

Implementing or Migrating to an IBM Gen 5 b-type SAN

Mirza Baig

Liam Dowds

Silviano Gaona

Paulo Neto

Gaston Rius

Megan Gilge



 **Cloud**

Storage



International Technical Support Organization

Implementing or Migrating to an IBM Gen 5 b-type SAN

August 2016

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

First Edition (August 2016)

This edition applies to the IBM Gen 5 b-type SAN products.

© Copyright International Business Machines Corporation 2016. 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	ix
Trademarks	x
IBM Redbooks promotions	xi
Preface	xiii
Authors	xiii
Now you can become a published author, too!	xv
Comments welcome	xv
Stay connected to IBM Redbooks	xv
Chapter 1. Introduction	1
1.1 IBM SAN b-type Gen 5 technology overview	2
1.1.1 Hardware features	2
1.1.2 Hardware naming convention: IBM and Brocade	2
1.1.3 Fabric OS features	3
1.1.4 Fabric Operating System hardware support	5
1.1.5 Management	5
1.1.6 IBM Network Advisor	5
1.1.7 Monitoring	6
1.1.8 Brocade Fabric Vision technology	6
1.1.9 IBM Spectrum Control	7
1.2 Product descriptions	7
1.2.1 Enterprise SAN directors	7
1.2.2 Mid-range SAN switches	7
1.2.3 Entry SAN switches	7
1.2.4 Specialty switches	8
Chapter 2. Fabric design overview	9
2.1 Topologies	10
2.1.1 Edge-core topology	10
2.1.2 Edge-core-edge topology	11
2.1.3 Full-mesh topology	11
2.2 Gen 5 Fibre Channel technology	11
2.2.1 Condor3 ASIC	11
2.2.2 Fabric Vision	12
2.3 Standard features	20
2.3.1 Zoning	20
2.3.2 ISL Trunking	20
2.3.3 Dynamic Path Selection	22
2.3.4 In-flight encryption and compression	24
2.3.5 NPIV	24
2.3.6 Dynamic Fabric Provisioning	25
Chapter 3. Management and monitoring tools	27
3.1 Web Tools	28
3.1.1 Web Tools introduction and features	28
3.1.2 System requirements	32
3.1.3 Java plug-in configuration	34

3.1.4 Value line licenses	34
3.1.5 Opening Web Tools	35
3.2 Command-line interface	40
3.3 Storage Management Initiative Agent	40
3.4 Fabric Vision	41
3.4.1 Base Fabric OS and Fabric Vision enabled features	41
3.4.2 Fabric OS and Fabric Vision licensing considerations	44
3.5 IBM Network Advisor	45
3.6 IBM Spectrum Control	47
Chapter 4. IBM Network Advisor	49
4.1 Planning for server and client system requirements	50
4.1.1 Operating system and hardware requirements for Server and Client	50
4.1.2 Browser requirements for IBM Network Advisor 12.4.2	53
4.1.3 IBM Network Advisor Server and concurrent client connections	53
4.2 IBM Network Advisor 12.4.2 upgrade path	53
4.3 Downloading the software	54
4.4 Pre-installation requirements	56
4.4.1 Additional pre-installation requirements for UNIX systems	56
4.4.2 Mapping a loopback address to the local host	57
4.5 Syslog troubleshooting	57
4.5.1 Finding the process	57
4.5.2 Stopping the process	57
4.6 IBM Network Advisor Version 12.4.2 installation	58
4.7 Upgrading to IBM Network Advisor V12.4.2 from an existing IBM Network Advisor installation	75
4.8 IBM Network Advisor web client	87
4.9 User, device discovery, and dashboard management	87
4.9.1 User management	88
4.9.2 Discovering and adding SAN fabrics	93
4.10 New features of IBM Network Advisor V12.4.2	98
4.11 IBM Network Advisor Dashboard overview	99
4.12 Scheduling daily or weekly backups for the fabric configuration	101
4.13 Call Home	102
4.13.1 System requirements	104
4.13.2 Editing an email Call Home center	105
Chapter 5. Product hardware	107
5.1 Hardware overview	108
5.1.1 Entry level, midrange, and director models	108
5.1.2 IBM Gen 5 SAN b-type 16 Gbps family	108
5.2 IBM Gen 5 SAN b-type family	108
5.2.1 IBM SAN24B-5 (2498-F24, 2498-X24, and 2498-24G)	109
5.2.2 IBM System Networking SAN48B-5 (2498-F48)	111
5.2.3 IBM System Networking SAN96B-5 (2498-F96 / 2498-N96)	114
5.2.4 IBM System Networking SAN384B-2 (2499-416) and IBM System Networking SAN768B-2 (2499-816)	117
5.2.5 IBM Fabric backbone blades	121
5.2.6 Optical UltraScale Inter-Chassis Links	129
5.2.7 IBM System Storage SAN42B-R Extension Switch	137
Chapter 6. B-type SAN monitoring and management with IBM Spectrum Control	141
6.1 Software prerequisites	142
6.2 Interoperability matrixes for supported switches and directors	142

6.2.1 Virtual Fabrics support	143
6.3 Monitoring agents for switches and fabrics	143
6.3.1 Monitoring and managing fabrics with the Storage Resource Agent (IBM Tivoli Storage Productivity Center 5.2.7 and earlier only)	145
6.3.2 Monitoring and managing fabrics with the SNMP Agent (out-of-band)	146
6.3.3 Monitoring and managing fabrics with the Brocade SMI Agent	146
6.3.4 Monitoring and managing fabrics with multiple agents	147
6.4 Adding switches and fabrics	148
6.4.1 Device discovery	148
6.4.2 Device probes	148
6.4.3 Adding switches by using SNMP Agent (out-of-band)	149
6.4.4 Adding switches with the SMI Agent	151
6.4.5 Adding switches by using Storage Resource Agent (IBM Tivoli Storage Productivity Center 5.2.7 and earlier)	153
6.5 Testing connectivity for a switch	159
6.6 Enabling the switch performance monitoring	160
6.6.1 Troubleshooting Performance Monitor data collection problems	161
6.7 Viewing Switches and Fabrics details	163
6.7.1 Viewing fabric details	163
6.8 Zoning	172
6.8.1 Non-standard zones	172
6.8.2 Zone control capabilities of IBM Spectrum Control	172
6.8.3 Setting the zoning policy (automatic zoning feature)	173
6.9 Alerting	175
6.9.1 Prerequisites for using alerts	175
6.9.2 Setting up the alert notification settings	176
6.9.3 Enabling the default alert definitions	177
6.9.4 Custom alert definitions	179
6.9.5 Managing and acknowledging alerts	183
6.10 Performance monitoring	184
6.10.1 Performance thresholds violations review	185
6.10.2 Performance review	186
Chapter 7. Fabric Vision	191
7.1 Fabric Vision	192
7.1.1 Flow Vision	192
7.2 ClearLink Diagnostics Port	193
7.2.1 Enabling D_port Diagnostics by using IBM Network Advisor	194
7.2.2 Selecting and running D_port diagnostics	195
7.3 Bottleneck detection	195
7.3.1 Enabling, displaying, and disabling bottleneck detection	196
7.3.2 Enabling bottleneck monitor in IBM Network Advisor	197
7.3.3 Configuring Bottleneck monitors	197
7.3.4 Bottleneck Monitor Suggested initial settings	198
7.3.5 Displaying bottleneck statistics	199
7.4 Buffer credit depletion and recovery	199
7.5 Fabric Performance Impact monitoring	200
7.6 Managing Forward Error Correction	202
7.6.1 Enabling, disabling, and viewing FEC status	202
7.7 Monitoring, Alerting, and Performance Suite	202
7.7.1 Enabling MAPS	203
7.7.2 Configuring MAPS with Fabric Watch Rules	204
7.7.3 MAPS Slow Drain Device Quarantine	206

7.7.4	MAPS port fencing	208
7.7.5	MAPS Toggling and decommissioning	209
7.7.6	MAPS Dashboard	210
Chapter 8.	Virtual Fabrics	213
8.1	IBM b-type Virtual Fabric	214
8.1.1	Virtual Fabrics introduction	214
8.1.2	Logical switches and logical fabrics	214
8.2	Virtual Fabric features	215
8.2.1	Logical switch	216
8.2.2	Logical fabric	216
8.2.3	ISL sharing	217
8.2.4	User accounts	217
8.3	Configuring Virtual Fabrics	217
8.3.1	Changing the context to a different logical switch	217
8.3.2	Enabling Virtual Fabrics	218
8.3.3	Disabling Virtual Fabrics	220
8.3.4	Logical switch management	221
8.3.5	Modifying the base switch	222
8.3.6	Creating a logical switch	223
8.3.7	Deleting a logical switch	228
8.3.8	Displaying the logical switch configuration	228
8.3.9	Changing the fabric ID of a logical switch	230
8.3.10	Changing a logical switch to a base switch	231
8.3.11	Configuring a logical switch for XISL use	232
8.3.12	Creating a logical fabric that uses XISLs	234
Chapter 9.	Implementation and migration strategies	235
9.1	Designing a storage area network	236
9.1.1	Redundancy and resiliency	237
9.1.2	Switch interconnections	237
9.1.3	UltraScale ICL connectivity for Gen 5 directors	238
9.1.4	SAN768B-2 and SAN384B-2 UltraScale ICL connection preferred practices	238
9.1.5	Device placement	238
9.2	Migration assessment	240
9.2.1	Assessing the existing fabric topology	240
9.2.2	Assessing the new fabric	241
9.2.3	Logistic planning of hardware installation	241
9.2.4	Preliminary migration planning	241
9.2.5	Gather infrastructure information	242
9.3	Migration strategy	242
9.3.1	Migration methods	243
9.4	Developing a migration plan	243
9.5	Preparing to migrate	243
9.6	Performing the migration and validation	244
9.6.1	Offline migration	244
9.6.2	Redundant and single fabric online migration	244
9.7	Completing the migration	244
9.8	Licensing	244
9.8.1	Available Fabric OS licenses	245
9.8.2	License administration	251
Chapter 10.	Fabric administration	253
10.1	Administration practices	254

10.2 Initial setup	254
10.3 Audit and syslog configuration	255
10.3.1 Audit log	255
10.3.2 Syslog	256
10.4 Network Time Protocol	259
10.5 Zoning	260
10.5.1 Zoning preferred practices	261
10.5.2 Peer Zoning	262
10.6 Trunking	264
10.6.1 Configuring trunk groups	264
10.6.2 Enabling Trunking	266
10.6.3 Types of trunking	267
10.7 Fibre Channel over distance	268
10.7.1 Buffer credits	268
10.7.2 Fabric interconnectivity over Fibre Channel at longer distances	269
10.8 Fibre Channel over IP	269
10.8.1 Configuring an FCIP tunnel	270
10.9 FC-FC routing overview	271
10.9.1 Setting up FC to FC routing	272
10.9.2 Logical SAN Zones	272
10.9.3 Fibre Channel routing and virtual fabrics	273
10.10 FCIP and FCR	274
10.10.1 Using EX_Ports and VEX_Ports	274
10.11 Access Gateway and N_Port ID Virtualization	277
10.12 Inter-chassis links	278
10.12.1 Supported topologies	278
10.12.2 QSFP-based ICL connection requirements	279
10.12.3 ICL trunking and trunk groups	280
10.12.4 ICL diagnostic tests	283
10.12.5 Summary	283
10.13 Fabric OS management	284
10.13.1 Firmware download enhancements	285
10.14 Upgrading firmware or rolling back to an earlier version	285
10.14.1 Preparing for upgrades	286
10.14.2 Staging the Fabric OS package for download to the switch	287
10.14.3 Upgrading firmware	289
Chapter 11. Security	291
11.1 Role-Based access controls	292
11.2 Default accounts	292
11.3 User accounts	292
11.4 Security protocols	293
11.5 Access control lists	294
11.5.1 SCC policy	294
11.5.2 DCC policy	294
11.5.3 FCS policy	294
11.5.4 IP Filter	295
11.5.5 Authentication protocols	295
11.6 Policy Database Distribution	295
11.7 In-flight encryption and compression: b-type (16 Gbps) platforms only	296
11.8 In-flight encryption and compression guidelines	297
Chapter 12. Troubleshooting	299

12.1	General problem determination	300
12.2	Errors and symptoms	300
12.3	Switch and port status	300
12.3.1	Displaying the switch status	300
12.3.2	Port status	303
12.4	Port errors	306
12.4.1	Viewing port statistics with IBM Network Advisor	306
12.4.2	Viewing the port statistics in the CLI	307
12.4.3	Resetting the port error statistic counters	308
12.4.4	Understanding error counters	308
12.4.5	SFP and optic levels	309
12.5	System messages and RAS logs	310
12.5.1	System message types	310
12.5.2	RASLog messages	310
12.5.3	Audit log messages	311
12.5.4	First-Failure data capture messages	311
12.6	SAN health	311
12.6.1	Installing Brocade SAN Health	312
12.6.2	Using SAN Health Diagnostics Capture	312
12.6.3	SAN Health Professional	313
12.7	Collecting support data	314
12.7.1	Saving comprehensive diagnostic files to the server	314
12.7.2	Scheduling technical support and event information collection by using IBM Network Advisor	314
12.7.3	Starting immediate technical support information collection	316
12.8	Using MAPS for problem determination	317
12.8.1	Port health and cyclic redundancy checks monitoring	317
12.8.2	Back-end port monitoring	319
12.8.3	FRU Health	320
12.8.4	Fibre Channel over IP (FCIP) Health	321
12.9	Flow Vision	322
12.9.1	Flow Monitor	323
12.9.2	Flow Generator	324
12.9.3	Flow Mirroring	324
	Related publications	325
	IBM Redbooks	325
	Online resources	325
	Help from IBM	325

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM Spectrum Control™	System i®
FICON®	IBM Spectrum Storage™	System Storage®
Global Technology Services®	IBM z Systems™	Tivoli®
IBM®	Netcool®	Tivoli Enterprise Console®
IBM Cloud Managed Services®	OS/390®	z Systems™
IBM SmartCloud®	Redbooks®	z/OS®
IBM Spectrum™	Redbooks (logo)  ®	

The following terms are trademarks of other companies:

Performance View, and Inc. device are trademarks or registered trademarks of Kenexa, an IBM Company.

Linux is a trademark of Linus Torvalds 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.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get personalized notifications of new content
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



Download
Now

Android



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks

About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

The IBM® b-type Gen 5 Fibre Channel directors and switches provide reliable, scalable, and secure high-performance foundations for high-density server virtualization, cloud architectures, and next generation flash and SSD storage. They are designed to meet the demands of highly virtualized private cloud storage and data center environments.

This IBM Redbooks® publication helps administrators learn how to implement or migrate to an IBM Gen 5 b-type SAN. It provides an overview of the key hardware and software products and explains how to install, monitor, tune, and troubleshoot your storage area network (SAN).


Read this publication to learn about fabric design, managing and monitoring your network, key tools such as IBM Network Advisor and Fabric Vision, and troubleshooting.

Authors

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



Mirza Baig is an Advisory Cloud Architect at IBM Cloud Managed Services® (CMS) in the US. He has over 15 years of experience in data center technologies, cloud, and open systems. Mirza is an IBM certified IT Architect and his expertise includes designing solutions around public, private, and hybrid clouds. Since 2008, he has held positions at IBM Technical Support Services and IBM Systems Lab Services. Before joining IBM, he worked for telecom providers and IT managed services firms in Europe and the Middle East.



Liam Dowds is a Remote Storage Support Representative for IBM Canada. He has 17 years of experience with IBM in various roles including Customer Support for the complete portfolio of IBM products, Open Systems customization, and storage support for the full line of IBM disk storage and SAN products. He holds numerous SAN certifications and currently focuses on providing remote support to IBM SAN clients. His areas of expertise include b-type switch problem determination and configuration, and total solution end-to-end problem determination.



Silviano Gaona is a Senior Systems Engineer at Brocade Communications supporting the IBM account. He has over 23 years of IT experience and 15 years at Brocade. He has worked in Customer Support, as a Technical Developer, as a Technical Trainer, and in System Engineering. He has extensive experience in SAN, including Fibre Channel fabric, Ethernet Fabrics, and IP storage.



Paulo Neto is a Senior Project Engineer at Johnson Controls in Germany. He has 27 years of IT experience. He worked at IBM Portugal for 25 years. He was an IBM Certified IT Specialist (Level 2). His last position at IBM was with the SSO PanIoT Storage Service Line as a Storage Technical Lead. His areas of expertise include SAN and Storage Virtualization, Design, Implementation, Disaster Recovery, and Customer Support. He has written extensively on IBM Tivoli® Storage Manager, IBM SAN Volume Controller, and SAN. He holds a Master of Science degree in Informatics from the University of Porto.



Gaston Rius is a Senior IT Specialist from Buenos Aires, Argentina with 16 years of IT experience. Before joining IBM Global Technology Services® (GTS) in 2004, he worked as a Data Center and UNIX Support IT Specialist for three years at a major telecommunication company and then two years as Pre-Sales Support and Implementor for an IBM Business Partner. His current role in IBM is as a Storage Subject Matter Expert (SME) and his focus is on field storage solutions design, implementation, and client support to multiple outsourced clients worldwide. His areas of expertise include a wide array of storage hardware and software products from multiple vendors, storage virtualization, storage area networks, storage tiering, remote copy services, IBM Tivoli Storage Productivity Center, IBM SmartCloud® Virtual Storage Center, and Brocade Network Advisor. He is also interested in patents and cloud.



Megan Gilge is a Project Leader in the IBM International Technical Support Organization. Before joining the ITSO four years ago, she was an Information Developer in the IBM Semiconductor Solutions and IBM System i® areas.

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

Jon Tate
International Technical Support Organization

Silviano Gaona
Brian Steffler
Brocade Communications

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>



Introduction

This chapter describes the IBM SAN b-type (Brocade) Gen 5 family of products. It includes the hardware naming conventions (IBM versus Brocade) and the associated software components.

This chapter includes the following sections:

- ▶ IBM SAN b-type Gen 5 technology overview
- ▶ Product descriptions

1.1 IBM SAN b-type Gen 5 technology overview

This section introduces the IBM SAN b-type Gen 5 technology and the features that are provided by Brocade Fabric Operating System (FOS). See the following website for the most up-to-date information:

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

1.1.1 Hardware features

The IBM b-type Gen 5 Fibre Channel directors and switches provide reliable, scalable, and secure high-performance foundations for mission-critical storage environments based on the new 16 Gbps Fibre Channel technology. They are designed to meet the demands of highly virtualized private cloud storage and data center environments. The portfolio consists of entry level 12-port fabric switches. It goes up to 3456 16 Gbps ports (or 4608 8 Gbps ports) when connecting nine backbone chassis in a full mesh topology through UltraScale Inter-Chassis Links (ICLs). These storage area network (SAN) platforms support 2, 4, 8, and 16 Gbps auto-sensing ports and deliver enhanced fabric resiliency and application uptime through advanced features.

The Condor3 application-specific integrated circuit (ASIC) enables support for Fabric Vision features (Forward Error Correction (FEC), Flow Monitoring, Flow Generator, and Flow Mirroring), native 10 Gbps Fibre Channel, in-flight encryption and compression, ClearLink diagnostic technology (supported only on the 16 Gbps ports), and increased buffers.

This new Gen 5 family allows a simple server deployment with dynamic fabric provisioning, which enables organizations to eliminate fabric reconfiguration when adding or replacing servers through the virtualization of the host worldwide names (WWNs).

1.1.2 Hardware naming convention: IBM and Brocade

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

Table 1-1 IBM b-type family product and Brocade equivalent names

IBM name	IBM machine type and model	Brocade name
IBM SAN24B-5	2498-24G, 2498-X24 2498-F24 (2 power supplies)	Brocade 6505
IBM SAN48B-5	2498-F48	Brocade 6510
IBM SAN96B-5	2498-F96 and 2498-N96	Brocade 6520
IBM SAN384B-2	2499-416	Brocade DCX 8510-4
IBM SAN768B-2	2499-816	Brocade DCX 8510-8
IBM SAN42B-R	2498-R42	Brocade 7840

1.1.3 Fabric OS features

FOS V7 with the b-type Gen 5 products offers a set of advanced features. Not all of these features are available for all switch models and some of them are offered as optional licenses. For a detailed description of available licenses, see 9.8, “Licensing” on page 244.

The following list introduces the features with a brief explanation:

- ▶ *Advanced Web Tools* enable graphical user interface (GUI) based administration, configuration, and maintenance of fabric switches and SANs.
- ▶ *Advanced Zoning* segments a fabric into virtual private SANs to restrict device communication and apply certain policies only to members within the same zone:
 - *Peer Zoning* is a new feature that was introduced in FOS 7.4 where the zone membership is differentiated into principal members and non-principal or peer members. For more information, see Chapter 10, “Fabric administration” on page 253.
 - *Target Driven Zoning* allows end devices to create Peer Zone configurations through inband commands. This feature enables zoning to be configured by management software on storage devices and reduces the manual configuration that is needed on switches.
- ▶ *Virtual Fabrics* allow a physical switch to be partitioned into independently managed Logical Switches, each with its own data, control, and management paths.
- ▶ *Full Fabric* allows a switch to be connected to another switch. It is required to enable expansion ports (E_Ports).
- ▶ The *Adaptive Networking* service is a set of features that provides users with tools and capabilities for incorporating network policies to ensure optimal behavior in a large SAN. FOS V7.0 supports two types of quality of service (QoS) features with the 16 Gbps fabric backbones: Ingress rate limiting and session ID (SID)/destination ID (DID)-based prioritization.
- ▶ *Server Application Optimization (SAO)* enhances overall performance and virtual machine scalability by extending b-type data center fabric technologies to the server infrastructure. SAO enables individual traffic flows to be configured, prioritized, and optimized, from end to end, throughout the data center.
- ▶ *Enhanced Group Management (EGM)* enables additional device-level management functions for IBM b-type SAN products when it is added to the element management. It also allows large consolidated operations, such as firmware downloads and configuration uploads and downloads for groups of devices.
- ▶ *Extended Fabrics* extend SAN fabrics beyond the Fibre Channel standard of 10 km by optimizing internal switch buffers to maintain performance on ISLs that are connected at extended distances.
- ▶ *Integrated Routing* allows any 16 Gbps Fibre Channel port to be configured as an EX_Port supporting Fibre Channel Routing.
- ▶ *Integrated 10 Gbps Fibre Channel Activation* enables Fibre Channel ports to operate at 10 Gbps.
- ▶ *IBM Fibre Channel connection (FICON®) with Control Unit Port (CUP) Activation* is designed to provide in-band management of the supported SAN b-type switch and director products by system automation for IBM z Systems servers. This support provides a single point of control for managing connectivity in active FICON I/O configurations. To enable in-band management on multiple switches and directors, each chassis must be configured with the appropriate FICON CUP feature. System automation for IBM OS/390® or z/OS® can now use FICON to concurrently manage IBM Enterprise Systems

Connection (ESCON) Director 3092, in addition to supported SAN b-type switch and director products.

- ▶ An *Inter-chassis license* with 16× (4×16 Gbps) QSFP provides connectivity up to 2 km from the switching backplane of one half of an eight-slot chassis to the other half, or to a 4-slot chassis.
- ▶ An *Enterprise ICL license* supports up to 3,840 16 Gbps universal Fibre Channel ports (using 16 Gbps 48-port blades), up to 5,120 8 Gbps universal Fibre Channel ports (using 8 Gbps 64-port blades), and ICL ports (32 or 16 per chassis, with optical QSFP). These ports can be connected up to nine chassis in a full-mesh topology or up to 10 chassis in a core-edge topology. Connecting five or more chassis through ICLs requires an Enterprise ICL license.
- ▶ An *Advanced Extension activation license* enables two advanced extension features, Fibre Channel over IP (FCIP) trunking and adaptive rate limiting (ARL), on the IBM System Networking SAN768B-2 or IBM System Networking SAN384B-2 systems. The FCIP trunking feature allows multiple IP source and destination address pairs (defined as FCIP circuits) through multiple 1 GbE interfaces to provide a high-bandwidth FCIP tunnel and failover resiliency. The ARL feature is designed to provide a minimum bandwidth guarantee for each tunnel with full usage of the available network bandwidth without affecting throughput performance under a high traffic load.
- ▶ An *Extension blade 10 GbE activation license* enables up to two 10 GbE ports on the 8 Gbps extension blades or eight 10 Gbps Fibre Channel ports on the first eight ports of a 16 Gbps port blade. With this license, two additional operating modes, in addition to a 1 GbE port mode, can be selected. Either two 10 GbE ports, or ten 1 GbE and one 10 GbE ports, can be configured on an 8 Gbps extension blade when this license is activated.
- ▶ An *FICON Accelerator activation license* uses advanced networking technologies, data management techniques, and protocol intelligence to accelerate FICON disk and tape read-and-write operations over geographically extended distances. It can do so while also maintaining the integrity of command and acknowledgment sequences. Ideal for data migration, disaster recovery, and business continuity solutions beyond 300 km, it supports emulation for IBM z/OS Global Mirror (formerly Extended Remote Copy (XRC)) and tape pipelining for FICON tape and virtual tape.
- ▶ *ISL Trunking* enables Fibre Channel packets to be distributed efficiently across multiple ISLs between two IBM b-type SAN fabric switches and directors while preserving in-order delivery. Both b-type SAN devices must have trunking activated.
- ▶ *Fabric Vision* provides end-to-end visibility and insight across the storage network through advanced diagnostic, monitoring, and management:
 - *Monitoring and Alerting Policy Suite (MAPS)* is an optional SAN health monitor that allows you to enable each switch to constantly monitor its SAN fabric for potential faults and automatically alerts you to problems long before they become costly failures. MAPS replaces Fabric Watch in FOS 7.4.
 - *Flow Vision* is a comprehensive tool that enables administrators to identify, monitor, and analyze specific application data flows to maximize performance, avoid congestion, and optimize resources. Flow Vision replaces Advanced Performance Monitoring (APM) in FOS 7.4
 - *Fabric Performance Impact (FPI) Monitoring* provides separate, easy-to-understand notifications that are based on severity and impact of device latency. It has simple configuration with zero user-configurable thresholds. It provides three simultaneous monitoring windows for concurrent monitoring of latency from brief, severe conditions to sustained, low-level conditions. FPI monitoring replaces the previous FOS Bottleneck Detection feature.

Note: MAPS and Flow Vision features are supported only on Gen 4 (8 Gbps) and Gen 5 (16 Gbps) and later running FOS v7.2.0 or later. Flow Vision includes Flow Monitoring, Flow Generator, and Flow Mirror.

1.1.4 Fabric Operating System hardware support

FOS V7.x supports only 8 Gbps and 16 Gbps hardware platforms. For the latest list of supported devices, see the IBM SAN b-type Firmware Version 7.x Release Notes at the following website:

<https://www.ibm.com/support/entry/portal/support>

1.1.5 Management

The IBM b-type Gen 5 Fibre Channel directors and switches can be managed in several ways:

- ▶ *IBM Network Advisor* is a software management platform that unifies network management for SAN and converged networks. It integrates the SAN management capabilities of IBM Data Center Fabric Manager (DCFM) and the IP network management capabilities of Brocade IronView Network Manager (INM). It provides users with a consistent user interface, proactive performance analysis, and troubleshooting capabilities across Fibre Channel (FC) and b-type Fibre Channel over Ethernet (FCoE) installations.
- ▶ *Web Tools* is a built-in web-based application that provides administration and management functions on a per-switch basis.
- ▶ *A command-line interface (CLI)* enables an administrator to monitor and manage individual switches, ports, and entire fabrics from a standard workstation. It is accessed through Telnet, SSH, or serial console.
- ▶ *SMI Agent* enables integration with SMI-compliant Storage Resource Management (SRM) solutions, such as IBM Spectrum™ Control. The SMI Agent is embedded in the IBM Network Advisor or it can be deployed separately.

Note: IBM System Storage® DCFM is not qualified with and does not support the management of switches operating with FOS V7.0 and later firmware versions. You must first upgrade DCFM to IBM Network Advisor V12.0 or later if you are planning to upgrade devices to FOS V7.1.0 or later.

1.1.6 IBM Network Advisor

IBM Network Advisor is the preferred tool for managing and monitoring the IBM b-type Gen 5 SANs. It is a software management tool that provides comprehensive management for data, storage, and converged networks.

IBM Network Advisor includes an intuitive interface, and provides an in-depth view of performance measures and historical data. It receives Simple Network Management Protocol (SNMP) traps, syslog event messages, and customizable event alerts, and contains the Advanced Call Home feature that enables you to automatically collect diagnostic information and send notifications to IBM Support for faster fault diagnosis and isolation.

IBM Network Advisor is the preferred tool to manage the IBM b-type Gen 5 fabrics. Chapter 4, “IBM Network Advisor” on page 49 provides detailed information about how to install and configure IBM Network Advisor.

Note: Because IBM Network Advisor is the preferred tool to manage the b-type Gen 5 fabrics, it is used throughout this publication to show how to perform most of the configuration, administration, and troubleshooting tasks. Chapter 4, “IBM Network Advisor” on page 49 provides detailed information about how to plan, deploy, and configure IBM Network Advisor on a Fibre Channel network installation.

1.1.7 Monitoring

There are several monitoring tools and notification methods that allow you to monitor your entire b-type Gen 5 fabric. These tools can be integrated with external applications.

Health monitors

Fabric Watch and MAPS are monitors that allow you to enable each switch to constantly monitor its SAN fabric for potential faults. They automatically alert you to problems long before they become costly failures. MAPS is available only in FOS V7.2.0 or later and it disables Fabric Watch after it is activated (they are mutually exclusive).

Performance monitors

Advanced Performance Monitoring and Flow Vision are performance monitors that integrate with IBM Network Advisor. Flow Vision is available only in FOS V7.2.0 or later. The Fabric Watch feature is no longer available in FOS v7.4.0 and later.

Notification methods

Several alert mechanisms that can be used including email messages, SNMP traps, and log entries.

An email alert sends information about a switch event to a one or multiple specified email addresses.

The SNMP notification method is an efficient way to avoid having to log in to each switch individually, which you must do for error log notifications.

The RASLog (switch event log) can be forwarded to a central station. IBM Network Advisor can be configured as a syslog recipient for the SAN devices.

1.1.8 Brocade Fabric Vision technology

Brocade Fabric Vision technology is an advanced hardware and software architecture that combines capabilities from FOS, b-type Gen 5 devices, and IBM Network Advisor. It helps administrators address problems before they affect operations, accelerate new application deployments, and reduce operational costs. It includes several critical hardware and software monitoring options, management, and diagnostic capabilities that help increase your fabric resiliency, reduce downtime, and optimize performance.

For more information about Brocade Fabric Vision technology, see 7.1, “Fabric Vision” on page 192.

1.1.9 IBM Spectrum Control

IBM Spectrum Control™ belongs to the IBM Spectrum Storage™ family of products. It is a comprehensive, end-to-end data and storage management solution that monitors, automates, and analyzes multivendor storage environments. IBM Spectrum Control provides a single point of control that allows administrators to manage every aspect of the storage infrastructure. It combines management of file, object, SAN, and server-based and software-defined storage. Storage can be viewed and monitored at the application, departmental, or server level.

IBM Spectrum Control enables monitoring of the entire data path including storage systems, devices, and SAN fabric components from multiple vendors. It also supports automatic resource and topology discovery, monitoring and alerts, zone control, and SAN error prediction capabilities.

Chapter 6, “B-type SAN monitoring and management with IBM Spectrum Control” on page 141 provide detailed information about IBM Spectrum Control.

For more information about IBM Spectrum Control, see the product specifications at the following website:

<http://www.ibm.com/software/tivoli/csi/cloud-storage/>

For more information about the IBM Spectrum Storage family, see the following website:

<http://www.ibm.com/systems/storage/spectrum/>

1.2 Product descriptions

IBM SAN products and solutions provide integrated small and medium business (SMB) and enterprise SAN solutions with multiprotocol local, campus, metropolitan, and global storage networking that provides reliable, scalable, and high-performance Fibre Channel connectivity.

The following products are available to meet those wide ranges of requirements.

1.2.1 Enterprise SAN directors

For highest availability and scalability enterprise solutions, these directors are available:

- ▶ SAN768B-2
- ▶ SAN384B-2

1.2.2 Mid-range SAN switches

For scalable, affordable SMB and enterprise solutions, these switches are available:

- ▶ SAN96B-5
- ▶ SAN48B-5

1.2.3 Entry SAN switches

For simple, affordable SMB solutions, these switches are available:

- ▶ SAN24B-4 Express
- ▶ SAN24B-5

1.2.4 Specialty switches

These speciality switches are available:

- ▶ SAN42B-R
- ▶ SAN06B-R

For full product descriptions, see Chapter 5, “Product hardware” on page 107.

Note: For the most current information about available products, see the IBM Fibre Channel Storage Area Networks (SAN) website:

<http://www.ibm.com/systems/storage/san/>



Fabric design overview

This chapter provides a high-level overview of common fabric designs based on IBM b-type Gen 5 16 Gbps products and features. The topics include various topologies along with benefits and limitations for each topology. The guidelines that are outlined in this chapter do not apply to every environment, but they can help guide you through the decisions that you must make for a successful storage area network (SAN) design.

This chapter includes the following sections:

- ▶ Topologies
- ▶ Gen 5 Fibre Channel technology
- ▶ Standard features

2.1 Topologies

This section describes the most common topologies for fabric connectivity, core-edge, or edge-core-edge fabrics. A topology is described in terms of how the switches are interconnected, such as ring, core-edge, and edge-core-edge or fully meshed.

The preferred SAN topology to optimize performance, management, and scalability is a tiered, core-edge topology (sometimes called core-edge or tiered core edge). This approach provides good performance without unnecessary interconnections. At a high level, the tiered topology has many edge switches that are used for device connectivity, and fewer core switches that are used for routing traffic between the edge switches, as shown in Figure 2-1.




Figure 2-1 Four scenarios of tiered network topologies (hops shown in bolded orange)

The scenarios have these characteristics:

- ▶ Scenario A has localized traffic, which can have small performance advantages but does not provide ease of scalability or manageability.
- ▶ Scenario B, also known as edge-core, separates the storage and servers, thus providing ease of management and moderate scalability.
- ▶ Scenario C, also known as edge-core-edge, has both storage and servers on edge switches, which provide ease of management and is much more scalable.
- ▶ Scenario D is a full-mesh topology, and server to storage is no more than one hop. Designing with UltraScale Inter-Chassis Links (ICLs) is an efficient way to save front-end ports, and users can easily build a large (for example, 1536-port or larger) fabric with minimal SAN design considerations.

2.1.1 Edge-core topology

The edge-core topology (Scenario B in Figure 2-1) places initiators (servers) on the edge tier and storage (targets) on the core tier. Because the servers and storage are on different switches, this topology provides ease of management and good performance, with most traffic traversing only one hop from the edge to the core.

The disadvantage of this design is that the storage and core connections are in contention for expansion.

Note: Adding an IBM SAN director as the SAN core can reduce the expansion contention.

2.1.2 Edge-core-edge topology

The edge-core-edge topology (Scenario C in Figure 2-1) places initiators on one edge tier and storage on another edge tier, leaving the core for switch interconnections or connecting devices with network-wide scope, such as Dense Wavelength Division Multiplexers (DWDMs), inter-fabric routers, storage virtualizers, tape libraries, and encryption engines.

Because servers and storage are on different switches, this design enables independent scaling of compute and storage resources, ease of management, and optimal performance. Traffic traverses only two hops from the edge through the core to the other edge. In addition, it provides an easy path for expansion because ports and switches can readily be added to the appropriate tier as needed.

2.1.3 Full-mesh topology

A full-mesh topology (Scenario D in Figure 2-1 on page 10) allows you to place servers and storage anywhere because the communication between source to destination is no more than one hop. With optical UltraScale ICLs, you can build a full-mesh topology that is scalable and cost-effective compared to the previous generation of SAN products.

Note: Hop count is not a concern if the total switching latency is less than the disk I/O timeout value.

2.2 Gen 5 Fibre Channel technology

Gen 5 Fibre Channel technology is designed for high-density server virtualization, cloud architectures, and next generation flash and SSD storage. Gen 5 provides Fabric Vision and 16 Gbps performance.

This section describes the new Gen 5 Fibre Channel technology and its features.

2.2.1 Condor3 ASIC

The Condor3 ASIC is the kernel of the Gen 5 switches. Condor3 ASIC provides unmatched performance compared to its predecessors. Condor3 ASIC increases the frames that are switched per second and the total throughput bandwidth, which is all done with increased energy efficiency. Here are some of the significant Condor3 ASIC specifications:

- ▶ Performance and compatibility:
 - 420 million frames switched per second
 - 768 Gbps of bandwidth
 - 16/10/8/4/2 Gbps speed
 - EX/E/F/M/“D” on any port
- ▶ Industry-leading efficiency with less than 1 watt/Gbps.
- ▶ More scalable across distance:
 - 8000 buffers (four times of what exists on Gen 4)
 - Up to 5000 km distance at 2 Gbps

- ▶ Unmatched investment protection that is compatible with over 30 million existing SAN ports.
- ▶ UltraScale Optical ICLs support optical connections to up to 10 chassis and distances up to 100 meters.
- ▶ Fabric Vision provides advanced diagnostic tests, monitoring, and management that maximizes availability, resiliency, and performance.
- ▶ ClearLink Diagnostic Ports ensure link-level integrity from the server adapter across fabrics and ICLs.
- ▶ Forward Error Correction provides automatic recovery of transmission errors, which enhances reliability of transmission, and in turn results in higher availability and performance.
- ▶ In-flight Encryption/Compression.
- ▶ Secure ISL connectivity and compression of ISL traffic for bandwidth optimization.
- ▶ Using 10 Gbps native Fibre Channel, you can configure any Condor3 port because 10 Gbps Fibre Channel eliminates the need for specialized ports for optical MAN (10 Gbps DWDM) connectivity.
- ▶ ASIC-Enabled Buffer Credit Loss Detection and Automatic Recovery at Virtual Channel Level.
- ▶ Auto Link Tuning for Back-end Ports.
- ▶ E_Port Top Talkers and Concurrency with Fibre Channel Routing.
- ▶ Monitors top bandwidth-consuming flows in real time on each individual ISL and EX_Ports.

2.2.2 Fabric Vision

Brocade Fabric Vision technology is an advanced hardware and software solution that combines capabilities from the Brocade Gen 5 Fibre Channel ASIC.

Fabric Vision is partially compatible with Gen 4 switches and fully supported on Gen 5 switches. It is a set of new software features that works with the new Gen 5 hardware capabilities to provide advanced diagnostic tests, improved monitoring, and management. It is designed to maximize availability, resiliency, and performance, as well as simplify SAN deployment and management.

Licensing: Most Fabric Vision features and capabilities, such as Brocade ClearLink Diagnostics, are included in Fabric OS. MAPS and Flow Vision are available with an optional Fabric Vision license.

If you have existing licenses for both Advanced Performance Monitoring and Brocade Fabric Watch, you will automatically receive the Fabric Vision capabilities when you upgrade to Fabric OS 7.2.0 or later. You do not need to purchase an additional license.

There are many technologies behind Fabric Vision:

- ▶ Switch, director, and adapter ASICs.
- ▶ Delivery in Brocade Fabric Operating System (FOS) begins primarily with Version 7.0, with some features available before Version 7.0.
- ▶ IBM Network Advisor V12.0 and later delivers aspects of the new architecture.

The following are the main Fabric Vision features:

- ▶ **ClearLink Diagnostic Ports:** Ensures optical and signal integrity for Gen 5 Fibre Channel optics and cables.
- ▶ **Fabric Performance Impact (FPI) Monitoring:** Uses predefined thresholds and alerts with MAPS to automatically detect and alert latency, identify slow drain devices, and pinpoint exactly which devices are causing and are affected by a bottlenecked port. FPI monitoring also provides the ability to automatically mitigate the effects of slow drain devices or even resolve the slow drain behavior at the source.

FPI functionality replaces the Bottleneck and Credit recovery functionality in Fabric OS v7.3.0 and later.

- ▶ **Latency Bottleneck Detection:** Enables proactive monitoring, alerting, and visualization of high latency devices and high latency ISLs that are affecting application performance. It simplifies SAN administration by narrowing troubleshooting efforts.
- ▶ **Forward Error Correction (FEC):** Automatically detects and recovers from bit errors, enhancing transmission reliability and performance.
- ▶ **Buffer Credit Recovery at the VC level:** Automatically detects and recovers buffer credit loss at the Virtual Channel level, providing protection against performance degradation and enhancing application availability.
- ▶ **Health and Performance Dashboards:** Provide integration with IBM Network Advisor, providing all the critical information in one window.
- ▶ **Monitoring and Alerting Policy Suite (MAPS):** Policy-based monitoring tool that simplifies fabric-wide threshold configuration and monitoring.
- ▶ **Flow Vision:** A comprehensive tool that enables administrators to identify, monitor, and analyze specific application data flows without using taps.

Note: MAPS and Flow Vision are available only with FOS V7.2 or later.

ClearLink Diagnostic Ports

ClearLink Diagnostic Ports identify and isolate optics and cable problems faster by reducing fabric deployment and diagnostic times. It has these main functions:

- ▶ Non-intrusively verifies transceiver and cable health
- ▶ Tests electrical and optical transceiver components
- ▶ Monitors and trends transceiver health based on uptime
- ▶ Conducts cable health checks
- ▶ Monitors and sets alerts for digital diagnostic tests
- ▶ Ensures predictable application performance over links
- ▶ Provides granular latency and distance measurement for buffer credit assignment
- ▶ Simulates application-level I/O profiles

The background that ensures the advanced diagnostic tests for 16G SFP+ and 16G links is a new diagnostic port type that is known as D_Port. D_Port is used to diagnose optics and cables, and is configured by the user to run diagnostic tests.

D_Port mode allows you to convert a Fibre Channel port into a diagnostic port for testing link traffic, electrical loopbacks, and optical loopbacks between a pair of switches, a pair of access gateways, and a switch. Support is also provided for running D_Port tests between a host bus adapter (HBA) and a switch. The test results that are reported can be useful in diagnosing various port and link problems.

Note: D_Port ports must use 10G or 16G Brocade-branded small form-factor pluggables (SFP).

Understanding D_Port

D_Port does not carry any user traffic, and is designed to run only specific diagnostic tests for identifying link-level faults or failures. To start a port in D_Port mode, you must configure both ends of the link between a pair of switches (or switches configured as Access Gateways), and disable the existing port before you can configure it as a D_Port.

Figure 2-2 illustrates an example D_Port connection between a pair of switches through SFPs (port assignments vary).




Figure 2-2 Example of a basic D_Port connection between switches

After the ports are configured as D_Ports, the following basic test suite is run in the following order, depending on the SFPs that are installed:

1. Electrical loopback (with 16G SFP+ only)
2. Optical loopback (with 16G SFP+ only)
3. Link traffic (with 10G SFPs and 16G SFPs+)
4. Link latency and distance measurement (with 10G SFPs and 16G SFPs+)

Note: Electrical and optical loopback tests are not supported for ICLs.

Figure 2-3 shows the D_port tests capabilities.

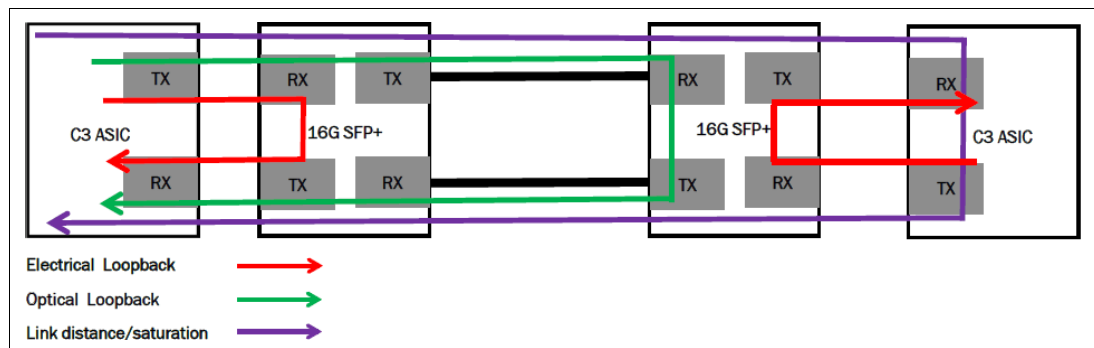


Figure 2-3 D_port tests

The following are the fundamentals of D_Port testing:

- ▶ The user configures the ports on both ends of the connection.
- ▶ After both sides are configured, a basic test suite is initiated automatically when the link comes online, conducting diagnostic tests in the following order:
 - a. Electrical loopback
 - b. Optical loopback
 - c. Link traffic
- ▶ After the automatic test is complete, the user can view results (through command-line interface (CLI) or graphical user interface (GUI)) and rectify any issues that are reported.
- ▶ The user can also start (and restart) the test manually to verify the link.

Advantages of ClearLink Diagnostic ports

Use the D_Port tests for the following situations:

- ▶ Testing a new link before you add it to the fabric
- ▶ Testing a trunk member before you join it with the trunk
- ▶ Testing long-distance cables and SFPs

Tests can be run with the following options:

- ▶ Number of test frames to transmit
- ▶ Size of test frames
- ▶ Duration of the test
- ▶ User-defined test payload
- ▶ Predefined pattern for use in the test payload
- ▶ Testing with FEC on or off (default is off)
- ▶ Testing with credit recovery (CR) on or off (default is off)

For more information about using D_Port, see 7.2, “ClearLink Diagnostics Port” on page 193.

Latency Bottleneck Detection

A *bottleneck* is a port in the fabric where frames cannot get through as fast as they should. The offered load in the port is greater than the achieved egress throughput. Bottlenecks can cause unwanted degradation in throughput on various links. When a bottleneck occurs at one place, other points in the fabric can experience bottlenecks as the traffic backs up.

Note: If you are running Fabric OS v7.3.0 or later, use Fabric Performance Impact Monitor, which replaces Bottleneck Detection. If you are running an earlier version, you can use the Bottleneck Detection feature. Bottleneck Detection and FPI are mutually exclusive.

The Latency Bottleneck Detection feature enables you to perform the following tasks:

- ▶ Prevent degradation of throughput in the fabric.

The bottleneck detection feature alerts you to the existence and locations of devices that are causing latency. If you receive alerts for one or more F_Ports, use the CLI to check whether these F_Ports have a history of bottlenecks.
- ▶ Reduce the time that it takes to troubleshoot network problems.

If you notice one or more applications that are slowing down, you can determine whether any latency devices are attached to the fabric and where they are. You can use the CLI to display a history of bottleneck conditions on a port. If the CLI shows above-threshold bottleneck severity, you can narrow the problem down to device latency rather than problems in the fabric.

A *latency bottleneck* is a port where the offered load exceeds the rate at which the other end of the link can continuously accept traffic, but does not exceed the physical capacity of the link. This condition can be caused by a device that is attached to the fabric that is slow to process received frames and send back credit returns. A latency bottleneck because of such a device can spread through the fabric and can slow down unrelated flows that share links with the slow flow.

A *congestion bottleneck* is a port that is unable to transmit frames at the offered rate because the offered rate is greater than the physical data rate of the line. For example, this condition can be caused by trying to transfer data at 8 Gbps over a 4 Gbps ISL.

