VCs get 10%.

SID/DID traffic priority assignment is done by means of zoning. Special zone name prefixes are used for assigning source and destination device worldwide names (WWNs):

- QOSH_ for high priority traffic
- QOSL_ for low priority traffic

Medium priority is the default. The SID/DID pairs not assigned into any of the QoS zones will be set to this priority level.

- ▶ Ingress rate limiting:

You can use this feature to throttle down the ingress port speeds, which can be useful in case of congested ISLs. Only the F and FL_ports support ingress rate limiting. In addition, this feature can only be used on 8 Gbps capable ports. Ingress rate limiting is unidirectional. It limits only the data transfer rate from the device to the switch port. The transfer rate from the switch port to the device is not affected.

7.1.6 Frame Based ISL Trunking

This feature enables the IBM Converged Switch B32 to have up to 8 ports between a pair of switches to be combined into a logical ISL with speeds of up to 64 Gbps (128 Gbps full duplex) for optimal bandwidth utilization and load balancing, and exchange-based load balancing across ISLs with DPS (included in Fabric OS).

7.1.7 Fabric Watch

Fabric Watch provides real-time monitoring of switch health, performance, and security. The information it provides enables the SAN administrator to act proactively and, therefore, to avoid unnecessary downtime. Numerous operational parameters of the switches in the fabric are tracked, and automatic alerting takes place whenever switches operate outside acceptable thresholds.

Fabric OS v6.1 also introduces Port Fencing in Fabric Watch, which disables automatically a port that operates outside of the defined thresholds.

7.1.8 Advanced Performance Monitoring

This licensed feature provides a comprehensive set of technologies for tracking performance and bandwidth usage in the SAN. It is based on Brocade Frame Filtering technology and performance counter engine. You can use the information that is gathered for performance bottleneck identification and capacity planning.

The Advanced Performance Monitoring feature provides these capabilities:

- ▶ End-to-end monitors that you can use to monitor the traffic between host/target pairs. Host and target are specified with Source ID (SID) and Destination ID (DID). This performance monitor counts the number of received (RX_COUNT) and transmitted (TX_COUNT) words in frames.
- ▶ Filter-based monitors can measure the amount of particular subset of total traffic. You use filters to specify the traffic type in which you are interested. There is a set of standard filters, such as SCSI Read commands, SCSI Write commands, SCSI traffic frames, and IP traffic frames. In addition, you can set up custom filters to gather the specific statistics suitable to your needs.
- ▶ ISL monitors measure the traffic across an inter-switch link (ISL) to all reachable destination domains. Then, you can see easily which destination domain consumes the highest amount of bandwidth.
- ▶ Top Talkers monitoring is a feature that was introduced in Fabric OS v6.0.0. These monitors can identify the SID/DID pairs that consume the most bandwidth. You can use the Top Talkers data to re-route the traffic to less busy ports, in case the original ports are becoming overloaded. You can also use the Top Talkers information to identify the SID/DID pairs that might require higher priority in order to have the adequate Quality of Service (QoS).

Attention: ISL monitors cannot be used when Top Talkers monitoring is installed. In addition, the new 8 Gbps products (SAN24B-4, SAN40B-4, SAN80B-4, SAN384B, SAN768B and 3758-B32) do not support ISL monitors.

7.1.9 Extended Fabrics

Normally, Fibre Channel standard allows up to 10 km Fibre Channel cable lengths in the fabric. The Extended Fabrics licensed feature provides SAN fabric extension to much longer distances. By optimally utilizing the internal switch buffers, it is possible to achieve ISLs at distances up to 500 km. Check with the vendor as to the exact distances that are supported.

7.1.10 ISL Trunking

You use this feature to enable configuration of multiple ISL connections into trunks, providing a much increased bandwidth level. You can establish the ISL trunk between any two supported 2, 4, or 8 Gbps switches, provided that both switches have this license installed. If an 8 Gbps capable switch is connected to a 2 or 4 Gbps switch, the ISL trunk operates at the lower speed. With 8 Gbps capable switches on both ends, you can establish trunks with up to eight ports, and the maximum bandwidth is up to 64 Gbps.

7.1.11 Integrated Routing

You can install this licensed feature on the new 8 Gbps capable products SAN384B, SAN768B, SAN80B-4, and SAN40B-4. It adds Fibre Channel Routing (FCR) support. You can configure the 8 Gbps ports as native EX_ports, thus you do not need to add a routing blade or use the SAN18B-R if you want to set up FCR. Additionally, this feature provides twice the bandwidth of existing FCR solutions when using 8 Gbps ports on both sides.

In the case of a SAN768B or SAN384B with a routing blade, you cannot use both the native EX_ports (available through the Integrated Routing licensed feature) and EX_ports on the routing blade. Only the VEX_ports can be used in conjunction with Integrated Routing EX_ports.

7.1.12 High Performance Extension over FCIP/FC

This feature contains the capabilities formerly included in FCIP Services license. It also provides additional features, such as FC-Fastwrite and IPSec. The license can be applied to the products supporting FCIP:

- ▶ SAN04B-R
- ▶ SAN18B-R
- ▶ SAN06B-R
- ▶ FCR Blade (FR4-18i)
- ▶ FCR Blade (FX8-24)

7.1.13 FICON Management Server

The FICON Management Server license is required to set up FICON CUP function. IBM/Brocade switches that support FICON CUP can appear as control units to IBM System z servers. FICON CUP is the protocol used by the management software on IBM System z to perform in-band management of the switches.

7.2 Using Web Tools to administer licenses

Web Tools provides all the functions that are required to administer licensed features. You can display a list of installed licenses, install additional licenses, and remove those licenses that are no longer required. In this section, we show only some basic tasks, such as verification of licensed ports, and a list of installed licenses.

7.2.1 Using Web Tools to check licensed ports

The Web Tools Port Administration window contains a column to tell whether a port is licensed. To display the Port Administration window, click **Port Admin**, as shown in Figure 7-1.

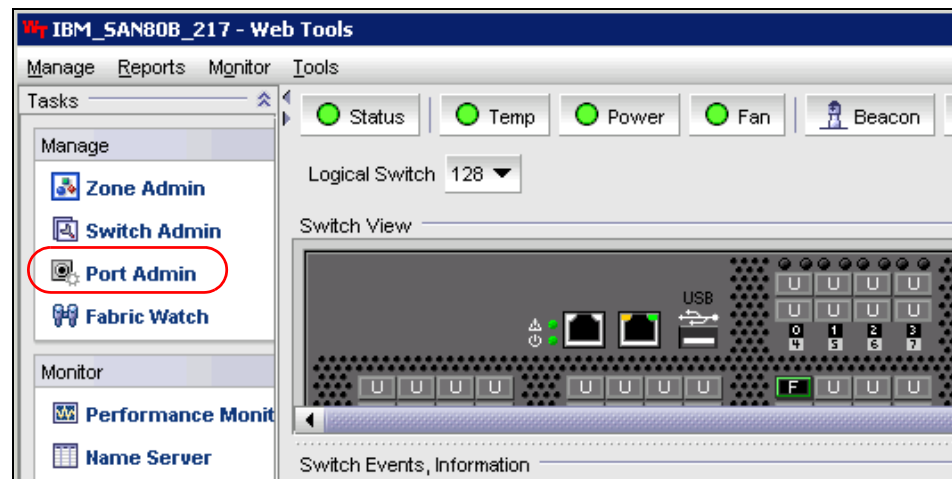


Figure 7-1 Web Tools: Launching the Port Administration window

The column that shows the licensed ports displays in advanced mode (see Figure 7-2).

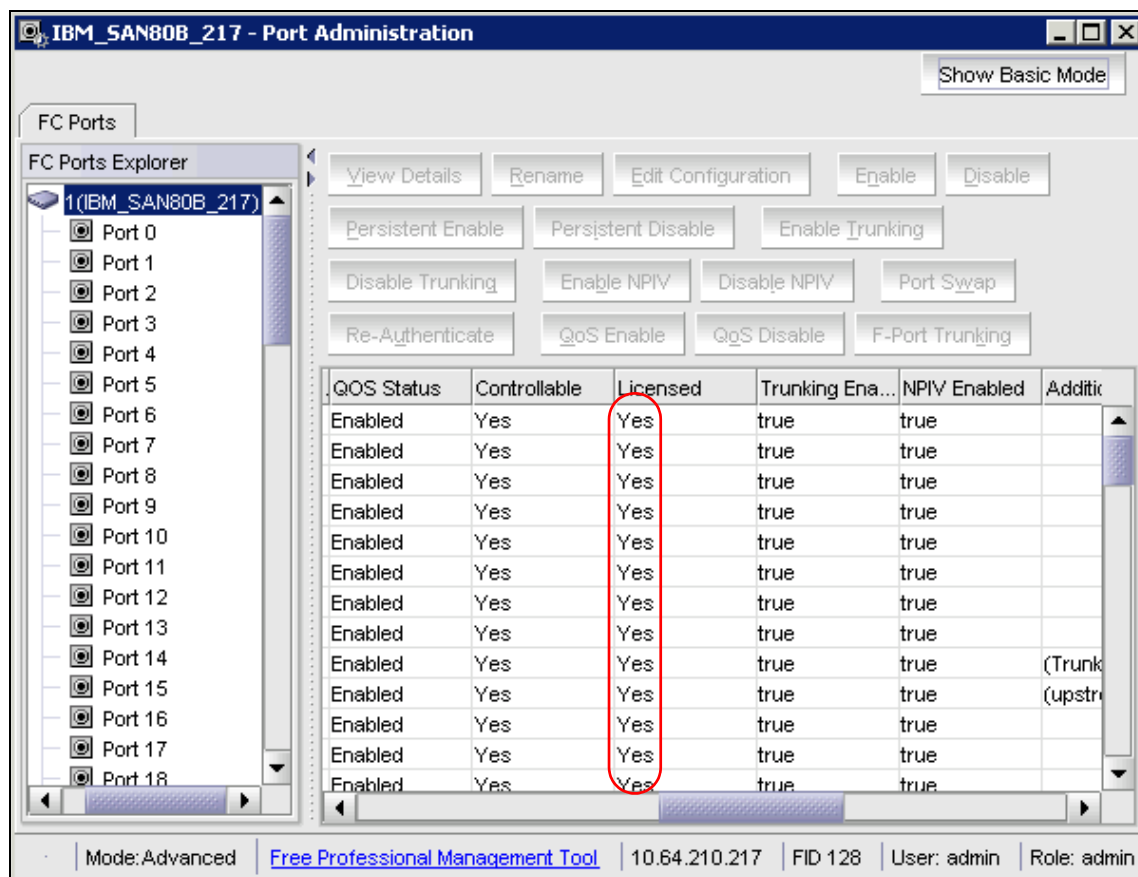


Figure 7-2 Port Administration window showing licensed ports

7.2.2 Installed licenses

To see the licenses that are installed on the switch, open the Switch Administration window by clicking **Switch Admin** task, as shown in Figure 7-3.

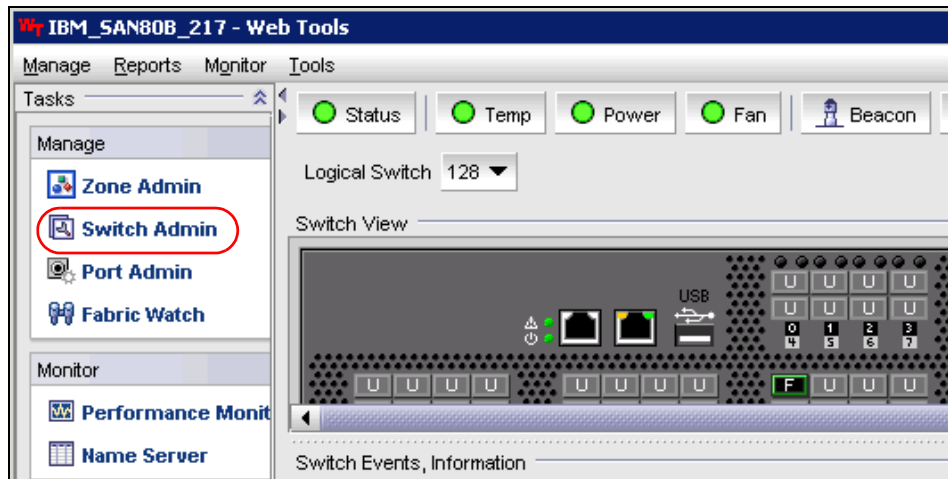


Figure 7-3 Web Tools - launching Switch Administration window

When the Switch Administration window displays, click the **License** tab, as shown in Figure 7-4.

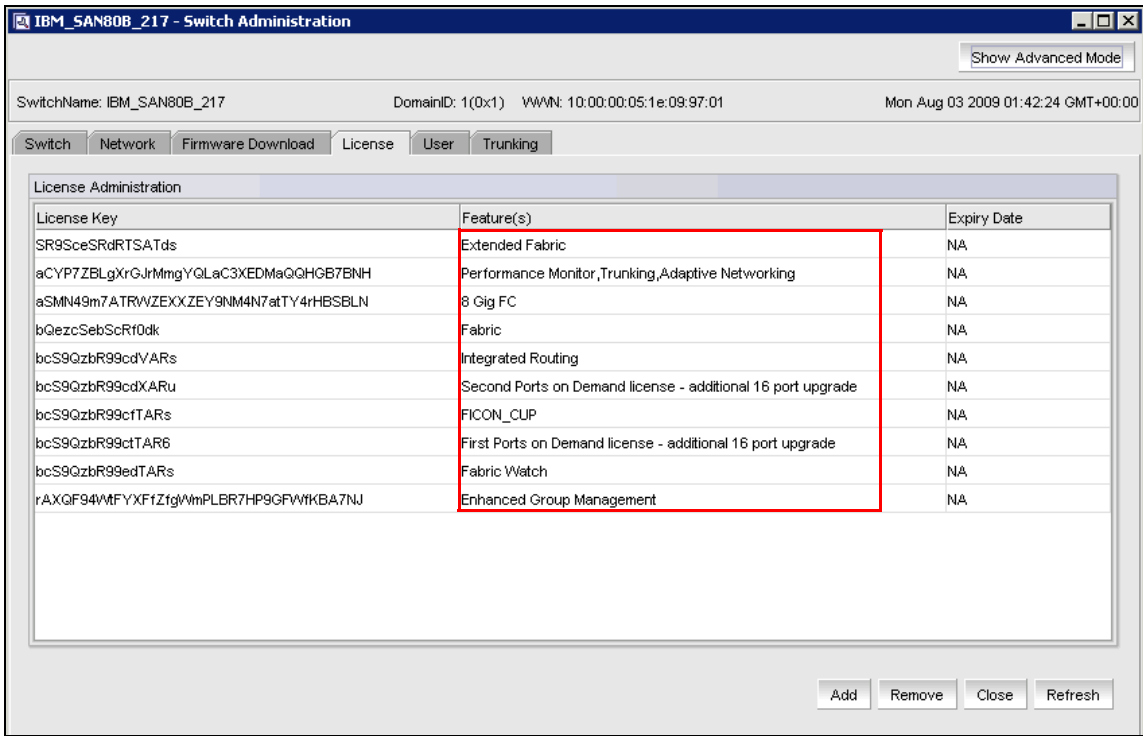


Figure 7-4 Installed licenses on this switch

In our example, we see the list of licenses installed on our switch.

7.3 Tips on solving licensing issues

Various things can go wrong with licensing keys. Switch or mainboard replacement, key mismatches, and typographical errors are the most common issues.

There are some simple steps you can take to solve issues with licensing keys:

1. A Switch/Mainboard was replaced, then the new Switch/Mainboard is missing licenses because licenses are bound to a switch's WWN.

Note the old and the new Switch/Mainboard WWN and use the following website to generate the keys:

<https://www-912.ibm.com/FruLicenseRequestClient/>

If this does not work, you will have to ask a Brocade authorized representative to open a ticket at Brocade.

Brocade will then transfer the license key from the old WWN to the new WWN.

2. You have received a transaction key and you made a typographical error when generating the license key.

Note the correct WWN and the incorrect WWN and ask a Brocade representative to open a ticket at Brocade.

Brocade will transfer the license key from the incorrect WWN to the correct one.

3. You received a transaction key and successfully generated the license key but lost the license key before you activated it on the switch.

Note the switch type and WWN and ask a Brocade representative to open a ticket at Brocade.

Brocade can check their database and see which licenses were activated for this switch WWN.



Web Tools

In this chapter, we discuss the features of Web Tools in greater detail. Although the Web Tools display has changed over time, the tools still have the same basic look and feel to them as they had in previous versions.

8.1 Web Tools walk-through

In this chapter, we describe the features of Web Tools in detail. However, before we begin that discussion, in this section we provide some basic information about getting started with Web Tools.

8.1.1 Web Tools, the EGM license, and DCFM

Beginning with Fabric OS version 6.1.1, Web Tools functionality is tiered and integrated with DCFM. If you are migrating from a Web Tools release prior to Fabric OS version 6.1.1, this might impact how you use Web Tools.

A Web Tools license is not required, and a basic version of Web Tools is available for free. Additional functionality can be added by obtaining the Enhanced Group Management (EGM) license. Table 8-1 compares Basic Web Tools features to Web Tools with the EGM license. The EGM license is only for 8 Gbps platforms, such as the IBM SAN768B, SAN384B enterprise-class platforms, and the IBM SAN80B-4, SAN40B-4 and SAN24B-4 switches. For non-8 Gbps platforms, all functions are available without the EGM license.

Table 8-1 Web Tools features enabled by the EGM license

Feature	Basic Web Tools	Web Tools with EGM license
Active Directory Support	Yes	Yes
AD Context Switching	No	Yes
AD Filtered Views	Yes	Yes
Admin Domain Management	No	Yes
AG Management	Yes	Yes
Analyze Zone Config	No	No
Basic Zoning and TI Zoning	Yes	Yes
Blade Management	Yes	Yes
Cloning a Zone	No	Yes
Config Upload/Download	Yes	Yes
Convenience Function from Tools Menu	No	No

Feature	Basic Web Tools	Web Tools with EGM license
Device Accessibility Matrix	No	No
Easy to configure iSCSI Wizard	Yes	Yes
Extended Fabric Management	No	Yes
F_Port Trunk Management	No	Yes
Fabric Events	No	No
Fabric Summary	No	No
Fabric Tree	Yes	Yes
FCIP Tunnel Configuration	No	No
FCIP Tunnel Display	Yes	Yes
FCR Management	Yes	Yes
FCR Port Configuration	Yes	Yes
FICON CUP Tab	No	Yes
FRU Monitoring	Yes	Yes
High Availability	Yes	Yes
IP Sec Policies	Yes	Yes
ISL Trunk Management	No	Yes
ISL Trunking Information	Yes	Yes
License Management	Yes	Yes
Long Distance	No	Yes
Logical Switch Context Switching	No	Yes
PDCM Matrix	No	Yes
Port Administration	Yes	Yes
Print Zone Database Summary	No	No
RBAC	Yes	Yes

Feature	Basic Web Tools	Web Tools with EGM license
Routing and DLS Configuration	No	Yes
Security Policies Tab (such as ACL)	Yes	Yes
Switch Info Tab	Yes	Yes
Switch Status	Yes	Yes
Switch View right-click options	Yes	Yes
Trace Dump	Yes	Yes
USB Management	Yes	Yes
User Management	Yes	Yes
Verify and troubleshoot accessibility between devices	Yes	Yes

Also beginning with Fabric OS version 6.1.1, some Web Tools capabilities are moved from Web Tools to DCFM. Table 8-2 summarizes these changes.

The functionality that was moved from Web Tools into DCFM is applicable to both DCFM Professional and DCFM Enterprise.

Table 8-2 Web Tools functionality moved to DCFM

Function	Web Tools 6.1.0	DCFM	Comments
Add Un-Zoned Devices	Zone Admin	Configure → Zoning Reverse Find in the Zoning dialog provides the view of the zoned and unzoned devices in the fabric if all zone members are selected for Find.	

Function	Web Tools 6.1.0	DCFM	Comments
Analyze Zone Config	Zone Admin	<ol style="list-style-type: none"> 1. Configure → Zoning Reverse Find in the Zoning dialog provides the view of the zoned and unzoned devices in the fabric if all zone members are selected for Find. 2. Device Tree and Topology: Connected End Devices - Custom Display from the top level in the main frame provides the device tree and topology view for all the zoned devices if all zones are selected in the active zone configuration. 	
Define Device Alias	Zone Admin	Configure → Zoning	
Device Accessibility Matrix	Zone Admin	Configure → Zoning The Compare dialog provides the Storage-Host and Host-Storage view in a tree representation that is comparable to the Device Accessibility Matrix when all devices are selected.	

Function	Web Tools 6.1.0	DCFM	Comments
Fabric Events	Monitor → Fabric Events	Monitor → Logs → Events?	
Fabric Summary	Reports → Fabric Summary	Monitor → Reports → Fabric Summary Report?	
FCIP Tunnel Configuration	Port Admin Module; GigE Tab	Configure → FCIP Tunnel	Viewing FCIP tunnels is still supported in Web Tools 6.1.1, but New, Edit Config, and delete are only available in DCFM.
GigE Ports Interface	Port Admin Module; GigE Tab	Configure → FCIP Tunnel	
GigE Ports Route	Port Admin Module; GigE Tab	Configure → FCIP Tunnel	
Non-local switch ports display in zoning tree	Zone Admin Admin Domain Switch Admin → DCC Policies Performance Monitoring	Configure → Zoning	In Web Tools, non-local switch port id/WWN can be added using text box.
Remove Offline or Inaccessible Devices	Zone Admin	Configure → Zoning Replace/Replace All zone members by selecting the offline devices from the zone tree. Offline devices have an unknown overlay badge with good visibility.	
Zone database summary print	Zone Admin	Configure → Zoning Zoning report for both online and offline database.	

8.1.2 System requirements

Web Tools requires that your browser conform to HTML v4.0, JavaScript v1.0, and Java plug-in v1.6.0_16 or higher.

Brocade has certified and tested Web Tools on the platforms shown in Table 8-3.

Table 8-3 Supported Operating Systems and Browsers

Operating system	Browser
Windows 7	Internet Explorer 7.0/8.0
Windows Server 2008 Standard	Internet Explorer 7.0/8.0
Windows Vista Business	Internet Explorer 7.0/8.0
RedHat Enterprise Server 5 Advanced Platform	Firefox 2.0
SUSE Linux Enterprise Server 10	Firefox 2.0/3.0
Linux Red Hat AS 3.0	Firefox 2.0
Linux Red Hat AS 4.0	Firefox 2.0/3.0
Windows 2000	Firefox 2.0, Internet Explorer 6.0
Windows 2003 Server, SP2	Firefox 2.0/3.0, Internet Explorer 7.0/8.0
Windows XP, SP3	Firefox 2.0/3.0, Internet Explorer 7.0/8.0
Windows XP, SP2	Firefox 2.0, Internet Explorer 7.0/8.0
Solaris 10 (SPARC only)	Firefox 2.0
Solaris 9 (SPARC only)	Firefox 2.0

For Windows systems, a minimum of 256 MB of RAM for fabrics comprising up to 15 switches, 512 MB of RAM for fabrics comprising more than 15 switches, and a minimum of 8 MB of video RAM are preferable. A DCX with a fully populated FC8-64 blade requires a minimum of 512MB RAM.

Setting the refresh frequency for Internet Explorer

Correct operation of Web Tools with Internet Explorer requires specifying the appropriate settings for browser refresh frequency and process model. Refresh the browser pages frequently to ensure the correct operation of Web Tools:

1. Click **Tools** → **Internet Options** in the browser.
2. Click the **General** tab and click **Settings** under “Temporary Internet Files.”
3. Click **Every time I visit the webpage** under “Check for newer versions of stored pages,” as shown in Figure 8-1.

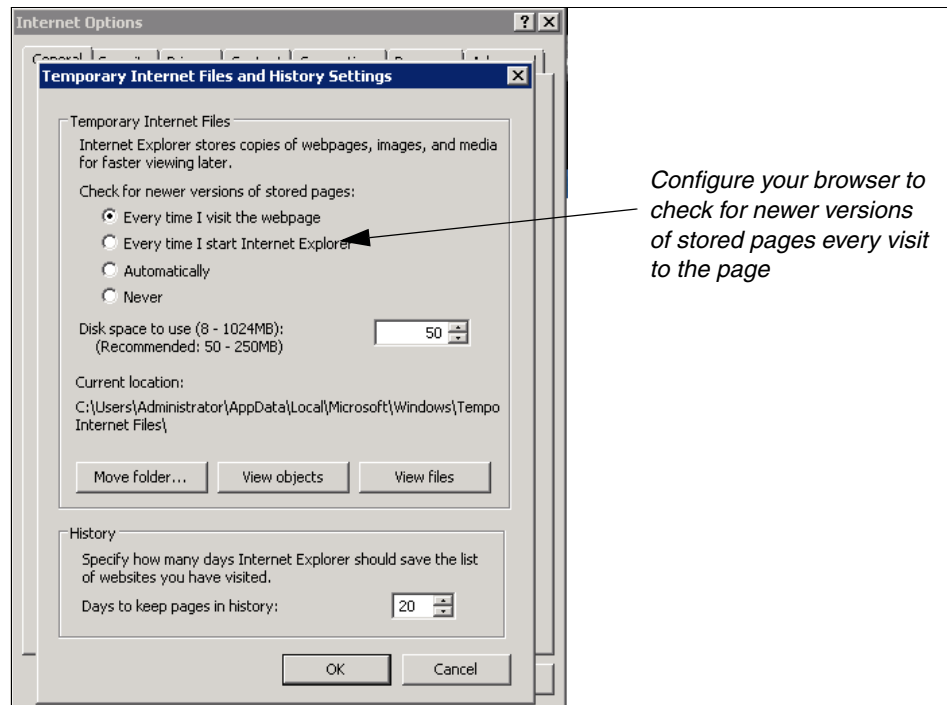


Figure 8-1 Configuring Internet Explorer

Deleting temporary Internet files used by Java applications

For Web Tools to operate correctly, you must delete the temporary Internet files used by Java applications:

1. From the Control Panel, open Java.
2. Click the **General** tab and click **Settings**.
3. Click **Delete Files** to remove the temporary files used by Java applications (see Figure 8-2).

4. Click **OK** on the confirmation dialog box.
You can clear the Trace and Log files check box if you want to keep those files.
5. Click **OK**.
6. On the Java Control Panel, click **View** to review the files that are in the Java cache. (If you have deleted all the temporary files, the list is empty).

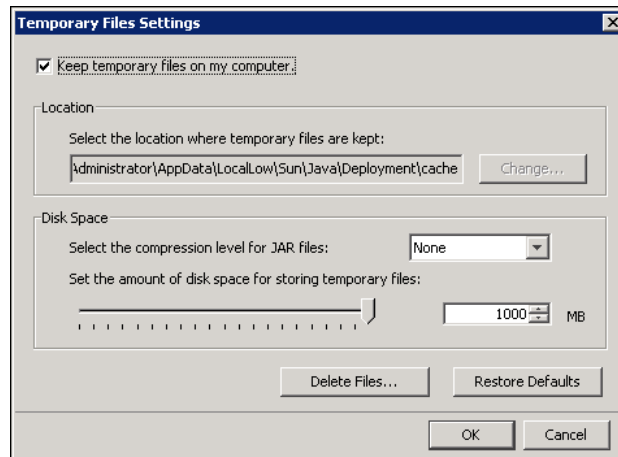


Figure 8-2 Temporary Internet Settings dialog box

8.1.3 Java installation on the workstation

The Java Plug-in must be installed on the workstation. If you attempt to open Web Tools without any Java Plug-in installed:

- ▶ Internet Explorer automatically prompts and downloads the proper Java Plug-in.
- ▶ Firefox downloads the most recently released Java Plug-in.

If you attempt to open Web Tools with an earlier version Java Plug-in installed:

- ▶ Internet Explorer might prompt for an upgrade, depending on the existing Java Plug-in version.
- ▶ Firefox uses the existing Java Plug-in.

Installing the JRE on your Linux client workstation

Follow these steps:

1. Locate the JRE on the Internet, at the following URL:
http://java.sun.com/products/archive/j2se/5.0_13/index.html

Attention: This URL points to a non-IBM website and is subject to change without notice.

2. Select JRE 5.0 Update 13.
3. Follow the instructions to install the JRE.
4. Create a symbolic link: From this location:
\$FIREFOX/plugins/libjavaplugin_oji.so
To this location:
\$JRE/plugin/\$ARCH/ns600/libjavaplugin_oji.so

Installing the Java plug-in on Windows

Follow these steps:

1. Click **Start Menu** → **Control Panel** and select the Java Plug-in Control Panel.
2. Click the **About** tab.
3. Determine whether the correct Java Plug-in version is installed:
 - If the correct version is installed, Web Tools is ready to use.
 - If no Java Plug-in is installed, point the browser to a switch running Fabric OS v5.2.0 or later to install JRE 1.6.0. Web Tools guides you through the steps to download the proper Java Plug-in.
 - If an outdated version is currently installed, uninstall it, reboot your computer, re-open the browser, and enter the address of a switch running Fabric OS v5.2.0 or later to install JRE 1.6.0. update16. Web Tools guides you through the steps to download the proper Java Plug-in.

8.1.4 Java plug-in configuration

If you are managing fabrics with more than 10 switches or 1000 ports, or if you are using the iSCSI Gateway module extensively, increase the default heap size to 256 MB to avoid out-of-memory errors.

If you are using a Mozilla family browser (Firefox, Netscape, and so on), set the default browser in the Java control panel.

The following procedures instruct you in increasing the default heap size in the Java Control Panel and in setting the default browser.

Configuring the Java plug-in for Windows

Follow these steps:

1. From the Start menu, select **Control Panel** → **Java**.
2. Click the **Java** tab (see Figure 8-3).

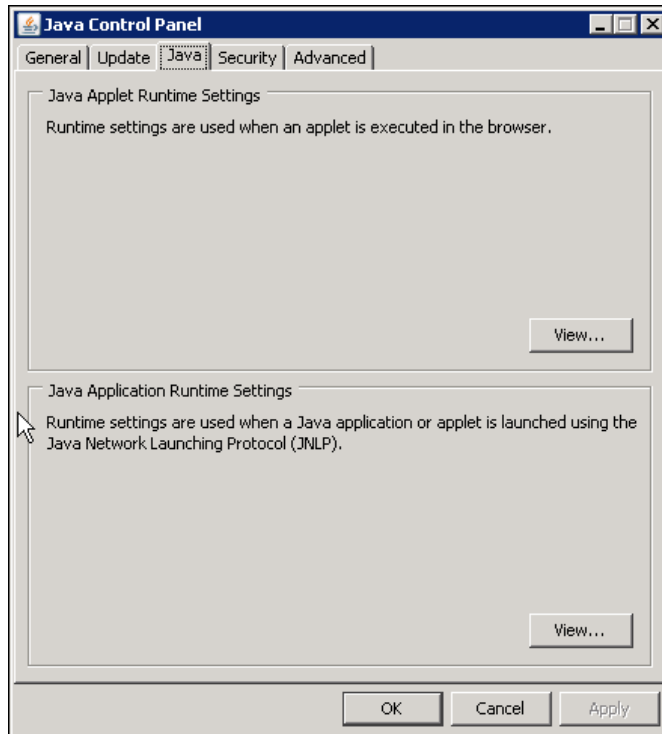


Figure 8-3 Java Control Panel

3. In the section Java Applet Runtime Settings, click **View**.
The Java Applet Runtime Settings dialog box displays (Figure 8-4).




Figure 8-4 Java Runtime Settings

4. Double-click in the Java Runtime Parameters field and type the following information to set the minimum and maximum heap size:

`-Xms256m -Xmx256m`

In this example, the minimum and maximum sizes are both 256 MB.

5. Click **OK**.
6. Click **Apply** to apply your settings and close the Java Control Panel.

Configuring the Java plug-in for Mozilla family browsers

Follow these steps:

1. From the Start menu, select **Control Panel** → **Java**.
2. Click the **Advanced** tab and expand the **Default Java for browsers** option (Figure 8-5).

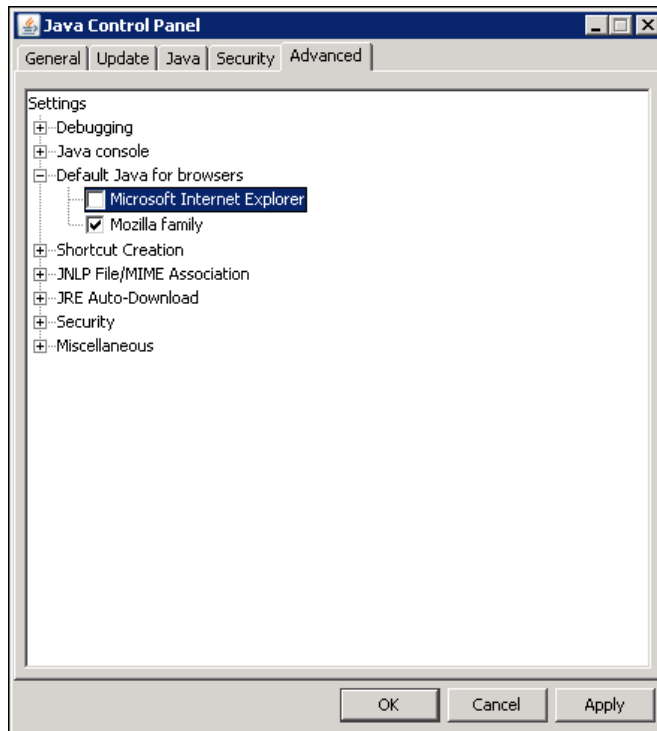


Figure 8-5 Default Java for browsers option

3. Select **Mozilla family** and click **OK**.
4. Click **Apply** to apply your settings and close the Java Control Panel.

8.1.5 Value line licenses

If your fabric includes a switch with a limited switch license and you are opening Web Tools using that switch, if the fabric exceeds the switch limit indicated in the license, Web Tools allows a 30-day “grace period” in which you can still monitor the switch through Web Tools. However, Web Tools will display warning messages periodically.

These messages warn you that your fabric size exceeds the supported switch configuration limit and tells you how long you have before Web Tools will be disabled. After the 30-day grace period, you will no longer be able to open Web Tools from the switch with the limited switch license if that switch is still exceeding the switch limit.

Web Tools is part of the Fabric OS of a switch. When you open Web Tools on a switch, you can manage other switches in the fabric that have lower or higher firmware versions. It is important to note that when accessing these switches you are opening the remote switch’s version of Web Tools, and the functionality available for those switches might vary.

8.1.6 Opening Web Tools

You can open Web Tools on any workstation with a compatible Web browser installed. For a list of Web browsers compatible with Fabric OS v6.4.0, see Table 8-3 on page 187 Table 3. Web Tools supports both HTTP and HTTPS protocol.

1. Open the Web browser and type the IP address of the device in the Address field:

`http://10.77.77.77`

or

`https://10.77.77.77`

2. Press Enter.

A browser window opens to open Web Tools. A Login dialog box opens. See “Logging in” on page 194, for more information. If you are using Firefox, the browser window is left open. You can close it anytime after the Login dialog box displays. If you are using Internet Explorer, the browser window automatically closes when the login dialog box displays. When you have successfully logged in the Web Tools interface opens (see Figure 8-6).

If you have installed EZSwitchSetup on your workstation, the EZSwitchSetup Switch Manager displays the first time you access the device. EZSwitchSetup provides an easy to use wizard interface that can be used to simplify the initial setup procedure for smaller switches. See the EZSwitchSetup Administrator's Guide for information about the EZSwitchSetup interface.

If you want to use Web Tools instead of EZSwitchSetup, click **Advanced Management** in the lower-left corner of the window to open the Web Tools interface. This section describes only the Web Tools interface.

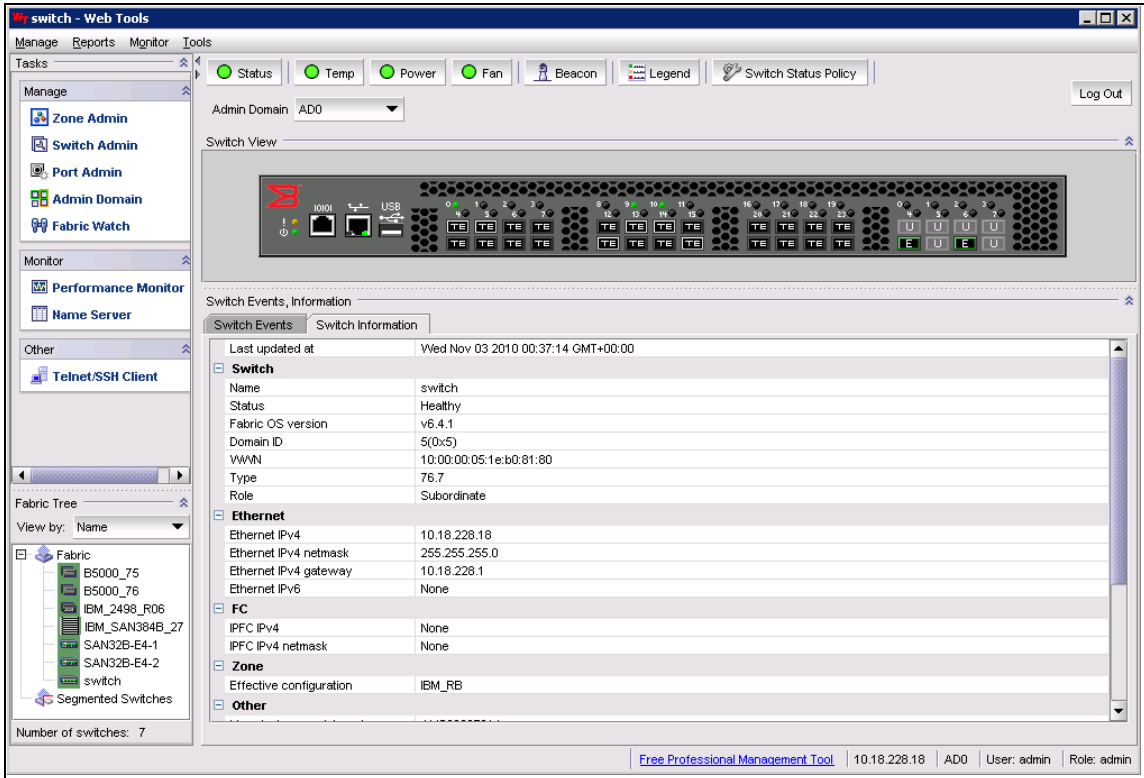


Figure 8-6 Web Tools interface

Logging in

When you use Web Tools, you must log in before you can view or modify any switch information. This section describes the login process.

Prior to displaying the login window, Web Tools displays a security banner (if one is configured for your switch), which you must accept before logging in. The security banner displays every time you access the switch.

When you are presented with the login window, you must provide a user name and a password. Your home Admin Domain is automatically selected. You can choose to log into an Admin Domain other than your home domain:

1. Click **Run** on the signed certificate applet.

If you select the check box **Always trust content from this publisher**, the dialog box is not displayed when you open Web Tools again, as shown in Figure 8-7.

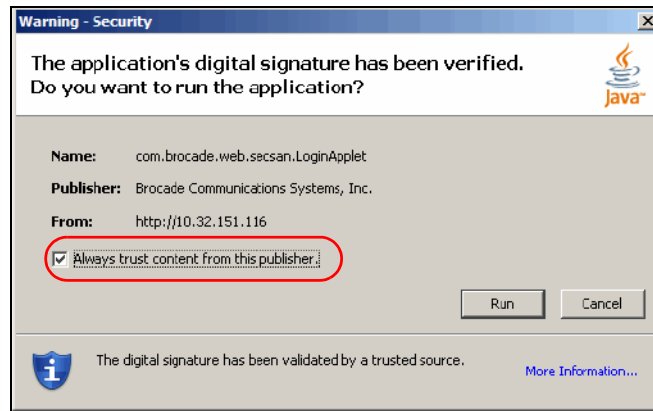


Figure 8-7 Signed applet certificate

2. Click **OK** in the security banner window, if one displays.
3. In the login dialog box, as shown in Figure 8-8, type your user name.
4. Type the password.

If your current password has expired, you must also provide a new password and confirm the new password.

5. Click **OK**.

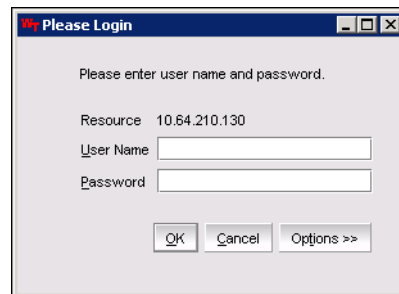


Figure 8-8 Login dialog box

Logging in to a Virtual Fabric

If you are logging in to a platform that is capable of supporting Virtual Fabrics, the log in dialog box provides the option of logging in to a virtual fabric. The following platforms support virtual fabrics:

- ▶ IBM SAN768B and IBM SAN387B
- ▶ IBM SAN80B
- ▶ IBM SAN40B

1. Select **Options** to display the Virtual Fabric options.

You are given a choice between **Home Logical Fabric** and **User Specified Virtual Fabric** as shown in Figure 8-9. Home Logical Fabric is the default. This option logs in to the physical switch, and displays the physical switch configuration. It is given a default fabric ID number of 128.




Figure 8-9 Virtual Fabric login option

2. Log in to a logical fabric:

- ▶ To log in to the home logical fabric, select the **Home Logical Fabric** radio button and click **OK**.
- ▶ To log in to a logical fabric other than the home logical fabric, select the **User Specified Logical Fabric** radio button, type in the fabric ID number, and click **OK**.

Logging in to an Admin Domain

If you are logging in to a platform that is capable of supporting Admin Domains, the log in dialog box provides the option of logging in to an Admin Domain.

You do not have an Admin Domain option if the Access Gateway or Interoperability mode is enabled. Admin Domains and Virtual Fabrics are mutually exclusive.

Follow these steps:

1. Select **Options** to select an Admin Domain other than your default home domain.

You are given a choice of **Home Domain** (the default), or **User Specified Domain**, as shown in Figure 8-10.




Figure 8-10 Login dialog box with Admin Domain options

2. Log in to an Admin Domain:
 - ▶ To log in to the home domain, select the **Home Domain** radio button and click **OK**.
 - ▶ To log in to an Admin Domain other than the home domain, select the **User Specified Domain** radio button, type in the Admin Domain name or number, and click **OK**.

If the user name or password is incorrect, a dialog box displays indicating an authentication failure.

If you entered valid credentials, but specified an invalid Admin Domain, a dialog box displays from which you can choose a valid Admin Domain or click **Cancel** to log in to your home domain, as shown in Figure 8-11.

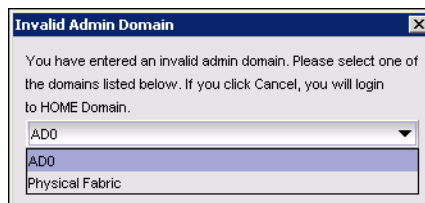


Figure 8-11 Invalid Admin Domain dialog box

Logging out

You can end a Web Tools session either by logging out or by closing the **Switch Explorer** window. You might be logged out of a session involuntarily, without explicitly clicking the **Logout** button, under the following conditions:

- ▶ A physical fabric administrator changes the contents of your currently selected Admin Domain.
- ▶ Your currently selected Admin Domain is removed or invalidated.
- ▶ Your currently selected Admin Domain is removed from your Admin Domain list.
- ▶ You initiate a firmware download from the Web Tools Switch Administration window. In this case, you are logged out a few minutes later when the switch reboots.
- ▶ Your session times out.

Role-Based Access Control

Role-Based Access Control (RBAC) defines the capabilities that a user account has based on the role the account is assigned. For each role, there is a set of pre-defined permissions on the jobs and tasks that can be performed on a fabric and its associated fabric elements.

When you log in to a switch, your user account is associated with a pre-defined role. The role that your account is associated with determines the level of access you have on that switch and in the fabric. Following is a description of each of the roles:

- ▶ admin:
You have full access to all of the Web Tools features.
- ▶ operator:
You can perform any actions on the switch that do not affect the stored configuration.
- ▶ securityadmin:
You can perform actions that do not affect the stored configuration.
- ▶ switchadmin:
You can perform all actions on the switch, with the following exceptions:
 - You cannot modify zoning configurations.
 - You cannot create new accounts.
 - You cannot view or change account information for any accounts. You can only view your own account and change your account password.

- ▶ zoneadmin:
You can only create and modify zones.
- ▶ fabricadmin:
You can do everything the Admin role can do except create new users.
- ▶ basicswitchadmin:
You have a subset of Admin level access.
- ▶ user:
You have non-administrative access and can perform tasks such as monitoring system activity.

Session management

A Web Tools session is the connection between the Web Tools client and its managed switch. A session is established when you log in to a switch through Web Tools. When you close Switch Explorer, Web Tools ends the session.

A session remains in effect until one of the following conditions happens:

- ▶ You log out.
- ▶ You close the Switch Explorer window.
- ▶ The session ends due to inactivity (time out).

A session automatically ends if no information was sent to the switch for more than two hours. Because user key strokes are not sent to the switch until you apply or save the information, it is possible for your session to end while you are entering information in the interface. For example, entering a zoning scheme in the Zoning module does not require you to send information to the switch until you save the scheme.

Web Tools does not display a warning when the session is about to time out. If your session ends due to inactivity, all Web Tools windows become invalid and you must restart Web Tools and log in again.

Web Tools enables sessions to both secure and non-secure switches.

Access rights for your session are determined by your role-based access rights and by the contents of your selected Admin Domain. After you log in, you can change to a different Admin Domain at any time; however, you cannot change your role-based permissions.

Ending a Web Tools session

To end a Web Tools session, perform one of the following actions:

- ▶ Click **Logout** in Switch Explorer.

- ▶ Click the **X** in the upper-right corner of Switch Explorer window to close it.
- ▶ Close all open Web Tools windows.

Attention: If you click Logout in Switch Explorer, and Web Tools leaves the Temperature, Fan, Power, and the Switch status windows open, you must manually close them.

Requirements for IPv6 support

The following list provides requirements for Web Tools IPv6 support:

- ▶ In a pure IPv6 environment, you must configure DNS maps to the IPv6 address of the switch.
- ▶ The switch name is required to match the DNS name that is mapped to the IPv6 address.
- ▶ If both IPv4 and IPv6 addresses are configured, Web Tools uses the IPv4 address to launch the switch.
- ▶ Use a switch with v5.3.0 or higher firmware to manage a mixed fabric of IPv4 and IPv6 switches.
- ▶ Switches running on version 5.2.0 do not discover IPv6 address-only switches in the same fabric until the IPv4 address is configured.

Web Tools system logs

The log4j framework is used to write Web Tools log files.

Web Tools automatically creates the log directories the first time in this directory:

- ▶ <webtools> directory
- ▶ Web Tools switch support save directory with name format as <core switch name-IP address- Switch WWN>, which has the following files:
 - Log4j.XML - This configuration file can be edited with a compatible XML editor if data at startup is to be collected.
 - webtools.log - This log file for Web Tools is maintained at 2 MB size limit.
 - switchinfo.txt - This file contains basic switch information such as switch name, FOS version, switch type, Ethernet configuration with IP, subnet mask, and gateway.

8.1.7 Requirements for the examples in this chapter

In the examples in this chapter, we use the SAN80B switch, SAN256B director, and SAN384B backbone to describe the Web Tools graphical user interface (GUI), although the functions are identical on any of the IBM System Storage SAN switch family. We use Fabric Operating System v6.4.1 (Fabric OS).

8.1.8 Overview of the Web Tools user interface

In this section, we provide a brief overview of the Web Tools GUI. To open the interface, start the Web browser if it is not already active and enter the switch name or IP address in the Location/Address field. Enter login credentials when prompted.

Tip: When managing a multi-switch fabric, enter the switch name or IP address of the switch with the largest port count and the highest firmware level.

The first thing you see when you log in to a switch with Web Tools is Switch Explorer, shown in Figure 8-12. Switch Explorer is divided into areas that provide access to, and information about, the switch and fabric.

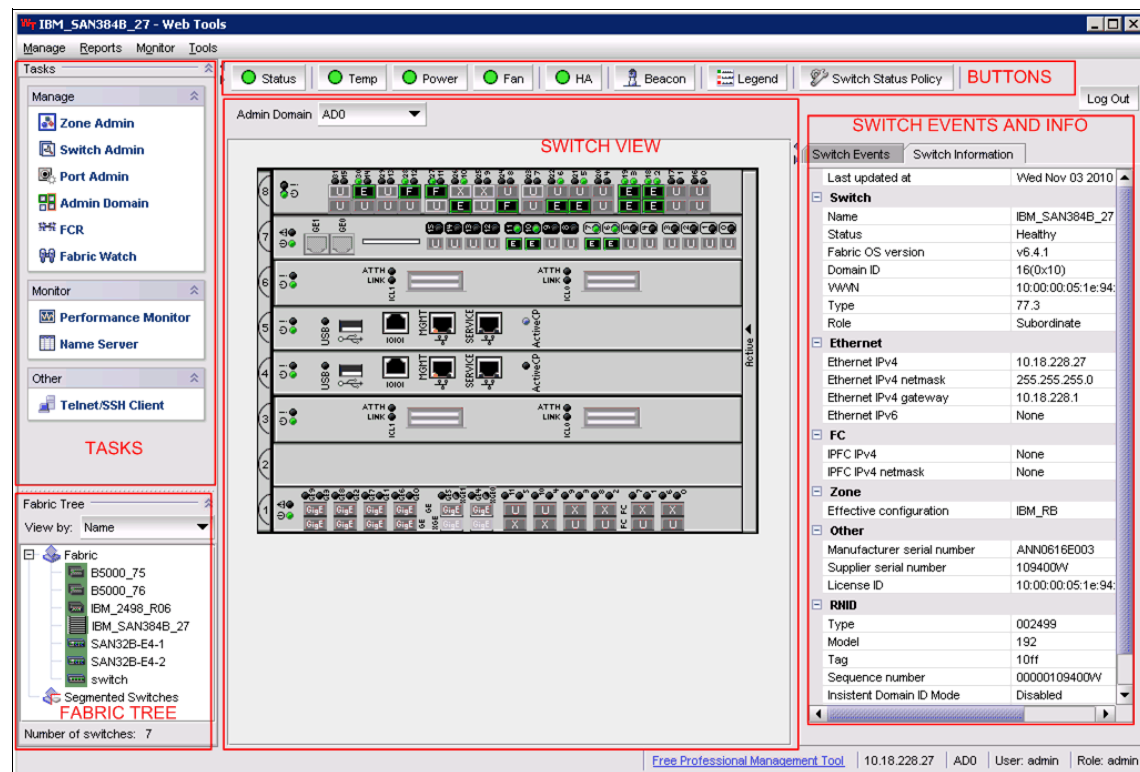


Figure 8-12 Web Tools Switch Explorer window

As highlighted in Figure 8-12, the Web Tools window includes the following panels:

- ▶ Tasks panel: Includes shortcuts to the tasks that are available in the Web Tools interface. These tasks are also accessible through pull-down menu bar options.
- ▶ Fabric Tree panel: Lists the switches in the fabric. In our examples, we have two switches in the fabric. You can launch Web Tools for another switch by selecting its icon.
- ▶ Switch View panel: Displays a picture of the switch. You can click any of the switch ports to launch the Port Administration window.
- ▶ Switch Events, Information panel: Contains two tabs:
 - Switch Information tab, which displays basic information about the switch.
 - Switch Events tab, which shows the event log.

The window also includes an area with buttons, which we discuss in detail in 8.2, “Web Tools buttons” on page 204.

The information that is displayed in these panels differs depending on the switch that is accessed and the licenses that are applied. For example, Figure 8-13 shows the Tasks panel on a SAN80B (left), a SAN256B (center) and a SAN384B (right).

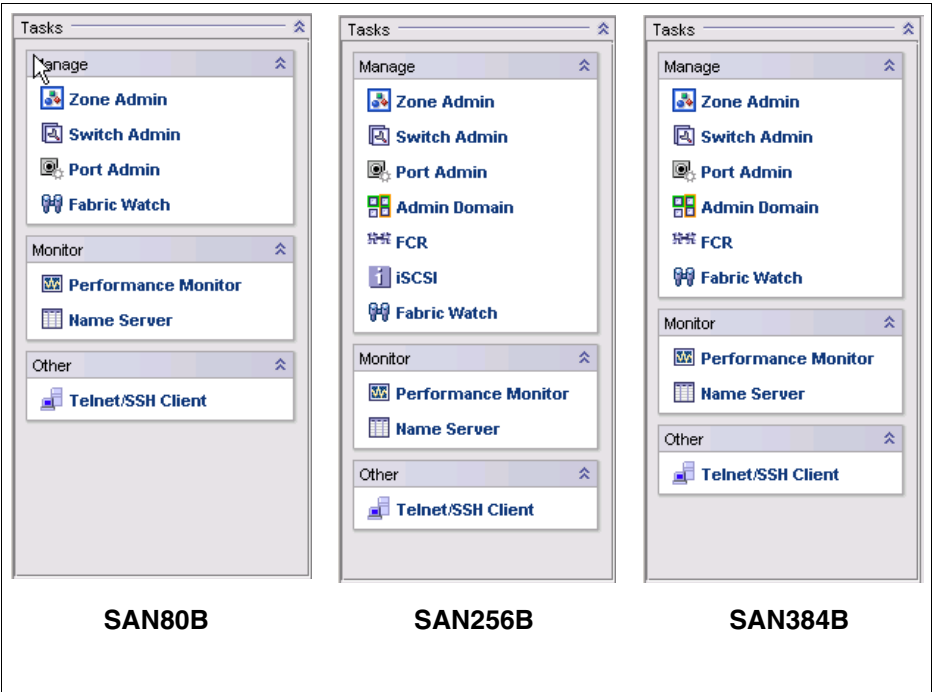


Figure 8-13 Tasks panel on SAN32B-3, SAN256B and SAN768B

The Switch View panel differs significantly between storage area network (SAN) switches (such as SAN80B-3) and SAN directors or backbones (SAN256B or SAN384B). For example, Figure 8-14 shows the Switch View panel for the SAN80B-3, and Figure 8-15 shows a very different panel for the SAN384B.




Figure 8-14 Switch View panel for a SAN80B

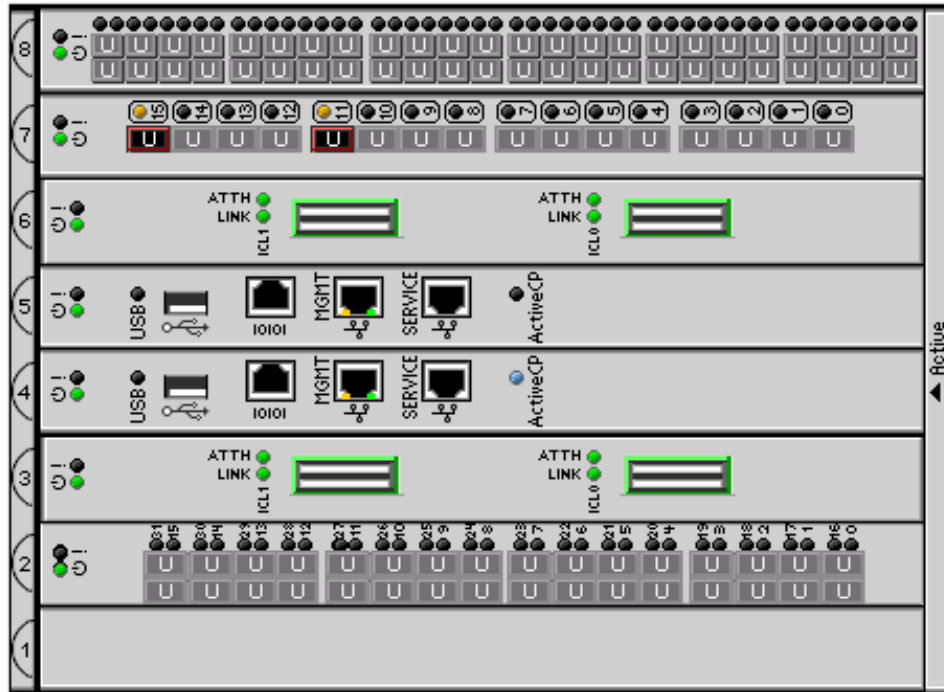


Figure 8-15 Switch View panel for the SAN384B

You access the Web Tools functions by clicking various items in the Web Tools window, such as:

- ▶ Web Tools buttons
- ▶ Tasks in the Tasks panel
- ▶ Ports in Switch View panel
- ▶ Elements in the Fabric View panel

In the remaining sections of this chapter, we describe the functions of Web Tools buttons and tasks in detail.

8.2 Web Tools buttons

Web Tools buttons display above the Switch View panel and provide quick access to hardware status and environmental information. The buttons use the following color coding:

- ▶ A green mark indicates an optimal state of components or parameters.
- ▶ A yellow mark is a sign of a degraded state.
- ▶ A red mark means that the monitored component is down.

Figure 8-16 shows the Web Tools buttons, which we explain in the sections that follow.



Figure 8-16 Web Tools buttons

8.2.1 Status button

The Status button is available on all IBM System Storage SAN Switch models. Clicking **Status** opens the Switch Health Report window, as shown in Figure 8-17, which shows the health of the switch.

Action

Report

- Switch Health
- Port Detail
 - Healthy
 - Marginal
 - Faulty
 - All
- SAM

Switch Health Report Report Time: Fri Jul 17 2009 10:32:28 GMT

Switch Name: IBM_SAN384B_213
IP Address: 10.64.210.213
Switch State: **HEALTHY**
Duration (H:M): 17: 44

Switch State Contributors	State
Power supplies monitor	HEALTHY
Temperatures monitor	HEALTHY
Fans monitor	HEALTHY
WWN servers monitor	HEALTHY
Standby CP monitor	HEALTHY
Blades monitor	HEALTHY
Core Blades monitor	HEALTHY
Flash monitor	HEALTHY
Marginal ports monitor	HEALTHY
Faulty ports monitor	HEALTHY
Missing SFPs monitor	HEALTHY

All ports are healthy.

Figure 8-17 SAN384B switch status view from Web Tools

From here, you can navigate to obtain information about the health of different ports on the switch. Under Port Detail, you can view the ports in *HEALTHY*, *MARGINAL*, and *FAULTY* status. Clicking **All** displays details about all the ports.

Figure 8-18 shows the details for just the healthy ports on a SABN80B, which is helpful information in understanding the port states.

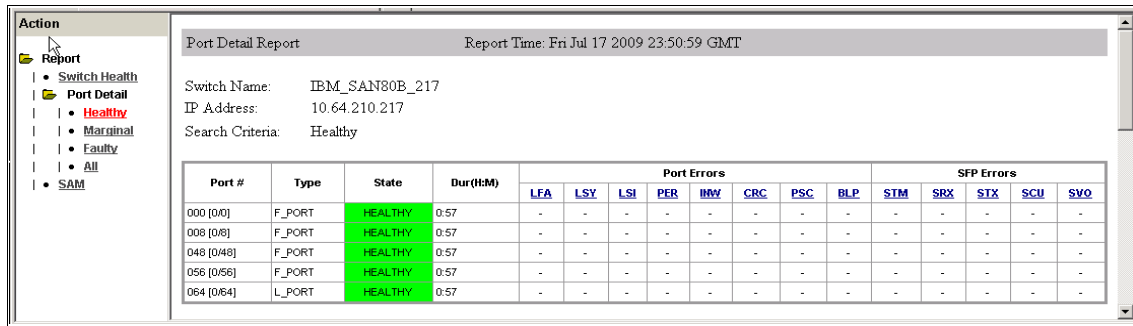


Figure 8-18 SAN80B Healthy Port Detail Report

The report view provides a full description of each of the columns. Information about this report is detailed in Table 8-4 for reference.

Table 8-4 Port Detail Report: Error interpretation

Error	Description and suggested action
LFA (Link Loss)	Description: Number of link loss occurrences exceeded range for time period.
	Action: Troubleshoot transmitters, receivers, and fibers, and verify that all cables connect properly.
LSY (Sync Loss)	Description: Number of sync loss occurrences exceeded range for time period.
	Action: Check for problems with the appropriate SFP and cable. If you continue to experience sync loss errors, troubleshoot your HBA and contact your support representative.
LSI (Signal Loss)	Description: Number of signal loss occurrences exceeded range for time period.
	Action: Troubleshoot transmitters, receivers, and fibers, and verify that all cables connect properly.
PER (Protocol Error)	Description: Number of protocol errors exceeded range for time period.
	Action: Check both ends of your connection, and verify that your cable and SFP are not faulty.

Error	Description and suggested action
INW (Invalid Word)	Description: Number of invalid words exceeded range for time period.
	Action: Verify that your cable is not faulty and check both ends of your connection. Troubleshoot your SFP to verify that it is not faulty.
CRC (Invalid CRC)	Description: Number of invalid CRC errors exceeded range for time period.
	Action: Check your SFPs, cables, and connections for faulty hardware. Clean all optical hardware.
PSC (Port State)	Description: Port hardware state changed too often due to fabric reconfiguration.
	Action: All State Changes messages are informational. Respond to this message as is appropriate to the particular policy of the user installation.
BLP (Buffer Limited Port)	Description: Port buffer credit was not large enough.
	Action: Reset the buffer credit.
STM (SFP Temperature)	Description: SFP temperature is out of specifications.
	Action: Temperature-related messages usually indicate that you must replace the SFP.
SRX (SFP RX)	Description: SFP receive power is out of specification.
	Action: Replace the SFP.
STX (SFP TX)	Description: SFP transmit power is out of specifications.
	Action: If the current rises above the high boundary, you must replace the SFP.
SCU (SFP Current)	Description: SFP current is out of specifications.
	Action: If the current rises above the high boundary, you must replace the SFP.
SVO (SFP Voltage)	Description: SFP voltage is out of specifications.
	Action: Frequent messages indicate that you must replace the SFP.
-	Meaning: Monitoring value is within the threshold.
X	Meaning: Monitoring value is over the threshold.

You can display the same information at the Telnet prompt by entering the **switchStatusShow** command, as shown in Example 8-1.

Example 8-1 The switchStatusShow command output

```
IBM_SAN384B_27:admin> switchstatusshow
Switch Health Report                                Report time: 11/03/2010
11:33:05 AM
Switch Name:    IBM_SAN384B_27
IP address:     10.18.228.27
SwitchState:    HEALTHY
Duration:       17:44

Power supplies monitor  HEALTHY
Temperatures monitor   HEALTHY
Fans monitor           HEALTHY
WWN servers monitor    HEALTHY
CP monitor             HEALTHY
Blades monitor         HEALTHY
Core Blades monitor    HEALTHY
Flash monitor          HEALTHY
Marginal ports monitor HEALTHY
Faulty ports monitor   HEALTHY
Missing SFPs monitor   HEALTHY

All ports are healthy
IBM_SAN384B_27:admin>
```

Selecting **SAM** from the menu displays the SAM (Switch Availability Monitoring Report), as shown in Figure 8-19.

Action	SAM Report Report Time: Fri Jul 17 2009 17:47:35 GMT				
	Switch Name: IBM_SAN80B_217 IP Address: 10.64.210.217				
Report					
• Switch Health					
• Port Detail					
• Healthy					
• Marginal					
• Faulty					
• All					
• SAM					

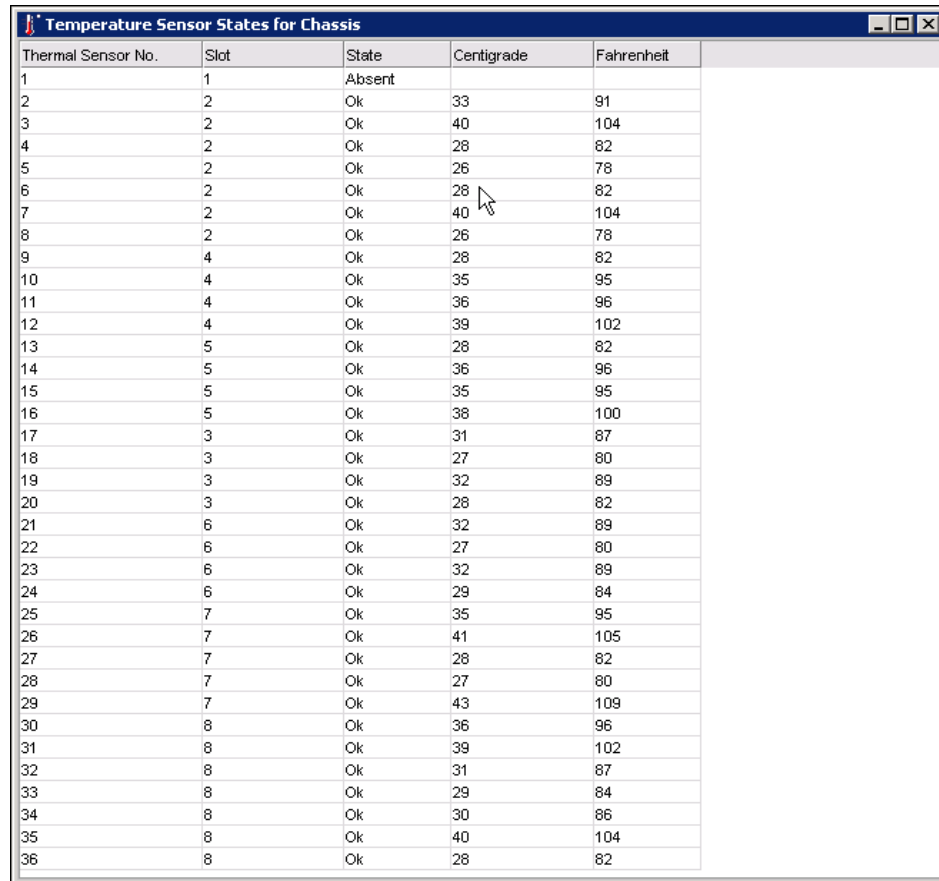
Port #	Type	Total Up Time (Percent)	Total Down Time (Percent)	Down Occurrence (Times)	Total Offline Time (Percent)
001 [0/1]	U	0	0	0	100
002 [0/2]	U	0	0	0	100
003 [0/3]	U	0	0	0	100
004 [0/4]	U	0	0	0	100
005 [0/5]	U	0	0	0	100
006 [0/6]	U	0	0	0	100
007 [0/7]	U	0	0	0	100
009 [0/9]	U	0	0	0	100
010 [0/10]	U	0	0	0	100

Figure 8-19 Extract from the SAM Report

8.2.2 Temp button

The Temp button changes color from green to show that all temperatures are within the defined limits or changes to yellow or red, depending on the policy thresholds.

Clicking **Temp** displays detailed temperature information. Figure 8-20 shows an example for a SAN384B chassis.



Thermal Sensor No.	Slot	State	Centigrade	Fahrenheit
1	1	Absent		
2	2	Ok	33	91
3	2	Ok	40	104
4	2	Ok	28	82
5	2	Ok	26	78
6	2	Ok	28	82
7	2	Ok	40	104
8	2	Ok	26	78
9	4	Ok	28	82
10	4	Ok	35	95
11	4	Ok	36	96
12	4	Ok	39	102
13	5	Ok	28	82
14	5	Ok	36	96
15	5	Ok	35	95
16	5	Ok	38	100
17	3	Ok	31	87
18	3	Ok	27	80
19	3	Ok	32	89
20	3	Ok	28	82
21	6	Ok	32	89
22	6	Ok	27	80
23	6	Ok	32	89
24	6	Ok	29	84
25	7	Ok	35	95
26	7	Ok	41	105
27	7	Ok	28	82
28	7	Ok	27	80
29	7	Ok	43	109
30	8	Ok	36	96
31	8	Ok	39	102
32	8	Ok	31	87
33	8	Ok	29	84
34	8	Ok	30	86
35	8	Ok	40	104
36	8	Ok	28	82

Figure 8-20 SAN384B Temperature status window

To display similar information at a Telnet command line, issue the **tempShow** command as shown in Example 8-2.

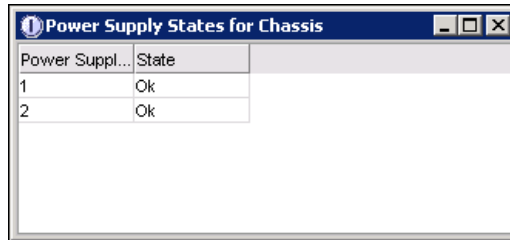
Example 8-2 SAN384B tempShow output

```
IBM_SAN384B_213:admin> tempShow
```

Sensor ID	Slot	State	Centigrade	Fahrenheit
=====				
1	1	Absent		
2	2	Ok	33	91
3	2	Ok	40	104
4	2	Ok	28	82
5	2	Ok	26	78
6	2	Ok	28	82
7	2	Ok	40	104
8	2	Ok	26	78
9	4	Ok	28	82
10	4	Ok	35	95
11	4	Ok	36	96
12	4	Ok	39	102
13	5	Ok	28	82
14	5	Ok	36	96
15	5	Ok	35	95
16	5	Ok	38	100
17	3	Ok	31	87
18	3	Ok	27	80
19	3	Ok	32	89
20	3	Ok	28	82
21	6	Ok	32	89
22	6	Ok	27	80
23	6	Ok	32	89
24	6	Ok	29	84
25	7	Ok	35	95
26	7	Ok	41	105
27	7	Ok	28	82
28	7	Ok	27	80
29	7	Ok	43	109
30	8	Ok	36	96
31	8	Ok	39	102
32	8	Ok	31	87
33	8	Ok	29	84
34	8	Ok	30	86
35	8	Ok	40	104
36	8	Ok	29	84

8.2.3 Power button

The color of the Power button indicates the overall health of the power supply status. Clicking **Power** displays the window shown in Figure 8-21.

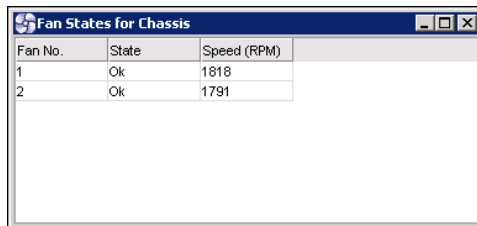
A screenshot of a window titled "Power Supply States for Chassis". It contains a table with two columns: "Power Suppl..." and "State".

Power Suppl...	State
1	Ok
2	Ok

Figure 8-21 SAN348B power status

8.2.4 Fan button

If all conditions are normal according to the switch policy settings, the Fan button is green. Clicking **Fan** displays an informational window that describes the state of each fan, as shown in Figure 8-22.

A screenshot of a window titled "Fan States for Chassis". It contains a table with three columns: "Fan No.", "State", and "Speed (RPM)".

Fan No.	State	Speed (RPM)
1	Ok	1818
2	Ok	1791

Figure 8-22 SAN384B fan status

You can gather the same information from a Telnet command line by entering the **fanShow** command as shown in Example 8-3.

Example 8-3 SAN384B fanShow command

```
IBM_SAN384B_213:admin> fanShow
Fan 1 is Ok, speed is 1846 RPM
Fan 2 is Ok, speed is 1818 RPM
```

8.2.5 HA button

The SAN768B, SAN384B and SAN256B support High Availability (HA) features. The color of the HA button indicates the overall high availability status of the switch. This button enables you to perform tasks such as CP failover or synchronization services on the CP.

Clicking **HA** launches the High Availability window shown in Figure 8-23. The first tab shows the status of the services for the switch. Notice that in the upper, right corner, the HA Status field is green and displays the message Non-Disruptive Failover Ready. If the HA Status field was not green, then you need to synchronize the services before attempting to initiate failover. When the HA Status field shows Non-disruptive Failover Ready, a failover can be initiated without disrupting the fabric.

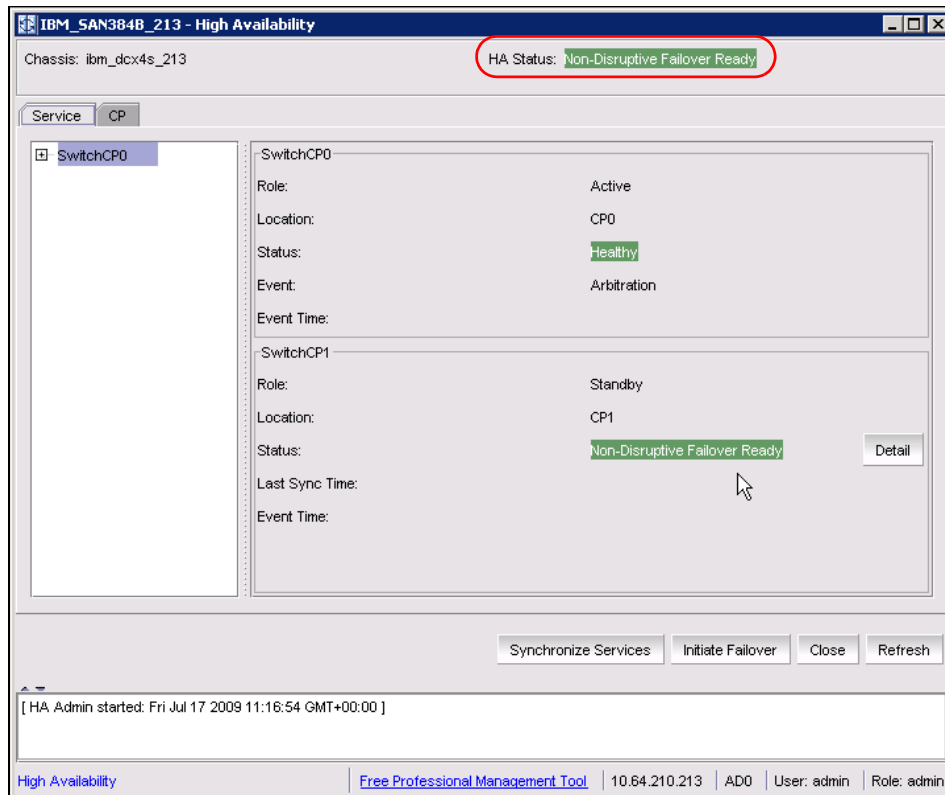


Figure 8-23 SAN384B High Availability status

When you select **Synchronize Services** (shown in Figure 8-23), you are prompted with a warning to confirm your actions (as shown in Figure 8-24). Click **Yes**.

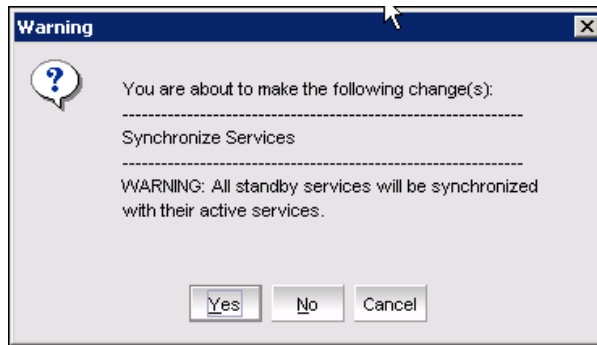


Figure 8-24 Synchronize services warning

From the same panel in Figure 8-23, you can initiate the failover and monitor the status by clicking **Initiate Failover**. A warning displays, as shown in Figure 8-25.

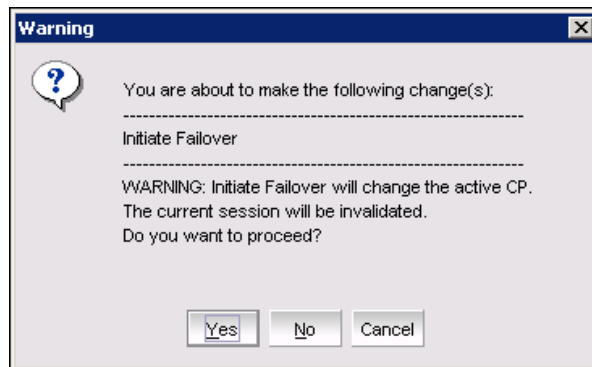


Figure 8-25 Initiate failover warning

After clicking **Yes**, failover is initiated. The HA status field changes to red with the message Non-Redundant Failover to indicate that failover is taking place. Just before it completes the failover, HA status shows yellow and indicates Disruptive Failover Ready. When failover is complete, the CPs have changed as shown in Figure 8-26, and the HA status returns to Non-Disruptive Failover Ready.

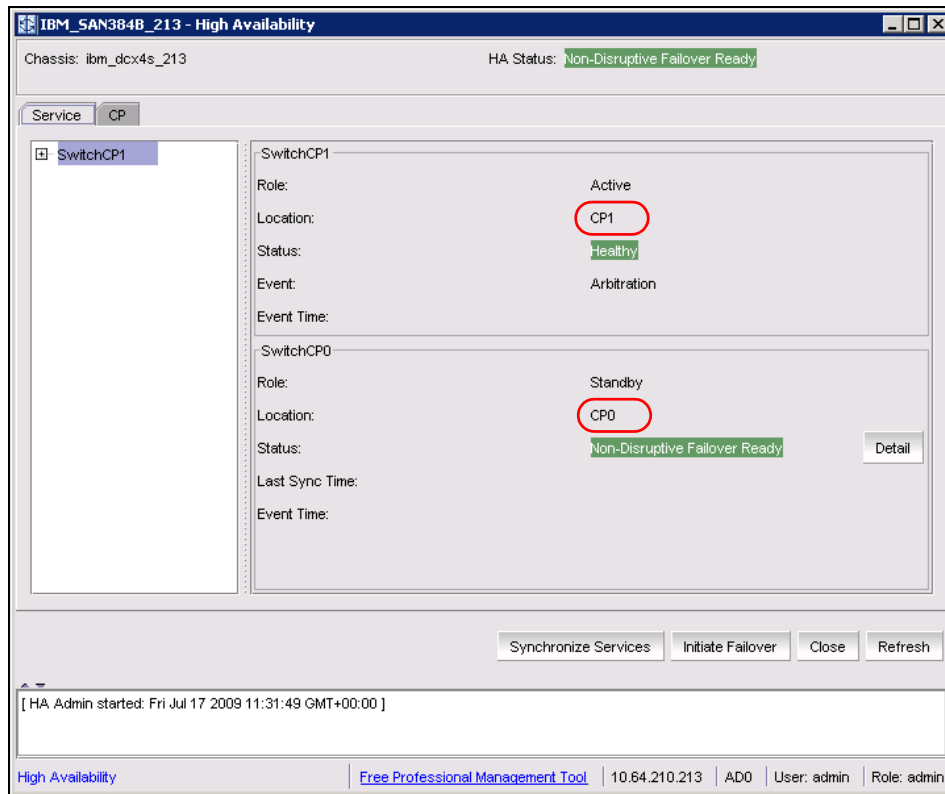


Figure 8-26 SAN384B: Failover complete

Failover: A non-disruptive failover can take a few minutes to complete. It is possible that the connection to the switch might be lost during that time.

8.2.6 Beacon button

The beaconing function allows you to locate a switch physically by sending a signal to the specified switch, which causes a yellow LED pattern to flash from side to side on the switch. This flashing pattern makes the switch very easy to find.

To activate beaconing, click **Beacon** as shown in Figure 8-27.




Figure 8-27 SAN384B showing the Beacon button

You can toggle this function on and off when the switch is identified.

8.2.7 Switch Status Policy button

This enhancement in Web Tools 6.4.0 and later defines the policy for switch status notifications, and is used to define what constitutes a healthy switch status and also the parameters to define for marginal/down state of a switch. The entire switch status changes if any one of the parameter values changes. Figure 8-28 shows the Web Tools options available to define the switch status policy of a SAN384B.

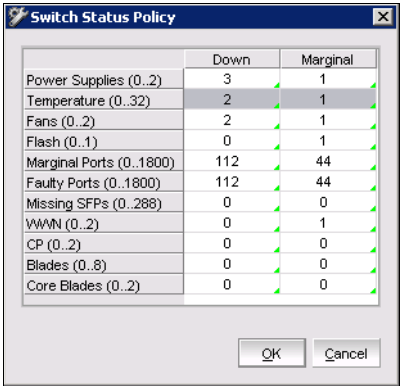


Figure 8-28 Switch status policy

This is similar to defining the policy with command `switchstatuspolicyset`.

8.2.8 Legend button

Clicking **Legend** displays the explanation of color-coded status icons, as shown in Figure 8-29.



Figure 8-29 Web Tools legend

8.3 Name Server task

Select this task to display the Name Server table, as shown in Figure 8-30. The table includes all name server entries for the fabric, not just those name server entries that are local to the host domain. Each row in the table represents a different device that has logged in to the fabric. The Name Server table provides a good cross reference of WWPN/WWN and the port position on the switch. It also lists the zones of which the port is a member and, therefore, can be a very useful problem determination tool.

IBM_SAN256B_130 - Name Server									
<input type="checkbox"/> Auto Refresh		Auto-Refresh Interval: 15 seconds		Number of Devices: 8					
Domain	Port #	Port ID	Port Type	Device Port WWN	Device Node WWN	Device Name	Capability	FDML Host Name	WWN Company ID
1(0x1)	111(0x6F)	016f00	N	50:06:06:98:04:50:c7:38	50:06:06:98:04:50:c7:39	IPAddr: 0.0.0...	NS		Brocade Communications Systems, Inc.
1(0x1)	105(0x69)	016900	N	50:06:06:98:04:50:c7:08	50:06:06:98:04:50:c7:09	IPAddr: 0.0.0...	NS		Brocade Communications Systems, Inc.
1(0x1)	109(0x6D)	016d00	N	50:06:06:98:04:50:c7:28	50:06:06:98:04:50:c7:29	IPAddr: 0.0.0...	NS		Brocade Communications Systems, Inc.
1(0x1)	107(0x6B)	016b00	N	50:06:06:98:04:50:c7:18	50:06:06:98:04:50:c7:19	IPAddr: 0.0.0...	NS		Brocade Communications Systems, Inc.
1(0x1)	110(0x6E)	016e00	N	50:06:06:98:04:50:c7:30	50:06:06:98:04:50:c7:31	IPAddr: 0.0.0...	NS		Brocade Communications Systems, Inc.
1(0x1)	106(0x6A)	016a00	N	50:06:06:98:04:50:c7:10	50:06:06:98:04:50:c7:11	IPAddr: 0.0.0...	NS		Brocade Communications Systems, Inc.
1(0x1)	104(0x68)	016800	N	50:06:06:98:04:50:c7:00	50:06:06:98:04:50:c7:01	IPAddr: 0.0.0...	NS		Brocade Communications Systems, Inc.
1(0x1)	108(0x6C)	016c00	N	50:06:06:98:04:50:c7:20	50:06:06:98:04:50:c7:21	IPAddr: 0.0.0...	NS		Brocade Communications Systems, Inc.

Figure 8-30 SAN256B Name Server table (part 1 of 3)

The Name Server table contains the following parameters:

Domain	Domain ID of the switch to which the device is connected
Port #	Port number of the switch to which the device is connected
Port ID	The Fibre Channel Port address of the device (basically, a 24-bit hexadecimal number)

Port Type	Shows whether the port is a public loop port (NL) or a switch fabric port (N)
Device Port WWN	Worldwide name for the device port (WWPN)
Device Node WWN	Worldwide name of the device node (WWNN)
Device Name	Name of the device according to the SCSI INQUIRY, such as FCP or IP
Capability	The Name Server Capability
FDMI Host Name	Displays the FDMI host name of the device
WWN Company ID	Displays vendor company based on device WWN

Scroll to the right to see the remaining parameters in the Name Server table. Figure 8-31 displays the next set.

The screenshot shows a window titled "IBM_SAN256B_130 - Name Server". It includes an "Auto Refresh" checkbox, an "Auto-Refresh Interval" of 15 seconds, and a "Number of Devices: 8". The table below lists device details:

NPIV(or)Virtual(or)Physical	Host vs. Target	Member Of Zones	Member Of Aliases	FC4 Type	Class Of Service	Fabric Port Name	Fabric Port WWN	Port IP Address
Virtual	Unknown(initiator/target)			FCP	3		20:6f:00:60:69:80:45:0c	N/A
Virtual	Unknown(initiator/target)			FCP	3		20:69:00:60:69:80:45:0c	N/A
Virtual	Unknown(initiator/target)			FCP	3		20:6d:00:60:69:80:45:0c	N/A
Virtual	Unknown(initiator/target)			FCP	3		20:6b:00:60:69:80:45:0c	N/A
Virtual	Unknown(initiator/target)			FCP	3		20:6e:00:60:69:80:45:0c	N/A
Virtual	Unknown(initiator/target)			FCP	3		20:6a:00:60:69:80:45:0c	N/A
Virtual	Unknown(initiator/target)			FCP	3		20:68:00:60:69:80:45:0c	N/A
Virtual	Unknown(initiator/target)			FCP	3		20:6c:00:60:69:80:45:0c	N/A

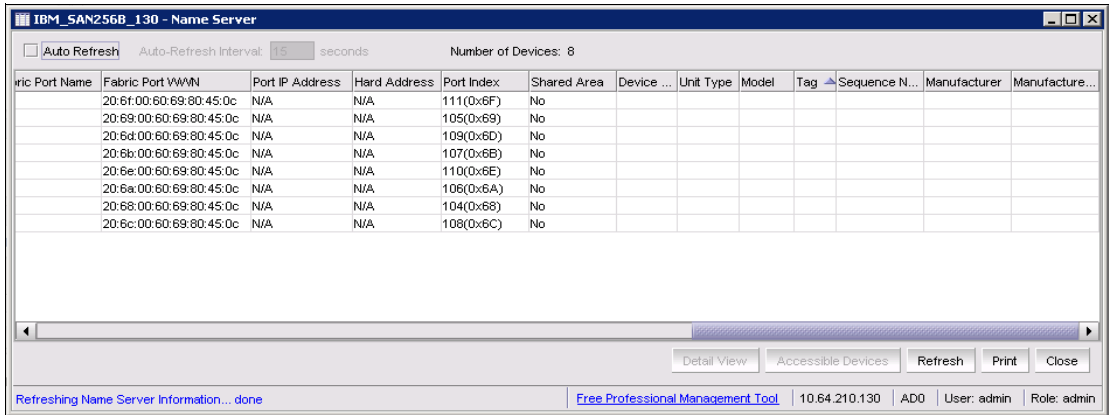
At the bottom, there are buttons for "Detail View", "Accessible Devices", "Refresh", "Print", and "Close". A status bar at the very bottom shows "Refreshing Name Server Information... done", a link to "Free Professional Management Tool", and system information: "10.64.210.130 | ADO | User: admin | Role: admin".

Figure 8-31 SAN256B Name Server table (part 2 of 3)

The remaining parameters include these:

NPIV or Virtual or Physical	Identifies type of device, virtual or physical
Host versus Target	Identifies type of device, host, or target
Member of Zones	List of zones to which the device belongs
Member of Aliases	List of aliases for this device
FC4 Type	Fibre Channel FC4 layer types supported by device, such as FCP or IP
Class of Service	Class of service that the device supports
Fabric Port Name	Displays the name of the port
Fabric Port WWN	The worldwide name of the fabric port
Port IP Address	IP address of the fabric port

Figure 8-32 shows the final set of Name Server table parameters.

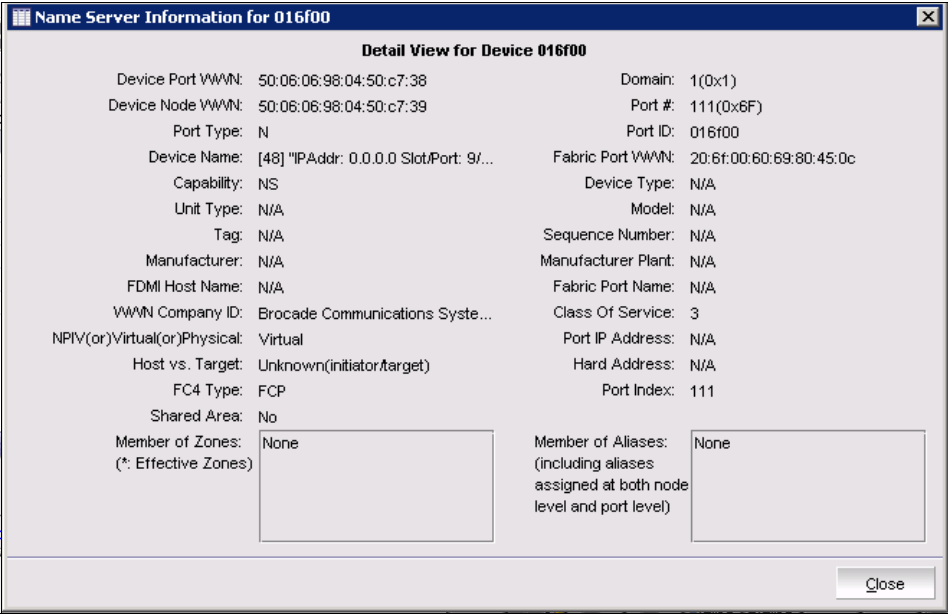


The screenshot shows a window titled "IBM_SAN256B_130 - Name Server". It has a toolbar with "Auto Refresh" (unchecked), "Auto-Refresh Interval: 15 seconds", and "Number of Devices: 8". Below the toolbar is a table with the following columns: Fabric Port Name, Fabric Port WWN, Port IP Address, Hard Address, Port Index, Shared Area, Device ..., Unit Type, Model, Tag, Sequence N..., Manufacturer, and Manufacture... The table contains 8 rows of data. At the bottom of the window, there is a status bar with the text "Refreshing Name Server Information... done", a link "Free Professional Management Tool", and a footer with "10.64.210.130 | AD0 | User: admin | Role: admin".

Fabric Port Name	Fabric Port WWN	Port IP Address	Hard Address	Port Index	Shared Area	Device ...	Unit Type	Model	Tag	Sequence N...	Manufacturer	Manufacture...
	20:6f:00:60:69:80:45:0c	N/A	N/A	111(0x6F)	No							
	20:69:00:60:69:80:45:0c	N/A	N/A	105(0x69)	No							
	20:6d:00:60:69:80:45:0c	N/A	N/A	109(0x6D)	No							
	20:6b:00:60:69:80:45:0c	N/A	N/A	107(0x6B)	No							
	20:6e:00:60:69:80:45:0c	N/A	N/A	110(0x6E)	No							
	20:6a:00:60:69:80:45:0c	N/A	N/A	106(0x6A)	No							
	20:68:00:60:69:80:45:0c	N/A	N/A	104(0x68)	No							
	20:6c:00:60:69:80:45:0c	N/A	N/A	108(0x6C)	No							

Figure 8-32 SAN256B Name Server table (part 3 of 3)

To view all of the details for a given device in the Name Server table, highlight the device in which you are interested and click **Detail View** to open the detailed view window as shown in Figure 8-33.



The screenshot shows a window titled "Name Server Information for 016f00". It has a close button in the top right corner. The window displays detailed information for a specific device. The title bar also includes "Detail View for Device 016f00". The information is organized into two columns. The left column contains: Device Port WWN: 50:06:06:98:04:50:c7:38, Device Node WWN: 50:06:06:98:04:50:c7:39, Port Type: N, Device Name: [48] "IPAddr: 0.0.0.0 Slot/Port: 9/...", Capability: NS, Unit Type: N/A, Tag: N/A, Manufacturer: N/A, FDMI Host Name: N/A, WWN Company ID: Brocade Communications Syste..., NPIV(or)Virtual(or)Physical: Virtual, Host vs. Target: Unknown(initiator/target), FC4 Type: FCP, Shared Area: No, Member of Zones: None (*: Effective Zones). The right column contains: Domain: 1(0x1), Port #: 111(0x6F), Port ID: 016f00, Fabric Port WWN: 20:6f:00:60:69:80:45:0c, Device Type: N/A, Model: N/A, Sequence Number: N/A, Manufacturer Plant: N/A, Fabric Port Name: N/A, Class Of Service: 3, Port IP Address: N/A, Hard Address: N/A, Port Index: 111. At the bottom right, there is a "Close" button.

Figure 8-33 Name Server table entry - detailed view

8.4 Zone Admin task

Selecting this task launches the Zone Administration window. We describe zone administration in detail in Chapter 12, “Basic zoning” on page 513.

8.5 Admin Domain task

An Administrative Domain (*Admin Domain* or *AD*) is a logical grouping of fabric elements that defines what switches, ports, and devices you can view and modify. An Admin Domain is a filtered administrative view of the fabric.

Admin Domains: If you do not implement Admin Domains, the feature has no impact on users, and you do not need to learn how to use this functionality.

Admin Domains permit access to a configured set of users. Using Admin Domains, you can partition the fabric into logical groups and allocate administration of these groups to different user accounts so that these accounts manage only the Admin Domains assigned to them and do not make changes to the rest of the fabric.

For example, you can put all the devices in a particular department in the same Admin Domain for ease of managing those devices. If you have remote sites, you can put the resources in the remote site in an Admin Domain and assign the remote site administrator to manage those resources.

Admin Domains and Virtual Fabrics are mutually exclusive and are not supported at the same time on a switch.

Do not confuse Admin Domains with zones:

- ▶ Zones define which devices and hosts can communicate with each other.
- ▶ Admin Domains define which users can manage which devices, hosts, and switches.

Attention: You do not use the Admin Domain window to assign Admin Domains to particular user accounts. These assignments are performed in the Switch Administration window User tab.

You can have up to 256 Admin Domains in a fabric (254 user-defined and 2 system-defined), numbered from 0 through 255. Admin Domains are designated by a name and a number. This document refers to specific Admin Domains using the format “ADn” where n is a number between 0 and 255.

The two predefined Admin Domains have the following meanings:

- ▶ AD0 is a system-defined Admin Domain that contains all online devices, switches, and ports that have not been assigned manually to any user-defined Admin Domain. However, you can assign members manually to AD0. In addition, AD0 contains devices from switches running Fabric OS earlier than v5.2.0.
- ▶ AD255 (physical fabric) contains all devices, switches, and ports in the fabric. It provides a full, unfiltered view of the fabric. You can manage other Admin Domains within AD255, but you cannot manage zones. AD255 is not associated with any zone database.

Domain: Do not confuse an Admin Domain number with the domain ID of a switch. They are two different identifiers. The Admin Domain number identifies the Admin Domain and has a range of 0 through 255. The domain ID identifies a switch in the fabric and has a range of 1 through 239.

An “AD-capable switch” is a switch that meets the following requirements:

- ▶ Runs Fabric OS v5.2.0 or later (on both CPs, if a dual CP switch)
- ▶ Has a valid Advanced Zoning license, for switches running Fabric OS v5.2.x through 6.2.x.

Zoning: Switches running Fabric OS v6.1.0 or later do not need an Advanced Zoning license because zoning is bundled with the Fabric OS.

A “non-AD-capable switch” is a switch that is running one of these possibilities:

- ▶ Fabric OS v5.1.x or earlier.
- ▶ Fabric OS v5.2.x through Fabric OS v6.0.x, but does not have an Advanced Zoning license.
- ▶ Fabric OS v5.2.0 or later on one CP but Fabric OS v5.1.x or earlier on the other (for dual-CP switches) and the HA state is “synchronized”.

An AD-aware switch is a switch that runs Fabric OS v5.2.0 or later (on both CPs, if a dual CP switch) and that has a valid Advanced Zoning license.

Admin Domains allow you to do the following actions:

- ▶ Define the scope of an Admin Domain to encompass ports and devices within a switch or a fabric.
- ▶ Share resources across multiple Admin Domains. For example, you can share array ports and tape drives between multiple departments.
- ▶ Have a separate zone database for each Admin Domain.
- ▶ Move devices from one Admin Domain to another without traffic disruption, cable reconnects, or discontinuity in zone enforcement.
- ▶ Provide strong fault and event isolation between Admin Domains.
- ▶ Have visibility of all physical fabric resources. All switches, E_Ports, and FRUs (including blade information) are visible.
- ▶ Implement Admin Domains in a fabric with some switches running AD-unaware firmware versions (that is, firmware versions earlier than Fabric OS v5.2.0).
- ▶ Continue to run existing third-party management applications. Prior and existing versions of third-party management applications continue to work with admin and user IDs.

Figure 8-34 shows a fabric with two Admin Domains: AD1 and AD2.




Figure 8-34 Fabric with two Admin Domains

Figure 8-35 shows how users get a filtered view of this fabric, depending on which Admin Domain they are in. As depicted in this diagram, users can see all switches and E_Ports in the fabric, regardless of their Admin Domain; however, the switch ports and end devices are filtered based on Admin Domain membership.

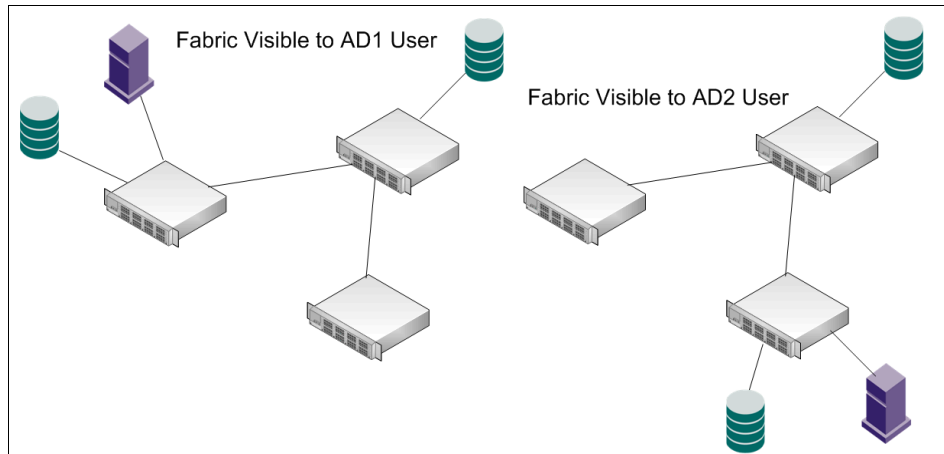


Figure 8-35 Filtered fabric views when using Admin Domains

8.5.1 Requirements for Admin Domains

This section lists the requirements for implementing Admin Domains in a fabric:

- ▶ Admin Domains are supported only on fabrics with one or more switches running Fabric OS v5.2.0 and later.
- ▶ If a switch runs Fabric OS earlier than v6.1.0, you must install Advanced Zoning license on it. On Fabric OS v6.1.0 or later, this license is no longer required (zoning became a part of base Fabric OS features).
- ▶ The default zone mode setting must be set to *No Access* before you create Admin Domains.
- ▶ Virtual Fabrics must be disabled before you create Admin Domains.
- ▶ To use Admin Domains and the FC-FC Routing Service in the same fabric, the switches connecting to the FC Router must be running Fabric OS v5.2.0 or later.
- ▶ If you are using LSAN zones:
 - Do not use LSAN zone names ending with `_AD n` (where n is the Admin Domain number).
 - Do not use LSAN zone names longer than 57 characters.
- ▶ You must be in the native operating mode to use Admin Domains because Admin Domains are not supported in interoperability mode.
- ▶ Gigabit Ethernet (GigE) ports cannot be members of an Admin Domain.

8.5.2 Creating an Admin Domain

The first step in creating an Admin Domain is to set the AD context to the physical fabric (AD255) this is done by selecting **Physical Fabric** in the Admin Domain drop-down menu as shown in Figure 8-36.



Figure 8-36 Admin Domain Selection menu

Click **Yes** in the confirmation dialog box (Figure 8-37).




Figure 8-37 Switch Admin Domain Context Dialog Box

To launch the Admin Domain window, click **Admin Domain** in the Web Tools Tasks panel. Figure 8-38 shows an example of the Admin Domain window.

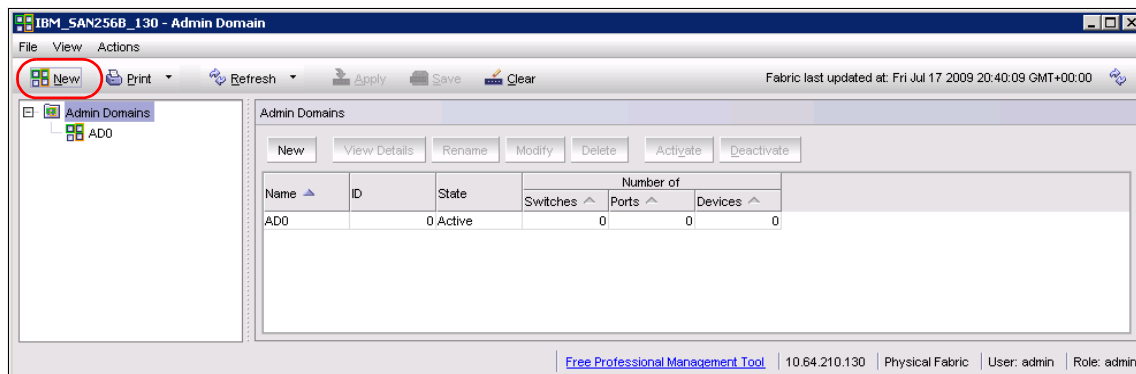


Figure 8-38 Admin Domain window

Only the system-defined AD0 is listed in the Admin Domain window. For this example, we want to create another Admin Domain and name it *AD2* using the following steps:

1. Click the **New** icon to launch the Admin Domain creation wizard. You need to provide basic information about the new Admin Domain (name and ID) as shown in Figure 8-39.

The screenshot shows the 'Create Admin Domain' wizard. The 'Steps' pane on the left indicates the current step is '1. Basic Information'. The main content area is titled 'Basic Information' and contains the following fields:

- Name:** The 'User Specified' radio button is selected, and the text box contains 'AD2'. The 'Auto Assigned (Based on numeric ID)' option is unselected.
- ID:** The 'User Specified' radio button is selected, and the text box contains '2'. The range '1 .. 254' is displayed next to the box. The 'Auto Assigned (Next available ID: 1)' option is unselected.
- State:** The 'Active' checkbox is checked.

At the bottom of the window, there are three buttons: 'Cancel', 'Previous', and 'Next'.

Figure 8-39 Create Admin Domain: Basic information

2. To add AD2 and, therefore, to populate the fields accordingly, continue by clicking **Next**.
3. Next, you select the membership. Select the switches, ports, and devices to form the new Admin Domain. As shown in Figure 8-40, we add ports 0-7 on Blade1 of the SAN256B to AD2 for this example. You need to highlight the members in the **Available Members** panel, hold down *Ctrl* to select multiple members, and then add these members to the **Selected Members** panel by clicking the **Add>** button.

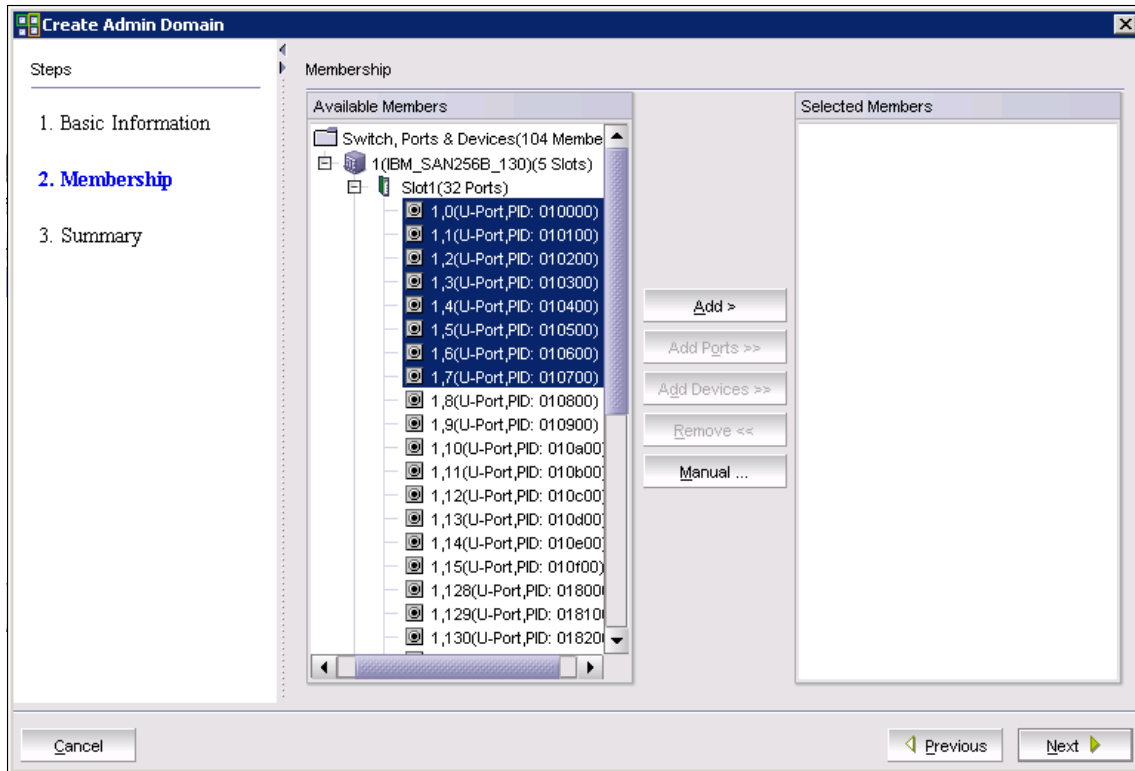


Figure 8-40 Create Admin Domain - membership selection

As shown in Figure 8-41, we can see that the members we chose have now been added to the Selected Members panel.

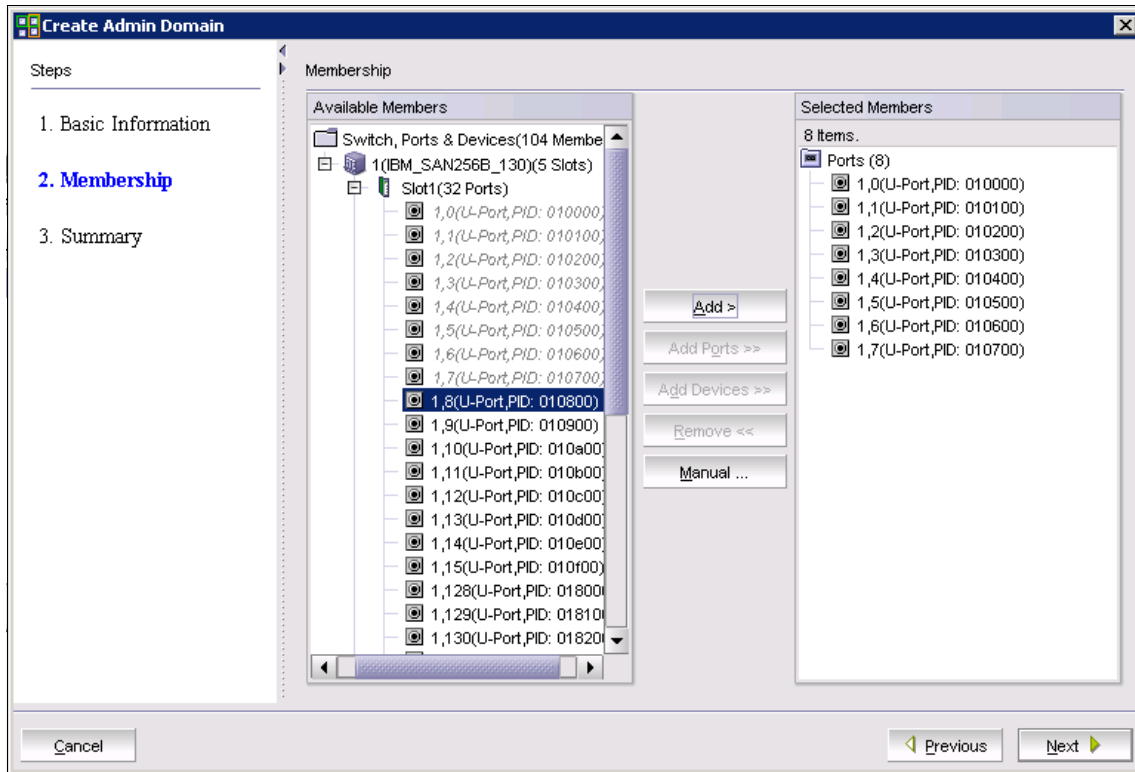


Figure 8-41 Create Admin Domain: Selected members

- Click **Next** to proceed to the Summary window, shown in Figure 8-42. You can now review the selections and go back to make corrections. When everything is correct, click **Finish**.

Create Admin Domain

Steps

1. Basic Information
2. Membership
- 3. Summary**

Summary

AD2

AD Name AD2 ID 2 State Active

Switches 0 Ports 8 Devices 0

Port Members (8)

Domain	Slot	Port	Port Index	Port Type	Port Name
1	1	0	0	U-Port	
1	1	1	1	U-Port	
1	1	2	2	U-Port	
1	1	3	3	U-Port	
1	1	4	4	U-Port	
1	1	5	5	U-Port	
1	1	6	6	U-Port	
1	1	7	7	U-Port	

Cancel Print Previous Finish

Figure 8-42 Create Admin Domain: Summary

5. The new Admin Domain AD2 is created and displays in the Admin Domain window (Figure 8-43). However, you need apply the AD configuration by clicking **Apply** icon to make the new Admin Domain effective.

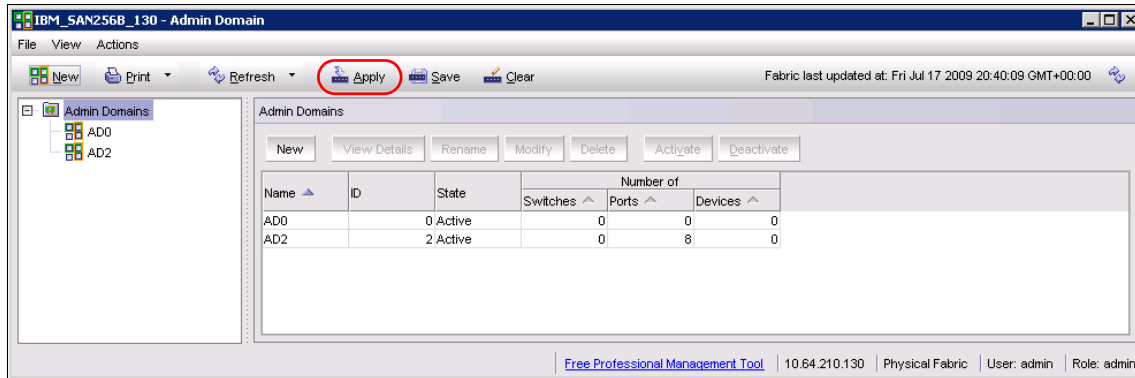


Figure 8-43 New Admin Domain AD2

6. Click **Yes** in the Apply AD Confirmation dialog box, Figure 8-44.

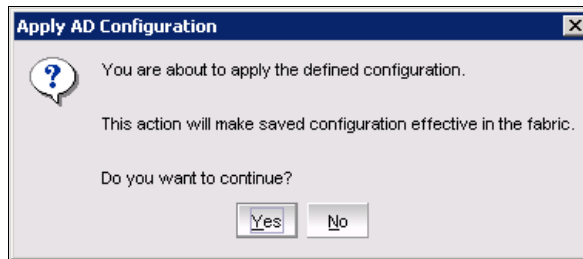


Figure 8-44 Apply AD Confirmation

7. When done, you can close the Admin Domain window to return to the main Web Tools window. The new Admin Domain AD2 is now available for selection in the Admin Domain pull-down menu, as shown in Figure 8-45.




Figure 8-45 Web Tools: Admin Domain selection for AD2

Admin Domains are used in Switch Administration User tab settings. See 8.7.7, “User tab” on page 276 for more details.

8.6 Port Admin task

To access the detailed port information, select the appropriate port on the switch (Switch View panel) or the Port Admin task (Tasks panel), as in Figure 8-46. This launches the Port Administration window Figure 8-47). From this window, you can select any of the switch ports to display the details.




Figure 8-46 Port Admin task

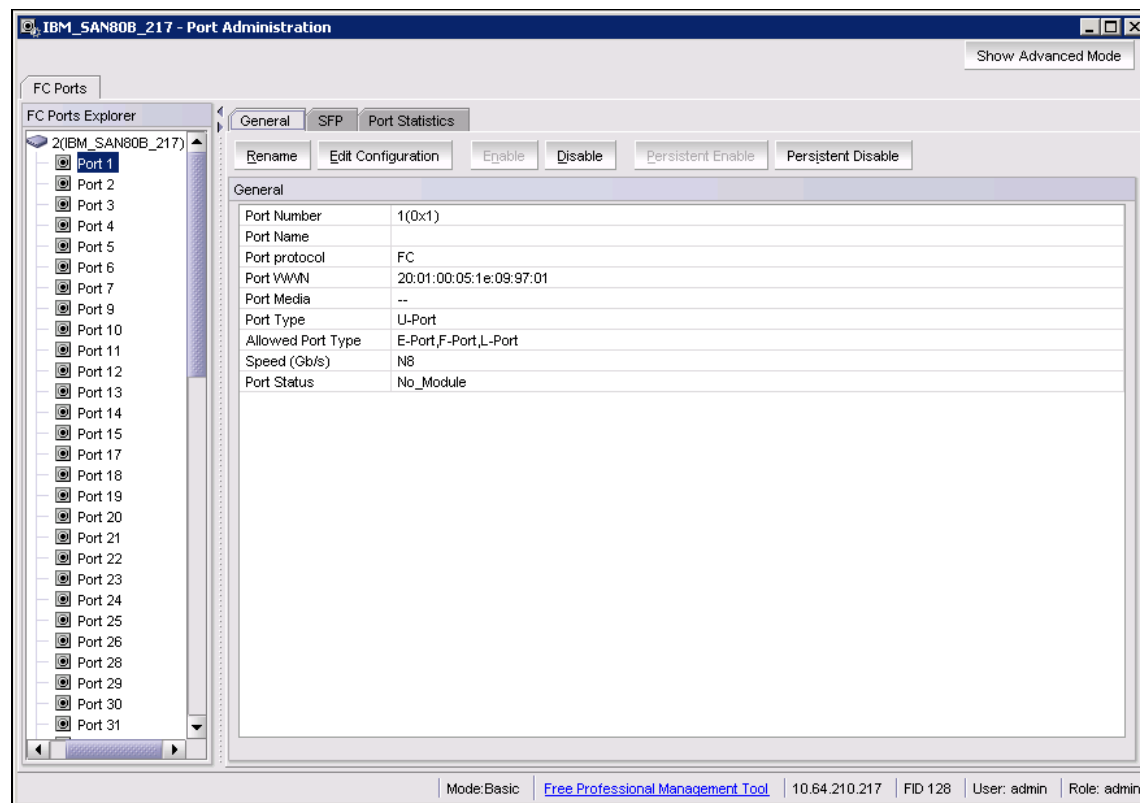


Figure 8-47 SAN80B Port Administration window

The window contains three tabs:

- ▶ General
- ▶ SFP
- ▶ Port Statistics

When the window opens, it displays the General tab. This tab shows basic information about the selected port. It also contains a set of buttons that you can use to perform a certain action on selected port. Buttons that are not applicable to the selected port are disabled. When multiple ports are selected, only the tasks that can be performed on all of the selected ports are displayed and the others are disabled.

Use the SFP tab to see detailed information about the SFP that is installed, as shown in Figure 8-48. This will not be visible if there is no SFP installed with the port in an empty state.

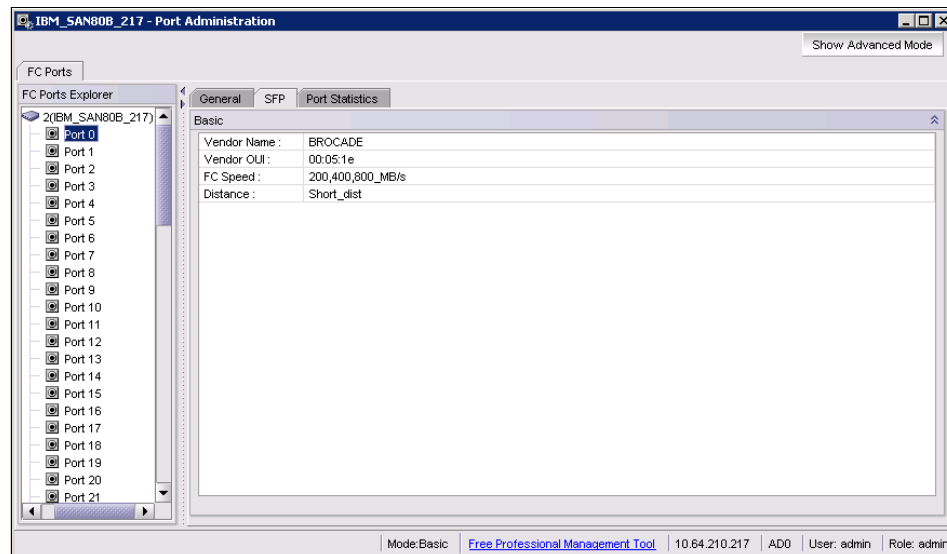


Figure 8-48 SAN32B-3 Port Administration SFP tab

Figure 8-49 shows the Port Statistics tab.

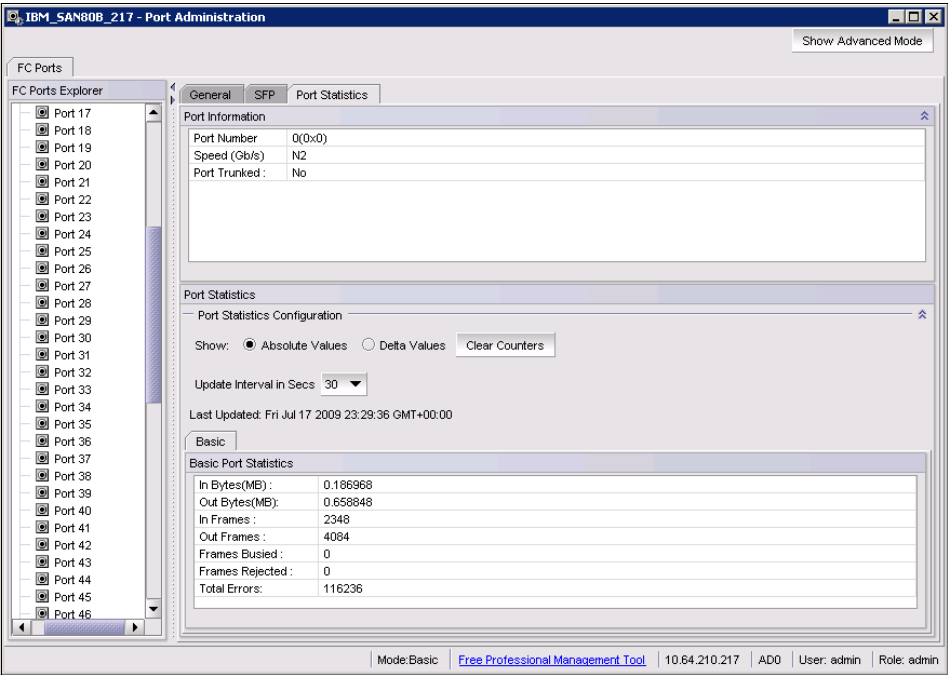


Figure 8-49 SAN32B-3 Port Administration - Port Statistics tab

You can view the port statistics for a specific port in the Basic or Advanced tabs, and all the errors display on the Error Details tab. Figure 8-50 shows these different tabs.

Basic

Advanced

Error Details

Basic Port Statistics

In Bytes(MB) :	0.186968
Out Bytes(MB):	0.658848
In Frames :	2348
Out Frames :	4084
Frames Busied :	0
Frames Rejected :	0
Total Errors:	116236

Basic

Advanced

Error Details

Advanced Port Statistics

C2 Frames Received	0
C3 Frames Received	2348
C3 Frames Discarded	0
Link Control Frames Received	0
Mcast Frames Received	0
Mcast Timeouts	0
Mcast Frames Transmitted	0
Time R_RDY Priority	0
Time BB_Credit Zero	0
Encd Errs Inside Frames	0
Encd Errs Outside Frames	116236
Frames with CRC Errs	0
Short Frames	0
Long Frames	0
Bad End-of-Frames	0

Basic

Advanced

Error Details

Error Details

Link Failure	0
Loss of Sync :	2
Loss of Signal	4
Protocol Error,	0
Invalid Transmitted Word	116236
Invalid CRC	0
Delimiter Error	0
Address Error	0
Inbound Link Reset	2
Outbound Link Reset	0
Inbound Offline Sequence	0
Outbound Offline Sequence	2

Figure 8-50 Port Statistics Basic, Advanced, and Error Details tabs

The SAN256B, SAN384B, and SAN768B have a slightly different display in which you can choose to see either the FC or the GigE ports. Figure 8-51 shows the Port Administration window on a SAN256B.

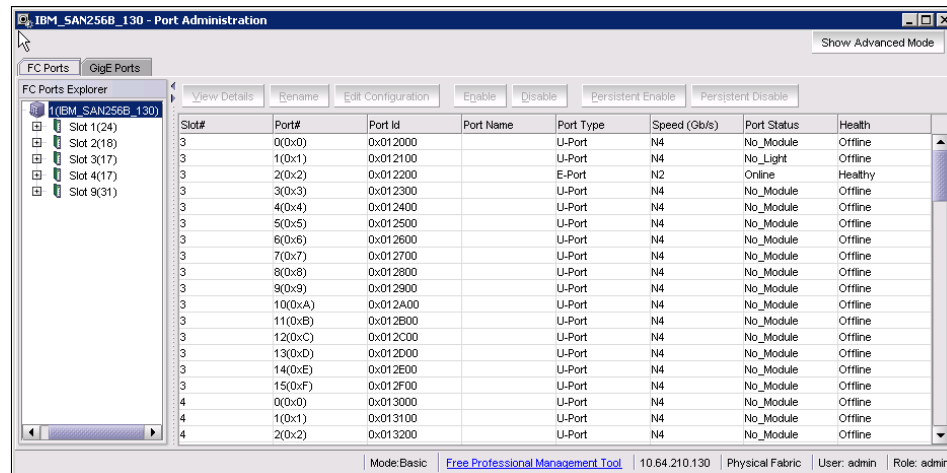


Figure 8-51 SAN256B Port Administration window

From the Port Administration window including the basic and advanced mode, you can perform the following functions:

- ▶ Rename a port
- ▶ Edit Configuration
- ▶ Enable or disable a specific port
- ▶ Persistent enable or persistent disable a port
- ▶ Enable or disable trunking for a specific port (default value is enabled)
- ▶ Enable or disable N_Port ID virtualization (NPIV)
- ▶ Port swap
- ▶ F_Port Trunking
- ▶ Re-Authenticate
- ▶ F_Port BB Credit
- ▶ QoS Enable/Disable
- ▶ Port Beacon enable /Disable
- ▶ WWN to N_Port mapping (Access Gateway)

(We explain the Configuration function in more detail in 8.6.2, “Editing the configuration” on page 237.)

Figure 8-52 shows the view of the General tab in advanced mode.

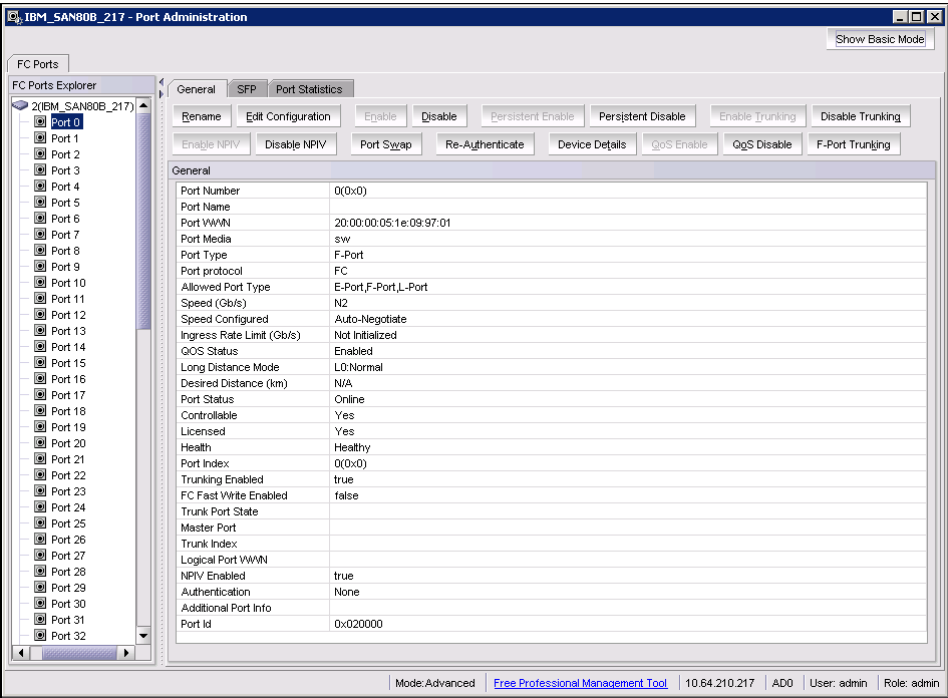


Figure 8-52 Advanced mode

Table 8-5 describes the fields for a specific port on the General Tab of the Port Administration window.

Table 8-5 Fields on the General Tab for a port

Field	Description
Port Number	The port number.
Port Name	Symbolic name assigned to the port.
Port Protocol	Type of protocol used on the port.
Port WWN	Port worldwide name.
Port Media	Type of Media connected to the port (sw or lw).
Port Type	Displays the current type of port.
Allowed Port Type	Configurable port types for the specific port.
Speed (Gbps)	Displays the actual speed at which the port is connected.
Speed Configured	Displays Speed the speed at which the port is configured.

Field	Description
Long Distance Mode	Shows the long distance mode selected for the port.
Desired Distance (km)	Shows the desired distance set for this port.
Port Status	Displays the current status of the port.
Controllable	Yes, if we can change port configuration settings.
Licensed	Shows whether the port is licensed, or requires POD license before it can be enabled and used.
Health	Displays the port health status.
Port Index	Shows whether the port has been swapped with another port.
Trunking Enabled	Displays trunking status.
FC Fast Write Enabled	Shows FC Fast Write status.
NPIV Enabled	Displays status of NPIV capability.
Additional Port Info	More information about the port.

Depending on the SAN switch type, additional fields might be present that contain information about Ingress Rate Limit (only supported on SAN768B and SAN384B), QoS, trunking attributes, and so on.

8.6.1 Renaming a port

Port naming is an optional feature to enable easier port management and identification. This option is available only for the FC and FCIP virtual ports and is not available for GE ports. From Port Admin Tasks, we can rename the port by clicking the **Rename** button. It can also be done by right-clicking the port from the switch view of Web Tools and clicking **Configure** → **Rename**, where we can enter a name for the port.

8.6.2 Editing the configuration

This function allows you to view and change various port parameters. Follow these steps:

1. Select the port in the left pane of the Port Administration window and then click **Edit Configuration**.

When editing the configuration of a port that is online, you are presented with a confirmation dialog box as shown in Figure 8-53.




Figure 8-53 Online Port Edit Configuration Confirmation

2. Click **Yes** to accept the confirmation, and you see the FC Port Configuration Wizard, shown in Figure 8-54. The first step is to configure the Allowed Port Types.

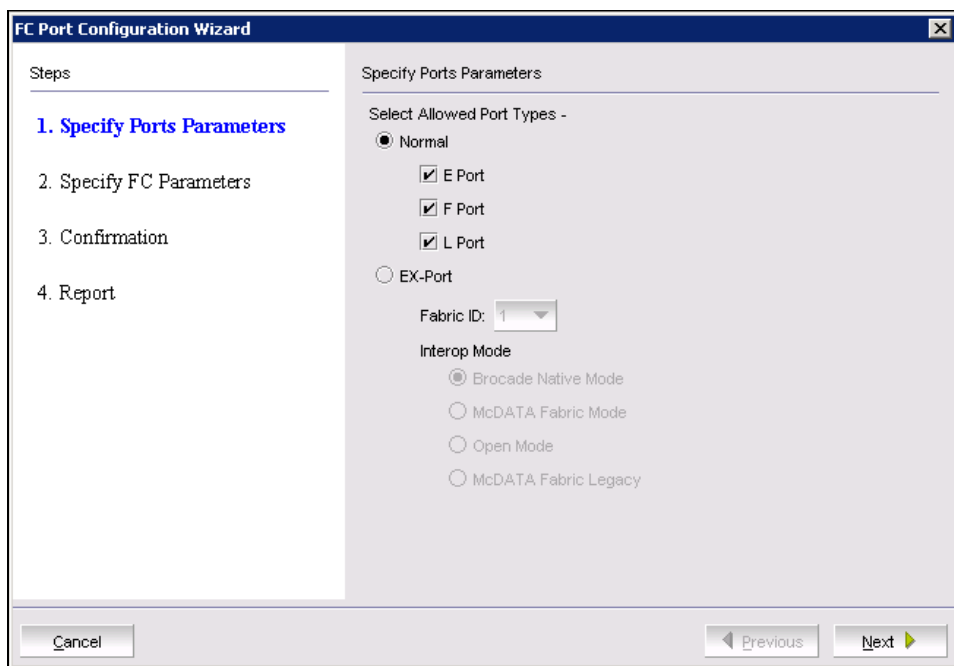


Figure 8-54 FC Port Configuration Wizard Port Parameters

3. Set the allowed port types as required, and click **Next** to continue.

4. Now, you set FC parameters. As shown in Figure 8-55, there are four parameters that can be set:
- Speed: Port speed can be set the speed to Auto, 1G, 2G, 4G, and 8G.
 - Ingress Rate Limit, Ingress rate limiting is a licensed feature that requires the Adaptive Networking license. Ingress rate limiting restricts the speed of traffic from a particular device to the switch port. The valid values measured in megabits per second (Mbps) are: 200, 400, 600, 800, 1000, 1500, 2000, 2500, 3000, 3500, 4000, 5000, 6000, 7000, and 8000.
 - Long Distance Mode: This sets the mode for long distance ports. The following values are valid:
 - L0: Normal
 - LE: <=10KM
 - LD: Auto
 - LS: Static
 - Desired Distance: This is used if the Long Distance mode is set to LD. LD calculates buffer credits based on the distance measured during port initialization. An upper limit is placed on the calculation by providing a desired distance value. If the measured distance is more than the desired distance, the desired distance is used in the calculation; otherwise, the measured distance is used.

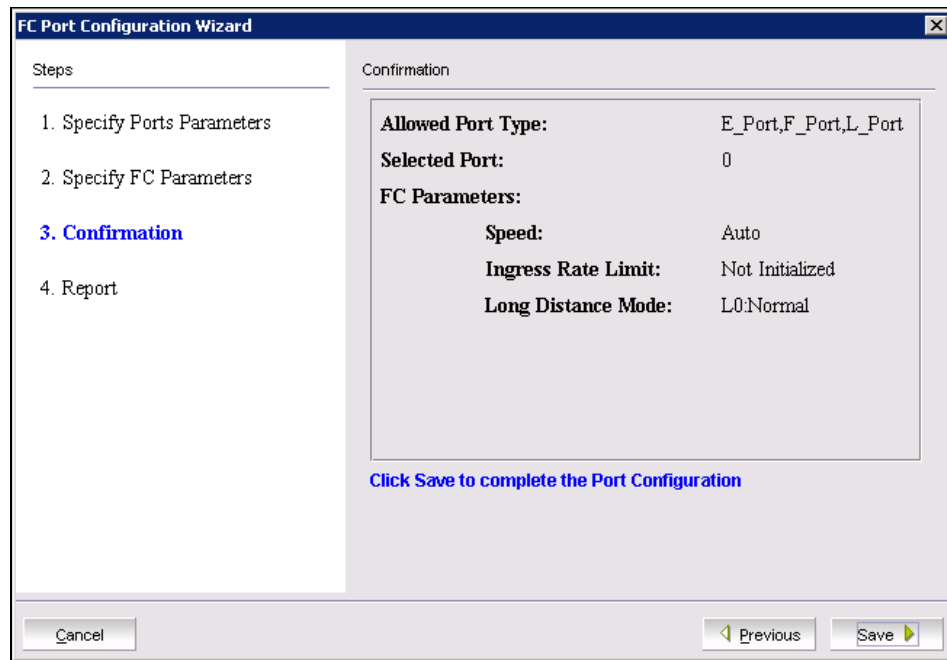
The screenshot shows a window titled "FC Port Configuration Wizard" with a close button in the top right corner. On the left, a "Steps" pane lists four steps: "1. Specify Ports Parameters", "2. Specify FC Parameters" (highlighted in blue), "3. Confirmation", and "4. Report". The main area is titled "Specify FC Parameters" and contains four configuration items, each with a label and a control:

Parameter	Value
Speed	Auto (dropdown menu)
Ingress Rate Limit(Mb/s)	Not Initialized (dropdown menu)
Long Distance Mode	L0:Normal (dropdown menu)
Desired Distance (km)	N/A (text input field)

At the bottom of the window, there are three buttons: "Cancel" on the left, and "Previous" and "Next" on the right, with arrow icons.

Figure 8-55 FC Port Configuration Wizard FC Parameters

5. After changing the appropriate fields, click **Next** to display the Confirmation window, as shown in Figure 8-56.



The image shows a software window titled "FC Port Configuration Wizard". It has a "Steps" pane on the left and a "Confirmation" pane on the right. The "Steps" pane lists four steps: "1. Specify Ports Parameters", "2. Specify FC Parameters", "3. Confirmation" (which is highlighted in blue), and "4. Report". The "Confirmation" pane displays the following configuration details:

Allowed Port Type:	E_Port,F_Port,L_Port
Selected Port:	0
FC Parameters:	
Speed:	Auto
Ingress Rate Limit:	Not Initialized
Long Distance Mode:	L0:Normal

Below the table, there is a blue text instruction: "Click Save to complete the Port Configuration". At the bottom of the window, there are three buttons: "Cancel", "Previous", and "Save".

Figure 8-56 FC Port Configuration Wizard Confirmation

6. Review the configuration changes and click **Save** to complete the port configuration. You are presented with a success page, as shown in Figure 8-57. Click **Close**.

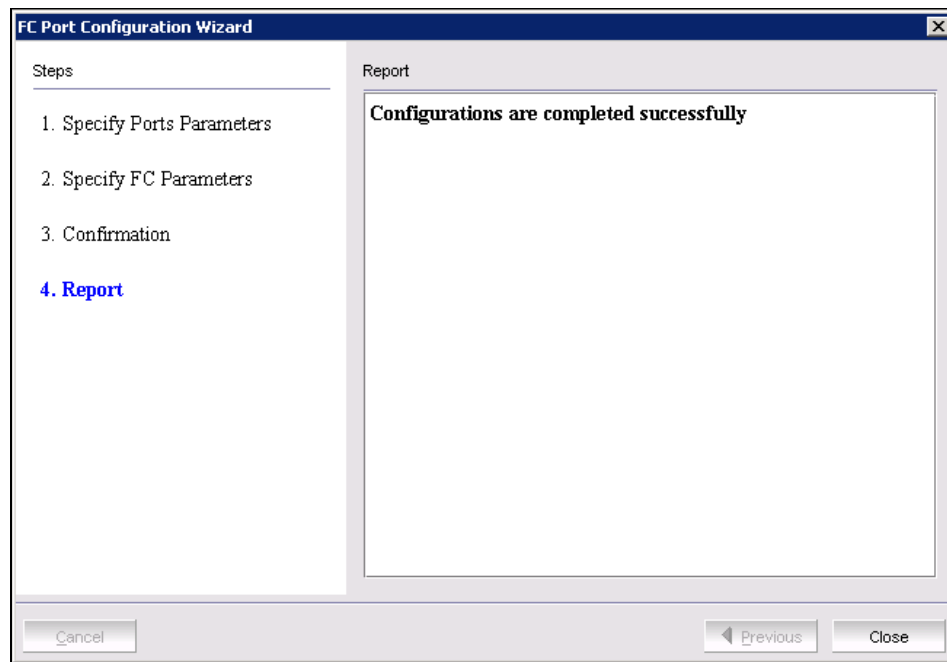


Figure 8-57 FC Port Configuration Wizard Success

8.6.3 Enabling and disabling a port

From port admin tasks we can enable or disable a port by clicking the **Port enable** or **Port disable** buttons. Some configuration changes, such as port swapping, require that you disable and enable the port. This can also be done from the switch view in Web Tools by right-clicking the port and then **Configure** → **Enable/Disable**.

8.6.4 Persistent enable and persistent disable options for a port

The Persistent disable and enable option makes it possible to have ports in an enabled or disabled state across restarts. This can be done from the switch view of Web Tools by right-clicking the port, then selecting **Configure** → **Persistent Enable/Disable**. From the Port Admin Tasks view, we can also do this by clicking the **Persistent enable** or **Persistent disable** button.

8.6.5 Enabling or disabling trunking for a specific port

Trunking is enabled by default for a port, which enables an ISL connected from the same port group to form a trunk. This is described more detail in Chapter 13, “Multiple switches and fabrics” on page 587. If required, we can manually either enable or disable the trunking feature on a port from the port admin tasks advanced mode view by clicking **Enable Trunking** or **Disable Trunking**.

8.6.6 Enabling or disabling N_Port ID virtualization (NPIV)

This feature is available when the NPIV license is installed. From Web Tools, you can either enable or disable NPIV for a port by clicking the **Enable NPIV** or **Disable NPIV** button in the port admin task in advanced mode. This is a mandatory feature when Access Gateway mode is enabled.

8.6.7 Port swap

We might need to perform a port swap without affecting the host end device configurations. In such cases this port swap option allows us to maintain the same port index or port ID in the new port so that we can change the connectivity to the new port without having to change the port ID. While doing this in a backbone or director, we need to input the new slot number along with the new port number. Figure 8-58 shows the option to input the new port number when we click **Port Swap** from the port administration window.

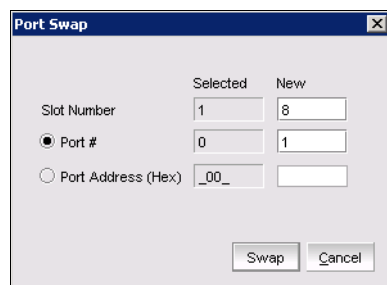
A screenshot of a 'Port Swap' dialog box. The dialog has a title bar with 'Port Swap' and a close button. Inside, there are two columns: 'Selected' and 'New'. Under 'Selected', there is a 'Slot Number' field with the value '1'. Below that, there are two radio buttons: 'Port #' (which is selected) and 'Port Address (Hex)'. The 'Port #' radio button has a corresponding field with the value '0'. The 'Port Address (Hex)' radio button has a corresponding field with the value '_00_'. Under the 'New' column, there is a field with the value '8' corresponding to the 'Slot Number' field, and a field with the value '1' corresponding to the 'Port #' radio button. At the bottom right, there are two buttons: 'Swap' and 'Cancel'.

Figure 8-58 Port swap task

8.6.8 F_Port Trunking

This is a preferred option used for Nodes and targets connected by switches in Access Gateway mode. F_Port trunking requires the following criteria to be met:

- ▶ The Enhanced Group Management (EGM) license is installed.
- ▶ Trunking needs to be enabled in the port.

- ▶ Access gateway switch must have a trunking license.
- ▶ The port must not be configured for Long distance connection.
- ▶ The port must not be swapped.

When we create an F_Port trunk, a logical unit called a Trunk Index (TI) is formed that represents the physical ports in the trunk. From Advanced mode, when we select **F_Port trunking**, it will prompt us to select the ports as shown in Figure 8-59.




Figure 8-59 F_Port Trunk group creation

Then click **OK** to Create Trunk Group and select the **Trunk Index** as shown in Figure 8-60.




Figure 8-60 F_Port Trunk Index

By selecting the **Trunk Index** and clicking **OK**, we get the trunk group as shown in Figure 8-61. From here we click **OK** and then **Yes** to confirm it.

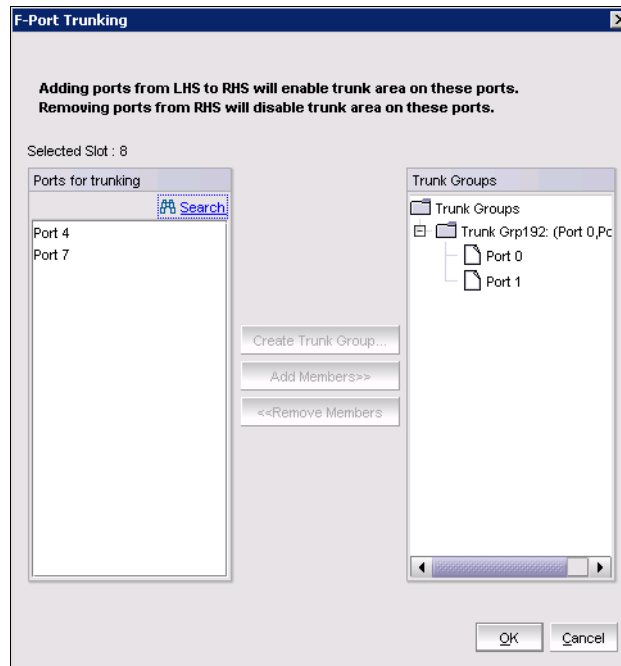


Figure 8-61 F_Port trunking

8.6.9 Re-authenticating

Devices or switches connected by F_Port or E_Port can be re-authenticated from the Port Admin Tasks view if DH-CHAP secrets or PKI certificates are set with switch level security policy. If the security features are not defined and available, this test of authentication will fail and the port will be disabled. This is described more in Chapter 14, “Security” on page 623

8.6.10 F_Port BB credit

BB credit for an F_Port can be defined from this advanced mode option. The default value is 8. This option allows us to change the default value if required for an F_Port by clicking the **F port BB credit** button, and it will prompt for a new value.

8.6.11 QoS Enable/Disable

This option requires the Adaptive Networking licensed feature. We can enable or disable QoS for a port by clicking **Qos Enable** or **Qos Disable** from Port Admin Advanced mode. More details are provided in Chapter 15, “Adaptive Networking” on page 695.

8.6.12 Port beaconing

Port beaconing enables easy port identification and can be enabled or disabled from the port admin task by clicking **Port Beacon Enable** or **Port Beacon Disable**. When enabled, the port blinks amber and green for 2.5 seconds each, which will be visible from the switch view.

8.6.13 WWN to N_Port mapping

In Access Gateway mode, the facility to map a WWN to an N_Port has been introduced. When a switch is in Access Gateway mode, the port admin task has two options to “Configure N-Port groups” and “Configure F-N port mappings.” For this feature of WWN to an N_Port, we click **Configure F-N port mappings** as shown in Figure 8-62. All other port configuration options in Port Admin are disabled because the switch is in Access Gateway mode.

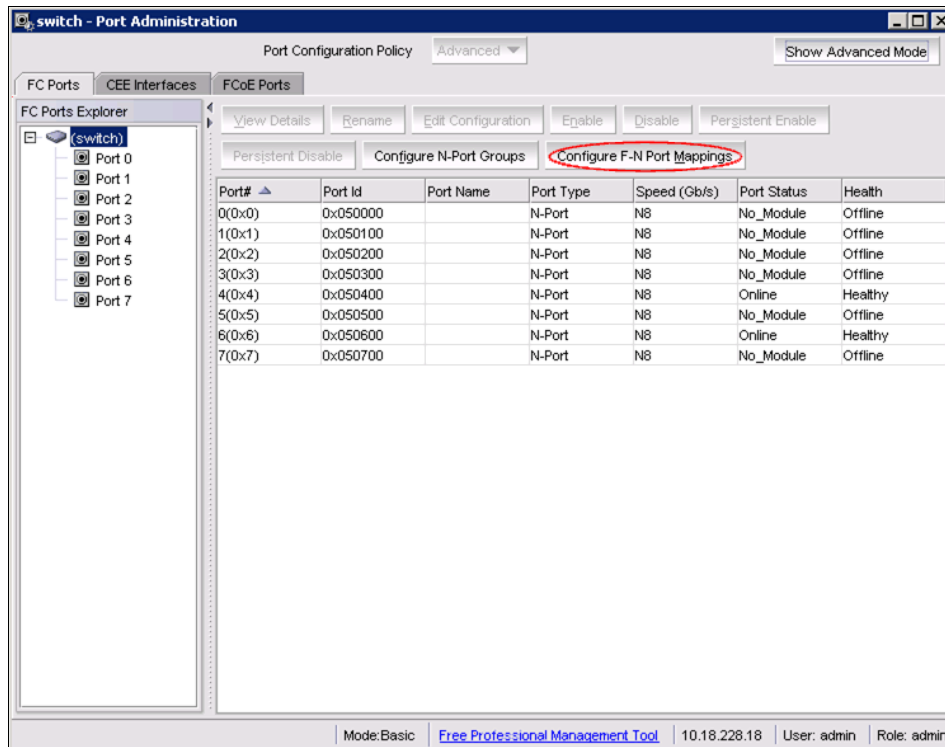


Figure 8-62 Port admin tasks for Access Gateway mode

When we click **Configure F-N port Mappings**, it gives us the N_Port mapping configuration with a list of the mapping groups defined as shown in Figure 8-63.

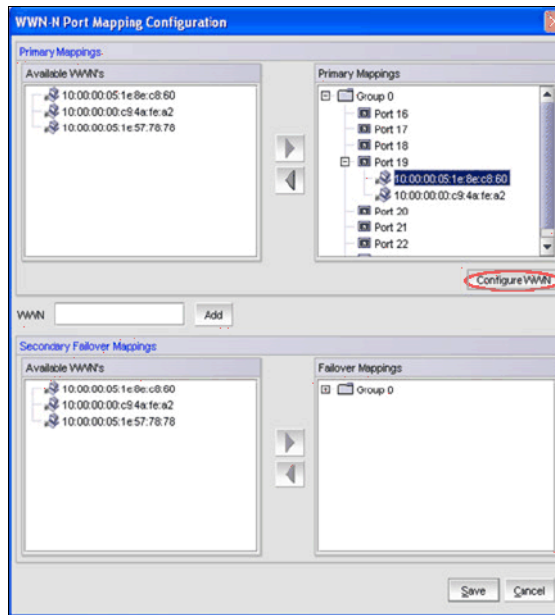


Figure 8-63 WWN N_Port mapping

Here we can enable WWN mapping by selecting a WWN and then click **Configure WWN**, which gives us the option to select **Enable Rule for WWN configuration** as shown in Figure 8-64.




Figure 8-64 WWN config enable

8.6.14 Port Administration window on the SAN256B and SAN768B

While most of the buttons in the Port Administration window perform the same function on all switches, the SAN256B, SAN384B, and SAN768B have two tabs on the left pane. Figure 8-65 shows the first tab, which is for FC ports.

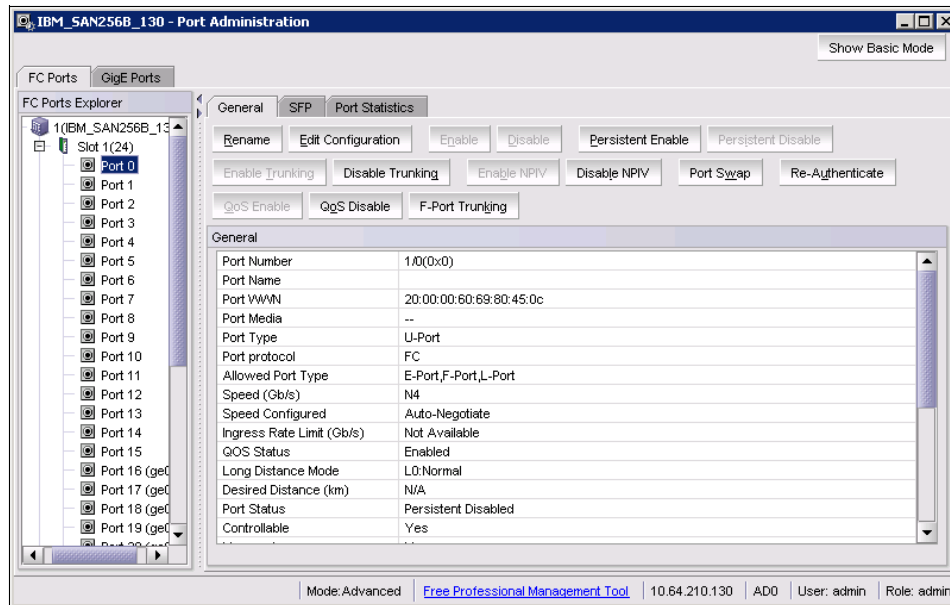


Figure 8-65 SAN256B Port Administration FC Ports tab

Figure 8-66 shows the GigE ports tab in the Port Administration window of SAN256B.

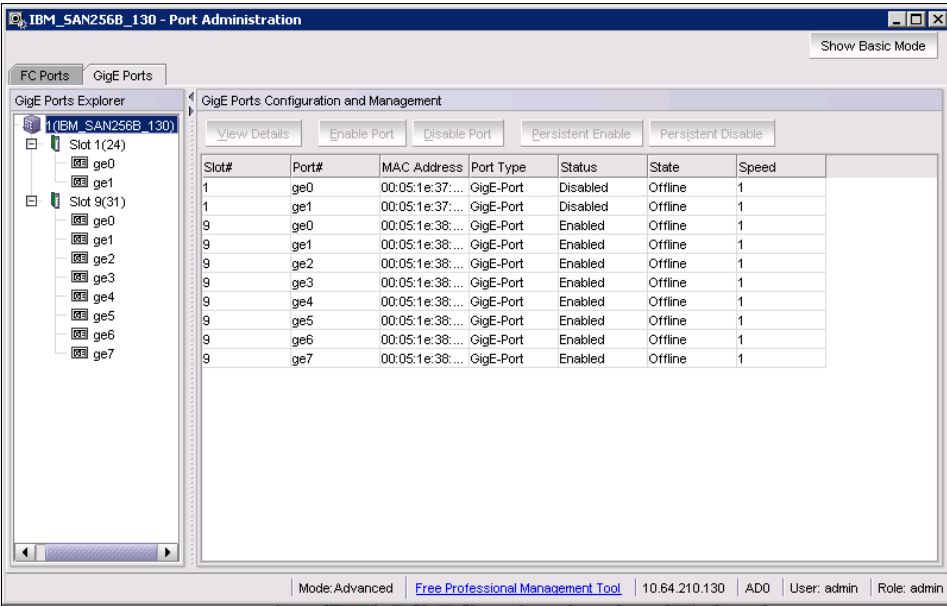


Figure 8-66 SAN256B Port Administration GigE Ports tab

8.6.15 Port Administration for the FCOE switch

This topic is described in the book, *IBM Converged Switch B32*, SG24-7935-00, available at this website:

<http://www.redbooks.ibm.com/redpieces/abstracts/sg247935.html?Open>

8.6.16 Port Administration for the IBM System Storage SAN06B-R

This topic is described in the book, *IBM System Storage b-type Multiprotocol Routing: An Introduction and Implementation*, SG24-7544-03, available at this website:

<http://www.redbooks.ibm.com/abstracts/sg247544.html?Open>

8.7 Switch Admin task

The Switch Admin task on the Tasks panel (Figure 8-67) is used to launch the Switch Administration window.




Figure 8-67 Switch Admin task

In this section, we discuss the Switch Administration in detail.

8.7.1 Switch Administration window layout

In this example, we explore the Switch Administration on a SAN80B. When the administration window opens, it is composed of five areas (labeled A, B, C, D, and E), as shown in Figure 8-68.

Tip: If you hover the mouse over buttons and other areas of the window, information displays about their function.

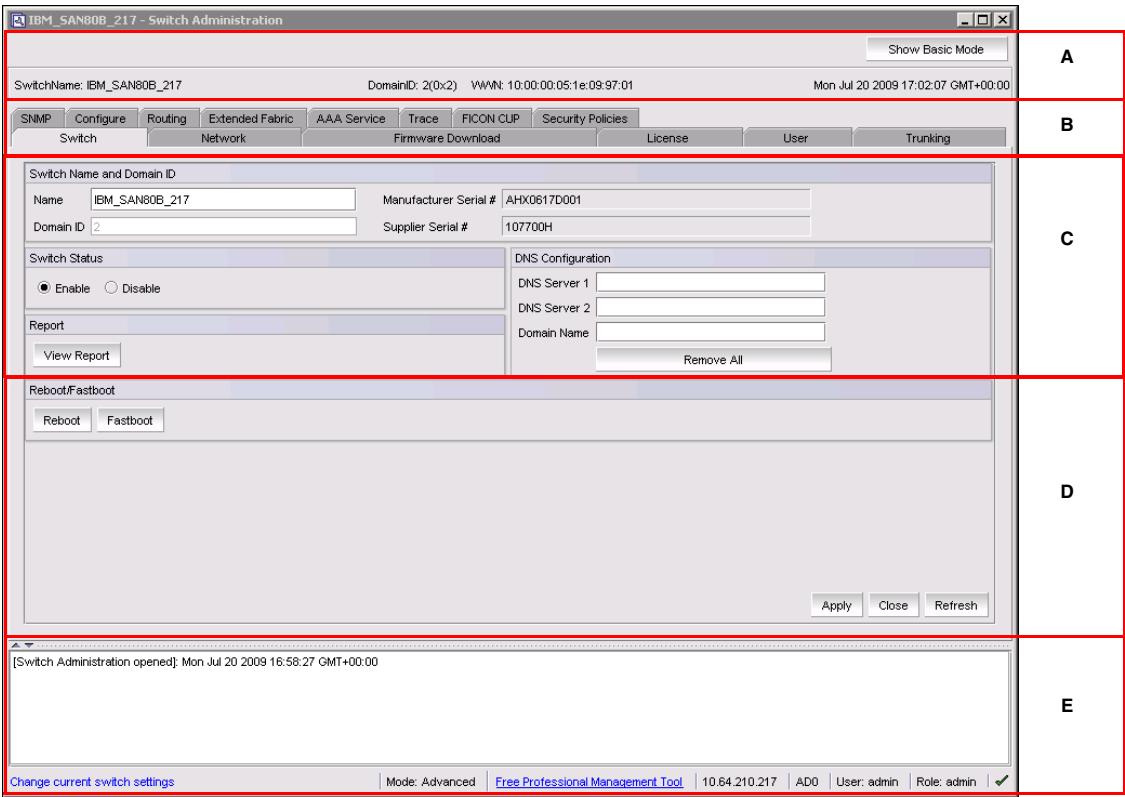


Figure 8-68 SAN80B administration window layout

Here we describe the five areas of the window:

- ▶ Area A: Displays summary information, switch name, domain ID, date, and time. You can use the button on the right side to switch between Basic and Advanced Mode display.
- ▶ Area B: Allows navigation through the different management panels (by clicking the desired tab). The content of this area depends on the licenses installed on the switch. In addition, content also depends on switch administration mode:
 - In basic mode, you see only the seven most commonly used tabs (Switch, Network, Firmware Download, License, User, Blade (on chassis based systems), and Trunking).
 - Advanced mode displays all available tabs.
- ▶ Area C: Contains parameters to be set in the current panel.
- ▶ Area D: Contains the button bar.
- ▶ Area E: Contains the report window that allows viewing of the switch report upon operation completion.

We describe the Switch Administration tabs in the sections that follow.

8.7.2 Switch tab

When the administration window first opens, the Switch tab displays by default, as shown in Figure 8-68 on page 251.

On the first tab, you can define the switch name and the domain ID, set the DNS configuration, enable or disable the entire switch, and view a detailed report of the switch.

You can also perform the following actions:

- ▶ Reboot switch
- ▶ Fastboot switch

Table 8-6 describes the fields on the Switch tab.

Table 8-6 Switch Information tab

Field	Description
Name	Enter data for the switch name. Enter a new name to change a name in this field.
Domain ID	Displays or sets switch domain ID. Domain IDs must be unique within a fabric. To change domain ID, enter new domain ID in this field. Use a number from 1 to 239 for normal operating mode (FCSW compatible) and a number from 0 to 31 for VC encoded address format mode (backward compatible to SilkWorm 1000 series). Note: The switch needs to be disabled to change this value.
Manufacturer Serial #	Physical serial number of the switch.
Supplier Serial #	Supplier serial number of switch for display only.
(Status) Enable	Click the radio button to enable the switch.
(Status) Disable	Click the radio button to disable the switch.
DNS Server 1	Enter the primary DNS server in this field.
DNS Server 2	This is the field for secondary DNS server.
Domain Name	Enter the Domain Name.
Reboot	Click to reboot the switch.
Fastboot	Use this button to perform fastboot.
Apply	Click to save any changes made to this tab and remain in the current tab. You can make additional changes and click Apply when making changes incrementally.
Close	Click to exit the Switch Admin view. If you make changes but do not commit them by clicking Apply , a dialog box is presented to allow you to commit or delete the changes.
Refresh	Click to retrieve current values from the switch.

Click **View Report** to display a window as shown in Figure 8-69. The detailed report includes a list of all the types of switches connected to the local switch, the inter-switch links (ISLs), list of ports, the Name Server information, details on the configured zones, and SFP serial ID information.

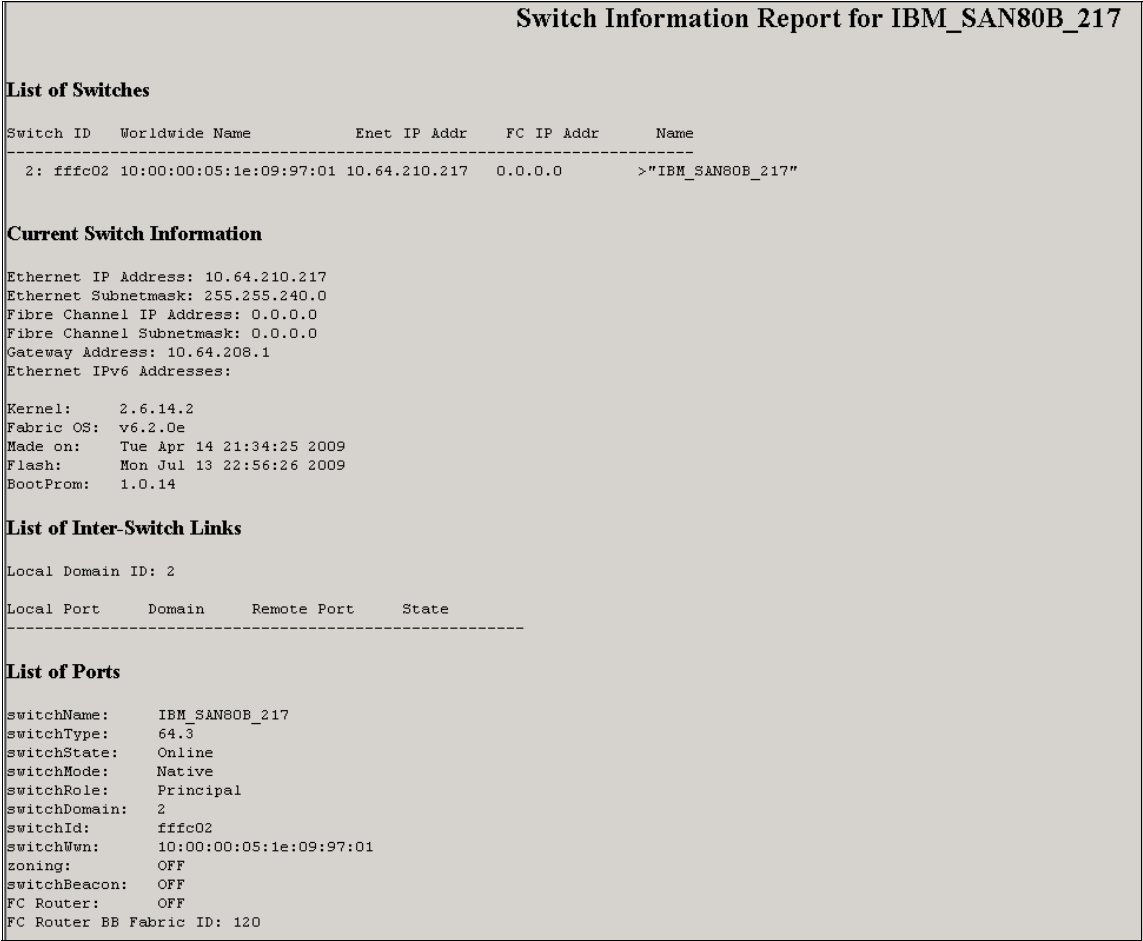


Figure 8-69 SAN80B Switch report

8.7.3 Network tab

Use the Network tab to modify the IP settings of the switch. Figure 8-70 displays the network tab for the SAN80B switch.