You can set alert thresholds for the severity and duration of the bottleneck. If a bottleneck is reported, you can then investigate and optimize the resource allocation for the fabric. Using the zone setup and Top Talkers, you can also determine which flows are destined to the affected F_Ports.

You configure bottleneck detection on a per-fabric or per-switch basis, with per-port exclusions.

Note: Bottleneck detection is disabled by default. The preferred practice is to enable bottleneck detection on all switches in the fabric, and leave it on to gather statistics continuously.

Supported configurations for bottleneck detection

Remember the following configuration rules for bottleneck detection:

- ▶ The switch must be running FOS V6.4.0 or later.
- ▶ Bottleneck detection is supported on Fibre Channel ports and FCoE F_Ports.
- ▶ Bottleneck detection is supported on the following port types:
 - E_Ports
 - EX_Ports
 - F_Ports
 - FL_Ports
- ▶ F_Port and E_Port trunks are supported.
- ▶ Long-distance E_Ports are supported.
- ▶ FCoE F_Ports are supported.
- ▶ Bottleneck detection is supported on 4 Gbps, 8 Gbps, and 16 Gbps platforms.
- ▶ Bottleneck detection is supported in Access Gateway mode.
- ▶ Bottleneck detection is supported whether Virtual Fabrics is enabled or disabled. In VF mode, bottleneck detection is supported on all fabrics, including the base fabric.

For more information about how bottlenecks are configured and displayed, see the *Fabric OS Administrator's Guide*, which you can find at the following website:

<http://my.brocade.com/>

Forward Error Correction

Forward Error Correction (FEC) provides a data transmission error control method by including redundant data (error-correcting code) to ensure error-free transmission on a specified port or port range. When FEC is enabled, it can correct one burst of up to 11-bit errors in every 2112-bit transmission, whether the error is in a frame or a primitive.

FEC is enabled by default. It is supported on E_Ports on 16 Gbps-capable switches and on the N_Ports and F_Ports of an access gateway by using RDY, Normal (R_RDY), or Virtual Channel (VC_RDY) flow control modes. It enables automatically when negotiation with a switch detects FEC capability. This feature is enabled by default and persists after driver reloads and system reboots. It functions with features such as QoS, trunking, and BB_Credit recovery.

Limitations

Here are the limitations of this feature:

- ▶ FEC is configurable only on Gen 5 16 Gbps-capable switches.
- ▶ FEC is supported only on 1860 and 1867 Brocade Fabric Adapter ports operating in HBA mode that are connected to 16 Gbps Gen 5 switches running FOS V7.1 and later.

FEC is not supported in the following scenarios:

- ▶ When the HBA port speed changes to less than 16 Gbps, this feature is disabled.
- ▶ For HBA ports that operate in loop mode or in direct-attach configurations.
- ▶ On ports with DWDM.

Buffer credit recovery at the Virtual Channel level

The management of buffer credits in wide-area SAN architectures is critical. Furthermore, many issues can arise in the SAN network whenever buffer credit starvation or buffer credit loss occurs.

Buffer credit loss detection and recovery is part of the Gen 5 Fibre Channel diagnostic and error recovery technologies. It helps you avoid a “stuck” link condition or an extended lack of buffer credits for an extended time period, resulting in loss of communication across the link.

The IBM b-type Gen 5 16 Gbps Fibre Channel network implements a multiplexed ISL architecture called Virtual Channels (VCs), which enables efficient usage of E_Port to E_Port ISL links.

Virtual Channels create multiple logical data paths across a single physical link or connection. They are allocated their own network resources, such as queues and buffer-to-buffer credits.

Virtual Channels are divided into three priority groups. P1 is the highest priority, which is used for Class F, F_RJT, and ACK traffic. P2 is the next highest priority, which is used for data frames. The data Virtual Channels can be further prioritized to provide higher levels of Quality of Service (QoS). P3 is the lowest priority and is used for broadcast and multicast traffic.

QoS is a licensed traffic shaping feature that is available in FOS. QoS allows the prioritization of data traffic based on the SID and DID of each frame.

Through the usage of QoS zones, traffic can be divided into three priorities: High, medium, and low, as shown in Figure 2-4. The seven data VCs, VC8 through VC14, are used to multiplex data frames based on QoS zones when congestion occurs.




Figure 2-4 Virtual Channel on a QoS enabled ISL

IBM Gen 5 Fibre Channel switches can detect buffer credit loss at the VC level. If the application-specific integrated circuits (ASICs) detect only a single buffer credit lost, can restore the buffer credit without interrupting the ISL data flow. If the ASICs detect more than one buffer credit lost or if they detect a “stuck” VC, they can recover from the condition by resetting the link. This process requires retransmission of frames that were in transit across the link at the time of the link reset.

When a switch automatically detects and recovers buffer credit loss at the VC level, it provides protection against performance degradation and enhances application availability.

Health and Performance Dashboards

The IBM b-type Gen 5 16 Gbps switches, integrated with IBM Network Advisor V12.x and later, can provide all the critical information about the health and performance of a network in a single window. With a customizable dashboard, it is possible to define what is critical and what to monitor.

MAPS also provides a CLI dashboard, which is available when you do not have IBM Network Advisor.

For more information about dashboards, see Chapter 4, “IBM Network Advisor” on page 49.

Monitoring and Alerting Policy Suite

MAPS is an optional SAN health monitor that supported on all switches that are running FOS V7.2.0 or later. It allows you to enable each switch to constantly monitor itself for potential faults and automatically alerts you to problems before they become costly failures.

MAPS tracks various SAN fabric metrics and events. Monitoring fabric-wide events, ports, and environmental parameters enables early fault detection and isolation, and performance measurements.

MAPS provides a set of predefined monitoring policies that allow you to immediately use MAPS on activation.

In addition, MAPS provides customizable monitoring thresholds. These thresholds allow you to configure specific groups of ports or other elements so that they share a common threshold value. You can configure MAPS to provide notifications before problems arise, for example, when network traffic through a port is approaching the bandwidth limit. MAPS lets you define how often to check each switch and fabric measure, and specify notification thresholds. Whenever fabric measures exceed these thresholds, MAPS automatically provides notification by using several methods, including email messages, SNMP traps, and log entries.

The MAPS dashboard provides you with the ability to quickly view what is happening on the switch. This insight helps administrators dig deeper to see details about exactly what is happening on the switch (for example, the kinds of errors and the error count).

MAPS provides a seamless migration of all customized Fabric Watch thresholds, thus allowing you to take advantage of the advanced capabilities of MAPS. MAPS provides extra advanced monitoring, such as monitoring for the same error counters across different periods, or having more than two thresholds for error counters. MAPS also provides support for you to monitor the statistics that are provided by the Flow Monitor feature of Flow Vision.

Note: MAPS is the next generation monitoring tool that replaces Fabric Watch. MAPS cannot coexist with Fabric Watch. MAPS was introduced in FOS v7.2.0. Fabric Watch is no longer available in FOS v7.4.0 or later.

Flow Vision

Introduced in FOS V7.2, Flow Vision¹ is a comprehensive tool that enables administrators to identify, monitor, and analyze specific application data flows.

Flow Vision provides these features:

- ▶ Flow Monitor: Provides comprehensive visibility into application flows in the fabric, including the ability to learn (discover) flows automatically.
- ▶ Flow Mirror: You can use this function to nondisruptively create copies of the application flows, which can be captured for deeper analysis (only mirroring to processor is supported in FOS V7.2).
- ▶ Flow Generator: Test traffic generator for pre-testing the SAN infrastructure (including internal connections) for robustness before deploying the applications.

¹ Available only with FOS V7.2 or later

Note: Using Flow Vision features requires a Fabric Vision license or both Fabric Watch and Advanced Performance Monitor (APM) licenses. Flow Vision is the next generation Performance monitoring tool and it replaces APM. APM cannot coexist with Flow Vision. Flow Vision was introduced in FOS v7.2.0. APM is no longer available in FOS v7.4.0 or later.

For more information, see Chapter 7, “Fabric Vision” on page 191.

2.3 Standard features

This section describes some of the standard features that are available.

2.3.1 Zoning

Zoning is a fabric-based service that enables you to partition your SAN into logical groups of devices that can access each other.

For example, you can partition your SAN into two zones, winzone and unixzone, so that your Windows servers and storage do not interact with your UNIX servers and storage. You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions. For example, you can create a temporary zone to back up nonmember devices.

A device in a zone can communicate only with other devices that are connected to the fabric within the same zone. A device not included in the zone is not available to members of that zone. When zoning is enabled, devices that are not included in any zone configuration are inaccessible to all other devices in the fabric. For more information about this topic, see *Introduction to Storage Area Networks*, SG24-5470.

2.3.2 ISL Trunking

ISL Trunking is an optional software product that is available for all FOS-based Fibre Channel switches, directors, and fabric backbones. ISL Trunking technology optimizes the usage of bandwidth by allowing a group of links to merge into a single logical link, called a trunk group. Traffic is distributed dynamically over this trunk group, achieving greater performance with fewer links. Within the trunk group, multiple physical ports appear as a single port, which simplifies management. Trunking also improves system reliability by maintaining in-order delivery of data and avoiding I/O retries if one link within the trunk group fails.

Figure 2-5 shows the ISL with and without trunking.




Figure 2-5 ISL Trunking

The first example in Figure 2-5 on page 20 (on the left) shows a fabric without trunking. When the trunk is not enabled, there is no traffic optimization, so a link can become congested even when there is bandwidth available on other ISL links.

When the trunking feature is activated, all physical ISLs become a single logical ISL, so the performance is optimized by balancing the traffic across all physical links automatically. Trunking is frame-based instead of exchange-based. Because a frame is much smaller than an exchange, frame-based trunks are more granular and better balanced than exchange-based trunks and provide maximum usage of links.

Note: An ISL Trunking license is required for any type of trunking, and must be installed on each switch that participates in trunking.

Port groups for trunking

To establish a trunk, several conditions must be met, one of which is that all of the ports in a trunk group must belong to the same port group. A port group is a group of eight ports, the members of which are based on the user port number, such as 0 - 7, 8 - 15, 16 - 23, and so on up to the number of ports on the switch. The maximum number of port groups is platform-specific.

Ports in a port group are usually contiguous, but they might not be. For information about which ports can be used in the same port group for trunking, see the appropriate *Hardware Reference Manual* for your switch.

Figure 2-6 shows the port group for the SAN96B-5.




Figure 2-6 SAN96B-5 port group

Supported configurations for trunking

Here are the supported configurations for trunking:

- ▶ Trunk links can be 2 Gbps, 4 Gbps, 8 Gbps, 10 Gbps, or 16 Gbps, depending on the b-type platform.
- ▶ The maximum number of ports per trunk and trunks per switch depends on the b-type platform.
- ▶ You can have up to eight ports in one trunk group to create high-performance ISL trunks between switches, providing up to 128 Gbps (based on a 16 Gbps port speed).
- ▶ If in-flight encryption/compression is enabled, you can have a maximum of two ports per trunk.
- ▶ An E_Port or EX_Port trunk can be up to eight ports wide. All the ports must be next to each other, in the clearly marked groups on the front of the product.

Trunks operate best when the cable length of each trunked link is roughly equal to the length of the others in the trunk. For optimal performance, no more than 30-meters difference is recommended. Trunks are compatible with both short wavelength (SWL) and long wavelength (LWL) fiber-optic cables and transceivers.

Trunking is performed according to the QoS configuration on the ports. That is, in a trunk group, if there are some ports with QoS enabled and some with QoS disabled, they form two different trunks: One with QoS enabled and the other with QoS disabled.

Requirements for trunk groups

The following requirements apply to all types of trunking:

- ▶ The Trunking license must be installed on every switch that participates in trunking.
- ▶ All of the ports in a trunk group must belong to the same port group.
- ▶ All of the ports in a trunk group must meet the following conditions:
 - They must be running at the same speed.
 - They must be configured for the same distance.
 - They must have the same encryption, compression, QoS, and FEC settings.
- ▶ Trunk groups must be between b-type switches. Trunking is not supported on M-EOS or third-party switches.
- ▶ There must be a direct connection between participating switches.
- ▶ Trunking cannot be done if ports are in ISL R_RDY mode. You can disable this mode by using the `portCfgIslMode` command.
- ▶ Trunking is supported only on FC ports. Virtual FC ports (VE_ or VEX_Ports) do not support trunking.

2.3.3 Dynamic Path Selection

Available as a standard FOS feature, exchange-based routing or Dynamic Path Selection (DPS) optimizes fabric-wide performance by automatically routing data to the most efficient available path in the fabric.

DPS is where exchanges or communication between end devices in a fabric are assigned to egress ports in ratios that are proportional to the potential bandwidth of the ISL or trunk group. When there are multiple paths to a destination, the input traffic is distributed across the different paths in proportion to the bandwidth that is available on each of the paths. This configuration improves usage of the available paths, reducing possible congestion on the paths. Every time there is a change in the network (which changes the available paths), the input traffic can be redistributed across the available paths. This is an easy and nondisruptive process when the exchange-based routing policy is engaged.

DPS augments ISL Trunking to provide more effective load balancing. With DPS, traffic loads are distributed at the exchange level across independent ISLs or trunks, and in-order delivery is ensured within the exchange. The combination of trunking and DPS provides immediate benefits to network performance, even in the absence of 16 Gbps devices. DPS in particular can provide performance advantages when connecting to lower-speed 4 Gbps switches. As a result, this combination of technologies provides the greatest design flexibility and the highest degree of load balancing.

Figure 2-7 shows DPS balancing data flow between different ISL trunk paths.




Figure 2-7 Dynamic Path SelectionPort types

The following port types can be part of a b-type device:

- ▶ **D_Port:** A diagnostic port that lets an administrator isolate the ISL to diagnose link-level faults. This port runs only specific diagnostic tests and does not carry any fabric traffic. For more information, see “ClearLink Diagnostic Ports” on page 13.
- ▶ **E_Port:** An expansion port that is assigned to ISL links to expand a fabric by connecting it to other switches. Two connected E_Ports form an ISL. When E_Ports are used to connect switches, those switches merge into a single fabric without an isolation demarcation point. ISLs are non-routed links. For more information, see 2.3.2, “ISL Trunking” on page 20.
- ▶ **EX_Port:** A type of E_Port that connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, an EX_Port appears as a normal E_Port. It follows applicable Fibre Channel standards like other E_Ports. However, the router terminates EX_Ports rather than allowing different fabrics to merge, which happens on a switch with regular E_Ports. An EX_Port cannot be connected to another EX_Port.
- ▶ **F_Port:** A fabric port that is assigned to fabric-capable devices, such as SAN storage devices.
- ▶ **G_Port:** A generic port that acts as a transition port for non-loop fabric-capable devices.
- ▶ **L_/FL_Port:** A loop or fabric loop port that connects loop devices. L_Ports are associated with private loop devices, and FL_Ports are associated with public loop devices.
- ▶ **M_Port:** A mirror port that is configured to duplicate (mirror) the traffic that passes between a specified source port and destination port. This configuration is supported only for pairs of F_Ports. For more information about port mirroring, see the *Fabric OS Troubleshooting and Diagnostics Guide*, which you can find at the following website:
<http://my.brocade.com/>
- ▶ **U_Port:** A universal Fibre Channel port. This is the base Fibre Channel port type, and all unidentified or uninitiated ports are listed as U_Ports.
- ▶ **VE_Port:** A virtual E_Port that is a gigabit Ethernet switch port that is configured for an FCIP tunnel.
- ▶ **VEX_Port:** A virtual EX_Port that connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, a VEX_Port appears as a normal VE_Port. It follows the same Fibre Channel protocol as other VE_Ports. However, the router terminates VEX_Ports rather than allowing different fabrics to merge, which is what happens on a switch with regular VE_Ports.

2.3.4 In-flight encryption and compression

The in-flight encryption and compression features of FOS allow frames to be encrypted or compressed at the egress point of an ISL between two IBM b-type switches. They can then be decrypted or extracted at the ingress point of the ISL. These features use port-based encryption and compression. You can enable the encryption and compression feature for both E_Ports and EX_Ports on a per-port basis. By default, this feature is initially disabled for all ports on a switch.

The purpose of encryption is to provide security for frames while they are in flight between two switches. The purpose of compression is for better bandwidth usage on the ISLs, especially over long distance. An average compression ratio of 2:1 is provided, but your compression ratios will depend on the compressibility of your data. Frames are never left in an encrypted or compressed state when delivered to an end device. Both ends of the ISL must terminate in 16G-capable FC ports.

Encryption and compression can be enabled at the same time for an ISL, or you can enable either encryption or compression selectively. Figure 2-8 shows an example of 16 Gbps links connecting three Brocade switches. One link is configured with encryption and compression, one with just encryption, and one with just compression.

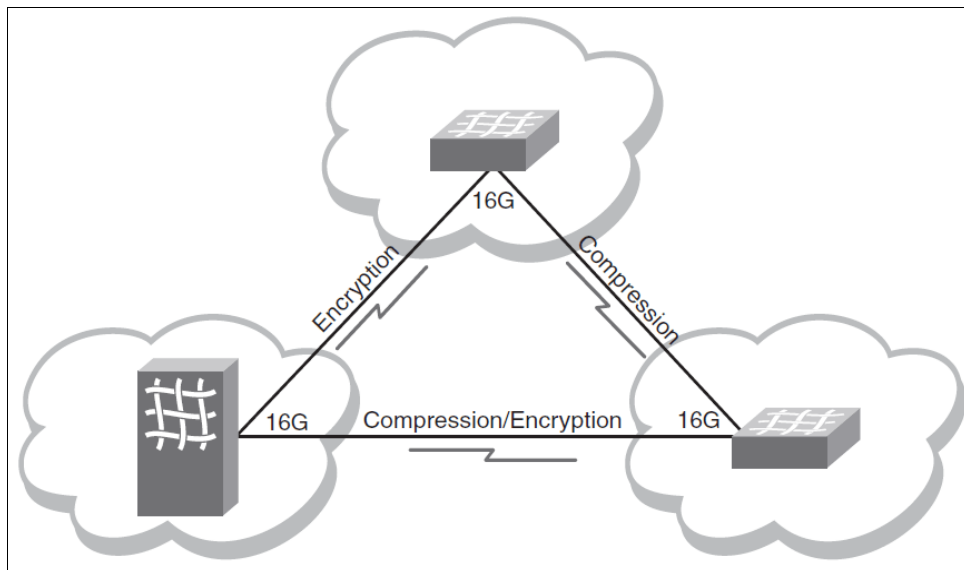


Figure 2-8 Encryption and compression on 16 Gbps ISLs

Note: No license is needed to configure and enable in-flight encryption or compression.

For more information, see the *Fabric OS Administrator's Guide*, which you can find at the following website:

<http://my.brocade.com/>

2.3.5 NPIV

N_Port ID Virtualization (NPIV) enables a single Fibre Channel protocol port to appear as multiple, distinct ports, providing separate port identification within the fabric for each operating system image behind the port (as though each operating system image had its own unique physical port).

NPIV assigns a different virtual port ID to each Fibre Channel protocol device. NPIV enables you to allocate virtual addresses without affecting your existing hardware implementation. The virtual port has the same properties as an N_Port, and can register with all services of the fabric.

Each NPIV device has a unique device PID, Port worldwide name (WWN), and Node WWN, and behaves the same as all other physical devices in the fabric. Multiple virtual devices that are emulated by NPIV appear no different from regular devices that are connected to a non-NPIV port.

The same zoning rules apply to NPIV devices as non-NPIV devices. Zones can be defined by domain, port notation; by WWN zoning; or both. However, to perform zoning to the granularity of the virtual N_Port IDs, you must use WWN-based zoning.

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

For more information, see the *Fabric OS Administrator's Guide*, which you can find at the following website:

<http://my.brocade.com/>

2.3.6 Dynamic Fabric Provisioning

Introduced in FOS V7.0, Dynamic Fabric Provisioning (DFP) simplifies server deployment in a Fibre Channel SAN (FC SAN) environment.

Server deployment typically requires that multiple administrative teams (for example, server and storage teams) coordinate with each other to perform configuration tasks, such as zone creation in the fabric and LUN mapping and masking on the storage device. These tasks must be complete before the server is deployed. Before you can configure WWN zones and LUN masks, you must discover the physical port worldwide name (PWWN) of the server. This requirement means that administrative teams cannot start their configuration tasks until the physical server arrives (and its physical PWWN is known). Because the configuration tasks are sequential and interdependent across various administrative teams, it might take several days before the server is deployed in an FC SAN.

DFP simplifies and accelerates new server deployment and improves operational efficiency by using a fabric-assigned PWWN (FA-PWWN). An FA-PWWN is a “virtual” port WWN that can be used instead of the physical PWWN to create zoning, and LUN mapping and masking. When the server is later attached to the SAN, the FA-PWWN is then assigned to the server.

The FA-PWWN feature allows you to perform the following tasks:

- ▶ Replace one server with another server, or replace failed HBAs or adapters within a server, without having to change any zoning or LUN mapping and masking configurations.
- ▶ Easily move servers across ports or Access Gateways by reassigning the FA-PWWN to another port.
- ▶ Use the FA-PWWN to represent a server in boot LUN zone configurations so that any physical server that is mapped to this FA-PWWN can boot from that LUN, thus simplifying boot over SAN configuration.

Note: For the server to use the FA-PWWN feature, it must be using a Brocade HBA or adapter. For more information, see the release notes for the HBA or adapter versions that support this feature.

Configuration of the HBA must be performed to use the FA-PWWN.

For more information, see the *Fabric OS Administrator's Guide*, which you can find at the following website:

<http://my.brocade.com/>



Management and monitoring tools

This chapter introduces and describes the built-in and external management tools that are available for b-type switches. Detailed explanations for some of these tools, including IBM Network Advisor, Fabric Vision, and IBM Spectrum Control, are detailed in later chapters.

This chapter includes the following sections:

- ▶ Web Tools
- ▶ Command-line interface
- ▶ Storage Management Initiative Agent
- ▶ Fabric Vision
- ▶ IBM Network Advisor
- ▶ IBM Spectrum Control

3.1 Web Tools

Web Tools is a built-in web-based application that enables administrators to monitor and manage single or small fabrics, switches, and ports.

3.1.1 Web Tools introduction and features

The graphical user interface (GUI) of Web Tools can be accessed through any Java capable web browser from any workstation with Ethernet network access to the switches, or from IBM Network Advisor.

For Fabric OS version 6.1.1 and later, Web Tools functionality is tiered and therefore some capabilities are moved to IBM Network Advisor.

If you are migrating from a Web Tools release before Fabric OS version 6.1.1, see Table 3-1 for information about what features are available and a list of changes when you plan for your Fabric OS upgrades.

Table 3-1 Web Tools functionality moved to IBM Network Advisor

Function	FOS 6.1.0 and earlier - Web Tools	FOS 6.1.1 and later - IBM Network Advisor	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.	
Analyze Zone Config	Zone Admin	<ul style="list-style-type: none">▶ 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.▶ 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	

Function	FOS 6.1.0 and earlier - Web Tools	FOS 6.1.1 and later - IBM Network Advisor	Comments
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.	
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 by 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.	

Table 3-2 shows a list of Web Tools functions that were moved to IBM Network Advisor.

Table 3-2 Web Tools functionality moved to IBM Network Advisor

Function	FOS 6.1.0 and earlier - Web Tools	FOS 6.1.1 and later - IBM Network Advisor
Add Un-Zoned Devices	Zone Admin	Configure → Zoning
Analyze Zone Config	Zone Admin	Configure → Zoning
Define Device Alias	Zone Admin	Configure → Zoning
Device Accessibility Matrix	Zone Admin	Configure → Zoning
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
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 → DCC policies → Performance Monitoring	Configure → Zoning
Remove Offline or Inaccessible Devices	Zone Admin	Configure → Zoning
Zone database summary print	Zone Admin	Configure → Zoning

No license is required for most of the fabric-related tasks that can be done with Web Tools. However, additional Web Tools functionality can be added when you obtain the Enhanced Group Management (EGM) license. The EGM license is only required by 8 Gbps platforms. For non-8 Gbps platforms, all functionalities are available without the EGM license. Table 3-3 provides a comparison between features that are included in Basic Web Tools and the features that are enabled by the EGM license.

Table 3-3 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

Feature	Basic Web Tools	Web Tools with EGM license
Cloning a Zone	No	Yes
Config Upload/Download	Yes	Yes
Convenience Function from Tools Menu	No	No
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
Routing and DLS Configuration	No	Yes
Security Policies Tab (such as ACL)	Yes	Yes
Switch Info Tab	Yes	Yes
Switch Status	Yes	Yes

Feature	Basic Web Tools	Web Tools with EGM license
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

3.1.2 System requirements

At the time of writing, Web Tools requires that your browser conform to HTML v4.0, JavaScript v1.0, and Java runtime environment (JRE) 1.7.0_80 or JRE 1.8.0_51 update or later. Read the Fabric Operating System (FOS) release notes for Web Tools requirements.

Note: If there are multiple JRE versions installed, go to the Java Control Panel and clear the lower JRE versions so that Web Tools can launch using the latest JRE version. If you are working with a different Fabric OS version than 7.4.1 and you are having issues with Web Tools and the Java Runtime Environment, see the Web Tools Administrator's Guide of the corresponding Fabric OS version to confirm which JRE version to use.

The operating systems that are shown in Table 3-4 have been tested and certified by for Web Tools (FOS v7.4.1) use.

Table 3-4 Certified and tested platforms

Operating system	Browser
Oracle Enterprise Linux 7.0	Firefox 34.0
Red Hat Enterprise Linux 6.6 Adv	Firefox 34.0
Red Hat Enterprise Linux 7.0 Adv	Firefox 34.0
SUSE Linux Enterprise Server 12	Firefox 34.0
Windows 2008 R2 Enterprise	Firefox 34.0, Internet Explorer 9.0
Windows 8.1	Firefox 34.0, Internet Explorer 11.0
Windows 2012 R2	Firefox 34.0, Internet Explorer 11.0

The platforms shown in Table 3-5 are supported by Brocade for Web Tools (FOS v7.4.1) use.

Table 3-5 Supported platforms

Operating system	Browser
Oracle Enterprise Linux 7.0	Firefox 34.0
SUSE Linux Enterprise Server 11 (SP2) (32-Bit)	Firefox 34.0
SUSE Linux Enterprise Server 12	Firefox 34.0
Windows 2008 Standard	Firefox 34, Internet Explorer 11.0

Operating system	Browser
Windows 7 Professional (32-Bit)	Firefox 34, Internet Explorer 8.0/9.0/10.0/11.0
Windows 7 SP1	Firefox 34, Internet Explorer 8.0/9.0/11.0
Windows 2008 (SP2) Enterprise (64-Bit)	Firefox 34, Internet Explorer 9.0/11.0
Windows 8 Enterprise(64-Bit)	Firefox 34, Internet Explorer 10.0/11.0
Windows Server 2008 R2 (SP1) Enterprise (64-Bit)	Firefox 34, Internet Explorer 9.0/10.0/11.0
Windows Server 2012 Standard (64-Bit)	Firefox 34, Internet Explorer 10.0/11.0

For Microsoft Windows systems, a minimum of 1 GB of RAM for fabrics comprising up to 15 switches, 2 GB of RAM for fabrics comprising more than 15 switches, or an IBM SAN768B/SAN768B-2 with a fully populated blade is required.

Setting the refresh frequency for Microsoft 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.

To set the Internet Explorer options, complete the following steps:

1. Click **Tools** → **Internet Options** in the browser.
2. Click the General tab and click **Settings** in the “Browsing history” section.
3. Click **Every time I visit the webpage** under “Check for newer versions of stored pages,” as shown in Figure 3-1.

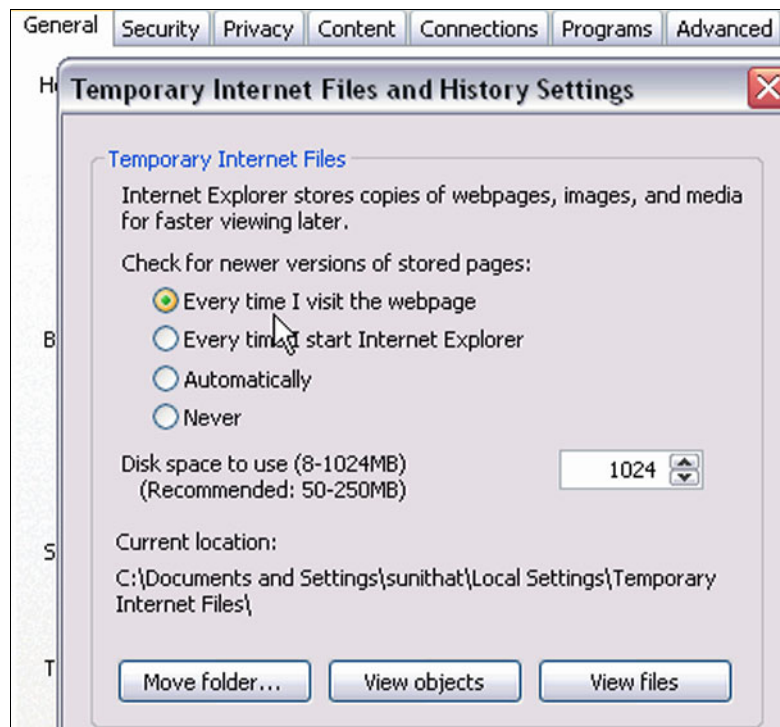


Figure 3-1 Temporary Internet Files and history Settings

Deleting temporary Internet files that are used by Java applications

For Web Tools to operate correctly, you must delete the temporary Internet files used by Java applications. To do so, complete these steps:

1. From the Control Panel, open Java.
2. Click the General tab and click **Settings**.
3. Click **Delete Files** to remove the temporary files that are used by Java applications (see Figure 3-2).

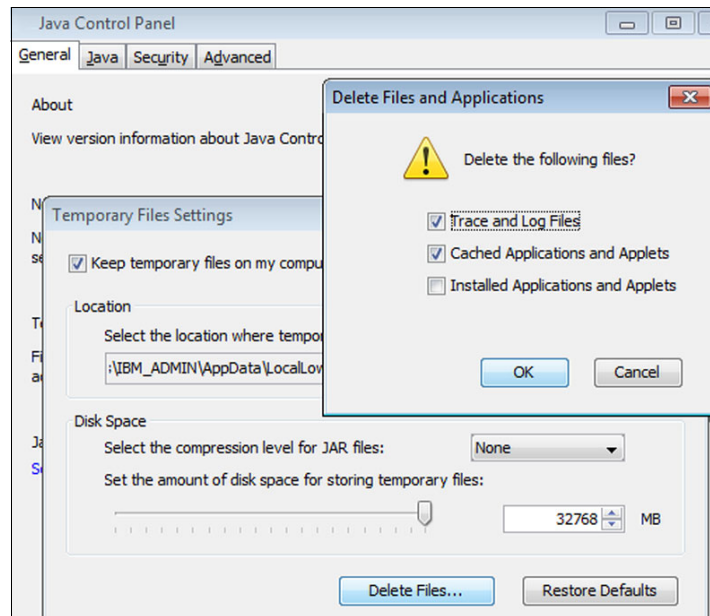


Figure 3-2 Java Temporary Internet Files delete dialog

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 successfully deleted all the temporary files, this list will be empty.

3.1.3 Java plug-in configuration

If you are managing fabrics with more than 10 switches or 1000 ports, 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.

3.1.4 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. When you access these switches, you are opening the remote switch's version of Web Tools, and the functionality available for those switches might vary.

3.1.5 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 v7.4.1, see Table 3-4 and Table 3-5. Web Tools supports both the HTTP and HTTPS protocols.

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**.

The Web Tools login dialog box displays asking for user credentials. When you are successfully logged in, the Web Tools main window appears as shown in Figure 3-3.

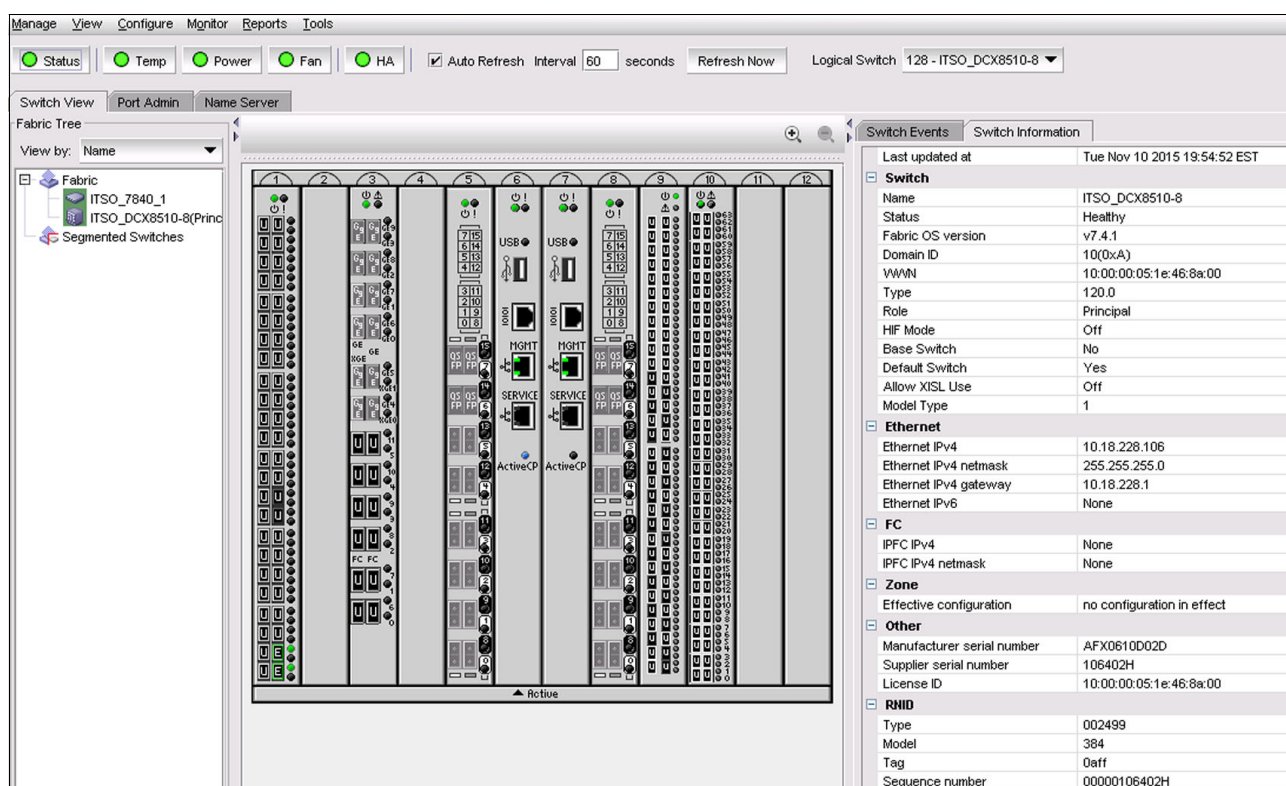


Figure 3-3 Web Tools main window

If you are using Firefox, the browser window is left open. You can close it anytime after the Login dialog box is displayed. If you are using Internet Explorer, the browser window automatically closes when the login dialog box is displayed.

If you have installed EZSwitchSetup on your workstation, the EZSwitchSetup Switch Manager is displayed the first time that 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. Refer to 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 guide describes only the Web Tools interface.

Note: To avoid a potential POODLE attack and establish a secure connection, disable the SSL 3.0 protocol option from your web browser settings.

Logging in to a Virtual Fabric

If you are logging in to a platform with Virtual Fabrics enabled, the log in window provides the option for whether logging in to a specific Logical Switch or to the Home Logical Fabric. To do so, complete these steps:

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 3-4. 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.

The image shows a web browser window titled "Please Login". Inside the window, there is a text prompt "Please enter user name and password." Below this, there is a "Resource" field with the value "10.18.228.106". There are two input fields: "User Name" and "Password". Below these is a section titled "Virtual Fabric" containing two radio button options: "Home Logical Fabric" (which is selected) and "User Specified Logical Fabric" (which is unselected). Next to the "User Specified Logical Fabric" option is an empty input field for a fabric ID number. At the bottom of the window are three buttons: "OK", "Cancel", and "Options <<".

Figure 3-4 Web Tools login window

2. Use one of the following methods to log in to a logical fabric:
 - To log in to the home logical fabric, select **Home Logical Fabric** and click **OK**.
 - To log in to a logical fabric other than the home logical fabric, select **User Specified Logical Fabric**, enter 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 window 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.

To log in to an Admin Domain, complete 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 3-5.




Figure 3-5 Login dialog box with Admin Domain options

2. Use one of the following methods to log in to an Admin Domain:
 - To log in to the home domain, select **Home Domain** and click **OK**.
 - To log in to an Admin Domain other than the home domain, select **User Specified Domain**, enter 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 window displays from which you can choose a valid Admin Domain or click **Cancel** to log in to your home domain, as shown in Figure 3-6.

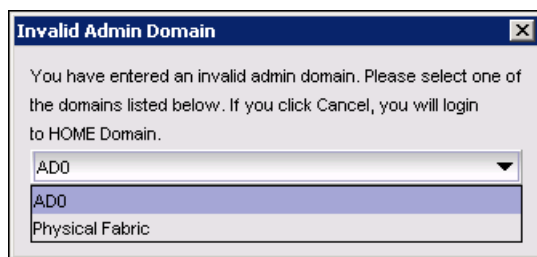


Figure 3-6 Invalid Admin Domain dialog box

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 clicking **Logout**, in 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 start 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. Each role has a set of predefined 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 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 following roles are available:

- ▶ **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.

Note: 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 in 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 later 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 directories with name format `<core switch name-IP address-Switch WWN>`, which includes 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.

For additional Web Tools information, such as use cases and administration, see the *Web Tools Administrator's Guide 7.4.1* at the following link:

<http://www.brocade.com/content/html/en/administration-guide/fos-741-webtools>

Specifications of the Fabric environment in this chapter: The SAN768B-2 fabric backbone and FOS code version 7.4.1 were used to write the Web Tools contents of this chapter. No major changes have occurred since Fabric OS version 6.1.1. If you require Web Tools information that is associated with a different Fabric OS code, see the *Web Tools Administrator's Guide* for that version.

3.2 Command-line interface

The command-line interface (CLI) enables an administrator to monitor and manage individual switches, ports, and entire fabrics from a standard workstation. It is accessed through the Ethernet network through either Telnet or SSH connections or physically through the serial console on the switch.

Fabric OS uses RBAC to control access to all Fabric OS operations. Each feature is associated with an RBAC role. You need to know which role is allowed to run a command, make modifications to the switch, or view the output of the command. Consider these requirements when you create users.

To determine which RBAC role you need to run a command, review the “Role-Based Access Control” section in the *Fabric OS Administrator's 7.4.1 Guide* at the following link:

<http://www.brocade.com/content/html/en/administration-guide/fos-741-adminguide/>

When you use the CLI, remember that both commands and options are case-sensitive. See the Fabric OS command reference for a description of all of the available CLI commands, their options, and their uses.

You can download the Fabric OS command reference for Fabric OS v7.4.1 at the following link:

<http://www.brocade.com/content/html/en/command-reference-guide/fos-741-commandref/>

3.3 Storage Management Initiative Agent

The Brocade Storage Management Initiative (SMI) Agent is a proxy agent that allows applications to manage IBM b-type SAN infrastructures from a single access point by communicating with multiple switches and multiple fabrics.

The SMI Agent is integrated with IBM Network Advisor, and is therefore included in any IBM Network Advisor standard installation. The SMI agent provides greater scalability, reduced overhead of deploying and managing multiple applications, and fewer applications contacting devices for information over the host-based predecessor agents.

You can also deploy the SMI agent separately. To do so, you must download IBM Network Advisor installation package and select **SMI Agent Only** as the installation type during the wizard.

3.4 Fabric Vision

This section describes Fabric Vision. It provides the following information:

- ▶ Base Fabric OS and Fabric Vision enabled features
- ▶ Fabric OS and Fabric Vision licensing considerations

3.4.1 Base Fabric OS and Fabric Vision enabled features

The following features are included in the base Fabric OS (FOS) releases:

- ▶ ClearLink Diagnostics: Helps to ensure optical and signal integrity for Gen 5 Fibre Channel optics and cables by using the *ClearLink Diagnostic Port (D_Port)* technology. It can perform a complete optical, electrical, and link saturation test to ensure reliable connections to validate configuration before deployment and support of high-performance fabrics.
- ▶ Bottleneck Detection: Identifies and provides alerts about device or ISL congestion and abnormal levels of latency in the fabric.

Note: Bottleneck detection is no longer available in Fabric OS release 7.4.0 or later. In version 7.4.0 and later, use Fabric Performance Impact (FPI) instead.

- ▶ Forward Error Correction (FEC): Enables recovery from bit errors in links, enhancing transmission reliability and performance.
- ▶ Credit Loss Recovery: Helps overcome performance degradation and congestion due to buffer credit loss.
- ▶ Monitoring and Alerting Policy Suite (MAPS) Basic Monitoring: Monitors system resources through MAPS basic monitoring policy. In releases of Fabric OS before 7.4.0, Fabric Watch did environmental monitoring and switch status policy monitoring, even if the Fabric Watch license was not active. Similarly, in Fabric OS 7.4.0 and later, many of these features are monitored by MAPS even if the Fabric Vision license is not active.

MAPS Basic Monitoring will look switch status overall health and switch resources changes (field-replaceable unit (FRU) state, Flash memory space, Temperature, CPU usage, Memory usage, and Ethernet management port).

The following advanced technologies and capabilities are available with the optional Fabric Vision technology license:

- ▶ Monitoring and Alerting Policy Suite (MAPS): Provides a new, easy-to-use solution for policy-based threshold monitoring and alerting. MAPS proactively monitors the health and performance of the SAN infrastructure to ensure application uptime and availability. By leveraging pre-built rule/policy-based templates, MAPS simplifies fabric-wide threshold configuration, monitoring, and alerting. Administrators can configure the entire fabric (or multiple fabrics) at one time by using common rules and policies, or customize policies for specific ports or switch elements with IBM Network Advisor. MAPS provides the following features:
 - Policy-based monitoring, including the following features:
 - Pre-defined monitoring groups and pre-validated monitoring policies that users can use. Pre-defined monitoring groups include switch ports attached to servers, switch ports attached to storage, E_Ports, short-wavelength small form-factor pluggables (SFPs), long-wavelength SFPs, and more. Pre-defined monitoring policies include aggressive, moderate, and conservative policies based on monitoring thresholds and actions.

- Flexibility to create custom monitoring groups, such as switch ports attached to high-priority applications and another group of switch ports attached to low-priority applications, and monitor each group according to its own unique rules.
 - Flexible monitoring rules to monitor a specific counter for different threshold values, then take different actions when each threshold value is crossed. For example, users can monitor a CRC error counter at a switch port and generate a RASlog when the error rate reaches two per minute, send an email notification when the error rate is at five per minute, and fence a port when the error rate exceeds ten per minute.
 - Ability to monitor both sudden failures and gradually deteriorating conditions in the switch. For example, MAPS can detect and alert users if a CRC error counter suddenly increases to five per minute, or gradually increases to five per day.
 - Support for multiple monitoring categories, enabling monitoring of the overall switch status, switch ports, SFPs, port blades, core blades, switch power supplies, fans, temperature sensors, security policy violations, fabric reconfigurations, CPU and memory utilization, traffic performance, Fibre Channel over IP (FCIP) health, scalability limits, and so on.
 - Support for multiple alerting mechanisms (RASlogs, Simple Network Management Protocol (SNMP) traps, email notifications). Administrators can tailor the frequency of alert messages to reduce duplicate notifications.
- Fabric Performance Impact (FPI) Monitoring: FPI can be referred as the “next generation” bottleneck detection. It uses predefined thresholds and alerts with MAPS to automatically detect and alert administrators to severe levels or transient spikes of latency. It also identifies slow drain devices that might affect the network. This feature uses advanced monitoring capabilities and intuitive MAPS dashboard reporting to indicate various latency severity levels, pinpointing exactly which devices are causing or affected by a slow drain device scenario. FPI monitoring also allows you to automatically mitigate the effects of slow drain devices or even resolve the slow drain behavior at the source. Available ports actions are port fencing, port toggle, and Slow Drain Device Quarantine (SDDQ):
- Port fencing: MAPS supports port fencing for both E_Ports and F_Ports. This action automatically takes ports offline when configured thresholds in a rule are exceeded. Port fencing immediately takes ports offline, which might cause loss of traffic.
 - Port toggle: The toggle action temporarily disables a port and then reenables it, allowing the port to reset and recover from some device-based issues. If the issue does not get resolved, the port toggling action will suspend the port for a longer period, forcing the port traffic to switch over to a different path if one is available.
 - Slow Drain Device Quarantine (SDDQ): SDDQ with quality of service (QoS) monitoring, allows MAPS to identify a slow-draining device and quarantine it by automatically moving all traffic that is destined to the F_Port that is connected to the slow-draining device to a low-priority virtual channel so that the traffic in the original virtual channel does not experience back pressure.
- CLI dashboard: This dashboard of health and error statistics provides at-a-glance views of switch status and various conditions that are contributing to the switch status, enabling you to get instant visibility into any hot spots at a switch level and take corrective actions. Provides overall status of the switch health and the status of each monitoring category, including any out-of-range conditions and the rules that were triggered. The dashboard also provides historical information about the switch status for up to the last seven days. It automatically provides raw counter information for various error counters.

- Configuration and Operational Monitoring Policy Automation Services Suite (COMPASS): Simplifies deployment, safeguards consistency, and increases operational efficiencies of larger environments with automated switch and fabric configuration services. Administrators can configure a template or adopt an existing configuration as a template and seamlessly deploy the configuration across the fabric. In addition, they can ensure that settings do not drift over time with COMPASS configuration and policy violation monitoring within IBM Network Advisor dashboards.
- Proactive flow monitoring using MAPS: MAPS can monitor flows and generate alerts based on user-defined rules, enabling users to monitor and be alerted when established thresholds are exceeded.
- Flow Vision: Enables administrators to identify, monitor, and analyze specific application flows to simplify troubleshooting, maximize performance, avoid congestion, and optimize resources. Flow Vision includes these features:
 - Flow Monitor: Provides comprehensive visibility into flows within the fabric, including the ability to automatically learn flows and non-disruptively monitor flow performance. Administrators can monitor all flows from a specific host to multiple targets/LUNs, from multiple hosts to a specific target/LUN, or across a specific ISL. Additionally, they can perform LUN-level monitoring of specific frame types to identify resource contention or congestion that is impacting application performance. Flow Monitor provides the following capabilities:
 - Comprehensive visibility into application flows in the fabric, including the ability to learn (discover) flows automatically.
 - Monitoring of application flows within a fabric at a given port.
 - Pre-defined flow to discover all application flows going through all device ports on a switch for network provisioning and planning.
 - Statistics that are associated with the specified flows to gain insights into application performance, such as transmit frame count, receive frame count, transmit throughput, receive throughput, SCSI Read frame count, SCSI Write frame count, number of SCSI Reads and Writes per second (IOPS), and so on.
 - When NPIV is used on the host, users can monitor virtual machine (VM)-to-LUN-level performance.
 - Monitoring of various frame types at a switch port to provide deeper insights into the storage I/O access pattern at the LUN level, reservation conflicts, and I/O errors. Examples of frame types include SCSI Read, SCSI Write, SCSI Reserve, ABTS, and BA_ACC.
 - Flow Monitor is integrated with MAPS to enable threshold-based monitoring and alerting of flows.
 - Flow Generator: A built-in traffic generator for pre-testing and validating the data center infrastructure, including route verification and integrity of optics, cables, ports, back-end connections, and ISLs, for robustness before deploying applications. Flow Generator allows users to perform these actions:
 - Configure a Gen 5 Fibre Channel-capable port as a simulated device that can transmit frames at a 16 Gbps line rate.
 - Emulate a Gen 5 Fibre Channel SAN without actually having any hosts, targets, or SAN testers, and pre-test the entire SAN fabric.
 - Flow Mirror: As storage networks get larger and more complex, non-intrusive diagnostic tools are becoming increasingly important to help identify problems without disturbing the operating fabric. Flow mirroring is a diagnostic feature within Flow Vision that addresses this need.

Flow Mirror provides you with the ability to perform these actions:

- Non-disruptively create copies of application flows that can optionally be captured for deeper analysis.
- Define a traffic pattern and create a real-time copy of this traffic, allowing you to analyze a live system without disturbing existing connections. You can also use this feature as a way to view traffic that is passing through a port.
- Use a MAPS logical port group for either an ingress or egress port.
- Use a predefined flow to mirror SCSI command, first response, status, and ABTS frames from all the Gen 5 F_Ports on a switch.

Flow Mirror duplicates the specified frames in a user-defined flow, and sends them to one of these destinations:

- The local switch CPU. This form is called CPU flow mirroring (CFM), and has a limit of 256 frames per second.
- An analyzer/packet sniffer connected through a port in the SAN. The bandwidth limit for this flow mirroring technique is the bandwidth of the mirror destination port. This form is called Local flow mirroring (LFM), and mirrors the flow to a port on the same physical switch.

Note: Any mirroring possible in CFM is also possible in LFM, but LFM and CFM are mutually exclusive. Only one form of mirroring can be active at a time.

3.4.2 Fabric OS and Fabric Vision licensing considerations

Fabric Vision license replaces both the Advanced Performance Monitoring (APM) and Fabric Watch (FW) licenses. Any switch that is currently licensed for both APM and FW are Fabric Vision capable switches. Fabric Vision features become available on these switches when they are upgraded to Fabric OS code versions 7.2.x or later.

Note: In cases where a switch is already licensed for either Fabric Watch or APM (but not both), the missing license can be procured to enable the Fabric Vision capabilities.

Several Fabric Vision features, such as MAPS and Flow Vision, remained optional across the 7.2.X and 7.3.X family code releases. However, for Fabric OS version 7.4.0, Fabric Watch has been deprecated (removed) and MAPS basic monitoring replaces it for health monitoring.

Similarly, APM has been replaced by Flow Vision, which has been enhanced to support monitoring features equivalent to APM monitors, such as Top Talkers. Table 3-6 summarizes these monitoring features changes.

Table 3-6 Monitoring features changes

Existing monitoring feature	Replaced by
Fabric Watch	Monitoring and Alerting Policy Suite (MAPS)
APM	Flow Vision
Bottleneck detection	Fabric Performance Impact (FPI)

With the removal of APM, a number of commands have either been removed from Fabric OS entirely, or have been modified. See the *Fabric OS Command Reference* document at the following link for information about specific commands that have been removed or changed:

<http://www.brocade.com/content/html/en/command-reference-guide/fos-741-commandref/>

If you do not want to migrate to these new tools and want to continue using your existing Fabric Watch, bottleneck detection and Advanced Performance Monitoring configurations, you must stay on Fabric OS code version 7.3.X or earlier.

Attention: Upgrading to FOS 7.4.0 or later without converting the existing Fabric Watch configuration to MAPS custom policies will result in the loss of any Fabric Watch custom threshold configurations.

If you want to use your existing Fabric Watch threshold rules in MAPS, you must convert them to a custom MAPS policy before you install Fabric OS 7.4.0. When you complete the conversion, you can enable and configure MAPS to use the custom converted policy.

If you do not install the Fabric Vision license, Fabric Watch rules to MAPS custom policy conversion is not possible. In this case, if you still want to activate MAPS (or you are forced to if you are upgrading to FOS 7.4.0 or later), you will only get the MAPS basic monitoring capabilities. See 3.4.1, “Base Fabric OS and Fabric Vision enabled features” on page 41 for more details about MAPS basic monitoring.

Note: Fabric Vision was introduced in Fabric OS Version 7.2.0. If you are upgrading to a target FOS code either 7.2.X or 7.3.X, you can still choose to use Fabric Watch as your switch health monitoring solution. You can also migrate the customized Fabric Watch rules to a custom MAPS policy, enable MAPS, and activate that policy.

For more information about Fabric Vision technology, see Chapter 7, “Fabric Vision” on page 191.

3.5 IBM Network Advisor

IBM Network Advisor is a management application that provides easy and centralized management of the network, and quick access to all product configuration applications. You can configure, manage, and monitor a network with ease by using this application.

Figure 3-7 illustrate the various areas of the management application's SAN tab.

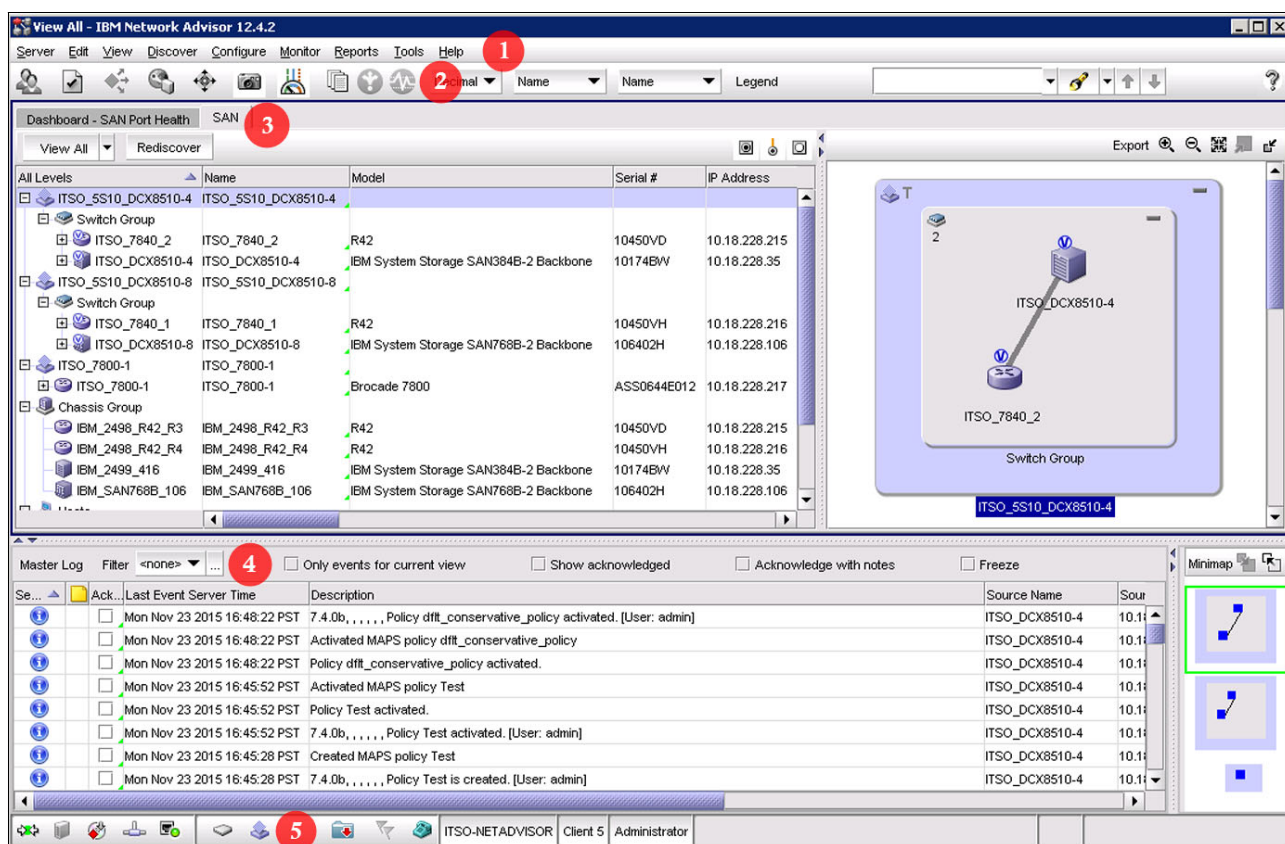


Figure 3-7 IBM Network Advisor SAN tab main window

The numbers in Figure 3-7 coorespond with these items:

1. Menu bar: Lists commands that you can perform on the Management application. The available commands vary depending on which tab (SAN or Dashboard) you select.
2. Toolbar: Provides buttons that enable quick access to windows and functions. The available buttons vary depending on which tab (SAN or Dashboard) you select.
3. Tabs: Provides quick access to the following views:
 - Dashboard tab: Provides a high-level overview of the network managed by the Management application server.
 - SAN Tab: Displays the Master Log, Minimap, Connectivity Map (topology), and Product List.
4. Master Log: Displays the Master Log.
5. Status bar: Displays the connection, port, product, fabric, special event, Call Home, and backup status, as well as Server and User data.

Note: Because IBM Network Advisor is the preferred tool to manage the b-type Gen 5 fabrics, it is used throughout this publication to show how to perform most of the configuration, administration, and troubleshooting tasks. Chapter 4, “IBM Network Advisor” on page 49 provides detailed information about how to plan, deploy, and configure IBM Network Advisor on a Fibre Channel network installation.

3.6 IBM Spectrum Control

IBM Spectrum Control data management and storage management solutions deliver comprehensive monitoring, automation, and analytics capabilities to your storage environments. They help to address the two most significant storage costs: Efficient physical capacity management and storage administration. You can get important insights into cloud, virtualized, and software-defined storage environments to make more informed and better business decisions.

IBM Spectrum Control solutions can help you to accomplish these tasks:

- ▶ Identify and categorize storage assets for file, block, and object data.
- ▶ Generate departmental and application views of storage.
- ▶ Optimize data placement within the infrastructure.
- ▶ Identify unused and wasted storage space.

The IBM Spectrum control offerings are:

- ▶ On-premises: IBM Virtual Storage Center provides a richer experience with an active administration console, automated storage provisioning, and tier optimization. Virtual Storage Center includes a storage service catalog, drivers for OpenStack and VMware, and application-aware snapshot management.

For more information about b-type SAN monitoring and management with IBM Spectrum Control (formerly Tivoli Storage Productivity Center), see Chapter 6, “B-type SAN monitoring and management with IBM Spectrum Control” on page 141.

- ▶ On the Cloud: IBM Storage Insights provides similar function with no hardware to install and only one VM to deploy when one or more IBM storage systems are installed. This offering is delivered as a SaaS solution of the IBM Service Engage offerings catalog.

Note: IBM Storage Insights does not provide any level of integration with Fibre Channel switches and fabrics, for which reason all of the IBM Spectrum Control content that is described on this book just focuses on the IBM Virtual Storage Center solution.

For more information about the IBM Spectrum Control offerings, see the following link:

<http://www.ibm.com/software/tivoli/csi/cloud-storage/>



IBM Network Advisor

This chapter describes how to deploy IBM Network Advisor in your network. It also describes how to upgrade an existing IBM Network Advisor installation and the corresponding version upgrade path.

It also introduces user, fabric, and dashboard management along with some preferred practices such as call home and weekly backups.

This chapter includes the following sections:

- ▶ Planning for server and client system requirements
- ▶ IBM Network Advisor 12.4.2 upgrade path
- ▶ Downloading the software
- ▶ Pre-installation requirements
- ▶ Syslog troubleshooting
- ▶ IBM Network Advisor Version 12.4.2 installation
- ▶ Upgrading to IBM Network Advisor V12.4.2 from an existing IBM Network Advisor installation
- ▶ IBM Network Advisor web client
- ▶ User, device discovery, and dashboard management
- ▶ New features of IBM Network Advisor V12.4.2
- ▶ IBM Network Advisor Dashboard overview
- ▶ Scheduling daily or weekly backups for the fabric configuration
- ▶ Call Home

4.1 Planning for server and client system requirements

The following Fabric OS (FOS) platforms are supported by IBM Network Advisor V12.4.2:

- ▶ Fabric OS (FOS) V5.0 or later in a pure FOS fabric
- ▶ Fabric OS (FOS) V6.0 or later in a mixed fabric
- ▶ Fabric OS (FOS) V7.0 or later

For more information about the hardware and software that is supported for IBM Network Advisor V12.4.2, see the “About this document” section of the *Network Advisor 12.4.2 SAN User Manual* at the following link:

<http://www.brocade.com/content/html/en/user-guide/networkadvisor-1242-san-manual>

Two IBM Network Advisor storage area network (SAN) packages are available, depending on the size of your SAN network:

- ▶ IBM Network Advisor Professional Plus, licensed to manage up to 36 fabrics and 2,560 ports.
- ▶ IBM Network Advisor SAN Enterprise option, licensed to manage up to 100 fabrics and 15,000 ports.

Requirement: IBM Network Advisor Enterprise is required to manage the SAN768B-2 and SAN768B 8 slot directors.

For more information about features that are supported on each edition, see the “Edition feature support” section of the *Network Advisor Software Licensing Guide 12.4.2* at the following link:

<http://www.brocade.com/content/html/en/software-licensing-guide/networkadvisor-1242-licenseguide>

4.1.1 Operating system and hardware requirements for Server and Client

The following section describes server and client operating system requirements and selecting the upgrade path when upgrading to Version 12.4.2 from an earlier version of IBM Network Advisor.

To prevent any conflicts with other applications that use the same resources and ports (such as SNMP, web server, and so on), run IBM Network Advisor on a dedicated server.

Note: Enterprise edition and Professional Plus edition are not supported on 32-bit operating systems.

Table 4-1 summarizes the required operating systems for the IBM Network Advisor server deployment.

Table 4-1 Server operating system requirements

Operating System	Version
Windows	<ul style="list-style-type: none"> ▶ 2008 R2 Data Center Edition (x86 64-bit) ▶ 2008 R2 Standard Edition (x86 64-bit) ▶ 2008 R2 Enterprise Edition (x86 64-bit) ▶ 2012 Data Center Edition (x86 64-bit) ▶ 2012 Standard Edition (x86 64-bit) ▶ 2012 R2 Data Center Edition (x86 64-bit) ▶ 2012 R2 Standard Edition (x86 64-bit) ▶ 8 Enterprise (x86 64-bit) ▶ 8.1 Enterprise (x86 64-bit)
Linux	<ul style="list-style-type: none"> ▶ RedHat Enterprise 6.4 Advanced (x86 64-bit) ▶ RedHat Enterprise 6.5 Advanced (x86 64-bit) ▶ RedHat Enterprise 6.6 Advanced (x86 64-bit) ▶ RedHat Enterprise 7.0 Advanced (x86 64-bit) ▶ SuSE Enterprise Server 11.3 (x86 64-bit) ▶ SuSE Enterprise Server 12 (x86 64-bit) ▶ Oracle Enterprise 6.4 (x86 64-bit) ▶ Oracle Enterprise 6.5 (x86 64-bit) ▶ Oracle Enterprise 7.0 (x86 64-bit)
Guest VMs (supports all server OS versions available for Windows and Linux.)	<ul style="list-style-type: none"> ▶ VMware ESXi 5.11 ▶ VMware ESXi 5.5 ▶ Microsoft Hyper-V (Hyper-V Server 2008 R2, Windows Server 2012, Windows Server 2012 R2 Data Center) ▶ KVM (RH 6.5)

Table 4-2 summarizes the client operating system requirements. IBM Network Advisor clients are supported on 32-bit and 64-bit Windows and Linux systems.

Table 4-2 Client operating system requirements

Operating System	Version
Windows	<ul style="list-style-type: none"> ▶ 7 Enterprise (x86 32-bit) ▶ 8 Enterprise (x86 32-bit) ▶ 8.1 Enterprise (x86 32-bit) ▶ 2008 R2 Data Center Edition (x86 64-bit) ▶ 2008 R2 Standard Edition (x86 64-bit) ▶ 2008 R2 Enterprise Edition (x86 64-bit) ▶ 2012 Data Center Edition (x86 64-bit) ▶ 2012 Standard Edition (x86 64-bit) ▶ 2012 R2 Data Center Edition (x86 64-bit) ▶ 2012 R2 Standard Edition (x86 64-bit) ▶ 7 Enterprise (x86 64-bit) ▶ 8 Enterprise (x86 64-bit) ▶ 8.1 Enterprise (x86 64-bit)

Operating System	Version
Linux	<ul style="list-style-type: none"> ▶ RedHat Enterprise 6.4 Advanced (x86 32 bit) ▶ RedHat Enterprise 6.5 Advanced (x86 32 bit) ▶ RedHat Enterprise 6.6 Advanced (x86 32 bit) ▶ RedHat Enterprise 7.0 Advanced (x86 32 bit) ▶ SuSE Enterprise Server 11.3 (32 bit) ▶ SuSE Enterprise Server 12 (32 bit) ▶ Oracle Enterprise 6.4 (x86 32 bit) ▶ Oracle Enterprise 6.5 (x86 32 bit) ▶ Oracle Enterprise 7.0 (32 bit) ▶ RedHat Enterprise 6.4 Advanced (x86 64 bit) ▶ RedHat Enterprise 6.5 Advanced (x86 64 bit) ▶ RedHat Enterprise 6.6 Advanced (x86 64 bit) ▶ RedHat Enterprise 7.0 Advanced (x86 64 bit) ▶ SuSE Enterprise Server 11.3 (x86 64 bit) ▶ SuSE Enterprise Server 12 (x86 64 bit) ▶ Oracle Enterprise 6.4 (x86 64 bit) ▶ Oracle Enterprise 6.5 (x86 64 bit) ▶ Oracle Enterprise 7.0 (x86 64 bit)
Guest VMs (supports all client OS versions available for Windows and Linux.)	<ul style="list-style-type: none"> ▶ VMware ESXi 5.1 ▶ VMware ESXi 5.5 ▶ Microsoft Hyper-V (Hyper-V Server 2008 R2, Windows Server 2012, Windows Server 2012 R2 Data Center) ▶ KVM (RH 6.5)

Table 4-3 summarizes the minimum system requirements for running IBM Network Advisor on Windows and Linux.

Table 4-3 Minimum system requirements for IBM Network Advisor

Resources	Professional Edition	Professional Plus Edition
Installed memory	4 GB (32-bit) 6 GB (64-bit)	6 GB
Processor core count (including physical and logical ones)	2	2
Disk space	10 GB	20 GB

Table 4-4 summarizes the preferred system requirements for running IBM Network Advisor on Windows and Linux based on the size of the environment.

Table 4-4 Minimum system requirements for IBM Network Advisor

Resources	Small	Medium	Large
Installed memory	16 GB	16 GB	16 GB
Processor core count (including physical and logical ones)	2	4	8
Disk space	20 GB	80 GB	100 GB

Note: If you enable **supportsave** to run periodically or configure the IBM Network Advisor as the upload failure data capture location for monitored switches, then extra disk space is required. Each switch **supportsave** file is approximately 5 MB, and each upload failure data capture file is approximately 500 KB. To determine the disk space requirements, multiply the frequency of scheduled **supportsave** commands by 5 MB and the expected upload failure data capture files by 500 KB before the planned periodic purge activity.

4.1.2 Browser requirements for IBM Network Advisor 12.4.2

The use of IBM Network Advisor and the launch of Element Manager (Web Tools) from the application are supported from the following browsers with a Java plug-in:

- ▶ Browsers:
 - Windows Internet Explorer 11.0.9 on Windows
 - Firefox 24 and later on Windows or Linux
 - Google Chrome 33 on Windows
- ▶ Java Plug-ins: For the current supported Java runtime environment (JRE) version for Network Advisor and Web Tools, refer to the Release Notes.

Note: For higher performance, use a 64-bit JRE. If the minimum system requirement is not met, you will be blocked from the configuration and an error message will be displayed.

For information about JRE patches, see the following website:

<http://www.oracle.com/technetwork/java/javase/downloads/index.html>

4.1.3 IBM Network Advisor Server and concurrent client connections

Network Advisor has the following client and server system requirements:

- ▶ In the Professional edition, a single server supports a single client, which must be a local client only.
- ▶ In Professional Plus and Enterprise editions, a single server supports a maximum of 25 clients, which can be local or remote on 64-bit servers. To support more than eight clients, you must make the following changes to your configuration:
 - Increase the server memory size. You can configure the server memory size from the Options window, Memory Allocations pane. For more information, see the Network Advisor User Manual or online help.
 - Increase the PostgreSQL database shared buffers memory allocation to 1024 MB by editing the `Install_Home\data\databases\postgresql.conf` file.

4.2 IBM Network Advisor 12.4.2 upgrade path

The following upgrade path must be followed when upgrading from earlier product versions:

DCFM 10.4.X → Network Advisor 11.1.X → Network Advisor 12.0.X → Network Advisor 12.2.X → Network Advisor 12.3.X or Network Advisor 12.4.X

Note: Enterprise and Professional Plus editions are not supported on 32-bit servers. If you are upgrading from an earlier product version running on a 32-bit server and plan to migrate to Enterprise and Professional Plus editions, you will require a 64-bit server. For more information, see “Pre-migration requirements when migrating from one server to another” in the *Network Advisor Installation and Migration Guide 12.4.2*. The link is shown at the end of this section.

For more information about IBM Network Advisor installation and migration, see the *Network Advisor SAN Installation and Migration Guide 12.4.2*, which can be found on the following link:

http://ibm.brocadeassist.com/public/IBMNA_Release

4.3 Downloading the software

Complete the following steps to download the software and documentation from the Brocade IBM Assist website:

1. Go to the Brocade IBM Assist website.
http://ibm.brocadeassist.com/public/IBMNA_Release
2. Select the highest version number for the latest generally available code. This book shows version 12.4.2.
 - To download the IBM Network Advisor installation file, click **Brocade Network Advisor 12.4.x** and then **Brocade Network Advisor 12.4.2 (IBM) GA**.
 - To download the documentation, click **Brocade Network Advisor 12.4.2 Manuals** and then select the manual that you want to download.
3. Select one of the following links to download the software, as shown in Figure 4-1 on page 55:
 - Network Advisor 12.4.2 (IBM) GA for Windows
 - Network Advisor 12.4.2 (IBM) GA for Linux

Product Downloads

Download by: All Software Products or Search: Enter product or file name Search

Product Name	Release Date
▶ Brocade SDN Controller	
▶ Brocade DCFM Enterprise	
▶ Brocade API Toolkit	
▼ Brocade Network Advisor	
▼ Brocade Network Advisor 12.4.x	08-18-2015
▼ Brocade Network Advisor 12.4.2 (IBM) GA	08-18-2015
Network Advisor 12.4.2 Release Notes v1.2 (pdf, 1.04 MB)	
Network Advisor 12.4.2 md5 Checksum (md5, 0.72 KB)	
Network Advisor 12.4.2 for IBM Windows (zip, 1.2 GB)	
Network Advisor 12.4.2 for IBM Linux (gz, 2.26 GB)	
▶ Brocade Network Advisor 12.4.1 (IBM) GA	06-12-2015
▶ Data Center Fabric Manager (DCFM)	08-18-2015
▶ EFCM	
▶ Services Director	
▶ SteelApp	
▶ vRouter 5400	
▶ vRouter 5600	

Tip: Make a selection using the drop down menu.


Figure 4-1 IBM Network Advisor product download selection

- Complete the Verification Email window information and click **Submit**, as shown in Figure 4-2.

Verification Email

Email *

Retype Email *

 [Refresh](#)

Type the word you see in the above image. *

[Cancel](#)

Figure 4-2 Verification Email window

An email will be sent to the email address entered. You must follow the email instructions to continue with the next step.

- Read and complete the Export Compliance form, select **I Agree**, and click **Submit**.

6. Read the Brocade End User License Agreement and click **I Accept**.
7. Click **Save** in the File Download window.
8. Browse to the location where you want to save the software and click **Save**.

4.4 Pre-installation requirements

Before you install IBM Network Advisor, ensure that you meet the following requirements:

- ▶ For specific system requirements, see 4.1, “Planning for server and client system requirements” on page 50.
- ▶ To avoid errors, close all instances of the application before you begin the installation or uninstallation procedure.
- ▶ For UNIX systems, if you still receive error message after closing the application, enter the following commands:
 - a. `#ps -ef | grep -i ""` to list the process ID
 - b. `#kill -9 "Process_ID"`, where Process_ID is any management application process

4.4.1 Additional pre-installation requirements for UNIX systems

Ensure that you meet the following requirements for a UNIX system:

- ▶ Ensure that an X Server is available for display and is configured to permit X Client applications to display from the host on which they are installing the IBM Network Advisor server. Typically, this simply requires that the system console is present and running with a logged-in user on the X Server-based desktop session, such as KDE or GNOME.
- ▶ Ensure that the DISPLAY environment variable is correctly defined in the shell with a valid value. For example, to display to the local console, run **export DISPLAY=:0.0**, or to display to a remote system that has an X Server running, run **export DISPLAY=Remote_IP_address:0.0**. You might also need to configure your firewall because it might block the display to the X Server, which listens by default on TCP port 6000 on the remote host.
- ▶ To display to a remote system, you must permit the remote display of the X Server by running **xhost+IP**, where IP is the IP address of the IBM Network Advisor server host from the X-based desktop of the remote system.
- ▶ Make sure that you test the DISPLAY definition by running **xterm** from the same shell from which you run **install.bin**. A new X terminal window to the destination X Server display should open.
- ▶ For Linux OS with the SELinux security policy enabled, ensure that you complete the following steps:
 - a. Disable the SELinux security policy by running **setenforce 0**.
 - b. Install the application as described in 4.6, “IBM Network Advisor Version 12.4.2 installation” on page 58.
 - c. Enable the SELinux security policy by running **setenforce 1**.

4.4.2 Mapping a loopback address to the local host

To map the loopback address to the local host, complete the following steps:

1. Open the host file:
 - For Windows, the hosts file is in the `WINDOWS\system32\drivers\etc` directory.
 - For Linux, the host file is in the `/etc` directory.
2. Add the following entries:
 - For an IPV4 machine:
`127.0.0.1 localhost`
 - For an IPV6 machine:
`127.0.0.1 localhost`
`::1 localhost`
3. Save and close the file.

4.5 Syslog troubleshooting

If the default syslog port number is in use, you will not receive any syslog messages from the device. Use one of the following procedures (depending on your operating system) to determine which process is running on the syslog port and stop that process.

4.5.1 Finding the process

To find the process, complete the following steps:

1. Open a command window.
2. Choose one of the following options:
 - On Linux systems, enter `netstat -nap | grep 514` and press **Enter**.
 - The process running on port 514 is displayed.
 - Example output: `UDP 0 0 ::ffff:127:0:0:1:514 :::* 27397`
 - On Windows systems, enter `netstat -anb | find /i "514"` and press **Enter**.
 - The process running on port 514 is displayed.
 - Example output: `UDP 127:0:0:1:514 *:* 3328`

4.5.2 Stopping the process

Choose one of the following options:

- ▶ On Linux systems, enter `Kill -9 "<PID>"` and press **Enter**. For example, `kill -9 "27397"`.
- ▶ On Windows systems, enter `taskkill /F /PID "<PID>"` and press **Enter**. For example, `taskkill /F /PID "3328"`. You can also run the following procedure instead:
 - a. Press **Ctrl+Shift+Esc** to open the Windows Task Manager.
 - b. Click the Processes tab.
 - c. Click the PID column header to sort the processes by PID.
 - d. Select the process that you want to stop and click **End Process**.

4.6 IBM Network Advisor Version 12.4.2 installation

Before you install the application, ensure that your system meets the minimum pre-installation requirements that are described in 4.4, “Pre-installation requirements” on page 56. If you are migrating data (upgrading), see 4.7, “Upgrading to IBM Network Advisor V12.4.2 from an existing IBM Network Advisor installation” on page 75.

This section describes how to perform a new IBM Network Advisor installation on both Windows and UNIX platforms. To do so, you need these privileges:

- ▶ On Windows system, you must be an Administrator with read and write privileges.
- ▶ On UNIX systems, you must be the root user.

To install IBM Network Advisor, complete the following steps:

1. Choose one of the following options:
 - For a Windows system, navigate to the Download_Location\Application_Name\Windows, right-click `install.exe`, and select **Run as administrator**.
 - For a UNIX system, complete the following steps
 - i. On the management application server, go to Download_Location/Application_Name/UNIX_Platform/bin.
 - ii. Execute either of the following commands:
`./install.bin` or `sh install.bin`

Note: On a Linux system, if you double-click the `install.bin` file, select **run**. Do *not* select **run in Terminal**.

2. Figure 4-3 shows the Introduction window for the installation. Click **Next** to proceed or **Cancel** to exit the upgrade.

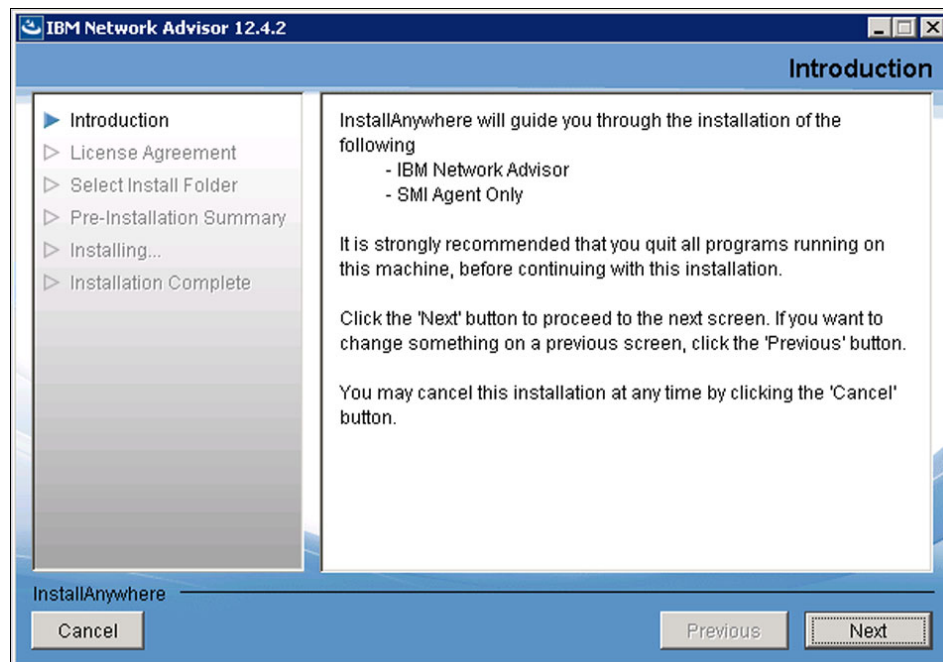


Figure 4-3 Introduction window for installation

3. A window with the license agreement opens. Accept the IBM Network Advisor license to proceed. After you accept the license agreement, you are prompted for the installation location. *Do not* install to the root directory (C:\Windows or / (UNIX). Figure 4-4 shows the options to select the installation location.

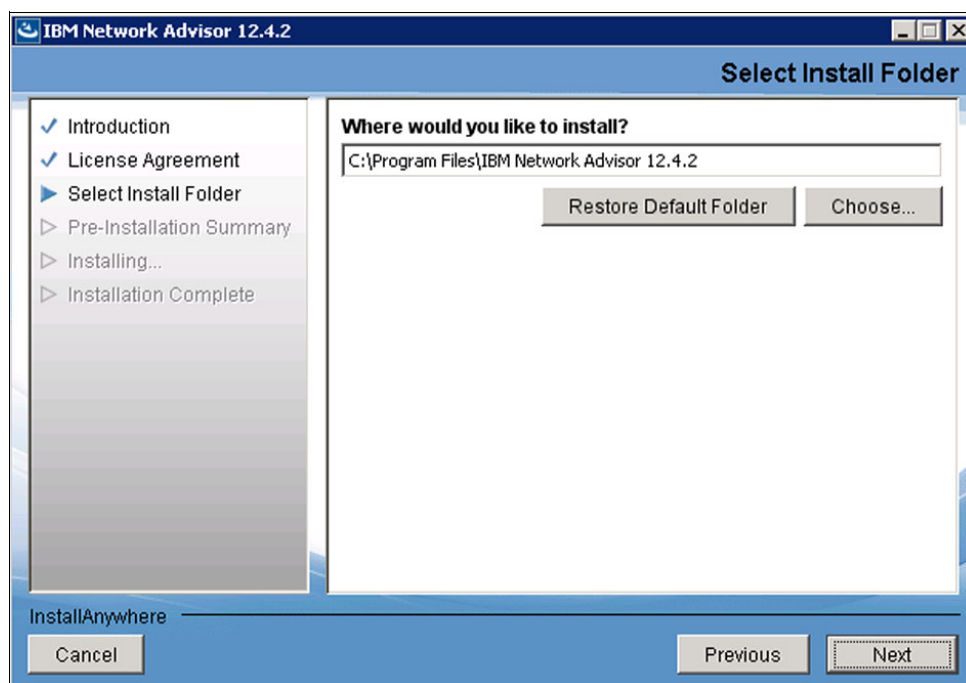


Figure 4-4 Installation folder options

4. After you select the target location, the Pre-Installation Summary window opens, as shown in Figure 4-5. This window describes the product and the target location.

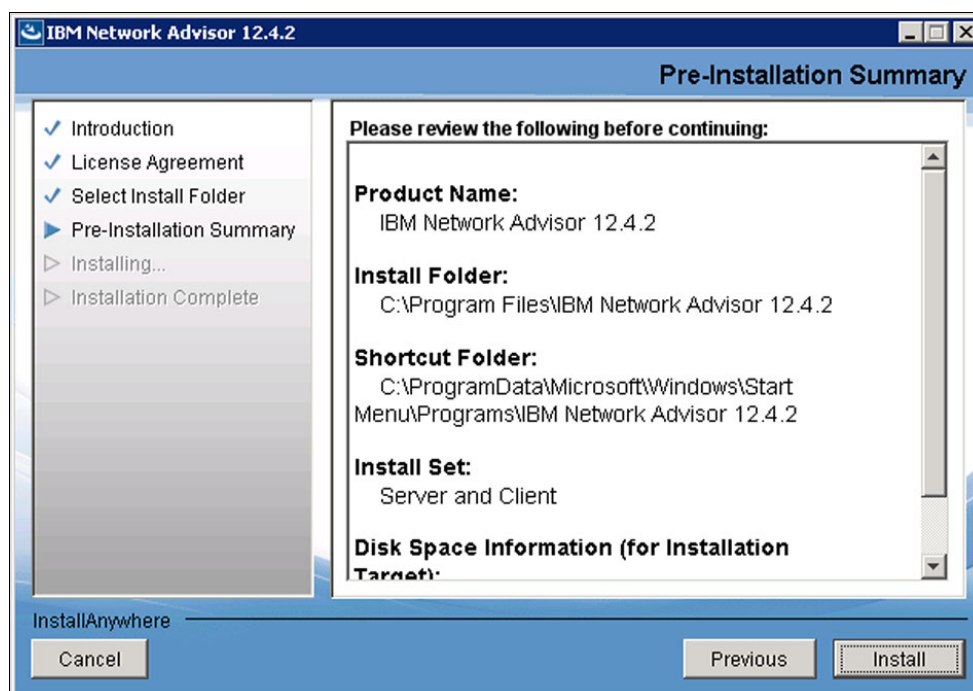


Figure 4-5 Pre-Installation Summary window

5. After you carefully review and agree with the pre-installation summary, click **Install** to proceed with installation. Figure 4-6 shows the start of the installation process.

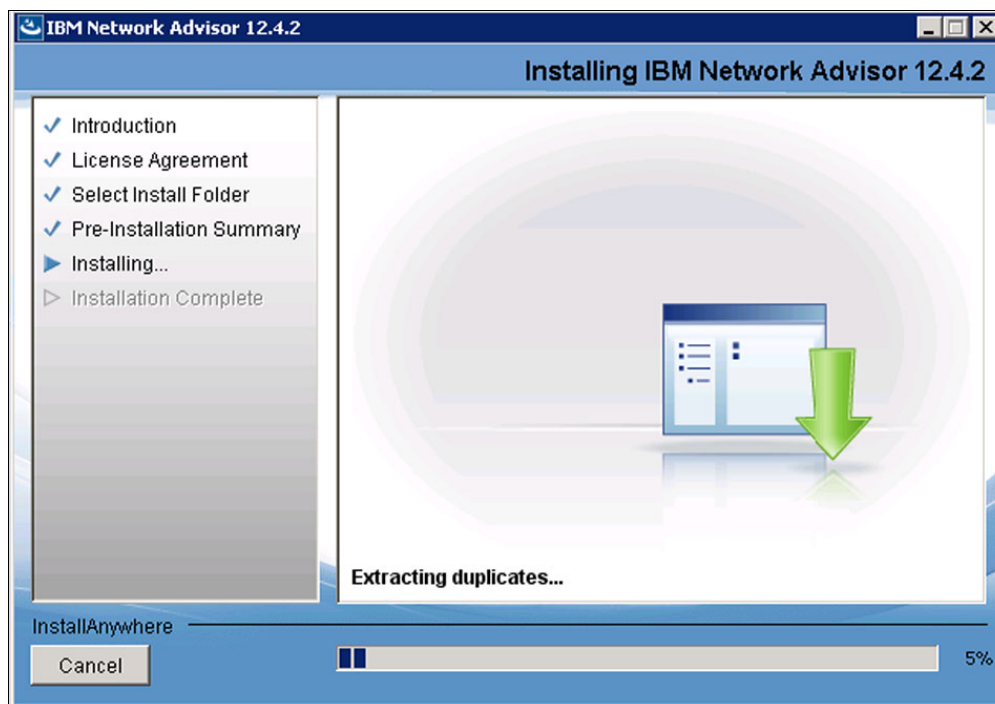


Figure 4-6 Installing IBM Network Advisor

6. After the completion of the installation, as shown in Figure 4-7, the Installation Complete window opens. Ensure that the **Launch IBM Network Advisor Configuration** check box is selected to proceed with the configuration (it is selected by default). Click **Done**.

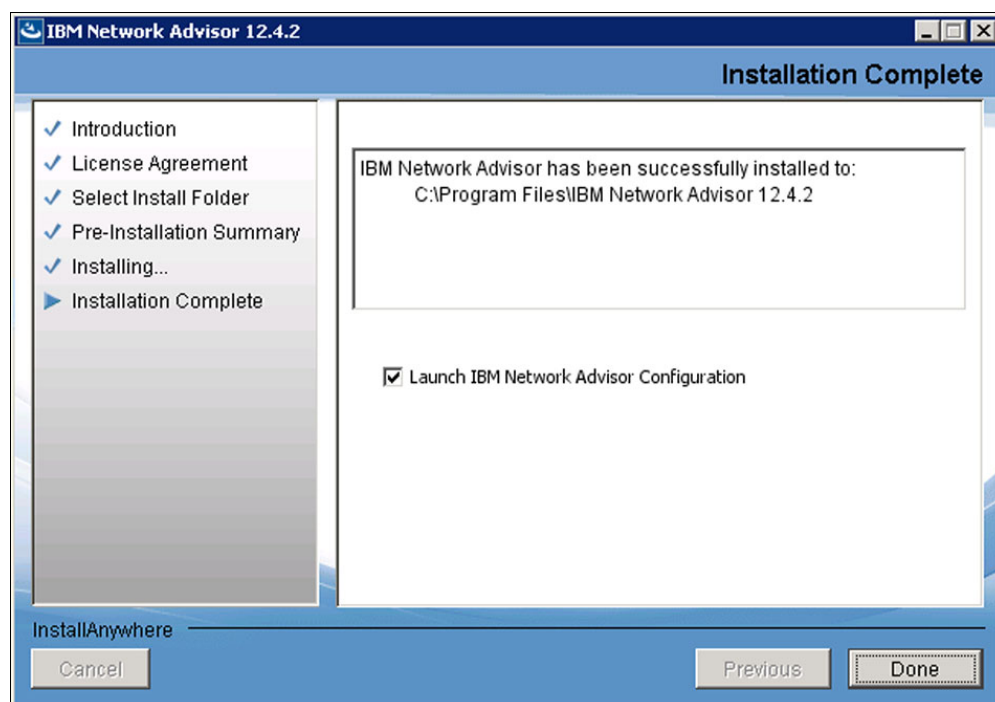


Figure 4-7 Installation Complete window

If the local host is not mapped to the loopback address, an error message is displayed. In this case, you must map the loopback address to the local host. To learn how to configure the loopback address, see 4.4.2, “Mapping a loopback address to the local host” on page 57.

If the Launch IBM Network Advisor Configuration check box is cleared, as shown in Figure 4-8, a window opens that prompts you to select the box.




Figure 4-8 Alert window for configuration check box

7. After you select the check box and click **Done**, the configuration wizard window opens listing the configuration activities like migrating data and settings, choosing the installation type, license, and FTP server, and so on (Figure 4-9). Click **Next**.

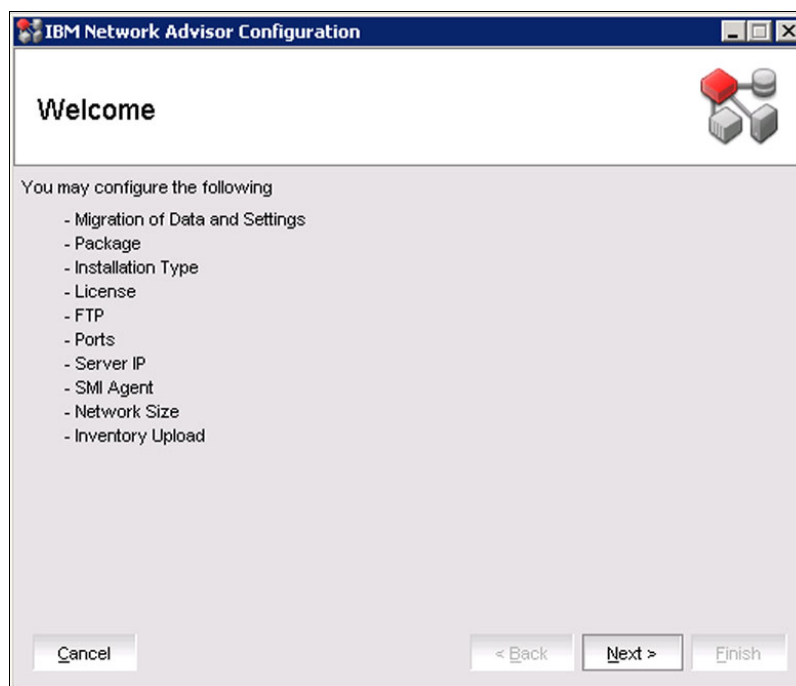


Figure 4-9 Configuration wizard main window

8. The Copy Data and Settings from previous releases window opens and prompts you for a copy of the data and settings from your previous installation. Because this is a new installation, select **No, don't copy any data and settings** (Figure 4-10). Click **Next**.

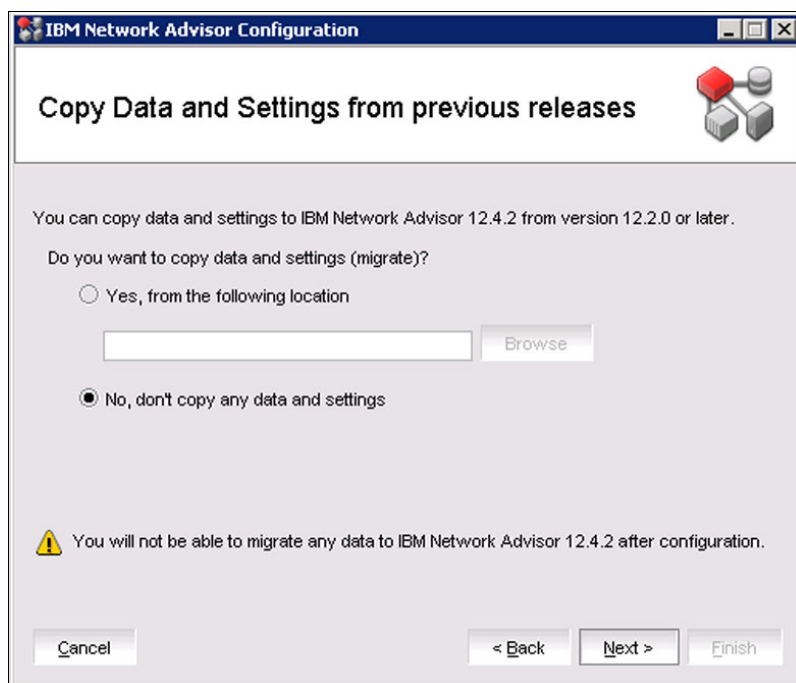


Figure 4-10 Copy data and settings from previous releases window

9. The Package window opens and prompts you to choose a package. IBM Network Advisor clients are not available in SMI Agent, so you must select **SAN with SMI Agent**, as shown in Figure 4-11. Click **Next**.

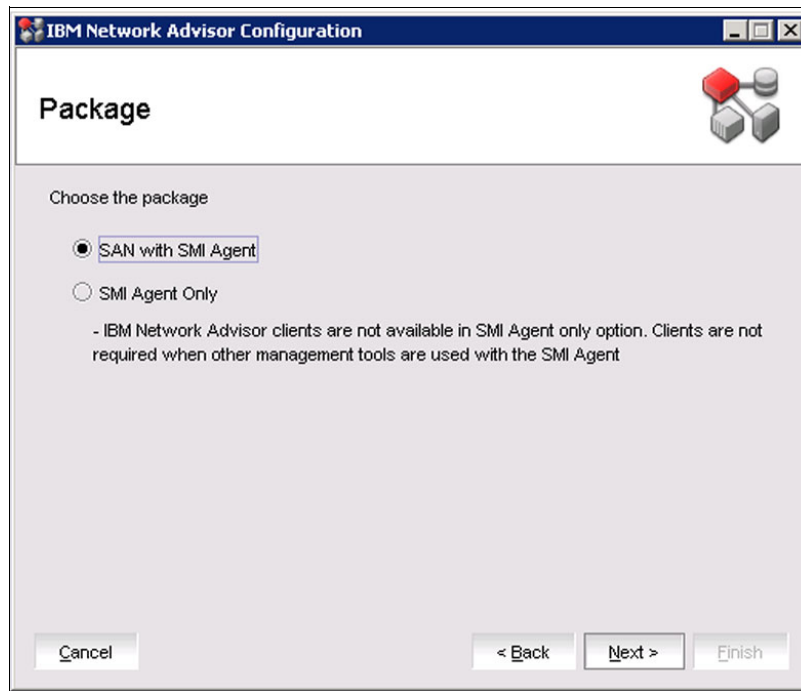


Figure 4-11 Package selection window

10. The Installation Type window opens and prompts you to choose an installation type, as shown in Figure 4-12.

Important: Obtain and store the license in a known secure location before you proceed with the upgrade.