The screenshot displays the 'Network' tab of the 'IBM_SAN80B_217 - Switch Administration' window. The window title bar includes a 'Show Advanced Mode' button. The top status bar shows 'SwitchName: IBM_SAN80B_217', 'DomainID: 2(0x2)', 'WWN: 10.00.00.05.1e.09.97.01', and the date/time 'Mon Jul 20 2009 17:31:07 GMT+00:00'. Below the status bar are tabs for 'Switch', 'Network', 'Firmware Download', 'License', 'User', and 'Trunking'. The 'Network' tab is active, showing the 'Ethernet and Fibre Channel IP Configuration' section. This section includes fields for 'IPv4 Address' (Ethernet IP: 10.64.210.217, Ethernet Mask: 255.255.240.0, Gateway IP: 10.64.208.1), 'Fibre Channel Net IP' (0.0.0.0), and 'Fibre Channel Net Mask' (0.0.0.0). There is also a 'DHCP: Disabled' checkbox. Below this is the 'IPv6 Address' section with an 'Ethernet IPv6' field and a 'State' dropdown. A checkbox for 'Enable IPv6 Auto Configuration' is present, followed by an 'Auto Configured IPv6 Addresses' table with columns 'IP Address', 'Prefix', and 'State'. Below that is an 'IPv6 Gateways' section with an 'IP Address' field. The 'Syslog IP's Configuration' section features a table with 'Syslog IP' and 'Current Value' columns, listing six IP addresses. To the right of this table is a 'New IP' field and buttons for 'Add', 'Remove', and 'Clear All'. At the bottom right of the configuration area are 'Apply', 'Close', and 'Refresh' buttons. The footer bar contains a link 'Configure IP Addresses for Ethernet and Fibre channels.', 'Mode: Basic', a link 'Free Professional Management Tool', the IP '10.64.210.217', 'AD0', 'User: admin', 'Role: admin', and a green checkmark icon.

Syslog IP	Current Value
1	10.64.223.12
2	10.64.210.36
3	10.64.210.35
4	10.64.210.101
5	10.64.210.102
6	10.64.210.100

Figure 8-70 SAN80B Network tab

Table 8-7 describes the fields in the Network tab (which is shown in Figure 8-71 on page 258).

Table 8-7 Network tab

Field	Description
Ethernet IP	Display or set the Ethernet IP address.
Ethernet Mask	Display or set the Ethernet IP Subnet Mask.
Gateway IP	Display or set the Gateway IP address.
Fibre Channel Net IP	Display or set the Fibre Channel IP address.
Fibre Channel Net Mask	Display the Fibre Channel Subnet Mask.
IPv6 Address	Set the IPv6 address, if required.
Enable IPv6 Auto Configuration Check box	When IPv6 auto-configuration is enabled, the platform will engage in stateless IPv6 auto-configuration. When IPv6 auto-configuration is disabled, the platform will relinquish usage of any auto-configured IPv6 addresses that it might have acquired while it was enabled.
Syslog IPs	Display the six syslog IP addresses for a user to configure.
Add	Add syslog IP address entered in field.
Remove	Remove syslog IP address.
Clear All	Remove all previous syslog IP entries.
Apply	Click to save the changes made to this tab and to stay in the current tab. You can make additional changes and click Apply when making changes incrementally.
Close	Click to exit the Admin window. If you make changes but do not commit them by clicking Apply , a dialog box displays to allow you to commit or delete the changes.
Refresh	Click to retrieve current values from the switch.

Upgrades: An upgrade from Fabric OS v6.1 or earlier, which does not support IPv6 auto-configuration, to a platform that does support IPv6 auto-configuration, such as Fabric OS v6.2 or later, will cause IPv6 auto-configuration to be enabled on the upgraded platform. In upgrades or downgrades between versions of Fabric OS that support auto-configuration, the enabled state of IPv6 auto-configuration will not be changed.

Overview of syslogd

The Fabric OS maintains an internal log of all error messages. However, the internal log buffers are limited in capacity. When they are full, new messages overwrite old messages.

You can configure the switch to send error log messages to a UNIX host system that supports **syslogd**. You can configure this host system to receive error or event messages from the switch and then store them in its file system, overcoming the size limitations of the internal log buffers on the switch.

The host system can be running UNIX, Linux, or any other operating system as long as it supports standard **syslogd** functionality. The switch by itself does not assume any particular operating system to be running on the host system.

To configure the syslog function, simply put the IP address of the host running the **syslogd** in the Syslog IP field, and click **Add**. After adding all logging host IP addresses to the list, you must click **Apply** to save the changes.

Network tab: SAN256B and SAN768B

The network configuration panel on a SAN256B director or SAN768B backbone looks slightly different. As shown in Figure 8-71, the panel contains an additional section called *Advanced IP Configuration*. This section allows you to configure the Ethernet management ports on each CP.

Figure 8-71 shows the Network tab for the SAN256B.

IBM_SAN256B_130 - Switch Administration

Show Advanced Mode

SwitchName: IBM_SAN256B_130 DomainID: 1(0x1) WWN: 10:00:00:60:69:80:45:0c Mon Jul 20 2009 17:40:03 GMT+00:00

Switch Network Firmware Download License User Blade Trunking

Ethernet and Fibre Channel IP Configuration

IPv4 Address

Ethernet IP 10.64.210.130 Fibre Channel Net IP 0.0.0.0

Ethernet Mask 255.255.240.0 Fibre Channel Net Mask 0.0.0.0

Gateway IP 10.64.208.1

IPv6 Address

Ethernet IPv6 State :

Advanced IP Configuration

IPv4 Address

CP0 Ethernet IPv4 10.64.210.131 CP1 Ethernet IPv4 10.64.210.132

CP0 IPv4 Subnet Mask 255.255.240.0 CP1 IPv4 Subnet Mask 255.255.240.0

IPv6 Address

CP0 Ethernet IPv6 State :

CP1 Ethernet IPv6 State :

☐ Enable IPv6 Auto Configuration

Auto Configured IPv6 Addresses

IPv6 Gateways

Syslog IP's Configuration

Syslog IP	Current Value
1	10.64.210.105
2	10.64.210.100
3	10.64.210.45
4	10.64.210.103

New IP

Add

Remove

Clear All

Apply Close Refresh

Configure IP Addresses for Ethernet and Fibre c... Mode: Basic [Free Professional Management Tool](#) 10.64.210.130 ADO User: admin Role: admin ✓

Figure 8-71 Network tab of the SAN256B

258 Implementing an IBM b-type SAN with 8 Gbps Directors and Switches

8.7.4 Firmware Download tab

We use the Firmware Download tab to upgrade the Fabric OS. The firmware upgrade procedure normally requires an FTP server that stores the Fabric OS files. Additionally, if the switch is equipped with a USB port, you can use a Brocade-branded USB memory key as the source for firmware download. Figure 8-72 shows an example of the Firmware Download panel on a IBM Converged B32 switch.

The screenshot shows the 'switch - Switch Administration' web interface. The 'Firmware Download' tab is selected. The interface displays the following information and controls:

- Switch Information:** SwitchName: switch, DomainID: 5(0x5), VVNN: 10:00:00:05:1e:b0:81:80, Mon Nov 08 2010 23:55:56 GMT+00:00.
- Navigation Tabs:** SNMP, Configure, Routing, Extended Fabric, AAA Service, Trace, FICON CUP, Security Policies, Switch, Network, Firmware Download (selected), License, User, Trunking, CEE.
- Current Firmware Information:** Primary partition: 6.4.0b, Secondary partition: 6.4.0b.
- Firmware Key Information:** 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".
 - Host Name or IP: 10.18.228.37
 - User Name: admin
 - Password: ••••••••
 - Protocol Type: File Transfer Protocol (FTP)
 - Specify Firmware Path: Firmware/switches/6.x/6.4.x/6.4.1/
- Buttons:** Download, Close, Refresh.
- Log:** [Switch Administration opened]: Mon Nov 08 2010 23:33:53 GMT+00:00, [Switch Administration closed]: Mon Nov 08 2010 23:34:53 GMT+00:00, [Switch Administration closed]: Mon Nov 08 2010 23:34:53 GMT+00:00, [Switch Administration opened]: Mon Nov 08 2010 23:53:56 GMT+00:00.
- Footer:** Firmware download, Mode: Advanced, [Free Professional Management Tool](#), 10.18.228.18, AD0, User: admin, Role: admin, ✓.

Figure 8-72 Firmware Download tab

Uploading switch configuration before changing firmware

The IBM Converged B32 switch has a USB port, therefore, the USB radio button is active. Fabric OS v6.4.0 allows you to download either the actual firmware or the public firmware key.

Always upload a copy of the switch configuration before performing any firmware change. The configuration upload function is available in the **Configure** tab (described in 8.7.8, "Configure tab" on page 281).

Upgrading the firmware using Web Tools

In this section, we show the firmware upgrade procedure with Web Tools. In our example, we upgrade the firmware of IBM Converged B32 switch from v6.4.0b to v6.4.1. The firmware versions are stored on an FTP server.

To upgrade the firmware, proceed as follows:

1. Launch Web Tools for IBM Converged B32 switch and select **Switch Admin** as shown in Figure 8-73.




Figure 8-73 IBM Converged SAN B32 switch Web Tools

- From the Switch Administration window. Select the **Firmware Download** tab, as shown in Figure 8-74.

switch - Switch Administration

Show Advanced Mode

SwitchName: switch DomainID: 5(0x5) VVNN: 10.00.00.05:1e:b0:81:80 Tue Nov 09 2010 00:36:56 GMT+00:00

Switch Network Firmware Download License User Trunking CEE

Current Firmware Information

Primary partition: 6.4.0b
Secondary partition: 6.4.0b

Firmware Key Information

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"

Host Name or IP: 10.18.228.37

User Name: admin

Password:

Protocol Type: File Transfer Protocol (FTP)

Specify Firmware Path: Firmware/Switches/6.x/6.4.x/6.4.1/v6.4.1

Download Close Refresh

[Firmware download failed]: Tue Nov 09 2010 00:36:56 GMT+00:00
[Firmware download started]: Tue Nov 09 2010 00:35:56 GMT+00:00
Initiating firmware download ...
From Host: 10.18.228.37
File Path: /Firmware/Switches/6.x/6.4.x/6.4.1/
[Error]: The server is inaccessible or firmware path is invalid. Please make sure the server name or IP address, the user/password and the firmware path are valid.
[Firmware download failed]: Tue Nov 09 2010 00:35:56 GMT+00:00

Start transport Mode: Basic [Free Professional Management Tool](#) 10.18.228.18 AD0 User: admin Role: admin ✓

Figure 8-74 Switch Administration Firmware Download tab

- Complete the fields as appropriate. Enter the IP address of FTP server, the user name and password, and the directory that contains the firmware files (in our case, the directory is Firmware/Switches/6.x/6.4.x/6.4.1/v6.4.1/).

4. Then, click **Download** to initiate the firmware upgrade process.

A confirmation window displays, see Figure 8-75. The information displayed in the window reminds you that the switch will reboot after the download is done. Note that on the Converged Switch (Elara) this will cause a cold/disruptive boot. This means that the management network connection to the switch will be lost and will need to be re-established.

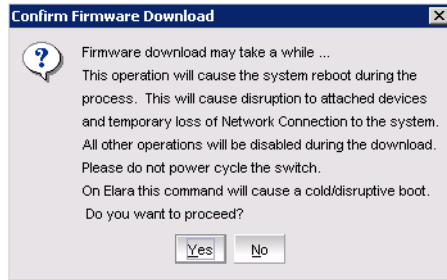


Figure 8-75 Firmware download confirmation

5. Continue by clicking **Yes**.
6. The firmware download and upgrade process takes a while (up to 30 minutes). When it completes, you have to close and reopen Web Tools and reconnect to the switch. As shown in Figure 8-76, the switch now runs Fabric OS v6.4.1 (Figure 8-76).




Figure 8-76 IBM Converged SAN B32 running Fabric OS v6.4.1

8.7.5 SNMP tab

We use the SNMP tab to administer the SNMP subsystem. From the SNMP tab, you specify the switch community string, location, trap level, and trap recipients. SNMPv3 is available from Fabric OS v4.4 onwards, as well as SNMPv1. As shown in Figure 8-77, you can set traps using either SNMPv1 or SNMPv3.

Apart from using Web Tools, SNMP parameters can also be set with Telnet commands or Data Center Fabric Manager (Figure 8-77).

IBM_SAN80B_217 - Switch Administration

SwitchName: IBM_SAN80B_217 DomainID: 2(0x2) VWN: 10:00:00:05:1e:09:97:01 Mon Jul 20 2009 18:52:25 GMT+00:00

SNMP Configure Routing Extended Fabric AAA Service Trace FICON CUP Security Policies

Switch Network Firmware Download License User Trunking

SNMP Information

Contact Name: Field Support.

Description: Fibre Channel Switch.

Location: End User Premise.

Enable/Disable Authentication Trap

☐ Enable Authentication Trap

SNMPv3 Inform / Trap Recipient

User Name	Recipient IP	Port Number	Trap Level
admin - RO	10.64.210.104	162	4 - Informational
snmpadmin2 - RW	0.0.0.0	162	0 - None
snmpadmin3 - RW	0.0.0.0	162	0 - None

SNMPv1 Community/Trap Recipient

Community String	Recipient	Port Number	Access Control	Trap Level
Secret C0de	10.64.210.74	162	Read Write	4 - Informational
OrigEquipMtr	dell-storage-x-074	162	Read Write	4 - Informational
private	127.0.0.1	162	Read Write	0 - None
public	10.64.210.72	162	Read Only	4 - Informational

Access Control List

Access Host	Access Control List
0.0.0.0	Read Write
0.0.0.0	Read Write
0.0.0.0	Read Write
0.0.0.0	Read Write

Apply Close Refresh

Configure SNMP parameters Mode: Advanced Free Professional Management Tool 10.64.210.217 AD0 User: admin Role: admin

Figure 8-77 SNMP tab

Creating a new SNMPv1 trap

To create a new trap, proceed as follows:

1. Double-click a community string in the SNMPv1 section, and enter a new community string.
2. Double-click a recipient IP address in the SNMPv1 section, and enter a new IP address.
3. Click **Apply**.

Creating a new SNMPv3 trap

To create a new trap, proceed as follows:

1. Select a user name from the User Name drop-down list in the SNMPv3 section.
2. Double-click a recipient IP address in the SNMP v3 section, and enter a new IP address.
3. Select a trap level from the Trap Level drop-down list.
4. Click **Apply**.

Table 8-8 describes the fields on the SNMP tab.

Table 8-8 SNMP tab

SNMP Basic Information	
Contact Name	Displays or sets contact information for switch. Default is Field Support.
Description	Displays or sets system description. Default is Fibre Channel Switch.
Location	Displays or sets the location of switch. Default is End User Premise.
Enable Authentication Trap	Check to enable authentication traps; uncheck to disable (preferable).
SNMPv1 Community/Trap Recipient	
Community String	Displays the community strings that are available to use. A community refers to a relationship between a group of SNMP managers and an SNMP agent, in which authentication, access control, and proxy characteristics are defined. A maximum of six community strings can be saved to the switch.
Recipient	Displays the IP address of the Trap Recipient. A trap recipient receives the message sent by an SNMP agent to inform the SNMP management station of a critical error.
Access Control	Displays the Read/Write access of a particular community string. Read only access means that a member of a community string has the right to view, but cannot be changed. Read/Write access means that a member of a community string can be both viewed and changed.
Trap Level	Sets severity level of switch events that prompt SNMP traps. Default is 0.

SNMPv3 Trap Recipient	
User Name	Displays user names that are available to use. The user names are predefined with different Read/Write or Read Only access. The predefined user names are snmpadmin1, snmpadmin2, snmpadmin3 with Read/Write access and snmpuser1, snmpuser2, snmpuser3 with Read Only access.
Recipient IP	Displays the IP address of the Trap Recipient. A trap recipient receives the message sent by an SNMP agent to inform the SNMP management station of a critical error.
Trap Level	Sets severity level of switch events that prompt SNMP traps. Default is 0
Access Control List Configuration	
Access Host	Displays the IP address of the host of the access list.
Access Control List	Displays the Read/Write access of a particular access list. Read only access means that a member of an access list has the right to view, but cannot make changes. Read/Write access means that a member of an access list can both view and make changes.
Apply	Click to save the changes made to this tab. You can make additional changes and click Apply when making changes incrementally.
Close	Click to exit the Admin window. If you have made changes but did not commit them clicking Apply , a dialog box displays.
Refresh	Click to retrieve current values from the switch.

You can also set SNMP parameters with Telnet using the **snmpConfig** command. In older Fabric OS releases, the following set of commands was used to set and view SNMP settings:

- **agtcfgSet**
- **agtcfgShow**
- **agtcfgDefault**

The functionality of these commands is now available through the **snmpConfig** command. You also use the **snmpConfig** command to view or set the MIB capability, instead of older commands **snmpMibCapSet** and **snmpMibCapShow**. Example 8-4 shows current MIB capability settings for the SAN80B switch.

Example 8-4 Using the snmpConfig command to verify MIB capability settings

```
IBM_SAN80B_217:admin> snmpConfig --show mibCapability
```

```
FE-MIB: YES
SW-MIB: YES
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
FCIP-MIB: NO
ISCSI-MIB: YES
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
    swFabricWatchTrap: YES
    swTrackChangesTrap: YES
FA-TRAP: YES
    connUnitStatusChange: YES
    connUnitEventTrap: YES
    connUnitSensorStatusChange: YES
    connUnitPortStatusChange: YES
SW-EXTTRAP: NO
FICON-TRAP: YES
    linkRNIDDeviceRegistration: YES
    linkRNIDDeviceDeRegistration: YES
    linkLIRListenerAdded: YES
    linkLIRListenerRemoved: YES
    linkRLIRFailureIncident: YES
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: YES
```

8.7.6 License tab

We use the License tab to install the license keys that you have purchased. You use license keys to enable additional features on a switch. You can also use the table within the License tab to remove a listed license from the switch.

Figure 8-78 displays the License tab for a SAN80B switch.

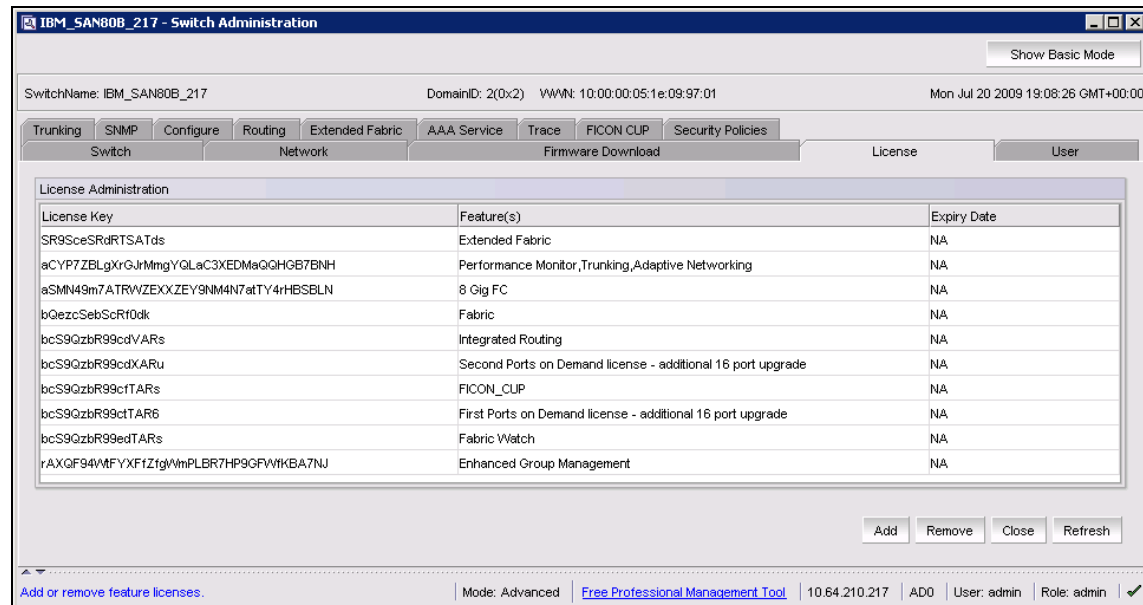


Figure 8-78 Switch Administration License tab

The following licenses are currently enabled on this switch:

- ▶ Extended Fabric
- ▶ Performance Monitor, Trunking, Adaptive Networking
- ▶ 8 Gig FC
- ▶ Fabric
- ▶ Integrated Routing
- ▶ Second Ports on Demand licence - additional 16 port upgrade
- ▶ FICON_CUP
- ▶ First Ports on Demand licence - additional 16 port upgrade
- ▶ Fabric Watch
- ▶ Enhanced Group Management

Certain licenses might be available only for a limited time as indicated, in the Expiry Date column. In our example, all currently enabled licenses never expire.

To enable additional licenses, you need the following items:

- ▶ A license transaction key, which is supplied in the documentation when purchasing a license.
- ▶ A license ID of the switch. You can obtain this ID in two ways:
 - In the Switch Information panel on main Web Tools window as shown in Figure 8-79.

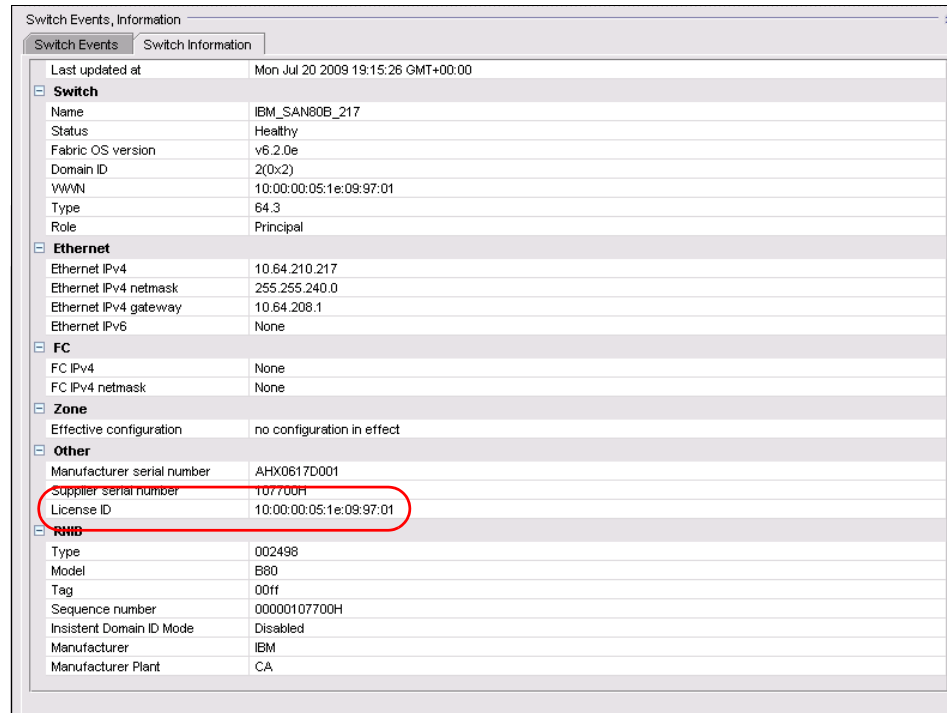


Figure 8-79 License ID field

- Using the **licenseIdShow** command, as shown in Example 8-5.

Example 8-5 The licenseIdShow command

```
IBM_SAN80B_217:admin> licenseIdShow
10:00:00:05:1e:09:97:01
```

Our SAN80B switch license ID is 10:00:00:05:1e:09:97:01, which is actually the same license ID as switch WWN.

You need the license ID and the transaction key from the documentation to obtain the *license activation key* on the Feature activation keys Web page, as discussed in the next section.

Obtaining the license activation keys

The Feature activation keys Web page is available at the following URL:

<http://www.ibm.com/storage/key>



Figure 8-80 displays the Web page. To obtain the license activation key:

1. Read the instructions carefully, then click **Generate one or more activation keys**.

Welcome to our enhanced web site for feature activation keys.

To make our licensing procedures easier for you to use, the site no longer requires the switch ID number. Older paper packs may request that you gather the switch ID number in step 1. You may ignore that instruction. You must still enter the switch World Wide Name to activate your paper pack.

If you have gone to any trouble to gather the switch ID, we apologize for the inconvenience.

If you need technical assistance related to obtaining your feature activation key(s), please call IBM support in your country ( 800-IBM-SERV  in the U.S.) and open a hardware support incident for your switch.

Retrieve license key(s) for an existing SAN Switch

License key values may be found using one of the following:

- World-Wide Name (*)
- Transaction Key

Proceed to retrieve license key(s)

Activate a SAN Switch feature

Step 1 of 2: Get a feature activation key

Please have the following information ready to obtain your key.

This information is confidential; the email address will only be used to email the license key to you as a backup.

- E-mail Address
- World-Wide Name (*)
- Transaction Key

Generate one or more activation keys

Figure 8-80 Feature activation keys Web page

2. Enter your email address, switch WWN/License ID, and transaction keys as shown in Figure 8-81. Complete the fields as appropriate, then click **Continue**.

Email Notification

* Email address

itso@us.ibm.com

* Verify email address

itso@us.ibm.com

Installation Site Information

* World Wide Names/License ID's and Transaction Keys

World Wide Name/License ID	Transaction Key
10:00:00:05:1e:09:73:fd	6848c56b06f9f7fb22a247
10:00:00:	
10:00:00:	
10:00:00:	
10:00:00:	

Add more rows

This data may be used by IBM or selected organizations, such as Lenovo, to provide you with information about other offerings. To receive this via email, check the first box below. Alternatively, if you would prefer not to receive such information by any means, check the second box.

☐

Please use e-mail to send me information about other offerings.

☐

Please do not use this data to send me information about other offerings.

By clicking **Continue** you agree that IBM may process your data in the manner indicated above and as described in [Privacy](#).

Continue

Back

Figure 8-81 Feature activation keys: Enter email address, WWN/License ID, and transaction keys

3. Next, verify the entries (see Figure 8-82).

IBM Systems > Systems Support > Storage support > Storage area
network (SAN) >

SAN Switch feature activation

Activate a SAN Switch feature

Please verify the following information, then click Submit to create feature activation key(s).

Email Notification

Email address **itso@us.ibm.com**

Installation Site Information

World Wide Name/License ID	Transaction Key
10:00:00:05:1e:09:73:fd	6848c56b06f9f7fb22a247




 **Submit**  Back  Cancel

Figure 8-82 Feature activation keys: Verify entered information

4. Click **Submit** to create license activation keys for the selected features.
5. Finally, the license activation keys are generated and presented to you, as shown in Figure 8-83. The license keys are also sent to your email address.

Activate a SAN Switch feature

Activation keys:

World Wide Name	Transaction Key	License Key	Feature Name
10:00:00:05:1E:09:73:FD	6848C56B06F9F7FB22A247	TFfSD7HPXgNCmCtZ4W4tZGrXXgmGK7ABBJRFL	Enterprise Bundle: ISL Trunking, Advanced Performance Monitoring & Adaptive Networking

Installation Guide

Note: The following instructions assume the switch has been attached to a network and is accessible from your PC or workstation.

Via telnet

1. Connect to the switch via telnet session from a terminal emulation program initiated on a PC or workstation on your network.
2. Type the command: telnet [switch IP address]. When prompted, enter Administrator User name and password.
3. After login, enter the command: licenseAdd "license-key" (quotes are needed for this command). Be sure to enter the license key value exactly as received.
4. Check the status of installed license features by entering the command: licenseShow. This command displays a list of all licensed features enabled on the switch.

From a web browser (via GUI)

1. Connect to the switch with your browser. From the Switch view, select the Admin button.
2. When prompted, enter administrator user name and password.
3. On successful login, you will be presented the Admin screen, select the License Admin tab, Enter the new License Key. Be sure to enter the license key value exactly as received.
4. Select the Add License button.

To validate that this LicenseKey has been installed

From front panel display - If your switch is a model with a front panel display, the WWN can be shown on the display. Select the following command sequence: Status Menu > WorldWide Menu.

Figure 8-83 Feature activation keys: Generated license keys

In addition to the license keys, this Web page also displays installation steps that you need to take to enable the licenses on the switch.

Removing a license key

To remove a license key, follow these steps:

1. Highlight the license key to remove and click **Remove**, as shown in Figure 8-84.

2. Click **Yes** to confirm that you are removing the license.
3. Click **Refresh** to show that the license was removed.

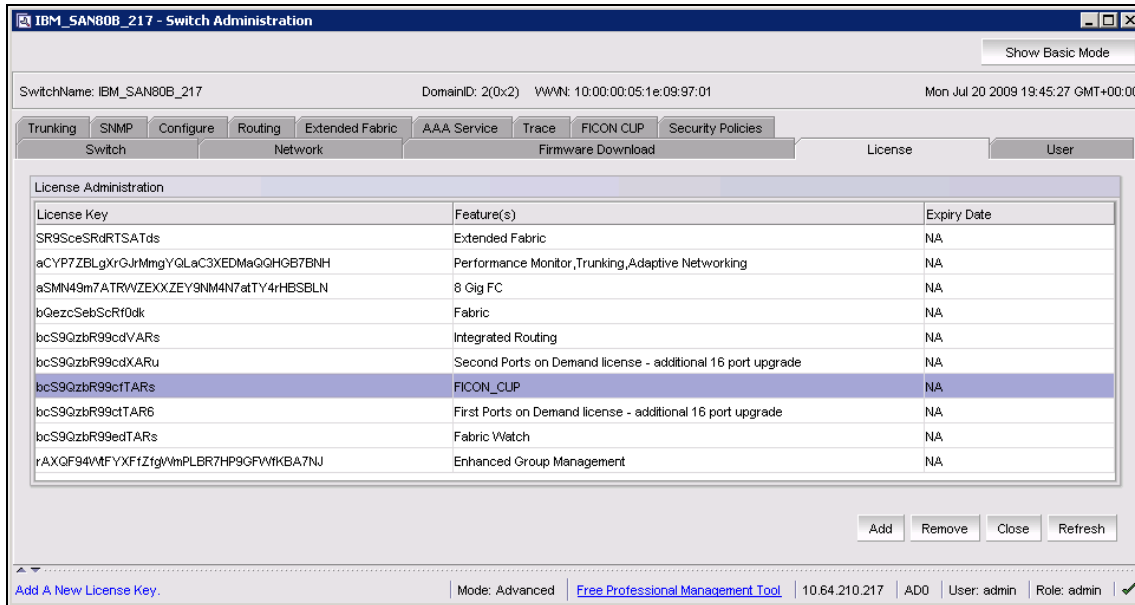


Figure 8-84 Removing a license

Installing a license key through the CLI

To install a license key feature using the CLI, perform the following steps:

1. From a command prompt, use the Telnet command to log in to the switch using an account that has administrative privileges. In the following example, we replace *address* with the switch IP address:

```
C:\telnet address
```

2. To determine which licenses are already installed on the switch, run the **licenseShow** command. A list displays of all the licenses currently installed on the switch, as shown in Example 8-6.

Example 8-6 The licenseshow CLI output from SAN40B-4

```
IBM_SAN80B_217:admin> licenseShow
bQezcSebScRf0dk:
    Fabric license
bcS9QzbR99cdTARs:
    Fabric Watch license
SR9SceSRdRTSATds:
    Extended Fabric license
```

```
bcS9QzbR99cdVARs:
    Integrated Routing license
bcS9QzbR99ctTAR6:
    First Ports on Demand license - additional 16 port upgrade license
bcS9QzbR99cdXARu:
    Second Ports on Demand license - additional 16 port upgrade license
rAXQF94WtFYXFfZfgWmPLBR7HP9GFWfKBA7NJ:
    Enhanced Group Management license
aSMN49m7ATRWZEXXZEY9NM4N7atTY4rHBSBLN:
    8 Gig FC license
aCYP7ZBLgXrGJrMmgYQLaC3XEDMaQQHGB7BNH:
    Performance Monitor license
    Trunking license
    Adaptive Networking license
```

3. To install a license key, enter the following command on the command line:

licenseAdd “key”

In this command, “key” is the license key that is provided to you, enclosed in double quotation marks. The license key is case sensitive, so you must enter it exactly as given as shown in Example 8-7.

Example 8-7 Adding a License Key

```
IBM_SAN80B_217:admin> licenseAdd "bcS9QzbR99cfTARs"
adding license-key [bcS9QzbR99cfTARs]
```

4. Use **licenseShow** again to verify that the license was added.

If the license is listed, the feature is installed and available. Otherwise, repeat step 3.

Example 8-8 adds the FICON_CUP license to the SAN80B switch.

Example 8-8 Adding licenses

```
IBM_SAN80B_217:admin> licenseShow
bQezcSebScRf0dk:
    Fabric license
bcS9QzbR99edTARs:
    Fabric Watch license
SR9SceSRdRTSATds:
    Extended Fabric license
bcS9QzbR99cdVARs:
    Integrated Routing license
bcS9QzbR99ctTAR6:
    First Ports on Demand license - additional 16 port upgrade license
```

bcS9QzbR99cdXARu:
Second Ports on Demand license - additional 16 port upgrade license
bcS9QzbR99cfTARs:
FICON_CUP license <-----New License
rAXQF94WtFYXFfZfgWmPLBR7HP9GFWfKBA7NJ:
Enhanced Group Management license
aSMN49m7ATRWZEXXZEY9NM4N7atTY4rHBSBLN:
8 Gig FC license
aCYP7ZBLgXrGJrMmgYQLaC3XEDMaQQHGB7BNH:
Performance Monitor license
Trunking license
Adaptive Networking license

8.7.7 User tab

The User tab allows you to perform user administration tasks, such as adding new users, changing properties of existing users, and deleting user accounts that are no longer needed.

Figure 8-85 displays the layout of the User tab.

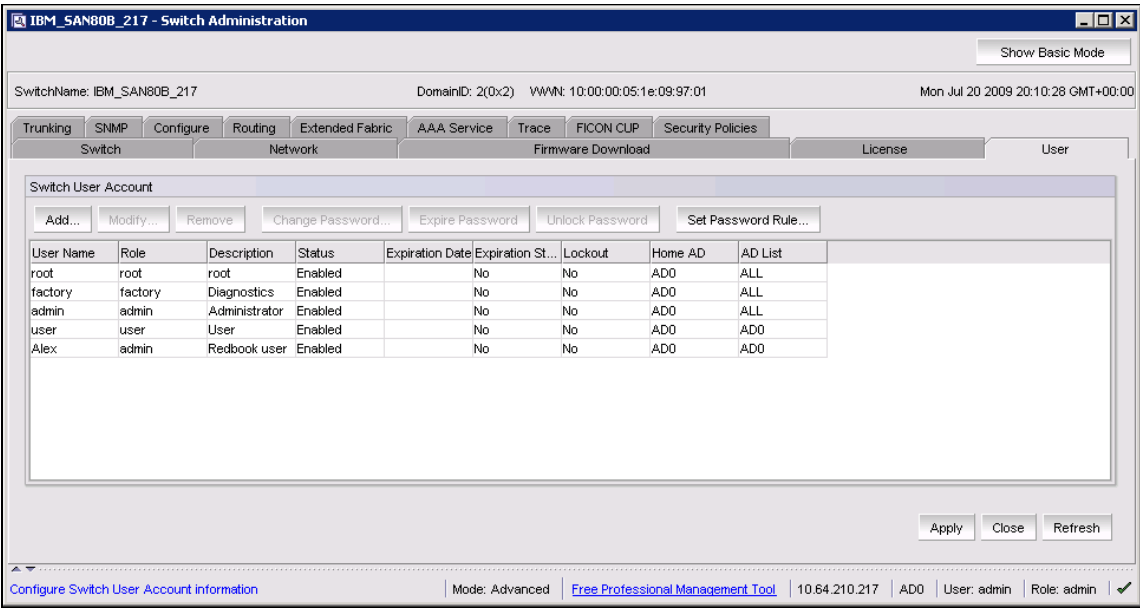
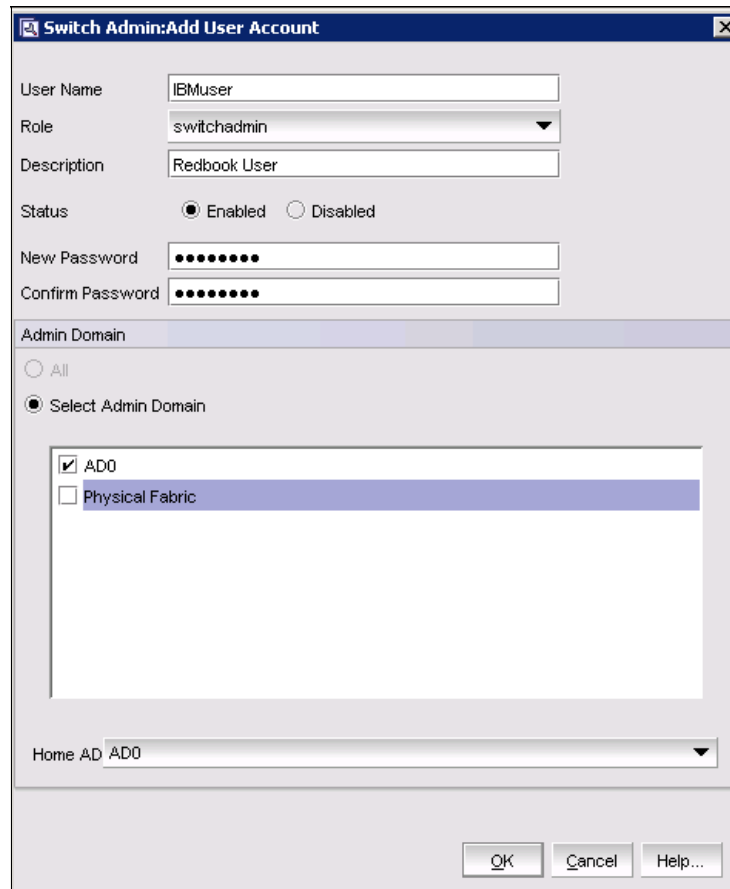


Figure 8-85 User tab

To add a new user:

1. Click **Add** to open the Switch Admin: Add User Account window, as shown in Figure 8-86.



The image shows a web-based window titled "Switch Admin: Add User Account". It contains several input fields and a list of domains. The "User Name" field is filled with "IBMuser". The "Role" dropdown menu is set to "switchadmin". The "Description" field is filled with "Redbook User". The "Status" section has two radio buttons: "Enabled" (which is selected) and "Disabled". Below this are two password fields: "New Password" and "Confirm Password", both filled with eight dots. A section titled "Admin Domain" has two radio buttons: "All" and "Select Admin Domain" (which is selected). Below the "Select Admin Domain" radio button is a list box containing two items: "AD0" (which is checked with a small square icon) and "Physical Fabric" (which is unchecked). At the bottom of the "Admin Domain" section is a dropdown menu labeled "Home AD" with "AD0" selected. At the very bottom of the window are three buttons: "OK", "Cancel", and "Help...".

Figure 8-86 Add User Account

- In this example, we add a user named *IBMuser* with the *switchadmin* role. The user is associated with *AD0*. After entering all the necessary information, click **OK**, and the new user is created (see Figure 8-87).

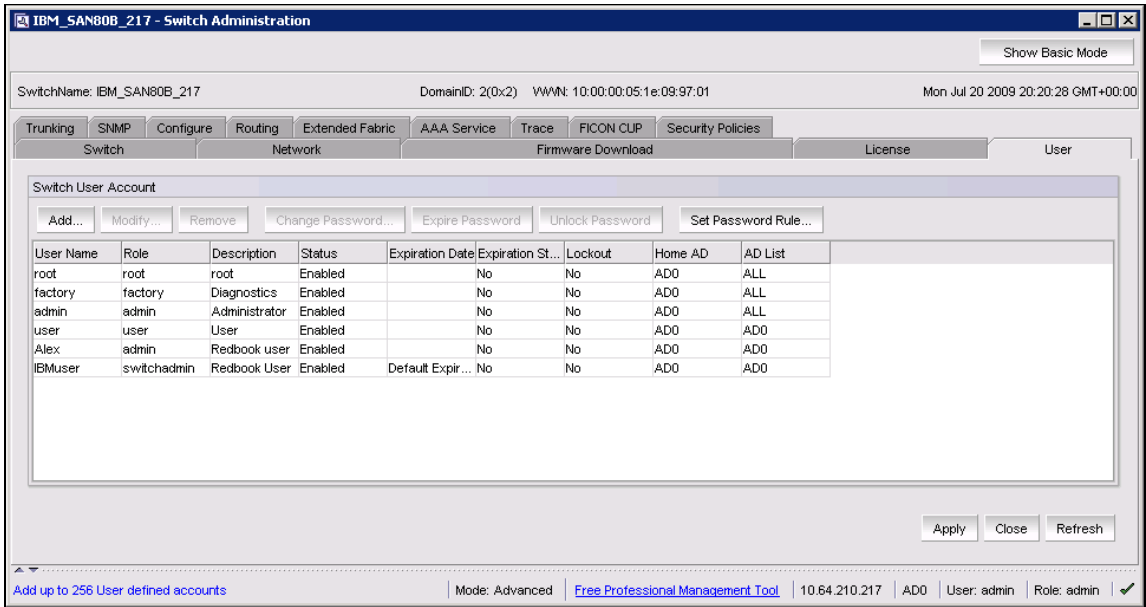


Figure 8-87 New user is created

- A set of buttons above the list of users allows you to perform the user administration tasks. We have seen the use of the **Add button**, Table 8-9 explains the remaining buttons in the User tab and their actions.

Table 8-9 User administration button actions

Button	Action
Modify	Use this button to change the user account properties.
Remove	Deletes the currently selected user.
Change Password	Set the user's password to a new value.
Expire Password	Set the user's password to expired state.
Unlock Password	Reset locked-out users.
Set Password Rule	Define the rules for user's password.

4. For the changes to be committed successfully to the switch, you must click **Apply** to open a window and confirm your actions, as shown in Figure 8-88.

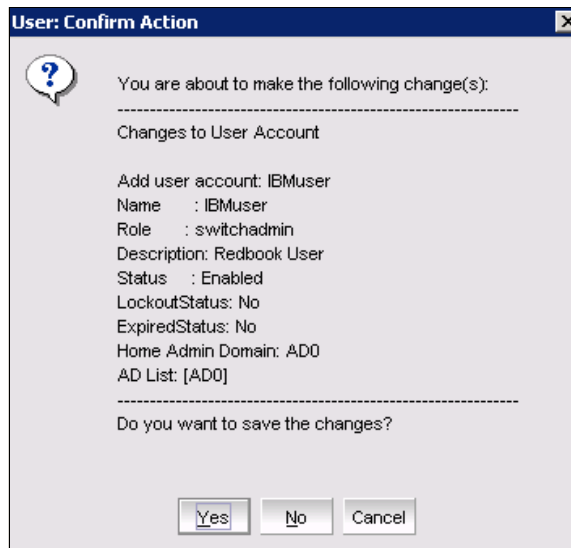


Figure 8-88 Confirm changes to user account

5. Click **Yes** to confirm and commit the changes and to complete the process of adding a new user.

In this example, we set the new user to be the *switchadmin*. Other available access levels include:

- ▶ **Admin:**
This access level allows change and view access to all functions. From Telnet access, the Admin level allows use of all available commands. Typically, most switch administration is performed at this level.
- ▶ **User:**
This access level provides view access only. Users cannot make zoning changes or any switch configuration changes. This level is best for monitoring switch activity.
- ▶ **SwitchAdmin:**
This new role has most of the existing permissions of the traditional Admin role but cannot create or change fabric security policies, cannot create or change fabric zoning policies, and cannot create or manage users.
- ▶ **FabricAdmin:**
The FabricAdmin role can perform administrative tasks, but cannot manage users. It also cannot perform AD management.
- ▶ **ZoneAdmin:**
This role allows zone management only.
- ▶ **BasicSwitchAdmin:**
BasicSwitchAdmin access level is a subset of Admin role. Most available tasks are for monitoring purpose and can perform limited local switch management.
- ▶ **Operator:**
This role can perform a set of tasks required to do routine maintenance.
- ▶ **SecurityAdmin:**
SecurityAdmin can exercise security related tasks. This includes user management tasks.

Attention: The User tab does not display or modify the RADIUS host server database.

8.7.8 Configure tab

Figure 8-89 shows the Configure tab. You cannot make changes to the settings on this tab if the switch is enabled; however, the configuration upload/download facility is available regardless of the switch status. In our example, we disable the switch so that we can make configuration changes.

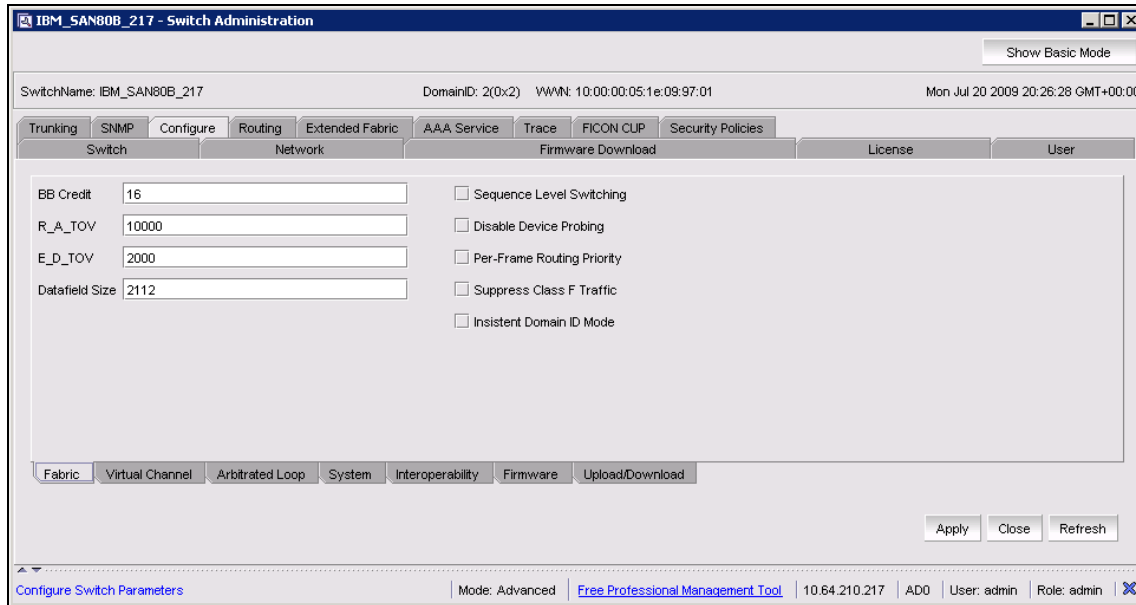


Figure 8-89 Switch Administration: Configure tab

The Configure tab includes the following tabs:

- ▶ Fabric
- ▶ Virtual Channel
- ▶ Arbitrated Loop
- ▶ System
- ▶ Interoperability
- ▶ Firmware
- ▶ Upload/Download

In the sections that follow, we describe the parameters that are configurable through each of these tabs.

Fabric parameters

The following Fabric parameters are available:

- ▶ **BB Credit:** The buffer-to-buffer (BB) credit represents the number of buffers available to attached devices for frame receipt. This value ranges from 1 to 27. Default value is 16.
- ▶ **R_A_TOV:** The Resource Allocation Time Out Value (R_A_TOV) is displayed in milliseconds. Allocated circuit resources with detected errors are not released until this timeout value has expired. If the condition is resolved prior to the timeout, the internal time out clock resets and waits for the next error condition.
- ▶ **E_D_TOV:** Error Detect Time Out Value (E_D_TOV) is displayed in milliseconds. This timer is used to flag a potential error condition when an expected response is not received (an acknowledgment or reply in response to packet receipt, for example) within the set time limit. If the time for an expected response exceeds the set value, then an error condition occurs.
- ▶ **Datafield Size:** The largest datafield size in bytes.
- ▶ **Sequence Level Switching:** When Sequence Level Switching is enabled, frames of the same sequence from a particular source are transmitted together as a group. When this is feature disabled, frames are transmitted interleaved among multiple sequences. Under normal conditions, Sequence Level Switching needs to be disabled for better performance.
- ▶ **Disable Device Probing:** When Disable Device Probing is enabled, devices that do not register with the Name Server are not present in the Name Server data base. Set this mode only if the switch N_Port discovery process (PLOGI, PRLI, INQUIRY) causes an attached device to fail.
- ▶ **Per-Frame Routing Priority:** In addition to the eight virtual channels used in frame routing priority, support is also available for per-frame-based prioritization when this value is set. When Per-Frame Routing Priority is enabled, the virtual channel ID is used in conjunction with a frame header to form the final virtual channel ID.
- ▶ **Suppress Class F Traffic:** When enabled, all class F interswitch frames are transmitted as class 2 frames to support remote fabrics which involve ATM gateways, which do not support class F traffic.
- ▶ **Insistent Domain ID Mode:** Setting this mode makes the current domain ID insistent across reboots, power cycles and failover, which required fabric wide to transmit FICON data.

Virtual Channel parameters

This feature enables fine tuning of ISLs by configuring parameters for the eight virtual channels. These parameters are used for congestion control. Use the default values for these parameters unless expert advice is available. Figure 8-90 displays the Virtual Channel tab.

The screenshot shows the 'IBM_SAN80B_217 - Switch Administration' web interface. At the top, there's a header with 'SwitchName: IBM_SAN80B_217', 'DomainID: 2(0x2)', 'WWN: 10:00:00:05:1e:09:97:01', and 'Mon Jul 20 2009 20:38:28 GMT+00:00'. Below the header is a navigation bar with tabs: Trunking, SNMP, Configure, Routing, Extended Fabric, AAA Service, Trace, FICON CUP, Security Policies, Switch, Network, Firmware Download, License, and User. The 'Virtual Channel' tab is selected. The main content area contains eight input fields for VC Priority 2 through VC Priority 7, with values 2, 2, 3, 2, 2, 3 respectively. At the bottom, there's a status bar with 'Configure Switch Parameters', 'Mode: Advanced', 'Free Professional Management Tool', '10.64.210.217', 'AD0', 'User: admin', 'Role: admin', and a refresh icon. Buttons for 'Apply', 'Close', and 'Refresh' are also present.

Figure 8-90 Virtual Channel tab

Arbitrated Loop parameters

The Arbitrated Loop parameters include these:

- ▶ Send Fan Frames: Specifies that fabric address notification (FAN) frames be sent to public loop devices with notification of their node ID and address. When enabled, frames are sent, and when disabled, frames are not sent.
- ▶ Always send RSCN: Following the completion of loop initialization, a Registered State Change Notification (RSCN) is issued when FL_Ports detect the presence of new devices or the absence of pre-existing devices. When this mode is enabled, a RSCN is issued upon completion of loop initialization, regardless of the presence or absence of new or preexisting devices.
- ▶ Do Not Allow AL_PA 0x00: This option disallows AL_PA values from being 0.

System parameters

The System tab lets you change the *Disable RLS probing* parameter. Use this setting to disable Read Link Error Status of the AL_PAs.

Interoperability parameters

The Interoperability tab allows you to set the switch to operate in one of the following modes:

- ▶ Brocade Native Fabric Mode
- ▶ McDATA Fabric Mode
- ▶ McDATA Open Fabric Mode

If you need to set either of the two McDATA modes, be aware of the following considerations:

- ▶ The McDATA Fabric Mode requires that the domain ID is in the range 1 through 31.
- ▶ The McDATA Open Fabric Mode domain ID range is 97 through 131.

If the domain ID is outside these values when you try to enable McDATA interoperability mode, Web Tools prompts you to change the domain ID first.

When enabling any McDATA interoperability mode, the zoning database is reset.

Firmware parameters

The Firmware tab contains only one parameter, which is *Enable Signed Firmware Download*. When this option is enabled, the system validates the firmware that is downloaded to the switch. Firmware validation cannot be done during the very first download; however, after the first firmware download is complete, the public key is downloaded, so that the validation works on subsequent firmware downloads.

Upload/download parameters

The Upload/Download tab enables you to manipulate the switch configuration. You can store (upload) the configuration to an FTP server or to a Brocade-branded USB memory key and download a previously stored configuration from the FTP server or USB key to the switch.

Figure 8-91 shows an example of the Upload/Download tab.

The screenshot displays the 'IBM_SAN80B_217 - Switch Administration' web interface. At the top, a status bar shows 'SwitchName: IBM_SAN80B_217', 'DomainID: 2(0x2)', 'VWWN: 10:00:00:05:1e:09:97:01', and the date 'Mon Jul 20 2009 20:47:29 GMT+00:00'. Below this is a navigation menu with tabs for 'SNMP', 'Configure', 'Routing', 'Extended Fabric', 'AAA Service', 'Trace', 'FICON CUP', and 'Security Policies'. The 'Configure' tab is active, and within it, the 'Upload/Download' sub-tab is selected. The main content area is titled 'Function' and contains two radio buttons: 'Config Upload' (selected) and 'Config Download to Switch'. Below these is a section for 'Select source of configuration file' with 'Network' (selected) and 'USB' options. A red instruction reads: 'Provide Host details, Transfer Protocol and Path for Configuration file' followed by '*Password is optional, if user name is "anonymous"'. The form fields include: 'Host Name or IP' (10.64.210.103), 'User Name' (ibm), 'Password' (masked with dots), 'Protocol Type' (File Transfer Protocol (FTP)), and 'Configuration File Name' (SAN80B_config_200709). An 'Upload/Download Progress' bar is shown below the form. At the bottom of the main area are tabs for 'Fabric', 'Virtual Channel', 'Arbitrated Loop', 'System', 'Interoperability', 'Firmware', and 'Upload/Download' (selected). The bottom of the interface features 'Apply', 'Close', and 'Refresh' buttons, and a footer bar with links like 'Enter the Password', 'Mode: Advanced', 'Free Professional Management Tool', and system information including '10.64.210.217', 'AD0', 'User: admin', and 'Role: admin'.

Figure 8-91 Upload/Download tab

To save the configuration file to an FTP server, proceed as follows:

1. Click **Config Upload**.
2. Provide the FTP server IP address, user name, password, and file name of the configuration file.

Naming: Remember to use a sensible naming convention for your configuration files to ensure that you are able to recover to the appropriate point as required.

3. Then, click **Apply**.
4. When prompted to verify that you want to perform this function (as shown in Figure 8-92), click **Yes** to continue.




Figure 8-92 Confirm configuration upload

When completed, the confirmation message for the upload displays on the report window.

8.7.9 Routing tab

We discuss the Routing tab in 13.1.5, “Routing policies” on page 601.

8.7.10 Extended Fabric tab

We discuss the Extended Fabric tab in “Extended Fabrics” on page 595.

8.7.11 AAA Service tab

Fabric OS v6.1.0 supports RADIUS and Active Directory server authentication. You can use the Switch Administration AAA Service tab to configure the RADIUS or Active Directory servers. Figure 8-93 shows the AAA tab.

IBM_SAN80B_217 - Switch Administration

Show Basic Mode

SwitchName: IBM_SAN80B_217 DomainID: 2(0x2) VVWN: 10:00:00:05:1e:09:97:01 Mon Jul 20 2009 20:52:29 GMT+00:00

SNMP Configure Routing Extended Fabric AAA Service Trace FICON CUP Security Policies

Switch Network Firmware Download License User Trunking

AAA Service

Primary AAA Service: Switch Database Secondary AAA Service: None

RADIUS Configuration

Server	Port	Timeout(s)	Authentication
--------	------	------------	----------------

ADLDAP Configuration

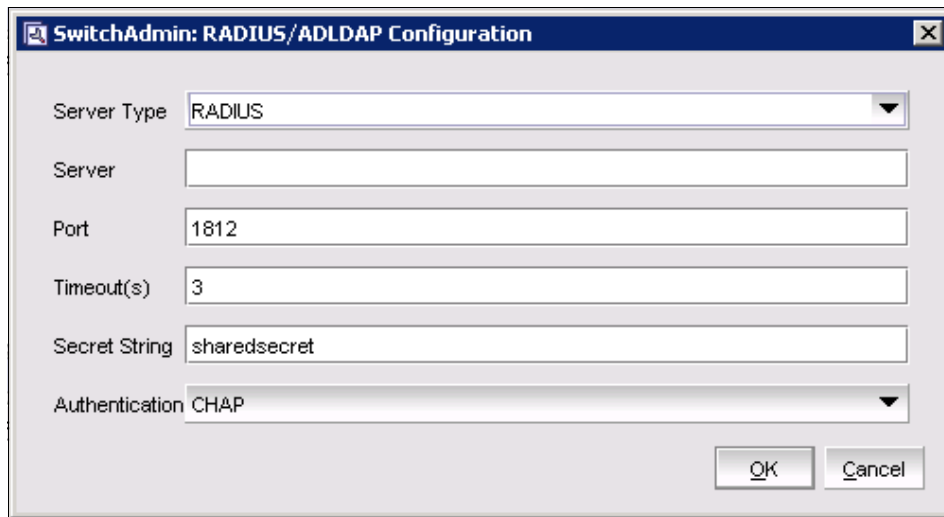
Server	Port	Timeout(s)	Domain
--------	------	------------	--------

Apply Close Refresh

Configure AAA Service and RADIUS servers Mode: Advanced [Free Professional Management Tool](#) 10.64.210.217 AD0 User: admin Role: admin

Figure 8-93 AAA Service tab

Click **Add** to configure RADIUS or Active Directory server. A dialog box displays, as shown in Figure 8-94. Enter the appropriate values and click **OK**.



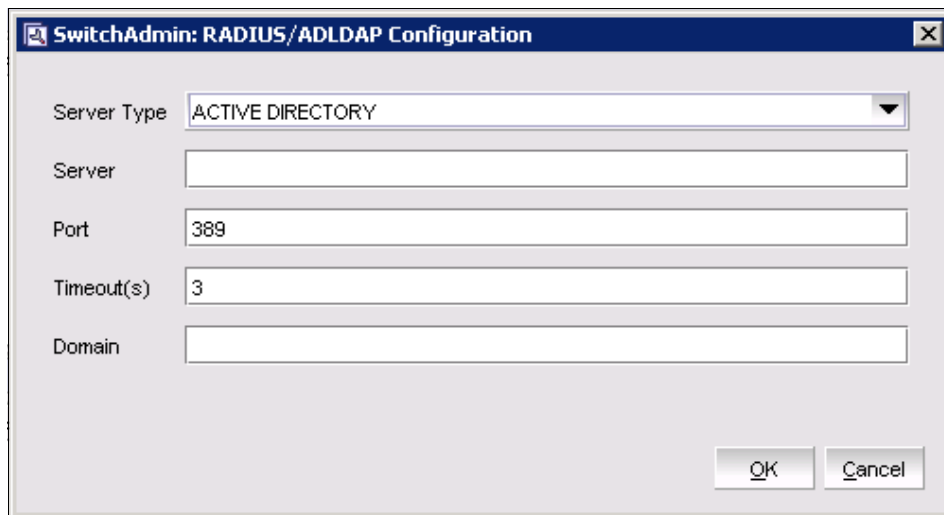
The dialog box titled "SwitchAdmin: RADIUS/ADLDAP Configuration" contains the following fields and controls:

- Server Type:** A dropdown menu with "RADIUS" selected.
- Server:** An empty text input field.
- Port:** A text input field containing "1812".
- Timeout(s):** A text input field containing "3".
- Secret String:** A text input field containing "sharedsecret".
- Authentication:** A dropdown menu with "CHAP" selected.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Figure 8-94 Adding RADIUS server

You can configure up to five RADIUS servers and reorder them at a later time on the AAA Service tab panel. RADIUS servers are contacted in the order that they display in the RADIUS Configuration list.

Figure 8-95 shows an example of configuring an Active Directory server.



The dialog box titled "SwitchAdmin: RADIUS/ADLDAP Configuration" contains the following fields and controls:

- Server Type:** A dropdown menu with "ACTIVE DIRECTORY" selected.
- Server:** An empty text input field.
- Port:** A text input field containing "389".
- Timeout(s):** A text input field containing "3".
- Domain:** An empty text input field.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Figure 8-95 Active Directory server configuration

Now that the servers are defined, you can modify or remove them by highlighting them and clicking either **Modify** or **Remove**. When you have finished listing all the servers in the configuration, you can change the order in which they are contacted for authentication by using the up and down arrow on the right of the window displaying the list of servers. Table 8-10 describes the details of the AAA tab functions.

Table 8-10 AAA tab functions

Function	Description
Primary AAA Service	Primary Service Engine
Secondary AAA Service	A Backup Service Engine
RADIUS configuration	A window displaying RADIUS servers in the configuration
ADLDAP configuration	Active Directory LDAP server
Port	Port for which RADIUS or ADLDAP server is defined
Timeout or timeouts	Timeout value in minutes
Authentication	Authentication protocol used
Up/Down Arrows	Navigate order for which servers are contacted
Add	Add a new RADIUS or ADLDAP server
Modify	Modify an existing RADIUS or ADLDAP server
Remove	Remove an existing RADIUS or ADLDAP server
Apply	Apply and commit changes to the switch
Close	Close the Administration window
Refresh	Refresh the view from the current switch data

8.7.12 Trace tab

If a switch experiences a serious problem, a trace dump can be generated to capture valuable information about the running state. IBM and Brocade technical support teams can use the trace dump to help understand and solve the problem. Tracing is enabled at all times. The tracing information is stored continuously into a circular buffer in switch system memory. Thus, old trace information eventually is overwritten with new information. Capturing a trace dump means storing the current content of the trace buffer. This way, you can preserve important troubleshooting information before it is overwritten.

Trace dump generation takes place in the following cases:

- ▶ You can trigger the trace dump manually using the CLI **traceDump** command.
- ▶ A trace dump can occur when a critical-level LOG message occurs.
- ▶ You can use the CLI **traceTrig** command to set up trace dump generation in case of another particular LOG message.
- ▶ A trace dump can occur if a kernel panic occurs.
- ▶ A trace dump can occur when the hardware watchdog timer expires.

After the trace dump is generated, it needs to be uploaded to an FTP server. Otherwise, the next trace dump will overwrite the existing one.

The Trace tab allows you to set the FTP server upload parameters for a trace dump as follows:

- ▶ FTP host server IP address, directory, and login credentials
- ▶ Automatic or manual trace dump upload
Use this option to enable automatic trace dump upload as soon as the trace dump is generated.

Figure 8-96 shows the trace dump upload settings. The Trace Dump Availability section of this panel displays information about the last trace dump taken and whether the dump was uploaded automatically to the FTP server.

The screenshot shows the 'IBM_SAN80B_217 - Switch Administration' web interface. The 'Trace' tab is selected, and the 'Trace Dump Availability' section is visible. The 'Trace FTP Host' section contains fields for Host IP (10.64.210.103), Remote Directory (SAN80B), User Name (ibm), and Password (masked with dots). The 'Trace Dump Availability' section shows the trace dump generation time as 'Mon Jul 20 21:07:08 2009' and a checkbox for 'Trace Auto FTP Uploaded' which is currently unchecked. The 'Auto FTP Upload' section has radio buttons for 'Enable' (selected) and 'Disable'. At the bottom right are 'Apply', 'Close', and 'Refresh' buttons. The status bar at the bottom shows 'Enable Auto FTP upload', 'Mode: Advanced', a link to 'Free Professional Management Tool', and user information: '10.64.210.217 | AD0 | User: admin | Role: admin | ✓'.

Figure 8-96 Switch Administration: Trace tab

8.7.13 Security Policies tab

You use the Security Policies tab to configure the Access Control List (ACL) policies. Fabric OS v6.1.0 supports the following ACL policy types:

- ▶ Fabric Configuration Server (FCS) policy
- ▶ Device Connection Control (DCC) policy
- ▶ Switch Connection Control (SCC) policy
- ▶ IP Filter policy

Figure 8-97 shows an example of the Security Policies tab.

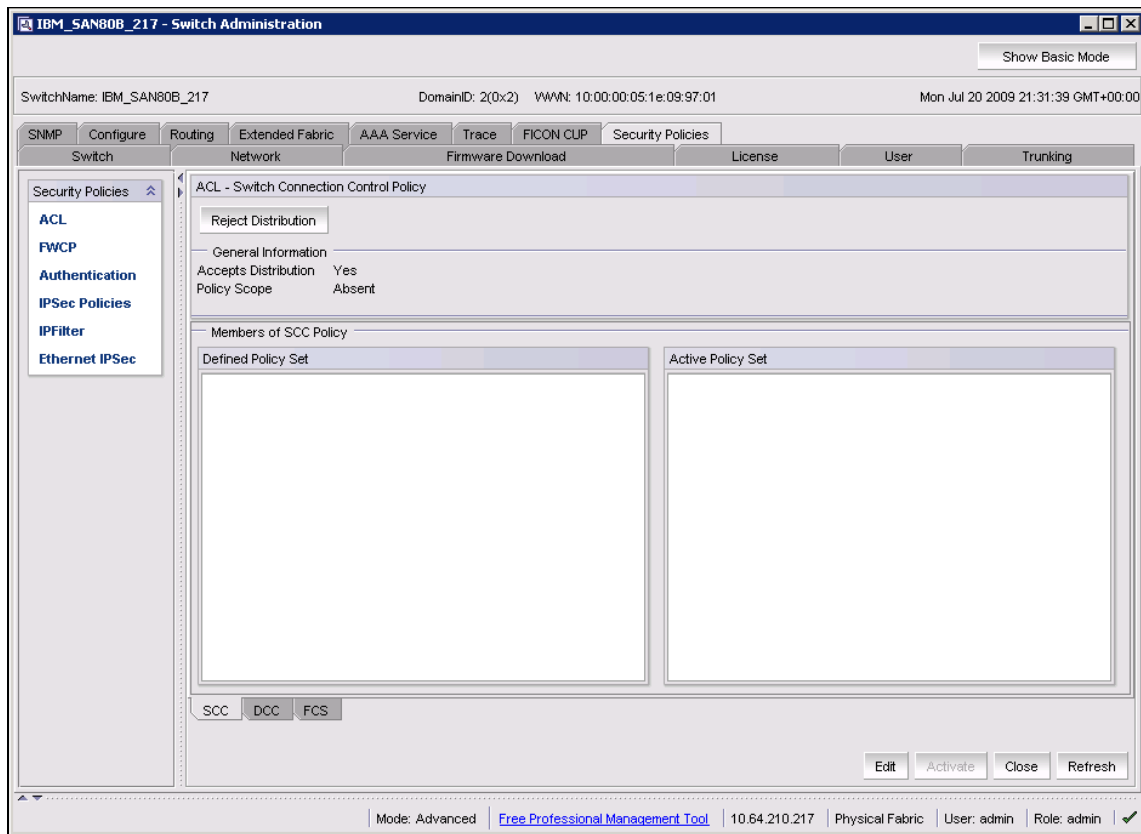


Figure 8-97 Switch Administration: Security Policies tab

From this tab, you can add new FCS, SCC, DCC, and IP Filter policies by clicking **Edit**. You must also activate the policies after you create them. We explain the different policies in the following sections.

Attention: All the actions in this section require the user to be logged in to Administrative Domain (AD) 255 with the suggested role. If Administrative Domains have not been implemented, log in to AD0.

FCS policy

The FCS policy is used to restrict which switches can perform the fabric-wide configuration changes. If you do not set up and activate this policy, then any switch can change fabric configuration.

If the FCS policy is active, then the following rules apply:

- ▶ If the FCS list contains only one switch, there will be no backup FCS switches. In case the FCS switch is unavailable, the fabric is left without an FCS switch.
- ▶ If multiple FCS switches are defined, the first switch in the list is designated as primary FCS switch. Others are backup FCS switches. In case of primary FCS switch failure, the next switch on the FCS list becomes the new primary.

Figure 8-98 shows an example of adding two switches to the FCS list.

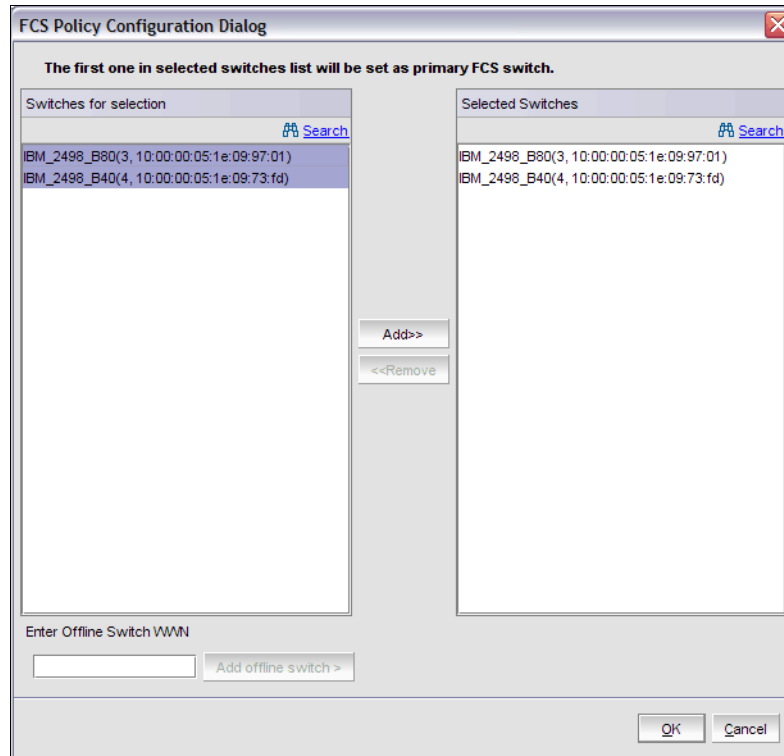


Figure 8-98 Adding switches to FCS policy list

DCC policy

You use DCC policies to specify which Fibre Channel devices can connect to which Fibre Channel switch ports. If no DCC policy is defined, then any device can attach to any switch port in the fabric. DCC policy is configured by specifying the device port WWN and the switch port to which it can connect. You can define multiple DCC policies, and you can use particular device WWNs and switch ports in several DCC policies. This way, you can create a set of ports that a certain device is allowed to use, and a set of devices that can connect to a certain port.

Devices that are not listed in any DCC policy can connect to any switch port that is not specified in a DCC policy. Proxy devices can always connect to any switch port in the fabric. Setting up a DCC policy has no effect on these devices.

DCC policies names must always have a prefix DCC_POLICY_. The total allowed length of names is 30 characters, including the mandatory prefix. This leaves up to 19 alphanumeric or underscore characters to select a unique DCC policy name.

Figure 8-99 shows an example of defining a DCC policy.

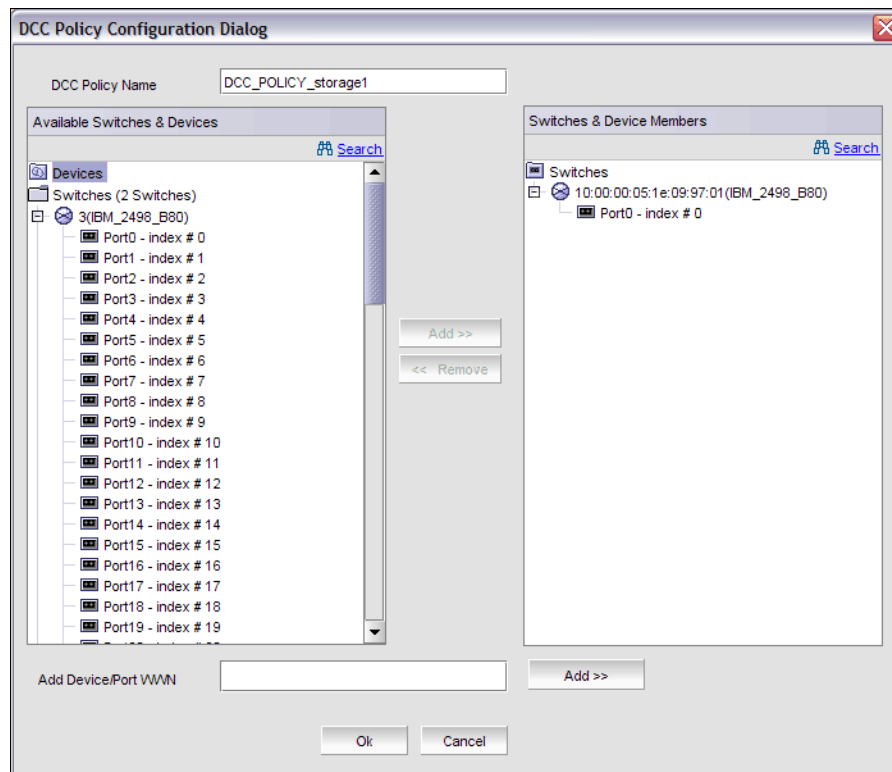


Figure 8-99 Defining a DCC policy

SCC policy

You use the SCC policy to control which switches can join the fabric. As opposed to DCC policy (where you can define multiple policies), there can be only one SCC policy, and its name must be SCC_POLICY.

If the SCC policy does not exist, then any switch can join the fabric; however, if you have defined the policy, then only the switches that are included in the policy can be fabric members. Use WWN, Domain ID, or switch names to indicate the member switches.

IP Filter policy

You can use the IP Filter policy to secure the IP management interfaces. By configuring these policies, you can set up a firewall which permits or denies the IP management traffic based on the policy rules. You can have up to 256 rules within an IP Filter policy, and each rule includes:

- ▶ The source IP address or an address group prefix
- ▶ The destination port number or name (for example, Telnet, SSH, HTTP, and so forth.)
- ▶ Protocol type (TCP or UDP)
- ▶ Filtering action for the rule (permit or deny)

For example, you can configure a policy to permit Telnet access only from a certain IP address.

To provide separate packet filtering for IPv4 and IPv6 addressing, two IP Filter policy types exist. Each policy type can have up to six policies defined, but only one policy per type can be activated.

Figure 8-100 shows an example of creating an IP Filter policy.

Create IP Filter Policy

Policy Details

Policy Name:

Policy Type:

IP Filter Rules

Item: 1 [Export](#) [Copy](#) [Search](#)

Rule Order ▲ ¹	Source IP	Service/Destination Port	Protocol	Permission
1	10.64.223.10	TELNET	tcp	Permit

Figure 8-100 Creating an IP Filter policy

8.7.14 FICON CUP tab

The FICON CUP tab within Web Tools allows for FICON configuration. This topic is beyond the intended scope of this book.

8.7.15 Trunking tab

We discuss the Trunking tab in 13.1.3, “ISL Trunking” on page 590.

8.8 Telnet/SSH Client task

This task allows you to connect to the switch using Telnet or the SSH client. To initiate connection:

1. Click **Telnet/SSH Client** in the Tasks panel (see Figure 8-101).



Figure 8-101 Telnet/SSH Client task

2. In the Preference Dialog box, shown in Figure 8-102, select either the Telnet or SSH client, and provide the path to the client utility (for example, PuTTY). Then, click **OK**.

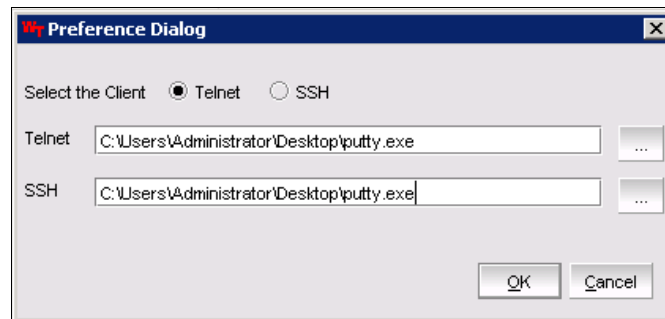


Figure 8-102 Telnet/SSH Preference Dialog box

3. In our example, we use PuTTY as the Telnet Client. The utility then launches, as shown in Figure 8-103. After entering IP address of the switch, PuTTY establishes connection and we are prompted to login.

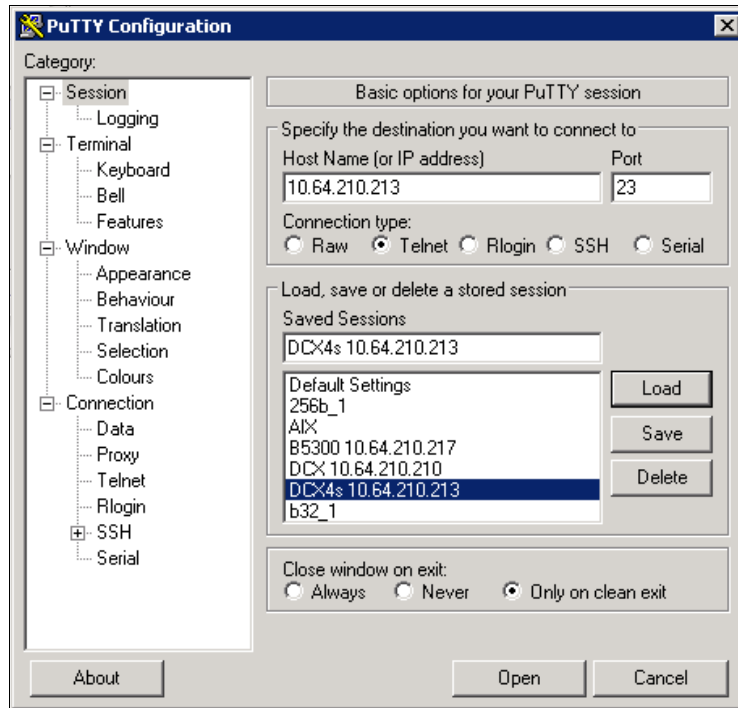


Figure 8-103 PuTTY as SSH client

8.9 Fabric Watch task

To access the Fabric Watch function, click **Fabric Watch** from the Tasks panel, as shown in Figure 8-104.



Figure 8-104 Fabric Watch button

The Fabric Watch window opens, as shown in Figure 8-105.

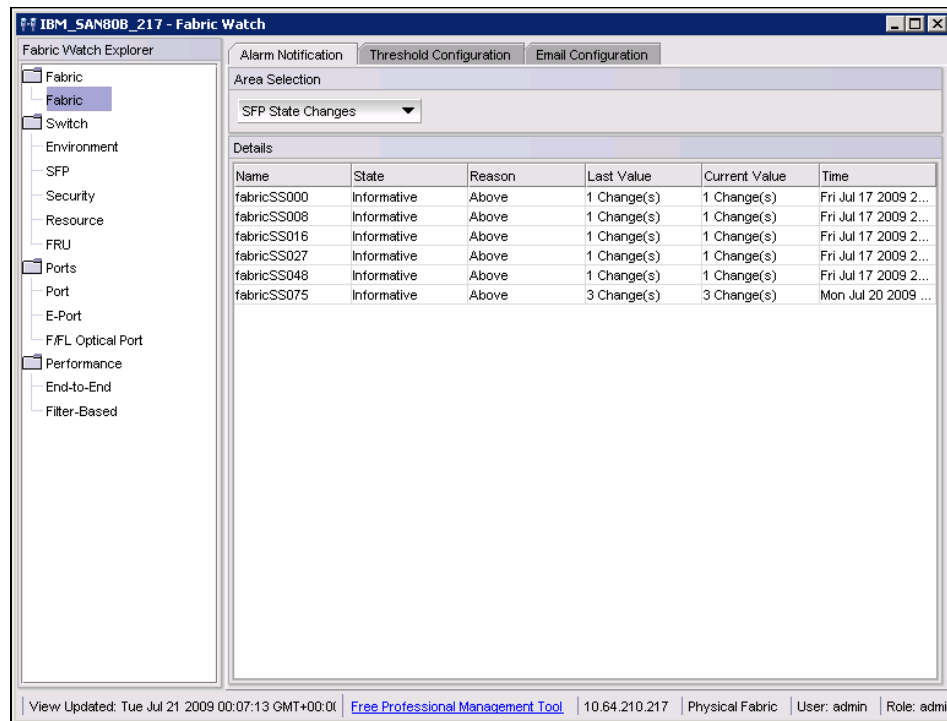


Figure 8-105 Fabric Watch initial view

The window is divided into two sections. The left-hand side has a tree structure that lists the *Classes* that can be monitored using Fabric Watch. If you expand the *Classes*, all the *Areas* that are associated with a particular *Class* are displayed.

The main part of the window on the right-hand side has a display with the following tabs:

- ▶ Alarm Notification
- ▶ Threshold Configuration
- ▶ Email Configuration

Also on the right-hand side is the Area Selection, which contains a context driven drop down menu which is used to select elements that are linked to the Area selected in the left-hand window.

8.9.1 Alarm Notification tab

Use the Alarm Notification tab to view the information for all elements of the Fabric Watch, Fabric, or Performance Monitor classes. The information displayed includes:

- ▶ The name of the fabric
- ▶ The last event state
- ▶ The last event reason
- ▶ The last event value
- ▶ The current value
- ▶ The last event time

The Alarm Notification tab refreshes the displayed information according to the threshold configuration.

Figure 8-106 shows the Alarm Notification tab.

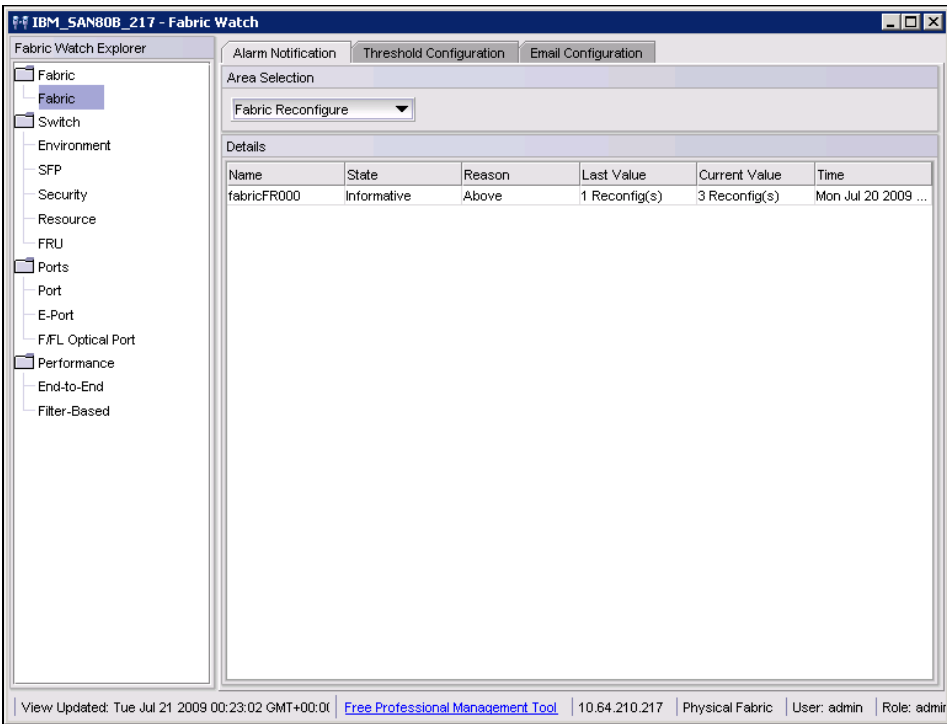


Figure 8-106 Fabric Watch Alarm Notification

8.9.2 Threshold Configuration tab

Use the Threshold Configuration tab to view and configure Fabric Watch thresholds for the Fabric Watch class currently selected in the organizational tree on the left side of the window. Figure 8-107 shows the Threshold Configuration tab.

The screenshot shows a web-based configuration interface for IBM SAN80B_217 - Fabric Watch. The window has a title bar and a menu bar with 'Alarm Notification', 'Threshold Configuration', and 'Email Configuration'. The 'Threshold Configuration' tab is active. On the left is a 'Fabric Watch Explorer' tree with nodes: Fabric, Switch, Environment, SFP, Security, Resource, FRU, Ports, Port, E-Port, F/FL Optical Port, Performance, End-to-End, and Filter-Based. The 'Fabric' node is selected. The main area has sub-tabs: 'Trait Configuration', 'Alarm Configuration', 'Element Configuration', and 'Configuration Report'. The 'Trait Configuration' sub-tab is active. It contains a 'System Default' and a 'Custom Defined' column for configuring thresholds. The 'Unit' is 'Segmentation(s)', 'Time Base' is 'None', 'Low Boundary' is '0', 'High Boundary' is '0', and 'Buffer Size' is '0'. The 'Activate Level' section has radio buttons for 'System Default' (selected) and 'Custom Defined'. At the bottom right are 'Apply' and 'Refresh' buttons. The status bar at the bottom shows: 'View Updated: Tue Jul 21 2009 00:26:52 GMT+00:00', a link to 'Free Professional Management Tool', IP '10.64.210.217', 'Physical Fabric', 'User: admin', and 'Role: admin'.

	System Default	Custom Defined
Unit	Segmentation(s)	Segmentation(s)
Time Base	None	None
Low Boundary	0	0
High Boundary	0	0
Buffer Size	0	0

Activate Level
☒ System Default ☐ Custom Defined

Apply Refresh

View Updated: Tue Jul 21 2009 00:26:52 GMT+00:00 | [Free Professional Management Tool](#) | 10.64.210.217 | Physical Fabric | User: admin | Role: admin

Figure 8-107 Configure Thresholds

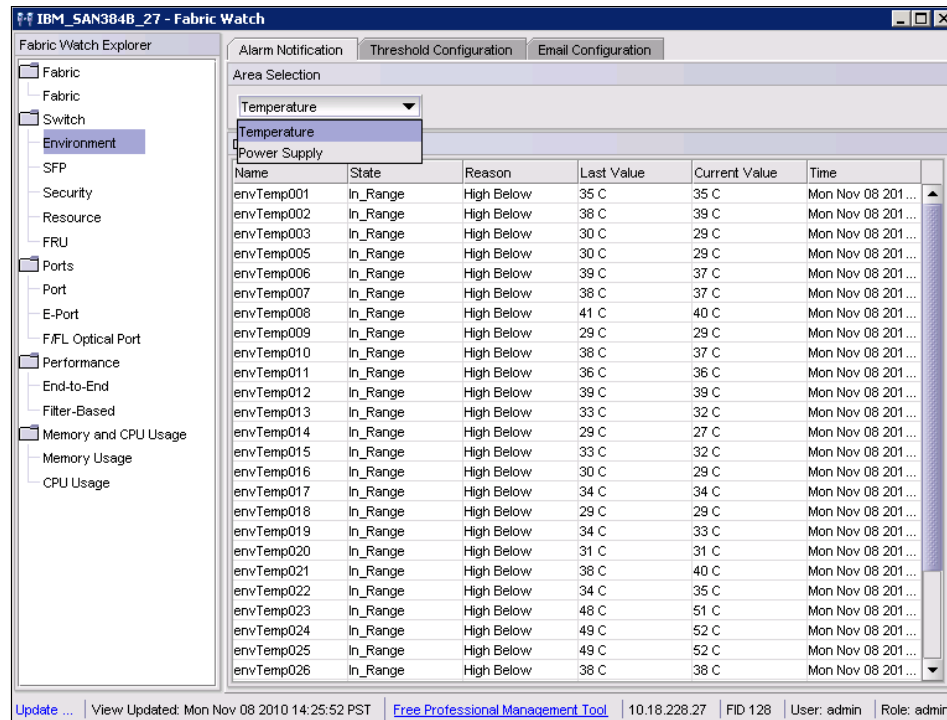
The Threshold Configuration display changes according to the Class and Area selected in the organizational tree. However, the Threshold Configuration tab always includes the same options, as follows.

- ▶ **System Default:** Click to return settings to default values.
- ▶ **Custom Defined:** Specify new settings.
- ▶ **Apply:** Click to apply the values specified in the current display.
- ▶ **Refresh:** Refresh view with current information from switch.

Important: When making changes in a given window, they are not saved until you click **Apply**.

Thresholds for the Environmental classes

The Environmental classes are displayed by highlighting **Environment** in the panel on the left and then clicking the Threshold Configuration tab as shown in Figure 8-108.



The screenshot shows the 'IBM_SAN384B_27 - Fabric Watch' application window. The 'Threshold Configuration' tab is active. On the left, the 'Fabric Watch Explorer' tree has 'Environment' selected. The main area shows a table of thresholds for the 'Temperature' area.

Name	State	Reason	Last Value	Current Value	Time
envTemp001	In_Range	High Below	35 C	35 C	Mon Nov 08 201...
envTemp002	In_Range	High Below	38 C	39 C	Mon Nov 08 201...
envTemp003	In_Range	High Below	30 C	29 C	Mon Nov 08 201...
envTemp005	In_Range	High Below	30 C	29 C	Mon Nov 08 201...
envTemp006	In_Range	High Below	39 C	37 C	Mon Nov 08 201...
envTemp007	In_Range	High Below	38 C	37 C	Mon Nov 08 201...
envTemp008	In_Range	High Below	41 C	40 C	Mon Nov 08 201...
envTemp009	In_Range	High Below	29 C	29 C	Mon Nov 08 201...
envTemp010	In_Range	High Below	38 C	37 C	Mon Nov 08 201...
envTemp011	In_Range	High Below	36 C	36 C	Mon Nov 08 201...
envTemp012	In_Range	High Below	39 C	39 C	Mon Nov 08 201...
envTemp013	In_Range	High Below	33 C	32 C	Mon Nov 08 201...
envTemp014	In_Range	High Below	29 C	27 C	Mon Nov 08 201...
envTemp015	In_Range	High Below	33 C	32 C	Mon Nov 08 201...
envTemp016	In_Range	High Below	30 C	29 C	Mon Nov 08 201...
envTemp017	In_Range	High Below	34 C	34 C	Mon Nov 08 201...
envTemp018	In_Range	High Below	29 C	29 C	Mon Nov 08 201...
envTemp019	In_Range	High Below	34 C	33 C	Mon Nov 08 201...
envTemp020	In_Range	High Below	31 C	31 C	Mon Nov 08 201...
envTemp021	In_Range	High Below	38 C	40 C	Mon Nov 08 201...
envTemp022	In_Range	High Below	34 C	35 C	Mon Nov 08 201...
envTemp023	In_Range	High Below	48 C	51 C	Mon Nov 08 201...
envTemp024	In_Range	High Below	49 C	52 C	Mon Nov 08 201...
envTemp025	In_Range	High Below	49 C	52 C	Mon Nov 08 201...
envTemp026	In_Range	High Below	38 C	38 C	Mon Nov 08 201...

At the bottom of the window, there is a status bar with the following information: Update ... View Updated: Mon Nov 08 2010 14:25:52 PST Free Professional Management Tool 10.18.228.27 FID 128 User: admin Role: admin

Figure 8-108 Environmental Thresholds

The panel contains tabs that you can use to define how you intend to monitor the environmental factors of the switch:

- ▶ Traits
- ▶ Alarm Configuration
- ▶ Element Configuration
- ▶ Configuration Report

Each tab contains an Area Selection pull-down menu to select the Fabric Watch area. In the example in Figure 8-108, we selected **Temperature**.

Table 8-11 describes the values and information about the Traits tab.

Table 8-11 Traits values and information

Value	Description
Unit	The string used to define the unit of measurement for the area
Time base	The time base for the area
Low Boundary	The low threshold for the event setting comparison
High Boundary	The high threshold for the event setting comparison
Buffer size	Size of the buffer zone in the event setting comparison
Activate level	Radio button to use System Default settings or Custom Defined settings
Apply	Apply the new values to the switch
Refresh	Refresh view with current information from the switch

Thresholds for the SFP classes

You display the SFP classes by highlighting **SFP** in the panel on the left and then clicking the **Threshold Configuration** tab. The **Area Selection** pull-down menu displays the classes to be configured, as shown in Figure 8-109.

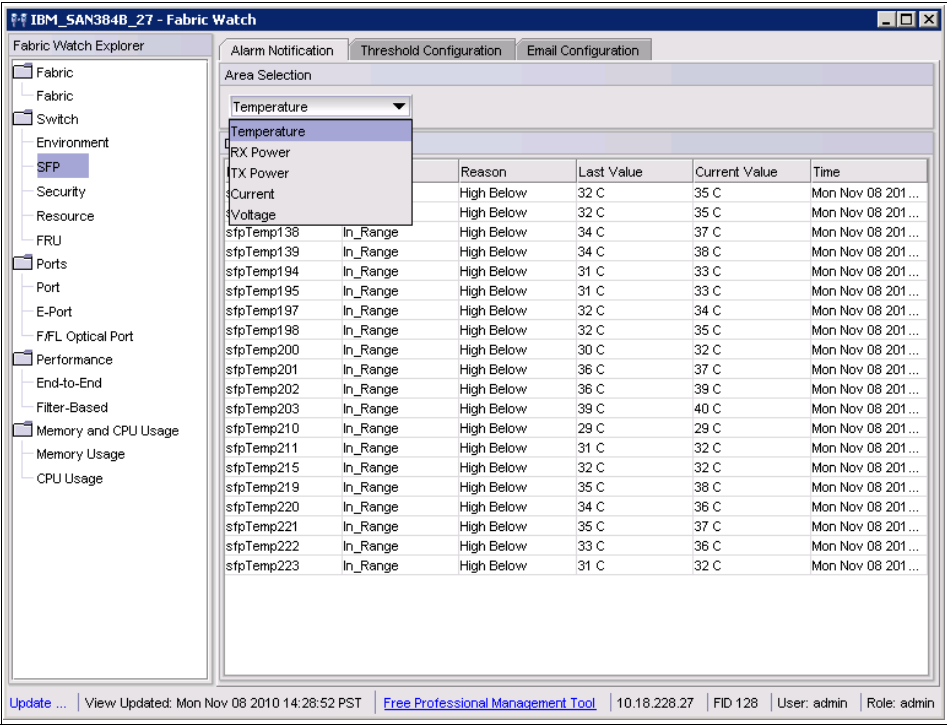


Figure 8-109 SFP thresholds

The available areas are Temperature, RX Power, TX Power, Current, and Voltage.

The Alarm Configuration tab has two areas to show the Default settings and the Customer defined settings as described in Table 8-12.

Table 8-12 Alarm Configuration settings

Value	Description
Changed	Event of counter changed
Below / Low Below	Event of counter fell below low boundary
Low Above	Event of counter fell above low boundary
Above / High Above	Event of counter fell above high boundary
In between / High Below	Event of counter is between the high/low boundaries
ERROR_LOG	Event notification to error log
SNMP_TRAP	Event notification through SNMP trap
RAPI_TRAP	Event notification through RAPI trap
EMAIL_ALERT	Event notification through email
System Default	Radio button indicating system defaults taken
Custom Defined	Radio button indicating custom defined
Apply	Apply the new values to the switch
Refresh	Refresh view with current information from the switch

Thresholds for Port classes

The Port, E_Port, F/FL Copper Port, F/FL Optical Port classes display the following fields for each area:

- ▶ Link Loss
- ▶ Sync Loss
- ▶ Signal Loss
- ▶ Protocol Error
- ▶ Invalid Words
- ▶ Invalid CRCs
- ▶ RX Performance
- ▶ TX Performance
- ▶ State Changes

Figure 8-110 shows the thresholds for the Port class.

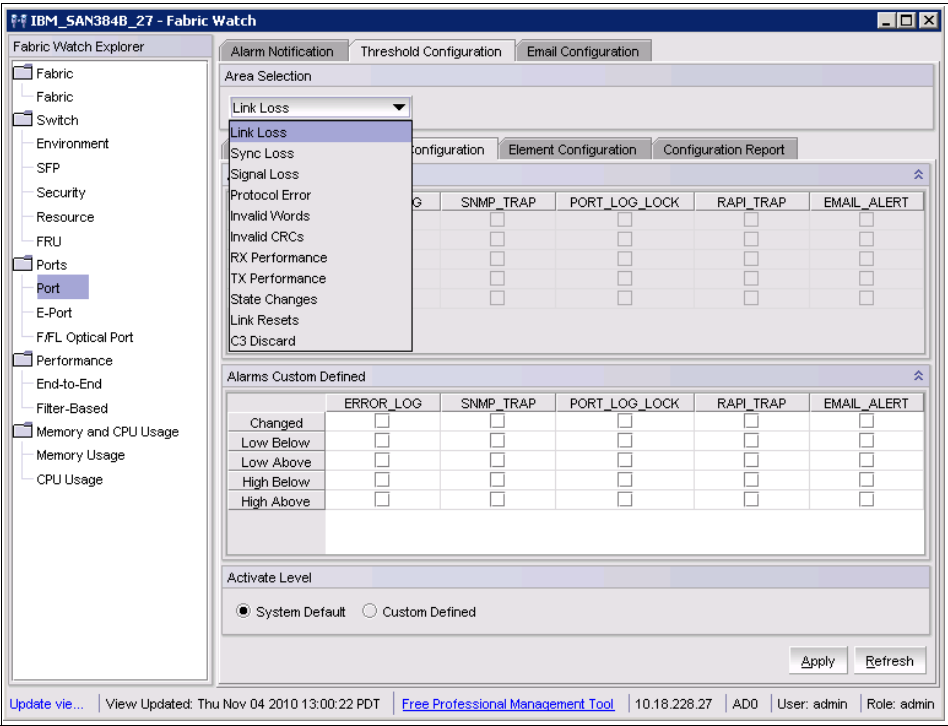


Figure 8-110 Port thresholds

Thresholds for Performance classes

Use the Threshold Configuration tab to view and configure End-to-End thresholds for the Performance class that is currently selected in the organizational tree on the left side of the window.

Be aware that you must define the SID/DID pair through the Performance Monitor before you can monitor the threshold in the End-to-End class. Figure 8-111 shows the Threshold Configuration tab for the End-to-End Thresholds.

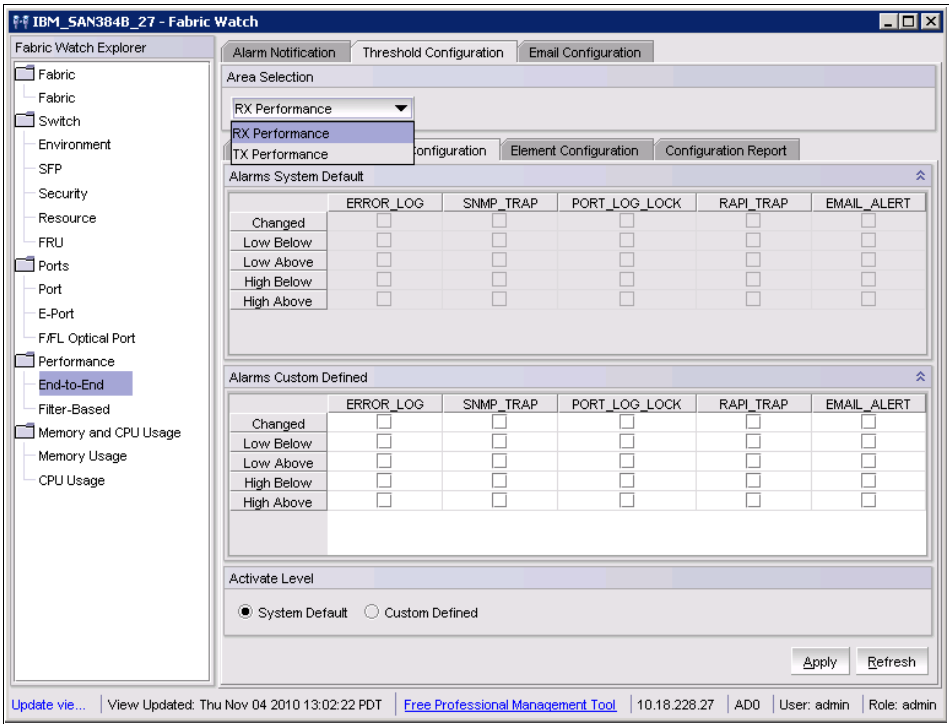


Figure 8-111 Threshold Configuration tab for End-to-End class

Use the Threshold Configuration tab to view and configure Filter-Based thresholds for the Performance class currently selected in the organizational tree on the left side of the window, as shown in Figure 8-112.

Thresholds: You must predefine the filter type in the Performance Monitor before you can use the Filter-Based thresholds.

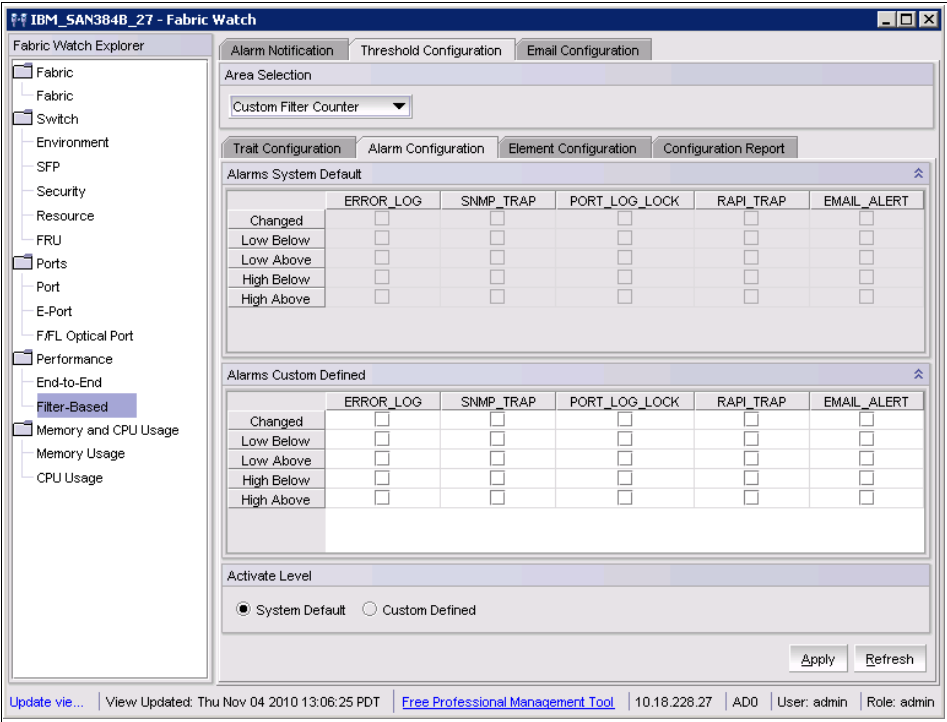


Figure 8-112 Threshold Configuration tab with Filter-Based class

8.9.3 Configuration Report tab

Use the Configuration Report tab to view the current Fabric Watch threshold parameters for the area selected in the Fabric Watch tree.

Figure 8-113 shows the Configuration Report tab for the Port class.

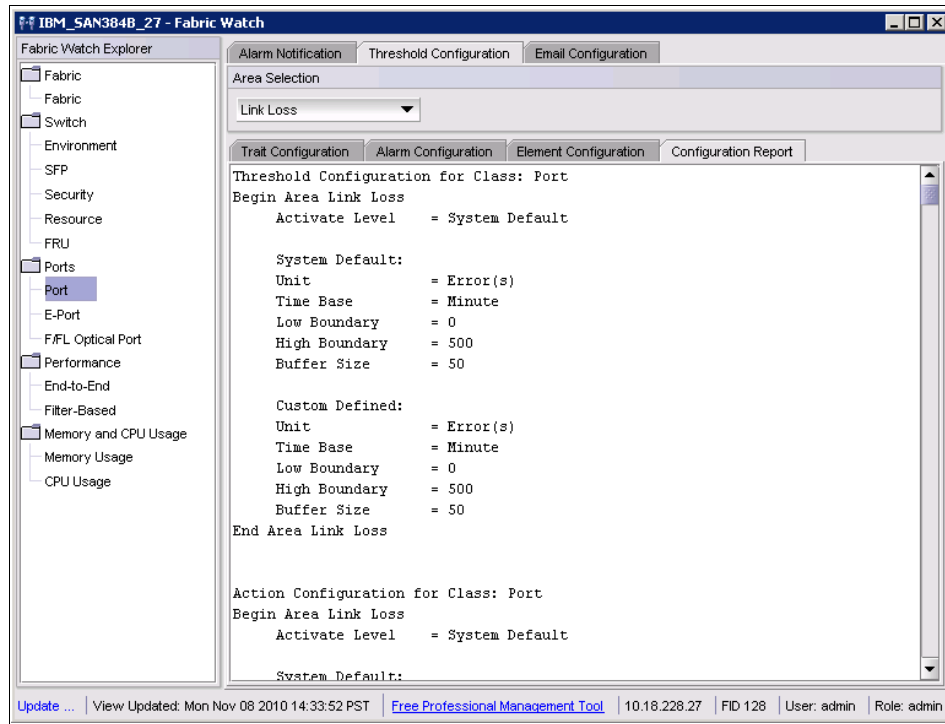


Figure 8-113 Port class -Configuration report

8.9.4 Memory and CPU Usage monitor with Fabric Watch

Fabric Watch can be defined to monitor and alert for memory and CPU usage. From the Fabric Watch window, click **Memory usage** or **CPU usage**, which has two tabs:

- ▶ Trait configuration: Defines the polling interval, threshold level, and number of retries to be done before alerting.
- ▶ Alarm Configuration: Defines the way or type of alert to be triggered.

Figure 8-114 and Figure 8-115 indicate these Fabric Watch configurations for Memory or CPU usage.

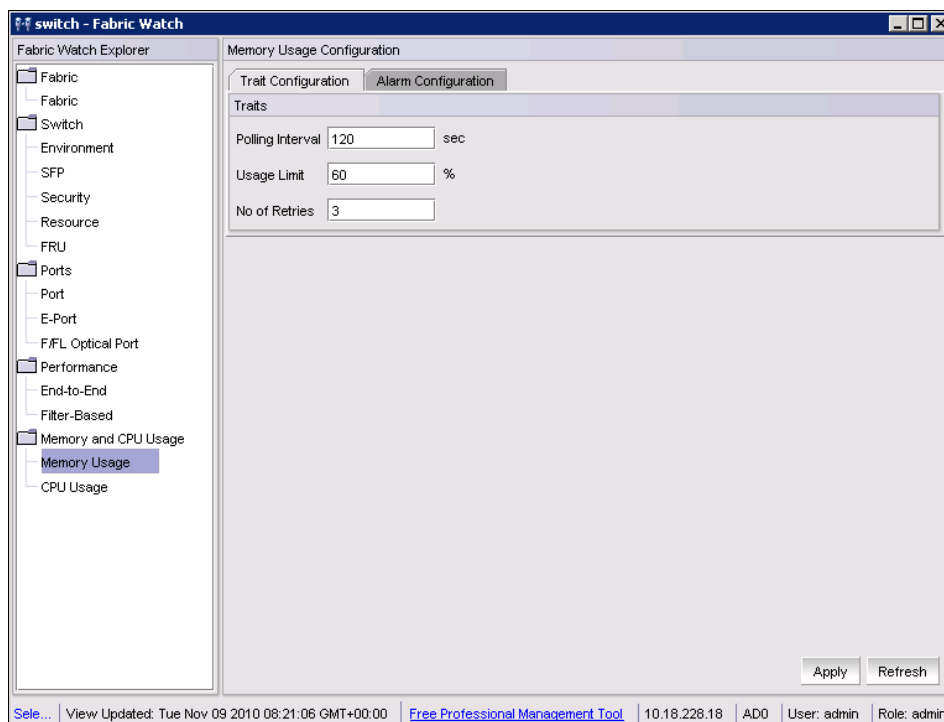


Figure 8-114 Fabric Memory usage trait configuration

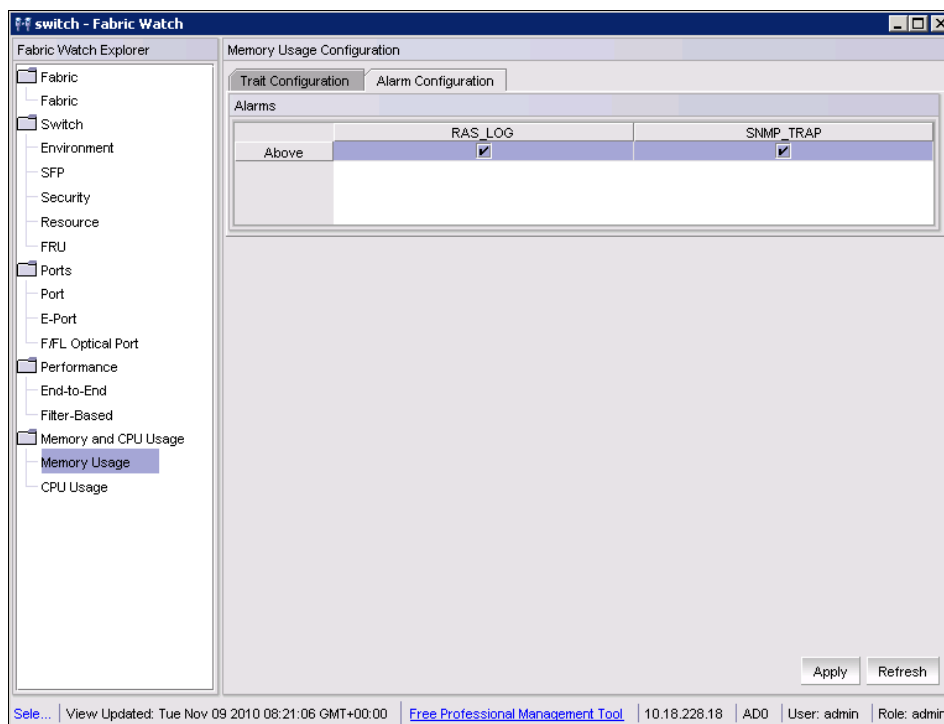


Figure 8-115 Fabric Watch alarm configuration - memory usage

8.9.5 Modifying settings for switches with one power supply

The IBM default settings for Fabric Watch cause a switch with a single power supply to appear yellow in the Web Tools, indicating a *MARGINAL* status. The status can also be obtained by clicking **Status** in the switch view to open a window that describes the cause of the marginal state, as shown in Figure 8-116.

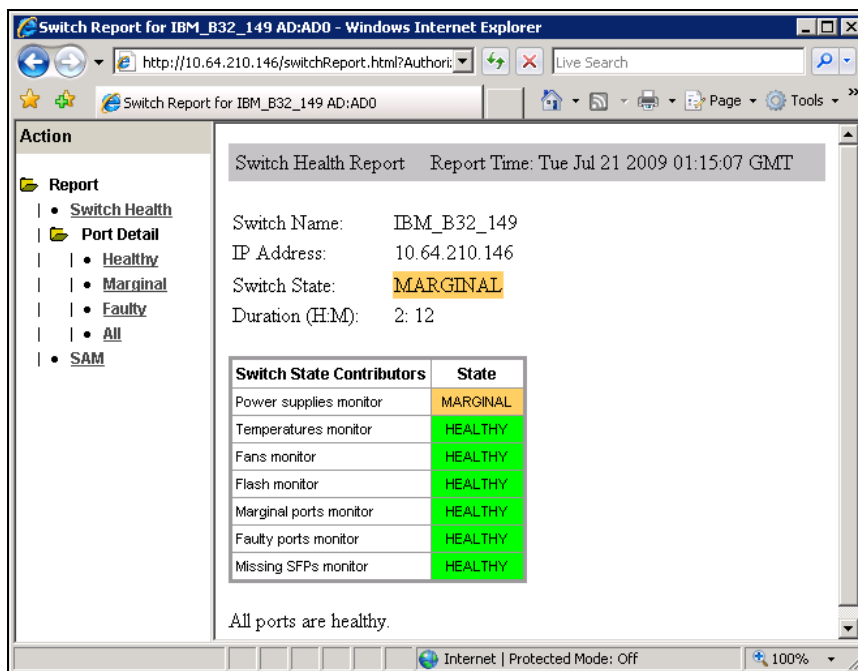


Figure 8-116 Checking the switch status

The switch status can be changed to *HEALTHY* using a Telnet connection. Use the **switchStatusShow** command to display the current health of the switch. After using **switchStatusPolicySet** to clear the current condition, again use **switchStatusShow** to demonstrate that a switch with only one power supply is then shown with a *HEALTHY* status. See Example 8-9 for details.

Example 8-9 Using switchStatusPolicySet to clear unnecessary marginal status

```
IBM_B32_149:admin> switchStatusShow
Switch Health Report                                Report time: 07/21/2009
01:18:09 AM
Switch Name:    IBM_B32_149
IP address:     10.64.210.146
SwitchState:    MARGINAL
Duration:       02:15

Power supplies monitor  MARGINAL
Temperatures monitor   HEALTHY
Fans monitor           HEALTHY
Flash monitor          HEALTHY
Marginal ports monitor  HEALTHY
Faulty ports monitor   HEALTHY
```

Missing SFPs monitor HEALTHY

All ports are healthy

IBM_B32_149:admin> switchStatusPolicySet

To change the overall switch status policy parameters

The current overall switch status policy parameters:

	Down	Marginal

PowerSupplies	2	1
Temperatures	2	1
Fans	2	1
Flash	0	1
MarginalPorts	2	1
FaultyPorts	2	1
MissingSFPs	0	0

Note that the value, 0, for a parameter, means that it is NOT used in the calculation.

** In addition, if the range of settable values in the prompt is (0..0),

** the policy parameter is NOT applicable to the switch.

** Simply hit the Return key.

The minimum number of

Bad PowerSupplies contributing to DOWN status: (0..2) [2] 0
Bad PowerSupplies contributing to MARGINAL status: (0..2) [1] 0
Bad Temperatures contributing to DOWN status: (0..5) [2]
Bad Temperatures contributing to MARGINAL status: (0..5) [1]
Bad Fans contributing to DOWN status: (0..3) [2]
Bad Fans contributing to MARGINAL status: (0..3) [1]
Out of range Flash contributing to DOWN status: (0..1) [0]
Out of range Flash contributing to MARGINAL status: (0..1) [1]
MarginalPorts contributing to DOWN status: (0..32) [2]
MarginalPorts contributing to MARGINAL status: (0..32) [1]
FaultyPorts contributing to DOWN status: (0..32) [2]
FaultyPorts contributing to MARGINAL status: (0..32) [1]
MissingSFPs contributing to DOWN status: (0..32) [0]
MissingSFPs contributing to MARGINAL status: (0..32) [0]


```
Policy parameter set has been changed
IBM_B32_149:admin> switchStatusShow
Switch Health Report                                Report time: 07/21/2009
01:20:28 AM
Switch Name:      IBM_B32_149
IP address:       10.64.210.146
SwitchState:      HEALTHY
Duration:         00:00

Power supplies monitor  HEALTHY
Temperatures monitor   HEALTHY
Fans monitor           HEALTHY
Flash monitor          HEALTHY
Marginal ports monitor HEALTHY
Faulty ports monitor   HEALTHY
Missing SFPs monitor   HEALTHY
```

All ports are healthy

To change the default settings, issue the **switchStatusPolicySet** command.

The first section of response to the command is the same as though you issue the **switchStatusPolicyShow** command and displays a list of the current settings. Here, you can see that the *Power Supplies* line is defined to be *MARGINAL* if the switch is powered by one power supply. These default settings assume that the switch has two power supplies and that one has failed. Obviously, for a switch purchased with a single power supply, this is not valid.

You are then prompted to enter the new values for each setting, starting with the *DOWN* value for the Faulty Ports, then the *MARGINAL* value for Faulty Ports. You can simply press Enter or type the same number to use default values. Then, you are prompted for the next setting, and eventually, for the Power supply *DOWN* and *MARGINAL* values.

Enter zero for the number of *bad power supplies contributing to the DOWN status* as well as zero for the number of *bad power supplies contributing to the MARGINAL status*. Indeed, because we are working with only one power supply, if it goes down, then the whole switch goes down. There is no marginal status.

At the bottom of the Telnet display, after running the **switchStatusShow** command, you can see that the chassis status has changed from *MARGINAL* to *HEALTHY*.

8.9.6 Email Configuration

Use the Email Configuration tab to configure the destination email ID to receive any alerts selected in the threshold configuration. The Email Configuration tab is shown in Figure 8-117. Also on this tab, you can enable or disable the email function for Fabric Watch alerts and send a test email to ensure that the function is working.

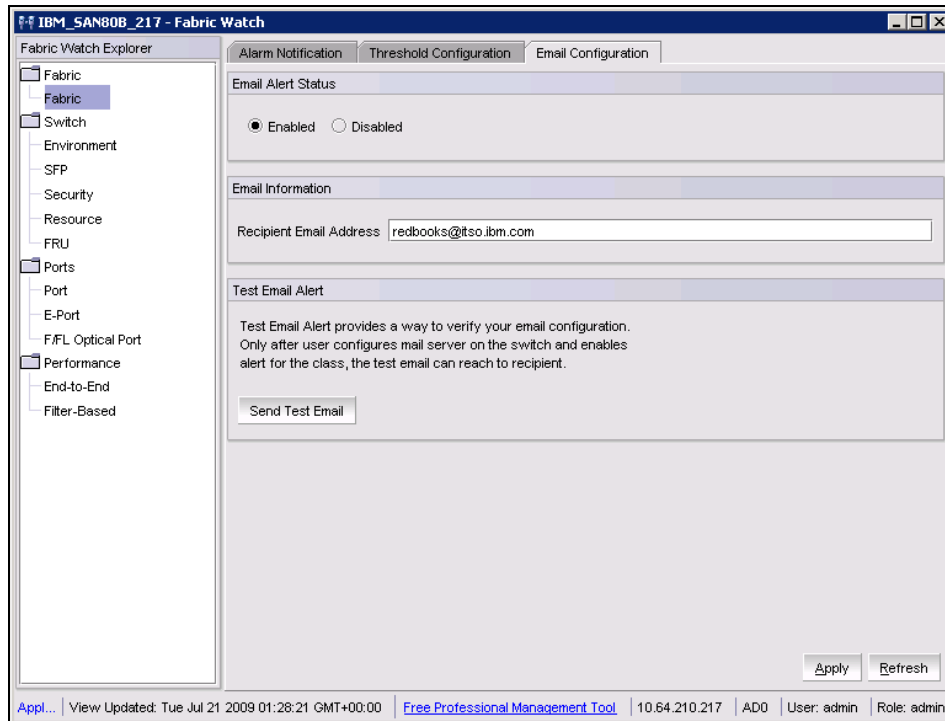


Figure 8-117 Email Configuration tab

8.10 IBM SAN ICL connectivity

In this section, we explain the procedure to interconnect two IBM SAN768B or IBM SAN384B chassis using the ICL cables. Web Tools is the tool that you can use to verify the state of fabric and both chassis after the ICL connectivity is established.

8.10.1 Before you begin

Prior to merging, look at each chassis with Web Tools.

In our example, the first chassis is named SAN384B_213. Figure 8-118 shows the Web Tools status of this chassis.

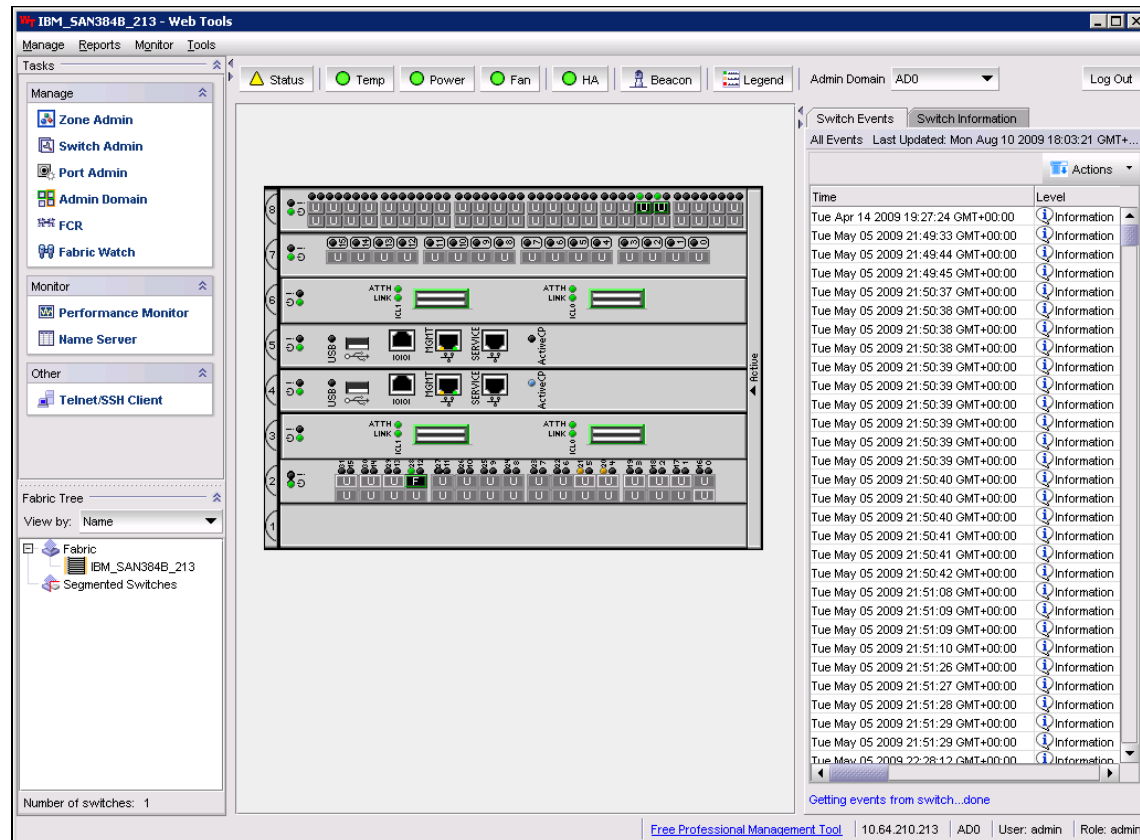


Figure 8-118 First chassis: SAN384B_213

The chassis is managed through IP address 10.64.210.213. This SAN384B has principle role in the fabric and is the only switch in the fabric.

Next, look at the second chassis, named SAN384B_215, as shown in Figure 8-119.

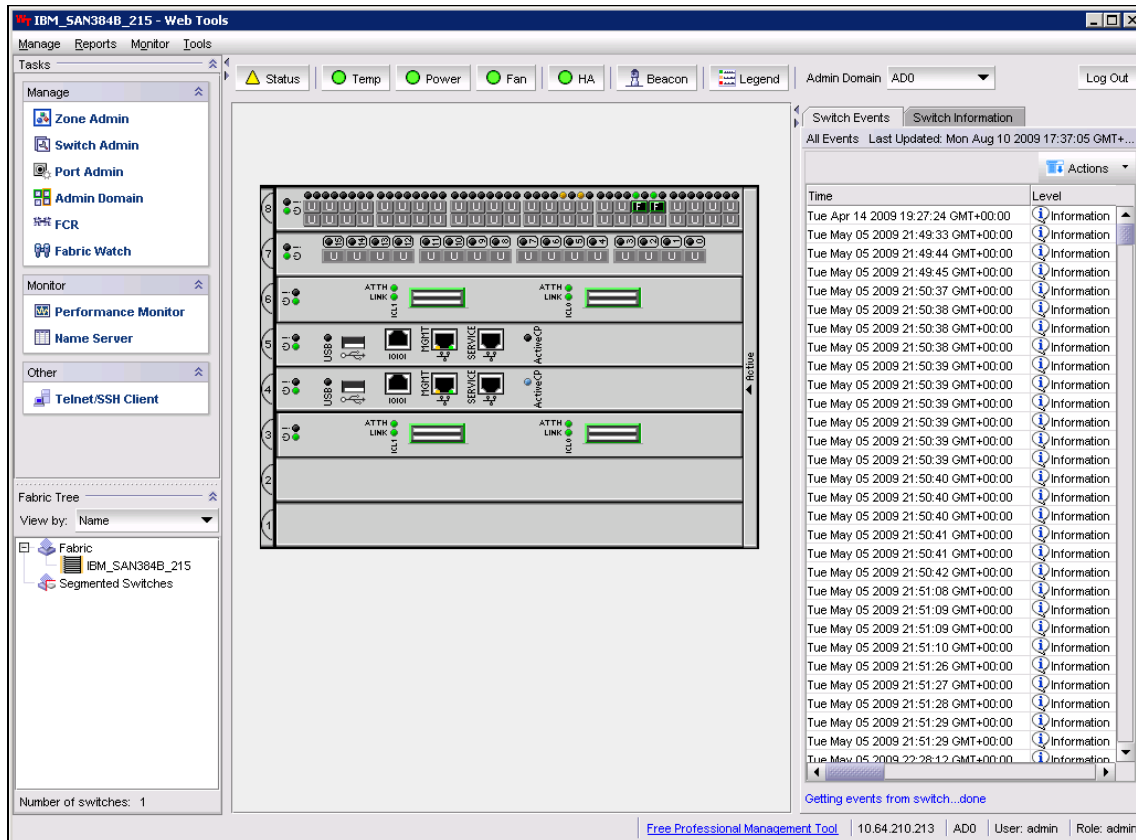


Figure 8-119 Second chassis: SAN768B_2

We use IP address 10.64.210.215 to manage this chassis. The SAN384B_215 has principal role in its fabric and is actually the only member of the fabric.

8.10.2 ICL cabling

The ICL cable connectors must be connected to corresponding ICL ports on the core (CR8) blades. The two core blades are installed in slots 3 and 6. Because each core blade has two ICL ports, we need four ICL cables.

The basic rules for proper ICL cabling are as follows:

- ▶ ICL0 ports must be cabled to ICL1 ports.
- ▶ ICL1 ports must be cabled to ICL0 ports. It is not allowed to attach ICL0 to ICL0 or ICL1 to ICL1.

Figure 8-120 shows an example of correct cabling. See Chapter 2, “Data Center Fabric” on page 21 for other ICL interconnection possibilities.




Figure 8-120 SAN384B: ICL cabling

Following these rules, you can interconnect the two chassis. After a few minutes, Web Tools indicates that the fabrics have merged, as shown in Figure 8-121.

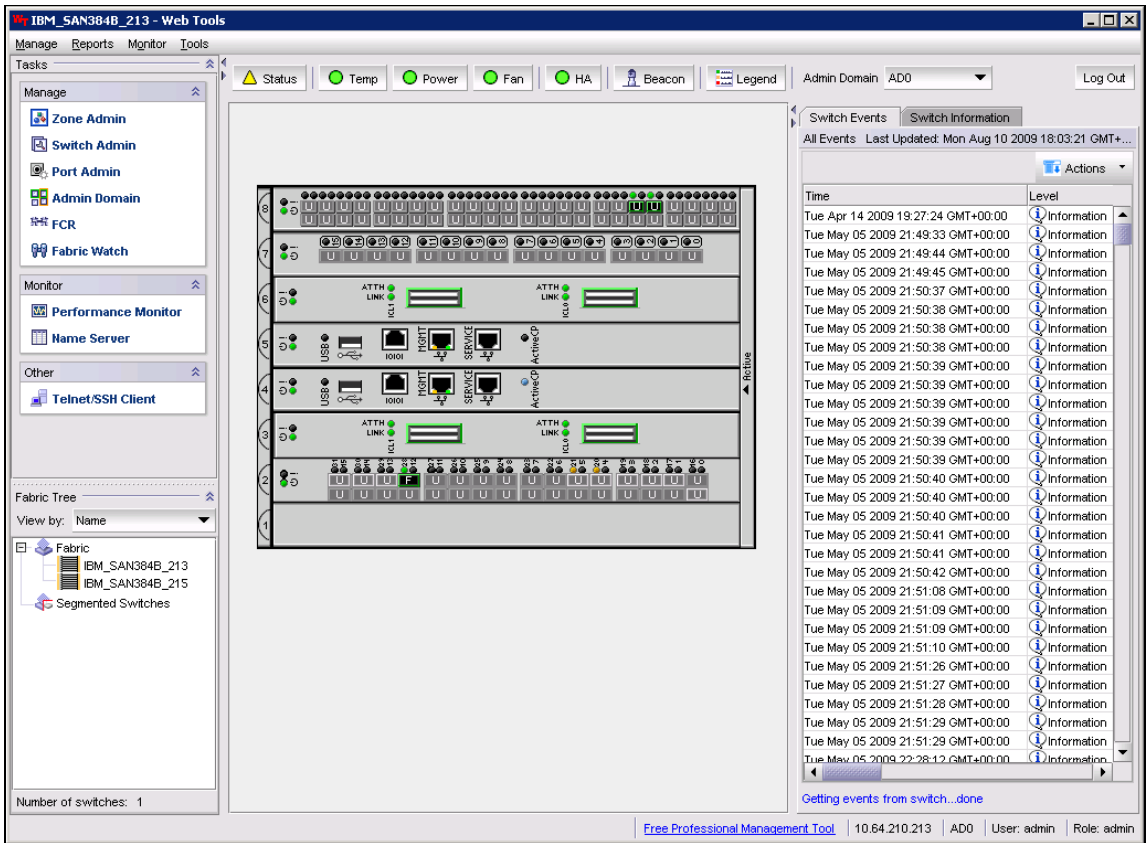


Figure 8-121 SAN384B_213 after connecting the ICL cables

SAN384B_213 is still in the principal role, but the fabric now has two members, as shown in Figure 8-122.




Figure 8-122 Fabric members

You can now check the status of the second chassis, SAN384B_215.
As shown in Figure 8-123, the chassis is now a member of the same fabric as SAN384B_213, and its role has been changed to subordinate.

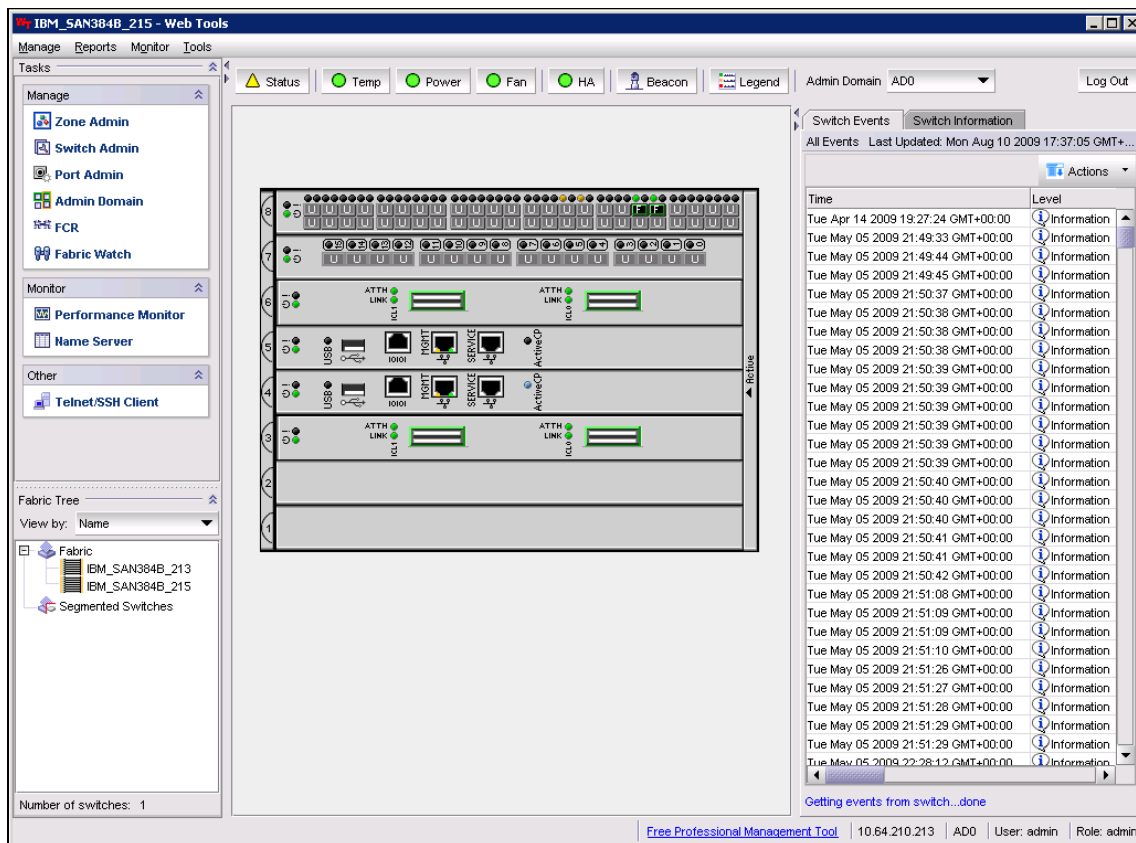


Figure 8-123 SAN384B_215 after connecting ICL cables

Finally, look at the ICL ports. In the Web Tools window for SAN384B_213, click **Port Admin** to launch the Port Administration applet.

The two core blades are installed in slots 3 and 6. Select ports on the core blade in slot 3, as shown in Figure 8-124.

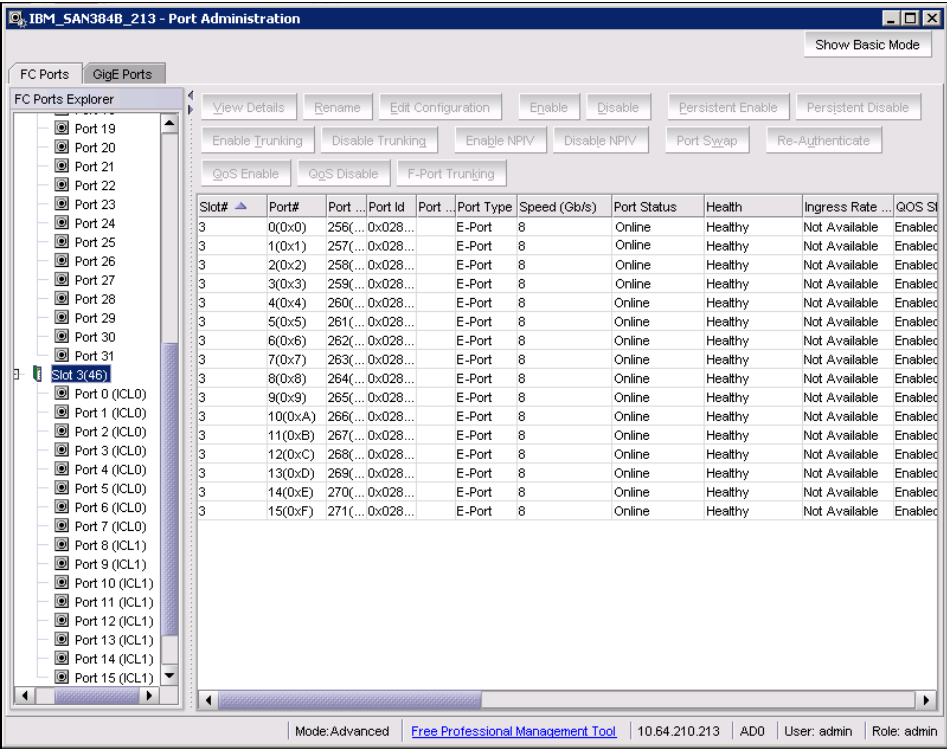


Figure 8-124 ICL ports

Now, the ICL ports all have E_Port type, are fixed to 8 Gbps, and are grouped in eight-port trunks.



IBM System Storage Data Center Fabric Manager

A key component of the IBM/Brocade DCF (Data Center Fabric) architecture is a new product, IBM System Storage Data Center Fabric Manager (DCFM). This is an end-to-end fabric management software platform that combines the capabilities of IBM/Brocade existing fabric management software:

- ▶ Enterprise Fabric Connectivity Manager (EFCM)
- ▶ Fabric Manager

The DCFM architecture integrates the best management features of EFCM and Fabric Manager. It is based on EFCM's Graphical User Interface (GUI) and Fabric Manager's messaging and data management design for improved performance and scalability.

In this chapter we explore several key capabilities, to help data center managers and administrators better understand the advantages of deploying IBM System Storage Data Center Manager (DCFM).

There are two types of DCFM: *DCFM Professional* and *DCFM Enterprise*. We briefly describe the differences between them and concentrate on the features and capabilities of DCFM Enterprise.

9.1 DCFM products

As stated before, DCFM is available as the following types:

- ▶ DCFM Professional: Free
- ▶ DCFM Enterprise: Licensed by server

9.1.1 DCFM Professional

DCFm Professional is a free product and is targeted at customers seeking a less extensive management solution for smaller SANs. This software, which is included with IBM/Brocade switches, provides these capabilities:

- ▶ Allows management of a *single* Fabric OS (FOS) fabric (up to 1,000 switch ports) at a time.
- ▶ Performs group switch management beyond the scope of Web Tools.

You can download a free version of DCFM Professional from the Brocade webpage.

Support: Be aware that IBM provides no support for the DCFM Professional version. IBM will provide support for the licensed DCFM Enterprise version only.

Features *not supported* that are available in Enterprise Edition are as follows:

- ▶ Full IBM/Brocade Backbone management with features such as QoS and end-to-end performance monitoring
- ▶ Support for up to 24 physical fabrics, 9,000 switch ports, and 20,000 end devices
- ▶ FICON management for mainframe environments
- ▶ Fabric-based encryption
- ▶ Comprehensive FCR (Fiber Channel Routing) and FCIP management
- ▶ Advanced Call Home Support
- ▶ Support for security schemes (RADIUS, LDAP, Active Directory, NIS/NIS+, and more)
- ▶ Historical performance data collection
- ▶ Data persistence for up to two years of data, out-of-box Open Database Connectivity (ODBC), and Java Database Connectivity (JDBC) access
- ▶ M-EOS support
- ▶ Remote clients

There is limited support for other features of Enterprise Edition.

Notes for DCFM Professional:

- ▶ DCFM Professional does not support the IBM System Storage SAN768B, Brocade DCX, and m-EOS.
- ▶ DCFM Professional supports IBM System Storage SAN384B and Brocade DCX-4s.
- ▶ DCFM Professional does not allow use as a Fusion Agent proxy for management applications such as IBM TPC.

9.1.2 DCFM Enterprise Edition

DCFMEnterprise is an enterprise-class product targeted at customers that demand a management software solution with comprehensive support for the following capabilities:

- ▶ IBM/Brocade Backbone switch (SAN768B) - Data Center Fabric (DCF)
- ▶ Fabric-based encryption support for data-at-rest solutions
- ▶ End-to-end manageability of the data center fabric from HBA ports through switch ports to storage ports

Fusion Agent: DCFMEnterprise Edition allows use as a Fusion Agent proxy for management applications such as IBM TPC.

DCFMEnterprise provides multi-protocol networking support for the following products:

- ▶ Fibre Channel
- ▶ Fiber Connectivity (FICON)
- ▶ Fibre Channel over IP (FCIP)
- ▶ Fibre Channel Routing (FCR)
- ▶ Internet SCSI (iSCSI)
- ▶ (Future) Fibre Channel over Ethernet (FCoE) and Converged Enhanced Ethernet (CEE)

9.1.3 Enhanced Group Management

Both products require the Enhanced Group Management (EGM) license, which is an FOS license that enables multi-switch operations for those switches that have EGM enabled. This license is already the default in most switches except the 2498-24E, for example.

It helps automate operations across multiple switches to save time and streamline repetitive operations, which are typically prone to error. Brocade EGM drives consistency across fabrics, while minimizing the risk associated with potential downtime due to configuration mismatches. Enhanced Group Management has the following features:

- ▶ It offers a FOS enabled client for SAN switch group management.
- ▶ It drives consistency across SAN by automating repetitive tasks:
 - Downloads firmware to multiple switches and meets upgrade windows
 - Provisions new switches by uploading an existing switches configuration
 - Sets up a switch configuration once, then replicates it to multiple others
- ▶ It minimizes risk through troubleshooting, diagnostics, and monitoring:
 - Data collection across multiple switches for holistic problem resolution
 - Snapshots of issues as they occur for fast, effective root cause analysis
 - Regular back-up of last known configurations for quick restores
 - Monitors performance to proactively address over-utilized devices/links
- ▶ It ensures compliant security settings:
 - Can be set up once, then replicates across multiple switches the settings for SCC/DCC policies, IPFilter, LDAP, and RADIUS Server configurations

9.1.4 DCFM Enterprise scalability

The scalability of DCFM Enterprise is limited by the used switch types.

- ▶ Pure FOS Fabrics:
 - Monitor up to 24 physical fabrics with support for:
 - 120 switches
 - 9,000 switch ports
 - 20,000 hosts or storage devices
 - 40 Access Gateways
 - 5 minutes performance monitoring polling

- ▶ Mixed Fabrics (FOS and M-EOS):
 - Monitor up to 24 physical fabrics with support for:
 - 60 switches
 - 5,000 switch ports
 - 10,000 hosts or storage devices
 - 40 Access Gateways
 - 5 minutes performance monitoring polling
- ▶ Pure M-EOS Fabrics:
 - Monitor up to 24 physical fabrics with support for:
 - 60 switches
 - 5,000 switch ports
 - 10,000 hosts or storage devices
 - 40 Access Gateways
 - 5 minutes performance monitoring polling

9.1.5 DCFM operating system support

In general, DCFM supports different configurations. The configuration support depends on the DCFM version and on the FOS and M-EOS level.

The following firmware platforms are supported by the release of DCFM 10.4.X:

- ▶ Fabric OS v5.0 or later in a pure Fabric OS fabric
- ▶ Fabric OS v6.0 or later in a mixed Fabric OS and M-EOS fabric
- ▶ M-EOS and M-EOSn 9.7 or later in a mixed Fabric OS and M-EOS fabric
- ▶ M-EOS and M-EOSn 9.9.2 or later in a pure M-EOS fabric

For details about requirements, see *Data Center Fabric Manager User Manual*, GC52-1304-03.

http://www-01.ibm.com/support/docview.wss?rs=1314&context=STBVU4&dc=DA400&uid=ssg1S7003231&loc=en_US&cs=utf-8&lang=en

9.2 DCFM installation

The installation of the DCFM Enterprise Edition is fairly straightforward. DCFM should be installed on a separate server. The requirements for the server depend on the platform and the size of the fabric and are described in the *Data Center Fabric Manager Installation Guide*, GA32-0786-00, which is available at this website:

<http://www-01.ibm.com/support/docview.wss?uid=ssg1S7003232>

9.2.1 Installation of DCFM Enterprise Edition on Windows platform

Follow these steps for the installation:

1. Insert the installation DVD into the DVD-ROM drive.
 - a. If autorun is enabled, the installer begins automatically.
 - b. If autorun is not enabled, open the following file:
`<DVD_drive>\DCFM_win\install.exe`

The Introduction panel displays (see Figure 9-1).

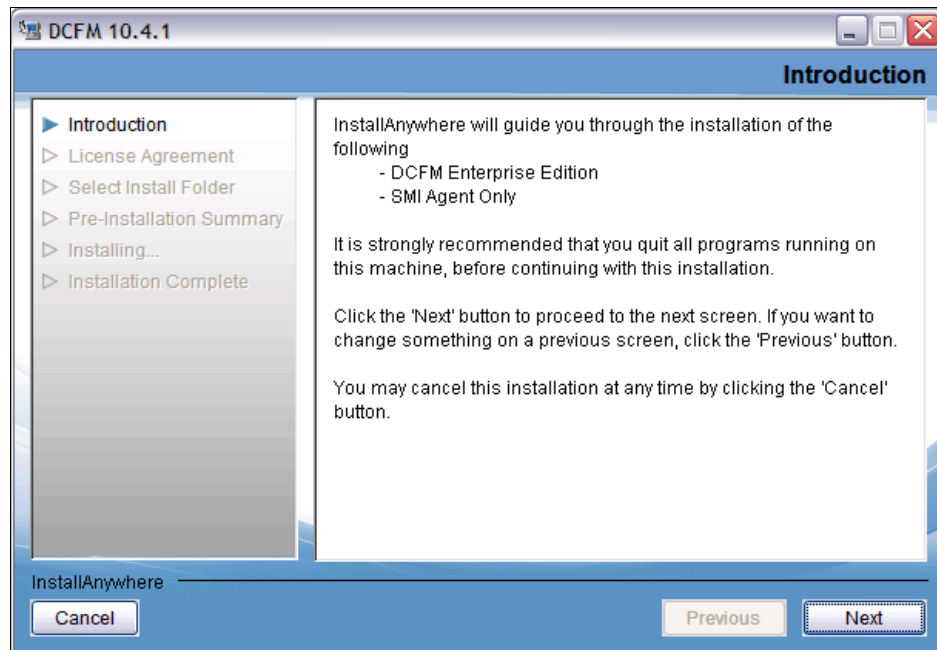


Figure 9-1 DCFM Introduction panel

2. Click **Next** on the Introduction panel.
3. Read the agreement on the **License Agreement** panel, select **I accept the terms of the License Agreement** and click **Next**.

The **Select Install Folder** displays (see Figure 9-2).

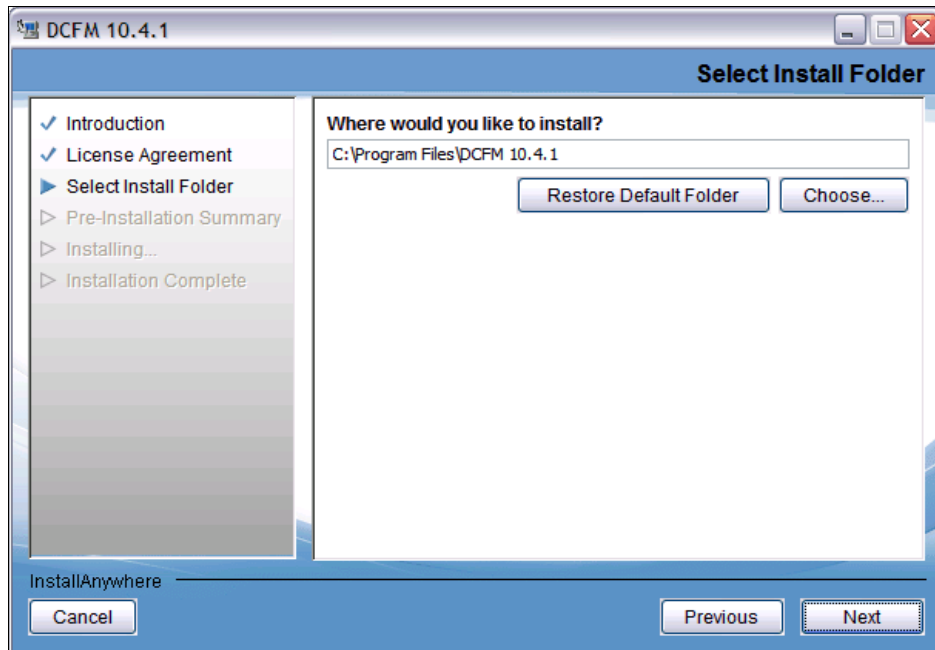


Figure 9-2 Select Install Folder dialog box

4. Select the usual location for your system's application files (for example, **C:\Program Files\DCFm 10.4.1**) on the Select Install Folder panel, and click **Next**.

Important: Do *not* install to the root directory (for example, C:\).

5. Review the displayed installation summary on the Pre-Installation Summary panel and click **Install** (see Figure 9-3).

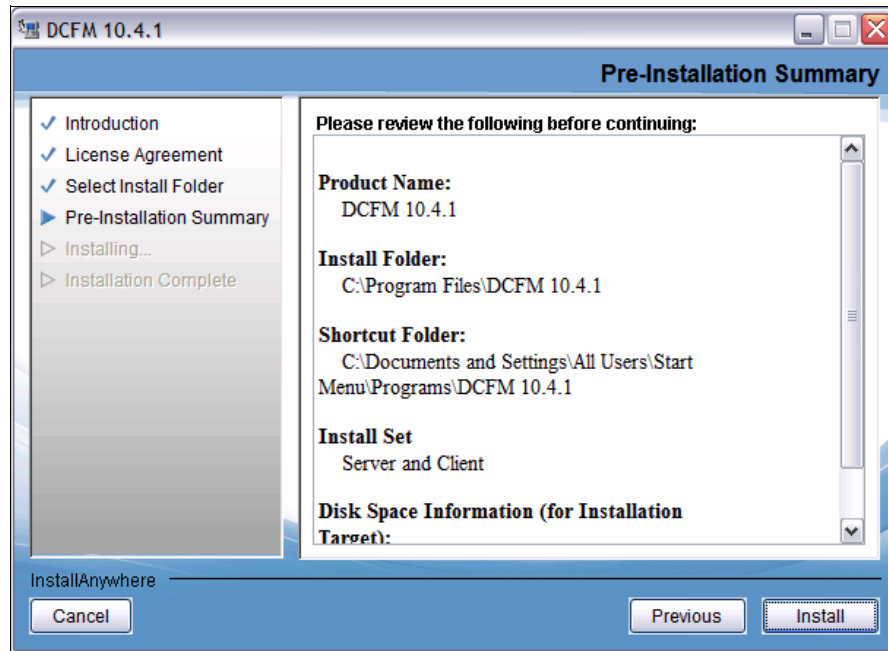


Figure 9-3 Pre-Installation Summary panel

6. The Installation Complete panel will be shown. Make sure that the Launch DCFM Configuration check box is selected (default) and click **Done**. This will start the DCFM itself.
7. The DCFM will start with a Welcome panel; click **Next**.
8. Select **No** on the Copy Data and Settings panel and then click **Next** (Figure 9-4).

Data migration: There is an option to migrate data from EFCM, FM, or an older DCFM version. We do not cover this topic here. Therefore, you can find more information at the following websites:

Data Center Fabric Manager Installation, Migration and Transition Guide:

http://www-01.ibm.com/support/docview.wss?rs=1314&context=STBVU4&dc=DA400&uid=ssg1S7003035&loc=en_US&cs=utf-8&lang=en

Data Center Fabric Manager Migration Guide:

http://www-01.ibm.com/support/docview.wss?rs=1314&context=STBVU4&dc=DA400&uid=ssg1S7003233&loc=en_US&cs=utf-8&lang=en

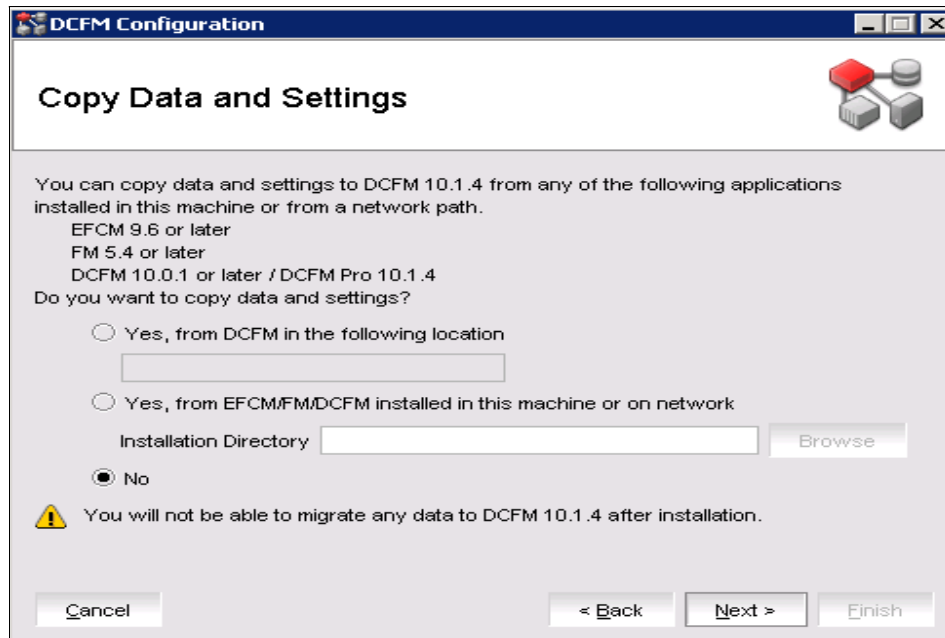


Figure 9-4 Copy Data and Settings dialog box

9. **Enter** the serial number (found on the DVD jewel case) and license key (on the Key Certificate) on the Server License panel, and click **Next**.

Server license: If your installation does not require a serial number and license key, the Server License panel does not display.

10. Select Internal FTP Server or External FTP Server on the FTP Server panel and click **Next** (see Figure 9-5).

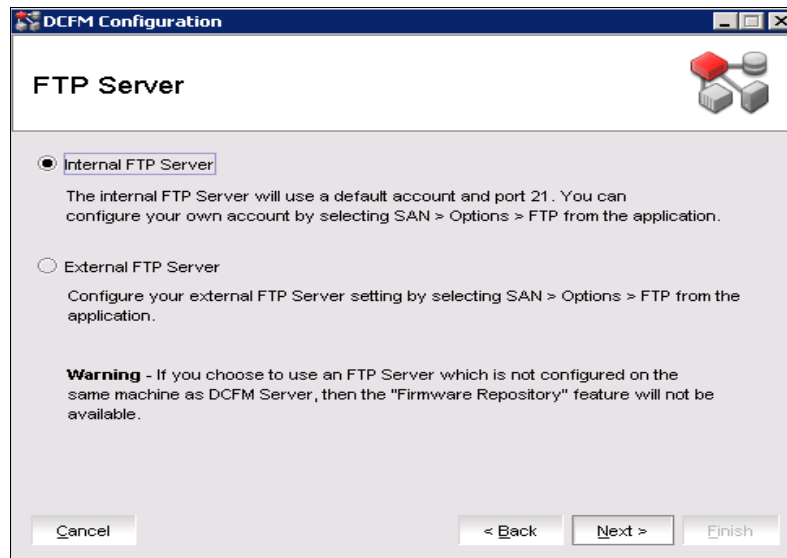


Figure 9-5 FTP Server Dialog box

Attention: If you use an FTP Server that is not configured on the same machine as the management application, the Firmware Repository feature will not be available.

11. Complete the following steps on the Server IP Configuration panel (see Figure 9-6) and click **Next**:
- Select an address from the Client - Server IP Configuration Return Address list.
 - Select an address from the Switch - Server IP Configuration Preferred Address list.

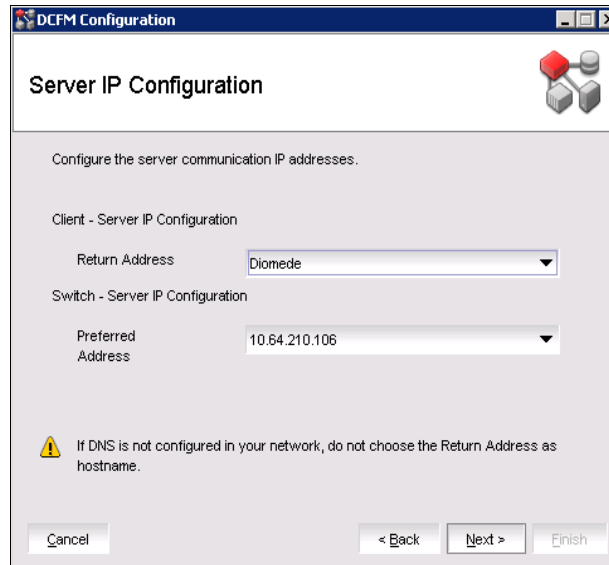


Figure 9-6 Server IP Configuration

12. Complete the following steps on the Server Port Configuration panel (see Figure 9-7):
 - a. Enter a port number in the Syslog Port Number field (default is 514).
 - b. *Enable SSL* by selecting the SSL Enabled check box.
 - c. Enter a port number in the Web Server Port Number field (default is 443 if SSL Enabled is selected; otherwise, the default is 80).
 - d. Enter a port number in the SNMP Port Number field (default is 162).
 - e. Enter a port number in the Starting Port Number field (default is 24600).
The server requires 16 consecutive free ports beginning with the starting port number.

Important: Do not use port 2638 for any of these port numbers. Port 2638 is used internally by the server.

DCFm Configuration

Server Port Configuration

DCFm requires Syslog, Web Server, and SNMP port numbers, as well as 16 consecutive port numbers from a "Starting port #".

Do not use port 2638 as it is required internally by the DCFM Server.

Syslog Port #

SSL Enabled ☐

Web Server Port #

SNMP Port #

Starting Port #

Change this configuration by selecting SAN > Options > Server Port from the application.

Figure 9-7 Server Port Configuration

13. Then click **Next**.
14. The SMI Agent Configuration panel come up. Enter the following data:
 - a. Enable the SMI Agent by selecting the Enable SMI Agent check box.
 - b. Enable the SLP by selecting the Enable SLP check box, if necessary.
 - c. Enable the SSL by selecting the Enable SSL check box, if necessary.
 - d. Enter the SMI Agent port number in the SMI Agent Port # field (default is 5989 if SSL Enabled is selected; otherwise, the default is 5988).
 - e. Click **Next**.
15. Select the option on the SAN Size panel. You have the following choices:
 - a. Small (managing up to 2000 ports, 1-20 domains)
 - b. Medium (managing up to 5000 ports, 21-60 domains)
 - c. Large (managing up to 9000 ports, 61-120 domains)

Ports: Port count is equal to the total number of switch ports across all fabrics.

16. Click **Next**.

17. Verify your configuration and license information on the Server License Summary panel and click **Next**.
16. Select the Start Client check box on the Start Server panel (see Figure 9-8).

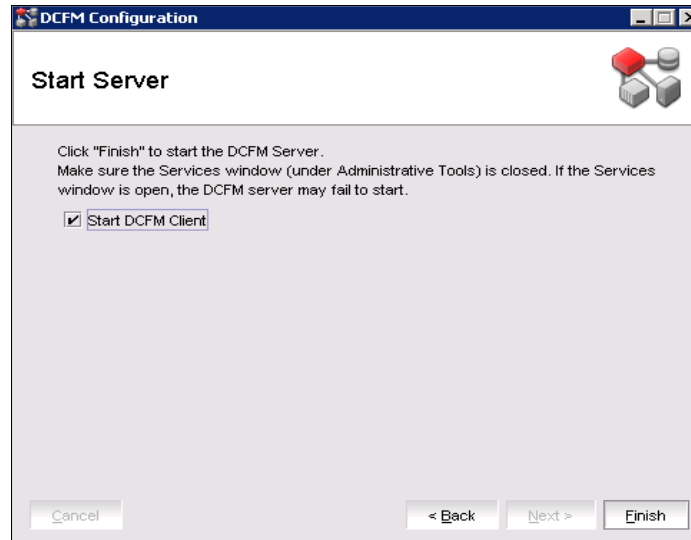


Figure 9-8 Start Server/Client dialog box

18. Click **Finish**.

After all of the DCFM services are started, the Log In dialog box displays. The default UserId is *Administrator* and the default password is *password*.

9.2.2 DCFM server and client

On a Windows platform, when experiencing problems with starting the DCFM client, you might need to use only one version of Java Runtime. As shown in Figure 9-9, you can select the needed version. If you have more than one version you can select it. The support version is 1.6.0_16 for DCFM.

Important: Use only one DCFM *server* for managing and monitoring your fabric. The use of more than one DCFM server for your fabric is *not supported*.

Start the DCFM *client* instead by typing in your browser:

http://<dcfm server ip address> on any host that you want to run the client.

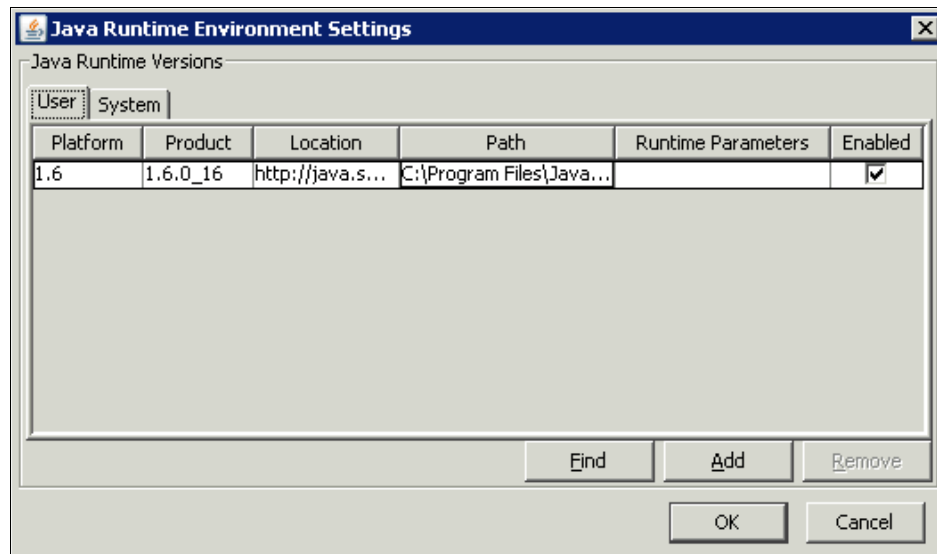


Figure 9-9 Java Runtime Environment Settings

9.3 DCFM GUI orientation

In the following sections we show the main features of the DCFM GUI.

Reference: For more details, see the *Data Center Fabric Manager User Manual*, GC52-1304-03, available at this website:

http://www-01.ibm.com/support/docview.wss?rs=1314&context=STBVU4&dc=DA400&uid=ssg1S7003231&loc=en_US&cs=utf-8&lang=en

9.3.1 Front panel

The management application's main window contains a number of areas. Figure 9-10 illustrates the various areas.

You can view all panels by selecting **View** → **Show Panels** → **All Panels**.

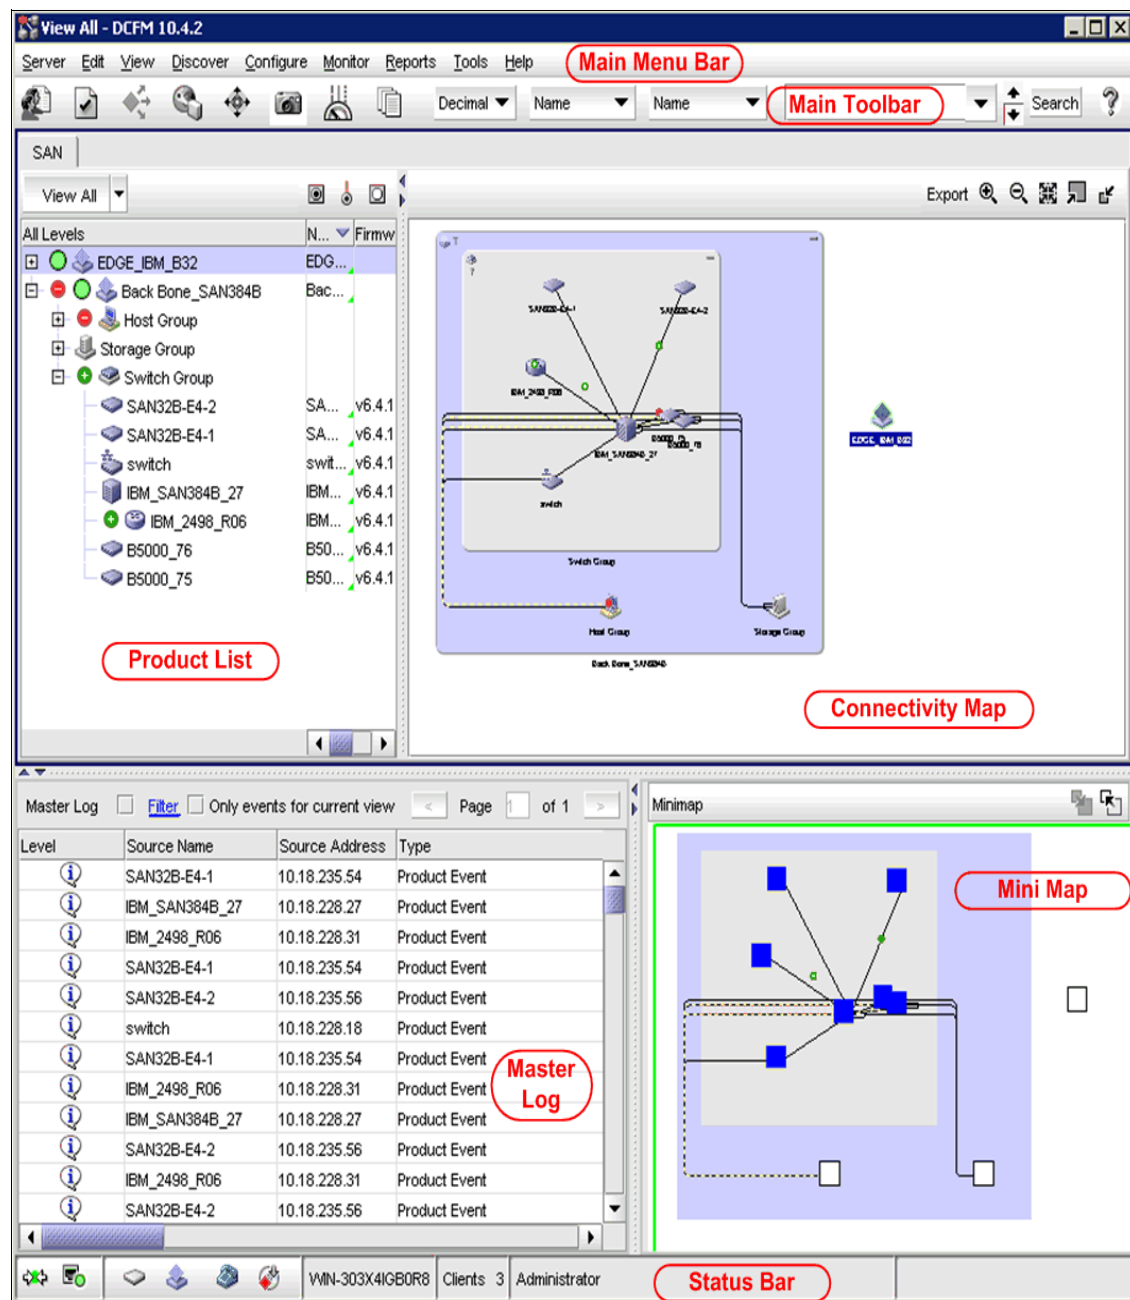


Figure 9-10 Front Panel

Table 9-1 shows the main components of the front panel and the functions they provide.

Table 9-1 DCFM Front Panel components

Component	Description
Menu Bar	Lists commands you can perform on the SAN.
Toolbar	Provides buttons that enable quick access to dialog boxes and functions.
Product List	Lists the devices discovered in the SAN.
Master Log	Displays all events that have occurred on the SAN.
Connectivity Map	Displays the SAN topology, including discovered and monitored devices and connections.
Minimap	Displays a “bird’s-eye” view of the entire SAN.
Status Bar	Displays data regarding the Server, connection, device, and fabric.

9.3.2 Main toolbar

The toolbar is located at the top of the main window and provides icons to perform the various functions as shown in Figure 9-11.

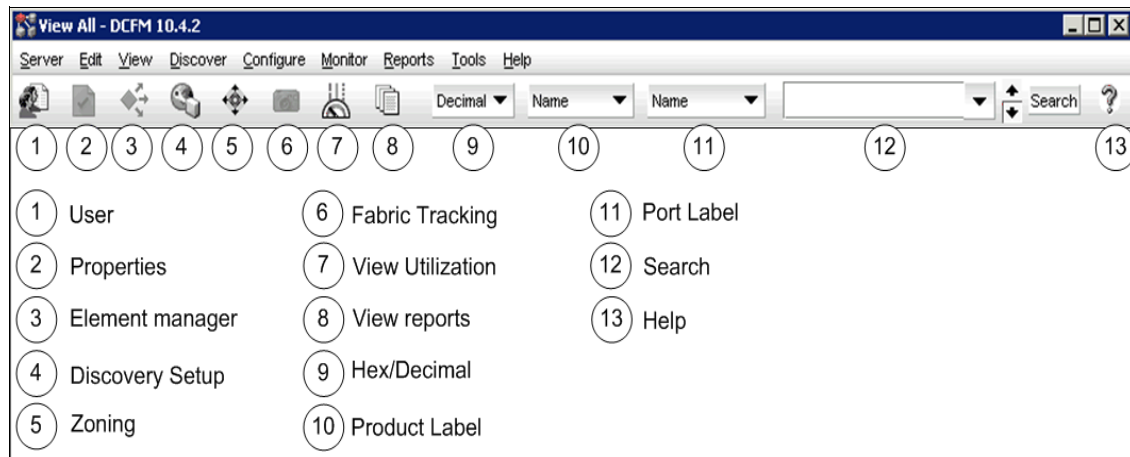


Figure 9-11 DCFM Main Toolbar

These functions are described in Table 9-2.

Table 9-2 Functions of the Main Toolbar

Icon	Description
Users	Displays the Server Users dialog box. Use to configure users, user groups, and permissions.
Properties	Displays the Properties dialog box of the selected device. Use to view or edit device properties.
Launch Element Manager	Launches the Element Manager of the selected device. Use to configure a device through its Element Manager.
Discovery Setup	Displays the Discover Setup dialog box. Use to configure discovery.
Zoning	Displays the Zoning dialog box. Use to configure zoning.
Fabric Tracking	Select to turn track fabric changes on and off for the selected device or group.
View Utilization	Displays or hides the utilization legend.
View Reports	Displays the View Reports dialog box. Use to view available reports.
Domain ID/Port #	Use to set the domain ID or port number to display as decimal or hex in the Connectivity Map.
Product Label	Use to set the product label for the devices in the Connectivity Map and product List.
Port Label	Use to set the port label for the devices in the Connectivity Map and Product List.
Product List Search	Use to search for a device in the product list.
Help	Displays the Online Help

9.3.3 Product list

The Product List displays an inventory of all discovered devices and ports. It is a quick way to look up product and port information, including serial numbers, firmware, WWN and IP addresses.

Figure 9-12 shows the Product List Panel, which can be displayed by selecting **View** → **Show Panels** → **Product List** or pressing **F9**.

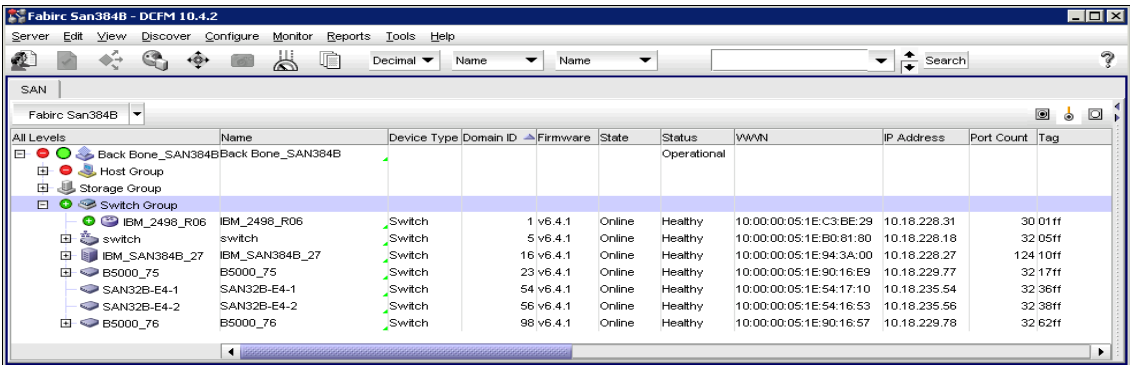


Figure 9-12 Product List Panel \

Status display in the product list

Figure 9-13 shows three types of icons: status, routing, and tracking.

Icon	Meaning
No icon	Operational Status
	Degraded Status
	Failed Status
	Unknown Status
	Routed In
	Routed Out
	Device Removed
	Device Added

Figure 9-13 Status, Routing, and Tracking Icons

- ▶ Status Icons (Degraded Status, Failed Status, Unknown Status):
Shown as tags these indicate the status of the displayed object. They are overlaid over the product icons. In a worse-case status they are rolled up up to closed objects on the Product List.
- ▶ Routing icons (Routed In, Routed Out):
Show participation in Routing)
- ▶ Tracking Icons (Device Removed and Device Added):
Show if an object or a connection has been added or removed

Displaying ports and products in the product list

The structure can be flattened to only display Products or Ports.

You can choose which properties you want to see. Next we show the number of possibilities to change the view in the Product List and Connectivity Panel using the Toolbar.

Figure 9-14 shows the Node WWN of the switch, which can be chosen from the *Product Label* menu located on the toolbar.

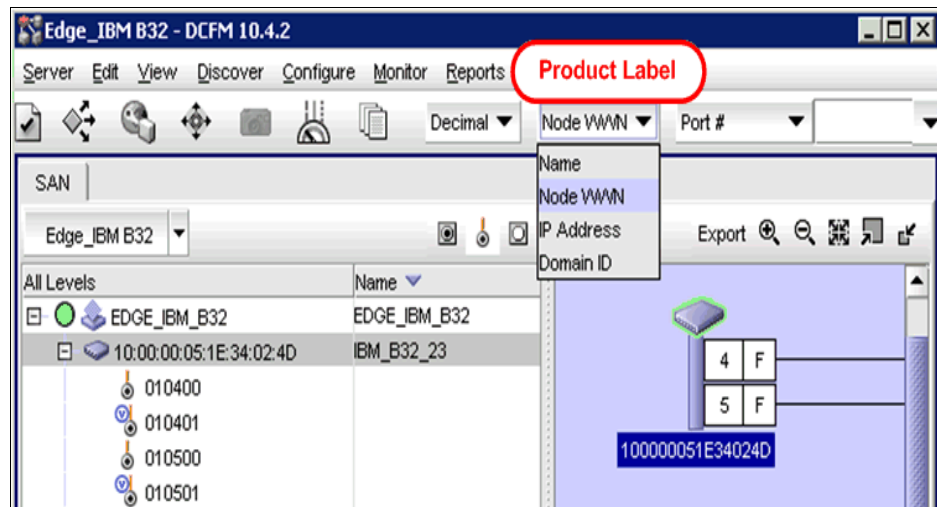


Figure 9-14 Showing options in Product Label menu

The *Port Label* menu on the toolbar allows you to choose how to display the ports (see Figure 9-15).

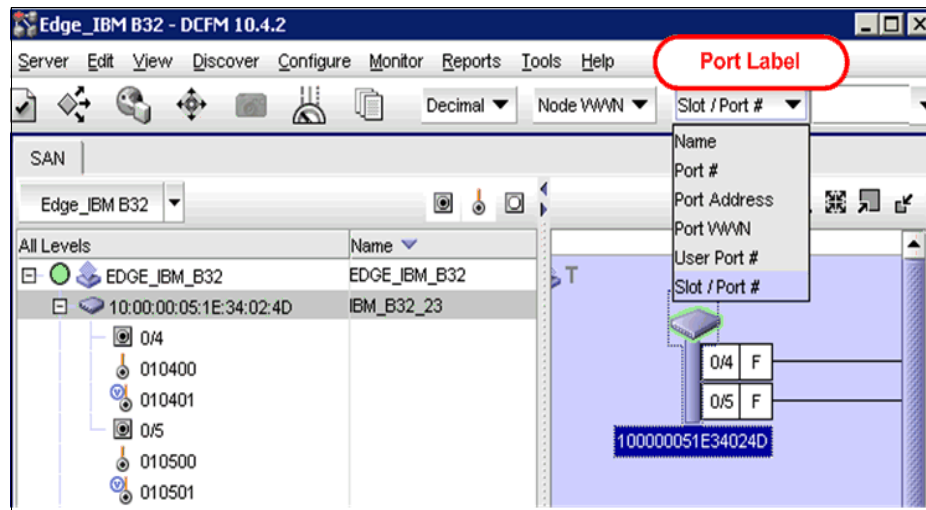


Figure 9-15 Showing options in Port Label menu

When you right-click the switch, a menu with various choices is presented, as shown in Figure 9-16. You can turn the following displays on or off:

- ▶ Occupied Ports
- ▶ Unoccupied Ports
- ▶ Attached Ports
- ▶ Switch to Switch Connections

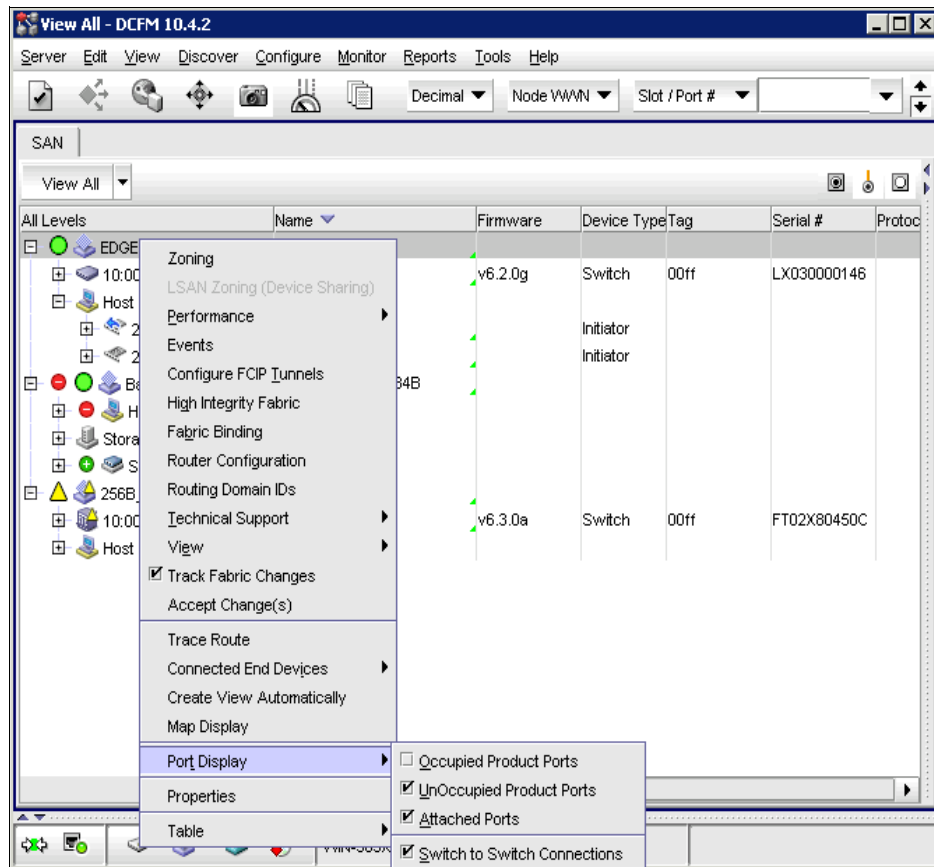


Figure 9-16 Port display choices

9.3.4 Connectivity Map

The Connectivity Map has the following functions:

- ▶ Displays physical and some logical connectivity of SAN products. The map grouping includes Fabrics, Switch, Host, and Storage Groups.
- ▶ Right-clicking in each group will bring up different options.
- ▶ Displays connection and product flyovers containing additional information determined by users

The legacy switch strategy to display a different icon for every model is replaced by two icons ('pizza box' and director).

Sample Connectivity Maps

Figure 9-17 shows the Connectivity Map in which you can see that different switches are configured. When you move the mouse pointer over the switch symbol, you get more information.

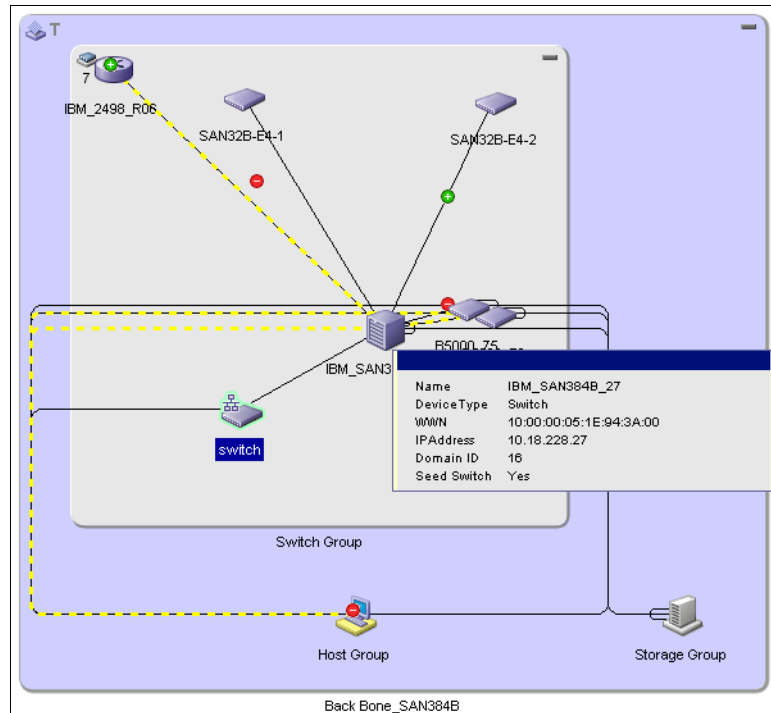


Figure 9-17 Connectivity map

Figure 9-18 shows the Toolbox, which is located at the top right hand side of the View window and provides tools to zoom in and out of the Connectivity Map, collapse, and expand groups, and fit the topology to the window.

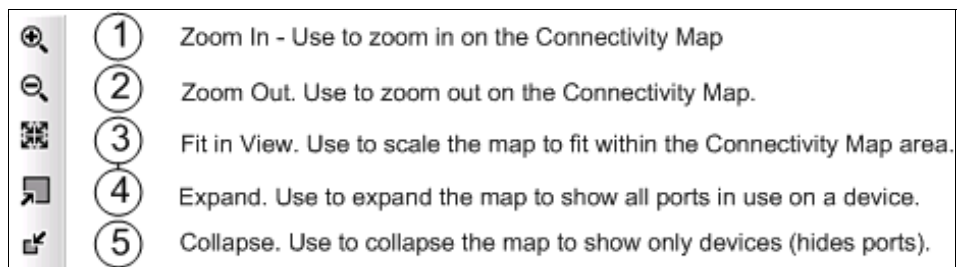


Figure 9-18 Toolbox

You can right-click in Fabrics, Switch, Host, and Storage Groups, and each group will bring up different options.

Fabric right-click menu

The initial DCFM view only displays the switches in a fabric. DCFM will allow you to hide or show all, as well as other zone-based collections of end devices.

When right-clicking the Fabric, you can see the possible choices of showing the devices as in Figure 9-19.

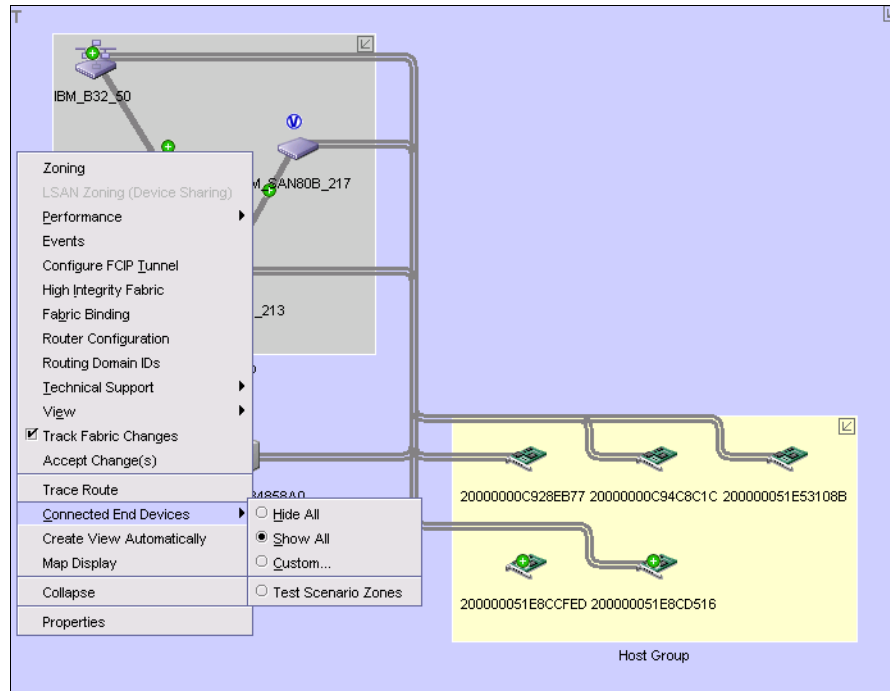


Figure 9-19 Connected End Devices - Show All view

Figure 9-19 shows all the connected devices. But you can choose to display only devices you want to see by clicking **Connected End Devices** → **Custom**.

The Dialog Box, in which you can see the list of currently active zones in the fabric, is displayed (see Figure 9-20).

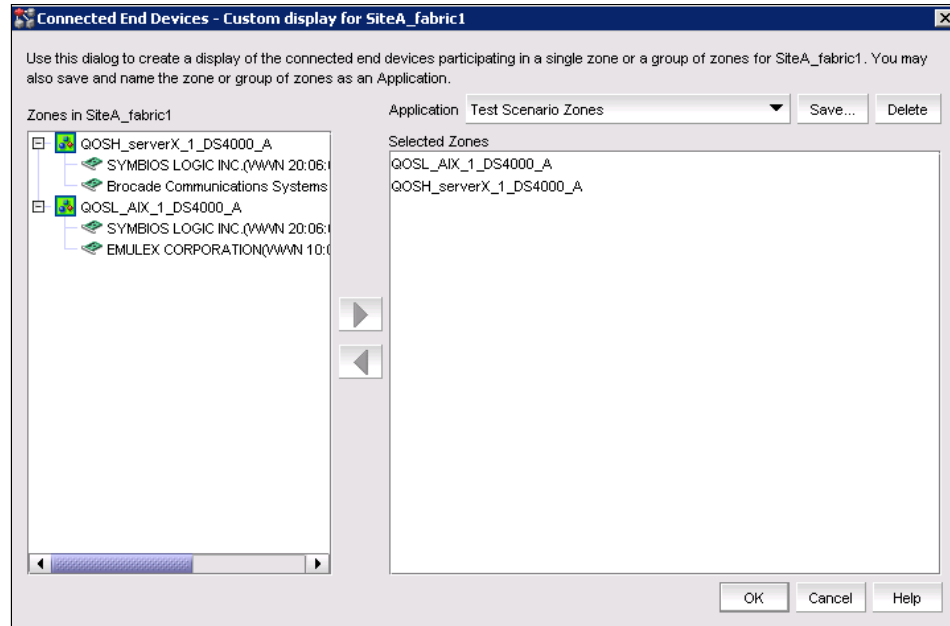


Figure 9-20 Custom Dialog

This allows you to select one or more zones and to name this zone collection (we called ours “Test Scenario”).

In this case, only end devices in the zones will be displayed when the name is selected from the custom menu (see Figure 9-21).

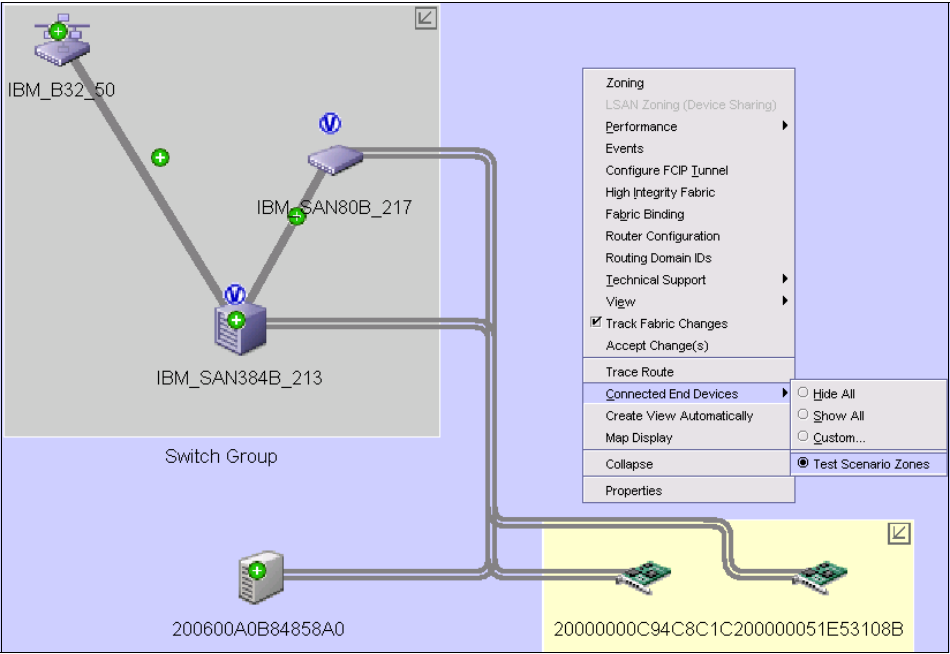


Figure 9-21 Displaying Test Scenario zoned devices Only

You can also choose the way the map is displayed by right-clicking any empty point in the Connectivity Map (see Figure 9-22).

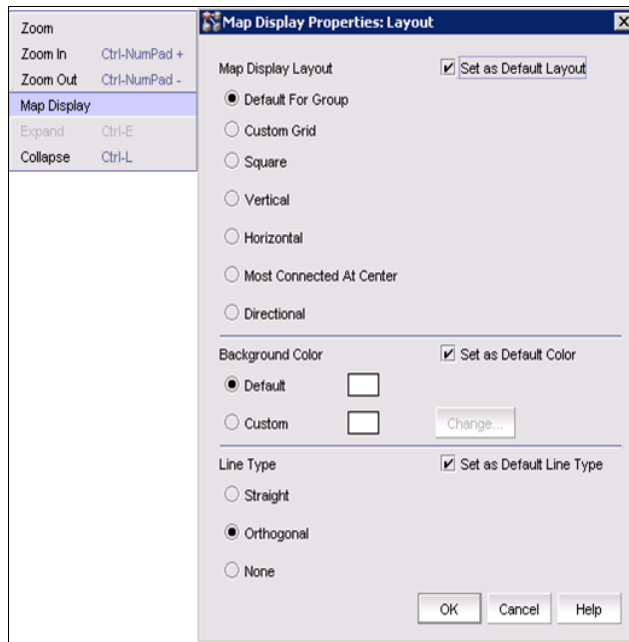


Figure 9-22 Map Display

You can set the Layout, Color, and Line Type displayed for the groups presented on the Map.

9.3.5 Master Log

The Master Log lists the events and alerts that have occurred on the SAN. The log can be filtered to display events by category, severity, or user-entered event content.

You can do any of the following tasks:

- ▶ Determine log size and enable or disable log paging
- ▶ Hide / Show selected events
- ▶ Display event details

In the first column, the severity of an event (Level) is displayed.

The tags describing severity are shown in Figure 9-23.




Event Icon	Description
	Informational
	Warning
	Error

Figure 9-23 Severity of an event

9.3.6 Performance Legend

Performance Legend displays ranges and colors used in the Connectivity Map display (see Figure 9-24). This legend is only available when you enable the View utilization option under **Monitor** → **Performance** → **View Utilization**.

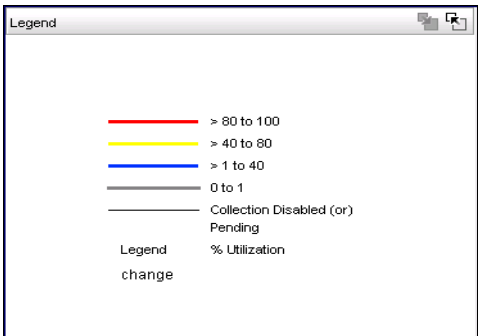


Figure 9-24 Performance Legend

When you select the **Change** button, the next display allows users to set ranges and colors that are going to be used in the Connectivity map (see Figure 9-25).

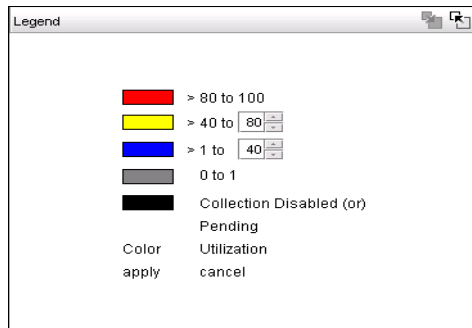


Figure 9-25 Changing range and colors in connectivity maps

9.3.7 Minimap

The Minimap, which displays in the lower right corner of the main window, is useful for getting a bird's-eye view of the SAN, or to quickly jump to a specific place on the Connectivity Map. To jump to a specific location on the Connectivity Map, click that area on the Minimap. A close-up view of the selected location displays on the Connectivity Map.

Use the Minimap (see Figure 9-26) to view the entire SAN and to navigate to more detailed map views. This feature is especially useful if you have a large SAN because it:

- ▶ Displays entire Connectivity Map
- ▶ Displays current Connectivity Map view (outlined in green)
- ▶ Allows users to drag or click to reposition the current Connectivity Map view




Figure 9-26 Minimap

9.3.8 Status bar

The status bar displays at the bottom of the main window. The status bar provides a variety of information about the SAN and the application. The icons on the status bar change to reflect different information, such as the current status of products, fabrics, and backup.

Figure 9-27 shows the status bar with descriptions of the icons.

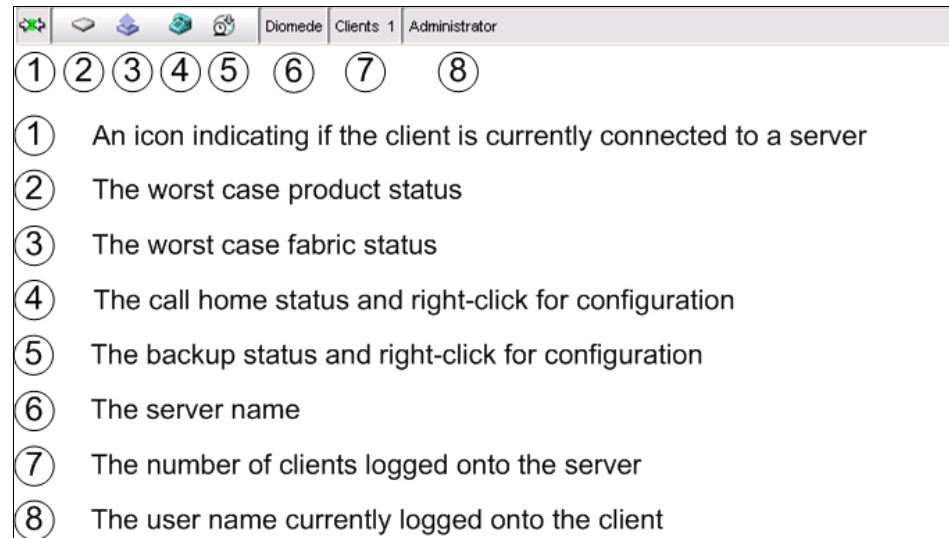


Figure 9-27 Status Bar

The Product status icons (number 2 in Figure 9-27 on page 352) are shown here in Figure 9-28.









No icon	Healthy/Operational
	Attention
	Degraded/Marginal
	Device Added
	Device Removed/Missing
	Down/Failed
	Routed In
	Routed Out
	Unknown/Link Down

Figure 9-28 Product Status Icons

9.3.9 Fabric tracking

Fabric Tracking indicates if any product or connection in the fabric has been added or removed.

A green plus-sign icon is displayed with products or connections that are added, and a red minus-sign icon is displayed with ones that are removed.

Fabric Tracking can be accessed by right-clicking a fabric in the Connectivity Map. Figure 9-29 shows Fabric Tracking.

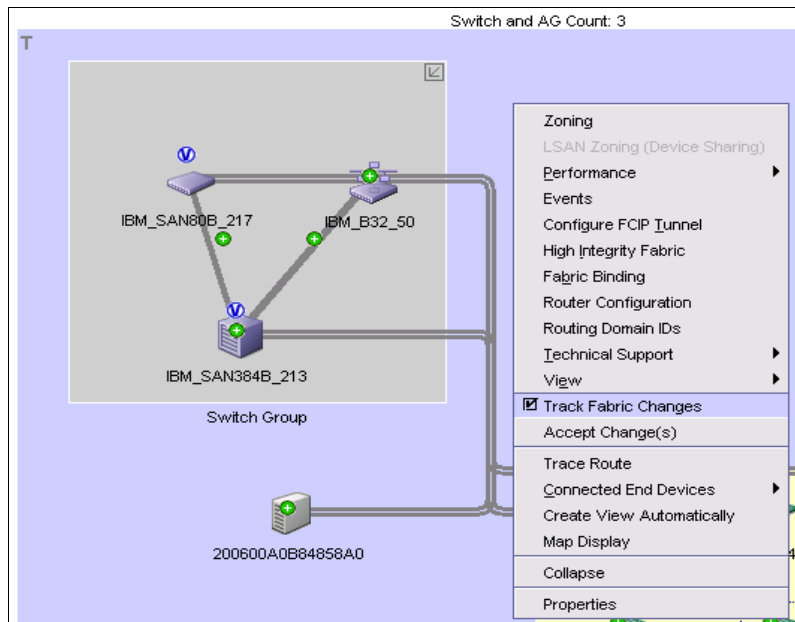


Figure 9-29 Fabric Tracking

If you click the option **Track Fabric Changes** to disable Fabric Tracking (as shown in Figure 9-29) you will switch off tracking for all devices, switches, and connections in this fabric.

Fabric Tracking: The default for Fabric Tracking for new fabrics is *Enabled*.

If there were changes, a DCFM Message box as shown in Figure 9-30 comes up. Here you have to confirm to reset the status for all devices, switches and connections. If you do **Accept Change(s)** instead of disabling the tracking, the same DCFM Message box comes up. Here you reset the tracking to a new baseline.

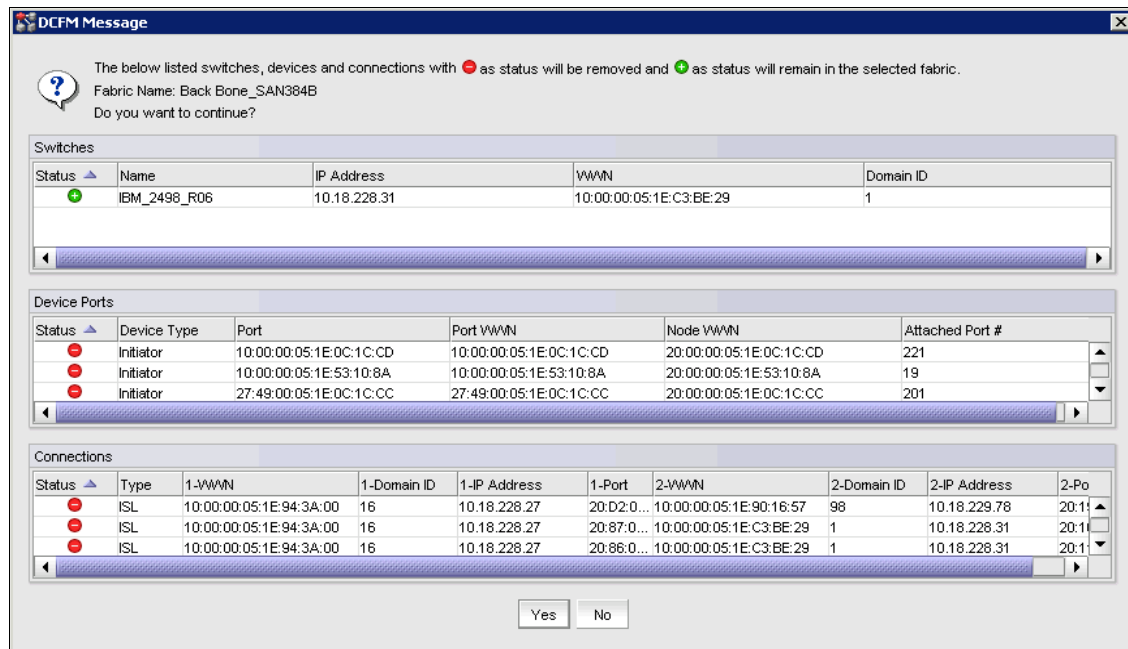


Figure 9-30 Track Fabric changes dialog box

9.3.10 WWN display

DCFm will display WWNs with colons and accept WWN input from users with or without the colons.

Exception: The spacing of the connectivity map cannot support the display of WWNs with colons without reducing the density of objects on the map.

WWNs will always be displayed with upper-case A, B, C, D, E, and F characters.

9.3.11 Object naming

Object Naming can be accessed by right-clicking a particular device and selecting **Properties** (see Figure 9-31).

Object Names: User defined Object Names are stored on managed devices as well as in the DCFM database.

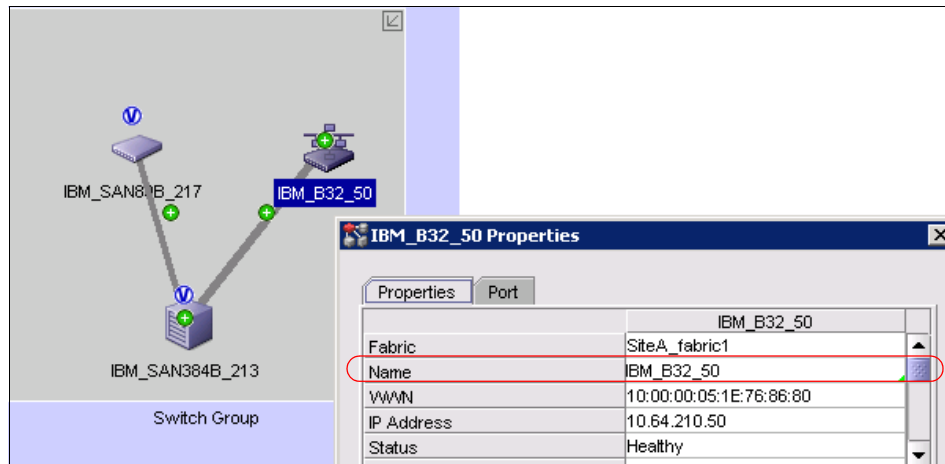


Figure 9-31 Setting Object Name

9.4 DCFM Fabric Discovery

Discovery is the process by which the management application contacts the devices in your SAN. When you configure discovery, the application discovers products connected to the SAN and illustrates each product and its connections on the Connectivity Map.

9.4.1 Seed switch

To run the discovery, DCFM needs the *seed switch*, which is the switch that will discover and populate DCFM using in-band communication with remaining fabric such as these:

- ▶ Name Server information
- ▶ Zoning information
- ▶ Fabric membership information

Firmware requirements for the seed switch are as follows:

- ▶ Pure FOS fabrics: FOS switch running 5.0 or later firmware.
- ▶ Mixed fabrics: FOS switch running 6.0 or later and M-EOS and M-EOSn 9.7 or later.
- ▶ Pure M-EOS fabrics: M-EOS and M-EOSn 9.9.2 or later
- ▶ DCFM expects the FOS Discovery switch to run the highest level of firmware in the fabric.

When you discover a fabric, the management application checks to confirm that the seed switch is running the latest Fabric OS version in the fabric, and if it is not, the management application prompts you to select a new seed switch.

For Fabric Operating System, seed switches depend on the size of the SAN:

- ▶ For Small Fabrics:
Use at least an entry-level switch (SAN24B-4).
- ▶ For Medium Fabrics:
Use at least a SAN40B-4 (Backbone switches or SAN256B are best).
- ▶ For Large Fabrics:
Use at least a SAN40B-4 (Backbone switches or SAN256B are best).

You have to have Fabric Operating System *admin* privilege (or equivalent accounts: *root*, *admin*, or *factory*).

Seed switch: Access Gateway-mode switches, switches connected to the fabric through EX/ VEX port types, and IBM Converged Switch B32 *cannot* be used as a seed switch for discovery.

You can change the seed switch as long as the new one follows the rules and is:

- ▶ HTTP-reachable from the management application
- ▶ Running the latest Fabric OS version in the fabric
- ▶ A primary Fabric Configuration Server (FCS)

Privilege: Discovery needs FOS *admin* privilege (or equivalent accounts: *root*, *admin*, or *factory*).

9.4.2 Setting up the discovery

The Discovery window can be reached by using **Discover** → **Setup**.

Figure 9-32 shows how to open the Discovery window from the Main Menu Bar.




Figure 9-32 Accessing Discovery Window

The Discovery window displays, as shown in Figure 9-33.

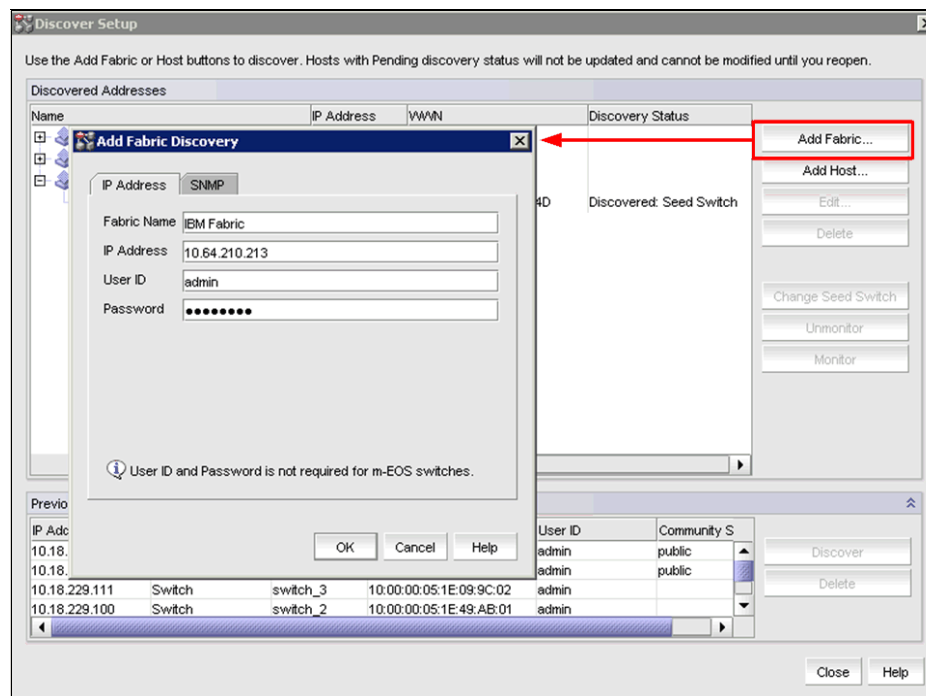


Figure 9-33 Discovery Setup Window

Select the **Add Fabric** button from the Discovery Setup window to input the seed switch information. You can also change the setting for SNMP (see Figure 9-34).

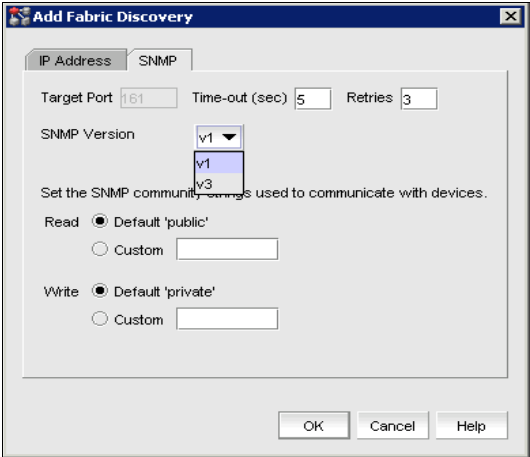


Figure 9-34 Discovery Setup Dialog for SNMP

Discovery status

You can determine the discovery status of products by looking at the Status column in the Product List, and also the operational status (Figure 9-35). Be aware of the following considerations:

- ▶ *Unknown* is equivalent to the discovery state *Offline*.
- ▶ *Healthy/Operational*, *Degraded/Marginal*, and *Down/Failed*, are equivalent to a discovery state of *Online*.

All Levels	Contact	Description	State	Status	User Colu...
SiteA_fabric1				Operational	
Host Group					
Switch Group					
IBM_B32_50	Field Support.	Fibre Chann...	Online	Healthy	
E Port Trunk 3					
20:00:00:05:1E:76:86:80					
20:07:00:A0:B8:48:58:A2					
20:08:00:05:1E:76:86:80					
10:00:00:05:1E:8C:D5:16					
10:00:00:05:1E:8C:CF:ED					
20:09:00:05:1E:76:86:80					
10:00:00:05:1E:8C:D5:16					
IBM_SAN384B_213	Field Support.	Fibre Chann...	Online	Healthy	
E Port Trunk 224					
E Port Trunk 221					
20:54:00:05:1E:94:3A:00			Online	Online	
20:39:00:05:1E:09:97:01			Online	Online	

Figure 9-35 Discovery Status in Product List window

Changing the discovery switch

It is possible to change the seed switch in the DCFM Discovery setup window. In case there is a new switch that has the desired Fabric Configuration Server (FCS) policy, you can do it from here also. The management application requires that the seed switch is the primary FCS switch at the time of discovery.

Monitoring fabrics

You can activate or suspend discovery on a fabric without losing any fabric information. By activating discovery on the fabric, the fabric will go to the state *Discovered Monitored*. By suspending discovery on the fabric, we mean that the fabric will go to the *Unmonitored Discovered* state as shown in Figure 9-36.

When two DCFM monitored fabrics merge, the following events will happen:

- ▶ Two representations of the single new fabric will exist.
- ▶ The fabric that was first discovered will remain active and the other fabric will automatically be placed in the Unmonitored state.
- ▶ You can then delete the Unmonitored fabric or swap the monitoring of the two representations.

In Figure 9-36, rounded rectangles point out the same switch IBM_SAN384_213 after fabric merge in two places. A second representation of the switch IBM_SAN384_213 is in the Unmonitored state.

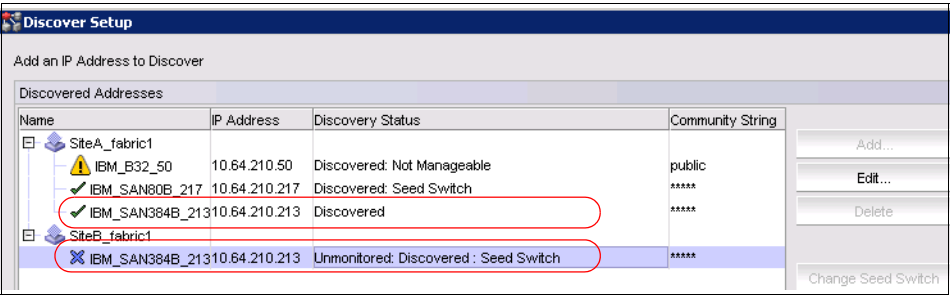


Figure 9-36 Unmonitored Fabric after Fabric merge

The management application enables you to view the fabric monitoring status through the Discover Setup dialog box. Figure 9-37 illustrates and describes the icons that indicate the current status of the discovered fabrics.




	Displays when the fabric is managed and the switch management status is okay.
	Displays when the fabric is managed and the switch management status is not okay.
	Displays when the fabric is not managed.

Figure 9-37 Icons indicating the current status of the discovered fabrics

Discovering missed switches

If a fabric has been discovered and if some of its switches segment into single or multiple new fabrics, or if the fabric was deleted, you can now easily re-discover those new fabrics without entering their credentials (see the rounded rectangle in Figure 9-38 in *Previously Discovered Addresses*).

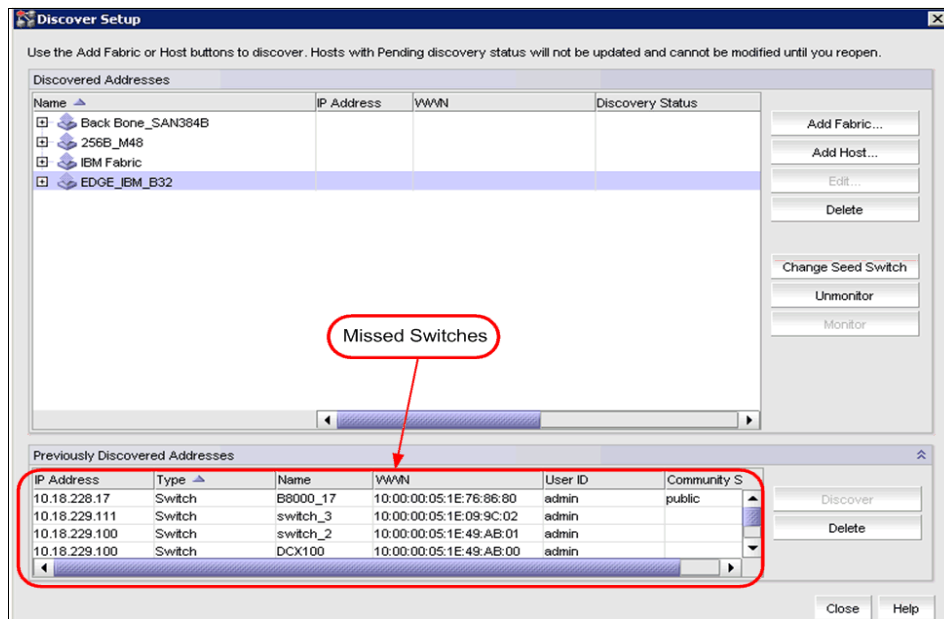


Figure 9-38 Missed Switches Rediscovery

Deleting a fabric from the management application

If you decide that you no longer want the management application to discover and monitor a specific fabric, you can delete it. Deleting a fabric also deletes the fabric data on the server except for user-assigned names for the device port, device node, and device enclosure information (see Figure 9-39).

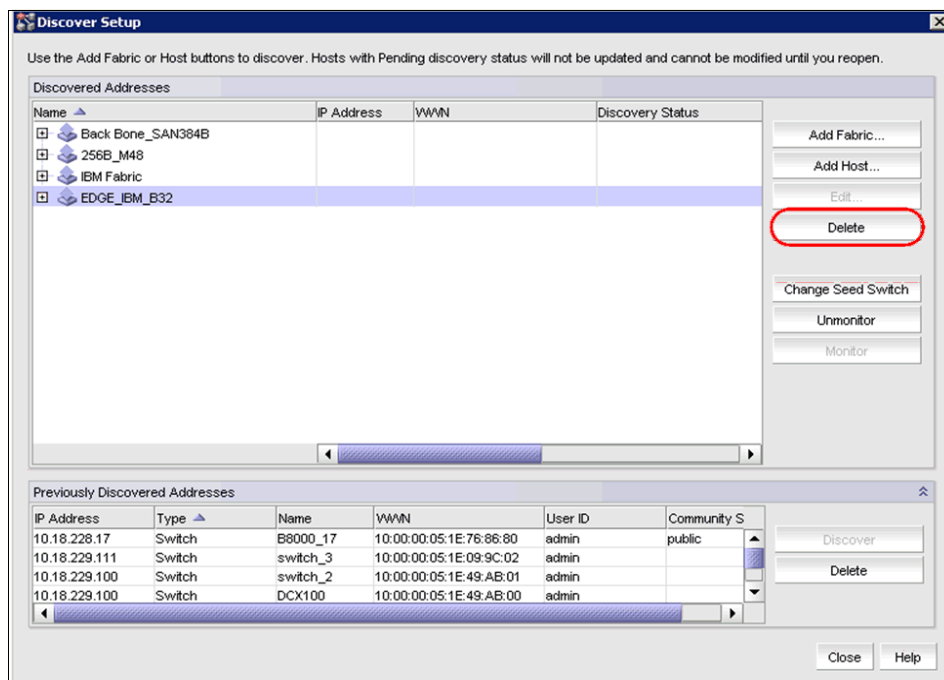


Figure 9-39 Deleting the fabric from Management Application

9.4.3 DCFM Discovery Verification

There is a limitation of monitored fabrics in DCFM Enterprise Edition.

For pure FOS fabrics:

- ▶ You can monitor up to 24 physical fabrics with support for:
120 switches, 9,000 switch ports, 20,000 hosts or storage devices

For mixed fabrics:

- ▶ You can monitor up to 24 physical fabrics with support for:
60 switches, 5,000 switch ports, 10,000 hosts or storage devices

For pure M-EOS fabrics:

- ▶ You can monitor up to 24 physical fabrics with support for:
60 switches, 5,000 switch ports, 10,000 hosts or storage devices

You can use the **Unmonitor** button to unmonitor fabrics from the Management Application and select a new Fabric to monitor as shown in Figure 9-40.

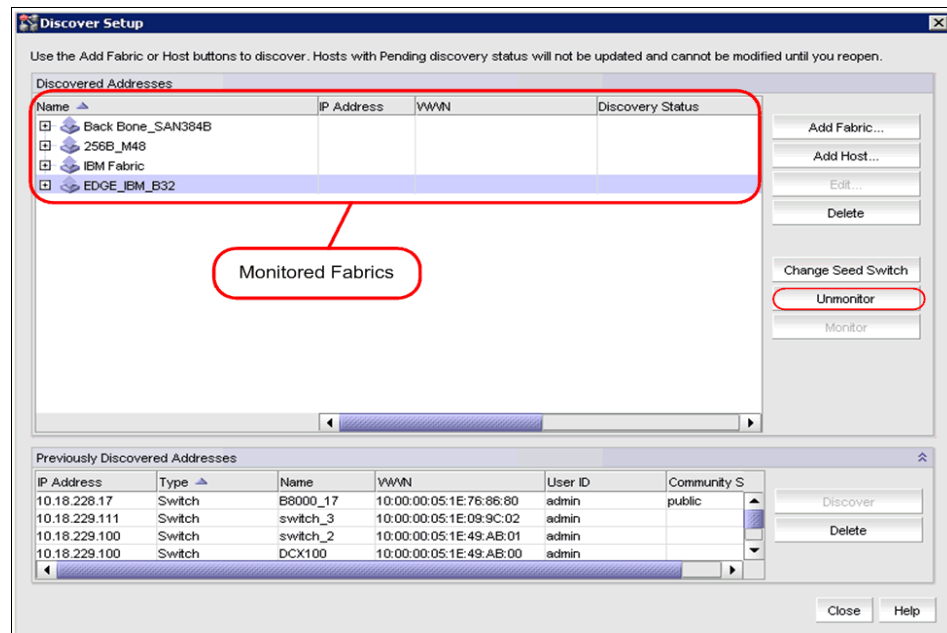


Figure 9-40 Unmonitor fabrics

9.5 DCFM reports

Presenting and archiving data about a SAN is equally as important as gathering the data.

The following standard report types are available to generate from the collected data out of the DCFM:

- ▶ Fabric Summary: Lists information about discovered fabrics including fabric and switch details, device information, and ISL and trunk summary.
- ▶ Fabric Ports: Lists discovered ports including used and unused ports. Port data for each fabric is divided into three parts: Fabric-wide port details, Switch-wide port details, and individual port details.

The following device specific reports are available through the Report menu:

- ▶ Performance: Lists historical performance-related data.
- ▶ Zone: Lists zoning objects.

Viewing reports: Reports can be viewed in the reports dialog, which supports the export of reports to various formats (html, pdf, xml).

9.5.1 Fabric Summary Report and Port Report

The reports can be generated from Main Fabric View. Go to **Report** → **Generate** (see Figure 9-41).




Figure 9-41 Reports menu

You can choose the **Fabric Summary Report** as shown in Figure 9-42.

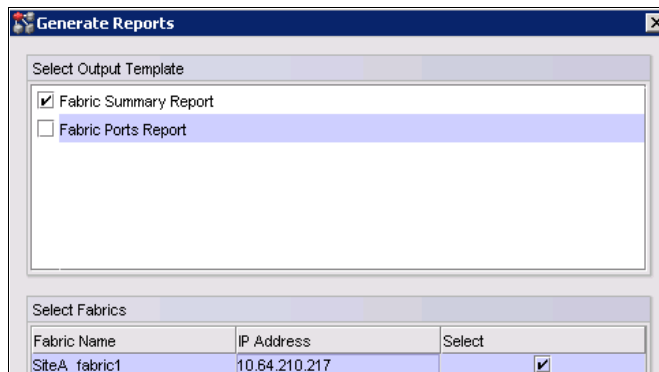


Figure 9-42 Dialog box for generating reports

Finally, the fabric report displays (see Figure 9-43).

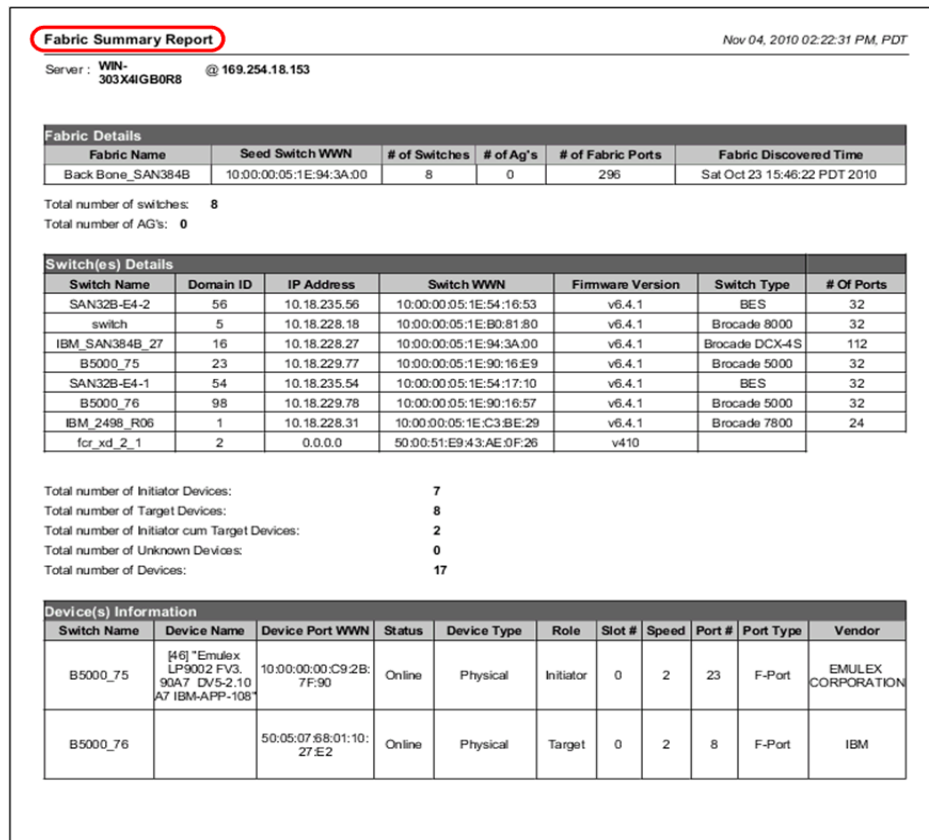


Figure 9-43 Fabric Summary Report

The reports contain the following information:

- Details regarding the Fabrics, Switches, and Devices discovered in the SAN:
 - Fabric Detail:

Fabric Name, WWN, Is Secure, Fabric Creation Time, Total Number of Switches, Total Number of AGs, Total number of FC Ports, AD Information
 - Switch Details:

Switch Name, Domain ID, IP Address, Switch WWN, Firmware Version, Switch Type, Total Number of Ports

- Device Information:
 - Summary: Total Number of Devices, Total Number of Initiators, Total Number of Targets, Total Number of Initiators + Targets
 - Device Name, Device Port WWN, Status, Physical/Virtual, Role (Initiator/Target), Switch Name, Slot Number, Switch Port Number, Port Type, Device Type, Vendor

The Fabric Ports Report, which can be seen in Figure 9-44, displays information about all the discovered ports in all discovered fabrics.

Fabric Ports Report

Nov 04, 2010 02:27:29 PM, PDT

Server : WIN-303X4IGB0R8 @ 169.254.18.153

Fabric : Back Bone_SAN384B

Director Utilization					Switch Utilization			
Total Fabric Ports	Total Number of Ports	Number of Ports connected	Number of Ports Free	Number of Ports allocated	Total Number of Ports	Number of Ports connected	Number of Ports Free	Number of Ports allocated
313	0	0	0	0	313	66	248	0

Details

IP Address	Switch Name	Domain/ Port #	Zone	Connected Device					Port Name	Port Speed (GBPS)	Port Status	Port State	Port Type	Physical/ Logical
				DeviceName	Vendor	Device Type	Model	PortWwn						
10.18.235.56	SAN32B-E4-2	D56 /P0	(Online)							8	No_Mod ule	Offline	U-Port	Physical
10.18.235.56	SAN32B-E4-2	D56 /P1	(Online)							8	No_Mod ule	Offline	U-Port	Physical
10.18.235.56	SAN32B-E4-2	D56 /P2	(Online)							8	No_Light	Offline	U-Port	Physical
10.18.235.56	SAN32B-E4-2	D56 /P3	(Online)							8	No_Mod ule	Offline	U-Port	Physical
10.18.235.56	SAN32B-E4-2	D56 /P4	(Online)							8	No_Mod ule	Offline	U-Port	Physical
10.18.235.56	SAN32B-E4-2	D56 /P5	(Online)							8	No_Mod ule	Offline	U-Port	Physical
10.18.235.56	SAN32B-E4-2	D56 /P6	(Online)							8	No_Mod ule	Offline	U-Port	Physical
10.18.235.56	SAN32B-E4-2	D56 /P7	(Online)							8	No_Mod ule	Offline	U-Port	Physical

Figure 9-44 Fabric Ports Report

9.5.2 Generating performance reports

In order to generate a historical performance report, you have to enable that you want to collect data constantly to receive the necessary historical data required for a meaningful report. To enable, select **Monitor** → **Performance** → **Historical Data Collection** → **Enable SAN Wide**.

To generate a report, select: **Monitor** → **Performance** → **Historical Graph**

The generated report can be seen in Figure 9-45.

Historical Performance Report

Server: (Diomedee)

@ IP Address: (10.64.210.106)

Report Configuration

Favorite Name

MyData_20090803

Main Measure

Tx % Utilization

Display

Top 5 of Tx % Utilization

From

All FC Ports

For

Last 1 Hour

Granularity

5 Minutes

Additional Measures

Rx % Utilization

Top 5 of All FC Ports by Tx % Utilization

#	Fabric	Source	Source Port	Port Type	Destination	Destination Port	Tx % Utilization	Rx % Utilization
1	SiteA_fabric1	IBM_SAN80B_217	9	E-Port	IBM_SAN384B_213	8/29	0	0
2	SiteA_fabric1	IBM_SAN80B_217	57	E-Port	IBM_SAN384B_213	2/20	0	0
3	SiteA_fabric1	IBM_SAN80B_217	19	F-Port	20:00:00:00:C9:4C:8C:1C	10:00:00:00:C9:4C:8C:1C, 10:00:00:00:C9:4C:8C:1C	0	0
4	SiteA_fabric1	IBM_SAN80B_217	56	E-Port	IBM_SAN384B_213	2/21	0	0
5	SiteA_fabric1	IBM_SAN80B_217	1	U-Port			0	0

Figure 9-45 Historical Performance Report

9.5.3 Generating zoning reports

You can also generate a report for the current zone DB in the fabric.

- ▶ To generate a report for the edited zone DB, you must save it to the fabric first.
- ▶ To generate a zoning report, select **Configure** → **Zoning** → **Fabric**.
- ▶ You can also right-click the device and select **Zoning**.

The Zoning dialog box displays, as shown in Figure 9-46.

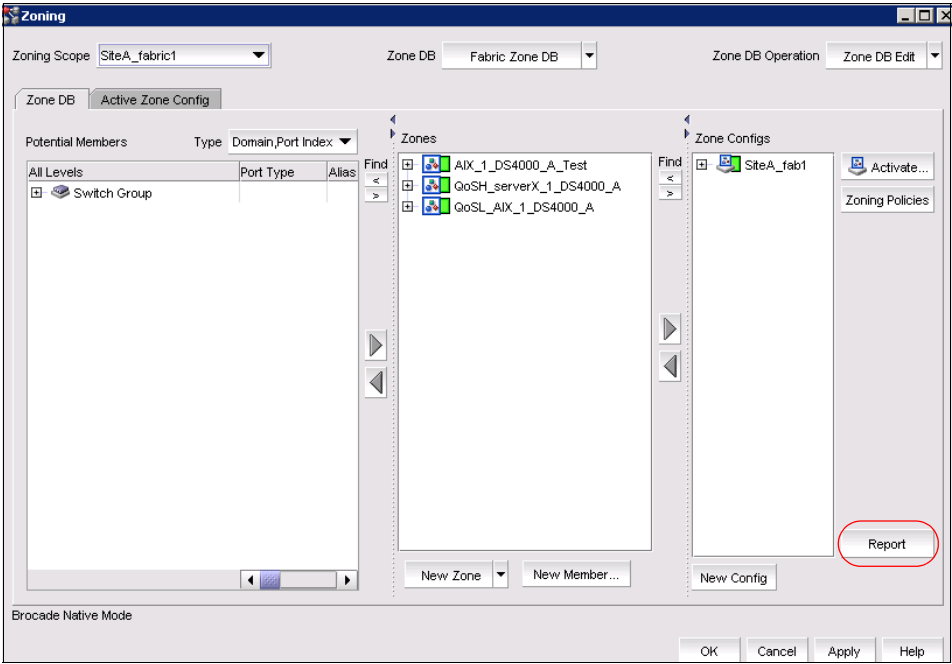


Figure 9-46 Zoning window

Click **Report** and a DCFM message is issued. Confirm with **OK** and the selected report automatically displays in the View Reports dialog box (see Figure 9-47).

Zoning Summary Report

Jun 25, 2009 01:37:10 PM, PDT

Fabric: SiteA_fabric1

Server: Diomeda@10.64.210.106

Zone DB: (Online)

Size: 626 bytes

Zone DB\Zone Configs

SiteA_fab1

Zone Name	Number of Members	Members Logged In	Active
AIX_1_DS4000_A_Test	2	2	Yes
QoSH_serverX_1_DS4000_A	2	1	Yes
QoSL_AIX_1_DS4000_A	2	2	Yes

Zone DB\Zones

AIX_1_DS4000_A_Test (member of 1 Zone Config)

Alias Members - DS4000_A

WWN Members

Member Properties				Attached to Switch Properties			
Logged In	Port Name	Port WWN	Node Name	Node WWN	Switch Name	Domain ID	Port Index
Yes		20:06:00:A0:B8:48:58:A1		20:06:00:A0:B8:48:58:A0	IBM_SAN384B_213	2	92

Alias Members - AIX_1

WWN Members

Member Properties				Attached to Switch Properties			
Logged In	Port Name	Port WWN	Node Name	Node WWN	Switch Name	Domain ID	Port Index
Yes		10:00:00:00:C9:4C:8C:1C		20:00:00:00:C9:4C:8C:1C	IBM_SAN80B_217	1	19

QoSH_serverX_1_DS4000_A (member of 1 Zone Config)

WWN Members

Member Properties				Attached to Switch Properties			
Logged In	Port Name	Port WWN	Node Name	Node WWN	Switch Name	Domain ID	Port Index
Yes		20:06:00:A0:B8:48:58:A1		20:06:00:A0:B8:48:58:A0	IBM_SAN384B_213	2	92
No		10:00:00:05:1E:53:10:8B		20:00:00:05:1E:53:10:8B	IBM_SAN80B_217	1	63

Figure 9-47 Zoning reports

9.6 Event logs

You can view all events that take place in the SAN through the Master Log at the bottom of the main window. You can also view a specific log by selecting an option from the **Monitor** → **Logs** menu option (see Figure 9-48).

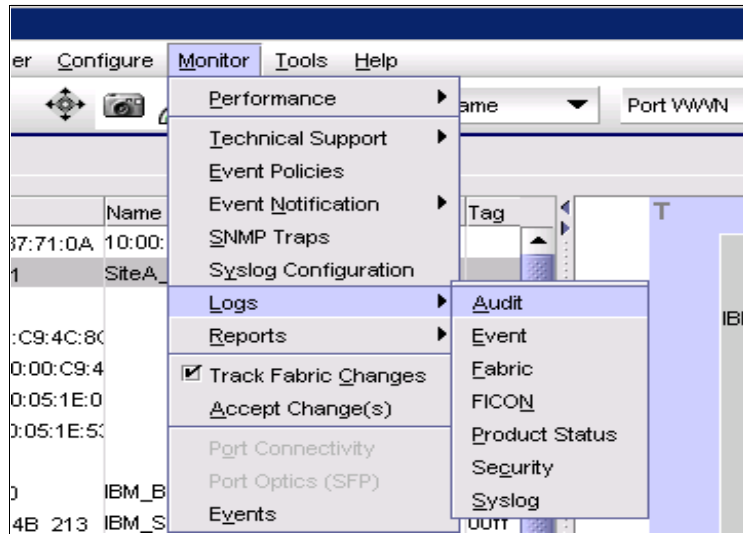


Figure 9-48 Viewing Event Logs

As shown in Figure 9-48, you have the option to filter the logs and only display the area of interest. The following options for displaying logs are available:

- ▶ Audit
- ▶ Event
- ▶ Fabric
- ▶ FICON
- ▶ Product Status
- ▶ Security
- ▶ Syslog

Table 9-3 shows the logs and their descriptions.

Table 9-3 Logs

Name	Description
Event Log	Displays all “Product Events” type events from all discovered switches
Fabric Log	Displays “Product Events,” “Device Status,” and “Product Audit” type events for all discovered fabrics
FICON Log	Displays all the “LIR” and “RLIR” type events, for example, “link incident” type events
Product Status Log	Displays events which indicate a change in Switch Status for all discovered switches

Name	Description
Security Log	Displays all security events for the discovered switches
Syslog Log	Displays syslog messages from switches

The master log contains the columns as shown in Table 9-4.

Table 9-4 Master Log Columns

Column name	Description
Level	The severity of the event
Source Name.	The product on which the event occurred
Source Address.	The IP address (IPv4 or IPv6 format) of the product on which the event occurred
Type	The type of event that occurred (for example, client/server communication events)
Description	A description of the event
First Occurrence Host Time	The time and date the event first occurred on the host
Last Occurrence Host Time	The time and date the event last occurred on the host
First Occurrence Switch Time	The time and date the event first occurred on the switch
Last Occurrence Switch Time	The time and date the event last occurred on the switch
Operational Status	The operational status of the product on which the event occurred
Count	The number of times the event occurred
Module Name	The name of the module on which the event occurred
Message ID	The message ID of the event
Contributor	The name of the contributor on which the event occurred
Node WWN	The world wide name of the node on which the event occurred
Fabric Name	The name of the fabric on which the event occurred

Figure 9-49 shows how it is possible to hide events by selecting the desired events and choosing the **Hide Selection** option when you right-click in the Master Log window.

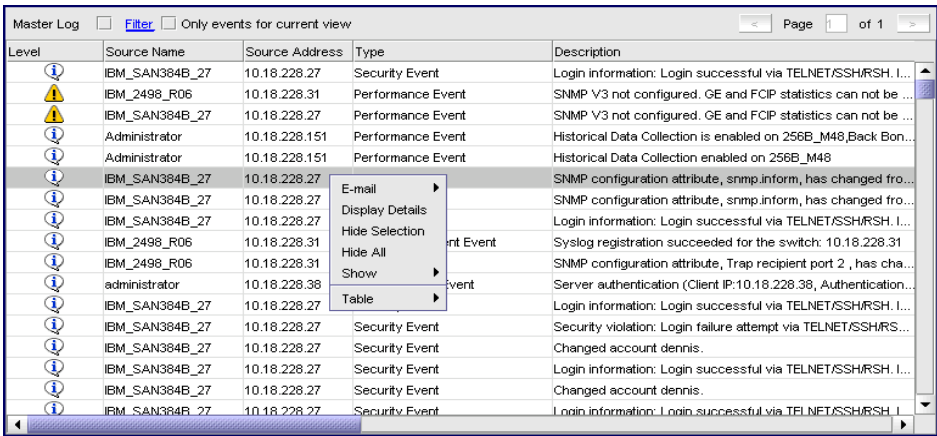


Figure 9-49 Hiding highlighted events

You can display the details for a particular event as shown in Figure 9-50 by selecting the event, right clicking it, then selecting **Display Details**.

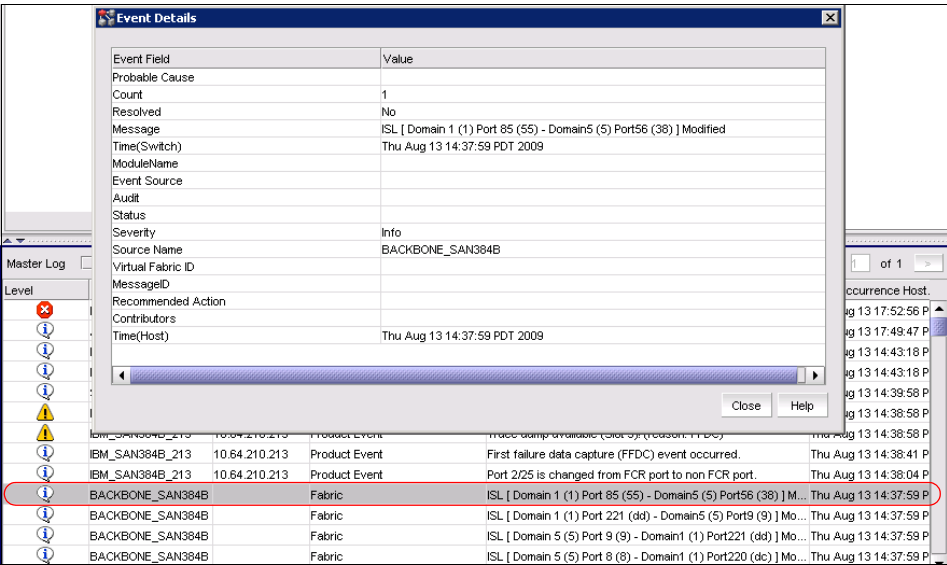


Figure 9-50 Event Details

Figure 9-51 displays how to define a filter for the events. To do this, select the option **Filter** on the top of the window Master Log.

- ▶ To include an event type in the filter, select the event from the Available Events table and click the right arrow.
- ▶ To exclude an event type from the filter, select the event from the Selected Events table and click the left arrow.

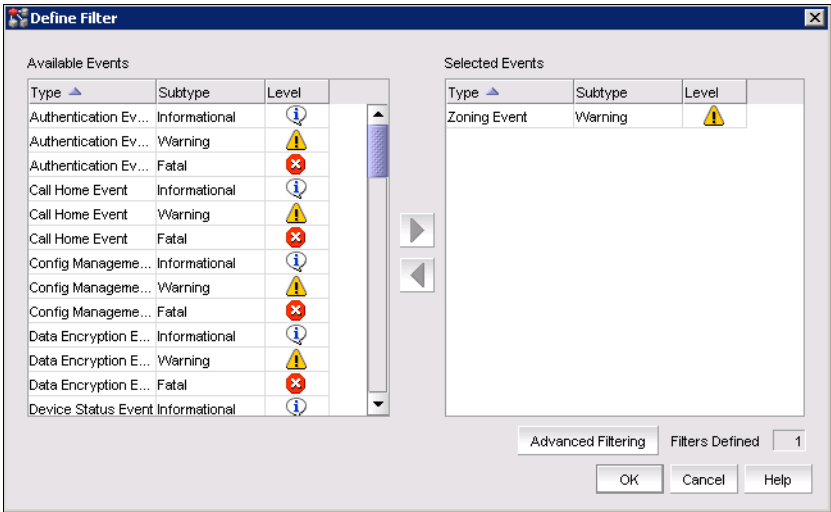


Figure 9-51 Defining a filter for the event

Figure 9-52 shows the results of defining a filter for only Zoning Events in the Master Log.

Master Log <input checked="" type="checkbox"/> Filter <input type="checkbox"/> Only events for current view						
Level	Source Name	Source Address	Type	Description	First Occurrence Host...	
Warning	IBM_SAN80B_217	10.64.210.217	Zoning	Failed to delete zone DB content from Fabric SiteA_fab1.	Mon Jun 15 04:50:32 PDT...	
Warning	IBM_SAN80B_217	10.64.210.217	Zoning	Failed to activate zone config SiteA_fab1 and save the zon...	Mon Jun 15 00:53:08 PDT...	
Warning	IBM_SAN80B_217	10.64.210.217	Zoning	Failed to activate zone config SiteA_fab1 and save the zon...	Mon Jun 15 00:44:27 PDT...	
Warning	IBM_SAN80B_217	10.64.210.217	Zoning	Failed to activate zone config SiteA_fab1 and save the zon...	Mon Jun 15 00:42:41 PDT...	

Figure 9-52 Master Log with the filter (zoning events only)

9.7 Performance management

In this DCFM performance management section, we concentrate only on the performance management features in DCFM, explaining how you can become familiar with their use. However, we do not describe the theory behind it and also do not provide real life examples.

For more information about Performance monitoring, see Chapter 16, “Performance monitoring” on page 739.

Performance monitoring provides details about how much traffic and errors a specific port or switch generates on the fabric over a specific time frame. You can also use performance to indicate the switches that create the most traffic and to identify the ports that are most congested.

You can monitor the performance of your SAN using the following methods:

- ▶ Gather and display real time performance data, monitoring, and graphical display (FC ports, GigE, and FCIP).
- ▶ Persist and display historical performance data (FC and GigE (GE) ports as well as FCIP tunnels) for selected fabrics or the entire SAN.
- ▶ Support End-to-End monitors for real time and historical performance data.
- ▶ Enforce user-defined performance thresholds and notification when thresholds are exceeded.
- ▶ Display percentage utilization on the client for FC and FCIP links.
- ▶ Provide user-defined aging scheme (5 minutes, 30 minutes, 2 hours and 1 day granularity).
- ▶ Provide enhanced performance reports.

In Figure 9-53 you can see the menu with possible options that you can use when managing performance data with DCFM.

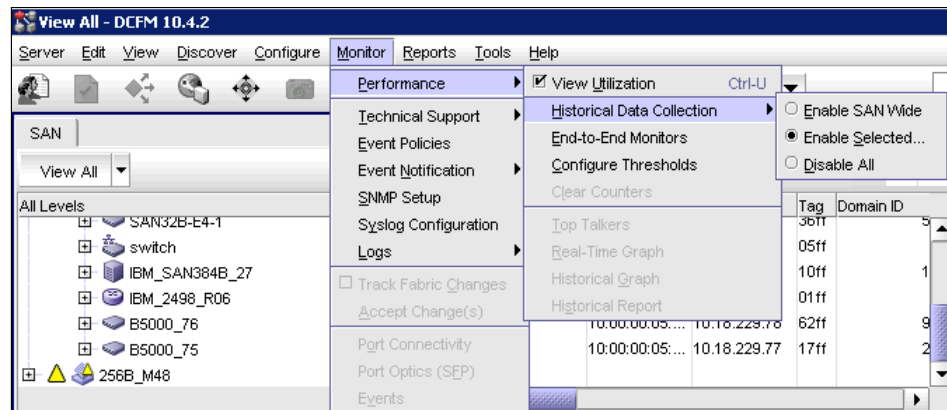


Figure 9-53 Pull-down menu

9.7.1 Performance measures

There are several performance measures available to you, depending on the object type from which you want to gather performance data (see Table 9-5).

Table 9-5 Performance measures

Parameter	Description
Tx% Utilization	Available for FC and GE ports, FCIP tunnels, and End-to-End monitors.
Rx% Utilization	Available for FC and GE ports, FCIP tunnels, and End-to-End monitors.
Tx MB/Sec	Available for FC and GE ports, FCIP tunnels, and End-to-End monitors.
Rx MB/Sec	Available for FC and GE ports, FCIP tunnels, and End-to-End monitors.
CRC Errors	Available for FC ports and End-to-End monitors
Signal Losses	Available for FC ports only.
Sync Losses	Available for FC ports only.
Link Failures	Available for FC ports only
Sequence Errors	Available for FC ports only
Invalid Transmissions	Available for FC ports only
Rx Link Resets	Available for FC ports only
Tx Link Resets	Available for FC ports only
Dropped Packets	Available for FCIP tunnels only.
Compression Ratio	Available for FCIP tunnels only.
Latency	Available for FCIP tunnels only
Link Retransmits	Available for FCIP tunnels only

9.7.2 Collecting performance data

Data collected through Advanced Performance Monitoring is deleted when the switch is rebooted. Using the Data Center Fabric Manager (DCFM) Enterprise Edition, you can store performance data persistently.

9.7.3 Real time performance data

Real time performance enables you to collect data from managed switches in your SAN, and is only supported on the following managed objects:

- ▶ FC (E_ and F_ports) and
- ▶ GE_ports as well as FCIP tunnels.

You can use real time performance to configure the following options:

- ▶ Select the polling rate from 10 seconds up to 1 minute.
- ▶ Select up to 32 ports from up to a maximum of 10 switches for graphing performance.
- ▶ Choose to display the same Y-axis range for each displayed object per measure type for easier comparison of graphs.

To generate a real time performance graph for a switch, complete the following steps:

1. Select the fabric, switch, or port for which you want to generate a performance graph.
2. Select **Monitor** → **Performance** → **Real-Time Graph**.

The Real time Port Selector dialog box displays (see Figure 9-54).

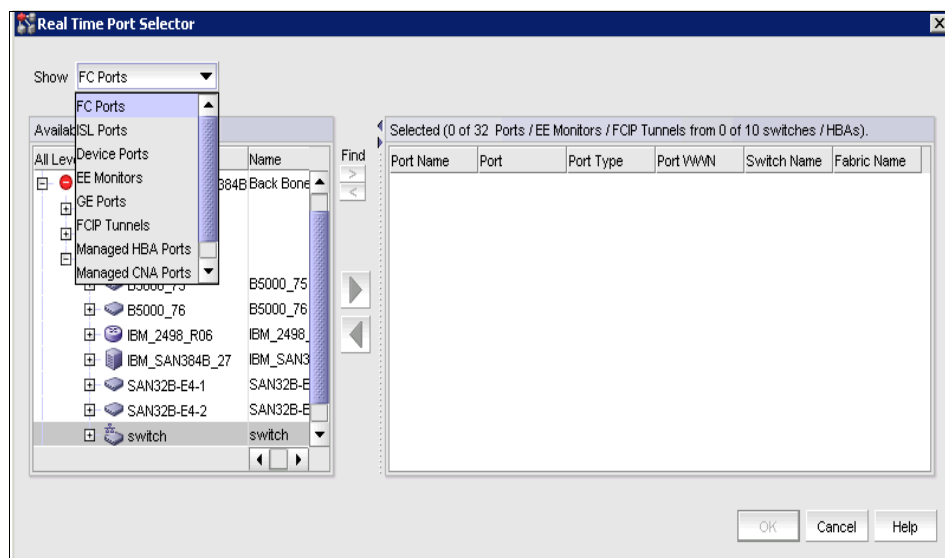


Figure 9-54 Real Time Port Selector

3. Select the object type (FC Ports, ISL Ports, Device Ports, EE Monitors, GE Ports, FCIP Tunnels, Managed HBA ports, Managed CNA ports and 10 GE ports) for which you want to graph performance from the Show list (see the rounded rectangle in Figure 9-54).
4. Click the right arrow to move the selected ports from the Available to the Selected table.

The Real Time Performance Graphs dialog box displays (see Figure 9-55).




Figure 9-55 Real Time Performance Graph

You can select the measure by which you want to gather performance data from the Measures list. To select more than one measure, click the *Additional Measures* expand arrows and select the check box for each additional measure (see Figure 9-56).

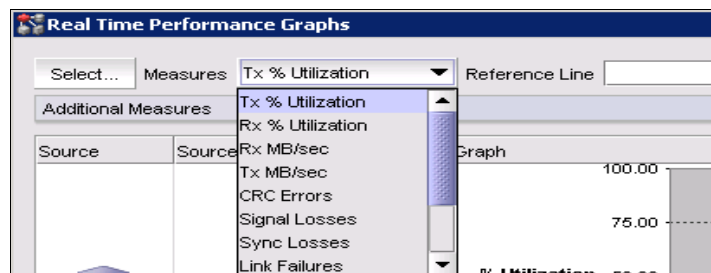


Figure 9-56 Additional Measures menu

The performance measures are described in detail in Table 9-5 on page 375.

9.7.4 Historical performance data

Performance should be enabled constantly to receive the necessary historical data required for a meaningful report. The following options and features are available for obtaining historical performance data, which:

- ▶ Enables you to collect historical performance data from the entire SAN or from selected switch.
- ▶ Persists data on every polling cycle (5 minutes).
- ▶ Stores up to 3456 records (maximum) for each port. Most ports require 600 KB disk space; however, the 256-Port Director requires 7GB disk space.
- ▶ Uses the Round Robin Database (RRD) style aging scheme.
- ▶ Enables 5 minute, 30 minute, 2 hour, and 1 day granularity.
- ▶ Supports interpolation for up to 6 data points.
- ▶ Generates reports.

To enable historical performance collection for all fabrics in the SAN, select **Monitor** → **Performance** → **Historical Data Collection** → **Enable SAN Wide**.

To enable historical performance collection for selected fabrics, select **Monitor** → **Performance** → **Historical Data Collection** → **Enable Selected**. (see Figure 9-57).

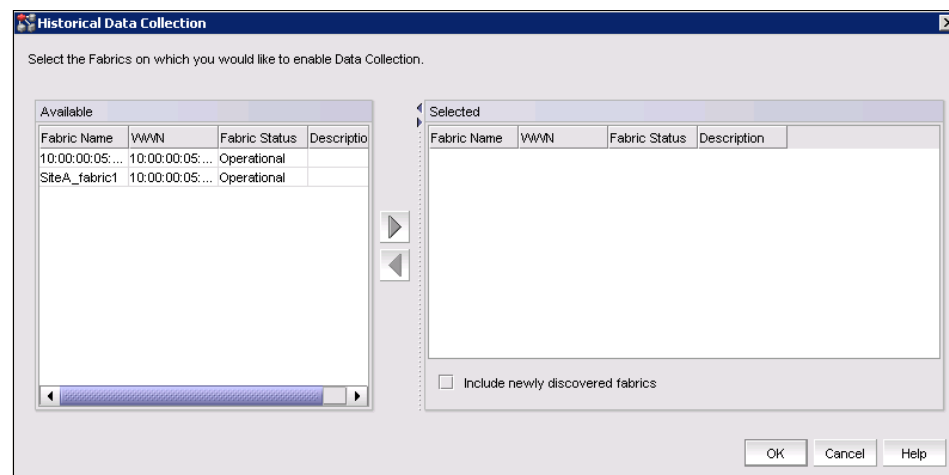


Figure 9-57 Historical performance collection for selected fabrics

Generating a historical performance graph

To generate a historical performance graph for a switch, complete these steps:

1. Select the switch for which you want to generate a performance graph.
2. Select **Monitor** → **Performance** → **Historical Graph**.
 - You can also right-click the switch and select **Performance** → **Historical Graph**.

The Historical Performance Graph dialog box displays (see Figure 9-58).

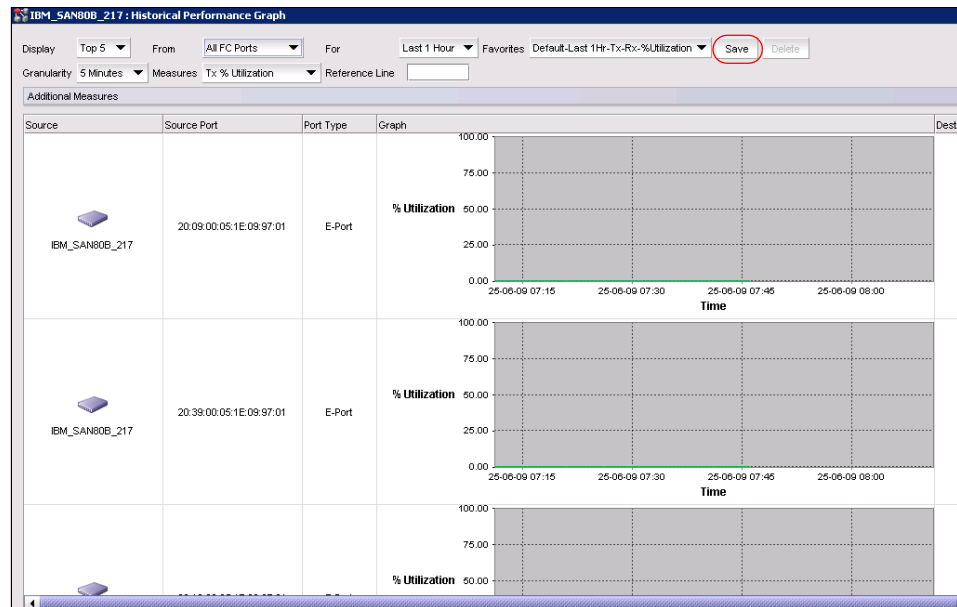


Figure 9-58 Historical Performance Graph

Historical Performance Graph:

- ▶ Display: The Top N performers (5, 10, 15, 20) for a selected object type.
- ▶ From Object: FC, Device, ISL Ports, FCIP Tunnels, EE Monitors, Custom
- ▶ For: Last hour/day/week/custom
- ▶ Granularity: 5/30/120/1920 minutes
- ▶ Measures: Many different measurements

You can filter the historical data by:

- ▶ Filtering data by ports
- ▶ Filtering data by time

You can save the historical performance graph by selecting the button **Save** (see the rounded rectangle in Figure 9-58).

Historical reports/tables

To generate a historical performance report for a device, proceed as follows:

- ▶ Select the device for which you want to generate a performance report.
- ▶ Select **Monitor** → **Performance** → **Historical Report**.
or Right-click the device and select **Performance** → **Historical Report**.

The Historical Performance Table dialog box displays (see Figure 9-59).




Figure 9-59 Historical Performance table

Click **Apply** to display Historical Performance Report (see Figure 9-60).

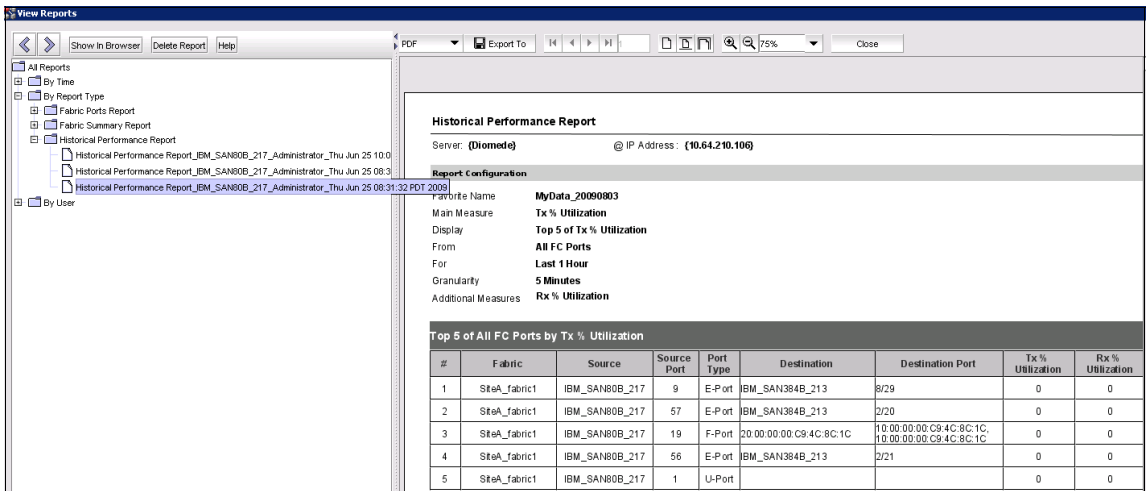


Figure 9-60 Historical Performance Report

9.7.5 Performance thresholds

Performance allows you to apply thresholds and event notification to real-time performance data, as well as historical performance data depending on the setting. A performance monitor process (thread) monitors the performance data against the threshold setting for each port and issues an appropriate alert to notify you when the threshold is exceeded.

License: A Threshold Policy requires a Fabric Watch License.

To create a threshold policy, select **Monitor** → **Performance** → **Configure Thresholds**. The Set Threshold Policies dialog box displays (see Figure 9-61).

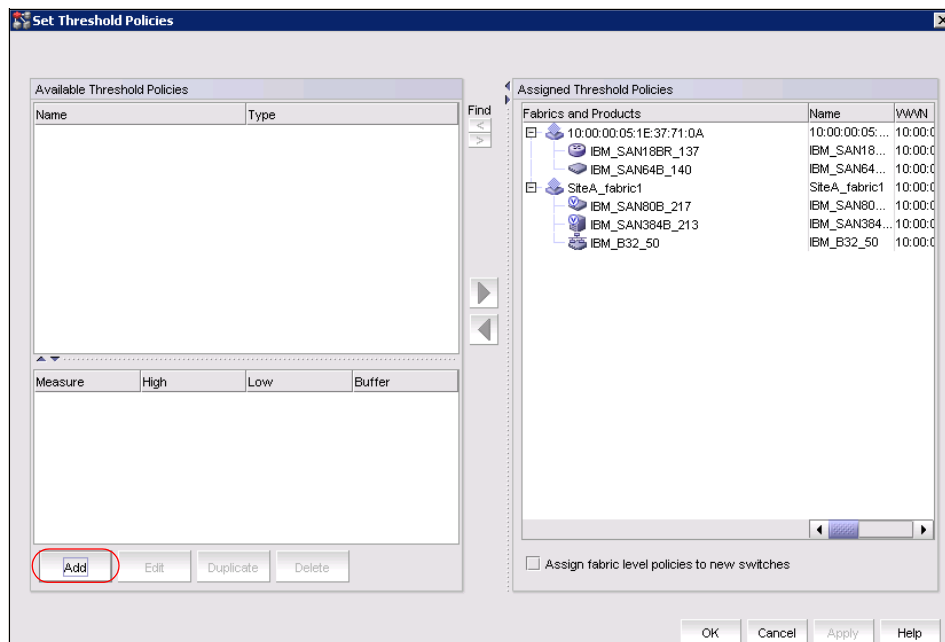
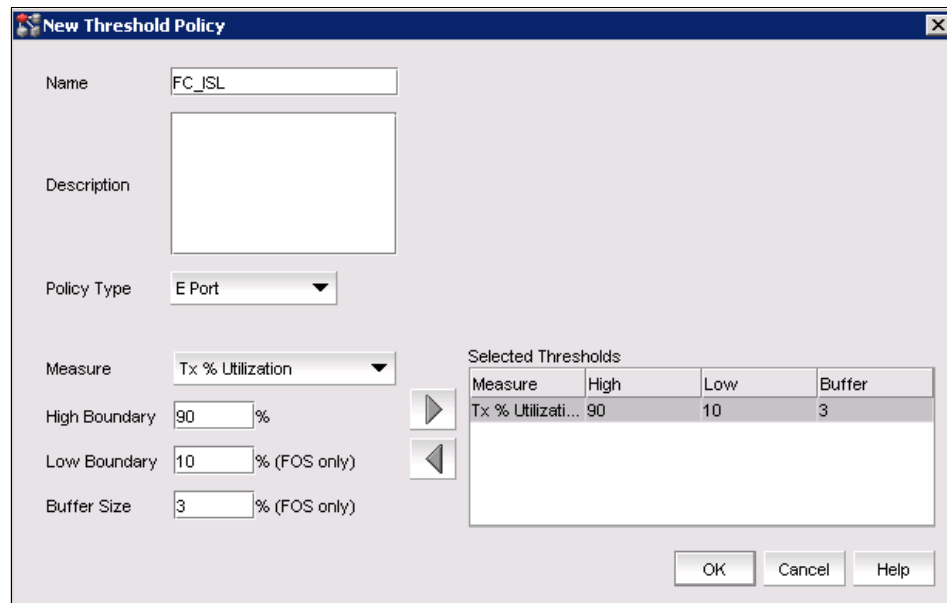


Figure 9-61 Set Threshold Policies dialog box

To add a policy, click the **Add** button and the New Threshold Policy dialog box displays (see Figure 9-62).



The dialog box titled "New Threshold Policy" contains the following fields and controls:

- Name:** A text box containing "FC_ISL".
- Description:** An empty text box.
- Policy Type:** A dropdown menu showing "E Port".
- Measure:** A dropdown menu showing "Tx % Utilization".
- High Boundary:** A text box containing "90" followed by a "%" symbol.
- Low Boundary:** A text box containing "10" followed by "% (FOS only)".
- Buffer Size:** A text box containing "3" followed by "% (FOS only)".
- Selected Thresholds:** A table with the following data:

Measure	High	Low	Buffer
Tx % Utilizati...	90	10	3

Navigation buttons (OK, Cancel, Help) are located at the bottom right.

Figure 9-62 New Threshold Policy

You can choose the following parameters:

- **Policy Type:**
Set for either E Port or F/FL Port
- **Measure:**
Choose Tx% Utilization, Rx% Utilization

Attention: You cannot add the same measure more than once. If you try to add another threshold with the same measure, the new values overwrite the older threshold values in the Selected Thresholds table.

You can create an SNMP trap and an event in the Master Event Log when thresholds are exceeded:

- ▶ High Boundary threshold <Measure, value set for high boundary> exceeded for <switch name>
- ▶ Low Boundary threshold <Measure, value set for low boundary> exceeded for <switch name>
- ▶ <Measure, value set for high or low boundary> has returned to normal for <switch name>

You can Edit, Duplicate, and Delete the policy as shown in Figure 9-63.

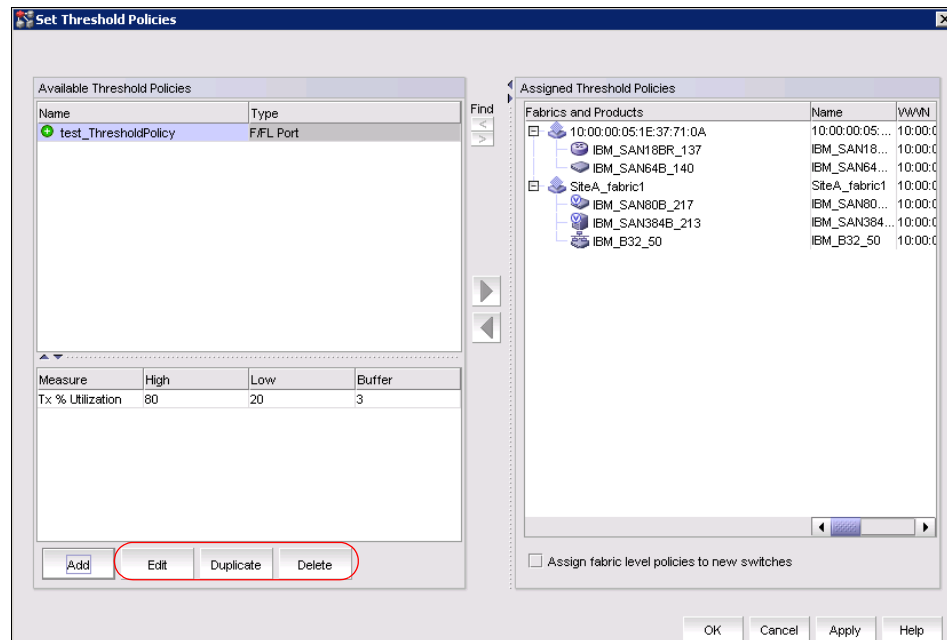


Figure 9-63 Set Threshold Policy

To assign a Threshold Policy, do the following steps (see Figure 9-64).

1. Select one or more threshold policies you want to assign to a fabric or switch in the Available Threshold Policies table.

Press Ctrl or Shift and then click to select multiple policies.

2. Click the right arrow button to apply the selected policies to the selected fabrics and switches.

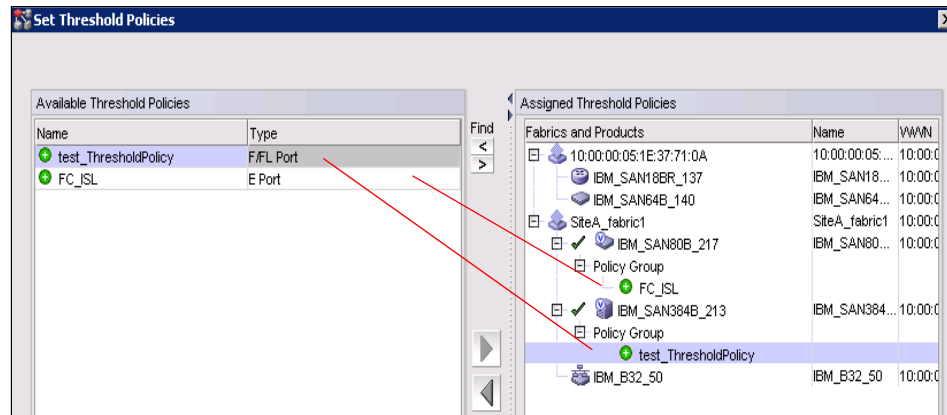


Figure 9-64 Assigning Threshold Policy

9.7.6 Connection utilization

Attention: Connection utilization is only supported on the following managed objects: E_ports, F_ports, N_ports, and FCIP tunnels.

Performance connection utilization for switch ports provides the following features:

- ▶ Enables you to turn the utilization display on and off from the menu and tool bar.
- ▶ Displays moving dotted colored lines that originate from a port.
- ▶ Displays two lines in the topology (when turned on); one represents percentage utilization for transmit and the other percentage utilization for receive. The movement of the line determines if it is a transmit or a receive:
 - Receive (Rx): The line moves into a port.
 - Transmit (Tx): The line moves out of a port.
- ▶ Displays different colors to represent the percentage utilization range

Connections: Fabrics where performance data collections are not enabled display connections as thin black lines (normal display).

The Utilization Legend display and configuration is shown in Figure 9-65.

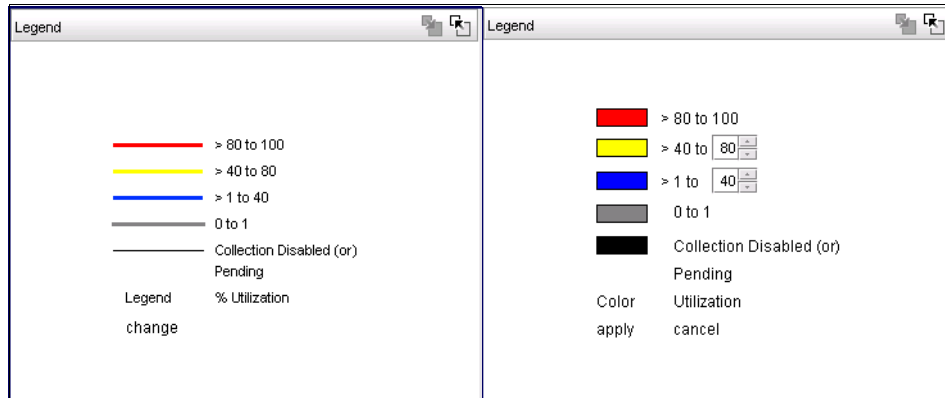


Figure 9-65 Utilization Legend display and configuration

The range is configurable:

- ▶ 0 to 1% Utilization = gray line.
- ▶ 1 to T1, where T1 is the first threshold, = dashed blue line.
- ▶ T1 to T2, where T2 is the second threshold, = dashed yellow line.
- ▶ T2 to 100% = dashed red line.

If Data Collection is Disabled for a Fabric, standard black connection lines are displayed.

To enable connection utilization select you can do one of the following actions:

- ▶ Select **Monitor** → **Performance** → **View Utilization**.
- ▶ Press CTRL + U.
- ▶ Click the Utilization icon in the ToolBar (see Figure 9-66).

Figure 9-66 shows the “Marching Ants,” the colored animated dashed line in the View Utilization view, which display utilization for FC and FCIP.

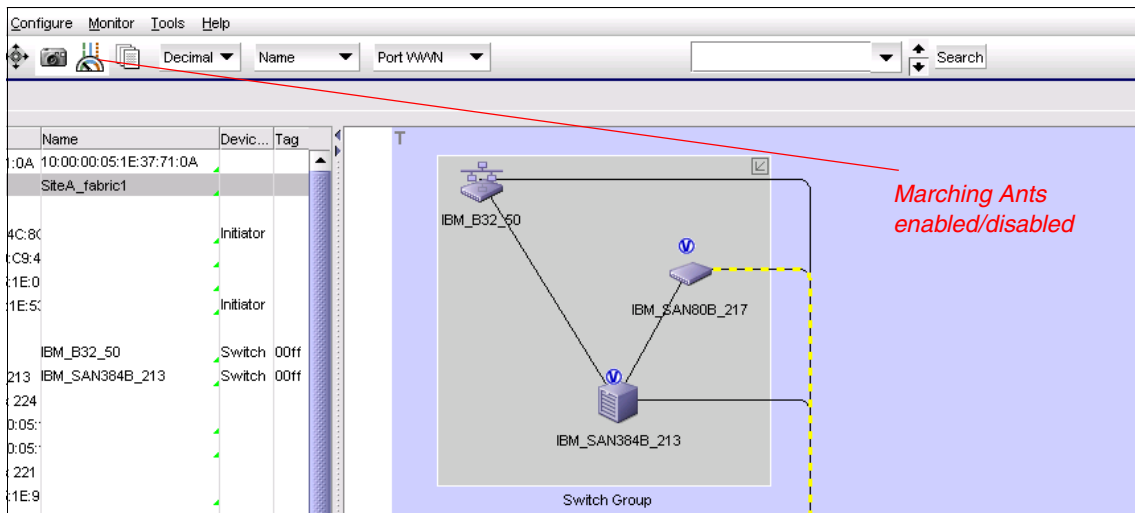


Figure 9-66 Displaying colored animated dashed line

9.8 Encryption configuration

You can configure encryption switches from DCFM. Go to the Encryption Center and select **Configure** → **Encryption** as shown in Figure 9-67 to go to the Encryption Center.

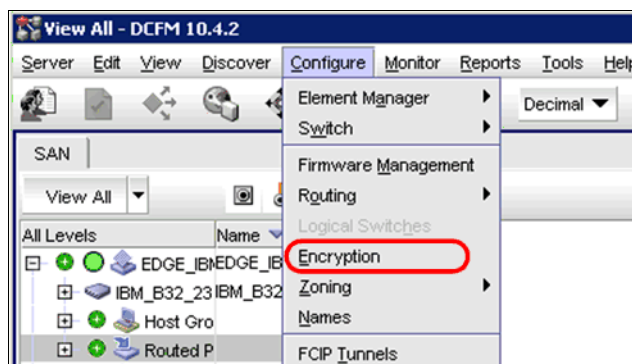


Figure 9-67 Display the Encryption Center

The Encryption Center will start. From here you have the possibility to configure an encryption switch or an encryption blade (SAN32B-E4 or Encryption Blade). See Figure 9-68.

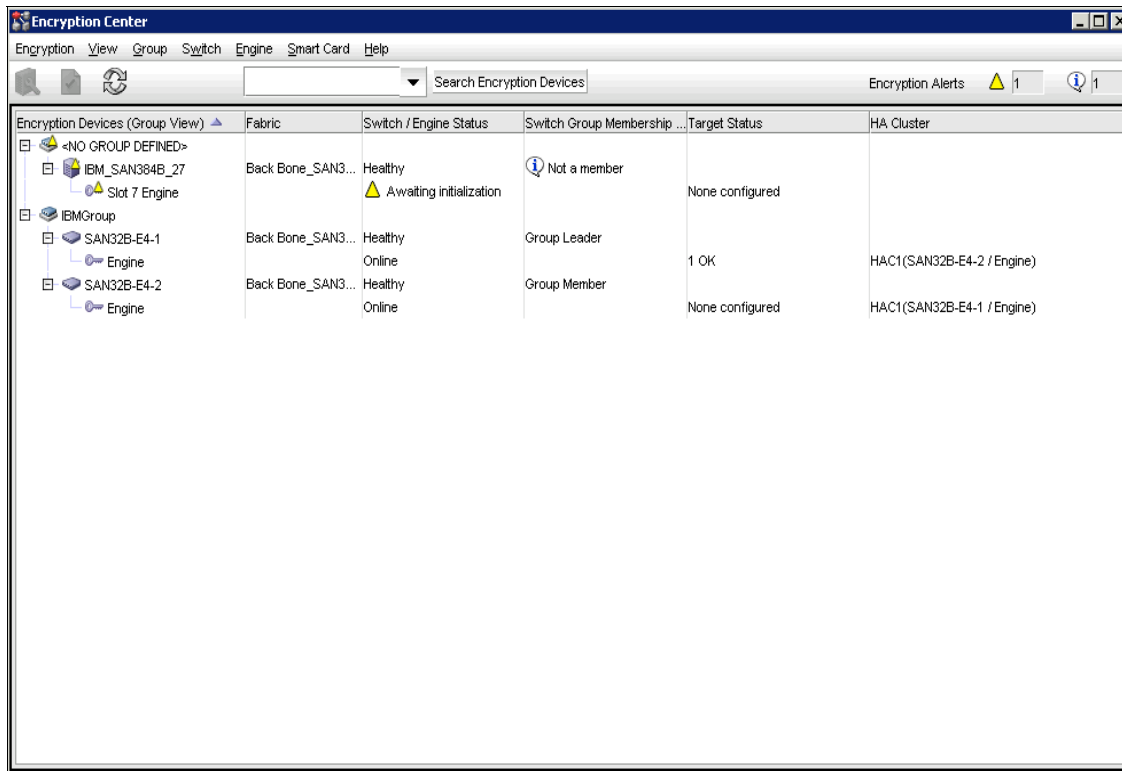


Figure 9-68 Encryption Center

You can perform the following configurations:

- ▶ Storage Encryption Configuration:
 - Launch the Configure Encryption dialog.
 - View switch, group, or engine properties.
 - View the Encryption Group Properties Security tab.
 - View encryption targets, hosts, and LUNs.
 - View LUN centric view.
 - View all re-key sessions.
 - Add/remove paths and edit LUN configuration on LUN centric view.
 - Rebalance encryption engines.
 - Decommission LUNs.
 - Edit smart card.

- Create a new encryption group or add a switch to an existing encryption group.
- Edit group engine properties (except for the Security tab).
- Add targets.
- Select encryption targets and LUNs to be encrypted or edit LUN encryption settings.
- Edit encryption target hosts configuration.
- ▶ Storage Encryption Key Operations:
 - Launch the Configure Encryption dialog.
 - View switch, group, or engine properties,
 - View the Encryption Group Properties Security tab.
 - View encryption targets, hosts, and LUNs.
 - Initiate manual LUN re-keying.
 - Enable and disable an encryption engine.
 - Zeroize an encryption engine.
 - Restore a master key.
 - Edit key vault credentials.
- ▶ Storage Encryption Security:
 - Launch the Configure Encryption dialog.
 - View switch, group, or engine properties.
 - View encryption targets, hosts, and LUNs.
 - Create a master key.
 - Backup a master key.
 - View and modify settings on the Encryption Group Properties Security tab (quorum size, authentication cards list and system card requirement).
 - Establish link keys for LKM key managers.

For more details, see *Implementing the IBM System Storage SAN32B-E4 Encryption Switch*, SG24-7922, available at this website:

<http://www.redbooks.ibm.com/abstracts/sg247922.html?Open>

9.9 User management

DCFM allows you to manage users and to start select **Server** → **User** as shown in Figure 9-69.

Privilege: You must have the User Management privilege to perform this task. This user are System Administrator and Security Officer.




Figure 9-69 Display user management

You can now manage the different users. Figure 9-70 shows the possible actions:

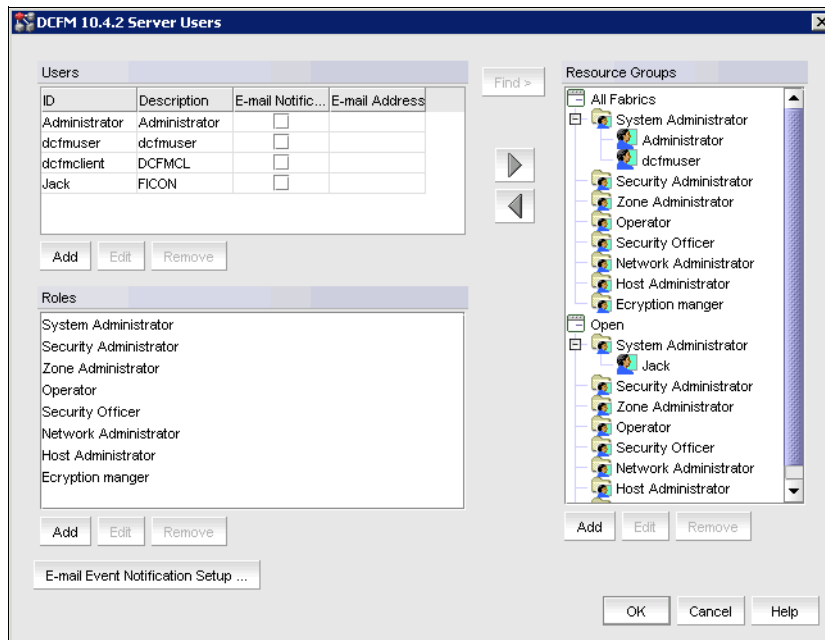


Figure 9-70 Server User management

- ▶ User definitions:
 - Add new user
 - Delete existing user
 - Edit a user definition
- ▶ Resource Groups definitions:
 - Add new resource group
 - Delete existing resource group
 - Edit the resource group
- ▶ Roles definition:
 - Add new role
 - Delete an existing role
 - Edit a role
- ▶ Email event notification setup

9.10 DCFM Server Management Console

The DCFM Server Management Console (SMC) is an automatically installed, stand-alone application for managing the management application server.

From Windows you can launch the management console from the **Start → All Programs** menu and then select the DCFM folder. Select the **Server Management Console**.

The SMC panel has six tabs, as shown in Figure 9-71.

You can perform the following tasks using the SMC:

- ▶ From the Services tab, you can start, stop, and restart services on the server.
- ▶ From the Ports tab, you can change the management application server or Web Server port number.
- ▶ From the Authentication tab, you can configure an authentication server (LDAP or Radius server), and establish authentication policies.
- ▶ From the Database tab, you can restore server application data.
- ▶ From the Technical Support Information tab, you can collect information for technical support.
- ▶ The HCM Upgrade tab enables you to upgrade the management application to include a new version of HCM.

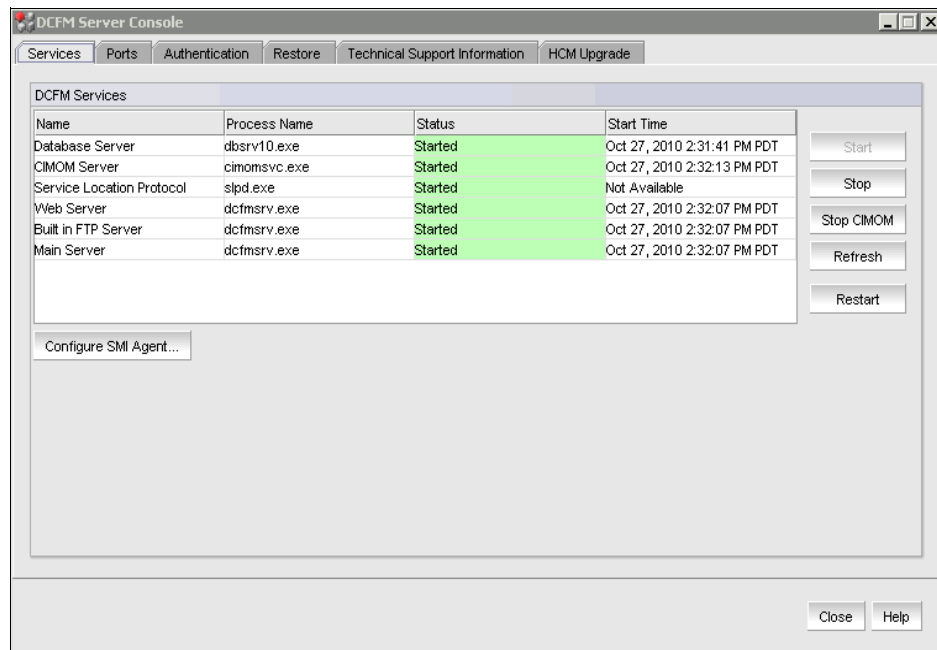


Figure 9-71 DCFM Server Management Console

9.10.1 Changing server port numbers

Use the Ports tab of the Server Management Console to change the management application server and Web Server port numbers (see Figure 9-72). The default Web Server port number is 80. The management application server default port number is 24600.

Server restart: The server automatically restarts if you change the server port number. You must manually restart the server if you change only the Web Server port number.

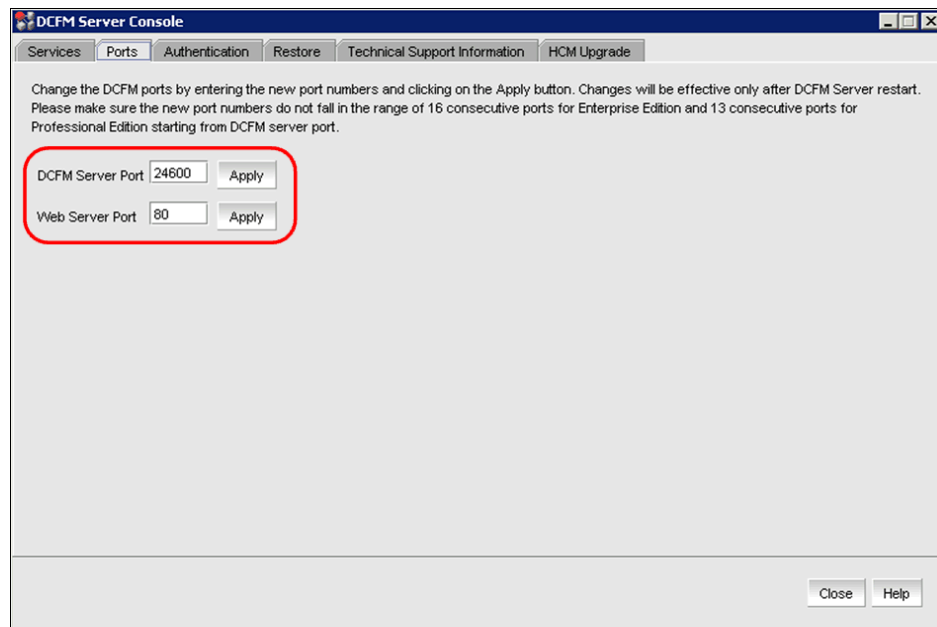


Figure 9-72 DCFM Server Management Console Ports

9.10.2 Restoring the database

First, configure the backup options for the DCFM server database as well as the path to the backup files. Go to the DCFM main menu under **SAN** → **Options**. The Options dialog box will display as shown in Figure 9-73.

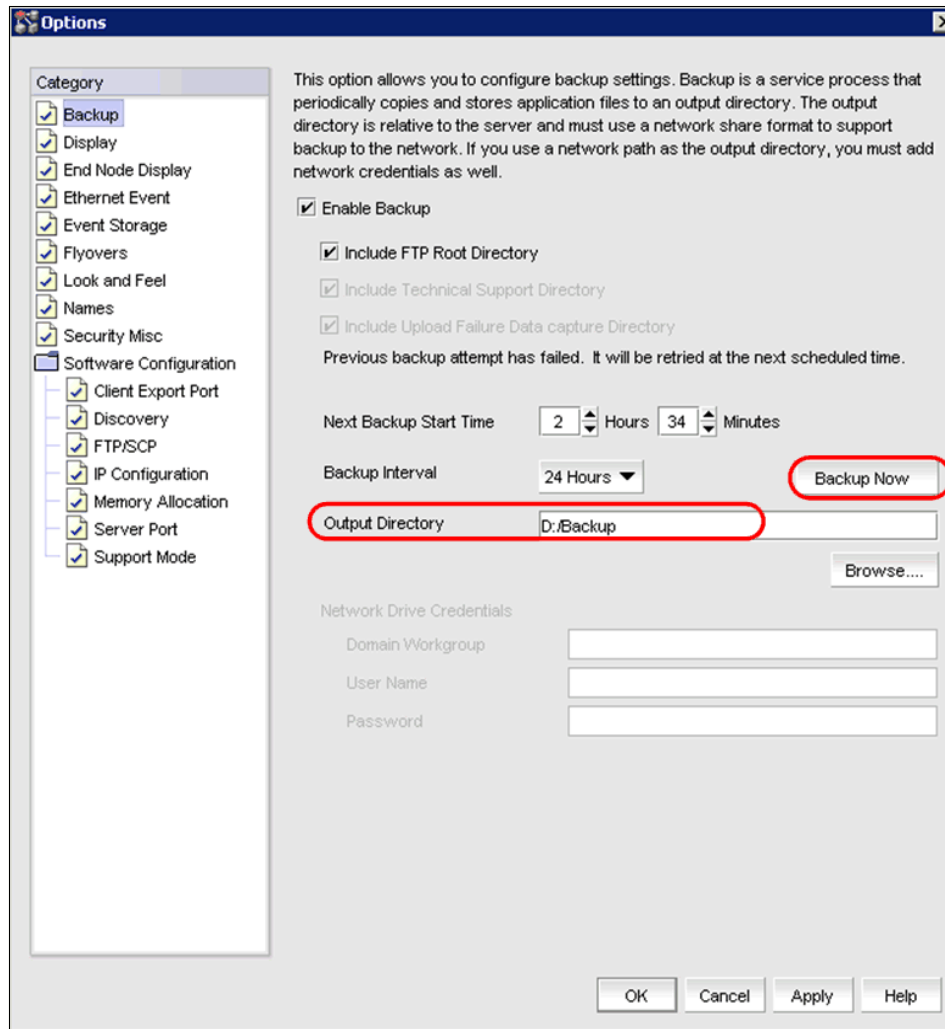


Figure 9-73 Backup Options dialog box

From the dialog box, you can choose **Backup** on the left pane. Set the Backup Interval, Output Directory, and Backup Now options. When a backup is available you can restore this backup from DCFM Server Management Console.

Go to the Restore Database option in DCFM Server Management Console. To perform this operation, you must know the path to the backup files (see rounded rectangle in Figure 9-74).

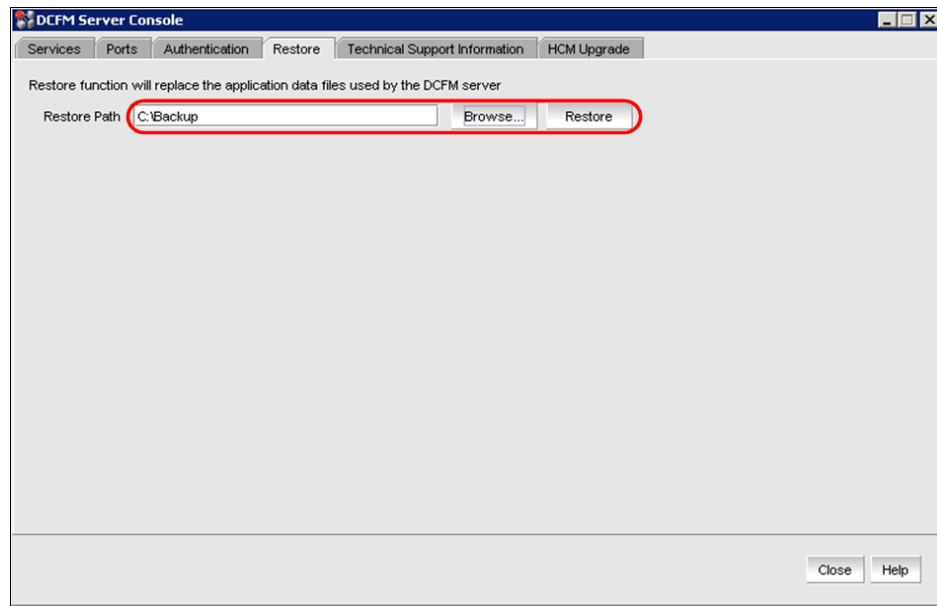


Figure 9-74 DCFM Server Management Console - Restore

Insert the path where the backup files are stored and then click **Restore** to start the process of restoring the DCFM Server Database.

9.10.3 Configuring authentication

You can configure the method that the DCFM server will use for authentication as shown in Figure 9-75.

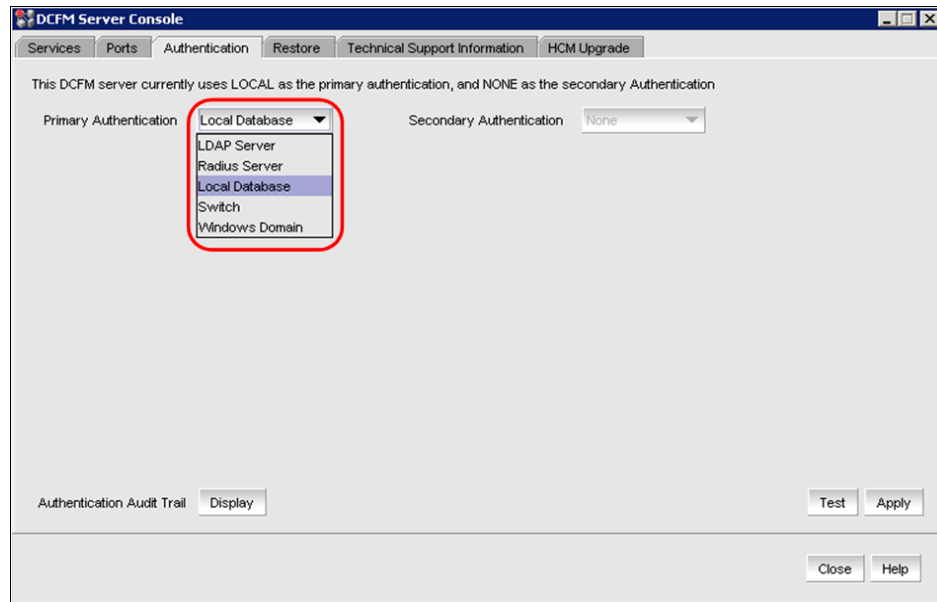


Figure 9-75 DCFM Server Console Authentication

All responses to authentication requests coming from clients are logged to an audit trail log file.

This file is automatically backed up on the first day of every month.

Select the **Authentication** tab and click **Display** next to the Authentication Audit Trail. The Login dialog box displays. Enter your username and password in the appropriate fields and click **OK**.

The Authentication Audit Trail log displays as in Figure 9-76.

The audit trail shows user names that have attempted to log in to the management application, and changes to user authentication.




Figure 9-76 Authentication Audit Trail

9.10.4 Capturing technical support information

The Technical Support Information tab of the SMC allows you to capture technical support information, such as server data for all services. This information is saved in a zip file in a location that you specify (see Figure 9-77).

If you do not specify an output path, the management application automatically saves the data to the <Install DIR>/support directory.

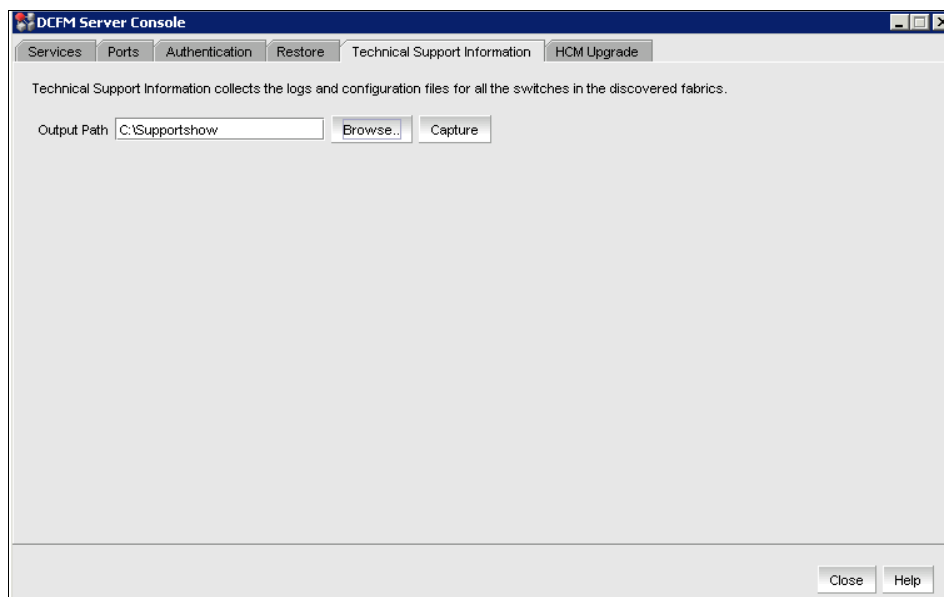


Figure 9-77 Technical Support Information

It allows users to capture data collection of the DFCM server (Figure 9-78).

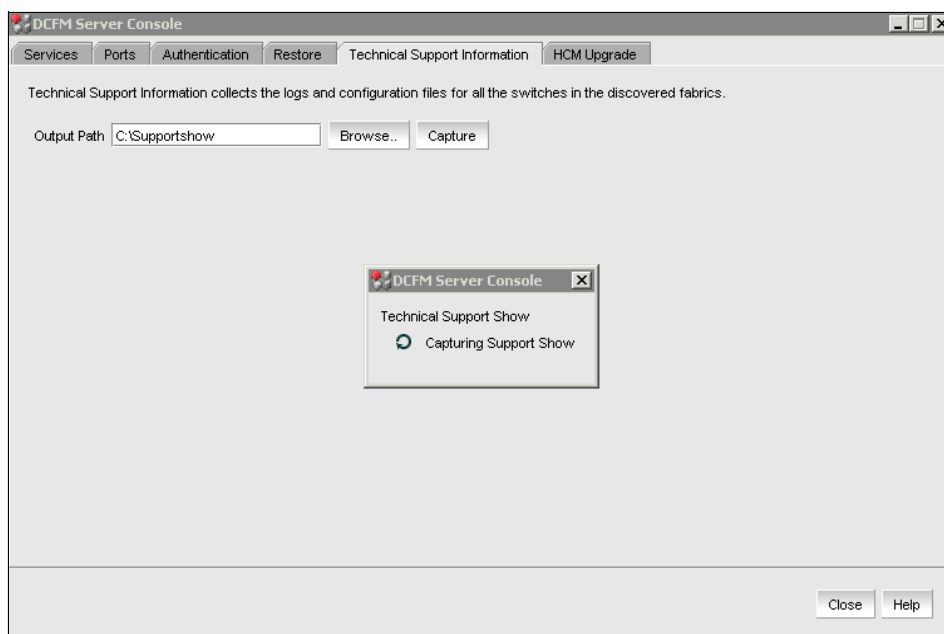


Figure 9-78 Technical Support - Capturing Support Show

9.10.5 Gathering switch information for support

To gather switch information for support (it will activate the **supportsave** command), right-click the switch, then choose from DCFM main menu, **Technical Support** → **Switch / Host SupportSave** (see Figure 9-79).

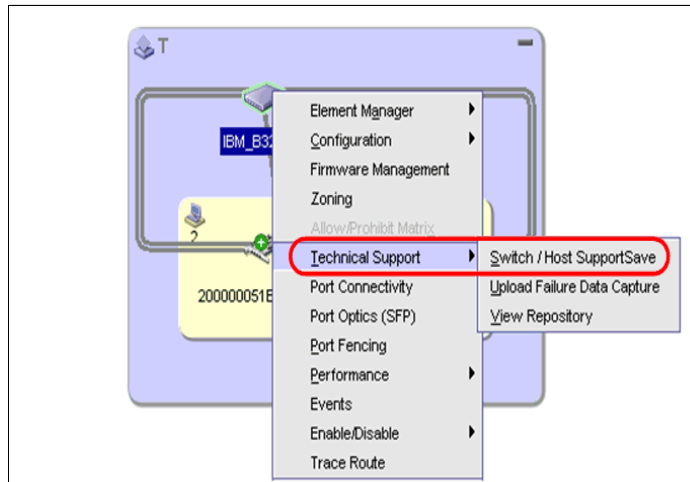


Figure 9-79 Collecting data for switch

The Technical SupportSave dialog box displays as shown in Figure 9-80. Select the switches you want to collect data for in the Available Core Switches table and click the right arrow to move them to the Selected Switches table.

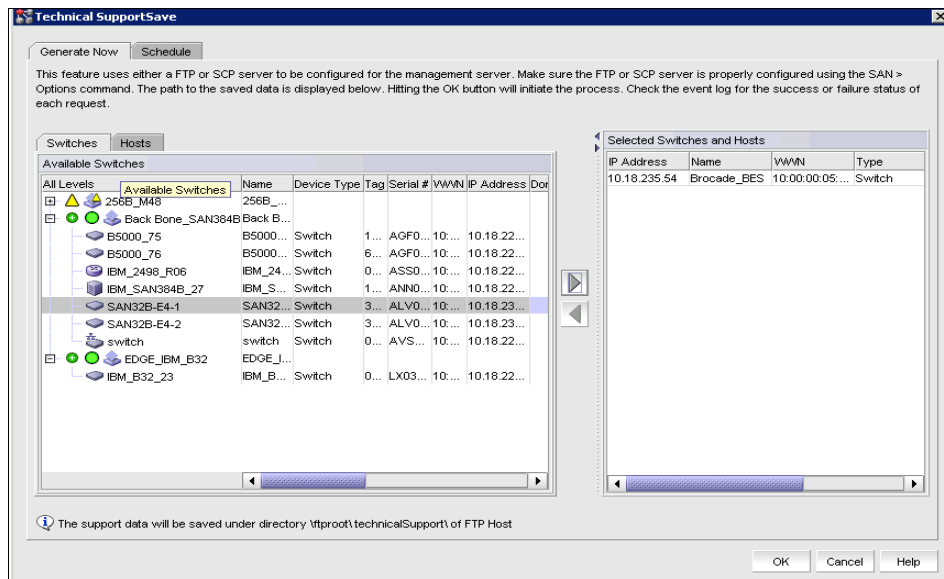


Figure 9-80 Technical Support Data

Note the location where the **supportsave** data is being written to, at the bottom of the window. Click **OK** to start the collection. A message box will display as shown in Figure 9-81.




Figure 9-81 DCFM supportsave message box

9.10.6 Viewing technical support information

To view the technical support information, select **Monitor** → **Technical Support** → **View Repository**.

The **supportSave** repository will be displayed as shown in Figure 9-82. From here you can ftp, email, delete, or view the available **supportsave**. For email or ftp, you have to configure this first.

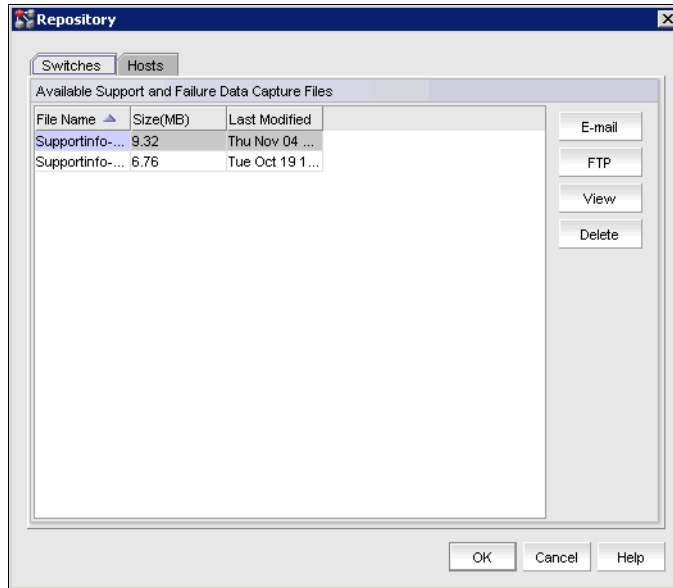


Figure 9-82 Supportsave repository

9.10.7 HMC upgrade

As you see in Figure 9-83 you can also update the HMC from the DCFM server Console. Select the location where the update is stored and click **Upgrade**.




Figure 9-83 HMC upgrade



Host Connectivity Manager

The Host Connectivity Manager (HCM) is a management software application for configuring, monitoring, and troubleshooting Brocade HBAs and Converged Network Adapters (CNAs) in a storage area network (SAN) environment.

The management software has two components:

- ▶ The agent, which runs on the host
- ▶ The management console, which is the graphical user interface client used to manage the HBA or CNA

You can manage the software on the host or remotely from another host. The communication between the management console and the agent is managed using Java Script Object Notation - Remote Procedure Call (JSON-RPC) over https.

Reference: All HCM, utility, SMI-S Provider, boot software, and driver installation packages, as well as the Driver Update Disk (DUD), are described in the *Brocade Adapters Installation and Reference Manual*, 53-1001254-05, available at this website:

<http://www.brocade.com/services-support/drivers-downloads/HBA/index.page>

10.1 HCM features

In this section we discuss software features and the tree node menus.

10.1.1 Software features

Common HBA and CNA management software features include these:

- ▶ Discovery using the agent software running on the servers attached to the SAN, which enables you to contact the devices in your SAN
- ▶ Configuration management, which enables you to configure local and remote systems. With HCM you can configure the following items:
 - Local host
 - Brocade 4 Gbps and 8 Gbps HBAs
 - HBA ports (including logical ports, base ports, remote ports, and virtual ports)
 - Brocade 10 Gbps single-port and 10 Gbps dual-port converged network adapters (CNAs)
 - CEE ports
 - FCoE ports (CNA only)
 - Ethernet ports (CNA only)
- ▶ Diagnostics, which enable you to test the adapters and the devices to which they are connected:
 - Link status of each adapter and its attached devices
 - Loopback test, which is external to the adapter, to evaluate the ports (transmit and receive transceivers) and the error rate on the adapter
 - Read/write buffer test, which tests the link between the adapter and its devices
 - FC protocol tests, including echo, ping, and traceroute
 - Monitoring, which provides statistics for the SAN components listed in Table 10-1.
 - Security, which enables you to specify a CHAP secret and configure authentication parameters
 - Event notifications, which provide asynchronous notification of various conditions and problems through a user-defined event filter

10.1.2 Tree node pop-up menus

You can use the HCM GUI main menu or the Command Line Utility to configure, monitor, and troubleshoot your SAN components. The instructions for using each feature are detailed in subsequent sections of this chapter. For each SAN component, you can optionally right-click its icon and a pop-up menu displays, showing the features available for that component (see Table 10-1).

The HCM GUI consists of three layers, and the features display differently depending on the configuration. There are three possible configuration scenarios, as follows:

- ▶ Both the storage driver and the link layer driver are installed.
- ▶ Only the storage driver is installed.
- ▶ Only the link layer driver is installed.

Whether the FCoE Port node or the Ethernet node are presented in the tree depends on the drivers that are installed.

Table 10-1 HCM tree pop-up menus

SAN component	Pop-up menu feature
Host	Refresh All Start Polling Upload Boot Code Image Change Agent Password Configure Names Basic Port Configuration Persistent Binding Statistics → Teaming Statistics Authentication Teaming Support Save Restore VLAN and Team Tree → Copy Search Collapse All Expand All
Brocade HBA 4 Gbps or 8 Gbps	Refresh Define Name Upload Boot Code Image Basic Port Configuration Persistent Binding Port Statistics Diagnostics Authentication Enable Adapter Tree → Copy Search Collapse All Expand All

SAN component	Pop-up menu feature
HBA Port	Refresh Define Name Port Configuration → Basic Advanced Persistent Binding Virtual Port → Create Delete Statistics → Port statistics FCP IM Module statistics Fabric Statistics IOC Statistics QoS Statistics Diagnostics FC-SP → Authentication Authentication Statistics Enable Port Beacon → Port Link Tree → Copy Search Collapse All Expand All
Converged Network Adapter (CNA), 10 Gbps	Refresh Define Name Upload Boot Code Image Basic Port Configuration Persistent Binding Port Statistics Diagnostics Authentication Enable Adapter Tree → Copy Search Collapse All Expand All
FCoE Port	Refresh Persistent Binding Virtual Port → Create Delete Statistics → Fabric IOC FCOE FC-SP → Authentication Authentication Statistics Enable Port Tree → Copy Search Collapse All Expand All
Ethernet Port	Refresh Statistics → Eth Eth IOC VLAN VLAN Configuration Tree → Copy Search Collapse All Expand All

10.2 Getting started with HCM software

In the topics that follow, we show some of the key features of HCM.

10.2.1 HCM software launch

The following procedures describe how to launch the HCM application in Windows and Linux.

Launching the application on Windows platforms

After installing the HCM software, locate Brocade HCM on the Windows platform by selecting **Start → Programs → Brocade Adapter Software → Host Connectivity Manager** or click the desktop Host connectivity manager icon to launch the application.

The Login Dialog box (Figure 10-1) displays when the HCM software is first launched.

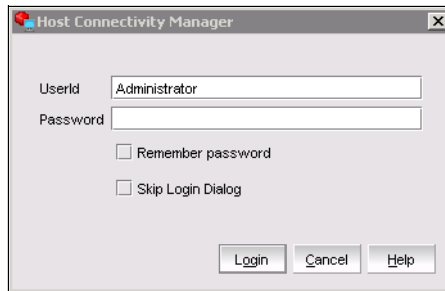


Figure 10-1 HCM Login Dialog box

The factory default user ID and password are *Administrator* and *password*. After you log in for the first time, change the default password to a new one using the HCM GUI.

Launching the application on Linux platforms

After installing the HCM software, locate Brocade HCM on the Linux platform:

- ▶ If using a GNOME shell, double-click the **Host connectivity manager** icon to launch the application.
- ▶ If using a KDE shell, single-click the **Host connectivity manager** icon to launch the application.

Or start the application from the GUI console terminal command prompt using the following commands:

- ▶ `[root@dpe2950228036 ~]# cd /opt/brocade/adapter/client/`
- ▶ `[root@dpe2950228036 client]# ./hcm.sh`

Attention: This command will not work from the SSH terminal, it is required to be executed from the GUI console connected locally or, for example, by VNC viewer.

Launching the application from web browser

After installing the HCM software in the server with the Brocade HBA/CNA, Host Connectivity Manager can also be connected using the web browser by entering the URL:

`https://server-host:34568/index.html`

Where:

- ▶ `server-host` is the hostname or IP address of the server with the Brocade HBA/CNA adapter with the driver installed and the HCM agent running.
- ▶ `34568` is the TCP/IP port through which the HCM agent communicates with the HCM server.

10.2.2 Command line utility

The HCM has a command line utility known as the Brocade CLI Utility (BCU). This utility can be started from the shortcut in the Desktop or from the command line prompt (Example 10-1).