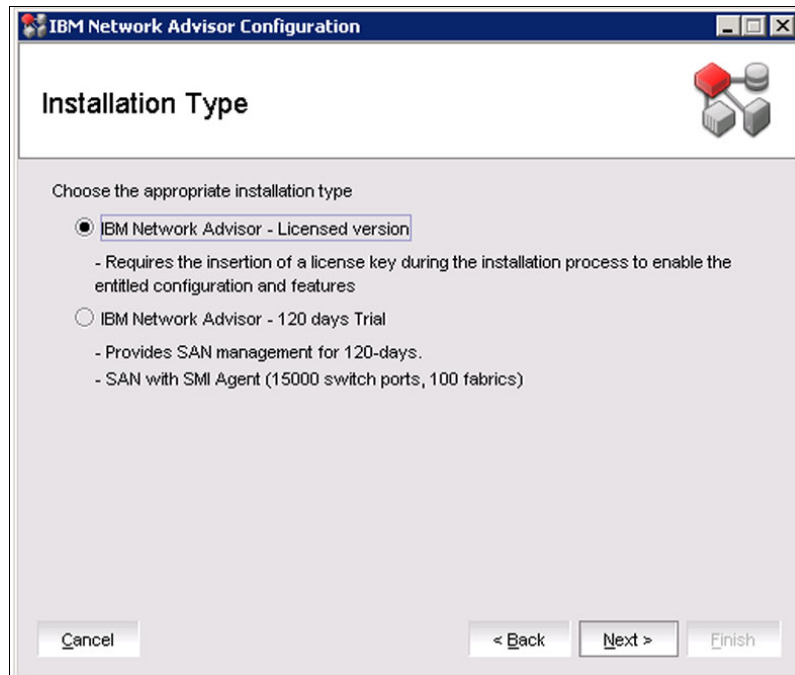


Figure 4-12 Choosing the installation type

There are two options: Licensed version and 120 days trial. Select **IBM Network Advisor - Licensed version** and click **Next**.

11. The Server License window opens and prompts you to enter the license details, as shown in Figure 4-13. Input the serial number and license key by clicking **Browse** and navigate to the location of the file that contains the information. Click **Next**.



The screenshot shows a window titled "IBM Network Advisor Configuration" with a sub-header "Server License". Below the sub-header is a small icon of three cubes. The main area contains the instruction "Enter the License Key or browse and select the license file." followed by two input fields: "Serial #" and "License Key". A "Browse" button is positioned below the "License Key" field. At the bottom of the window are four buttons: "Cancel", "< Back", "Next >", and "Finish".

Figure 4-13 Providing license details

12. The FTP / SCP / SFTP Server window opens and prompts you to configure the FTP server, as shown in Figure 4-14.

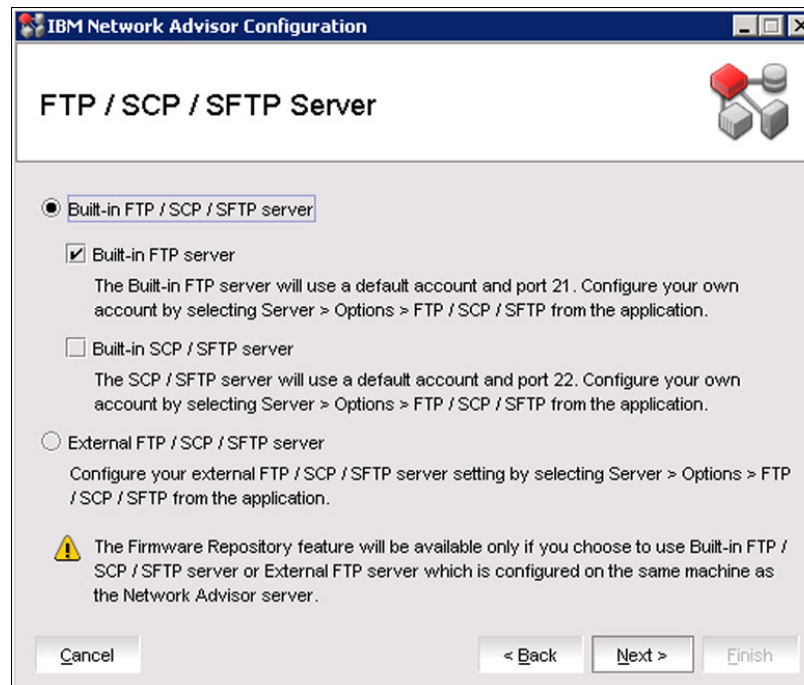


Figure 4-14 Setting up the servers and protocols for remote file transfers

This window shows the options of Built-in FTP/SCP/SFTP server and External FTP/SCP/SFTP server. The server where IBM Network Advisor is installed can also act as an FTP server, so it is a preferred practice to choose the Built-in FTP/SCP/SFTP option.

Note: If you choose External FTP/SCP/SFTP, ensure that the server where the IBM Network Advisor is installed is also configured as an FTP server because if you do not do so, the Firmware repository feature will not be available.

Select the **Built-in FTP/SCP/SFTP** option and click **Next**.

13. The Database Administrator Password (dcmadmin) window opens and prompts you to provide the password for the database, as shown in Figure 4-15.



IBM Network Advisor Configuration

Database Administrator Password (dcmadmin)

Choose the database password option.

☒ Default password

☐ New password

Password

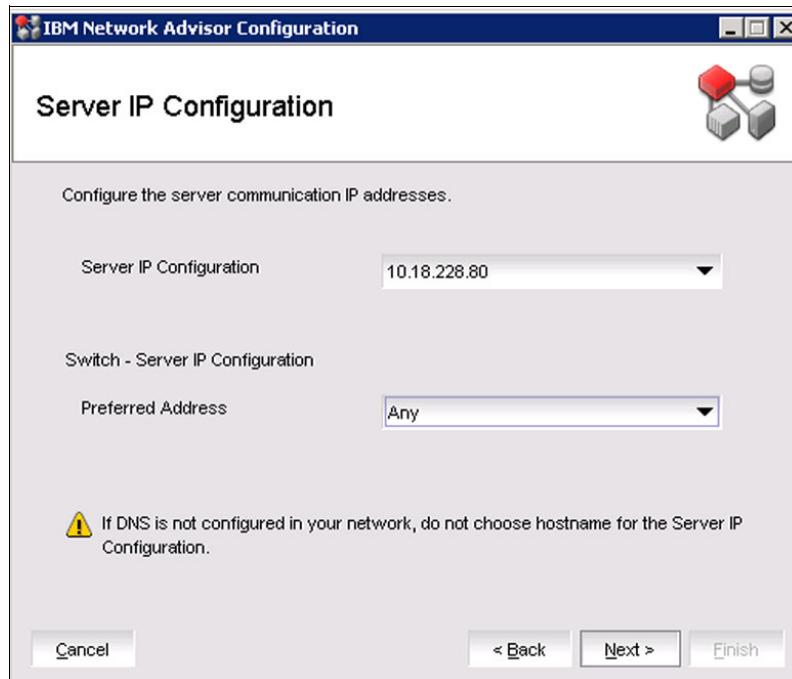
Confirm Password

 Database password can be changed later using Server Management Console.

Figure 4-15 Database Administrator password settings

You are presented with two options: **Default password** and **New password**. If you know the new password, select that option and provide the password. If you are not sure about the new password, select the **Default password** option and proceed. You can change the Database password later by using the Server Management Console. Click **Next**.

14. The Server IP Configuration window opens and prompts you to enter the server IP configuration details, as shown in Figure 4-16. In the Server IP Configuration drop-down box, you can choose a server name only if DNS is configured in the environment. If DNS is not configured, choose the IP address. Click **Next**.



The screenshot shows a window titled "IBM Network Advisor Configuration" with a sub-header "Server IP Configuration". Below the sub-header is a network icon. The main area contains the instruction "Configure the server communication IP addresses." followed by two dropdown menus. The first dropdown, labeled "Server IP Configuration", shows the IP address "10.18.228.80". The second dropdown, labeled "Switch - Server IP Configuration", shows the option "Any". Below these is a warning icon and text: "If DNS is not configured in your network, do not choose hostname for the Server IP Configuration." At the bottom are four buttons: "Cancel", "< Back", "Next >", and "Finish".

Figure 4-16 Server IP configuration details

If you select **Any** on the Preferred Address configuration, the server will listen to all the network requests through any of the IP addresses configured on the server.

15. The Server Configuration window opens and prompts you to provide the port details, as shown in Figure 4-17.

IBM Network Advisor Configuration

Resource Configuration and Validation

IBM Network Advisor requires Web Server, Database, TFTP, Syslog and SNMP port numbers, as well as 11 consecutive port numbers from a Starting port #. On enabling HTTP redirection, port # 80 is used to redirect the HTTP requests to HTTPS. Minimum system requirements will be validated.

Web Server Port # (HTTPS)

Redirect HTTP Requests to HTTPS ☒

Database Port #

Starting Port #

Syslog Port #

SNMP Port #

TFTP Port #

[Change this configuration by selecting Server > Options > Server Port from the application.](#)

Figure 4-17 Server's ports details

Complete the following steps:

- Enter a port number in the Web Server Port# (HTTPS) field (the default is 443):
 - Enable HTTP redirection to HTTPS by selecting the **Redirect HTTP Requests to HTTPS** check box.
 - When you enable HTTP redirection, the server uses port 80 to redirect HTTP requests to HTTPS. You can configure the server port settings by using the Options dialog box in the Server Port pane.
- Enter a port number in the **Database Port#** field (Default is 5432). Do not use a port number below 1024.
- Enter a port number in the **Starting Port Number** field (the default is 24600):
 - For Professional software, the server requires 15 consecutive free ports beginning with the starting port number.
 - For Trial and Licensed software, the server requires 18 consecutive free ports beginning with the starting port number.
- Enter a port number in the **Syslog Port Number** field (the default is 514). If the default Syslog port is already in use, you do not receive any syslog messages from the device. To find and stop the process that is running on the default Syslog port number, see 4.5, "Syslog troubleshooting" on page 57.
- Enter a port number in the **SNMP Port Number** field (the default is 162).

Click **Next** to continue.

If you enter a port number that is already in use, a warning displays next to the associated port number field. Edit that port number and click **Next**.

The IBM Network Advisor installation process validates your server resources against the prerequisites. A notification window like the one shown in Figure 4-18 is displayed if these validations fail. See 4.1.3, “IBM Network Advisor Server and concurrent client connections” on page 53 to ensure all the IBM Network Advisor installation prerequisites are met.

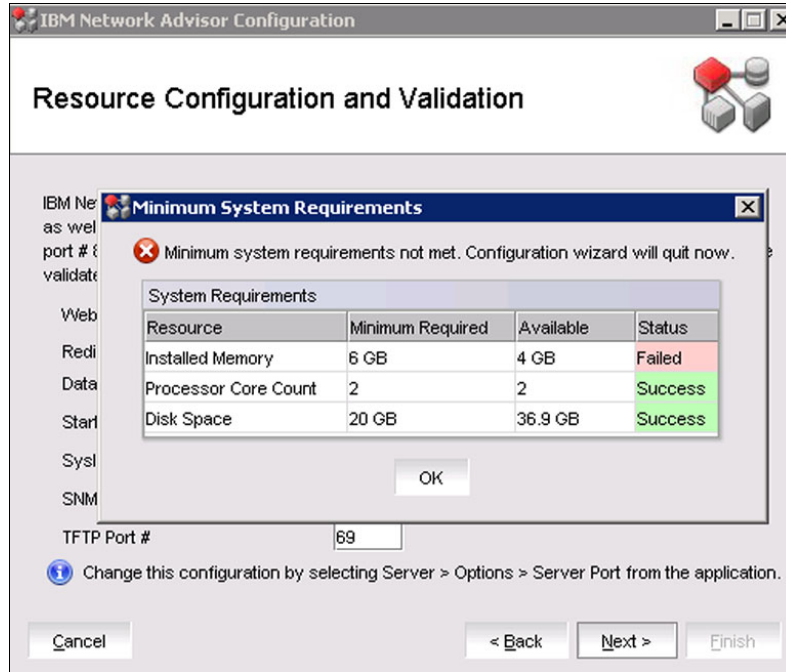


Figure 4-18 Installation pre-requisites not met notification window

- The SMI Agent Configuration window opens and prompts you to configure the SMI Agent, as shown in Figure 4-19.

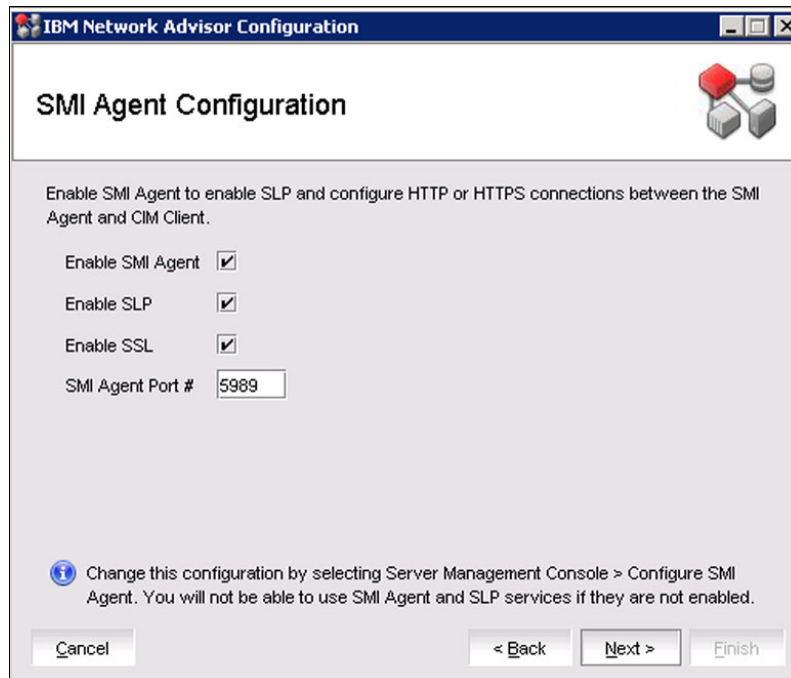


Figure 4-19 SMI Agent configuration window

Select **Enable SMI Agent**, then **Enable SSL** and enter 5989 as the port number (the default is 5988). Click **Next**.

17. The SAN Network Size window opens and prompts you to configure the SAN network, as shown in Figure 4-20. Choose the option that best suits your network size and click **Next**.

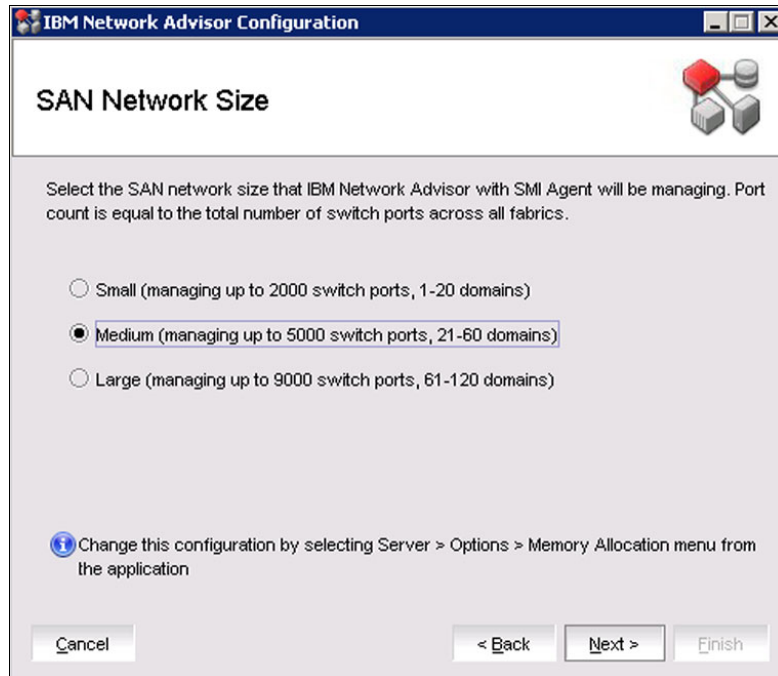
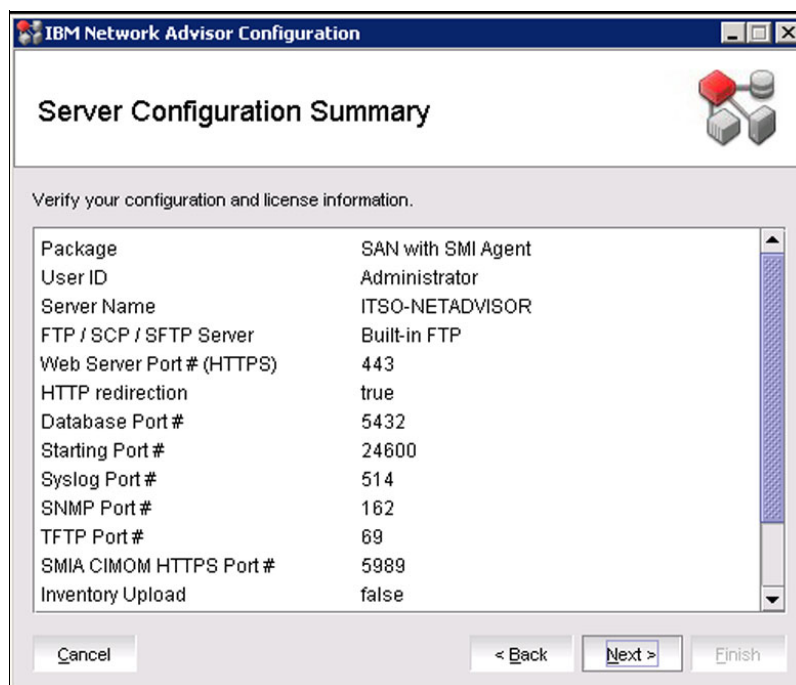


Figure 4-20 Network size selection window

Note: Port count is equal to the total number of switch ports to be managed across all fabrics.

18. The Server Configuration Summary window opens, as shown in Figure 4-21. Review the details in the window. If you are satisfied, click **Next**.



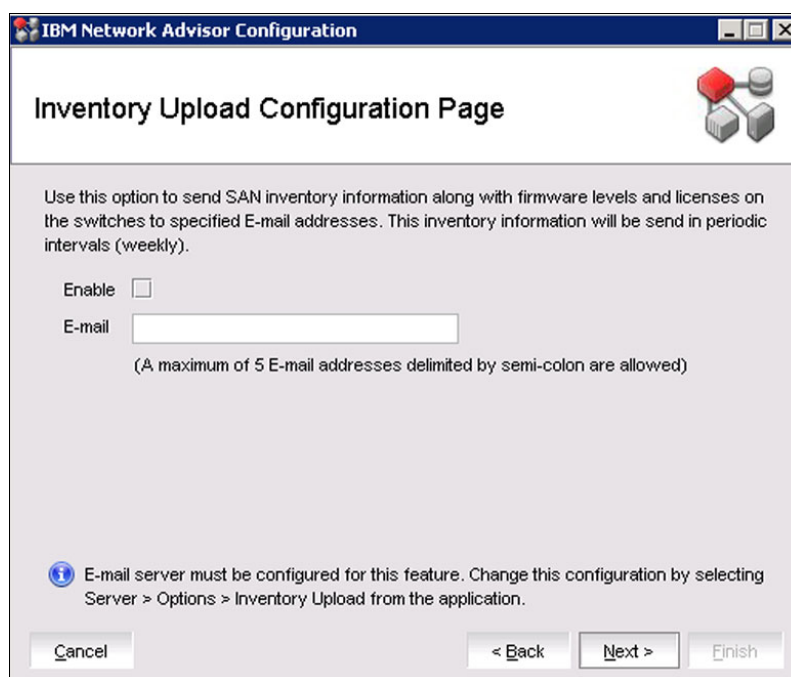
The window titled "IBM Network Advisor Configuration" displays the "Server Configuration Summary". It contains a table with the following configuration details:

Package	SAN with SMI Agent
User ID	Administrator
Server Name	ITSO-NETADVISOR
FTP / SCP / SFTP Server	Built-in FTP
Web Server Port # (HTTPS)	443
HTTP redirection	true
Database Port #	5432
Starting Port #	24600
Syslog Port #	514
SNMP Port #	162
TFTP Port #	69
SMIA CIMOM HTTPS Port #	5989
Inventory Upload	false

At the bottom, there are buttons for "Cancel", "< Back", "Next >", and "Finish".

Figure 4-21 Server configuration summary window

19. The Inventory Upload configuration window is displayed, as shown in Figure 4-22. For IBM Network Advisor 12.4.0 and later, a feature called Inventory Upload allows users to send the managed switches' inventory, their FOS code levels, and licenses to up to five email addresses.




The window titled "IBM Network Advisor Configuration" displays the "Inventory Upload Configuration Page". It includes the following configuration options:

Use this option to send SAN inventory information along with firmware levels and licenses on the switches to specified E-mail addresses. This inventory information will be send in periodic intervals (weekly).

Enable ☐

E-mail

(A maximum of 5 E-mail addresses delimited by semi-colon are allowed)

 E-mail server must be configured for this feature. Change this configuration by selecting Server > Options > Inventory Upload from the application.

At the bottom, there are buttons for "Cancel", "< Back", "Next >", and "Finish".

Figure 4-22 SAN Inventory upload configuration window

If you want to receive this information, select **Enable** and enter the list of email recipients delimited by a semi-colon (;). When you are done, click **Next**.

20. The Start Server window opens and prompts you to start the server and client, as shown in Figure 4-23. You can start the client after the installation by selecting the **Start Client** check box, or you can leave it clear to start the client at a later stage.

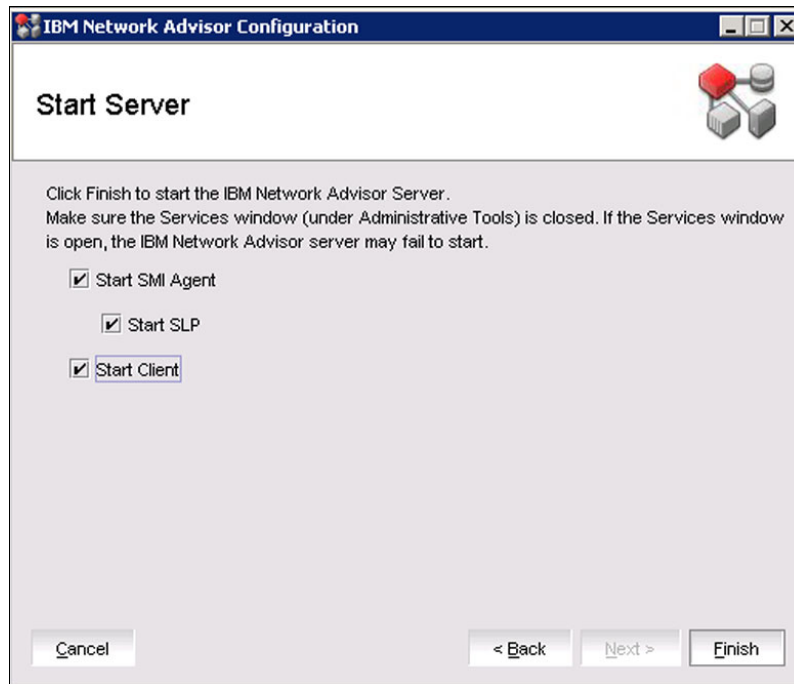


Figure 4-23 Starting the server and client

Ensure that the Service window (under Administrative Tools) is closed. If the Services window is open, IBM Network Advisor might fail to start.

Select **Finish**. Services will be started, as shown in Figure 4-24.

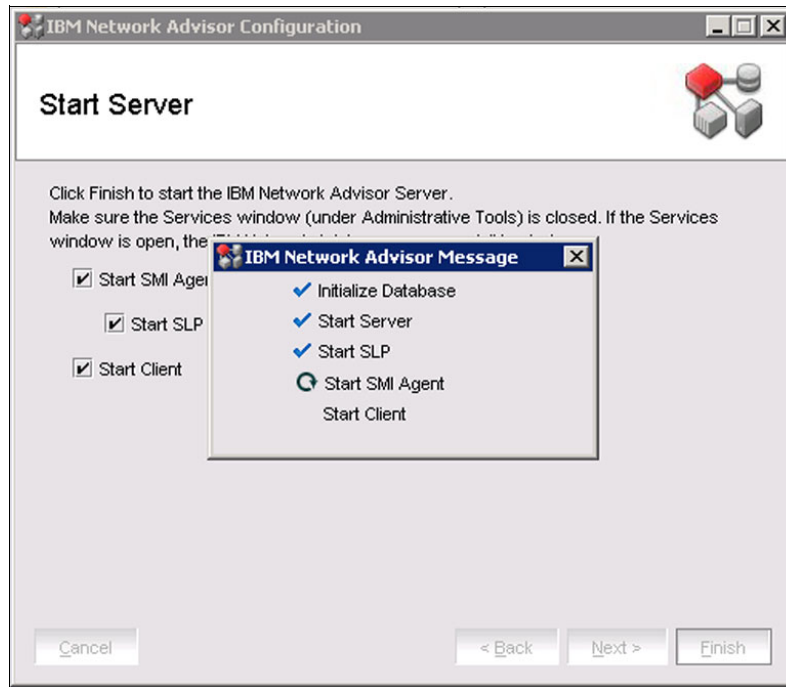


Figure 4-24 IBM Network Advisor servers start status window

21. A Security Alert window opens, as shown in Figure 4-25, which prompts you to permit the traffic by accepting the security settings.

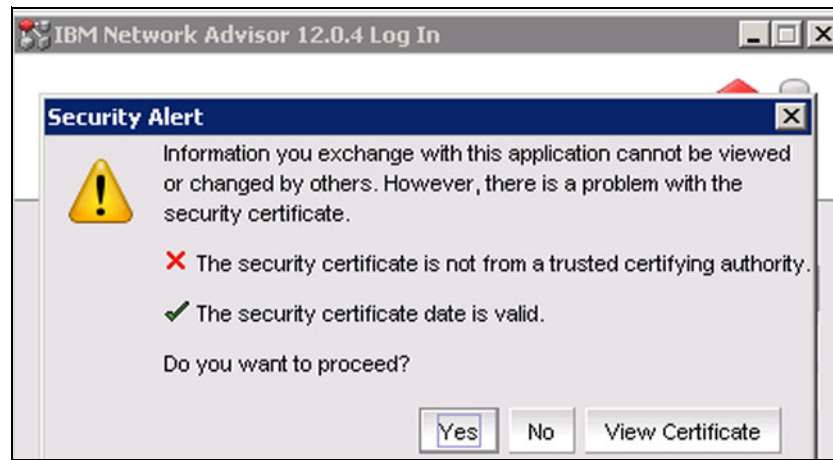


Figure 4-25 Security Alert

Select **Yes**.

22. A login window opens and prompts you to provide the login credentials, as shown in Figure 4-26. The default credentials are administrator/password. After you provide the credentials, click **Login**.

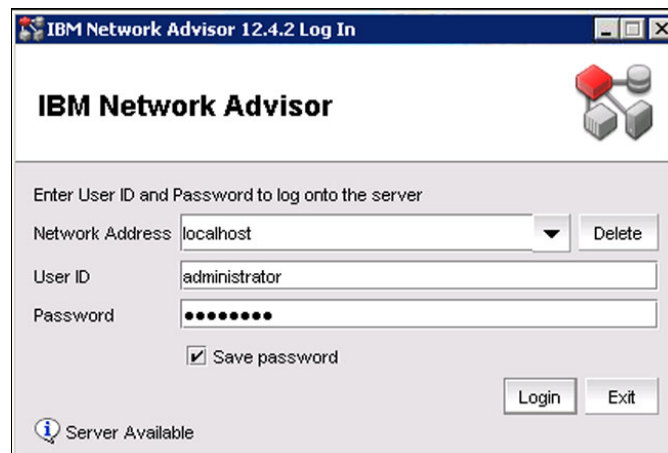


Figure 4-26 IBM Network Advisor login window

4.7 Upgrading to IBM Network Advisor V12.4.2 from an existing IBM Network Advisor installation

This section describes upgrading an existing IBM Network Advisor installation to Version 12.4.2. For the upgrade path reference, see 4.2, “IBM Network Advisor 12.4.2 upgrade path” on page 53. During the upgrade process, the old version is uninstalled and the new version with an upgraded database is installed.

Before you proceed with the upgrade, back up your configuration. To back up your configuration, go to installation location and copy the IBM Network Advisor 12.4.2 folder.

This section describes upgrading to IBM Network Advisor V12.4.2 on both Windows and UNIX platforms. To do so, you must have these privileges:

- ▶ On Windows systems, you must be an administrator with read and write privileges.
- ▶ On UNIX systems, you must be the root user.

To upgrade the new application version, complete the following steps:

1. Choose one of the following options:
 - For a Windows system, open the `Download_Location\Application_Name\Windows\install.exe` file.
 - For a UNIX system, complete the following steps:
 - i. On the management application server, go to the `Download_Location/Application_Name/UNIX_Platform/bin` directory.
 - ii. Run one of the following commands:
`./install.bin` or `sh install.bin`

Note: On a Linux system, if you double-click the `install.bin` file, select **RUN**. Do *not* select **RUN in Terminal**.

2. The Introduction window opens, as shown in Figure 4-3 on page 58. Click **Next** to proceed or **Cancel** to exit the upgrade.
3. A window with the license agreement opens. Accept the license agreement to proceed to the next step. A window opens and prompts you for the installation folder location, as shown in Figure 4-4 on page 59. Select the default location because it is the current IBM Network Advisor installed folder location. If you want to choose a different location, do so now. Click **Next**.
4. The Pre-Installation Summary window opens, as shown in Figure 4-6 on page 60. Select **Install** to proceed with the installation.
5. After the installation completes, the Installation Complete window opens, as shown in Figure 4-8 on page 61. Select **Launch IBM Network Advisor Configuration** and click **Done** to start the configuration.
6. The Welcome window opens and describes the migrate data and settings, choosing the installation type, license, FTP server, ports, server IP SMI agent, and network size, as shown in Figure 4-10 on page 62. Click **Next**.

7. The Copy Data and Settings from previous releases window opens and prompts you for a copy of the data and settings from your previous installation, as shown in Figure 4-27. You must provide the folder location where the previous version was installed. This window will also show the minimum IBM Network Advisor version from where you can migrate the data from. For upgrade path information, see 4.2, “IBM Network Advisor 12.4.2 upgrade path” on page 53. Remember that it is not possible to migrate the data from a previous version after the installation completes.

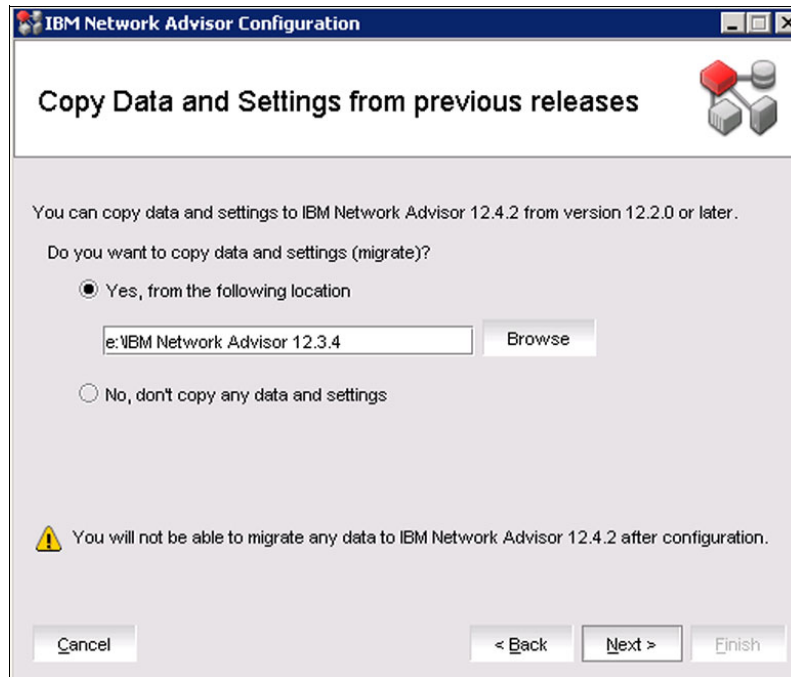


Figure 4-27 Copy Data and Settings from previous releases window

Click **Next**.

8. The Resource Validation and Data Migration window opens, as shown in Figure 4-28. In this stage, the IBM Network Advisor installation validates all of the server resources that are required for the product installation and starts the data migration process.

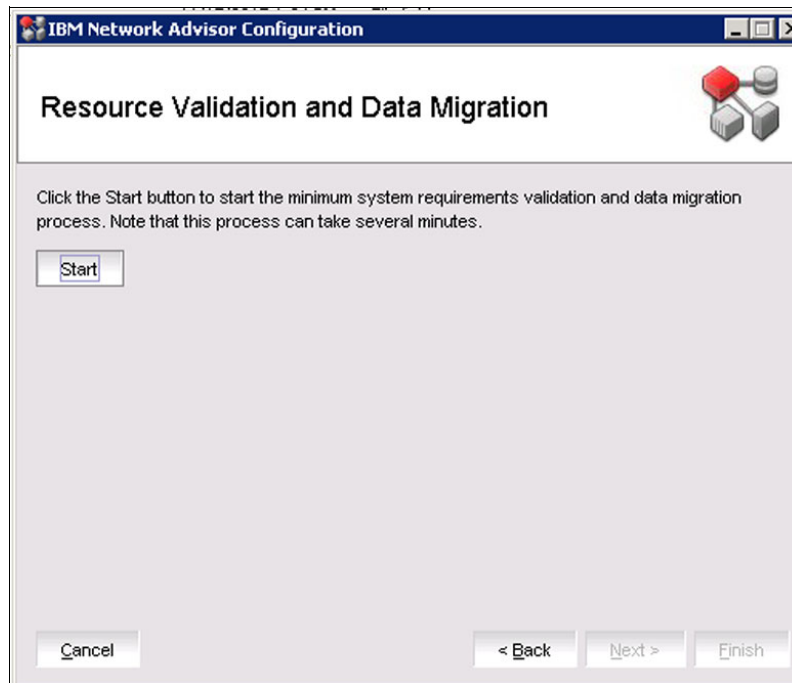


Figure 4-28 Resource Validation and Data Migration window

Click **Start**.

9. A message is displayed to confirm that the data has been successfully migrated, as shown in Figure 4-29.


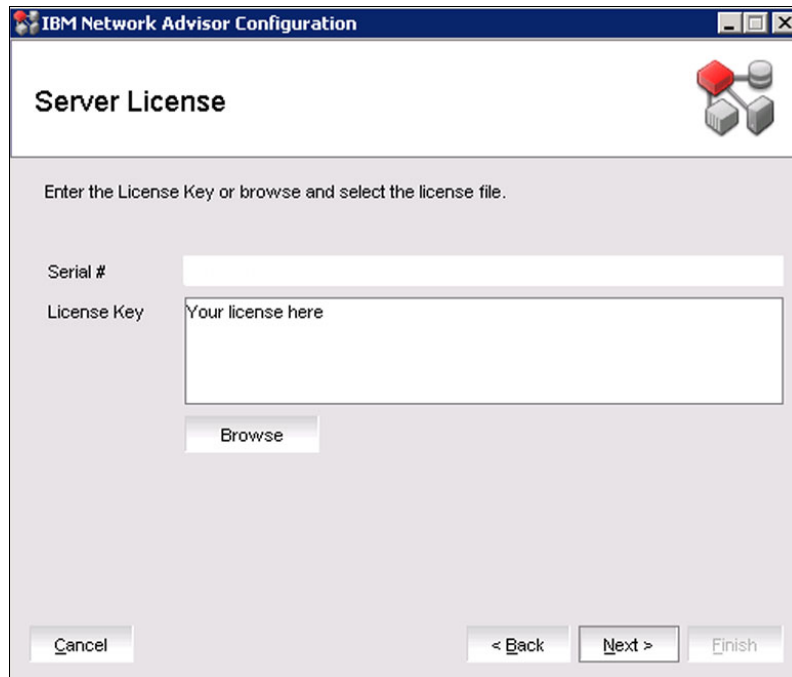


Figure 4-29 Data migration confirmation message

Click **Start**.

10. The next window asks for the product license to be entered, as shown in Figure 4-30. If you are upgrading an IBM Network Advisor instance that was already licensed, this window will autopopulate the license information for you. Otherwise, you can enter a new license if you are simultaneously upgrading and licensing the product.



The image shows a Windows-style dialog box titled "IBM Network Advisor Configuration". Inside the dialog, the title "Server License" is displayed at the top left, and a small icon of three server racks is at the top right. Below the title, there is a text prompt: "Enter the License Key or browse and select the license file." The main area contains two input fields: "Serial #" with an empty text box, and "License Key" with a text box containing the placeholder text "Your license here". Below the "License Key" field is a "Browse" button. At the bottom of the dialog, there are four buttons: "Cancel" on the left, and "< Back", "Next >", and "Finish" on the right.

Figure 4-30 IBM Network Advisor License window

11. The FTP / SCP / SFTP Server window opens and prompts you to configure the FTP server, as shown in Figure 4-31.

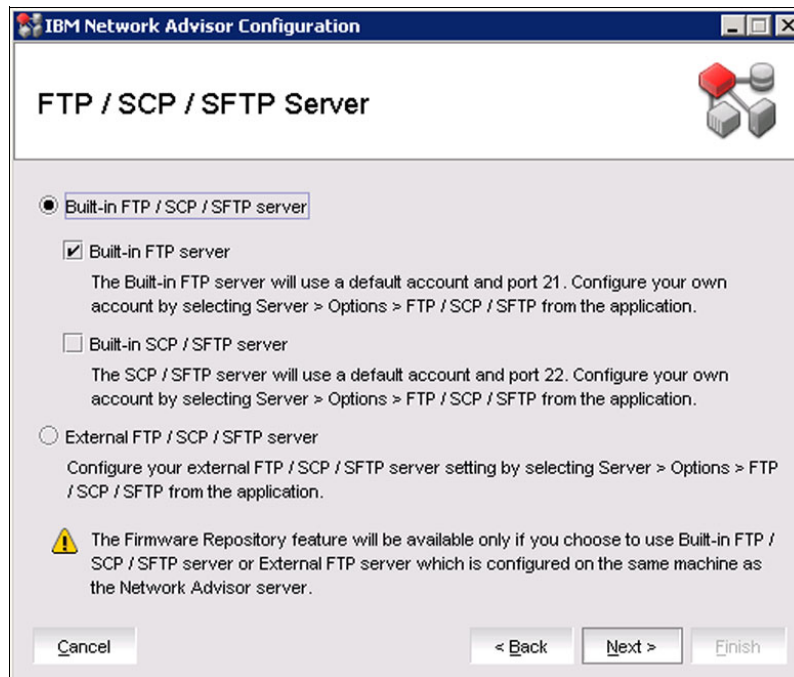


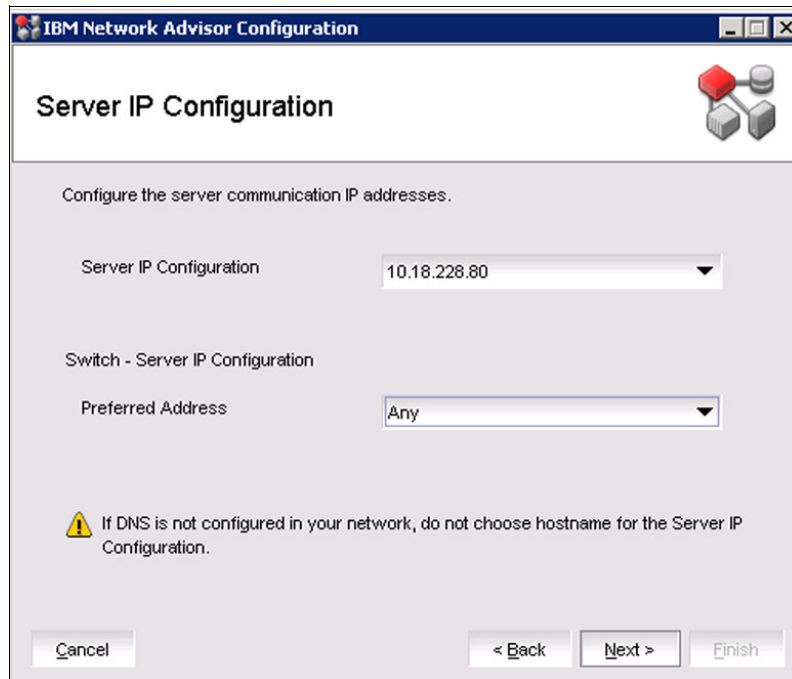
Figure 4-31 Setting up the servers and protocols for remote file transfers

This window shows the options of **Built-in FTP/SCP/SFTP server** and **External FTP/SCP/SFTP server**. The server where IBM Network Advisor is installed can also act as an FTP server, so it is a preferred practice to choose the **Built-in FTP/SCP/SFTP** option.

Note: If you choose **External FTP/SCP/SFTP**, ensure that the server where the IBM Network Advisor is installed is also configured as an FTP server because if you do not do so, the Firmware repository feature will not be available.

Select **Built-in FTP/SCP/SFTP** and click **Next**.

12. The Server IP Configuration window opens and prompts you to enter the server IP configuration details, as shown in Figure 4-32. In the **Server IP Configuration** drop-down box, you can choose a server name only if DNS is configured in the environment. If DNS is not configured, choose the IP address. Click **Next**.



The screenshot shows a window titled "IBM Network Advisor Configuration" with a sub-header "Server IP Configuration". Below the sub-header is a network icon. The main area contains the instruction "Configure the server communication IP addresses." followed by two dropdown menus. The first dropdown, labeled "Server IP Configuration", shows the value "10.18.228.80". The second dropdown, labeled "Switch - Server IP Configuration" and "Preferred Address", shows the value "Any". Below these is a warning icon and text: "If DNS is not configured in your network, do not choose hostname for the Server IP Configuration." At the bottom are four buttons: "Cancel", "< Back", "Next >", and "Finish".

Figure 4-32 Server IP configuration details

If you select **Any** on the Preferred Address configuration, the server will listen to all the network requests through any of the IP addresses configured on the server.

13. The Server Configuration window opens and prompts you to provide the port details, as shown in Figure 4-33.

IBM Network Advisor Configuration

Resource Configuration and Validation

IBM Network Advisor requires Web Server, Database, TFTP, Syslog and SNMP port numbers, as well as 11 consecutive port numbers from a Starting port #. On enabling HTTP redirection, port # 80 is used to redirect the HTTP requests to HTTPS. Minimum system requirements will be validated.

Web Server Port # (HTTPS)

Redirect HTTP Requests to HTTPS ☒

Database Port #

Starting Port #

Syslog Port #

SNMP Port #

TFTP Port #

[Change this configuration by selecting Server > Options > Server Port from the application.](#)

Figure 4-33 Server's ports details

Complete the following steps:

- Enter a port number in the **Web Server Port# (HTTPS)** field (the default is 443):
 - Enable HTTP redirection to HTTPS by selecting **Redirect HTTP Requests to HTTPS**.
 - When you enable HTTP redirection, the server uses port 80 to redirect HTTP requests to HTTPS. You can configure the server port settings by using the Options dialog box in the Server Port pane.
- Enter a port number in the **Database Port#** field (Default is 5432). Do not use a port number below 1024.
- Enter a port number in the **Starting Port Number** field (the default is 24600):
 - For Professional software, the server requires 15 consecutive free ports beginning with the starting port number.
 - For Trial and Licensed software, the server requires 18 consecutive free ports beginning with the starting port number.
- Enter a port number in the **Syslog Port Number** field (the default is 514). If the default Syslog port is already in use, you do not receive any syslog messages from the device. To find and stop the process that is running on the default Syslog port number, see 4.5, "Syslog troubleshooting" on page 57.
- Enter a port number in the **SNMP Port Number** field (the default is 162).

Click **Next**.

If you enter a port number that is already in use, a warning displays next to the associated port number field. Edit that port number and click **Next**.

14. The SAN Network Size window opens and prompts you to configure the SAN network, as shown in Figure 4-34. Choose the option that best suits your network size and click **Next**.

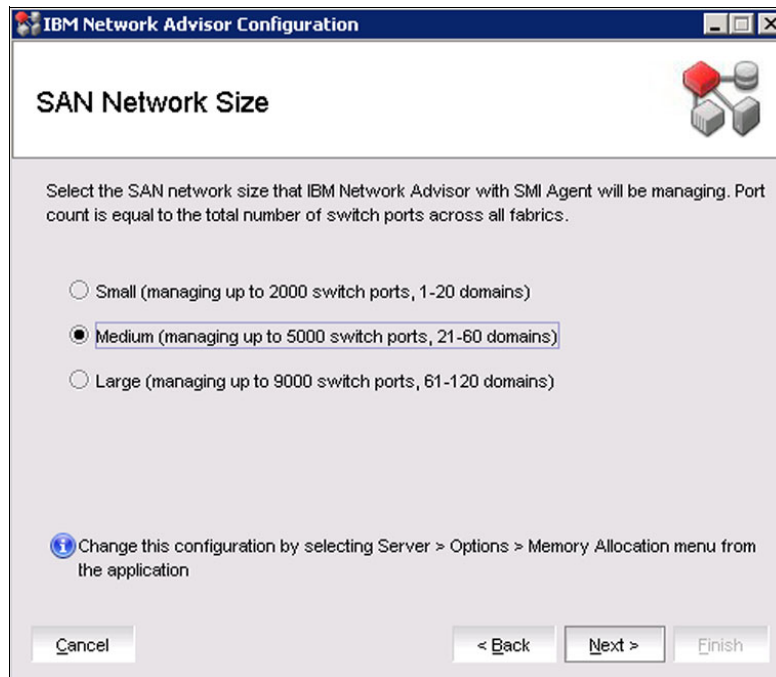
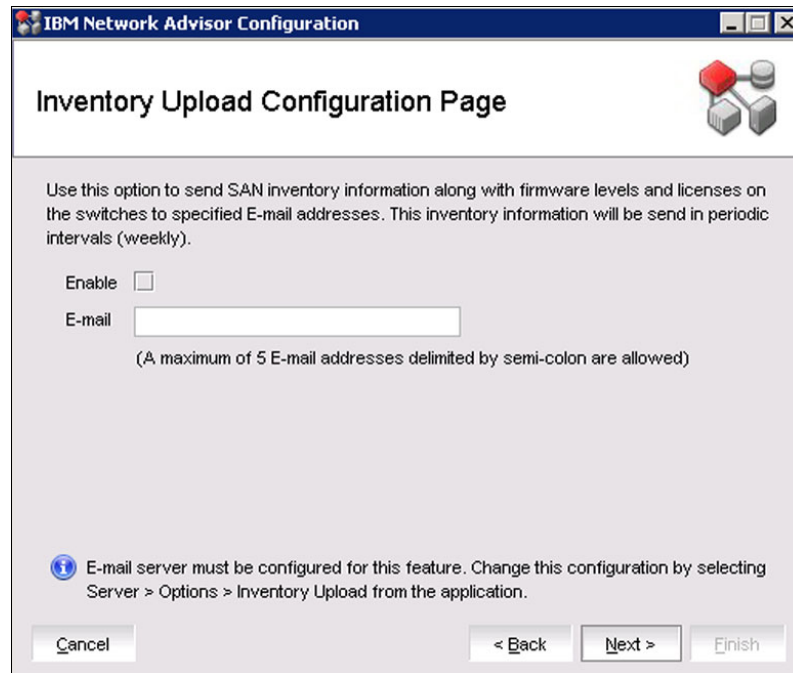


Figure 4-34 Network size selection window

Note: Port count is equal to the total number of switch ports to be managed across all fabrics.