Example 10-1 BCU - Brocade CLI utility

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu --version
Brocade CLI utility
Version:FCHBA2.2.0.2
```

10.2.3 HCM configuration data

The HCM configuration data files hold values defined by user for HCM. The default data folder depends on the platform:

Windows:

`C:\Users\Administrator\HCM\data (<user home dir>\hcm\data)`

Linux:

`/root/hcm/data (<user home dir>/hcm/data)`

10.2.4 Remembering the password

The Login dialog has a check box to remember the password. If you check the **Remember password** check box, you do not need to enter the password the next time you launch the application.

10.2.5 Skipping login

Take one of the following actions to manage the Skip Login feature:

- ▶ Enable *Skip Login* by checking the **Skip Login Dialog** check box:
If the Skip Login check box is checked, it automatically disables the **Remember password** option.
- ▶ Disable **Skip Login** by setting *hba-application.skip-login=false* in the file:
<user home Dir >HCM/data/HBAAApplication.properties

Select the **Skip Login** check box if you do not want the Login dialog box to display the next time the application is started.

10.2.6 Changing an HCM application password

You can change the default password of the application to a different password using the Change HCM Password dialog.

Note the following considerations when you change a password:

- ▶ You must validate your user identity by supplying your old password before you can change to a new password. The new password must be different than the old password.
- ▶ The password can begin with an alphabetic, numeric, or special character.
- ▶ The default minimum and maximum length of the password is 8 and 64 characters. You can configure the password length in the file:
<user home Dir >/HCM/data/HBAAApplication.properties

```
# min chars for the application password
password_min=8
#max chars for the application password
password_max=64
```
- ▶ The password is encrypted and stored in the *noitacitnehtua.properties* file.

Follow these steps:

1. From the Host Connectivity Manager, select **Configure** → **Change Password** → **Change Password for HCM User**; the Change HCM Password dialog box displays (Figure 10-2).

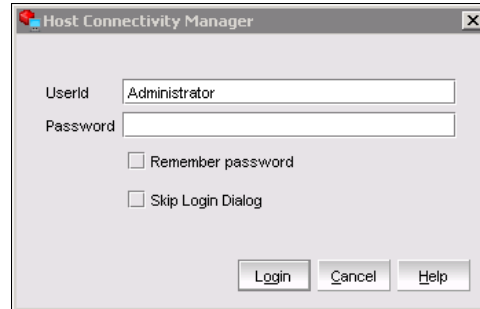


Figure 10-2 Change HCM Password dialog box

2. Type the current password for the account. The default user name and password are *Administrator* and *password*.
3. Type the old password for the account.
4. Type the new password of the account.
The new password must have at least one character different from the old password.
5. Retype the new password in the Confirm New password field.
6. Click **OK**.

Case: Both the user name and passwords are case-sensitive.

10.2.7 Changing an HCM agent password

You can change the default password of the agent to a different password using the **Change HCM Agent Password** dialog.

Note the following considerations when you change a password:

- You must validate your user identity by supplying your old password before you can change to a new password. The new password must be different than the old password.

Follow these steps:

1. From the Host Connectivity Manager, click **Configure** → **Change Password** → **Change Agent Password**; the Change HCM Agent Password dialog box displays (Figure 10-3).

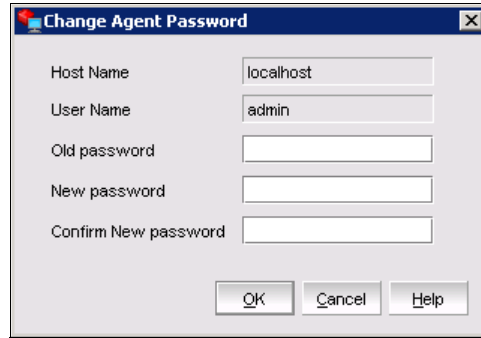
The image shows a Windows-style dialog box titled "Change Agent Password". It has a standard title bar with a close button (X). The dialog contains five text input fields: "Host Name" (pre-filled with "localhost"), "User Name" (pre-filled with "admin"), "Old password", "New password", and "Confirm New password". At the bottom right, there are three buttons: "OK", "Cancel", and "Help".

Figure 10-3 Change HCM Agent password dialog box

2. Type the current password for the account. The default user name and password are *admin* and *password*.
3. Type the new password of the account.
The new password must have at least one character different from the old password.
4. Retype the new password in the *Confirm New password* field.
5. Click **OK**.

Case: Both the user name and passwords are case-sensitive.

10.2.8 Resetting a password or restoring a factory default password

After a successful installation, copy this file to your personal folder:

```
<user home Dir> /HCM/data/noitacitnehtua.properties
```

You do this so that in case the password is lost, you can overwrite the `noitacitnehtua.properties` file in the data folder with the local copy. This restores the factory default user name (*Administrator*) and password (*password*).

10.2.9 Backing up data after an uninstall

If you uninstall the Brocade HCM software, you are prompted to back up the application configuration data that was created during installation.

The following application configuration files are backed up in the data directory:

- ▶ HBAApplication.properties
- ▶ SetupDiscovery.properties
- ▶ HbaAliasdb.properties
- ▶ log4j.xml
- ▶ noitacitnehtua.properties

To restore the backed-up configuration data when you re-install the HCM, you must manually overwrite the new data directory contents with the backed-up data. This restores your previous settings. The restore can be also done by using HCM as described in 10.2.10, “Backing up HCM data using HCM”.

10.2.10 Backing up HCM data using HCM

To back up HCM data whenever required using HCM, select the **Host** and then click **Tools** → **Backup HCM Data**, the Backup HCM Data window displays (Figure 10-4).

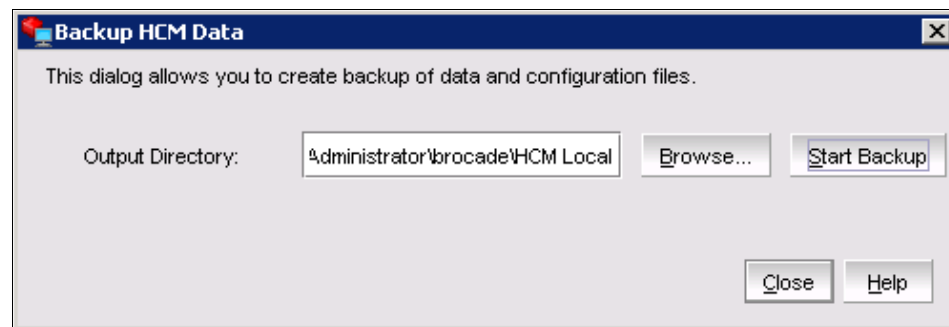


Figure 10-4 HCM Data backup

We need to enter a Backup directory and then click **Start Backup**, which will complete and indicate its success with a Backup completion message (Figure 10-5).




Figure 10-5 Backup completion

Click **OK** and close the Backup HCM Data window, and the required config data can be now seen in the backup directory mentioned, which will be used later when a restore is required.

10.2.11 Restoring HCM data using HCM

Restoring HCM data is performed from the Host Connectivity Manager by selecting the host and then click **Tool** → **Restore** → **HCM Data**. The **HCM Data** window displays (Figure 10-6).

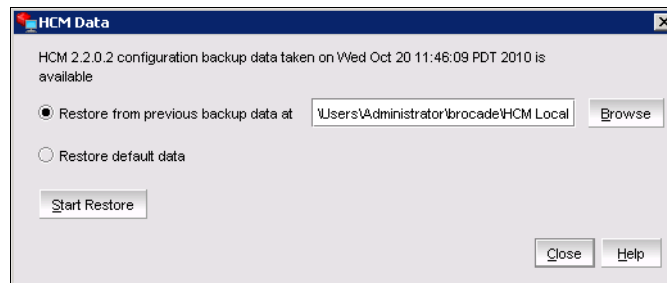


Figure 10-6 HCM Data restore

We need to enter the backup directory and click **Start Restore**, which will complete the restore, and success will be indicated with the completion message (Figure 10-7).




Figure 10-7 HCM Data restore completion

HCM has to be restarted to use the restored data config.

10.2.12 HCM main window

From the Host Connectivity Manager main window (Figure 10-8), you can manage all the adapters installed in this computer. Alternatively, you can manage adapters installed in remote computers, if the computers are networked. Only one host can be managed at a time; multiple host management is not supported.

For instructions on how to install both the driver and GUI, the driver only, or the GUI only, see the *Brocade Adapters Installation and Reference Manual*, 53-1001254-05, available at this website:

<http://www.brocade.com/services-support/drivers-downloads/HBA/index.page>

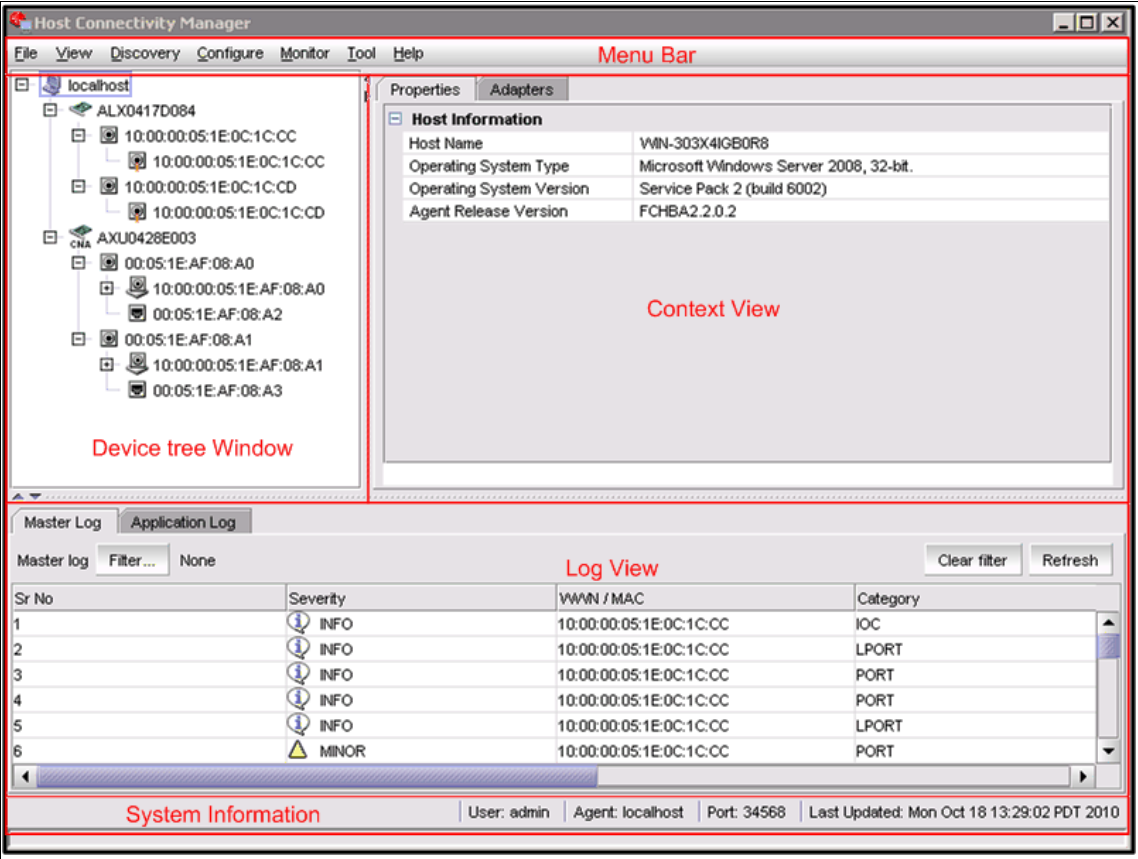
















Figure 10-8 Host Connectivity Manager main window








10.2.13 HCM product icons

On the left side of the Host Connectivity Manager, there is a navigation tree for representing the managed host with adapters and ports. Each tree node has an icon to represent the type of node. If the operational status is offline, link-down, or error, a small red diamond displays on the upper right corner of the icon.

Table 10-2 shows the product icons that represent the components that HCM manages.

Table 10-2 HCM product icons

Item	Icon	Item	Icon
Host (Agent UP)		Remote Port (Initiator) online	
Host (Agent Down)		Remote Port (Initiator) offline	
HBA Online		Remote Port (Target) online	
HBA Offline		Remote Port (Target) offline	
CNA Online		LUN	
CNA Offline		Ethernet Port	
Port (with SFP) link up		Base Port (link up)	

Item	Icon	Item	Icon
Port (with SFP) link down		Base Port (link down)	
Port (without SFP) link up		Virtual Port (online)	
Port (without SFP) link down		Virtual Port (offline)	
FCoE Port			

10.2.14 Discovery

Discovery enables you to contact the adapters present in a specified host in your SAN. The setup discovery profile is saved in the SetupDiscovery.properties file to remember the history of each host and related attributes of discovered hosts.

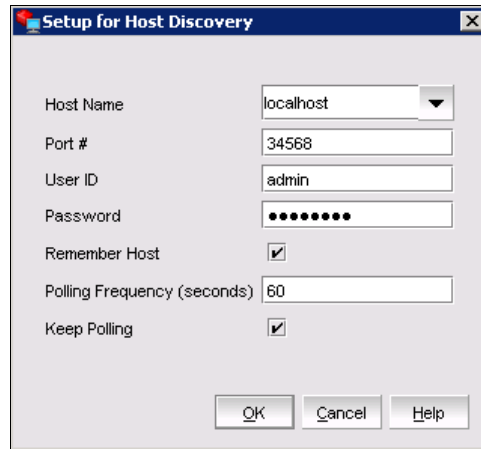
When you log in to HCM, the specified host is automatically contacted (discovered) and displayed on the navigation tree. By default, the local host is automatically contacted (discovered) and displayed on the navigation tree. When you configure and turn on discovery, the application discovers Brocade adapters in that host, connected to the SAN.

Discovery: The HCM application enables you to discover Brocade adapters, ports, virtual ports, remote ports, and LUNs using out-of-band discovery only.

10.2.15 Setting up out-of-band discovery for an adapter

When performing out-of-band discovery, you are managing the adapter remotely. The application connects to the agent running on the host server over the IP network and product information is copied back from the Brocade adapter to the server. If you do not configure the application to directly discover the devices, the connections and attached devices might not display correctly.

1. From the Host Connectivity Manager, click **Discovery** → **Setup**. The Setup for Discovery dialog box displays (Figure 10-9).

The image shows a Windows-style dialog box titled "Setup for Host Discovery". It contains several input fields and checkboxes. The "Host Name" field is a dropdown menu with "localhost" selected. The "Port #" field is a text box containing "34568". The "User ID" field is a text box containing "admin". The "Password" field is a text box with ten dots. The "Remember Host" checkbox is checked. The "Polling Frequency (seconds)" field is a text box containing "60". The "Keep Polling" checkbox is checked. At the bottom right are three buttons: "OK", "Cancel", and "Help".

Host Name	localhost
Port #	34568
User ID	admin
Password	••••••••••
Remember Host	<input checked="" type="checkbox"/>
Polling Frequency (seconds)	60
Keep Polling	<input checked="" type="checkbox"/>

Figure 10-9 Setup for Discovery dialog box

2. From the *Host Name* list, select the host name from where you will discover the adapter.

For the first time, the Host Name list will contain only the Local host. You must specify the Hostname or the IP address for discovering the remote servers. Only previously-discovered servers are available in the Host Name list.

3. Type the port number in the *Port Number* text box. The default is 34568.
4. Type in the user ID and password that will authenticate the SAN product with the agent. The default user ID and password are *admin/password*.

Change the agent password on the host for security reasons.

Tip: Click the **Remember Host** check box if you do not want to type it in each time you set up discovery.

5. In the *Polling Frequency (Seconds)* text box, specify the value for how frequently the application has to poll for newly discovered devices.

All parameters related to the adapters that are installed in that server are refreshed each time the poll occurs.

Polling: If the Keep Polling check box is checked, polling occurs after the specified polling interval. If the check box is not checked, polling stops.

6. Click **OK**.

10.2.16 Logging off HCM

End the HCM session using one of the following methods:

- ▶ From the Host Connectivity Manager, click **File** → **Exit**.
- ▶ Click the **X** in the upper-right corner of the HCM window to close it.

10.3 Host configuration

In the topics that follow we discuss host configuration aspects.

10.3.1 Host security authentication

Use the HCM GUI or the command line utility to display the authentication settings and status. There are five well-known DH groups; however, only DH-CHAP group 0, called NULL DH, is supported in this release.

Solaris: Security authentication is not supported on Solaris platforms.

10.3.2 Configuring security authentication using the GUI

You can access the Fibre Channel Security Protocol Configuration dialog box by selecting the Host, an HBA, or an HBA port from the device tree.

Solaris: FC-SP is not available for Solaris platforms.

1. Select the appropriate device based on how you want to configure security authentication:
 - From the host level, select the host from the device tree.
 - From the HBA level, select the adapter from the device tree.
 - From an HBA port, select a port from the device tree.
2. Select **Configure** → **FC-SP** from the main menu, or perform the appropriate following step to open the security authentication dialog box:
 - From the host level, right-click the host and select **FC-SP** from the list.
The Fibre Channel Security Protocol Configuration (host level) dialog box displays
 - From the adapter level, right-click the adapter and select FC-SP from the list.

The Fibre Channel Security Protocol Configuration (adapter level) dialog box displays. This dialog box is identical to the Fibre Channel Security Protocol Configuration (host level) dialog box.

- From the adapter port level, right click a port and select **FC-SP** → **Authentication** from the list. The port level Fibre Channel Security Protocol Configuration dialog box displays, as shown in Figure 10-10.

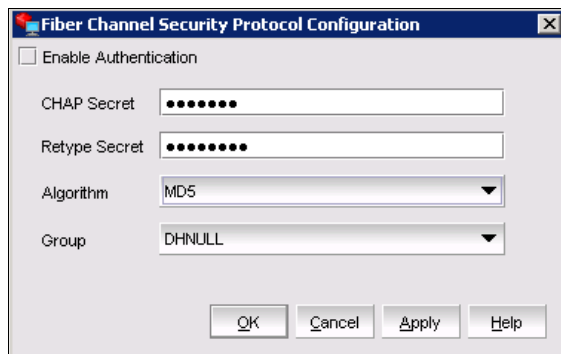


Figure 10-10 Port level Fibre Channel Security Protocol Configuration dialog box

3. Configure the following parameters on the Port Security Authentication tab:
 - a. Select the **Enable Authentication** check box to enable or disable the authentication policy.

If authentication is enabled, the port attempts to negotiate with the switch. If the switch does not participate in the authentication process, the port skips the authentication process.
 - b. Type and retype the secret.

The maximum length of the secret is 63 bytes. The default secret for each interface is its port world wide name (PWWN) without the colons; for example, 0102030405060708.
 - c. Select the algorithm type from the list:
 - MD5: A hashing algorithm that verifies a message's integrity using Message Digest version 5.
 - SHA1: A secure hashing algorithm that computes a 160-bit message digest for a data file that is provided as input.
 - MD5SH1: Similar to the MD5 hashing algorithm, but used for DH-CHAP authentication
 - SHA1MD5: Similar to the SHA1 hashing algorithm, but used for DH-CHAP authentication

- d. Select DHNULL as the group value (this is the only group that is supported).
4. Click **Apply** to apply the changes.
5. Click **OK** to save the changes and close the dialog box.

10.3.3 Configuring security authentication using the CLI

The following commands are used to configure and display the security authentication of ports.

1. Check the port state of the Brocade adapter as shown in Example 10-2.

Example 10-2 Port state check

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu port --list
```

Port#	Type	PWWN/MAC	FC Addr/ Eth dev	Media	State	Spd
<hr/>						
1/0	fc	10:00:00:05:1e:0c:1c:cc	10c900	sw	Linkup	8G
1/1	fc	10:00:00:05:1e:0c:1c:cd	10dd00	sw	Linkup	8G
2/0	cee	00:05:1e:af:08:a0	--	sw	CEE Linkup	10G
	fcoe	10:00:00:05:1e:af:08:a0	051701		Linkup	
	eth	00:05:1e:af:08:a2	LAC# 5		Linkup	
2/1	cee	00:05:1e:af:08:a1	--	sw	CEE Linkup	10G
	fcoe	10:00:00:05:1e:af:08:a1	051401		Linkup	
	eth	00:05:1e:af:08:a3	LAC# 6		Linkup	

2. Set the authentication algorithm for the port as shown in Example 10-3 where we set the algorithm for port 1/0.

Example 10-3 Authentication algorithm for Brocade adapter port

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu auth --algo 1/0 md5
Authentication algorithm set
```

3. Enable the authentication secret as shown in Example 10-4 for port 1/0.

Example 10-4 Set the auth secret

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu auth --secret 1/0
"sec2ibmsw"
Successfully set the auth secret
```

4. Display the authentication settings. Example 10-5 displays the Auth policy state for port 1/0.

Example 10-5 Authentication state for HBA port 1/0

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu auth --show 1/0
```

Port	Port Status	Hash Type	Group Type	Auth Status
1/0	Linkup	MD5	DH-NULL	uninit

10.3.4 Buffer credits

Buffer-to-buffer credit flow control is implemented to limit the amount of data a port sends, based on the number and size of the frames sent from that port. This scheme allows Fibre Channel to be self-throttling, thereby allowing it to establish a reliable connection without the need to accommodate dropped frames due to congestion. Buffer credit limits between each device and the fabric are communicated at the time of fabric login. One buffer credit allows a device to send one frame of data (typically 1 or 2 KB). Buffer credits cannot be configured on an adapter.

The default BB Credit is 1. The baseline for the calculation is one credit per kilometer at 2 Gbps. This yields the following values for 10 km:

- ▶ 5 credits per port at 1 Gbps
- ▶ 10 credits per port at 2 Gbps
- ▶ 20 credits per port at 4 Gbps
- ▶ 40 credits per port at 8 Gbps

See *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116 for more information about buffer credits.

10.3.5 Basic port configuration

For each port, you can configure the following parameters using the Basic Port Configuration dialog box, the Command line Utility or both. Table 10-3 lists the features and configuration options.

Table 10-3 Basic port configuration options

Port configuration parameter	Configurable using the GUI	Configurable using the CLI	For more Information
Port logging level ¹	Yes	Yes	10.3.7, “Port logging level” on page 423
Configure speed (HBA only)	Yes	Yes	10.3.8, “Port speed” on page 426
Frame data field size	Yes	Yes	10.3.9, “Frame data field size” on page 428
Persistent Binding Note: The persistent binding option is available on Windows platforms only.	Yes	Yes	10.3.10, “Persistent binding” on page 428
QoS (HBA only)	Yes	Yes	10.3.11, “QoS (HBA only)” on page 430
Path Time Out	Yes	Yes	10.3.12, “Path Time Out” on page 433
Target Rate Limiting	Yes	Yes	10.3.13, “Target rate limiting” on page 435
¹ If an Ethernet port is selected, the Basic Port Configuration dialog displays an additional feature called “Eth Logging Level.”			

10.3.6 Opening the Basic Port Configuration dialog box

There are slight changes in HCM’s basic port configuration dialog, depending on the operating system. You can access the Basic Port Configuration dialog box by selecting the Host, an HBA, an HBA port, a CNA, or an FCoE port from the device tree.

Follow these steps:

1. Select a device from the device tree.
2. Select **Configure** → **Basic Port Configuration** from the main menu. The Basic Port Configuration dialog box displays (Figure 10-11).

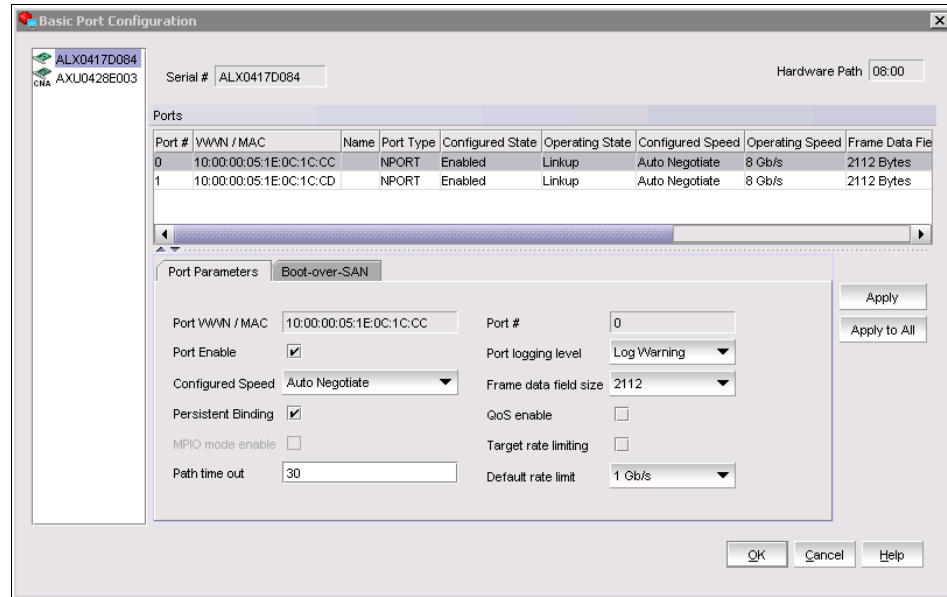


Figure 10-11 Basic Port Configuration dialog box - Windows, Linux, and VmWare

10.3.7 Port logging level

The number of messages logged by the host depends on the predetermined logging level. Although the adapter might generate many messages, only certain types of messages are logged based on the specified logging level.

Configuring the port logging level using the GUI

Follow these steps:

1. Select **Configure** → **Basic Port Configuration** from the Host Connectivity Manager. The Basic Port Configuration dialog box displays.
2. Select a value from the *Port Logging Level* list as shown in Figure 10-12. Supported values are Log Invalid, Log Critical, Log Error, Log Warning, and Log Info.




Figure 10-12 Adapter Port Logging level

3. Click **Apply** to apply the changes.
4. Click **OK** to save the changes and close the window.

Configuring the port logging level using the CLI

Follow these steps:

1. Check the state of the port logging level (Example 10-6)

Example 10-6 Display log level for Brocade adapter port

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu log --level 1/0
FW log level is Warning
HAL log level is Warning
FCS log level is Warning
DRV log level is Warning
AEN log level is Warning
```

2. Change the log level (Example 10-7)

Example 10-7 Change log level of Brocade adapter port

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu log --level 1/0
info
Log level set to Info
```


3. Display the changed the log level state (Example 10-8).

Example 10-8 Displaying the changed log level state

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu log --level 1/0
FW log level is Info
HAL log level is Info
FCS log level is Info
DRV log level is Info
AEN log level is Info
```

Ethernet logging level

The number of messages logged by the host depends on the predetermined logging level. Although the Ethernet port might generate many messages, only certain types of messages are logged based on the specified logging level.

Configuring the Ethernet logging level using the GUI

Follow these steps:

1. Select an Ethernet port from the Host Connectivity Manager device tree.
2. Select **Configure** → **Port Configuration** → **Basic** from the Host Connectivity Manager, the Basic Port Configuration dialog box displays.
3. Select the Eth Parameter tab and select a value from the Eth Logging Level list on the Eth Parameters Tab (Figure 10-13).

Supported values are Log Invalid, Log Critical, Log Error, Log Warning, and Log Info.




Figure 10-13 Basic Port Configuration Eth Parameters Tab

4. Click **Apply** to apply the changes.
5. Click **OK** to save the changes and close the window.

Configuring the Ethernet logging level using the CLI

1. Check the present state of the ethloglevel (Example 10-9).

Example 10-9 State of ethlog level

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu ethlog --level 2/0
Log level is Info
```

2. Change the state of the ethloglevel (Example 10-10).

Example 10-10 Changing ethloglevel

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu ethlog --level 2/0
warning
Log level set to Warning
```

10.3.8 Port speed

Port speed is the maximum amount of data that can pass through the port at a given second. The unit of measurement is in gigabits per second (Gbps).

Attention: For the 10 Gbps CNA, the only option is auto-negotiate.

Configuring the port speed using the GUI

Follow these steps:

1. Select **Configure** → **Basic Port Configuration** from the Host Connectivity Manager.

The Basic Port Configuration dialog box displays.

2. Select a value from the Configured Speed list.

Speed options for the HBA are 1 Gbps, 2 Gbps, 4 Gbps, and 8 Gbps. The available speed options depend on the HBA's speed and the port's SFP. Auto-negotiate is the preferable setting and it is the default as shown in Figure 10-14.

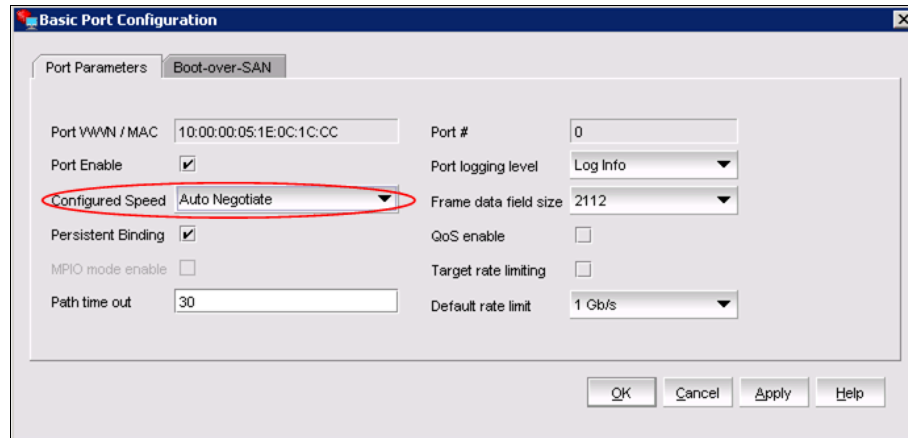


Figure 10-14 HBA Port config indicating speed

3. Click **Apply** to apply the changes.

A port disable/enable configuration dialog displays, confirming the configured speed, which will take effect when the port is disabled or enabled (Figure 10-15).




Figure 10-15 Speed Change Confirmation Dialog box

4. Click **Yes** to continue, or **No** to cancel the operation.
5. If the port can be disabled and enabled, click **Yes** and it will complete the Speed change and indicate the completion state (Figure 10-16).




Figure 10-16 Disable and enable of HBA port after speed change

6. Click **OK** to close the window.

Configuring the port speed using the CLI

The **bcu port** command is used to set the speed from the CLI (Example 10-11).

Example 10-11 Speed setting of HBA port from CLI

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu port --speed 1/0 auto  
Setting will be enforced after port --disable and --enable
```

10.3.9 Frame data field size

Buffer credits determine the maximum amount of frame data. If the number of buffer credits is not large enough to handle the link distance and speed, performance can be severely limited.

See “Buffer credits” on page 421 for information about buffer credits.

Specifying the maximum frame size using the GUI

Follow these steps:

1. Select **Configure** → **Basic Port Configuration** from the Host Connectivity Manager, and the Basic Port Configuration dialog box displays (Figure 10-17 on page 429).
2. Select the frame size from the **Frame Data Field Size** list. Options include 512, 1024, 2048, and 2112 Mbps. The default value is 2112.
3. Click **Apply** to apply the change.
4. Click **OK** to close the window.

Specifying the maximum frame size using the CLI

The command **bcu port --dfsize** is used to Set the maximum data frame size using CLI (Example 10-12).

Example 10-12 Specifying the maximum frame size

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu port --dfsize 1/0 2112  
Setting will be enforced after port --disable and --enable
```

10.3.10 Persistent binding

Persistent binding enables you to permanently assign a system SCSI target ID to a specific FC device. Persistent binding can be achieved by binding to world wide port name (WWPN), world wide node name (WWNN), or device ID (DID).

You can access the Persistent Binding dialog box by selecting the Host, an HBA, a CNA, a CEE port, or an FCoE port from the device tree.

Enabling and disabling persistent binding using the GUI

Persistent binding can be enabled or disabled from the HCM GUI using the following steps:

1. Launch the Basic Port Configuration dialog at the port level.
2. Check or uncheck the persistent binding check box in the Basic Port Configuration dialog (Figure 10-17).

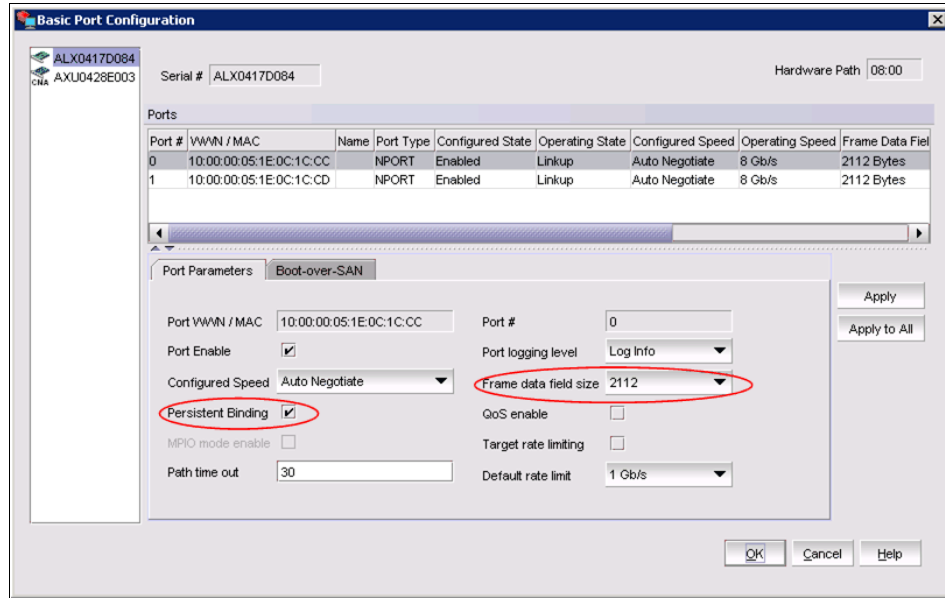


Figure 10-17 Basic Port Configuration dialog box, Frame Size and Persistent Binding

Windows: The Persistent Binding option is available only for Windows hosts.

3. Click **Apply** to apply the change.
4. Click **OK** to close the window.

Persistent binding using the CLI

The `bcu pbind` and `bcu drvconf` commands are used to list and configure persistent binding. The `bcu drvconf --pbind_enable` option is used to enable persistent binding from the CLI. This option enables persistent binding in all Brocade adapters in host driver level and have no option to enable or disable for individual adapters as the HCM GUI have.

Example 10-13 shows the CLI options for enabling persistent binding.

Example 10-13 Persistent binding with CLI

```
C:\Users\Administrator>bcu pbind --list 2/1
Persistent mapping status: Disabled
pbind map:
local port pwn:      10:00:00:05:1e:af:08:a1
remote port pwn:      20:36:00:a0:b8:47:39:b0
Bus:                  0
Target:               4

local port pwn:      10:00:00:05:1e:af:08:a1
remote port pwn:      20:37:00:a0:b8:47:39:b0
Bus:                  0
Target:               5

C:\Users\Administrator>bcu drvconf --key pbind_enable --val 1
Successfully set value = 1 for key = pbind_enable

C:\Users\Administrator>bcu pbind --list 2/1
Persistent mapping status: Enabled
pbind map:
local port pwn:      10:00:00:05:1e:af:08:a1
remote port pwn:      20:36:00:a0:b8:47:39:b0
Bus:                  0
Target:               4

local port pwn:      10:00:00:05:1e:af:08:a1
remote port pwn:      20:37:00:a0:b8:47:39:b0
Bus:                  0
Target:               5
```

10.3.11 QoS (HBA only)

QoS: The QoS feature is not supported on the converged network adapter (CNA).

The QoS feature is not supported on the converged network adapter (CNA). Quality of Service (QoS) works in conjunction with the QoS feature on switch F_Ports. The Fabric operating system (FOS) provides a mechanism to assign traffic priority (high, medium, or low) for a given source and destination traffic flow. By default, all flows are marked as medium.

This feature is supported only on 8 Gbps HBA ports installed on specific switch models that use Fabric OS v6.2 and later. The following licenses must be installed on the switch connected to each HBA port (edge switch):

- ▶ Adaptive Networking (AN) license
- ▶ Server Application Optimization (SAO) license

To determine if these licenses are installed on the connected switch, execute the Fabric OS **show licenses** command. For more information about QoS, see *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116.

Configuring QoS on the switch side

On the switch side, you can create QoS zones using the PWWNs that correspond to devices in a source/destination traffic flow. You need a Server Application Optimization (SAO) license installed on the switch to enable QoS. In addition, an Adaptive Network (AN) license is required on the switch to enable QoS on the switch ports.

You enable or disable QoS settings on ports with the **portCfgQos** command. For more information about configuring QoS, see *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116.

Configuring QoS on the HBA side

There are three possible QoS states:

- ▶ Enabled, online: QoS is established with the switch.
- ▶ Enabled, offline: QoS negotiation failed and QoS was not established with the switch. Possible reasons for failure might be that the license is not installed on the switch or QoS is not enabled on the port.
- ▶ Disabled.

Tip: You must first disable the port and enable the port before QoS changes.

Follow these steps:

1. In the **Tree View**, select the adapter for configuration.
2. Select **Configure** → **Basic Port Configuration** from the Host Connectivity Manager.

The Basic Port Configuration dialog box displays (Figure 10-18).

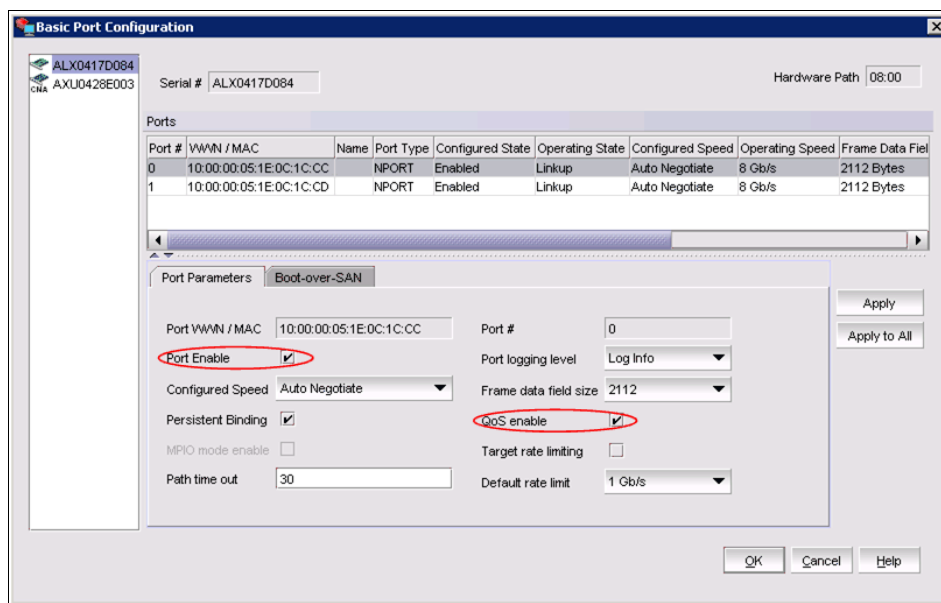


Figure 10-18 Basic Port Configuration

3. Disable the port by unchecking the **Port enable** check box.
A confirmation dialog box displays (Figure 10-19).




Figure 10-19 Port Disable confirmation dialog box

4. Click **OK**.
5. Click the **QoS enable** check box.
6. Click the **Port enable** check box to re-enable the port.
7. Click **OK** to apply the changes.

A confirmation box displays (Figure 10-20).




Figure 10-20 Update confirmation dialog box

8. Click **Yes**, which will perform a port disable and then an enable, and will indicate success with a confirmation, Click **OK**.
9. Click **OK** in the *Basic Port Configuration* window to return to the HCM main window.

Enabling and disabling QOS using the CLI

The **bcu port** is the command used to enable or disable QOS. First enter the **bcu port --disable <port_id>** command, followed by the **bcu port --enable <port_id>** command, before the **bcu qos --enable** or **bcu qos --disable** commands take effect. The sequence is listed in Example 10-14.

Example 10-14 Enabling QOS with the CLI

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu port --disable 1/0
port disabled
```

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu port --enable 1/0
port enabled
```

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu qos --enable 1/0
qos for port id 1/0 enabled
Setting will be applied after port --disable and port --enable
```

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu port --disable 1/0
port disabled
```

```
C:\Program Files\BROCADE\Adapter\driver\util>bcu port --enable 1/0
port enabled
```

10.3.12 Path Time Out

The Path Time Out value is used to handle path failover when the timeout value is reached. The value range is 0 to 60. If The Path time Out value is set to 0, it forces an immediate failover.

Specifying Path Time Out using the GUI

Follow these steps:

1. Select **Configure** → **Basic Port Configuration** from the Host Connectivity Manager.

The Basic Port Configuration dialog box displays (Figure 10-21).

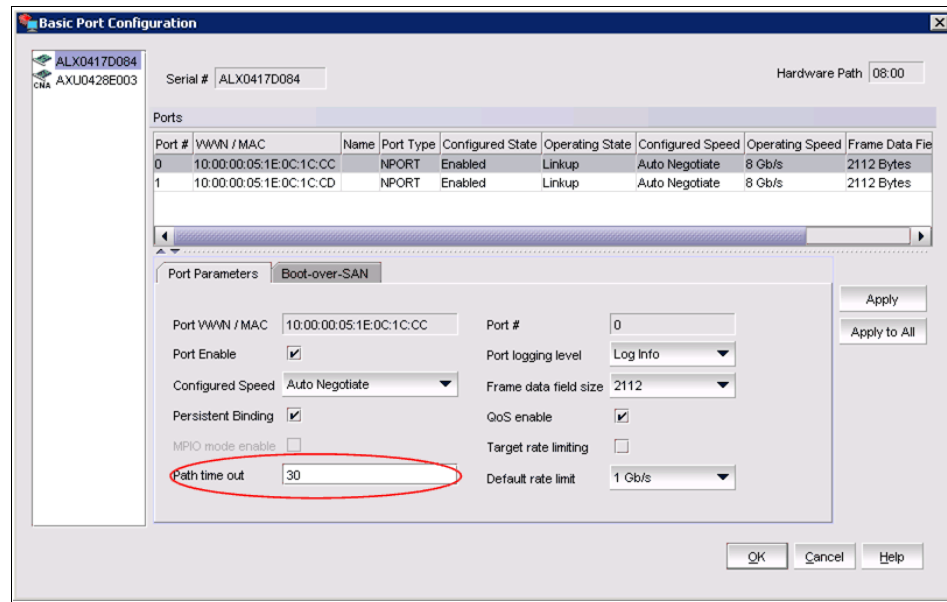


Figure 10-21 Path Time Out

2. Type a value in the *Path Time Out* text field. A timeout value of 30 is the default.
3. Click **OK** to close the window.

Specifying the Path Time Out using the CLI

The Path Time Out value is set by using the command **bcu fcpim** as shown in Example 10-15.

Example 10-15 Path Time Out value with CLI

```
C:\Users\Administrator>bcu fcpim --pathtov 1/0 20
path timeout is set to 20
```

10.3.13 Target rate limiting

The target rate limiting feature is used to minimize congestion at the adapter port caused by a slow drain device operating in the fabric at a slower speed. A remote port's operating speed is determined from the fabric. Traffic destined to the remote port is limited to its current operating speed.

Enabling and disabling rate limiting using the GUI

Target rate limiting is supported only when the adapter port is connected to the fabric. Therefore, target rate limiting is not supported when the port is directly connected with another device. Follow these steps:

1. Select **Configure** → **Basic Port Configuration** from the Host Connectivity Manager; the Basic Port Configuration dialog box displays (Figure 10-22).

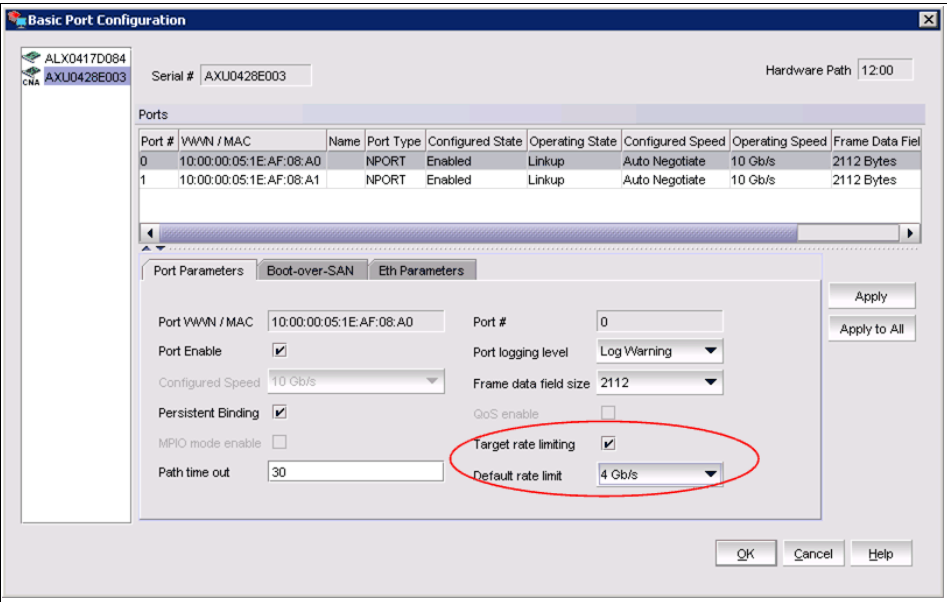


Figure 10-22 Target Rate Limiting

2. Enable the Target Rate Limiting feature by clicking the corresponding check box.

Attention: Target Rate Limiting and QoS cannot be enabled at same time for an adapter.

3. Select the default rate limit from the list. Options include 1 Gbps, 2 Gbps, and 4 Gbps; the default is 2 Gbps.

4. Click **OK** to close the window.

A disable port/enable port warning message displays (Figure 10-23).




Figure 10-23 Disable Port Warning Message

4. Click **Yes** to continue.

A disable port/enable port completion message displays (Figure 10-24).




Figure 10-24 Disable/Enable port completion message

5. Click **OK**.

Enabling and disabling rate limiting using the CLI

The **bcu ratelim** command is used to enable the Target rate limit.

Example 10-16 shows the configuration sequence where we change the target rate limit for the CNA port 2/1 from 1G to 4G.

Example 10-16 Target Rate Limit with CLI

```
C:\Users\Administrator>bcu port --disable 2/1
port disabled

C:\Users\Administrator>bcu port --enable 2/1
port enabled

C:\Users\Administrator>bcu ratelim --enable 2/1
ratelim for port id 2/1 enabled
Setting will be enforced after port --disable and --enable

C:\Users\Administrator>bcu ratelim --query 2/1
Target Rate Limiting: enabled
Default TRL Speed is: 1G
```

```
C:\Users\Administrator>bcu ratelim --defspeed 2/1 4  
Setting will be enforced after port --disable and --enable
```

```
C:\Users\Administrator>bcu ratelim --query 2/1  
Target Rate Limiting:   enabled  
Default TRL Speed is:   4G
```

```
C:\Users\Administrator>bcu port --disable 2/1  
port disabled
```

```
C:\Users\Administrator>bcu port --enable 2/1  
port enabled
```

10.3.14 Boot over SAN

The boot over SAN feature allows you to target remote boot devices (LUNs on SAN storage arrays) from which to boot the host system. When the host's operating system and adapter driver are installed on the remote device, the adapter BIOS and user-configurable boot instructions stored in adapter flash memory allow the host to boot from the device.

SAN boot: Various operating systems require you to follow specific guidelines to enable servers to boot from a SAN. Understanding these requirements is key to a successful deployment of a boot over SAN environment.

Boot LUNs are identified to adapter ports using the BIOS Configuration Utility and CLI commands. These utilities also allow you to enable or disable BIOS for booting the host system over SAN, set boot options, and set the port speed. For instructions, see the *Brocade Adapters Installation and Reference Manual*, 53-1001254-05, available at this website:

<http://www.brocade.com/services-support/drivers-downloads/HBA/index.page>

After you have configured boot devices using the BIOS Configuration Utility, you can enable or disable BIOS for boot over SAN, set boot options, and set port speed using the HCM GUI. All configuration is stored in flash memory.

10.3.15 Configuring Boot over SAN

The boot-LUN table lists the vendor information, LUN capacity, and whether the LUNs are accessible. These fields are not editable.

You can access the *Boot over SAN* dialog box by selecting the Host, an HBA, or CNA from the device tree.

Attention: Boot over SAN configuration using the *Basic Port Configuration* dialog box is enabled on all platforms if the agent version is 1.1 or higher.

1. Select **Configure** → **Basic Port Configuration** from the Host Connectivity Manager.

The Basic Port Configuration dialog box displays.

2. Click the **Boot-over-SAN** tab. The Boot-over-SAN dialog box displays (Figure 10-25).

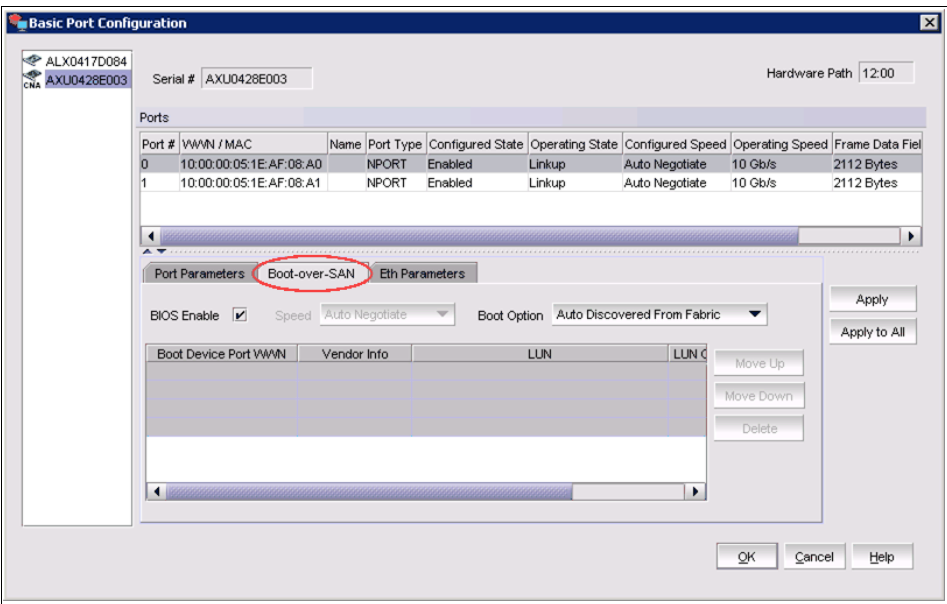


Figure 10-25 Boot over SAN View

3. Click the **BIOS Enable** check box to enable Boot Over SAN.
4. From the *Boot Option* list, select one of the following choices:
 - **Auto Discovered from Fabric:** Enables Boot over SAN using boot LUN information stored in the fabric. This is the default setting.

- **First Visible LUN:** Enables Boot over SAN from the first discovered LUN in the SAN.
 - **User Configured LUNs:** Allows the user to select and prioritize the remote target and LUN for booting over SAN.
5. Select the Boot Device Port WWN row in the table, then click the up and down arrows to move the row up or down in the table or use “move up” and “move down” buttons in the side of the rows. The host will attempt to boot from the first LUN in the table, and then move on to succeeding LUNs.
 - You can delete a row using the Delete button.
 - For *User Configured LUNs* click the **Boot Device Port WWN** and **LUN** fields to manually enter boot LUNs to the table. These LUNs must be visible to the adapter to be accessible as boot LUNs.
 6. Click **OK**.

The Vendor Info, LUN Capacity, and Accessible status that corresponds to the selected boot device and LUN displays automatically.

10.3.16 Boot code image upload

You can upload a boot code image on the local host or on an HBA. The boot-over-SAN feature is not supported on the converged network adapter (CNA). Follow these steps to upload the latest boot code using the HCM GUI.

Solaris: On Solaris systems, the Boot Code Image Upload menu is disabled if the host does not have a Fibre Channel HBA card or if the driver version is lower than 1.1.0.7 (the version must be 1.1.0.7 or higher for Solaris).

10.3.17 Updating the boot code using the GUI

1. Download the boot code (brocade_adapter_boot_fw_v1-1-0-6) from the following website, to a folder on your local drive:
<http://www.brocade.com/services-support/drivers-downloads/HBA/index.page>
2. Launch HCM.
3. Right-click a host or adapter from the device tree and select **Upload Boot Code Image** from the list.
 - Right-clicking a host uploads the boot code image to all adapters that are installed on the host.

- Right-clicking an adapter uploads the boot code image to the selected adapter only.

The Boot Code Image Upload dialog displays (Figure 10-26).




Figure 10-26 Boot Code Image Upload dialog box

4. Click the **Browse** button and navigate to the location of the boot code image.
5. Select the boot code image and click **Open**.

The selected file uploads. If an error occurs during the uploading process, an error message displays.

10.3.18 Virtual port configuration

Virtual ports (V_Ports) appear to the hosts as physical ports in the data network. One or more virtual ports are assigned to each host, and a host can access storage at a virtual port only if the virtual port has been assigned to the host.

V_Port: You cannot create a V_Port that already exists in the Names dialog. If you need to re-create a V_Port that has been deleted through an interface other than the currently managing HCM or the V_Ports deleted on Linux servers reboot, you must first manually remove the V_Port's WWN from the Names dialog box in HCM. If you do not manually remove the V_Port from HCM, an error message displays that the V_Port already exists. See ““Removing a name entry” on page 449 for instructions.

10.3.19 Creating a virtual port

You create virtual ports on HBA ports and FCoE ports only; virtual ports are not supported on an adapter. Virtual ports are not supported for VMware and Solaris agents. Follow these steps:

1. Select a physical HBA port or an FCoE port from the device tree.
2. Select **Configure** → **Virtual Port** → **Create** from the main menu.

or

Right-click the physical port and select **Virtual Port** → **Create** from the list.

The Virtual Port Creation dialog box displays as shown in Figure 10-27.

The following fields are system-generated:

- Physical port world wide name.
- Virtual port world wide name: This WWN must be unique.
- Virtual node world wide name: The system returns the default node WWN, which is the physical port node WWN.

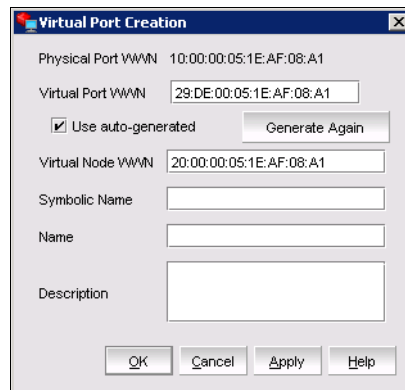


Figure 10-27 Virtual Port Creation dialog box

Attention: By default, the Use auto-generated check box is selected and the Generate Again button is enabled. You can still edit the Virtual Port WWN field if Use auto-generated is selected.

3. (Optional). Provide a symbolic name for the virtual port.
4. (Optional). Provide an alias name for the virtual port. By creating an alias, you can assign a familiar name to a device or group multiple devices into a single name. This can simplify cumbersome data entry and allows an intuitive naming structure.
5. (Optional). Enter descriptive information about the virtual port into the Description text box.
6. Click **OK** to apply the changes and close the window.

The new virtual port will now display in the Tree View as shown in Figure 10-28.

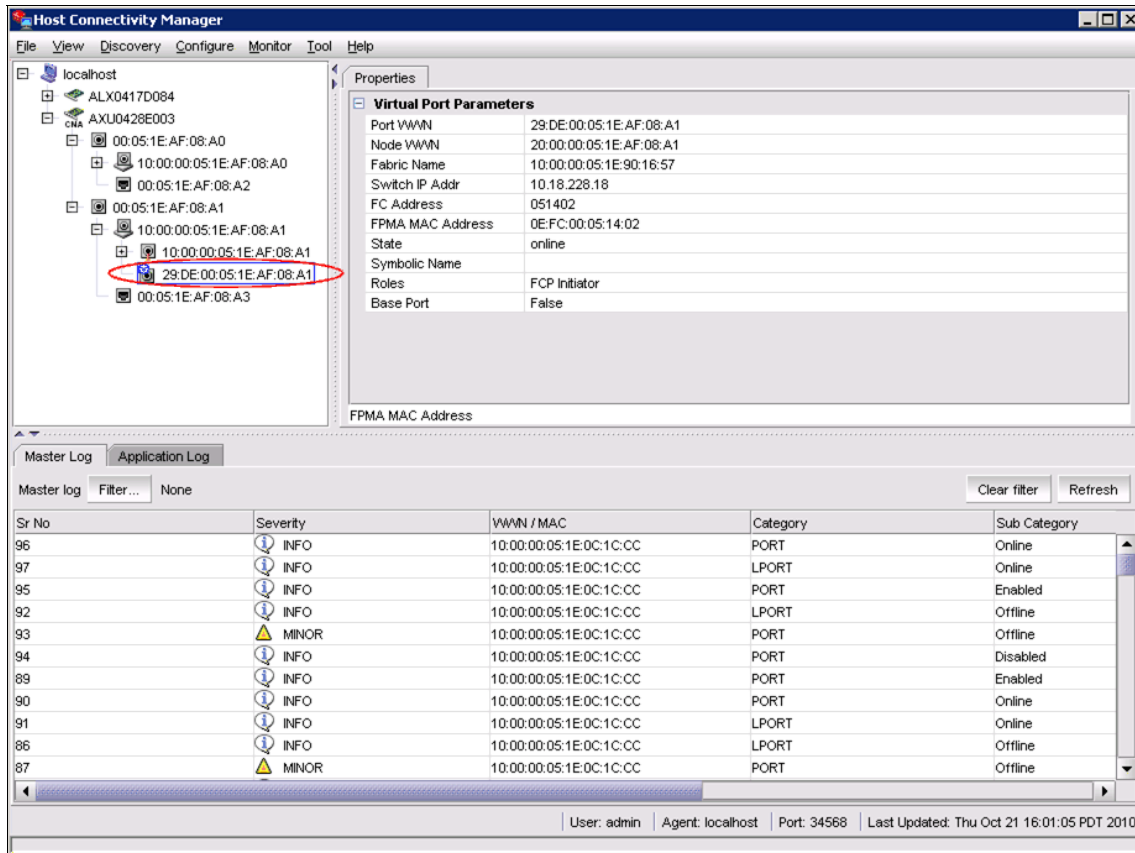


Figure 10-28 Newly Created virtual port on the CNA

10.3.20 Deleting a virtual port

Follow these steps:

1. Select a virtual port from the device tree.
2. Select **Configure** → **Virtual Port** → **Delete** from the main menu.

or

Right-click the virtual port and select **Delete** from the list.

A warning message displays, asking for confirmation (Figure 10-29).




Figure 10-29 Virtual Port Deletion

3. Click **OK** to continue.

10.3.21 HCM logging levels

In this section we discuss considerations regarding logging levels.

Setting the logging level for modules

You can set the logging level for the following modules:

- ▶ Agent communication log, where all messages are exchanged between the HCM GUI application and the HCM agent.
- ▶ HCM debug log, where messages are logged locally.

If you do not set an HCM logging level, then TRACE, which is the most verbose and the default setting, is used.

Configuring the HCM logging level using the GUI

Follow these steps:

1. Select **Configure** → **HCM Logging Levels** from the Host Connectivity Manager.

The Configure HCM Logging Levels dialog box displays (Figure 10-30).




Figure 10-30 Configure HCM Logging Levels dialog box

2. From both the Agent Communication Log and the HCM Debug Log lists, select one of the following choices:
 - Trace, which is the most verbose.
 - Debug

- Info
 - Warning
 - Error
 - Fatal, which is the least verbose.
3. Click **Apply** to apply the change.

10.3.22 Advanced port configuration

You can access the Advanced Port Configuration dialog box by selecting an HBA port or an FCoE port from the device tree.

For each port, you can configure the following parameters using the Advanced Port Configuration dialog box, the Command Line utility, or both. Table 10-4 lists the features and configuration options.

Table 10-4 Features and configuration options

Port Configuration parameter	Configurable using the GUI	Configurable using the CLI
Interrupt Control Coalesce	Yes	Yes
Interrupt Control Latency	Yes	Yes
Interrupt Control Delay	Yes	Yes

10.3.23 Opening the Advanced Port Configuration dialog box

Follow these steps:

1. Select an HBA port or FCoE port from the device tree.
2. From the Host Connectivity Manager, select **Configure** → **Port Configuration** → **Advanced**.

The Advanced Port Configuration dialog box displays (Figure 10-31).

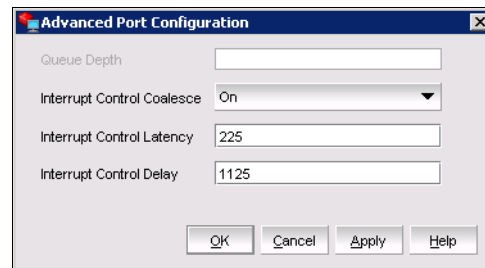


Figure 10-31 Advanced Port Configuration dialog box

Interrupt Control Coalesce

Interrupt control coalescing allows the system to change CPU utilization by varying the number of interrupts generated. Increasing the latency monitor timeout value should result in a lower interrupt count and less CPU utilization, which might result in higher throughput.

Configuring the Interrupt Control Coalesce using the GUI

Follow these steps:

1. Select a port from the device tree.
 - From the Host Connectivity Manager, select **Configure** → **Advanced Port Configuration**.
 - or
 - Right-click a port and select **Port Configuration** → **Advanced**.

The Advanced Port Configuration dialog box displays (Figure 10-32).




Figure 10-32 Advanced Port Configuration dialog box CNA

2. Set the latency and delay values:
 - Select **On** from the Interrupt Control Coalesce list.
 - Specify the latency monitor timeout value in microseconds, if coalesce is set to on. Latency timeout values supported are 0-225 microseconds. Setting the latency time out value to 0 disables the latency monitor time out interrupt.
 - Specify the delay timeout value in microseconds, if coalesce is set to on. Delay timeout values supported are 0-1125 microseconds. Setting the delay timeout value to 0 disables the latency monitor time out interrupt.
3. Click **OK**.

10.3.24 NPIV

N-Port ID Virtualization (NPIV) enables a single Fibre Channel protocol port to appear as multiple, distinct ports. NPIV provides separate port identification within the fabric for each operating system image (partition) behind the port, as if each operating system image had its own unique physical port.

Each NPIV device has a unique virtual port ID (PID), port WWN, and node WWN. The virtual port has the same properties as an N_Port and is therefore capable of registering with all services of the fabric. In other words, multiple virtual devices emulated by NPIV appear no different than regular devices connected to a non-NPIV port. The maximum number of virtual PIDs for an N_Port on a FC switch is 255. For a CEE switch, the maximum number of V_Ports is 64.

NPIV is available at the physical port level or at the virtual fabric level. If virtual fabric ports are detected, then you cannot configure NPIV parameters at the physical port level. If virtual fabric ports are deleted on the switch port side, the NPIV parameters can then be configured at the physical port level. No settings are available for V_Ports from basic port configuration.

Solaris: NPIV is not supported on Solaris platforms.

10.3.25 Name configuration

The Host Connectivity Manager allows you to configure names as a method of providing familiar, simple names to world wide names for adapters, ports, virtual ports, and remote ports in the SAN. (A logical port can be a base port or a virtual port.) Only unique names are allowed.

You can access the *Configure Names* dialog box by selecting an HBA, an HBA port, a Virtual Port, a CNA, or a CEE port from the device tree.

You can perform the following name tasks using either the *Configure Names* dialog or the *Define Names* dialog:

- ▶ Associate a name that represents an adapter, port, virtual port, or remote port. Note the following considerations about names:
 - Among all adapters, two cannot have duplicate names.
 - Among all the ports, two cannot have duplicate names.
 - A port and adapter can have the same name.
 - You cannot associate a name for a storage device.
 - Name changes on remote ports and virtual ports are sent to the *.properties file local to the HCM application but are not sent to the agent.

- ▶ Add a detached WWN and an associated name with Type and operational status as Unknown.
- ▶ Remove or disassociate a name from a WWN.

Define Names: You can launch the Define Names dialog by right-clicking an adapter, port, remote port, or V_Port.

Name validation

Observe the following considerations when you define a name:

- ▶ The name cannot begin with a number.
- ▶ The name cannot begin with an underscore (`_`) or hyphen (`-`), however an underscore or hyphen character is allowed within the name; for example, `name1_name-2`.
- ▶ No special characters are allowed, except for an underscore or hyphen.
- ▶ The maximum length of the name is 15 characters.
- ▶ The maximum length of the description is 80 characters.

Editing the name fields

Only the name, the world wide name (WWN), and the description fields are editable. Depending on the component, the following occurs when you edit the name fields:

- ▶ Name changes on the adapter and ports are sent to the agent and stored in the *HbaAliasdb.properties* file.
- ▶ Name changes on remote ports and virtual ports are sent to the *HbaAliasdb.properties* file local to the HCM application but are not sent to the agent.

1. Select an HBA, an HBA port, a Virtual Port, a CNA or a CEE port from the device tree.
 2. Select **Configure** → **Names** from the Host Connectivity Manager.
- or

Right-click a device from the device tree and select **Define Name**.

The *Configure Names* dialog box displays all the discovered and detached (undiscovered) names (Figure 10-33).

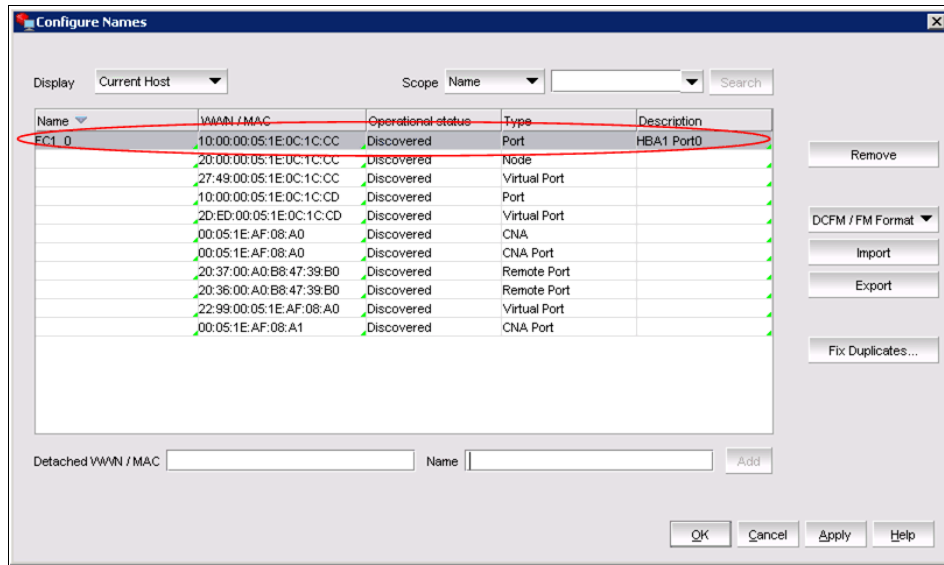


Figure 10-33 Configure Names dialog box

3. Select a row and edit the name, the WWN, and the description, as needed.
4. Click **OK**.

A Name Change Confirmation dialog box displays (Figure 10-34).




Figure 10-34 Name Change Confirmation dialog box

5. Click **Yes**.

Adding name entries

You can add up to 2000 names, which are then stored in the *HbaAliasdb.properties* file. The entries persist during reboot.

The WWN types are as follows:

- ▶ Node
- ▶ Port
- ▶ Remote Port
- ▶ V_Port

- Dual Role (port type that acts as initiator and target)
- Unknown

Follow these steps:

1. Select an HBA, an HBA port, a Virtual Port, a CNA or a CEE port from the device tree.
2. Select **Configure** → **Names** from the Host Connectivity Manager.

The Configure Names Dialog box opens as in Figure 10-35.

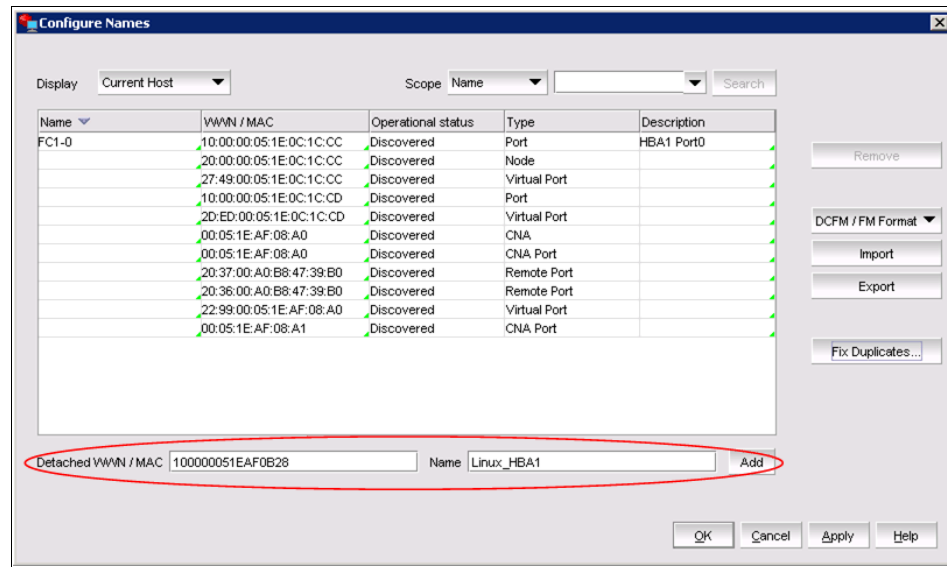


Figure 10-35 Configure names dialog

3. Type a name that represents an adapter, port, or storage device into the **Name** text box.
4. Type a valid WWN that corresponds to the name in the **Detached WWN/MAC** text box.
5. Click **Add**.
6. Click **OK** to close the window.

The new component is added to the Name list.

Removing a name entry

The Remove functionality clears the name and description values of a selected detached WWN:

Follow these steps:

1. Select an HBA, an HBA port, a Virtual Port, a CNA, or a CEE port from the device tree.
2. Select **Configure** → **Names** from the Host Connectivity Manager.

or

Right-click a device from the device tree and select **Define Name**.

The Configure Names dialog box displays all the names available at the host.

3. Select one of the following choices from the Display list:

- Current Host
- All WWNs/MACs
- Only Nodes
- Only Ports
- Only Logical Ports
- Only Virtual Ports
- Only Remote Ports

A list of names for the devices you selected displays (Figure 10-36).

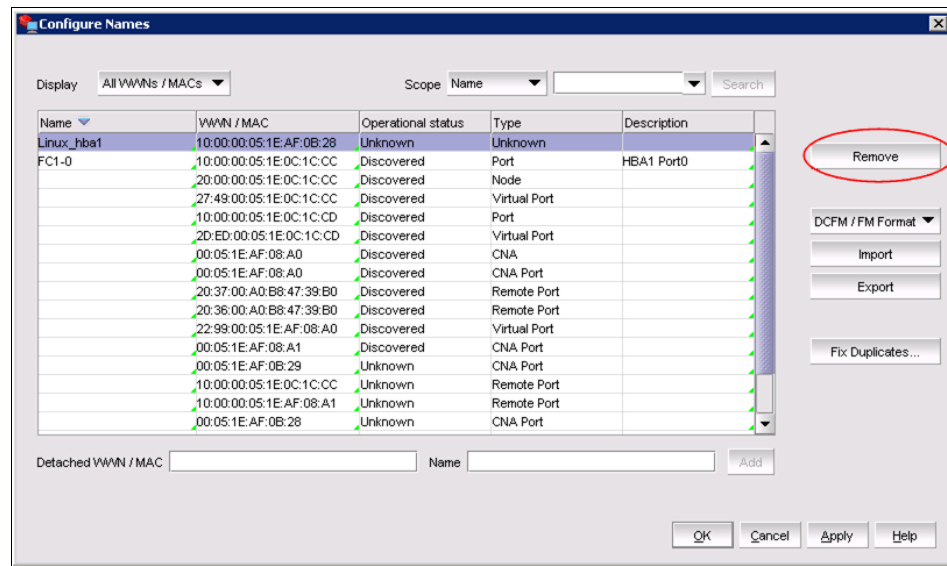


Figure 10-36 Remove a Name Entry

4. Select a device to highlight it and click the **Remove** button to remove the discovered device from the list. The Remove button clears the names of the discovered WWN and the entire row of the detached (undiscovered) WWN.
5. Click **OK** to close the window.

10.3.26 Exporting the properties for a WWN

You can export the properties for a world wide name in .csv, .properties, or .txt file format.

Follow these steps:

1. Select an HBA, an HBA port, a Virtual Port, a CNA or a CEE port from the device tree.
2. Select **Configure** → **Names** from the Host Connectivity Manager.

The Configure Names dialog box displays as in Figure 10-37.

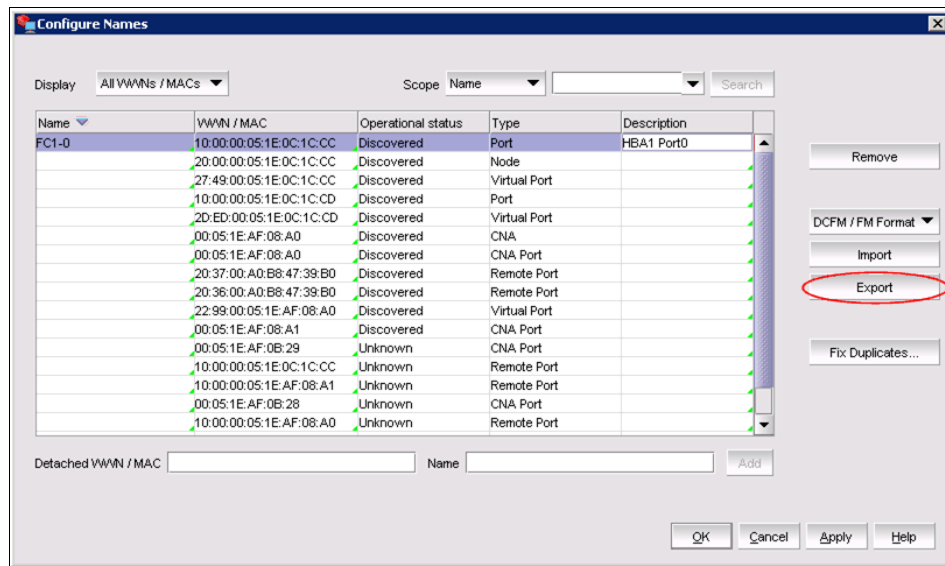


Figure 10-37 Export Properties of a WWN

3. Select one of the following choices from the Display list:
 - Current Host
 - All WWNs/MACs
 - Only Nodes
 - Only Ports
 - Only Logical Ports
 - Only Virtual Ports
 - Only Remote Ports

4. Click the **Export** button.

The Save dialog box displays. You can save the properties file in .txt, .csv, or .properties format (Figure 10-38).

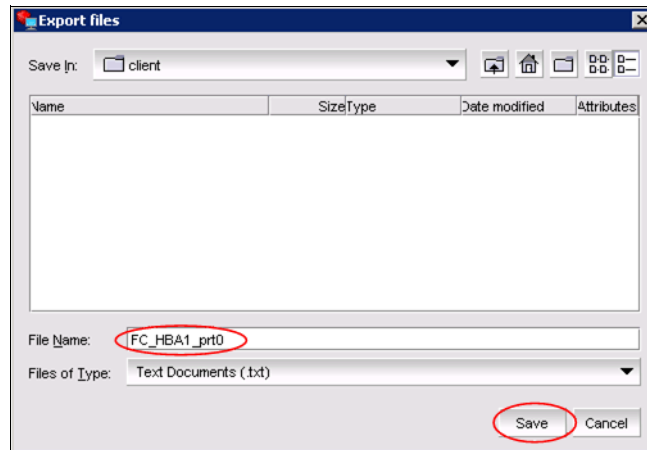


Figure 10-38 Export Save As dialog box

5. Name the file, and click **Save**.
6. Click **OK** to close the window.

10.3.27 Importing the properties for a WWN

1. Select **Configure** → **Names** from the Host Connectivity Manager.
The Configure Names dialog box displays.
2. Select one of the following choices from the Display list:
 - Current Host
 - All WWNs/MACs
 - Only Nodes
 - Only Ports
 - Only Logical Ports
 - Only Virtual Ports
 - Only Remote Ports

3. Click the **Import** button.

The Import dialog box displays (Figure 10-39).

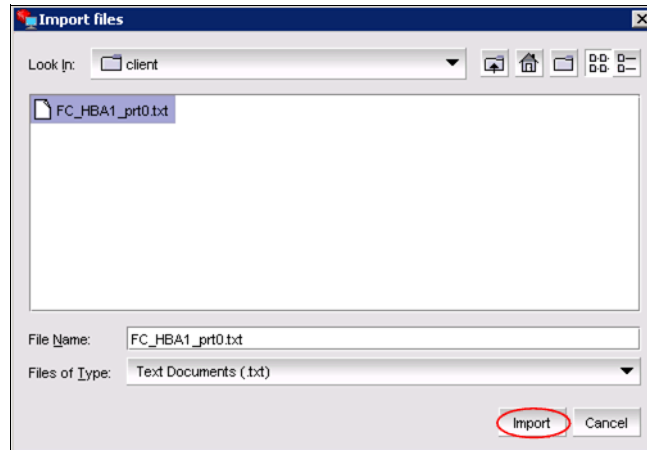


Figure 10-39 Import Properties dialog box

4. Navigate to the location of the .properties, .csv, or .txt file from which you will import properties for the selected device.
5. Click the file, and click **Import**.
6. Click **OK** to close the window.

10.3.28 Importing properties in EFCM format

You can use this procedure to import properties in Enterprise Fabric Connectivity Manager (EFCM) format.

Follow these steps:

1. In the Configure Names dialog box, select **EFCM Format** and then select **Import** (Figure 10-40).

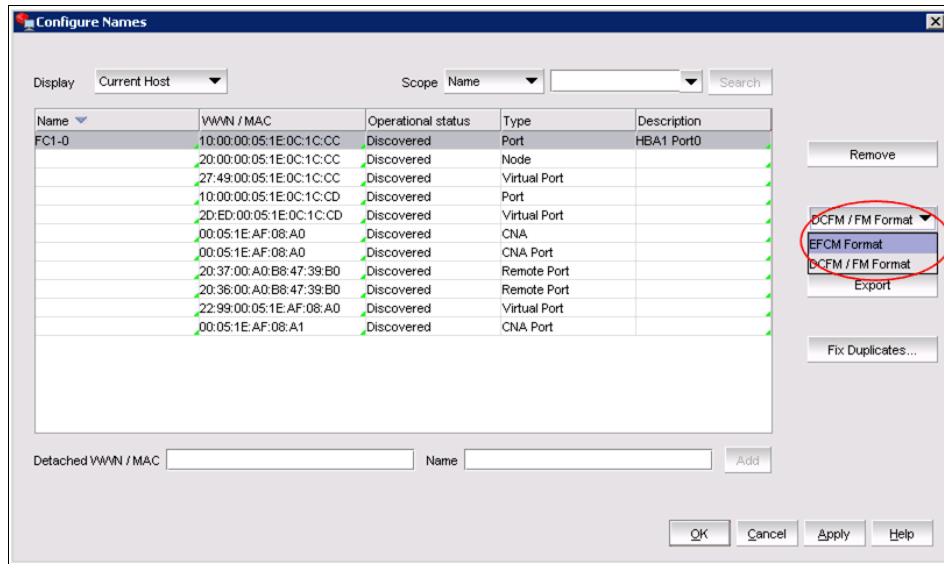


Figure 10-40 Import from ECFM format

The Import dialog box displays (Figure 10-41).

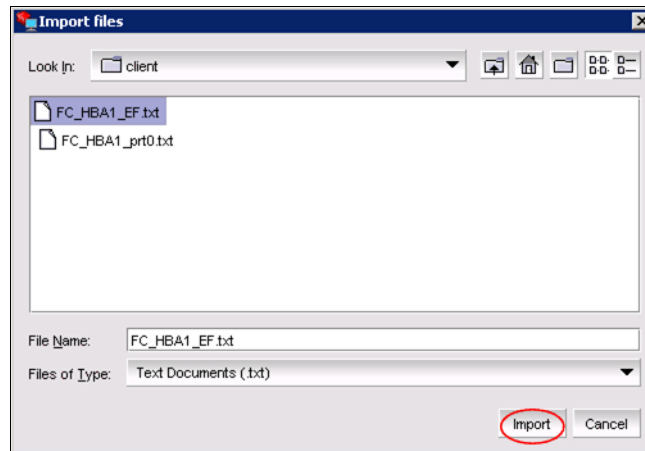


Figure 10-41 Import ECFM format dialog box

2. Navigate to the file and click **Import**.
3. Click **OK**.

Attention: The same procedure can be followed for importing adapter properties in DCFM or Fabric Manager (FM) format. We need to select DCFM/FM format in step 1 instead of ECFM format.

10.3.29 VLAN configuration

A Virtual LAN (VLAN) is a collection of network nodes that share the same broadcast domain regardless of their physical location or connection point to the network. A VLAN serves as a logical workgroup with no other physical barriers and allows users to share information and resources as though located on the same LAN.

Attention: VLAN configuration is a Windows-only feature.

There are three types of VLANS:

- ▶ Regular VLAN:
A regular VLAN is identified using a VLAN ID (with a range of 1-4094, where 0 is used for an untagged VLAN) and a VLAN name.
- ▶ Passthru VLAN:
A Passthru VLAN has VLAN ID 0 and PASSTHRU as its VLAN Name. It can be created or deleted at any time and is treated as a regular VLAN; however, a Passthru VLAN is not editable.
- ▶ Port VLAN (PVID):
You create a Port VLAN using Windows Device Manager. The VLAN ID is assigned when it is created and the VLAN name is PORT VLAN. You cannot create, edit, or delete a Port VLAN using the Host Connectivity Manager (HCM).

VLANS: If a PORT VLAN exists in the VLAN configuration, you cannot perform any add, delete, or edit operations on any VLAN. In addition, you cannot view statistical information about any VLAN.

10.3.30 Adding a VLAN

You can access the VLAN Configuration dialog box by selecting an Ethernet port from the device tree. This procedure provides instructions about how to add a VLAN to an Ethernet port.

Attention: You can create a regular VLAN or a passthru VLAN only if a Port VLAN ID (PVID) does not exist. You cannot name a regular VLAN “PORT LAN” or “Passthru.”

Follow these steps:

- 1. From the Ethernet port level, select an Ethernet port from the device tree.
- 2. Select **Configure** → **VLAN Configuration** from the main menu.

or

Right-click an Ethernet port and select **VLAN Configuration** from the list (Figure 10-42).

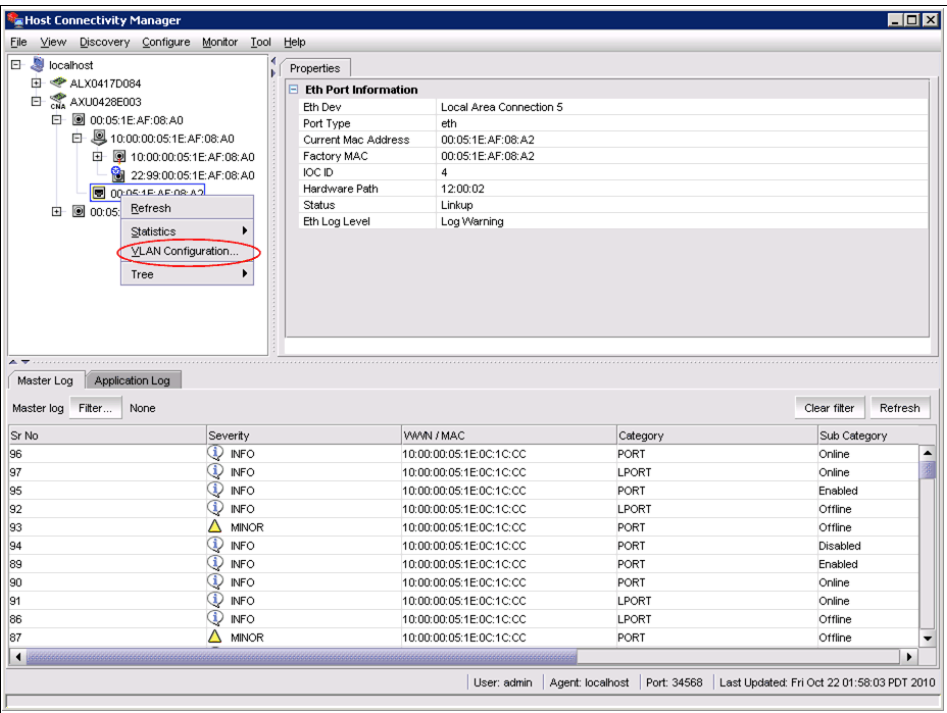


Figure 10-42 Vlan Configuration Selection

The VLAN Configuration dialog displays (Figure 10-43).




Figure 10-43 VLAN Configuration dialog box

3. Click **Add** on the VLAN Configuration dialog (Figure 10-44). The Add VLAN dialog displays.




Figure 10-44 Add VLAN dialog box

4. Enter a VLAN identifier in the **VLAN ID** text box. The range is 1 to 4094.
5. Enter a VLAN name in the **VLAN Name** text box. The VLAN name must not exceed 31 characters.
6. (Optional) Click the **Create Passthru** checkbox to designate the VLAN as a Passthru VLAN.
7. Click **OK**.

10.3.31 Editing a VLAN

You can access the VLAN Configuration dialog box by selecting an Ethernet port from the device tree. This procedure provides instructions about how to edit an existing VLAN. (You cannot edit a PORT VLAN or a Passthru VLAN.)

Follow these steps:

1. From the Ethernet port level in the Tree View, select an Ethernet port from the device tree.
2. Select **Configure** → **VLAN Configuration** from the main menu.

or

Right-click an Ethernet port and select **VLAN Configuration** from the list.

The VLAN Configuration dialog displays (Figure 10-45).

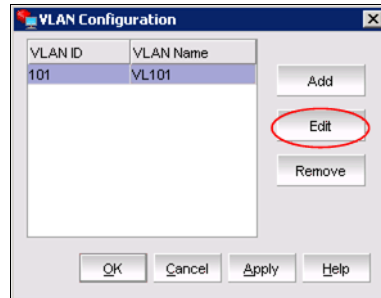


Figure 10-45 Edit VLAN Configuration

3. Select a VLAN from the list.
4. Click **Edit** on the VLAN Configuration dialog.

The Edit VLAN dialog displays (Figure 10-46).




Figure 10-46 Edit VLAN dialog box

5. Type a new name in the **VLAN Name** text box.
6. Click **OK**.

10.3.32 Removing a VLAN

You can access the VLAN Configuration dialog box by selecting an Ethernet port from the device tree. This procedure provides instructions about how to remove an existing VLAN.

1. From the Ethernet port level, select an Ethernet port from the device tree.
2. Select **Configure** → **VLAN** Configuration from the main menu.

or

Right-click an Ethernet port and select **VLAN Configuration** from the list.

The VLAN Configuration dialog displays (Figure 10-47).

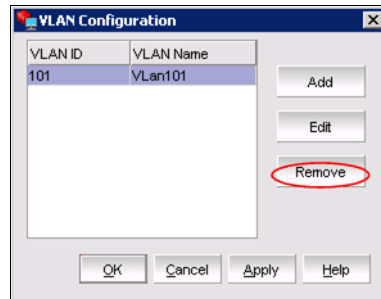


Figure 10-47 Remove VLAN Configuration

3. Select the VLAN for removal
4. Click **Remove** on the VLAN Configuration dialog.

A warning dialog displays (Figure 10-48).




Figure 10-48 Remove VLAN warning dialog box

5. Click **OK** to remove the VLAN from the configuration.

10.4 Monitoring

In the following topics we discuss the various monitoring capabilities.

10.4.1 Performance monitoring

The Host Connectivity Manager (HCM) Port Statistics window enables you to monitor the performance of the CNA and the traffic between the CNA and the LUNs. You can use the information to isolate and troubleshoot areas that impact application performance.

The components listed in Table 10-5 display statistics when the FCoE port node is selected.

Table 10-5 Statistics monitored by component

Component	Statistics monitored
Local Host	► Port, Teaming
HBA	► Port
HBA Port	<ul style="list-style-type: none"> ► Port ► FCP IM Module ► Fabric ► IOC ► QoS
CNA	► Port
CEE Port	<ul style="list-style-type: none"> ► Port ► CEE ► FCP IM Module
Ethernet Port	<ul style="list-style-type: none"> ► Eth IOC ► VLAN ► Eth
FCoE Port	<ul style="list-style-type: none"> ► Fabric ► IOC
Logical Port and base port	► Logical Port
Virtual Port	<ul style="list-style-type: none"> ► Logical Port ► Virtual Port
Device (HCM does not have a statistics monitor for LUNs)	► Remote Port → Target IFCP IM

10.4.2 Polling frequency rate

The faster the polling rate, the more quickly the HCM GUI receives indications from the host. However, faster polling rates consume more of your system's CPU and network resources and can therefore slow the system.

Controlling the polling frequency rate: To control port statistics polling, do one of the following actions from any of the Statistics dialog boxes:

1. To enter a Statistics dialog box, right-click the required component in the Tree View and select the statistics option from the menu as shown in Figure 10-49.

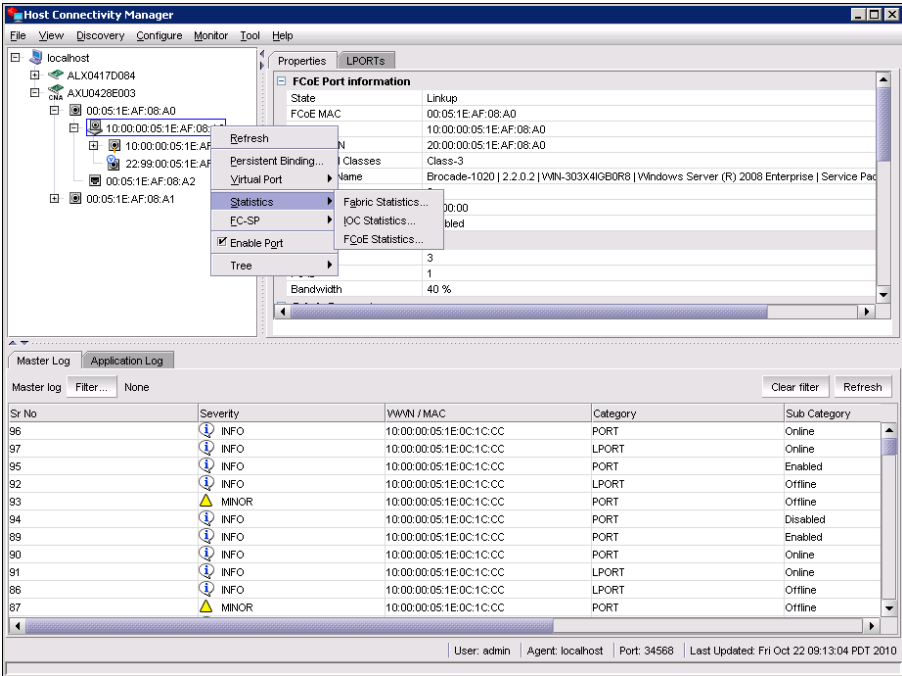


Figure 10-49 Statistics from HCM

The statistics dialog box for the selected element opens (Figure 10-50).

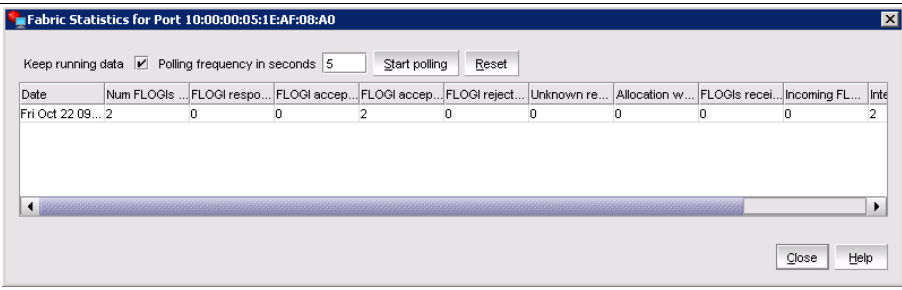


Figure 10-50 Port Statistics dialog box

Statistics: The statistics gathered and displayed will vary depending on which component and element are selected.

2. Click the **Start Polling** box to manually poll the port statistics.
3. Type the polling rate in the **Polling Frequency in Seconds** text box. The range is between 5 and 3600 seconds. The default is 5 seconds.
4. Click the **Stop Polling** box to stop port statistics polling.
5. Check the **Keep Running Data** check box to see the trend.

10.4.3 Resetting statistics

Follow these steps:

1. Click the **Reset** button on any of the Statistics dialog boxes (Figure 10-51).

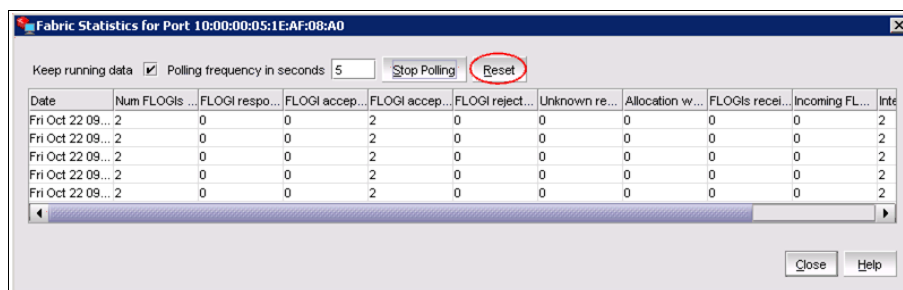


Figure 10-51 Reset the statistics

A warning dialog displays (Figure 10-52).

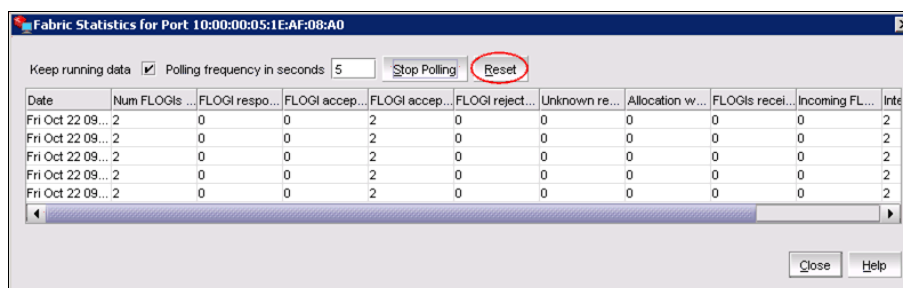


Figure 10-52 Reset Statistics warning message

2. Click **Yes**.
All of the statistics are reset to 0.

10.4.4 Master Log

The Master Log Properties dialog box, described in Table 10-6, displays a list of all events that have occurred. Event monitoring enables early fault detection and isolation on a selected adapter. You can filter the events based on the user-defined criteria shown in Figure 10-54 on page 465.

Table 10-6 Master Log fields

Field	Description
Filter button	Click to launch the Master Log Filter dialog box.
Clear Filter button	Click to clear the master log filter option set.
Sr No column	Displays a numbering sequence in ascending order.
Severity column	Displays the event severity (informational, minor, major, or critical).
WWN/MAC column	Displays the world wide name (WWN) or the media access control (MAC) address of the device on which the event occurred.
Category column	Displays the category of event; for example, Rport or ITNIM.
Subcategory column	Displays the subcategory of the main category.
Description column	Displays a brief description of the event.
Date/Time column	Displays the date and time when the event occurred.

Table 10-7 describes the icons that represent the four event types.

Table 10-7 HCM Master Log events

Description
Critical-level messages indicate that the software has detected serious problems that will eventually cause a partial or complete failure of a subsystem if not corrected immediately; for example, a power supply failure or rise in temperature must receive immediate attention.
Major messages represent conditions that do not impact overall system functionality significantly. For example, timeouts on certain operations, failures of certain operations after retries, invalid parameters, or failure to perform a requested operation.
Minor messages highlight a current operating condition that should be checked or it might lead to a failure in the future. For example, a power supply failure in a redundant system relays a warning that the system is no longer operating in redundant mode and that the failed power supply needs to be replaced or fixed.
Information-level messages report the current non-error status of the system components; for example, the online and offline status of a fabric port.

10.4.5 Filtering event log entries

Event filtering enables you to block events based on user-defined criteria (severity or type of log). Events that have been filtered out do not display in the Master Log:

1. Click the **Filter** button in the Master Log section of the bottom pane of the HCM main window (Figure 10-53).

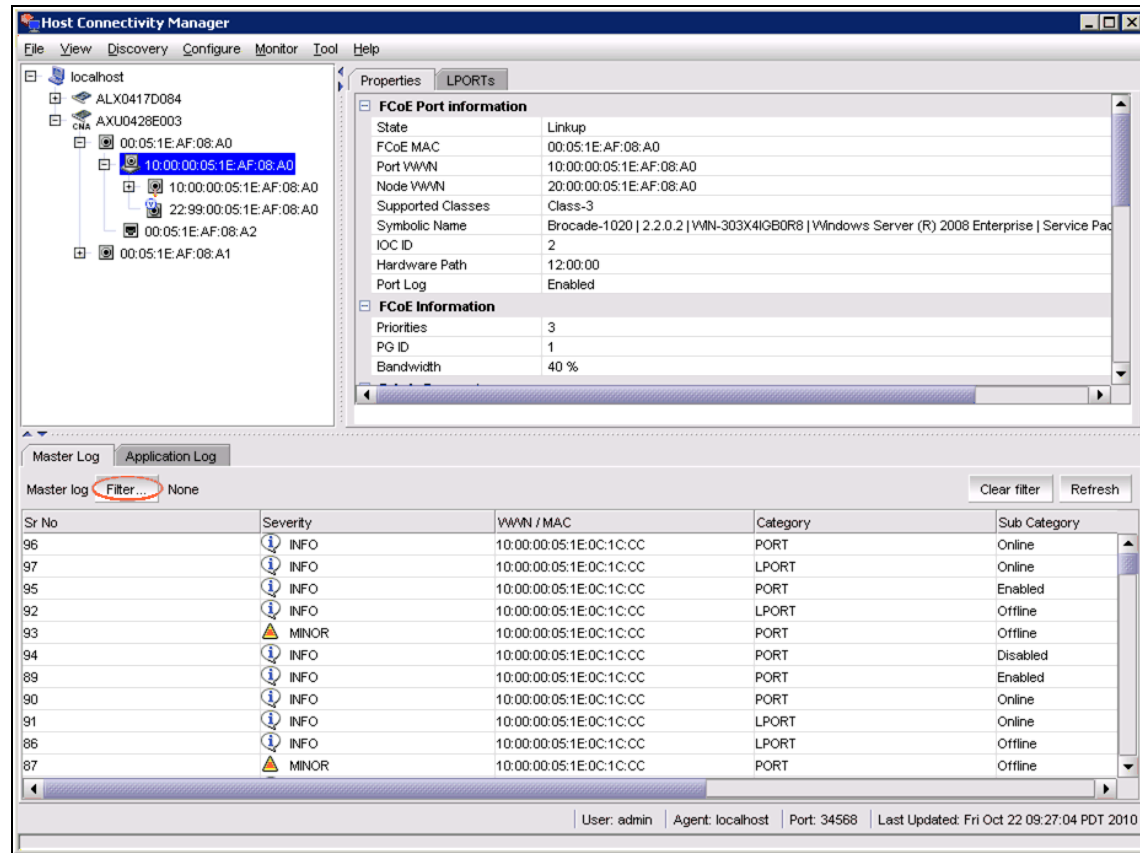


Figure 10-53 Filter Log Entries

The Master Log Filter dialog box displays (Figure 10-54).

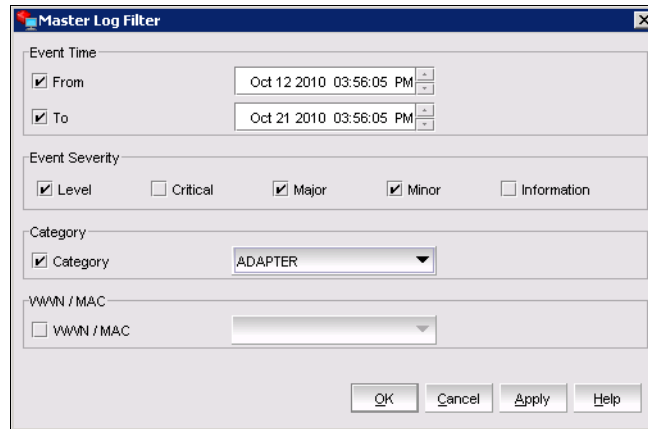
The image shows a Windows-style dialog box titled "Master Log Filter". It contains four main sections: "Event Time" with "From" and "To" date/time pickers set to Oct 12 2010 03:56:05 PM and Oct 21 2010 03:56:05 PM respectively; "Event Severity" with checkboxes for Level (checked), Critical, Major (checked), Minor (checked), and Information; "Category" with a checked "Category" checkbox and a dropdown menu showing "ADAPTER"; and "WWN / MAC" with an unchecked checkbox and an empty dropdown menu. At the bottom are "OK", "Cancel", "Apply", and "Help" buttons.

Figure 10-54 Master Log Filter dialog box

2. Filter the events using one or a combination of the criteria shown in Figure 10-54.

Events: The Category is the type of event; for example, an adapter, port, or audit.

3. Click **Apply** to save your changes, or click **Cancel** to exit the window.
or
Click **OK** to save the changes and exit the window.

10.4.6 Application log

The application log displays all application-related informational and error messages, as well as the following attributes (Figure 10-55):

- ▶ Date and time the message occurred
- ▶ Severity of the message
- ▶ Description of the message
- ▶ The agent IP address

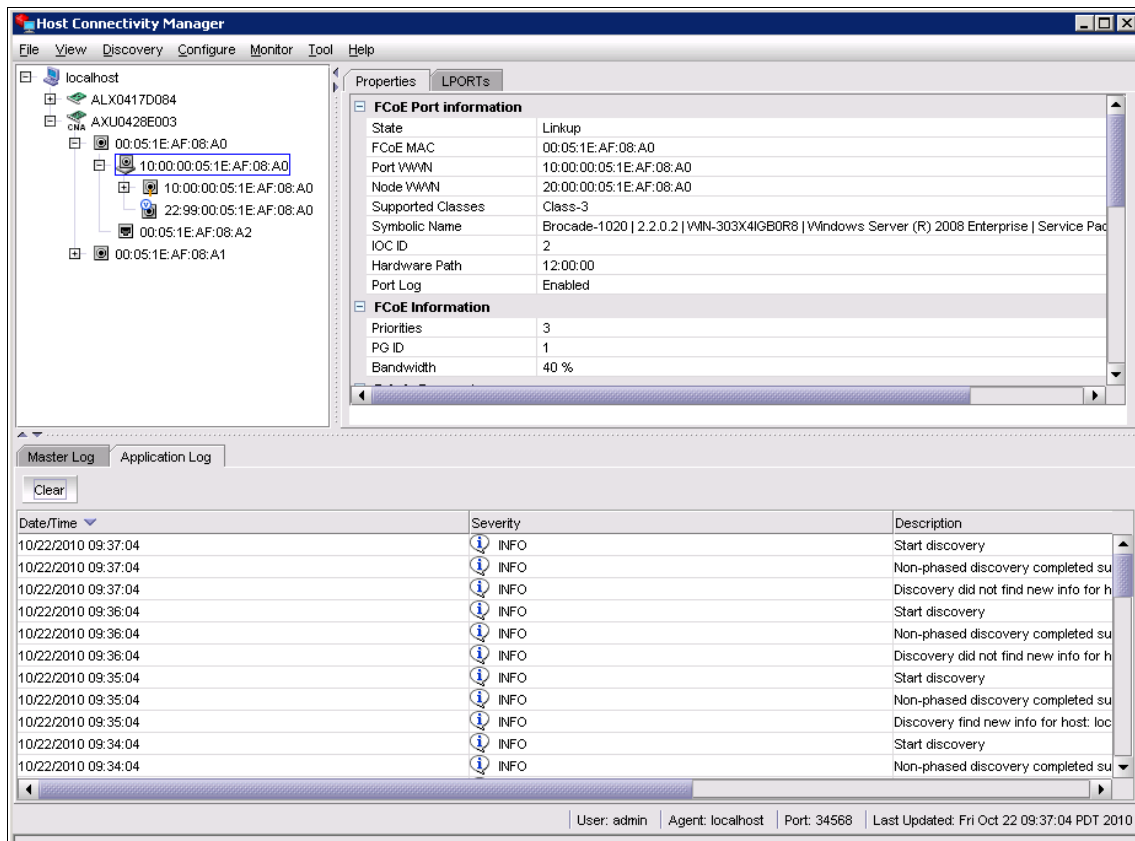


Figure 10-55 HCM Application Log

10.4.7 Syslog support

Syslog forwarding is the process by which you can configure the Host Connectivity Manager (HCM) agent to send Syslog messages to other computers through port 514. You can configure the HCM agent to forward events to a maximum of three Syslog destinations. These events will display in the operating system logs.

10.4.8 Opening the Syslog Server Configuration dialog box

Follow these steps:

1. Select an adapter from the device tree.
2. Select **Configure** → **Syslog** from the main menu.

The Syslog Server Configuration dialog box displays (Figure 10-56).

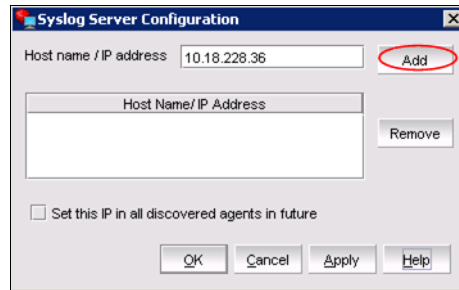


Figure 10-56 Syslog Server Configuration dialog box

3. Enter the host name or IP address of the destination device into the *Host Name/IP Address* field.
4. Click **Add** to register the host as a Syslog destination.
5. Click **OK** to close the dialog box.

10.4.9 Removing a host server

Follow these steps:

1. Select an adapter from the device tree.
2. Select **Configure** → **Syslog** from the main menu.

The Syslog Server Configuration dialog box displays (Figure 10-57).

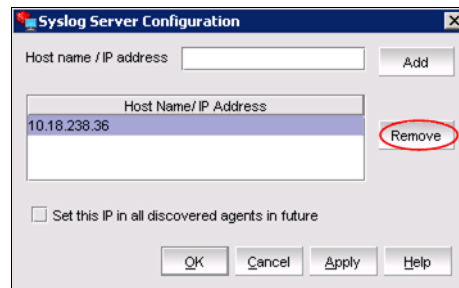


Figure 10-57 Syslog Server remove dialog box

3. Select the Server from the list.
4. Click **Remove** to remove the host as a Syslog destination.
5. Click **OK** to close the dialog box.



Virtual Fabrics

Virtual Fabrics is an architecture to virtualize hardware boundaries. Traditionally, SAN design and management is done at the granularity of a physical switch. The Virtual Fabrics feature allows SAN design and management to be done at the granularity of a port.

In this chapter we discuss Virtual Fabrics and provide examples of how to implement this feature.

11.1 IBM/Brocade Virtual Fabric

In this section we explain what the IBM/Brocade Virtual Fabrics feature is, and how it is configured to an operational state in a live environment.

11.1.1 Virtual Fabrics introduction

IBM/Brocade Virtual Fabrics allows IT organizations to manage IT assets by corporate function, utilize different permission levels for SAN administrators, and maintain required levels of data and fault isolation without increasing cost and complexity. In addition, Virtual Fabrics can reduce hardware costs by optimizing resource utilization.

With the release of FOS v6.2, organizations can utilize an ANSI standard-based implementation of Virtual Fabrics. The Virtual Fabrics feature includes two new capabilities: Logical switches and logical fabrics, both available in the base FOS firmware.

Physical switches can be partitioned into independently managed logical switches, each with their own data, control, and management paths. In addition, they can be configured in any mode, including McDATA Fabric or McDATA Open Fabric modes.

Logical switches can allocate fabric resources “by the port” rather than by the switch. They also provide a way to simplify charge-back for storage by customer, department, or application while cost-effectively consolidating SAN resources. Because logical switches do not need to be enabled on every switch in a SAN, deployment is simple and non-disruptive in existing environments.

11.1.2 Logical switches and logical fabrics

Next we describe some of the capabilities of these components.

Figure 11-1 introduces logical switches and logical fabrics.

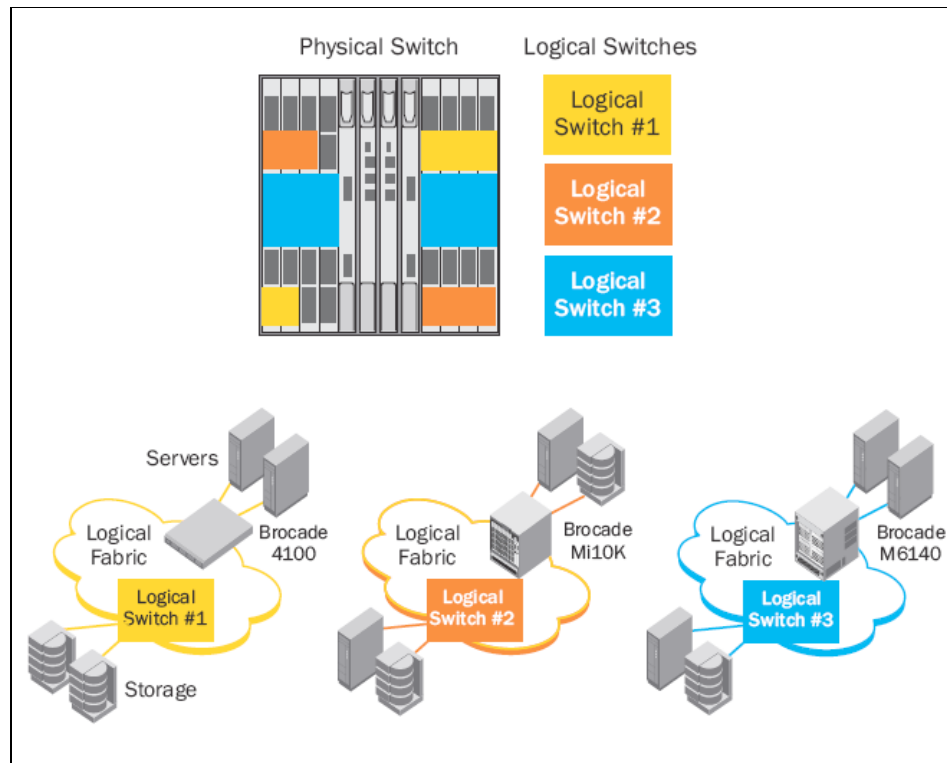


Figure 11-1 Logical switches and logical fabrics

Layer-2 traffic isolation is available with a special eXtended ISL (XISL) shared by multiple logical fabrics, or with dedicated ISL connections between Logical switches in the same logical fabric. Both ISL and XISL connections can use front ports or Inter-Chassis Link (ICL) connections with frame trunking and Dynamic Path Selection (DPS) for full bandwidth utilization. The logical fabrics capability supports Integrated Routing at Layer 3. Routing connections attach an integrated backbone fabric to multiple edge fabrics. Zoning allows traffic to flow between specific devices in any edge fabric.

Virtual Fabrics is available on 8 Gbps products that are “Virtual Fabrics-capable” such as the IBM SAN768B, IBM SAN384B, IBM SAN80B and the IBM SAN40B switches. For investment protection, products that are not Virtual Fabrics capable (such as earlier 2 and 4 Gbps FOS and M-EOS products) can seamlessly connect to Virtual Fabrics-capable products without requiring a reconfiguration of the existing switches. See Table 11-1.

Table 11-1 Supported logical switch creation limits

Platform	Maximum Logical Switches*/Chassis
SAN768B	8
SAN384B	8
SAN80B-4	4
SAN40B-4	3
SAN24B-4	Currently not supported

*Numbers include the Default Switch and Base Switch.

Attention: On the SAN80B-4 and the SAN40B-4, the Default Switch and Base Switch can be the same.

To simplify Virtual Fabrics management, organizations can use DCFM. After they are created, Logical switches and fabrics are managed the same as their physical counterparts. Alternatively, organizations can use the standard FOS CLI to enter commands or script configuration and management operations for Virtual Fabrics.

The Virtual Fabrics (VF) feature is easy to set up and simple to manage for “port-level” partitioning of physical switches into independent logical switches. It does not reduce fabric or chassis scalability, preserving ROI and seamlessly supporting advanced FOS features such as frame trunking, DPS, Fibre Channel Routing, Adaptive Networking, Top Talkers, Access Gateway, Access Gateway trunking, and FCIP for extension.

11.2 What Virtual Fabrics are

In this section we discuss the different features of Virtual Fabrics. For more detailed information see the Fabric OS version 6.4.+ Administrators Guide only available through the Partner Network website at (navigate to Product Documentation and register or login):

<http://www.brocade.com/data-center-best-practices/resource-center/index.page>

This section describes the logical switch and logical fabric features.

The Virtual Fabrics suite consists of the following specific features:

- ▶ Logical switch
- ▶ Logical fabric
- ▶ Device sharing

Virtual Fabrics is supported on the following platforms:

- ▶ IBM System Storage SAN40B
- ▶ IBM System Storage SAN80B
- ▶ IBM System Storage SAN384B
- ▶ IBM System Storage SAN768B

Other non-Virtual Fabric capable switches can connect to Virtual Fabrics without any reconfiguration.

11.2.1 Logical switch

A logical switch is the fundamental component of a Virtual Fabric. When enabled on a VF-capable switch, Virtual Fabrics allows users to divide the switch into multiple logical switches. Ports in the physical switch can be dynamically allocated to any logical switch in the chassis and can be reallocated to Logical switches as needed. Port, switch, and fabric management are performed in the same way as for physical switches or fabrics.

Default logical switch

The default logical switch (default switch) is automatically created when Virtual Fabrics is enabled on a VF-capable switch. Initially, the default switch contains all the physical switch resources and ports. For Director switches, the ports on any blade inserted into the chassis initially belong to the default switch. Ports required by user defined logical switches are dynamically allocated from the default switch by the chassis administrator. As long as the Virtual Fabrics feature is enabled, there is a default switch, even when all ports in the default switch have been allocated to other logical switches. The default switch supports all the same port types as the physical switch.

Base switch

Base Switch is a logical switch used to communicate among different logical switches. The legacy EX_port is connected to the base logical switch. Also, Inter-Switch Links (ISLs) connected to the Base Switch are used to communicate among different fabrics. The default logical switch supports E_ and EX_ports.

Logical Switch

Logical Switch is a collection of zero or more ports, that act as a single Fibre Channel (FC) switch. When Virtual Fabrics is enabled on the chassis, there is always at least one default logical switch instance. You must assign each logical switch (default or general) in the same chassis to a different logical fabric. The logical switch supports all E_ and F_ports.

Attention: EX_ports are only allowed on the Base Switch.

11.2.2 Logical fabric

The Fabric ID (FID) assigned to a logical switch identifies its traffic as belonging to a specific logical fabric. Logical switches in other chassis with the same FID can join into a logical fabric. Logical switches within a logical fabric can be directly connected with ISLs (front ports and/or ICL connections), supporting frame trunking and DPS. As is the case in a physical fabric, the ISL connection carries traffic for a single fabric. An alternative to dedicated ISL connections at Layer 2 uses the base fabric to carry traffic for multiple logical fabrics on the same physical connection, maintaining fabric isolation.

11.2.3 ISL sharing

When a base switch is connected to another base switch, an XISL connection is created. When logical switches with the same FID are configured to use the XISL, the base switches automatically create a Logical ISL (LISL) within the XISL. The LISL isolates traffic from multiple fabrics: each LISL is dedicated to traffic for a single fabric. Think of it this way: the physical XISL connection between two base switches automatically forms an LISL “tunnel” dedicated to the traffic to and from logical switches, as shown by the dashed lines in Figure 11-21 on page 497.

11.2.4 Administrative Domains

An Administrative Domain (Admin Domain or AD) is a logical grouping of fabric elements that defines which switches, ports, and devices can be viewed and modified. An Admin Domain is a filtered administrative view of the fabric.

Basically Admin Domains define which users can manage which devices, hosts, and switches.

Important: Virtual Fabrics and Admin Domains are mutually exclusive and are not supported at the same time on a switch:

- To use Admin Domains, you must first disable Virtual Fabrics.
- To use Virtual Fabrics, you must first delete all Admin Domains.

11.2.5 User accounts

Table 11-2 lists the predefined user accounts offered by Fabric OS available in the local switch user database.

Table 11-2 Default local user accounts

Account name	Logical fabric	Description
admin	LF1-128 home: 128	Observe-modify permission.
factory	LF1-128 home: 128	Reserved.
root	LF1-128 home: 128	Reserved.
user	LF1-128 home: 128	Observe-only permission.

The password for all default accounts should be changed during the initial installation and configuration for each switch.

11.3 Configuring Virtual Fabrics

In this section we present a limited set of instructions and commands for configuring and managing Virtual Fabrics. For complete explanations, read the Fabric OS version 6.2.0 Administrators Guide and Fabric OS Command Reference Manual available only available through the Partner Network website at (navigate to Product Documentation and register or login):

<http://www.brocade.com/data-center-best-practices/resource-center/index.page>

Virtual Fabrics (VF) can be managed with Data Center Fabric Manager (DCFM), in this section we demonstrate how to configure VF using the standard Fabric OS v6.4.+ Command Line Interface (CLI), and DCFM.

11.3.1 Changing the context to a different logical switch

When Virtual Fabric is enabled, you want to move between the defined virtual switches. This is done using either Webtools or the **setcontext** command from CLI.

This is how we did it:

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. In Example 11-1 we show how to switch to FID 20.

Example 11-1 To change the logical switch context to FID 20:

```
IBM_SAN384B_27:admin> setcontext 20
```

11.3.2 Enabling Virtual Fabrics

Virtual Fabrics is disabled by default on switches that you *upgrade* to Fabric OS v6.2.0 or later. Virtual Fabrics is enabled by default on a new chassis. Before you can use the Virtual Fabrics features, such as logical switch and logical fabric, you must enable Virtual Fabrics.

Attention: When you enable Virtual Fabrics, the CPs are rebooted and all EX_Ports are disabled after the reboot.

DCFM

To enable virtual fabrics (VF) all you need is to do is to select the switch that VF is going to be configured on, and then right click to get the drop down menus for the switch and select the Enable Virtual Fabric option, as shown in Figure 11-2.

Requirement: SNMP V3 must be enabled and configured for management of Virtual Fabrics,

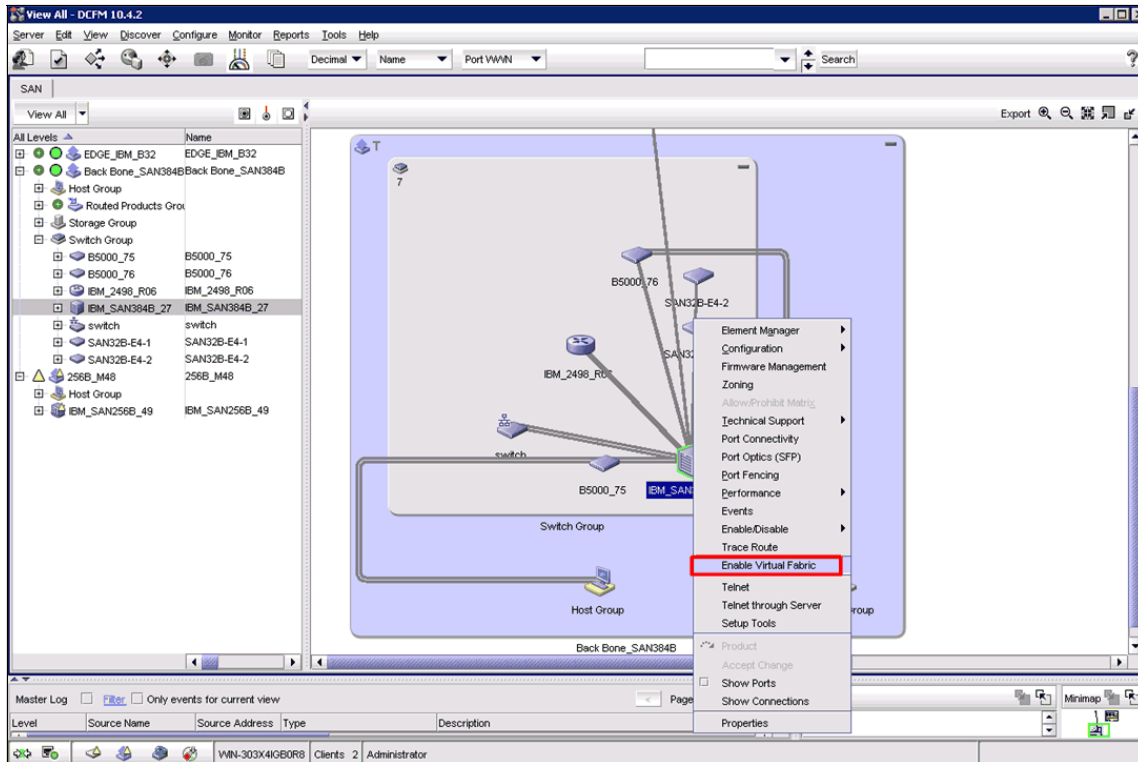


Figure 11-2 Enable Virtual Fabric

The warning message will display, shown in Figure 11-3. Read the warning message and select the **OK** button.




Figure 11-3 VF warning message

When this operation is completed and the reboot is done, you will see a **V** symbol above the VF enabled switch, and there will be a chassis group in the product list, as shown in Figure 11-4.

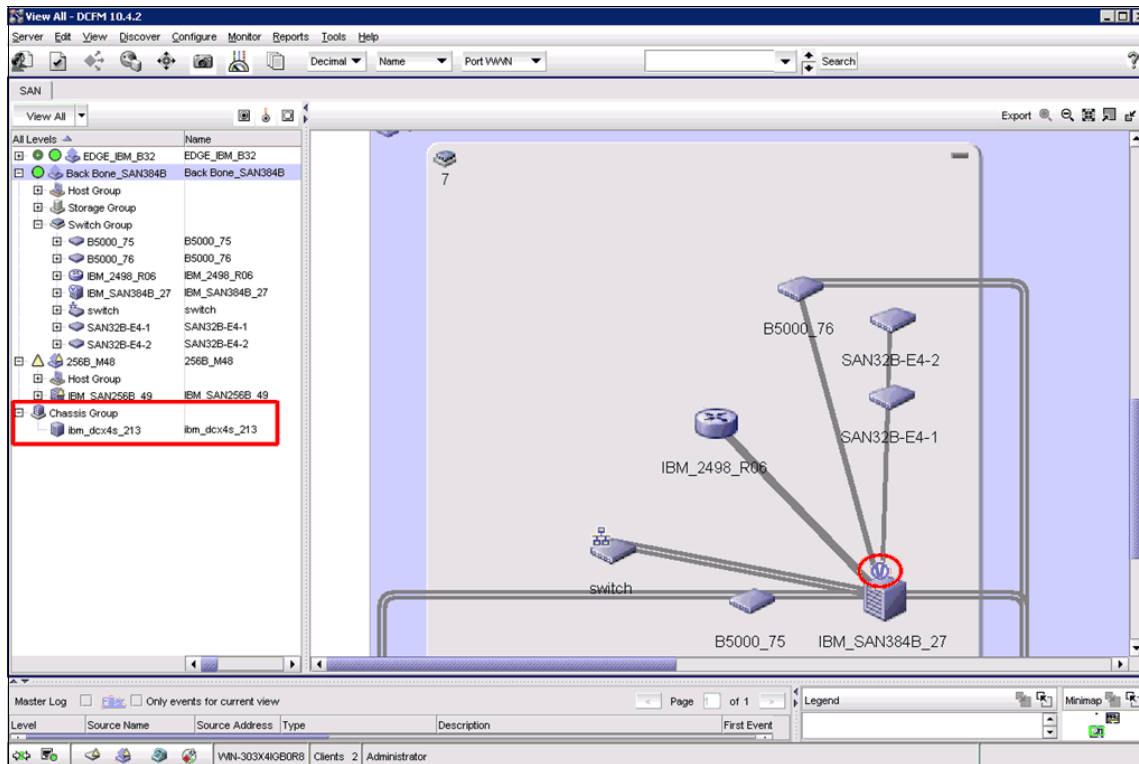


Figure 11-4 VF enabled switch

CLI

To perform management of virtual fabrics, you must have admin privileges on the switch chassis.

Example 11-2 checks whether Virtual Fabrics is enabled or disabled and then enables it.

Example 11-2 Enabling Virtual Fabrics

```
IBM_SAN384B_27:admin> fosconfig --show
FC Routing service:          enabled
iSCSI service:               Service not supported on this Platform
iSNS client service:         Service not supported on this Platform
Virtual Fabric:              disabled
Ethernet Switch Service:     disabled
```

```
IBM_SAN384B_27: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
```

11.3.3 Disabling Virtual Fabrics

In this section we discuss how to disable Virtual Fabrics.

DCFM

To disable Virtual Fabrics, select the switch in the chassis group displayed in the product list, right-click to open the drop down menu options and select the option to disable Virtual Fabrics, shown in Figure 11-5.

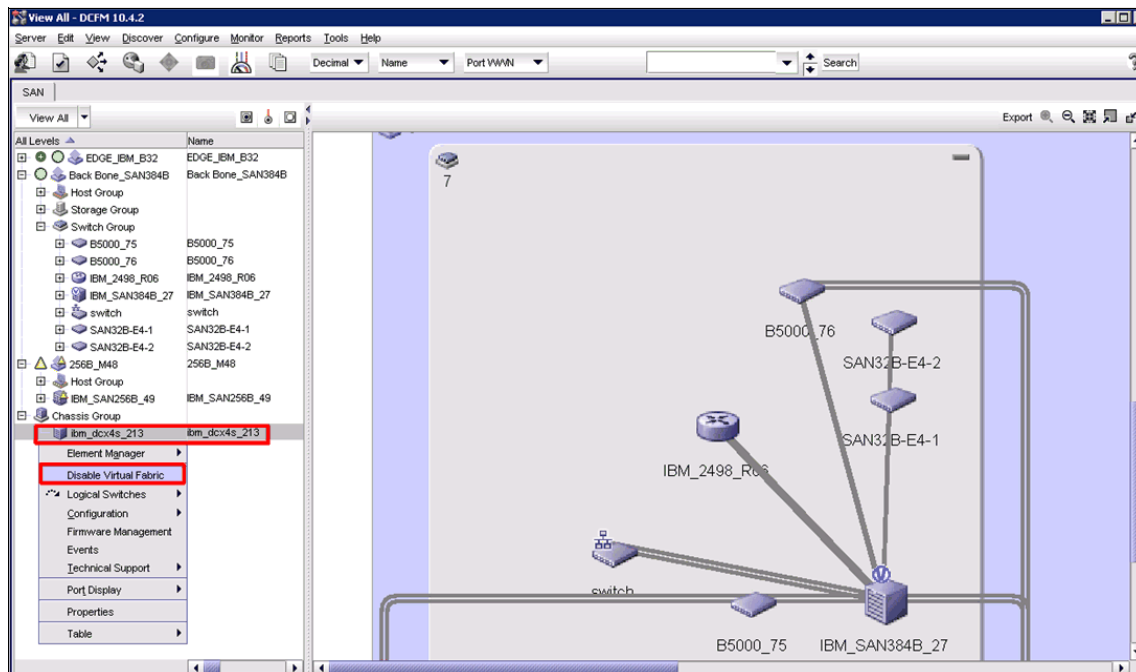


Figure 11-5 Disable Virtual Fabric

The same warning message is shown as in enabling the VF, as shown in Figure 11-3 on page 477. Read the warning and select the **OK** button if applicable.

CLI

Example 11-3 checks whether Virtual Fabrics is enabled or disabled and then disables it.

Example 11-3 Disabling Virtual Fabrics

```
IBM_SAN384B_27: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

IBM_SAN384B_27:FID128: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
VF has been disabled. Your system is being rebooted.
```

Attention: Enabling and disabling Virtual Fabrics is disruptive and will reboot the switch.

11.3.4 Logical switch management

DCFM is used to manage logical switches after Virtual Fabrics has been enabled. From the DCFM Configure drop-down menu, select the **Logical Switches** option, shown in Figure 11-6.

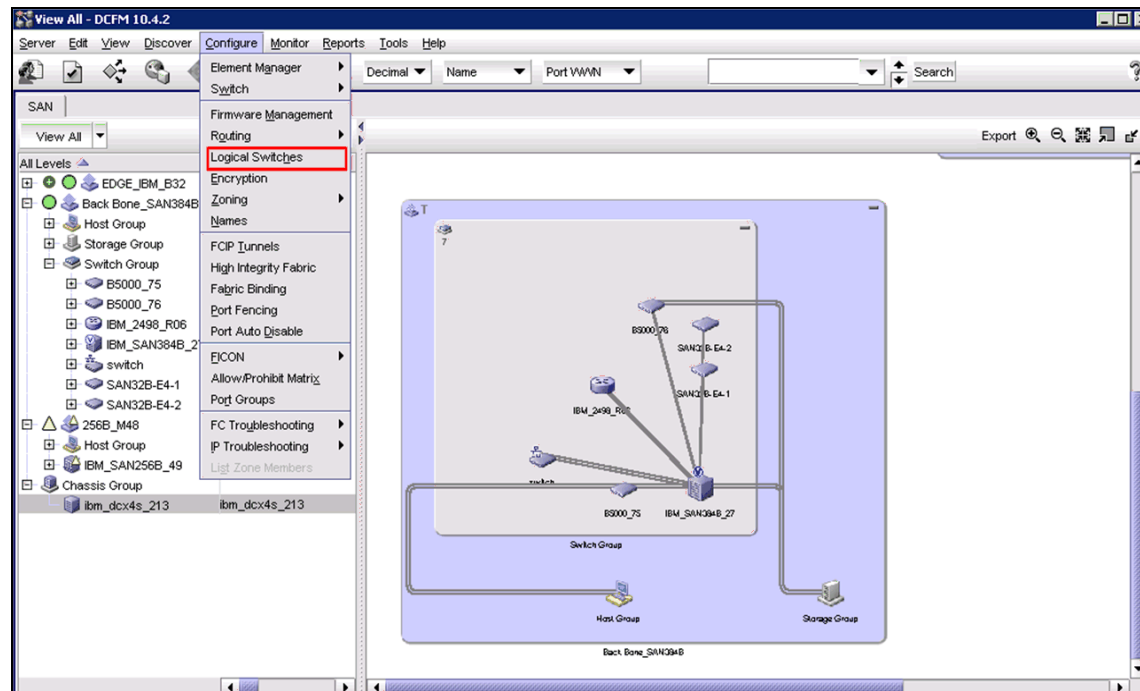


Figure 11-6 Logical switches DCFM

This opens the Logical Switches management window, shown in Figure 11-7. When Virtual Fabrics is enabled, a base switch is automatically created with an FID of 128, the same as the backbone switch, and all ports in the switch are placed into this base switch.

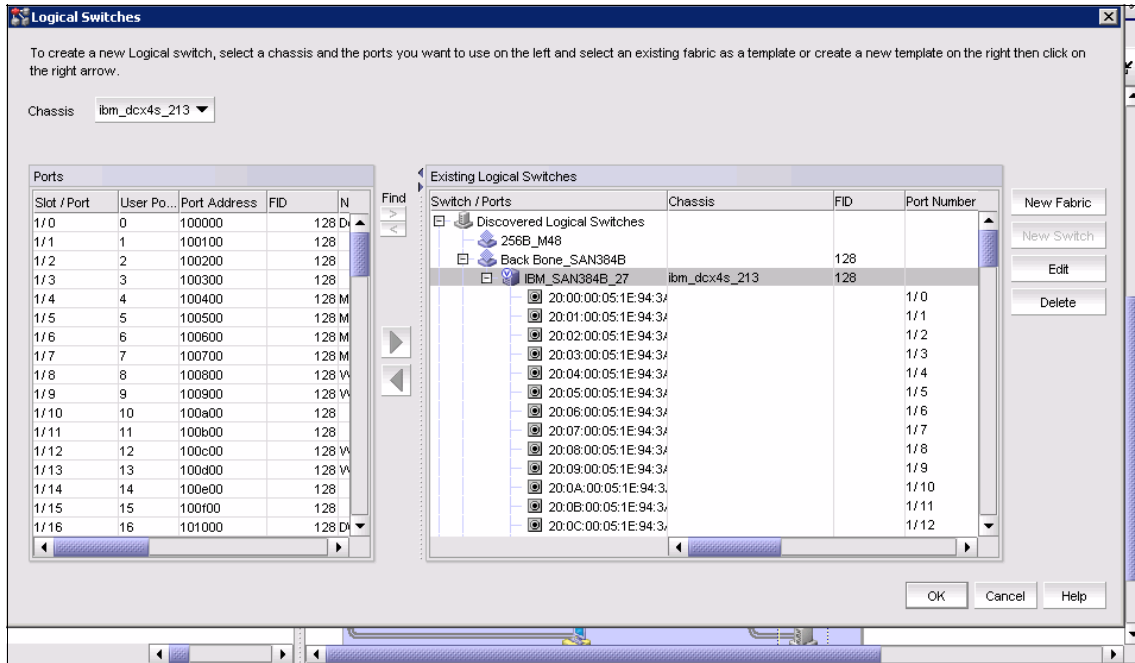


Figure 11-7 Logical Switch management

11.3.5 Modifying the base switch

To modify the base switch, select the base switch from the Logical Switches window and select the **Edit** button. This will allow the modification of all base switch parameters, as shown in Figure 11-8.

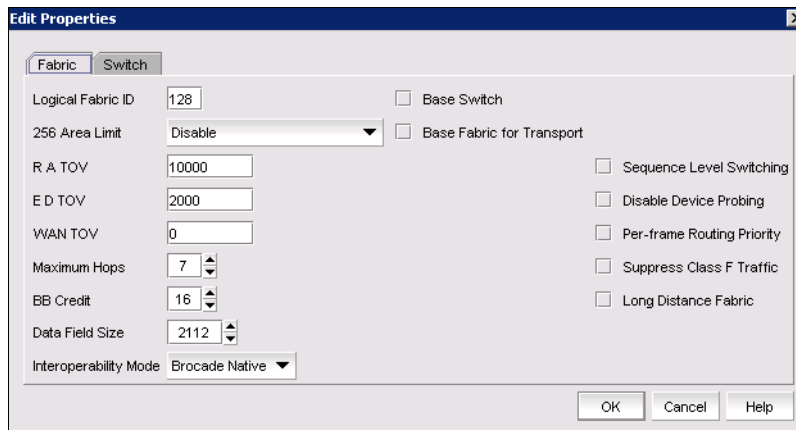


Figure 11-8 Edit Properties

After the configuration edit is complete, select the **OK** button from the Edit screen and then select the **OK** button from the Logical Switches management window.

This action opens a confirmation window. Read the message on the window and select **OK**. This performs a configuration operation and displays the progress of the command under the status field, as shown in Figure 11-9.

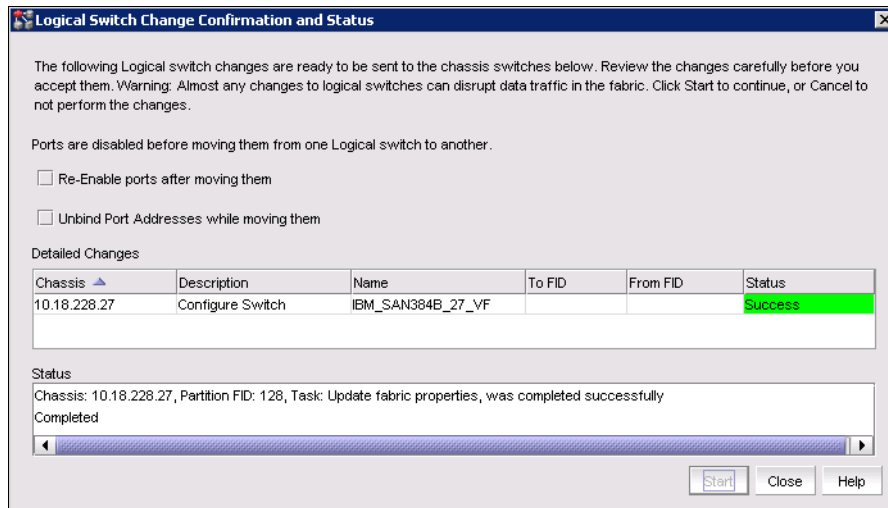


Figure 11-9 Confirmation window

11.3.6 Creating a logical switch

When the logical switch is created, it is automatically enabled and it has no ports assigned.

To create a logical switch, open the Logical Switches view, and select the **New Fabric** option. This action brings up the New Logical Fabric template. Select the options required for the new fabric, shown in Figure 11-10, and when done, select the **OK** button.

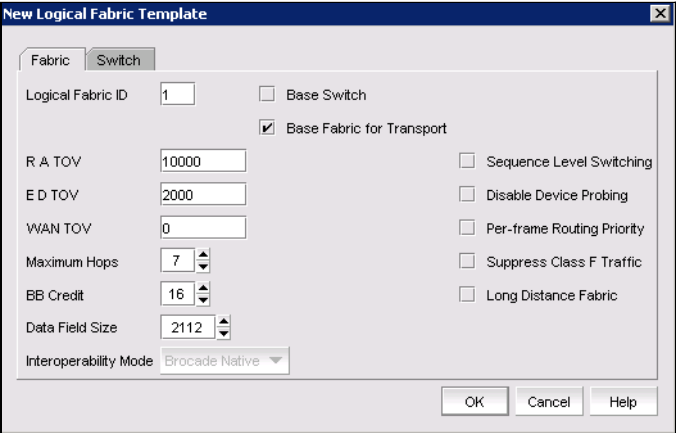


Figure 11-10 New Logical Fabric Template

The new logical fabric displays in the Logical Switches window. Select the new fabric and then select the **New Switch** button, as shown in Figure 11-11

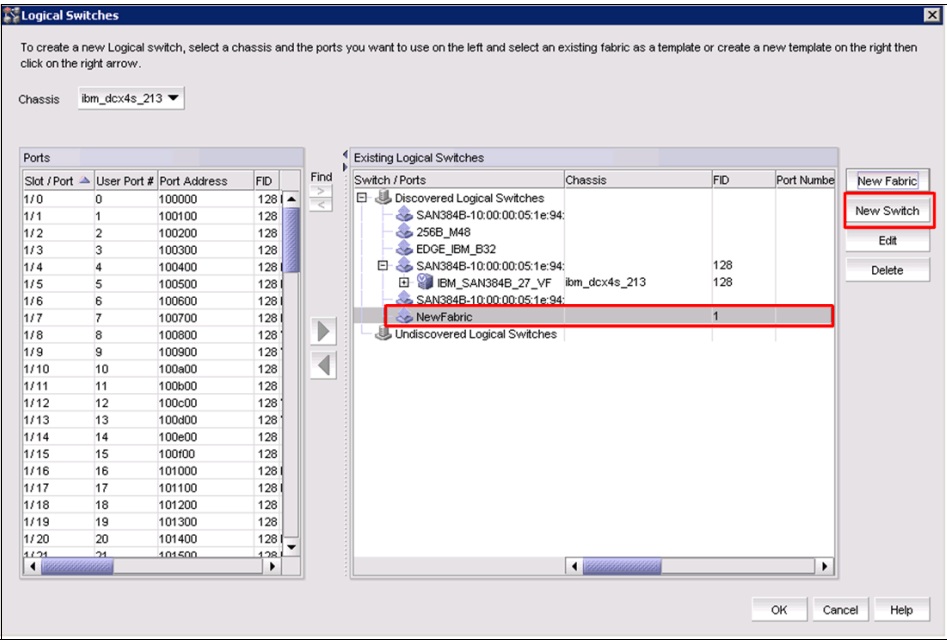


Figure 11-11 Add new switch

The new logical switch dialog frame opens. Configure the new logical switch as required by modifying the fields as shown in Figure 11-12.

New Logical Switch

Fabric | **Switch**

Logical Fabric ID: 1

256 Area Limit: Disable

R A TOV: 10000

E D TOV: 2000

WAN TOV: 0

Maximum Hops: 7

BB Credit: 16

Data Field Size: 2112

Interoperability Mode: Brocade Native

☐ Base Switch

☐ Base Fabric for Transport

☐ Sequence Level Switching

☐ Disable Device Probing

☐ Per-frame Routing Priority

☐ Suppress Class F Traffic

☐ Long Distance Fabric

OK Cancel Help

Figure 11-12 New logical switch Fabric parameters

Under the Switch option, you can change the switch name and domain ID, as shown in Figure 11-13.

New Logical Switch

Fabric | **Switch**

Name: ITSO_SW2

Preferred Domain ID: 25

☐ Insistent Domain ID

OK Cancel Help

Figure 11-13 New logical switch

Select the **OK** button to add the switch.

From the logical switch window, select the new logical switch and add the ports that are required for this switch by selecting them and adding them to the newly created logical switch as shown in Figure 11-14. This process can be used at any time to add or delete ports from the logical switch.

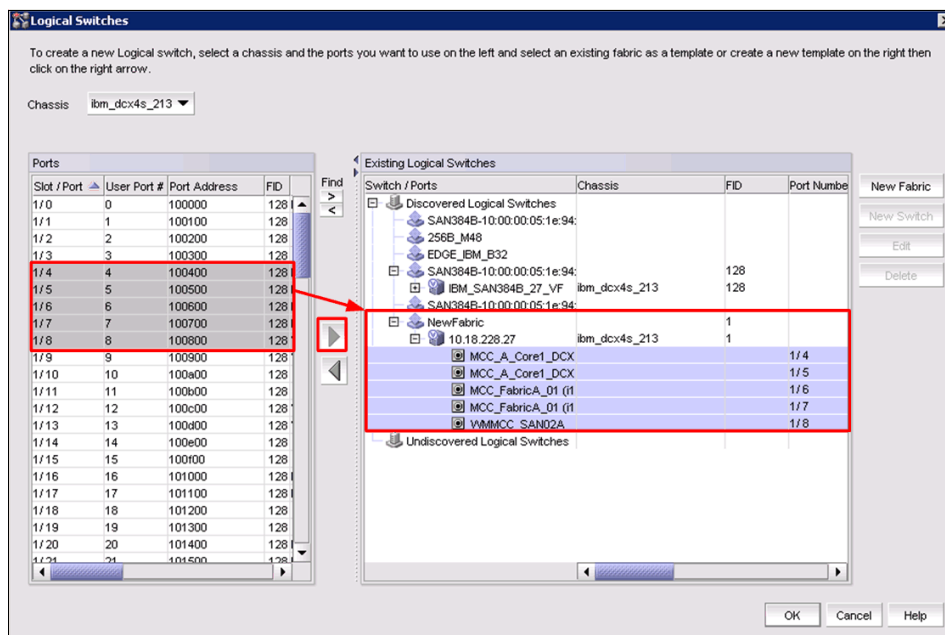


Figure 11-14 Add ports to logical switch

Now select the **OK** button to process the new configuration. The Logical Switch Change Confirmation and Status window displays. Read the information in the window and then select the **Start** button to complete the addition of the logical switch, as shown in Figure 11-15.

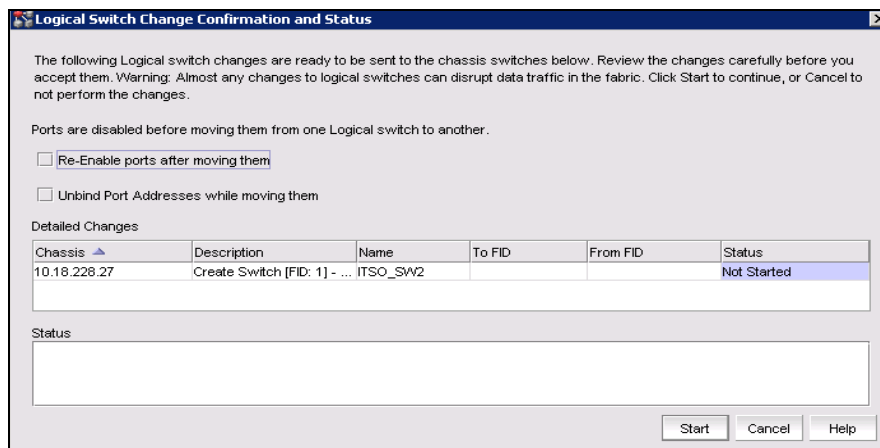


Figure 11-15 Logical switch change confirmation and status

The status bar displays the status of the activation. It will change to *Successful* when completed, and the newly created fabric and switch will display in DCFM, as shown in Figure 11-16.

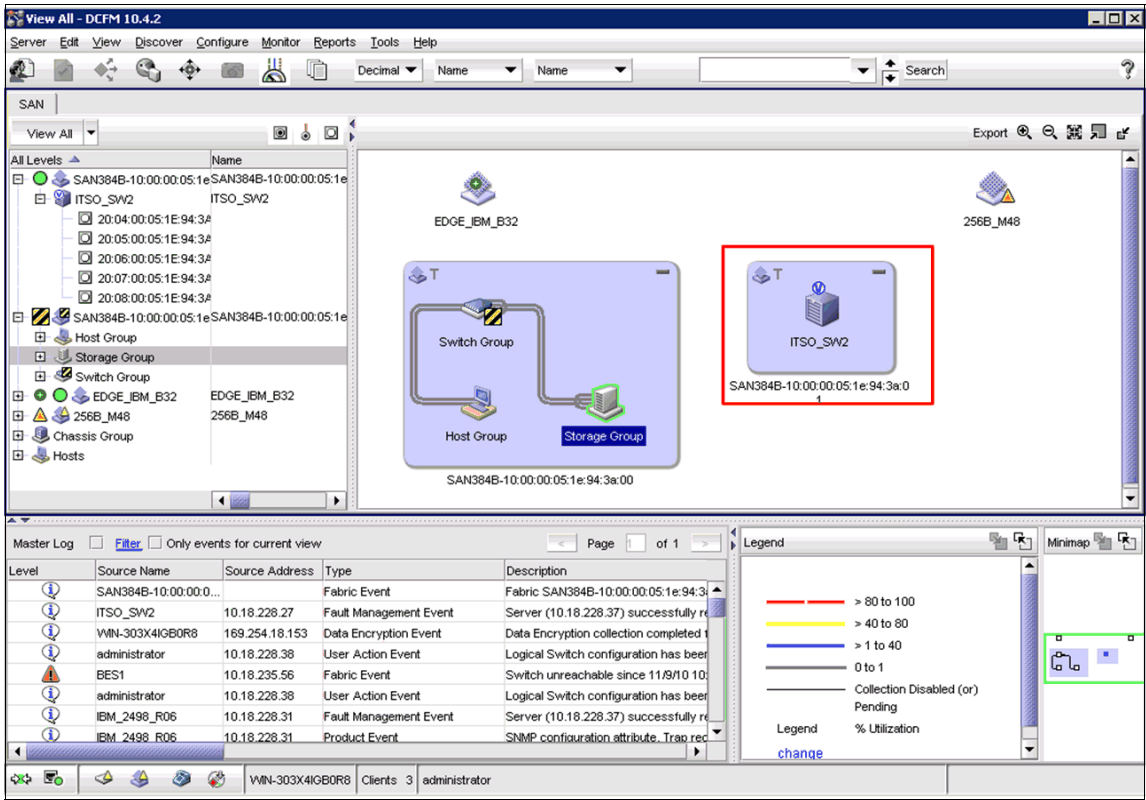


Figure 11-16 DCFM logical switch

11.3.7 Deleting a logical switch

To delete a logical switch, open the Logical Switches configuration window, select the switch you want to delete, and select the **Delete** button. A warning message will pop up. Read the warning and click the **Yes** button, as shown in Figure 11-17.

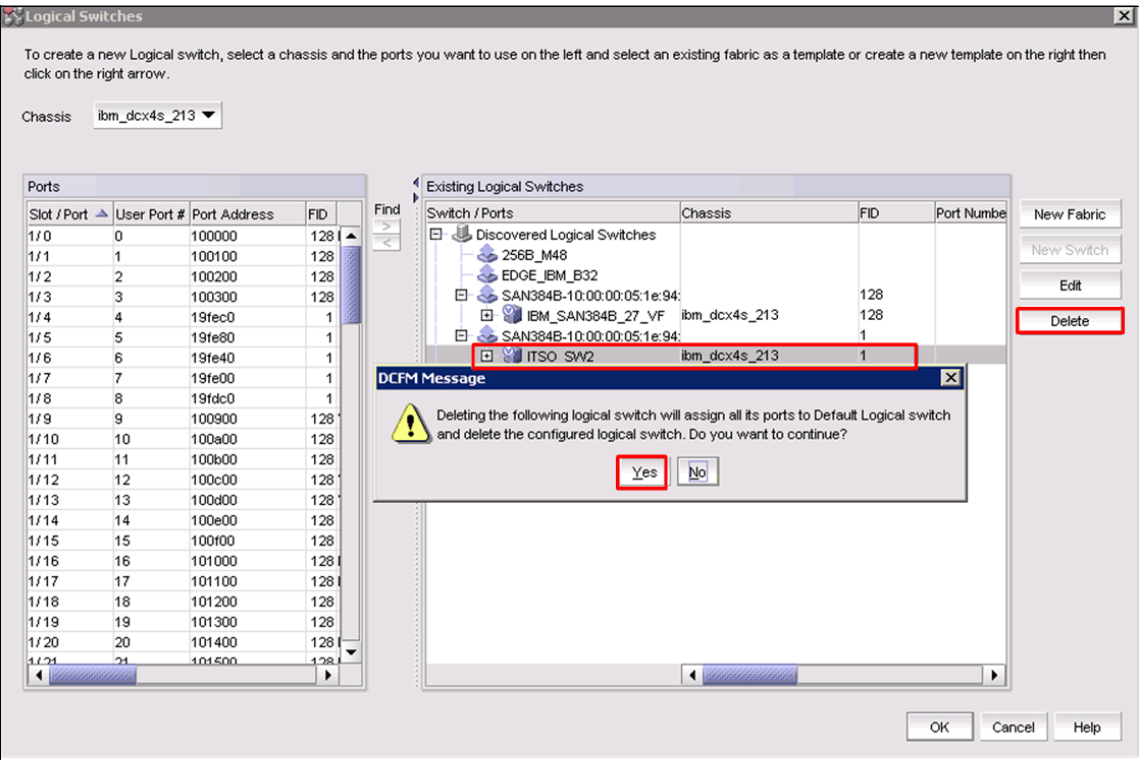


Figure 11-17 Delete switch

When the delete is completed, then select the **OK** button to activate the new configuration.

Attention: The default logical switch cannot be deleted.

11.3.8 Displaying the logical switch configuration

Example 11-4 shows the configuration created using the CLI.

Example 11-4 To display the logical switch configuration

```
IBM_SAN384B_27_VF:FID128:admin> lscfg --show
```

Created switches: 128(ds) 1

Slot	1	2	3	4	5	6	7	8

Port								
0	128		128			128	128	128
1	128		128			128	128	128
2	128		128			128	128	128
3	128		128			128	128	128
4	1		128			128	128	128
5	1		128			128	128	128
6	1		128			128	128	128
7	1		128			128	128	128
8	1		128			128	128	128
9	128		128			128	128	128
10	128		128			128	128	128
11	128		128			128	128	128
12	128		128			128	128	128
13	128		128			128	128	128
14	128		128			128	128	128
15	128		128			128	128	128
16	128							128
17	128							128
18	128							128
19	128							128
20	128							128
21	128							128
22	128							128
23	128							128
24	128							128
25	128							128
26	128							128
27	128							128
28	128							128
29	128							128
30	128							128
31	128							128

11.3.9 Changing the fabric ID of a logical switch

To change the fabric ID of an existing logical switch, select the **Logical Switches** window and select the **Edit** button. This opens the Edit Properties, where you can change the logical fabric ID, as shown in Figure 11-18.

The screenshot shows the 'Edit Properties' dialog box with the 'Switch' tab active. The 'Logical Fabric ID' field contains the value '35'. Below it, the '256 Area Limit' is set to 'Disable'. Further down, 'R A TOV' is '10000', 'E D TOV' is '2000', and 'WAN TOV' is '0'. 'Maximum Hops' is set to '7', 'BB Credit' to '16', and 'Data Field Size' to '2112'. The 'Interoperability Mode' is 'Brocade Native'. On the right side, there are seven unchecked checkboxes: 'Base Switch', 'Base Fabric for Transport', 'Sequence Level Switching', 'Disable Device Probing', 'Per-frame Routing Priority', 'Suppress Class F Traffic', and 'Long Distance Fabric'. At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

Figure 11-18 Change Logical Fabric ID

The fabric ID indicates in which fabric the logical switch participates. By changing the fabric ID, you are moving the logical switch from one fabric to another.

On the Logical Switches window, the switch will display under the new fabric ID, as shown in Figure 11-19.

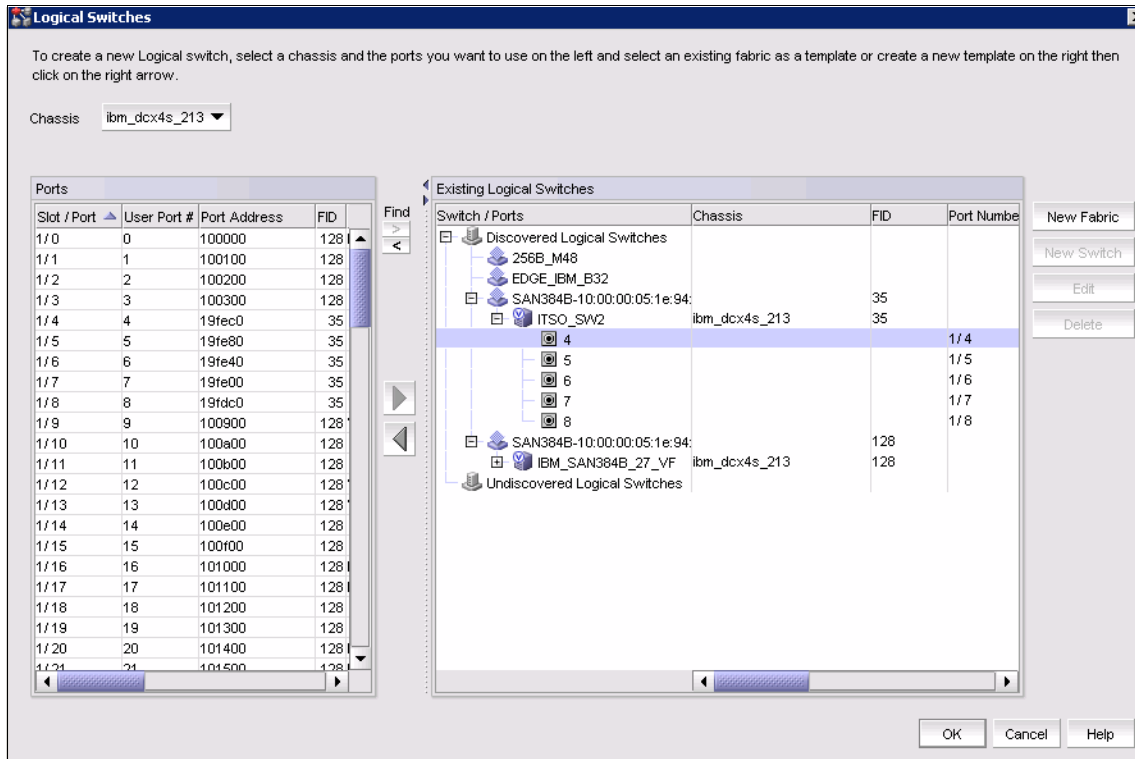


Figure 11-19 Logical Switch view with changed ID

To activate the change, select the **OK** button on the logical switches window, then read the confirmation message and select the **Start** button to complete the operation.

Attention: It might be necessary to delete the fabric from DCFM and then rediscover the fabric for the new logical switch to be shown.

11.3.10 Changing a logical switch to a base switch

Only the base switch can be used for Inter Switch Links. If there is no base switch already, you might want to change one of the logical switches to a base switch. To do this, select the logical switch in the Logical Switch View window and then select the **Edit** button. The Edit Properties window will display, where you can check the base switch button and then the **OK** button, as shown in Figure 11-20.

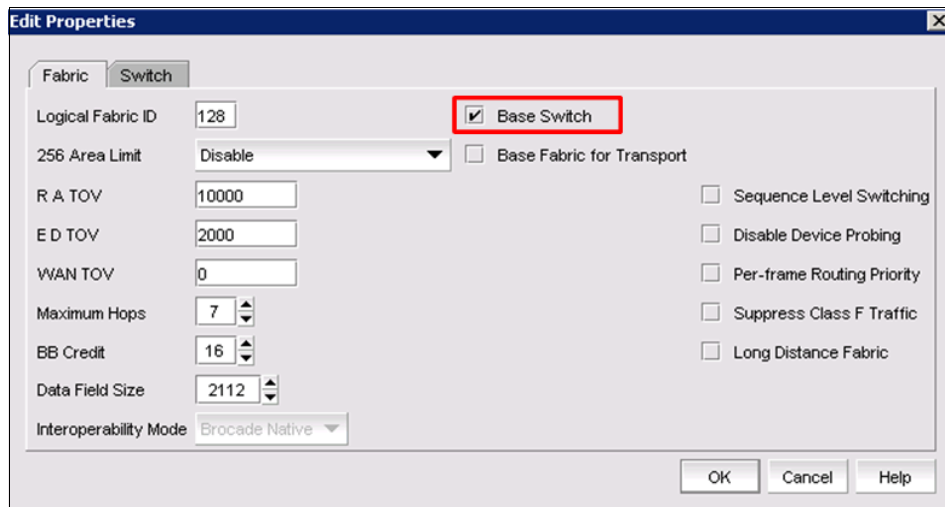


Figure 11-20 Edit Properties base unit

To activate the change, select the **OK** button in the Logical Switches window, then read the confirmation message and select the **Start** button to complete the operation.

Important: Trunk areas must be disabled to change a switch into a base switch. You can do this using the `porttrunkarea --disable all` command. The switch must be disabled to do this.

11.3.11 Configuring a logical switch for XISL use

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. Check with the `switchshow` command whether the switch is enabled for XISL use as shown in Example 11-5.

Example 11-5 Check XISL with switchshow

```
IBM_SAN384B_213:FID128:admin> switchshow
switchName:      IBM_SAN384B_213
switchType:      77.3
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:    1
```

switchId: fffc01
switchWwn: 10:00:00:05:1e:94:3a:00
zoning: OFF
switchBeacon: OFF
FC Router: OFF
Allow XISL Use: OFF
LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes]

Index	Slot	Port	Address	Media	Speed	State	Proto
0	1	0	010000	--	N4	No_Module	
1	1	1	010100	--	N4	No_Module	
2	1	2	010200	--	N4	No_Module	

- 2. Enter the **switchdisable** command to disable the switch (no output returned)
- 3. Use the **configure** command to configure the switch as shown in Example 11-6:

Example 11-6 use configure to allow or disallow XISL use

```
switch_100:FID100:admin> configure

Configure...

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

Domain: (1..239) [1] 100
Allow XISL Use (yes, y, no, n): [yes]
Enable a 256 Area Limit
(0 = No,
 1 = Zero Based Area Assignment,
 2 = Port Based Area Assignment): (0..2) [0]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
WAN_TOV: (0..30000) [0]
MAX_HOPS: (7..19) [7]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..27) [16]
Disable FID Check (yes, y, no, n): [no]
```

```
Insistent Domain ID Mode (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
Portlog events enable (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
webtools attributes (yes, y, no, n): [no]
```

WARNING: The domain ID will be changed. The port level zoning may be affected

4. Respond to the remaining prompts or press **Ctrl-d** to accept the other settings and exit.
5. Enter the **switchenable** command to re-enable the switch (no output is returned).

11.3.12 Creating a logical fabric using XISLs

The following procedure describes the *flow* in how to create a logical fabric using multiple chassis and XISLs. We discuss this topic in more detail in 11.4, “A real life example of Virtual Fabrics” on page 495. Follow these steps:

1. Set up the base switches in each chassis:
 - a. Connect to the physical chassis and log in using an account assigned to the admin role with the chassis-role permission.
 - b. Enable the Virtual Fabrics feature, if it is not already enabled. This automatically creates the default logical switch, with FID 128. All ports in the chassis are assigned to the default logical switch.
 - c. Create a base switch and assign it a fabric ID that will become the FID of the base fabric.
 - d. Assign ports to the base switch.
 - e. Repeat the prior steps for all chassis that are to participate in the logical fabric.
2. Physically connect ports in the base switches to form XISLs.
3. Enable all of the base switches. This forms the base fabric.

4. Configure the logical switches in each chassis:
 - a. Connect to the physical chassis and log in using an account assigned to the admin role with the chassis-role permission.
 - b. Create a logical switch and assign it a fabric ID for the logical fabric. This FID must be different from the FID in the base fabric.
 - c. Assign ports to the logical switch.
 - d. Physically connect devices and ISLs to the ports on the logical switch.
 - e. (Optional) Configure the logical switch to use XISLs, if it is not already XISL-capable. By default, newly created logical switches are configured to allow XISL use.
 - f. Repeat the prior steps for all chassis that are to participate in the logical fabric, using the same fabric ID whenever two switches need to be part of a single logical fabric.
5. Enable all logical switches by using the **switchenable** command.

Now the logical fabrics are formed.

The **fabricShow** command displays all logical switches configured with the same fabric ID as the local switch and all non-Virtual Fabric switches connected through ISLs to these logical switches.

The **switchShow** command displays logical ports as E_Ports, with -1 for the slot and the user port number for the slot port.

11.4 A real life example of Virtual Fabrics

In this section we demonstrate how to set up Virtual Fabrics in a live environment. The case is built in a lab environment and can be scaled larger or smaller according to customer needs.

In order to demonstrate the ability to share ISLs, we build four logical switches that merge into two logical fabrics. The shared ISLs, called XISLs, will carry all traffic between the two data centers. This includes the two fabrics we build in the example, as well as other future departments or users who might buy parts of a partitioned switch, but do not have their own ISLs.

New users or departments can use the existing XISLs and still have their own individual switched fabrics spanning over sites.

11.4.1 The scenario

In this scenario we assume that a customer has two data centers and wants to build a switched fabric that spans the data centers. The customer has a mix of UNIX servers and Windows servers and wants to have only one chassis at each site. The customer also wants two redundant fabrics in order to apply an IBM/Brocade best practice, and so that a failure in one fabric will not cause total system down situations.

For this purpose we are partitioning switches building two redundant fabrics where server HBAs and storage controllers connect to separate switched fabrics. UNIX systems will access “just a bunch of disks” (JBOD) at site B and Windows systems will access IBM DS4000 at site A.

The customer has provided two ISL connections and wants these to be shared among current and future users. For this purpose we create base switches on each site which can only be used for ISL traffic. These extended ISLs are called XISLs.

Because we are only going to have two fabrics, then two ISLs might be sufficient, but making the ISLs into XISLs gives us the flexibility of having separate fabrics in the future to make use of the existing infrastructure.

To perform SAN management, we create a user with the credentials to administer *only* the hardware ports that belong to their respective SAN’s.

The customer is purchasing two Virtual Fabric capable SAN-switches; one is the IBM System Storage SAN768B Director for site A and the other is an IBM System Storage SAN80B for site B.

The Virtual Fabrics that we are building are logically shown in Figure 11-21.

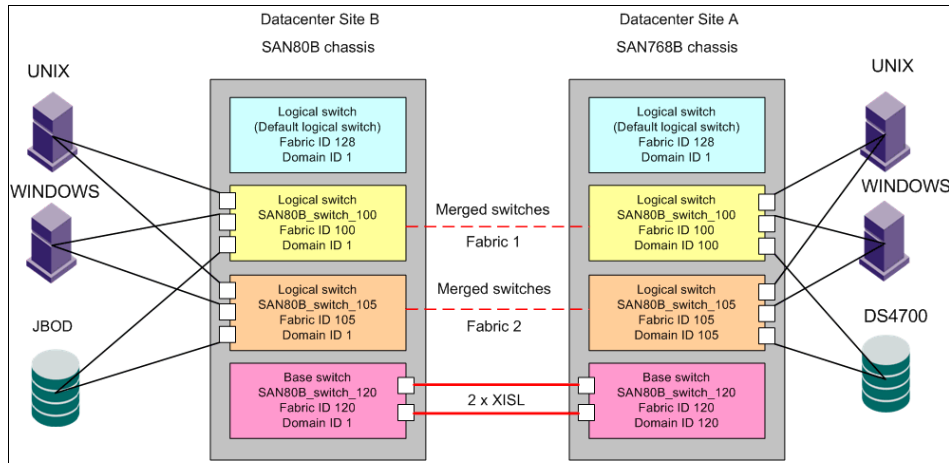


Figure 11-21 Logical setup

The Virtual Fabrics that we are building are shown physically in Figure 11-22.




Figure 11-22 Virtual Fabric Lab. Setup

11.4.2 Enabling Virtual Fabric on the switches

By default, Virtual Fabric is enabled on the VF-capable switches. However, if the system is being upgraded from Fabric OS v5.3, then Virtual Fabric is disabled and will need to be enabled.

Example 11-7 shows that VF is disabled.

Example 11-7 Virtual Fabric is disabled

```
IBM_SAN80B_4_217: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
```

Note that at this stage without Virtual Fabric enabled, Admin Domain will be available. Virtual Fabric and Admin Domain cannot work at the same time and Admin Domain will be disabled when enabling Virtual Fabric.

Figure 11-23 shows that the Admin Domain is enabled while Virtual Fabric is disabled.

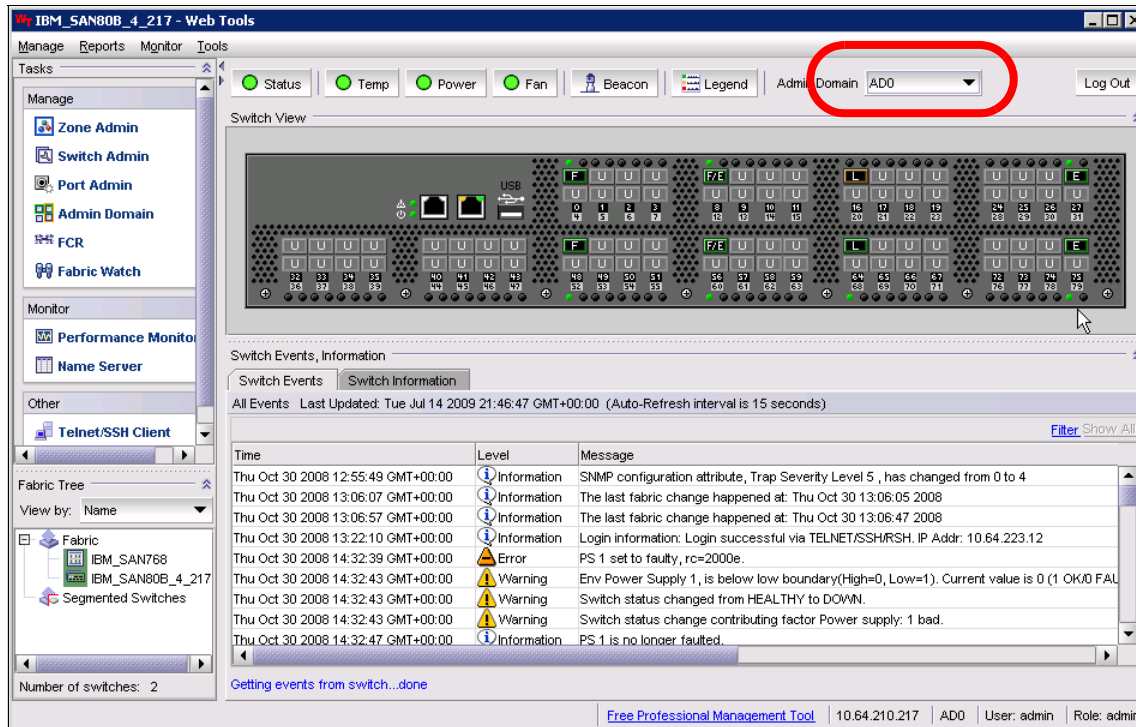


Figure 11-23 Admin Domain enabled

Example 11-8 shows VF enabled.

Example 11-8 Enabling Virtual Fabric

```
IBM_SAN80B_4_217: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.
```

This operation will be performed at both switches. Example 11-9 shows Virtual Fabric as enabled.

Example 11-9 After reboot Virtual Fabric is enabled

```
IBM_SAN80B_4_217: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
```

We now enable Virtual Fabric on both our switch chassis. At this time only the default logical switch with Fabric ID (FID) 128 exists. All ports in the two switches belong to the default switch, and because these have the same FID, the switches will merge into a single fabric. Example 11-10 shows this situation.

Example 11-10 Fabrics have merged

```
IBM_SAN80B_217:FID128:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr Name
-----
1: fffc01 10:00:00:05:1e:46:8a:00 10.64.210.210 "IBM_SAN768B_210"
2: fffc02 10:00:00:05:1e:09:97:01 10.64.210.217 "IBM_SAN80B_217"
```

The Fabric has 2 switches
COMMAND OUTPUT REMOVED FOR CLARITY

Attention: Depending on the product and FOS version, the switch ports on the newly created Virtual Fabric might be disabled or persistently disabled, and they will need to be re-enabled.

Zoning: If different zoning configurations exist on the switches that are being interconnected, they can merge if there is no zoning conflict. This can be avoided by persistently disabling the ports before enabling the Virtual Fabric.

- Command example: **portcfgpersistentdisable 3/8**
- Command example: **portcfgpersistentenable 3/8**

Later we create additional logical switches.

Figure 11-24 shows that Admin Domain is disabled when Virtual Fabric is enabled.

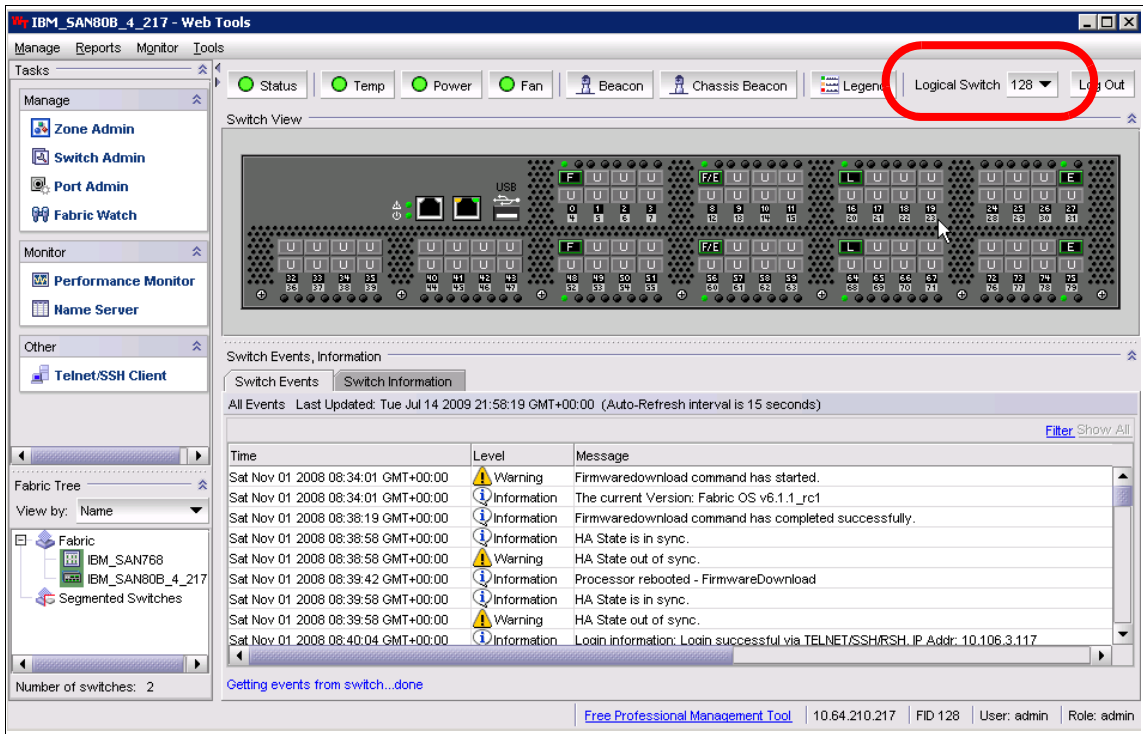


Figure 11-24 Virtual Fabric enabled

11.4.3 Creating logical switches

In this section we create six logical switches. Two of these will be base switches that are only used for carrying ISL traffic between chassis. The four remaining logical switches will be two at each site. Example 11-11 shows an example of a logical switch being created.

Example 11-11 A logical switch is created on the SAN80B switch

```
IBM_SAN80B_217:FID128:admin> lscfg --create 100
About to create switch with fid=100. Please wait...
Logical Switch with FID (100) 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.

Example 11-12 shows that all the ports belong to the default switch.

Example 11-12 All ports still belong to the default switch

```
IBM_SAN80B_217:FID128:admin> lscfg --show
```

```
Created switches: 128(ds) 100
```

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	128	128
Port	20	21	22	23	24	25	26	27	28	29
FID	128	128	128	128	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
Port	40	41	42	43	44	45	46	47	48	49
FID	128	128	128	128	128	128	128	128	128	128
Port	50	51	52	53	54	55	56	57	58	59
FID	128	128	128	128	128	128	128	128	128	128
Port	60	61	62	63	64	65	66	67	68	69
FID	128	128	128	128	128	128	128	128	128	128
Port	70	71	72	73	74	75	76	77	78	79
FID	128	128	128	128	128	128	128	128	128	128

We see that now one *additional* switch with FID 100 is present, but all resources still belong to FID 128 (the default switch).

11.4.4 Assigning ports to the newly created switch

In Example 11-13 we have logged on to the new default switch.

Example 11-13 Switch to the newly created logical switch

```
IBM_SAN80B_217:FID128:admin> setcontext 100  
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.
```

The commands we show are all issued from the switch with FID 100.

Now we want to add resources to the newly added switch. This can be done with or without the **force** option. SAN administrators might want to use the **-force** option in order to save time.

In Example 11-14 we are adding ports 0, 8, and 16 to the logical switch. These are disabled by default and need to be enabled (Figure 11-15).

Example 11-14 Ports are added to the logical switch

```
switch_100:FID100:admin> lscfg --config 100 -port 0 -force  
Making this configuration change. Please wait...  
Configuration change successful.  
Please enable your ports/switch when you are ready to continue.
```

```
switch_100:FID100:admin> lscfg --config 100 -port 8 -force  
Making this configuration change. Please wait...  
Configuration change successful.  
Please enable your ports/switch when you are ready to continue.
```

```
switch_100:FID100:admin> lscfg --config 100 -port 16 -force  
Making this configuration change. Please wait...  
Configuration change successful.  
Please enable your ports/switch when you are ready to continue.
```

```
switch_100:FID100:admin> portenable 0  
switch_100:FID100:admin> portenable 8  
switch_100:FID100:admin> portenable 16
```

Example 11-15 The switch now has 3 ports

```
switch_100:FID100:admin> switchshow
switchName:      switch_100
switchType:      64.3
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:05:1e:09:97:02
zoning:          OFF
switchBeacon:    OFF
FC Router:       OFF
Allow XISL Use:  ON
LS Attributes:   [FID: 100, Base Switch: No, Default Switch: No]
```


Area	Port	Media	Speed	State	Proto	
0	0	id	N2	Online	F-Port	10:00:00:00:c9:28:ec:1a
8	8	id	N2	Online	F-Port	10:00:00:00:c9:32:a8:65
16	16	id	N2	Online	L-Port	9 public

At this point the switch must be disabled to set a unique Domain ID and to allow the use of XISL (allow is the default setting). Domain IDs and switchnames will be configured as shown in Figure 11-25.

Switchname	SAN80B			SAN768B		
	FID 100	FID 105	FID 120	FID 100	FID 105	FID 120
	Domain ID	Domain ID	Domain ID	Domain ID	Domain ID	Domain ID
SAN80B_switch_100	1					
SAN80B_switch_105		1				
SAN80B_switch_120			1			
SAN768B_switch_100				100		
SAN768B_switch_105					105	
SAN768B_switch_120						120

Figure 11-25 Switch names and Domain IDs

Attention: If switch Domain IDs are not unique, the switches will segment, and a fabric merge will not happen.

Example 11-16 shows how we are setting the Domain ID.

Example 11-16 Setting the Domain ID

```
switch_100:FID100:admin> configure
```

Configure...

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

```
Domain: (1..239) [1] 100  
Allow XISL Use (yes, y, no, n): [yes]  
Enable a 256 Area Limit  
  (0 = No,  
   1 = Zero Based Area Assignment,  
   2 = Port Based Area Assignment): (0..2) [0]  
R_A_TOV: (4000..120000) [10000]  
E_D_TOV: (1000..5000) [2000]  
WAN_TOV: (0..30000) [0]  
MAX_HOPS: (7..19) [7]  
Data field size: (256..2112) [2112]  
Sequence Level Switching: (0..1) [0]  
Disable Device Probing: (0..1) [0]  
Suppress Class F Traffic: (0..1) [0]  
Per-frame Route Priority: (0..1) [0]  
Long Distance Fabric: (0..1) [0]  
BB credit: (1..27) [16]  
Disable FID Check (yes, y, no, n): [no]
```

```
Insistent Domain ID Mode (yes, y, no, n): [no]  
Virtual Channel parameters (yes, y, no, n): [no]  
F-Port login parameters (yes, y, no, n): [no]  
Zoning Operation parameters (yes, y, no, n): [no]  
RSCN Transmission Mode (yes, y, no, n): [no]  
Arbitrated Loop parameters (yes, y, no, n): [no]  
System services (yes, y, no, n): [no]  
Portlog events enable (yes, y, no, n): [no]  
ssl attributes (yes, y, no, n): [no]  
rpcd attributes (yes, y, no, n): [no]  
webtools attributes (yes, y, no, n): [no]
```

WARNING: The domain ID will be changed. The port level zoning may be affected

Now enable the switch using the command **switchenable**.

At this time, using the same steps as before, we create one additional logical switch on the SAN80B, and we create two new logical switches on the SAN768B switch. We assign ports to the switches as indicated in Figure 11-26.

Device	SAN80B			SAN768B		
	FID 100	FID 105	FID 120	FID 100	FID 105	FID 120
	Fabric 1	Fabric 2	Base sw	Fabric 1	Fabric 2	Base sw
WIN-A HBA1				3/15		
WIN-A HBA2					10/15	
UNIX-A HBA1				3/14		
UNIX-A HBA2					10/14	
DS4000 ctrl 1				3/0		
DS4000 ctrl 1					10/0	
ISL 1			27			3/8
ISL 2			75			10/8
WIN-B HBA1	0					
WIN-B HBA2		48				
UNIX-B HBA1	8					
UNIX-B HBA2		56				
JBOD ctrl 1	16					
JBOD ctrl 2		64				

Figure 11-26 Switch connections

At this point we have created four switches with individual names and unique Domain IDs. The switches are configured to allow XISL.

11.4.5 Creating the base switch

The base switch is used for Inter Switch Link traffic called XILS only. We assign our two current ISL connections to the base switch, and because our logical switches are configured for XISL use, the switches will merge into a single fabric.

In Example 11-17, before creating the base switch, the other individual logical switches are shown as stand-alone switches.

Example 11-17 The logical switches are stand alone switches

```
SAN80B_switch_100:FID100:admin> fabricshow
Switch ID    Worldwide Name          Enet IP Addr Name
-----
1: fffc01 10:00:00:05:1e:09:97:02 10.64.210.217 "SAN80B_switch_100"

COMMAND OUTPUT REMOVED FOR CLARITY
```

Now we create the base switches on both of the chassis in our setup. Example 11-18 shows how it is done on the SAN80B.

Switches: The logical switches can be created from any of the other switches, as long as the user has administrator privileges on the switch chassis.

Example 11-18 The base switch is created and ports assigned

```
IBM_SAN80B_217:FID128:admin> lscfg --create 120 -base -force  
About to create switch with fid=120. Please wait...  
Logical Switch with FID (120) 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.

```
IBM_SAN80B_217:FID128:admin> lscfg --config 120 -port 27  
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.
```

```
IBM_SAN80B_217:FID128:admin> lscfg --config 120 -port 75 -force  
Making this configuration change. Please wait...  
Configuration change successful.  
Please enable your ports/switch when you are ready to continue.
```

After enabling the switch ports, the base switch as well as the remaining switches will merge into logical fabrics. The switches that merge will be the ones that have similar FIDs. A logical switch cannot merge with another logical switch if the FIDs are different. Example 11-19 shows that the switches have merged.

Example 11-19 Switches now merge successfully

```
SAN80B_switch_100:FID100:admin> switchshow  
switchName:    SAN80B_switch_100  
switchType:    64.3  
switchState:    Online  
switchMode:    Native  
switchRole:    Principal  
switchDomain:    1  
switchId:    fffc01  
switchWwn:    10:00:00:05:1e:09:97:02  
zoning:    OFF  
switchBeacon:    OFF  
FC Router:    OFF
```

```
Allow XISL Use: ON
LS Attributes: [FID: 100, Base Switch: No, Default Switch: No]

Area Port Media Speed State      Proto
=====
  0  0  id    N2    Online        F-Port  10:00:00:00:c9:28:ec:1a
  8  8  id    N2    Online        F-Port  10:00:00:00:c9:32:a8:65
 16 16  id    N2    Online        L-Port  9 public
 80 80  --    --    Online        E-Port  10:00:00:05:1e:46:8a:01
"SAN768B_switch_100" (downstream)

SAN80B_switch_100:FID100:admin> fabricshow
Switch ID    Worldwide Name      Enet IP Addr Name
-----
  1: fffc01 10:00:00:05:1e:09:97:02 10.64.210.217 >"SAN80B_switch_100"
100: fffc64 10:00:00:05:1e:46:8a:01 10.64.210.210 "SAN768B_switch_100"

The Fabric has 2 switches

COMMAND OUTPUT REMOVED FOR CLARITY
```

11.4.6 Creating a user to manage the Virtual Fabric

In this section we create a user to manage the Virtual Fabrics that we created in the previous sections. The user will have administrator privileges to perform management on Virtual Fabrics FID100 and FID105, but will not have any privileges to perform management on the switch chassis.

We use WebTools to perform this task.

Figure 11-27 shows the admin user privileges for all Virtual Fabrics.

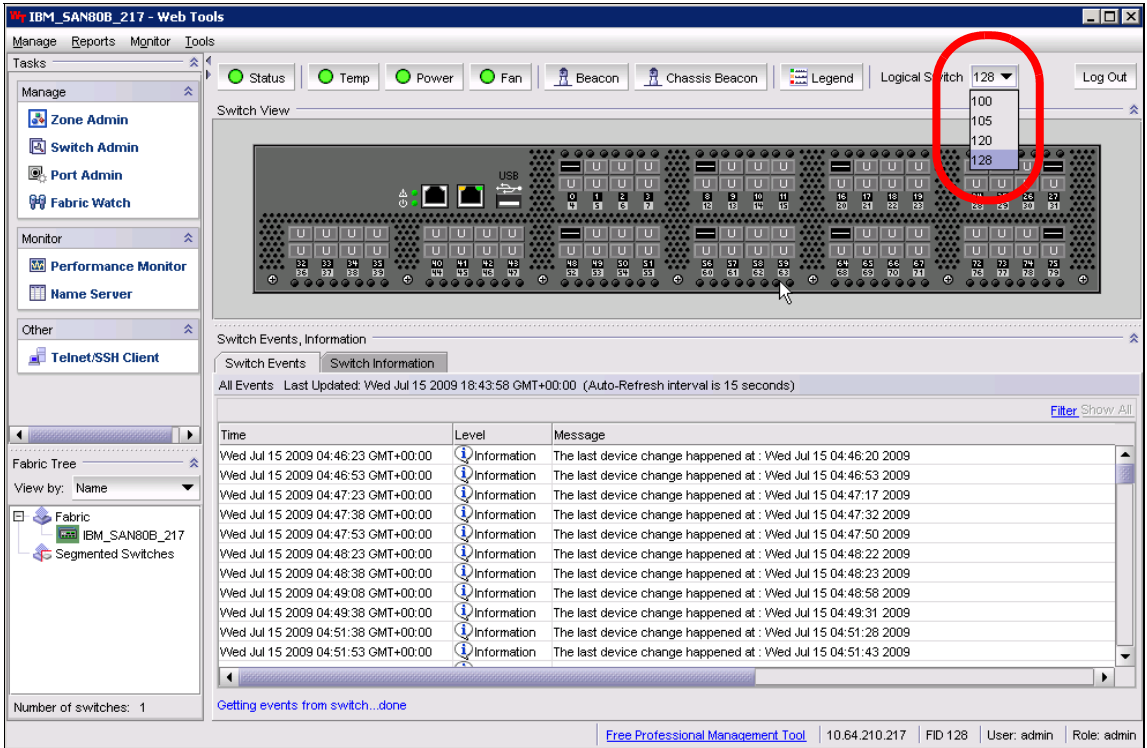


Figure 11-27 Logged in as admin user

We create a user named “Alex” with permission to only manage FID 100 and 105.

Figure 11-28 shows how the user Alex is created.

Logical Fabric ID	User Role
100	admin
101	No Access
102	No Access
103	No Access
104	No Access
105	No Access
106	No Access

Home Logical Fabric Id: 100

Chassis Access Role: No Access

Figure 11-28 Create user Alex

After applying the new user, Alex will be able to log in to the switch. Alex will have admin rights for FID 100 and FID 105 and will be able to only perform management tasks on resources that are applied to the user Alex.

Figure 11-29 shows which ports the user Alex can manage.

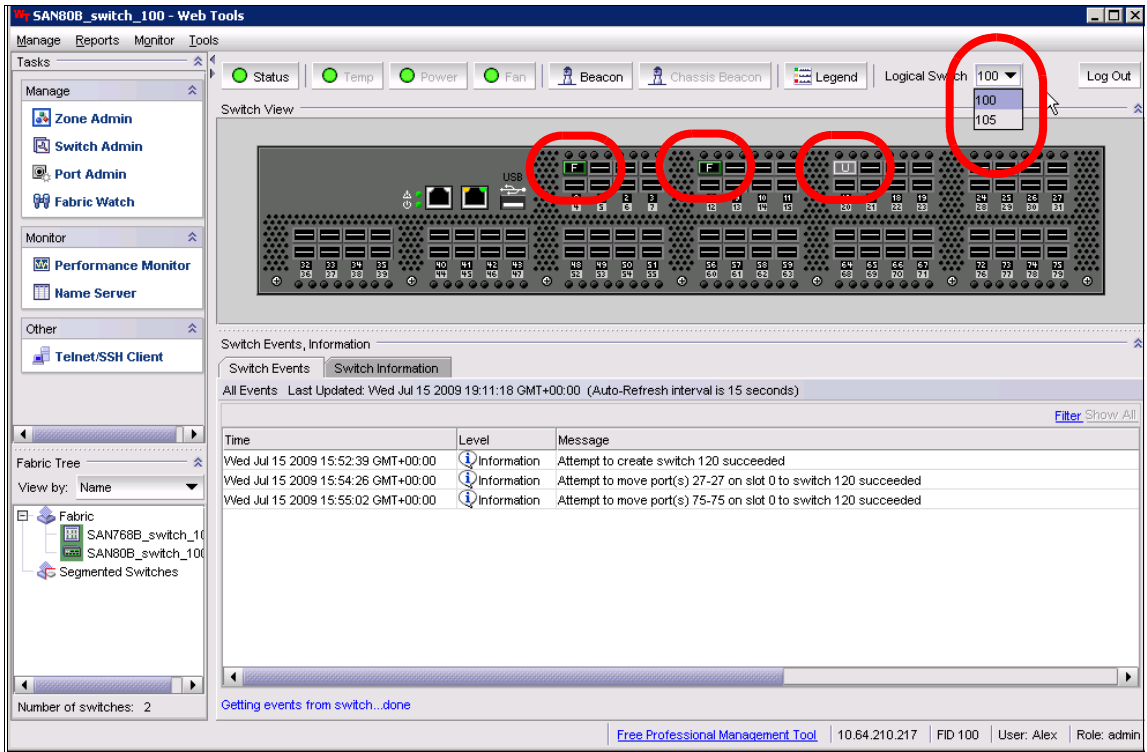


Figure 11-29 Logged in as user Alex

The next step for user Alex is to create zones for each server HBA. Zoning is covered in greater detail in Chapter 12, “Basic zoning” on page 513.

We have now finished creating Virtual Fabrics and have a working setup where SAN switches have been partitioned to separate individual fabrics.

Our scenario shows how to add just a few ports to a logical switch. In a real life situation, it is more likely that an entire switch blade will be used as a logical switch, and a separate switch blade will be used for ISL connections.



Basic zoning

In this chapter, we discuss the basics of zoning, which allows you to define specific groups of fabric-connected devices to ensure that the access between them is controlled. Zoning enables you to partition a storage area network (SAN) into logical groups of devices that can access each other.

12.1 Zoning in general

Zoning allows you to define specific groups of fabric-connected devices to ensure that the access between them is controlled.

Attention: Be aware that any devices that are not configured in a zone will not be accessible.

Zoning enables you to partition a storage area network (SAN) into logical groups of devices that can access each other. Zoning is critical even in Fabrics with storage based LUN masking. Often storage based LUN masking is viewed as a replacement for zoning, however this is *not* the case. In a heterogeneous server environment, zoning is another layer of security to existing storage LUN masking. It allows administrators to secure not just their storage, but also allows them to isolate servers and even adapter cards from each other.

Zones can be configured dynamically. They can vary in size, depending on the number of fabric-connected devices, and devices can belong to more than one zone. Because zone members can access only other members of the same zone, a device not included in a zone is not available to members of that zone.

12.1.1 Mixed fabrics

When using a mixed fabric—that is, a fabric that contains two or more switches running different fabric operating systems— use the switch with the highest Fabric Operating System (Fabric OS) level to perform zoning tasks. If the switch is running Fabric OS v6.0.x or earlier, it must have an Advanced Zoning license enabled.

If the fabric includes a third-party switch product, only worldwide name (WWN) zoning is supported. Other types of zoning, including QuickLoop, are not supported.

When zone or Fabric Assist (FA) zone members are specified by fabric location only (domain or area), or by device name only (node name or port WWN), zone boundaries are enforced at the hardware level, and the zone is referred to as a *hard zone*.

When zone members are specified by fabric location (domain or area) and other members of the same zone are specified by device name (node name or port WWN), zone enforcement depends on Name Server lookups, and the zone is referred to as a *soft zone*.

12.1.2 Zone configurations

A zone configuration is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is in effect, all zones that are members of that configuration are in effect.

Several zone configurations can reside on a switch at once, and you can quickly alternate between them. For example, you might want to have one configuration enabled during the business hours and another enabled overnight. However, only one zone configuration can be enabled at a time.

The different types of zone configurations are as follows:

- ▶ **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.

Important: Ensure that only one person is making configuration changes to your environment at any one time. Using the `killtelnet` command provides a view of who is logged in to the switch and a method for removing any sessions that should not be in place:

```
BDPOC01L01:admin> killtelnet
Collecting login information....
List of sessions (2 found)
```

Session No	USER	TTY	IDLE	LOGIN@	FROM
0	admin0	pts/0	1:50	17:52	9.155.66.103
1	admin0	pts/1	0.00s	18:53	9.155.66.205

Enter Session Number to terminate (q to quit)

12.2 Zoning using DCFM

In DCFM, one way to launch the Zone Admin is by right-clicking a fabric in the View panel, as shown in Figure 12-1, and selecting **Zoning** → **Fabric**.

Attention: Zone Admin displays only if an Advanced Zoning licence is installed on the switch.

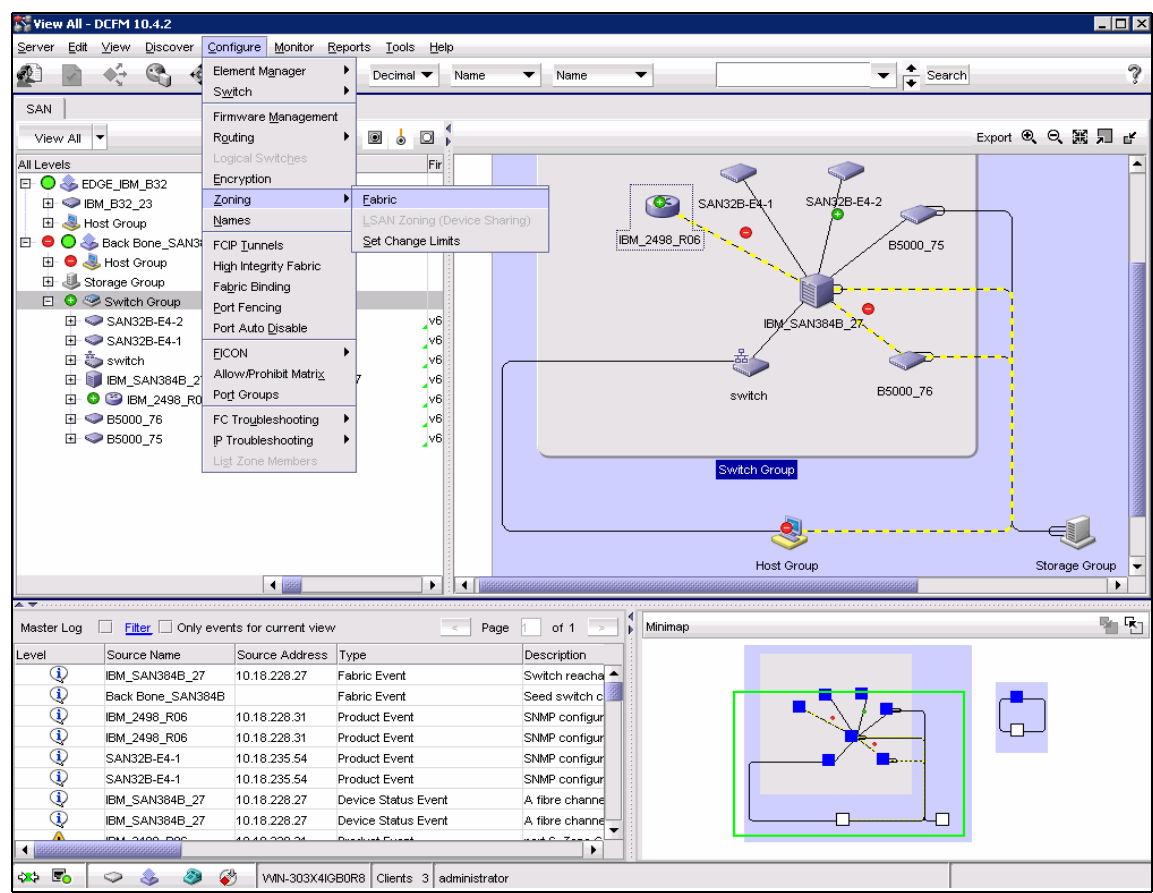


Figure 12-1 Opening Zone Admin

The *Zoning dialog box* is used to set up, maintain, and activate the zones across the fabric. From here, you can also define aliases for members in a zone and create the zones that form the active configuration across the fabric, as shown in Figure 12-2.

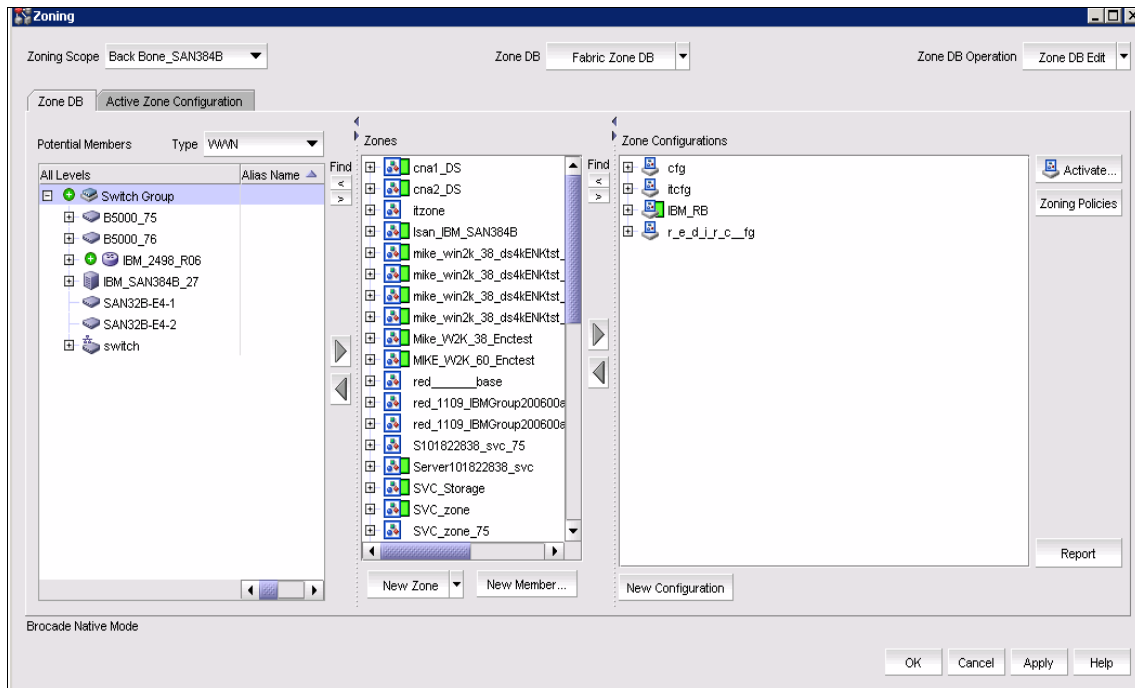


Figure 12-2 Zoning dialog box

Administrative privileges are required to access this function. When administering zoning on an IBM System Storage SAN Switch, follow these steps:

1. Define aliases for device WWPN.
2. Define zones to establish groupings.
3. Add zone members (using the aliases you defined in Step 1).
4. Place zones into one or more zone configurations.
5. Enable one of the zone configurations (only one can be enabled at a time).

Zoning: When configuring zones for encryption, Alias zoning is not supported in containers. You must use the real WWPN for the zoning configuration.

You can choose how zoning elements are displayed in the Zoning dialog box. The zoning view that you select determines how members are displayed in the Alias Selection List. The views filter the fabric and device information that is displayed in the Selections for the selected view, making it easier for you to create and modify zones, especially when creating hard zones as in Figure 12-3.

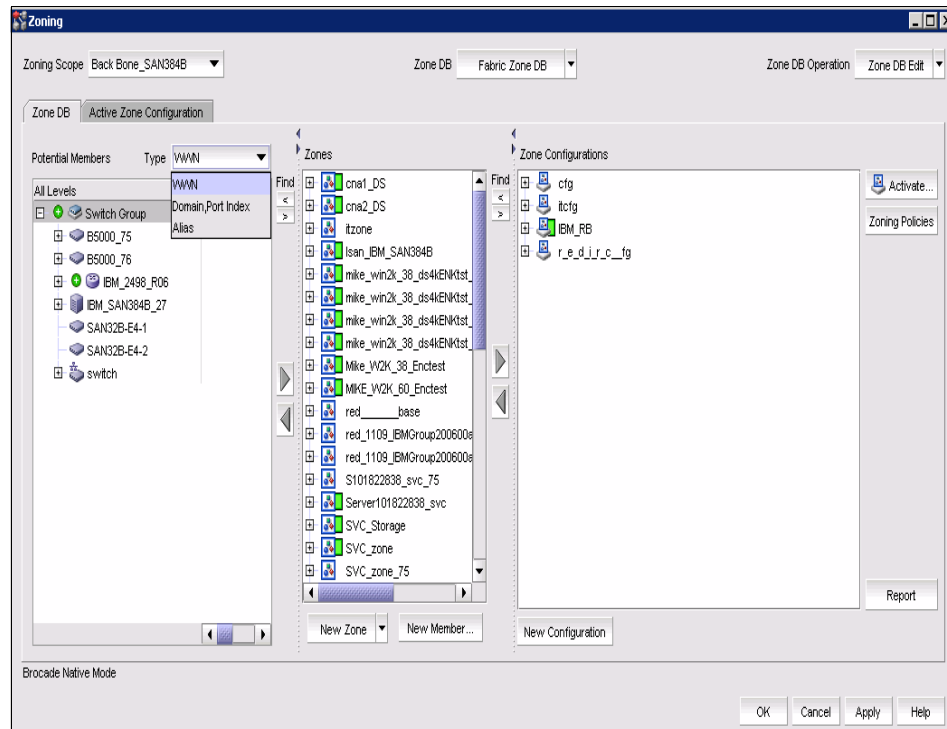


Figure 12-3 Zoning dialog box view

12.2.1 Administrative Domains

An Administrative Domain (*Admin Domain* or *AD*) is a logical grouping of fabric elements that defines what switches, ports, and devices you can view and modify. An Admin Domain is a filtered administrative view of the fabric.

Authorization: To manage Admin Domains, you must be a physical fabric administrator. A physical fabric administrator is a user with the admin role and access to all Admin Domains (AD0 through AD255). Only a physical fabric administrator can perform Admin Domain configuration and management.

Admin Domains permit access to a configured set of users. Using Admin Domains, you can partition the fabric into logical groups and allocate administration of these groups to different user accounts so that these accounts manage only the Admin Domains that are assigned to them and do not make changes to the rest of the fabric.

For example, you can put all the devices in a particular department in the same Admin Domain for ease of managing those devices. If you have remote sites, you can put the resources in the remote site in an Admin Domain and assign the remote site administrator to manage those resources.

Admin Domains and zones: Do not confuse Admin Domains with zones:

- ▶ *Zones* define which devices and hosts can communicate with each other.
- ▶ *Admin Domains* define which users can manage which devices, hosts, and switches.

You can have up to 256 Admin Domains in a fabric (254 user-defined and two system-defined), numbered from 0 through 255. Admin Domains are designated by a name and a number. In this book, we refer to specific Admin Domains using the format AD*n*, where *n* is a number between 0 and 255.

The Admin Domain is used mainly in terms of fabric administration perspective. Each Admin Domain has its own zone database, with both defined and effective zone configurations and all related zone objects (zones, zone aliases, and zone members). Within an Admin Domain, you can configure zoning only with the devices that are present in that Admin Domain (direct members).

If you upgrade a fabric to Fabric OS v5.2.0 or higher, the zone database from the pre-v5.2.0 fabric is referred to as the *root zone database* and is owned by Admin Domain 0 (AD0). Each zone database has its own name space.

Fabric OS v6.1 adds support for distributing the Defined Zone Configuration database in InteropMode 2.

12.2.2 Implementing Administrative Domains

If you implement Admin Domains, you must set the default zoning mode to **No Access** before you create Admin Domains. To do so, click **Zoning Policies** (see the rounded rectangle in Figure 12-4).

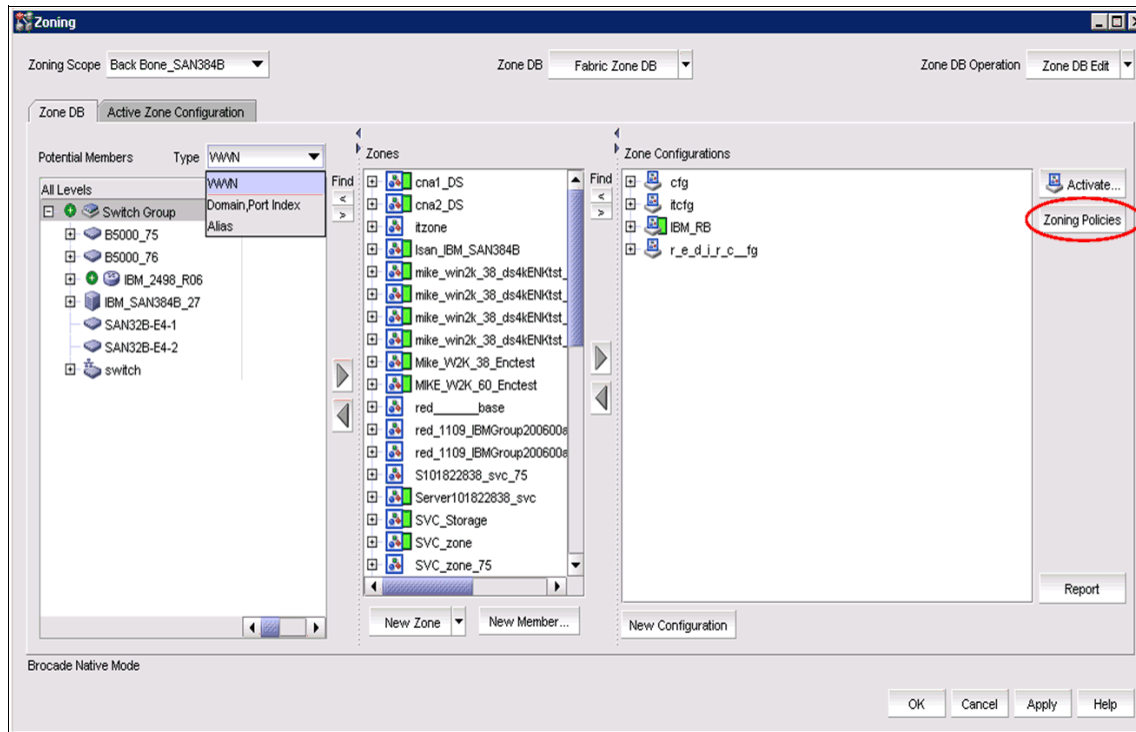


Figure 12-4 Disable default zoning mode

The Zoning Policies dialog box displays. Click **Disable** (No Access) as shown in Figure 12-5 and click **OK**.




Figure 12-5 Zoning Policies

12.3 Implementing zoning

In this section, we provide details about how to manage zoning.

12.3.1 Managing zoning

You can monitor and manage zoning with the Web Tools Zone Admin, CLI and with DCFM. The information is collected from the selected switch.

In DCFM you must be logged into with a user name that has the following privileges:

- ▶ Zoning Activation
- ▶ Zoning Online
- ▶ Zoning Offline

The Role, *System Administrator*, for example, has these privileges; see Figure 12-6.

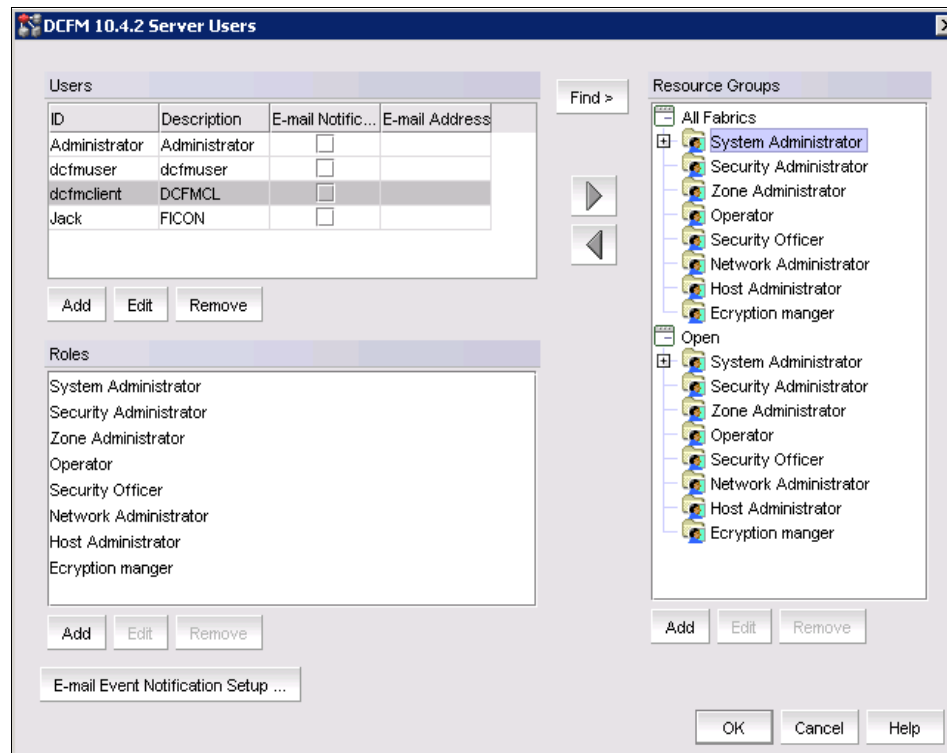


Figure 12-6 DCFM User Management

All other roles only allow *view* or *read-only* access. Most of the zoning operations are disabled in *read-only* mode. A user can be set up for zone administration, such as the zone administrator, see Figure 12-7 which shows the rules set up for this user.

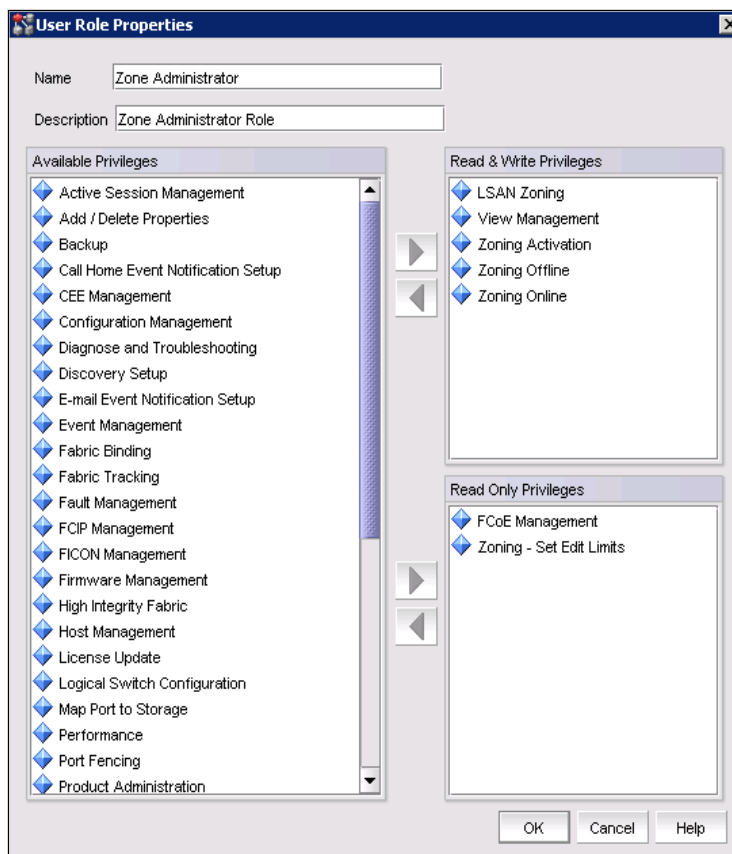


Figure 12-7 Zone Administrator

12.3.2 Creating an alias

By defining an alias to a port (or ports) or WWN (or WWNs), you can simplify your understanding of the device that you are working with on the other tabs. Using a sensible naming convention also assists with troubleshooting at a later stage, making it easier to find specific devices, especially when a SAN grows in complexity. Assign aliases and ensure that they are maintained to help to identify SAN components correctly using the Alias tab.

Methods for creating an alias

You can specify members of an alias using the following methods:

- ▶ A switch domain and port index number pair (for example, 2 and 20)
- ▶ Device node and device port WWNs

In this section we describe methods for creating an alias using DCFM.

Important: When configuring zones for encryption, alias zoning is not supported in containers. You must use the real WWPN in order for frame redirection to be applied, regular zones for hosts and targets must be defined in the effective configuration. Hosts and targets must be zoned together by worldwide port name (WWPN) rather than worldwide node name (WWNN) in configurations where frame redirection will be used. If hosts or targets are zoned together using worldwide node name, frame redirection will not occur properly.

Using DCFM to create an alias

To use DCFM to create an alias, follow these steps:

1. From the DCFM Main window, click **Configure** → **Zoning** → **Fabric**. to open the Zoning dialog box,
2. Select Alias in the **Type** pull-down menu, as shown in Figure 12-8.

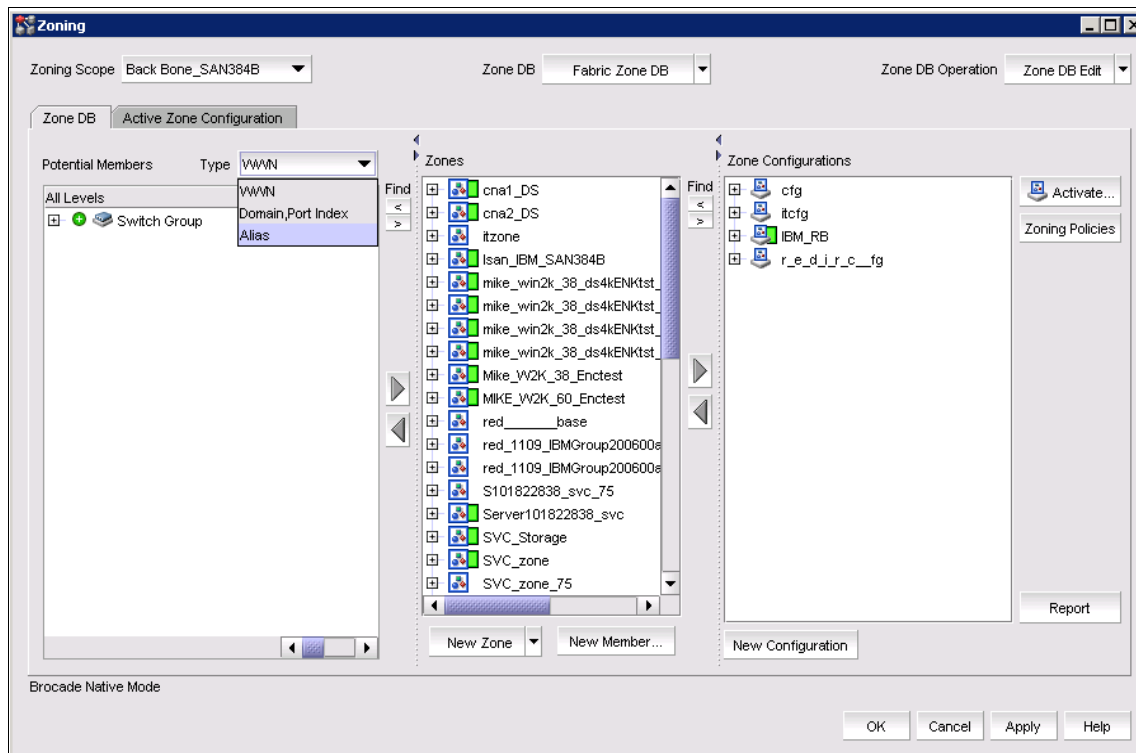


Figure 12-8 Choose display format

3. Click **New Alias** (see Figure 12-9).

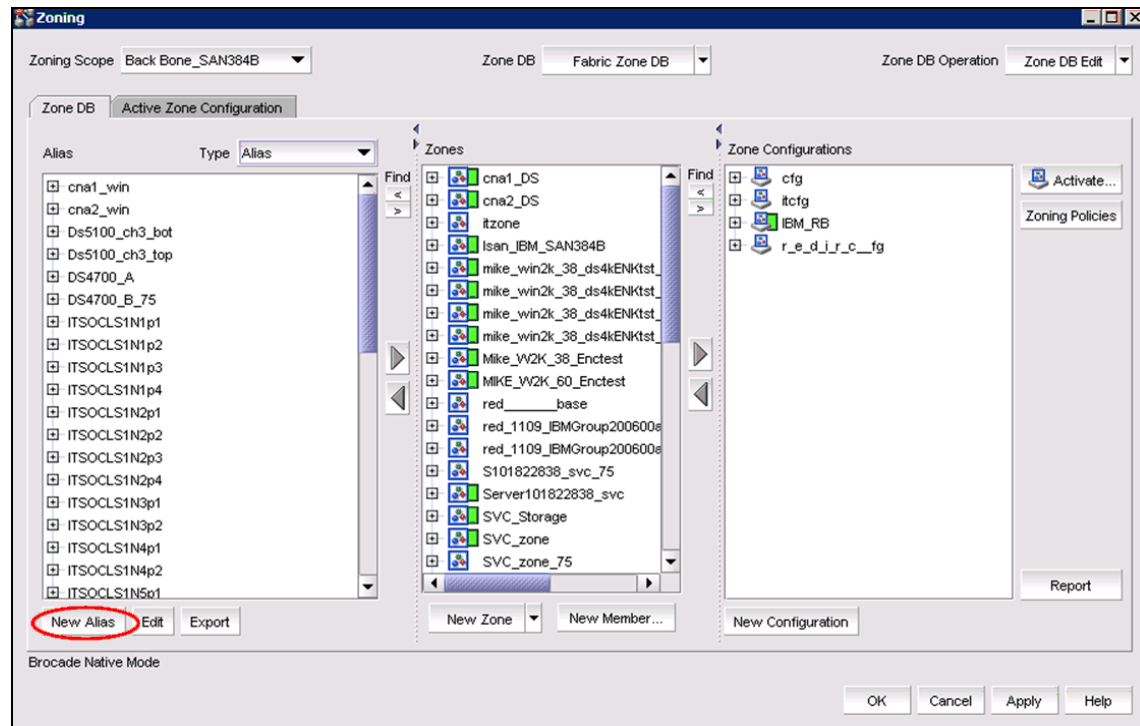


Figure 12-9 New Alias

4. The New Alias dialog box displays, as shown in Figure 12-10.

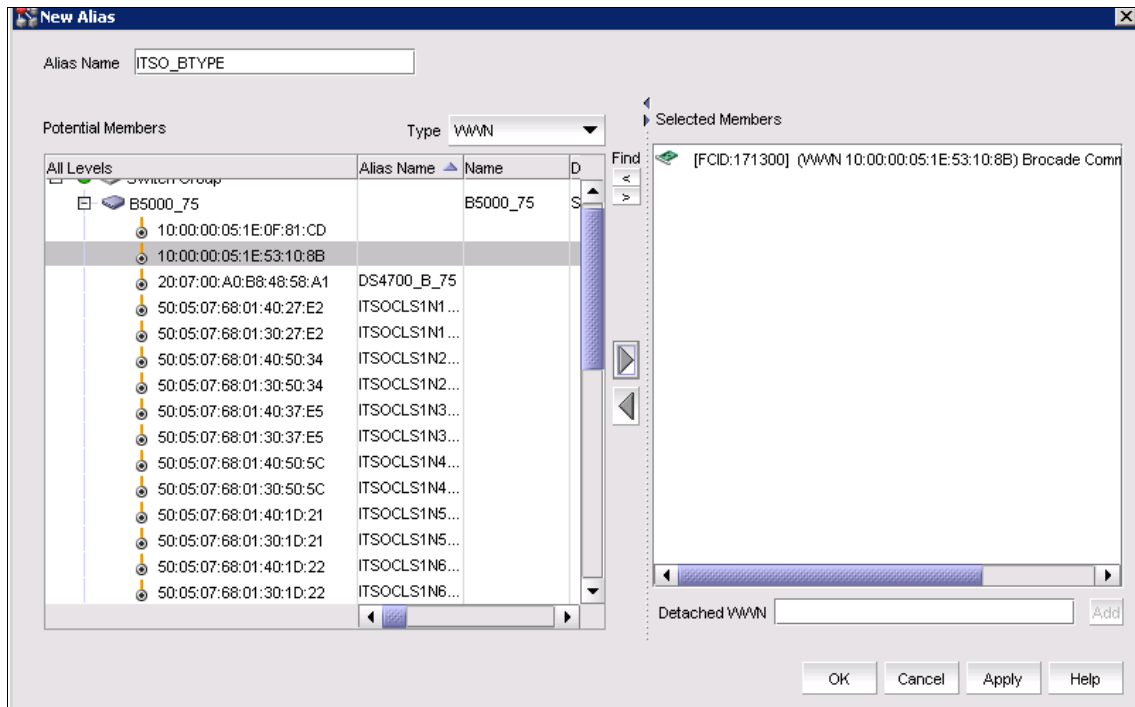


Figure 12-10 New Alias dialog box

Enter a name for the new alias, and click the WWN you want to attach to the alias. Click the right arrow to move the WWN into the Selected Members section.

If the device you want to alias is not connected to the fabric jet and is therefore not visible in the Potential Members section, you can also manually type in the WWN of the device in the Detached WWN field on the right bottom and click **Add**. Click **OK**.

The new alias displays in the Alias section of the Zoning dialog box (see the rounded rectangle in Figure 12-11).

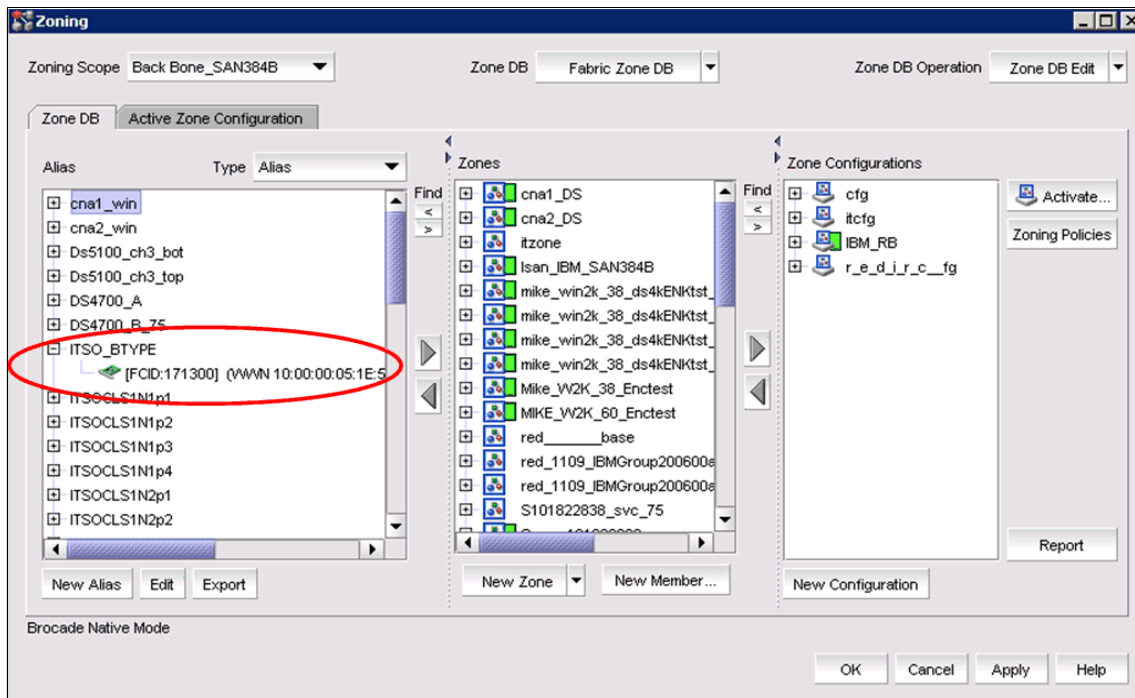


Figure 12-11 Zoning dialog box with new alias

5. Click **Apply** or **OK** or repeat Steps 3-4 to add additional aliases. Save without enabling.
6. When this is completed, click **Apply** and/or **OK**. DCFM will save the changes in the fabric without activating it in the active config (Figure 12-12).

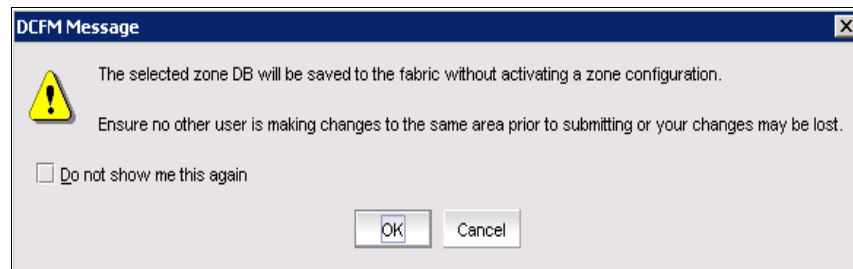


Figure 12-12 DCFM Message box

12.3.3 Creating a zone

A *zone* is a region within the fabric in which specified devices can communicate. A device can communicate only with other devices that are connected to the fabric within its specified zone.

You use the Zoning dialog box to create and manage zones. A zone can have one or multiple members and can include ports, WWNs, aliases, AL_PAs, or Quickloop.

Support: Quickloop is no longer supported from v4.4.x Fabric OS onwards.

Specifying members of a zone

You can specify members of a zone using the following methods:

- ▶ Alias names
- ▶ Switch domain and port index number pair (for example, 2 and 20)
- ▶ WWN (device)

Important:

- ▶ Create individual zones of each host to the disk storage subsystems. Also, hosts need a separate HBA for tape communication and, again, must be in another individual host/tape zone.
- ▶ Best practice is to have only one initiator (host HBA) in a zone, unless there is a specific requirement, such as encryption zones.
- ▶ This small granularity of zoning removes unnecessary PLOGI activity from host to host, as well as removing the risk of issues caused by a faulty HBA affecting others.

Using DCFM to create a zone

Follow these steps to create a zone with DCFM:

1. From the DCFM Main window, click **Configure** → **Zoning** → **Fabric** to open the Zoning dialog box.
2. Select Alias in the **Type** pull-down menu.
3. Go to the Zone tab, and click **New Zone**, as shown in Figure 12-13.

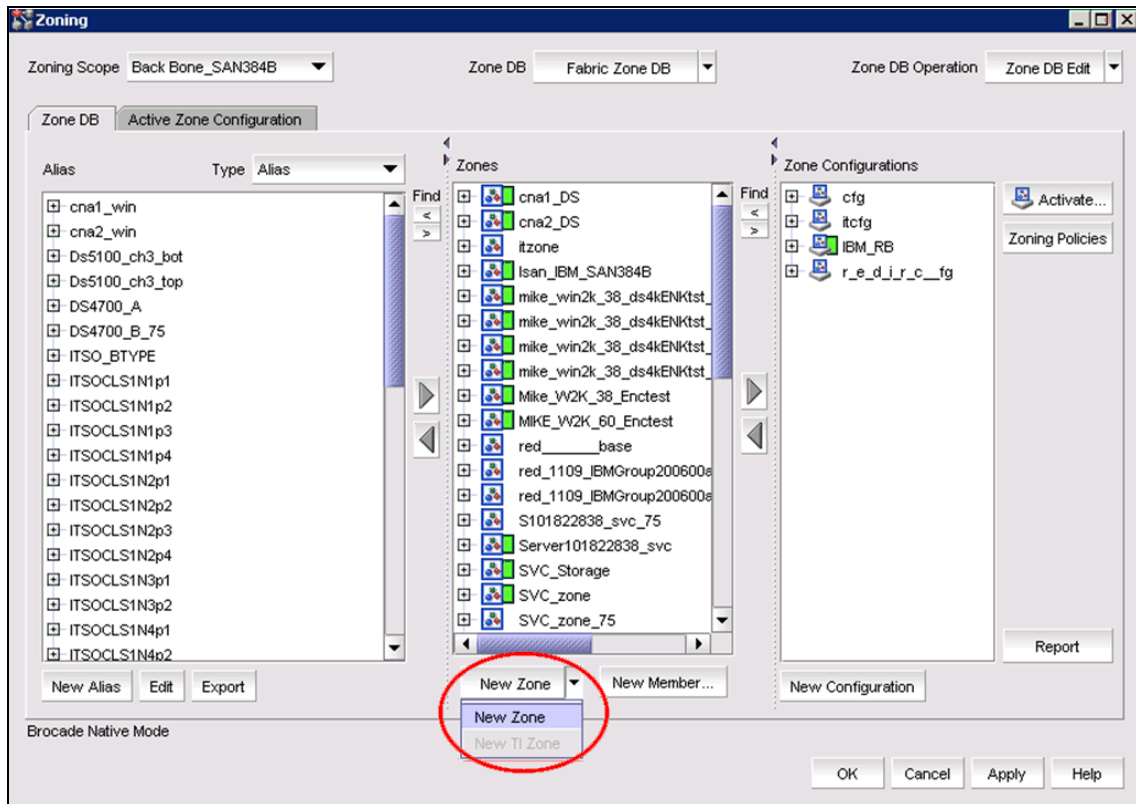


Figure 12-13 New Zone

4. Enter a name for the new zone, and click outside of the naming field
Figure 12-14. The new zone is displayed in the Zones section.

LSAN: If you are creating an LSAN zone, the zone name must begin with the letters, LSAN_.

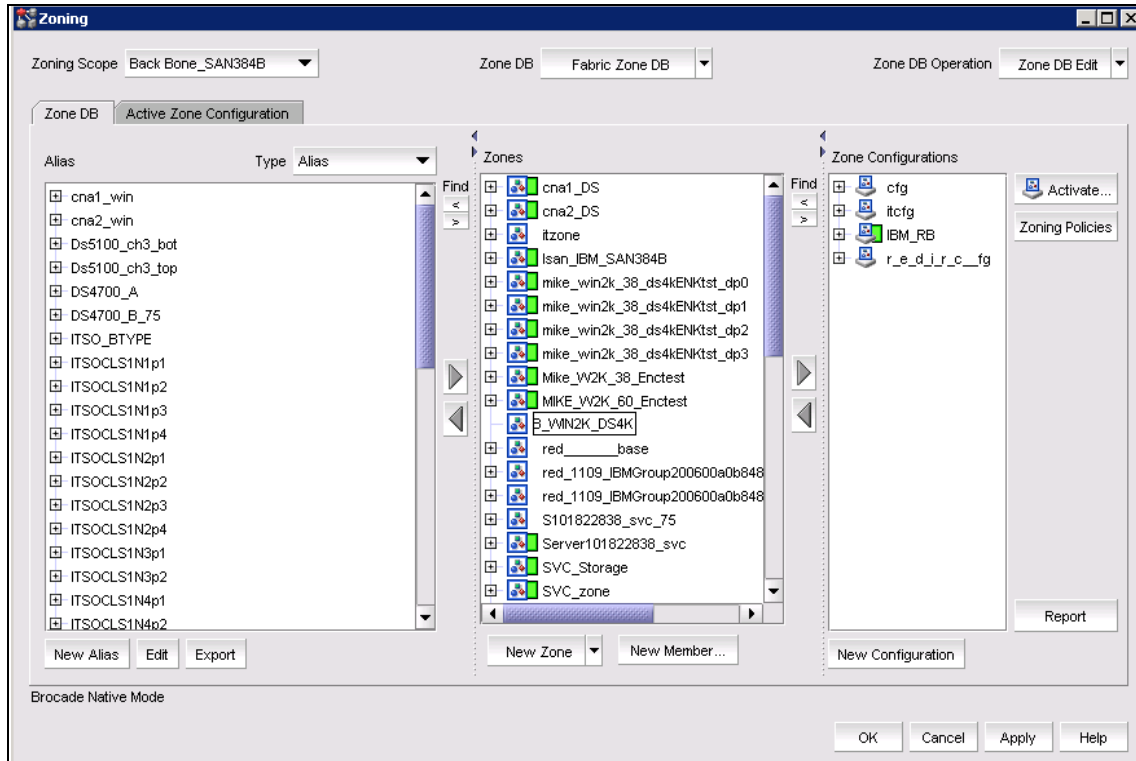


Figure 12-14 Create new zone name

5. Now highlight the aliases that you want to have in the zone and move these aliases into the new zone by clicking the right arrow. See Figure 12-15.
In the example we add a host **ITSO_BTTYPE** and a target **DS4700_A** to our new zone called **ITSO_SANB_WIN2K_DS4K**.




Figure 12-15 Move aliases to a zone

6. The Zones Section now has the aliases in the zone **ITSO_SANB_WIN2K_DS4K**. See Figure 12-16.

Repeat step 5 to add more aliases to the zone if required.




Figure 12-16 Aliases in zone

7. To save the configuration changes, click **Apply** or **OK**, a DCFM Message window opens that warns you that the changes made will only be saved to the fabric. See Figure 12-17.

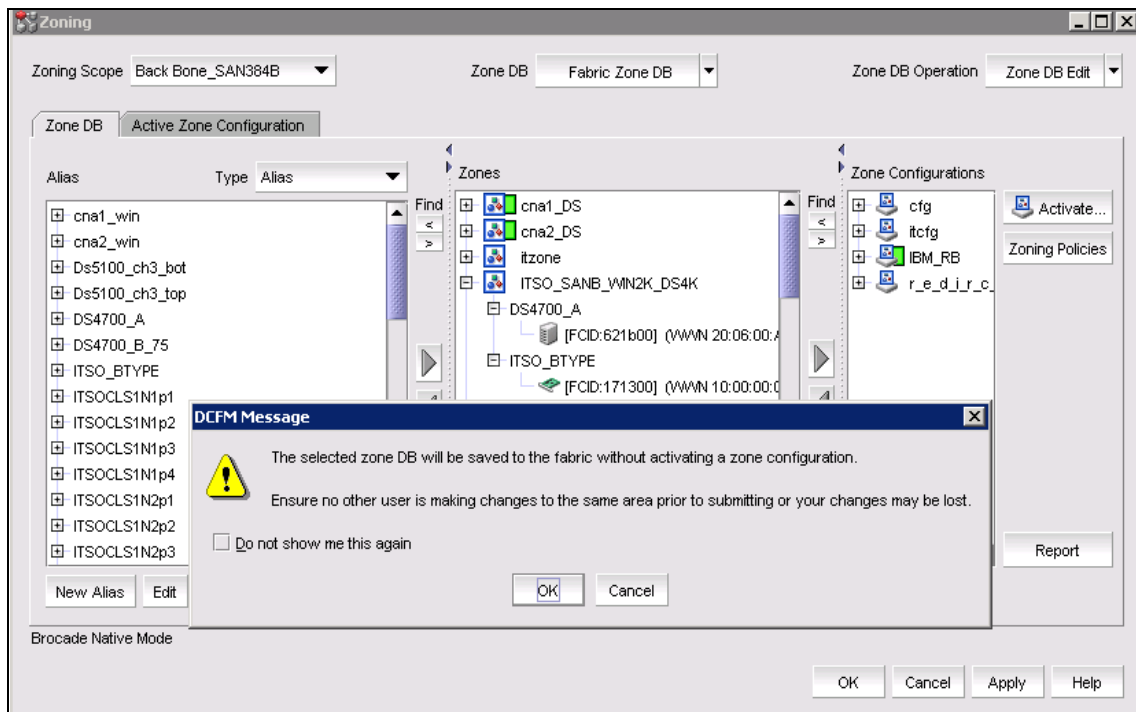


Figure 12-17 Save the configuration

12.3.4 Creating a zone configuration

To create a zone configuration:

1. From the DCFM Main window, click **Configure** → **Zoning** → **Fabric** to open the Zoning dialog box
2. Click **New Config** and name the new config in the Zone Configs section (see Figure 12-18).

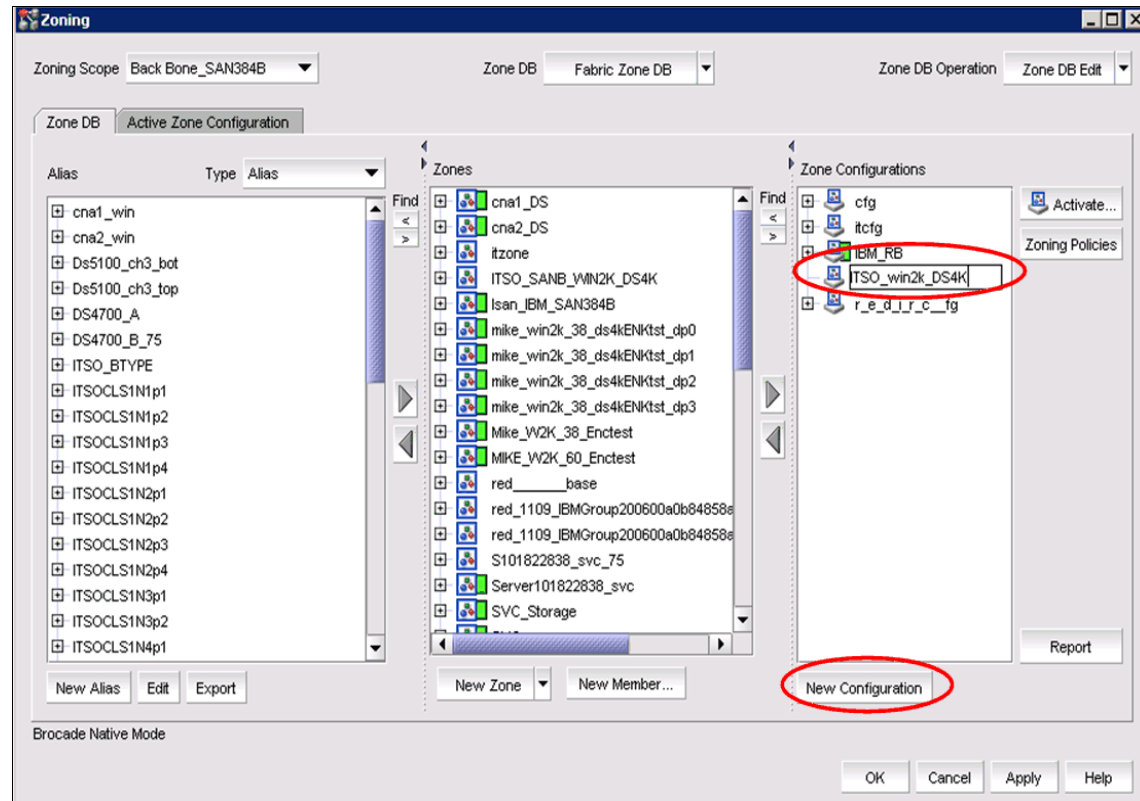


Figure 12-18 Name the new config

3. Highlight the Zone or Zones that you want to move to the ITSO_Backup_cfg and move the Zone or Zones by clicking the right arrow. See Figure 12-19.




Figure 12-19 Move zone to the new config

4. Click **OK** or **Apply** to save the config to the fabric without activating it.

12.3.5 Enabling zone configurations

Several zone configurations can reside on a switch at the same time, and you can alternate between them quickly. For example, you might want to have one zone configuration enabled during the business hours and another enabled for backups overnight. However, only one zone configuration can be enabled at a time.

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

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

To enable a zone configuration, follow these steps:

1. From the DCFM Main window, click **Configure** → **Zoning** → **Fabric** to open the Zoning dialog box.
2. To activate a configuration, highlight the configuration in the Zone Config section and click **Activate...** See Figure 12-20.

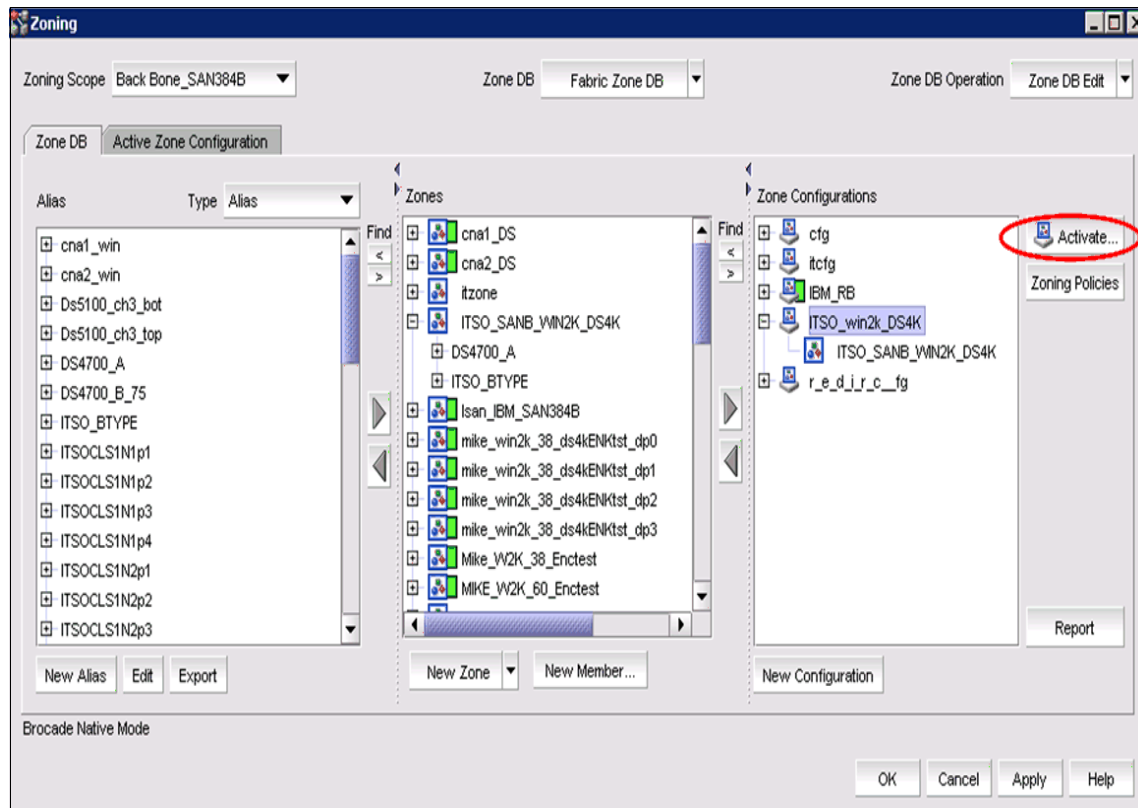


Figure 12-20 Enable Config

The Activate Zone Configuration window displays. This window gives detailed information about the changes that will take place by activating the new configuration. Carefully read the information given. To activate the new configuration, click **OK**, as shown in Figure 12-21.

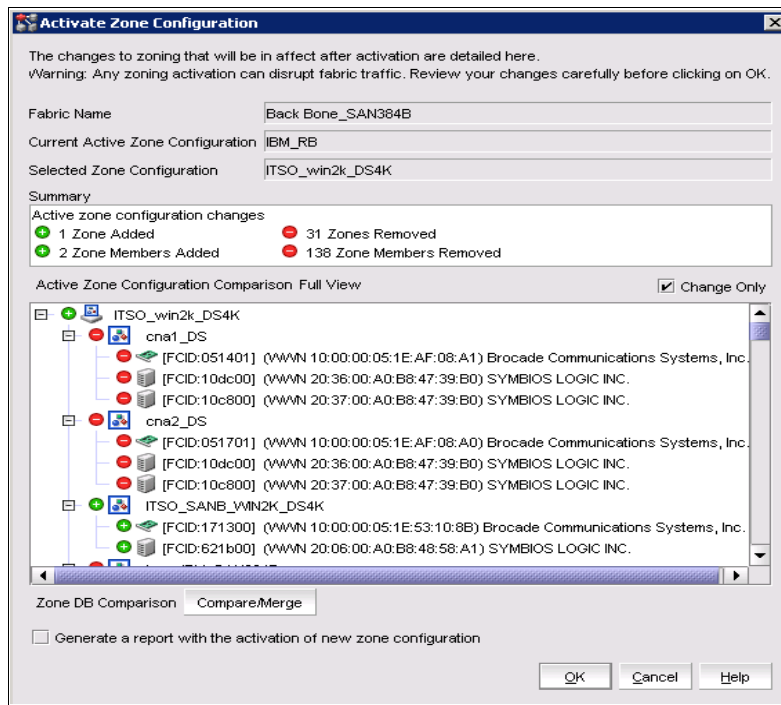


Figure 12-21 Activate Zone Configuration window

3. DCFM saves the zone database to the fabric, and enables the zone configuration that replaces the old one. A message box displays to inform you that these changes were successfully made (see Figure 12-22). Click **OK**.

Precautions:

- ▶ Remember to back up your configuration prior to making any configuration changes so that you can always get back to your starting point if there are any problems.
- ▶ Take care when enabling zone configurations. Adding new zones does not impact any currently running definitions, although removing a zone might have a large impact to the current environment.




Figure 12-22 DCFM Message

12.3.6 Adding a zone to a existing zone configuration

The process used to add a new zone into an existing zone configuration is the same as adding it to an existing zone configuration (Figure 12-23).

Select the zone that you require to add, and by selecting the active zone configuration, place this zone into the active configuration using the arrow buttons. This zone is added to the active zone configuration and will show in this configuration without the green button to show that it is not yet active.

The active zone configuration, in the example, *IBM_RB*, will no longer have the green button next to the name to indicate that there are zones not active in the configuration. Select the active zone configuration and click the **Activate** button.

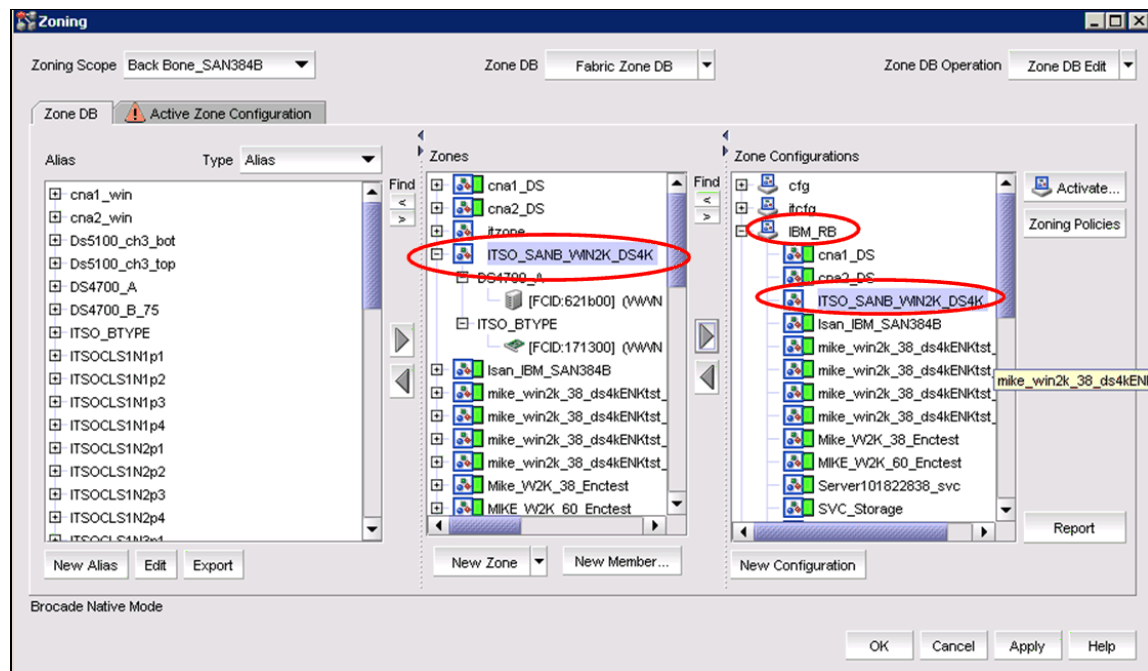


Figure 12-23 Adding zone to active zone configuration

The next window shows you what is about to change; this must only be the new zone that was added (Figure 12-24).

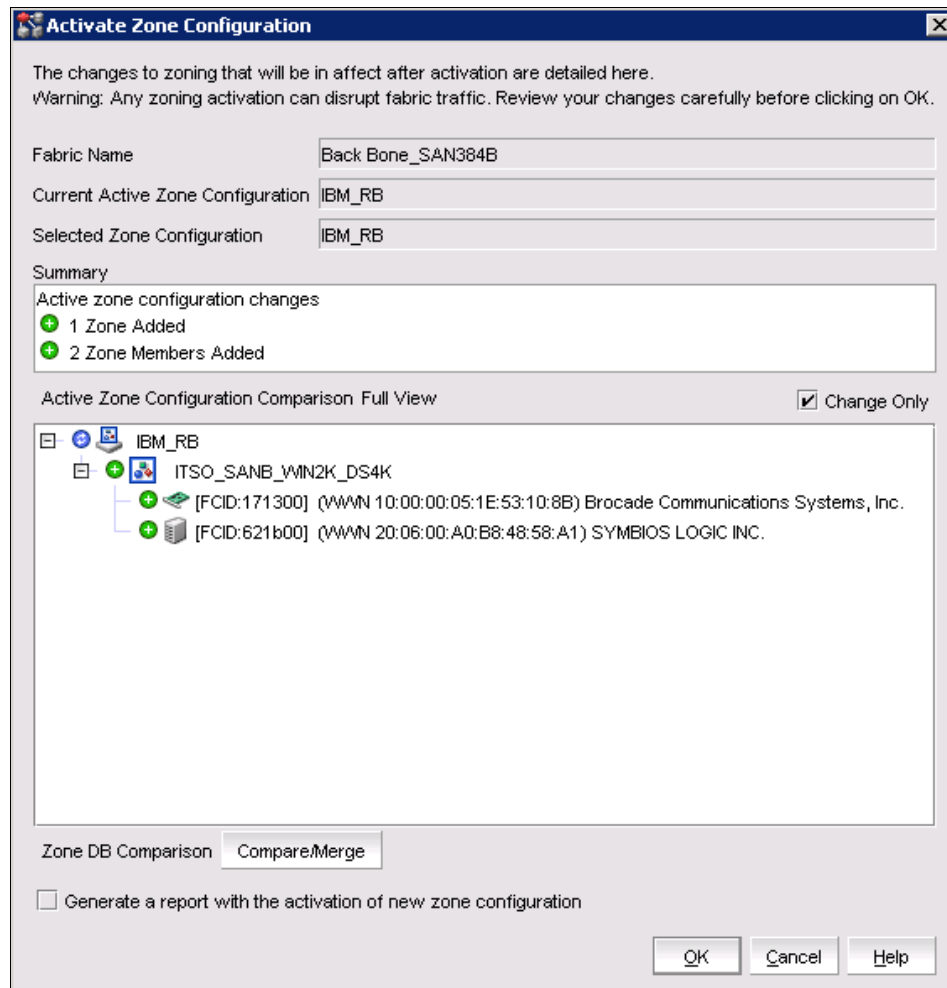


Figure 12-24 Add to existing zone configuration

Click the **OK** button and this will activate the zones displayed in the active zone configuration comparison full view window. This process is concurrent and will not affect any other configured zones.

12.3.7 Analyzing a zone configuration

After any configuration is saved, you can right-click the Fabric Name in the Connectivity Map → select **Connected End Devices** → **Custom...**
See Figure 12-25.

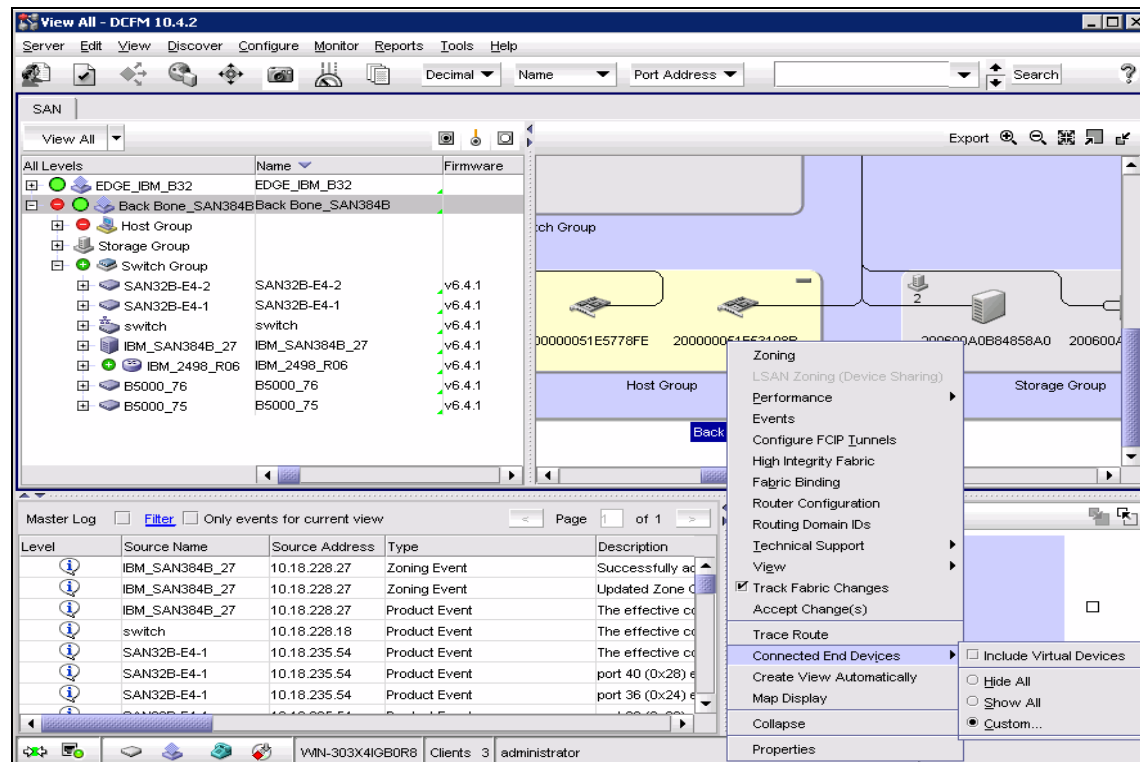


Figure 12-25 Analyze Zone Config

The Connected End Devices Window opens (Figure 12-26). Select the zones you want to show in the connectivity map → Click the right arrow to move the zones to **Selected Zones**.

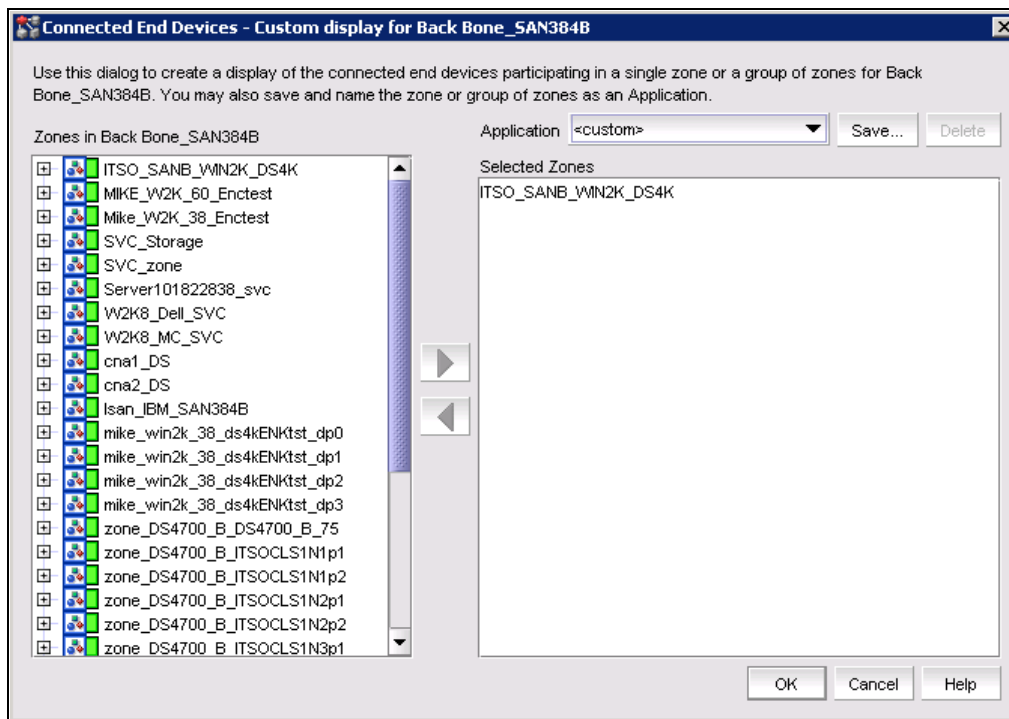


Figure 12-26 Connected End Devices window

In the connectivity map, now only the selected zone displays. See Figure 12-27.

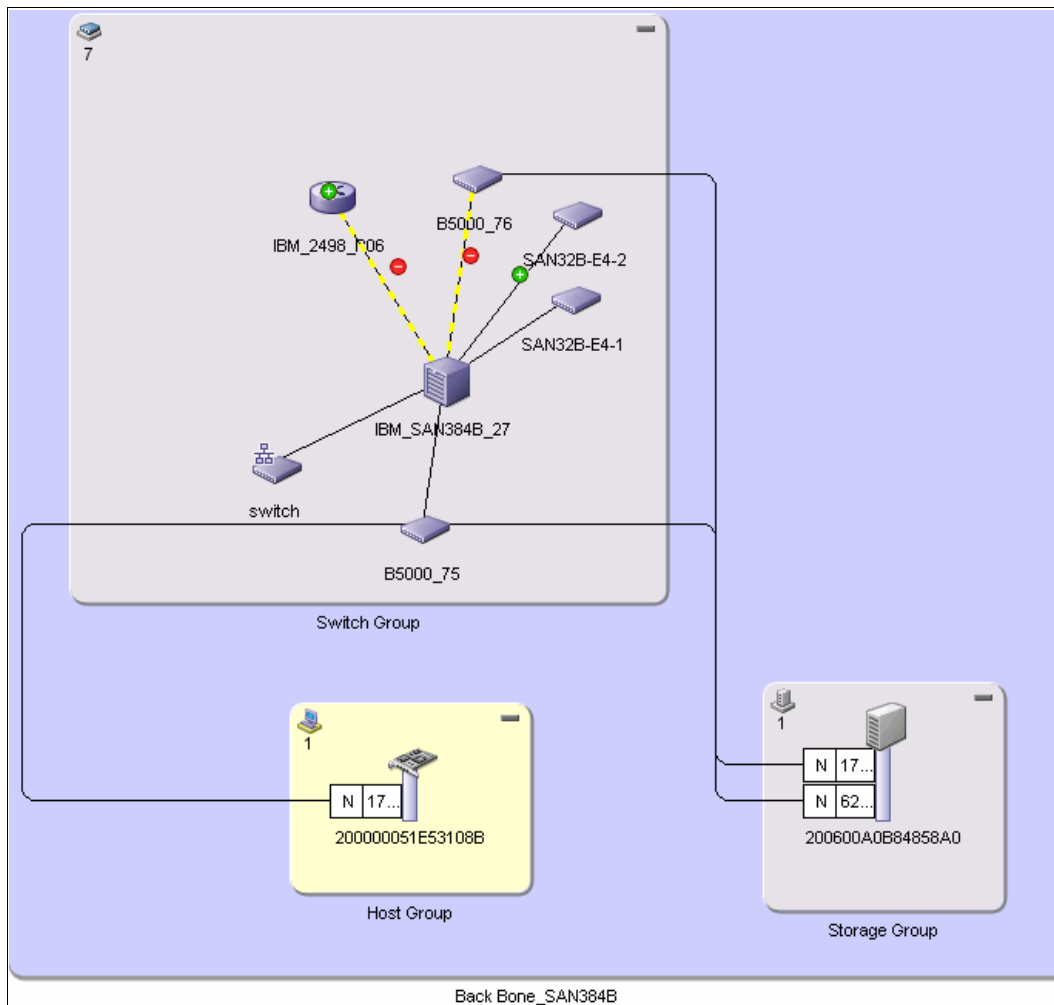


Figure 12-27 Connectivity map

From the Connectivity map, you can now gather more information. For example, select the link between an adapter card and a switch by right-clicking it, then click **Properties**. The Connection Properties window opens (see Figure 12-28).

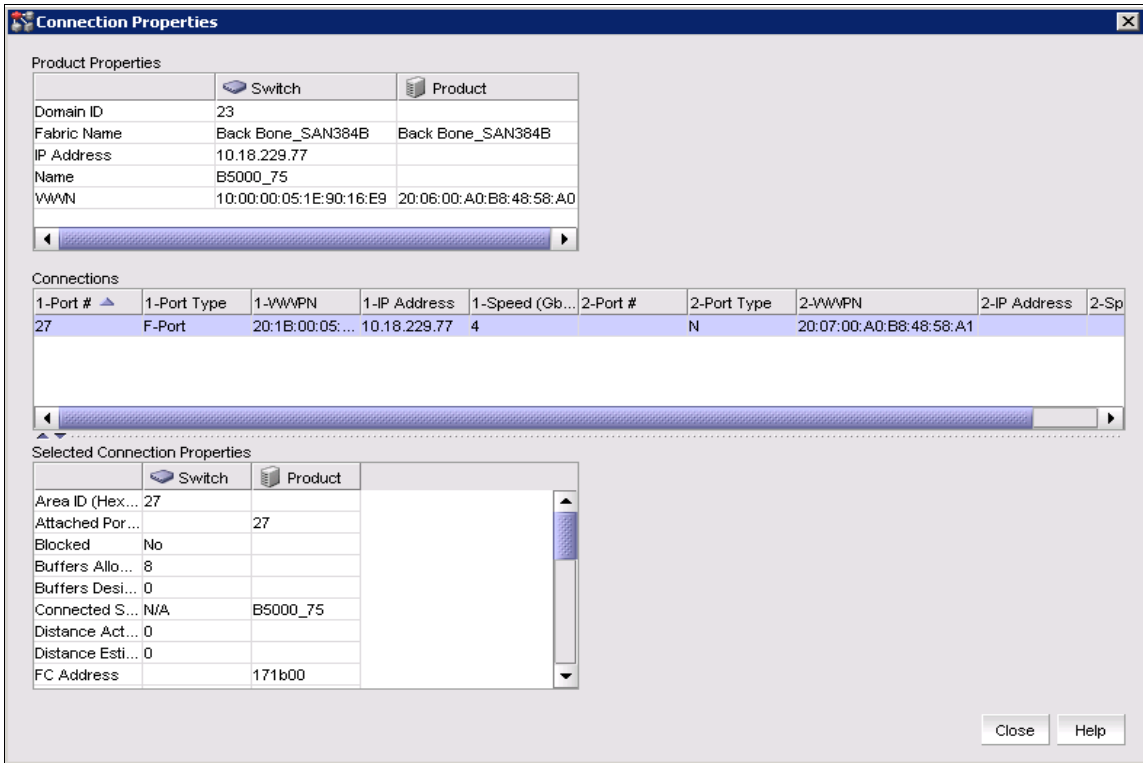


Figure 12-28 Connection Properties

12.4 Basic zoning using Web Tools

Web Tools is an easy-to-use interface that enables organizations to monitor and manage *single* Brocade Fabric OS (FOS) switches. (For fabric-wide monitoring, management, and zone administration, see 12.2, “Zoning using DCFM” on page 516.)

Tool: DCFM is the preferred tool for all zone configuration and administration. The basic version, called DCFM Professional, does not require a licence and can be downloaded from Brocade, however it is not supported by IBM.

DCFm Professional is designed for organizations that have a FOS-only environment and want a management solution for smaller SANs based on a single fabric. DCFm professional can be download from the following link:

http://www.brocade.com/forms/jsp/dcfmdownload/dcfm_download.jsp

Features removed: Starting with Fabric OS version 6.1.1, the following features related to fabric configuration and management have been removed from the Web Tools management interface and implemented in Data Center Fabric Manager (DCFm):

- ▶ Add Un-Zoned Devices
- ▶ Analyze Zone Config
- ▶ Define Device Alias
- ▶ Device Accessibility Matrix
- ▶ Fabric Events
- ▶ Fabric Summary
- ▶ FCIP Tunnel Configuration
- ▶ GigE Ports Interface
- ▶ GigE Ports Route
- ▶ Non-Local Switch Ports Display in Zoning Tree
- ▶ Remove Offline or Inaccessible Devices
- ▶ Zone Database Summary Print

Tasks can be performed by using a Java-capable Web browser from standard laptop, desktop PCs, or workstations from any location within the enterprise.

12.4.1 To start zoning with Web Tools

Follow these steps:

1. Open the Web browser and type the IP address of the device in the Address field, for example, `http://9.155.77.77` and press Enter. Click **Run** on the signed certificate applet. If you select the check box, **Always trust content from this publisher**, the dialog box is not displayed when you open Web Tools again.
2. A browser window opens to open Web Tools and a Login dialog box opens.

Browser window: If you are using Firefox, the browser window is left open. You can close it anytime after the Login dialog box displays. If you are using Internet Explorer, the browser window automatically closes when the login dialog box displays.

3. Log in using your User name and Password (see Figure 12-29).




Figure 12-29 Login dialog box

Be aware that Web Tools uses Role-Based Access Control (RBAC). See Table 12-1 for the capabilities of the various roles.

Table 12-1 Roles and their capabilities

Role	Capabilities
admin	You have full access to all of the Web Tools features.
operator	You can perform any actions on the switch that do not affect the stored configuration.
securityadmin	You can perform actions that do not affect the stored configuration.
switchadmin	You can perform all actions on the switch, except that: <ul style="list-style-type: none">▶ You cannot modify zoning configurations.▶ You cannot create new accounts.▶ You can only view your own account and change your account password.
zoneadmin	You can only create and modify zones.
fabricadmin	You can do everything the Admin role can do except create new users.
user	You have non-administrative access and can perform tasks such as monitoring system activity.
basicswitchadmin	You have a subset of Admin level access.

4. The first thing that you see when you log in to a switch with Web Tools is the Switch Explorer, shown in Figure 12-30.

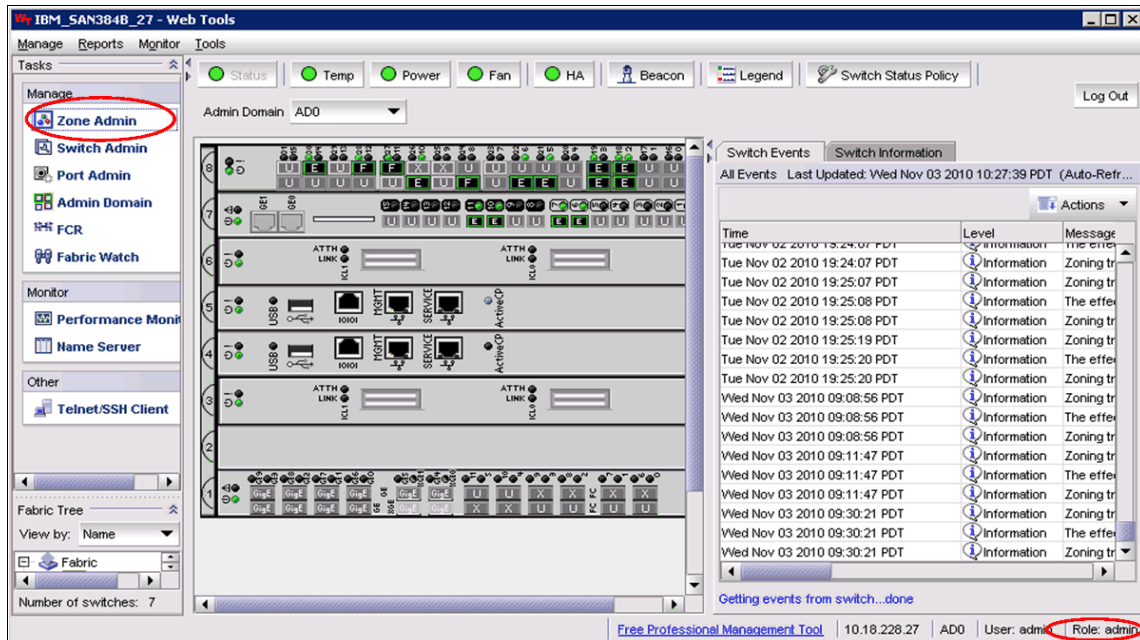


Figure 12-30 Switch Explorer

5. Be sure that the role assigned to you has the rights to perform zoning (see Table 12-1 on page 545). The role assigned to you is shown in the Switch Explorer in Figure 12-30 (see the rounded rectangle). Click **Zone Admin**.

6. The Zone Admin window opens, as shown in Figure 12-31.

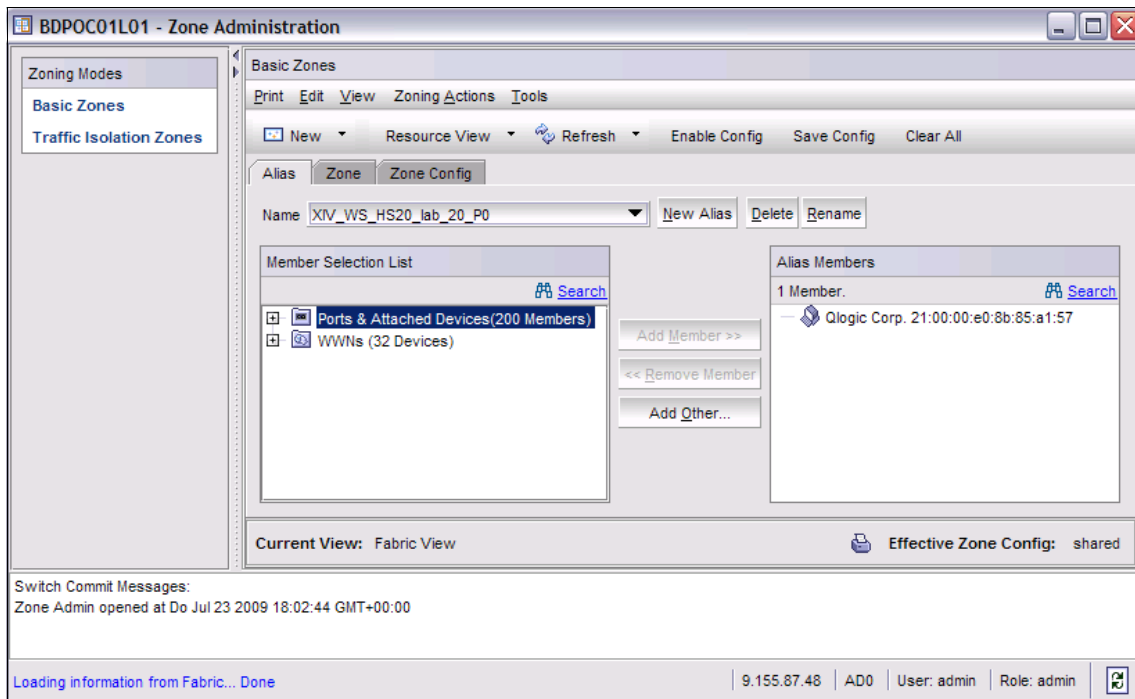


Figure 12-31 Zone Administration window

Changes: Any changes you make in the Zone Administration window are held in a buffered environment and are not updated in the zoning database until you save the changes. If you close the Zone Administration window without saving your changes, your changes are lost. Consider the following terms:

- *Saving:* Updates the zoning database on the switch with the local changes from the Web Tools buffer.
- *Refreshing:* Copies the current state of the zoning database on the switch to the Web Tools buffer, overwriting its current contents.

12.4.2 Creating an alias

By defining an alias to ports or WWNs, you can simplify the configuration of a device that is being zoned. Using a sensible naming convention also assists with troubleshooting at a later stage, making it easier to find specific devices, especially when a SAN grows in complexity. Assign aliases and ensure that they are maintained to help to identify SAN components correctly using the Alias tab.

You can specify members of an alias using the following methods:

- ▶ A switch domain and port index number pair (for example, 2 and 20)
- ▶ Device node and device port WWNs

In this section we describe methods for creating an alias:

1. Select a format to display zoning members in the Member Selection by clicking **View** → **Choose Fabric Resources View**, then. Select **Fabric View** or **Devices Only** as shown in Figure 12-32.




Figure 12-32 Select a format

- Click the **Alias** tab, and click **New Alias**. The Create New Alias dialog box displays, as shown in Figure 12-33.

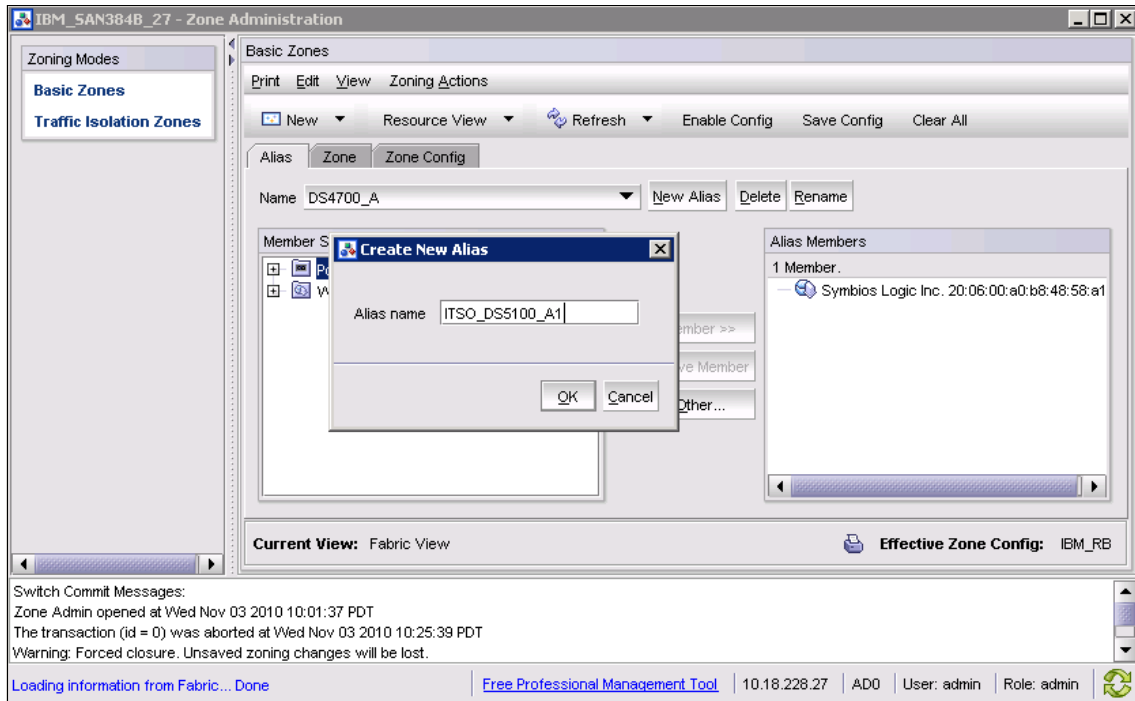


Figure 12-33 New Alias

- Enter a name for the new alias, and click **OK**. The new alias displays in the Name drop-down list. In this example, we create an alias with the name **ITSO_DS5100_A1**.
- 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.
- Click the elements in the Member Selection List that you want to include in the alias. The **Add Member** button becomes active (Figure 12-34).
- Click **Add Member** to add an alias member. The selected member is added to the Alias (Figure 12-34).
- Optionally, 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 you can save and enable your changes (Figure 12-34).
- Click **Save Config** to save the configuration changes. Click **OK** in the window that opens. Remember that this is not applying to the fabric (Figure 12-34).

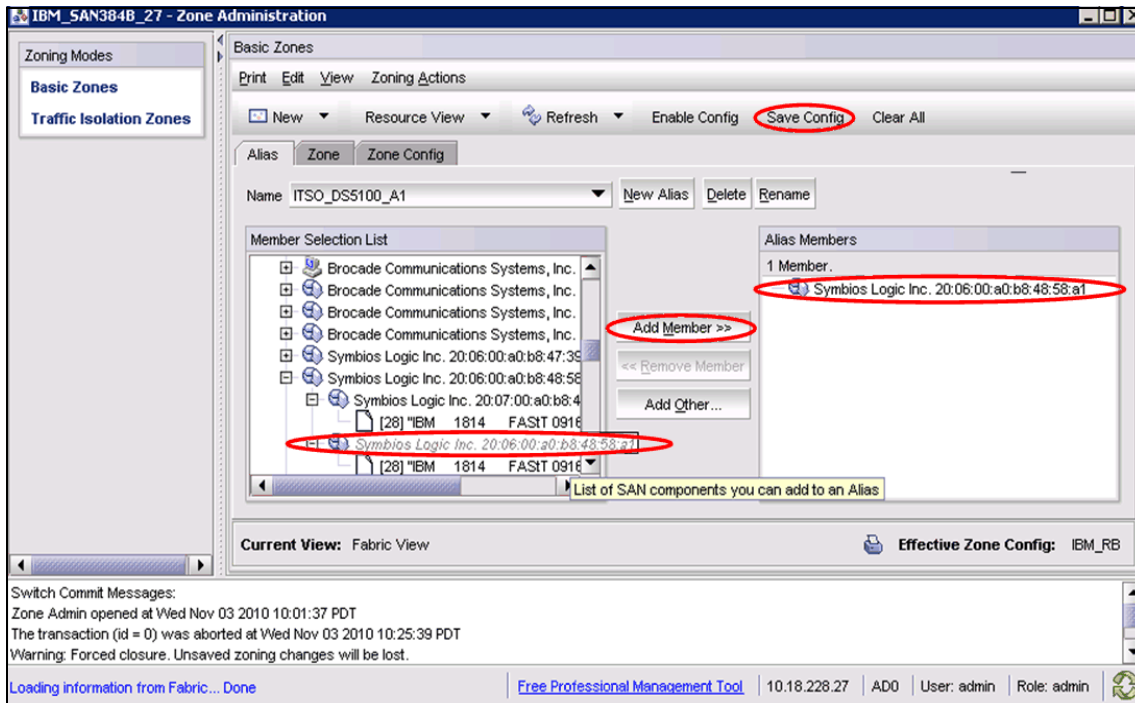


Figure 12-34 Save without enabling

12.4.3 Creating a zone

A *zone* is a region within the fabric in which specified devices can communicate. A device can communicate only with other devices that are connected to the fabric within its specified zone.

You use the Zone tab to specify which devices or switch ports are in the selected zone and to create and manage zones. A zone can have one or multiple members and can include ports, WWNs, aliases, AL_PAs, or Quickloop.

Quickloop: Quickloop is no longer supported from v4.4.x Fabric OS onwards.

You can specify members of a zone using the following methods:

- ▶ Alias names
- ▶ Switch domain and port index number pair (for example, 2 and 20)
- ▶ WWN (device)

Important: Create individual zones of each host to the disk storage subsystems. Also, hosts need a separate HBA for tape communication and, again, must be in another individual host/tape zone.

This small granularity of zoning removes unnecessary PLOGI activity from host to host, as well as removing the risk of issues caused by a faulty HBA affecting others.

12.4.4 Using Web Tools to create a zone

1. From the Zone Administration Main window, go to the Zone tab, and click **New Zone**. The Create New Zone dialog box displays. Enter a name for the new zone, and click **OK**, as shown in Figure 12-35.

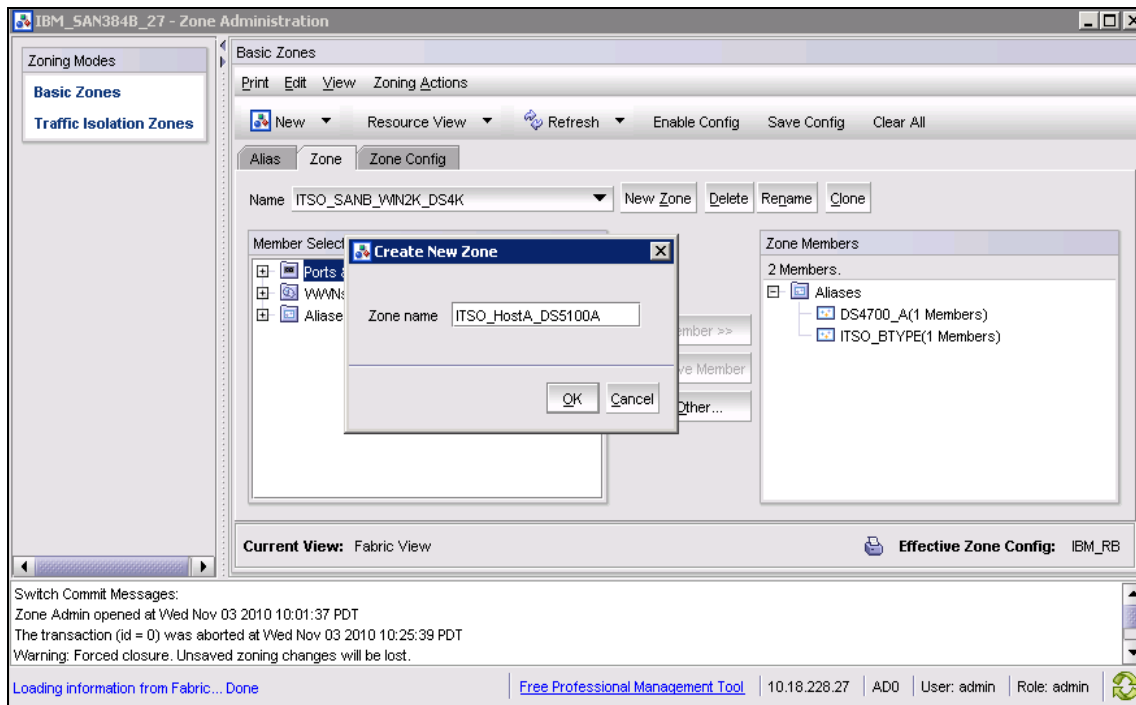


Figure 12-35 Create new zone

The new zone displays in the Name drop-down list.

LSAN: If you are creating an LSAN zone, the zone name must begin with the letters, LSAN_.

2. Expand the Member Selection List to view the nested elements. The choices that are available in the list depend on the selection made in the View menu.
3. 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. Click **Add Member** to add the zone member. The selected member is moved to the Zone Members window (Figure 12-36).
4. Optionally, 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 (Figure 12-36).
5. Click **Save Config** to save the configuration changes. without applying them to the fabric. Click **OK** at the next window (Figure 12-36).

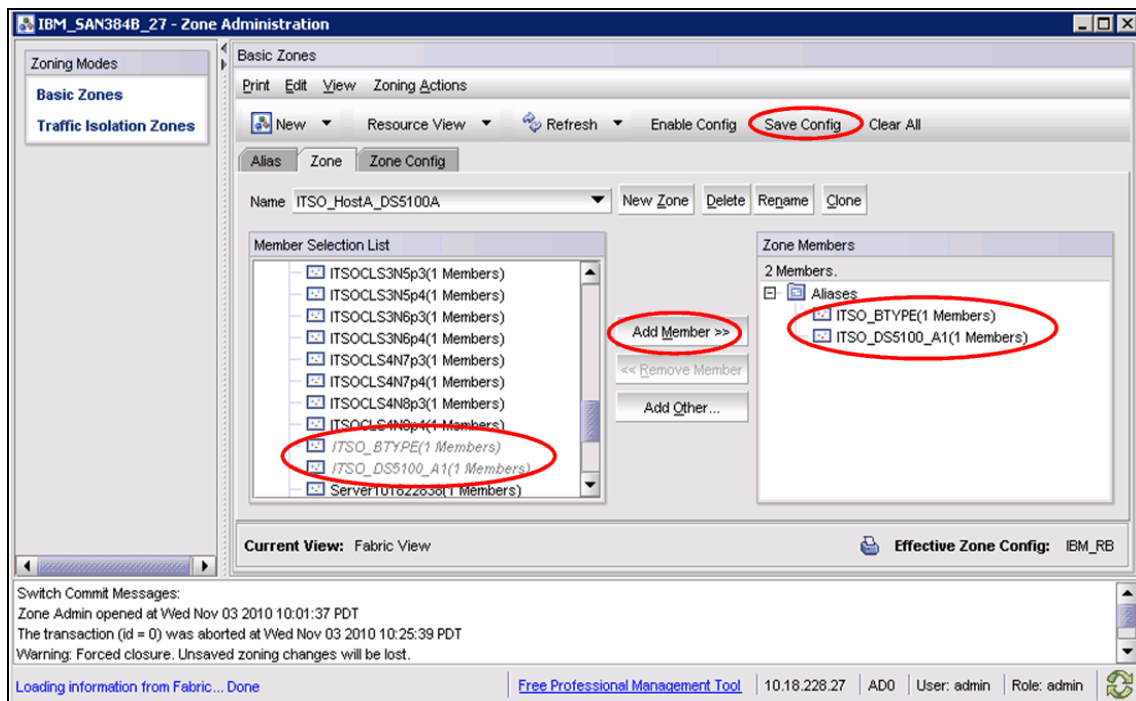


Figure 12-36 Zone members

12.4.5 Creating a zone configuration

Use the Zone Config tab to create or to update a zone configuration. You use zone configurations to enable or disable a group of zones at the same time. To create a zone configuration, follow these steps:

1. Click the **Zone Config** tab, and click **New Zone Config**. The Create a New Config dialog box opens. Enter a name for the new configuration and click **OK**, as shown in Figure 12-37.

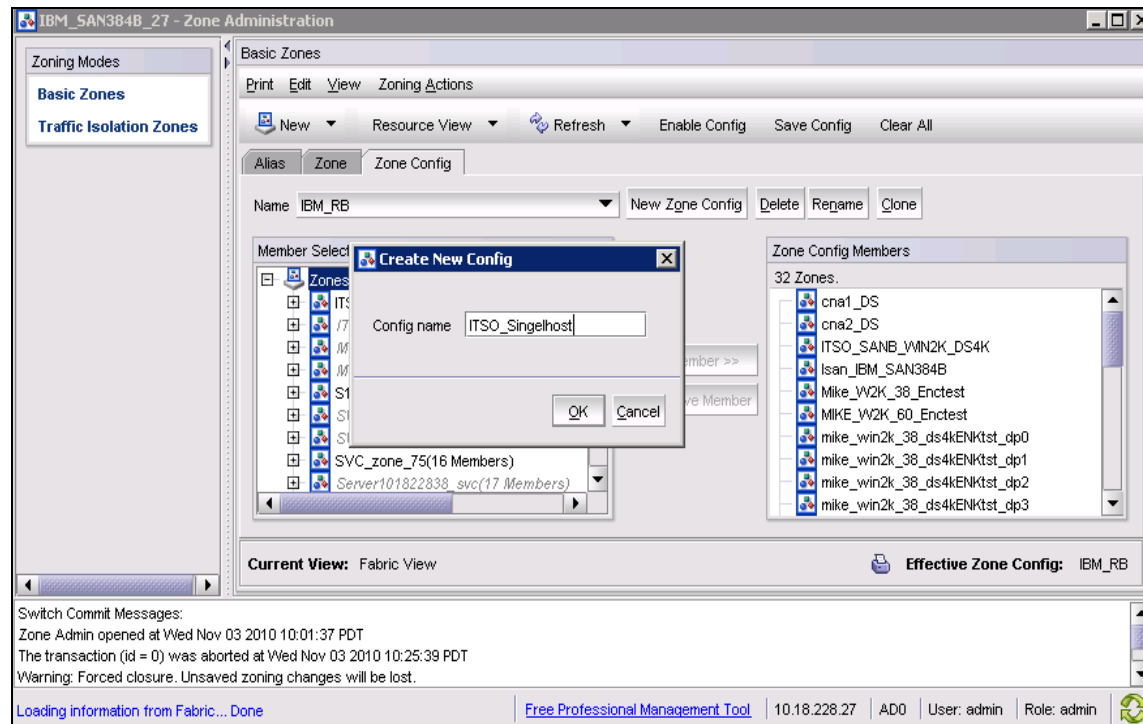


Figure 12-37 Create New Zone Config

- Expand the Member Selection List to view the nested elements.
The choices that are available in the list depend on the selection made in the View menu.
- Select an element in the Member Selection List that you want to include in your configuration.
The **Add Member** button then becomes active.
- Click **Add Member** to add zone configuration members. The selected members are moved to the Config Members window, as shown in Figure 12-38.

Click **Save Config** to save the zone configuration changes. without applying them to the fabric, as shown in Figure 12-38.

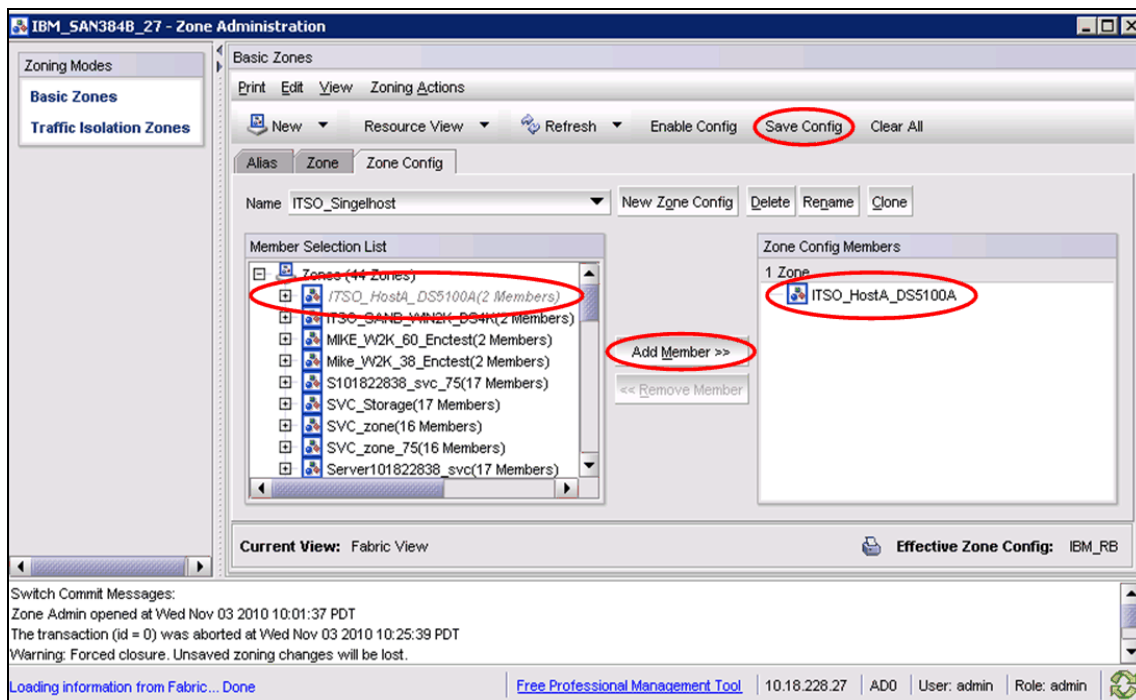


Figure 12-38 Zone Config

12.4.6 Enabling zone configurations

Several zone configurations can reside on a switch at the same time, and you can alternate between them quickly. For example, you might want to have one zone 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 saved automatically, and then the selected zone configuration is enabled.

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

To enable a zone configuration, follow these steps:

1. In the Zone Administration window, click **Enable Config**.
2. Select the zone configuration to be enabled from the drop-down menu and click **OK**, as shown in Figure 12-39.

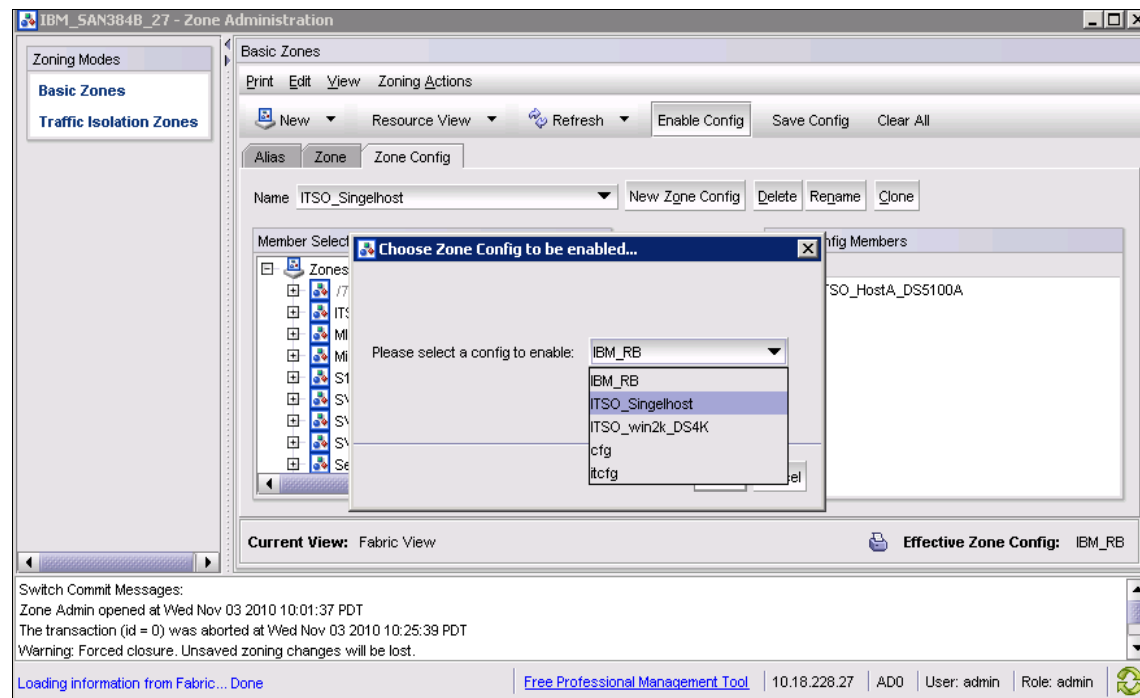


Figure 12-39 Choosing the right config to enable

This saves the zone database to the fabric, and enables the zone configuration that replaces the old one. A message box displays to inform you that this save can result in temporary disruption to I/O. Click **Yes**, as shown in Figure 12-40.

Precautions:

- ▶ Remember to back up your configuration prior to making any configuration changes so that you can always get back to your starting point if there are any problems.
- ▶ Take care when enabling zone configurations. Adding new zones does not impact any currently running definitions, although removing a zone might have a large impact to the current environment.




Figure 12-40 Enabling the zone configuration

3. The commit process can be monitored and you can wait for the commit to be successful. When the save is complete, the Effective Zone Config is now ITSO_Singelhost (Figure 12-41).




Figure 12-41 Effective zone config

12.4.7 Analyzing a zone configuration

After any configuration is saved, select **Analyze Config**, as shown in Figure 12-42, to check the validity of the zone configuration. Analyzing a zone configuration can alert you to ports and WWNs that you have not included.




Figure 12-42 Analyze Zone Config

You are prompted to refresh the current configuration from the switch so that the analyze operation checks the most recent information from the fabric, as shown in Figure 12-43.

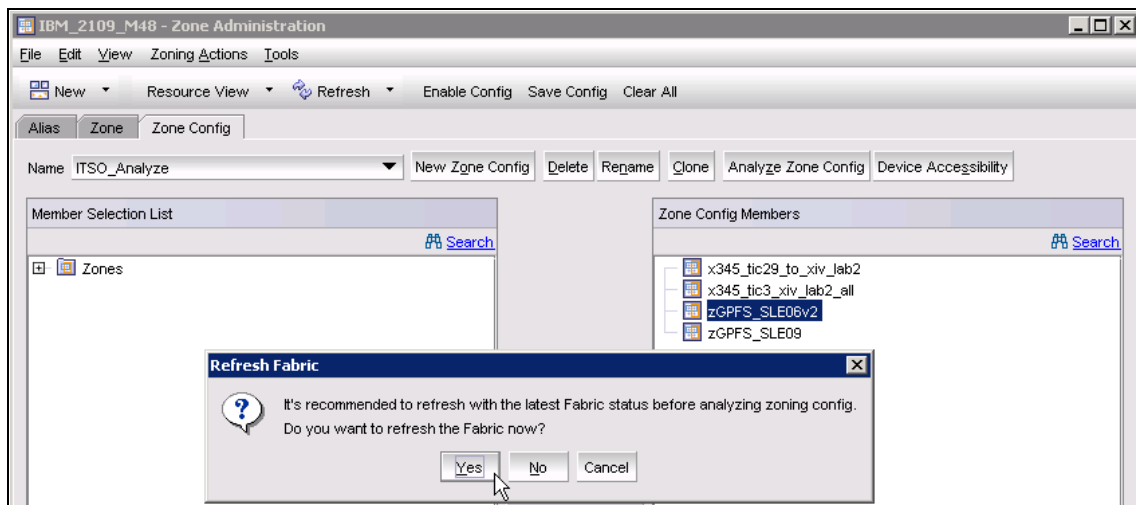


Figure 12-43 Refresh status

Remember to review the output of the analysis and make adjustments (if appropriate) before activating the configuration. Figure 12-44 shows an example of the output, which indicates the WWNs listed are not members of the selected configuration.




Figure 12-44 Sample of Analyze Config output

The Zoning Configuration Analyze window displays a summary of the saved configuration and attempts to point out some of the zoning conflicts before applying the changes to the switch. Some of the potential errors it might catch include these:

- ▶ Ports, WWNs, or devices that are part of the selected configuration but that are not part of the fabric
- ▶ Zones with only a single member

12.4.8 Zoning and E_Ports

When creating a zone, you only work with device ports or host ports (F_Ports, FL_Ports, and L_Ports). Any ISL ports (E_Ports) should not be included in zone definitions. Consider the example presented in Figure 12-45.




Figure 12-45 Zoning implementation: Zoning and E_Ports

To create Zone A, we include these:

- ▶ Domain ID 4, Port 6 (4,6)
- ▶ Domain ID 5, Port 9 (5,9)

However, we do *not* include any ISL ports:

- ▶ Domain ID 4, Port 7 (4,7)
- ▶ Domain ID 5, Port 8 (5,8)

Similarly, to create Zone B, we include these:

- ▶ Domain ID 4, Port 5 (4,5)
- ▶ Domain ID 5, Port 10 (5,10)

Zones do not affect data traffic across ISLs in cascaded switch configurations. Because hard zoning enforcement is performed at the destination, an ISL can carry data traffic from all zones.

Therefore, when dealing with zoning, the fabric should be seen as a “cloud” to which devices are attached. That is, define the end-to-end destinations, and do not include the path to get there.

12.4.9 Broadcast zone

Fibre Channel allows sending broadcast frames to all Nx_Ports if the frame is sent to a broadcast well-known address (FFFFFF); however, many target devices and HBAs cannot handle broadcast frames. To control which devices receive broadcast frames, you can create a special zone, called a *broadcast zone*, that restricts broadcast packets to only those devices that are members of the broadcast zone.

If there are no broadcast zones or if a broadcast zone is defined but not enabled, broadcast frames are not forwarded to any F_Ports. If a broadcast zone is enabled, broadcast frames are delivered only to those logged-in Nx_Ports that are members of the broadcast zone and are also in the same zone (regular zone) as the sender of the broadcast packet.

A broadcast zone can have domain, port, WWN, and alias members.

You can set up and manage broadcast zones using the standard zoning commands, which we describe in 12.3, “Implementing zoning” on page 521.

Broadcast zoning is enforced only for Fabric OS v5.3.x or later switches. If the fabric contains switches running Fabric OS versions earlier than v5.3.x, then all devices that are connected to those switches receive broadcast packets, even if they are not members of a broadcast zone.

12.5 Backing up a zone configuration

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

You can copy the configuration backup to an FTP server or the USB drive. IBM/Brocade 8 Gbps switches support taking configuration backup in USB drive. However, the USB drive must be a Brocade-branded USB drive.

12.5.1 Backing up a zone configuration to an FTP server

To back up the zone configuration to an FTP server, follow these steps:

1. Click **Switch Admin** on the Switch Main Page as shown in Figure 12-46.

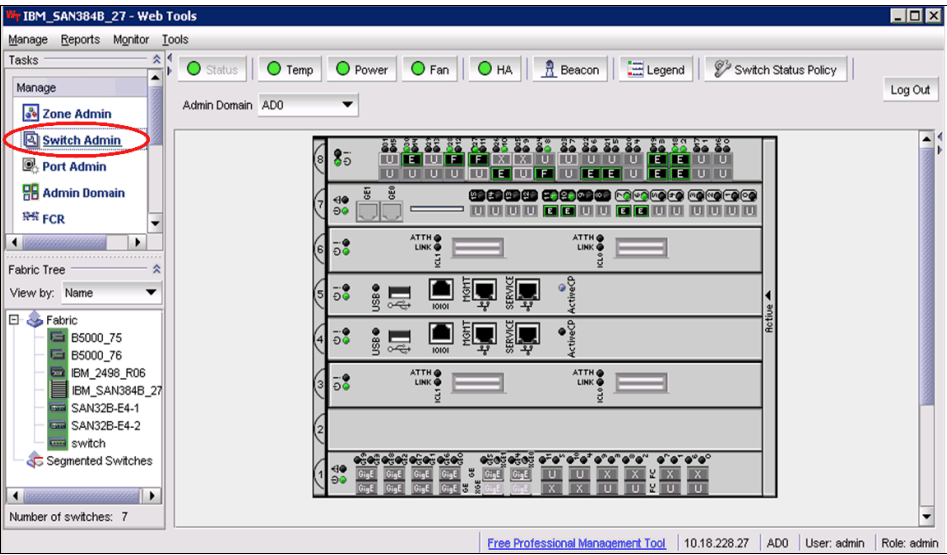


Figure 12-46 Switch Admin

2. Then, click **Show Advanced Mode** as shown in Figure 12-47.

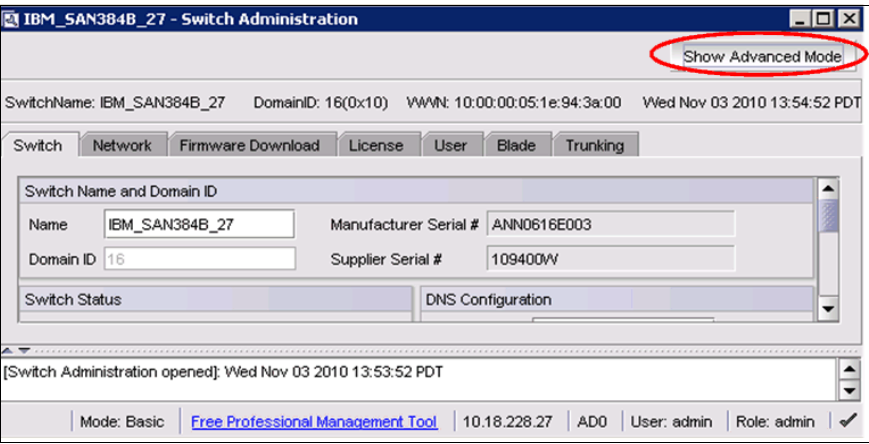


Figure 12-47 Advanced mode

3. Select the **Configure** tab on the top tab bar, and click the **Upload/download** tab at the bottom tab bar.
4. Select both **Config Upload** and **Network**.

Enter the details of the FTP OR SCP server to receive the zone configuration backup to the FTP server, as shown in Figure 12-48.

IBM_SAN384B_27 - Switch Administration

Show Basic Mode

SwitchName: IBM_SAN384B_27 DomainID: 16(0x10) WWW: 10:00:00:05:1e:94:3a:00 Wed Nov 03 2010 13:39:13 PDT

Configure Routing Extended Fabric AAA Service Trace FICON CUP Security Policies

Switch Network Firmware Download License User Blade Trunking SNMP

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: 10.18.228.36

User Name: root

Password: ••••

Protocol Type: Secure Copy Protocol (SCP)

Configuration File Name: SAn384b_27

Upload/Download Progress:

Fabric Virtual Channel Arbitrated Loop System Interoperability Firmware Upload/Download

Apply Close Refresh

[Switch Administration opened]: Wed Nov 03 2010 13:30:12 PDT

Select Protocol Mode: Advanced Free Professional Management Tool 10.18.228.27 AD0 User: admin Role: admin ✓

Figure 12-48 Completing the FTP server details for the configuration backup

The details to complete in this window include these:

- **Host Name or IP:** The host name or IP address of the FTP server where you want to store the configuration backup (for example, 10.64.228.36).
 - **User Name:** The user name of the FTP Server to upload the config file (for example, root).
 - **Password:** The password for the user name.
 - **Protocol Type:**
 - File Transfer Protocol (FTP): The default setting.
 - Secure Copy Method (SCP)
 - **Configuration File Name:** The name of the configuration file.
5. When the information is complete, click **Apply**. Confirm the configuration upload by clicking **Yes**, as shown in Figure 12-49.




Figure 12-49 Confirmation window

When the configuration upload completes successfully, a message displays as shown in Figure 12-50.

The screenshot shows the 'IBM_SAN384B_27 - Switch Administration' window. At the top, there's a 'Show Basic Mode' button. Below it, the switch name 'IBM_SAN384B_27', DomainID '16(0x10)', WWN '10:00:00:05:1e:94:3a:00', and the date/time 'Wed Nov 03 2010 14:03:52 PDT' are displayed. A series of tabs are visible: 'Configure', 'Routing', 'Extended Fabric', 'AAA Service', 'Trace', 'FICON CUP', 'Security Policies', 'Switch', 'Network', 'Firmware Download', 'License', 'User', 'Blade', 'Trunking', and 'SNMP'. The 'Configure' tab is active, and within it, the 'Upload/Download' sub-tab is selected. The 'Function' section has 'Config Upload' selected. The 'Select source of configuration file' section has 'Network' selected. Below this, a message says 'Provide Host details, Transfer Protocol and Path for Configuration file' with a note '*Password is optional, if user name is "anonymous"'. Fields for 'Host Name or IP' (10.18.228.36), 'User Name' (root), 'Password' (masked with dots), 'Protocol Type' (Secure Copy Protocol (SCP)), and 'Configuration File Name' (SAN384B_27) are present. An 'Upload/Download Progress' bar is empty. At the bottom of the form are 'Apply', 'Close', and 'Refresh' buttons. A log window at the bottom shows the following messages: '[Switch Administration opened]: Wed Nov 03 2010 13:53:52 PDT', '[upload started]: Wed Nov 03 2010 14:02:52 PDT', 'Initiating configuration file [SAN384B_27] upload to 10.18.228.36', 'ConfigUpload completed successfully.', and '[upload completed]: Wed Nov 03 2010 14:02:52 PDT'. The status bar at the very bottom shows 'upload completed successfully.', 'Mode: Advanced', a link to 'Free Professional Management Tool', the IP '10.18.228.27', 'AD0', 'User: admin', 'Role: admin', and a green checkmark.

Figure 12-50 Configuration upload successful message

12.5.2 Backing up a zone configuration to a Brocade USB device

Only IBM/Brocade 8 Gbps switches and DCX Backbone, with Fabric OS v6.1x or later support backup of the configuration using the Brocade supplied USB drive.

Here are some points to consider:

- ▶ In the case of the DCX Backbone, insert the USB drive on the active CP only.
- ▶ Enable the USB device. An LED lights after a successful enable.
- ▶ Disable the device before unplugging it to prevent data corruption on the USB drive, and the indicator LED stops glowing.
- ▶ You need to purchase the USB drive separately. The USB drive does not ship with the switch.

To back up a zone configuration to a USB device, follow these steps:

1. Plug in the USB device to the USB port. Open Web Tools, and click the USB icon, as shown in the Figure 12-51.

USB: Make sure to plug in the USB device to the switch (the Active CP in the case of the DCX Backbone).

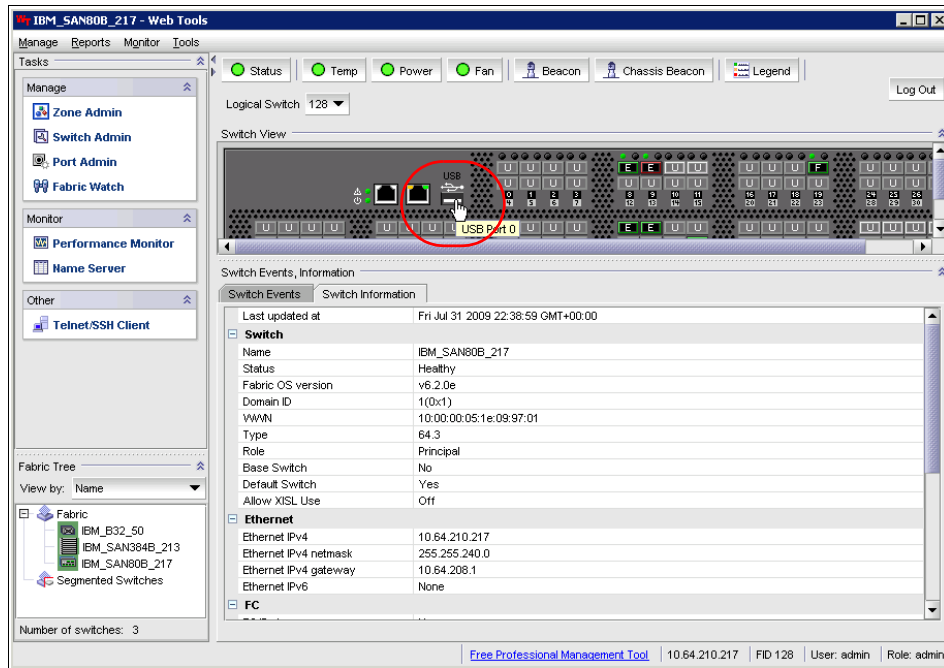


Figure 12-51 Configuration backup to the USB device

2. Click **Mount USB Device** as shown in Figure 12-52.

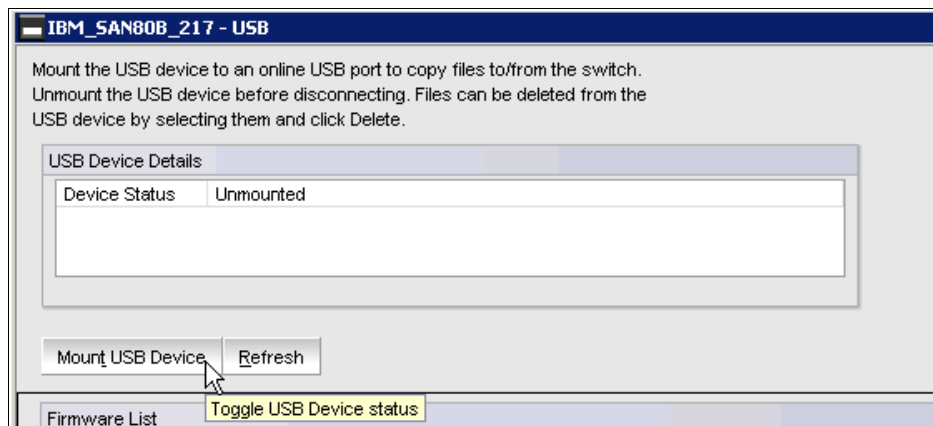


Figure 12-52 Mounting the USB device

3. Confirm the mount by clicking **Yes** in the message dialog box as shown in Figure 12-53.

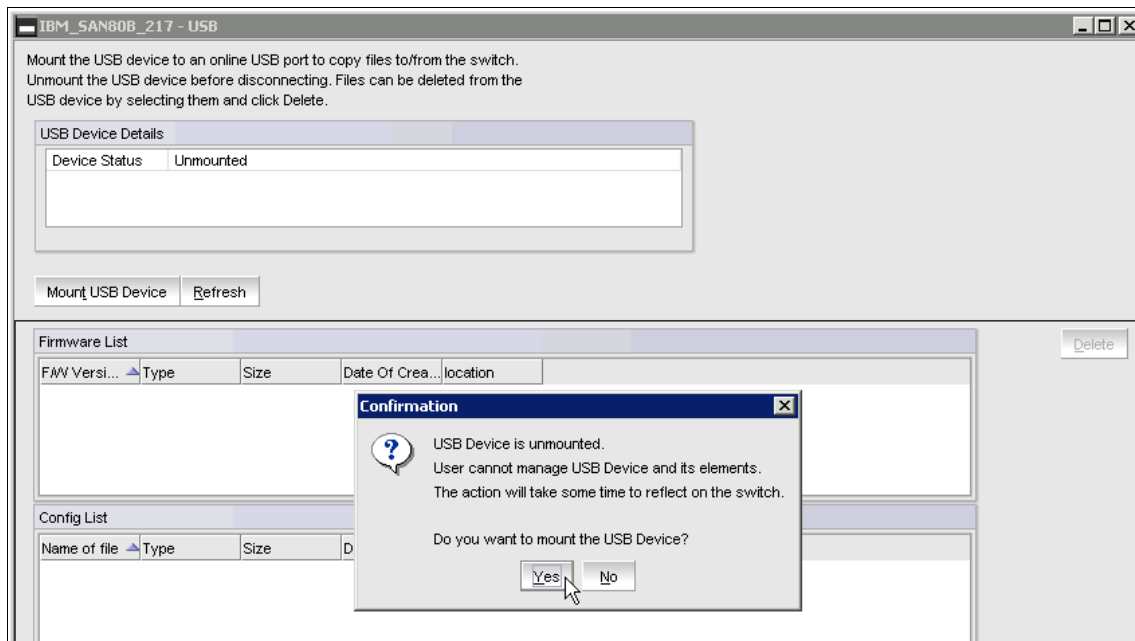


Figure 12-53 Confirmation for mounting the USB device

When the USB device is mounted, you might see backup files as highlighted in Figure 12-54.

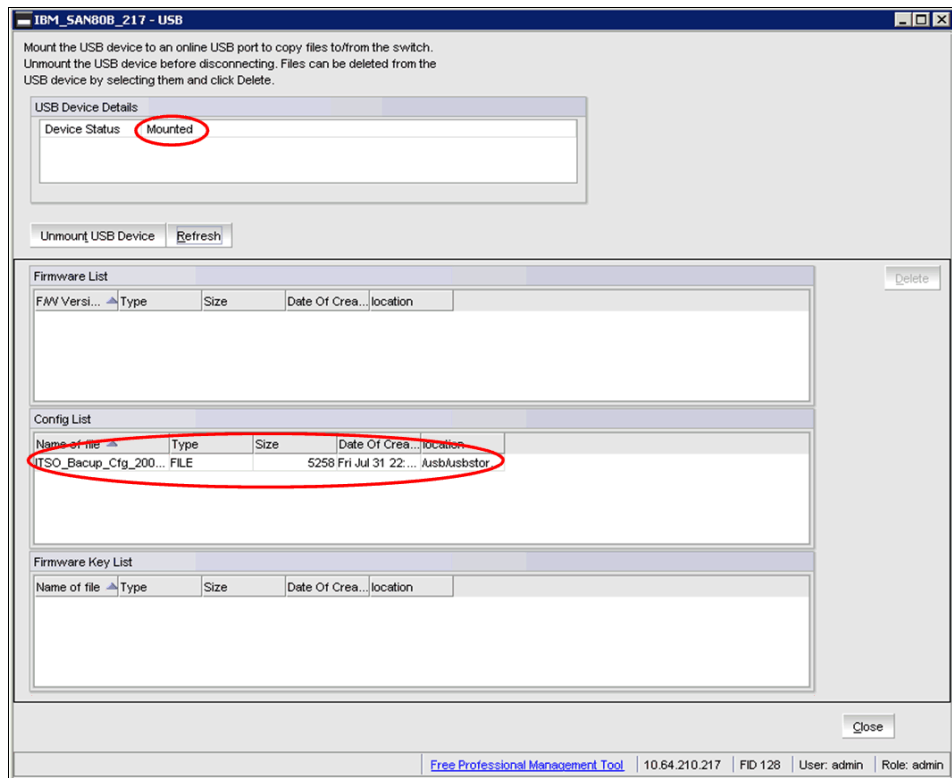


Figure 12-54 List the contents of the USB Device

4. Next, follow the steps that are similar to those described in 12.5.1, “Backing up a zone configuration to an FTP server” on page 562 to back up the configuration, except that you need to select the USB check box in order to select the USB drive as the source (as highlighted in Figure 12-55). Then, select the Configuration File Name and click **Apply**.

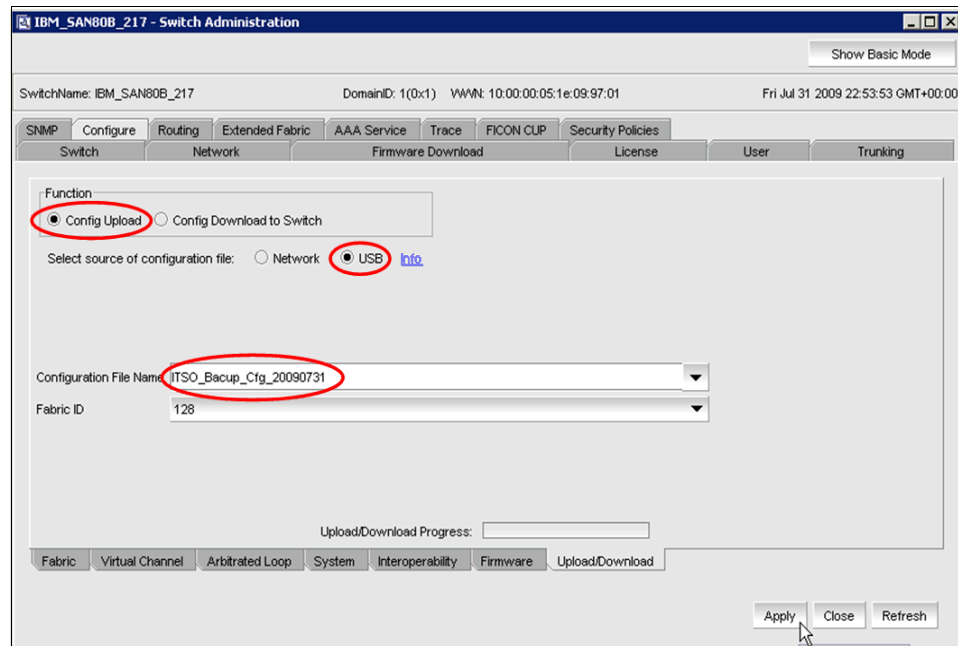


Figure 12-55 Config Upload to a USB device

5. Select **Yes** in the confirmation window, as shown in Figure 12-56.




Figure 12-56 Confirmation for the configuration upload

When the configuration upload completes successfully, a message displays as shown in Figure 12-57.

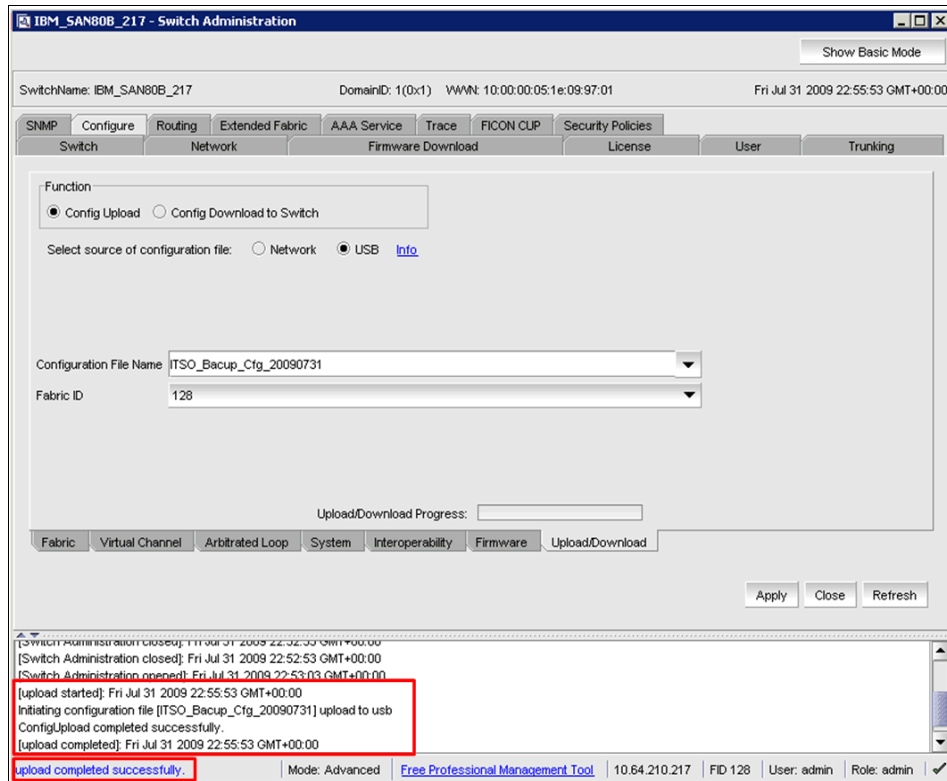


Figure 12-57 Config Upload completed successfully

6. Unmount the USB device before unplugging it to prevent data corruption, as shown in Figure 12-58.




Figure 12-58 Unmount USB device

12.5.3 Downloading a zone configuration from a USB device

Follow these steps:

1. See 12.5.2, “Backing up a zone configuration to a Brocade USB device” on page 566 for directions about how to mount the USB device and how to check the status of the USB device after mounting.
2. Open the switch Administration window and select the **Switch** tab as shown in Figure 12-59. For the Switch Status, select **Disable** and click **Apply**.

IBM_SAN80B_217 - Switch Administration

Show Basic Mode

SwitchName: IBM_SAN80B_217 DomainID: 1(0x1) VVNN: 10:00:00:05:1e:09:97:01 Sun Aug 02 2009 04:03:17 GMT+00:00

SNMP Configure Routing Extended Fabric AAA Service Trace FICON CUP Security Policies

Switch Network Firmware Download License User Trunking

Switch Name and Domain ID

Name: IBM_SAN80B_217 Manufacturer Serial #: AHX0617D001

Domain ID: 1 Supplier Serial #: 107700H

Switch Status

☐ Enable ☒ Disable

DNS Configuration

DNS Server 1:

DNS Server 2:

Domain Name:

Remove All

Report

View Report

Reboot/Fastboot

Reboot Fastboot

Apply Close Refresh

Apply the changes

Preparing configuration file [ITSO20090801] upload to user
ConfigUpload completed successfully.
[upload completed]: Sun Aug 02 2009 03:34:39 GMT+00:00
[download started]: Sun Aug 02 2009 03:35:39 GMT+00:00
Initiating configuration file [ITSO20090801] download to Switch
ConfigDownload completed successfully.
[download completed]: Sun Aug 02 2009 03:35:39 GMT+00:00

No changes sent to Switch. Mode: Advanced [Free Professional Management Tool](#) 10.64.210.217 FID 128 User: admin Role: admin

Figure 12-59 Switch Disable

3. Select **Yes** in the confirmation window, as shown in Figure 12-60.

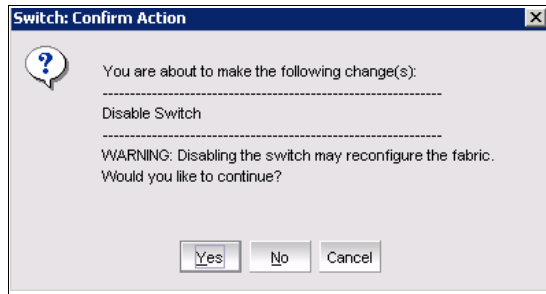


Figure 12-60 Confirmation window

4. Next, follow steps that are similar to those described in 12.5.2, “Backing up a zone configuration to a Brocade USB device” on page 566, except that you need to select the *Config Download to Switch* check box in order to download the config from the USB drive (as highlighted in Figure 12-61). Then, select the Configuration File Name, and click **Apply**.

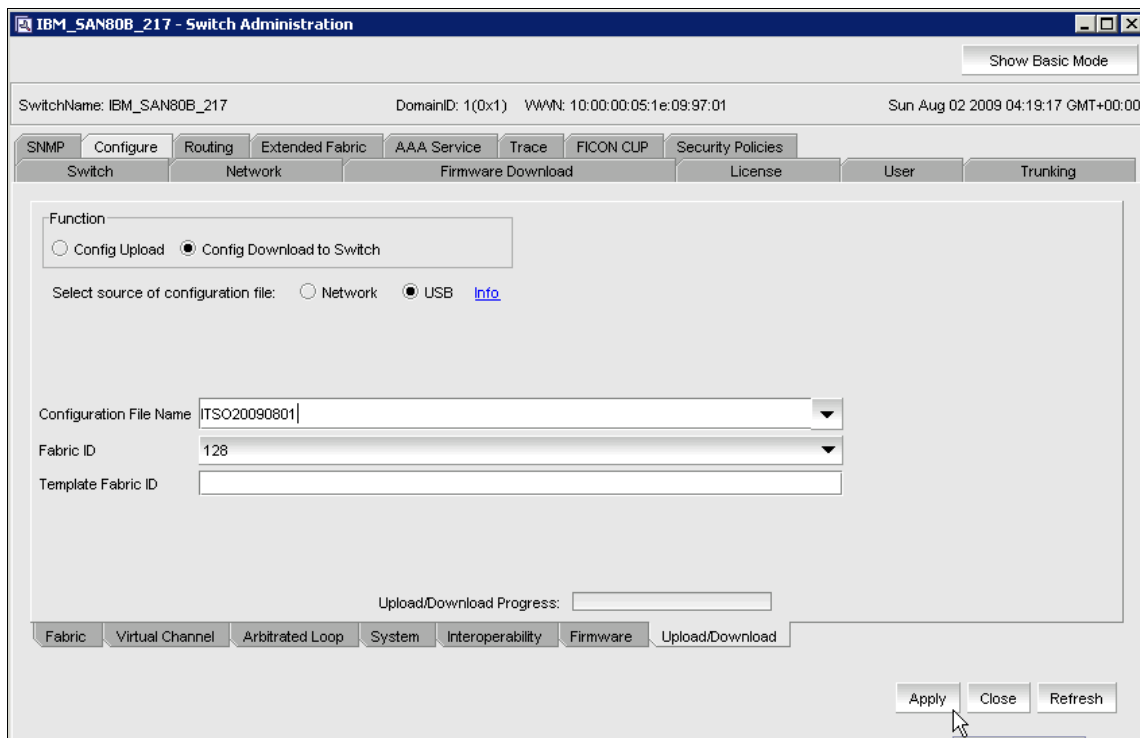


Figure 12-61 Config Download to Switch

5. Select **Yes** in the confirmation window, as shown in Figure 12-62.

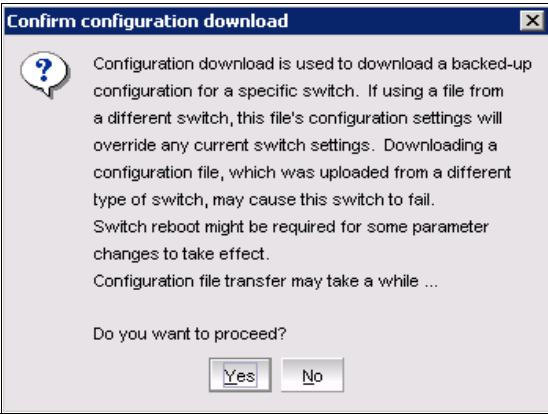


Figure 12-62 Confirm configuration download

6. When the configuration download completes successfully, a message displays as shown in Figure 12-63. Click the **Switch** tab.

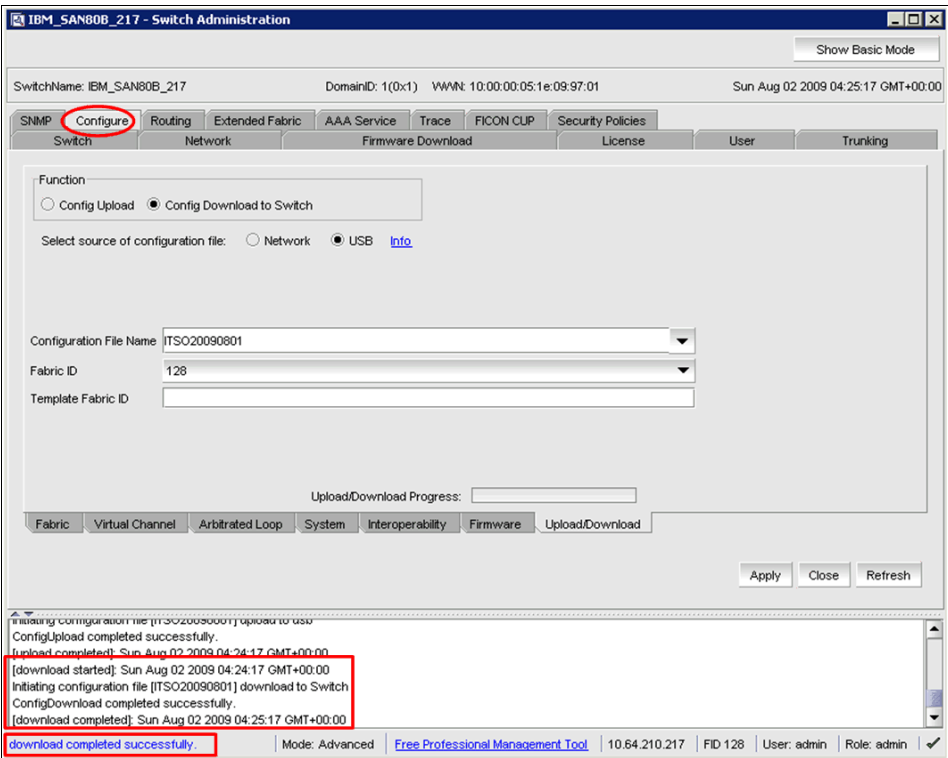


Figure 12-63 Config Download completed successfully

7. Select Switch Status Enable as shown in Figure 12-64 and click **Apply**.

The screenshot shows the 'IBM_SAN384B_27 - Switch Administration' window. At the top, there's a 'Show Basic Mode' button. Below it, the switch name 'IBM_SAN384B_27', Domain ID '16(0x10)', WWN '10:00:00:05:1e:94:3a:00', and date 'Wed Nov 03 2010 14:52:56 PDT' are displayed. A series of tabs are visible: Configure, Routing, Extended Fabric, AAA Service, Trace, FICON CUP, Security Policies, Switch, Network, Firmware Download, License, User, Blade, Trunking, and SNMP. The 'Switch' tab is active. Under 'Switch Name and Domain ID', the Name is 'IBM_SAN384B_27', Domain ID is '16', Manufacturer Serial # is 'ANN0616E003', and Supplier Serial # is '109400W'. In the 'Switch Status' section, the 'Enable' radio button is selected and circled in red, while the 'Disable' radio button is unselected. To the right, the 'DNS Configuration' section has fields for 'DNS Server 1', 'DNS Server 2', and 'Domain Name', with a 'Remove All' button below. Below the status section is a 'Report' section with a 'View Report' button. At the bottom of the main area is a 'Reboot/Fastboot' section with 'Reboot' and 'Fastboot' buttons. A log window at the bottom shows the following text: '[download started]: Wed Nov 03 2010 14:51:54 PDT', 'Initiating configuration file [SAN384B_27] download to Switch', 'Doing configDownload on switch...', 'Activating configDownload: Switch is disabled', and 'ConfigDownload completed successfully.'. The bottom status bar includes links for 'Change current switch settings', 'Mode: Advanced', a link to 'Free Professional Management Tool', IP '10.18.228.27', AD0, User: admin, Role: admin, and a close button.

Figure 12-64 Switch Status Enable

8. Select **Yes** in the confirmation window, as shown in Figure 12-65.

The screenshot shows a 'Switch: Confirm Action' dialog box. It features a question mark icon in a speech bubble. The text inside reads: 'You are about to make the following change(s):', followed by 'Enable Switch' on a line with a dashed underline. Below this is a 'WARNING: Enabling the switch may reconfigure the fabric. Would you like to continue?' message. At the bottom, there are three buttons: 'Yes', 'No', and 'Cancel'.

Figure 12-65 Confirm Action

Check that status has changed to Enable Switch and the changes have been saved to the switch, as shown in Figure 12-66.

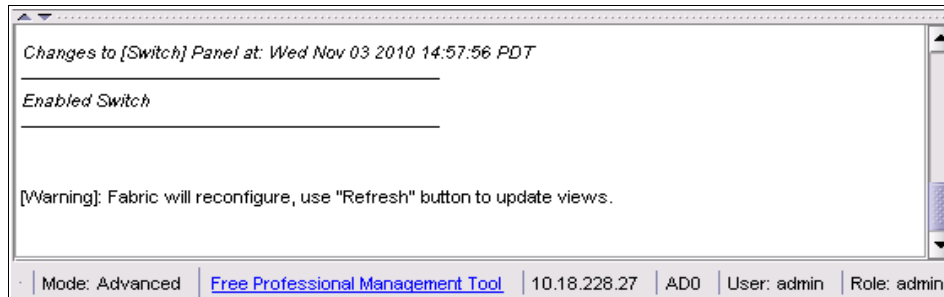


Figure 12-66 Changes saved to switch

9. Unmount the USB device as shown before in Figure 12-58 on page 571.

12.6 Zoning using CLI

To use the CLI or Telnet to create an alias, follow these steps:

1. Connect to the switch and log in as admin.
2. Enter the **aliCreate** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example 12-1 creates an alias called **DS8000_p1** for WWN **50:05:07:63:04:18:03:16** and another alias called **DS8000_p2** for WWN **50:05:07:63:04:08:c3:16**.