15. The Inventory Upload configuration window appears, as shown in Figure 4-35. When you start IBM Network Advisor 12.4.0, a new feature called Inventory Upload allows users to upload the managed switches' inventory as well as their FOS code levels and licenses to up to five email addresses.



The screenshot shows a window titled "IBM Network Advisor Configuration" with a sub-header "Inventory Upload Configuration Page". The page contains a description: "Use this option to send SAN inventory information along with firmware levels and licenses on the switches to specified E-mail addresses. This inventory information will be send in periodic intervals (weekly)." Below this is an "Enable" checkbox, which is currently unchecked. To the right of the checkbox is an "E-mail" text input field. Below the input field is a note: "(A maximum of 5 E-mail addresses delimited by semi-colon are allowed)". At the bottom, there is an information icon followed by the text: "E-mail server must be configured for this feature. Change this configuration by selecting Server > Options > Inventory Upload from the application." The window has a standard Windows-style title bar with minimize, maximize, and close buttons. At the bottom of the window are four buttons: "Cancel", "< Back", "Next >", and "Finish".

Figure 4-35 SAN Inventory upload configuration window

If you want to receive this information, select the **Enable** check-box and enter the list of email recipients delimited by a semi-colon (;). When you are done, click **Next**.

16. The Server Configuration Summary window opens, as shown in Figure 4-36. Review the details in the window. If you are satisfied, click **Next**.

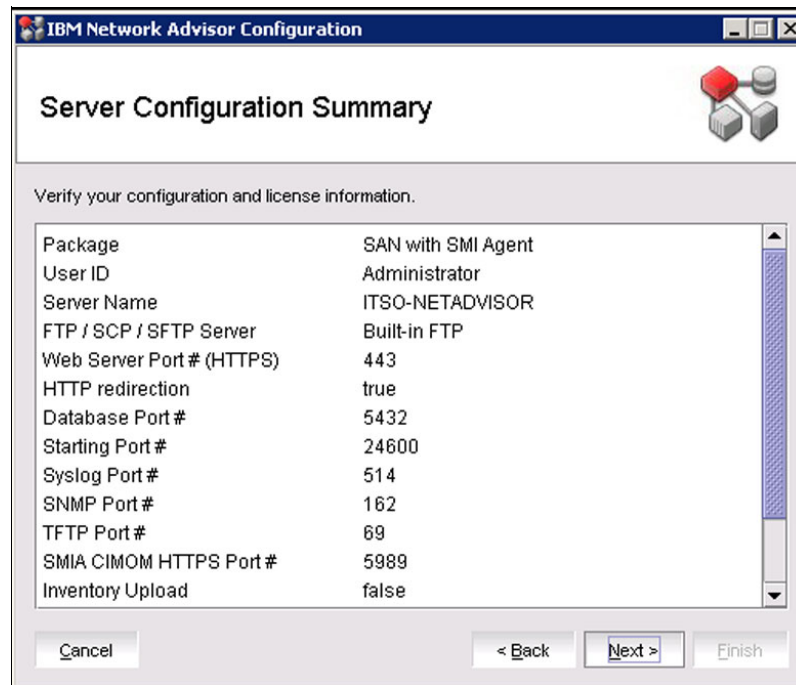


Figure 4-36 Server configuration summary window

17. The Start Server window opens and prompts you to start the server and client, as shown in Figure 4-37. You can start the client after the installation by selecting the **Start Client** check box, or you can leave it clear to start the client at a later stage.



Figure 4-37 Starting the server and client

Ensure that the Service window (under Administrative Tools) is closed. If the Services window is open, IBM Network Advisor might fail to start.

18. Select **Finish**.

The upgrade process finishes by migrating and initializing the database, uninstalling the older IBM Network Advisor version, and starting all of the IBM Network Advisor services, as shown in Figure 4-38.

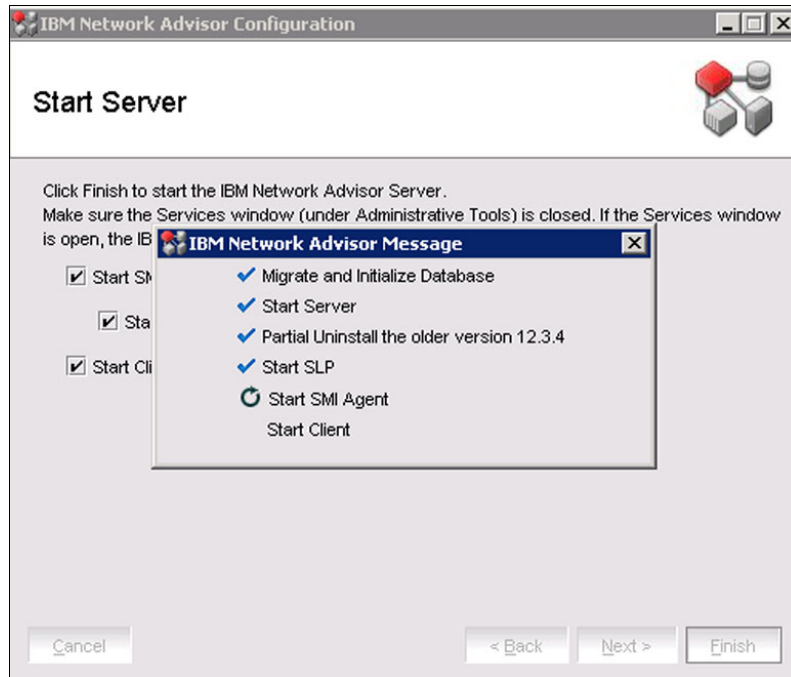


Figure 4-38 Completing the IBM Network Advisor upgrade installation

19. When the upgrade process completes, a login window opens and prompts you to provide the login credentials, as shown in Figure 4-39. After you provide the credentials, click **Login** to start using the upgraded IBM Network Advisor.



Figure 4-39 IBM Network Advisor login window

4.8 IBM Network Advisor web client

The IBM Network Advisor web client provides a high-level overview of your network and quick access to dashboard monitors and reports. However, it cannot be used for configuration and management of your fabrics.

Figure 4-40 shows the web client's main window with its various areas.

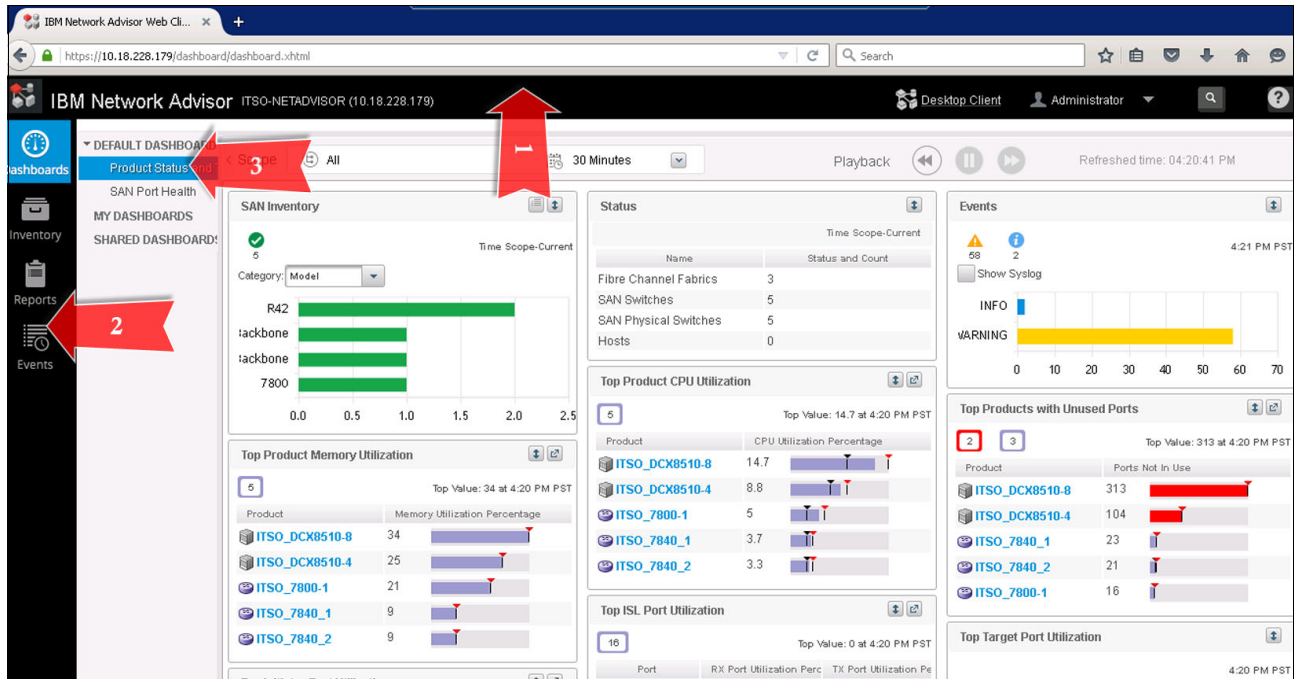


Figure 4-40 IBM Network Advisor web client main window

The numbers in Figure 4-40 correspond to these items:

1. Web client banner and toolbar: Displays the Management application server name and status as well as buttons to perform various functions.
2. Left pane: Contains the Expand navigation bar that provides a list of features you can access.
3. Right pane: Displays the detail for the feature that is selected in the left pane.

For more information about the IBM Network Advisor web client, see the “Web Client” section of the *Network Advisor 12.4.2 SAN User Manual* at the following link:

<http://www.brocade.com/content/html/en/user-guide/networkadvisor-1242-san-manual>

4.9 User, device discovery, and dashboard management

The following section describes the most common user account management activities, dashboard management, and fabric discovery and management.

4.9.1 User management

The user management application contains the information about the IBM Network Advisor users and their privileges, roles, and assigned Areas of Responsibilities (AORs).

Privileges define which features the users have access to. A role is a group of predefined privileges that is assigned to multiple users who needs access to the same menu options.

An AOR contains selected fabrics and devices, and when applied to users it defines which devices those users are able to manage.

Creating users

By default, the Administrator account is created when you are installing or upgrading IBM Network Advisor. The first time you log in, use the default administrator credentials administrator (for the user ID) and password (for the password). For security reasons, change the administrator default password. To create a user, click **Server** → **Users**.

As shown in Figure 4-41 a window opens with the Users, Policy, and Authentication Server Groups tabs. Click the Users tab and then click **Add** to create a user.

The screenshot displays the 'Users Management' application window. At the top, there are three tabs: 'Users', 'Policy', and 'Authentication Server Groups'. The 'Users' tab is active. Below the tabs, there are fields for 'Authentication-Primary' (set to 'Local Database'), 'Secondary' (set to 'None'), and 'Authorization' (set to 'Local Database').

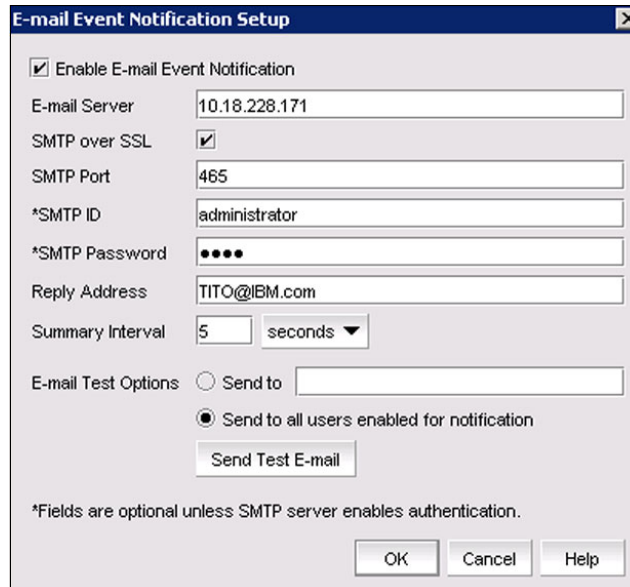
The main area is divided into three panes:

- Users:** A table with columns: User ID, Full Name, Roles, Area Of Responsi..., E-mail Notification, Account Enabled, Policy Violations, Account State, and Description. The first row shows 'Administrator' with roles 'SAN System Adm...', 'All Hosts', and 'All Fabri...'. Below the table are buttons: Add, Edit, Duplicate, Delete, Enable, Disable, Unlock, Copy Preferences, Paste Preferences, Export, and Import.
- Roles:** A table with columns: Name and Description. It lists roles like 'Host Administrator', 'Network Administrator', 'Operator', 'SAN System Administrator', 'Security Administrator', 'Security Officer', and 'Zone Administrator'. Below the table are buttons: Add, Edit, Duplicate, and Delete.
- AOR:** A table with columns: Name and Description. It lists 'All Fabrics' and 'All Hosts'. Below the table are buttons: Add, Edit, Duplicate, and Delete.

Figure 4-41 Users Management main window

In the Users tab, you can see three main panes: Users, Roles, and AOR. In the Users pane, you can see all the users. When you select a user, you can see their roles and responsibilities in the Roles and AOR panes.

Before you create users, define the email Event Notification settings by clicking the **Email Events Notification Setup** button on the lower end of the Users main window. Enter your SMTP (email) server IP address, port, and credentials, as shown in Figure 4-42. If these things are not configured, users will not be able of receive Event notifications by using email.



The image shows a Windows-style dialog box titled "E-mail Event Notification Setup". It contains several configuration fields and options. At the top, there is a checked checkbox labeled "Enable E-mail Event Notification". Below this, the "E-mail Server" field contains the IP address "10.18.228.171". The "SMTP over SSL" checkbox is also checked. The "SMTP Port" field is set to "465". The "*SMTP ID" field contains "administrator", and the "*SMTP Password" field is masked with four dots. The "Reply Address" field contains "TITO@IBM.com". The "Summary Interval" is set to "5" with a dropdown menu showing "seconds". Under "E-mail Test Options", there are two radio buttons: "Send to" (unselected) and "Send to all users enabled for notification" (selected). A "Send Test E-mail" button is located below the radio buttons. At the bottom, there is a note: "*Fields are optional unless SMTP server enables authentication." and three buttons: "OK", "Cancel", and "Help".

<input checked="" type="checkbox"/> Enable E-mail Event Notification	
E-mail Server	10.18.228.171
SMTP over SSL	<input checked="" type="checkbox"/>
SMTP Port	465
*SMTP ID	administrator
*SMTP Password	••••
Reply Address	TITO@IBM.com
Summary Interval	5 seconds ▼
E-mail Test Options	
<input type="radio"/> Send to	
<input checked="" type="radio"/> Send to all users enabled for notification	
<input type="button" value="Send Test E-mail"/>	
*Fields are optional unless SMTP server enables authentication.	
<input type="button" value="OK"/> <input type="button" value="Cancel"/> <input type="button" value="Help"/>	

Figure 4-42 Email event notification setup

After you define the email Event Notification, click **Add** in the Users pane to start creating users. Provide all the user details, as shown in Figure 4-43.

Add User

User ID: ITSO_SAN Full Name: ITSO SAN Administration

Password: Description: Administrator

Confirm Password: Phone Number:

Account Status: ☒ Enable E-mail Notification: ☐ Enable [Filter](#)

Account State: Active E-mail Address: TITO@IBM.com

Assign the Roles and AOR for this user

Available Roles / AOR

- AOR
- Roles

Selected Roles / AOR

- AOR
 - All Hosts
 - All Fabrics
- Roles
 - SAN System Administrator
 - Network Administrator
 - Security Administrator
 - Zone Administrator
 - Operator
 - Security Officer
 - Host Administrator

OK Cancel Help

Figure 4-43 Adding a user

When you are satisfied with the user information input, click **OK**.

While you create users, you can define roles for the SAN administrators by using the Roles tab. You can assign both roles and AOR. Select the roles that you want to assign from the Available Roles / AOR pane and click the right arrow to move them to Selected Roles / AOR pane, and then click **OK**. For information about modifying the user, see “Modifying user accounts” on page 91.

Modifying user accounts

To modify a user account, click **Servers** → **Users**, go to the Users pane, click the account you want to modify, and select **Edit**. A window opens as shown in Figure 4-44. Select the roles that you want to add / remove and click **OK**. You can also change the password, email address, and email notification filter.

Edit User

User ID: ITSO_SAN Full Name: ITSO SAN Administrator

Password: Description: Administrator

Confirm Password: Phone Number:

Account Status: ☒ Enable E-mail Notification: ☐ Enable [Filter](#)

Account State: Active E-mail Address: TITO@IBM.com

Assign the Roles and AOR for this user

Available Roles / AOR

- AOR
 - Roles

Selected Roles / AOR

- AOR
 - All Hosts
 - All Fabrics
- Roles
 - SAN System Administrator
 - Network Administrator
 - Security Administrator
 - Zone Administrator
 - Operator
 - Security Officer
 - Host Administrator

OK Cancel Help

Figure 4-44 Modifying an existing user

Disabling user accounts

To disable a user account, click **Server** → **Users** and then go to the Users pane. Select the user account that you want to disable and click **Disable**. After you click **Disable**, a window opens with a warning message that states that if the user is logged in they will be logged out, as shown in Figure 4-45. To enable the account again, select the disabled account and click **Enable**.

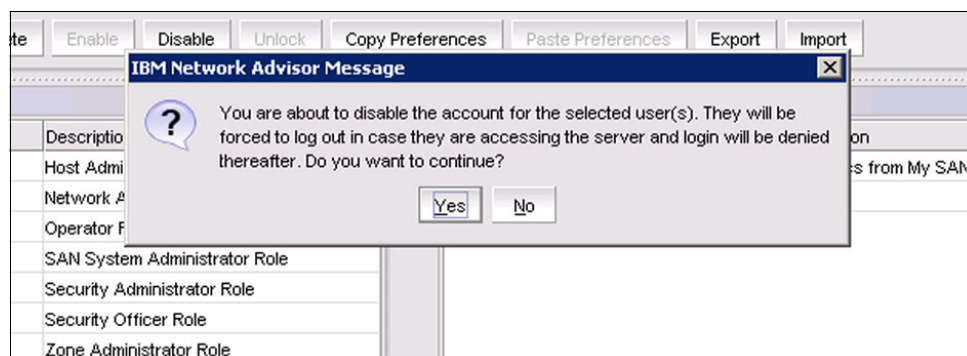


Figure 4-45 Disabling a user account

Deleting user accounts

To delete an account permanently, click **Server** → **Users** and then go to the Users pane. Select the user account that you want to delete and click **Delete**. After you click **Delete**, a dialog box opens to confirm the deletion, as shown in Figure 4-46. Click **Yes** to delete the account or **No** to cancel the deletion.

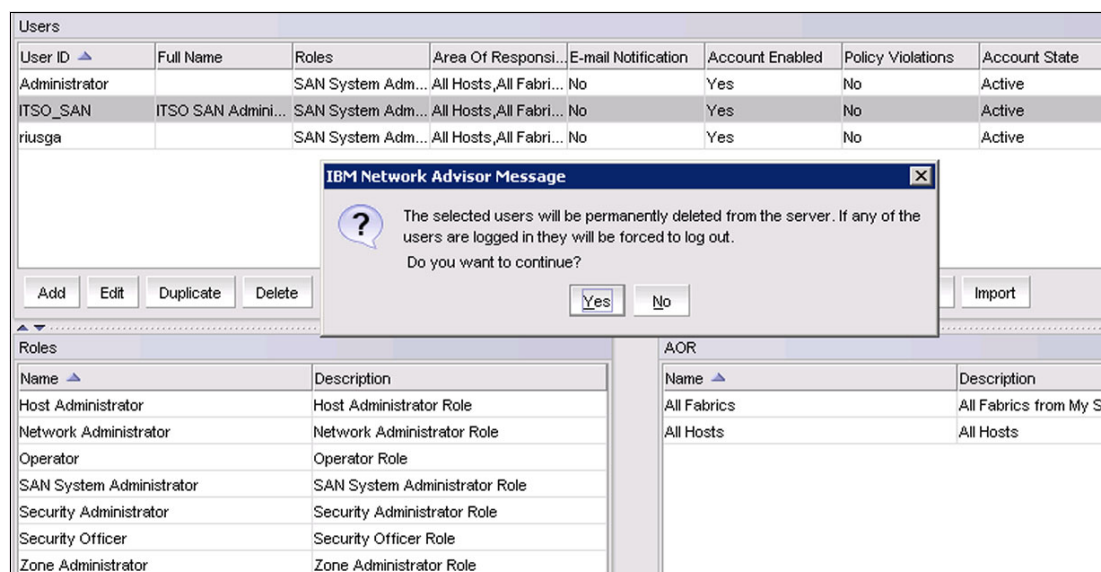


Figure 4-46 Deleting a user account

Defining password policies for user accounts

As shown in Figure 4-47, you can define different password policies for a user account. You can set the password age and the warning period to alert about the expiration of the current password, and the password history. Also, you can set a policy to define a password with a lockout threshold and the lockout duration. Click **View Policy Violators** to see any user accounts that were violated.

The screenshot displays the 'Policy' tab in the IBM Network Advisor user management interface. It contains several sections for configuring password and login policies:

- Password Expiration:** Includes 'Password Age' set to 60 days (range 0-999) and 'Warning Period' set to 45 days (range 0-998).
- Password History:** Includes 'History Count' set to 1 (range 1-24).
- Password Format:** Includes a checkbox for 'Empty Password' (unchecked), and fields for 'Minimum Length' (8, range 4-127), 'Upper Case Characters' (2, range 0-127), 'Lower Case Characters' (1, range 0-127), 'Number of Digits' (2, range 0-127), 'Punctuations Required' (0, range 0-127), 'Maximum Repeat' (2, range 1-127), and 'Maximum Sequence' (1, range 1-127).
- Lockout Support:** Includes 'Lockout Threshold' set to 2 times (range 0-999) and 'Lockout Duration' set to 30 minutes (range 0-99999).
- Login Policy:** Includes 'Login Mode' set to 'Concurrent Login' and 'Action' set to 'Reject New Sessions'.

Figure 4-47 Customizing users policies

For more information about IBM Network Advisor user management, see the “User Account Management” section in the *Network Advisor 12.4.2 SAN User Manual* on the following link:

<http://www.brocade.com/content/html/en/user-guide/networkadvisor-1242-san-manual>

4.9.2 Discovering and adding SAN fabrics

This section describes how to discover and add fabrics to IBM Network Advisor. It also explains terms such as seed switch, Fabric Configuration Switches (FCS), and Fabric Discovery.

Seed switch

The *seed switch* is the switch that IBM Network Advisor uses to communicate with the fabrics. The seed switch uses in-band communication to obtain fabric-wide information about the

name server, zoning, and fabric membership from all other switches. There must be one seed switch present in each of the managed fabrics.

The seed switch must ideally be the one running the highest FOS code level in the fabric, prioritizing the DCX backbone chassis for this role. It must also be HTTP reachable by the IBM Network Advisor server.

Note: If ipfilter is implemented on the switch, the IBM Network Advisor server IP address must be granted with http (port 80) connectivity to the switch.

Sometimes, the seed switch is auto-selected, such as when a fabric segments or when two fabrics merge. Other times, you are prompted (when an event is triggered) to change the seed switch, such as in the following cases:

- ▶ If during a fabric discovery the management application detects that the seed switch is not running a supported version, you are prompted to change the seed switch.
- ▶ When one or more switches join the fabric or if the switch firmware is changed on any of the switches in the fabric, the management application checks to make sure that the seed switch is still running a supported version. If it is not, then you are prompted to either upgrade the firmware on the seed switch or to change the seed switch to a switch that is running supported firmware.

If a new fabric is created as a result of a pre-existing fabric segmentation, the management application continues to monitor that new fabric. However, if any switch with a later FOS version joins the fabric, an event is triggered that informs you that the seed switch is not running the latest firmware. The event will suggest that you change the seed switch role to this new switch that is running the highest level of firmware.

Note: If a seed switch is segmented or merged, historical data, such as the offline zone database, profile, reports, and Firmware Download profile, can be lost. Segmentation of a seed switch does not result in formation of a new fabric. If a merge occurs, the historical data is lost only from the second fabric.

You can change the seed switch if the following conditions are met:

- ▶ The new seed switch is HTTP-reachable from the management application.
- ▶ The new seed switch is a primary FCS.
- ▶ The new seed switch is running the latest FOS version in the fabric.

This operation preserves historical and configuration data, such as performance monitoring and user-customized data for the selected fabric.

If during the seed switch change, the fabric is deleted but the rediscovery operation fails (for example, if the new seed switch becomes unreachable using HTTP), then you must discover the fabric again. If you rediscover the fabric by using a switch that was present in the fabric before the change seed switch operation was performed, then all of the historical and configuration data is restored to the rediscovered fabric. If you rediscover the fabric by using a switch that was added to the fabric after the fabric was deleted, then the historical and configuration data is lost.

If multiple users try to change the seed switch of the same fabric simultaneously, only the first change seed switch request is run. Subsequent requests that are initiated before the first request completes fail.

If another user changes the seed switch of a fabric you are monitoring, and if you have provided login credentials for only that seed switch in the fabric, then you lose connection to that seed switch.

Seed switch failover

The management application collects fabric-wide data (such as fabric membership, connectivity, name server information, and zoning) by using the seed switch. Therefore, when a seed switch becomes unreachable or there is no valid seed switch, the fabric becomes unmanageable.

When the seed switch cannot be reached for three consecutive fabric refresh cycles, the management application looks for another valid seed switch in the fabric, verifies that it can be reached, and has valid credentials. If the seed switch meets this criteria, the management application automatically fails over to this new switch, which becomes the new seed switch.

It is possible that auto-failover might occur to a seed switch that is not running the latest firmware version. In this instance, any function that has a direct dependency on the firmware version of the seed switch is affected and restricted by the failover seed switch capabilities.

Changing the seed switch

When you change the seed switch for a fabric, the management application performs these checks in the following order:

- ▶ Identifies all switches and removes switches running unsupported firmware versions.
- ▶ Identifies which of the remaining switches are running the latest firmware versions.
- ▶ Filters out switches that are not reachable.
- ▶ Identifies which switches are Virtual Fabric-enabled switches (FOS only). If there are Virtual Fabric-enabled switches, the management application uses only these switches as preferred seed switches candidates. If there are no Virtual Fabric-enabled switches, it continues with the next check.

To change the seed switch, complete the following steps:

1. Click **Discovery** → **Fabrics**. The Discover Fabrics window opens.
2. Select the fabric for which you want to change the seed switch from the Discovered Fabrics table. If a device joins or merges with a fabric and fabric tracking is active, you must accept changes to the fabric before the new devices display in the Seed Switch window.
3. Click **Seed Switch**.
If the fabric contains other switches that are running the latest version and are also HTTP-reachable from the management application, the Seed Switch window opens. Otherwise, a message displays that you cannot change the seed switch.
4. Select a switch to be the new seed switch from the Seed Switch window.
You can select only one switch. Only switches that are running the latest FOS version in the fabric are displayed. The current seed switch is not displayed in this list.
5. Click **OK** in the Seed Switch window.
If you are not already logged in to the seed switch, the Fabric Login window opens. If you are successfully authenticated, the fabric is deleted from the management application without purging historical data, and the same fabric is rediscovered with the new seed switch.
6. Click **Close** in the Discover Fabrics window.

Fabric Configuration Server policies

The FCS policy in the base FOS can be set on a local switch basis and can be set on any switch in the fabric. The FCS policy is not present by default. It must be created. When the FSC 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.

For more information about FCS, see the “Configuring Security Policies” section in the *Brocade Fabric OS 7.4.1 Administrator Guide*, found at the following link:

<http://www.brocade.com/content/html/en/administration-guide/fos-741-adminguide/>

While it discovers and adds new fabrics, the management application checks to confirm that the seed switch is running a supported FOS version in the fabric. If it is not, the management application prompts you to select a new seed switch.

For a FOS fabric, the seed switch must be the primary FCS. If you use a non-primary FCS to discover the fabric, the management application displays an error and does not allow the discovery to proceed. If the management application has already discovered the fabric, but you then create the FCS policy and the seed switch is not a primary FCS, an event is generated during the next poll.

The management application cannot discover a fabric that is in the process of actively configuring to form a fabric. Wait until the fabric is formed and stable, then reattempt the fabric discovery.

After fabric discovery successfully completes, all clients are updated to display the newly discovered fabric. During fabric discovery, you can define an IPV4 or IPV6 address. However, the management application uses the preferred IP format to connect to the devices.

Note: Discovery of a secure FOS fabric in strict mode is not supported.

FCS policy and seed switches

The management application requires that the seed switch is the primary FCS switch at the time of discovery.

Setting the time on the fabric sets the time on the primary FCS switch, which then distributes the changes to other switches.

When the FCS policy is defined, running the **configdownload** command is allowed only from the primary FCS switch. However, the management application does not check at the time of download that the switch is the primary FCS switch.

Discovering specific IP addresses or subnets

To discover specific IP address or subnets, complete the following steps:

1. Click **Discover** → **Fabric**. Figure 4-48 shows the discovery procedure.

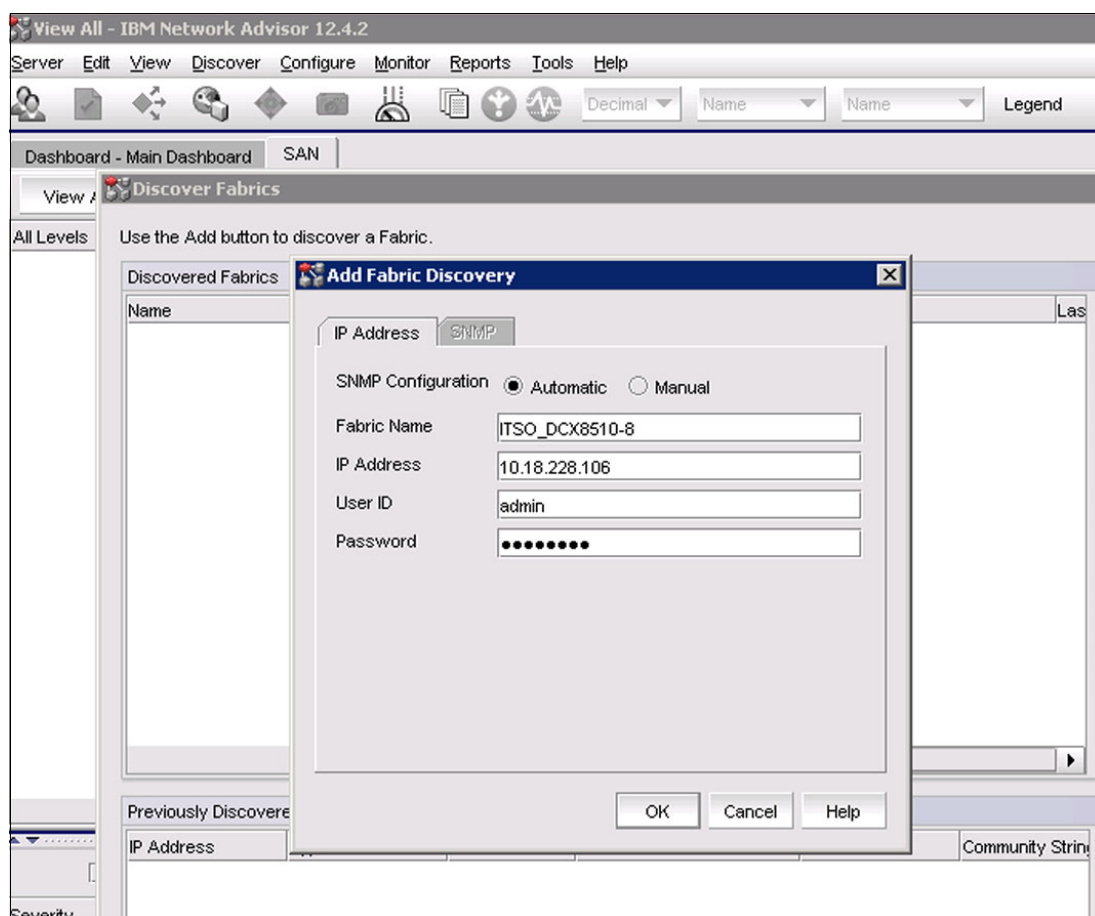


Figure 4-48 Discovering new fabrics

2. After you provide the required information, click **OK** to discover and add the fabric.
3. To configure the SNMP settings, set the SNMP Configuration to **Manual** and click the SNMP tab to customize the SNMP protocol version and credentials.

Deleting a fabric from IBM Network Advisor

To delete a fabric from IBM Network Advisor, click **Discover** → **Fabric**, select the fabric that you want to delete, and click **Delete**. A window opens and prompts you to confirm the deletion, as shown in Figure 4-49. Click **Yes** to delete or **No** to cancel the deletion.

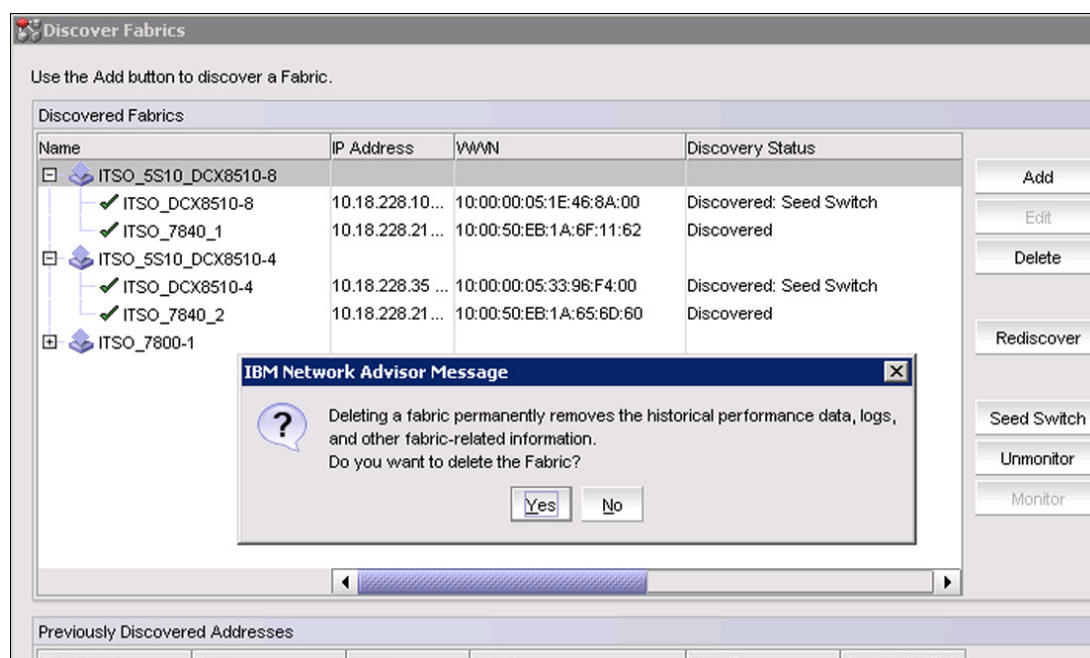


Figure 4-49 Deleting a fabric on IBM Network Advisor

4.10 New features of IBM Network Advisor V12.4.2

The following features were introduced with IBM Network Advisor V12.4.2:

- ▶ FOS 7.4 Platform support
- ▶ COMPASS
 - Custom configuration blocks and templates
 - Template configuration from existing switch configuration
 - Create/Edit/Delete user-defined switch groups
 - Link templates to fabrics or user-defined switch groups
 - Monitor configuration drifts on linked switches
 - Dashboard widget to monitor configuration drifts
- ▶ IP Storage
 - Display in property sheets
 - Dashboard support for monitoring Port Health
 - Storage port details in Network OS cluster reports
- ▶ IP Extension
 - Enhancements to Topology, FCIP Tunnel Config dialog
 - Support to configure FC and IP compression modes
 - Support for quality of service (QoS) distribution settings
 - Adaptive Rate Limiting configuration
 - HA config support for existing circuits

- ▶ Web Client Enhancements
 - Real Time Graph widget support
 - Dashboard Playback support
 - Event page enhancements
 - REST API enhancements
- ▶ Dashboard Enhancements
 - Enhancements to Network Scope Zones and Zone Alias
- ▶ MAPS Enhancements
 - Bottleneck detection indication
 - FPI support for new actions that include FMS, Toggle, and SDDQ
 - Selective distribution of policies
 - Clear E-mail support
- ▶ Zoning Enhancements
 - Support for Peer, LSAN peer, and Target Driven Peer Zones
- ▶ FICON
 - Enhancements to Configure Cascaded FICON Fabric dialog
 - FICON Merge Wizard enhancements
 - Encryption and Compression configuration
- ▶ Fault Management Enhancements
 - New KPI widget in Dashboard
 - Master Log enhancements
 - Audit Log enhancements
- ▶ Other Enhancements
 - Flow Vision - Zone Alias support
 - Port Decommission enhancements
 - SAN42B-R Base Switch support
 - Fabric Watch deprecation for Fabric OS V7.4.0 and later
 - Embedded server JRE upgraded to version 1.7u80
 - Remote client JRE support for 1.8u51
 - Jserver upgrade to 3.10

For more information about features that are introduced in this and earlier IBM Network Advisor 12 product family versions, see *IBM Network Advisor 12 Release Notes* at the following link:

ftp://public.dhe.ibm.com/storage/san/networkadvisor/IBM_Network_Advisor_v12.4.2_ReleaseNotes.pdf

4.11 IBM Network Advisor Dashboard overview

The IBM Network Advisor dashboard can be accessed by selecting the Dashboard - Product Status and Traffic tab on the main window. It provides a high-level overview of the network and the current state of the management devices. You can use the dashboards to easily check the status of the devices in the network. The dashboards also provide several features to help you quickly access reports, device configurations, and system logs.

The dashboards update regardless of the currently selected tab (SAN or Dashboard) or the SAN size. However, data might become momentarily out of sync between the dashboards and other areas of the applications. For example, if you remove a product from the network while another user navigates from the dashboard to a more detailed view of the product, the product might not appear in the detail view immediately.

Figure 4-50 illustrates the IBM Network Advisor Dashboard main window and its components.

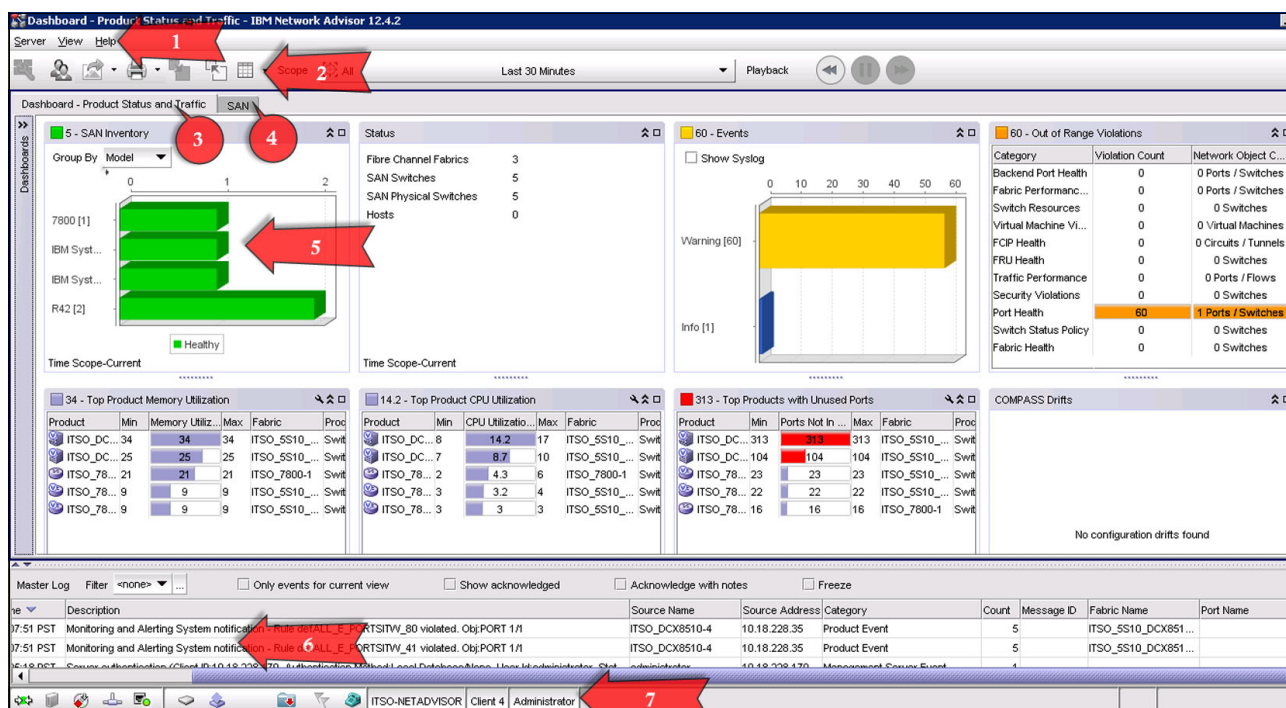


Figure 4-50 IBM Network Advisor Dashboard main window

The numbers correspond to the following items:

1. Menu bar: Lists commands that you can run on the dashboard. The dashboard also provides a menu to reset the dashboard back to the defaults. You can reset the dashboard back to the default settings by right-clicking and selecting Reset to Default.
2. Toolbar: Provides buttons that enable quick access to windows and functions.
3. Dashboard tab: Provides a high-level overview of the network that is managed by the management application server.
4. SAN: Displays the master log, Minimap, Connectivity map (topology), and product list.
5. Widgets: Displays the operational status, inventory status, event summary, performance monitors, and overall network or fabric status.
6. Master log: Displays all events that have occurred on the management application.
7. Status bar: Displays the connection, port, product, fabric, special event, Call Home, backup status, and server and user data.

The Dashboard be customized according to your needs:

- Network Scope: Sets the network scope to your AOR or custom subset of objects (fabrics, devices, or groups).

- ▶ **Time Interval:** Setting the time interval configures the data display time range for all the applicable widgets. Time Interval in the Scope list allows you to select a specific time range for which you want to display data in the Dashboards or Events page.
- ▶ **Dashboard Display:** You can set the dashboard display to minimize or expand all status widgets and performance monitors and return to the default settings.
- ▶ **Dashboard Playback:** Allows you to use the dashboard control buttons (**Rewind**, **Pause**, and **Forward**) to configure the dashboard to play back the recorded status and performance widget data incrementally or to pause playback.
- ▶ **Default Dashboards:** IBM Network Advisor provides preconfigured dashboards, which provide high-level overview of the network, the current states of managed devices, and performance of devices, ports, and traffic in the network.
- ▶ **My Dashboards:** The My Dashboards list includes all dashboards that you create in IBM Network Advisor. The My Dashboards list does not display until you save a dashboard to My Dashboards in IBM Network Advisor. If you share a dashboard, you created the shared icon displays next to the dashboard name in the My Dashboards list.
- ▶ **Shared Dashboards:** The Shared Dashboards list includes all user-defined dashboards that have been shared with other users in IBM Network Advisor. Shared dashboards display in the following format: dashboard_name (user_name). The Shared Dashboards list does not display until a dashboard is shared with other users.

For more information about these IBM Network Advisor dashboard features and the multiple dashboard widgets, see the “Dashboard customization” section of the *Network Advisor 12.4.2 SAN User Manual*, at the following link:

<http://www.brocade.com/content/html/en/user-guide/networkadvisor-1242-san-manual>

4.12 Scheduling daily or weekly backups for the fabric configuration

As a good practice, you should back up the configuration of all your fabrics in your SAN environment on a frequent basis. If the backup is configured at the SAN level instead of selecting individual switches, any new fabric that is discovered is automatically added to the list of fabrics to be backed up to the IBM Network Advisor repository.

To schedule the daily backup of the configuration by using IBM Network Advisor, click **Configure** → **Configuration File** → **Schedule Backup**.

As shown in Figure 4-51 on page 102, select the **Enable Scheduled backup** check box and schedule the frequency. You can choose **Daily**, **Weekly**, or **Monthly**, and as mentioned before the good practice is to create a daily backup schedule. Choose the time window to schedule the backup, and also choose the period of days you want to keep the configuration backups.

The Purge Backup settings range from 7 through 90 days; the default is 30 Days. You can choose all fabrics in the network by selecting **Backup all fabrics**, or by selecting some of the switches in the network. You should collect a backup of all the fabrics in the environment.

Figure 4-51 shows the scheduling of the backup of all fabrics and discovered switches.

Scheduled Backup of Switch Configurations

☒ Enable scheduled backup

Schedule

Frequency: Daily

Day: Wednesday

Hour: 23 Minute: 00

Purge Backups: 30 days and older

Scope - Includes all switches discovered at time of backup

☒ Backup all fabrics

Selected Fabrics

Backup	Fabric Name ▲	Status	# of Switches
<input checked="" type="checkbox"/>	ITSO_5S10_DCX851...	Healthy	2
<input checked="" type="checkbox"/>	ITSO_5S10_DCX851...	Healthy	2
<input checked="" type="checkbox"/>	ITSO_7800-1	Healthy	1

OK Cancel Help

Figure 4-51 Schedule backups of switch configurations window

Note: Older 8 Gbps b-type Gen 4 switch platforms require the Enhanced Group Management (EGM) license for this configuration backup procedure and to use the supportSave module.

For more information about switch configuration backups, see the “Configuration file management” section of the *Network Advisor 12.4.2 SAN User Manual* at the following link:

<http://www.brocade.com/content/html/en/user-guide/networkadvisor-1242-san-manual>

4.13 Call Home

Call Home notification allows you to configure the management application server to automatically send an email alert or dial in to a support center to report system problems with specified devices (FOS switches, routers, and directors). If you are upgrading from a previous release, all of your Call Home settings are preserved.

If you are installing IBM Network Advisor for the first time in your environment, click **Monitor** → **Event Notification** → **Call Home**, as shown in Figure 4-52, where you see that Call Home is enabled for different call home centers.