Example 12-1 The aliCreate command

```

magic_cl:admin> alicreate DS8000_p1, 50:05:07:63:04:18:03:16
magic_cl:admin> alicreate host_p1, 10:00:00:00:c9:2f:bd:5f
magic_cl: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

```

12.6.1 Using CLI to create a zone

Attention: Before executing the `cfgDisable`, `cfgEnable`, or `cfgSave` commands, execute the `rcsDisabled` command to check whether your fabric has Reliable Commit Service (RCS) enabled (`rcsDisabled=0`). If RCS is disabled (`rcsDisabled=1`), check for older switches in the fabric. After the older switches are upgraded, RCS is enabled by default.

RCS is available on all switches running Fabric OS v4.1 and later. RCS guarantees that either all or none of the switches receive the new zone configuration. Use RCS to secure a reliable propagation of the latest zone configuration.

If you use non-RCS mode, you must log in to every switch to monitor the status of the zone configuration.

Broadcast: To create a broadcast zone, use the reserved name *broadcast*. Do not give a regular zone the name of *broadcast*.

To create a zone using CLI or Telnet, follow these steps:

1. Connect to the switch and log in as admin.
2. Enter the `zoneCreate` command.
To create a broadcast zone, use the reserved name *broadcast*.
3. Enter the `cfgSave` command to save the change to the defined configuration.

Example 12-2 creates a zone called `host_p1_to_DS8000_p1` and adds alias `host_p1` as one of the `host_p1_to_DS8000_p1` members.

Example 12-2 The zoneCreate command

```
magic_c1:admin> zonecreate "host_p1_to_DS8000_p1", "host_p1"
magic_c1: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 ...
magic_c1:admin>
```

Adding devices (members) to a zone

To add a device to a zone, follow these steps:

1. Connect to the switch and log in as `admin`.
2. Enter the **zoneAdd** command.
3. Enter the **cfgSave** command to save the change to the defined configuration

Example 12-3 adds the alias **DS8000_p1** to the **host_p1_to_DS8000_p1** zone.

Example 12-3 The zoneAdd and cfgSave commands

```
magic_c1:admin> zoneadd "host_p1_to_DS8000_p1", "DS8000_p1"
magic_c1: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 ...
magic_c1:admin>
```

Removing devices (members) from a zone

To remove devices (members) from a zone, follow these steps:

1. Connect to the switch and log in as `admin`.
2. Enter the **zoneRemove** command.
3. Enter the **cfgSave** command to save the change to the defined configuration

Example 12-4 remove **DS8000_p1** from the **host_p1_to_DS8000_p1** zone.

Example 12-4 The zoneRemove command

```
magic_c1:admin> zoneremove "host_p1_to_DS8000_p1", "DS8000_p1"
magic_c1: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
```

Deleting a zone

To delete a zone, follow these steps:

1. Connect to the switch and log in as admin.
2. Enter the **zoneDelete** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

Example 12-5 deletes **host_p1_to_DS8000_p1** from the configuration.

Example 12-5 The zoneDelete command

```
magic_c1:admin> zonedeldelete "host_p1_to_DS8000_p1"
magic_c1: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
```

12.6.2 Using CLI to create a zone configuration

To use CLI or Telnet to create a zone configuration, follow these steps:

1. Connect to the switch and log in as admin.
2. Enter the **cfgCreate** command.
3. Enter the **cfgSave** command to save the change to the defined configuration (Example 12-6).

Example 12-6 The cfgCreate command

```
magic_c1:admin> cfgcreate "ITS0_Cfg", "host_p1_to_DS8000_p1"
magic_c1: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
```

Adding zones (members) to a zone configuration

To add zones (members) to a zone configuration, follow these steps:

1. Connect to the switch and log in as admin.
2. Enter the **cfgAdd** command.
3. Enter the **cfgSave** command to save the change to the defined configuration (Example 12-7).

Example 12-7 The cfgAdd command

```
magic_c1:admin> cfgadd "ITS0_Cfg", "host_p1_to_DS8000_p2"
magic_c1: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
```

Removing zones (members) from a zone configuration

To remove zones (members) from a zone configuration, follow these steps:

1. Connect to the switch and log in as admin.
2. Enter the **cfgRemove** command.
3. Enter the **cfgSave** command to save the change to the defined configuration (Example 12-8).

Example 12-8 The cfgRemove command

```
magic_c1:admin> cfgremove "ITS0_Cfg", "host_p1_to_DS8000_p2"
magic_c1: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
```

Deleting a zone configuration

To delete a zone configuration, follow these steps:

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

3. Enter the **cfgSave** command to save the change to the defined configuration (Example 12-9).

*Example 12-9 The **cfgDelete** command*

```
magic_cl:admin> cfgdelete "ITS0_Cfg"
magic_cl: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
```

Clearing changes to a zone configuration

To clear changes to a zone configuration:

1. Enter the **cfgTransAbort** command.
2. When this command is executed, all changes since the last save operation (performed with the **cfgSave** command) are cleared.

Example 12-10 clears the removal of a member from zone1, which was done in error with the **zoneRemove** command.

*Example 12-10 The **cfgTransAbort** command (after a **zoneRemove** command)*

```
magic_cl:admin> zoneremove "host_p1_to_DS8000_p1", "host_p1"
magic_cl:admin>
magic_cl:admin> cfgtransabort
```

12.6.3 Backing up a zone configuration using the CLI

You can back up a zone configuration using the **configUpload** command, as follows:

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.

Respond to the prompts as follows:

- **Protocol (SCP or FTP):** If your site requires the use of Secure Copy, specify SCP. Otherwise, specify FTP. If you leave this prompt blank, then the default specified in the square brackets ([]) is used.

- **Server Name or IP Address:** Enter the name or IP address of the server where the file will be stored (for example, **9.155.66.102**). You can enter a server name if DNS is enabled.
- **User name:** Enter the user name of your account on the server (for example, **admin**).
- **File name:** Specify a file name for the backup file (for example, **magic_c1_bkp20090721**). You specify absolute path names using a forward slash (/). Relative path names create the file in the user's home directory on UNIX servers and in the directory where the FTP server is running on Windows servers.
- **Password:** Enter your account password for the server. The Password field is a required field even if you are logged in as an anonymous user. In such cases, the value can be ignored by the FTP service.

Example 12-11 shows the **configUpload** command run on a switch without Admin Domains.

Example 12-11 The configUpload command to the FTP server

```

magic_c1:admin> configupload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 9.155.66.102
User Name [user]: ftpuser
File Name [config.txt]: magic_c1_bkp20090721
Password:
Upload complete
magic_c1:admin>

```

12.6.4 Backing up a zone configuration using a USB drive

Only IBM/Brocade 8 Gbps switches and DCX Backbone, with Fabric OS v6.1x or later support backup of the configuration using the Brocade supplied USB drive.

Here are some points to consider:

- ▶ In the case of the DCX Backbone, insert the USB drive on the active CP only.
- ▶ Enable the USB device. An LED lights after a successful enable.
- ▶ Disable the device before unplugging it to prevent data corruption on the USB drive, and the indicator LED stops glowing.
- ▶ You need to purchase the USB drive separately. The USB drive does not ship with the switch.
- ▶ Restore the backed up configuration using the **configDownload** command.

Example 12-12 shows how to enable the USB device.

USB: Make sure to plug in the USB device to the switch (the Active CP in the case of the DCX Backbone).

Example 12-12 Enable the USB device

```
magic_c1:admin> usbstorage -e
Trying to enable USB device. Please wait...
USB storage enabled
magic_c1:admin>
```

Example 12-13 shows the commands to back up the configuration to the USB device.

Example 12-13 The configUpload command

```
magic_c1:admin> configupload -U ITS0_Backup_20090730
configUpload complete: All config parameters are uploaded
magic_c1:admin> usbstorage -l
firmware\                0B      2007 Sep 28 15:33
config\                  23kB    2009 Jul 31 23:15
    ITS0_Bacup_Cfg_20090722 5kB    2009 Jul 22 21:49
    ITS0_Bacup_Cfg_20090731 5kB    2009 Jul 31 22:56
    ITS0_Backup_20090730 12kB    2009 Jul 30 23:15
support\                 0B      2007 Sep 28 15:33
firmwarekey\             0B      2007 Sep 28 15:33
Available space on usbstorage 99%
magic_c1:admin>
```

Unmount the USB device before unplugging it to prevent data corruption; see Example 12-14.

Example 12-14 Unmount the USB device

```
magic_c1:admin> usbstorage -d
USB storage disabled
magic_c1:admin>
```

The configuration is now backed up.

12.6.5 Downloading a zone configuration from an FTP server

In case the configuration is lost, unintentional changes have occurred, or replacement of the switch is necessary, you can download the configuration from the latest backup.

Restoring a configuration involves overwriting the configuration on the switch by downloading a previously saved backup configuration file. Make sure that the configuration file that you download is compatible with your switch model, because configuration files from other model switches or firmware versions might cause your switch to fail.

You can download configuration files to a switch while the switch is enabled. You do not need to disable the switch.

Switch: For some Admin Domain configurations, the switch must be disabled.

To download a zone configuration from an FTP server, follow these steps:

1. Telnet to the switch and login as `admin`.
2. Enter the **configDownload** command.
3. When prompted, respond as follows:
 - **Protocol (SCP or FTP):** If your site requires the use of Secure Copy, specify **SCP**. Otherwise, specify **FTP**. If you leave this prompt blank, then the default specified in the square brackets (`[]`) is used.
 - **Server Name or IP Address:** Enter the name or IP address of the server where the file will be stored (for example, **9.155.66.102**). You can enter a server name if DNS is enabled.
 - **User name:** Enter the user name of your account on the server (for example, `admin`).
 - **File name:** Specify the file name for the backup file to be downloaded. You can specify absolute path names using a forward slash (`/`). Relative path names create the file in the user's home directory on UNIX servers and in the directory where the FTP server is running on Windows servers.
 - **Password:** Enter your account password for the server. The Password field is a required field even if you are logged in as an anonymous user. In such cases, the value can be ignored by the FTP service.
4. When prompted with the message **Do you want to continue [y/n]**, enter **y**.
5. Wait for the configuration to be restored.

Example 12-15 shows the **configDownload** command run on a switch without Admin Domains.

Example 12-15 The configDownload command

```
magic_c1:admin> configdownload
Protocol (scp, ftp, local) [ftp]:
Server Name or IP Address [host]: 9.155.66.102
User Name [user]: admin
File Name [config.txt]: magic_c1_bkp20090720
```

***** 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:

Activating configDownload: Switch is disabled

configDownload complete: All config parameters (except any AD Headers, SFOS and Security parameters) are downloaded to ADO
magic_c1:admin>

Reboot: Because some configuration parameters require a reboot in order to take effect, after you download a configuration file, you must reboot to be sure that the parameters are enabled. Before the reboot, this type of parameter is listed in the configuration file, but it is not effective until after the reboot.

12.6.6 Downloading a zone configuration from a USB device

See 12.6.4, “Backing up a zone configuration using a USB drive” on page 582 for directions about how to mount the USB device and how to check the status of the USB device after mounting.

Enter the **configDownload -U <Filename>** command to download the config from the USB device.

Example 12-16 shows the **configDownload** command from a USB device.

Example 12-16 The configDownload command

```
magic_c1:admin> switchdisable
magic_c1:admin> configdownload -U ITS0_Backup_20090730
```

***** 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

Activating configDownload: Switch is disabled

configDownload complete: All config parameters (except any AD Headers,SFOS and Security parameters) are downloaded to ADO
magic_c1:admin>

Reboot: Because some configuration parameters require a reboot in order to take effect, after you download a configuration file, you must reboot to be sure that the parameters are enabled. Before the reboot, this type of parameter is listed in the configuration file, but it is not effective until after the reboot.

Unmount the USB device before unplugging it to prevent data corruption see Example 12-14 on page 583.

Example 12-17 Unmount the USB device

```
magic_c1:admin> usbstorage -d
USB storage disabled
magic_c1:admin>
```



Multiple switches and fabrics

In this chapter, we discuss the considerations for multiple switch environments, such as merging fabrics, duplicate domain IDs, zoning configuration conflicts, and operating parameter conflicts.

13.1 Multiple switch environments

In this section, we focus on multiple switch environment considerations.

13.1.1 Gateway links

A gateway merges SANs into a single fabric by establishing point-to-point E_Port connectivity between two Fibre Channel switches that are separated by a network with a protocol such as IP or SONET.

Except for link initialization, gateways are transparent to switches; the gateway simply provides E_Port connectivity from one switch to another.

By default, switch ports initialize links using the Exchange Link Parameters (ELP) mode 1. However, gateways expect initialization with ELP mode 2, also referred to as ISL R_RDY mode. Therefore, to enable two switches to link through a gateway, the ports on both switches must be set for ELP mode 2.

Any number of E_Ports in a fabric can be configured for gateway links, provided that you follow these guidelines:

- ▶ All switches in the fabric must be upgraded to Fabric OS v5.2.0 or later.
- ▶ All switches in the fabric are using the core PID format.
- ▶ The switches connected to both sides of the gateway are included when determining switch count maximums.
- ▶ Extended links (those created using the Extended Fabrics licensed feature) are not supported through gateway links.

Example 13-1 shows how to enable R_RDY on port 8/47 using the **portcfgislmode** command. The example is performed on a IBM SAN384B switch. Commands are slightly different for the non-director type switches.

For more detailed information, see the Fabric OS Administrator's Guide, available at the following website:

<http://www.brocade.com>

Example 13-1 Enable ISL R_RDY mode using portcfgislmode

```
IBM_SAN384B_213:FID128:admin> portcfgislmode
Usage: portCfgISLMode [SlotNumber/]PortNumber      Mode
Mode:    1 - Enable ISL R_RDY Mode on port
         0 - Disable ISL R_RDY Mode on port

IBM_SAN384B_213:FID128:admin> portcfgislmode 8/47 1
ISL R_RDY Mode is enabled for port 239. Please make sure the PID
formats are consistent across the entire fabric.
```

After running the command in the example, the ISL link is now operational.

13.1.2 Buffer credit recovery

This feature is available only with IBM 8 Gbps switches. Buffer recovery credit allows links to recover after frames and R_RDYs are lost when the credit recovery logic is enabled. Buffer credit recovery maintains performance because as soon as one credit is lost, it attempts to recover. During link reset, the frame and credit loss counters are reset without performance degradation. This feature is supported only on long distance E_Ports that are connected between GoldenEye2 and Condor2-based ports. Buffer credit recovery does not require any configuration.

If a long distance E_Port is connected to any other type of Application Specific Integrated Circuit (ASIC), the buffer credit recovery feature is disabled. Virtual E_Ports and Virtual EX_Ports do not support long distance. The buffer credit recovery feature is enabled for Normal and Virtual Channel flow control modes.

A port that supports BB_Credit recovery maintains the following BB_Credit recovery values:

- ▶ **BB_SC_N**: The log2 of BB_Credit Recovery modulus.
- ▶ **BB_RDY_N**: Counts the number of R_RDY primitives received modulo 2BB_SC_N.
- ▶ **BB_FRM_N**: Counts the number of frames that are received modulo 2BB_SC_N(VC) mode, and Extended VC mode.

No configuration is required because the configuration is done internally by the Fabric OS and the ASIC.

13.1.3 ISL Trunking

ISL Trunking is an optionally licensed product on the IBM b-type family of switches. ISL Trunking requires you to purchase and install a separate ISL Trunking license key, which has to be installed on all switches that participate in the trunk.

The ISL Trunking feature allows up to eight ISLs to merge logically into a single link. An ISL-link is a connection between two switches through an Expansion Port (E_Port).

When using ISL Trunking to aggregate bandwidth of up to eight ports, the speed of the ISLs between switches in a fabric is multiplied correspondingly up to eight times.

For example, at 4 Gbps speeds, trunking 4 ports between two SAN-24B switches delivers an ISL throughput of up to 16 Gbps. Trunking at 8 Gbps with 8 ISL-links forms 8-port trunks that can deliver up to 64 Gbps. ISL trunking is extended to N_Ports where trunks are formed when the edge switch is running Fabric OS v6.2.0 or later.

Figure 13-1 shows some examples of ISL Trunking.

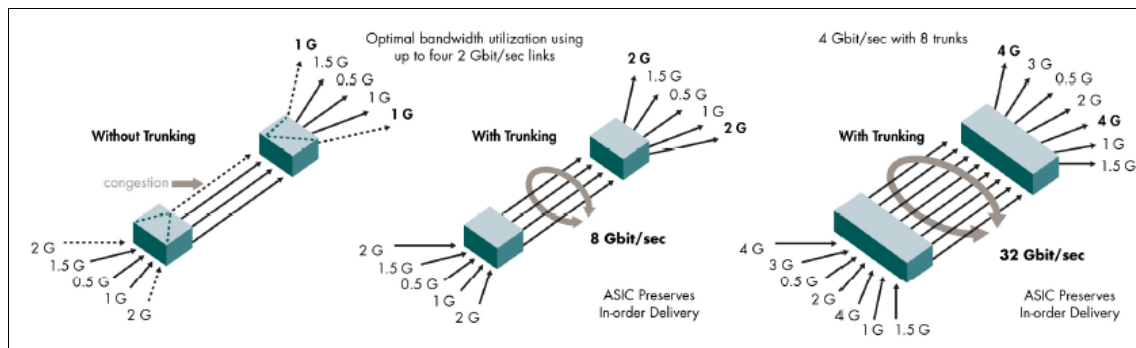


Figure 13-1 Trunking examples

You can manage ISL Trunking using Telnet commands or the Web Tools interface.

The ISL Trunking feature has many advantages. It supports high-bandwidth, large-scale SANs that include core switches. ISL Trunking provides a high bandwidth path between switches in a fabric, while balancing the traffic across the individual links and maintaining *in-order delivery* of data packets to their destination.

Attention: *In-order delivery* is the preferred setting in an IBM fabric. However, the user can change this setting.

ISL Trunking uses frame-level load balancing. You can use it with Exchange Based Routing, to achieve faster fabric convergence, as well as higher availability.

In the remainder of this section we discuss the distinct advantages of using ISL Trunking.

ISL: The 6-port 10 Gbps blade for the IBM SAN Director type switches can only be used for ISL connectivity. This blade has no support for ISL trunking.

Trunk groups, ports, and masters

ISL Trunking performs load balancing dynamically, at the frame level, across a set of available links between two adjacent switches. Ports on a switch are grouped in port groups. Trunks can only form from port group to port group. When a trunk group is formed, one port is referred to as the *trunk master*. In Fabric OS v5.x and later, if a master link goes offline, a new master is selected automatically with no disruption to traffic. Using trunking with previous versions of Fabric OS or over an EX_Port, if a master goes offline, there is a temporary disruption to traffic as the link is rebuilt.

Trunk groups

A *trunk group* is identified by the trunk master that represents the entire group. The remainder of the group members are referred to as *subordinate links* that help the trunk master direct traffic across ISLs, allowing efficient and balanced in-order communication.

The trunking groups are based on the user port number with contiguous eight ports as one group, such as, 0-7, 8-15, and 16-23. You can enable and disable trunking and set trunk port speeds (for example, 2 Gbps, 4 Gbps, 8 Gbps, or autonegotiate) for entire switches or for individual ports.

Trunk ports

Observe the following criteria for standard distance trunking:

- ▶ There must be a direct connection between participating switches.
- ▶ Trunk ports must reside in the same port group.
- ▶ Trunk ports must run at the same speed (either 2 Gbps, 4 Gbps, or 8 Gbps).
- ▶ Trunk ports must be set to the same ISL mode (L0 is the default).
- ▶ If the switch is in Access Gateway mode, the trunk ports must be F_Ports.
- ▶ The switch must be set to interopMode 0 for Brocade Native mode.
- ▶ The port ISL mode must be disabled (using the **portCfgIslMode** command).

Trunk masters

The *trunk master* implicitly defines the trunk group. All ports with the same master are considered to be part of the same group. Each trunk group includes a single trunk master and several trunk subordinate links. The first ISL established in a trunk group is assigned to be the trunk master, also known as the *principal ISL*. After the trunk group is fully established, all data packets that are intended for transmission across the trunk are distributed dynamically at the frame level across the ISLs in the trunk group, while preserving in-order delivery.

Masterless EX port trunking

Fabric OS v6.2.0 implemented dynamic port binding, this was so an area can be dynamically assigned to a port in a virtual fabric. Fabric OS v6.3.0 uses this dynamic binding to reassign the area from a old master to a new master. When the mater goes offline, a slave port becomes the master by assuming the area of the old master. This means that all slave ports will remain online and when the old master becomes online it will be given a new unused area.

Installing an ISL Trunking license

The IBM b-type family of switches require that you install an ISL Trunking license on both switches at either end of an ISL trunk, in order to enable trunking.

Administering ISL Trunking

The ISL Trunking feature is managed by performing some administration tasks. These tasks among other include:

- ▶ Enabling or disabling the trunking
- ▶ Enabling and disabling ports of a switch
- ▶ Setting the speed of a port
- ▶ Debugging a trunking link failure

The ISL Trunking feature is administered using Telnet commands.

Enabling an ISL Trunking license

After you unlock the ISL Trunking license, trunking is enabled automatically across all ports, but you must re-initialize the ports that are used for ISLs so that they recognize that trunking is enabled. You perform this procedure only once.

To initialize the ports again, you can either disable and then enable the switch again using **switchDisable** and then **switchEnable**, or you can disable and then enable the affected ports again using **portDisable [slot/]port** and **portEnable [slot/]port**. By disabling and enabling the switch itself, all ports are available for trunking.

Managing trunking using the CLI

Example 13-2 is an example of how to enable trunking using the Fabric OS v6.2.0 Command Line Interface (CLI). An IBM b-type switch and its ports have trunking enabled by default. As such, trunks will form automatically if more than one ISL is connected within a port group of the switches.

Example 13-2 Enabling and Managing trunking

```
IBM_SAN384B_213:FID128:admin> trunkshow  
No trunking links
```

```
IBM_SAN384B_213:FID128:admin> switchcfgtrunk 1  
Configuration applied to all ports except the following VE/VEX_Ports  
(ports 16 - 31).
```

```
IBM_SAN384B_213:FID128:admin> portcfgtrunkport 2/18 1  
IBM_SAN384B_213:FID128:admin> portcfgtrunkport 2/19 1
```

```
IBM_SAN384B_213:FID128:admin> trunkshow  
1: 82-> 75 10:00:00:05:1e:09:97:01 2 deskew 15 MASTER  
83-> 74 10:00:00:05:1e:09:97:01 2 deskew 15
```

```
IBM_SAN384B_213:FID128:admin> trunkdebug 218 219  
port 218 and 219 connect to the switch 10:00:00:05:1e:09:97:01
```

```
IBM_SAN384B_213:FID128:admin>
```

Managing trunking using the GUI

As an alternative to the CLI, you can select the individual ports from the Web Tools administrative interface and enable ISL Trunking from there.

You disable or enable trunking using the Ports tab, as shown in Figure 13-2. Select either **Enable Trunking** or **Disable Trunking**.

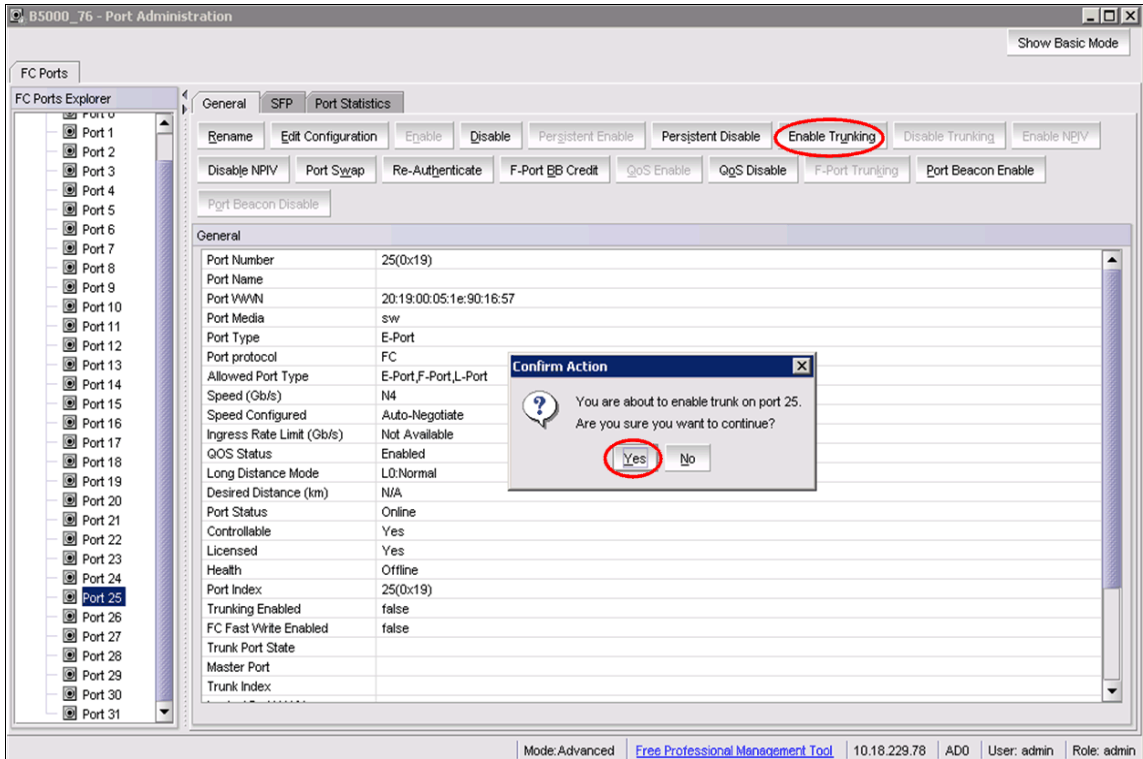


Figure 13-2 Enable or disable trunking on port

Figure 13-3 shows the additional items that display in the Ports tab window when you scroll the window. In this example, trunking is enabled on port 24, it is configured as a subordinate (slave) trunk port, and Port 25 is chosen as master trunk port.

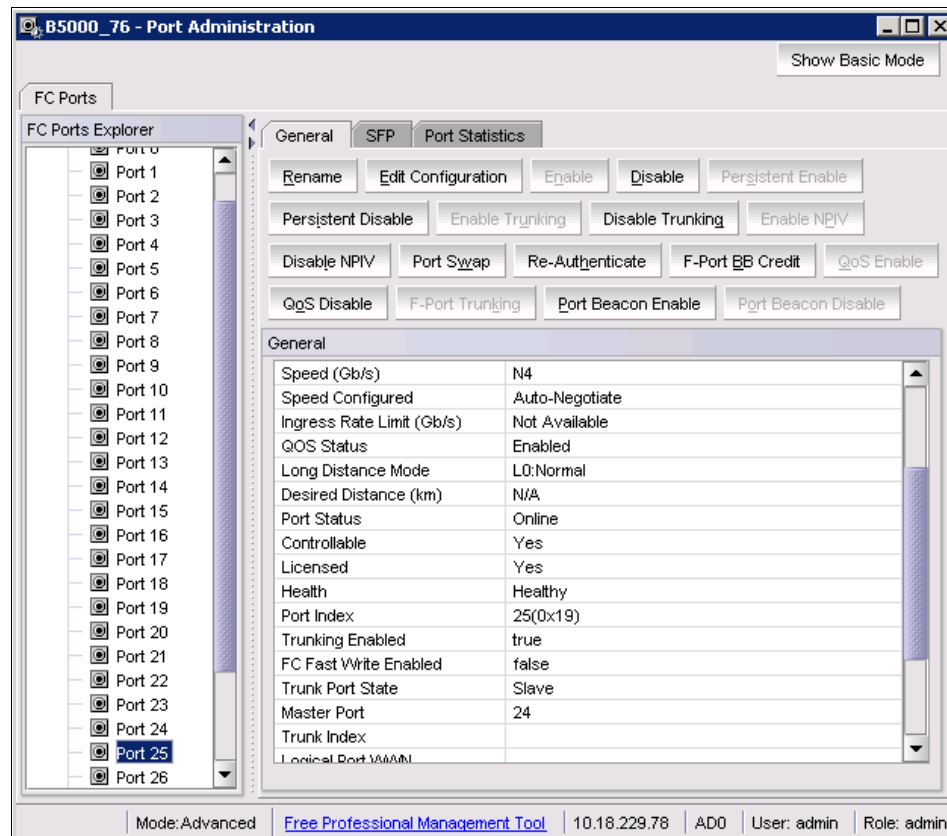


Figure 13-3 Web Tools Port tab additional details

13.1.4 Connecting switches over distance

In these sections we show how to connect switches and fabrics over extended distances.

Extended Fabrics

Extended Fabrics software optimizes switch buffering to ensure the highest possible performance on ISLs. When Extended Fabrics is installed on gateway switches, the ISLs (E_Ports) are configured with a large pool of buffer credits.

The enhanced switch buffers help ensure that data transfer can occur at near-full bandwidth to efficiently utilize the connection over the extended links.

The Extended Fabrics feature extends the distance the ISLs can reach over a dark fiber or wave division multiplexing (WDM) connection. This is accomplished by providing enough buffer credits on each side of the link to compensate for latency introduced by the extended distance.

Licensing

A Brocade Extended Fabrics license is required before you can implement long distance dynamic (LD) and long distance static (LS) distance levels. The LD and LS settings are necessary to achieve maximum performance results over Inter-Switch Links (ISLs) that are greater than 10 km.

Distance: Performance can vary depending on the condition of the fiber optic connections between the switches. Losses due to splicing, connectors, tight bends, and other degradation can affect the performance over the link and the maximum distance that is possible.

Configuring Extended Fabrics

You can configure ports to support long distance links through Telnet or using Web Tools interfaces.

There are seven possible long distance levels for a port (shown in Table 13-1). Fabric OS v6.x and later only supports modes L0, LE, LD, and LS.

Ports are arranged in port groups (different than port groups for trunking), with a common pool of buffer credits to draw from. Certain buffers are dedicated for each port, and others are shared among the ports. In L0 mode, which is normal port mode, ports are usually given 8 buffer credits, which satisfies most distances within a data center. In LE mode, ports reserve a set amount of buffer credits depending on link speed to support distances up to 10 km. L0 and LE modes do not require an Extended Fabric license.

In Extended Fabric mode, one port is given an increase of dedicated buffers from this pool. Modes L0.5, L1, and L2 reserve a dedicated number of increased buffer credits depending on link speed to support a defined distance. Mode LD has the port calculate dynamically how many buffer credits to allocate itself based on distance calculated during port initialization. You can set an upper limit on distance. Mode LS calculates a static number of buffer credits to allocate a port based on a desired distance value.

The total number of frame buffers in a port group is limited, and the Extended ISL Modes matrix introduces a combination of long distance modes that are available, as shown in Table 13-1.

Table 13-1 Extended ISL Modes

Mode	Buffer Allocation			Distance	Distance	Distance	Distance	License
	1 Gbps	2 Gbps	4 Gbps	@1 Gbps	@2 Gbps	@4 Gbps	@8 Gbps	Required
L0	5(26)	5(26)	5(26)	10 km	5 km	5 km	5 km	No
LE	11	16	26	10km	10 km	10 km	10 km	No
L0.5	18	31	56	25 km	25 km	25 km	NA	Yes
L1	31	56	106	50 km	50 km	50 km	NA	Yes
L2	56	106	206	100 km	100 km	100 km	NA	Yes
LD	Auto	Auto	Auto	Auto (Max 500 km)	Auto (Max 250 km)	Auto (Max 100 km)	Auto Max 100 Km Depends of Switch Model	Yes
LS	varies	varies	varies	varies (Max 500 km)	varies (Max 250 km)	varies (Max 100 km)	Auto Max 100 Km Depends of Switch Model	Yes

Support: Long distance modes L0.5, L1, and L2 are not supported on Fabric OS v6.x.

The buffer allocation and distance vary in this table based upon user specified distances.

For dynamic long distance links using mode LD and LS, you can approximate the number of buffer credits that are reserved using the following formula:

$$(\text{Reserved Buffer for Distance Y}) = (X * \text{LinkSpeed} / 2) + 6$$

Where:

- ▶ X = the distance in kilometers.
- ▶ LinkSpeed = the link speed in Gbps
- ▶ 6 = the number of buffer credits reserved for Fabric Services, Multicast, and Broadcast traffic. This is a static number.

Example 13-3 shows the calculation.

Example 13-3 Calculating reserved buffers for extended links

Distance = 50km

Link speed = 4Gbps

Formula: $(50 * 4 / 2) + 6 = 106$

106 buffers will be reserved for the given port when a 50km cable is connected and longdistance mode LD or LS is configured

Distance: For IBM 8 Gbps switches, the number of free or reserved buffers is not the same in all models. So the maximum long distance varies for each switch model. Consult your switch vendor for the maximum distance that is supported for your switch model.

Configuring the port for extended distance using CLI

You can configure a port to support long distance links using the Telnet command **portcfglongdistance** or by using Web Tools as in Example 13-4.

Example 13-4 Setting the port long distance parameter

```
IBM_SAN80B_217:FID128:admin> portcfglongdistance 75 LD 1 50
Reserved Buffers = 106
```

```
IBM_SAN80B_217:FID128:admin> portshow 74
portName:
portHealth: HEALTHY
```

```
Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x103 PRESENT ACTIVE E_PORT G_PORT U_PORT
portType: 18.0
POD Port: Port is licensed
portState: 1 Online
portPhys: 6 In_Sync
portScn: 64 Segmented Flow control mode 0
port generation number: 70
portId: 024a00
portIfId: 43020809
portWwn: 20:4a:00:05:1e:09:97:01
portWwn of device(s) connected:
```

Distance: auto (desired = 50 Km)
portSpeed: 4Gbps

The example set the port for a distance of 50 kilometers at 4Gbps speed.

Configuring the port for extended distance using GUI

As shown in Figure 13-4, the Extended Fabric tab within Web Tools allows you to configure long distance ports. The director type switches have slot subtabs when configuring a given port.

For all other models, you just highlight the given port that you want to configure as long distance.

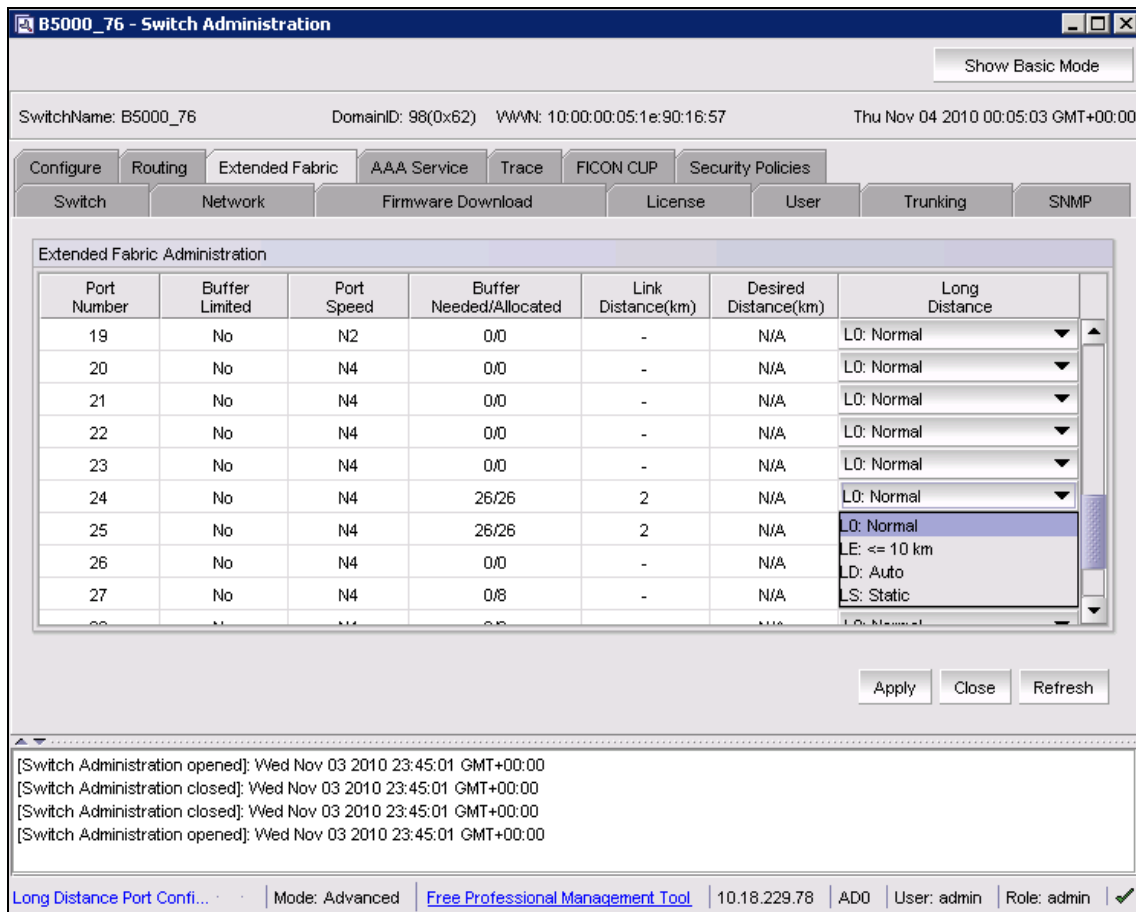


Figure 13-4 Steps to change the Extended Distance Mode

After highlighting the port to configure, go to the Long Distance column on the far right-hand side, and click the down arrow to show the options available for configuration. Table 13-2 lists the details with the Extended Fabric tab.

Table 13-2 Extended Fabric configuration

Port Number	Port Number for all switch models
Buffer Limited	If large distances are set onto various ports within an 8-port block, the remaining ports within that block might have to have their allocated buffer count reduced to enable the long distance configuration.
Port Speed	1G, 2G, 4G, 8G as set speeds. N1, N2 N4 as negotiated speeds.
Buffer Needed/Allocated	Actual buffer usage of port.
Link Distance	Real distance in kilometers.
Desired Distance	Desired distance in kilometers for the port based on port speed.
Long Distance	<p>L0 = Normal value, long distance disabled LE = Extended normal enabled (up to 10 km)</p> <p>The following items require Extended Fabric License:</p> <ul style="list-style-type: none"> ▶ LD = Dynamic link enabled, operates at distances up to 500 km for 1 Gbps, 250 km for 2 Gbps, or 125 km for 4 Gbps and 100 km for 8 Gbps depending upon frame buffer availability within the port group and the switch model. ▶ LS = Static setting enabled. Buffer credits statically configured based on link distance, operates at distances up to 500 km for 1 Gbps, 250 km for 2 Gbps, or 125 km for 4 Gbps and 100 km for 8 Gbps depending upon frame buffer availability within the port group and switch model.
Slot Number tab	Tab for the slots in the director type switches that display the ports on the given slot for the logical switch.
Apply	Apply and commit changes to the switch.
Close	Close Administrator window.
Refresh	Refresh the view with current data from the switch.

13.1.5 Routing policies

This section discusses the routing policies that are available to tune routing performance.

Attention: For most configurations, the default routing policy is optimal and provides the best performance. Therefore, change the routing policy *only* if there is a performance issue that is of concern or if a particular fabric configuration requires it.

Routing can be configured and monitored using the GUI or CLI. Next we show an example of how to view the current setting using the GUI as well as CLI.

Figure 13-5 shows the Routing tab with the default Exchange-Based-Routing policy enabled. You can alternatively select Port-Based-Routing. Changing this setting requires the switch to be disabled.

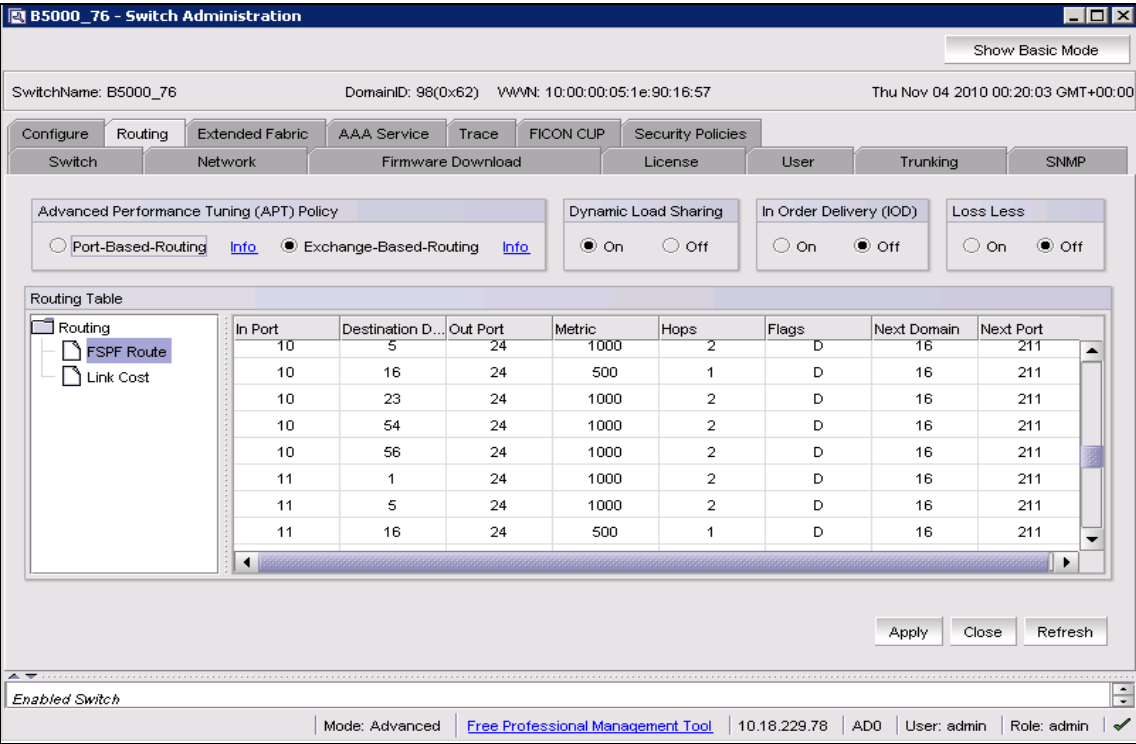


Figure 13-5 Routing tab

Example 13-5 is the same setting as viewed from the CLI.

Example 13-5 Viewing the routing policy

```
IBM_SAN80B_217:FID128:admin> aptpolicy  
Current Policy: 3 0(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
```

```
IBM_SAN80B_217:FID128:admin>
```

Exchange-based routing

The choice of routing path is based on the Source ID (SID), Destination ID (DID), and Fibre Channel originator exchange ID (OXID), optimizing path utilization for the best performance.

In exchange-based routing, every exchange can take a different path through the fabric. Exchange-based routing requires the use of the Dynamic Load Sharing (DLS) feature; when this policy is in effect, you cannot disable the DLS feature.

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

- ▶ AP shared link policy (default)
- ▶ AP dedicated link policy

The AP dedicated link policy dedicates some links to egress traffic and some to ingress traffic.

Port-based routing

The choice of routing path is based only on the incoming port and the destination domain. To optimize port-based routing, DLS can be enabled to balance the load across the available output ports within a domain.

Using port-based routing, you can assign a static route, in which the path chosen for traffic does not change when a topology change occurs unless the path becomes unavailable. If the static route violates FSPF, it is not used. In contrast, Exchange Based Routing policies always employ dynamic path selection.

Support: Static routing is a legacy setting and is currently only supported on the SAN40B and SAN80B switches. For other products, as an alternative, you can use the traffic isolation feature to create a dedicated path for interswitch traffic.

Dynamic Load Sharing

Routing is based generally on the incoming port and the destination domain. Thus, all the traffic coming in from a port (either E_Port or EX_Port) directed to the same remote domain is routed through the same output E_Port.

To optimize fabric routing when there are multiple equivalent paths to a remote switch, traffic is shared among all the paths. Load sharing is recomputed when a switch is booted up or every time a change in the fabric occurs. A change in the fabric is defined as an E_Port going up or down, or an EX_Port going up or down.

In an IBM fabric, if Dynamic Load Sharing (DLS) is turned off, load sharing is performed only at boot time or when an E_Port comes up. Optimal load sharing is rarely achieved with DLS disabled.

If DLS is turned on, routing changes can affect working ports. For example, if an E_Port goes down, another E_Port can be rerouted from one E_Port to a different E_Port. The switch minimizes the number of routing changes, but some are necessary in order to achieve optimal load sharing.

Turning on DLS can affect performance when using it in conjunction with the in-order delivery option.

In-order delivery

You can use the in-order delivery option to enforce in-order delivery of frames during a fabric topology change. In a stable fabric, frames are always delivered in-order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for example, a link goes down), traffic is rerouted around the failure, which can cause frames to be delivered out of order. This option ensures that frames are not delivered out of order, even during fabric topology changes by implementing a timeout value after a fabric change before sending or dropping the next frame.

In an IBM fabric, the in-order delivery option is by default set to *on*.

Use this option with care, because it can cause a delay in the establishment of a new path when a topology change occurs. Only use this option if there are devices connected to the fabric that cannot tolerate the occasional out of order delivery of frames.

You can change the routing policy using the **aptPolicy** command, but you must disable the switch first.

Example 13-6 shows the steps to change the routing policy from the default Exchange Based Routing to port-based routing.

Example 13-6 shows how to change the routing policy from default to port-based routing.

Example 13-6 Policy check and change

```
IBM_SAN80B_217:FID128:admin> aptpolicy  
Current Policy: 3 0(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
```

```
IBM_SAN80B_217:FID128:admin> switchdisable
```

```
IBM_SAN80B_217:FID128:admin> aptpolicy 1  
Policy updated successfully.
```

```
IBM_SAN80B_217:FID128:admin> switchenable
```

Lossless

Lossless DLS enables Dynamic Load Sharing for optimal utilization of the ISLs without causing any frame loss. Note that frame loss can be guaranteed only when a new additional path is used to do load rebalancing. Frame loss cannot be guaranteed on an existing data path that encounters failure. FOS v6.4.0 adds support for the Lossless DLS with DPS (Exchange based routing).

The In Order Delivery (IOD) capability can be enabled optionally for both Port Based Routing and Exchange Based Routing policies. In pre-FOS v6.4.0 versions the Lossless DLS feature was supported only for Port Based Routing and IOD was always enabled. This feature is enabled using the Lossless option shown in Figure 13-5 on page 601.

In Virtual Fabrics, lossless DSL can be enabled on a per logical switch basis. It is best that the logical switch be defined at an ASIC boundary so that ports from the same ASIC are not assigned to a different logical switch.

Restriction: The entire path in the switching fabric must be 8 Gbps ASIC's and is not supported on the advanced function blades and switches.

FSPF Route

As shown in Figure 13-6, when you select the FSPF Route option (highlighted) under the *Routing* tree, the main area of the window then displays the FSPF routing table, including the destination domain and port, hop count, and the metric being the cost assigned to that link.

The screenshot shows the 'B5000_76 - Switch Administration' window. The 'Routing' tab is selected, and the 'FSPF Route' option is highlighted in the left-hand tree. The main area displays the 'Routing Table' with the following data:

In Port	Destination D...	Out Port	Metric	Hops	Flags	Next Domain	Next Port
10	5	24	1000	2	D	16	211
10	16	24	500	1	D	16	211
10	23	24	1000	2	D	16	211
10	54	24	1000	2	D	16	211
10	56	24	1000	2	D	16	211
11	1	24	1000	2	D	16	211
11	5	24	1000	2	D	16	211
11	16	24	500	1	D	16	211

At the bottom of the window, the status bar shows 'Enabled Switch', 'Mode: Advanced', 'Free Professional Management Tool', '10.18.229.78', 'AD0', 'User: admin', 'Role: admin', and a green checkmark.

Figure 13-6 FSPF Routing Table Details

Table 13-3 defines the different columns.

Table 13-3 FSPF Route field descriptions

Field	Description
In Port	Displays the Port number where the frames enter the switch.
Destination Domain	Displays the destination domain ID for the participating static routes for a particular In Port. The destination domain is the target of the out port.
Out Port	Displays the Out port. It should be within the range of ports that are available for static routes for the current domain. More than one out port can be used for any In port with a different domain ID. Each domain ID requires an out port.
Metric	Displays the calculated cost of reaching the destination domain.
Hops	Displays the number of hops in the “shortest path” route.
Flags	Displays whether the route is Static (<i>S</i>) or Dynamic (<i>D</i>).
Next Domain	Displays the next domain ID in the routing path. The Next Domain is the switch that the “Out Port” is connected to.
Next Port	Displays the next Port in the routing path. The Next Port is the port number that the “Out Port” is physically connected to.

Static Route

A *static route* is a route that defines a specific path and does not change when a topology change occurs, unless the path that is defined by the route becomes unavailable.

A static route can be assigned only when the active routing policy is port-based routing. When exchange-based routing is active, you cannot assign static routes.

A reason for configuring static routes is that some devices (can be legacy storage devices) do not tolerate out-of-order exchanges; in such cases, use the port-based routing policy.

Support: Static routes are supported only on the IBM SAN40B and SAN80B platforms.

Link cost

This next option under the Routing tree allows you to view the link cost for a specific link, as shown in Figure 13-7. By double-clicking in the Cost field for the specific port, you can modify the cost. This setting has an effect on the cost value that the local switch has for this link. It uses this value to calculate the lowest cost path to a destination on other switches within the fabric. For a 1 Gbps ISL, the default cost is 1000. For a 2/4/8 Gbps ISL, the default cost is 500. Valid values for link cost are from 1 to 9999.

The screenshot shows the 'B5000_76 - Switch Administration' window. The 'Routing' tab is selected, and the 'Link Cost' option is chosen from the left-hand tree. The main area displays a table with 'Port Number' and 'Cost' columns. Ports 0 through 7 are listed, all with a cost of 500. Port 6 is highlighted. Below the table are 'Apply', 'Close', and 'Refresh' buttons. At the bottom, a status bar shows 'Enabled Switch' and various system details.

Port Number	Cost
0	500
1	500
2	500
3	500
4	500
5	500
6	500
7	500

Mode: Advanced | [Free Professional Management Tool](#) | 10.18.229.78 | AD0 | User: admin | Role: admin | ✓

Figure 13-7 Routing link cost

13.2 Merging fabrics

Merging a fabric occurs where two or more separate fabrics are combined, as shown in Figure 13-8.

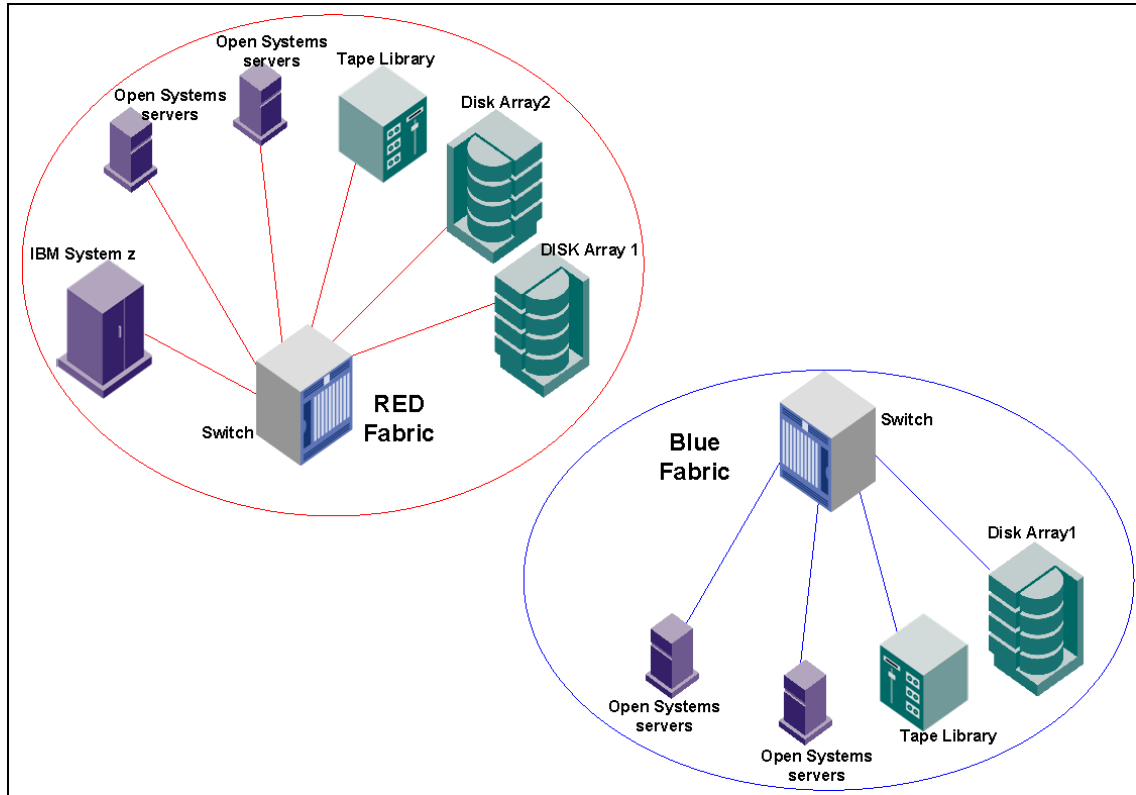


Figure 13-8 Two separate SAN fabrics

These separate SAN fabrics can be merged to form a larger SAN fabric by connecting the switches using an ISL, as shown in Figure 13-9.




Figure 13-9 A merged fabric

The zoning information of the different fabric is merged when the different fabrics are connected together, assuming that there are no conflicting definitions.

A merge might occur when an organization acquires another company or when two business units within one company merge. The result is that a SAN fabric is extended through the addition of another complete fabric.

Important: A SAN switch should be disabled before adding it to an existing fabric.

Some conflicts might occur when two fabrics are merged. The most common sources of conflict are:

- ▶ Duplicate domain ID
- ▶ Zoning configuration conflicts
- ▶ Operating parameters inconsistency (for example, core PID format)
- ▶ InteropMode when merging IBM b-type switches with IBM m-type switches

When issues occur, part of the SAN fabric is said to be *segmented*. You can identify a segmentation using the **switchShow** commands or by the flashing orange LED on the ISL port.

The following sections describe these conflicts and possible solutions.

13.2.1 Duplicate domain IDs

Domain IDs are used to uniquely identify a switch within a fabric. Therefore, each switch within the same fabric must have a unique domain ID. Duplicate domains causes the ISL between the two switches to be segmented, as shown in Figure 13-10.

Switch Events, Information		
Switch Events Switch Information		
All Events Last Updated: Fri Jul 24 2009 19:54:30 GMT+00:00 (Auto-Refresh interval is 15 seconds)		
Filter, Show		
Time	Level	Message
Fri Jul 24 2009 18:54:43 GMT+00:00	Warning	Effective Insistent domain ID for the fabric changed from OFF to ON
Fri Jul 24 2009 18:54:46 GMT+00:00	Information	Switch status changed from MARGINAL to HEALTHY.
Fri Jul 24 2009 18:54:54 GMT+00:00	Error	Port 8 Disabled: Insistent Domain ID 1 could not be obtained. Principal Assigned Domain ID = 2.
Fri Jul 24 2009 18:54:54 GMT+00:00	Error	Port 9 Disabled: Insistent Domain ID 1 could not be obtained. Principal Assigned Domain ID = 2.
Fri Jul 24 2009 18:54:54 GMT+00:00	Error	Port 56 Disabled: Insistent Domain ID 1 could not be obtained. Principal Assigned Domain ID = 2.
Fri Jul 24 2009 18:54:54 GMT+00:00	Error	Port 57 Disabled: Insistent Domain ID 1 could not be obtained. Principal Assigned Domain ID = 2.

Figure 13-10 Domain ID segmentation error log

To solve this overlap, change the domain ID of one of the switches participating in the ISL using the Web Tools interface in the Switch Admin tab or using the CLI **configure** command. An overview of used fabric ID's can be retrieved from DCFM or the **fabricshow** CLI command.

You can avoid domain ID overlap easily by disabling the Insistent Domain ID function. This is done with CLI using the **switchDisable** command. When bringing back the switches online, the domain ID automatically is negotiated and set to a valid value as shown in Example 13-7.

Example 13-7 Steps to disable Insistent Domain ID mode

```
IBM_SAN80B_217:FID128:admin> switchdisable
IBM_SAN80B_217:FID128:admin> configure
```

Configure...

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

Domain: (1..239) [1]

```

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

```

```

Insistent Domain ID Mode (yes, y, no, n): [no] n
Virtual Channel parameters (yes, y, no, n): [no] CTRL-D to exit

```

```

IBM_SAN80B_217:FID128:admin> switchenable

```

```

IBM_SAN80B_217:FID128:admin> fabricshow

```

```

Switch ID      Worldwide Name      Enet IP Addr Name
-----
1: fffc01 10:00:00:05:1e:94:3a:00 10.64.210.213 "IBM_SAN384B_213"
2: fffc02 10:00:00:05:1e:09:97:01 10.64.210.217 "IBM_SAN80B_217"
4: fffc04 10:00:00:05:1e:76:68:00 10.64.210.51 >"IBM_B32_51"
10: fffc0a 10:00:00:05:1e:76:86:80 10.64.210.50 "IBM_B32_50"

```

```

The Fabric has 4 switches

```

The example shows that even if the Domain ID is set to 1 for the SAN80B switch, the Insistent Domain ID function selects Domain ID 2 for the switch, which allows the switch to merge with the other switches in the fabric.

13.2.2 Zoning configuration conflicts

When merging two fabrics, zoning information from the two previously separate fabrics is merged into the new fabric (as much as possible). If there are effective configurations active in both fabrics, they must match exactly. In a defined or effective configuration, if zone objects have the same name in each fabric, their type, content, and order of content must also match.

Figure 13-11 shows an example of segmentation due to zoning.

Switch Events, Information		
Switch Events Switch Information		
All Events Last Updated: Fri Jul 24 2009 20:16:44 GMT+00:00 (Auto-Refresh interval is 15 seconds)		
Filter Show		
Time	Level	Message
Fri Jul 24 2009 20:08:18 GMT+00:00	Error	Port 8 Disabled: Insistent Domain ID 1 could not be obtained. Principal Assigned Domain ID = 2.
Fri Jul 24 2009 20:08:18 GMT+00:00	Error	Port 4 Disabled: Insistent Domain ID 1 could not be obtained. Principal Assigned Domain ID = 2.
Fri Jul 24 2009 20:08:18 GMT+00:00	Error	Port 56 Disabled: Insistent Domain ID 1 could not be obtained. Principal Assigned Domain ID = 2.
Fri Jul 24 2009 20:08:18 GMT+00:00	Error	Port 57 Disabled: Insistent Domain ID 1 could not be obtained. Principal Assigned Domain ID = 2.
Fri Jul 24 2009 20:08:24 GMT+00:00	Information	The last fabric change happened at: Fri Jul 24 20:08:18 2009
Fri Jul 24 2009 20:08:39 GMT+00:00	Information	The last device change happened at: Fri Jul 24 20:08:36 2009
Fri Jul 24 2009 20:14:43 GMT+00:00	Information	The effective configuration has changed to SiteA_fab1A.
Fri Jul 24 2009 20:15:14 GMT+00:00	Information	The last fabric change happened at: Fri Jul 24 20:15:12 2009
Fri Jul 24 2009 20:15:19 GMT+00:00	Warning	Switch status changed from HEALTHY to MARGINAL.
Fri Jul 24 2009 20:15:19 GMT+00:00	Warning	Switch status change contributing factor Switch offline.
Fri Jul 24 2009 20:15:53 GMT+00:00	Warning	Effective insistent domain ID for the fabric changed from ON to OFF
Fri Jul 24 2009 20:15:54 GMT+00:00	Warning	port 9, Zone Conflict.
Fri Jul 24 2009 20:15:55 GMT+00:00	Warning	port 57, Zone Conflict.
Fri Jul 24 2009 20:15:55 GMT+00:00	Information	Switch status changed from MARGINAL to HEALTHY.
Fri Jul 24 2009 20:16:04 GMT+00:00	Information	The last fabric change happened at: Fri Jul 24 20:15:55 2009
Fri Jul 24 2009 20:16:24 GMT+00:00	Information	The last device change happened at: Fri Jul 24 20:16:13 2009

Figure 13-11 Zoning conflict

13.2.3 Merging fabrics example

In this section we demonstrate how to merge two fabrics, each containing defined zoning configurations. The scenario is that two individual SANs are going to be interconnected and merged into a single fabric.

In our example, we merge fabric 1 in each site. We call the two configurations SiteA_fab1 (SAN80B switch) and SiteB_fab1 (SAN384B switch). We build in a non-valid zone configuration in order to demonstrate segmentation of the switches, and to show the resolution. We start with Example 13-8.

Example 13-8 Configuration SiteA before merge

```
IBM_SAN80B_217:FID128:admin> cfgshow
Defined configuration:
cfg:  SiteA_fab1
      z1_BL1_DS4000; z1_AIX_hba1_DS4000
zone:  z1_AIX_hba1_DS4000
      AIX_hba1; DS4000
zone:  z1_BL1_DS4000
      Blade1_hba1; DS4000
alias: AIX_hba1
      1,20
```

```

alias: Blade1_hba1
      1,21
alias: DS4000 1,22

Effective configuration:
cfg:   SiteA_fab1
zone:  z1_AIX_hba1_DS4000
      1,20
      1,22
zone:  z1_BL1_DS4000
      1,21
      1,22

```

```
IBM_SAN80B_217:FID128:admin>
```

We have similar alias names in both fabrics, and their content is different. If similar alias names exist in two fabrics to be merged, they must have the same content, or the fabrics will segment and the merge will fail. Example 13-9 shows SiteB configuration before the merge.

Example 13-9 Configuration SiteB before merge

```

IBM_SAN384B_213:FID128:admin> cfgshow
Defined configuration:
cfg:   SiteB_fab1
      z1_TSM_DS4000; z1_AIX_DS4000
zone:  z1_AIX_DS4000
      AIX_hba0; DS4000
zone:  z1_TSM_DS4000
      TSMserver_hba1; tape_lib_drive1; DS4000
alias: AIX_hba0
      2,10
alias: DS4000 2,11
alias: TSMserver_hba1
      2,12
alias: tape_lib_drive1
      2,13

Effective configuration:
cfg:   SiteB_fab1
zone:  z1_AIX_DS4000
      2,10
      2,11
zone:  z1_TSM_DS4000
      2,12

```

2,13
2,11

IBM_SAN384B_213:FID128:admin

We interconnect the two switches at this point by enabling port 57 on our switch. Port 57 in our example is one end of an ISL-connection between the two switches.

The DS4000 alias exist in both fabrics, hence the two fabrics will segment and merge fails. This can be seen with the output of **switchshow** or **portshow**. Example 13-10 is an example of checking with **switchshow**.

Example 13-10 Switchshow shows segmented port 57

```
IBM_SAN80B_217:FID128:admin> switchshow
switchName:      IBM_SAN80B_217
switchType:      64.3
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:    1 (unconfirmed)
switchId:        fffc01
switchWwn:       10:00:00:05:1e:09:97:01
zoning:          ON (SiteA_fab1)
switchBeacon:    OFF
FC Router:       OFF
Allow XISL Use:  OFF
LS Attributes:   [FID: 128, Base Switch: No, Default Switch: Yes]
```

```
Area Port Media Speed State      Proto
=====
  0   0   --   N8   No_Module
.
.
57  57   id   N8   Online          E-Port  segmented, (zone
conflict) (Trunk master)
```

We now have to fix any zoning configuration errors that might exist, and in our example we do this by renaming the DS4000 alias in the Site A fabric as shown in Example 13-11.

Example 13-11 Rename the DS4000 alias

```
IBM_SAN80B_217:FID128:admin> zoneobjectrename DS4000, DS4000_A
```

```
IBM_SAN80B_217:FID128: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 ...

```
IBM_SAN80B_217:FID128:admin>
```

Next we have to disable the zoning configuration on one of the switches. By not doing so, the fabrics would segment and the merge would fail. We therefore choose which zoning configuration we want to disable. In Example 13-12 we disable the zoning configuration for SiteA, so that the effective zoning configuration will be the one from SiteB.

Example 13-12 Disable zoning configuration

```
IBM_SAN80B_217:FID128:admin> cfgdisable
```

You are about to disable zoning configuration. This action will disable any previous zoning configuration enabled.

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

Updating flash ...

```
IBM_SAN80B_217:FID128:admin> portenable 57
```

```
IBM_SAN80B_217:FID128:admin> switchdisable
```

```
IBM_SAN80B_217:FID128:admin> switchenable
```

```
IBM_SAN80B_217:FID128:admin> fabricshow
```

Switch ID	Worldwide Name	Enet IP	Addr Name
-----------	----------------	---------	-----------

1:	fffc01 10:00:00:05:1e:09:97:01	10.64.210.217	>"IBM_SAN80B_217"
2:	fffc02 10:00:00:05:1e:94:3a:00	10.64.210.213	"IBM_SAN384B_213"

The Fabric has 2 switches

Our two fabrics have successfully merged, and the configuration from SiteB is now the effective configuration. The combined configuration is containing all the zoning elements from both fabrics. However, zones from the previously disabled configuration in Site A have to be added to the effective configuration.

In Example 13-13 we show the fabric after the merge.

Example 13-13 Fabric after merge

```
IBM_SAN384B_213:FID128:admin> cfgshow
Defined configuration:
  cfg:  SiteA_fab1
        z1_BL1_DS4000; z1_AIX_hba1_DS4000
  cfg:  SiteB_fab1
        z1_TSM_DS4000; z1_AIX_DS4000
  zone: z1_AIX_DS4000
        AIX_hba0; DS4000
  zone: z1_AIX_hba1_DS4000
        AIX_hba1; DS4000_A
  zone: z1_BL1_DS4000
        Blade1_hba1; DS4000_A
  zone: z1_TSM_DS4000
        TSMserver_hba1; tape_lib_drive1; DS4000
  alias: AIX_hba0
        2,10
  alias: AIX_hba1
        1,20
  alias: Blade1_hba1
        1,21
  alias: DS4000 2,11
  alias: DS4000_A
        1,22
  alias: TSMserver_hba1
        2,12
  alias: tape_lib_drive1
        2,13

Effective configuration:
  cfg:  SiteB_fab1
  zone: z1_AIX_DS4000
        2,10
        2,11
  zone: z1_TSM_DS4000
        2,12
        2,13
        2,11

IBM_SAN384B_213:FID128:admin>
```

We now add the zoning elements from the disabled configuration SiteA_fab1 to the effective configuration SiteB_fab1. This can be accomplished using Web Tools, or using the CLI as in Example 13-14 on page 617.

Example 13-14 Adding zones to the effective configuration

```
IBM_SAN80B_217:FID128:admin> cfgadd SiteB_fab1, "z1_AIX_hba1_DS4000"
```

```
IBM_SAN80B_217:FID128:admin> cfgadd SiteB_fab1, "z1_BL1_DS4000"
```

```
IBM_SAN80B_217:FID128:admin> cfgdelete SiteA_fab1
```

```
IBM_SAN80B_217:FID128:admin> cfsave
```

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 ...

```
IBM_SAN80B_217:FID128:admin> cfgenable SiteB_fab1
```

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 'SiteB_fab1' configuration (yes, y, no, n): [no] y

zone config "SiteB_fab1" is in effect

Updating flash ...

```
IBM_SAN80B_217:FID128:admin>
```

The two fabrics are now merged into one single configuration. The zone objects in the SiteA_fab1 configuration were added to the SiteB_fab1 configuration, and SiteB_fab1 becomes the effective configuration. Example 13-15 shows the final configuration.

Example 13-15 The final merged configuration

```
IBM_SAN80B_217:FID128:admin> cfgshow
```

Defined configuration:

```
cfg:  SiteB_fab1
```

```
      z1_TSM_DS4000; z1_AIX_DS4000; z1_AIX_hba1_DS4000;
```

```
      z1_BL1_DS4000
```

```

zone:  z1_AIX_DS4000
        AIX_hba0; DS4000
zone:  z1_AIX_hba1_DS4000
        AIX_hba1; DS4000_A
zone:  z1_BL1_DS4000
        Blade1_hba1; DS4000_A
zone:  z1_TSM_DS4000
        TSMserver_hba1; tape_lib_drive1; DS4000
alias: AIX_hba0
        2,10
alias: AIX_hba1
        1,20
alias: Blade1_hba1
        1,21
alias: DS4000  2,11
alias: DS4000_A
        1,22
alias: TSMserver_hba1
        2,12
alias: tape_lib_drive1
        2,13

```

Effective configuration:

```

cfg:  SiteB_fab1
zone:  z1_AIX_DS4000
        2,10
        2,11
zone:  z1_AIX_hba1_DS4000
        1,20
        1,22
zone:  z1_BL1_DS4000
        1,21
        1,22
zone:  z1_TSM_DS4000
        2,12
        2,13
        2,11

```

IBM_SAN80B_217:FID128:admin>

Attention: When two switches are being interconnected, the zoning configurations on one of the switches must be disabled, or the merge will fail. This includes the default-zone if set to no access. This hidden zone must also be disabled. To do this, use the command **defZone --allaccess**.

The configuration must be applied with the **cfgsave** command.

13.2.4 Merging with a configuration cleared switch

Another solution is to make sure that the switch you are adding to the fabric is cleared of any zoning information, by following this process:

1. Issue the **switchDisable** command to disable the switch.
2. Disable the active configuration using the **cfgDisable** command.
3. Issue the **cfgClear** command to clear all zoning information.
4. Issue the **defZone -allaccess** command to set the default mode to all access.
5. Issue the **switchEnable** command to enable the switch, as shown in Example 13-16.

Example 13-16 Merging the fabric by clearing the configuration of a fabric

```
IBMIBM_SAN80B_217:FID128:admin> switchdisable
```

```
IBM_SAN80B_217:FID128:admin> cfgdisable
```

```
You are about to disable zoning configuration. This
action will disable any previous zoning configuration enabled.
Do you want to disable zoning configuration? (yes, y, no, n): [no] y
Updating flash ...
Effective configuration is empty. "No Access" default zone mode is ON.
```

```
IBM_SAN80B_217:FID128: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]
y
```

```
IBM_SAN80B_217:FID128:admin> defzone --allaccess
```

```
You are about to set the Default Zone access mode to All Access
Do you want to set the Default Zone access mode to All Access ? (yes,
y, no, n): [no] y
```

```

IBM_SAN80B_217:FID128: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 ...

IBM_SAN80B_217:FID128:admin> switchenable

IBM_SAN80B_217:FID128:admin>

```

13.2.5 Operating parameter conflicts

Conflicts due to fabric-wide operating parameters are less common because default values for these settings suit most requirements. However, conflicts can occur when dealing with a multivendor environment or distance solution installations.

Error log messages can vary quite a bit, depending on the source of the problem. It is beyond the scope of this book to discuss all the possible error log messages; however, Figure 13-12 shows an example.

Switch Events		
Switch Information		
All Events Last Updated: Tue Jul 28 2009 16:11:12 GMT+00:00 (Auto-Refresh interval is 15 seconds)		
		Actions
Time	Level	Message
Tue Jul 28 2009 16:10:43 GMT+00:00	Warning	Switch status changed from MARGINAL to HEALTHY.
Tue Jul 28 2009 16:10:46 GMT+00:00	Warning	port 84, ELP rejected by the other switch.
Tue Jul 28 2009 16:10:49 GMT+00:00	Information	The last fabric change happened at: Tue Jul 28 16:10:41 2009
Tue Jul 28 2009 16:10:51 GMT+00:00	Information	Switch status changed from MARGINAL to HEALTHY.
Tue Jul 28 2009 16:11:09 GMT+00:00	Information	The last device change happened at : Tue Jul 28 16:10:59 2009

Figure 13-12 ELP conflict

In this example, there is an exchange link parameter mismatch, which has caused the segmentation.

One solution to this problem is to make sure that the fabric-wide operating parameters are consistent throughout all the participating switches.

Use the **configure** command to set the correct R_A_TOV parameter and other specific parameters and ensure that all parameters, except the domain ID, are identical throughout all the switches in the fabric.

With Fabric OS v6.1.x and later, the Core PID mode is set as the default. When connecting 1-Gbps switches to 8 Gbps switches, the core PID mode in the 1 Gbps switch must be in core PID mode. Devices connected to the 1 Gbps switch must be taken offline while changing the PID mode of the switch.

Reboot: Rebooting the switch might be required in some cases when changing system parameters. Reboot helps to prevent inconsistencies.

13.2.6 InteropMode

You need to consider InteropMode when connecting IBM m-type switches to IBM b-type switches. The different types of InteropMode that you can select include:

- ▶ InteropMode 0: For Brocade Native mode, which supports all stand-alone Brocade fabrics, but no interoperability support
- ▶ InteropMode 1: No longer supported; was the original Open Fabric mode; replaced by InteropMode 3
- ▶ InteropMode 2: For McDATA Fabric mode, which supports M-EOS switches v9.6.2 and later running in McDATA Fabric mode
- ▶ InteropMode 3: For McDATA Open Fabric mode, which supports M-EOS switches v9.6.2 and higher running in Open Fabric mode

Example 13-17 shows the output of the **interopmode** command.

Example 13-17 Output of interopmode

```
IBM_SAN384B_213:FID128: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

IBM_SAN384B_213:FID128:admin>
```

Having a fabric in InteropMode 2 or 3 can exclude some newer features. See the following link for the interoperability guide for your switch model:

<https://www-304.ibm.com/systems/support/supportsite.wss/allproducts?taskind=2&brandind=5000031>

To change InteropMode, see Example 13-18. The switch must be disabled using the **switchDisable** command before issuing the **interopmode** command. The switch is rebooted automatically. Therefore, devices that are connected to the switch must be taken offline also.

Example 13-18 Changing InteropMode

```
IBM_SAN384B_213:FID128:admin> switchdisable
IBM_SAN384B_213:FID128:admin> interopmode 2
The switch effective configuration will be lost.

Interop Mode Will Be Changed and switch will be Enabled

Do you want to continue? (yes, y, no, n): [no] y
InteropMode: McDATA Fabric
      Default Zone: Off
      Safe Zone: Off

IBM_SAN384B_213:FID128:admin>
```

In this specific situation, reboot is not needed, and InteropMode is now shown as McDATA Fabric.



Security

In this chapter we provide information and procedures for configuring basic and advanced Fabric OS v6.4.1 security features such as user account management, Access Control Lists (ACL) policies, authentication policies, and IP Filtering for Brocade's Fibre Channel switches.

Important: The Secure Fabric OS licensed feature is no longer supported or available on Fabric Operating System v6.x and later. All the security features are available in the base Fabric OS starting in Fabric OS v5.3.0 and later.

14.1 User accounts 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.

Each user-defined account is associated with the following elements:

- ▶ Admin Domain list: Specifies the Administrative Domains that a user account is allowed to log in to.
- ▶ Home Admin Domain: Specifies the Admin Domain that the user is logged in to by default. The home Admin Domain must be a member of the user's Admin Domain list.
- ▶ Virtual Fabric list: Specifies the Virtual Fabric that a user account is allowed to log in to.
- ▶ Home Virtual Fabric: Specifies the Virtual Fabric that the user is logged in to by default. The home Virtual Fabric must be a member of the user's Virtual Fabric list.
- ▶ Role: Determines functional access levels within the bounds of the user's current Admin Domain.

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

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

14.1.1 User authentication

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. The local user database is manually synchronized using the **distribute** command to push a copy of the switch's local user database to all other Fabric OS v5.3.0 and later switches in the fabric.

For setting up user authentication through RADIUS or LDAP, we refer you to the *Fabric OS Administrator's Guide*, which is only available through the Partner Network website at the following location (navigate to the Product Documentation and register or login):

<http://www.brocade.com/data-center-best-practices/resource-center/index.page>

14.1.2 Role-Based Access Control

Role-Based Action Control (RBAC) defines the capabilities that a user account has, based on the role that the account has been assigned. For each role, there is a set of predefined permissions on the jobs and tasks that can be performed on a fabric and its associated fabric elements. Fabric OS v6.1.0 and later uses RBAC to determine which commands a user can issue.

When you log in to a switch, your user account is associated with a predefined role. The role that your account is associated with determines the level of access you have on that switch and in the fabric. The chassis-role permission is not a role like the other role types, but a permission that is applied to a user account. You can use the **userConfig** command to add this permission to a user account. For clarity, this permission has been added to Table 14-1, which outlines the Fabric OS predefined roles.

Table 14-1 Fabric OS roles

Role name	FOS version	Duties	Description
Admin	All	All administration	All administrative commands.
BasicSwitch Admin	v5.2.0 and later	Restricted switch administration	Administrative use with a subset of admin-level commands, mostly for monitoring with limited switch (local) access
FabricAdmin	v5.2.0 and later	Fabric and switch administration	All switch and fabric commands; excludes user management and Admin Domains commands
Operator	v5.2.0 and later	General switch administration	A subset of administrative tasks for routine switch maintenance

Role name	FOS version	Duties	Description
SecurityAdmin	v5.3.0 and later	Security administration	All switch security and user management functions
SwitchAdmin	v5.0.0 and later	Local switch administration	Administrative use excluding security, user management, and zoning
User	All	Monitoring only	Non-administrative use, such as monitoring system activity
ZoneAdmin	v5.2.0 and later	Zone administration	Zone management commands only

Admin Domain considerations: Legacy users with no Admin Domain specified and whose current role is admin will have access to AD 0 through 255 (physical fabric admin); otherwise, they will have access to AD0 only.

14.1.3 Local database user accounts

User add, change, and delete operations are subject to the subset rule: An admin with AD_list 0-10 or LF_list 1-10 cannot perform operations on an admin, user, or any role with an AD_list 11-25 or LF_list 11-128. The user account being changed must have an AD_list or LF_list that is a subset of the account that is making the change.

Types of accounts

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

Default accounts

Table 14-2 is a list of the predefined accounts offered by Fabric OS available in the local switch user database. The password for all default accounts must be changed during the initial installation and configuration for each switch.

Table 14-2 lists the default user-accounts.

Table 14-2 Default user accounts

Account name	Role	Admin Domain	Logical fabric	Description
admin	admin	AD0-255 home: 0	LF1-128 home: 128	Observe-modify permission.
factory	factory	AD0-255 home: 0	LF1-128 home: 128	Reserved.
root	root	AD0-255 home: 0	LF1-128 home: 128	Reserved.
user	user	AD0 home: 0	LF1-128 home: 128	Observe-only permission.

14.2 Account management

In the topics that follow, we cover account management.

14.2.1 Displaying account information

To display the account information, you can use Web Tools or the CLI.

Web Tools

From the main Webtools menu select **Switch Admin**; this opens the Switch Administrator window. Select the **User** tab and you get a list of all defined users and roles, as shown in Figure 14-1.

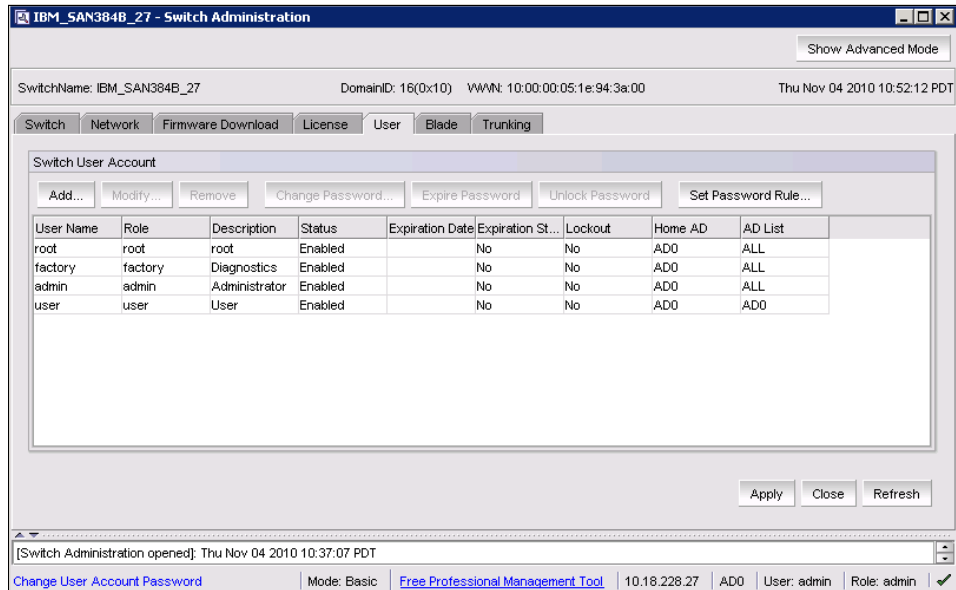


Figure 14-1 User tab

CLI

Follow these steps:

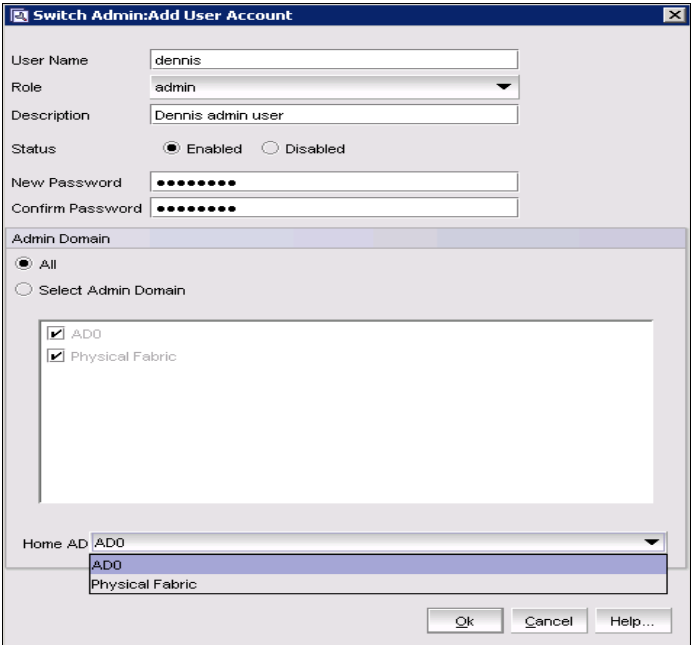
1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the appropriate **show** operands for the account information that you want to display:
 - **userConfig --show -a** to show all account information for a logical switch
 - **userConfig --show username** to show account information for the specified account
 - **userConfig --showad -a adminDomain_ID** to show all accounts permitted to select the specified adminDomain_ID
 - **userConfig --showlf -l logicalFabric_ID** for each LF in an LF_ID_list, which displays a list of users that include that LF in their LF permissions.

14.2.2 Creating an account

In the next sections we give examples of user management, such as creating, deleting, and modifying a local switch user. To modify or add a user, you can use either Web Tools or the CLI.

Web Tools

To create a new user, select the **Add** button in the user window, fill in the user details, and select the role and default admin domain, as shown in Figure 14-2.



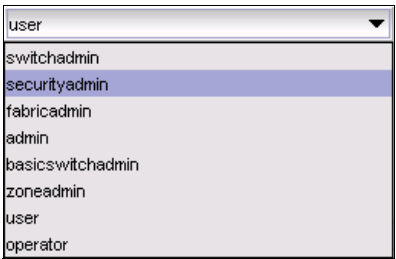
The dialog box titled "Switch Admin: Add User Account" contains the following fields and options:

- User Name: dennis
- Role: admin (dropdown menu)
- Description: Dennis admin user
- Status: ☒ Enabled ☐ Disabled
- New Password: [masked]
- Confirm Password: [masked]
- Admin Domain:
 - ☒ All
 - ☐ Select Admin Domain
 - AD0 (checked)
 - Physical Fabric (checked)
- Home AD: AD0 (dropdown menu with AD0 and Physical Fabric options)
- Buttons: Ok, Cancel, Help...

Figure 14-2 Add user account

As shown in Figure 14-3, under the **Role** tab, you can select the role of the new user.

Password: The default password minimum length is 8 characters.



A dropdown menu showing the following roles:

- user
- switchadmin
- securityadmin
- fabricadmin
- admin
- basicswitchadmin
- zoneadmin
- user
- operator

Figure 14-3 Role

This user will be displayed in the switch administration user window as shown in Figure 14-4.

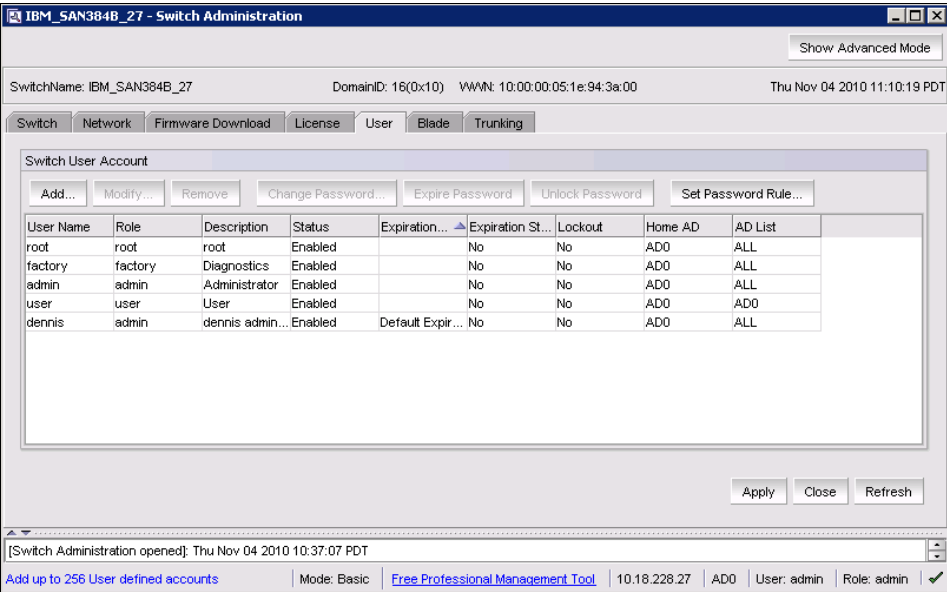


Figure 14-4 New user add

Select the **Apply** button to activate the new user, and you will receive a confirmation window with all details on the new user, as shown in Figure 14-5.

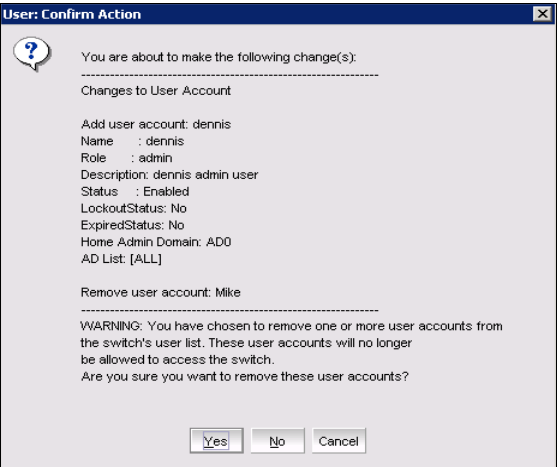


Figure 14-5 New user confirmation

Select the **Yes** button and the new user is active.

CLI

Example 14-1 shows the creation of an account using CLI

Example 14-1 Creating an account

```
IBM_SAN256B_130:admin> userconfig --add dennis -r admin
Setting initial password for dennis
Enter new password:
Re-type new password:
Account dennis has been successfully added.
IBM_SAN256B_130:admin>
```

14.2.3 Modifying User and Account settings

Figure 14-6 shows the options available for switch user account modification.

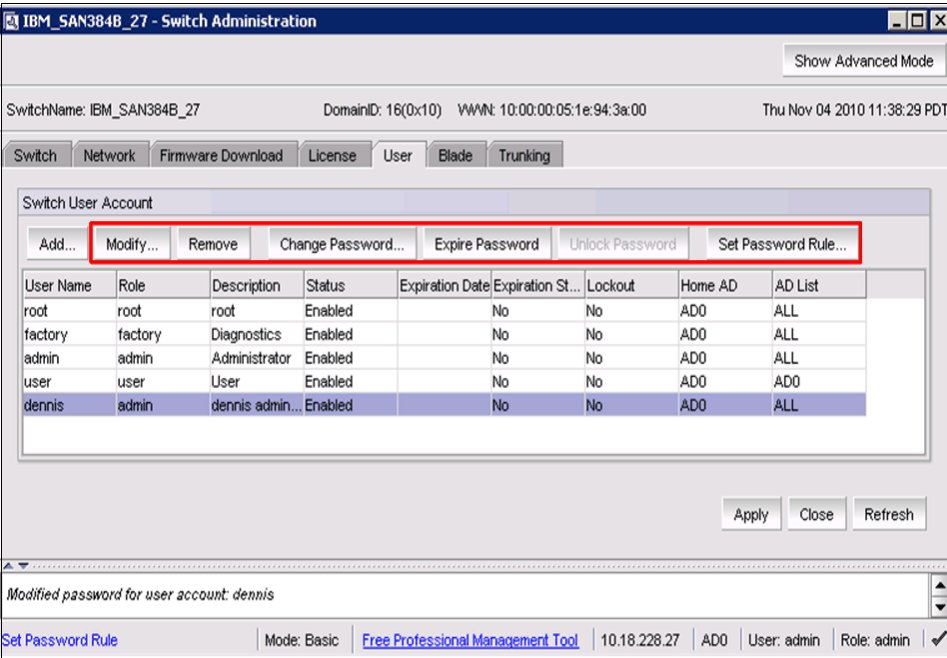
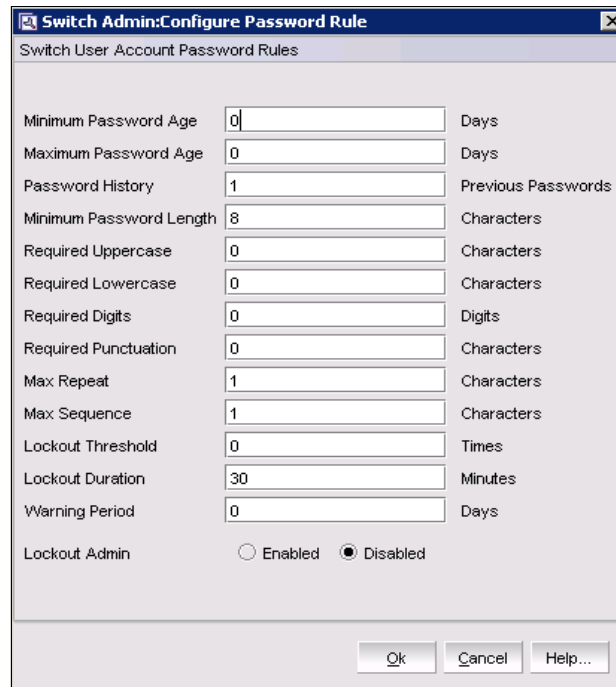


Figure 14-6 User modification tabs

Attention: You are not able to modify a default account.

Changing the password rules

Select the **Set Password Rule** button from the Web Tools User window. This action will open a window where you can change the password rules for all accounts on this switch, as shown in Figure 14-7.



The dialog box titled "Switch Admin: Configure Password Rule" contains the following settings:

Setting	Value	Unit
Minimum Password Age	0	Days
Maximum Password Age	0	Days
Password History	1	Previous Passwords
Minimum Password Length	8	Characters
Required Uppercase	0	Characters
Required Lowercase	0	Characters
Required Digits	0	Digits
Required Punctuation	0	Characters
Max Repeat	1	Characters
Max Sequence	1	Characters
Lockout Threshold	0	Times
Lockout Duration	30	Minutes
Warning Period	0	Days
Lockout Admin	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled	

Buttons: Ok, Cancel, Help...

Figure 14-7 Password rules

Fill out the dialog box for the password rules you want to enforce. Choose whether to enable or disable the lockout administration features.

Lockout: If you choose to disable the lockout administration, the user is never locked out of the system.

Click **OK** to close the dialog box and then click **Apply** and the **Yes** button in the confirmation window to activate your changes. The new user is now active.

Changing the password for a user

From the Web Tools User window, select the user that you want to perform the password change on, and press the **Change Password** button. This will open the Set User Account Password window, shown in Figure 14-8.

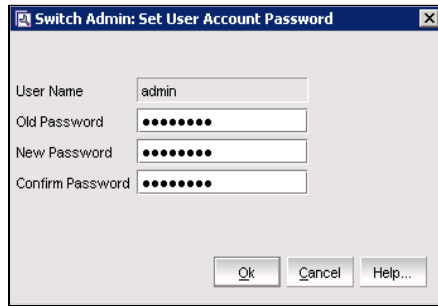


Figure 14-8 Change Password

Enter the old password and then the new password, with confirmation. Click **OK** to close the dialog box and then click **Apply** and the **Yes** button in the confirmation window to activate your changes.

Password: The new password must comply with the password rules set.

Modifying an account

Any user defined account can be modified using the **Modify** button. This will bring up the modify user account window, as shown in Figure 14-9.

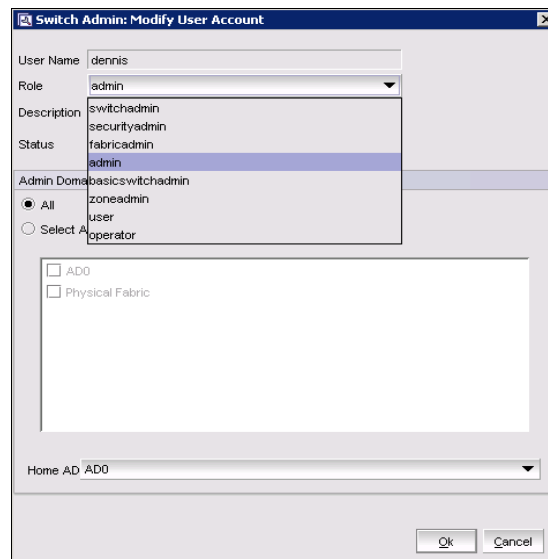


Figure 14-9 Modify user

Make the changes to the user and then Click **OK** to close the dialog box and then click **Apply** and the **Yes** button in the confirmation window to activate your changes.

Removing a user

From the Web Tools User window, select the user that you want to remove and press the **Remove** button. This will remove the user from the users list. Click **Apply** and the **Yes** button in the confirmation window to activate your changes.

Attention: You are not able to delete a default account.

Expiring a password

From the Web Tools User window, select the user that you want to expire the password and press the **Expire Password button**. This will set the user password to expired. Click **Apply** to activate your changes. The user state will change to expired as shown in Figure 14-10.

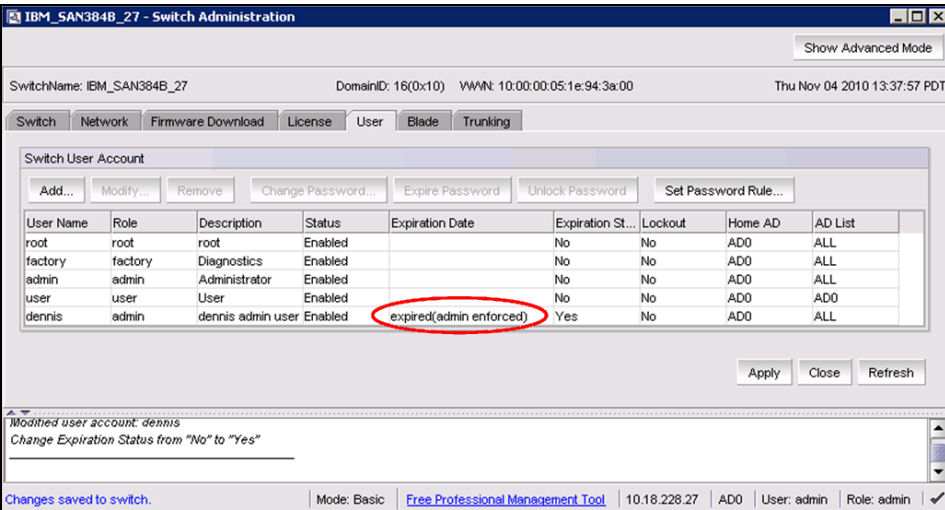


Figure 14-10 Expired user password

At the next login, this user will be requested to change the password before login is accepted. The user password can be changed from the switch administration user menu using the **Change Password** button, or when the user logs in for the first time, as shown in Example 14-2. This option can be used to force a user to change their login password.

Example 14-2 Login

```
IBM_SAN384B_27 login: dennis
Password:
-----
Your password has expired. Please change your password now.
Changing password for dennis
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.
IBM_SAN384B_27:dennis>
```

14.3 Security protocols

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

14.3.1 Security protocol support

Fabric OS supports the security protocols shown in Table 14-3.

Table 14-3 Security protocol support

Protocol	Description
HTTPS	HTTPS is a Uniform Resource Identifier scheme used to indicate a secure HTTP connection. Web Tools supports the use of hypertext transfer protocol over secure socket layer (HTTPS).
LDAPS	Lightweight Directory Access Protocol over SSL uses a certificate authority (CA). By default, LDAP traffic is transmitted unsecured. You can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL) / Transport Layer Security (TLS) technology in conjunction with LDAP.
SCP	Secure Copy (SCP) is a means of securely transferring computer files between a local and a remote host or between two remote hosts, using the Secure Shell (SSH) protocol. Configuration upload and download support the use of SCP.
SNMP	Supports SNMPv1, v2, and v3. SNMP is used in network management systems to monitor network-attached devices for conditions that warrant administrative attention.
SSH	Secure Shell (SSH) is a network protocol that allows data to be exchanged over a secure channel between two computers. Encryption provides confidentiality and integrity of data. SSH uses public-key cryptography to authenticate the remote computer and allow the remote computer to authenticate the user, if necessary.
SSL	Supports SSLv3, 128-bit encryption by default. Fabric OS uses secure socket layer (SSL) to support HTTPS. A certificate must be generated and installed on each switch to enable SSL.

Table 14-4 describes additional software or certificates that you must obtain to deploy secure protocols.

Table 14-4 Items needed to deploy secure protocols

Protocol	Host side	Switch side
SSHv2	Secure shell client	None
HTTPS	No requirement on host side except a browser that supports HTTPS	Switch IP certificate for SSL
SCP	SSH daemon, scp server	None
SNMPv1,SNMPv2,SNMPv3	None	None

The security protocols are designed with the four main use cases in Table 14-5.

Table 14-5 Usage cases

Fabric	Management interfaces	Comments
Nonsecure	Nonsecure	No special setup is needed to use Telnet or HTTP.
Nonsecure	Secure	Secure protocols can be used. An SSL switch certificate must be installed if HTTPS is used.
Secure	Secure	Switches running earlier Fabric OS versions can be part of the secure fabric, but they do not support secure management. Secure management protocols must be configured for each participating switch. Nonsecure protocols can be disabled on nonparticipating switches. If SSL is used, then certificates must be installed. For more information about installing certificates, see “Installing a switch certificate” on page 651.
Secure	Nonsecure	You must use SSH because Telnet is not allowed with some features. Nonsecure management protocols are necessary under these circumstances: <ul style="list-style-type: none">▶ The fabric contains switches running Fabric OS v3.2.0.▶ There are software tools that do not support secure protocols, for example, Fabric Manager v4.0.0.▶ The fabric contains switches running Fabric OS versions earlier than v4.4.0. Nonsecure management is enabled by default.

14.3.2 Secure file copy

In this section we discuss considerations regarding secure file copy.

Using the `configure` command

You can use the **`configure`** command to specify that secure file copy (SCP) is used for configuration uploads and downloads, as shown in Example 14-3.

Setting up SCP for configuration uploads and downloads

Follow these steps for configuration uploads 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] n
```

14.4 Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) is a standard method for monitoring and managing network devices. Using SNMP components, you can program tools to view, browse, and manipulate IBM switch variables and set up enterprise-level management processes.

Every IBM switch carries an SNMP agent and management information base (MIB). The agent accesses MIB information about a device and makes it available to a network management station. You can manipulate information of your choice by trapping MIB elements using the Fabric OS command line interface (CLI), Web Tools, or DCFM.

The SNMP access control list (ACL) provides a way for the administrator to restrict SNMP get and set operations to certain hosts and IP addresses. This is used for enhanced management security in the storage area network.

For details on Brocade MIB files, naming conventions, loading instructions, and information about using Brocade's SNMP agent, see the *Fabric OS MIB Reference*, 53-1001156-01.

You can configure SNMPv3 and SNMPv1 for the automatic transmission of SNMP information to management stations.

The configuration process involves configuring the SNMP agent and configuring SNMP traps. Use the **snmpConfig** command to configure the SNMP agent and traps for SNMPv3 or SNMPv1 configurations, and the security level. You can specify no security, authentication only, or authentication and privacy.

The SNMP trap configuration specifies the MIB trap elements to be used to send information to the SNMP management station. There are two main MIB trap choices:

- ▶ Brocade-specific MIB trap:
Associated with the Brocade-specific MIB (SW-MIB), this MIB monitors IBM/Brocade switches specifically.
- ▶ FibreAlliance MIB trap:
Associated with the FibreAlliance MIB (FA-MIB), this MIB manages SAN switches and devices from any company that complies with FibreAlliance specifications.

If you use both SW-MIB and FA-MIB, you might receive duplicate information. You can disable the FA-MIB, but not the SW-MIB.

You can also use these additional MIBs and their associated traps:

- ▶ FICON-MIB
This MIB is for FICON environments.
- ▶ SW-EXTTRAP
This MIB includes the Software Serial Number (swSsn) as a part of Brocade SW traps.

For information about Brocade MIBs, see the *Fabric OS MIB Reference*, 53-1001156-01.

For information about the specific commands used in these procedures, see the online help or the *Fabric OS Command Reference*, 53-1001186-01.

The SNMP Agent configuration interface is interactive for all parameters except mibCapability, which can be configured both interactively and with command-line options on platforms running Fabric OS v6.4.0 and later. The enhanced command-line interface supports enabling or disabling a single MIB or all MIBs, configuring a single trap only, and managing traps in excess of 32.

In Fabric OS v6.3.0 and later, the SNMPv3 configuration supports sending inform requests as an alternative to trap requests. Traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. The sender cannot determine if the trap was received. However, an SNMP manager that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the manager does not receive an inform request, it does not send a response. If the sender never receives a response, the inform request can be sent again. Thus, informs are more likely to reach their intended destination.

All values successfully changed by this command take effect immediately and are persistent across power cycles and reboots.

14.4.1 SNMP and Virtual Fabrics

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

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

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

Filtering ports

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

Switch and chassis context enforcement

All attributes are classified into two categories:

- ▶ Chassis-level attributes
- ▶ Switch-level attributes

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

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

14.4.2 Security level

Use the **snmpConfig --set seclevel** command as shown in Example 14-4 to set the security level. You can specify no security, authentication only, authentication and privacy, or off. You need to set the security for the **GET** command and the **SET** command, for example, to configure for authentication and privacy for both commands.

Example 14-4 Configuring the SNMP GET and SET commands

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

14.4.3 snmpConfig command

Use the **snmpConfig --set** command as shown in Example 14-5 to change either the SNMPv3 or SNMPv1 configuration. You can also change access control, MIB capability, and system group.

In Fabric OS v6.3.0 and later, the **--set snmpv3** command supports an interactive option to enable or disable informs by setting the parameter **SNMP Informs Enabled** to true or false. If informs are enabled, all trap destinations receive inform requests. If informs are disabled, all trap destinations receive trap requests. When informs are enabled, the engine ID must be set to correspond to the management engine IP address. Informs are by default disabled. IPv6 Informs are currently not supported.

Traps can be received at the default port 162; this can be modified from the default port during the setup process.

Example 14-5 Example SNMPv3 configuration

```
IBM_SAN384B_27:admin> snmpconfig --set snmpv3
SNMP Informs Enabled (true, t, false, f): [false] f
SNMPv3 user configuration(snmp user not configured in FOS user database
will have physical AD and admin role as the default):
User (rw): [snmpadmin1] adminuser
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 1
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol
[DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES192(5)/AES256(6)]: (1..6) [2] 2
User (rw): [snmpadmin2] shuser
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 1
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol
[DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES192(5)/AES256(6)]: (1..6) [2] 1
New Priv Passwd:
Verify Priv Passwd:
User (rw): [snmpadmin3] nosec
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 3
Priv Protocol
[DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES192(5)/AES256(6)]: (2..2) [2] 2
User (ro): [snmpuser1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 3
Priv Protocol
[DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES192(5)/AES256(6)]: (2..2) [2] 2
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 3
Priv Protocol
[DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES192(5)/AES256(6)]: (2..2) [2] 2
User (ro): [admin]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 3
Priv Protocol
[DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES192(5)/AES256(6)]: (2..2) [2] 2
SNMPv3 trap recipient configuration:
Trap Recipient's IP address : [0.0.0.0] 10.64.210.103
UserIndex: (1..6) [1] 1
Trap recipient Severity level : (0..5) [0] 4
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [0.0.0.0] 10.64.210.103
UserIndex: (1..6) [2] 2
Trap recipient Severity level : (0..5) [0] 2
Trap recipient Port : (0..65535) [162]
```

```
Trap Recipient's IP address : [0.0.0.0]
Trap Recipient's IP address : [0.0.0.0]
Trap Recipient's IP address : [0.0.0.0]
Trap Recipient's IP address : [10.127.140.95] 0.0.0.0
Committing configuration.....done.
IBM_SAN384B_27:admin>
```

Example 14-6 shows how to set the SNMP access list configuration.

Example 14-6 Example of accessControl configuration

```
IBM_SAN384B_27:admin> snmpconfig --set accessControl
SNMP access list configuration:
Access host subnet area : [0.0.0.0] 10.64.210.0
Read/Write? (true, t, false, f): [true]
Access host subnet area : [0.0.0.0] 10.64.210.0
Read/Write? (true, t, false, f): [true]
Access host subnet area : [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area : [0.0.0.0] 10.64.210.0
Read/Write? (true, t, false, f): [true]
Access host subnet area : [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area : [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration.....done.
IBM_SAN384B_27:admin>
```

Example 14-7 demonstrates how to show MIB capability.

Example 14-7 Example of mibCapability configuration

```
IBM_SAN384B_27:admin> snmpconfig --show mibcapability
FE-MIB: YES
SW-MIB: YES
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
FCIP-MIB: YES
ISCSI-MIB: YES
IF-MIB: YES
BD-MIB: YES
SW-TRAP: YES
    swFault: YES
    swSensorScn: YES
```

```

swFCPortScn: YES
swEventTrap: YES
    DesiredSeverity:None
swFabricWatchTrap: YES
    DesiredSeverity:None
swTrackChangesTrap: YES
swIPv6ChangeTrap: YES
swPmgrEventTrap: YES
swFabricReconfigTrap: NO
swFabricSegmentTrap: NO
swExtTrap: NO
FA-TRAP: YES
    connUnitStatusChange: YES
    connUnitDeletedTrap: YES
    connUnitEventTrap: YES
    connUnitSensorStatusChange: YES
    connUnitPortStatusChange: YES
FICON-TRAP: YES
    linkRNIDDeviceRegistration: YES
    linkRNIDDeviceDeRegistration: YES
    linkLIRRListenerAdded: YES
    linkLIRRListenerRemoved: YES
    linkRLIRFailureIncident: YES
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: YES
ISCSI-TRAP: YES
    iscsiTgtLoginFailure: YES
    iscsiIntrLoginFailure: YES
    iscsiInstSessionFailure: YES
IF-TRAP: YES
    linkDown: YES
    linkUp: YES
BD-TRAP: YES
    bdTrap: YES
    bdClearTrap: YES
IBM_SAN384B_27:admin>

```

Example 14-8 shows how to reset the system group configuration.

Example 14-8 Example of systemGroup configuration (default)

```

IBM_SAN384B_27:admin> snmpconfig --default systemGroup
*****

```

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

```
sysDescr = MCC_A_BB_DCX
sysLocation = 1320 Denison Street, Markham, Tile: 1L41
sysContact = Conntact Name: 416-956-6886 dlitssan@cibc.ca
authTraps = 0 (OFF)
```

```
Are you sure? (yes, y, no, n): [no] y
IBM_SAN384B_27:admin>
```

14.5 Secure Sockets Layer protocol

Secure Sockets Layer (SSL) protocol provides secure access to a fabric through Web-based management tools like Web Tools. SSL support is a standard Fabric OS feature.

Switches configured for SSL grant access to management tools through hypertext transfer protocol over SSL links (which begin with `https://`) instead of standard links (which begin with `http://`).

SSL uses public key infrastructure (PKI) encryption to protect data transferred over SSL connections. PKI is based on digital certificates obtained from an Internet Certificate Authority (ICA) that acts as the trusted key agent.

Certificates are based on the switch IP address or fully qualified domain name (FQDN), depending on the issuing CA. If you change a switch IP address or FQDN after activating an associated certificate, you might have to obtain and install a new certificate. Check with the ICA to verify this possibility, and plan these types of changes accordingly.

14.5.1 Browser and Java support

Fabric OS supports the following Web browsers for SSL connections:

- ▶ Internet Explorer v7.0 (Microsoft Windows)
- ▶ Mozilla Firefox v2.0 (Solaris and Red Hat Linux)

In countries that allow the use of 128-bit encryption, use the latest version of your browser. For example, Internet Explorer 7.0 and later supports 128-bit encryption by default. You can display the encryption support (called “cipher strength”) using the Internet Explorer **Help** → **About** menu option. If you are running an earlier version of Internet Explorer, you might be able to download an encryption patch from the Microsoft website at <http://www.microsoft.com>.

Preferably, upgrade to the Java 1.6.0 Plug-in on your management workstation. To find the Java version that is currently running, open the Java console and look at the first line of the window.

14.5.2 SSL configuration overview

You configure for SSL by obtaining, installing, and activating digital certificates for SSL support. Certificates are required on all switches that are to be accessed through SSL. Also, you must install a certificate in the Java Plug-in on the management workstation, and you might need to add a certificate to your Web browser.

Configuring for SSL involves the following main steps, which we describe in detail in the next topics:

- 1. Choose a certificate authority (CA).
- 2. Generate the following items on each switch:
 - a. A public and private key, by using the **secCertUtil genkey** command
 - b. A certificate signing request (CSR), by using the **secCertUtil genscr** command
- 3. Store the CSR on a file server by using the **secCertUtil export** command.
- 4. Obtain the certificates from the CA (Table 14-6).

You can request a certificate from a CA through a Web browser. After you request a certificate, the CA either sends certificate files by email (public) or gives access to them on a remote host (private). Typically, the CA provides the certificate files listed in Example 14-3 on page 638.

Table 14-6 SSL certificate files

Certificate file	Description
<i>name.crt</i>	The switch certificate.
<i>nameRoot.crt</i>	The root certificate. Typically, this certificate is already installed in the browser, but if not, you must install it.
<i>nameCA.crt</i>	The CA certificate. It must be installed in the browser to verify the validity of the server certificate, or server validation fails.

5. On each switch, install the certificate. After the certificate is loaded on the switch, HTTPS starts automatically.
6. If necessary, install the root certificate to the browser on the management workstation.
7. Add the root certificate to the Java Plug-in keystore on the management workstation.

14.5.3 Certificate authorities

To ease maintenance and allow secure out-of-band communication between switches, consider using one certificate authority (CA) to sign all management certificates for a fabric. If you use different CAs, management services operate correctly, but the DCFM Master Log is unable to retrieve events for the entire fabric.

Each CA (for example, Verisign or GeoTrust) has slightly different requirements; for example, some generate certificates based on IP address, while others require an FQDN, and most require a 1024-bit public/private key, while others might accept a 2048-bit key. Consider your fabric configuration, check CA websites for requirements, and gather all the information that the CA requires.

Generating a public and private key

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter the **seccertutil genkey** command to generate a public/private key pair as shown in Example 14-9.

Example 14-9 genkey

```
IBM_SAN384B_27:admin> seccertutil genkey
Generating a new key pair will automatically do the following:
1. Delete all existing CSRs.
2. Delete all existing certificates.
3. Reset the certificate filename to none.
4. Disable secure protocols.
Continue (yes, y, no, n): [no] y
Select key size [1024 or 2048]: 1024
Generating new rsa public/private key pair
Done.
IBM_SAN384B_27:admin>
```

The system reports that this process will disable secure protocols, delete any existing CSR, and delete any existing certificates. Respond to the prompts to continue and select the key size. Because CA support for the 2048-bit key size is limited, select 1024 in most cases.

Generating and storing a CSR

After generating a public/private key, perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter the command **seccertutil gencsr** and enter the requested information as shown in Example 14-10:

Example 14-10 gencsr

```
IBM_SAN384B_27:admin> seccertutil gencsr
Country Name (2 letter code, eg, US):US
State or Province Name (full name, eg, California):California
Locality Name (eg, city name):San Jose
Organization Name (eg, company name):IBM
Organizational Unit Name (eg, department name):STG
Common Name (Fully qualified Domain Name, or IP address):ibm.com
Generating CSR, file name is: 10.18.228.27.csr
Done.
IBM_SAN384B_27:admin>
```

Your CA might require specific codes for Country, State or Province, Locality, Organization, and Organizational Unit names. Make sure that your spelling is correct and matches the CA requirements. If the CA requires that the Common Name be specified as an FQDN, make sure that the fully qualified domain name is set on the domain name server. The IP address or FQDN will be the server on which the certificate will be put.

3. Enter the command **seccertutil export** to store the CSR:
4. Enter the requested information. You can use either FTP or SCP, as shown in Example 14-11.

Example 14-11 export

```
IBM_SAN384B_27:admin> seccertutil export
Select protocol [ftp or scp]: scp
Enter IP address: 10.18.228.36
Enter remote directory: ./
Enter Login Name: root
root@10.18.228.36's password:
Success: exported CSR.
IBM_SAN384B_27:admin>
```

If you are set up for secure file copy protocol, you can select it; otherwise, select **ftp**. Enter the IP address of the switch on which you generated the CSR. Enter the remote directory name of the FTP server to which the CSR is to be sent. Enter your account name and password on the server.

Obtaining certificates

Check the instructions on the CA website; then, perform this procedure for each switch:

1. Generate and store the CSR as described in “Generating and storing a CSR” on page 648.
2. Open a Web browser window on the management workstation and go to the CA website.

Follow the instructions to request a certificate. Locate the area in the request form into which you are to paste the CSR.

3. Through a Telnet window, connect to the switch and log in as admin.
4. Enter the command **seccertutil showcsr** as shown in Example 14-12.

Example 14-12 showcsr

```
IBM_SAN384B_27:admin> seccertutil showcsr
verify OK
Certificate Request:
  Data:
    Version: 0 (0x0)
    Subject: C=US, ST=California, L=San Jose, O=IBM, OU=STG,
CN=ibm.com
    Subject Public Key Info:
      Public Key Algorithm: rsaEncryption
      RSA Public Key: (1024 bit)
        Modulus (1024 bit):
          00:bc:bd:59:56:17:d8:8b:b9:cf:8c:ce:2d:48:17:
          ec:15:28:70:61:08:89:71:74:d9:45:b6:b5:8b:b9:
          f9:3c:d4:18:97:68:fc:8a:17:55:b3:e7:f5:00:d5:
          52:ff:da:cd:fc:ae:ae:a2:90:ec:1c:b4:0e:f2:26:
          43:6c:e2:e4:f5:5d:3c:de:82:ab:9d:b0:41:c2:09:
          91:73:05:40:59:22:a1:a5:bc:23:03:de:ce:d9:04:
          ed:0c:0e:cf:25:d8:b8:aa:c0:c5:19:7e:51:74:4b:
          d4:c5:5e:55:c0:c9:2a:2c:03:d0:9f:af:93:95:88:
          95:f2:e0:d4:3e:5d:35:a9:61
        Exponent: 65537 (0x10001)
    Attributes:
      serialNumber          :10:00:00:05:1e:94:3a:00
  Requested Extensions:
```

```

X509v3 Subject Alternative Name:
DNS:ibm.com, IP Address:10.18.228.27
Signature Algorithm: sha1WithRSAEncryption
83:75:30:90:09:73:f4:d9:ec:b7:bb:6f:c0:b7:74:49:77:da:
8b:a4:28:b2:41:82:8c:b3:d3:7b:22:47:57:82:48:87:69:8e:
80:05:4e:43:57:bb:45:dc:54:15:a6:16:2d:ac:f3:a4:b3:de:
05:65:10:d3:92:23:34:f9:ab:6f:d3:3e:ed:63:fb:ad:99:8c:
03:37:5b:6c:25:4b:e2:41:ed:cb:2a:8b:51:45:61:af:43:1d:
b7:6a:f8:84:af:43:17:ae:d1:08:00:59:33:fa:8c:bd:90:d3:
94:f1:a0:af:0f:45:76:af:2d:a3:2b:fb:dd:74:7b:97:0b:49:
0c:5e

```

CSR contents in base64 format

```

-----BEGIN CERTIFICATE REQUEST-----
MIIB8DCCAVKCAQAwYzELMAKGA1UEBhMCVVMxEzARBgNVBAGTCkNhbg1mb3JuaWEx
ETAPBgNVBACTCFhbiBkb3N1MQwwCgYDVQQKEwNJQK0xDDAKBgNVBAStA1NURzEQ
MA4GA1UEAxMHawJtLmNvbTCBnzANBgkqhkiG9w0BAQEFAAOBjQAwgYkCgYEAuL1Z
VhfYi7nPjM4tSBfsFShwYQiJcXTZRba1i7n5PNQYl2j8ihdVs+f1ANVS/9rN/K6u
opDsHLQ08iZDbOLk9V083oKrnBBWgmRcwVAWSKhpbwjA9702QTtDA7Pjdi4qsDF
GX5RdEvUxV5VwMkqLAPQn6+TlYiV8uDUP101qWECAwEAAaBNMCAGA1UEBTEZExcX
MDowMDowMDowNToxZTo5NDoxYTowMDApBgkqhkiG9w0BCQ4xHDAaMBGGA1UdEQQR
MA+CB21ibS5jb22HBAoS5BswDQYJKoZIhvcNAQEFBQADgYEAg3Uwka1z9Nnst7tv
wLd0SXfai6QoskGCjLPtYJHV4JIh2m0gAVOQ1e7RdxUFaYWLazzpLPeBWUQ05Ij
NPmrb9M+7WP7rZmMAzdbbCVL4kHtyyqLUUVhr0Mdt2r4hK9DF67RCABZM/qMvZDT
1PGgrw9Fdq8toyv73XR7lwtJDF4=
-----END CERTIFICATE REQUEST-----

```

IBM_SAN384B_27:admin>

The contents of the CSR are displayed.

5. Locate the section that begins with “BEGIN CERTIFICATE REQUEST” and ends with “END CERTIFICATE REQUEST”.
6. Copy and paste this section (including the BEGIN and END lines) into the area provided in the request form; then, follow the instructions to complete and send the request.

It might take several days to receive the certificates. If the certificates arrive by email, save them to an FTP server. If the CA provides access to the certificates on an FTP server, make note of the path name and make sure that you have a login name and password on the server.

Installing a switch certificate

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter the command **seccertutil import**.
3. Select a protocol, enter the IP address of the host on which the switch certificate is saved, and enter your login name and password, as shown in Example 14-13.

Example 14-13 import

```
IBM_SAN384B_27:admin> seccertutil import
Select protocol [ftp or scp]: scp
Enter IP address: 10.18.228.36
Enter remote directory: ./
Enter certificate name (must have ".crt" or ".cer" ".pem" or ".psk"
suffix)::SAN384B.crt
Enter Login Name: root
root@10.18.228.36's password:
Success: imported certificate [SAN384B.crt].
IBM_SAN384B_27:admin>
```

After the certificate is loaded on the switch, HTTPS starts automatically.

The browser

If the root certificate is not already installed on your browser, you must install it. To see whether it is already installed, check the certificate store on your browser.

The next procedures are guides for installing root certificates to Internet Explorer and Mozilla Firefox browsers. For more detailed instructions, see the documentation that came with the certificate.

Checking and installing root certificates on Internet Explorer

Follow these steps:

1. Select **Tools** → **Internet Options**.
2. Click the **Content** tab.
3. Click **Certificates**.
4. Click the **Intermediate** or **Trusted Root** tabs and scroll the list to see if the root certificate is listed. Take the appropriate following action based on whether you find the certificate:
 - If the certificate is listed, you do not need to install it. You can skip the rest of this procedure.
 - If the certificate is not listed, click **Import**.
5. Follow instructions in the Certificate Import wizard to import the certificate.

Checking and installing root certificates on Mozilla Firefox

Follow these steps:

1. Select **Tools** → **Options**.
2. Click **Advanced**.
3. Click the **Encryption** tab.
4. Click **View Certificates** → **Authorities** tab and scroll the list to see if the root certificate is listed.

For example, its name might have the form *nameRoot.crt*. Take the appropriate following action based on whether you find the certificate:

- If the certificate is listed, you do not need to install it. You can skip the rest of this procedure.
 - If the certificate is not listed, click **Import**.
5. Browse to the certificate location and select the certificate. For example, select *nameRoot.crt*.
 6. Click **Open** and follow the instructions to import the certificate.

Root certificates for the Java Plug-in

For information about Java requirements, see 14.5.1, “Browser and Java support” on page 645.

Java plug-in considerations

This procedure is a guide for installing a root certificate to the Java Plug-in on the management workstation. If the root certificate is not already installed to the plug-in, then install it. For detailed instructions, see the documentation that came with the certificate and to the Sun Microsystems website:

<http://www.sun.com>

Installing a root certificate to the Java plug-in

Follow these steps to install the certificate:

1. Copy the root certificate file from its location on the FTP server to the Java Plug-in bin. For example, the bin location might be:
`C:\program files\java\j2re1.6.0\bin`
2. Open a Command Prompt window and change the directory to the Java Plug-in bin.
3. Enter the **keytool** command and respond to the prompts:
`C:\Program Files\Java\j2re1.6.0\bin> keytool -import -alias RootCert
-file RootCert.crt -keystore ..\lib\security\RootCerts
Enter keystore password: changeit`

```
Owner: CN=Brocade, OU=Software, O=Brocade Communications, L=San
Jose,
ST=California, C=US
Issuer: CN=Brocade, OU=Software, O=Brocade Communications, L=San
Jose,
ST=California, C=US
Serial number: 0
Valid from: Thu Jan 15 16:27:03 PST 2007 until: Sat Feb 14 16:27:03
PST 2007
Certificate fingerprints:
MD5: 71:E9:27:44:01:30:48:CC:09:4D:11:80:9D:DE:A5:E3
SHA1: 06:46:C5:A5:C8:6C:93:9C:FE:6A:C0:EC:66:E9:51:C2:DB:E6:4F:A1
Trust this certificate? [no]: yes
Certificate was added to keystore
```

In the example, **changeit** is the default password and **RootCert** is an example root certificate name.

Summary of certificate commands

Table 14-7 summarizes the commands for displaying and deleting certificates. For details about the commands, see the *Fabric OS Command Reference*, 53-1001186-01.

Table 14-7 Commands for displaying and deleting SSL certificates

Command	Description
secCertUtil show	Displays the state of the SSL key and a list of installed certificates.
secCertUtil show <i>filename</i>	Displays the contents of a specific certificate.
secCertUtil showcsr	Displays the contents of a CSR.
secCertUtil delete <i>filename</i>	Deletes a specified certificate.
secCertUtil delcsr	Deletes a CSR.

14.6 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) that runs on the switch and supports a wide variety of encryption algorithms, such as Blowfish-Cipher block chaining (CBC) and Advanced Encryption Standard (AES).

Security: To maintain a secure network, avoid using Telnet or any other unprotected application when you are working on the switch.

The File Transfer Protocol (FTP) is also not secure. When you use FTP to copy files to or from the switch, the contents are in clear text. This includes the remote FTP server's login and password. This limitation affects the following commands: **saveCore**, **configUpload**, **configDownload**, and **firmwareDownload**.

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

Fabric OS v6.2.0 supports SSH protocol v2.0 (ssh2). For more information about SSH, see the SSH IETF website:

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

14.6.1 SSH public key authentication

OpenSSH public key authentication provides password-less logins, known as SSH authentication, which use public and private key pairs for incoming and outgoing authentication. This feature allows only one *allowed-user* to be configured to utilize OpenSSH public key authentication.

Authentication protocols

Using OpenSSH RSA and DSA, the authentication protocols are based on a pair of specially generated cryptographic keys, called the private key and the public key. The advantage of using these key-based authentication systems is that in many cases, it is possible to establish secure connections without having to manually type in a password. RSA and DSA asynchronous algorithms are FIPS-compliant.

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 public keys related to the old *allowed-user* are lost.

14.6.2 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.

Setup steps

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 (see your host's documentation as necessary), and typing the following command (see Example 14-14):

```
sshutil -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 14-14 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 following command to import the key (Example 14-15):

sshUtil importpubkey

Respond to the prompts as follows:

IP Address	Enter the IP address of the switch. IPv6 is supported by sshUtil .
Remote directory	Enter the path to the remote directory where the public key is stored.
Public key name	Enter the name of the public key.
Login name	Enter the name of the user granted access to the host.
Password	Enter the password for the host.

Example 14-15 Example of adding the public key to the switch

```
switch:alloweduser> sshutil importpubkey
Enter IP address:10.64.210.130
Enter remote directory:~ausser/.ssh
Enter public key name(must have .pub suffix):id_dsa.pub
Enter login name:ausser
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 following command (see Example 14-16):

sshUtil genkey

Enter a passphrase for additional security.

Example 14-16 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 following command to export the key (see Example 14-17):

sshUtil exportpubkey

Respond to the prompts as follows:

IP Address	Enter the IP address of the remote host. IPv6 is supported by sshUtil .
Remote directory	Enter the path to the remote directory where the public key will be stored.
Login name	Enter the name of the user granted access to the remote host.
Password	Enter the password for the remote host.

Example 14-17 Example of exporting a public key from the switch

```
switch:alloweduser> sshutil exportpubkey  
Enter IP address:10.64.210.103  
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 might have 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

Follow these steps to delete the keys when necessary:

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.

For more information about IP Filter policies, see 14.9.26, "IP Filter policy" on page 686.

14.7 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.

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

14.7.1 Blocking Telnet

Follow these steps for blocking Telnet:

1. Connect to the switch and log in as admin.
Connect through some means other than Telnet: for example, through SSH.
2. Create a policy by typing the following command (see Example 14-18):

```
ipfilter --create polycname -type < ipv4 | ipv6 >
```

Where *polycname* is the name of the new policy and **-type** specifies an IPv4 or IPv6 address.

Example 14-18 Example of creating a policy

```
ipfilter --create block_telnet_v4 -type ipv4
```

3. Add a rule to the policy, by typing the following command (see Example 14-19):

```
ipfilter --addrule <polycname> -rule rule_number -sip source_IP -dp dest_port -proto protocol -act <deny>
```

Where the **-sip** option can be given as **any**, **-dp** is the port number for Telnet (23), and **-proto** is **TCP**.

Example 14-19 Example of adding a rule

```
ipfilter --addrule block_telnet_v4 -rule 1 -sip any -dp 23 -proto tcp -act deny
```

4. Save the new ipfilter policy by typing the following command (see Example 14-20 on page 659):

```
ipfilter --save polycname
```

Where *polycname* is the name of the policy and is optional.

Example 14-20 Example of saving a policy

```
ipfilter --save block_telnet_v4
```

5. Activate the new ipfilter policy by typing the following command:
(Example 14-21):

```
ipfilter --activate polycyname
```

Where *polycyname* is the name of the policy that you created in step 2.

Example 14-21 Example of activating a policy

```
ipfilter --activate block_telnet_v4
```

14.7.2 Unblocking Telnet

Follow these steps:

1. Connect to the switch through a means other than Telnet (for example, SSH) and log in as admin.
2. Type the following command:

```
ipfilter --delete telnet_polycyname
```

Where *telnet_polycyname* is the name of the Telnet policy.

3. To permanently delete the policy, type the following command:

```
ipfilter --save
```

14.7.3 Listener applications

Brocade switches block Linux subsystem listener applications that are not used to implement supported features and capabilities. Table 14-8 lists the listener applications that Brocade switches either block or do not start.

Table 14-8 Blocked listener applications

Listener Application	IBM SAN256B, IBM SAN768B and IBM SAN384B	IBM SAN04BR, SAN18BR, SAN24B, SAN40B and SAN80B switches and FA4-18, FC10-6, FC4-48, FC4-16IP, FC8-16/32/48, and FR4-18i blades
chargen	Disabled	Disabled
echo	Disabled	Disabled

Listener Application	IBM SAN256B, IBM SAN768B and IBM SAN384B	IBM SAN04BR, SAN18BR, SAN24B, SAN40B and SAN80B switches and FA4-18, FC10-6, FC4-48, FC4-16IP, FC8-16/32/48, and FR4-18i blades
daytime	Disabled	Disabled
discard	Disabled	Disabled
ftp	Disabled	Disabled
rexec	Block with packet filter	Disabled
rsh	Block with packet filter	Disabled
rlogin	Block with packet filter	Disabled
time	Block with packet filter	Disabled
rstats	Disabled	Disabled
rusers	Disabled	Disabled

14.8 Ports and applications used by switches

If you are using the FC-FC Routing Service, be aware that the `secModeEnable` command is not supported in Fabric OS v6.0.0 and later.

14.8.1 Access defaults

Table 14-9 lists the defaults for accessing hosts, devices, switches, and zones.

Table 14-9 Access defaults

Item	Access default
Hosts	Any host can access the fabric by SNMP.
	Any host can Telnet to any switch in the fabric
	Any host can establish an HTTP connection to any switch in the fabric.
	Any host can establish an API connection to any switch in the fabric.
Devices	All devices can access the management server.

Item	Access default
	Any device can connect to any FC port in the fabric.
Switches	Any switch can join the fabric.
	All switches in the fabric can be accessed through a serial port.
Zoning	No zoning is enabled.

14.8.2 Port configuration

Table 14-10 provides information about ports that the switch uses. When configuring the switch for various policies, take into consideration firewalls and other devices that might sit between switches in the fabric and your network or between the managers and the switch.

Table 14-10 Switch port usage

Port	Type	Common use	Comment
22	TCP	SSH	
23	TCP	Telnet	Use the ipfilter command to block the port.
80	TCP	HTTP	Use the ipfilter command to block the port.
111	TCP	sunrpc	This port is used by Platform API. Use the ipfilter command to block the port.
123	TCP	NTP	
161	UDP	SNMP	Disable the SNMP service on the remote host if you do not use it, or filter incoming UDP packets going to this port.
443	TCP	HTTPS	Use the ipfilter command to block the port.
	TCP	exec	
	TCP	login	
	TCP	shell	
	TCP		
			This port is used by the Platform API. Disable this port using the configure command.

14.9 Security policies

In the following topics we discuss security policies and their management.

14.9.1 ACL policies overview

Each supported Access Control List (ACL) policy listed here is identified by a specific name, and only one policy of each type can exist, except for DCC policies. Policy names are case-sensitive and must be entered in all uppercase. Fabric OS provides the following policies:

- ▶ Fabric configuration server (FCS) policy: Used to restrict which switches can change the configuration of the fabric.
- ▶ Device connection control (DCC) policies: Used to restrict which Fibre Channel device ports can connect to which Fibre Channel switch ports.
- ▶ Switch connection control (SCC) policy: Used to restrict which switches can join with a switch.
- ▶ IP filter policy (IPFilter) policy: Used to filter traffic based on IP addresses.

14.9.2 ACL policy management

All policy modifications are temporarily stored in volatile memory until those changes are saved or activated. You can create multiple sessions to the switch from one or more hosts. Make changes from one switch only to prevent multiple transactions from occurring. Each logical switch will have its own access control list.

ACL policy considerations

The FCS, SCC, and DCC policies in Secure Fabric OS are not interchangeable with Fabric OS FCS, SCC, and DCC policies. Uploading and saving a copy of the Fabric OS configuration after creating policies is highly desirable.

You can view the active and defined policy sets at any time. Additionally, in a defined policy set, policies created in the same login session also display, but these policies are automatically deleted if you log out without saving them.

Displaying ACL policies

Follow these steps to display the policies:

1. Connect to the switch and log in using an account assigned to the admin role.
2. To display all security policies, type the **secPolicyShow** command as shown in Figure 14-22 on page 663.

```
BDPOC01L01:admin> secPolicyShow
```

ACTIVE POLICY SET

DEFINED POLICY SET

```
BDPOC01L01:admin>
```

14.9.3 FCS policies

Fabric Configuration Server (FCS) policy in base Fabric OS can be performed on a local switch basis and can be performed on any switch in the fabric with Fabric OS v6.0.0 or later. Any switch with a pre-v5.3.0 version of Fabric OS cannot be included in the FCS list.

FCS policy creation

The FCS policy is not present by default, but must be created. When the FCS policy is created, the WWN of the local switch is automatically included in the FCS list. Additional switches can be included in the FCS list. The first switch in the list becomes the Primary FCS switch.

Only the Primary FCS switch is allowed to modify and distribute the database within the fabric. Automatic distribution is supported and you can either configure the switches in your fabric to accept the FCS policy or manually distribute the FCS policy. Changes made to the FCS policy are saved to permanent memory only after the changes have been saved or activated; they can be aborted later if you have set your fabric to distribute the changes manually. See Table 14-11.

Table 14-11 FCS Policy states

Policy state	Characteristics
No active policy	Any switch can perform fabric-wide configuration changes.
Active policy with one entry	A Primary FCS switch is designated (local switch), but there are no backup FCS switches. If the Primary FCS switch becomes unavailable for any reason, the fabric is left without an FCS switch.

Policy state	Characteristics
Active policy with multiple entries	A Primary FCS switch and one or more backup FCS switches are designated. If the Primary FCS switch becomes unavailable, the next switch in the list becomes the Primary FCS switch.

The FCS policy is designed to accommodate mixed fabric environments that contain switches with pre-v5.3.0 and later versions of Fabric OS. By setting the configuration parameters to accept fabric distribution, Fabric OS v6.0.0 and later switches can enforce FCS policy and perform database distribution among v5.3.0 and v6.0.0 and later switches while still allowing pre-v5.3.0 switches to join the fabric. The following items describe distribution behavior for pre-Fabric OS v5.3.0:

- ▶ Distribution to pre-v5.3.0 switches with specific domain IDs:
When specific domain IDs are given for the distribution, all domains must be on a switch with Fabric OS v5.3.0 or later. If one of the domains is pre-v5.3.0 the distribution operation will fail.
- ▶ Distribution to pre-v5.3.0 switches using the wild card (*) character:
When the wild card character is specified, distribution succeeds even if the fabric contains pre-v5.3.0 switches. However, the FCS database will be sent only to switches with a Fabric OS of v5.2.0 or later in the fabric and not to pre-v5.2.0 switches. Fabric OS v5.2.0 switches receive the distribution and will ignore the FCS database.

FCS policy restrictions

The backup FCS switches normally cannot modify the policy. However, if the Primary FCS switch in the policy list is not reachable, then a backup FCS switch is allowed to modify the policy.

After an FCS policy is configured and distributed across the fabric, only the Primary FCS switch can perform certain operations. Operations that affect fabric-wide configuration are allowed only from the Primary FCS switch. Backup and non-FCS switches cannot perform security, zoning and AD operations that affect the fabric configuration. The following error message is returned if a backup or non-FCS switch tries to perform these operations:

Can only execute this command on the Primary FCS switch.

Operations that do not affect the fabric configuration, such as show or local switch commands, are allowed on backup and non-FCS switches.

FCS enforcement applies only for user-initiated fabric-wide operations. Internal fabric data propagation because of a fabric merge is not blocked. Consequently, a new switch that joins the FCS-enabled fabric can still propagate the AD and zone database.

14.9.4 Overview of FCS policy management

Whether your intention is to create new FCS policies or manage your current FCS policies, you must follow certain steps to ensure that the domains throughout your fabric have the same policy.

The local-switch WWN cannot be deleted from the FCS policy.

Follow these steps:

1. Set the pre-v5.3.0 switches in the fabric to accept the FCS policy using the **fddCfg --localaccept** or **fddCfg --localreject** command.
2. Create the FCS policy using the **secPolicyCreate** command.
3. Activate the policy using the **secPolicyActivate** command.

If the command is not entered, the changes are lost when the session is logged out.

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

14.9.5 Creating an FCS policy

To create an FCS policy, follow these steps:

1. Connect to the switch and log in using an account assigned to the admin role.
2. Type **secPolicyCreate "FCS_POLICY", "member;...;member"**.

Where member indicates a switch that is eligible to become a primary or backup FCS switch. Specify switches by WWN, domain ID, or switch name. Enter the wild card (*) character to indicate all the switches in the fabric.

Example 14-23 shows how to create an FCS policy that allows a switch with domain ID 3 to become a primary FCS and domain ID 5 to become a backup FCS.

Example 14-23 Domain ID 3 becomes primary and 5 become backup FCS

```
BDPOC01L01:admin> secpolicycreate "FCS_POLICY", "3; 5"  
FCS_POLICY has been created.
```

3. To save or activate the new policy, enter either the **secPolicySave** or the **secPolicyActivate** command. After the policy has been activated, you can distribute the policy.

Attention: FCS policy must be consistent across the fabric. If the policy is inconsistent in the fabric, then you will not be able to perform any fabric-wide configurations from the primary FCS.

14.9.6 Modifying the order of FCS switches

To modify the order of FCS switches, follow these steps:

1. Log in to the Primary FCS switch using an account assigned to the admin role.
2. Type **secPolicyShow "Defined", "FCS_POLICY"**.
This displays the WWNs of the current Primary FCS switch and backup FCS switches.
3. Type **secPolicyFCSMove**; then provide the current position of the switch in the list and the desired position at the prompts.

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

To move a backup FCS switch from position 2 to position 3 in the FCS list, using interactive mode do the following, see Example 14-24.

Example 14-24 Moving a backup switch

```
BDPOC01L01:admin> secpolicyfcsmove  
PosPrimary WWN DId swName.  
=====
```

1	Yes	10:00:00:60:69:10:02:181	switch5.
2	No	10:00:00:60:69:00:00:5a2	switch60.
3	No	10:00:00:60:69:00:00:133	switch73.

```
Please enter position you'd like to move from : (1..3) [1] 2  
Please enter position you'd like to move to : (1..3) [1] 3
```

```
DEFINED POLICY SET  
FCS_POLICY
```

```
PosPrimaryWWN DId swName
```

```
1Yes 10:00:00:60:69:10:02:181 switch5.  
2No  10:00:00:60:69:00:00:133 switch73.  
3No  10:00:00:60:69:00:00:5a2 switch60.
```

4. Type the **secPolicyActivate** command to activate and save the new order.

14.9.7 FCS policy distribution

The FCS policy can be automatically distributed using the **fddCfg --fabwideset** command or it can be manually distributed to the switches using the **distribute -p** command. Each switch that receives the FCS policy must be configured to receive the policy.

Distributions: The default value for the distribution configuration parameter is *accept*, which means that the switch accepts all database distributions and is able to initiate a distribute operation for all databases.

To verify that a switch is configured to receive the policy, do the following steps:

1. Log in to the switch using an account assigned to the admin role.
2. Type **fddcfg --showall** to list the local switch configuration as shown in Example 14-25.

Example 14-25 List local switch configuration

```
BDPOC01L01:admin> fddcfg --showall  
Local Switch Configuration for all Databases:-  
    DATABASE - Accept/Reject  
-----  
          SCC -      reject  
          DCC -      accept  
          PWD -      accept  
          FCS -      accept  
          AUTH -      accept  
          IPFILTER -      accept  
  
Fabric Wide Consistency Policy:- ""
```

3. In Example 14-25, you can see that SCC is rejected. To accept SCC, enter the commands shown in Example 14-26.

Example 14-26 Accept SCC

```
BDPOC01L01:admin> fddcfg --localaccept "SCC"
Local Switch Configured to accept policies.
BDPOC01L01:admin>
BDPOC01L01:admin>
BDPOC01L01:admin>
BDPOC01L01:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
      DATABASE - Accept/Reject
-----
      SCC - accept
      DCC - accept
      PWD - accept
      FCS - accept
      AUTH - accept
      IPFILTER - accept
```

Fabric Wide Consistency Policy:- ""

```
BDPOC01L01:admin>
```

Switches in the fabric are designated as either a Primary FCS, backup FCS, or non-FCS switch. Only the Primary FCS switch is allowed to distribute the database. The FCS policy might need to be manually distributed across the fabric using the command **distribute -p** if there is no support for automatic distribution in a mixed environment with v5.3.0 and pre-v5.3.0 switches. Because this policy is distributed manually, the command **fddCfg --fabwideset** is used to distribute a fabric-wide consistency policy for FCS policy in an environment consisting of only Fabric OS v6.0.0 and later switches.

FCS enforcement for the **distribute** command is handled differently for FCS and other databases in an FCS fabric:

- ▶ For an FCS database, the enforcement allows any switch to initiate the distribution. This is to support FCS policy creation specifying a remote switch as Primary.
- ▶ For other database distributions, only the Primary FCS switch can initiate the distribution.

There is an FCS enforcement at the receiving switch, so the switch will verify whether the distribution is coming from the Primary FCS switch before accepting it. Distribution is accepted only if it is coming from a Primary FCS switch. Distribution of FCS policy can still be accepted from a backup FCS switch if the Primary is not reachable or from a non-FCS switch if the Primary FCS and none of the backup FCS switches are reachable.

14.9.8 DCC policies

Multiple DCC policies can be used to restrict which device ports can connect to which switch ports. The devices can be initiators, targets, or intermediate devices such as SCSI routers and loop hubs. By default, all device ports are allowed to connect to all switch ports; no DCC policies exist until they are created.

Each device port can be bound to one or more switch ports; the same device ports and switch ports can be listed in multiple DCC policies. After a switch port is specified in a DCC policy, it permits connections only from designated device ports. Device ports that are not specified in any DCC policies are allowed to connect only to switch ports that are not specified in any DCC policies.

When a DCC violation occurs, the related port is automatically disabled and must be re-enabled using the **portEnable** command. See Table 14-12.

Table 14-12 DCC policy state

Policy state	Characteristics
No policy	Any device can connect to any switch port in the fabric.
Policy with no entries	Any device can connect to any switch port in the fabric. An empty policy is the same as no policy.
Policy with entries	If a device WWN is specified in a DCC policy, that device is only allowed access to the switch if connected by a switch port listed in the same policy. If a switch port is specified in a DCC policy, it only permits connections from devices that are listed in the policy. Devices with WWNs that are not specified in a DCC policy are allowed to connect to the switch at any switch ports that are not specified in a DCC policy. Switch ports and device WWNs can exist in multiple DCC policies. Proxy devices are always granted full access and can connect to any switch port in the fabric.

Virtual Fabric considerations: The DCC policies that have entries for the ports that are being moved from one logical switch to another will be considered stale and will not be enforced. You can choose to keep stale policies in the current logical switch or delete the stale policies after the port movements. Use the **secPolicyDelete** command to delete stale DCC policies.

14.9.9 DCC policy restrictions

The following restrictions apply when using DCC policies:

- ▶ Some older private-loop HBAs do not respond to port login from the switch and are not enforced by the DCC policy. This does not create a security problem because these HBAs cannot contact any device outside of their immediate loop.
- ▶ DCC policies cannot manage or restrict iSCSI connections, that is, an FC Initiator connection from an iSCSI gateway.
- ▶ You cannot manage proxy devices with DCC policies. Proxy devices are always granted full access, even if the DCC policy has an entry that restricts or limits access of a proxy device.

14.9.10 Creating a DCC policy

DCC policies must follow the naming convention “**DCC_POLICY_nnn**,” where **nnn** represents a unique string. The maximum length is 30 characters, including the prefix **DCC_POLICY_**.

Device ports must be specified by port WWN. Switch ports can be identified by the switch WWN, domain ID, or switch name followed by the port or area number. To specify an allowed connection, enter the device port WWN, a semicolon, and the switch port identification.

The following methods of specifying an allowed connection are possible:

- ▶ `deviceportWWN;switchWWN` (port or area number)
- ▶ `deviceportWWN;domainID` (port or area number)
- ▶ `deviceportWWN;switchname` (port or area number)

Follow these steps:

1. Connect to the switch and log in using an account assigned to the admin role.
2. Type the **secPolicyCreate** “**DCC_POLICY_nnn**”, “**member;...;member**” command.

DCC_POLICY_nnn is the name of the DCC policy; **nnn** is a string consisting of up to 19 alphanumeric or underscore characters to differentiate it from any other DCC policies.

The member contains device or switch port information:
deviceportWWN;switch(port) as shown in Table 14-13.

Table 14-13 Switch port information

deviceportWWN	The WWN of the device port
switch	<p>The switch WWN, domain ID, or switch name. The port can be specified by port or area number. Designating ports automatically includes the devices currently attached to those ports. The ports can be specified using any of the following syntax methods:</p> <p>(*) Selects all ports on the switch.</p> <p>(1-6) Selects ports 1 through 6.</p> <p>[*] Selects all ports and all devices attached to those ports.</p> <p>[3, 9] Selects ports 3 and 9 and all devices attached to those ports.</p> <p>[1-3, 9] Selects ports 1, 2, 3, 9, and all devices attached to those ports.</p> <p>“**” This method can be used to indicate DCC lockdown. It creates a unique policy for each port in the fabric, locking it down to the device connected or creating an empty policy to disallow any device to be connected to it. This method can be done only when there are no other DCC policies defined on the switch.</p>

3. To save or activate the new policy, enter the appropriate command:
- To save the policy, enter the **secPolicySave** command.
 - To save and activate the policy, enter the **secPolicyActivate** command.
 - If neither of these commands is entered, the changes are lost when the session is logged out.

14.9.11 Creating a device policy

You can create a device policy to allow a DS8000® storage port and an AIX® host port to attach to Domain 97 ports 5 and 6, as shown in Example 14-27.

Where:

- The port WWN of the DS8000 port is 50:05:07:63:04:03:03:16
- The port WWN of the DIX host port is 10:00:00:00:c9:2a:f3:d5

Example 14-27 Creating a device policy

```
BDPOC01L01:admin> secpolicycreate "DCC_POLICY_Storage01", \  
"50:05:07:63:04:03:03:16;10:00:00:00:c9:2a:f3:d5;97(5,6)"  
DCC_POLICY_Storage01 has been created.  
  
BDPOC01L01:admin> secPolicyShow
```

```

ACTIVE POLICY SET

DEFINED POLICY SET
DCC_POLICY_Storage01
  Type      WWN                      DId swName
  -----
Switch  10:00:00:05:1e:36:05:42  97 BDPOC01L01.
=Index=> 5,6.
Device  50:05:07:63:04:03:03:16
Device  10:00:00:00:c9:2a:f3:d5

BDPOC01L01:admin>

```

14.9.12 Deleting a device policy

To delete the policy you just created, follow Example 14-28.

Example 14-28 Deleting the policy

```

BDPOC01L01:admin> secPolicyDelete DCC_POLICY_Storage01
About to delete DCC_POLICY_Storage01
ARE YOU SURE (yes, y, no, n): [no] y
BDPOC01L01:admin>

BDPOC01L01:admin> secPolicyShow

```

```

ACTIVE POLICY SET

DEFINED POLICY SET

BDPOC01L01:admin>

```

14.9.13 Activating policy changes

To activate the policy changes on all switches in the fabric, issue the commands shown in Example 14-29.

Example 14-29 Activate policy

```
BDPOC01L01:admin> secpolicyactivate
  About to overwrite the current Active data.
  ARE YOU SURE (yes, y, no, n): [no] y
  secpolicyactivate command was completed successfully.
BDPOC01L01:admin>
```

14.9.14 SCC policies

The switch connection control (SCC) policy is used to restrict which switches can join the fabric. Switches are checked against the policy each time an E_Port-to-E_Port connection is made. The policy is named SCC_POLICY and accepts members listed as WWNs, domain IDs, or switch names. Only one SCC policy can be created.

By default, any switch is allowed to join the fabric; the SCC policy does not exist until it is created. When connecting a Fibre Channel router to a fabric or switch that has an active SCC policy, the front domain of the Fibre Channel router must be included in the SCC policy (see Table 14-14).

Table 14-14 SCC policy states

Policy state	SCC policy enforcement
No active policy	All switches can connect to the switch with the specified policy.
Active policy that has no members	All neighboring switches are segmented.
Active policy that has members	The neighboring switches not specified in the SCC policy are segmented.

Virtual Fabric considerations: In a logical fabric environment, the SCC policy enforcement is not done on the logical ISL. For a logical ISL-based switch, the SCC policy enforcement is considered as the reference and the logical ISL is formed if the SCC enforcement passes on the extended ISL. The following functionality changes:

- ▶ A logical switch supports an SCC policy. You can configure and distribute an SCC policy on a logical switch.
- ▶ SCC enforcement is performed on an ISL based on the SCC policy present on the logical switch.

14.9.15 Creating an SCC policy

Follow these steps to create the policy:

1. Connect to the switch and log in using an account assigned to the admin role.
2. Type **secPolicyCreate** “**SCC_POLICY**”, “**member;...;member**”.

Where *member* indicates a switch that is permitted to join the fabric. Specify switches by WWN, domain ID, or switch name. Enter an asterisk (*) to indicate all the switches in the fabric.

Example 14-30 shows how to create an SCC policy that allows switches that have domain IDs 97 and 4 to join the fabric:

Example 14-30 Create policy

```
BDPOC01L01:admin> secpolicycreate "SCC_POLICY", "97;4"  
SCC_POLICY has been created.  
BDPOC01L01:admin>
```

3. Save or activate the new policy by entering either the **secPolicySave** or the **secPolicyActivate** command. If neither of these commands is entered, the changes are lost when the session is logged out.

14.9.16 Authentication policy for fabric elements

By default, Fabric OS v6.1.0 and later use DH-CHAP or FCAP protocols for authentication. These protocols use shared secrets and digital certificates, based on switch WWN and public key infrastructure (PKI) technology, to authenticate switches. Authentication automatically defaults to FCAP if both switches are configured to accept FCAP protocol in authentication. To use FCAP on both switches, PKI certificates have to be installed.

You can configure a switch with Fabric OS v5.3.0 or later to use DH-CHAP for device authentication. Use the **authUtil** command to configure the authentication parameters used by the switch. When you configure DH-CHAP authentication, you also must define a pair of shared secrets known to both switches as a secret key pair. Figure 14-11 illustrates how the secrets are configured. A secret key pair consists of a local secret and a peer secret. The local secret uniquely identifies the local switch. The peer secret uniquely identifies the entity to which the local switch authenticates. Every switch can share a secret key pair with any other switch or host in a fabric.

To use DH-CHAP authentication, a secret key pair has to be configured on both switches. You can use the command **authUtil --set -a <fcap|dhchap>** to set the authentication protocol, which can then be verified using the command **authUtil --show CLI**.



Figure 14-11 DH-CHAP authentication

If you use DH-CHAP authentication, then a secret key pair must be installed only in connected fabric elements. However, as connections are changed, new secret key pairs must be installed between newly connected elements. Alternatively, a secret key pair for all possible connections can be initially installed, enabling links to be arbitrarily changed while still maintaining a valid secret key pair for any new connection.

The switch authentication (AUTH) policy initiates DH-CHAP/FCAP authentication on all E_Ports. This policy is persistent across reboots, which means authentication will be initiated automatically on ports or switches brought online if the policy is set to activate authentication. The AUTH policy is distributed using the **distribute** command. The automatic distribution of the AUTH policy is not supported.

The default configuration directs the switch to attempt FCAP authentication first, DH-CHAP second. The switch can be configured to negotiate FCAP, DH-CHAP, or both.

The DH group is used in the DH-CHAP protocol only. The FCAP protocol exchanges the DH group information, but does not use it.

The AUTH policy is designed to accommodate mixed fabric environments that contain Fabric OS v6.0.0 and later along with pre-v6.0.0 switches. The policy states that PASSIVE and OFF allow connection from Fabric OS v6.0.0 and later switches to pre-v6.0.0 switches. These policy states do not allow switches to send the authentication negotiation and therefore continue with the rest of port initialization.

FOS v6.4.0 adds support for FCAP authentication using third-party self signed certificates. Starting with FOS v6.4.0 both Brocade issued certificates and/or third-party self signed certificates can be used for FCAP authentication. Prior to FOS v6.4.0 only Brocade issued certificates were supported.

Virtual Fabric considerations: If a Virtual Fabric is enabled, all AUTH module parameters such as shared secrets, as well as shared switch and device policies, are logical switch-wide. This means that you must configure shared secrets and policies separately on each logical switch, and the shared secrets and policies must be set on each switch prior to authentication. On logical switch creation, authentication takes default values for policies and other parameters.

14.9.17 E_Port authentication

The authentication (AUTH) policy allows you to configure DH-CHAP authentication on the switch. By default, the policy is set to PASSIVE and you can change the policy using the `authUtil` command. All changes to the AUTH policy take effect during the next authentication request. This includes starting authentication on all E_Ports on the local switch if the policy is changed to ON or ACTIVE, and clearing the authentication if the policy is changed to OFF. The authentication configurations will be effective only on subsequent E_ and F_Port initialization.

Virtual Fabric considerations:

- ▶ The switch authentication policy applies to all E_Ports in a logical switch. This includes ISLs and extended ISLs. Authentication of extended ISLs between two base switches is considered peer-chassis authentication. Authentication between two physical entities is required, so the extended ISL which connects the two chassis needs to be authenticated. The corresponding extended ISL for a logical ISL authenticates the peer-chassis, therefore the logical ISL authentication is not required.
- ▶ Because the logical ISLs do not carry actual traffic, they do not need to be authenticated. Authentication on re-individualization is also blocked on logical ISLs. The following error message is printed on the console when you execute the **authUtil --authinit** command on logical-ISLs:

Failed to initiate authentication. Authentication is not supported on logical ports <port#>.

A secret key pair has to be installed prior to changing the policy. The policy can be configured as shown in Example 14-31.

Example 14-31 Configuring the policy

```
switch:admin> authutil --policy -sw <ON|ACTIVE|PASSIVE|OFF>
```

The command arguments are explained in Table 14-15.

Status of data if failover occurs:

- ▶ If data input has not been completed and a failover occurs, the command is terminated without completion, and your entire input is lost.
- ▶ If data input has completed, the Enter key was pressed, and a failover occurs, data might or might not be replicated to the other CP depending on the timing of the failover. Log in to the other CP after the failover is complete and verify that the data was saved. If the data was not saved, run the command again.

Table 14-15 *authutil* command reference

authutil --policy -sw Parameter	Explanation:
ON	<p>Setting the AUTH policy to ON means that strict authentication is enforced on all E_Ports. If the connecting switch does not support authentication or the policy is switched to the OFF state, the ISL is disabled.</p> <p>During switch initialization, authentication begins automatically on all E_Ports. To enforce this policy fabric-wide, the fabric needs to have Fabric OS v5.3.0 and later switches only. The switch disables the port if it is connected to a switch that does not support authentication. Regardless of the policy, the E_Port is disabled if the DH-CHAP or FCAP protocol fails to authenticate the attached E_Port.</p>
OFF	<p>This setting turns off the policy. The switch does not support authentication and rejects any authentication negotiation request from another switch. A switch with the policy turned OFF cannot be connected to a switch with the policy turned ON. The ON state is strict and disables the port if any switch rejects the authentication. DH-CHAP shared secrets must be configured before changing the policy from the OFF to the ON state.</p> <p>The behavior of the policy between two adjacent switches is defined as follows. If the policy is ON or active, the switch sends an authentication negotiation request to the connecting switch. If the connecting switch does not support authentication or the policy is OFF, the request is rejected. After the authentication negotiation succeeds, the DH-CHAP authentication is initiated. If DH-CHAP authentication fails, the port is disabled and this is applicable in all modes of the policy.</p>

authutil --policy -sw Parameter	Explanation:
ACTIVE	<p>In this state the switch is more tolerant and can connect to a switch with any type of policy. During switch initialization, authentication begins on all E_Ports, but the port is not disabled if the connecting switch does not support authentication or the AUTH policy is turned to the OFF state.</p> <p>The authentication begins automatically during the E_Port initialization. A switch with this policy can safely connect to pre-v6.0.0 switches, because it continues E_Port initialization if the connecting switch does not support authentication. The switches with firmware pre-v3.2.0 do not support FCAP or DH-CHAP authentication, so an E_Port initializes without authentication. The switches with firmware version v3.2.0 and later respond to authentication negotiation and participate in FCAP and DH-CHAP handshaking. Regardless of the policy, the E_Port is disabled if the DH-CHAP or FCAP protocol fails to authenticate the attached E_Port.</p>
PASSIVE	<p>In the PASSIVE state, the switch does not initiate authentication, but participates in authentication if the connecting switch initiates authentication. The switch does not start authentication on E_Ports, but accepts the incoming authentication requests, and does not disable if the connecting switch does not support authentication or the policy is turned to the OFF state. This is the safest policy for switches connecting to pre-v5.3.0 switches.</p> <p>That means v5.3.0 and later switches can have authentication enabled and this will not impact the pre-v5.3.0 switches. By default, the pre-v5.3.0 switches act as passive switches, because they accept incoming authentication requests. Regardless of the policy, E_Port is disabled if the DH-CHAP or FCAP protocol fails to authenticate the attached E_Port.</p>

Virtual Fabric considerations: Because the device authentication policy has switch and logical switch-based parameters, each logical switch is set when Virtual Fabrics is enabled. Authentication is enforced based on each logical switch's policy settings.

14.9.18 AUTH policy restrictions

Fabric OS v5.1.0 implementation of DH-CHAP/FCAP does not support integration with RADIUS. All fabric element authentication configurations are performed on a local switch basis.

Device authentication policy supports devices that are connected to the switch in point-to-point manner and is visible to the entire fabric. The following devices are not supported:

- ▶ Public loop devices
- ▶ Single private devices
- ▶ Private loop devices
- ▶ Mixed public and private devices in loop
- ▶ NPIV devices
- ▶ FICON channels
- ▶ Configupload and download will not be supported for the following AUTH attributes: auth type, hash type, group type.

Supported HBAs

The following HBAs support authentication:

- ▶ Emulex LP11000 (Tested with Storport Miniport v2.0 windows driver)
- ▶ Qlogic QLA2300 (Tested with Solaris v5.04 driver)

Authentication protocols

Use the **authUtil** command to perform the following tasks:

- ▶ Display the current authentication parameters.
- ▶ Select the authentication protocol used between switches.
- ▶ Select the DH (Diffie-Hellman) group for a switch.

Run the **authUtil** command on the switch you want to view or change. Here are the different options to specify which DH group you want to use:

- ▶ 00 – DH Null option
- ▶ 01 – 1024 bit key
- ▶ 02 – 1280 bit key
- ▶ 03 - 1536 bit key
- ▶ 04 – 2048 bit key

This section illustrates using the **authUtil** command to display the current authentication parameters and to set the authentication protocol to DH-CHAP. For more details about the **authUtil** command, see the *Fabric OS Command Reference*, 53-1001186-01.

14.9.19 Viewing current authentication parameter settings for a switch

1. Log in to the switch using an account assigned to the admin role.
2. On a switch running Fabric OS v6.0.0 or later, type **authUtil --show** as demonstrated in Example 14-32.

Example 14-32 Viewing current auth parameters

```
BDPOC01L01:admin> authUtil --show
AUTH TYPE      HASH TYPE      GROUP TYPE
-----
fcap,dhchap    sha1,md5       0,1,2,3,4

Switch Authentication Policy: PASSIVE
Device Authentication Policy: OFF
BDPOC01L01:admin>
```

14.9.20 Setting authentication protocol used by the switch to DH-CHAP

Follow these steps to set the authentication protocol:

1. Log in to the switch using an account assigned to the admin role.
2. On a switch running Fabric OS v4.x or v5.x, type **authUtil --set -a dhchap**. On a switch running Fabric OS v3.x, type **authUtil "--set -a dhchap"**. (See Example 14-33.)

Example 14-33 Setting the authentication protocol

```
BDPOC01L01:admin> authUtil --set -a dhchap
Authentication is set to dhchap.
BDPOC01L01:admin>
BDPOC01L01:admin> authutil --show
AUTH TYPE      HASH TYPE      GROUP TYPE
-----
dhchap         sha1,md5       0,1,2,3,4

Switch Authentication Policy: PASSIVE
Device Authentication Policy: OFF
BDPOC01L01:admin>
```

DH-CHAP considerations:

- ▶ When using DH-CHAP, make sure that you configure the switches at both ends of a link.
- ▶ If you set the authentication protocol to DH-CHAP, you have not yet configured shared secrets, and authentication is checked (for example, you enabled the switch), then switch authentication fails.

14.9.21 Re-authenticating E_Ports

Use the command **authUtil** to re-initiate the authentication on selected ports. It provides flexibility to initiate authentication for specified E_Ports, a set of E_Ports, or all E_Ports on the switch. This command does not work on private, loop, NPIV, and FICON devices. The command **authUtil** can re-initiate authentication only if the device was previously authenticated. If the authentication fails because shared secrets do not match, the port is disabled.

This command works independently of the authentication policy; this means you can initiate the authentication even if the switch is in PASSIVE mode. This command is used to restart authentication after changing the DH-CHAP group, hash type, or shared secret between a pair of switches.

Important: This command might bring down E_Ports if the DH-CHAP shared secrets are not installed correctly.

Follow these steps:

1. Log in to the switch using an account assigned to the admin role.
2. On a switch running Fabric OS v5.3.0 and later, type the following commands as shown in Example 14-34, Example 14-35, and Example 14-36.

Example 14-34 Example for specific ports on the switch

```
BDPOC01L01:admin> authutil --authinit 2,3,4
BDPOC01L01:admin>
```

Example 14-35 Example for all E_Ports on the switch

```
BDPOC01L01:admin> authutil --authinit allE
BDPOC01L01:admin>
```

Example 14-36 Example for enterprise-class platforms using the slot/port format

```
BDPOC01L01:admin> authutil --authinit 1/1, 1/2
```

14.9.22 Secret key pairs

When you configure the switches at both ends of a link to use DH-CHAP for authentication, you must also define a secret key pair—one for each end of the link. Use the **secAuthSecret** command to perform the following tasks:

- ▶ View the WWN of switches with a secret key pair.
- ▶ Set the secret key pair for switches.
- ▶ Remove the secret key pair for one or more switches.

Notice the following characteristics of a secret key pair:

- ▶ The secret key pair must be set up locally on every switch. The secret key pair is not distributed fabric-wide.
- ▶ If a secret key pair is not set up for a link, authentication fails. The “Authentication Failed” (reason code 05h) error will be reported and logged.
- ▶ The minimum length of a shared secret is 8 bytes and the maximum length is 40 bytes.

This section illustrates using the **secAuthSecret** command to display the list of switches in the current switch’s shared secret database and to set the secret key pair for the current switch and a connected switch. For more details about the **secAuthSecret** command, see the *Fabric OS Command Reference*, 53-1001186-01.

Security: When setting a secret key pair, note that you are entering the shared secrets in plain text. Use a secure channel (for example, SSH or the serial console) to connect to the switch on which you are setting the secrets.

14.9.23 Viewing a list of secret key pairs in the current switch database

Follow these steps:

1. Log in to the switch using an account assigned to the admin role.
2. On a switch running Fabric OS v4.x or later, type **secAuthSecret --show**.
On a switch running Fabric OS v3.x, type **secAuthSecret “--show”**.
See Example 14-37.

Example 14-37 Viewing key pairs

```
BDPOC01L01:admin>secauthsecret --show
```

WWN	DId	Name

10:00:00:60:69:80:07:52		Unknown
10:00:00:60:69:80:5b:e8	1	BDPOC01L01

```
BDPOC01L01:admin>
```

The output displays the WWN, domain ID, and name (if known) of the switches with defined shared secrets.

14.9.24 Setting a secret key pair

Follow these steps to set the secret key pair:

1. Log in to the switch using an account assigned to the admin role.
2. On a switch running Fabric OS v4.x or later, type **secAuthSecret --set**, see Example 14-38 on page 684. On a switch running Fabric OS v3.x, type **secAuthSecret "--set"**.

The command enters interactive mode. The command returns a description of itself and the necessary input; then it loops through a sequence of switch specification, peer secret entry, and local secret entry. To exit the loop, press Enter for the switch name; then type **y**. See Example 14-38.

Example 14-38 Setting a secret pair

```
BDPOC01L01:admin>secAuthSecret --set
```

This command sets up secret keys for the DH-CHAP authentication. The minimum length of a secret key is 8 characters and maximum 40 characters. Setting up secret keys does not initiate DH-CHAP authentication. If switch is configured to do DH-CHAP, it is performed whenever a port or a switch is enabled.

Warning: Please use a secure channel for setting secrets. Using an insecure channel is not safe and may compromise secrets.

Following inputs should be specified for each entry.

1. WWN for which secret is being set up.
2. Peer secret: The secret of the peer that authenticates to peer.
3. Local secret: The local secret that authenticates peer.