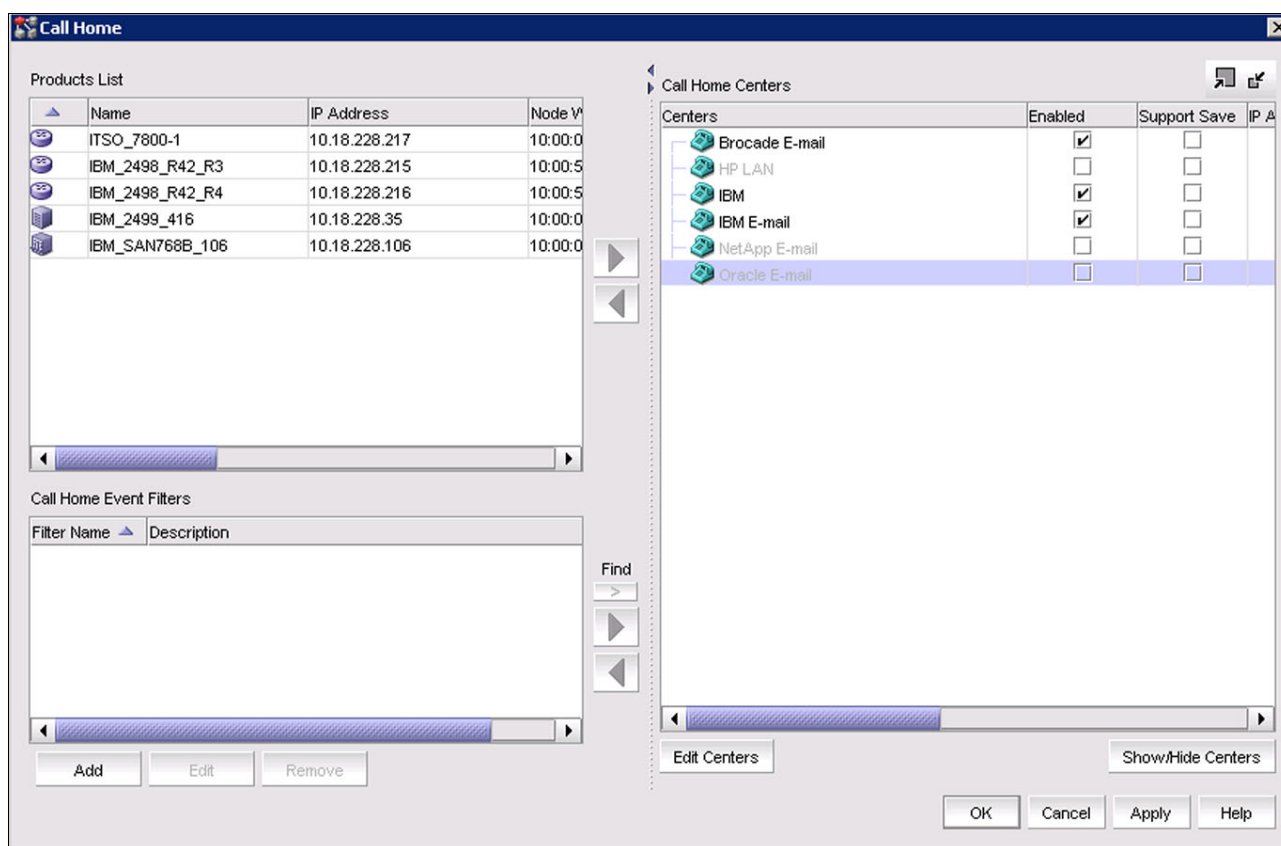


Figure 4-52 Call Home configuration main window

Clear the **Enabled** check boxes for the HP LAN, NetApp email, and Oracle email servers to disable or to hide them. When you clear one of these check boxes, a message is displayed: “Call Home center will be disabled. Do you want to continue?” Select **Yes** to continue or **No** to cancel.

Note: Call Home is supported on Windows systems for all modem and email Call Home centers and is supported on UNIX for the email Call Home centers.

Call Home allows you to automate tasks that occur when the Call Home event trigger is fired. When a Call Home trigger event occurs, the management application generates the following actions:

1. Sends an email alert to a specified recipient or dials in to a support center.
2. Triggers the **supportsave** command on the switch (if the **supportsave** command is enabled on the switch) before sending an alert. The **supportsave** location is included in the alert.
3. Adds an entry to the master log file and window display.
4. Generates an HTML report for email-based Call Home centers.

Call Home allows you to perform the following tasks:

- ▶ Assign devices to and remove devices from the Call Home centers.
- ▶ Define filters from the list of events that are generated by FOS devices.
- ▶ Edit and remove the filters that are available in the Call Home event filters table.
- ▶ Apply filters to and remove filters from the devices individually or in groups.
- ▶ Edit individual Call Home center parameters to dial a specified phone number or email a specific recipient.
- ▶ Enable and disable individual devices from contacting the assigned Call Home centers.
- ▶ Show or hide Call Home centers on the display.
- ▶ Enable or disable Call Home centers.

4.13.1 System requirements

Call Home through modem requires the following hardware equipment:

- ▶ Any Windows server with an internal or external modem connection
- ▶ An analog phone line

Open the Call Home main window by selecting **Monitor** → **Event Notification** → **Call Home**. Select **IBM** on the Call Home Centers list, and then click **Edit Centers**. After you select **Edit Centers**, the Configure Call Home Centers window opens, as shown in Figure 4-53. Select **IBM** from the **Call Home Centers** drop-down menu and configure all the required fields, such as **Primary Connection**, **Backup Connection**, and **Phone Number**. Click **OK**.

The screenshot shows the 'Configure Call Home Center' window. At the top, there's a title bar with the text 'Configure Call Home Center' and a close button. Below the title bar, the 'Call Home Centers' section has a dropdown menu currently showing 'IBM' and an 'Enable' checkbox that is checked. The 'Set heartbeat interval at' section has a checked checkbox, a text box with the number '1', and the text 'days (1-28)'. The 'Time Out' section has a text box with '60' and the text 'Sec'. The 'Retry Interval' section has a text box with '10' and the text 'Sec'. The 'Maximum Retries' section has a text box with '3'. The 'Call Home Center' section contains two text boxes: 'Primary Connection' and 'Backup Connection'. Below these is a 'Send Test' button. The 'Local Server' section has two text boxes: 'Phone Number' and 'Server ID', with 'Server ID' containing the text '55667788'. At the bottom of the window are four buttons: 'OK', 'Cancel', 'Apply', and 'Help'.

Figure 4-53 IBM Call Home Center (modem) configuration settings

4.13.2 Editing an email Call Home center

Email Call Home centers are available for IBM and Brocade. To edit one of these Call Home centers, click **Monitor** → **Event Notification** → **Call Home**, select either **IBM E-mail** or **Brocade**, and then select **Edit centers**. A window to configure the email parameters opens and prompts you to provide your information, as shown in Figure 4-54. After you enter the information, click **OK** to enable the email Call Home function.

Configure Call Home Center

Call Home Centers: IBM E-mail ☒ Enable

Customer Details

Name: ITSO Organization

Company: IBM

Phone (Office):

Phone (Mobile):

SMTP Server Settings

Server Name: 10.18.228.171

SMTP over SSL: ☐

Port: 25

*Username: administrator

*Password:

E-mail Notification Settings

Reply Address: TITO@IBM.com

Send To Address: TITO@IBM.com

Send Test

*Fields are optional unless the SMTP server enables authentication

OK Cancel Apply Help

Figure 4-54 IBM Call Home Center (email) configuration settings

For more information about IBM Network Advisor, see the *Network Advisor 12.4.2 SAN User Manual* at the following link:

<http://www.brocade.com/content/html/en/user-guide/networkadvisor-1242-san-manual>



Product hardware

This chapter contains product hardware features for IBM b-type 16 Gbps directors and switches.

This chapter includes the following sections:

- ▶ Hardware overview
- ▶ IBM Gen 5 SAN b-type family

5.1 Hardware overview

The IBM b-type family of products provides a range of entry level and midrange switches, and enterprise-class directors.

5.1.1 Entry level, midrange, and director models

The entry level, midrange, and director models provide 2, 4, 8, and 16 Gbps port-to-port non-blocking throughput. Hub-based Fibre Channel Arbitrated Loop (FC-AL) solutions reduce performance as devices are added by sharing the bandwidth. However, an IBM b-type 16 Gbps SAN Fabric throughput continues to increase as more ports are interconnected.

These models are fully interoperable with previous IBM b-type 8 Gbps SAN switches, and can be added to existing fabrics, enabling transition from existing Fibre Channel storage networks to the newer technology.

For the latest information, see IBM SAN b-type website:

<http://www.ibm.com/systems/networking/switches/san/b-type/index.html>

5.1.2 IBM Gen 5 SAN b-type 16 Gbps family

Table 5-1 lists the Gen 5 16 Gbps SAN switches and directors model number with speed and port capabilities, the current (at the time of writing) version of FOS, and the type of application-specific integrated circuit (ASIC).

Table 5-1 IBM Gen 5 and b-type family

Switch type	Number of ports	Port speed	FOS	ASIC version
SAN768B-2	Up to 384 16 Gbps Up to 512 8 Gbps	2, 4, 8, 10 ^a or 16 Gbps	Version 7.0 or later	Condor3
SAN384B-2	Up to 192 16 Gbps Up to 256 8 Gbps	2, 4, 8, 10 ^a or 16 Gbps	Version 7.0 or later	Condor3
SAN96B-5	Up to 96 16 Gbps	2, 4, 8, 10 ^a or 16 Gbps	Version 7.1 or later	Condor3
SAN48B-5	Up to 48 16 Gbps	2, 4, 8, 10 ^a or 16 Gbps	Version 7.0 or later	Condor3
SAN24B-5	Up to 24 16 Gbps	2, 4, 8, or 16 Gbps	Version 7.0.1 or later	Condor3
SAN42B-R	Up to 24 16 Gbps Up to 16 1/10 GbE Up to 2 40 GbE	2, 4, 8, or 16 Gbps, 1/10 GbE, 40 GbE, 10 Gbps FC	Version 7.3 or later	Condor3

a. Active 10 Gbps Fibre Channel support, which provides integrated dense wavelength division multiplexing (DWDM) metro connectivity

5.2 IBM Gen 5 SAN b-type family

The Gen 5 platform is designed to support a long-term solution for mission-critical applications that require secure, high-performance, high-density server virtualization, cloud architectures, and low-latency storage networks.

At the time of writing, six IBM b-type Gen 5 16 Gbps switches are in the portfolio:

- ▶ IBM SAN24B-5 (2498-F24, 2498-X24, and 2498-24G)
- ▶ IBM SAN48B-5 (2498-F48)
- ▶ IBM SAN96B-5 (2498-F96 / 2498-N96)
- ▶ IBM SAN42B-R (2498-R42)
- ▶ IBM Fabric backbones:
 - IBM System Networking SAN384B-2 Backbone (2499-416)
 - IBM System Networking SAN768B-2 Backbone (2499-816)

5.2.1 IBM SAN24B-5 (2498-F24, 2498-X24, and 2498-24G)

The SAN24B-5 is a 24-port entry level enterprise switch that combines flexibility, simplicity, and 16 Gbps Fibre Channel technology.

The SAN24B-5 requires Fabric Operating System (FOS) V7.0.1 or later. The Advanced Web Tools, Advanced Zoning, Full Fabric, and Enhanced Group Management features are part of the base FOS and do not require an extra license. Additional features such as Adaptive Networking, Advanced Performance Monitor, Fabric Watch, inter-switch link (ISL) Trunking, Extended Fabrics, Server Application Optimization, and 12-port Activation are available as optional licenses. Furthermore, an Enterprise Package is available as a bundle that includes one license for each of the optional licenses, except for the Extended Fabrics one.

SAN24B-5 provides the following features and benefits:

- ▶ Up to 24 auto-sensing ports of high-performance 16-Gbps technology in a single domain.
- ▶ Ports on Demand scaling (12 - 24 ports).
- ▶ 2, 4, 8, and 16 Gbps auto-sensing Fibre Channel switch and router ports:
 - 2, 4, and 8 Gbps performance is enabled by 8 Gbps SFP+ transceivers.
 - 4, 8, and 16 Gbps performance is enabled by 16 Gbps SFP+ transceivers.
- ▶ Enterprise features that maximize availability with redundant, hot-pluggable components and nondisruptive software upgrades and RAS functioning to help minimize downtime.
- ▶ Universal ports self-configure as E, F, or M ports. EX_Ports can be activated on a per-port basis with the optional Integrated Routing license. The D-port function is also available for diagnostic tests.
- ▶ Airflow is set for port side exhaust.
- ▶ ISL Trunking, which allows up to eight ports (at 2, 4, 8, or 16 Gbps speeds) between a pair of switches to combine to form a single, logical ISL with a speed of up to 128 Gbps (256 Gbps full duplex) for optimal bandwidth usage and load balancing. The base model permits one 8-port trunk plus one 4-port trunk.
- ▶ Dynamic Path Selection (DPS), which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- ▶ Brocade-branded SFP+ optical transceivers that support any combination of short wavelength (SWL), long wavelength (LWL), and extended long wavelength (ELWL) optical media among the switch ports.
- ▶ Extended distance support enables native Fibre Channel extension up to 7,500 km at 2 Gbps.
- ▶ Support for unicast traffic type.
- ▶ Delivers SAN technology within a flexible, simple, and easy-to-use solution dynamic fabric provisioning, critical monitoring, and advanced diagnostic features provide streamlined deployment and troubleshooting time.

- ▶ Dual functions as either a full-fabric SAN switch or an N_Port ID Virtualization (NPIV)-enabled access gateway.
- ▶ Support for an Access Gateway configuration, where server ports that are connected to the fabric core are virtualized.
- ▶ Extensive diagnostic and system-monitoring capabilities for enhanced high reliability, availability, and serviceability (RAS).
- ▶ The EZSwitchSetup wizard makes SAN configuration a three-step point-and-click task.
- ▶ Real-time power monitoring enables users to monitor real-time power usage of the fabric at a switch level.
- ▶ The base unit includes one (249824G/2498-X24) or two (2498-F24) integrated power supplies and fans.

Figure 5-1 shows a front view of the 2498-F24.



Figure 5-1 IBM SAN24B-5 (2498-F24)

For more information, see the IBM System Storage SAN24B-5 topic that is found at the following website:

<http://www.ibm.com/systems/networking/switches/san/b-type/san24b-5/index.html>

Hardware layout

The port side of the SAN24B-5 includes the system status LED, the console port, the Ethernet port and accompanying LEDs, the USB port, and the Fibre Channel ports and corresponding port status LEDs.

Figure 5-2 shows the port side of the SAN24B-5.




Figure 5-2 SAN24B-5 port side

Figure 5-3 shows the non-port side of the SAN24B-5, which contains the power supply (including the AC power receptacle and AC power switch) and fan assemblies. The base model configuration with a single assembly is shown.




Figure 5-3 SAN24B-5 non-port side

5.2.2 IBM System Networking SAN48B-5 (2498-F48)

The SAN48B-5 is a flexible, easy-to-use Enterprise-Class SAN switch for private cloud storage. It is a 1U form factor unit that is configurable in 24, 36, or 48 ports and supports auto-sensing 2, 4, 8, or 16 Gbps and 10 Gbps speeds. It can be deployed as a full-fabric switch or as an NPIV-enabled Access Gateway. It is also enhanced with enterprise connectivity that adds support for IBM FICON. It includes dual, hot-swappable redundant power supplies with integrated system cooling fans.

The SAN48B-5 requires FOS V7.0 or later. The Advanced Web Tools, Advanced Zoning, Enhanced Group Management, Fabric Watch, Full Fabric, and Virtual Fabrics features are embedded in the base FOS and do not require an additional license. Additional features, such as 12-port Activation, FICON with Control Unit Port (CUP) Activation, Adaptive Networking, Advanced Performance Monitoring, Extended Fabrics, Integrated Routing, ISL Trunking, Server Application Optimization (SAO), and Integrated 10 Gbps Fibre Channel Activation, are available as optional licenses. The Enterprise Advanced Bundle includes one license for each of the Extended Fabric, Advanced Performance Monitoring, Trunking Activation, Adaptive Networking, and SAO functions.

SAN48B-5 provides the following features and benefits:

- ▶ Up to 48 16 Gbps auto-sensing ports in an energy-efficient 1U form factor.
- ▶ Ports on Demand (PoD) licensing capabilities for scaling (24 - 48 ports in 12-port increments).
- ▶ Supports 2, 4, 8, and 16 Gbps auto-sensing Fibre Channel switch and router ports. 2, 4, and 8 Gbps performance is enabled by 8 Gbps SFP+ transceivers:
 - 4, 8, and 16 Gbps performance is enabled by 16 Gbps SFP+ transceivers.
 - 10 Gbps manual set capability on FC ports (requires the optional 10 Gigabit FCIP/Fibre Channel license).
 - 10 Gbps performance is enabled by 10 Gbps SFP+ transceivers.
 - Ports can be configured for 10 Gbps for metro connectivity.¹
- ▶ Gen 5 16 Gbps optimized ISL. ISL Trunking² allows up to eight ports (at 2, 4, 8, or 16 Gbps speeds) between a pair of switches to combine to form a single logical ISL with a speed of up to 128 Gbps (256 Gbps full duplex) for optimal bandwidth usage and load balancing.
- ▶ Universal ports self-configure as E, F, M, or D ports. EX_Ports can be activated on a per port basis with the optional Integrated Routing license.
- ▶ The Diagnostic Port (D-Port) feature provides physical media diagnostic, troubleshooting, and verification services.
- ▶ In-flight data compression and encryption on up to two ports provides efficient link usage and security.
- ▶ Options for port side exhaust (default) or non-port side exhaust airflow for cooling.
- ▶ Virtual Fabric (VF) support to improve isolation between different VFs.
- ▶ Fibre Channel Routing (FCR) service, available with the optional Integrated Routing license, provides improved scalability and fault isolation.
- ▶ FICON, FICON Cascading, and FICON Control Unit Port ready.
- ▶ Dynamic Path Selection (DPS), which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- ▶ Brocade-branded SFP+ optical transceivers, which support any combination of SWL, LWL, or ELWL optical media among the switch ports.
- ▶ Extended distance support enables native Fibre Channel extension up to 7,500 km at 2 Gbps.
- ▶ Support for unicast, multicast (255 groups), and broadcast data traffic types.
- ▶ Support for an Access Gateway configuration where server ports that are connected to the fabric core are virtualized.
- ▶ Extensive diagnostic and system-monitoring capabilities for enhanced high RAS.
- ▶ 10G Fibre Channel integration on the same port provides for DWDM metro connectivity on the same switch (can be done on the first eight ports only).
- ▶ The EZSwitchSetup wizard makes SAN configuration a three-step point-and-click task.

¹ Only the first eight ports can be used as Metro Mirror.

² Trunking is a licensable feature.

- Real-time power monitoring enables users to monitor real-time power usage of the fabric at a switch level.
- Multi-tenancy in cloud environments through Virtual Fabrics, Integrated Routing, Quality of Service (QoS), and fabric-based zoning features.

Figure 5-4 shows the SAN48B-5 (2498-F48) front view.



Figure 5-4 IBM System Storage SAN48B-5 (2498-F48) front view

For more information, see the IBM System Storage SAN48B-5 topic that is found at the following website:

<http://www.ibm.com/systems/networking/switches/san/b-type/san48b-5/index.html>

Hardware layout

The port side of the SAN48B-5 includes the system status LED, the console port, the Ethernet port and LEDs, the USB port, and the Fibre Channel ports and corresponding port status LEDs.

Figure 5-5 shows the SAN48B-5 port side.

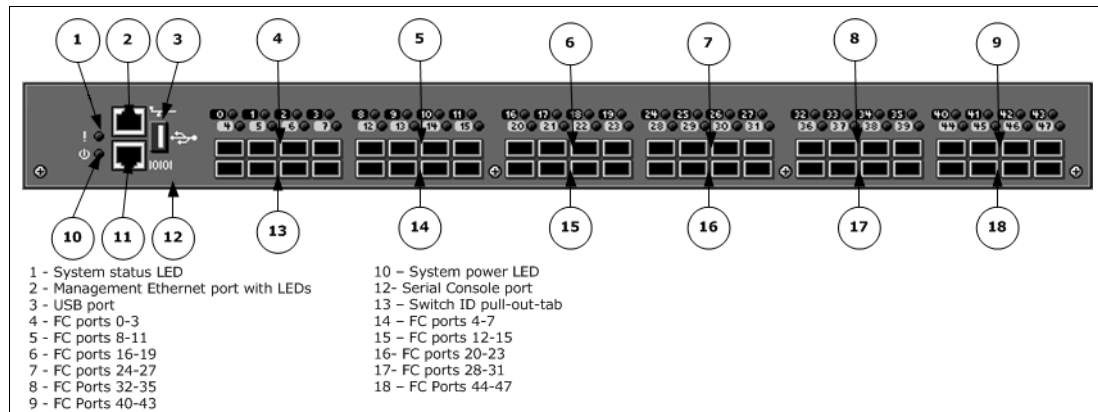


Figure 5-5 SAN48B-5 port side

The SAN48B-5 non-port side contains the power supplies, on/off switches, and the power plug receptacle. Figure 5-6 shows the SAN48B-5 non-port side.




Figure 5-6 SAN48B-5 non-port side

5.2.3 IBM System Networking SAN96B-5 (2498-F96 / 2498-N96)

The SAN96B-5 is a high-density enterprise-class switch for large and growing storage area network (SAN) infrastructures. It is designed to provide highly resilient, scalable, and simplified network infrastructure for storage that delivers 16 Gbps performance and Gen 5 capabilities. With up to 96 ports in a 2U form factor, SAN96B-5 is an enterprise-class Fibre Channel SAN switch that is designed for maximum flexibility. It offers a “pay-as-you-grow” scalability with PoD scaling (48 - 96 ports in 24-port increments).

SAN96B-5 requires FOS V7.1 or later. The Advanced Web Tools, Advanced Zoning, Virtual Fabrics, Full Fabric, Adaptive Networking, Server Application Optimization, and Enhanced Group Management features are embedded in the base FOS. These features do not require an extra license. Additional features, such as 24-port Activation, Advanced Performance Monitor, Fabric Watch, Extended Fabrics, Integrated Routing, Trunking Activation, and Integrated 10 Gbps Fibre Channel Activation, are available as optional licenses. The optional Enterprise Advanced Bundle includes one license for each of the Fabric Watch, Extended Fabric, Advanced Performance Monitor, and Trunking Activation features.

IBM Network Advisor V12.0 (or later) is the base management software for the SAN96B-5.

SAN96B-5 provides the following features and benefits:

- ▶ Up to 96 auto-sensing ports of high-performance 16 Gbps technology in a single domain.
- ▶ Supports 2, 4, 8, and 16 Gbps auto-sensing Fibre Channel switch and router ports. 2, 4, and 8 Gbps performance is enabled by 8 Gbps SFP+ transceivers:
 - 4, 8, and 16 Gbps performance is enabled by 16 Gbps SFP+ transceivers.
 - 10 Gbps manual set capability on FC ports (requires the optional 10 Gigabit FCIP/Fibre Channel license).
 - 10 Gbps performance is enabled by 10 Gbps SFP+ transceivers.
 - Ports can be configured for 10 Gbps for metro connectivity.³
- ▶ PoD licensing capabilities for scaling (48 - 72 or 96 ports).
- ▶ FC ports self-configure as E_ports and F_ports. EX_ports can be activated on a per-port basis with the optional Integrated Routing license.
- ▶ Mirror ports (M_ports) and diagnostic ports (D_ports) must be manually configured.
- ▶ The Brocade Diagnostic Port (D_port) feature provides physical media diagnostic, troubleshooting, and verification services.

³ Only the first eight ports can be used as Metro Mirror.

- ▶ In-flight data compression and encryption on up to 16 ports (up to eight ports at 16 Gbps) provides efficient link usage and security.
- ▶ Options for port side exhaust (default) or non-port side exhaust airflow for cooling.
- ▶ VF supports to improve isolation between different VFs.
- ▶ The FCR service, available with the optional Integrated Routing license, provides improved scalability and fault isolation.
- ▶ ISL Trunking (licensable), which allows up to eight ports (at 2, 4, 8, or 16 Gbps speeds) between a pair of switches to combine to form a single, logical ISL with a speed of up to 128 Gbps (256 Gbps full duplex) for optimal bandwidth usage and load balancing. There is no limit to how many trunk groups can be configured.
- ▶ DPS, which optimizes fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric.
- ▶ Brocade-branded SFP+ optical transceivers that support any combination of SWL, LWL, or ELWL optical media among the switch ports.
- ▶ Extended distance support enables native Fibre Channel extension up to 7,500 km at 2 Gbps.
- ▶ Support for unicast data traffic types.
- ▶ Dual redundant power supplies and integrated fans that support optional airflow configurations.
- ▶ Provides up to eight in-flight encryption and compression ports, delivering data center-to-data center security and bandwidth savings.
- ▶ Optimizes link and bandwidth usage with ISL Trunking and DPS.
- ▶ Extensive diagnostic and system-monitoring capabilities for enhanced high RAS.
- ▶ 10 Gbps Fibre Channel integration on the same port provides for DWDM metro connectivity on the same switch (can be done on first eight ports only with the appropriate license).
- ▶ The EZSwitchSetup wizard makes SAN configuration a three-step point-and-click task.
- ▶ Real-time power monitoring enables users to monitor real-time power usage of the fabric at a switch level.

Note: The only difference between the two models is airflow options. 2498-F96 is the “regular” version with air intake on non-port side and exhaust on port side. 2498-N96 has port to non-port side airflow.

Figure 5-7 shows the SANB96B-5 front view.




Figure 5-7 IBM System Storage SAN96B-5 (2498-F96) front view

For more information, see the IBM System Storage SAN96B-5 topic that is found at the following website:

<http://www.ibm.com/systems/networking/switches/san/b-type/san96b-5/index.html>

Hardware layout

The port side of the SAN96B-5 includes the system status LED, the console port, the Ethernet port and accompanying LEDs, the USB port, and the Fibre Channel ports and corresponding port status LEDs.

Figure 5-8 shows the SAN96B-5 port side view.




Figure 5-8 SAN96B-5 port side view

The SAN96B5 non-port side contains the power supplies (including the AC power receptacle) and fans.

Figure 5-9 shows the non-port side of the SAN96B-5.




Figure 5-9 SAN96B-5 non-port side view

5.2.4 IBM System Networking SAN384B-2 (2499-416) and IBM System Networking SAN768B-2 (2499-816)

IBM System Networking fabric backbone directors are designed to interconnect Fibre Channel based storage devices and servers. SAN768B-2 and SAN384B-2 Gen 5 Fibre Channel fabric backbones deliver reliable, scalable, and high-performance foundations for mission-critical storage.

SAN768B-2 and SAN384B-2 are designed to perform the following tasks:

- ▶ Increase business agility while providing nonstop access to information and reducing infrastructure and administrative costs.
- ▶ Deliver breakthrough performance with 16 Gbps Fibre Channel connectivity.
- ▶ Provide a long-term solution for mission-critical applications that require secure, high-performance, and low-latency storage networks.
- ▶ Take advantage of proven reliability and new technology to deliver enterprise-class reliability, availability, and serviceability.
- ▶ Simplify and centralize end-to-end SAN management with comprehensive diagnostic tests, monitoring, and automation.
- ▶ Improve energy efficiency by combining high bandwidth with low power consumption.
- ▶ Provide 99.999 percent uptime capabilities.
- ▶ Maximize investment protection.

Both directors provide the following features:

- ▶ Support 2, 4, 8, and 16 Gbps auto-sensing Fibre Channel switch and router ports. 2, 4, and 8 Gbps performance is enabled by 8 Gbps SFP+ transceivers.
 - 4, 8, and 16 Gbps performance is enabled by 16 Gbps SFP+ transceivers.
 - 10 Gbps manual set capability on FC ports (requires the optional 10 Gigabit FCIP/Fibre Channel license).
 - 10 Gbps performance is enabled by 10 Gbps SFP+ transceivers.
 - Supports 10 Gbps⁴ FC-type SFPs in 16 Gbps port blades only and also supports 10 GbE SFPs in the FX8-24 application blade. The two types of SFPs are not interchangeable.
- ▶ In-flight encryption and compression.
- ▶ 10 Gbps Fibre Channel ISL connections in metro optical connectivity or 10 Gbps dense wavelength division multiplexing (DWDM) devices.
- ▶ Optical ICL.
- ▶ Simplify scale-out network design to reduce network complexity, management, and costs.
- ▶ Support up to 100 m cable length for ICL links.
- ▶ Up to nine chassis in a full-mesh topology and up to 10 chassis core/edge.
- ▶ Up to 2.1 Tbps bandwidth on SAN768B-2 and up to 1.0 Tbps bandwidth on SAN384B-2.
- ▶ Optimize link and bandwidth usage with ISL Trunking and DPS.
- ▶ More scalable across distance:
 - 8000 buffers (four times of what exists on Gen 4).
 - Up to 5000 km distance at 2 Gbps.

⁴ The 10 Gbps ports can be configured manually on only the first eight ports of the 16 Gbps port blades.

SAN384B-2 (2499-416) hardware specification

SAN384B-2 (2499-416) is a powerful Gen 5 fabric backbone director in an 8U rack height chassis (plus 1U for the exhaust shelf).

The base version includes the following items:

- ▶ Eight-slot horizontal card cage
- ▶ Two Control processor (CP) blades
- ▶ Two Core (CR) switching blades
- ▶ Four slots for port and specialty blades
- ▶ Two standard power supplies in two bays
- ▶ Two cooling fan field-replaceable units (FRUs)

It has the following performance capabilities:

- ▶ Up to 192 ports at 16 Gbps in a single chassis
- ▶ 4.1 Tbps chassis bandwidth
- ▶ 3.1 Tbps Fibre Channel/FICON ports
- ▶ 1.0 Tbps ICL bandwidth
- ▶ 512 Gbps bandwidth per slot

Figure 5-10 shows the front view of a fully populated SAN384B-2.



Figure 5-10 Front-side angle of SAN384B-2

SAN768B-2 (2499-816) hardware specification

SAN768B-2 (2499-816) is the most scalable Gen 5 fabric backbone director, providing high-port density. It supports midrange to enterprise level SAN applications. The SAN768B-2 fabric backbone director integrates the new Gen 5 hardware into a 14U rack height.

It has the following features:

- ▶ 12-slot vertical card cage
- ▶ Two Control Processor (CP) blades
- ▶ Two Core (CR) switching blades
- ▶ Eight slots for port and specialty blades
- ▶ Two standard power supplies in four bays
- ▶ Three cooling fan FRUs with a minimum of two required for operation

It has the following performance capabilities:

- ▶ Up to 384 ports at full 16 Gbps speed or up to 512 8 Gbps Fibre Channel ports
- ▶ Up to 32 optical UltraScale ICL ports
- ▶ 8.2 Tbps total chassis bandwidth:
 - 6.1 Tbps Fibre Channel/FICON ports
 - 2.1 Tbps UltraScale ICL bandwidth
 - 512 Gbps bandwidth per slot

Figure 5-11 shows the front view of a fully populated SAN768B-2.




Figure 5-11 Front-side angle of SAN768B-2

For the latest information about SAN384B-2 and SAN768B-2, go to the following website:

<http://www.ibm.com/systems/networking/switches/san/b-type/san768b/index.html>

Blade support matrix

The IBM System Storage fabric backbone architectures support various blades and provide flexibility for different port density needs, multiprotocol capabilities, and fabric-based applications. Data center administrators can easily mix the blades to address specific business requirements and optimize cost/performance ratios.

Requirement: IBM Network Advisor Enterprise is required to manage the SAN768B-2 and SAN768B 8 slot directors.

The SAN768B-2 is a 12-slot chassis that consists of 8-port blades, two CP8 blades, and two CR8 blades, as shown in Figure 5-12.




Figure 5-12 SAN768B-2 front view

The SAN384B-2 is an 8-slot chassis that consists of 4-port blades, two CP8 blades, and two CR8 blades, as shown in Figure 5-13.




Figure 5-13 SAN384B-2 front view

The Gen 5 and Gen 4 directors have different blade compatibility, as described 5.2.5, “IBM Fabric backbone blades” on page 121.

Table 5-2 show the blade compatibility matrix for the directors at the time of writing.

Table 5-2 SAN768B-2 and SAN384B-2 blade compatibility

Blade	SAN768B-2	SAN384B-2	SAN768B (Gen4)	SAN384B (Gen4)
FC16-32	Yes	Yes	N/A	N/A
FC16-48	Yes	Yes	N/A	N/A
FC8-32E	Yes	Yes	N/A	N/A
FC8-48E	Yes	Yes	N/A	N/A
FC8-16	N/A	N/A	Yes	Yes
FC8-32	N/A	N/A	Yes	Yes
FC8-48	N/A	N/A	Yes	Yes
FC8-64	Yes	Yes	Yes	Yes
FC10-6	N/A	N/A	Yes	Yes
FR4-18i	N/A	N/A	Yes	Yes
FCOE10-24	Yes, only in FOS Version 7.3.0	No	Yes	Yes
FS8-18	Yes	Yes	Yes	Yes
FX8-24	Yes	Yes	Yes	Yes
CP8 Control Processor	Yes	Yes	Yes	Yes
CR16-8 Core Switching	Yes	N/A	Yes	N/A
CR16-4 Core Switching	N/A	Yes	N/A	Yes

5.2.5 IBM Fabric backbone blades

This section describes the Gen 5 fabric backbone blades and their features and technical details.

Control processor blade (CP8)

Supported by all Gen 4 and Gen 5 fabric backbones, the CP8 blades manage the overall functioning of the chassis. Each backbone director has two redundant control processors that are highly available and run FOS. The control processor functions are redundant active-passive (hot-standby). The blade with the active control processor is known as the “active control processor blade”, but either processor can be active or on standby. Additionally, on each processor there is a USB port and two network ports. The USB port is only for use with a USB storage device that is branded by Brocade. The dual IP ports allow a customer to potentially fail over internally on the same control processor without the loss of an IP connection, rather than fail over to the standby control.

Figure 5-14 shows the CP8 blade.




Figure 5-14 CP8 blade

Core routing blades (CR16-8 and CR16-4)

The Gen 5 fabric backbone includes two core routing blades for SAN384B-2 (CR16-4, which is shown in Figure 5-15 on page 123) or SAN768B-2 (CR16-8, which is shown in Figure 5-16 on page 123). These blades provide core switching and routing of the frames either from blade to blade or from the fabric backbone chassis through the ICL ports. The CR16-8 and CR16-4 blades work as an active-active cluster.

The CR16-8 has four Condor3 ASICs. The CR16-4 has two Condor3 ASICs. Each ASIC has dual connections to each ASIC group on each line card.

CR16-4 blades have these hardware features:

- ▶ Two Condor3 ASICs
- ▶ Eight Quad Small Form Factor Pluggable (QSFP) (ICL) ports
- ▶ Up to 1 Tbps backplane throughput

Figure 5-15 shows the CR16-4 core blade.




Figure 5-15 CR16-4 Core Blade

The CR16-8 blades have these hardware features:

- ▶ Three Condor3 ASICs
- ▶ Sixteen QSFP (ICL) ports
- ▶ Up to 2 Tbps backplane throughput

Figure 5-16 shows the CR16-8 core blade.




Figure 5-16 CR16-8 core blade

Note: The CR core blades are also responsible for routing ICL frames from backbone chassis. ICLs are described in 5.2.6, “Optical UltraScale Inter-Chassis Links” on page 129.

FC16-32 32-port 16 Gbps blade

The FC16-32 port blade is equipped with two Condor3 ASICs and can operate at a 16-Gbps full-line rate through the backplane or with local switching with no oversubscription (1:1).

It has the following hardware features:

- ▶ Two Condor3 ASICs
- ▶ 16x 16-Gbps ports per ASIC
- ▶ 512 Gbps backplane throughput
- ▶ No oversubscription at 16 Gbps (1:1)

Figure 5-17 shows the FC16-32 port blade.




Figure 5-17 FC16-32 port blade

FC16-48 48-port 16 Gbps blade

The FC16-48 port blade is equipped with two Condor3 ASICs, with one ASIC for each 24-port group. Although the backplane connectivity of this blade is identical to the FC16-32 blade, the FC16-48 blade uses 24 user-facing ports per ASIC rather than 16.

Oversubscription occurs only when the first 32 ports are fully used (16 Gbps) with no local switching.

It has the following hardware features:

- ▶ Two Condor3 ASICs
- ▶ 24x 16-Gbps ports per ASIC
- ▶ 512 Gbps backplane throughput
- ▶ 1.5:1 oversubscription at 16 Gbps

Figure 5-18 shows the FC16-48 port blade.




Figure 5-18 FC16-48 port blade

FC8-64 64-port 8 Gbps blade

Equipped with four Condor2 ASICs, the FC8-64 port blade offers 64x 8-Gbps ports (16 ports per ASIC) and provides a 2:1 oversubscription ratio at 8 Gbps switching through the backplane and no oversubscription with local switching.

It has the following key features:

- ▶ Four Condor2 ASICs
- ▶ 16x 8-Gbps ports per ASIC
- ▶ 256 Gbps backplane throughput
- ▶ 64 front-end ports and 32 back-end ports
- ▶ 2:1 oversubscription, and no oversubscription with local switching

Figure 5-19 shows the FC8-64 port blade.




Figure 5-19 FC8-64 port blade

FC8-32E 32-port Enhanced 8 Gbps blade

Equipped with two Condor3 ASICs, the FC8-32E port blade offers 328 Gbps front-end ports and 32 back-end ports and no oversubscription at 8 Gbps switching through the backplane. FC8-32E supports E, F, M, and EX Fibre Channel ports, and can operate at 2, 4, and 8 Gbps.

Figure 5-20 shows the FC8-32E port blade.




Figure 5-20 FC8-32E port blade

FC8-48E 48-port Enhanced 8 Gbps blade

Equipped with two Condor3 ASICs, the FC8-48E port blade offers 48x 8-Gbps front-end ports and 32 back-end ports and no oversubscription at 8 Gbps switching through the backplane. FC8-48E supports E, F, M, and EX Fibre Channel ports, and can operate at 2, 4, and 8 Gbps.

Figure 5-21 shows the FC8-48E port blade.




Figure 5-21 FC8-48E port blade

FX8-24 Extension Blade

The IBM FX8-24 Extension Blade accelerates and optimizes replication, backup, and migration over any distance. The twelve 8 Gbps Fibre Channel ports, ten 1 GbE ports, and up to two optional 10 GbE ports provide Fibre Channel and FCIP bandwidth, port density, and throughput for maximum application performance over IP wide area network (WAN) links.

Figure 5-22 shows the FX8-24 Extension Blade.




Figure 5-22 FX8-24 Extension Blade

5.2.6 Optical UltraScale Inter-Chassis Links

ICLs are high-performance ports for interconnecting multiple IBM System Storage fabric backbone chassis, enabling industry-leading scalability while preserving ports for server and storage connections. ICLs are designed to maximize SAN performance and scalability, minimize latency between chassis, and maximize load balancing and availability while simplifying network topologies.

UltraScale ICLs are based on optical QSFP. Each QSFP port combines four 16 Gbps links, providing up to 64 Gbps of throughput within a single cable.

Figure 5-23 shows an optical ICL cable and QSFP.



Figure 5-23 UltraScale ICSL cable and QSFP

UltraScale ICLs have these highlights and features:

- ▶ 64 Gbps (4x16 Gbps) bandwidth per each UltraScale ICL port
 - 32 UltraScale ICL ports per SAN768B-2 chassis
 - 16 UltraScale ICL ports per SAN384B-2 chassis
- ▶ Up to 2 Tbps UltraScale ICL bandwidth (four times than what exists for Gen 4)
- ▶ Up to 100 m OM4 optical cables
- ▶ Lowest-latency switching through the backplane versus ISLs
- ▶ Reduces the number of ISL cables that are required (a four to one reduction compared to traditional ISLs)
- ▶ Up to 33% more FC ports available for server and storage connectivity⁵
- ▶ Up to a 9-chassis full-mesh design with only a single hop between any two points within the fabric

Figure 5-24 shows the SANB384B-2 ICL on core blades.




Figure 5-24 SAN384B-2 ICL core blades

⁵ QSFP-based UltraScale ICL connections are on the core routing blades instead of consuming traditional ports on the port blades.

Figure 5-25 shows the ICL core blades.




Figure 5-25 CR16-8 and CR16-4 core blades

UltraScale ICL licensing

An ICL PoD license can be applied to the SAN768B-2 and SAN384B-2. Descriptions of the applicable licensing for the IBM System Storage fabric backbones running FOS V7.x are noted below.

Note: The ICL copper-based licensing that is present on the Gen 4 platforms is different from the Gen 5, and the following information does not apply to the Gen 4 fabric backbone directors.

ICL POD license: SAN768B-2 with Gen 5 Fibre Channel

One ICL PoD license on the SAN768B-2 enables the first 16 QSFP UltraScale ICL ports (enabling ICL ports 0 - 7 on each core blade). This is equivalent to 16× 64 Gbps, or 1 Tbps of bandwidth.

Two ICL PoD licenses enable the remaining 16 QSFP UltraScale ICL ports (enabling ICL ports 8 - 15 on each core blade), so all 32 QSFP ports across both core routing blades are enabled.

Figure 5-26 shows the CR16-8 core blade and ICL ports.




Figure 5-26 CR16-8 core blade

ICL PoD license: SAN384B-2 with Gen 5 Fibre Channel

Only one ICL PoD license is required to enable all 16 QSFP UltraScale ICL ports that are available on the *two* core blades of the SAN384B-2. This is equivalent to 16x 64 Gbps, or 1 Tbps of bandwidth.

Enterprise ICL license: SAN768B-2 and SAN384B-2 with Gen 5 Fibre Channel

The Enterprise ICL (EICL) license is required on each IBM System Storage fabric backbone chassis that connects to four or more chassis through UltraScale ICLs. This license requirement does not depend upon the total number of chassis that exist in a fabric, but only on how many chassis are directly connected through ICLs. This license is in addition to the ICL PoD license requirements, which enable the actual ICL ports.

Figure 5-27 shows examples where the EICL license is not necessary.

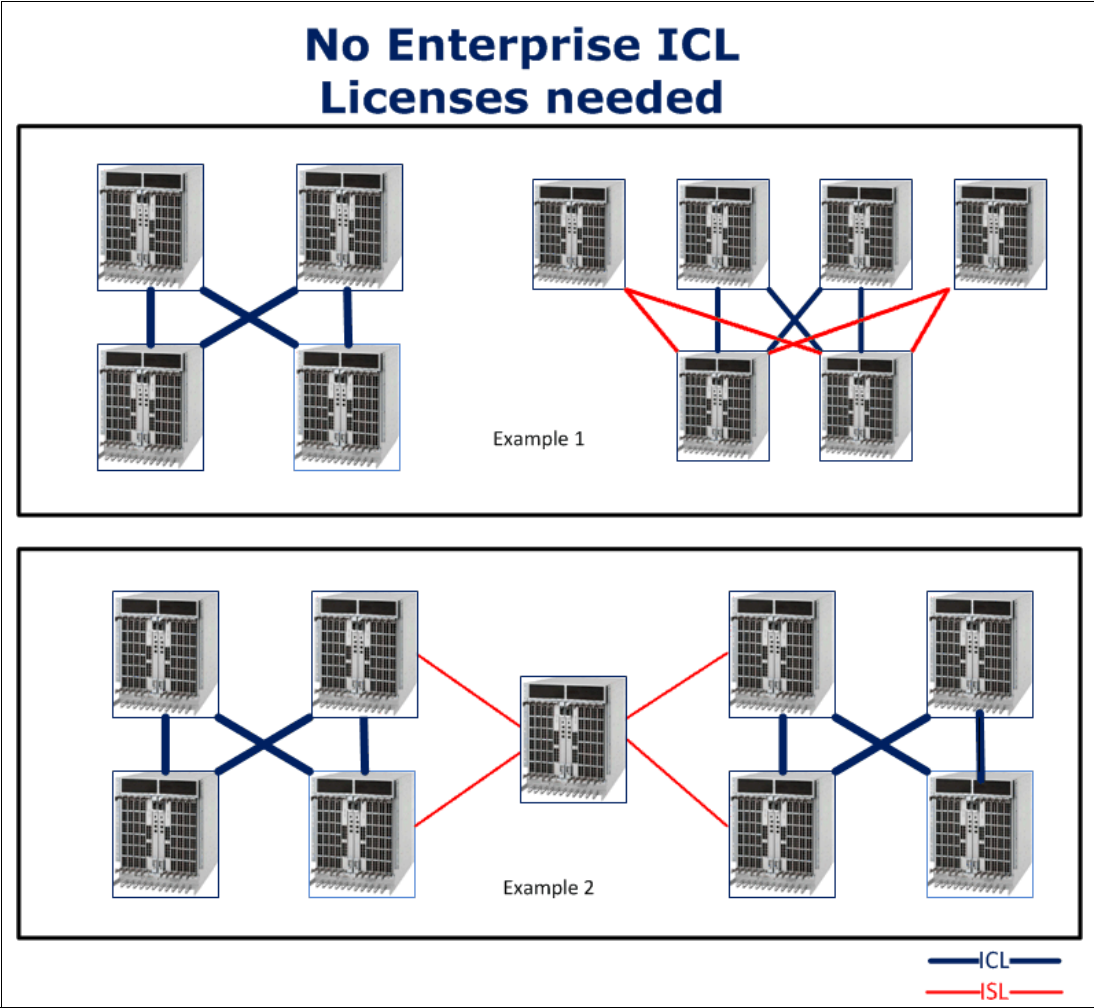


Figure 5-27 ICL and ISL examples

Figure 5-28 shows an EICL license usage example.

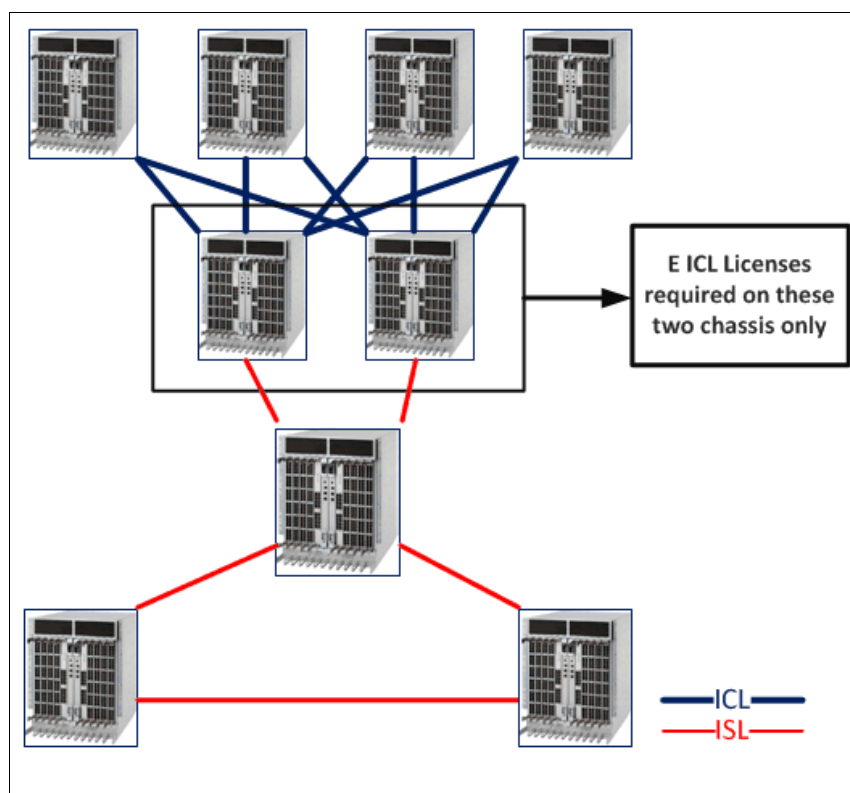


Figure 5-28 EICL license example

UltraScale ICL connections

To connect multiple IBM System Storage fabric backbone chassis through optical ICLs, a minimum of four ICL ports (two on each core blade) must be connected between each chassis pair. Figure 5-29 shows a diagram of the minimum connectivity between a pair of IBM System Storage fabric backbone chassis.




Figure 5-29 ICL connection

Note: If more than four ICL connections are required between a pair of IBM System Storage fabric backbones, extra ICL connections should be added in pairs (one on each core blade).

SAN768B-2 has 32 ICL ports available with both ICL PoD licenses installed. This configuration supports ICL connectivity to up to eight other chassis and at least 256 Gbps of bandwidth to each connected SAN768B-2.

SAN384B-2 has 16 ICL ports available, which supports up to four chassis and at least 128 Gbps of bandwidth to each connected SAN384B-2.

With 32 ICL ports available on the SAN768B-2 (with both ICL PoD licenses installed) and 16 ICL ports on SAN384B-2, these configurations support ICL connectivity to up to eight other chassis and at least 256 Gbps of bandwidth to each connected SAN768B-2.

A maximum of 16 UltraScale ICL connections or ICL trunk groups between any pair of IBM System Storage fabric backbone chassis is supported, unless they are deployed by using Virtual Fabrics, where a maximum of 16 ICL connections or trunks can be assigned to a single Logical Switch. This limitation is because of the maximum supported number of connections for Fabric Shortest Path First (FSPF) routing. Effectively, there should never be more than 16 ICL connections or trunks between a pair of SAN768B-2 chassis, unless Virtual Fabrics is enabled, and the ICLs are assigned to two or more Logical Switches. The exception to this situation is if eight port trunks are created between a pair of SAN768B-2 chassis. Details about this configuration are described in “Ultrascale ICL trunking and trunk groups”.

Note: QSFP-based UltraScale ICLs and traditional ISLs are not concurrently supported between a single pair of IBM System Storage fabric backbone chassis. All inter-chassis connectivity between any pair of IBM System Storage fabric backbone chassis must be done by using either ISLs or UltraScale ICLs.

Ultrascale ICL trunking and trunk groups

Trunking involves taking multiple physical connections between a chassis or switch pair and forming a single “virtual” connection and aggregating the bandwidth for traffic to traverse. Multiple hardware-based trunking solutions are available, including the ISL Trunking for traditional ISLs, trunking for Integrated Routing (FCR connectivity), trunking for Access Gateway, and also trunking for UltraScale ICLs. This section describes the trunking capability that is used with the QSFP-based UltraScale ICL ports on the IBM System Storage fabric backbone. For ISL Trunking, see 2.3.2, “ISL Trunking” on page 20.

Each optical ICL port has four independent 16-Gbps links, each of which terminates on one of four ASICs on each SAN768B-2 core blade, or two ASICs on each SAN384B-2 core blade. Trunk groups can be formed by using any of the ports that comprise contiguous groups of eight links on each ASIC.

Figure 5-30 shows the core blade ICL ports trunk.




Figure 5-30 Core blade ICL port trunk

Because there are four separate links for each QSFP-based UltraScale ICL connection, each of these ICL port groups can create up to four trunks, with up to eight links in each trunk. A trunk can never be formed by links within the same QSFP ICL port. This limit is because each of the four links within the ICL port terminates on a different ASIC for the SAN768B-2 core blade, or on either different ASICs or different trunk groups within the same ASIC for the SAN384b-2 core blade. Thus, each of the four links from an individual ICL is always part of independent trunk groups.

Figure 5-31 shows how ICL trunks are grouped.




Figure 5-31 ICL Trunks

5.2.7 IBM System Storage SAN42B-R Extension Switch

The IBM SAN42B-R extension switch is a purpose-built solution capable of moving data across data centers in Fibre Channel and IBM FICON environments. It is designed to replicate and back up large amount of data over WAN in a secure and reliable fashion while minimizing operation and capital expenses. It relies on any type of inter-data-center WAN link to extend storage applications over distance. SAN42B-R utilizes the same FOS that supports the entire IBM b-type Fibre Channel family of products.

The IBM SAN42B-R is a 2U unit with 24 16-Gbps Fibre Channel ports and 16 FCIP Ethernet interfaces with 1, 10, or 40 GbE ports.

Figure 5-32 shows the IBM SAN42B-R extension switch.



Figure 5-32 IBM SAN42B-R

Management and administrative tasks on SAN24b-R can be completed through familiar management tools, including IBM Network Advisor, web tools, the SAN Health utility, and the command-line interface (CLI). Optional FICON Control Unit Port capabilities enable legacy management applications to support IBM FICON environments.

Fabric services on SAN42B-R include Simple Name Server (SNS), registered state change notification (RSCN), NTP, RADIUS, and Reliable Commit Service (RCS). Optional services include FCR, FICON CUP, and FICON Management Server.

Platform features

SAN42B-R switches provide these features:

- ▶ Support for 2, 4, 8, and 16 Gbps auto-sensing Fibre Channel ports. The trunking technology groups up to eight 16 Gbps ports to create high performance 128 Gbps ISL trunks between switches. There is no limit on the number of trunk groups.
- ▶ Supports 1 GbE, 10 GbE, and 40 GbE FCIP Ethernet interfaces.
- ▶ Offers protection against WAN link failures, with tunnel redundancy for lossless path failover and guaranteed in-order data delivery. This advanced trunking feature allows multiple network paths to be used simultaneously. During a network path failure, extension trunking retransmits the lost packets and maintains data integrity without disruption.
- ▶ Scalable to a full fabric architecture of up to 254 switches by using SFP and SFP+ transceivers. SWL, LWL, ELWL for Fibre Channel and short reach wavelength, long reach wavelength, and ELWL transceivers for Ethernet are supported.
- ▶ Single fabric supported by up to seven hops and multiprotocol routing fabric with up to 19 hops.
- ▶ Unicast, multicast, and broadcast traffic types are supported.
- ▶ Support for high-performance port blades running at 2, 4, 8, 10, or 16 Gbps, enabling flexible system configuration.
- ▶ Redundant and hot-swappable power supplies, and components that enable a high availability platform and enable nondisruptive software upgrades.

Universal ports that self-configure as E_Ports, F_Ports, EX_Ports, D_Port (diagnostic), and M_Ports (mirror ports) and self discovery based on switch type (U_Port), VE Port (FCIP and virtual E Port).

IBM TS7700 Grid resiliency and disaster recovery solution

If you are using IBM z Systems™ with an IBM TS7700 Grid, you can use the IBM SAN42B-R Extension Switch as part of an enhanced solution for resiliency and business continuity.

A TS7700 grid refers to two to six physically separate TS7700 clusters that are connected to the IP network that you supply. The clusters can be separated by more than 1,000 km. A grid can be used to form disaster recovery and high availability solutions. A disaster recovery solution is achieved when multiple clusters are geographically distant from one another. A high availability solution is achieved when multiple clusters are located close to one another.

The clusters in a TS7700 Grid can, but do not need to be, geographically dispersed. In a multiple-cluster grid configuration, two TS7700 clusters are often located within 100 km of one another, whereas the remaining clusters can be located more than 1,000 km away. This configuration provides both a highly available and redundant regional solution while also providing a remote disaster recovery solution outside of the region.

The IBM SAN42B-R features, including extension hot code load (HCL), extension trunking, WAN-optimized TCP (WO-TCP), IP security (IPsec), Adaptive Rate Limiting (ARL), and data compression, help optimize and enhance the connection. You can use Fabric Vision and IBM Network Advisor to manage and monitor the connection.

For more information about the SAN42B-R Extension Switch and the TS7700 Grid, see the following links:

- ▶ *Enhanced Resilient Solutions for Business Continuity*

<http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=WH&infotype=SA&htmlfid=TSW03352USEN&attachment=TSW03352USEN.PDF>

- ▶ TS7700 Grid information in the IBM Knowledge Center

http://www.ibm.com/support/knowledgecenter/STFS69_3.3.0/ts7740_virtualized_composite_library.html


More information

For more information about the IBM System Storage SAN42B-R Extension Switch, see the following links:

- ▶ *IBM System Storage SAN42B-R Extension Switch, TIPS1209*

- ▶ IBM System Storage SAN42B-R product page

<http://www.ibm.com/systems/storage/san/b-type/san42b-r>



B-type SAN monitoring and management with IBM Spectrum Control

This chapter describes the IBM Spectrum Control 5.2.9 b-type SAN monitoring and management capabilities. It provides guidance on how to integrate your b-type fabric with the tool itself and how to perform the most common monitoring and management activities with IBM Spectrum Control.

This chapter includes the following sections:

- ▶ Software prerequisites
- ▶ Interoperability matrixes for supported switches and directors
- ▶ Monitoring agents for switches and fabrics
- ▶ Adding switches and fabrics
- ▶ Testing connectivity for a switch
- ▶ Enabling the switch performance monitoring
- ▶ Viewing Switches and Fabrics details
- ▶ Zoning
- ▶ Alerting
- ▶ Performance monitoring

Note: IBM Spectrum Control was formerly known as IBM Tivoli Storage Productivity Center (versions 5.2.7 and earlier).

6.1 Software prerequisites

Note: This chapter does not cover the software installation for the products described or their initial configuration. Instead, it describes integration of b-type switch devices with IBM Spectrum Control, and the available functions and capabilities that are supported for this type of devices.

Before proceeding any further with the topics that are described in this chapter, ensure that you have the following software prerequisites:

- ▶ A running instance of IBM Network Advisor 12.4.X and Storage Management Initiative (SMI) Agent. Fabrics and switches should be already discovered and managed by IBM Network Advisor before the SMI Agent integration. See the links in 6.2, “Interoperability matrixes for supported switches and directors” on page 142 for detailed information about supported versions.
- ▶ A running instance of *IBM Spectrum Control 5.2.9*.

Important:

- ▶ IBM Network Advisor and IBM Spectrum Control cannot be installed on the same server.
- ▶ If you are planning to integrate your b-type switches without the SMI Agent implementation (not recommended), then IBM Network Advisor is not required.

For information about IBM Network installation and configuration, see Chapter 4, “IBM Network Advisor” on page 49.

For installation and configuration instructions for IBM Spectrum Control 5.2.8 and later, see the IBM Knowledge Center at the following link:

<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>

For information about installing IBM Tivoli Productivity Center 5.2.7 and earlier releases, see the following link:

<http://www.ibm.com/support/knowledgecenter/SSNE44/welcome>

6.2 Interoperability matrixes for supported switches and directors

As a part of the prerequisites, interoperability between the b-type switches and IBM Spectrum Control 5.2.9 needs to be checked before you add and manage the b-type devices with the tool. See the following link for the corresponding compatibility matrix:

<http://www.ibm.com/support/docview.wss?uid=swg27047050>

For overall hardware, products and platform, and Spectrum Control interoperability, see the following link:

<http://www.ibm.com/support/docview.wss?uid=swg21386446>

IBM Spectrum Control supports homogeneous fabrics for Brocade, Cisco, and QLogic vendors. Heterogeneous fabrics are not supported. Specifically, fabrics with switches from two (or more) vendors are not supported unless all switches of one of the vendors are in NPV (N_Port Virtualization) mode.

Access Gateway mode switches are supported in Spectrum Control 5.2.8 and higher, but with some limitations. Refer to the following link for details:

<http://www.ibm.com/support/docview.wss?uid=swg21971011>

Notes:

- ▶ FICON support for b-type switches is available in IBM Network Advisor integrated SMI agent. However, there is no differentiation of FICON from FC in IBM Spectrum Control.
- ▶ Spectrum Control 5.2.9 now displays the switch blade number, port number, and blade names just as they were defined in other tools, such as IBM Network Advisor.

6.2.1 Virtual Fabrics support

Spectrum Control 5.2.8 and later supports Virtual Fabrics, with the following limitations:

- ▶ If b-type switches are partitioned into virtual fabrics, the virtual fabrics and switches are displayed in the IBM Spectrum Control GUI, but the physical fabrics and physical switches are not.
- ▶ The virtual fabrics and switches are displayed as though they were physical fabrics and switches. Therefore, there is no indication on which virtual fabrics are actually on the same physical fabric, and you cannot see which virtual switches are on the same physical switch.
- ▶ You must use a CIM agent to collect performance metrics for b-type switches that are partitioned into virtual fabrics. The simple network management protocol (SNMP) agent will not collect these performance metrics.

6.3 Monitoring agents for switches and fabrics

Depending on the type of information you want to collect or the functions that you want to use on your b-type switches and fabrics, you can use these items as valid data sources communication methods to set up your switch devices in IBM Spectrum Control:

- ▶ SMI (CIM) Agent
- ▶ SNMP agent
- ▶ Storage Resource agents (IBM Tivoli Storage Productivity Center v5.2.7 and earlier)
- ▶ A combination of these agents

Table 6-1 summarizes the fabric and switches monitoring capabilities of IBM Spectrum Control 5.2.9 and the agents supported for each one of those functions.

Table 6-1 Spectrum Control 5.2.9 Monitoring functions versus supported agents

Function	Agents
Monitor performance	SMI Agent or SNMP agent
Collect information about switches and switch ports	<i>Preferred:</i> SMI Agent Also supported: SNMP agent, Storage Resource agent (IBM Tivoli Storage Productivity Center v5.2.7 and earlier only)
Collect information about topology connectivity	<i>Preferred:</i> SMI Agent Also supported: SNMP agent, Storage Resource agent (IBM Tivoli Storage Productivity Center v5.2.7 and earlier only)
Collect information about zoning information and zone control	<i>Preferred:</i> SMI Agent
Generate alerts	<i>Preferred:</i> SMI Agent Also supported: SNMP agent
Collect information about hosts, endpoint devices, and device-centric and host-centric information	Storage Resource agent (IBM Tivoli Storage Productivity Center v5.2.7 and earlier only)

The following agents are available:

- ▶ The Storage Resource Agent (in-band - IBM Tivoli Storage Productivity Center 5.2.7 and earlier only) is deployed on a per server basis. It provides in-band fabric and switch self-discovery by using the server's host bus adapter (HBA) connections to the storage network. When the agent is deployed, it will autopopulate the Switches and Fabrics views of Tivoli Storage Productivity Center. A schedule can be then defined to retrieve information that is related to the switch configuration changes. However, it will not allow you to gather performance statistics, zoning information or to generate alerts. In addition to these limitations, if multiple fabrics are present, at least one server connection on each fabric is required to manage them.
- ▶ The SNMP Agent v1/v3 (out-of-band) allows the user to add and manage b-type switches in IBM Spectrum Control by using the SNMP v1/v3 protocol. However, it will not collect any zoning information or information that is related to hosts and endpoint devices. Both configuration probe and performance data gathering are supported when you use the SNMP protocol to communicate with the switches. To model your fabric environment with this method, you need to manually add each switch to IBM Spectrum Control. Protocol versions v1 and v3 are supported. However, v3 is the typical use and v1 is included only for existing support purposes.
- ▶ The SMI Agent (out-of-band) is deployed by default along with the IBM Network Advisor installation. It is also the preferred communication method setup between IBM Spectrum Control and your b-type switches. Both configuration probe and performance data gathering are supported by the SMI Agent. This agent is the most complete connectivity method in terms of fabric monitoring and management delivered capabilities. The configuration within Spectrum Control is as easy as inputting the SMI Agent's connection parameters (IP address, user name, and password). When it is deployed, Spectrum Control will communicate with all of your IBM Network Advisor managed fabrics by using the SMI Agent.

Note: For IBM Spectrum Control 5.2.8 and later, Storage Resource agents cannot be used to manage switches and fabrics.

6.3.1 Monitoring and managing fabrics with the Storage Resource Agent (IBM Tivoli Storage Productivity Center 5.2.7 and earlier only)

The Storage Resource Agent is deployed by default during the Tivoli Storage Productivity Center server installation. The agent provides host-based information to IBM Tivoli Storage Productivity Center. In the case of IBM Tivoli Storage Productivity Center server counts with HBA connections to the SAN, the agent can gather detailed topology information for the entire fabric using in-band connection through the HBA. To gather host-level and detailed HBA information for hosts other than the IBM Tivoli Storage Productivity Center server, the agent must be installed on each host where that information is wanted.

The Storage Resource agent provides the following functions:

- ▶ Gathers information about the SAN by querying switches and devices for topology and identification information.
- ▶ Gathers host-level information for the local system.
- ▶ Gathers information about the zoning of the fabric.
- ▶ Gathers information about the HBAs installed on the local system, including make, model, and driver versions.
- ▶ Gathers event information that is detected by the HBAs and forwards it to the Device server.

Storage Resource agents can complete the following tasks:

- ▶ Gather host-level and HBA information about the host that contains the Storage Resource agent.
- ▶ Provide detailed identification for devices that are in the same zone as the Storage Resource agent. The Storage Resource agent identifies the device by worldwide name (WWN) if the Storage Resource agent is not in the same zone as the device rather than being able to uniquely identify the device by device type (for example, host, and subsystem).
- ▶ Identify endpoint devices that are based on the endpoint device that is responding to Request Node Identification Data (RNID) or other queries.
- ▶ Gather a subset of the switch attributes. Some switch attributes can be collected only by using out-of-band Fabric agents.
- ▶ Tivoli Storage Productivity Center does not collect information about aliases through the Storage Resource agents.

When you deploy Storage Resource agents to hosts, the agents automatically run the discovery process and provide information about the fabrics or switches to which the hosts are attached. For example, if you have b-type switches and have Storage Resource agents deployed on the computers that are attached to these switches, the switches are automatically discovered when you deploy the Storage Resource agents.

Discovering fabrics with the Storage Resource agent provides the following information:

- ▶ Fabric WWN, information about fabric-to-switch relationships, and key attributes
- ▶ Host and device information (collected from Storage Resource agents only), including HBA information
- ▶ Basic information to identify the host and devices in the fabric

After a discovery is run, you can run probes to collect topology and zoning information for the fabrics. For b-type switches, use SMI Agents to collect all fabric data and enable performance monitoring, and use Storage Resource agents to collect information about HBAs.

For information about how to add switches and fabrics with the Storage Resource Agent, see 6.4, “Adding switches and fabrics” on page 148.

6.3.2 Monitoring and managing fabrics with the SNMP Agent (out-of-band)

IBM Spectrum Control uses SNMP (v1 or v3) queries to discover information about the switches and fabrics. Management Information Base (MIB) information is collected from the switches by the out-of-band fabric agent. Switches and directors are added as out-of-band agents and contacted from the IBM Spectrum Control Device server by SNMP.

The SNMP fabric agent provides the following functions:

- ▶ Gathers information about the fabric by querying the switch or director for topology information.
- ▶ Gathers virtual SAN information for Cisco switches.

Note the following limitations:

- ▶ Topology information is only gathered for the switch that was added as an out-of-band fabric agent. The agent cannot gather the topology information for any other switches that are connected to it unless they are also defined as out-of-band fabric agents. If you are doing out-of-band discovery on a fabric with several switches, you must install an out-of-band fabric agent for each switch in the fabric to discover the whole fabric.
- ▶ Device information is limited. Most devices are unknown with a type equal to “Other” and identified by their WWN.
- ▶ A working Ethernet connection must exist between the switch and the IBM Spectrum Control server.
- ▶ To enable events from the switch to the IBM Spectrum Control server, the switch must be configured to send SNMP traps to it.

For information about how to add switches and fabrics by using the out-of-band (SNMP) agent, see 6.4, “Adding switches and fabrics” on page 148.

6.3.3 Monitoring and managing fabrics with the Brocade SMI Agent

The Brocade SMI Agent is deployed by default with any IBM Network Advisor Professional Plus or IBM Network Advisor Enterprise edition installation. If you do not want to install IBM Network Advisor, select the option during the installation process to deploy only the SMI agent (“SMI Agent Only”). This option does not require a license, and deploys only the integrated SMI agent without the IBM Network Advisor management interface.

When you add fabrics and switches for monitoring, you can specify the SIM agent that manages them. IBM Spectrum Control connects to the SMI agent and automatically discovers the fabrics and switches that it manages. You can then add the discovered fabrics and switches, and configure data collection.

SMI agents provide the following benefits for fabric management:

- ▶ Probes do not occur in the data path.
- ▶ Propagation of alerts for real-time SAN events is not degraded by using fabric agents.

For information about how to add switches and fabrics by using the SMI Agent, see 6.4, “Adding switches and fabrics” on page 148.

6.3.4 Monitoring and managing fabrics with multiple agents

Sometimes it is not possible to deploy a single agent to cover all of the monitoring and management needs, or there is a specific need of redundancy to losing communication with the switch devices. For such cases, IBM Spectrum Control allows multiple connection methods to be defined for your switches and fabrics.

As mentioned earlier, using multiple agent types allows the redundant collection of information if one type of agent fails. Additionally, some types of agents provide certain features and information that the other agents do not. However, using multiple agent types can increase network traffic and the management load on the switches.

As an example, consider a set of switches monitored only by the Storage Resource Agent (SRA) that is deployed on a server that is attached by Fibre Channel to your fabric. The SRA, which uses in-band communication with the fabric, does not collect performance data. You can add this functionality by adding the switches by using an SNMP or SMI Agent.

Figure 6-1 shows the IBM Spectrum Control confirmation window when you add a switch by using SNMP. This switch was already managed by IBM Tivoli Storage Productivity Center by using the SRA. The request is successfully completed but a message is displayed to indicate that the switch was already managed. However, as a result of this new capability, you can enable Performance Monitoring for the switch.

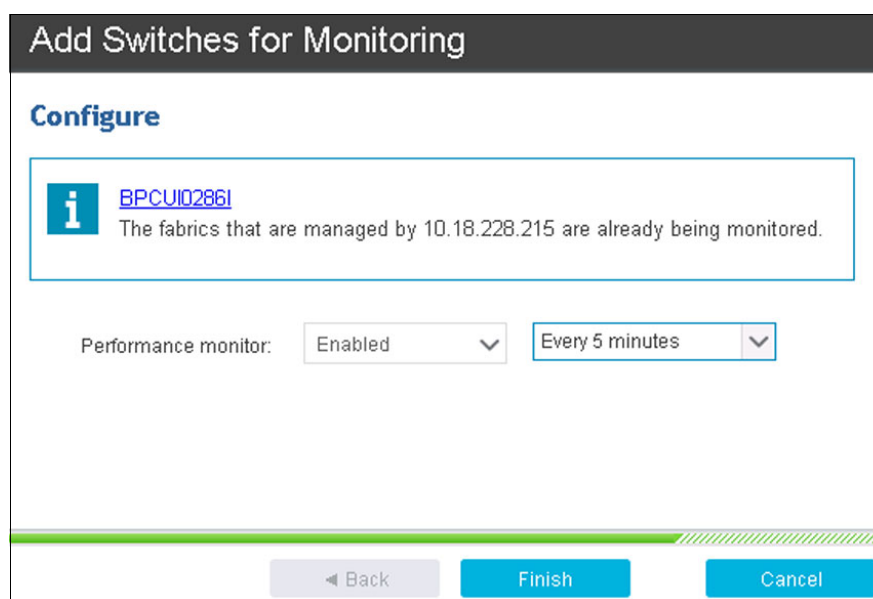


Figure 6-1 Adding a switch to IBM Virtual Storage Center by using SNMP

Figure 6-2 shows the confirmation of the Performance Monitoring data collection activation on the switch.

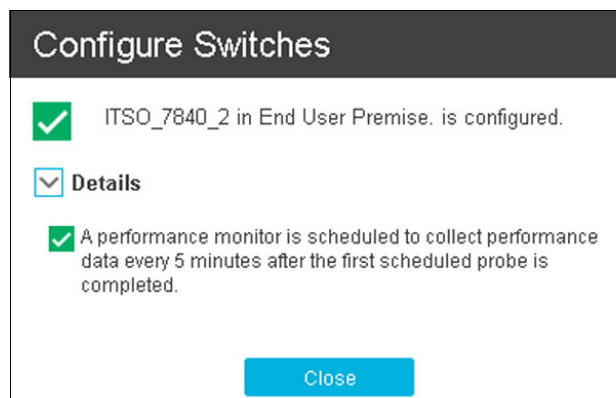


Figure 6-2 Switch configuration by using the SNMP confirmation window

After the additional Performance Monitoring capabilities are added, enable the monitor as described in 6.6, “Enabling the switch performance monitoring” on page 160.

6.4 Adding switches and fabrics

IBM Spectrum Control can discover devices in the SAN and collect data about the performance of those devices. This data collection can be performed either by using SNMP agents, SMI agents, or a combination of both. For versions 5.2.8 and later, the Storage Resource Agent is no longer available as connection method.

6.4.1 Device discovery

In IBM Spectrum Control, discovery is no longer available as a separate job as in earlier releases. Discovery now runs as part of adding a device or as part of probing a device.

For example, a switch that is attached to a server is discovered during a server probe. If you add a device to a SMI Agent, running the Add Device wizard again by using the SMI Agent will discover the new device. You will then be able to configure the device. Even if you do not run the Add Device wizard again in this situation, the next probe iteration will discover the new device. You will be able to see the new device in the list of switches and fabrics.

6.4.2 Device probes

In IBM Spectrum Control, a probe process is scheduled for a switch device, not for a data source (SNMP or SMI) or for a Monitoring Group.

When you start or schedule a probe for a switch, IBM Spectrum Control will probe the complete fabric, not just that switch. In this sense, the entire fabric can be thought of as a device. Thinking of the complete fabric makes sense because you need the information from the complete fabric to show any data paths.

Note: If multiple data sources are available, IBM Spectrum Control selects the best for a particular switch. For example, a SMI agent switch data source will always be used as preferred communication method, even if an SNMP agent data source is available.

6.4.3 Adding switches by using SNMP Agent (out-of-band)

If you want to add the switch to IBM Spectrum Control by using SNMP (for example, because you do not have a CIM Agent that is managing this switch), complete the following steps:

1. From the Spectrum Control top menu bar, select **Network** and then **Switches**. Then click **Add Switch**.
2. Select **Brocade** when asked for the switch vendor, as shown in Figure 6-3.



Figure 6-3 Switch vendor selection

3. Select **Monitor without Network Advisor** as shown in Figure 6-4 and click **Next**.

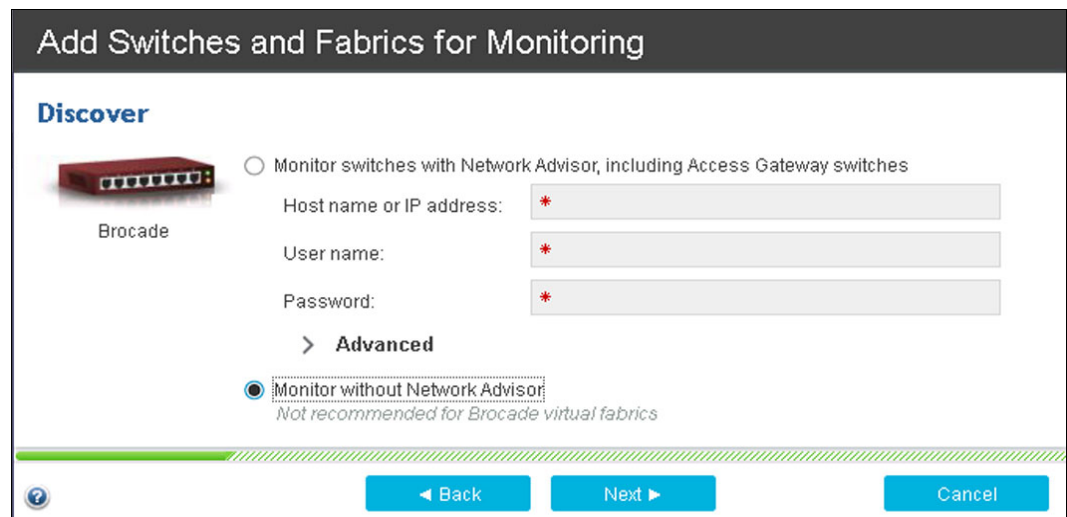


Figure 6-4 Selecting the switch discovery method

4. Select which SNMP protocol version to use (**v1** or **v3**) and complete the information for the SNMP data source as shown in Figure 6-5, then click **Next** to continue.

Figure 6-5 SNMP data sources configuration

5. IBM Spectrum Control will now verify that it can communicate with the switch by using SNMP as well as discover which devices are available through this data source.

Note: If you select SNMPv1 as the communication protocol and the discovery process fails, ensure that these conditions are true:

- ▶ The switch has SNMPv1 enabled,
- ▶ The SNMP v1 community that is being used is correct.
- ▶ If SNMP v1 control access is enabled, that the IP of the Spectrum Control server is added to the list of addresses that are allowed to communicate by SNMPv1 with the switch.

6. After the switches are discovered, the wizard will ask to set up the configuration probe and performance data schedules.
7. The wizard completes and a configuration probe is run to collect information about the switches.
8. Configure the SNMP trap notifications. SNMP traps are generated by the switch and directed to IBM Spectrum Control as an indication that something in the fabric changed and that a discovery must occur to identify the changes. The default configuration for handling switch traps is to send them from the switch to port 162 on the IBM Spectrum Control system. To successfully generate and receive traps, there are some configuration requirements:
 - The trap destination parameter on the switch must be set. This parameter is the host that receives the trap and sends it to IBM Spectrum Control. The parameter is set on the switch.
 - The destination port parameter on the switch must be set. IBM Spectrum Control listens on port 162 by default. This parameter is set on the host.

- The traps must be sent as SNMPv1. This parameter is set on the switch.
 - The trap severity level must be set to generate traps for change conditions. This level typically is set to send error level traps and anything more severe. This parameter is set in IBM Spectrum Control.
9. To complete the switch setup, enable the device's default or custom alerts so you are automatically notified when certain conditions are detected. For more information, see 6.9, "Alerting" on page 175.

6.4.4 Adding switches with the SMI Agent

If the switch that is being added to IBM Spectrum Control is available through a SMI Agent, complete the following steps:

1. From the Spectrum Control top menu bar, select **Network** → **Switches** → **Add Switch**.
2. Select **Brocade** when asked for the switch vendor, as shown in Figure 6-6.



Figure 6-6 Switch vendor selection

3. Select **Monitor with Network Advisor** as shown in Figure 6-7 and complete the connection information. Depending on the SMI agent connection parameters that you entered during the IBM Network Advisor installation, you might need to change the default communication protocol and port.

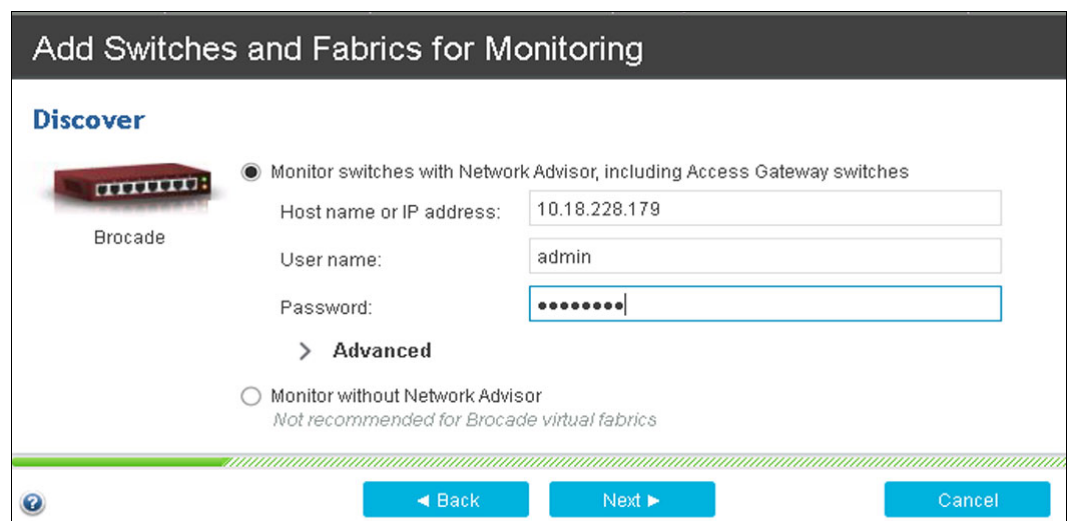


Figure 6-7 SMI agent connection information

Click **Next**.

4. A connection will be established with the remote SMI agent to query about its managed fabrics and devices. When this polling is complete, a new configuration window is displayed that lists the switches that are managed by the SMI Agent, as shown in Figure 6-8.

Name	Location	Host name or IP address
ITSO_DCX8510-4	Not Specified	10.18.228.35
ITSO_DCX8510-8	Not Specified	10.18.228.106

Figure 6-8 Review the discovered switches information

If wanted, make the corrections on the discovered switches' name and location fields and then click **Next**.

5. When the switches are discovered, the wizard will ask to set up the configuration probe and performance data schedules.
6. The wizard completes and a configuration probe is run to collect information about the switches.
7. If you want to use SNMP for data source redundancy, add the discovered switches by using SNMP as described in 6.4.3, "Adding switches by using SNMP Agent (out-of-band)" on page 149.
8. Configure the SNMP trap notifications. SNMP traps are generated by the switch and directed to IBM Spectrum Control as an indication that something in the fabric changed and that a discovery must occur to identify the changes. The default configuration for handling switch traps is to send them from the switch to port 162 on the IBM Spectrum Control system. To successfully generate and receive traps, there are some configuration requirements:
 - The trap destination parameter on the switch must be set. This parameter is the host that receives the trap and sends it to IBM Spectrum Control. The parameter is set on the switch.
 - The destination port parameter on the switch must be set. IBM Spectrum Control listens on port 162 by default. This parameter is set on the host.
 - The traps must be sent as SNMPv1. This parameter is set on the switch.
 - The trap severity level must be set to generate traps for change conditions. This level typically is set to send error level traps and anything more severe. This parameter is set in IBM Spectrum Control.
9. To complete the switch setup, enable the device's default or custom alerts so you are automatically notified when certain conditions are detected. For more information, see 6.9, "Alerting" on page 175.

6.4.5 Adding switches by using Storage Resource Agent (IBM Tivoli Storage Productivity Center 5.2.7 and earlier)

To discover a new switch and its fabric with this method, the Storage Resource agent must be deployed on at least one of the servers that are connected to the SAN fabric to be discovered. Complete the following tasks to implement the Storage Resource agents.

Implementing Storage Resource agents

The Storage Resource Agent can be deployed either by using the Tivoli Storage Productivity Center GUI or by a comma-delimited file. This section guides you through the GUI deployment procedure.

For more information about all the available Storage Resource agent deployment methods or about the Storage Resource agent itself, see the Tivoli Storage Productivity Center documentation on the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.7/com.ibm.tpc_V527.doc/fqz0_t_installing_agents.html?lang=en

Adding and managing a Storage Resource agent is not a new function for the Tivoli Storage Productivity Center. The functionality was migrated from the stand-alone GUI to the web-based GUI.

Use the web-based GUI to add servers by deploying Storage Resource agents. This process enables full server monitoring, which gathers the following server information:

- ▶ Asset information
- ▶ File and file system attributes
- ▶ Database application information
- ▶ Network-attached storage (NAS) device information
- ▶ Topology information
- ▶ Information about zoning and the fabrics that are visible to the server

Note: You must have Administrator privileges to deploy an SRA.

Before you deploy a Storage Resource Agent on a server, you need to disable the firewall for inbound connections (just for deployment purposes) on both the server where the agent will be deployed and the Tivoli Storage Productivity Center server. Otherwise the process might fail.

Adding Storage Resource agents manually by using the GUI

This section shows an example of adding a Storage Resource Agent manually on a Windows System. To do so, complete these steps:

1. From the Tivoli Storage Productivity Center GUI left navigation pane, select **Server** → **Resources** → **Servers** as shown in Figure 6-9.

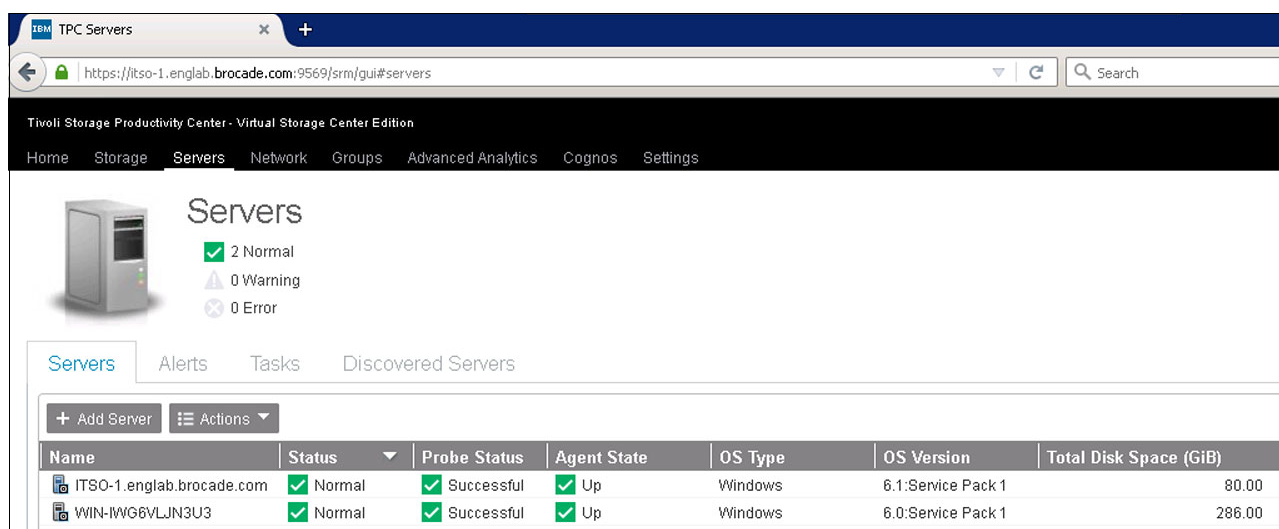


Figure 6-9 Server view main window

2. On the left of the top bar, click **Add Server**. The Add Server window appears as shown in Figure 6-10. Make sure that **Deploy an agent for full server monitoring** is selected, and select **Manually**.

Note: The **From Discovered Servers** option is not available for adding Storage Resource Agents.

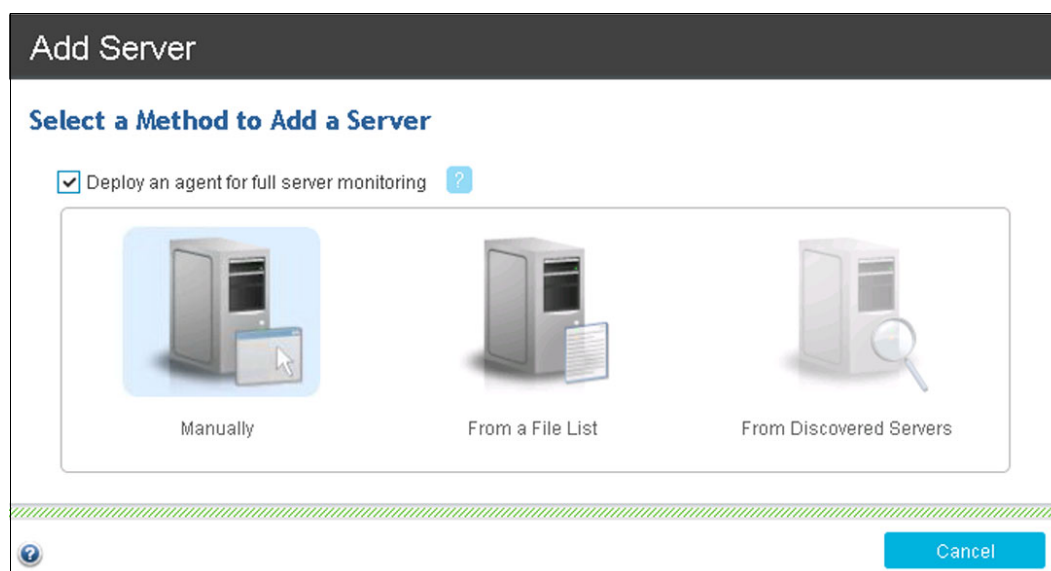
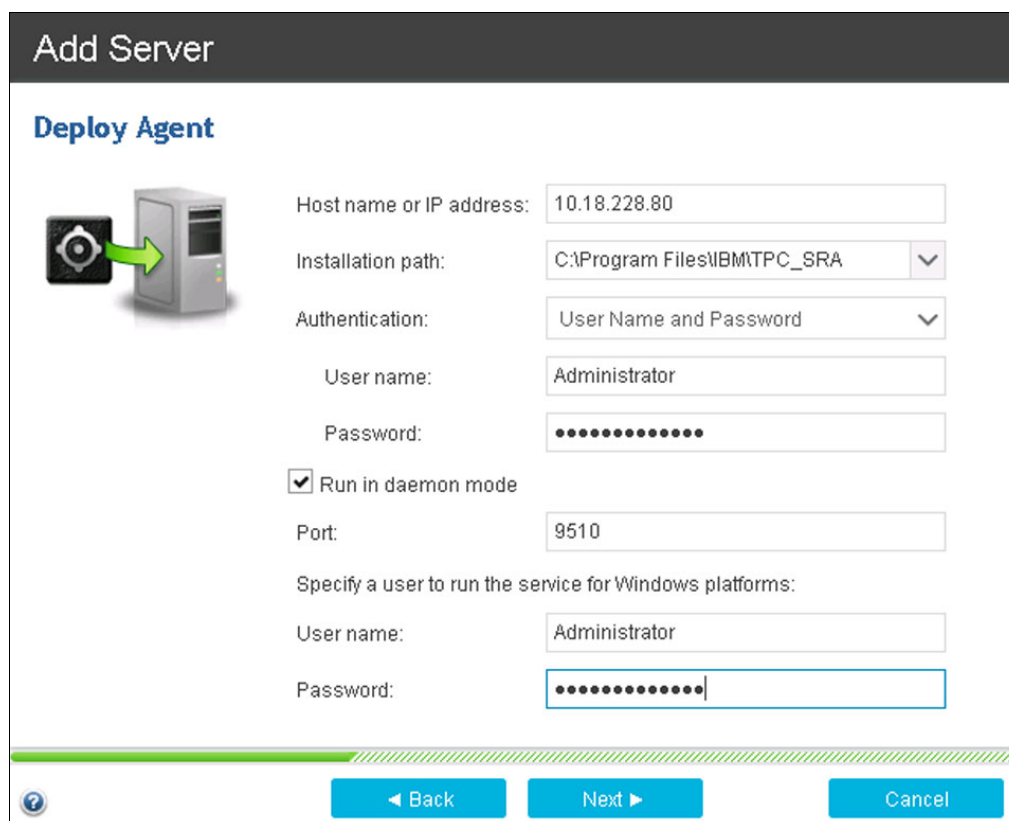


Figure 6-10 Add Server window

3. Create a Storage Resource Agent deployment job as shown in Figure 6-11. This example shows the simplest authentication method, which is User Name/Password. Secure Shell (SSH) is also available. You also need to provide the fully qualified installation path, for example "C:\Program Files\IBM\TPC_SRA". Additionally, Storage Resource agent will run in a non-daemon mode. You can also choose to overwrite previous installation agents to force installation of the Storage Resource agent.

Click **Next** to continue.



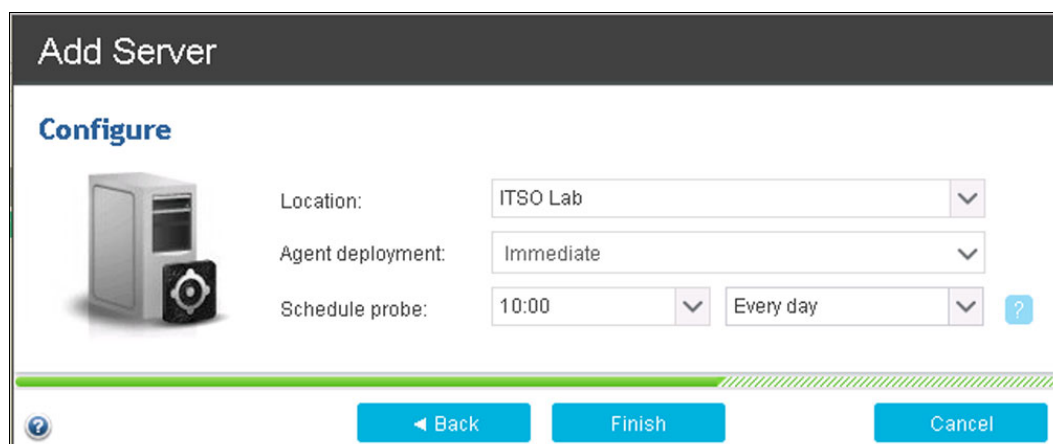
The 'Add Server' dialog window has a 'Deploy Agent' tab. On the left is an icon of a server with a green arrow pointing to it. The form contains the following fields and options:

- Host name or IP address: 10.18.228.80
- Installation path: C:\Program Files\IBM\TPC_SRA (dropdown)
- Authentication: User Name and Password (dropdown)
- User name: Administrator
- Password: (masked with dots)
- ☒ Run in daemon mode
- Port: 9510
- Specify a user to run the service for Windows platforms:
 - User name: Administrator
 - Password: (masked with dots)

At the bottom are buttons for '?', '< Back', 'Next >', and 'Cancel'.

Figure 6-11 Deploy agent details dialog window

4. Configure a schedule for the agent deployment job on Figure 6-12. In this example, after you click **Finish**, the deployment job runs immediately.



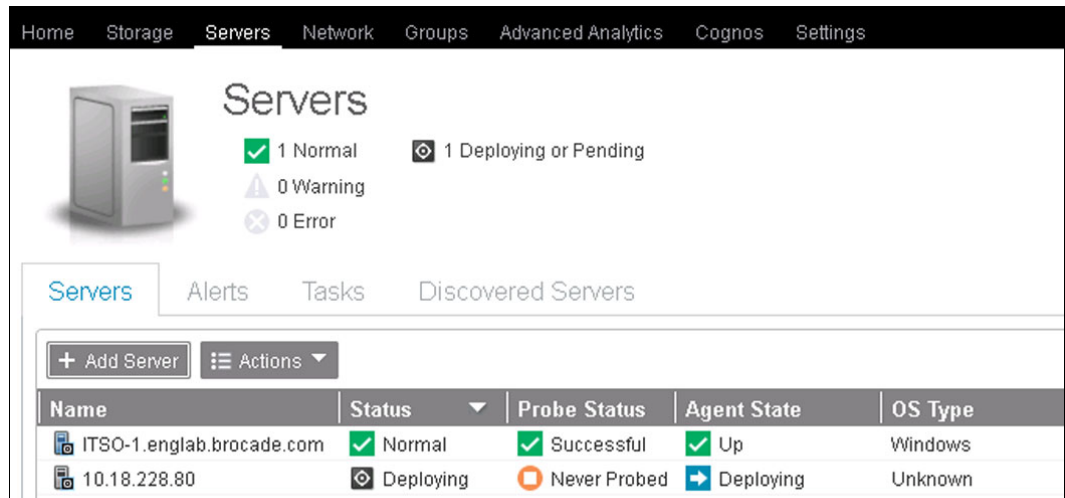
The 'Add Server' dialog window has a 'Configure' tab. On the left is an icon of a server. The form contains the following fields and options:

- Location: ITSO Lab (dropdown)
- Agent deployment: Immediate (dropdown)
- Schedule probe: 10:00 (dropdown) and Every day (dropdown)

At the bottom are buttons for '?', '< Back', 'Finish', and 'Cancel'.