```
Press Enter to start setting up shared secrets >

Enter WWN, Domain, or switch name (Leave blank when done):
10:00:00:60:69:80:5b:e8
Enter peer secret: <hidden>
Re-enter peer secret: <hidden>
Enter local secret: <hidden>
Re-enter local secret: <hidden>

Enter WWN, Domain, or switch name (Leave blank when done):
Are you done? (yes, y, no, n): [no] y
Saving data to key store... Done.
```

3. Disable and enable the ports on a peer switch using the **portDisable** and **portEnable** commands.

14.9.25 Distributing the local ACL policies

Follow these steps:

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the following command (see Example 14-39) to distribute the policies:
distribute -p <database_id> -d <switch_list>
3. Table 14-16 describes the arguments for the command.

Table 14-16 *Distribute policies command*

database_id	A semicolon-separated list of the local databases to be distributed: SCC, DCC, or both.
switch_list	A semicolon-separated list of switch Domain IDs, switch names, or switch WWN addresses of the target switches that will receive the distribution. Use an asterisk (*) to distribute the database to all Fabric OS v5.2.0 and later switches in the fabric. For example, entering the command distribute -p SCC -d “*” distributes the SCC policy to all v5.2.0 and later switches in the fabric.

To distribute the Switch Connection Control Policy and Device Connection Control Policy to domains 3 and 5 in the fabric, use the command shown in Example 14-39.

Example 14-39 Using the distribute command

```
BDPOC01L01:admin> distribute -p "SCC;DCC" -d "3;5"
```

To distribute SCC, FCS, and the Password database to all domains in the fabric that support the distribute feature, use the command in Example 14-40.

Example 14-40 Using the distribute command with wildcard

```
BDPOC01L01:admin> distribute -p "SCC;FCS;PWD" -d "*"
Wildcard domains are:
1 3 5
```

14.9.26 IP Filter policy

The IP Filter policy is a set of rules applied to the IP management interfaces as a packet filtering firewall. The firewall permits or denies the traffic to go through the IP management interfaces according to the policy rules.

Fabric OS supports multiple IP Filter policies to be defined at the same time. Each IP Filter policy is identified by a name and has an associated type. Two IP Filter policy types, IPv4 and IPv6, exist to provide separate packet filtering for IPv4 and IPv6. It is not allowed to specify an IPv6 address in the IPv4 filter, or specify an IPv4 address in the IPv6 filter. There can be up to six different IP Filter policies defined for both types. Only one IP Filter policy for each IP type can be activated on the affected management IP interfaces.

Audit messages will be generated for any changes to the IP Filter policies.

The rules in the IP Filter policy are examined one at a time until the end of the list of rules. For performance reasons, the most important rules must be specified at the top.

On a chassis system, changes to persistent IP Filter policies are automatically synchronized to the standby CP when the changes are saved persistently on the active CP. The standby CP will enforce the filter policies to its management interface after policies are synchronized with the active CP.

Virtual Fabric considerations: Each logical switch cannot have its own different IP Filter policies. IP Filter policies are treated as a chassis-wide configuration and are common for all the logical switches in the chassis.

14.9.27 Creating an IP Filter policy

You can create an IP Filter policy specifying any name and using type IPv4 or IPv6. The policy created is stored in a temporary buffer, and is lost if the current command session logs out. The policy name is a unique string composed of a maximum of 20 alpha, numeric, and underscore characters. The names `default_ipv4` and `default_ipv6` are reserved for default IP filter policies. The policy name is case-insensitive and always stored as lowercase. The policy type identifies the policy as an IPv4 or IPv6 filter. There can be a maximum of six IP Filter policies created for both types. Follow these steps:

1. Log in to the switch using an account assigned to the admin role.
2. Use the following command:

```
ipfilter --create <polycyname> -type < ipv4 | ipv6 >
```

Tip: To set a IP filter, IPSEC must be enabled. This is done using the `ipseconfig --enable` command.

3. Table 14-17 describes the arguments for the command.

Table 14-17 IP filter policy

polycyname	The name of the new policy
-type	Specified as an IPv4 or IPv6 address.

14.9.28 Cloning an IP Filter policy

You can display the IP Filter policy content for the specified policy name, or all IP Filter policies if a policy name is not specified.

For each IP Filter policy, the policy name, type, persistent state and policy rules are displayed. The policy rules are listed by the rule number in ascending order. There is no pagination stop for multiple screens of information. Pipe the output to the `|more` command to achieve this result.

If a temporary buffer exists for an IP Filter policy, the `--show` subcommand displays the content in the temporary buffer, with the persistent state set to no.

1. Log in to the switch using an account assigned to the admin role.
2. Type the following command, `ipfilter --show [polycyname]` where `[polycyname]` is the name of the policy.

14.9.29 Saving an IP Filter policy

You can save one or all IP Filter policies persistently in the defined configuration. The policy name is optional for this subcommand. If the policy name is given, the IP Filter policy in the temporary buffer is saved; if the policy name is not given, all IP Filter policies in the temporary buffer are saved. Only the CLI session that owns the updated temporary buffer can run this command.

Modification to an active policy cannot be saved without being applied. So, the **--save** subcommand is blocked for the active policies. Use **--activate** instead.

1. Log in to the switch using an account assigned to the admin role.
2. Type the following command, **ipfilter --save [polycyname]** where **[polycyname]** is the name of the policy and is optional.

14.9.30 Activating an IP Filter policy

IP Filter policies are not enforced until they are activated. Only one IP Filter policy per IPv4 and IPv6 type can be active. If there is a temporary buffer for the policy, the policy is saved to the defined configuration and activated at the same time. If there is no temporary buffer for the policy, the policy existing in the defined configuration becomes active. The activated policy continues to remain in the defined configuration. The policy to be activated replaces the existing active policy of the same type. Activating the default IP Filter policies returns the IP management interface to its default state. An IP Filter policy without any rule cannot be activated. This subcommand prompts for a user confirmation before proceeding.

Follow these steps to activate the policy:

1. Log in to the switch using an account assigned to the admin role.
2. Type the following command, **ipfilter --activate <polycyname>** where **<polycyname>** is the name of the policy.

14.9.31 Deleting an IP Filter policy

You can delete a specified IP Filter policy. Deleting an IP Filter policy removes it from the temporary buffer. To permanently delete the policy from the persistent database, run **ipfilter --save**. An active IP Filter policy cannot be deleted.

1. Log in to the switch using an account assigned to the admin role.
2. Type the following command, **ipfilter --delete <polycyname>** where **<polycyname>** is the name of the policy.
3. To permanently delete the policy, type the following command:
ipfilter --save

14.9.32 IP Filter policy rules

An IP Filter policy consists of a set of rules. Each rule has an index number identifying the rule. There can be a maximum of 256 rules within an IP Filter policy. Each rule contains the following elements:

- ▶ Source Address: A source IP address or a group prefix.
- ▶ Destination Port: The destination port number or name, such as: Telnet, SSH, HTTP, HTTPS.
- ▶ Protocol: The protocol type. Supported types are TCP or UDP.
- ▶ Action: The filtering action taken by this rule, either Permit or Deny.

For an IPv4 filter policy, the source address has to be a 32-bit IPv4 address in dot decimal notation. The group prefix has to be a CIDR block prefix representation. For example, 208.130.32.0/24 represents a 24-bit IPv4 prefix starting from the most significant bit. The special prefix 0.0.0.0/0 matches any IPv4 address. In addition, the keyword any is supported to represent any IPv4 address.

For an IPv6 filter policy, the source address has to be a 128-bit IPv6 address, in a format acceptable in RFC 3513. The group prefix has to be a CIDR block prefix representation. For example, 12:AB:0:0:CD30::/64 represents a 64-bit IPv6 prefix starting from the most significant bit. In addition, the keyword any is supported to represent any IPv6 address.

For the destination port, a single port number or a port number range can be specified. According to IANA (<http://www.iana.org>), ports 0 to 1023 are well-known port numbers, ports 1024 to 49151 are registered port numbers, and ports 49152 to 65535 are dynamic or private port numbers. Well-known and registered ports are normally used by servers to accept connections, while dynamic port numbers are used by clients.

For an IP Filter policy rule, you can only select port numbers in either the well-known or the registered port number range, between 0 and 49151, inclusive. This means that you have the ability to control how to expose the management services hosted on a switch, but not the ability to affect the management traffic that is initiated from a switch. A valid port number range is represented by a dash, for example, 7-30. Alternatively, service names can also be used instead of port numbers.

Table 14-18 lists the supported service names and the corresponding port number for each.

Table 14-18 Supported service names

Service name	Port number
http	443
rpcd	897
securerpcd	898
snmp	161
ssh	22
sunrpc	111
telnet	23
www	80

TCP and UDP protocols are valid selections. Fabric OS v5.3.0 and later do not support configuration to filter other protocols. Implicitly, ICMP type 0 and type 8 packets are always allowed to support ICMP echo request and reply on commands such as **ping** and **traceroute**. For the action, only **permit** and **deny** are valid.

For every IP Filter policy, the two rules listed in Table 14-19 on page 690 are always assumed to be appended implicitly to the end of the policy. This ensures that TCP and UDP traffic to dynamic port ranges is allowed, so that management IP traffic initiated from a switch, such as syslog, radius, and ftp, is not affected.

Table 14-19 Implicit IP Filter rules

Source address	Destination port	Protocol	Action
Any	1024-65535	TCP	Permit
Any	1024-65535	UDP	Permit

A switch with Fabric OS v5.3.0 or later will have a default IP Filter policy for IPv4 and IPv6. The default IP Filter policy cannot be deleted or changed. When an alternative IP Filter policy is activated, the default IP Filter policy becomes deactivated.

Table 14-20 lists the rules of the default IP Filter policy.

Table 14-20 Default IP policy rules

Rule number	Source address	Destination port	Protocol	Action
1	Any	22	TCP	Permit
2	Any	23	TCP	Permit
3	Any	897	TCP	Permit
4	Any	898	TCP	Permit
5	Any	111	TCP	Permit
6	Any	80	TCP	Permit
7	Any	443	TCP	Permit
9	Any	161	UDP	Permit
10	Any	111	UDP	Permit
11	Any	123	UDP	Permit
12	Any	600-1023	UDP	Permit

14.9.33 IP Filter policy enforcement

An active IP Filter policy is a filter applied to the IP packets through the management interface. IPv4 management traffic passes through the active IPv4 filter policy, and IPv6 management traffic passes through the active IPv6 filter policy. The IP Filter policy applies to the incoming (ingress) management traffic only. When a packet arrives, it is compared against each rule, starting from the first rule. If a match is found for the source address, destination port, and protocol, the corresponding action for this rule is taken, and the subsequent rules in this policy are ignored. If there is no match, then it is compared to the next rule in the policy. This process continues until the incoming packet is compared to all rules in the active policy.

If none of the rules in the policy match the incoming packet, the two implicit rules are matched to the incoming packet. If the rules still do not match the packet, the default action, which is to deny, is taken.

When the IPv4 or IPv6 address for the management interface of a switch is changed through the **ipAddrSet** command or manageability tools, the active IP Filter policies automatically become enforced on the management IP interface with the changed IP address.

NAT server: If a switch is part of a LAN behind a Network Address Translation (NAT) server, depending on the NAT server configuration, the source address in an IP Filter rule might have to be the NAT server address.

14.9.34 Adding a rule to an IP Filter policy

There can be a maximum of 256 rules created for an IP Filter policy. The change to the specified IP Filter policy is not saved to the persistent configuration until a save or activate subcommand is run.

1. Log in to the switch using an account assigned to the admin role.
2. Type the following command, **ipfilter --addrule <policyname> -rule <rule_number> -sip <source IP> -dp <dest port> -proto <protocol> -act <permit | deny>**
3. Table 14-21 describes the arguments for the **ipfilter addrule** command.

Table 14-21 IP filter add rule

policyname	Specifies the policy name which is a unique string composed of a maximum of 20 alphanumeric and underscore characters. The names default_ipv4 and default_ipv6 are reserved for the default IP Filter policies. The policy name is case-insensitive and always stored as lowercase.
-rule rule number	Specifies a valid rule number between 1 and the current maximum rule number plus one.
-sip source IP	Specifies the source IP address. For IPv4 filter type, the address must be a 32-bit address in dot decimal notation, or a CIDR block IPv4 prefix. For IPv6 filter type, the address must be a 128-bit IPv6 address in any format specified by RFC, or a CIDR block IPv6 prefix.
-dp destination port	Specifies the destination port number, or a range of port numbers, or a service name.
-proto protocol	Specifies the protocol type, either TCP or UDP.
-act <permit deny>	Specifies the permit or deny action associated with this rule.

14.9.35 Deleting a rule in an IP Filter policy

Deleting a rule in the specified IP Filter policy causes the rules following the deleted rule to shift up in rule order. The change to the specified IP Filter policy is not saved to persistent configuration until a save or activate subcommand is run.

1. Log in to the switch using an account assigned to the admin role.
2. Type the following command, **ipfilter --delrule <policyname> -rule <rule number>**

14.9.36 Aborting a transaction associated with IP Filter

A transaction is associated with a command line or manageability session. It is opened implicitly when the **--create**, **--addrule**, **--delrule**, **--clone**, and **--delete** subcommands are run. The **--transabort**, **--save**, or **--activate** subcommands explicitly end the transaction owned by the current command line or manageability session. If a transaction is not ended, other command line or manageability sessions are blocked on the subcommands that would open a new transaction.

1. Log in to the switch using an account assigned to the admin role.
2. Type the following command, **ipfilter --transabort**

14.9.37 IP Filter policy distributions

The IP Filter policy is manually distributed using the **distribute -p "IPFILTER"** command. The distribution includes both active and defined IP Filter policies. All policies are combined as a single entity to be distributed and cannot be selectively distributed. However, you can choose the time at which to implement the policy for optimization purposes. If a distribution includes an active IP Filter policy, the receiving switches activate the same IP Filter policy automatically. When a switch receives IP Filter policies, all uncommitted changes left in its local transaction buffer are lost, and the transaction is aborted.

Switches with Fabric OS v5.3.0 or later have the ability to accept or deny IP Filter policy distribution, through the commands **fddCfg --localaccept** or **fddCfg --localreject**. However, automatic distribution of IP Filter policy through Fabric Wide Consistent Policy is not supported in Fabric OS v6.2.0.

14.9.38 IP Filter policy restrictions

In a mixed fabric with Fabric OS v5.3.0 or later and pre-v5.3.0 switches, IP Filter policies cannot be distributed from a Fabric OS v6.2.0 switch to a pre-v5.3.0 switch. This means that the sending switch will fail a **distribute -p "IPFILTER"** operation, if the specified receiving domain list contains switches with Fabric OS v5.2.0 and earlier. When the asterisk (*) is used as the receiving domain, the sending switch distributes the IP Filter policies only to switches with Fabric OS v5.3.0 or later.



Adaptive Networking

Adaptive Networking is a suite of tools and capabilities that enables you to ensure optimized behavior in the SAN. Even under the worst conditions of congestion, Adaptive Networking features can maximize the fabric behavior and provide necessary bandwidth for high-priority, mission-critical applications, and connections.

Adaptive Networking is not a single feature but a suite of tools and capabilities provided for the SAN optimization.

In this chapter we discuss the following features in the Adaptive Networking suite:

- ▶ Traffic Isolation
- ▶ Quality of service (QoS) Ingress Rate Limiting
- ▶ QoS SID/DID Traffic Prioritization

Top Talkers, which is another Adaptive Networking feature, requires the Advanced Performance Monitoring license, and is described in Chapter 16, “Performance monitoring” on page 739.

15.1 Traffic Management

Traffic Management consists of Ingress Rate Limiting and Traffic Isolation (see Table 15-1). One feature allows you to control the flow of interswitch traffic and the second restricts the speed of traffic on a particular port. Only Ingress Rate Limiting requires Adaptive Networking License.

Table 15-1 Traffic Management

Service	Service name	Licence required
Traffic Management	Ingress Rate Limiting	Adaptive Networking
	Traffic Isolation	None

Licensing: With Fabric OS v6.4.0 it is possible to have slot-based licensing.

15.1.1 Committed rate considerations on FCIP

Starting with Fabric OS v6.4.0 in the case of an FCIP configuration, when you are configuring the traffic limit, you must have the same committed rate configuration on each end of a circuit. In previous releases you can have different committed rates on each side. This committed rate will be enforced at circuit initialization, and if the committed rates do not match, an error will be shown in the CLI or log.

Figure 15-1 illustrates this configuration.

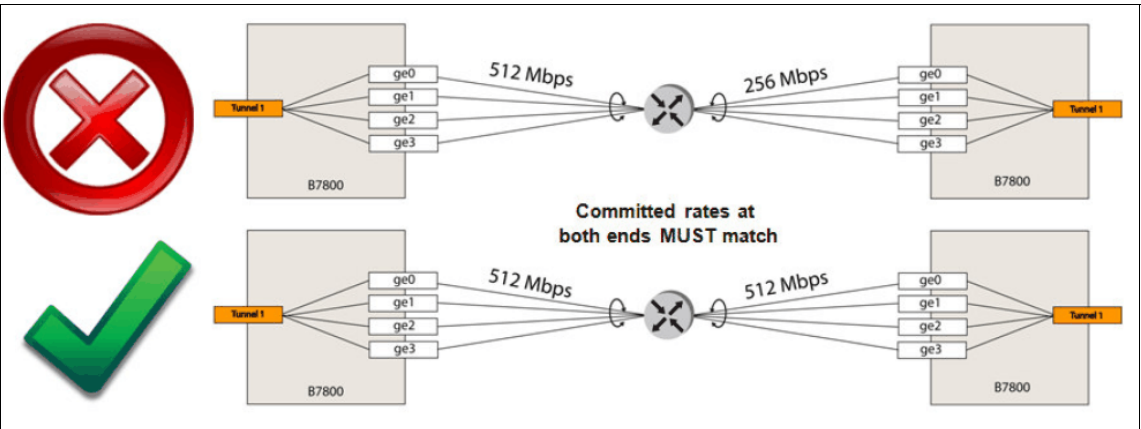


Figure 15-1 Committed rates at both ends must match

This includes tunnels that exist already and are upgraded to Fabric OS v6.4.0. After upgrade, the tunnel will not be able to go online, and an error will be generated.

Validation: Remember to validate the committed rates in your configuration, especially in the case of a Fabric OS update.

15.1.2 Adaptive Rate Limiting considerations

When using Adaptive Rate Limiting (ARL), consider the following guidelines:

- ▶ The maximum committed rate cannot be larger than five times the minimum committed rate, as an example, this would mean:
 - A minimum of 100 Mbps and a maximum of 500 Mbps is allowed.
 - A minimum of 10 Mbps and a maximum of 500 Mbps will not be allowed.
- ▶ The CLI will produce an error if the configuration request does not meet the preceding guidelines.

Figure 15-2 shows an example of this ARL guidelines.

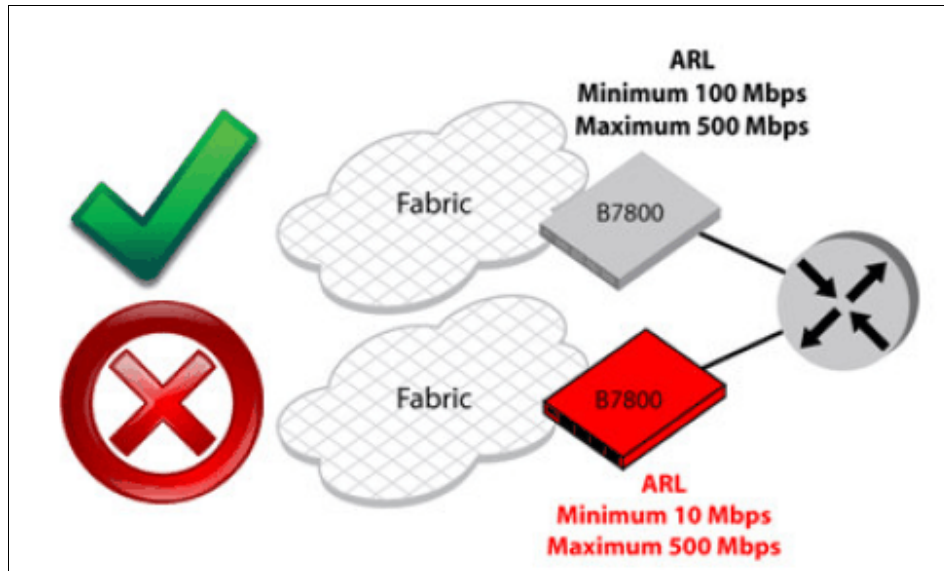


Figure 15-2 ARL limits

When updating an existing tunnel, the tunnel will continue to function using an invalid configuration. The administrator will not be able to make additional changes to the tunnel configuration until the ARL delta is compliant. Consider this factor when updating and later validating your configuration to avoid compliance issues.

15.1.3 Trunking across multiple FCIP circuits

When trunking across multiple FCIP circuits, the delta bandwidth between the circuits must be no greater than a factor of four. As an example, this would mean:

- ▶ Trunking between a circuit running on an OC3 (155.52 Mbps) and another running on an OC12 (622.08 Mbps) is allowed.
- ▶ Trunking between a 10 Mbps circuit and a 500 Mbps circuit is *not* desirable.

This rule will not be enforced with the CLI, but it is not supported, so consider it when defining your trunking. If the factor is greater than four, the tunnel might not fully utilize all the bandwidth available for the circuits, and you will not be using the optimal configuration.

This restriction only includes circuits with the same metric values (standby circuits, metric 1, are not included in this calculation).

In general, the minimum committed rate of a circuit will be 10 Mbps, and will be enforced by the CLI. A configuration attempt lower than this will fail.

Rate: With Fabric OS v6.3, the minimum committed rate was of 1.544 Mbps.

When upgrading an existing tunnel, the tunnel will continue to function using an invalid configuration. The administrator will not be able to make additional changes in the tunnel configuration until the minimum commit rate is compliant. Remember that some configurations are not supported, even if they seem to work.

15.1.4 Supported packet loss and delay

In certain cases, the tunnel might have tolerance to packet loss, and support a certain delay. This is well documented, and in Table 15-2 we can see the supported values for the latest two releases of FabricOS.

The following table shows the supported packet loss and delay in the two latest releases of FabricOS.

Table 15-2 Supported packet loss and delay

Tunnels	Fabric OS v6.3	Fabric OS v6.4.0
Both ends 1GbE	<ul style="list-style-type: none"> ▶ 200 ms latency ▶ 1% packet loss 	<ul style="list-style-type: none"> ▶ 200 ms latency ▶ 1% packet loss
One or both ends 10GbE	<ul style="list-style-type: none"> ▶ 50 ms latency ▶ 0.1% packet loss 	<ul style="list-style-type: none"> ▶ 100 ms latency ▶ 0.1% packet loss

15.1.5 Scalability considerations

When planning your network, scalability should be considered. It is very common to start with a small to medium configuration, and plan an upgrade in the future. In Fabric OS release 6.4.0, there is support for up to four FCoE 10GbE blades in a chassis.

Important: Downgrading to Fabric OS v6.3 will fail if there are more than two FCoE 10 GbE blades in the chassis.

15.2 Ingress rate limiting

Ingress rate limiting restricts the speed of traffic from a particular device to the switch port:

- ▶ It allows you to reduce existing congestion in the network or proactively avoid congestion.
- ▶ It enables you to offer flexible bandwidth limit services based on requirements.
- ▶ It allows you to let more important devices use the network bandwidth during specific services, such as network backup.

This feature is only available on 8 Gbps platforms/blades, which can be running at any supported speed (8,4,2,1 Gbps). By restriction of the speed from a particular device, we mean the following characteristics:

- ▶ ASIC delays the return of R_RDYs to external device by throttling back the *ingress speed*.
- ▶ The throughput is limited on the ingress side of the port.

The usage limitation is that it is restricted to F/FL_Ports only.

Support: Ingress Rate Limiting is supported on F/FL ports on the SAN24B-4, SAN40B-4, SAN06B-R, SAN80B-4, SAN 768B, or SAN384B switches. It is not supported on E/EX_Ports.

Ingress Rate Limiting will help you if you will experience “choke points” in the fabric, which can be caused by:

- ▶ Slow draining devices
- ▶ Congested ISLs

Virtual Fabrics: If the Virtual Fabrics feature is enabled, the rate limit configuration on a port is on a per-logical switch basis. That is, if a port is configured to have a certain rate limit value, and the port is then moved to a different logical switch, it would have no rate limit applied to it in the new logical switch. If that same port is moved back to the original logical switch, it would have the original rate limit take effect again.

Figure 15-3 shows the configuration of two servers sending and receiving traffic from one storage device.

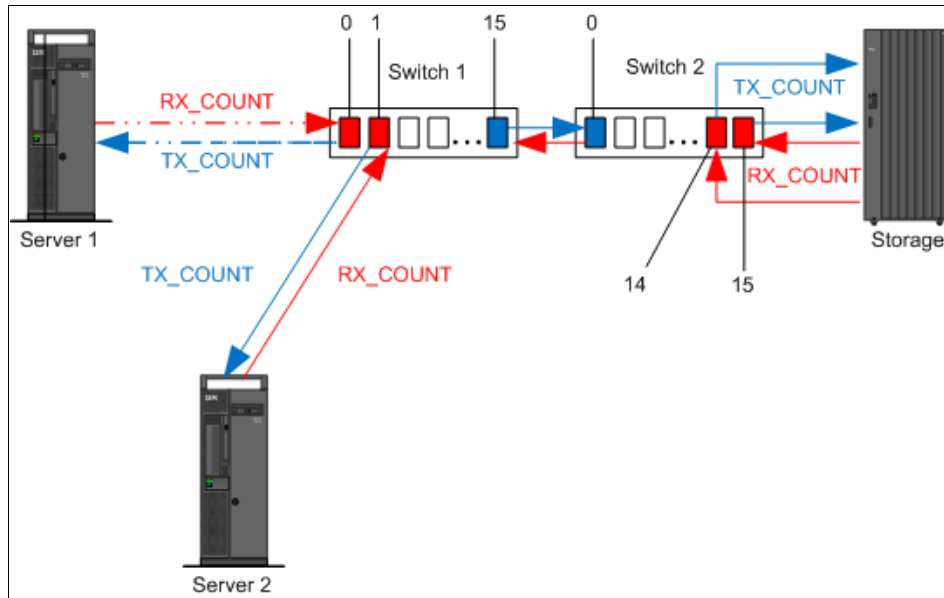


Figure 15-3 Two servers and one storage configuration

Referring to Figure 15-3 only the following ports can participate in Ingress Rate Limit settings:

- ▶ Port 0 switch 1
- ▶ Port 1 switch 1
- ▶ Port 14 switch 2
- ▶ Port 15 switch 2

It cannot be enabled for port 15 switch 1, and port 0 switch 2.

The settings for Ingress Rate Limiting are unidirectional. We describe what we mean by this in the following example:

- ▶ Ingress Rate Limiting is enabled only on port 0 switch 1 in Figure 15-3.
- ▶ Traffic returning from the target (port 14 and 15 on switch 2) would travel at full line speed to both servers, unless the ingress side of the target's ports (port 14 and 15 on switch 2) are also limited if both ports are in the same zone with port 0 server 1.
- ▶ When the ingress side of the target's ports are also throttled back, then traffic would be rate limited in both directions. In that case it can affect the transmission back to server 2 as well.

15.2.1 Ingress Rate limiting with the CLI

To set Ingress Rate Limiting on a given port, use the command in Example 15-1. The *rate* parameter of the command in Example 15-1 is set in Mbps.

Example 15-1 Setting Ingress Rate Limiting on a port

```
portcfgqos -setratelimit [slot/]port rate
```

To set the Ingress rate limit from a given port, use the command in Example 15-2.

Example 15-2 Setting Ingress Rate Limiting from a port

```
portcfgqos --resetratelimit [slot/]port
```

To show a port with Ingress Rate Limiting, use the command in Example 15-3.

Example 15-3 Showing Ingress Rate Limiting

```
portcfgshow 2/28
```

Example 15-4 shows real values when setting and displaying port speed limit.

Example 15-4 Setting Ingress rate Limiting on a port 2/28

```
IBM_SAN384B_213:FID128:admin> portcfgqos --setratelimit 2/28 200
```

```
IBM_SAN384B_213:FID128:admin> portcfgshow 2/28
```

Area Number:	92
Speed Level:	AUTO(HW)
Fill Word:	0(Idle-Idle)
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
ISL R_RDY Mode	OFF
RSCN Suppressed	OFF
Persistent Disable	OFF
NPIV capability	ON
QOS E_Port	ON
Port Auto Disable:	OFF
Rate Limit	0.2G
EX Port	OFF
Mirror Port	OFF
Credit Recovery	ON
F_Port Buffers	OFF

To set the Ingress Rate Limit *from* port 2/28, issue the command shown in Example 15-5.

Example 15-5 Setting Ingress Rate Limiting from port 2/28

```
IBM_SAN384B_213:FID128:admin> portcfgqos --resetratelimit 2/28
```

```
IBM_SAN384B_213:FID128:admin> portcfgshow 2/28
```

Area Number:	92
Speed Level:	AUTO(HW)
Fill Word:	0(Idle-Idle)
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

ISL R_RDY Mode	OFF
RSCN Suppressed	OFF
Persistent Disable	OFF
NPIV capability	ON
QOS E_Port	ON
Port Auto Disable:	OFF
Rate Limit	OFF
EX Port	OFF
Mirror Port	OFF
Credit Recovery	ON
F_Port Buffers	OFF

15.2.2 Ingress Rate Limiting with Web Tools

You can set up Ingress Rate Limiting for a particular port using WebTools. Click the port, then **Edit Configuration** (see Figure 15-4).

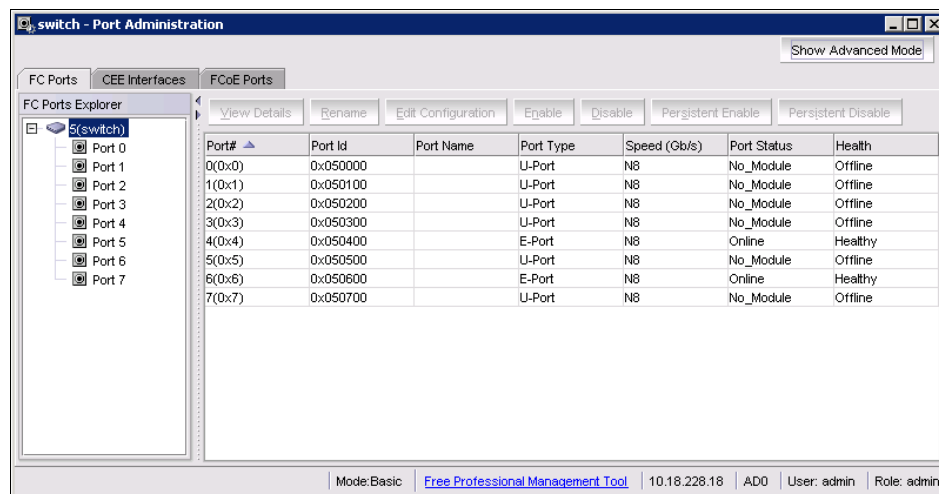


Figure 15-4 Port Administration in Web Tools

The dialog box in Figure 15-5 displays.

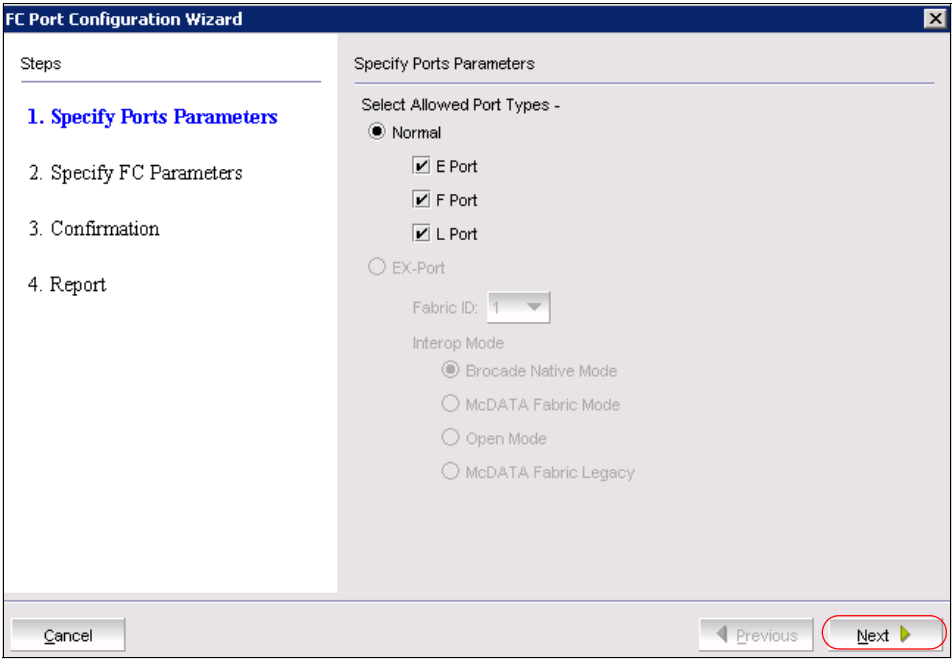


Figure 15-5 FC Port Configuration Wizard

Click **Next**, and on the next page of the dialog box, you see a list box with the possible changes of the speed on a port, as shown in Figure 15-6.

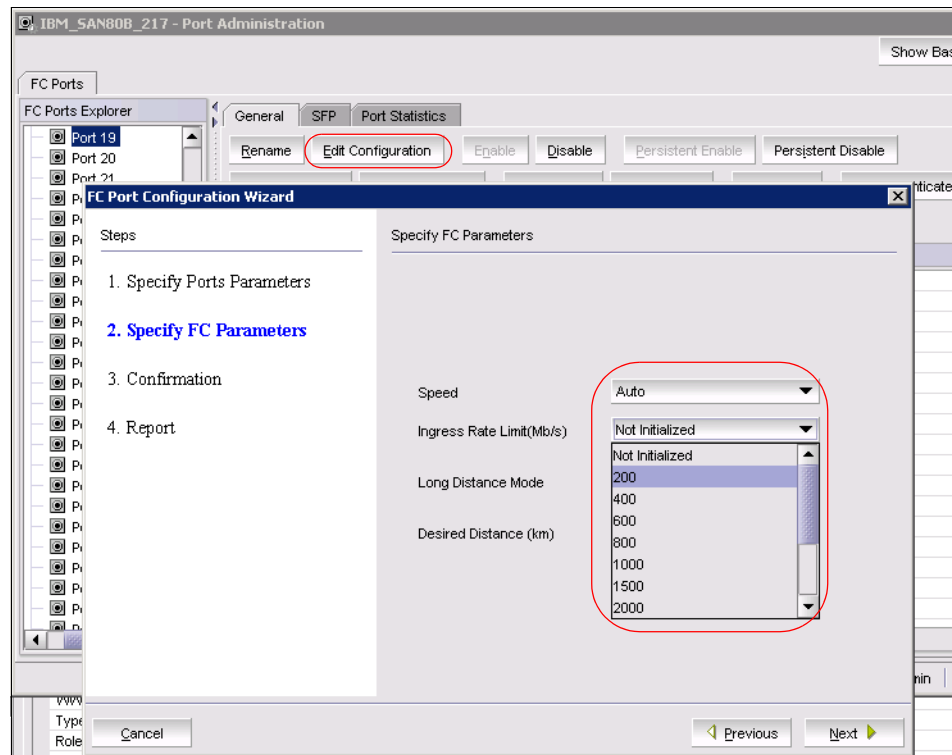


Figure 15-6 Ingress Rate Limiting with WebTools

15.3 Traffic Isolation

The second subject in Traffic Management is Traffic Isolation.

You can start to consider Traffic Isolation when:

- ▶ The host/target pair needs to have high priority traffic with no interruptions.
- ▶ The host/target pair is exchanging very high volumes of data and the data flow should have a low priority to avoid congestion in the SAN.
- ▶ The host/target pair needs to have a dedicated connection.

Traffic isolation is implemented using special zones called Traffic Isolation zones (TI zones). TI zones have the following characteristics:

- ▶ They use special zoning commands:
 - They control the routing of frames between zone members.
 - They do not control access to devices.
- ▶ They use a standard zoning configuration that must be in effect:
 - TI zones do not modify the routing table.
 - Routes are not being modified.
 - Existing routes are being dedicated for use by host/target pairs.
- ▶ The TI zone contains the set of N_Ports and E_Ports to be used for a specific traffic flow.

Figure 15-7 shows a TI zone configuration that consists of the following ports:

- ▶ N_Ports 1,0; 3,15
- ▶ E_Ports 1,14; 2,0; 2,14; 3,0

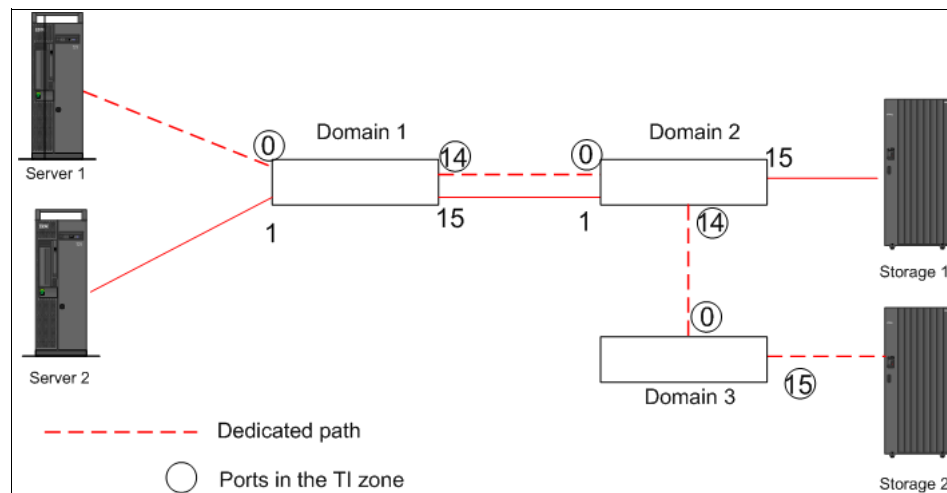


Figure 15-7 TI zone

The dotted line in Figure 15-7 indicates the dedicated path from Server 1 to Storage 2:

- ▶ Traffic entering Domain 1 from N_Port 0 is routed through E_Port 14.
- ▶ Traffic entering Domain 2 from E_Port 0 is routed to E_Port 14
- ▶ Traffic entering Domain 3 from E_Port 0 is routed to N_Port 15
- ▶ Traffic coming from port 1 in Domain 1 would not use E_Port 14, but would use E_Port 15 instead.

Traffic Isolation: The new features of the Traffic Isolation zones have been explained in Chapter 4, “Fabric Operating System” on page 91. Consider them when implementing Traffic Isolation.

15.3.1 TI zone failover

A TI zone can have failover enabled or disabled. In Table 15-3 we show the comparison of traffic behavior when failover is enabled and disabled in TI zone.

Table 15-3 Comparison of traffic behavior - failover disabled and enabled

Zones	Failover enabled	Failover disabled
TI zones	If the dedicated path cannot be used, the TI zone traffic will use a non-dedicated path instead.	If the dedicated path cannot be used, traffic for that TI zone is halted until the dedicated path is fixed, but the communication is maintained.
Non - TI Zones	Non-TI zone traffic will use the dedicated path if no other paths through the fabric exist, or if the non-dedicated paths are not the shortest paths.	Non-TI zone traffic will never use the dedicated path, even if there are no other paths through the fabric.

We assume the following conditions for Figure 15-3:

- ▶ The dedicated ISL between Domain 1 and Domain 2 goes offline, then the following events occur, depending on the failover option:
 - If failover is enabled for the TI zone:
 - The traffic is routed from Domain 1 to Domain 2 through E_Ports “1,15” and “2,1”.
 - When the failed TI zone is restored, traffic will be automatically failed back to the original route.

Failback: Failback is not a configurable feature.

- If failover is disabled for the TI zone:
 - The traffic is halted until the ISL between Domain 1 and Domain 2 is back online.

RSCN: An RSCN will be generated noting the failure of a path.

- When the TI zoned route is restored, traffic will be automatically rerouted back onto the TI zone route.

RSCN: An RSCN will be generated noting the restoration of a path.

- ▶ If the non-dedicated ISL between Domain 1 and Domain 2 goes offline, then the following events occur, depending on the failover option:
 - If failover is enabled for the TI zone, non-TI zone traffic is routed from Domain 1 to Domain 2 through the dedicated ISL.
 - If failover is disabled for the TI zone, non-TI zone traffic is halted until the non-dedicated ISL between Domain 1 and Domain 2 is back online.

Disabled failover considerations

When disabling failover, keep in mind the following considerations:

- ▶ Ensure that there are non-dedicated paths through the fabric for all devices that are not in a TI zone.
- ▶ Ensure that there are multiple paths between switches.
- ▶ Disabling failover locks the specified route so that only TI zone traffic can use it. Non-TI zone traffic, including domain controller frames, are excluded from using the dedicated path.

In Figure 15-8, if failover is disabled, Domain 4 cannot send domain controller frames to Domain 2 and 3. Domain controller frames include zone updates and Name Server queries.

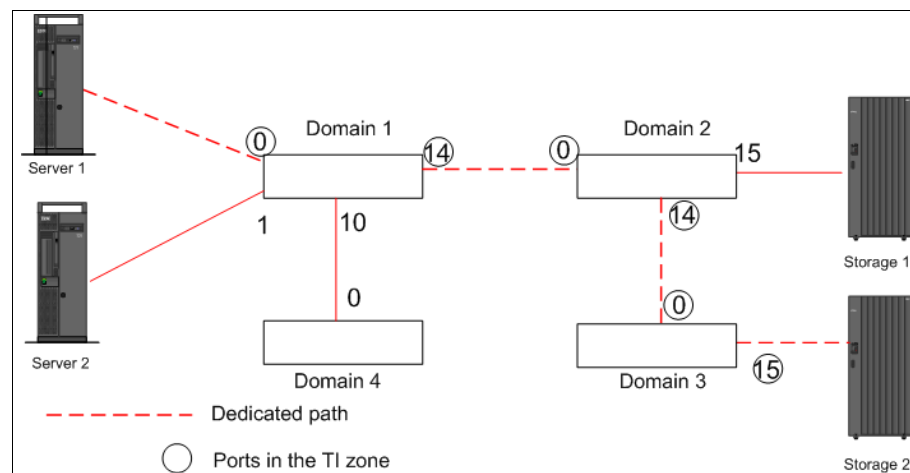


Figure 15-8 Failover disabled - Isolating Domain 4

Zoning considerations

Take the following considerations into account with zoning:

- ▶ Ensure that regular zone definitions match the TI zone definition. This is because the routing rules imposed by TI zones with failover disabled, will override regular zone definitions.
- ▶ Regular zone definitions should use Domain, Index (D,I) notation and not WWN notation; otherwise, RSCN notifications are not sent to the devices if the dedicated path is broken.
- ▶ Ensure that the insistent Domain ID feature is enabled; if a switch changes its active domain ID, the route is broken.
- ▶ A given port (N_Port or E_Port) used in a TI zone should not be a member of more than one TI zone.
- ▶ TI zones reside only in the defined configuration and not in the effective configuration. When you make any changes to TI zones, including creating or modifying them, you must enable the effective configuration for the changes to take effect, even if the effective configuration is unchanged.

15.3.2 FSPF routing rules and traffic isolation

FSPF, or Fabric Shortest Path First, is a path selection protocol for Fibre Channel Fabrics.

Table 15-4 shows the FSPF actions with TI zones failover.

Table 15-4 FSPF actions

Dedicated ISL	Failover enabled	Failover disabled
dedicated ISL is not the shortest path ISL	If failover is enabled, the traffic path for the TI zone is broken, and TI zone traffic uses the lowest cost path instead.	If failover is disabled, the TI zone traffic is blocked.
dedicated ISL is the only shortest path ISL	If failover is enabled, non-TI zone traffic as well as TI zone traffic uses the dedicated ISL.	If failover is disabled, non-TI zone traffic is blocked because it cannot use the dedicated ISL, which is the lowest cost path.

The considerations described in Table 15-4 are illustrated in the following figures.

Figure 15-9 shows the situation when the dedicated path is the only shortest path:

- There is a dedicated path between Domain 1 and Domain 2 and a non-dedicated, path that passes through Domain 4:
 - If failover is enabled:
 - All traffic will use the dedicated path, because the non-dedicated path is not the shortest path.
 - Server 2 can reach Storage 1 using dedicated path 1,14; 2,0.
 - If failover is disabled:
 - Non-TI zone traffic is blocked because the non-dedicated path is not the shortest path.
 - Server 2 cannot reach Storage 1 at all.

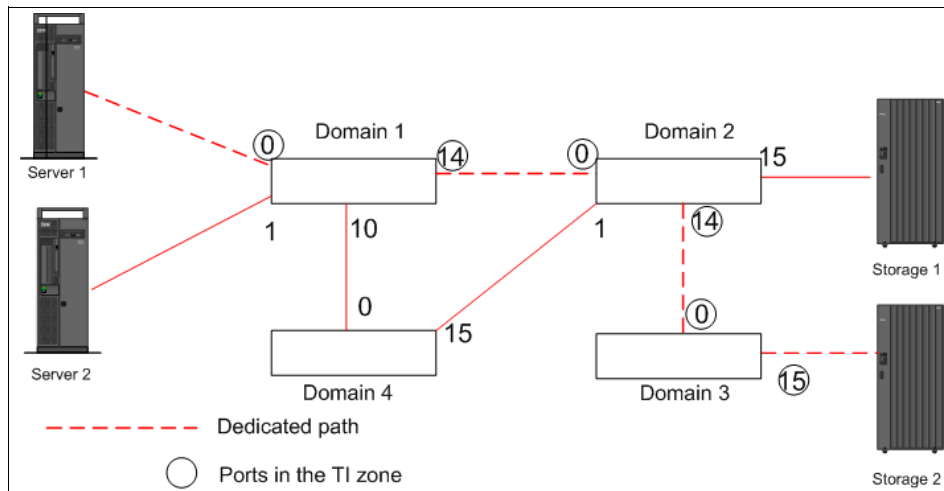


Figure 15-9 Dedicated path is the only shortest path

Figure 15-10 shows the situation when the dedicated path is not the shortest path:

- The dedicated path between Domain 1 and Domain 4 exists, but is not the shortest path.
 - If failover is enabled:
 - The TI zone traffic uses the shortest path, even though the E_Ports are not in the TI zone.
 - Server 1 reaches Storage 2 using the non-dedicated path 1,14; 2,0.

- If failover is disabled:
 - The TI zone traffic stops until the dedicated path is configured to become the shortest path.
 - Server 1 cannot reach Storage 1 at all.

Figure 15-10 shows the situation when the dedicated path is not the shortest path.

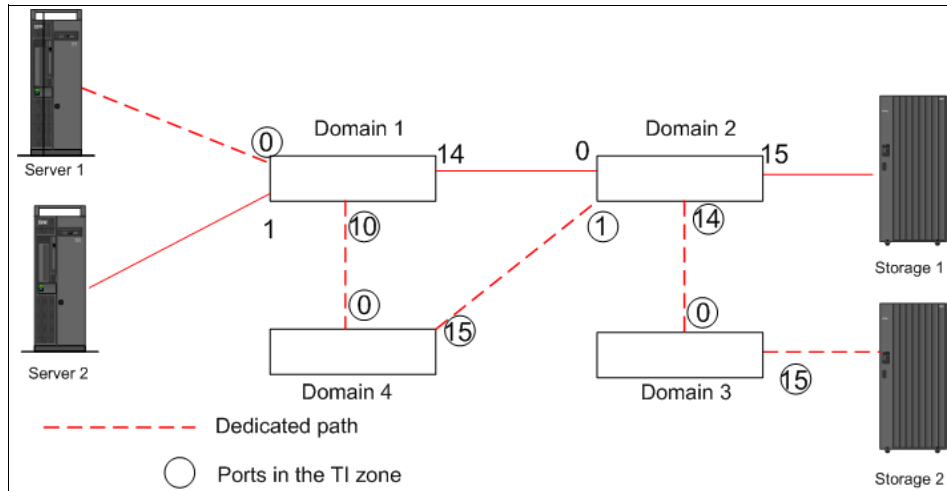


Figure 15-10 Dedicated path is not the shortest path

15.3.3 TI zone misconfiguration example

Figure 15-11 shows the following configurations:

- ▶ Two servers: Server 2 and Server 3 share the non-dedicated link to both storage devices: Storage 1 and Storage 2.
- ▶ To ensure that we have maximum throughput from Server 1 to both storage devices (Storage 1 and Storage 2), a dedicated link is used.
- ▶ The E_Port 0 in Domain 2 switch was erroneously omitted from the TI zone.
- ▶ The Domain 2 switch assumes that traffic coming from E_Port 9 is not part of the TI zone:
 - If failover is enabled:
 - Traffic is routed to E_Port 7 to reach Storage 2 and to port 15 to reach Storage 1.
 - If failover is disabled:
 - The route is broken and traffic stops.

- The net result is that we have three servers on non-dedicated paths and only after we have enabled the failover.




Figure 15-11 Misconfiguration of TI zones

Topology: Each TI zone is interpreted by each switch and each switch considers only the routing required for its local ports. No consideration is given to the overall topology.

15.3.4 Supported configurations

TI zones are supported on the following platforms:

- IBM SAN24B-4, SAN40B-4, SAN80B-4, Brocade 7500, 7500E, 7600 switches, IBM SAN256B, and Backbone platforms, all configured in Brocade Native Mode (interopmode 0)
- Switches running Fabric OS v6.0 or later

15.3.5 Virtual Fabric configuration

TI zones can be created in a logical fabric as in regular fabrics, except that:

- The disable failover option is not supported in logical fabrics that use XISLs.

- To create a TI zone for a logical fabric that uses XISLs, you must create two TI zones: one in the logical fabric and one in the base fabric. The combination of TI zones in the base fabric and logical fabric sets the path through the base fabric for logical switches.

15.3.6 TI zones using CLI

The general command syntax for TI zones is as follows:

```
Synopsis zone      --copy [source_AD.] source_zone_object
                    [dest_zone_object] [-f]
                    --expunge "zone_object"
                    --validate [[-f l] [-m mode] ["zone_object"]]
```

To create and manage traffic Isolation zones:

```
--create -t objecttype [-o optlist] name -p portlist
--add [-o optlist] name -p portlist
--remove name -p portlist
--delete name
--activate name
--deactivate name
--show [name]
--operation -t objtype [-o optlist] name -p portlist
```

Operands:

-t objecttype - Specifies the zone object type. This operand is supported only with the *--create* option. To create a TI zone, the value is *ti*.

-o optlist - Specifies list of options to control activation, deactivation, and failover mode.

If this option is not specified, the zone is created, by default, with failover enabled, and the zone will be activated. This operand is supported only with the *--create* and *--add* options.

Valid values for *optlist* are:

- a* - Activates the specified zone.
- d* - Deactivates the specified zone.
- n* - Disables failover mode.

Commands: The **cfgenable** command is required to commit all commands.

Example 15-6 shows the creation of a TI zone called “bluezone.”

Example 15-6 Creation of TI zone

```
IBM_SAN384B_213:FID128:admin> zone --create -t ti "bluezone" -p "1,19;  
1,56; 1,57; 2,92; 2,85; 2,21"
```

```
IBM_SAN384B_213:FID128:admin> cfgenable SiteA_fab1
```

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 'SiteA_fab1' configuration (yes, y, no, n): [no]
y

zone config "SiteA_fab1" is in effect

Updating flash ...

```
IBM_SAN384B_213:FID128:admin> zone --show
```

Defined TI zone configuration:

TI Zone Name: bluezone

Port List: 1,19; 1,56; 1,57; 2,92; 2,85; 2,21

Configured Status: Activated / Failover-Enabled

Enabled Status: Activated / Failover-Enabled

The **zone --remove** command allows the removal of previously defined ports in the TI zone (see Example 15-7).

Example 15-7 Removing ports from the TI zone “bluezone”

```
IBM_SAN384B_213:FID128:admin> zone --remove bluezone -p "2,92; 1,57"
```

```
IBM_SAN384B_213:FID128:admin> cfgenable SiteA_fab1
```

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 'SiteA_fab1' configuration (yes, y, no, n): [no]
y

zone config "SiteA_fab1" is in effect

Updating flash ...

```
IBM_SAN384B_213:FID128:admin> zone --show
Defined TI zone configuration:
```

TI Zone Name: bluezone

Port List: 1,19; 1,56; 2,85; 2,21

Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled

To add ports to the TI zone, use **zone --add** (see Example 15-8).

Example 15-8 Adding ports to TI zone

```
IBM_SAN384B_213:FID128:admin> zone --add bluezone -p "2,92; 1,57"
IBM_SAN384B_213:FID128:admin> cfm enable SiteA_fab1
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 'SiteA_fab1' configuration (yes, y, no, n): [no]
y
zone config "SiteA_fab1" is in effect
Updating flash ...
IBM_SAN384B_213:FID128:admin> zone --show
Defined TI zone configuration:
```

TI Zone Name: bluezone

Port List: 2,92; 1,57; 1,19; 1,56; 2,85; 2,21

Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled

If you want to disable failover, use the commands in Example 15-9.

Example 15-9 Adding ports to TI zone and disabling failover

```
IBM_SAN384B_213:FID128:admin> zone --add -o n "bluezone" -p "2,92"
IBM_SAN384B_213:FID128:admin> zone --show
Defined TI zone configuration:
```

TI Zone Name: bluezone

Port List: 2,92; 1,19; 1,56; 2,85; 2,21

Configured Status: Activated / Failover-Disabled

Enabled Status: Activated / Failover-Enabled

```
IBM_SAN384B_213:FID128:admin> cfgenable SiteA_fab1
```

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 'SiteA_fab1' configuration (yes, y, no, n): [no]

y

zone config "SiteA_fab1" is in effect

Updating flash ...

```
IBM_SAN384B_213:FID128:admin> zone --show
```

Defined TI zone configuration:

TI Zone Name: bluezone

Port List: 2,92; 1,19; 1,56; 2,85; 2,21

Configured Status: Activated / Failover-Disabled

Enabled Status: Activated / Failover-Disabled

In Example 15-8 on page 715 we can see that before enabling zoning we had “**Failover-Enabled**” even though we disabled it in the previous command.

This is correct behavior. Remember to issue the **cfgenable** command always after the **zone** command to force enabled configuration changes.

To enable the failover, run the command with the **zone --add** with the option **-o f** as shown in Example 15-10.

Example 15-10 Adding ports to TI zone and enabling failover

```
IBM_SAN384B_213:FID128:admin> zone --add -o f "bluezone" -p "1,57"
```

```
IBM_SAN384B_213:FID128:admin> zone --show
```

Defined TI zone configuration:

TI Zone Name: bluezone

Port List: 1,57; 2,92; 1,19; 1,56; 2,85; 2,21

Configured Status: Activated / Failover-Enabled

Enabled Status: Activated / Failover-Disabled

```
IBM_SAN384B_213:FID128:admin> cfgenable SiteA_fab1
```

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 'SiteA_fab1' configuration (yes, y, no, n): [no]

y

zone config "SiteA_fab1" is in effect

Updating flash ...

```
IBM_SAN384B_213:FID128:admin> zone --show
```

Defined TI zone configuration:

TI Zone Name: bluezone

Port List: 1,57; 2,92; 1,19; 1,56; 2,85; 2,21

Configured Status: Activated / Failover-Enabled

Enabled Status: Activated / Failover-Enabled

To deactivate the TI zone, use the commands shown in Example 15-11.

Example 15-11 Deactivating the TI zone

```
IBM_SAN384B_213:FID128:admin> zone --deactivate bluezone
```

```
IBM_SAN384B_213:FID128:admin> cfgenable SiteA_fab1
```

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 'SiteA_fab1' configuration (yes, y, no, n): [no]

y

zone config "SiteA_fab1" is in effect

Updating flash ...

```
IBM_SAN384B_213:FID128:admin> zone --show
```

Defined TI zone configuration:

TI Zone Name: bluezone

Port List: 1,57; 2,92; 1,19; 1,56; 2,85; 2,21

Configured Status: Deactivated / Failover-Enabled
Enabled Status: Deactivated

To activate the TI zone, use the commands shown in Example 15-12.

Example 15-12 Activating TI zone

```
IBM_SAN384B_213:FID128:admin> zone --activate bluezone
IBM_SAN384B_213:FID128:admin> cfgenable SiteA_fab1
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 'SiteA_fab1' configuration (yes, y, no, n): [no]
y
zone config "SiteA_fab1" is in effect
Updating flash ...
IBM_SAN384B_213:FID128:admin> zone --show
Defined TI zone configuration:
```

TI Zone Name: bluezone

Port List: 1,57; 2,92; 1,19; 1,56; 2,85; 2,21

Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled

Finally, to delete the TI zone, use the commands in Example 15-13.

Example 15-13 Deleting TI zone "bluezone"

```
IBM_SAN384B_213:FID128:admin> zone --delete bluezone
IBM_SAN384B_213:FID128:admin> cfgenable SiteA_fab1
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 'SiteA_fab1' configuration (yes, y, no, n): [no]
y
zone config "SiteA_fab1" is in effect
Updating flash ...
IBM_SAN384B_213:FID128:admin> zone --show
Defined TI zone configuration:
no TI zone configuration defined
```

15.3.7 Other zoning CLI commands

Attention: Because the TI zone is only part of the defined configuration, the following commands will generate an error if used with a TI zone:

- ▶ **cfgadd** and **cfgcreate**
- ▶ **zonecreate**, **zoneadd**, **zonedel**, and **zoneremove**

Example 15-14 shows the usage of the command **zoneshow** with the TI zone.

Example 15-14 Usage of zoneshow command with a TI zone

```
IBM_SAN384B_213:FID128:admin> zone --create -t ti "bluezone" -p "1,19; 1,56; 1,57; 2,92; 2,85; 2,21"
IBM_SAN384B_213:FID128:admin> cfgenable SiteA_fab1
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 'SiteA_fab1' configuration (yes, y, no, n): [no]
y
zone config "SiteA_fab1" is in effect
Updating flash ...
IBM_SAN384B_213:FID128:admin> zone --show
Defined TI zone configuration:
```

TI Zone Name: bluezone

Port List: 1,19; 1,56; 1,57; 2,92; 2,85; 2,21

Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled

```

IBM_SAN384B_213:FID128:admin> zoneshow
Defined configuration:
cfg:   SiteA_fab1
      AIX_1_DS4000_A; serverX_1_DS4000
cfg:   t_r_a_f_f_i_c_i_s_o_c_fg
      bluezone
zone:   AIX_1_DS4000_A
      AIX_1; DS4000_A
zone:   bluezone
      1,19; 1,56; 1,57; 2,92; 2,85; 2,21
zone:   serverX_1_DS4000
      serverX_1; DS4000_A
zone:   t_r_a_f_f_i_c_i_s_o_prop_zn
      1,3
alias:  AIX_1    10:00:00:00:c9:4c:8c:1c
alias:  DS4000_A 20:06:00:a0:b8:48:58:a1
alias:  serverX_1 10:00:00:05:1e:53:10:8b

Effective configuration:
cfg:   SiteA_fab1
zone:   AIX_1_DS4000_A
      10:00:00:00:c9:4c:8c:1c
      20:06:00:a0:b8:48:58:a1
zone:   serverX_1_DS4000
      10:00:00:05:1e:53:10:8b
      20:06:00:a0:b8:48:58:a1

```

As you can see, there is no TI zone in the effective configuration as shown by the **zoneshow** command.

Notes:

- The existing commands **cfgshow** and **zoneshow** can be used to display TI zones and their members.
- Failover attributes and status will not be displayed.

15.3.8 TI zones with DCFM

You can configure TI zones with DCFM. Click the zoning icon in the DCFM Main Toolbar (see Figure 15-12).




Figure 15-12 Zoning icon in DCFM Main Toolbar

The zoning window displays as shown in Figure 15-13. As you can see, the **New TI Zone** is grayed out and cannot be chosen because of the Alias with WWN which is displayed on the left.

Tip: The TI zone can only be created using D,I (Domain, Index) notation.

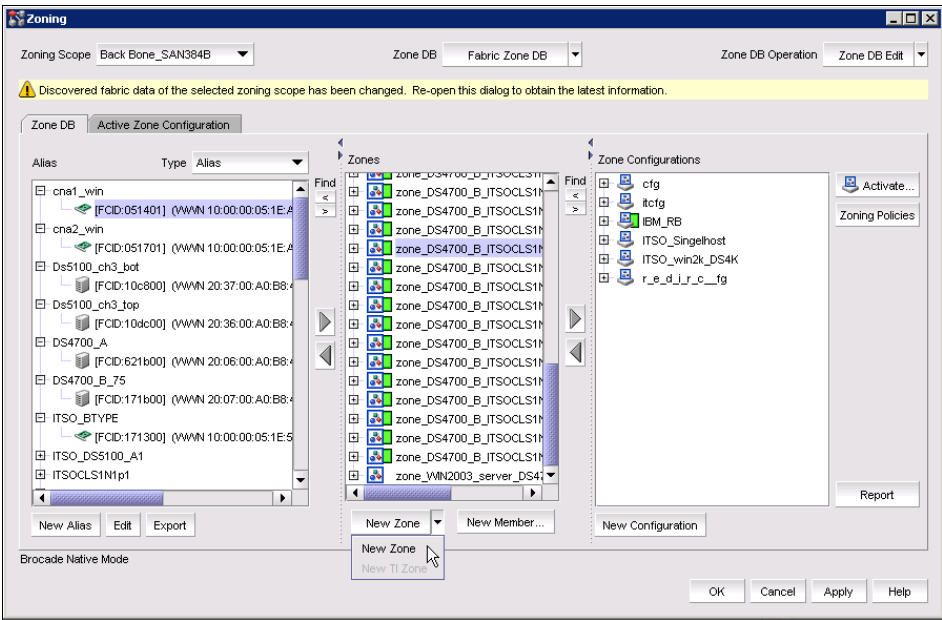


Figure 15-13 Grayed out menu option, New TI Zone

To create a TI zone, use the notation D,I (Domain, Index) in Figure 15-14.

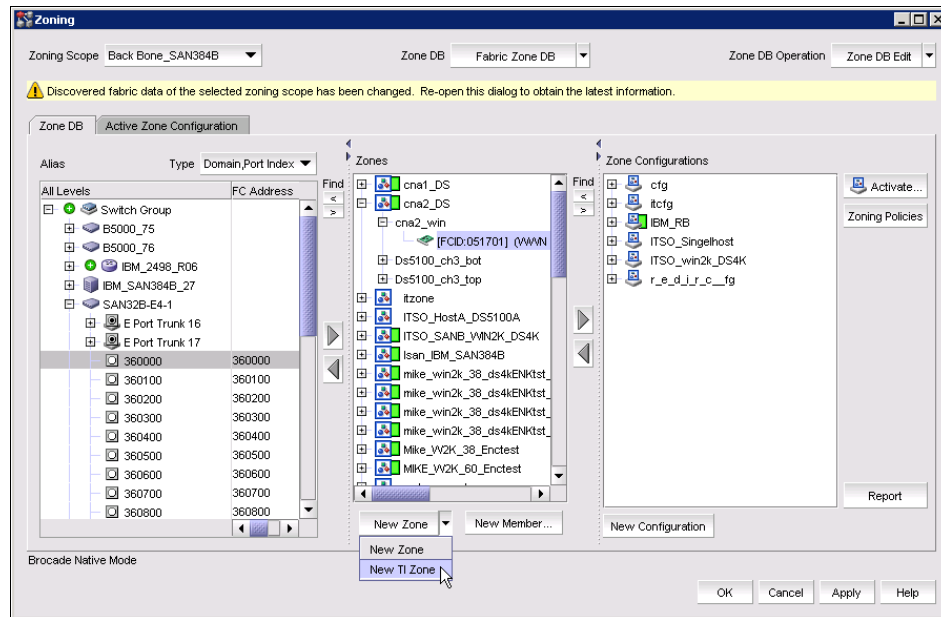


Figure 15-14 Creation of TI zone

We created a TI zone named *Test_TI_zone* based on the D,I notation (adding E_Ports and F_Ports).

A TI zone has its own set of operations and properties. By right-clicking a TI zone, the menu is displayed as in Figure 15-15.

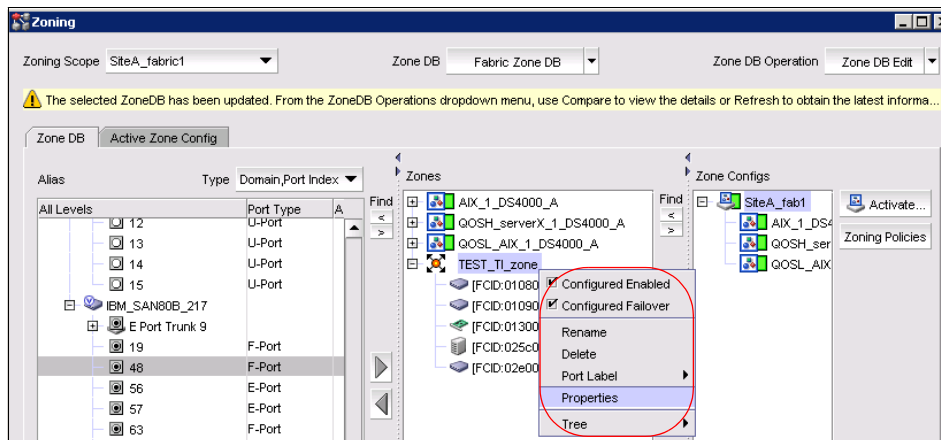


Figure 15-15 TI zone menu

The menu shows the possible choices:

- ▶ Configured Enabled
- ▶ Configured Failover
- ▶ Rename
- ▶ Delete
- ▶ Port Label (port # or port name)
- ▶ Properties
- ▶ Tree (showing options)

The properties dialog box is shown in Figure 15-16.

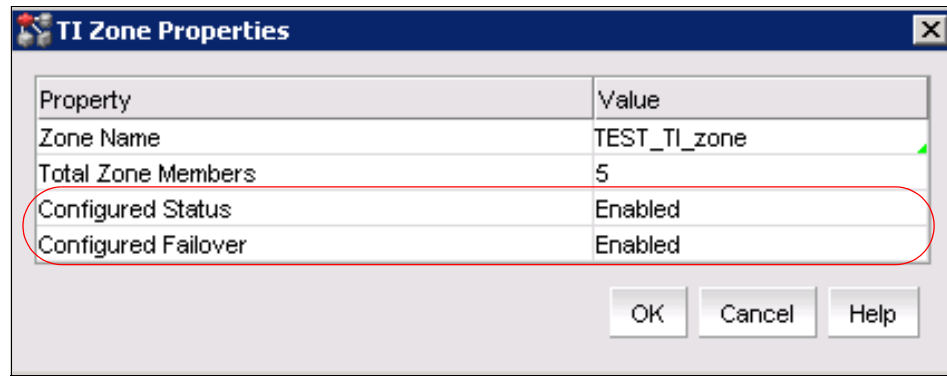


Figure 15-16 Properties Dialog Box

The TI zone cannot be added to zone config. A DCFM Message box is displayed as shown in Figure 15-17.




Figure 15-17 Adding TI zone to Zone config

15.4 QoS: SID/BID traffic prioritization

SID/DID traffic prioritization allows you to categorize the traffic flow between a given host and target as having a high or a low priority. This feature is based on the IBM/Brocade Virtual Channel Model, which is shown in Figure 15-18.

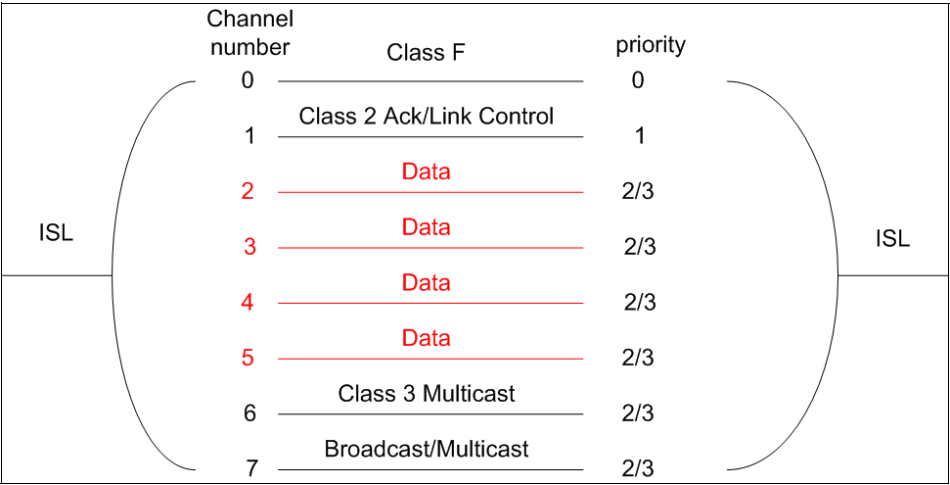


Figure 15-18 Original Virtual Channel Model(pre condor2/GoldenEye2 ASICs)

The key features of this model are as follows;

- ▶ It consists of 8 channels (from 0 to 7).
- ▶ It applies only to ISLs:
 - Each channel has its own buffer-to-buffer credits.
 - ISL can be part of the trunk group.
- ▶ There are four priority levels (0 to 3):
 - Level 0 is the highest priority.
 - Level 3 is the lowest priority.

In this channel, all data of different priorities (channel 2, 3, 4 and 5) can travel the link at the same time. In reality it means that:

- ▶ Traffic is not disrupted.
- ▶ Traffic will not disrupt other traffic.

For the new 8 Gbps platform/blades, we can also prioritize the traffic on the link by assigning to it priority levels.

The new ASIC (Condor2/GoldenEye2) has 16 Virtual Channels as shown in Figure 15-19.




Figure 15-19 New Virtual Channel Model

The new model in Figure 15-19 has these key features:

- ▶ 16 Virtual Channels
- ▶ Three priority levels (low, medium, high)
 - High priority gets 60% of total data bandwidth
 - Medium priority gets 30% of total data bandwidth
 - Low priority gets 10% of total data bandwidth

The new Virtual Channel model is used for QoS SID/DID Traffic Prioritization.

Considerations:

- ▶ The switch must run Fabric OS v6.0 or later.
- ▶ Hosts and targets must be connected to 8 Gbps capable switches.
- ▶ QoS enabled E_Port will form a QoS capable ISL with the neighboring switch only if the connecting E_Port on the neighboring switch is also QoS capable.

The data flow with the priority of high, medium or low depends on the numbers of SID/DID pairs in the Virtual Channel link.

For example, if there is a single low priority flow to a destination ID (DID) and several medium priority flows to that same DID, then it is possible that the medium priority flows would have less bandwidth.

15.4.1 QoS zones

Prioritization is accomplished by the use of QoS zones. A QoS zone is a special zone that indicates the priority of the traffic flow between a given host/target pair:

- ▶ The members of a QoS zone are WWNs of the host/target pairs.
- ▶ QoS zones can contain only WWN members.
- ▶ To distinguish the QoS zones from normal WWN zones, special prefixes are used:
 - QOSH_ is used to set high priority.
 - QOSL_ is used to set low priority.
 - The switch automatically sets the priority for the “host, target” pairs specified in the zones based on the priority level in the zone name.
- ▶ The default setting is medium priority:
 - This setting is used when no QoS zones are specified.

Figure 15-20 shows the QoS zones.




Figure 15-20 QoS zones

Assuming that you want to have a low priority from Server 1 to Storage, and high priority from Server 2 to Storage (as shown in Figure 15-20), set the QoS zones as follows:

- QOSL_Server1_Storage (with members Server1, Storage)
- QOSH_Server2_Storage (with members Server2, Storage)

Where:

QOSL_Server1_Storage, QOSH_Server2_Storage are the names of the zones.

Path selection between the *host*, *target* pairs is governed by FSPF rules, which means that switch 3 will not take part in the data flow. The considerations for switch 3 are covered in 15.4.2, “QoS E_Ports” on page 728.

QoS: QoS can be used for device pairs that exist within the same fabric only. QoS priority information is not passed over EX_ or VEX_Ports and should not be used for devices in separate fabrics.

15.4.2 QoS E_Ports

QoS zoning enables zoning between a given host/target pair that are connected to *F_Ports*.

In addition to configuring the hosts and targets in a zone, you must also enable QoS on individual *E_Ports* that might carry traffic between the given host and target pairs.

Figure 15-21 shows that two *E_Ports* are enabled for QoS traffic.




Figure 15-21 QoS with E_Ports enabled

With reference to Figure 15-20:

- ▶ You need to enable QoS on the E_Ports on ISL between switch 1 and switch 2.
- ▶ You do not need to enable QoS on the E_Ports on the ISLs between switch 1 and switch 3, and between switch 3 and switch 2, because these are not the shortest paths between the hosts and the targets.
 - However, if the ISL between switch 1 and switch 2 is broken, then the path through switch 3 would be used.

This is why you need to enable QoS on all possible E_ports (including ports 0 and 15 on switch 3, port 14 on switch 1 and port 1 on switch 2 (not shown) if you want to guarantee traffic priority.

A QoS enabled E_Port will form a QoS capable ISL with the neighboring switch only if the connecting E_Port on the neighboring switch is also QoS capable. Otherwise, the fabric module will negotiate down or up to non_QoS mode which is medium priority.

For Figure 15-21 on page 728, if the QoS will be not enabled on port 15 of switch 1 or port 0 of switch 2 (or both), the traffic will look as follows:

- ▶ Low priority from server 1 to switch 1
- ▶ High priority from server 2 to switch 1
- ▶ Medium priority from switch 1 to the target devices on Storage.

Considerations:

- ▶ If QoS is not enabled on an E_Port, the traffic prioritization stops at that point and the default of medium priority is used instead.
- ▶ You can prioritize flows between devices in a logical fabric. The rules for enabling QoS on E_Ports are the same as for physical fabrics.

15.4.3 Supported configurations and limitations

The supported configuration includes the following settings:

- ▶ Hosts and targets must be connected to 8 Gbps switches:
 - SAN24B-4
 - SAN40B-4
 - SAN80B-4
 - FC8-16, FC8-32, or FC8-48 port blade in the SAN768B and SAN384 backbone switches
- ▶ For larger fabrics with the intermediate switches:
 - All intermediate switches should be 8 Gbps capable OR
 - Must be running Fabric OS v6.0 or later
- ▶ If a host and target are included in two or more QoS zones with different priorities, the zone with the lowest priority takes precedence.
- ▶ Traffic prioritization is not supported on 10 Gbps ISLs.
- ▶ Traffic prioritization is not supported on mirrored ports.
- ▶ Traffic prioritization is enforced on the egress ports only, not on the ingress ports.

The supported configuration is shown in Figure 15-22. The circled QoS zone will have low or high priority traffic preserved across the fabric.

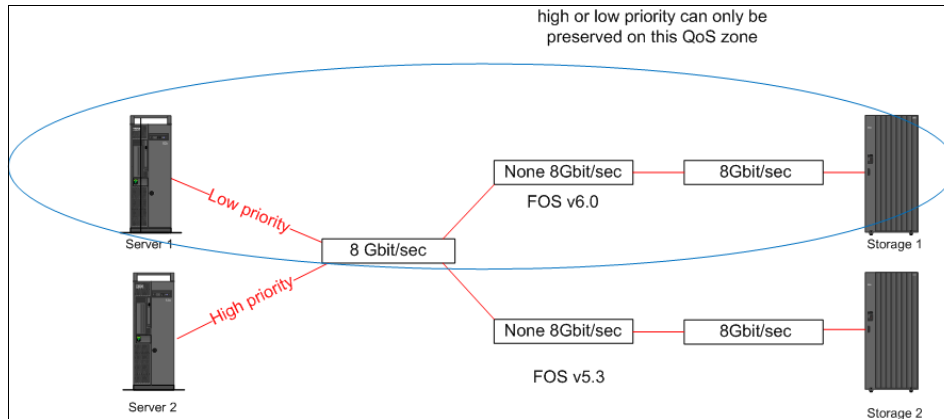


Figure 15-22 Supported configuration of QoS traffic prioritization

15.4.4 QoS with CLI

Because the QoS zones differ from normal zones only with the prefix in the zone name, use the CLI zoning commands to create QoS zones as described in Chapter 12, “Basic zoning” on page 513.

Commands

To enable QoS on a given port, use the command in Example 15-15.

Example 15-15 Enabling QoS on a given port

```
portcfgqos --enable [slot/]<port>
```

To disable QoS on a given port use the command in Example 15-16.

Example 15-16 Disabling QoS on a given port

```
portcfgqos --disable [slot/]<port>
```

To set the configuration to default, use the command in Example 15-17.

Example 15-17 Setting QoS on a given port to default

```
portcfgqos --default [slot/]<port>
```

To show the E_Port with QoS enabled, use the command in Example 15-18.

Example 15-18 Showing the port with QoS enabled

```
portcfgshow [slot/]<port>
```

Real life example

We disabled the configuration and deleted all zones as shown in Example 15-19.

Example 15-19 Zoneshow

```
IBM_SAN80B_217:FID128:admin> zoneshow
Defined configuration:
  alias: AIX_1    10:00:00:00:c9:4c:8c:1c
  alias: DS4000_A
                    20:06:00:a0:b8:48:58:a1
  alias: serverX_1
                    10:00:00:05:1e:53:10:8b
```

Effective configuration:

No Effective configuration: (No Access)

Now we are adding new zones as shown in Example 15-20.

Example 15-20 Adding QoS zones

```
IBM_SAN80B_217:FID128:admin> zonecreate
"QOSH_serverX_1_DS4000_A","10:00:00:05:1e:53:10:8b;
20:06:00:a0:b8:48:58:a1"
IBM_SAN80B_217:FID128:admin> zonecreate
"QOSL_AIX_1_DS4000_A","10:00:00:00:c9:4c:8c:1c;
20:06:00:a0:b8:48:58:a1"
IBM_SAN80B_217:FID128:admin> zoneshow
Defined configuration:
  zone:  QOSH_serverX_1_DS4000_A
                    10:00:00:05:1e:53:10:8b; 20:06:00:a0:b8:48:58:a1
  zone:  QOSL_AIX_1_DS4000_A
                    10:00:00:00:c9:4c:8c:1c; 20:06:00:a0:b8:48:58:a1
  alias: AIX_1    10:00:00:00:c9:4c:8c:1c
  alias: DS4000_A
                    20:06:00:a0:b8:48:58:a1
  alias: serverX_1
                    10:00:00:05:1e:53:10:8b
```

Effective configuration:

No Effective configuration: (No Access)

We are adding config "SiteA_fab1" with QoS zones (see Example 15-21).

Example 15-21 Adding config with Qos zones

```
IBM_SAN80B_217:FID128:admin> cfgcreate
"SiteA_fab1","QOSH_serverX_1_DS4000_A; QOSL_AIX_1_DS4000_A"
IBM_SAN80B_217:FID128:admin> zoneshow
Defined configuration:
  cfg:  SiteA_fab1
        QOSH_serverX_1_DS4000_A; QOSL_AIX_1_DS4000_A
  zone:  QOSH_serverX_1_DS4000_A
        10:00:00:05:1e:53:10:8b; 20:06:00:a0:b8:48:58:a1
  zone:  QOSL_AIX_1_DS4000_A
        10:00:00:00:c9:4c:8c:1c; 20:06:00:a0:b8:48:58:a1
  alias: AIX_1 10:00:00:00:c9:4c:8c:1c
  alias: DS4000_A
        20:06:00:a0:b8:48:58:a1
  alias: serverX_1
        10:00:00:05:1e:53:10:8b
```

Effective configuration:

No Effective configuration: (No Access)

The configuration is saved as shown in Example 15-22.

Example 15-22 Saving Configuration

```
IBM_SAN80B_217:FID128: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 ...
```

The configuration is enabled in Example 15-23.

Example 15-23 Cfgenable

```
IBM_SAN80B_217:FID128:admin> cfgenable SiteA_fab1
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 'SiteA_fab1' configuration (yes, y, no, n): [no]
y
zone config "SiteA_fab1" is in effect
Updating flash ...
IBM_SAN80B_217:FID128:admin> zoneshow
Defined configuration:
  cfg:  SiteA_fab1
        QOSH_serverX_1_DS4000_A; QOSL_AIX_1_DS4000_A
  zone:  QOSH_serverX_1_DS4000_A
        10:00:00:05:1e:53:10:8b; 20:06:00:a0:b8:48:58:a1
  zone:  QOSL_AIX_1_DS4000_A
        10:00:00:00:c9:4c:8c:1c; 20:06:00:a0:b8:48:58:a1
  alias: AIX_1    10:00:00:00:c9:4c:8c:1c
  alias: DS4000_A
        20:06:00:a0:b8:48:58:a1
  alias: serverX_1
        10:00:00:05:1e:53:10:8b

Effective configuration:
  cfg:  SiteA_fab1
  zone:  QOSH_serverX_1_DS4000_A
        10:00:00:05:1e:53:10:8b
        20:06:00:a0:b8:48:58:a1
  zone:  QOSL_AIX_1_DS4000_A
        10:00:00:00:c9:4c:8c:1c
        20:06:00:a0:b8:48:58:a1
```

15.4.5 Web Tools and QoS Zones

You can check the status and enable/disable QoS on the port as shown in Figure 15-23.

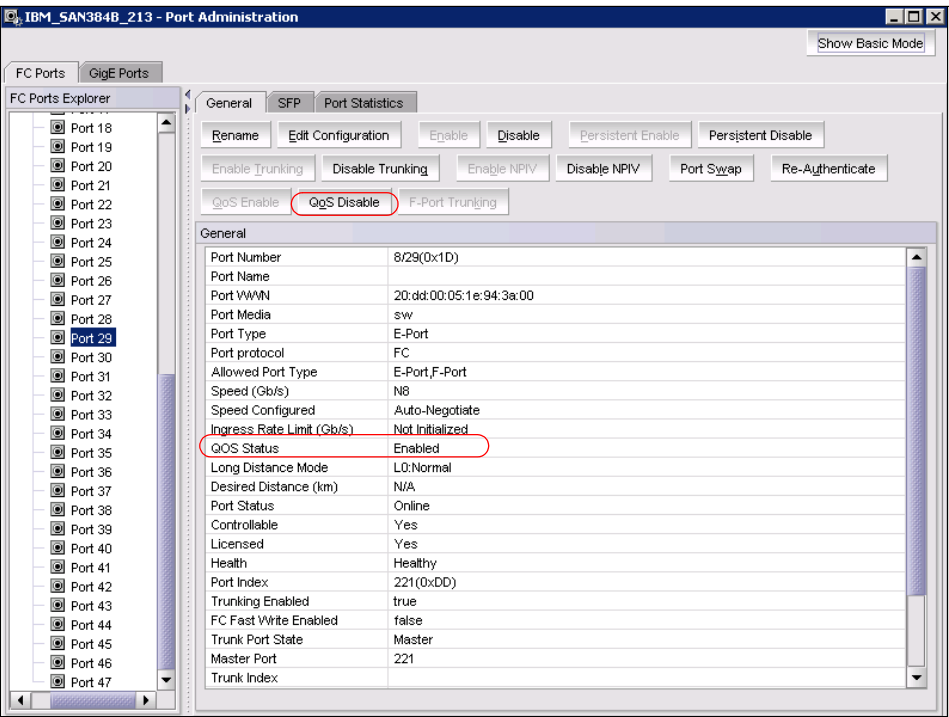


Figure 15-23 Enabling QoS on the port

To create QoS zones, choose **Manage** → **Zone admin** and follow the rules for zone administration (see Figure 15-24). Zone administration is covered in Chapter 12, “Basic zoning” on page 513.

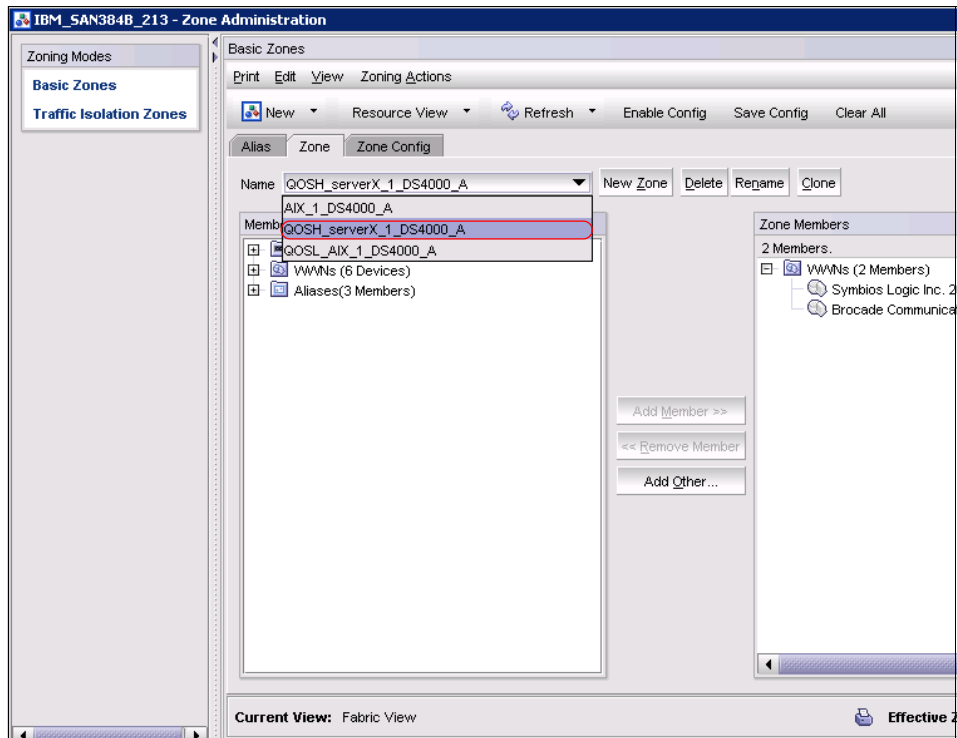


Figure 15-24 Creating QoS zones with Web Tools

15.4.6 DCFM and QoS zones

DCFm provides convenient access to change the priority of the QoS zone, which can be set with a right-click menu option, as shown in Figure 15-25.

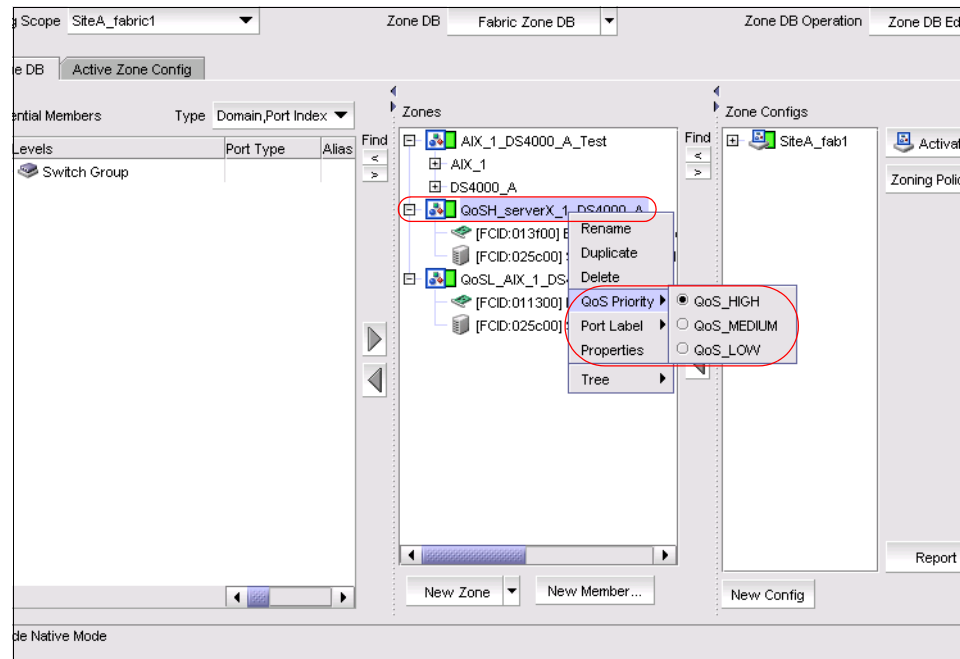


Figure 15-25 QoS priority With DCFM

You can change the priority of the normal zone. In this case, DCFM adds the prefix **QOSX** to the zone name, where **X** denotes the priority:

- ▶ **H** for High priority
- ▶ **L** for Low priority

As you can see in Figure 15-26, we changed the priority to *Low* for the zone **AIX_1_DS4000_A_Test**. The zone name was changed to **QoSLAIX_1_DS4000_A_Test**.

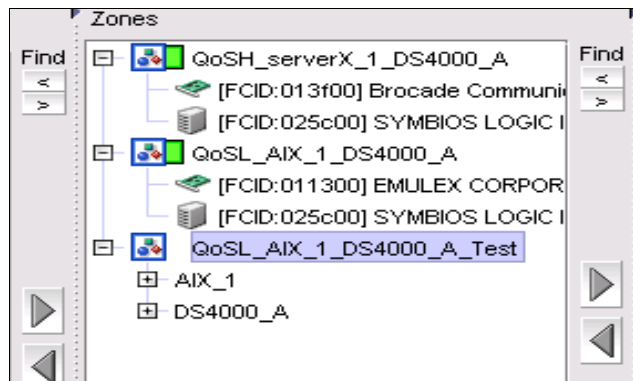


Figure 15-26 Changing priority to low for normal zone “AIX_1_DS4000_A_Test”



Performance monitoring

Basic performance monitoring tasks are available on all IBM/Brocade SAN products. These tasks include basic statistics, such as port throughput, switch throughput, utilization, and error count. With the purchase of the Advanced Performance Monitoring licensed feature, you can perform more advanced tasks, such as:

- ▶ End-to-end monitoring
- ▶ Filter-based monitoring
- ▶ ISL monitoring
- ▶ Top Talkers monitoring
- ▶ Tracking the SCSI commands rate
- ▶ SCSI versus IP traffic statistics

You can administer performance monitoring through either Telnet commands, Web Tools, or DCFM. However, many of the advanced monitoring tasks are only available through the command-line interface CLI (Telnet).

In this chapter we provide an overview of performance monitoring and how to use Basic and Advanced Performance Monitoring features.

16.1 Performance monitoring with Web Tools

The Web Tools Performance Monitor task provides real-time information about basic performance parameters, such as switch and port throughput, switch utilization, and port error rate. The Advanced Performance Monitor licensed feature adds additional monitoring capabilities in Web Tools, such as SID/DID throughput, SCSI read and write statistics, SCSI versus IP traffic, and AL_PA error rate.

To access Performance Monitor, click the **Performance Monitor** task in the Tasks panel (see Figure 16-1).




Figure 16-1 Performance Monitor task

The Performance Monitoring window opens, as shown in Figure 16-2.

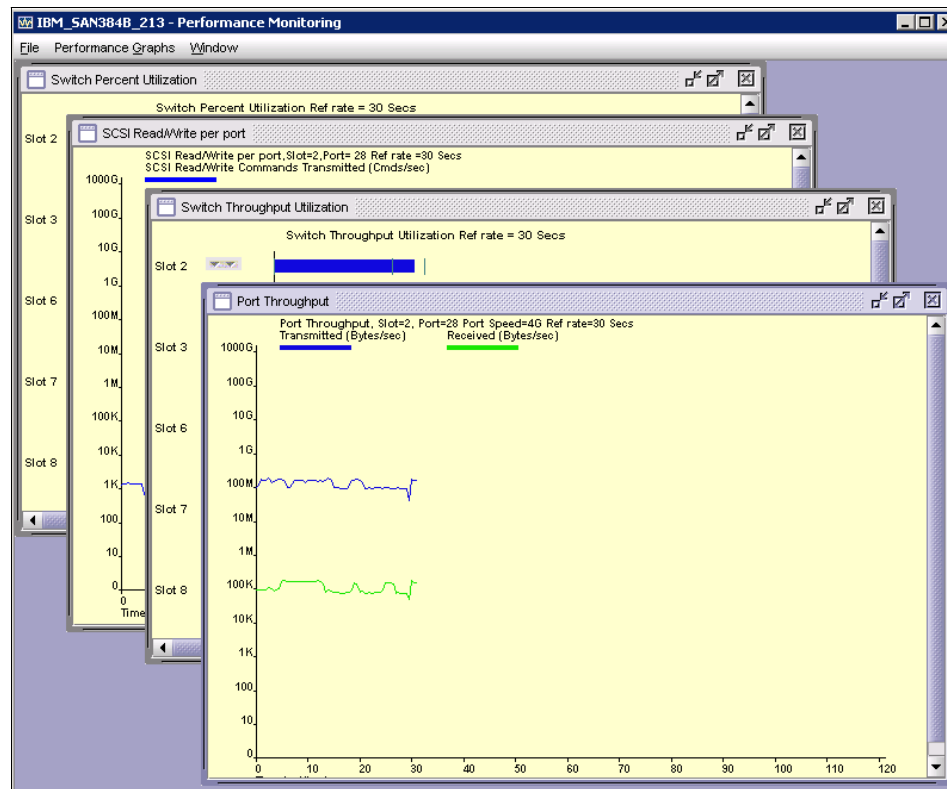


Figure 16-2 Performance Monitoring window

When the window opens, the Switch Throughput Utilization graph displays on the canvas. You can add the performance monitors in which you are interested. In Figure 16-3, several performance graphs are added. The canvas holds a maximum of eight graphs. All the graphs show real-time information and are updated every 30 seconds.

You can print the graph by right-clicking inside it and selecting **Print** on the pop-up menu (see Figure 16-3).




Figure 16-3 Printing Canvas

When you add the graphs that you want to monitor, it is possible to save the current layout of canvas to the switch by selecting **File** → **Save Current Canvas Configuration** on the menu bar. Provide the name and description for canvas configuration (Figure 16-4), and then click **Save Canvas**.

The figure shows a 'Save Canvas Configuration' dialog box. It has a title bar with a close button. Inside, there are two text input fields. The first is labeled 'Name' and contains the text 'IBM_SAN384B'. The second is labeled 'Description' and contains the text 'IBM_SAN384B_30.07.2009'. At the bottom right, there are two buttons: 'Save Canvas' and 'Cancel'.

Figure 16-4 Save canvas configuration

The canvas configuration is saved to the switch.

You can load the saved canvas configurations by selecting **File → Display Canvas Configurations** on the menu bar. The **Canvas Configuration List** opens, as shown in Figure 16-5.




Figure 16-5 Canvas Configuration List

Select one of the saved canvas configurations, then click **Load**. The graphs display on the canvas.

Apart from loading the canvas configuration, you can also **Edit** or **Copy** it, as shown in Figure 16-6.

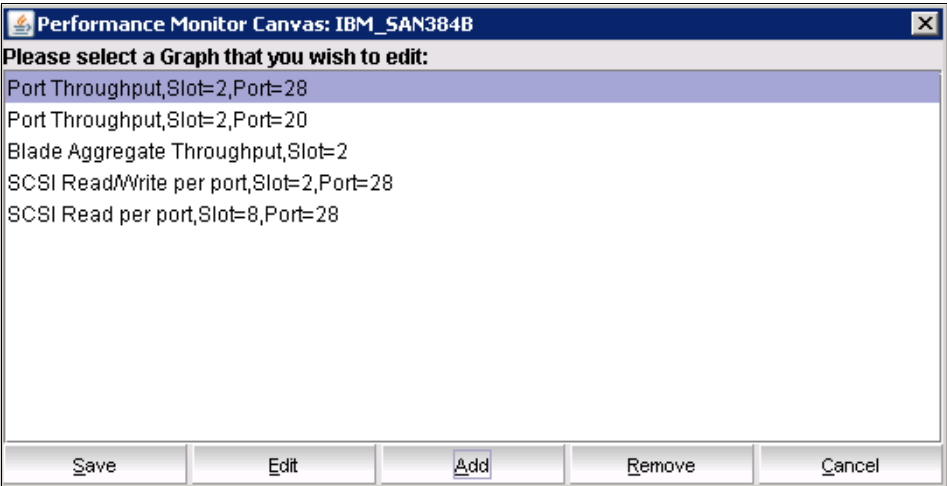


Figure 16-6 Edit canvas configuration

16.2 Basic Performance Monitoring

With Basic Performance Monitoring, you can measure the parameters according to Table 16-1.

Table 16-1 Basic performance monitoring parameters

Basic performance	Description
Port Throughput	Displays the performance of a port in bytes/second for frames received and transmitted
Switch Aggregate Throughput	Displays the aggregate performance of all ports on a switch
Switch Throughput Utilization	Displays the port throughput at the time the sample is taken
Port Error	Displays a line of CRC errors for a given port
Switch Percent Utilization	Displays the percentage of usage of a chosen switch at the time the sample is taken
Port Snapshot™ Error	Displays the CRC error count between sampling periods for all the ports on a switch

16.2.1 Basic Performance Monitoring with Web Tools

The Basic Performance Monitors are standard in Web Tools and do not require any additional license. You can access these monitors by clicking **Performance Graphs** → **Basic Monitoring** on the menu bar (Figure 16-7).

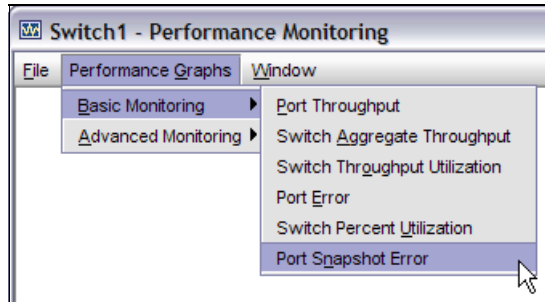


Figure 16-7 Basic performance monitors

The following graphs are available:

- ▶ Port Throughput
- ▶ Switch Aggregate Throughput
- ▶ Switch Throughput Utilization
- ▶ Port Error
- ▶ Switch Percent Utilization
- ▶ Port Snapshot Error

The SAN768B, SAN384B backbones and SAN256B director contain one additional option: the Blade Aggregate Throughput graph.

We explain these graphs in the following sections.

16.2.2 Throughput examples

In this section we describe several types of throughput.

Port Throughput

For the Port Throughput graph, you first need to specify the port to monitor (see Figure 16-8). Either enter the port number in the field, or drag and drop it from the port selection list on the left. Then, click **OK** to continue.

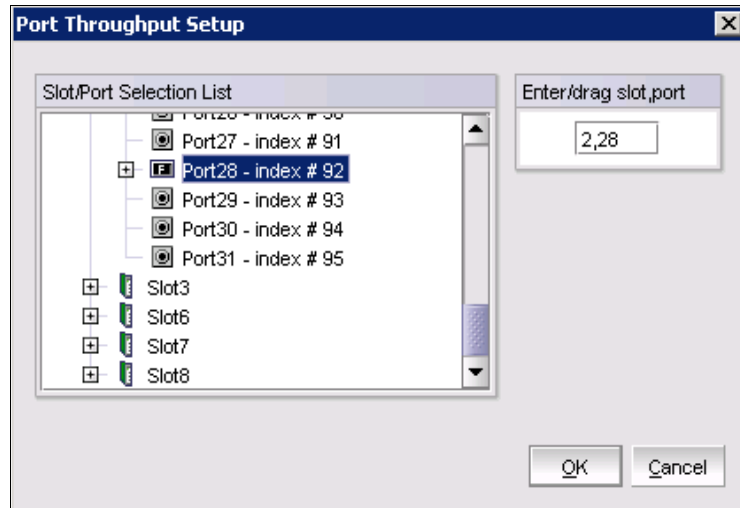


Figure 16-8 Port Throughput Setup

The Port Throughput graph displays on the canvas, as shown in Figure 16-9.



Figure 16-9 Port Throughput graph

We are copying files to the storage and the port 2/28 is connected to storage. As you can see, the rate of Bytes Transmitted are much bigger than the rate of Bytes Received which is correct.

Blade Aggregate Throughput

The Blade Aggregate Throughput graph is only available on the SAN768B, SAN384B, and SAN256B. It allows you to display a graph of total throughput for a certain blade. To use this graph, you need to select the blade, and then the graph displays on the canvas as shown in Figure 16-10.




Figure 16-10 Blade Aggregate Throughput for slot 2

We still copy the file to storage and the blade is on the core switch, which is connected to storage. There is no difference in Bytes Transmitted and Bytes Received, which is also correct (comparing to port throughput). The explanation is as follows:

- We have two switches in the fabric:
 - Edge switch with the servers connected
 - Core switch with the storage connected

- We have two trunked connections between two switches — but only one of the connection goes to slot number 2.
 - This is why we are receiving on blade 2, 50 MBps from our servers.
 - The other 50 MBps is received by blade number 8 (see Figure 16-11).




Figure 16-11 Blade Aggregate Throughput for slot 8

The storage is connected only to slot 2, which causes the following results:

- In Figure 16-11 you can see:
 - Transmit rate to Storage — 100 MB
 - Receive rate from Servers — 50 MB
- You can observe that the receive rate from servers is 50 MB — but all the traffic from slot 8 is routed to storage connected to slot 2.

Keep in mind that this is the overall Blade Throughput so it counts *all* the traffic to/from a particular blade. This is why graphs for the total blade throughput differ from port throughput.

We describe this situation in further detail in 16.3, “Advanced Performance Monitoring” on page 753.

Switch Aggregate Throughput

The Switch Aggregate Throughput graph shows the real-time total throughput on all switch ports, as shown in Figure 16-12.




Figure 16-12 Switch Aggregate Throughput graph

The RX and TX values are the same because they are gathered for the overall switch, and are not split between blades in slot 2 and 8.

Switch Throughput Utilization

Figure 16-13 shows an example of the Switch Throughput Utilization graph on a SAN32B-3. This graph shows the throughput on each switch port at the time the sample is taken. Because the SAN256B and the SAN768B can have a very large number of ports, the graph displays the throughput for each slot on these two products.




Figure 16-13 Switch Throughput Utilization graph

You can customize the graph by right-clicking inside it and choosing **Select Ports**. A window displays (shown in Figure 16-14) that allows you to select the ports that you want to see on the graph.

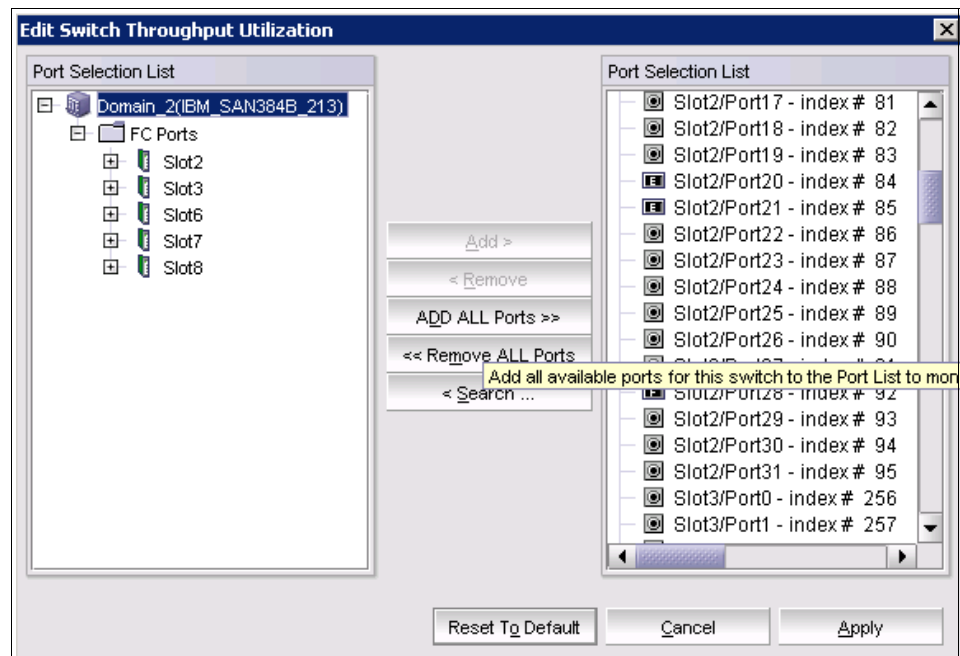


Figure 16-14 Edit Switch Throughput Utilization

Port Error

The Port Error graph shows the number of CRC errors for the selected port. To use the graph, you must select the port. Then, the actual graph displays. You can use the graph to detect and troubleshoot ports that are not performing up to expectations.

Switch Percent Utilization

The Switch Percent Utilization graph looks similar to the Switch Throughput Utilization graph, but the utilization information is displayed in percentages. You can select the ports to be displayed on the graph by right-clicking in the graph, then choosing **Select Ports**.

Figure 16-15 shows an example of the Switch Percent Utilization graph.




Figure 16-15 Switch Percent Utilization graph

Port Snapshot Error

The Port Snapshot Error graph shows the number of CRC errors that have occurred in the last sampling period for all the ports. On the SAN256B, SAN768B, and SAN384B, the number of CRC errors per slot is displayed. As with other similar graphs, you can customize the ports that you want to see displayed.

Figure 16-16 shows the Port Snapshot Error graph for SAN384B.




Figure 16-16 Port Snapshot Error graph for SAN384B

16.3 Advanced Performance Monitoring

Advanced Performance Monitoring is the licensed feature that provides comprehensive tools for monitoring the performance of the networked storage resources. This feature provides the following benefits:

- ▶ Supports direct-attach, loop, and switched FC SAN topologies
- ▶ Monitors transaction performance from source to destination (end-to-end monitoring)
- ▶ Reports Cyclic Redundancy Check (CRC) error measurement statistics
- ▶ Measures ISL performance and resource usage

16.3.1 Virtual Fabrics considerations

Attention: For Virtual Fabrics, each logical switch can have its own set of performance monitors. The installation of monitors is restricted to the ports that are present in the respective logical switch.

The number of logical switches that can be configured with monitors is also restricted, as we can see in Table 16-2.

Table 16-2 Number of logical switch that supports performance monitors

Platform	Maximum number of logical switches supported	Maximum number of logical switches on which monitors are supported
SAN768B	8	4
SAN384B	8	4
SAN40B-4	3	3
SAN80B-4	4	3

16.3.2 Performance Monitors

Advanced Performance Monitoring provides the following monitors:

- ▶ End-to-End monitors measure the traffic between a host/target pair.
- ▶ Filter-based monitors measure the traffic transmitted through a port with specific values in the first 64 bytes of the frame.
- ▶ ISL monitors measure the traffic transmitted through an Inter Switch Link (ISL) to different destination domains.
- ▶ Top Talkers monitors measure the flows that are major consumers of bandwidth on a switch or port.

The type of monitors supported depends on the ASIC. Table 16-3 shows the monitors supported on different switches.

Table 16-3 Monitors to Product support table

Product	EE	FILTER	ISL	Top Talker
SAN24B-4,SAN80B-4	YES	YES	NO	YES
SAN256B	YES	YES	YES	YES
SAN40B-4, SAN768B, SAN384B	YES	YES	NO	YES

We describe each of these monitors in the following sections.

16.3.3 Displaying Performance Monitors with the CLI

The command **perfmonitorshow** with the syntax shown in Example 16-1 displays end-to-end (EE), filter-based (FLT), and inter switch link (ISL) performance monitors on a port.

Example 16-1 perfmonitorshow command

```
perfmonitorshow --class monitor_class [slotnumber/]portnumber  
[interval]
```

We refer to this command when we describe each of the performance monitors.

16.3.4 SID/DID Performance Monitor

End-to-end performance monitoring counts the number of words in Fibre Channel frames for a specified Source ID (SID) and Destination ID (DID) pair. An end-to-end performance monitor includes these counts:

- ▶ RX_COUNT (words in frames received at the port)
- ▶ TX_COUNT (words in frames transmitted from the port)

An end-to-end monitor must be configured on the specific port, specifying the SID-DID pair (in hexadecimal).

Frames: The monitor counts only those frames with matching SID and DID.

Each SID or DID has the following three fields:

- ▶ Domain ID (DD)
- ▶ Area ID (AA)
- ▶ AL_PA (PP)

For example, the SID 0x213000 denotes: DD=0x21, AA=0x30, PP=0x00.

Traffic: End-to-end performance monitoring looks at traffic on SID/DID pairs in any direction.

In Figure 16-17 we show where to add end-to-end monitors on a port.




Figure 16-17 Setting end-to-end monitors on a port

The rules are as follows:

- ▶ For direction of traffic on SID/DID pairs:
 - Traffic on SID/DID pairs can be in any direction.
 - If the traffic is monitored in both directions, the Tx/Rx counters are reversed.
- ▶ When monitoring the traffic from Server to Storage:
 - Add monitor 0 on port 0 on switch 1.
 - Specify SID as 0x021300 and DID as 0x015c00.
 - For monitor 0, RX_COUNT is the number of words from Server to Storage, and TX_COUNT is the number of words from Storage to Server.

Word counts: It can be confusing as to why RX_COUNT is the number of words from Server to Storage. RX_COUNT means received words. But we are sending the words from Server to Storage. When you look from the point of view of port 0, it becomes clear. Port 0 is receiving the words from the server, and port 0 is the one from which you need to look at your traffic.

- ▶ When monitoring the traffic from Storage to Server:
 - Add monitor 1 on port 15 on switch 2.
 - Specify SID as 0x015c00 and DID as 0x021300.
 - For monitor 1, RX_COUNT is the number of words from Storage to Server, and TX_COUNT is the number of words from Server to Storage.

In Figure 16-18 we show the RX/TX dependency for the Server to Storage monitor.




Figure 16-18 RX/TX dependency for Server to Storage monitor

We have monitor 0 established for port 0 switch 1 for Server to Storage traffic as shown in Figure 16-18:

- ▶ The Server transmits frames.
- ▶ Port 0 receives frames (RX_COUNT)
 - For frames received at the port with the end-to-end monitor installed, the frame SID is the same as “SourceID” and the frame DID is the same as “DestID”.
 - The RX_COUNT is updated accordingly.
- ▶ Port 0 transmits frames(TX_COUNT)
 - For frames transmitted from the port with the end-to-end monitor installed, the frame DID is the same as “SourceID” and the frame SID is the same as “DestID”.
 - The TX_COUNT updated accordingly.

For our simple scenario as shown in Figure 16-18, when we copy one large file from Server to Storage with Monitor 0 enabled, we can observe the following numbers:

- ▶ RX_COUNT = 100 MB: Sending file from Server to Storage.
- ▶ TX_COUNT = 100 KB: Receiving confirmations from Storage to Server.

You can see the graphs of the previous example in “SID/DID Performance using Web Tools” on page 758.

If we have more than two switches, Figure 16-19 shows the correct placement of the end-to-end performance monitors.




Figure 16-19 Placement of the end-to-end monitors

As we can see, the rules are the same as for the two-switch configuration.

SID/DID Performance using Web Tools

Prior to displaying the actual graph, the SID/DID Performance Setup window opens, as shown in Figure 16-20.

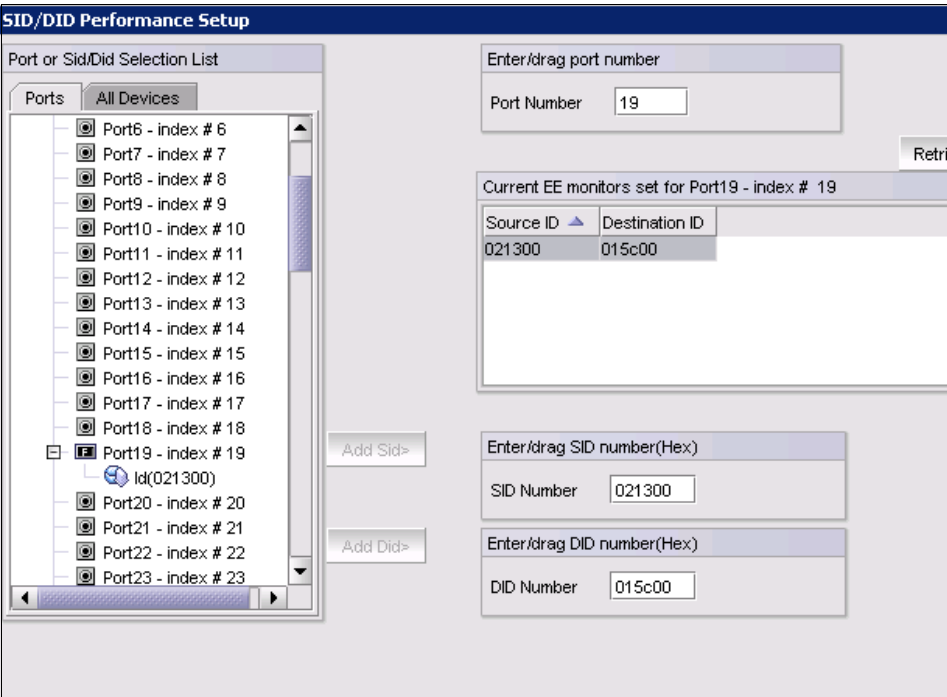


Figure 16-20 SID/DID Performance Setup window

Complete the Port, SID, and DID number fields as appropriate by either entering the values or dragging them from the selection list on the left. Then, click **OK**.

Alternatively, you can display a list of currently configured end-to-end (EE) monitors for the selected port by clicking **Retrieve preset EE monitors**. You can use the list of current EE monitors to select one or more EE SID/DID pairs. If you select multiple pairs, then a separate graph displays for each pair.

Next, in Figure 16-21, you can see the implementation of the scenario shown in Figure 16-18 on page 757, when we copy one large file from Server to Storage, having Monitor enabled as shown in Figure 16-20 on page 758.

The peaks on the graph Figure 16-21 show the time when the file is transmitted from Server to Storage.




Figure 16-21 SID/DID Performance graph (Server to Storage monitor)

Figure 16-22 shows that TX/RX counters are reversed because we measure the traffic from Storage to Server on the storage site FC switch and we exchanged SID/DID ports. The command to do this is shown in Example 16-3 on page 761.



Figure 16-22 SID/DID Performance graph (Storage to Server monitor)

SID/DID using the CLI

Adding SID/DID monitors

To add the Performance Monitor, do the following steps:

1. Connect to the switch and log in as admin.
2. Monitor the traffic from Server to Storage as shown in Example 16-2.

Example 16-2 Adding SID/DID monitor from Server to Storage

```
IBM_SAN80B_217:FID128:admin> perfaddeemonitor 19 0x021300 0x015c00
End-to-End monitor number 0 added.
IBM_SAN80B_217:FID128:admin> perfmonitorshow --class EE 19
```

There are 1 end-to-end monitor(s) defined on port 19.

```
KEY SID DID OWNER_APP TX_COUNT RX_COUNT OWNER_IP_ADDR
-----
0 0x021300 0x015c00 TELNET 0x0000000000000001b 0x00000000000000423 N/A
```

- To monitor the traffic from Storage to Server enter the command as shown in Example 16-3.

Example 16-3 Adding SID/DID monitor from Storage to Server

```
IBM_SAN384B_213:FID128:admin> perfaddeemonitor 2/28 0x015c00 0x021300
End-to-End monitor number 0 added.
```

```
IBM_SAN384B_213:FID128:admin> perfmonitorshow --class EE 2/28
```

There are 1 end-to-end monitor(s) defined on port 92.

```
KEY SID DID OWNER_APP TX_COUNT RX_COUNT OWNER_IP_ADDR
-----
0 0x015c00 0x021300 TELNET 0x00000000000000429 0x0000000000000001b N/A
```

Monitor: The monitor must be placed properly for the **perfAddEEMonitor** command to work successfully.

Deleting SID/DID monitors

We show how to delete the monitor in Example 16-4.

Example 16-4 Deleting SID/DID monitor

```
IBM_SAN80B_217:FID128:admin> perfdeleemonitor 19 0
End-to-End monitor number 0 deleted
```

16.3.5 End-to-end monitoring with DCFM

End-to-end monitors persist in the database and are enabled on one of the F_ports on the connected switch.

You can use these monitors to view both real time and historical performance data.

Licenses: Both the initiator switch and the target switch must have Performance Monitor licenses configured to create an end-to-end monitor.

To establish End-to-end monitors, choose **Select Monitor** → **Performance** → **End-to-End Monitors**.

The Set End-to-End Monitors dialog box displays (see Figure 16-23).

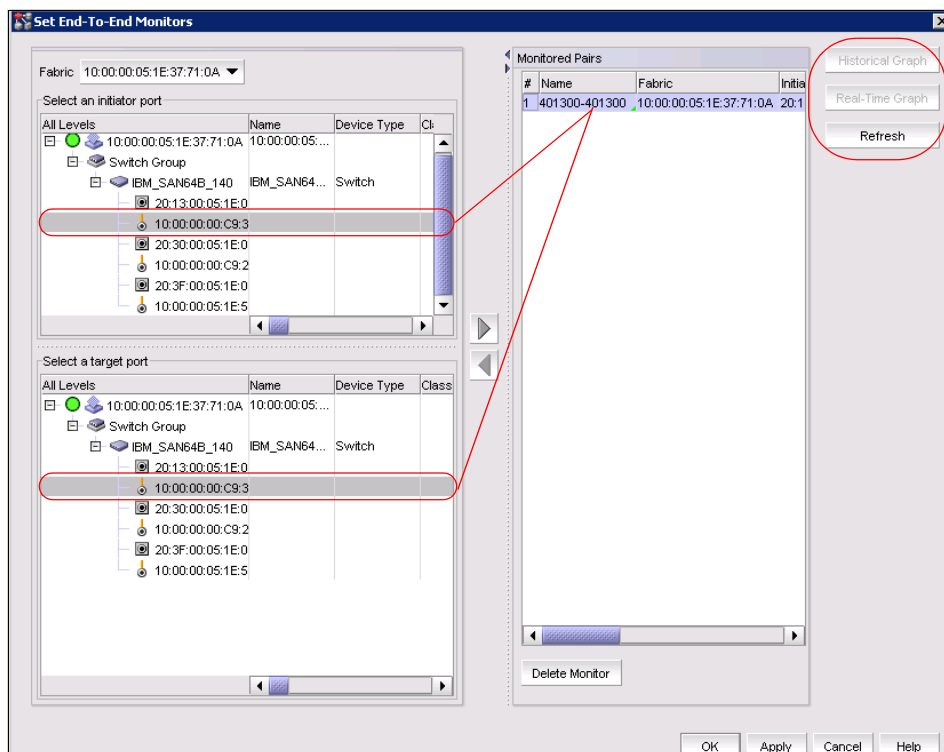


Figure 16-23 End-to-End Monitors dialog

An End-to-end Monitor consists of two pairs (see rounded rectangles in Figure 16-23).

You can select either an Initiator or Target and DCFM will automatically search for its counterpart.

You can display end-to-end monitors pairs in a real-time graph, historical time graph, and refresh them (see the rounded rectangles in the right corner in Figure 16-23). If you deleted end-to-end monitors with CLI or Web Tools, you can rewrite them back to the switch by using the **Refresh** button.

16.3.6 Filter-based performance monitoring

Filter-based monitoring allows you to count the number of frames with a particular pattern on a certain port. There is a set of predefined standard filters, and you can also configure custom-defined filters. Examples of standard filter-based monitors are the SCSI read or write commands count and the IP traffic frames count.

The maximum number of filters on most IBM/Brocade SAN products is 12 per port, in any combination of standard filters and user-defined filters. Some exceptions include certain entry-level SAN switches (where the maximum is eight filters per port) and the FC4-48 port blade on SAN256B. FC4-48 supports the following maximum values:

- ▶ Ports 0 through 15 support a maximum of 12 filter monitors per port, and 15 offsets per port for used defined monitors.
- ▶ Ports 16 through 31 have a maximum of 6 filter monitors per port, and 11 offsets per port for used defined monitors.
- ▶ Ports 32 through 47 do not support filter monitors.

Notes:

- ▶ For trunked ports, the filter is configured on the trunk master.
- ▶ For Virtual Fabrics, filter-based monitors are not supported on logical ISLs (LISLs), but are supported on ISLs and extended ISLs (XISLs).

Adding standard filter-based monitors using CLI

This section describes how to add standard filter-based monitors to a port. Use the Telnet commands listed in Table 16-4 to define filter-based monitors on a port.

Table 16-4 Add filter-based monitor commands

Command	Description
perfAddReadMonitor	Count the number of SCSI Read commands
perfAddWriteMonitor	Count the number of SCSI Write commands
perfAddRWMonitor	Count the number of SCSI Read and Write commands
perfAddSCSIMonitor	Count the number of SCSI traffic frames
perfAddIPMonitor	Count the number of IP traffic frames

Example 16-5 adds several filter monitors to port 19 on the switch SAN80B.

Example 16-5 Adding filter monitors to a port on switch SAN80B

```
IBM_SAN80B_217:FID128:admin> perfaddreadmonitor 19
SCSI Read filter monitor #0 added
IBM_SAN80B_217:FID128:admin> perfaddwritemonitor 19
SCSI Write filter monitor #1 added
IBM_SAN80B_217:FID128:admin> perfaddrwmonitor 19
SCSI Read/Write filter monitor #2 added
IBM_SAN80B_217:FID128:admin> perfaddscsimonitor 19
```

```
SCSI traffic frame monitor #3 added
IBM_SAN80B_217:FID128:admin> perfaddipmonitor 19
IP traffic frame monitor #4 added
```

In Example 16-6 we add several filter monitors to port 19 on the switch SAN384B.

Example 16-6 Adding filter monitors to a port on switch SAN384B

```
IBM_SAN384B_213:FID128:admin> perfaddreadmonitor 2/28
SCSI Read filter monitor #0 added
IBM_SAN384B_213:FID128:admin> perfaddwritemonitor 2/28
SCSI Write filter monitor #1 added
IBM_SAN384B_213:FID128:admin> perfaddrwmonitor 2/28
SCSI Read/Write filter monitor #2 added
IBM_SAN384B_213:FID128:admin> perfaddscsimonitor 2/28
SCSI traffic frame monitor #3 added
IBM_SAN384B_213:FID128:admin> perfaddipmonitor 2/28
IP traffic frame monitor #4 added
```

Displaying filter-based monitors

Use the **perfMonitorShow --class FLT** command to see the list of configured filters on a switch port, as shown in Example 16-7 through Example 16-10.

Example 16-7 Displaying a list of filter-based monitors on switch SAN80B

```
IBM_SAN80B_217:FID128:admin> perfMonitorShow --class FLT 19
```

There are 5 filter-based monitors defined on port 19.

KEY	ALIAS	OWNER_APP	FRAME_COUNT	OWNER_IP_ADDR
0	SCSIRead	TELNET	0x0000000000000000	N/A
1	SCSIWrite	TELNET	0x0000000000000000	N/A
2	SCSIR/W	TELNET	0x0000000000000000	N/A
3	SCSIFrame	TELNET	0x0000000000008ed45	N/A
4	IPFrame	TELNET	0x0000000000007fa9c	N/A

Example 16-8 Displaying a list of filter-based monitors on switch SAN80B with 5 second interval

```
IBM_SAN80B_217:FID128:admin> perfMonitorShow --class FLT 19 5
```

Showing filter monitors 19, 5

0	1	2	3	4
#Frames	#Frames	#Frames	#Frames	#Frames
0	0	0	0	0
0	0	0	12k	12k
0	0	0	12k	12k
0	0	0	12k	12k
0	0	0	12k	12k
0	0	0	12k	12k
0	0	0	12k	12k
0	0	0	13k	13k
0	0	0	12k	13k
0	0	0	12k	13k
0	0	0	12k	12k

Example 16-9 Displaying a list of filter-based monitors on switch SAN384B

```
IBM_SAN384B_213:FID128:admin> perfMonitorShow --class FLT 2/28
```

There are 5 filter-based monitors defined on port 92.

KEY	ALIAS	OWNER_APP	FRAME_COUNT	OWNER_IP_ADDR
0	SCSIRead	TELNET	0x00000000000026c31	N/A
1	SCSIWrite	TELNET	0x00000000000024f18	N/A
2	SCSIR/W	TELNET	0x00000000000022c08	N/A
3	SCSIFrame	TELNET	0x00000000000836664	N/A
4	IPFrame	TELNET	0x000000000007eef7a	N/A

Example 16-10 Displaying a list of filter-based monitors on switch SAN384B with 5 second interval

```
IBM_SAN384B_213:FID128:admin> perfMonitorShow --class FLT 2/28 5
```

Showing filter monitors 92, 5

0	1	2	3	4
#Frames	#Frames	#Frames	#Frames	#Frames
0	0	0	0	0
6.3k	6.3k	6.3k	202k	208k
6.5k	6.5k	6.5k	210k	217k
6.4k	6.4k	6.4k	207k	214k
6.2k	6.2k	6.2k	201k	207k
6.4k	6.4k	6.4k	206k	213k

6.4k	6.4k	6.4k	206k	212k
6.9k	6.9k	6.9k	223k	230k
6.5k	6.5k	6.5k	210k	217k
6.4k	6.4k	6.4k	207k	214k
6.4k	6.4k	6.4k	207k	214k

Custom filter-based monitors

In addition to the standard filters (SCSI read or write, SCSI, or IP frame count), you can create custom filters to qualify frames for statistics gathering to fit your own special requirements.

When using the custom filter-based monitors, you need to have knowledge of the FC frame structure because you must specify a series of *offsets, masks, and values*. We show the FC frame in Figure 16-24.




Figure 16-24 FC frame

The rules are as follows:

- ▶ The value of the offset must be between 0 and 63, in decimal format.
- ▶ Byte 0 indicates the first byte of the Start of Frame (SOF),
- ▶ Byte 4 is the first byte of the frame header,
- ▶ Byte 28 is the first byte of the payload (Data).

Only the following items can be selected as part of a filter definition:

- ▶ SOF
- ▶ Frame header
- ▶ First 36 bytes of payload (Data)

The following actions are performed by switch for all transmitted frames:

1. The byte in the frame at the specified offset is located.
2. The mask is applied to the byte found in the frame.

3. The value is compared with the given values in the **perfAddUserMonitor** command.
4. If a match is found, the filter counter is incremented.

On most IBM/Brocade SAN products, you can specify up to 15 different offsets for each port, and up to four values to compare against each offset. Certain entry level switches support up to seven different offsets per port.

Adding custom filter-based monitors using the CLI

Follow these steps:

1. Connect to the switch and log in as admin.
2. Enter the **perfaddusermonitor** command as shown in Example 16-11.

Example 16-11 perfaddusermonitor command syntax

```
perfaddusermonitor [slotnumber/]portnumber "group1ist" [ alias ]
```

Where:

- ▶ Slotnumber for bladed systems only; specifies the slot number of the port.
- ▶ Portnumber specifies the port number.
- ▶ Group1ist specifies up to six sets of offset, mask, and value.
- ▶ Alias is an optional name for the monitor.

Example 16-12 shows how to add a filter-based monitor for SOFi3 to a port.

Example 16-12 Adding a filter-based monitor

```
switch:admin> perfaddusermonitor 1/4 "0, 0xff, 6"  
User monitor #1 added
```

Where:

- ▶ 1/4 - slot/port number
- ▶ 0 - offset
- ▶ 0xff - mask
- ▶ 6 - value

The predefined values for SOf for Example 16-12 (offset 0) are described in Table 16-5.

Table 16-5 Predefined values at offset 0

Value	Start Of Frame (SOF)	Description
0	SOFf	SOF Fabric
1	SOFc1	SOF Connect Class 1
2	SOFi1	SOF Initiate Class 1
3	SOFn1	SOF Normal Class 1
4	SOFi2	SOF initiate Class 2
5	SOFn2	SOF Normal Class 2
6	SOFi3	SOF Initiate Class 3
7	SOFn3	SOF Normal Class 3

Deleting filter-based monitors

To delete a filter-based monitor, first list the valid monitor numbers using the **perfMonitorShow --class FLT** command, then use the **perfDelFilterMonitor** command to delete a specific monitor. If you do not specify the monitor number to delete, you are asked if you want to delete all entries as shown in Example 16-13 and Example 16-14.

Example 16-13 Deleting monitors on switch SAN384B

```
IBM_SAN384B_213:FID128:admin> perfdelfiltermonitor 2/28
This will remove ALL monitors on port 92, continue? (yes, y, no, n):
[no] y
IBM_SAN384B_213:FID128:admin> perfMonitorShow --class FLT 2/28
Filter Monitor is not present
```

Example 16-14 Deleting monitors on switch SAN80B

```
IBM_SAN80B_217:FID128:admin> perfdelfiltermonitor 19
This will remove ALL monitors on port 19, continue? (yes, y, no, n):
[no] y
```

16.3.7 ISL performance monitoring

ISL monitoring counts traffic to all reachable destination domains through an ISL. You can use the output to identify the destination domain consuming the largest portion of bandwidth. In contrast with end-to-end and filter-based monitors, you do not have to add these monitors because Fabric OS v4.4.0, ISL monitoring is enabled on E_Ports automatically.

Support: ISL monitoring is not supported on the newer ASIC, which includes the SAN24B-4, SAN40B-4, SAN80B-4, SAN768B, and SAN384B.

Use the **perfMonitorShow --class ISL** command to display the ISL traffic counters, as shown in Example 16-15.

Example 16-15 ISL monitor display

```
SAN32B_3:admin> perfMonitorShow --class ISL 1

Total transmit count for this ISL: 16904
Number of destination domains monitored: 2
Number of ports in this ISL: 1
Domain 2: 16904 Domain 4: 0
```

You can clear the ISL counters with the **perfMonitorClear --class ISL** command.

16.3.8 Top Talker monitors

The Top Talker feature is an enhancement to Advanced Performance Monitoring end-to-end monitors. They differ from end-to-end monitors in the following ways:

- ▶ End-to-end monitors cannot determine the “busiest” SID/DID pairs.
- ▶ Top Talkers monitors determine which SID/DID pairs are the major users of switch F_Port bandwidth.
- ▶ Top Talkers can be enabled on specific switch E_Ports or F_Ports.

The comparison between Top Talkers and end-to-end monitors is shown in Table 16-6.

Table 16-6 Top Talker versus End-to-End Monitors

Top Talkers	End-to-End Monitors
All possible SID/DID flow on a given port.	Single SID/DID pair.
Can monitor up to 10000 flows. If there are more flows than the H/W resources can support, the Top Talker samples traffic by looking at a new 256/2048 flows every second and extrapolates the measurement.	Can fail if number of flows exceeds the hardware resources: Condor - 256 flows Condor2 - 2048 flows

Data: Top Talkers data is available after an initial stabilization period, which is the time taken by a flow to reach the maximum bandwidth:

- ▶ 14 seconds in the SAN384B, SAN768B, SAN40B-4, SAN80B-4, and SAN24B-4
- ▶ 82 seconds in the SAN256B

You use the Top Talkers monitors to identify the SID/DID pairs that consume the highest amount of bandwidth across a particular port or the entire switch. If the total amount of traffic is within acceptable limits, then this information might not be that important. However, when the traffic amount exceeds the acceptable bandwidth, then the information from Top Talkers monitors can be effectively used to take actions such as these:

- ▶ Traffic can be routed to less busy ports, in order to reduce the load on a particular port.
- ▶ The SID/DID pairs identified as the top bandwidth consumers can be configured with appropriate Quality of Service (QoS) attributes, so that their communication flow will receive adequate priority.

The Top Talkers monitors were introduced in Fabric OS v6.0.0. When you enable the Top Talkers monitor on a port, it remains persistent across switch power cycles.

Top Talkers monitoring operates in one of the following two mutually exclusive modes:

- ▶ Port mode: In port mode, the Top Talkers monitor is installed on an F_Port and counts the traffic through that port. You can monitor either incoming (ingress) or outgoing (egress) traffic.
- ▶ Fabric mode: In fabric mode, the Top Talkers monitors are installed on all E_Ports in the fabric. They monitor and count the traffic of all possible SID/DID pairs and can therefore identify the top bandwidth consumers on a switch.

Top Talkers:

- ▶ The Top Talkers monitors measure the ingress E_Port traffic only.
- ▶ For Administrative Domains, the Top Talkers monitors are always installed in AD255.

16.3.9 Top Talkers monitors in port mode

A set of commands is available to add, delete, and display Top Talkers monitors operating in port mode. They are explained in the following sections.

Adding a Top Talkers monitor

To add a Top Talkers monitor to an F_Port, use the **perfttmon --add** command:

- For incoming traffic (receive, or RX) see Example 16-16.

Example 16-16 Adding Top Talkers for ingress traffic

```
perfttmon --add ingress [slot/]/<port>
```

- For outgoing traffic (transfer, or TX) see Example 16-17.

Example 16-17 Adding Top Talkers for egress traffic

```
perfttmon --add egress [slot/]/<port>
```

Deleting a Top Talkers monitor

To delete a Top Talkers monitor, use the **perfttmon --delete** command. See Example 16-18.

Example 16-18 Deleting Top Talkers Monitor

```
perfttmon --delete [slot/]<port>
```

Displaying Top Talkers monitor information

Use the **perfttmon --show** command to display Top Talkers information for a particular slot/port. See Example 16-19.

Example 16-19 Showing Top Talkers Monitors

```
perfttmon --show [slot]/<port> [ww|pid] [# of TT flows]
```

where:

ww|pid - Specifies the output display as either WWN or PID format. This operand is optional. If omitted, the command displays in WWN format.

of TT flows - Specifies "n" top talking flows. Valid values are between 1 and 32. If a value greater than 32 is entered, Top Talker displays counters for only 32 flows and a warning message. This operand is optional; if omitted, the command displays the top 8 flows.

Installation: Top Talkers is installed on an F_Port to measure the traffic originating from the F_Port and flowing to different destinations. The output displays the data in a sorted order based on the data rate of each flow.

Real life example

See Figure 16-25 and follow the scenario that we have implemented.

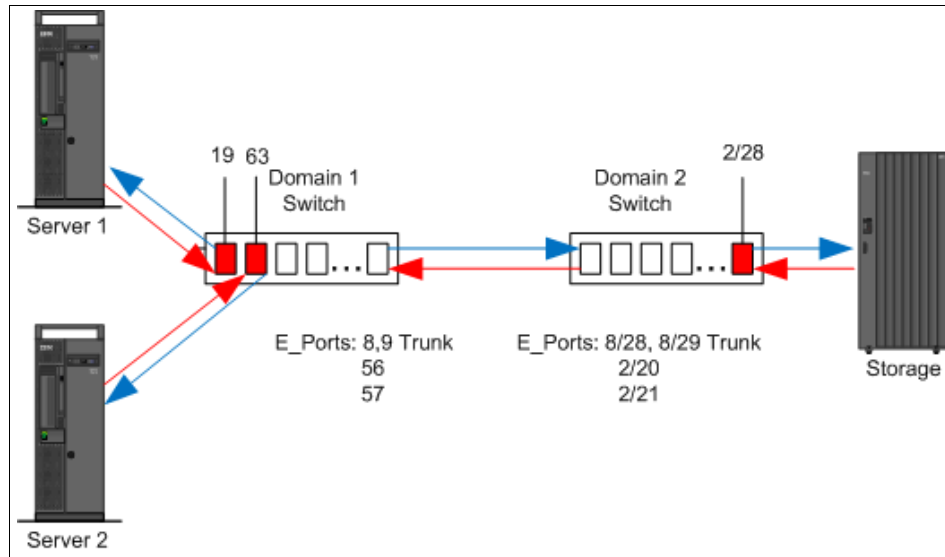


Figure 16-25 Real life example

Figure 16-25 shows the configuration with two servers and one storage device:

- ▶ Server 1 (WWNN of HBA:10:00:00:00:c9:4c:8c:1c)
- ▶ Server 2 (WWNN of HBA:10:00:00:05:1e:53:10:8b)
- ▶ Storage (WWNN of HBA: 20:06:00:a0:b8:48:58:a1)

In this configuration, Top Talkers Monitors can be configured differently, although not every configuration will bring the desired results:

- ▶ On the port 19 Domain 1 Switch, it will not bring the desired results because on port 19 there is only one talker (server 1).
- ▶ On the port 63 Domain 1 switch, there is also only one talker (server 2).
- ▶ On the port 2/28 Domain 2 switch, there is the proper placement of Top Talker; on that port both servers are talking to the storage.

Both servers are copying files to Storage 1. First we added the monitor for egress traffic on slot 2 port 28. This is very important to properly choose ingress/egress settings. See Example 16-20.

Example 16-20 Setting Top Talkers Monitors for egress traffic on a port

```
perfttmon --add egress 2/28
```

Example 16-21 shows the TopTalkers members on port 2/28.

Example 16-21 Showing Top Talkers on port (WWN and PID)

```
IBM_SAN384B_213:FID128:admin> perfttmon --show 2/28
```

Src_WWN	Dst_WWN	MB/sec
20:06:00:a0:b8:48:58:a1	10:00:00:00:c9:4c:8c:1c	87.859
20:06:00:a0:b8:48:58:a1	10:00:00:05:1e:53:10:8b	52.106


```
IBM_SAN384B_213:FID128:admin> perfttmon --show 2/28 pid
```

Src_PID	Dst_PID	MB/sec
0x025c00	0x011300	90.353
0x025c00	0x013f00	49.585

```
IBM_SAN384B_213:FID128:admin>
```

The SID/DID numbers shown in Example 16-21 are described in detail in 16.3.4, “SID/DID Performance Monitor”

To delete the monitor, enter the command in Example 16-22.

Example 16-22 Deleting Top Talkers monitor

```
IBM_SAN384B_213:FID128:admin> perfttmon --delete 2/28
IBM_SAN384B_213:FID128:admin> perfttmon --show 2/28 5
TT Monitor is not present
```

If you enable Top Talkers monitor on port 2/28 for Incoming traffic, the results might not be the ones you are expecting (see Example 16-23).

Example 16-23 Setting Top Talkers monitor for ingress traffic on a port

```
IBM_SAN384B_213:FID128:admin> perfttmon --add ingress 2/28
IBM_SAN384B_213:FID128:admin> perfttmon --show 2/28
```

Src_WWN	Dst_WWN	MB/sec
---------	---------	--------

20:06:00:a0:b8:48:58:a1	10:00:00:00:c9:4c:8c:1c	0.072
20:06:00:a0:b8:48:58:a1	10:00:00:05:1e:53:10:8b	0.002
IBM_SAN384B_213:FID128:admin> perfttmon --show 2/28		
=====		
Src_WWN	Dst_WWN	MB/sec
=====		
20:06:00:a0:b8:48:58:a1	10:00:00:00:c9:4c:8c:1c	0.073
20:06:00:a0:b8:48:58:a1	10:00:00:05:1e:53:10:8b	0.003

The traffic rate is correct because we are measuring the traffic from the storage, and the storage only sends the confirmation to both servers. We still copy files from servers to storage in this example.

16.3.10 Top Talkers monitors in fabric mode

Top Talkers monitors in fabric mode are mutually exclusive with Top Talkers monitors in port mode. When you enable fabric mode, you cannot add any Top Talkers monitor to F_ports. In addition, fabric mode Top Talkers monitors cannot coexist with any end-to-end monitors. Therefore, make sure that you delete the End-to-end monitors first.

Adding Top Talkers monitors on all switches in the fabric

To add fabric mode Top Talkers monitoring, use the command in Example 16-24.

Example 16-24 Adding Top Talkers Monitor on all E_Ports in the fabric

```
perfttmon --add fabricmode
```

You are reminded to remove all end-to-end monitors with the following message:

Before enabling fabric mode, please remove all EE monitors in the fabric.

continue? (yes, y, no, n):

If there are no end-to-end monitors in the fabric, continue by typing **y**. The command completes successfully if the local switch has no end-to-end monitors defined. If a remote switch does have end-to-end monitoring enabled, the command stills work on the local switch, but fabric mode fails on remote switch.

If you add a new switch to the fabric, the fabric mode Top Talkers configuration is not applied on it automatically. There is no automatic propagation. You must use the **perftTmon --add fabricmode** command on the new switch.

Deleting Top Talkers monitors in fabric mode

To delete all E_Ports Top Talkers monitors, use the command in Example 16-25.

Example 16-25 Delete all E-Port Top Talkers monitors in the fabric

```
perfttmon --delete fabricmode
```

Displaying Top Talkers monitors on a switch

To display a list of Top Talkers on a switch, use the command in Example 16-26.

Example 16-26 Displaying Top Talkers Monitors in fabric mode

```
perfttmon --show dom domainid [n] [wwn|pid]
```

where:

dom - Specifies the domain ID for the flow display

n - Specifies "n" Top Talking flows. Valid values are between 1 and 32. If a value greater than 32 is entered, Top Talker displays counters for only 32 flows and a warning message. This operand is optional; if omitted, the command displays the top 8 flows.

wwn|pid - Specifies display as either WWN or PID format. This operand is optional; if omitted, the command displays in WWN format.

Real life example continued

Example 16-27 shows how to enable fabric mode Top Talkers monitoring.

Example 16-27 Fabric mode TopTalkers

```
IBM_SAN384B_213:FID128:admin> perfttmon --add fabricmode
```

Before enabling fabric mode, please remove all EE monitors in the fabric

```
continue? (yes, y, no, n): [no] y
```

```
IBM_SAN384B_213:FID128:admin> perfttmon --show dom 1
```

```
IBM_SAN384B_213:FID128:admin> perfttmon --show dom 1
```

```
=====
```

Src_WWN	Dst_WWN	MB/sec	Potential	E-Ports
---------	---------	--------	-----------	---------

```
=====
```

20:06:00:a0:b8:48:58:a1	10:00:00:05:1e:53:10:8b	0.000	57	56
20:06:00:a0:b8:48:58:a1	10:00:00:05:1e:53:10:8b	0.000	9	
20:06:00:a0:b8:48:58:a1	10:00:00:00:c9:4c:8c:1c	0.000	9	
20:06:00:a0:b8:48:58:a1	10:00:00:00:c9:4c:8c:1c	0.000	57	56

In Example 16-28 you can see the results for traffic for the Domain 2 switch, which is a core switch.

Example 16-28 Showing core switch Top Talkers

```
IBM_SAN384B_213:FID128:admin> perfttmon --show dom 2
```

Src_WWN		Dst_WWN	MB/sec	
Potential E-Ports				
10:00:00:00:c9:4c:8c:1c	20:06:00:a0:b8:48:58:a1	53.474	2/20	2/21
10:00:00:00:c9:4c:8c:1c	20:06:00:a0:b8:48:58:a1	53.471	8/29	
10:00:00:05:1e:53:10:8b	20:06:00:a0:b8:48:58:a1	10.885	8/29	
10:00:00:05:1e:53:10:8b	20:06:00:a0:b8:48:58:a1	9.604	2/20	2/21

You can delete the monitor as shown in Example 16-29.

Example 16-29 Deleting Fabric Mode Tap Talkers monitor

```
IBM_SAN384B_213:FID128:admin> perfttmon --delete fabricmode
IBM_SAN384B_213:FID128:admin> perfttmon --show dom 2
TT Monitor is not present
```

For this simple example, we can see that the results are different for the particular domains:

- ▶ Domain 1 switch has a very low transfer on E_Ports, because there is no data passing through to the E_Ports in Domain 1.
- ▶ Domain 2 has a large data flow on the E_Ports, because this is a core switch and the E_Ports receive data from the edge switch.

Remember: Top Talkers monitors measure ingress E_Port traffic only.

16.3.11 Top Talkers monitoring considerations

Be aware of the following considerations regarding Top Talkers monitoring:

- ▶ Top Talker monitors cannot detect transient surges in traffic through a given flow.
- ▶ You cannot install a Top Talker monitor on a mirrored port.
- ▶ A Top Talker can monitor only 10,000 flows at a time.
- ▶ A Top Talker is not supported on VE_Ports, EX_Ports, and VEX_Ports.
- ▶ The maximum number of F_Port Top Talker monitors on an ASIC is 8. If Virtual Fabrics is enabled, the maximum number of F_Port Top Talker monitors on an ASIC is 4.

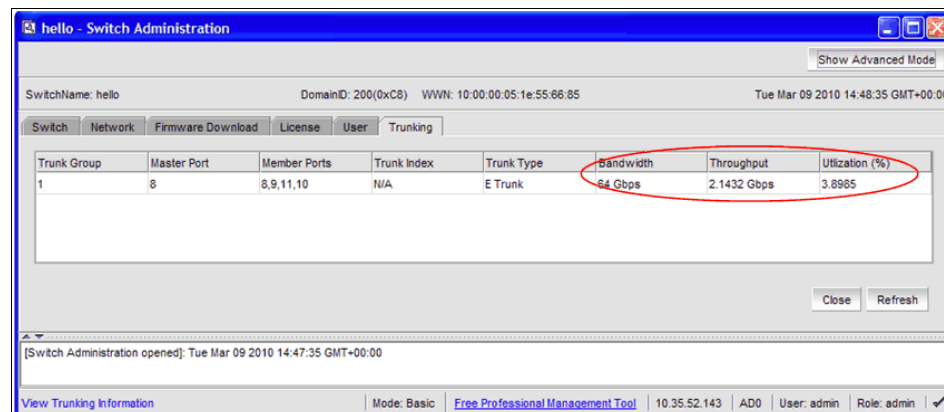
16.3.12 Trunk monitoring

If you want to monitor trunks, you can set the monitor only on the master port. The monitor will automatically move to a new master port, if it changes. Also, if a monitor is installed on a port which later becomes a subordinate port in the trunk, that monitor will move to the master port.

Note the following considerations:

- ▶ End-to-end monitors are not supported for ISLs.
- ▶ For F_Port trunks, end-to-end masks are allowed only on the F_Port trunk master. Unlike the monitors, if the master changes, the mask does not automatically move to the new master port.
- ▶ SAN24B-4 supports eight filter-based monitors for trunks.
- ▶ The SAN40B-4, SAN80B-4, SAN256B, SAN384B, SAN768B platforms support 12 filter-based monitors for trunks.

Recent enhancements in Web Tools provides the trunk bidirectional bandwidth, throughput and utilization of trunk as shown in Figure 16-26.



Trunk Group	Master Port	Member Ports	Trunk Index	Trunk Type	Bandwidth	Throughput	Utilization (%)
1	8	8,9,11,10	N/A	E Trunk	64 Gbps	2.1432 Gbps	3.8985

Figure 16-26 Trunk enhancements

16.3.13 Saving and restoring the monitoring configuration

You can save the current setup of end-to-end and filter monitors in the non-volatile memory on the switch by using the **perfCfgSave** command. This action overwrites the previously saved performance monitoring settings, so you are asked for confirmation.

If you want to apply the saved configuration, enter the **perfCfgRestore** command. This command overwrites the currently active performance monitoring configuration, and you are again asked to confirm the action.

Finally, you can clear the saved performance monitoring configuration in the non-volatile memory using the **perfCfgClear** command.

As the space in non-volatile memory is limited, the number of monitors saved is also limited as follows:

- ▶ Up to 16 end-to-end monitors per port can be saved.
- ▶ Up to 16 filter monitors per port can be saved.
- ▶ Up to 512 monitors per switch will be saved.
- ▶ If there are more than 512 monitors configured on the switch, the end-to-end monitors are saved first.

Memory: Monitors created by Web Tools are not saved in persistent memory.

16.4 SCSI commands with Web Tools

The SCSI Commands monitor shows the total number of read or write commands per switch port or per specific LUN on switch port. Figure 16-27 shows the available graph choices.

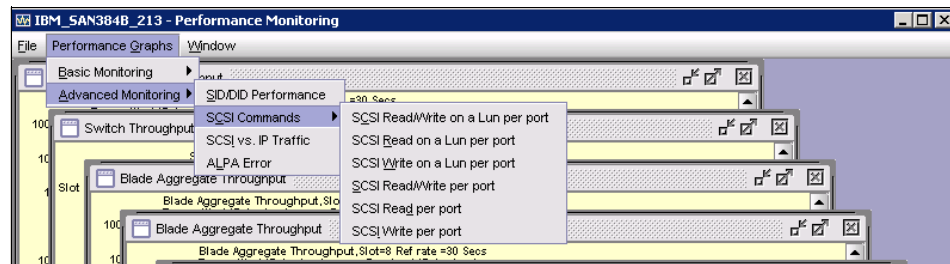


Figure 16-27 SCSI Commands monitor

If you select a graph showing the number of commands per port, you are prompted to specify the port number, as shown in Figure 16-28.

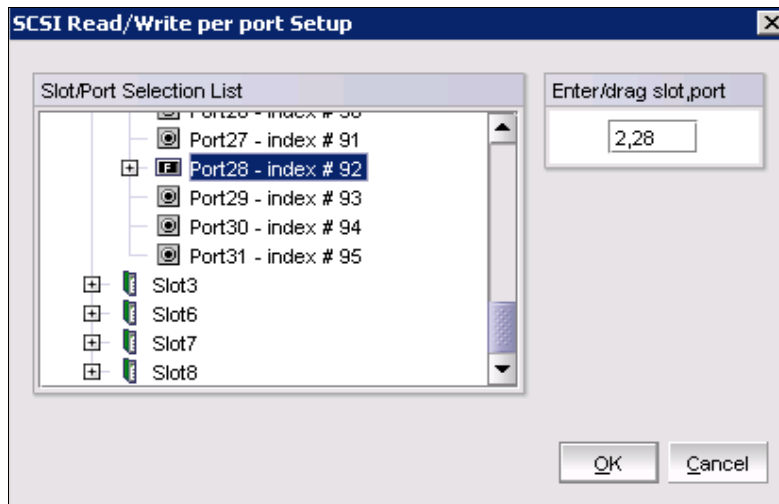


Figure 16-28 SCSI Read/Write per port Setup window

If you select any of the graphs displaying the number of commands on a LUN per port, then you need to enter the port number and also the LUN number (see Figure 16-29).

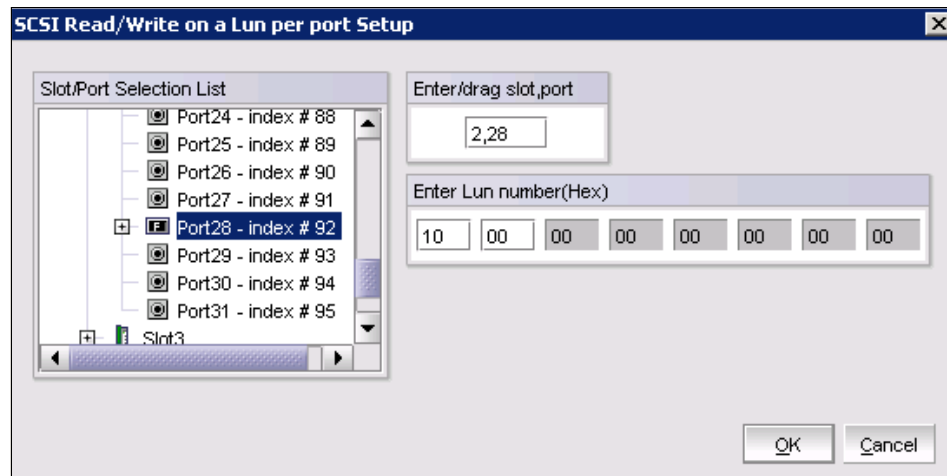


Figure 16-29 SCSI Read/Write on a LUN per port Setup

16.4.1 SCSI versus IP traffic

The SCSI versus IP Traffic graph monitor is accessible by selecting **Performance Graphs** → **Advanced Monitoring** → **SCSI vs. IP Traffic**.

First, you select the ports that you want to monitor, as shown in Figure 16-30. You add the ports that you want to monitor to the list on the right, and then click **Apply**.

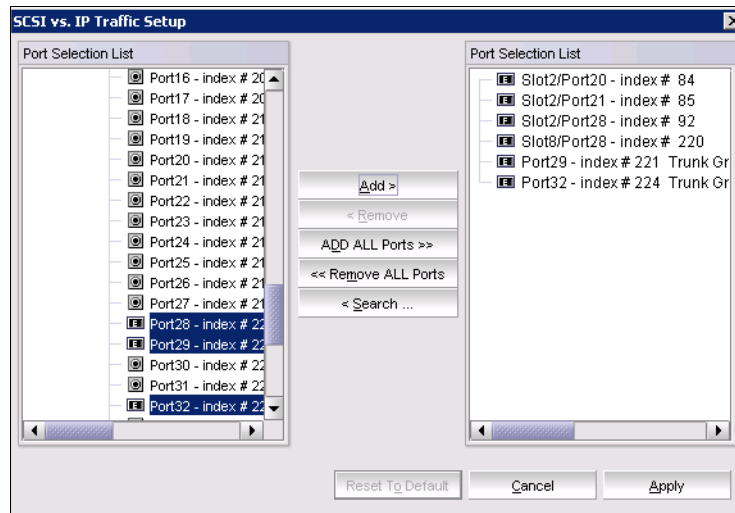


Figure 16-30 SCSI versus IP Traffic Setup window

The graph displays as shown in Figure 16-31.



Figure 16-31 SCSI vs. IP Traffic graph

This graph shows the percentage of IP and SCSI traffic on the current switch on a port basis.

16.4.2 ALPA error

This feature is available only on the older switches based upon the Bloom Application Specific Integrated Circuit (ASIC). Thus, we do not discuss it here.

16.5 Bottleneck detection

A bottleneck in the fabric is where frames cannot get through as fast as they should. In other words, a bottleneck is a port where the offered load is greater than the achieved egress throughput. These features detect two types of bottleneck detection.

16.5.1 Latency bottleneck

A port at which the offered load exceeds the rate at which the other end of the link can continuously accept the traffic, even if the physical link capacity is not exceeded. This can impact the other devices in the fabric sharing the link.

16.5.2 Congestion bottleneck

A port that is unable to transmit the frames at the offered rate because the offered rate is greater than physical data rate of the link.

In Fabric OS v6.3.x, bottleneck detection was configured on a per-port basis. Starting in Fabric OS v6.4.0, you configure bottleneck detection on a per-switch basis, with per-port exclusions. This bottleneck feature detects latency and congestion bottlenecks and reports the bottlenecks through RASlog alerts and SNMP traps. You can set alert thresholds for the severity and duration of the bottleneck.

Example 16-30 shows steps to enable bottleneck monitoring, disables alerts on port 1, excludes ports 2, 3, and 4 from bottleneck monitoring, and changes the alert settings on ports 2 and 3. The **bottleneckmon --status** command shows the settings for these ports. Note that this example changes the alert settings on ports 2 and 3, even though they are excluded from bottleneck detection

Bottleneck Detection

Example 16-30 Bottleneck monitoring