Figure 6-12 Schedule agent deployment

- On the Servers panel, the agent deployment job is displayed as shown in Figure 6-13.



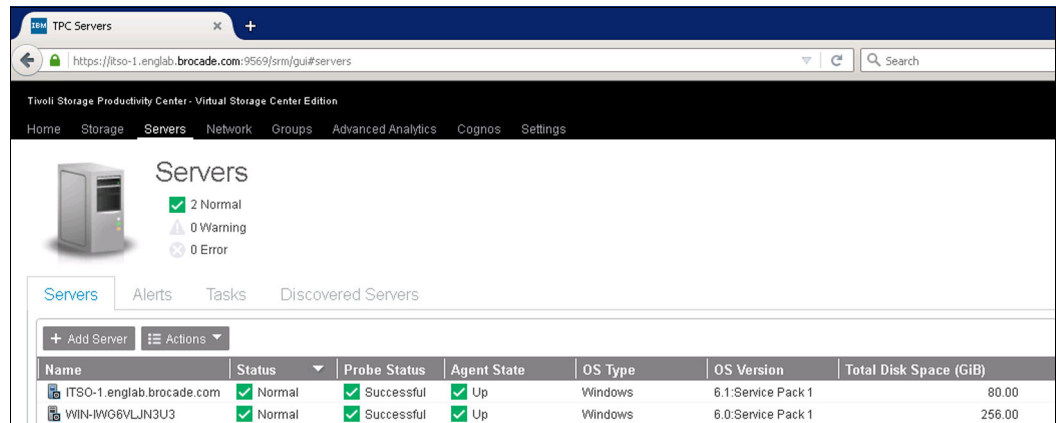
Servers

1 Normal, 0 Warning, 0 Error, 1 Deploying or Pending

Name	Status	Probe Status	Agent State	OS Type
ITSO-1.englab.brocade.com	Normal	Successful	Up	Windows
10.18.228.80	Deploying	Never Probed	Deploying	Unknown

Figure 6-13 Deploying the agent

When the deployment completes, the agent is ready to start collecting data as shown in Figure 6-14.



Servers

2 Normal, 0 Warning, 0 Error

Name	Status	Probe Status	Agent State	OS Type	OS Version	Total Disk Space (GiB)
ITSO-1.englab.brocade.com	Normal	Successful	Up	Windows	6.1 Service Pack 1	80.00
WIN-IWGBVLJN3U3	Normal	Successful	Up	Windows	6.0 Service Pack 1	256.00

Figure 6-14 Agent successfully deployed

Note: If there are any issues with the agent deployment, you can take action immediately by right-clicking the deployment and opening the following pane as shown in Figure 6-15.

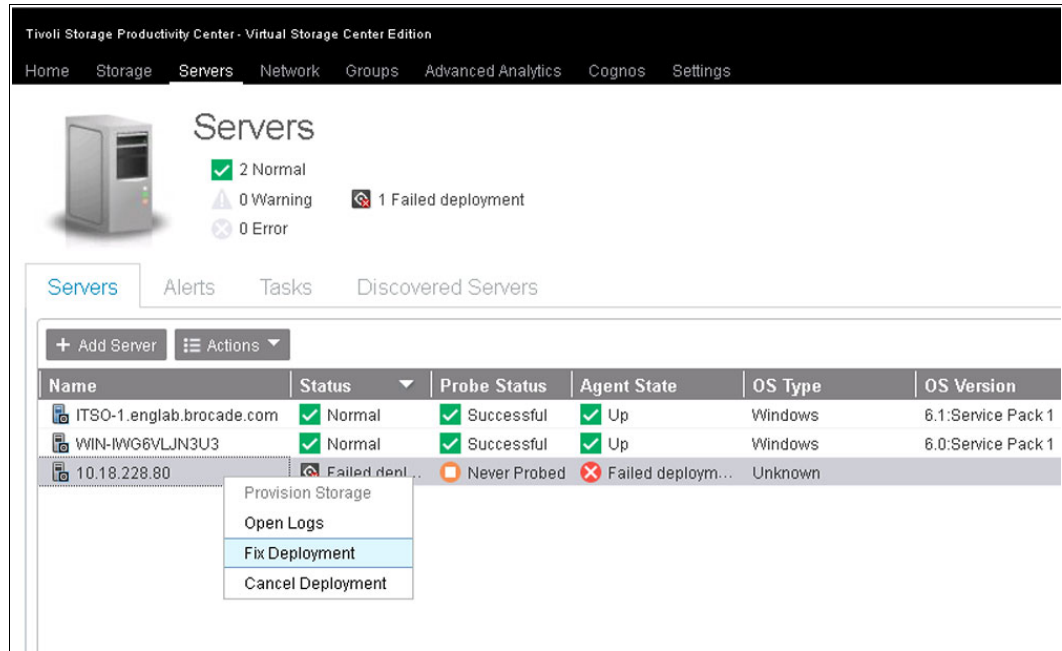


Figure 6-15 Failed agent deployment

Note: For Tivoli Storage Productivity Center V5.2, the Storage Resource agents that are installed with the Tivoli Storage Productivity Center server have the Fabric function disabled. This configuration is used because of potential fabric or switch firmware defects. In the past, it has been observed with products other than Tivoli Storage Productivity Center, that certain in-band fabric commands can cause a switch with a firmware defect to hang or reboot.

When Tivoli Storage Productivity Center does its initial probes, it might cause such an error in both fabrics at the same time, which can cause a large impact. Turning off the Fabric function prevents this scenario because you have the opportunity to schedule the probes so they do not run at the same time.

For more information about the Storage Resource agent deployment and limitations, see the following Tivoli Storage Productivity Center documentation link in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.7/com.ibm.tpc_v527.doc/fqz0_r_native_agents_deployment_considerations.html?cp=SSNE44_5.2.7&lang=en

Self-discovered switches and fabrics

During the Storage Resource agent's first probe run and in addition to the host-related server information, the agent will retrieve information pertaining the switches and fabrics where the server is connected.

When the Storage Resource Agent probe run is complete, the switches and fabrics will be automatically added to the Switches and Fabrics views of Tivoli Storage Productivity Center, as shown in Figure 6-16.

Switches

0 Normal
0 Warning
0 Error
4 Unknown

Switches Alerts Threshold Violations Performance

+ Add Switch Actions

Name	Status	Probe Status	Performance M...	Fabric	Principal Switch of Fabric
ITSO_7840_1	Unknown	Never Probed	Disabled	100000051E468A00	ITSO_DCX8510-8
ITSO_7840_2	Unknown	Never Probed	Disabled	100000053396F400	ITSO_DCX8510-4
ITSO_DCX8510-4	Unknown	Never Probed	Disabled	100000053396F400	ITSO_DCX8510-4
ITSO_DCX8510-8	Unknown	Never Probed	Disabled	100000051E468A00	ITSO_DCX8510-8

Figure 6-16 Storage Resource agent self discovered switches

To retrieve the complete information of your switches and fabrics, start each switch's configuration probe as shown in Figure 6-17.

Switches

5 Normal
0 Warning
0 Error

Switches Alerts Threshold Violations Performance

+ Add Switch Actions

Name	Status	Probe Status	Performance Monitor Status	IP Address	Fabric
ITSO_7800-1	Normal	Successful	Running	10.18.228.217	ITSO_7800-1
ITSO_7840_1	Normal	Successful	Running	10.18.228.216	ITSO_7800-1
ITSO_7840_2	Normal	Successful	Running	10.18.228.215	100050EB1A656D60
ITSO_DCX8510-4	Normal	Successful	Running	10.18.228.35	100050EB1A656D60
ITSO_DCX8510-8	Normal	Successful	Running	10.18.228.106	ITSO_7800-1

Context menu for ITSO_7840_2:

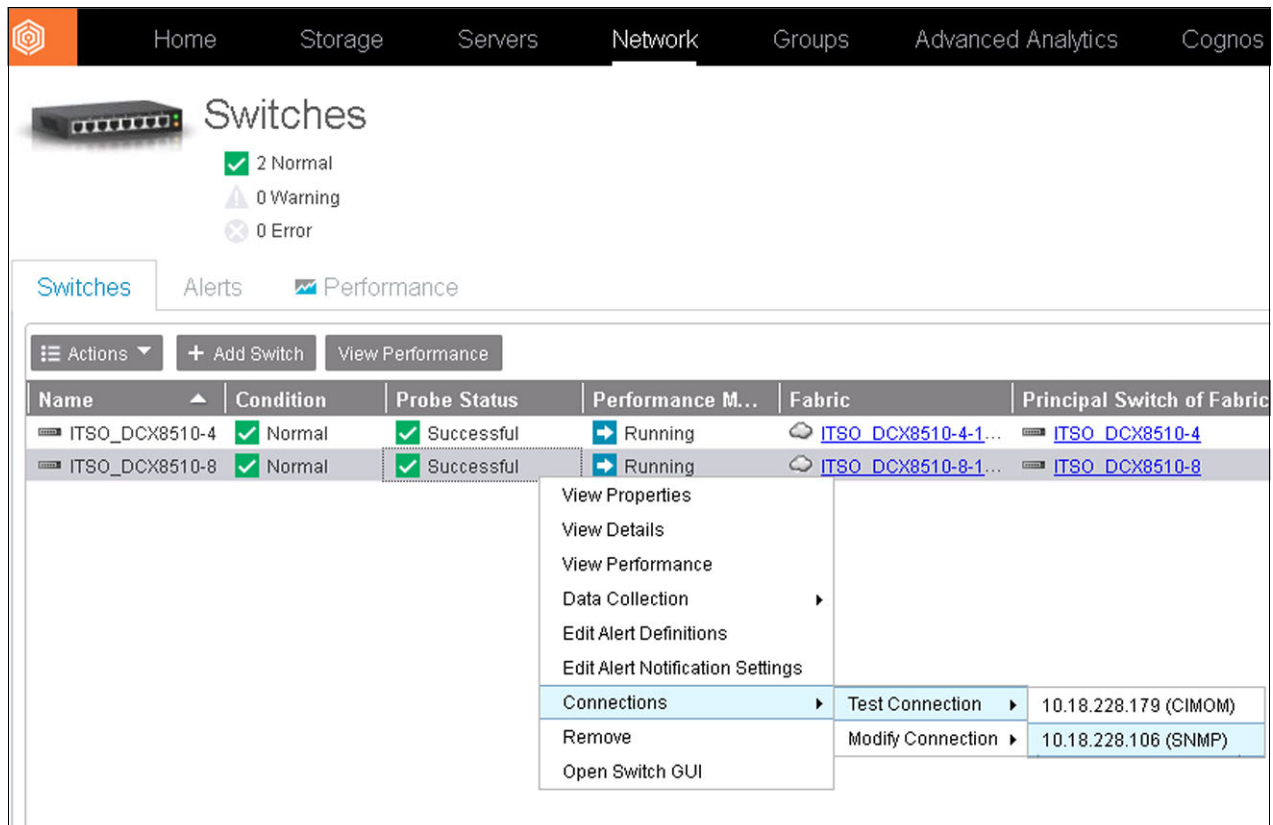
- View Properties
- View Details
- View Performance
- Data Collection
 - Start Probe
 - Open Probe Logs
 - Stop Performance Monitor
 - Open Performance Monitor Logs
 - Schedule
- Edit Alert Definitions
- Edit Alert Notification Settings
- Connections
- Remove
- Open Switch GUI

Figure 6-17 Start the switch configuration probe

6.5 Testing connectivity for a switch

To list and test the multiple connections for a switch, complete the following steps:

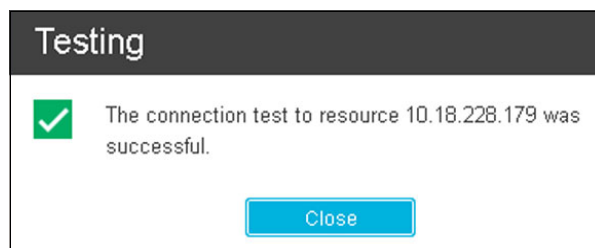
1. Select **Network** on the IBM Spectrum Control top menu, then select **Switches**.
2. In the Switches view, right-click the switch that you want, then select **Connections** and finally **Test Connection**. The multiple connections to that device will be listed as shown in Figure 6-18. To test any of these connection methods, click the one that you want.



Name	Condition	Probe Status	Performance M...	Fabric	Principal Switch of Fabric
ITSO_DCX8510-4	Normal	Successful	Running	ITSO_DCX8510-4-1...	ITSO_DCX8510-4
ITSO_DCX8510-8	Normal	Successful	Running	ITSO_DCX8510-8-1...	ITSO_DCX8510-8

Figure 6-18 Switch connections list and test

The connection test's results are displayed, as shown in Figure 6-19.



Testing

✓ The connection test to resource 10.18.228.179 was successful.

Close

Figure 6-19 Switch connection test results

6.6 Enabling the switch performance monitoring

If you did not configure the Performance Monitoring data collection at the time the switch device was added, you can do so at any time. To configure collection, follow the procedure that is described below, when the configuration data collection probe completes the first discovery run.

Note: Depending on the connection agents defined for the switch, Performance Monitoring might not be able to be enabled. Review Table 6-1 on page 144 to ensure that the available agents support this functionality.

To configure and enable Performance Monitor data collection, complete the following steps:

1. Select **Network** on the IBM Spectrum Control top menu, then select **Switches**.
2. In the Switches view, right-click the wanted switch, then select **Data Collection** and then **Schedule**. The data collection schedule configuration window is displayed as shown in Figure 6-20.

ITSO_7840_2

Data Collection Schedule

Probe: Enabled 01:30 Every day

[Open Logs](#)

Performance monitor: Disabled Every 5 minutes

Save Cancel

Figure 6-20 Data collection schedule configuration window

3. Click the drop down menu and enable the Performance monitor schedule, then select the sampling interval.

Note on collected samples: IBM Spectrum Control records performance statistics in a form of *samples*. Each sample is composed of the actual metric's value and a time stamp. The interval selection in the Performance monitor schedule represents the interval by which IBM Spectrum Control will generate a sample by averaging the performance data values that are gathered throughout that interval. The time stamp that is assigned to the sample represents the time that this averaging was done.

4. Click **Save**.

5. The result of the scheduling action appears, as shown in Figure 6-21.




Figure 6-21 Performance monitoring scheduling confirmation

6. In the Switches view, the new Performance Monitor Status for your switch should be Running, as shown in Figure 6-22.




Figure 6-22 Performance monitor status shown as running

6.6.1 Troubleshooting Performance Monitor data collection problems

Sometimes, the performance monitor data collection might be affected by problems that are related either with internal IBM Spectrum Control components, external switch components, or communication agents (SMI or SNMP) components. This situation will be reflected in the Performance Monitor Status under the IBM Spectrum Control **Switches** view.

Any status other than Running must be diagnosed. In order to access the Performance Monitor Logs, complete the following steps:

1. Select **Network** on the IBM Spectrum Control top menu, then select **Switches**.
2. In the Switches view, right-click the switch you are interested in, then select **Data Collection** and then **Open Performance Monitor Logs**, as shown in Figure 6-23.

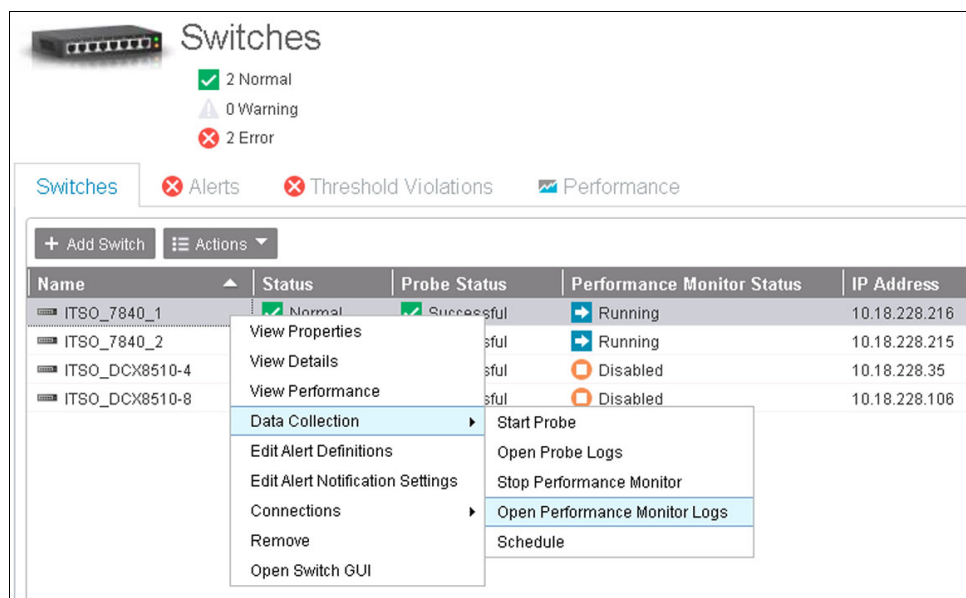


Figure 6-23 Opening performance monitor logs

3. A new window appears with all the performance monitor log entries for the selected switch device, as shown in Figure 6-24.

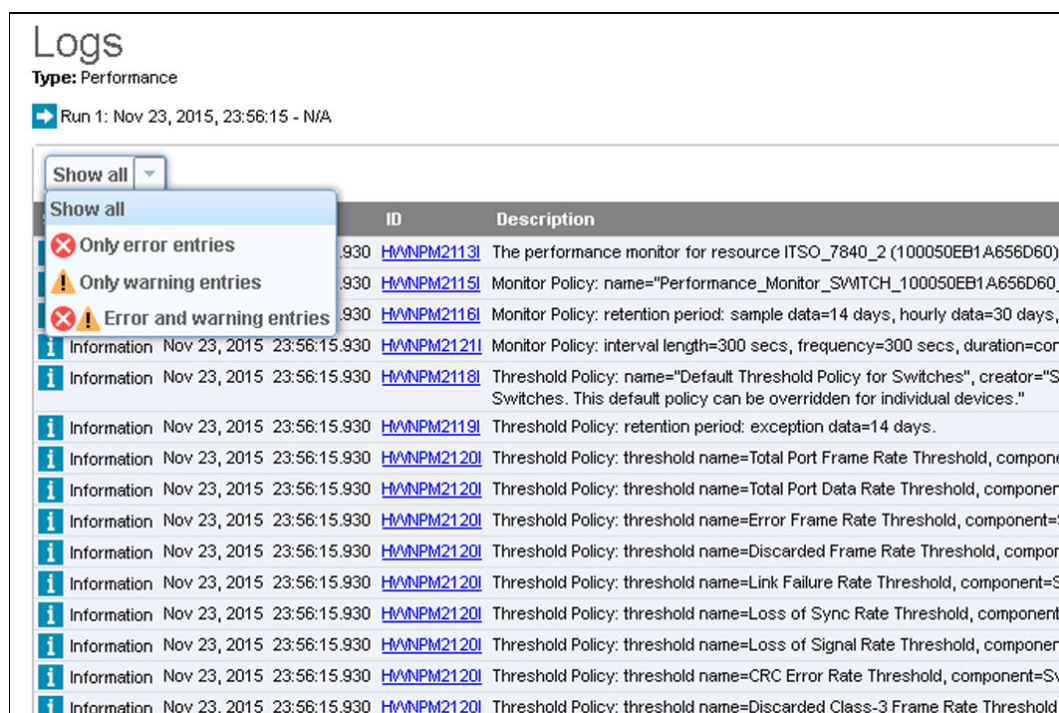


Figure 6-24 Performance monitor log window

The default window displays all of the performance monitor log entries. If you want to display errors, warnings, or both, you can apply the filter by clicking the filtering drop-down menu at the top left corner of the window, as shown in Figure 6-24 on page 162.

For information about errors in the performance monitor log, see the IBM Spectrum Control documentation in the IBM Knowledge Center at the following link:

<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>

6.7 Viewing Switches and Fabrics details

After the switch devices are successfully added to the IBM Spectrum Control instance and the probe has completed the corresponding switch and fabric discovery process, you can view the detailed information about the components and resources that are associated with a fabric or a switch.

6.7.1 Viewing fabric details

To access the fabric details, complete the following steps:

1. Select **Network** in the IBM Spectrum Control top menu, then select **Fabrics**.
2. In the Fabrics view, right-click the wanted fabric and then select **Details**, as shown in Figure 6-25.

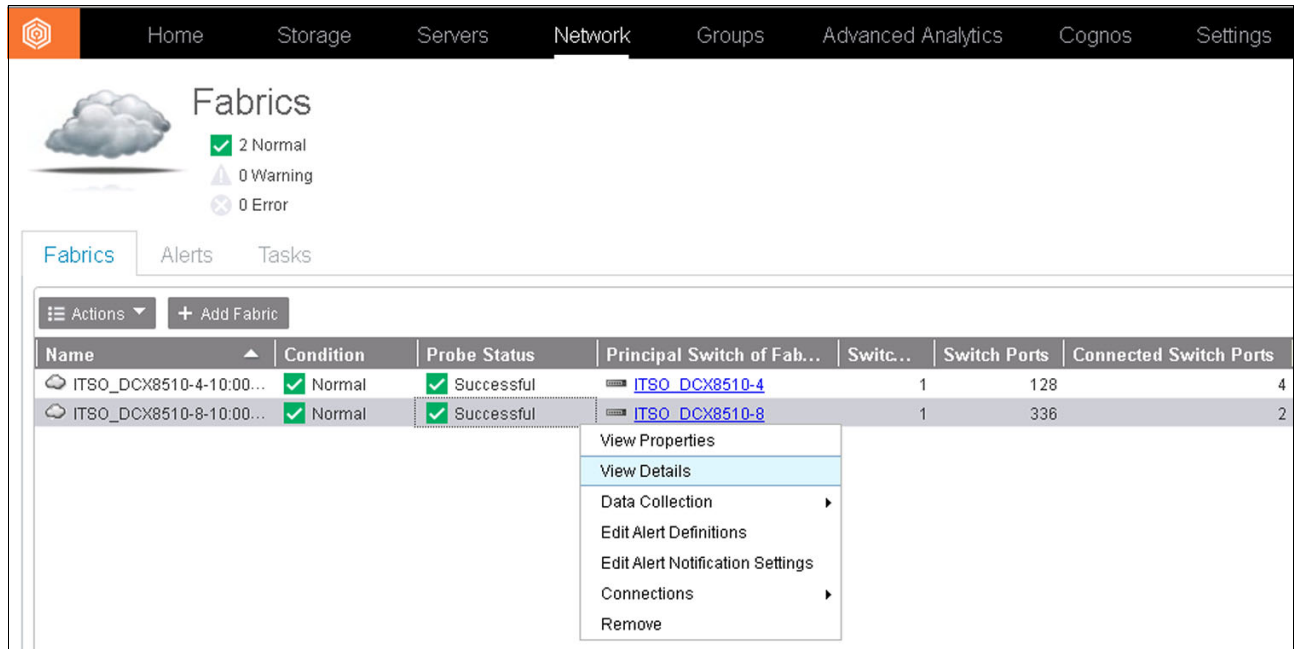


Figure 6-25 Showing fabric details

3. The fabric details main view will appear, as shown in Figure 6-26.




Figure 6-26 Fabric details main view

The upper left corner (region A of the view) shows the fabric's name. Click **Actions** to open a list of available actions for the fabric. You can also click the fabric's name to list of all of the fabrics that are managed by IBM Spectrum Control. If you select one of them, the window switches to the details view for this new fabric.

The left section of the view (region B) allows you to browse between different components and resources that are associated with the fabric. Each time that you click any of these items, the associated information will be presented in region C of the view.

The components and resources are grouped into three main sections:

1. General

The items in this section present key information about the fabric.

- Overview: When selected, this section presents you with multiple views displayed on the region C of the window that provide the user a quick insight of the overall fabric health and behavior. The default view selection is Summary. Figure 6-27 shows the Overview section.

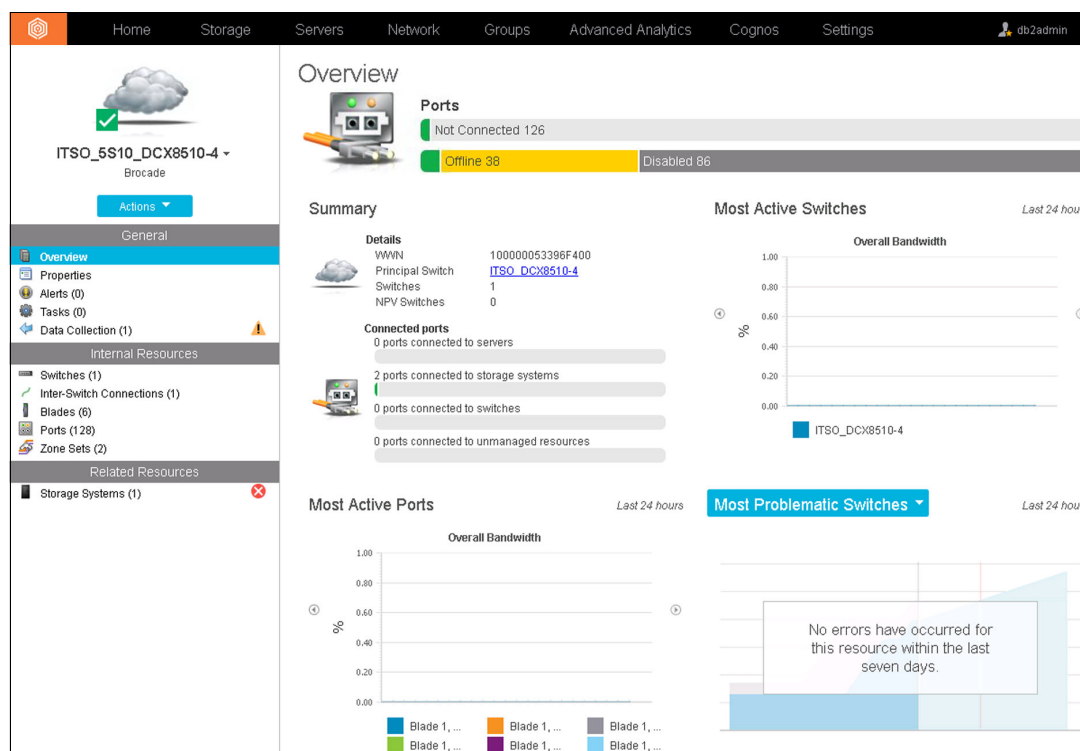


Figure 6-27 Main view of the Overview section

To switch between these views, click the view's name to select a different one. Figure 6-2 describes each of the available views of the Overview section.

Table 6-2 Available views in Overview section

View	Description
Summary	Shows the following information about the fabric: <ul style="list-style-type: none"> ▶ WWN: The World Wide Name (WWN) of the fabric. ▶ Principal Switch of Fabric: The name of the principal switch of the fabric at the time when the fabric was last probed. ▶ Switches: The number of physical and virtual switches in the fabric. ▶ NPV Switches: The number of Cisco switches in N_Port Virtualization (NPV) mode that are connected to the fabric. ▶ Ports: The total quantity of ports that are connected to servers, switches, storage systems, and unmanaged resources.

View	Description
Most Problematic Switches	Use the Most Problematic Switches view for visual summary of the switches with the largest number of errors. To switch between the error information for the last hour, last day, and last week, click the arrows icons next to the displayed period of time. Use the chart to quickly determine the switches with the largest number of errors and the types of errors that occurred. For more information, hover over areas in the charts.
Most Problematic Ports	Use the Most Problematic Ports view for visual summary of the switch ports with the largest number of errors.
Most Pervasive Errors	Use the Most Pervasive Errors section view for a visual summary of the types of errors that occur most frequently in the fabric. Each bar in the chart displays the total number of a type of error that occurred on all ports in the fabric.
Most Active Switches	Use the charts to view performance information for the most active switches in the fabric: <ul style="list-style-type: none"> ► Overall Bandwidth Percentage: A historical performance chart that shows the top six switches with the highest overall bandwidth percentage in the fabric. Each switch is represented by a line on the chart. ► Total Data Rate: A historical performance chart that shows the top six switches with the highest data rate in the fabric. ► Total Frame Rate: A historical performance chart that shows the top six switches with the highest frame rate in the fabric.
Most Active Ports	Use the charts to view performance information for the most active ports in the fabric: <ul style="list-style-type: none"> ► Overall Bandwidth Percentage: A historical performance chart showing the top six ports with the highest overall bandwidth percentage in the fabric. Each port is represented by a line on the chart. ► Total Data Rate: A historical performance chart that shows the top six ports with the highest data rate in the fabric. ► Total Frame Rate: A historical performance chart that shows the top six ports with the highest frame rate in the fabric.

Note: The charts show data collected hourly for the last 24 hours. If the charts are blank, ensure that the performance data collection for the fabrics is working.

To view different charts, click the name of a chart and select the chart that you want to view. Click the arrow icons on charts to switch between the different charts that are available for each chart type.

You can display multiple charts of the same type in different positions on the view. For example, you can display the Total Frame Rate and Total Data Rate charts for the Most Active Ports chart type at the same time in different positions on the view.

The selected charts are automatically displayed in the same positions the next time that you log on to the GUI and view the details of the resource.

- **Properties:** Displays the attributes of the fabric, including information about its status and connectivity, as shown in Figure 6-28.

The screenshot shows the 'Properties' section of the IBM Spectrum Control interface. The left sidebar contains a navigation menu with sections: General (Overview, Properties, Alerts (0), Tasks (0), Data Collection (1)), Internal Resources (Switches (1), Inter-Switch Connections (1), Blades (6), Ports (128), Zone Sets (2)), and Related Resources (Storage Systems (1)). The main content area is titled 'Properties' and has two tabs: 'General' (selected) and 'Connectivity'. The 'General' tab displays a table of attributes for the fabric 'ITSO_5S10_DCX8510-4' (Brocade). The attributes are as follows:

Attribute	Value
Name	ITSO_5S10_DCX8510-4
Status	Normal
Fabric Type	Brocade
Virtual	No
Probe Status	Normal
Probe Schedule	Daily. Next run at Jan 15, 2016, 01:30:00
Data Source Count	2
Location	—
Custom Tag 1	—
Custom Tag 2	—
Custom Tag 3	—

Figure 6-28 Main view of the Properties section

- **Alerts:** Allows you to display and manipulate the alerts definitions for the fabric. Figure 6-29 shows the Alerts section main view.

The screenshot shows the 'Alerts' section of the IBM Spectrum Control interface. The left sidebar is identical to the previous figure. The main content area is titled 'Alerts' and has three tabs: 'Alerts' (selected), 'Definitions', and 'Notification Settings'. The 'Alerts' tab displays a summary of alert counts: 0 Critical, 0 Warning, and 0 Informational. Below this is a table with columns: Condition, Severity, and Last Occurrence. The table is currently empty, displaying a message: 'No alert conditions were detected on monitored resources.'

Figure 6-29 Main view of the Alerts section

Table 6-3 describes the three main tabs in this section; Alerts, Definitions, and Notification Settings.

Table 6-3 Alerts section

Tab	Description
Alerts	Shows the alerts that are generated when certain conditions are detected on a fabric, conditions which are specified on the Definitions tab
Definitions	View and edit the alert definitions for a fabric. An alert definition includes a triggering condition and the notification settings for that condition. In general, the following types of conditions can trigger alerts: <ul style="list-style-type: none"> ► Data collection was not successful ► A change occurred in the configuration of a fabric
Notification Settings	View and edit the notification settings when alert conditions are detected on a fabric. These notification settings are applied globally to all the alert definitions in the Definitions tab. You can override the global settings by specifying different settings for each alert definition on the Definitions tab.

- Tasks: Enables you to view and manage the tasks that IBM Spectrum Control runs on the fabric, such as new zones creation during the block storage capacity provisioning. Figure 6-30 shows the Tasks main view.




Figure 6-30 Main view of the Tasks section

- Data Collection: Presents information about the different data collection jobs for the fabrics in your environment. You can view the status of the most recent run of a probe, start a probe run, schedule a probe, and view the logs for all the probe runs, as shown in Figure 6-31.




Figure 6-31 Main view of the Data Collection section

► Internal resources

Internal resources are all the components that exist in a fabric. Components that you can view include switches, inter-switch connections, blades, ports, and zone sets:

- Switches: To view information about the switches in the fabric, click **Switches**. The Switches section is displayed as shown in Figure 6-32. The number in parentheses shows the number of switches in the fabric.

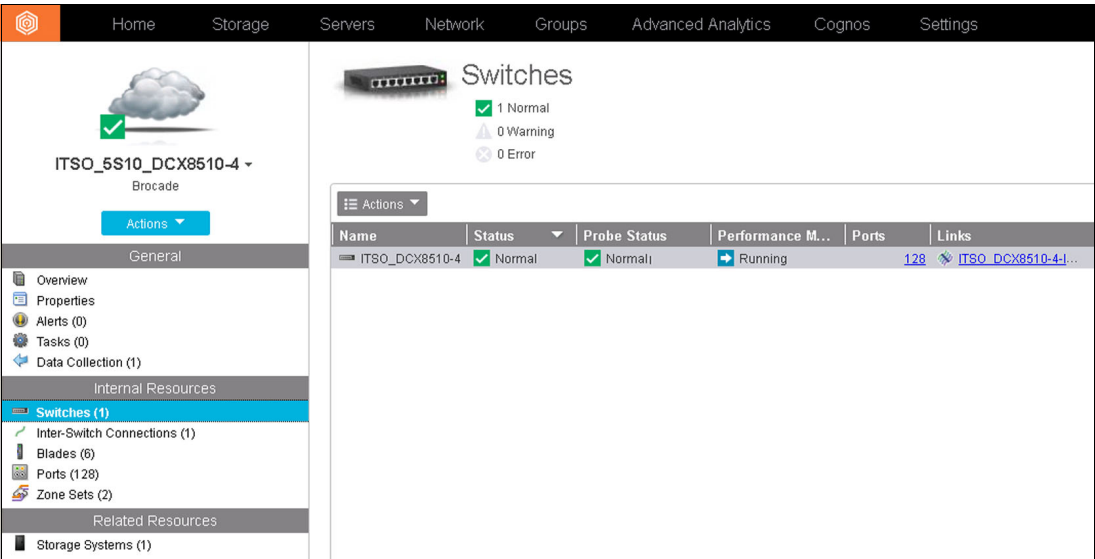


Figure 6-32 Internal resources - Switches section

- **Inter-switch Connections:** To view information about inter-switch connections, click **Inter-Switch Connections**. The number in parentheses shows the number of inter-switch connections that are associated with the fabric.

This section includes the following types of connections: Inter-switch link (ISL), ISL trunk, NPV link, and inter-chassis link (ICL) trunk.

The status values help you to decide whether you need to investigate or resolve issues with inter-switch connections. ISL trunks and ICL trunks have a warning status if less than half of the links have a port that has an error status. ISL trunks also have a warning status if half or more than half of the links have ports that have a warning status.

- **Blades:** Shows information about the switch blades that are associated with the fabric, as shown in Figure 6-33:
 - **Port Range:** The range of ports for the ports that are components of the blade. For example, 0-15, 16-31. This value represents the port index relative to the switch, not the port numbers relative to the blade.
 - **Ports:** The number of switch ports on a blade.
 - **Slot:** The physical slot on a switch to which a blade is attached.
 - **Status:** The status of a blade. Statuses include Normal, Warning, Error, Unreachable, and Unknown. Use the status to determine the condition of a blade, and if any actions must be taken. For example, if a blade has an Error status, take immediate action to correct the problem.

The screenshot displays the 'Blades' section in the IBM Spectrum Control interface. The top navigation bar includes links for Home, Storage, Servers, Network, Groups, Advanced Analytics, Cognos, and Settings. The main content area is titled 'Blades' and shows a summary of 6 Normal, 0 Warning, and 0 Error blades. Below this is a table listing the blades with columns for Slot, Switch, Status, Port Range, Vendor, Serial Number, and Blade Type. The table shows 6 blades, all with a status of Normal. The left sidebar contains a navigation menu with options like Overview, Properties, Alerts, Tasks, Data Collection, Internal Resources, and Related Resources.

Slot	Switch	Status	Port Range	Vendor	Serial Num...	Blade Type
5	ITSO_DCX8510-4	Standby		IBM	AHJ0429G08K	Control Processor
1	ITSO_DCX8510-4	Normal	0-31	IBM	BQA0448H04E	FC Port
3	ITSO_DCX8510-4	Normal	256-751	IBM	BQD0419G016	Core
4	ITSO_DCX8510-4	Normal		IBM	AHJ0423G01T	Control Processor
6	ITSO_DCX8510-4	Normal	272-799	IBM	BQD0428G01B	Core
8	ITSO_DCX8510-4	Normal	192-223	IBM	ATM0302F003	Unknown

Figure 6-33 Main view of the Blades section

- Ports: Shows information about all of the ports that are related to the fabric, as shown in Figure 6-34.

The screenshot displays the 'Ports' section for a Brocade switch named 'ITSO_5S10_DCX8510-4'. The interface includes a top navigation bar with tabs for Home, Storage, Servers, Network, Groups, Advanced Analytics, Cognos, and Settings. On the left, a sidebar shows a tree view of resources: Overview, Properties, Alerts (0), Tasks (0), Data Collection (1), Internal Resources (Switches (1), Inter-Switch Connections (1), Blades (6), Ports (128), Zone Sets (2)), and Related Resources (Storage Systems (1)). The 'Ports' section is highlighted in blue. The main content area shows a summary of 128 Normal ports, 0 Warning, and 0 Error. Below this is a table with columns: Switch, Blade Slot, Port Number, Status, Port Name, State, Speed, and WWPN. The table lists 16 rows of port information, all with a status of 'Operational'.

Switch	Blade Slot	Port Number	Status	Port Name	State	Speed	WWPN
ITSO_DCX8510-4	1	25	Operational		Online	16	201900053396F400
ITSO_DCX8510-4	1	10	Operational		Online	16	200A00053396F400
ITSO_DCX8510-4	1	21	Operational	F_Port	Online	8	201500053396F400
ITSO_DCX8510-4	1	20	Operational	F_Port	Online	8	201400053396F400
ITSO_DCX8510-4	6	31	Operational	G_Port	Disabled	16	500053396F47B31F
ITSO_DCX8510-4	3	31	Operational	G_Port	Disabled	16	500053396F47B2EF
ITSO_DCX8510-4	1	31	Operational	G_Port	Enable...	8	201F00053396F400
ITSO_DCX8510-4	6	30	Operational	G_Port	Disabled	16	500053396F47B31E
ITSO_DCX8510-4	3	30	Operational	G_Port	Disabled	16	500053396F47B2EE
ITSO_DCX8510-4	1	30	Operational	G_Port	Enable...	8	201E00053396F400
ITSO_DCX8510-4	6	29	Operational	G_Port	Disabled	16	500053396F47B31D
ITSO_DCX8510-4	3	29	Operational	G_Port	Disabled	16	500053396F47B2ED
ITSO_DCX8510-4	1	29	Operational	G_Port	Enable...	8	201D00053396F400
ITSO_DCX8510-4	6	28	Operational	G_Port	Disabled	16	500053396F47B31C
ITSO_DCX8510-4	3	28	Operational	G_Port	Disabled	16	500053396F47B2EC
ITSO_DCX8510-4	1	28	Operational	G_Port	Enable...	8	201C00053396F400
ITSO_DCX8510-4	3	27	Operational	G_Port	Disabled	16	500053396F47B2EB

Figure 6-34 Main view of the Ports section

- Zonesets: Shows information about all the zonesets in the fabric, whether they are active or inactive, as shown in Figure 6-35.

The screenshot displays the 'Zone Sets' section for a Brocade switch named 'ITSO_5S10_DCX8510-4'. The interface is similar to the previous one, with the 'Zone Sets' section highlighted in blue in the sidebar. The main content area shows a summary of 2 Zone Sets. Below this is a table with columns: Zone Set Name, Active, Zones, and Description. The table lists 2 rows of zone set information.

Zone Set Name	Active	Zones	Description
W2K_V7K_Zoning_cfg	No	1	
W2K_V7K_Zoning_cfg	Yes	1	

Figure 6-35 Main view of the Zone Sets section

► **Related resources**

This section displays information about resources that are related with the fabric. These resources include servers, storage systems, hypervisors, and resources that are not monitored but are associated, such as Discovered ports. Discovered ports are ports on resources that are not monitored by IBM Spectrum Control but are visible to a monitored fabric.

6.8 Zoning

In addition to the fabric and switches monitoring capabilities of IBM Spectrum Control, you can enable the automatic zoning feature through the web-based GUI. When automatic zoning is enabled, new zones will be created during the block storage provisioning tasks if new zones are needed to connect the storage system with the server. When you set the zoning policy, you can specify whether zone changes from automatic zoning are made to the active zone set, or to a new inactive active zone set.

6.8.1 Non-standard zones

B-type switches support some nonstandard zones such as quick loop zones, fabric assist zones, and protocol zones. If the switch configurations have these zones already defined, Fabric Manager preserves them and does not modify them in any way. If they are part of a zone set that is active at some time, the devices that are part of such zones that are also online are displayed in the topology Zone View.

You can create, change, and delete nonstandard zones by using the b-type switch management applications, like Web Tools or IBM Network Advisor.

Note: In IBM Tivoli Storage Productivity Center releases 5.1 and earlier, the stand-alone GUI allowed the user to manipulate the zoning (create, modify, and remove aliases, zones, and zonesets) through the Fabric Manager. This method provided an alternative to the devices management tools (such as Web Tools or IBM Network Advisor). In IBM Spectrum Control, this functionality is no longer available. IBM Spectrum Control can modify the fabric zoning only by using the automatic zoning feature during the storage provisioning task.

6.8.2 Zone control capabilities of IBM Spectrum Control

Table 6-4 shows the zone function capabilities of IBM Spectrum Control when managing b-type switches. Not all activities that are available with the switch management application are available in IBM Spectrum Control.

Table 6-4 IBM Spectrum Control zoning capabilities for b-type switches

Capability	Support
Zone control through Storage Resource Agent supported	No
Zone control through out-of-band SNMP agent supported	No
Zone control through CIMOM agent supported	Yes

Capability	Support
Domain/port zone members allowed	Yes
Port WWN zone members allowed	Yes
Node WWN zone members allowed	Yes
FCID zone members allowed	No
Zone Aliases supported	Yes

6.8.3 Setting the zoning policy (automatic zoning feature)

The IBM Spectrum Control automatic zoning feature is disabled by default, meaning that only storage systems with connectivity to the server requiring the capacity are candidates for provisioning.

To enable the automatic zoning feature, you need to define a zoning policy. When automatic zoning is enabled, IBM Spectrum Control can create zones during block storage provisioning to connect a server to a storage system.

When IBM Spectrum Control creates a provisioning task, it identifies the best location for the new storage that satisfies the requirements of the service class. If automatic zoning is enabled, then, during provisioning, existing zones are used if the server already has connectivity to the storage system. Otherwise, the required zone or zones are created between a host initiator ports and the storage controller ports.

To enable the zoning policy, complete the following steps:

1. Select **Advanced Analytics** on the IBM Spectrum Control top menu, then select **Provisioning**. The provisioning dialog will appear as shown in Figure 6-36.

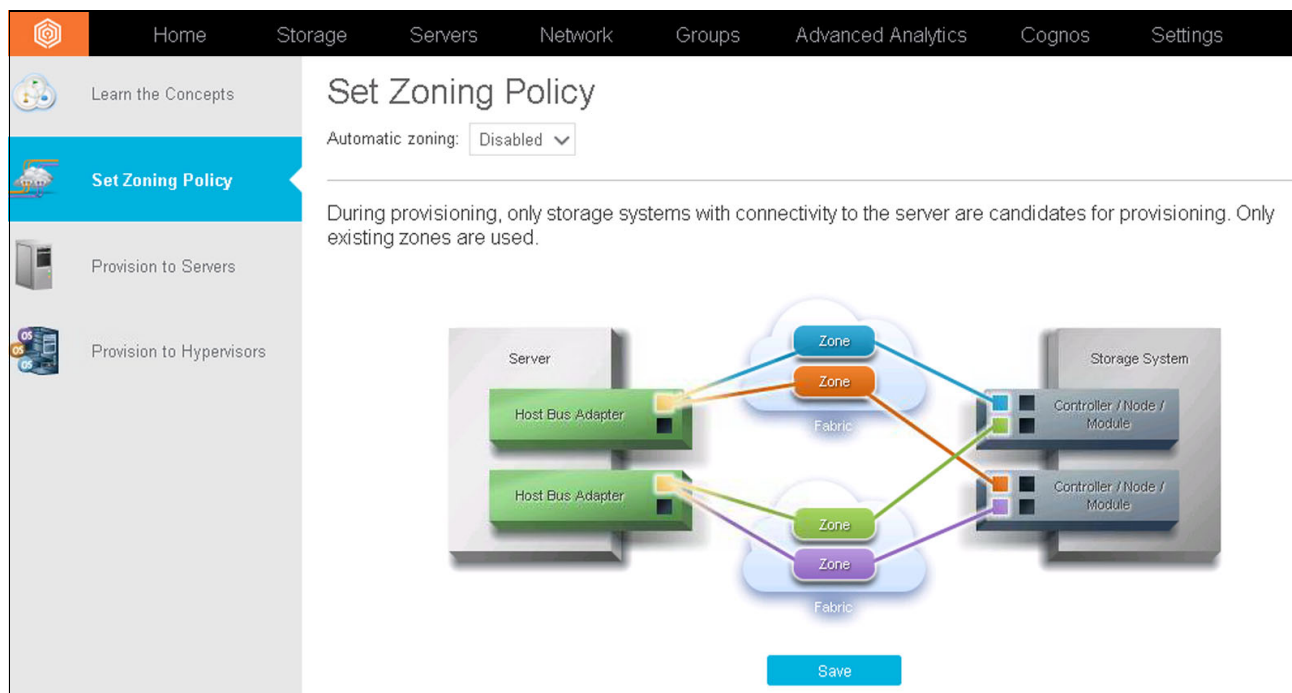


Figure 6-36 Setting the zoning policy

Click **Set Zoning Policy**.

2. Click the drop-down menu and select **Enabled**, as shown in Figure 6-37.

Set Zoning Policy

Automatic zoning: **Enabled**

During provisioning, new zones might be created to connect a server to the storage system. Existing zones are used if the server already has connectivity to the storage system. Otherwise, a new zone is created between a host initiator port and a controller, node, or module port.

Zone name prefix: ITSOTPC

The naming convention for new zones is ITSOTPC_host_storage-system_suffix-number.

☒ Make changes to the active zone set

Save

Figure 6-37 Zoning policy configuration

When it is enabled, IBM Spectrum Control can create zones during storage provisioning, if required.

Note: The naming convention for the zones to be created will be *prefix_host_storage-system_suffix-number*.

When you enable automatic zoning, you can specify the following options:

- **Zone name prefix:** If you specify a zone name prefix, all zones