```
switch:admin> bottleneckmon --enable
switch:admin> bottleneckmon --config -noalert 1
switch:admin> bottleneckmon --exclude 2-4
switch:admin> bottleneckmon --config -alert -lthresh .99 -ctresh .9
-time 4000 -qtime 600 2-3
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====
Switch-wide alerting parameters:
=====
Alerts - Yes
Latency threshold for alert - 0.100
Congestion threshold for alert - 0.800
Averaging time for alert - 300 seconds
Quiet time for alert - 300 seconds
Per-port overrides for alert parameters:
=====
Slot Port Alerts? LatencyThresh CongestionThresh Time(s) QTime(s)
```

```
=====
0 1 N -- -- -- --
0 2 Y 0.990 0.900 4000 600
0 3 Y 0.990 0.900 4000 600
Excluded ports:
=====
Port
=====
2
3
4
=====
```




Health and troubleshooting

In this chapter, we overview the steps that you can take to ascertain the health of the storage area network (SAN) fabric and to troubleshoot problems. We discuss SAN Health, a powerful tool that allows you to collect data and analyze this data for potential issues.

17.1 SAN Health

SAN Health is a very powerful tool that helps a SAN administrator or SAN user optimize the existing SAN. The tool allows you to collect data and to analyze this data for potential issues.

SAN Health provides a full status report on your SAN environment by the use of two mechanisms: a back-end reporting processor, and a front-end data collection agent. When the Front End (FE) has completed a scan of the SAN and collected all the appropriate data, the Back End (BE) analyzes this information for potential issues, and produces a Visio topology diagram of the SAN. The BE report covers fabrics, switches individual ports, and historical performance graphs. It also presents some best practice procedures.

17.1.1 New features of SAN Health

SAN Health 3.2.0b includes the following new features:

- ▶ Improved reporting for all Brocade M-series SAN solutions, including the Brocade Mi10k
- ▶ The ability to audit switches managed by Brocade EFCM
- ▶ Enhanced topology diagram layouts
- ▶ More detailed diagnostics information obtained from switches
- ▶ A redesign of the report content and layout
- ▶ FICON enhancements for mainframe environments

17.1.2 Implementing SAN Health

In this section, we explain how to download, install, and use SAN Health.

Installing Brocade SAN Health

To install Brocade SAN Health:

1. Go to the following link and download *SAN Health Diagnostic Capture*:
<http://www.brocade.com/services-support/drivers-downloads/san-health-diagnostics/index.page>
2. Unzip InstallSANHealth322.zip and run the file InstallSANHealth322.exe.
3. Follow the step-by-step instructions. Figure 17-1 shows the first step to install San Health. Click **Next**.

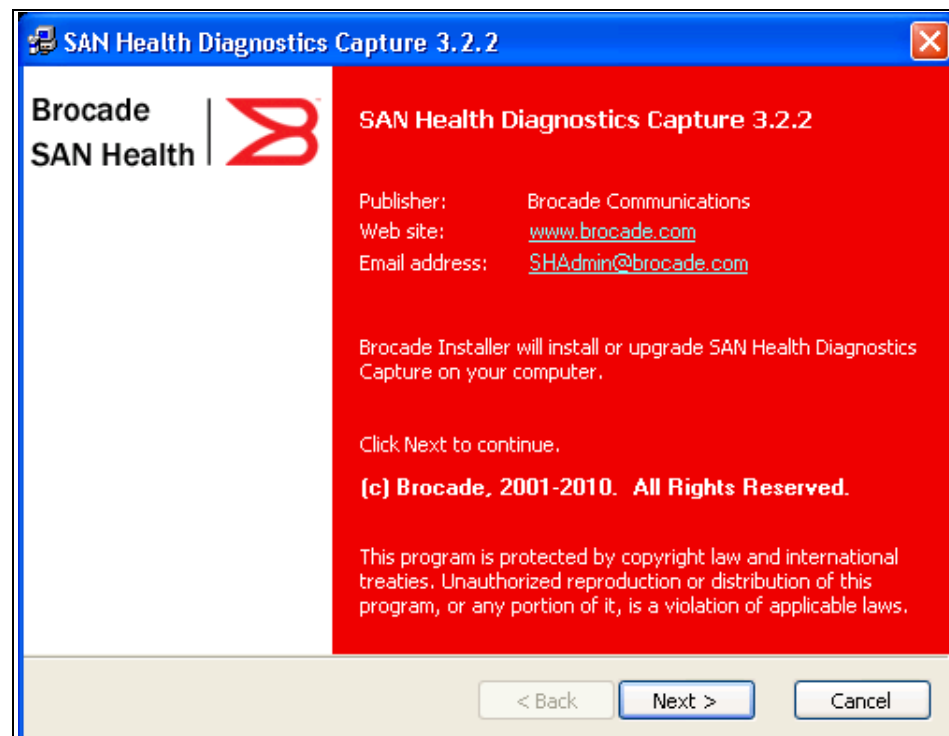


Figure 17-1 SAN Health Diagnostics Capture Installation panel

4. The license agreement displays, as shown in Figure 17-2. Read the license agreement. If you agree to the terms of the license, select **I agree to these terms and conditions**, and click **Next**. Otherwise, click **Cancel**.

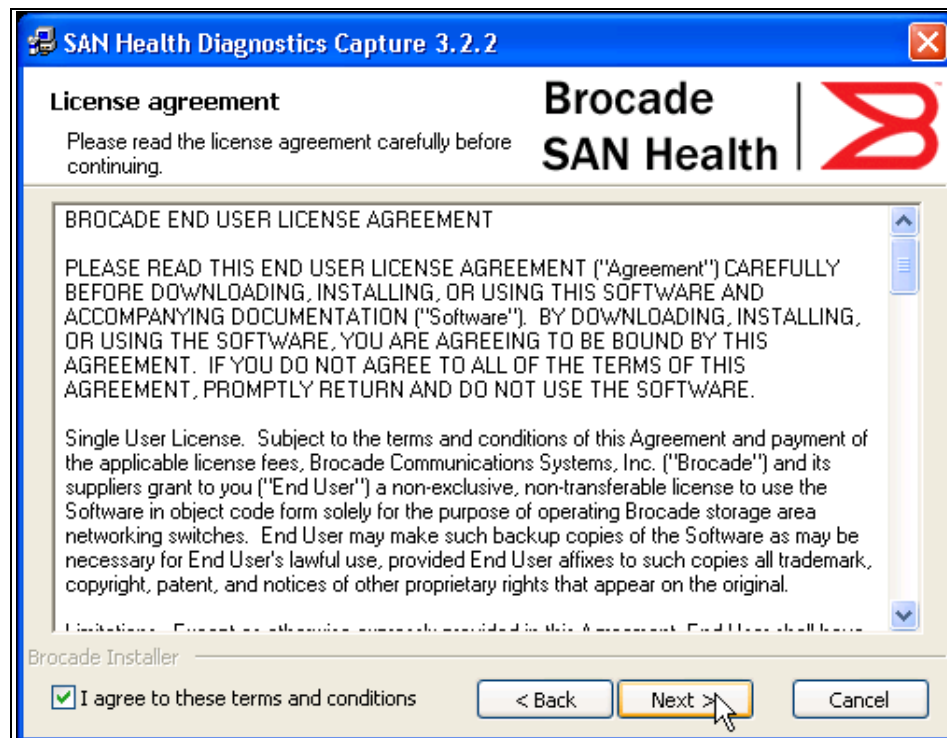


Figure 17-2 SAN Health license agreement

5. In the next panel, you select the installation folder and the audit and working folder. Check if you have enough space (10 MB) to load SAN Health. Click **Install** to install San Health as shown in Figure 17-3.

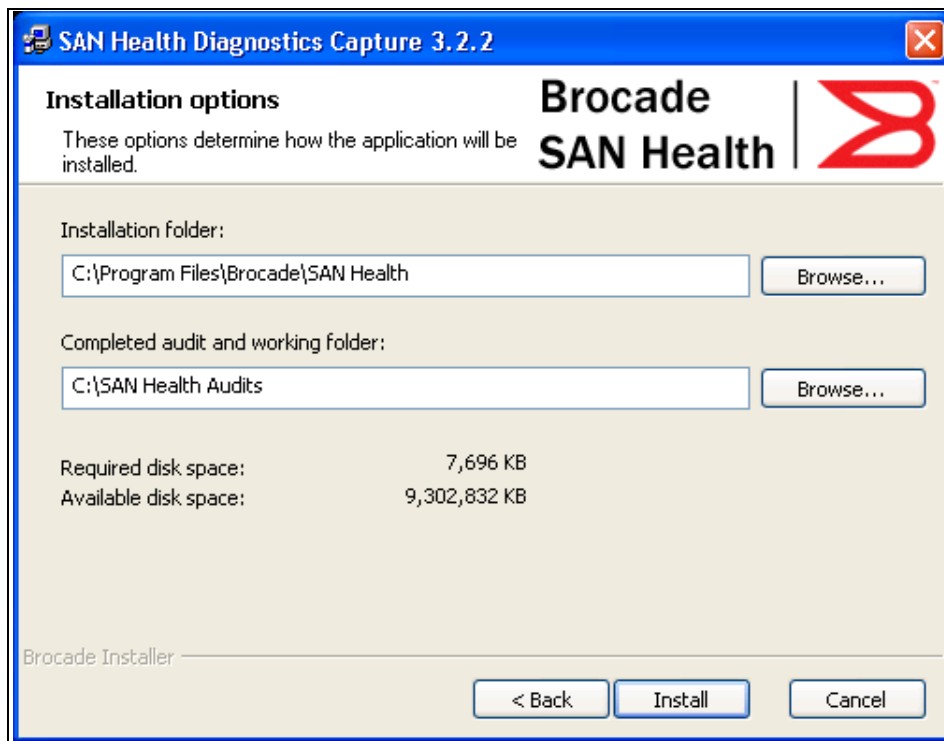


Figure 17-3 SAN Health installation options

SAN Health: If there is a previous version of SAN Health installed, the installation wizard will recognize this and will ask for permission to uninstall the older version.

6. Figure 17-4 shows the Installation completed panel. Select **Start the application** to start SAN Health as soon as you click **Finish**.



Figure 17-4 SAN Health installation complete

Using Brocade SAN Health

After you have downloaded, decompressed, and installed SAN Health, you can execute it using the desktop icon. The startup panel displays as shown in Figure 17-5. Click **New**.

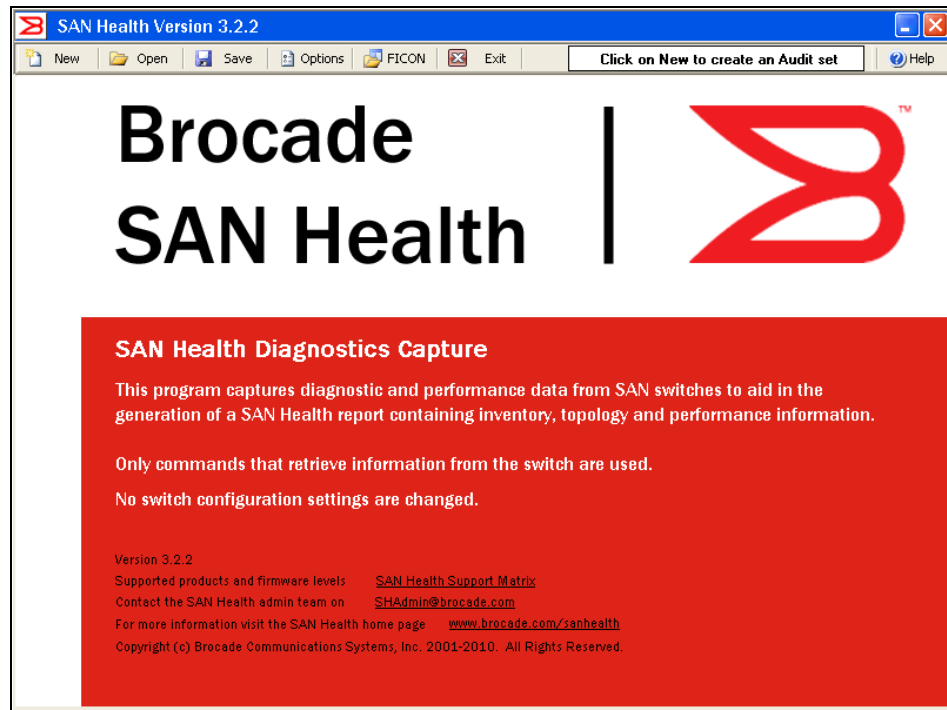


Figure 17-5 SAN Health startup panel

The interface might ask you to discard the current Audit Set if you are upgrading from a previous version. Depending on your needs, answer **Yes** or **No**.

To use SAN Health:

1. First you have to answer some basic questions in the **Site Details** tab. The information provided will be used on the title page of the SAN Health report. Mandatory fields are marked with an *. See Figure 17-6.

The screenshot shows the SAN Health Version 3.2.2 application window. The title bar is blue with the application name and version. The menu bar includes New, Open, Save, Options, FICON, Exit, and Help. A status bar at the top right indicates "No Switches Selected For Audit".

A yellow information box contains the following text:

To complete a SAN Health audit, first you need to create an audit set.

Step 1) Enter the Site Details, these are used on the title page of the SAN Health report.

Step 2) Enter the Report Return details, report processing is centralized at Brocade and reports are returned via a secure single sign on web page.

Step 3) Name the SAN that you are auditing, this is used as a title on the SAN Health report and in the report file name.

Step 4) Enter Switch or EFCM server IP Address and login credentials to add fabrics to the audit set.

Step 5) Click on the fabric(s) in the tree view to name the fabric, enter performance capture duration and specify the support provider for the fabric.

Step 6) Ensure that switch login credentials are correct and optionally set the Visio topology diagram position for each switch.

Step 7) Start the Audit, click on the "Preflight Check" button. If the audit set details have been entered correctly the "Start Audit" button is enabled.

While creating an audit set, SAN Health needs to connect to each switch to determine the switch model number and communication capabilities. Click on the "Test Connectivity and Get Switch Details" button at the SAN, fabric or individual switch levels to gather this information. During the connectivity test, fabric membership is determined and the switch is automatically moved into the appropriate fabric container.

The main window has a tabbed interface with the following tabs: Site Details, Report Return, SAN Details, Add Switches, Fabric Details, Switch Details, and Start Audit. The Site Details tab is active.

The Site Details tab contains the following fields:

Salutation*	<input checked="" type="radio"/> Mr <input type="radio"/> Ms*	Address1*	Gran Via, 1
First Name*	J	Address2	
Last Name*	Lainen	City*	Madrid
Job Title*	IT Architect	State/Province*	Not Applicable
Phone*	0344	Zip/Postal Code*	28001
Company Name*	IBM	Country*	Spain

A yellow information box on the right side of the Site Details tab contains the following text:

Site details are used on the title page of the SAN Health report
* Indicates a mandatory field

The status bar at the bottom shows the following information:

4:36:02 PM> New SAN Audit Set Started Using SAN Health Version 3.2

0 Warnings 0 Errors Telnet Activity

Figure 17-6 SAN Health Diagnostics Capture Site Details tab

2. Go to the **Report Return** tab and fill in an email Address for the report return as in Figure 17-7.

The screenshot shows the SAN Health Version 3.2.2 application window. The title bar reads "SAN Health Version 3.2.2". The menu bar includes "New", "Open", "Save", "Options", "FICON", "Exit", and "Help". A status bar at the top right says "No Switches Selected For Audit".

A large yellow box contains instructions for completing a SAN Health audit:

- To complete a SAN Health audit, first you need to create an audit set.
- Step 1) Enter the Site Details, these are used on the title page of the SAN Health report.
- Step 2) Enter the Report Return details, report processing is centralized at Brocade and reports are returned via a secure single sign on web page.
- Step 3) Name the SAN that you are auditing, this is used as a title on the SAN Health report and in the report file name.
- Step 4) Enter Switch or EFCM server IP Address and login credentials to add fabrics to the audit set.
- Step 5) Click on the fabric(s) in the tree view to name the fabric, enter performance capture duration and specify the support provider for the fabric.
- Step 6) Ensure that switch login credentials are correct and optionally set the Visio topology diagram position for each switch.
- Step 7) Start the Audit, click on the "Preflight Check" button. If the audit set details have been entered correctly the "Start Audit" button is enabled.

Below the instructions, it states: "While creating an audit set, SAN Health needs to connect to each switch to determine the switch model number and communication capabilities. Click on the 'Test Connectivity and Get Switch Details' button at the SAN, fabric or individual switch levels to gather this information. During the connectivity test, fabric membership is determined and the switch is automatically moved into the appropriate fabric container."

The main window has a tabbed interface with the following tabs: "Site Details", "Report Return" (selected), "SAN Details", "Add Switches", "Fabric Details", "Switch Details", and "Start Audit".

In the "Report Return" tab, there is a message: "The email address must be valid to ensure report return". Below this, there are two input fields: "Email*" and "Retype Email*", both containing "lainen@ibm.com".

To the right of the email fields, there is a section titled "Optional - Send a duplicate of the completed report to the following people?". It contains three checkboxes:

- ☒ A Brocade staff member you are working with Email Address [input field]
- ☐ Brocade Support (Case number required)
- ☐ Another company that you are working with

The bottom status bar shows: "4:36:02 PM> New SAN Audit Set Started Using SAN Health Version 3.2". On the right, it shows "0 Warnings", "0 Errors", and "Telnet Activity" with icons for a terminal, a printer, and a help icon.

Figure 17-7 Report Return tab

Reports: You can also send the report to additional readers by checking one of the check boxes to the right. If no check box is checked, only the email address provided will be used for returning the report.

3. Now, add the switches or fabrics into the data collection engine. Start by naming the SAN on the *SAN Details* tab as shown in Figure 17-8.

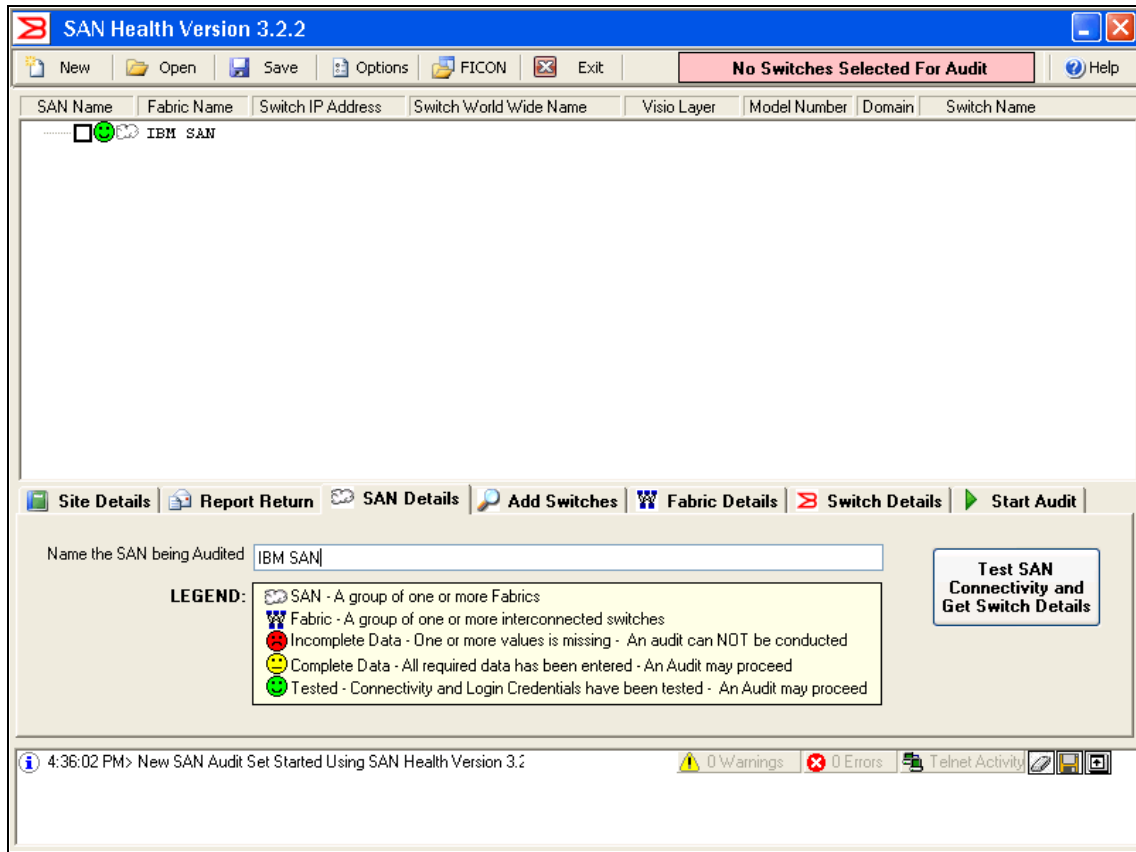


Figure 17-8 SAN Health Diagnostics Capture SAN Details

4. Next, add your switches using the *Add Switches* tab (Figure 17-9).

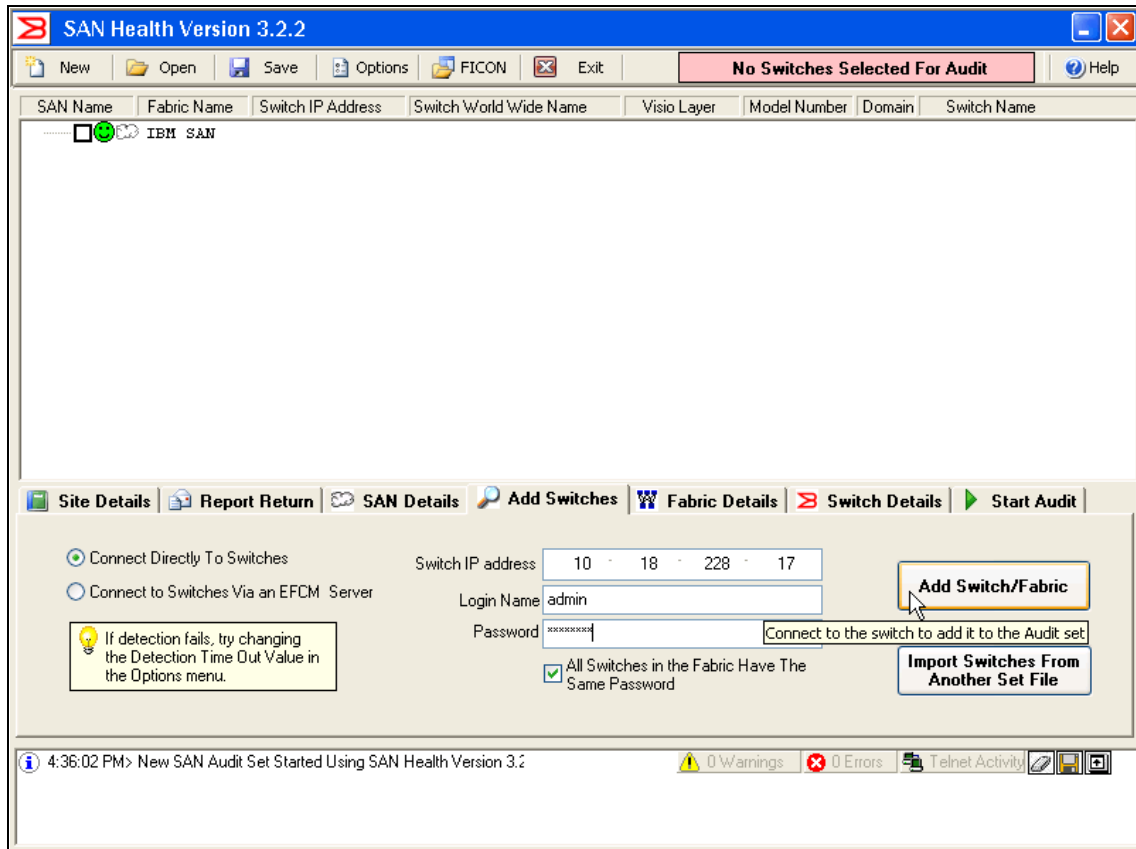


Figure 17-9 SAN Health Diagnostics Capture Add Switches

The software will start looking for the switches, as show in Figure 17-10.

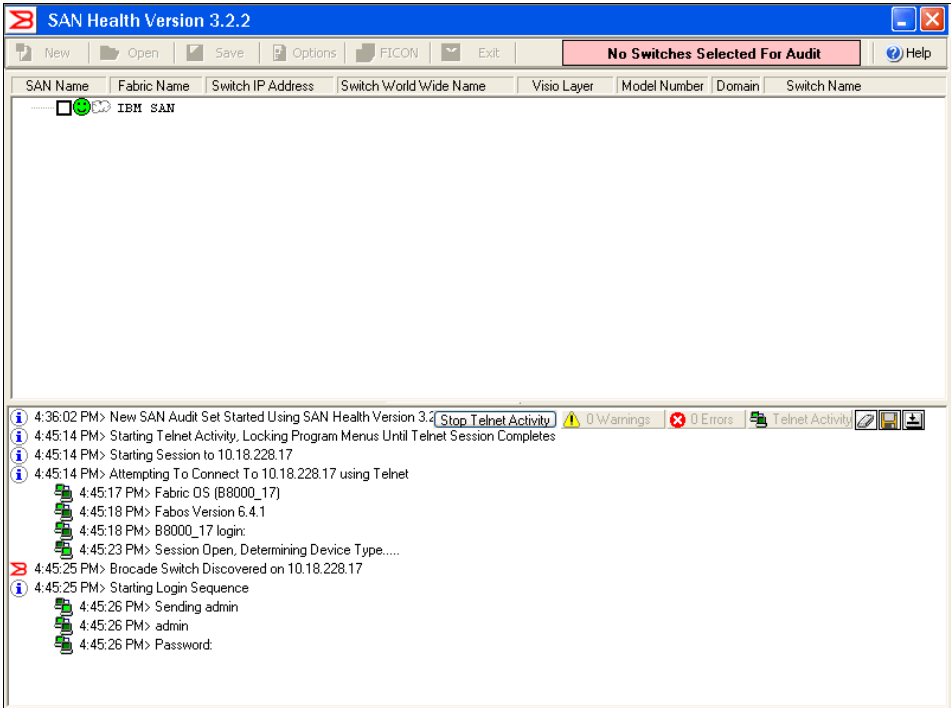


Figure 17-10 Adding a switch

- On the *Fabric* tab, provide details about the fabric. Then test the connectivity as shown in Figure 17-11.

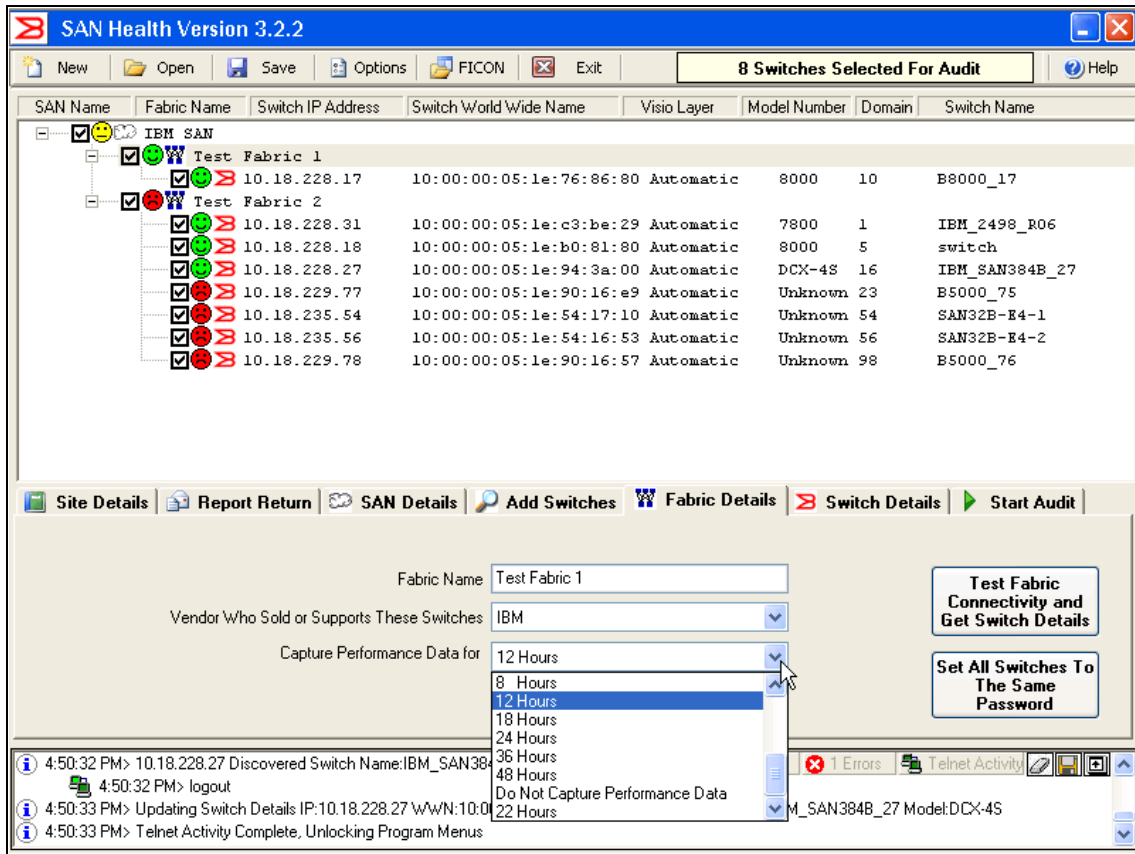


Figure 17-11 SAN Health Diagnostics Capture Fabric Details

- Go to the Start Audit tab and run the *Preflight check* as shown in Figure 17-12.

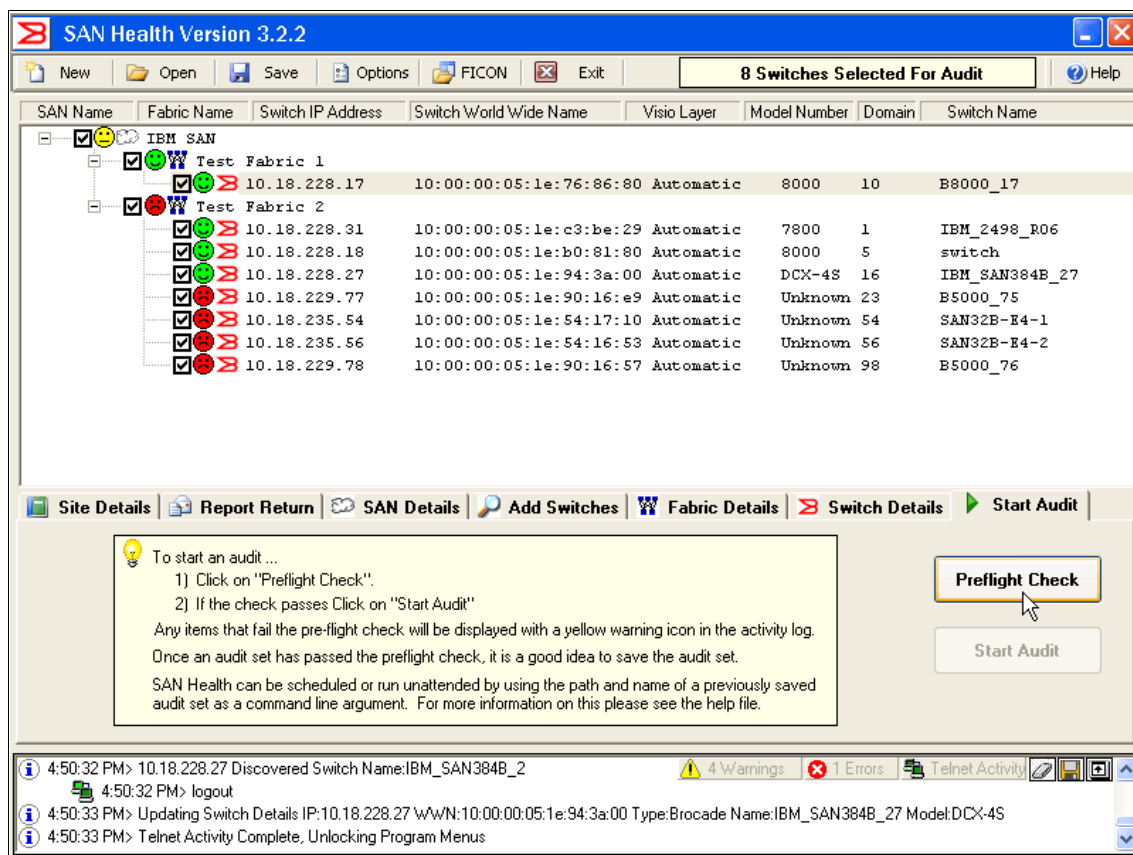


Figure 17-12 SAN Health Diagnostics Capture Start Audit

If the *Preflight check* did not pass, correct any error and rerun the check until it passes. Normally, you'll get a "green smiley icon" is all the tests are OK, as show in Figure 17-13.

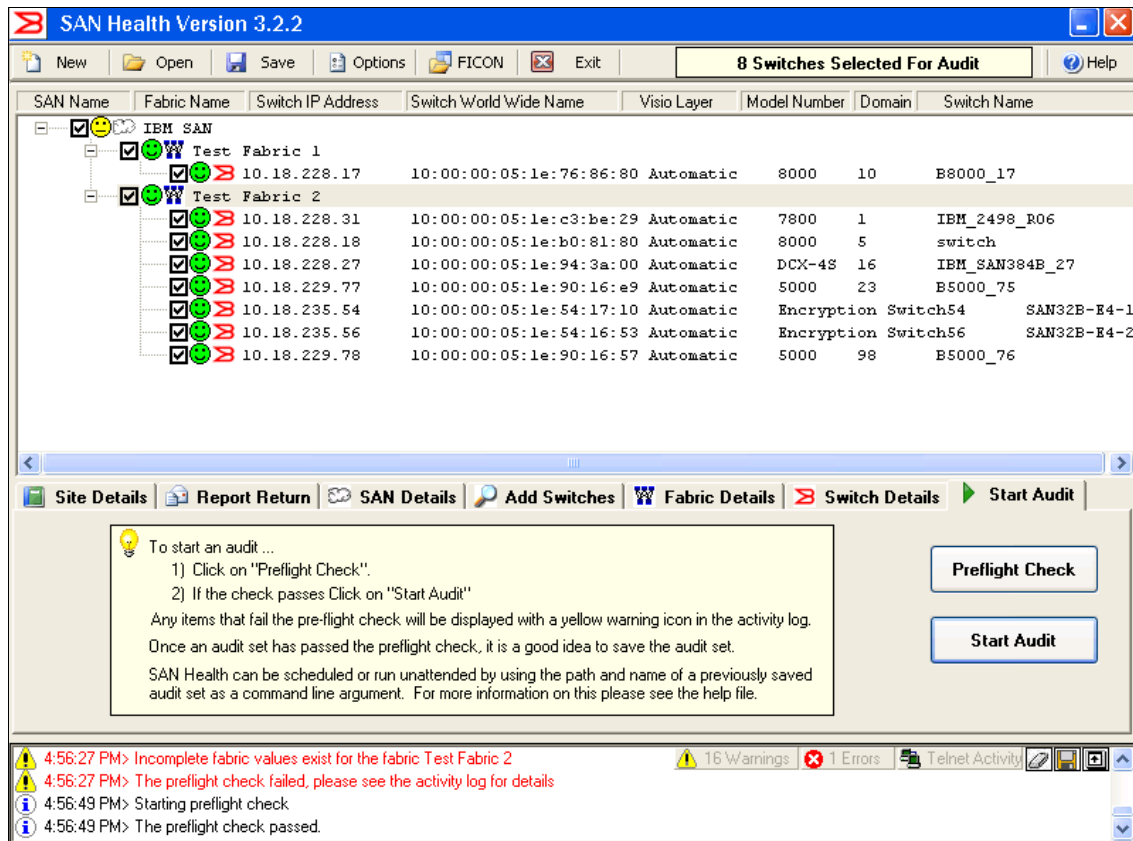


Figure 17-13 Preflight test successfully passed

7. The audit begins when you select **Start Audit**.

SAN Health gathers data. How long this process takes depends on the *capture performance data interval* that you set on the *Fabric* tab. You can watch the progress of the tool as it completes the checks.

8. Right-clicking a specific switch allows you to view its status details, as shown in Figure 17-14.




Figure 17-14 SAN Health Diagnostics Capture Show Details

If you decide to configure the audit for a long period, you'll see a status of the different checkpoints configured, as you can see in Example 17-15.

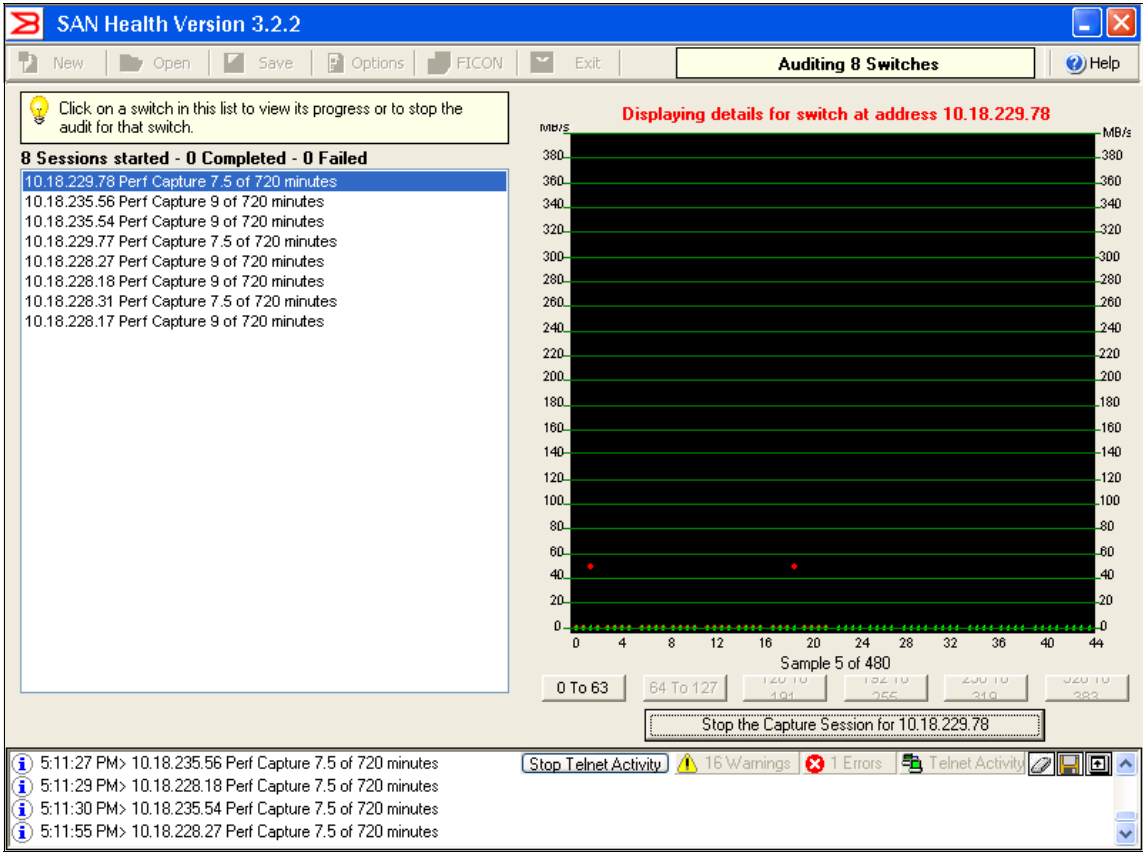
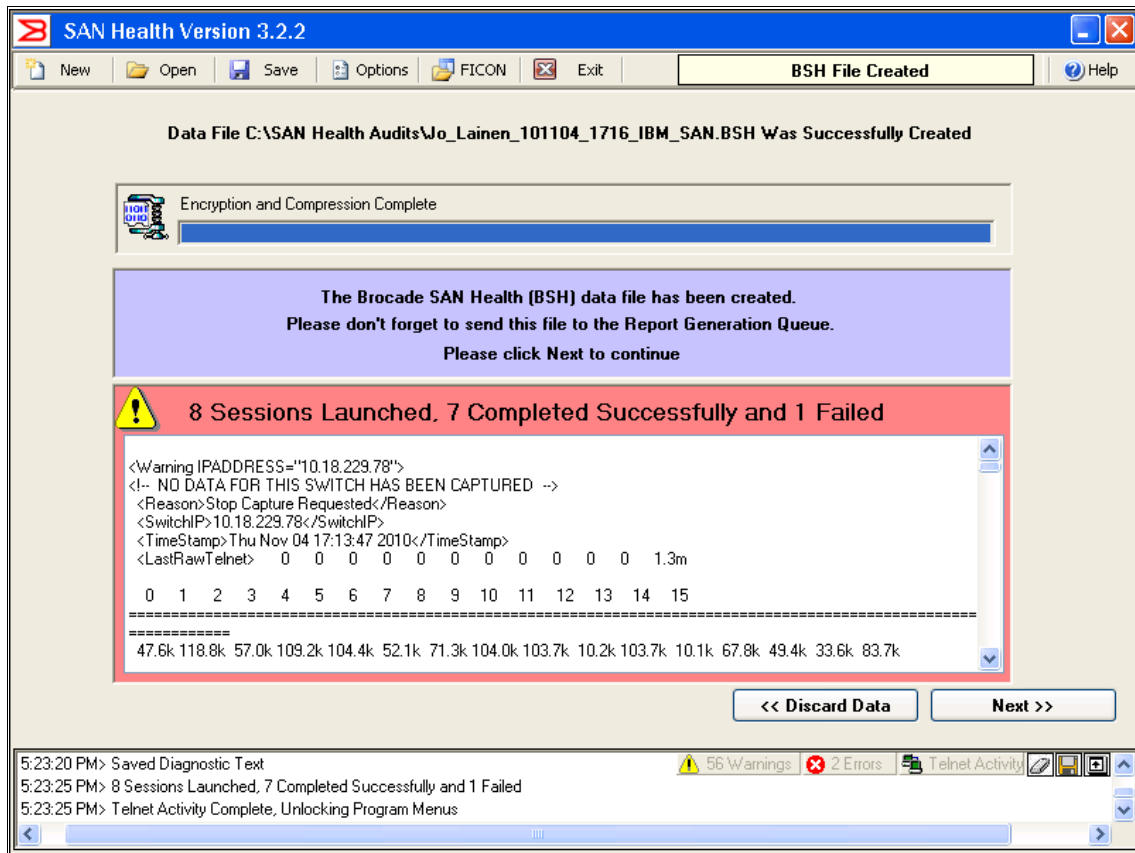


Figure 17-15 Status of the audit

9. When this process has completed, the output is an encrypted and compressed file that can be found in the C:\SAN Health Audits\ directory. You can also see the data file in Figure 17-16.



Error messages: In the screen capture, you can see one error message. In our test we stopped the audit in one of the switches to show this screen, and what would happen in case of an error. In a normal audit, you will not see this error message.

10. To complete the process, you have to send the encrypted SAN Health file (.BSH) to the Brocade report generator. You can either do that by clicking **Send to diagnostic data file to the report generation queue via HTTPS**, sending it as email attachment to SHUpload@brocade.com, or by uploading it manually to the Brocade URL as shown in Figure 17-17.




Figure 17-17 SAN Health Diagnostics Capture where to send the output

11. By return you will receive a link where you can download your analyzed data at the Brocade site as a .zip file. If you have an account at Brocade it will be stored there. If you do not have an account it will be generated automatically. The .zip file will contain two files. One is a Visio connection diagram of the SAN Layout, and the other is a thorough SAN analysis captured into an Excel spreadsheet. You must have Excel loaded on your workstation in order to view this report.

Zip file: Be aware that the .zip file is only available for download for 30 days.

The following figures show a selection of screen captures from this report.

Figure 17-18 shows the SAN Health Summary.

SAN SUMMARY DETAILS FOR IBM FABRIC														Table Of Contents				
SWITCHES IN SAN IBM Fabric																		
Fabric Name	Switch Name	Domain	IP Address	World Wide Name	Model	Speed	OS Ver	Ports	Unused									
IBM Fabric	IBM_2498_B40	4	10.64.210.183	10:00:00:05:1e:09:73:fd	5100	8G	6.1.0	40	38									
SUMMARY FOR 1 SWITCHES TOTALING 40 PORTS THAT ARE 5 % UTILIZED																		
Fabric Name	Switch Count				Port Count				Port Use Metrics									
	1G	2G	4G	Total	1G	2G	4G	Total	ISL Ports	Devices	Unused	Utilization						
IBM Fabric	0	0	0	1	0	0	0	40	2	0	38	5%						
DEVICE COUNT FOR ALL FABRICS																		
Device Description					Count	Device Description					Count							
NO DEVICES IN SAN																		
PORT USE																		
Fabric Name	Port Use						Fan Out Ratios			Port Long Distance Modes								
	Disk	Tape	Host	ISL	Free	Total	Host:Disk	Port:ISL	Device:ISL	10km	25km	50km	100km	Auto				
IBM Fabric	0	0	0	2	38	40	0:0	19:1	0:2	40	0	0	0	0				
BANDWIDTH UTILIZATION STATISTICS																		
Fabric Name	Device Bandwidth Utilization (per port)								ISL Bandwidth Utilization (per port)									
	Dev. Count	0 - 25%		25-75%		75-100%		Average MB/s	Max MB/s	ISL Count	0 - 25%		25-75%		75-100%		Average MB/s	Max MB/s
		Av	max	Av	Max	Av	Max				Av	max	Av	Max	Av	Max		
IBM Fabric	0	0	0	0	0	0	0	0	0	2	2	2	0	0	0	0	0	0
LICENSE SUMMARY																		
Full Fabric	1	WEB TOOLS			0	Zoning			0	Trunking			0	SES			0	
Perf. Monitoring	0	Fabric Watch			1	Extended Fabric			1	Remote Switch			0	Secure FOS			1	
Quick Loop	0	VL2 Upgrade			0	VL4 Upgrade			0									
ZONING METRICS																		
Fabric Name	Zone Database Use	Aliases Statistics				Zone Statistics				Config Statistics								
		Aliases	AvMem	MaxMem	Hanging	Zones	AvMem	MaxMem	Hanging	Configs	AvMem	MaxMem	Hanging					
IBM Fabric	0.3% of 1045k	27	1.3	5	27	8	2.2	6	8	4	3.8	6	4					

Figure 17-18 SAN Health Summary

Figure 17-19 shows a copy of the Visio diagram.

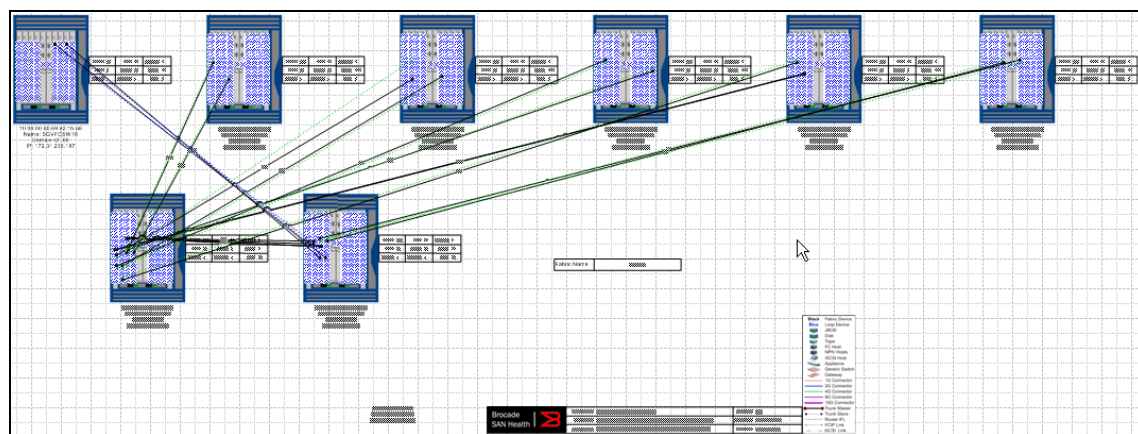


Figure 17-19 SAN Health Visio diagram

If you do not have Visio installed in your workstation, you can download a version called *Visio Viewer* that is free from:

<http://www.microsoft.com/downloads/details.aspx?familyid=D88E4542-B174-4198-AE31-6884E9EDD524&displaylang=en>

Figure 17-20 shows a fabric-specific summary.

SWITCH SUMMARY DETAILS FOR IBM_2498_B40															Table Of Contents				
IBM_2498_B40 IN FABRIC IBM FABRIC																			
Switch Name	IBM_2498_B40				Brocade Model	Brocade 5100 Switch				FOS Version	6.1.0								
IP Address	10.64.210.183				WwN	10:00:00:05:1e:09:73:fd				FOS Build Date	Tue Mar 11 23:05:38 2008								
FC IP address	-				Domain ID	4				Port Count	40								
Vendor	IBM				Serial Number	ALM0617D001				Unused Ports	38								
Switch Status	DOWN				MAC Address	-				ISL Ports	2								
Active Config	FabricA_Cfg				Ethernet Port	AUTO				Switch Speed	8G								
Zone DB Use	0.3 % of 1045 kbytes				Switch Date	Fri May 2 17:18:40 2008				Switch State	Online								
POST	Enabled				QuickLoop	Disabled				Switch Mode	Native								
Telnet Timeout	10 minutes				Fabric Watch	Alarms are enabled				Role in Fabric	Subordinate								
LICENSE SUMMARY																			
Full Fabric	Yes				WEB TOOLS	No				Zoning	No				SES	No			
Perf. Monitoring	No				Fabric Watch	Yes				Extended Fabric	Yes				Remote Switch	No	Secure FOS	Yes	
Quick Loop	No				VL2 Upgrade	-				VL4 Upgrade	-								
I/O PARAMETERS																			
Buffer Credits	16				Data Field Size	2112				Error Detect Time	2000				Resource Timeout	10000		IOD Setting	Enabled
Port ID Mode	1				Distance Mode	0				Interop Mode	No				Translative Mode	0		DLS Setting	NotSet
ENVIRONMENTAL STATUS																			
Switch Up Time				Last Boot Up At				Powered On For				Last Reboot Reason							
0 Days				-				-				-							
Fan Status or if OK Speed in RPM								Power Supply Status				Av CPU Load Prior to Audit							
Fan 1	Fan 2	Fan 3	Fan 4	Fan 5	Fan 6	PS 1	PS 2	PS 3	PS 4	1 Min	5 Min	15 Min							
OK	OK	-	-	-	-	OK	Faulty	-	-	0.34	0.32	0.27							
Switch Temperature Sensors						Director Blade Temperature Sensors													
Sen 1	Sen 2	Sen 3	Sen 4	Sen 5	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8	Slot 9	Slot 10					
C	35	22	30	21	-	-	-	-	-	-	-	-	-	-	-				
F	95	71	86	69	-	-	-	-	-	-	-	-	-	-	-				
Current Switch	Faulty Ports		Missing SFPs		Power Supplies		Temp Sensors		Faulty Fans		Port Status		ISL Status						
Alerting Policy	Marginal	Down	Marginal	Down	Marginal	Down	Marginal	Down	Marginal	Down	Marginal	Down	Marginal	Down					
Thresholds	4	10	0	0	1	1	1	2	1	2	4	10	-	-	-				
MONITORING AND ALERTING																			
FABRIC WATCH		SYSLOG SETTINGS																	
Status	Syslog Status	IP Address 1	IP Address 2	IP Address 3	IP Address 4	IP Address 5	IP Address 6												
enabled	In Use	-	-	-	-	-	-												
SNMP SETTINGS																			
Switch Description						Switch Location				Contact Information									
FibreChannelSwitch.						EndUserPremise.				FieldSupport.									
Community Strings : SecretCode(rw)				OrigEquipMfr(rw)		private(rw)		public(ro)		common(ro)		FibreChannel(ro)							
Access Control List : Not Configured				Not Configured		Not Configured		Not Configured		Not Configured		Not Configured							
Introduction		Summary		SAN Ports		Visio Topology Diagram		F IBM Fabric-468a00		Z IBM Fabric-468a00		S IBM							

Figure 17-20 A fabric-specific summary

17.2 Error logs

The b-type family of switches provide multiple sources of error logs and debug data. You can collect these logs from Web Tools or CLI or using automated tools that run when the switch experiences a critical problem. In addition there is also the possibility to collect data to analyze problems related to the DCFM server or client.

Some of these logs are:

- ▶ **TraceDump:** Dumps a copy of its memory and pointers into a trace file.
- ▶ **RASLOG:** Contains debug data from the switch.
- ▶ **supportShow:** Gathers configuration and status information from the switch.

17.2.1 Capturing a trace dump

When a switch “panics,” depending upon the circumstances, it might produce a trace dump, which can be uploaded automatically to an FTP server when the switch recovers from this failure.

From within Web Tools select the **Switch Admin** interface and expand the view to the switch with **Show Advanced Mode**. The Trace tab allows you to view and configure the FTP host target, enable or disable automatic trace uploads, and update a trace dump manually as shown in Figure 17-21.

Tracing is always on and generates a trace dump whenever there are certain actions within the switch, for example:

- ▶ Tracing is triggered manually through the **traceDump** command.
- ▶ A critical level log message occurs.
- ▶ A particular log message occurs because the **traceTrig** command has been used.
- ▶ A kernel panic occurs.
- ▶ A hardware watchdog timer expires.

The trace dump is maintained on the switch until it is uploaded through FTP, or until another trace dump is generated. Be aware that a new trace dump overwrites the previous trace dump.

SAN32B-E4-2 - Switch Administration

SwitchName: SAN32B-E4-2 DomainID: 56(0x38) VVWN: 10:00:00:05:1e:54:16:53 Fri Nov 05 2010 15:54:31 GMT+00:00

SNMP Configure Routing Extended Fabric AAA Service **Trace** FICON CUP Security Policies

Switch Network Firmware Download License User Trunking

Trace FTP Host

Host IP: 10.18.235.1 Remote Directory: /logs
 User Name: joselainen Password:

Trace Dump Availability

Trace dump generation time: Tue Oct 19 18:27:00 2010
 Trace Auto FTP Uploaded: ☐

Auto FTP Upload

☒ Enable ☐ Disable

Apply Close Refresh

[Switch Administration opened]: Fri Nov 05 2010 15:53:31 GMT+00:00

Enable Auto FTP upload Mode: Advanced [Free Professional Management Tool](#) 10.18.235.56 AD0 User: admin Role: admin ✓

Figure 17-21 Trace

17.2.2 The supportsave command

This command **supportsave** allows the manual upload of the following logs to an FTP server:

- ▶ **RASLOG**
- ▶ **TRACEdump**
- ▶ **supportshow**
- ▶ zone log
- ▶ RCS command log
- ▶ NS event log
- ▶ FSPF status log
- ▶ Any memory CORE files

The command structure, from the CLI, is as follows:

```
supportsave [-n] [-c] [-k] [-u] user_name [-p] password [-h] host_ip  
[-d] remote_dir [-I] protocol [-R] [-U] [-t] timeout-multiplier
```

With the new option **-t** you now have the possibility to extend the timeout value of the command **supportsave**.

Example 17-1 shows partial output from the **supportsave** command.

Example 17-1 Output from supportsave command

```
SAN32B-E4-1:admin> supportsave  
This command collects RASLOG, TRACE, supportShow, core file, FFDC data  
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] yes  
  
Host IP or Host Name: 10.18.228.151  
User Name: Uwe  
Password:  
Protocol (ftp or scp): ftp  
Remote Directory: .  
  
Saving support information for switch:SAN32B-E4-1, module:RAS...  
Saving support information for switch:SAN32B-E4-1, module:CTRACE_OLD...  
Saving support information for switch:SAN32B-E4-1, module:CTRACE_NEW...  
Saving support information for switch:SAN32B-E4-1, module:FABRIC...  
.....  
..... (some line delete for a better overview)  
.....  
Saving support information for switch:SAN32B-E4-1, module:MAPS...  
Saving support information for switch:SAN32B-E4-1,  
module:FABRIC_WATCH...  
Saving support information for switch:SAN32B-E4-1,  
module:DM_FTR_FFDC...  
Saving support information for switch:SAN32B-E4-1, module:PSDUMP...  
Saving support information for switch:SAN32B-E4-1, module:CORE_FFDC...  
No core or FFDC data files found!  
Saving support information for switch:SAN32B-E4-1, module:ENC_LOGGER...  
Saving support information for switch:SAN32B-E4-1, module:RAS_POST...  
  
SupportSave completed.
```

All files will be saved to the directory that you choose during the **supportsave** command as an ftp directory. You have to pack all files in a .zip file and upload it to the support center when needed.

To capture technical support information using DCFM, follow these steps:

1. Select **Monitor** → **Technical Support** → **Switch / Host Supportsave** (Figure 17-22).




Figure 17-22 Technical Support

2. Select a switch, (as shown in Figure 17-23). You can also choose whether to collect data from more than one switch.




Figure 17-23 Select switch

3. A confirmation panel displays, as shown in Figure 17-24, and warns that the capture might be time intensive. Click **OK**.

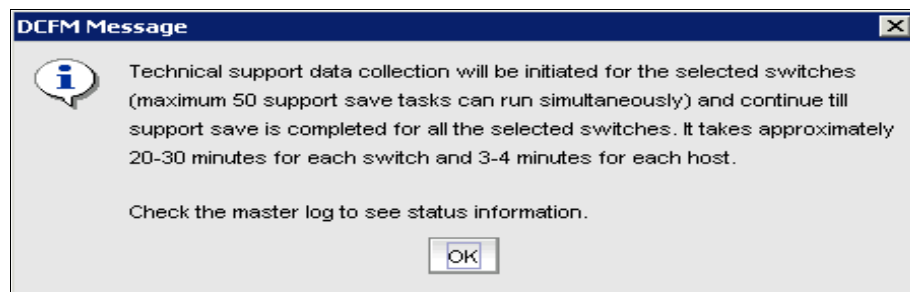
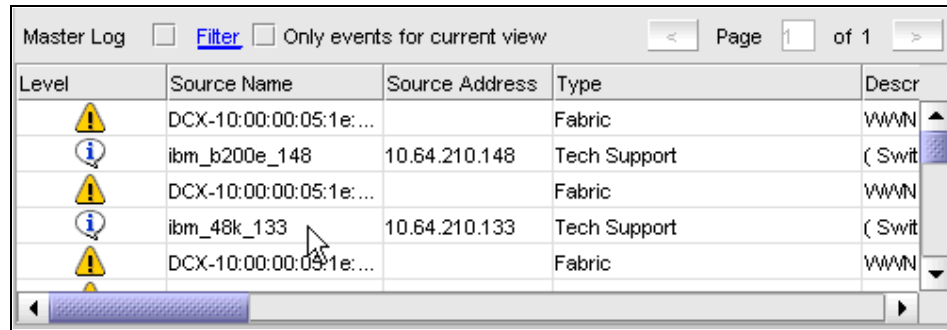


Figure 17-24 Confirmation panel

- To see if the process has completed successfully, go to the Master log tab and look for the information message, as shown in Figure 17-25.



Level	Source Name	Source Address	Type	Descr
Warning	DCX-10:00:00:05:1e:...		Fabric	WWN
Information	ibm_b200e_148	10.64.210.148	Tech Support	(Swit
Warning	DCX-10:00:00:05:1e:...		Fabric	WWN
Information	ibm_48k_133	10.64.210.133	Tech Support	(Swit
Warning	DCX-10:00:00:05:1e:...		Fabric	WWN

Figure 17-25 Information message

These logs can now be sent to the SAN hardware support team at IBM for further diagnosis.

- Viewing technical support information:

To view the captured information, select **Monitor** → **Technical Support** → **View Repository**.

The repository window opens (see Figure 17-26) and shows the captured data in zip files.

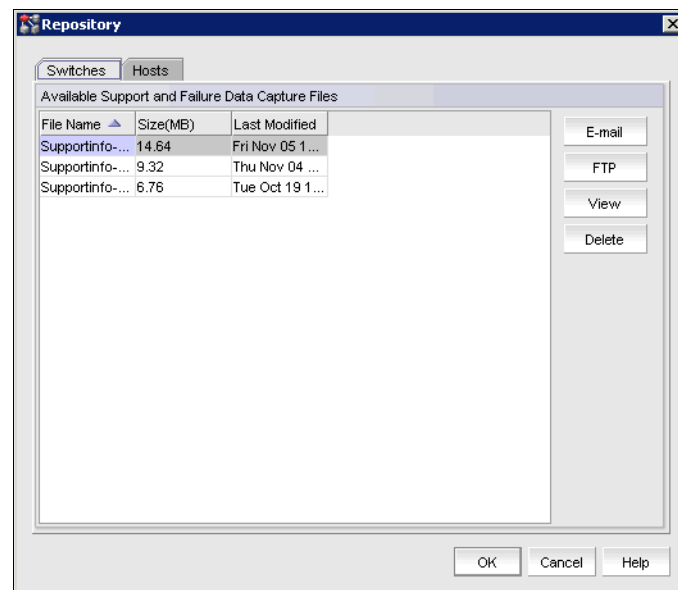


Figure 17-26 Repository window

17.2.3 DCFM support information

As already mentioned you can collect data to analyze a problem related to DCFM itself. Select **Monitor** → **Technical Support** → **Client SupportSave** as shown in Figure 17-27.

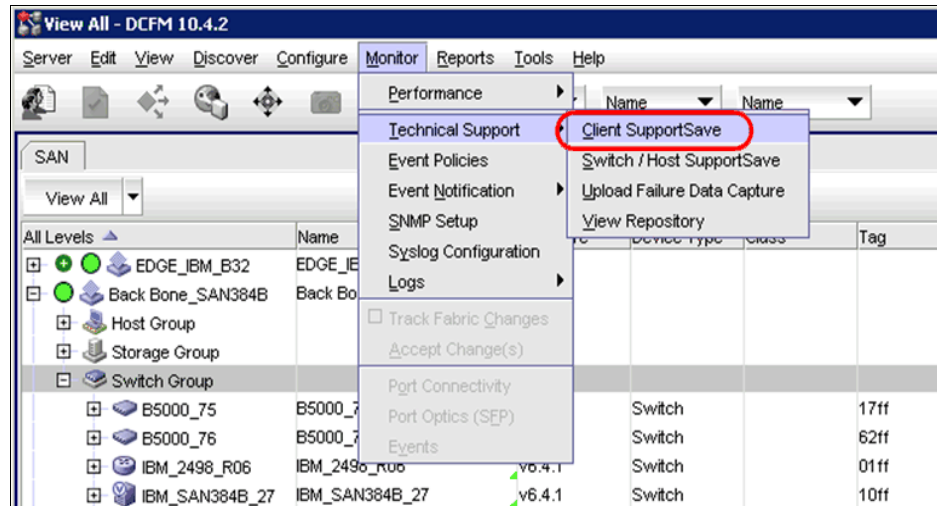


Figure 17-27 Support data from DCFM server

Be aware that this option is also available on the DCFM client side.

You will get an information message as to where the .zip file is stored as shown in Figure 17-28. You can now upload the file if needed.




Figure 17-28 DCFM location message

17.3 General troubleshooting

You can perform the following operations using FC troubleshooting:

- ▶ Trace Route (Path Information and FC Ping): Use to obtain the detailed routing information for any two selected device ports.

- **Device Connectivity Troubleshooting:** Use to identify any problems that might be preventing communication between the two selected device ports. The device ports can be selected from the same fabric or from two different fabrics.

17.3.1 Troubleshooting device connectivity

1. Select **Configure** → **FC Troubleshooting** → **Device Connectivity**. The Device Connectivity Troubleshooting dialog box displays as shown in Figure 17-29.

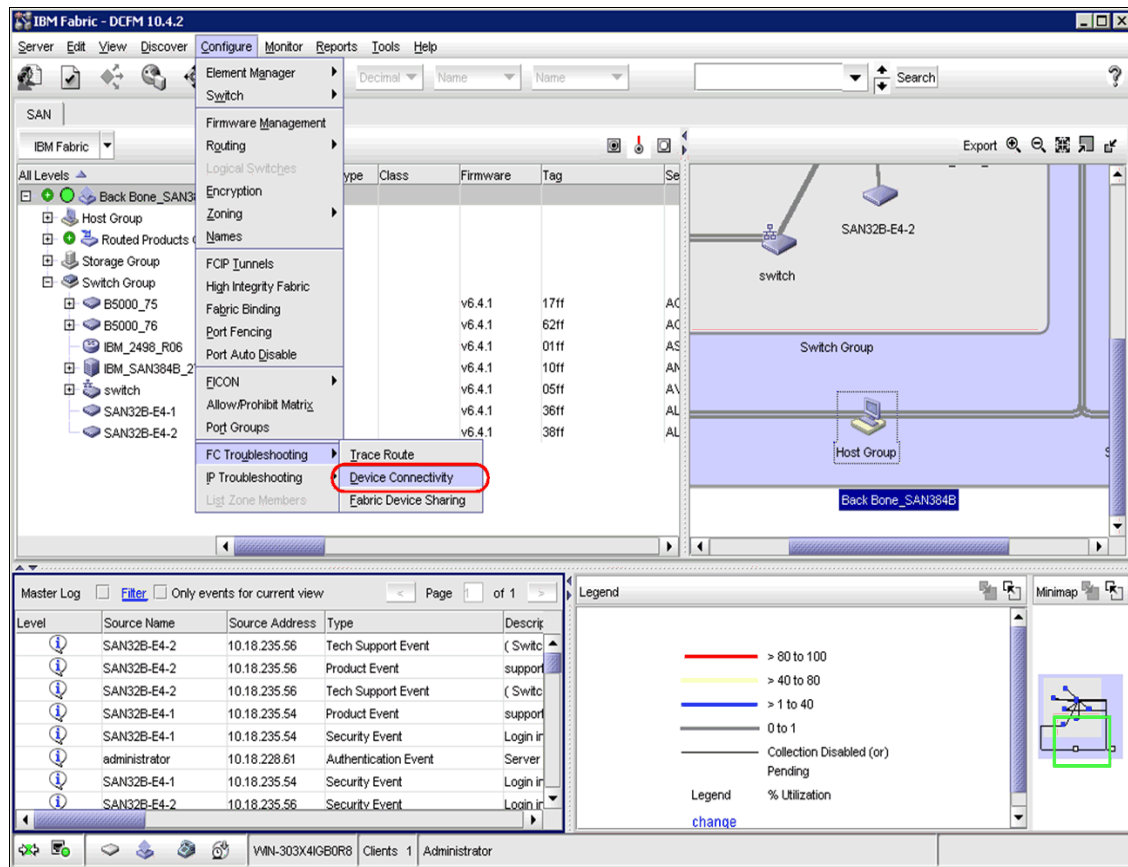


Figure 17-29 Selecting Device Connectivity

2. The wizard opens, as shown in Figure 17-30. The dialog is self explanatory. Select the device ports you want to troubleshoot, and click **OK**.
3. A panel displays, as shown in Figure 17-31, showing the Checks performed.

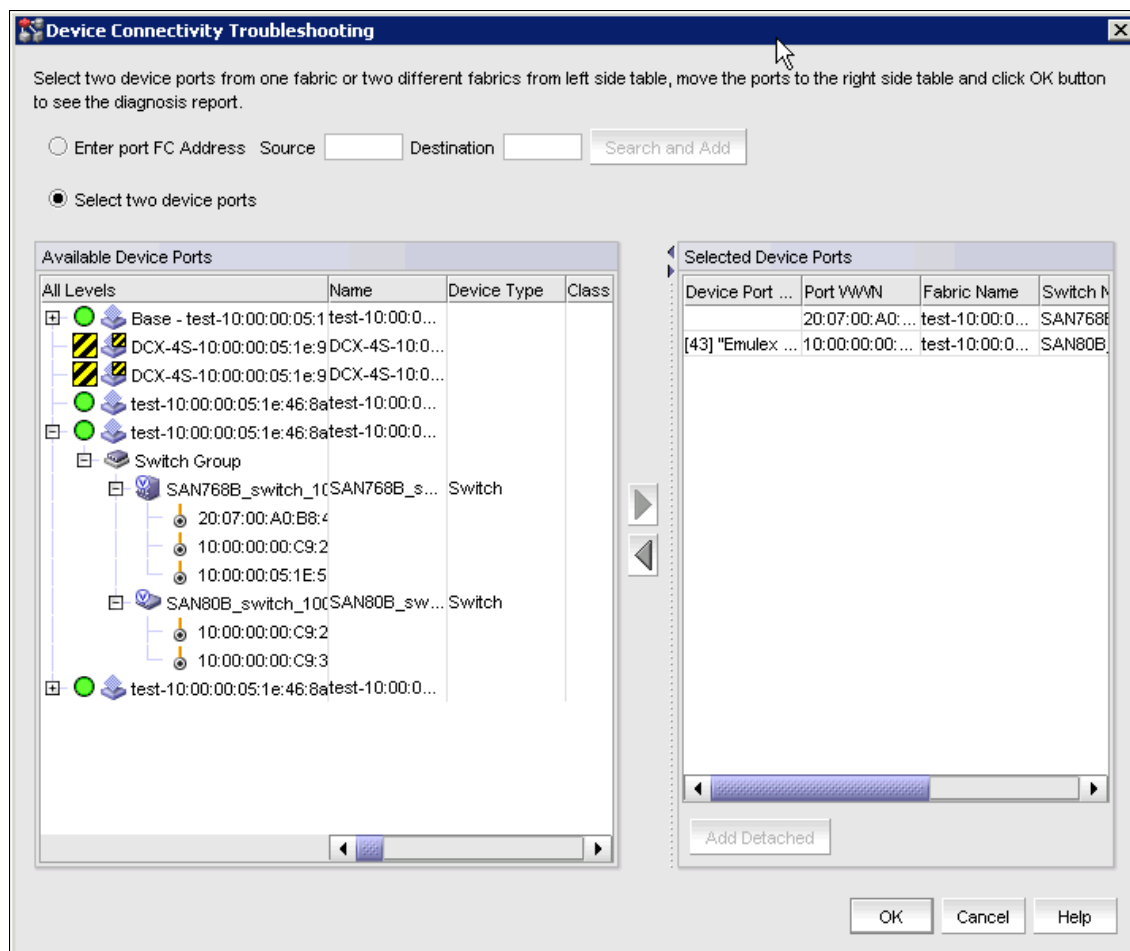


Figure 17-30 Wizard start




Figure 17-31 Performing Diagnosis

- The next panel displays a summary of the tests performed and their results, see Figure 17-32. You can either rerun the Checks, Trace Route the same ports, or Close to finish Troubleshooting.

Device Connectivity Troubleshooting Results

Summary

	Port ...	Port Name	Connected S...	Fabric
Device Port 1	1:07:00:A0:B8:48:58:A1		768B_switch_100	10:00:00:05:1e:46:8a:01
Device Port 2	1:00:00:00:C9:28:EC:1A	"Emulex LP952 FV3.92A2 DV5-5.31A0 4COMPAQ040"	180B_switch_100	10:00:00:05:1e:46:8a:01

Diagnosis Time: Sun Jun 07 00:38:41 PDT 2009

Failure: 2
Warning: 0
Success: 4
Information: 6
Total: 12

Details

Level	Diagnosis Test	Result	Suggested Resolution
✖	Zone configuration check	The zone configuration is not defined properly and the device port(s) are found to be members of different zones.	The device ports should be members of the same zone in the effective zone config.
✖	Zone configuration check	The zone configuration is not defined properly and the device port(s) are found to be members of different zones.	The device ports should be members of the same zone in the effective zone config.
ℹ	LSAN zone configuration check.	Selected Devices are from same fabric - test-10:00:00:05:1e:46:8a:01. No need for LSAN check.	None.
ℹ	LSAN zone configuration check.	Selected Devices are from same fabric - test-10:00:00:05:1e:46:8a:01. No need for LSAN check.	None.

Re-run Diagnosis Trace Route Close

Figure 17-32 Troubleshooting Results

Tip: The errors shown previously were forced by choosing devices not in the same zone.

17.3.2 Trace route

The trace route feature of DCFM combines the functionality of the **pathInfo** and **fcPing** commands.

Trace route information

Trace route displays detailed routing information from a source port or area on the local switch to a destination port or area on another switch. This routing information describes the exact path that a user data stream takes to go from the source port to the destination port, including all intermediate switches.

Attention: DCFM cannot capture the routing information if any of the switches in the path are running Fabric OS v2.x or XPath OS.

The route information depends on the state of the intermediate switches and their ports. The path obtained for two ports might not be the same at all times. Also, the reverse path might not be the same as the forward path.

If one of the ports is inactive, the path shown is the path if the port was active.

Trace route performs a zoning check between the source and destination ports and displays whether the selected device port worldwide names (WWNs) are part of an active zone configuration. Note that if the selected device port WWNs are part of a zone that is not active, then the trace route displays that the device ports are not zoned.

Trace route also displays the maximum, minimum, and average round trip time for the data between the device port WWNs and the domain controller.

Support: Trace route is only supported on Fabric OS switches running Fabric OS v5.2 or later.

Capturing trace route information

Follow these steps:

1. Select **Configure** → **FC Troubleshooting** → **Trace Route**, as shown in Figure 17-33.




Figure 17-33 Trace route

2. The *Trace Route* dialog box opens, as shown in Figure 17-34. Select two devices from the *Available Device Ports* panel and move them to the *Selected devices Ports* panel. Alternatively, you can search for devices either by device port WWN or device port name using the *Search and Add* panel. Click **OK** to start Trace Routing.

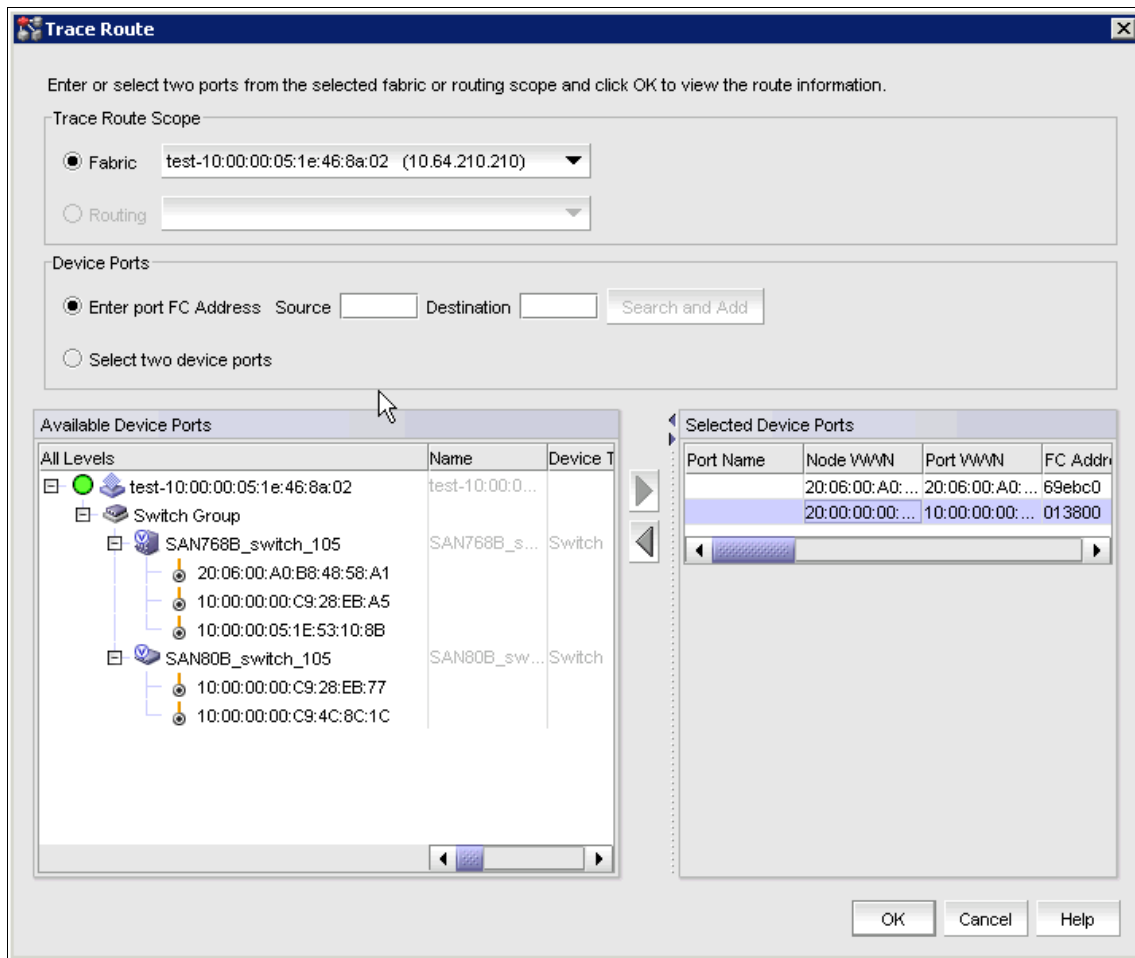


Figure 17-34 Select devices

3. The Trace Route Summary dialog box displays Figure 17-35 with the following information about the different tabs:

Trace Route Summary tab: This tab shows a brief summary of the trace including the port WWN, port name, FC address, switch name, whether ping was successful, round trip time (minimum, maximum, and average) and whether the device ports are in active zones.

Forward Route tab: This tab shows the path taken by data packets from the port belonging to the switch on which the trace route has been invoked (source port) to the port on the other switch (destination port).

Reverse Route tab: This tab shows the path from the destination port to the source port.

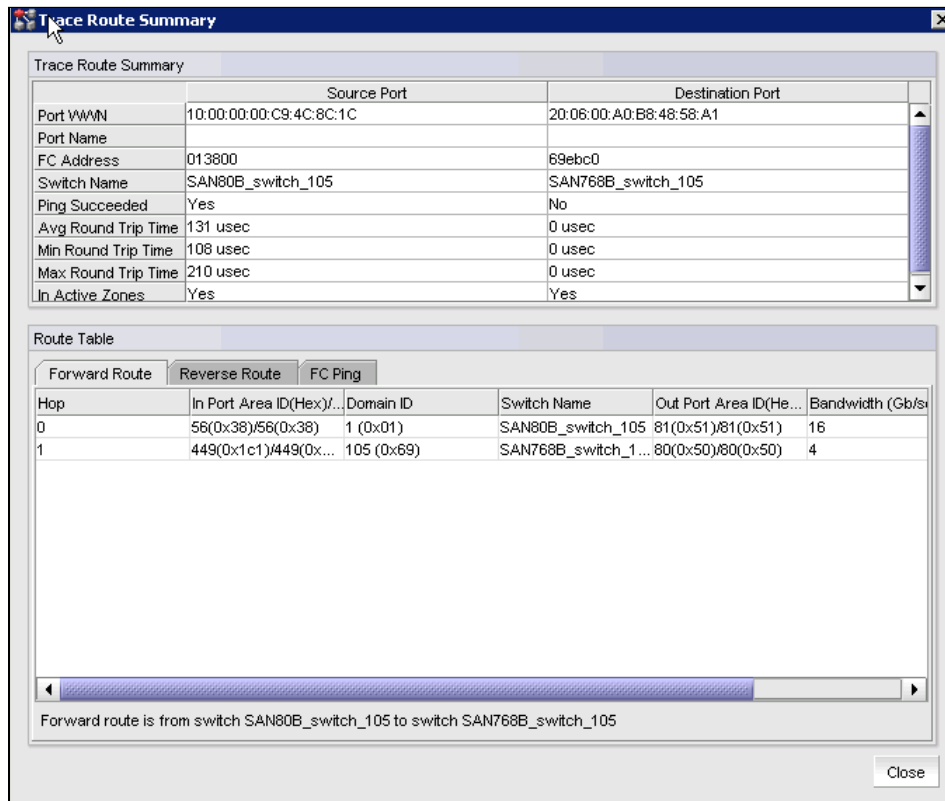


Figure 17-35 Trace Route Summary dialog box

17.4 Port Fencing

Fabric Operating System v6.1 (Fabric OS) adds a new feature called Port Fencing, which disables a port automatically when it is operating outside the bounds of normal operation.

Port Fencing allows you to protect your SAN from repeated operational or security problems experienced by ports. Use Port Fencing to set threshold limits for the number of specific port events permitted during a given time period on the selected object.

Port Fencing objects include the SAN, Fabrics, Directors, Switches (physical), Virtual Switches, Ports, as well as Port Types (E_port, F_port, and FX_port). Use Port Fencing to directly assign a threshold to these objects. When a switch does not support Port Fencing, a No Fencing Changes message displays in the Threshold field in the Ports table.

If the port detects more events during the specified time period, the device firmware blocks the port, disabling transmit and receive traffic until you investigate, solve the problem, and manually unblock the port.

The Port Fencing requirements include:

- ▶ Based on Fabric Watch Event Monitoring
- ▶ Requires Fabric Watch license to configure and use

17.4.1 Port Fencing using DCFM

In DCFM you can create thresholds, which you can then assign to available objects in the tree.

Port Fencing threshold types include these:

- ▶ Invalid CRCs (Fabric OS only):
Used to block a port when an Invalid CRCs violation meets the Fabric OS switch threshold.
- ▶ Invalid Words (Fabric OS only):
Use to block a port when an Invalid Words violation meets the Fabric OS switch threshold.
- ▶ Link (M-EOS and Fabric OS):
Used to block a port when a Link Level (Hot I/O) error meets the threshold. A Link Level (Hot I/O) occurs when an active loop port repeatedly receives a loop initialization primitive sequence error or an active non-loop port repeatedly receives a line repeater, offline sequence, or not operational sequence error.
- ▶ Link Reset (Fabric OS only):
Used to block a port when the link timeout errors meet the threshold.
- ▶ Protocol Errors (M-EOS and Fabric OS):
Used to block a port when one of these protocol errors meets the threshold:
 - ISL Bouncing-ISL has repeatedly become unavailable due to link down events.
 - ISL Segmentation (M-EOS only)-ISL has repeatedly become segmented.
 - ISL Protocol Mismatch-ISL has been repeatedly put into the Invalid Attachment state due to a protocol error.

► Security (M-EOS):

Used to block a port when one of the following security violations occurs:

- Authentication: The switch has repeatedly become unavailable due to authentication events.
- Fabric Binding: The switch has repeatedly become unavailable due to fabric binding events.
- Switch Binding: The switch has repeatedly become unavailable due to switch binding events. Switch Binding is enabled through a product's Element Manager.
- Port Binding: The switch has repeatedly become unavailable due to port binding events.
- ISL Security: (Generic Security Error) the switch on the other side of the ISL has detected a specific security violation, but is only able to indicate that a generic security violation has occurred or a security configuration mismatch was detected.
- N_port Connection Not Allowed-the switch has repeatedly become unavailable due to N_port connection not allowed events.

► Sync Loss (Fabric OS only):

Used this type of threshold to block a port when a sync loss violation type meets the Fabric OS switch threshold.

Adding Thresholds for Port Fencing

In this example we set a *CRC Threshold* of 1 CRC error per minute to the E-ports of a Fabric.

1. Select **Configure** → **Port Fencing** as shown in Figure 17-36. The Port Fencing dialog box displays.

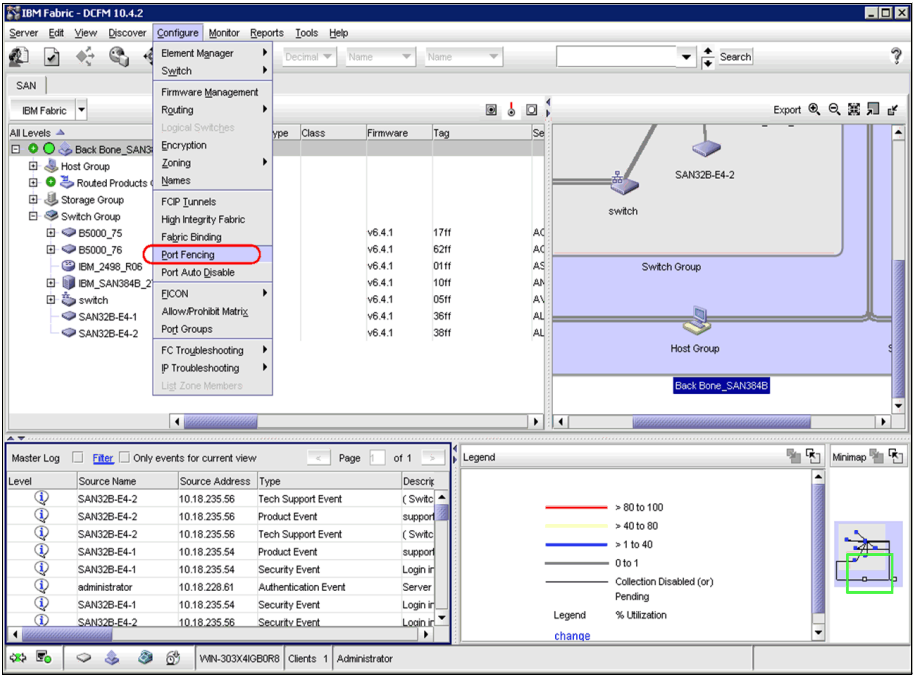


Figure 17-36 Select Port Fencing

2. Select **Invalid CRCs (FOS only)** from the Violation Type list in the The Port Fencing dialog box (see Figure 17-37). Click **Add**.

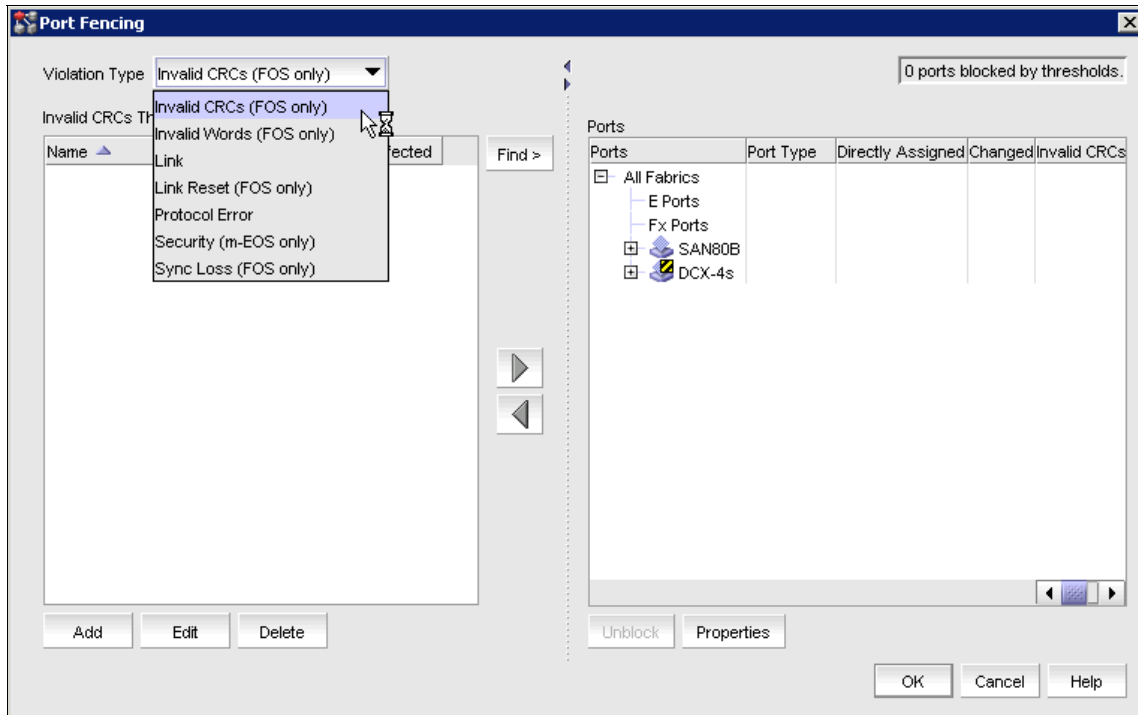


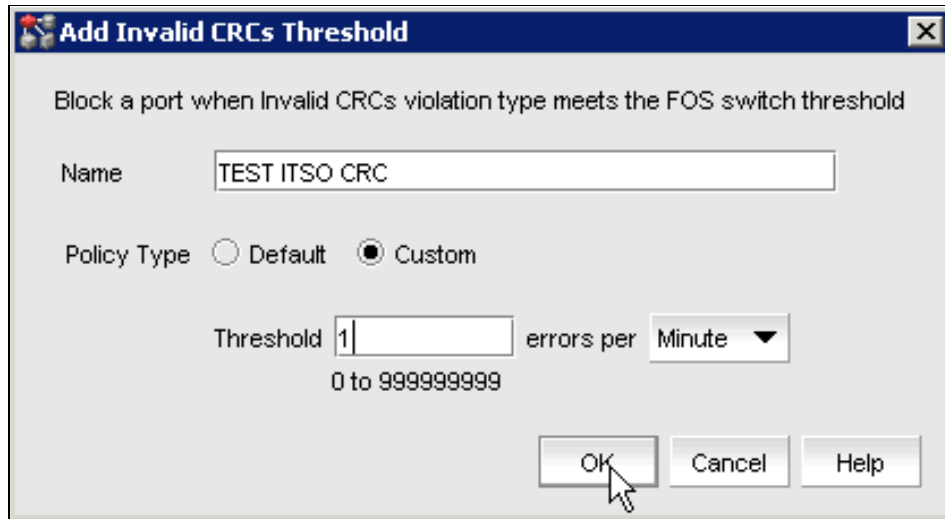
Figure 17-37 Port Fencing dialog box

The *Add Invalid CRCs Threshold* dialog box displays Figure 17-38. Enter a name for the threshold in the Name field (in our example *TEST ITSO CRC*) and select one of the following options:

- ▶ Default: Uses device defaults
- ▶ Custom: Uses your selections

Enter the number of invalid CRCs allowed for the threshold in the Threshold errors field. We use Custom in our example and use 1 error per minute.

Click **OK** to add the Invalid CRCs threshold to the table and close the *Add Invalid CRCs Threshold* dialog box.



The dialog box is titled "Add Invalid CRCs Threshold" and contains the following elements:

- A title bar with a close button (X) in the top right corner.
- A subtitle: "Block a port when Invalid CRCs violation type meets the FOS switch threshold".
- A "Name" label followed by a text input field containing "TEST ITSO CRC".
- A "Policy Type" section with two radio buttons: "Default" (unselected) and "Custom" (selected).
- A "Threshold" label followed by a text input field containing "1". Below this field is the range "0 to 999999999".
- An "errors per" label followed by a dropdown menu currently set to "Minute".
- Three buttons at the bottom right: "OK", "Cancel", and "Help". A mouse cursor is pointing at the "OK" button.

Figure 17-38 *Add Invalid CRCs Threshold* dialog box

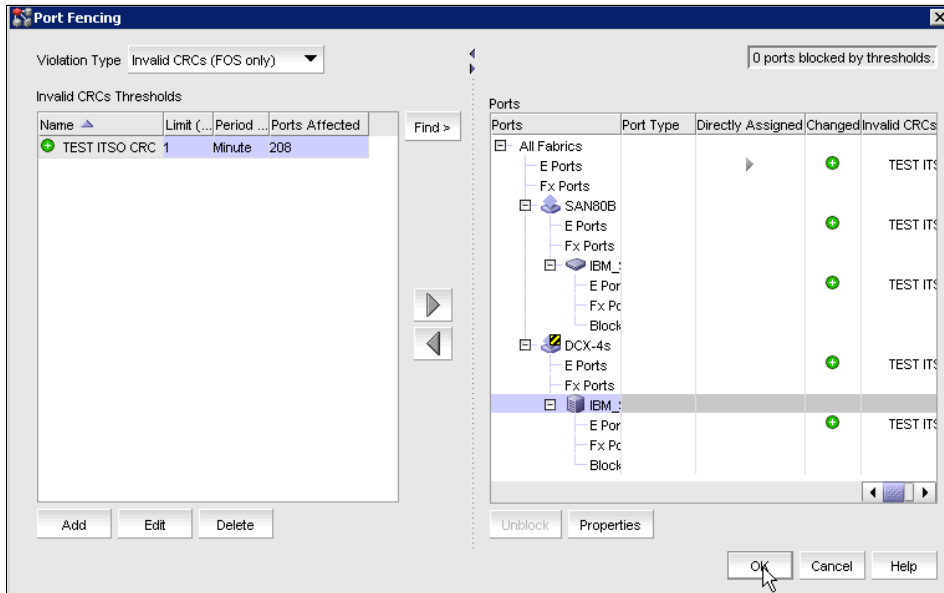


Figure 17-40 Port Fencing dialog box

Removing Thresholds for Port Fencing

In this example we remove Port Fencing thresholds:

1. Select **Configure** → **Port Fencing**. See Figure 17-36 on page 822. The Port Fencing dialog box displays.
2. Select a threshold type from the Violation Type list. See Figure 17-37 on page 823.
3. Select the object with the threshold that you want to remove in the Ports table of the Port Fencing dialog box (Figure 17-41). Click the left arrow.

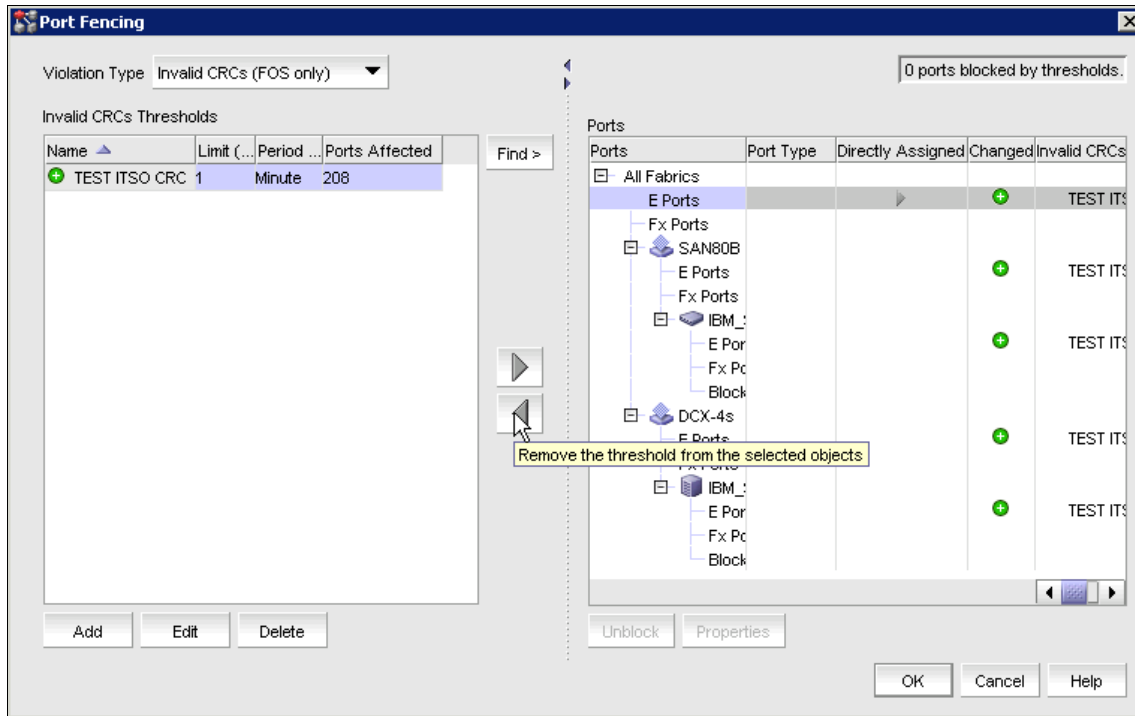


Figure 17-41 Port Fencing dialog box

- Click **OK**.

Unblocking a Fenced port

The management application allows you to unblock a port (only if it was blocked by Port Fencing) after the problem that triggered the threshold is fixed. When a port is blocked, an Attention icon displays next to the port node:



To unblock a port, complete the following steps:

- Select **Configure** → **Port Fencing** (see Figure 17-36 on page 822). The Port Fencing dialog box displays.
- Right-click anywhere in the Ports table and select **Expand**.
- Select a blocked port from the Ports table and click **Unblock**.
- Click **OK** on the message. If you did not solve the root problem, the threshold will trigger again.
- Click **OK** on the Port Fencing dialog box.

17.4.2 Port Fencing using CLI

In the examples that follow, we show the steps to configure Port Fencing for Link Loss on the E-Port class:

1. Telnet to the switch.
2. Enter **fwconfigure** as shown in Example 17-2.

Example 17-2 Step 1 to configure Port Fencing

```
b5000_147:admin> fwconfigure

1 : Environment class
2 : SFP class
3 : Port class
4 : Fabric class
5 : E-Port class
6 : F/FL Port (Optical) class
7 : Alpa Performance Monitor class
8 : EE Performance Monitor class
9 : Filter Performance Monitor
10 : Security class
11 : Resource class
12 : Quit
Select a class => : (1..12) [12]
```

3. Type 5 to select E_Port Class as shown in Example 17-3.

Example 17-3 Step 3 to configure port fencing

```
b5000_147:admin> fwconfigure

1 : Environment class
2 : SFP class
3 : Port class
4 : Fabric class
5 : E-Port class
6 : F/FL Port (Optical) class
7 : Alpa Performance Monitor class
8 : EE Performance Monitor class
9 : Filter Performance Monitor class
10 : Security class
11 : Resource class
12 : Quit
Select a class => : (1..12) [12] 5
```

Entered into the swthcfgEPort

```
1 : Link loss(E-port)
2 : Sync loss(E-port)
3 : Signal loss(E-port)
4 : Protocol error(E-port)
5 : Invalid words(E-port)
6 : Invalid CRCs(E-port)
7 : RXPerformance(E-port)
8 : TXPerformance(E-port)
9 : State Changes(E/VE-port)
10 : Link reset(E-port)
11 : Utilization(VE-port)
12 : Packet Loss(VE-port)
13 : C3 Discard(E-port)
14 : return to previous page
Select an area => : (1..14) [14]
```

4. Type 1 to select the Link Loss as shown in Example 17-4.

Example 17-4 Step 4 to configure port fencing

Entered into the swthcfgEPort

```
1 : Link loss(E-port)
2 : Sync loss(E-port)
3 : Signal loss(E-port)
4 : Protocol error(E-port)
5 : Invalid words(E-port)
6 : Invalid CRCs(E-port)
7 : RXPerformance(E-port)
8 : TXPerformance(E-port)
9 : State Changes(E/VE-port)
10 : Link reset(E-port)
11 : Utilization(VE-port)
12 : Packet Loss(VE-port)
13 : C3 Discard(E-port)
14 : return to previous page
Select an area => : (1..14) [14] 1
```

Index	ThresholdName	Port	CurVal	Status
	LastEvent	LasteventTime	LastVal	LastState
=====				
0	eportLink000		0 0 Error(s)/min	enabled
	inBetween	Tue May 6 16:44:19 2008	0 Error(s)/min	In_Range

```

1 eportLink001          1 0 Error(s)/min    enabled
inBetween    Tue May  6 16:46:25 2008 0 Error(s)/min    In_Range

```

```

1 : refresh
2 : disable a threshold
3 : enable a threshold
4 : advanced configuration
5 : return to previous page
Select choice => : (1..5) [5]

```

5. Type 4 to select advanced configuration as shown in Example 17-5.

Example 17-5 Step 5 to configure port fencing

```

1 : refresh
2 : disable a threshold
3 : enable a threshold
4 : advanced configuration
5 : return to previous page
Select choice => : (1..5) [5] 4

```

Index	ThresholdName	BehaviorType	BehaviorInt
0	eportLink000	Triggered	1
1	eportLink001	Triggered	1

Threshold boundary level is set at : Default

	Default	Custom
Unit	Error(s)	Error(s)
Time base	minute	minute
Low	0	0
High	5	5
BufSize	0	0

Threshold alarm level is set at : Default

Errlog-1, SnmpTrap-2, PortLogLock-4
 RapiTrap-8, EmailAlert-16, PortFencing-32

Valid alarm matrix is 63

	Default	Custom
Changed	0	0
Below	0	0
Above	0	0

```

InBetween          0          0

1 : change behavior type          11 : change threshold alarm level
2 : change behavior interval      12 : change changed alarm
3 : change threshold boundary level 13 : change below alarm
4 : change custom unit           14 : change above alarm
5 : change custom time base      15 : change inBetween alarm
6 : change custom low           16 : apply threshold alarm
changes
7 : change custom high          17 : cancel threshold alarm
changes
8 : change custom buffer        18 : return to previous page
9 : apply threshold boundary changes
10 : cancel threshold boundary changes
Select choice => : (1..18) [18]

```

6. Type 14 to enable port fencing as shown in Example 17-6.

Example 17-6 Step 6 to enable the port fencing

```

1 : change behavior type          11 : change threshold alarm level
2 : change behavior interval      12 : change changed alarm
3 : change threshold boundary level 13 : change below alarm
4 : change custom unit           14 : change above alarm
5 : change custom time base      15 : change inBetween alarm
6 : change custom low           16 : apply threshold alarm
changes
7 : change custom high          17 : cancel threshold alarm
changes
8 : change custom buffer        18 : return to previous page
9 : apply threshold boundary changes
10 : cancel threshold boundary changes
Select choice => : (1..18) [18] 14

```

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16, PortFencing-32

Valid alarm matrix is 63
Enter above alarm matrix => : (0..63) [0]

7. Type 32 to enable port fencing as shown in Example 17-7.

Example 17-7 Step 7 to enable port fencing

```

1 : change behavior type          11 : change threshold alarm level
2 : change behavior interval      12 : change changed alarm
3 : change threshold boundary level 13 : change below alarm

```

```

4 : change custom unit           14 : change above alarm
5 : change custom time base      15 : change inBetween alarm
6 : change custom low           16 : apply threshold alarm
changes
7 : change custom high          17 : cancel threshold alarm
changes
8 : change custom buffer         18 : return to previous page
9 : apply threshold boundary changes
10 : cancel threshold boundary changes
Select choice => : (1..18) [18] 14

```

```

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16, PortFencing-32

```

```

Valid alarm matrix is 63
Enter above alarm matrix => : (0..63) [0] 32

```

8. Type 3 to change the threshold boundary level to custom by selecting 2 as shown in Example 17-8.

Example 17-8 Step 8 to configure port fencing

```

Select choice => : (1..18) [18] 3
1 : Default
2 : custom
Enter boundary level type => : (1..2) [1] 2
Index ThresholdName BehaviorType BehaviorInt
0 eportLink000 Triggered 1
1 eportLink001 Triggered 1

```

Threshold boundary level is set at : Custom

	Default	Custom
Unit	Error(s)	Error(s)
Time base	minute	minute
Low	0	0
High	5	5
BufSize	0	0

Threshold alarm level is set at : Custom

```

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16, PortFencing-32

```

```

Valid alarm matrix is 63

```


	Default	Custom
Changed	0	0
Below	0	0
Above	0	32
InBetween	0	0


```

1 : change behavior type          11 : change threshold alarm level
2 : change behavior interval      12 : change changed alarm
3 : change threshold boundary level 13 : change below alarm
4 : change custom unit           14 : change above alarm
5 : change custom time base      15 : change inBetween alarm
6 : change custom low           16 : apply threshold alarm
changes
7 : change custom high          17 : cancel threshold alarm
changes
8 : change custom buffer        18 : return to previous page
9 : apply threshold boundary changes
10 : cancel threshold boundary changes
Select choice => : (1..18) [18]

```

9. Type 9 to apply the threshold boundary changes as shown in Example 17-9.

Example 17-9 Step 9 to configure port fencing

```

1 : change behavior type          11 : change threshold alarm level
2 : change behavior interval      12 : change changed alarm
3 : change threshold boundary level 13 : change below alarm
4 : change custom unit           14 : change above alarm
5 : change custom time base      15 : change inBetween alarm
6 : change custom low           16 : apply threshold alarm
changes
7 : change custom high          17 : cancel threshold alarm
changes
8 : change custom buffer        18 : return to previous page
9 : apply threshold boundary changes
10 : cancel threshold boundary changes
Select choice => : (1..18) [18] 9

```

10. Now press Enter until you reach the admin prompt.

17.4.3 Enabling Port Fencing for E_Port class link loss

To enable Port Fencing for E_Port class link loss, proceed as follows:

1. Telnet to the switch.
2. Enter the **fwconfigure** command.
3. Type 5 to select E_Port class.
4. Type 1 to select Link Loss (E_Port).
5. Type 4 to select advanced configuration.
6. Type 14 to select change alarm.
7. Type 32 to enable Port Fencing.
8. Type 3 to change the threshold boundary level, and type 2 to set the threshold boundary level to custom.
9. Type 9 to apply the threshold boundary level changes.
10. Press Enter until you reach the admin prompt.

17.4.4 Testing the configuration

You can simulate the link loss manually by unplugging the ISL cable multiple times. After the threshold has been reached/exceeded, the E_Port is disabled automatically.

Example 17-10 shows output for the **switchshow** command where port 1 is disabled automatically with the following error message:

“Disabled (Port Link Loss threshold exceeded)”

Now, you need to enable the port manually using the **Portenable <portnumber>** command.

Example 17-10 The switchshow command after the port is disabled automatically

```
b5000_147:admin> switchshow
switchName:    b5000_147
switchType:    58.1
switchState:    Online
switchMode:    Native
switchRole:    Principal
switchDomain:   4
switchId:      fffc04
switchWwn:     10:00:00:05:1e:90:14:c7
zoning:        OFF
```

switchBeacon: ON

Area	Port	Media	Speed	State	Proto
=====					
0	0	id	N4	No_Light	
1	1	id	N4	No_Sync	Disabled (Port Link Loss threshold exceeded)
2	2	id	N4	No_Light	
3	3	id	N4	No_Light	
4	4	--	N4	No_Module	
5	5	--	N4	No_Module	
6	6	--	N4	No_Module	
7	7	--	N4	No_Module	
8	8	--	N4	No_Module	
9	9	--	N4	No_Module	
10	10	--	N4	No_Module	
11	11	--	N4	No_Module	
12	12	id	N4	Online	F-Port 20:06:00:a0:b8:48:58:a1
13	13	id	N4	Online	F-Port 20:06:00:a0:b8:48:58:a2

17.4.5 Basic troubleshooting commands

You can use the following commands to perform basic troubleshooting of the SAN switch. Most of the commands are self-explanatory:

chassisShow	Displays all Field Replaceable Units (FRU).and their status.
diagShow	Displays diagnostics status for the switch port.
errDump	Dumps the external error log messages.
errshow	Displays the error log message.
fabStatsShow	Displays the fabric statistics information. Can be used to troubleshoot the Fabric merging issues
fanShow	Displays fan status and speed.
fcPing	Performs a zoning check between the source and destination WWN.
islShow	Displays ISL information.
portCfgShow	Displays port configuration settings.
portErrShow	Displays a summary of port errors.
portLogShow	Displays the port log with page breaks.
portShow	Displays the status of the specified port.
portStatsShow	Displays port hardware statistics.

psShow	Displays power supply status.
supportshow	Prints switch information for debugging purposes.
supportsave	Generate files for the support center for debugging purposes.
switchShow	Displays switch and port status.
topologyShow	Displays the fabric topology.

You can find more help with these commands by using **help <command name>** when logged in to the switch.

Example portErrShow

The **portErrShow** command is a good tool for a quick analysis of switch ports, as it provides an overview of useful (error-) counters.

The counters shown in Example 17-11 reflect the increased values over history of a port.

Example 17-11 output portErrShow

```

MagicC_1:admin> porterrshow
      frames  enc  crc  too  too  bad  enc  disc  link  loss  loss  frjt  fbsy
      tx   rx   in  err  shrt long  eof  out   c3  fail  sync  sig
=====
0:    34   37   0   10   0   0   0  108   0   0   2   3   0   0
1:   9.6m 5.4m   0   0   0   0   0   54   0   1   2   3   0   0
2:   10m 144m   4   0   4   0   9  211k   7   1   3   4   0   0
3:    34   37   0  100k   0   0   0  109   0   0   3   5   0   0
4:   8.1m 2.2m   0   0   0   0   0  85k   0   2  127  210   0   0
5:   80k  51k   0   0   0   0   0  91k   0   0   1   2   0   0
6:  112k 341k   0   0   0   0   0  3.4k   0  16  26  33   0   0
7:  134k 327k   0   0   0   0   0   1   0   8   9  10   0   0
8:    0    0   0   0   0   0   0   0   0   0  142k 284k   0   0
9:   15   15   0   0   0   0   0   4   0   0  43  82   0   0
10:  1.3k 1.0k   0   0   0   0   0   0   0   4   6   7   0   0
11:  1.9m 103k   0   0   0   0   0  31   0   6   7   8   0   0

```

Therefore these counters are not to be used for problem determination because they show errors since the counters have been cleared the last time.

To have a clear **baseline**, first all the counters have to be cleared.

1. Telnet to the switch.
2. Enter the command **statsclear** to clear all counters reflected by the **porterrshow** command. The counters should look similar to those in Example 17-12.

Example 17-12 output portErrShow after counters are cleared

```
MagicC_1:admin> statsclear
MagicC_1:admin> porterrshow
```

	frames		enc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbusy
	tx	rx	in	err	shrt	long	eof	out	c3	fail	sync	sig		
=====														
0:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2:	56	32	0	0	0	0	0	0	0	0	0	0	0	0
3:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
7:	3	3	0	0	0	0	0	0	0	0	0	0	0	0
8:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
9:	2	2	0	0	0	0	0	0	0	0	0	0	0	0
10:	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11:	0	0	0	0	0	0	0	0	0	0	0	0	0	0

3. With this baseline, the ports should be monitored over a proper amount of time (2-3 hours) without any maintenance action on it (such as resetting the link or setting it offline/online).

Explanation of porterrshow counters and their causes

Here we explain the **porterrshow** counters and their causes:

► **frames tx**

The number of transmitted frames of this port.

► **frames rx**

The number of received frames of this port.

► **enc_in**

Encoding error inside frame. The words inside a frame are encoded with the 8bit/10bit encoding scheme. Every 8 bit data byte can be represented by multiple 10 bit characters. If there is a corruption in this encoding in a received frame, the error rises. As the frames are newly encoded every time before they are sent out through a cable, the increasing of this counter means that you have physical problems between this port and the port at the other side of the cable.

► **crc_err**

CRC means *Cyclic Redundancy Check* and is a mechanism to detect corruptions in a frame. The CRC value is calculated in the source device (for example, the HBA) when the frame is composed and will not be altered until it reaches the destination device. Because of that, a corrupted frame

(calculated CRC does not match the CRC value in the frame) will cause an increasing of the `crc_err` counter of every switch it passes on its way to the destination device.

If there are increased `crc_err` on an ISL port without `enc_in` errors, there should be no physical problem on this ISL (at least on the direction towards this port). The action plan, then, is to look for increased `crc_err` on the connected switch and apply the same rule there.

► **too_long**

FC frames are 2148 bytes maximum. If an EOF (end of frame) is corrupted or data generation is incorrect, a `too_long` error is generated.

► **too_short**

The `too_short` is an error statistics counter that is incremental whenever a frame, bounded by an SOF (start of frame word) and EOF, is received and the number of words between the SOF and EOF is less than seven words (six words in the header plus a one word CRC), that is to say 38 bytes (not 48) including the SOF and EOF. This might be caused by the transmitter or an unreliable link.

► **bad_eof**

After a loss-of-synchronization error, continuous-mode alignment allows the receiver to reestablish word alignment at any point in the incoming bit stream while the receiver is Operational. Such realignment is likely (but not guaranteed) to result in Code Violations and subsequent loss of Synchronization. Under certain conditions, it might be possible to realign an incoming bit stream without loss of Synchronization. If such a realignment occurs within a received frame, detection of the resulting error condition is dependent upon higher-level function (for example, invalid CRC, missing EOF Delimiter).

► **enc_out**

This counter is similar to `enc_in`, but for transmission words outside normal FC frames. This error rises for 8 bit/10 bit encoding errors (invalid encoding, invalid disparity, and so on) in ordered sets such as IDLEs or R_RDYs. This counter can increase in huge amounts during speed negotiation and so should not be used alone for problem determination without clearing the statistics before. If you created a baseline first, an increased `enc_out` counter points to probable physical problems. It can be used even if no user-I/O is transferred over the port.

► **disc_c3**

The Discard class 3 counter should increase every time the switch has to drop (discard) a frame because of various reasons. Increases of this counter can point to congestion in the fabric, time-outs, busy destination devices,

unknown devices (for example, if a host sends to an address known to it but not to the fabric), a device sending frames without FLOGI first, an invalid destination address, or others. To determine the reason for the discards, the Fabric OS command **portstatshow** command can be used on current platforms and code versions.

► **link_fail**

If a Port remains in the LR Receive State (for example within the link reset or link initialization phase) for a period of time greater than R_T_TOV (Receiver-Transmitter Timeout Value, 100ms per default), a Link Reset Protocol Timeout will be detected that results in a Link Failure condition and the port enters the NOS (Not Operational State) Transmit State. The link failure also indicates that loss of signal or loss of sync lasting longer than the R_T_TOV value was detected while not in the Offline state. As for loss_sync, verify that the link was not brought down manually after setting the baseline.

► **loss_sync**

Synchronization failures on either bit or Transmission-Word boundaries are not separately identifiable and cause loss-of synchronization errors. There is an internal counter that rises with every invalid transmission word. An invalid transmission word is a word with an encoding error (-> enc_in and/or enc_out will be increased). If the port receives two consecutive valid transmission words, this internal counter is decreased by 1 (to a minimum of 0). If the internal counter reaches 4, a loss-of-synchronization error will be triggered and the loss-of-synchronization procedure will be started. If this happens after a baseline is set and without any maintenance action on the link/switch or on the device (such as reboot of the host), this counter indicates an unstable link and therefore physical problems.

► **loss_sig**

This occurs when a signal is transmitted but nothing is being received on the same port. Check if the connected device is powered on and cabled correctly. The Fabric OS command **switchshow** shows if there is no light on the port, **sfpshow** reveals the exact receive power values. Be aware that many devices require a configuration activation of their HBA in order to send light.

► **frjt**

If the fabric cannot process a class 2 frame, a F_RJT (fabric reject) is returned. Possible reasons can be (among others): class not supported, invalid source id, invalid destination id, N_Port permanently not available, N_Port temporary not available, Login required.

► **fbsy**

If the fabric cannot deliver a class 2 frame within E_D_TOV (Error Detect Timeout Value), the frame will be discarded and a F_BSY (fabric busy) is returned.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks publications

For information about ordering these publications, see “Help from IBM” on page 843. Note that some of the documents that we reference here might be available in softcopy only.

- ▶ *Introduction to Storage Area Networks*, SG24-5470
- ▶ *IBM TotalStorage: SAN Product, Design, and Optimization Guide*, SG24-6384
- ▶ *IBM System Storage/Brocade Multiprotocol Routing: An Introduction and Implementation*, SG24-7544
- ▶ *FICON Implementation Guide*, SG24-6497

Other resources

These publications are also relevant as further information sources:

- ▶ Clark, Tom. *IP SANs: An Introduction to iSCSI, iFCP, and FCIP Protocols for Storage Area Network*. Addison-Wesley Professional, first edition, December 2001. ISBN 0201752778.
- ▶ Farley, Marc. *Building Storage Networks*. McGraw-Hill/Osborne Media, first edition, January 2000. ISBN 0072120509.
- ▶ *Fabric OS Administrator's Guide*, 53-1000448
- ▶ *Secure Fabric OS Administrator's Guide*, 53-1000244

Referenced websites

These websites are also relevant as further information sources:

- ▶ IBM System Storage hardware, software, and solutions:
<http://www.storage.ibm.com>
- ▶ IBM System Storage, Storage Area Network:
<http://www.storage.ibm.com/snetwork/index.html>
- ▶ Brocade:
<http://www.brocade.com>
- ▶ Finisar:
<http://www.finisar.com>
- ▶ Veritas:
<http://www.veritas.com>
- ▶ Tivoli:
<http://www.tivoli.com>
- ▶ JNI:
<http://www.Jni.com>
- ▶ IEEE:
<http://www.ieee.org>
- ▶ Storage Networking Industry Association:
<http://www.snia.org>
- ▶ SCSI Trade Association:
<http://www.scsita.org>
- ▶ Internet Engineering Task Force:
<http://www.ietf.org>
- ▶ American National Standards Institute:
<http://www.ansi.org>
- ▶ Technical Committee T10:
<http://www.t10.org>
- ▶ Technical Committee T11:
<http://www.t11.org>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

10-Gbps interoperability 115
2005-B16 41, 87
2109-F16 License Administration 268
2109-M48 41, 78, 85
2498-B40 87
2498-B80 88
2499-192 73
8 Gbps Fibre Channel connectivity 12
8 Gbps license 171

A

AAA 17, 289
AAA Service 288
AAA Service tab 287
Aborting a transaction associated with IP Filter 693
Accelerator for FICON 96
Access Control List (ACL) 291
Access Gateway 93–94, 115, 196
Access Gateway mode 4
Access Gateway-mode switches 357
access level 280
account information 627
account management 627
ACL 638
ACL policies 662
ACL policy management 662
Activate policy changes 673
Activating an IP Filter policy 688
activating discovery 360
activation key 270
active configuration 517
active CP 137, 148
active CP blade 136–137, 139–140, 148, 151–152
Active Directory 287–288
Active Directory Services 17
Active Directory/LDAP 17
AD 474
AD255 770
Adaptive Networking 97, 114, 170, 172, 695
Adaptive Networking Services 13
AD-aware switch 221
Adding a rule to an IP Filter policy 692
Adding devices (members) to a zone 578
adding filter-based monitors 763
adding new users 276
Adding zones (members) to a zone configuration 580
additional logical switches 501
addresses assigned 126
adjacent ISLs 43
admin 198
Admin access level 280
Admin Domain 195–196, 199, 220, 223, 225, 474, 500, 519
 AD-aware switch 221
Admin Domain creation wizard 226
Admin Domain list 624
Admin Domain task 220
admin login 140
administration 17
Administration tools window 251
Administrative Domains 474
administrative privileges 274, 517
Advanced Encryption Standard 653
Advanced Feature 36
Advanced Feature summary 37
advanced mode 176, 236
Advanced Performance Monitoring 14–15, 78, 96, 118, 173, 739–740, 753
 ALPA Error 781
 SCSI commands 778
 SCSI versus IP Traffic 780
 SID/DID Performance 755
Advanced Zoning 13–15, 118
Advanced Zoning license 514
aggregate bandwidth 590
aging scheme 374
AL_PA (PP) 283
Alarm Notifications tab, Fabric Watch View 299
alerts 349
alias 218, 517, 528, 550
 creating 523–525, 547–549, 576
Alias Selection List 518
alias server 42–43
aliCreate command 576
ALPA Error monitor 781

- analog telephone line 149
- analyzing a zone configuration 539, 557, 559
- ANSI standard-based implementation 470
- Application Specific Integrated Circuit. *See* ASIC
- aptPolicy command 604
- Arbitrated Loop 40, 283
- ASIC 6, 40–42, 44–46, 50–51, 62, 78, 83–84, 88, 97, 589, 781
- ASIC interrupts 43
- Assigning ports 504
- ATM gateways 282
- attributes 641
- audit 799
- audit trail log file 395
- AUTH policy restrictions 679
- authentication 395, 635
- authentication parameter settings 681
- Authentication policy for fabric elements 674
- authentication protocol used by the switch 681
- Authentication protocols 680
- authentication requests 395
- authentication traps 265
- Authentication, Authorization, and Accounting. *See* AAA
- automated tools 806
- Automatic Distribution 115
- Automatic Port Configuration 115
- automatic trace dump 290
- automatically enabled 483
- auto-sensing 3, 41
- auto-sensing capability 40
- auto-sensing speed negotiation 41

B

- back up
 - zone configuration 561–564
 - zone configuration, CLI 581
 - zone configuration, USB device 566–567, 569–570
 - zone configuration, USB drive 582
- backbone 23, 25, 27, 33, 112, 257, 745
- Backbone Fabric ID change 114
- backbone models 27
- backbone-class 33
- backing up a zone configuration 561
- backplane 83
- backup copy 561–562
- backup FCS switches 293

- backup files 393
- backup options 393
- balanced paths 76
- bandwidth 33, 43–44, 96, 173, 590, 695, 726, 768, 770
- bandwidth utilization 471
- bandwidth, sharing 40
- base logical switch 473
- base switch 491, 494, 507
- Basic Performance Monitoring 745
- basic processor memories 45
- basic setup functions 128
- basic support software 92
- Basic Zoning using Web Tools 543
- basics of zoning 513
- BasicSwitchAdmin 280
- basicswitchadmin 199
- BB Credit 282
- BB_Credit Recovery 589
- BE data processor 802
- BE processor 802
- Beacon button 215
- beaconing 215
- Blade Aggregate Throughput 745, 747
- Blade Aggregate Throughput graph 747
- BladeCenter 171
- block chaining 653
- Blocking Telnet 658
- Bloom ASIC 781
- Blowfish-Cipher 653
- bottleneck identification 173
- broadcast address 42
- broadcast frames 561
- broadcast packets 561
- broadcast zone 561
- Brocade 561, 566, 582, 786–787, 791
 - Native Fabric Mode 284
 - Native mode 621
- Brocade HBA 93
- browser, Web 201
- BSH File 802
- buffer allocation 597
- buffer credits 596–597
- buffer recovery credit 589
- buffers 257, 282
- buffer-to-buffer 282
- bursty 82

C

- CA 647
- canvas
 - loading 743
 - saving 742
- canvas configuration 743–744
- capacity 173
- Capturing technical support information 396
- carbon emissions 26
- CEE 7, 26–27, 37
- Certificate authorities 647
- certificate authority 647
- cfgAdd command 580
- cfgClear command 619
- cfgCreate command 579
- cfgDelete command 580
- cfgDisable command 619
- cfgRemove command 580
- cfgSave command 576–581
- cfgTransAbort command 581
- Change Discovery Switch 360
- change the seed switch 357
- Changing a logical switch to a base switch 491
- Changing account parameters 631
- Changing server port numbers 392
- Changing the fabric ID 490
- Changing the password 632
- chassis-level attributes 640
- chassis-role permission 625
- class F interswitch frames 282
- Class of Service 218
- clearing changes from a zone configuration 581
- Clearing changes to a zone configuration 581
- Cloning an IP Filter policy 687
- CNA 7
- color coded connector 32
- color coding 204
- command-line interface (CLI) 43, 45, 131, 171, 290, 577, 579, 581
 - alias, creating 576
 - zone configuration, adding members 580
 - zone configuration, clearing changes 581
 - zone configuration, creating 579
 - zone configuration, deleting 580
 - zone configuration, removing members 580
 - zone, adding members 578
 - zone, creating 577
 - zone, deleting 579
 - zone, removing members 578
- commands
 - aliCreate 576
 - aptPolicy 604
 - cfgAdd 580
 - cfgClear 619
 - cfgCreate 579
 - cfgDelete 580
 - cfgDisable 619
 - cfgRemove 580
 - cfgSave 576–581
 - cfgTransAbort 581
 - configDownload 582, 584–586
 - configShow 143
 - configUpload 142, 581
 - configure 131, 141, 144, 610
 - date “MMDDhhmmYY” 144
 - defZone 619
 - fabricShow 141
 - fanShow 212
 - fastboot 144
 - fcPing 816
 - firmwareDownload 146
 - firmwareShow 147
 - haShow 137
 - help 836
 - interopmode 622
 - ipAddrSet 137
 - ipAddrShow 138
 - killTelnet 515
 - licenseAdd 171, 275
 - licenseIdShow 269
 - licenseShow 274
 - pathInfo 816
 - portCfgIsIMode 588
 - portDisable 593
 - portEnable 171, 593
 - rcsDisabled 577
 - snmpConfig 266–267
 - switchDisable 141, 144, 593, 610, 619, 622
 - switchEnable 141, 144, 593, 619
 - switchName 131, 141
 - switchShow 142, 610
 - switchStatusPolicySet 311, 313
 - switchStatusShow 208, 311
 - tempShow 211
 - traceDump 290
 - traceTrig 290
 - version 147
 - zoneAdd 578

- zoneCreate 577
- zoneDelete 579
- zoneRemove 578
- communications privacy 635
- community string 263, 265
- Condor 41, 46, 84
- Condor ASIC 55
- Condor2 4, 87, 725
- Condor2 ASIC 11, 36, 51, 53, 55, 61, 80
- configDownload command 582, 584–586
- configShow command 143
- configUpload command 142, 581
- configuration information 806
- configuration parameters 126
- configuration procedure, SAN768B 136
- configuration upload 284, 571
- configure command 131, 141, 144, 610
- configure RADIUS 288
- Configuring a logical switch for XISL 492
- Configuring Authentication 395
- Configuring SSH authentication 655
- Configuring the port for extended distance 599
- Configuring Virtual Fabrics 475
- congested ISLs 172, 700
- congestion 43, 97, 172, 695
- connection utilization 384–385
- Connectivity Map 344
- connectors 596
- console port 136
- consolidated SAN design 25
- consumers of bandwidth 754
- context enforcement 640
- control processor 13, 52
- control processor blade 52
- Converged Enhanced Ethernet (CEE) 7, 50
- converged network adapter 7
- cooling 26
- copper pin 31
- core 22–23, 30
- core blades 13, 53
- core PID 144, 621
- core PID format 144, 609
- core switching 11
- core switching blades 76
- core-to-edge 22
- cost 605
- counters 17, 112, 836
- CP blade 148
- CP8 52

- CP8 blade 53
- CR8 317
- CR8 blade 53
- CRC 48
- CRC errors 751–752
- creating a DCC policy 670
- creating a device policy 671
- creating a logical fabric 494
- creating a logical switch 481
- creating a zone 528–529, 550, 552, 577
- creating a zone configuration 533, 553–554, 579
- creating an account 628
- creating an Admin Domain 225
- creating an alias 523–525, 547–549, 576
- creating an FCS policy 665
- creating an IP Filter policy 687
- creating an SCC policy 674
- creating logical switches 502
- creating the base switch 507
- credentials 197
- credit recovery 589
- cryptographic keys 654
- cryptography 635
- CUP 78, 296
- CUP statistics 112
- current members 16
- current product range 46
- current switches 43
- custom filters 766
- cut-through 84

D

- Data Center Fabric 21, 25, 323
- Data Center Fabric Manager 121, 323
- Data Center Fabric Manager Enterprise 121
- data collection engine 794
- data field size 282
- data packets 590, 592
- data protection 50
- data traffic 560
- date “MMDDhhmmYY” command 144
- DCC 662
- DCC policy 291, 293–294, 669
- DCC policy name 294
- DCC policy restrictions 670
- DCC policy, defining 294
- DCC violation 669
- DCF 323

- DCFM 117, 121, 182
- DCFM and QoS zones 736
- DCFM architecture 323
- DCFM Compatibility 98
- DCFM Enterprise 184, 324
- DCFM Enterprise Edition 325
- DCFM Fabric Discovery 356
- DCFM GUI Orientation 337
- DCFM Installation 327
- DCFM Operating System Support 327
- DCFM Professional 184, 324
- DCFM Reports 363
- DCFM scalability 326
- DCFM server and 336
- DCFM Server Management Console 390
- DCFM to create a zone 529
- DCFM view 346
- debug data 806
- dedicated connection 705
- dedicated ISL 707
- default accounts 626
- default cost 607
- default IP address 131
- default logical switch 473
- default Web Server port number 392
- defined configuration 515, 709
- defZone command 619
- degraded 204
- deleting
 - zone 579
 - zone configuration 580
- Deleting a device policy 672
- Deleting a fabric 361
- Deleting a rule in an IP Filter policy 693
- Deleting a zone 579
- Deleting a zone configuration 580
- Deleting an account 633
- Deleting an IP Filter policy 688
- deleting user accounts 276
- Desired Distance 239
- destination domain 603
- Device connection control 662
- Device Connection Control policy. *See* DCC policy
- Device Connectivity Troubleshooting 813
- Device Information 366
- device level zoning 42
- device ports 559
- DH-CHAP 674–675, 681
- diagnostic commands 45

- diagnostics 45, 92, 126, 326
- DID mode 95
- digital certificates 646
- director model types 40
- director type 27
- Disable Device Probing 282
- Disabled Configuration 515
- disabling a port 819
- disabling failover 708
- disabling Virtual Fabrics 479
- discovered devices 340
- Discovery 356, 358
- Discovery Status 359
- Discovery switch 357
- Discovery Verification 362
- Displaying ACL 662
- disruption 44
- distance value 596
- Distributing the local ACL policies 685
- DLS 603
- DNS maps 200
- DNS name 200
- domain 76, 78, 110, 126, 131, 141
- Domain ID 126, 131, 141, 217, 252–253, 284, 295
- domain support 110
- downloading a zone configuration 584
- Downloading a zone configuration from a USB device 572, 585
- DPS 172, 471
- dump generation 290
- duplicate domain IDs 587, 610
- Dynamic Load Sharing. *See* DLS
- Dynamic Path 5
- Dynamic Path Selection 76, 89, 471
- Dynamic Ports On Demand 171

E

- E_D_TOV 282
- E_Port 559, 603
- E_Port authentication 676
- E_Ports 31, 33
- edge 22
- EEPROM test 45
- EFCM 323
- effective configuration 515, 617, 709, 720
- EGM 182
- EGM license 182
- elements 17

- email address 271
- email alerts 120
- e-mail configuration 314
- email configuration 314
- email notification 120
- enabling a zone configuration 534–536, 555–556
- enabling the switch 506
- Enabling Virtual Fabric on the switches 499
- Enabling Virtual Fabrics 476
- encryption 93, 645–646
- encryption enhancements 96
- Ending a Web Tools session 199
- End-to-end monitoring 739
- End-to-end monitoring with DCFM 761
- End-to-end monitors 754
- end-to-end monitors 173, 759
- energy efficiency 12, 26
- energy efficient 50, 89
- Enhanced Group Management 97, 182, 326
- enterprise data centers 25
- Enterprise Edition 326
- Enterprise Fabric Connectivity Manager 323
 - supported SAN hardware 470
- enterprise-class 325
- Environmental classes 301
- Error Detect Time Out Value. *See* E_D_TOV
- Error log 257, 610, 806
- error messages 139, 257
- errors 374
- Ethernet 126, 132
- Ethernet cable 154
- Ethernet protocol 7
- Event Log 369–370
- event type 373
- events 144, 349, 369
- Excel 803
- Exchange Based Routing 591, 601–602
- exchange-based load balancing 172
- Expansion Port. *See* E_Port
- Extended Fabric Activation 14–15, 78
- Extended Fabric mode 596
- Extended Fabric tab 600
- Extended Fabric, configuring 596
- Extended Fabrics 96, 174
- eXtended ISL 471
- EZSwitchSetup 118, 152–153, 156, 160, 163, 166–167, 194
 - troubleshooting 167
 - upgrading 167

F

- F_Ports 559
- fabric
 - merging 608, 611
 - segmented 610
- fabric address notification 283
- Fabric Assist (FA) 514
- Fabric Backbone 11
- Fabric Configuration Server 662–663
- Fabric Configuration Server policy. *See* FCS policy
- fabric core 23
- Fabric Detail 365
- Fabric Events task 217
- fabric health information 124
- Fabric ID 474
- fabric infrastructure 25
- Fabric Log 370
- Fabric Login 126
- Fabric Manager 17, 140, 145, 264, 323, 521
 - alias, creating 524–525, 548–549
 - trace route feature 816
 - Zone Admin 516–517
 - zone, adding a member 552
 - zone, creating 529, 552
- fabric mode 770, 774
- Fabric Operating System 37, 91
- Fabric Operating System v6.2.0 features 92
- Fabric OS 16–18, 40, 44–45, 92, 126, 257, 514, 519, 561, 566, 577, 582, 591, 595, 621, 623, 819
- Fabric parameters 141, 144, 282
- Fabric Port Name 218
- Fabric Port WWN 218
- fabric routing 603
- Fabric Summary Report 364
- Fabric Tracking 353
- Fabric Watch 14–15, 17, 78, 97, 118–121, 173, 297–298, 301, 310
 - alarm 120
 - email notification 120
 - Port Fencing 173
 - Port Log Lock 120
 - RAPI Trap 120
 - SNMP trap 120
 - Switch Event log 120
 - threshold parameters 308
- Fabric Watch View
 - Alarm Notifications tab 299
 - Threshold Configuration tab 300–301, 307
- FabricAdmin 280

- fabricadmin 199
- fabric-connected devices 514
- fabricShow command 141
- fabric-wide configuration changes 292
- failover 147, 214
- Fan button 212
- FAN. *See* Fabric Address Notification
- fanShow command 212
- fastboot command 144
- Fastboot switch 252
- FC Ping 812
- FC Routing 110
- FC4 Type 218
- FCoE 7, 26–27, 37
- FCoE. *See* Fibre Channel over Ethernet (FCoE)
- fcPing command 816
- FCR 114
- FCR and FCIP Enhancements 93–94
- FCR scalability 111
- FCS 662–663
- FCS Automatic Distribution 115
- FCS enforcement 668
- FCS policy 291–292, 663
- FCS policy distribution 667
- FCS policy management 665
- FCS policy restrictions 664
- FCS switches 293
- FDML host name 218
- Federal Information Processing Standards (FIPS) 18
- Fibre Channel 50, 171, 218, 561
 - Arbitrated Loop (FC-AL) 40
 - over Ethernet (FCoE) 50
 - Port address 217
 - Routing 174
- Fibre Channel over Ethernet 7
- Fibre Channel Routing (FCR) 174
- FICON 78, 115, 786
- FICON CUP 13, 78, 112, 115, 175, 296
- FICON CUP Cascading 115
- FICON Enhancements 93
- FICON Log 370
- FICON Management Server 97, 175
- FICON support 112
- FID 474
- filter 42, 373
- filter monitors 763
- filter type 307
- Filter-based monitoring 739
- filter-based monitors 173, 754, 763, 766
 - adding standard filter-based monitors 763
- filter-based thresholds 307
- filtered view 223
- Filtering ports 640
- filters 762–763
- firewall 158, 295, 686
- firmware 92, 340
- Firmware Download tab 259
- firmware files 261
- firmware update 145
- firmware upgrade 260
- firmware validation 284
- firmware versions 260
- firmwareDownload command 146
- firmwareShow command 147
- FL_Ports 559
- FLOGI. *See* Fabric Login
- flow 133
- Flow-Based QoS 114
- flows 754
- FOS 91
- frame 43, 591
- Frame Based ISL Trunking 172
- frame buffers 597
- frame filtering 42, 76, 173
- frame filtering, flow 42
- Frame Redirection 95, 113
- frame routing 51
- frame routing priority 282
- frame traffic 45
- frames 84
- FSPF 605, 709, 807
- FSPF cost 95
- FSPF Route 605
- FSPF routing rules 709
- FTP 290, 806
- FTP Server 332
- FTP server 284, 286, 806–807
 - back up zone configuration 562–564
 - downloading a zone configuration 584
- full bandwidth 44
- Full Fabric license 171

G

- gateway 588
- Gateway links 588
- Generate Reports 363

- Generating a public and private key 647
- GeoTrust 647
- GigE ports tab 249
- GoldenEye 41, 46
- GoldenEye2 6, 88, 725
- grace period 193
- graph 741, 747
 - Port Error 751
 - Port Snapshot Error 752
 - Port Throughput 745
 - printing 742
 - Switch Aggregate Throughput 747
 - Switch Percent Utilization 751
 - Switch Throughput Utilization 750

H

- HA button 213
- hard zone 514
- hardware components 8
- haShow command 137
- Hayes-compatible modem 147, 149
- HBA authentication 95
- health 212, 311, 785
- help command 836
- high availability 25, 50, 148, 213
- High Availability window 213
- High Performance Extension 96, 174
- historical performance 374
- historical performance data 378
- historical performance graph 379–380
- historical performance report 380
- Home Admin Domain 195, 624
- home domain 195
- Home Virtual Fabric 624, 640
- hop count 605
- Host Bus Adapters 122
- host ports 559
- Hot Code Load 114
- HTML 187
- HTTPS 636

I

- IBM Converged Switch B32 7, 172
- IBM default settings 310
- IBM System Storage and TotalStorage 40
- IBM System Storage SAN switch 126
- IBM System Storage SAN384B Director 9
- IBM System Storage SAN768B 11

- ICA 645
- ICL 31, 471
- ICL cable connector 50
- ICL Connectivity 314
- ICL ports 31, 53
- ICLs 97
- ingress 701
- Ingress Rate Limit 239
- Ingress Rate Limiting 97, 172, 699, 701
- Ingress Rate Limiting with Web Tools 703
- ingress side 97
- ingress speed 699
- initialization 42, 127, 131, 133
- initiate failover 213–214
- in-order delivery 44, 590, 603
- Insistent Domain ID Mode 282
- Installation of DCFM Enterprise Edition 328
- Integrated Routing 5–6, 37, 97, 112, 114, 170, 174
- Integrated Routing support 111
- Inter-Chassis Link 76, 471
- Inter-Chassis Link (ICL) 12, 170, 314
 - cables 172, 317
 - cabling 317
 - connectivity 314
 - license 172
 - ports 317
- internal log 257
- Internet Certificate Authority 645
- Interoperability 196
- interoperability 13, 37
- Interoperability mode 224
- Interoperability settings 284
- InteropMode 110, 112, 622
 - InteropMode 0 621
 - InteropMode 1 621
 - InteropMode 2 519, 621
 - InteropMode 3 621
- InteropMode 2 112
- interopmode command 622
- inter-switch link (ISL) 13, 76, 173, 254, 592, 607–608
 - monitoring 769
 - monitors 173
 - ports 559
- investment protection 73
- IP addresses 340
- IP Filter policy 291, 295, 662, 688
- IP Filter policy distributions 693
- IP Filter policy enforcement 691

- IP Filter policy restrictions 694
- IP Filter policy rules 689
- IP Filtering 623
- IP management interfaces 686
- IP settings 255
- IP traffic 739, 781
- ipAddrSet command 137
- ipAddrShow command 138
- IPFilter 662
- IPSec with IPv6 95
- IPv4 295
- IPv4 filter policy 689
- IPv6 295
- IPv6 Auto-configuration 95
- IPv6 filter policy 689
- IPv6 support 200
- ISL counters 769
- ISL monitoring 739
- ISL monitors 754
- ISL performance monitoring 768
- ISL sharing 474
- ISL Trunking 13–15, 43–45, 96, 118, 131, 174, 590
 - administering 592
 - group 591
 - license 592
 - license, enabling 593
 - long distance 596
 - trunk group 591–592
 - trunk master 591–592
 - trunk ports 592
 - trunk subordinate links 592

J

- Java 118, 158, 187–188
- Java plug-in 190
- JavaScript 187
- JRE 189

K

- kernel panic 806
- killTelnet command 515

L

- L_Ports 559
- Layer 2 fabrics 110
- Layer-2 traffic isolation 471
- LDAP 17

- LDAP Enhancements 96
- LDAPS 636
- LEDs 53
- Legend button 216
- level of access 625
- license activation keys 270
- license administration 169
- License ID 271
- license key 92, 170, 179, 268, 274, 590
- License tab 268
- licenseAdd command 171, 275
- licensed features 169
 - Adaptive Networking 170
 - Integrated Routing 170
 - Inter-Chassis Link (ICL) 170
- licensed port 175, 237
- licenseIdShow command 269
- licenseShow command 274
- licensing 17, 170
 - 8 Gbps 171
 - Full Fabric 171
 - Inter-Chassis Link (ICL) 172
 - Ports on Demand (PoD) 170
- Licensing Behavior 113
- licensing issues 179
- licensing keys 179
- lighthouse icon 216
- Lightweight Directory Access Protocol. *See* LDAP
- limited switch license 193
- line speed 701
- link cost 607
- link reset 589
- link speed 41
- Linux 92, 257
- Listener applications 659
- load balancing 89, 591
- load distribution 32
- load sharing 603
- Local database user accounts 626
- Local user database 625
- logging events 144
- Logging in to a Virtual Fabric 196
- Logging in to an Admin Domain 196
- logical fabric 471, 474, 712
- logical grouping 474, 519
- logical groups 514
- logical ISL 43–44
- logical switch 37, 473, 476
- logical switch configuration 489

- logical switches 470
- logically partition 27
- login window 195
- logs 370
- long distance 596, 598
- long distance levels 596
- Long Distance mode 239
- long distance ports 597
- loop initialization 283
- loop-back function 45
- low priority traffic 97
- lower provisioning time 44
- LSAN tagging 94
- LSAN zone names 224
- LSAN zones 224
- LUN 778
- LUN level zoning 42
- LUN per port 779

M

- M14 41
- management console 390
- management functions 17
- management information base 638
- Management tools 16, 18, 140
- managing the Virtual Fabric 509
- marching ants 386
- mask 131, 256
- master port 777
- master trunk 595
- McDATA 95
- McDATA Fabric 470, 622
- McDATA Fabric Mode 284
- McDATA Fabric mode 621
- McDATA interoperability 284
- McDATA Open Fabric 470
- McDATA Open Fabric Mode 284
- McDATA Open Fabric mode 621
- Member Selection List 549
- members, adding to a zone 552, 578
- members, adding to a zone configuration 580
- members, removing from a zone 578
- members, removing members from a zone configuration 580
- memory 778, 806
- merging fabrics 587, 608, 611
- Merging fabrics example 612
- methodologies 22

- metric 605
- MIB 638
- migration 50
- MIHPTO 112
- Minimap 351
- Missed Switches 361
- Missing Interrupt Handler Primary Time-out 112
- mixed fabrics 112, 514
- modem
 - cable 147, 149
 - connecting 149
 - connection 148
 - connection, verifying 151
 - lamps 151–152
 - Off Hook (OH) indicator 151
 - port 149
 - remote 150
 - Ring indicator 151
 - serial ports 147
 - setup 147, 150
- modes 284
- Modifying the order of FCS switches 666
- modular switching platform 25
- monitor 42, 139
- monitoring 326, 374
- Monitoring Fabrics 360
- monitoring switch activity 280
- multicast group 42
- multicast routing table 42
- multiple switch environments 587

N

- N80B 88
- name server 43
- Name Server lookups 514
- Name Server queries 708
- Name Server table 217, 219
- Name Server task 217
- naming convention 670
- native connectivity 110
- native operating mode 224
- Network Config 255
- network configuration panel 257
- Network tab 257
- new messages 257
- non-dedicated paths 708
- non-disruptive failover 213–214
- Nonsecure 637

- non-volatile memory 778
- number of frames 762
- numbering scheme 84–85

O

- Object Naming 356
- Obtaining certificates 649
- one power supply 313
- open fabric management 16
- Open Fabric mode 621
- Opening Web Tools 193
- OpenSSH public key 654
- operating parameter conflicts 620
- operator 198, 280
- optimal state 204
- optimized behavior 695
- Optionally Licensed Software 96
- overlap 610
- over-subscription 83

P

- packet filtering firewall 686
- panic 290
- partition a storage area network 514
- partitioning 472
- password 134, 137, 140, 151, 156
- path selection protocol 709
- Pathinfo 95
- pathInfo command 816
- pay-as-you-grow 17, 170
- perfCfgSave command 777
- performance 2, 43, 89, 173, 603
- Performance class 306
- performance counter engine 173
- Performance Data 375
- Performance Legend 350
- performance management 373
- performance management features 373
- performance measures 375
- Performance Monitor 299, 306–307, 740, 755
- performance monitoring 42–43, 78
 - basic 744
 - Switch Utilization Throughput 741
- performance monitoring tasks 739
- performance reports 367
- Performance Thresholds 381
- Per-Frame Routing Priority 282
- permissions 625

- persistent disable 235, 241
- persistent enable 235, 241
- PID 621
- Pipelining 94
- PKI 635, 645
- PLOGI activity 528, 551
- POD license 6, 87
- pointers 806
- point-to-point E_Port connectivity 588
- Port Admin task 231
- Port Administration 248
- port area numbers 60, 62, 71–72
- Port Based Routing 601
- port blades 58, 136
- port configuration 661
- port density 51
- Port Error 745
- Port Error graph 751
- Port Fencing 116, 173, 819
- Port Fencing for E_Port class link loss 834
- Port Fencing using DCFM 820
- Port Identifier (PID) 142–143
- port information 231
- Port IP Address 218
- port level zoning 42
- Port Log Lock 120
- Port Mirroring 116
- Port mode 770
- Port Name 218
- port number 217, 236
- port numbering 88
- port position 217
- Port Report 364
- port selection 745
- Port Snapshot Error 745
- Port Snapshot Error graph 752
- port speeds 172
- port states 206
- Port Throughput 745
- port throughput capability 46
- Port Throughput graph 745
- port-based routing 602
- portCfgIsIMode command 588
- portDisable command 593
- portEnable command 171, 593
- portErrShow 836
- porterrshow counters 837
- Ports On Demand 17
- Ports on Demand 96

- Ports on Demand (PoD)
 - enabling 170
 - license 170
 - licensing 170
- Ports tab 289
- POST 45, 126, 128, 136, 151
- Power button 212
- power consumption 12, 26, 47, 75, 89
- power supply 17, 84, 212, 310–311, 313
- power-on self tests 45
- predefined accounts 626
- predefined role 625
- Preflight check 798
- primary FCS switch 293, 663
- principal 319
- principal ISL 592
- principal switch 126
- Prioritization 114
- priority 724
- priority flows 726
- priority traffic 705
- problem determination 217
- Product List 342
- Product Status Log 370
- Professional Edition 327
- protocol 121
- protocol level zoning 42
- public key 284
- public key infrastructure 635, 645
- public loop 218, 283
- PuTTY 296

Q

- QoS Zones 726
- QoS zones 730
- Quality of Service 97, 172, 770
- quick setup 152
- Quickloop 550

R

- R_A_TOV 282
- RADIUS 17, 287
- RADIUS Enhancements 95
- RAPI Trap 120
- RASLOG 806–807
- rate limit configuration 700
- RBAC 198, 625
- rcsDisabled command 577

- real life example of Virtual Fabrics 495
- real time performance 374
- Real Time Performance Data 376
- real-time monitoring 173
- Re-authenticating E_Ports 682
- reboot 252
- recipient IP address 264
- recovery logic 589
- Redbooks Web site
 - Contact us xxii
- redundant control processors 76
- Refresh Frequency 188
- Registered State Change Notification. *See* RSCN
- Reliable Commit Service (RCS) 577
- Remote LDAP server 624
- Remote RADIUS server 624
- Remote Switch 603
- Removing devices (members) from a zone 578
- Removing Thresholds for Port Fencing 826
- replication 50
- report window 252, 286
- Requirements for Admin Domains 224
- Resource Allocation Time Out Value. *See* R_A_TOV
- Restoring the database 393
- RLS probing 284
- Role 624
- Role-Based Access Control 198, 625
- Role-Based Access Control (RBAC) 18
- role-based permissions 199
- root certificates 651
- round trip time 816
- routes 44
- routing 42, 84
- Routing icons 342
- routing information 816
- routing table 126, 605
 - multicast 42
- routing tables 42, 126
 - unicast 42
- RSCN 42, 118, 283
- RX Power 303

S

- S 653
- SAN design 22
- SAN Director 14–15
- SAN Health 785–787, 791, 799, 804

- SAN Layout 803
- SAN04B-R Upgrade 98
- SAN16B 171
- SAN24B-4 3, 170, 173
- SAN256B 14–15, 45, 78, 85, 90, 133, 136, 147, 213, 235, 251
- SAN256B architecture 80
- SAN32B-2 133, 171
- SAN32B-3 152
- SAN348B 92
- SAN384B 9, 76
- SAN40B-4 4, 87, 170, 173–174
- SAN64B 171
- SAN768B 11–13, 50, 131, 140–141, 143, 172–174, 213, 314
- SAN768B architecture 64, 66
- SAN80-B4 6
- SAN80B-4 170, 173–174
- save the configuration 286
- save the configuration changes 532
- Saved Configuration 515
- Saving an IP Filter policy 688
- scalability 110
- scalability limits 111
- scalable 3
- SCC 662
- SCC policies 673
- SCC policy 291, 294
- SCP 636–637
- SCSI 739, 762
- SCSI commands monitors 778
- SCSI commands rate 739
- SCSI INQUIRY 218
- SCSI traffic 781
- SCSI versus IP Traffic monitor 780
- SDRAM 45
- secret key pair 677, 683
- Secret key pairs 683
- Secure 637
- secure access 645
- Secure Fabric OS 18, 623
- Secure file 637
- secure file copy 637
- secure login channel 654
- secure network 654
- secure protocols 635
- secure shell 653
- Secure Shell protocol 653
- Secure Sockets Layer protocol 645
- security 16–17
 - external 17
 - frame filtering 17
 - physical access 17
 - policies 280
 - software based 17
 - within SAN 17
 - zoning 17
- Security Activation 14–15
- Security Enhancements 93, 95
- security features 623
- Security level 641
- Security Log 371
- security policies 291, 662
- security problems 819
- security protocols 635
- SecurityAdmin 280
- securityadmin 198
- Seed Switch 356
- seed switch 360
- SEEPROM 48
- segmented fabric 610
- Sequence Level Switching 282
- serial cable 132–133, 135–137
- serial communication programs 133, 136, 150
- serial connection 134, 136, 139
- serial numbers 331, 340
- serial port 126, 133, 135
- SerialLink 45
- Server Application Optimization 98
- service level agreements 121, 124
- Session management 199
- Setting a secret key pair 684
- setting the domain id 506
- Setting up SCP 637
- settings 131, 144
- setup 126
- SFP 46, 142
- SFP classes 303
- sharing 44
- shipping plug 133, 135, 140–141
- shortest path 710
- SID/DID 306
- SID/DID pairs 759, 770
- SID/DID Performance monitor 755
- SID/DID prioritization 172
- Simple Network Management Protocol 121, 638
- sions 198
- SLA 121

- SMC 390
- SNMP 17, 43, 121, 265–266, 636, 638
 - trap 304
- SNMP access control list 638
- SNMP and Virtual Fabrics 640
- SNMP tab 263, 265
- SNMP trap 120
- snmpConfig 641
- snmpConfig command 266–267
- SNMPv1 636
- SNMPv1 trap 264
- SNMPv2 636
- SNMPv3 636
- SNMPv3 trap 265–266
- soft zone 514
- SoTCP Enhancements 95
- speed 40, 42, 45, 239
- splicing 596
- SSH 636, 653
- SSH client 296
- SSH public key authentication 654
- SSHv2 636
- SSL 636, 645
- SSL configuration overview 646
- stabilization period 770
- standby CP blade 136, 139–140, 148, 151–152
- static route 606
- statistics gathering 766
- status 352
- Status Bar 352
- Status button 205, 310
- Status Icons 342
- status information 806
- subordinate 320
- subordinate port 777
- summary information 252
- SupportSave 807
- supportShow 806
- suspend discovery 360
- switch
 - health 205
- Switch Admin task 250
- switch administration 280
- Switch Administration window 251
- Switch Aggregate Throughput 745
- Switch Aggregate Throughput graph 747
- switch authentication 675
- Switch Availability Monitoring Report 209
- switch configuration 259, 280

- Switch connection control 662
- Switch Details 365
- Switch Event log 120
- switch functionality 45
- switch information for support 398
- Switch Manager utility 166
- switch model types 40
- switch name 131, 140, 201, 252–253
- Switch Percent Utilization 745
- Switch Percent Utilization graph 751
- Switch tab 252
- Switch Throughput Utilization 745
- Switch Throughput Utilization graph 750
- switch user database 475
- Switch Utilization Throughput 741
- switch WWN 271
- switchadmin 198
- SwitchAdmin access level 280
- switchDisable command 141, 144, 593, 610, 619, 622
- switchEnable command 141, 144, 593, 619
- switches
 - modify settings 310
- Switch-level attributes 640
- switchName command 131, 141
- switchShow command 142, 610
- switchStatusPolicySet command 311, 313
- switchStatusShow command 208, 311
- Synchronize Services 214
- Syslog Log 371
- syslogd 257
- system memory 290

T

- Technical Support Information 396
- Telco wiring 149
- Telnet 136, 140, 142–143, 145, 208, 211–212, 266, 274, 296, 311, 313, 577, 584, 590, 592, 596, 763
 - alias, creating 576
 - zone configuration, adding members 580
 - zone configuration, clearing changes 581
 - zone configuration, creating 579
 - zone configuration, deleting 580
 - zone configuration, removing members 580
 - zone, adding members 578
 - zone, creating 577
 - zone, deleting 579
 - zone, removing members 578

- Telnet protocol 658
- Temp button 210
- temperature 17
- temperature information 210
- temporary internet files 188
- Temporary License Support 98
- temporary licenses 93
- temporary use 98
- tempShow command 211
- terminal emulator application 132–134, 136, 150
- threshold 314
- threshold configuration 299
- Threshold Configuration tab, Fabric Watch View 300–301, 307
- threshold parameters 308
- Thresholds for the Environmental classes 301
- throttle 172
- throughput 43, 46
- TI zone failover 707
- TI zones 706
- TI zones with DCFM 721
- tight bends 596
- timeout value 145, 289, 603
- time-saving tools 123
- toolbar 339
- Toolbox 345
- Top Talker monitors 769
- Top Talkers 695
- Top Talkers monitoring 173, 739
- Top Talkers monitors 754, 770, 776
 - adding 771, 774
 - deleting 771, 775
 - displaying 775
 - displaying information 771
 - fabric mode 774
- top-of-rack 7, 27
- topology 42, 142
- topology changes 603, 606
- TopTalkers feature 96
- trace buffer 290
- trace dump 290–291, 806–807
- trace route 812, 816
- Trace Route Summary 818
- Trace tab 290
- traceDump command 290
- traceTrig command 290
- tracing 290
- Tracking Icons 342
- traffic 172, 374, 728

- traffic control 172
- traffic flow 726
- Traffic Isolation 114–115, 705
- Traffic Isolation zones 706
- traffic load 42
- Traffic Management 696, 705
- traffic patterns 81–82
- Traffic Prioritization 725
- Traffic prioritization 729
- traffic prioritization 724
- transaction 693
- transaction key 179, 271
- transmission 592
- transmitter negotiation 126
- trap level 263, 265–266
- trap recipients 263
- troubleshooting 326, 785, 835
- Troubleshooting device connectivity 813
- troubleshooting ports 751
- trunk group 591–592
- trunk master 591–592
- trunk master link 44
- trunk monitoring 777
- trunk ports 592, 595
- trunk speeds 89
- trunk subordinate links 592
- trunking
 - group 141
- Trunking tab 296
- trunks
 - monitoring 777
- trusted key agent 645
- TX Power 303

U

- Unblocking a Fenced port 827
- Unblocking Telnet 659
- unicast 126
- unicast routing table 42
- unidirectional 172
- unified management 122
- unmonitor fabrics 363
- updating a zone configuration 553–554
- upgrades 109
- upload 259
- USB 4, 11, 52, 259, 284
- USB drive 561, 566, 582
- USB memory key 259

- USB port 7, 259
- user 199
- User access level 280
- user accounts 475, 625
- User accounts overview 624
- user administration tasks 278
- User authentication 624
- User tab 276

V

- Value line licenses 193
- vendor company 218
- Verisign 647
- version command 147
- VF 94
- VF-capable switches 499
- View Report button 254
- Viewing technical support information 400
- Viewing the list of secret key pairs 683
- virtual channel ID 282
- Virtual Channels 725
- virtual channels (VC) 172, 282–283
- virtual channels parameters 283
- Virtual Fabric 5–6, 10, 37, 640
- Virtual Fabric configuration 712
- Virtual Fabric is disabled 499
- Virtual Fabric list 624
- Virtual Fabrics 92, 94, 110, 113, 196, 700, 753
- Virtual Fabrics introduction 470
- virtualization 11, 50
- Visio 786, 803–804
- Visio Viewer 805
- VxWorks 92

W

- watchdog 806
- watchdog timer 290
- Web browser 201
- Web Tools 14–15, 17, 78, 118, 140, 145, 170–171, 175, 260, 262, 284, 310, 314–315, 555, 566, 590, 594, 596, 598, 610
 - Beacon button 215
 - buttons 204
 - EZ error message 167
 - Fabric Tree panel 202
 - Fan button 212
 - features 118
 - HA button 213

- High Availability window 213
- Legend button 216
- Performance Monitor 740
- Power button 212
- Status button 205
- Switch Events, Information panel 202
- Switch View panel 202
- Tasks panel 202
- Temp button 210
- Zone Admin 516–517, 521
- Web Tools license 182
- Web Tools to create a zone 551
- workload peaks 43
- world wide name 48
- worldwide name. *See* WWN
- WWN 48, 76, 179, 340
 - zoning
- WWN cards 74
- WWN Display 355
- WWNN 218
- WWPN 217–218

X

- XISL 471, 474
- XISLs 492

Z

- Zone Admin 516–517
- Zone Admin task 220
- zone configurations 515, 534
- zone database 519
- zoneAdd command 578
- ZoneAdmin 280
- zoneadmin 199
- zoneCreate command 577
- zoneDelete command 579
- zoneRemove command 578
- zones 217–218
- zoning 41–42, 172, 514, 516, 609
 - adding a member 552
 - adding members 578
 - adding members to a configuration 580
 - administrative privileges 517
 - Advanced 161
 - analyzing a configuration 557
 - analyzing a zone configuration 559
 - back up a configuration 561–564, 566–567, 569–570, 581–582

- check 816
- clearing changes from a configuration 581
- configuration 161, 517, 611
- conflicts 559
- creating 577
- creating a configuration 553–554, 579
- creating a zone 528–529, 550, 552
- Custom 161
- deleting 579
- deleting a configuration 580
- downloading a configuration 584
- enabling a configuration 535–536, 555–556
- hard zone 514
- icon 220
- implementing 521
- information 611
- license 170, 517
- managing 521
- removing members 578
- removing members from a configuration 580
- segmentation 612
- soft zone 514
- Typical 161
- updating a zone configuration 553–554
- Zoning Activation 521
- zoning configuration 706
- zoning configuration conflicts 587
- Zoning Consideration 709
- zoning elements 518
- Zoning Offline 521
- Zoning Online 521
- zoning reports 367



Redbooks

Implementing an IBM b-type SAN with 8 Gbps Directors and Switches

(1.5" spine)

1.5" <-> 1.998"

789 <-> 1051 pages



Implementing an IBM b-type SAN with 8 Gbps Directors and Switches

Learn about the latest additions to the IBM b-type portfolio

Refresh and enhance your skills and awareness

Increase your SAN knowledge

“Do everything that is necessary and absolutely nothing that is not.”

This IBM Redbooks publication, written at a Data Center Fabric Manager v10.1.4 and Fabric Operating System v6.4 level, consolidates critical information while also covering procedures and tasks that you are likely to encounter on a daily basis when implementing an IBM b-type SAN.

The products that we describe in this book have more functionality than we can possibly cover in a single book. A storage area network (SAN) is a powerful infrastructure for consolidation, distance solutions, and data sharing. The quality applications that the IBM SAN portfolio provides can help you take full advantage of the benefits of the SAN.

In this book, we cover the latest additions to the IBM b-type SAN family and show how you can implement them in an open systems environment. In particular, we focus on the Fibre Channel Protocol (FCP) environment. We address the key concepts that these products bring to the market and, in each case, we provide an overview of the functions that are essential to building a robust SAN environment.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks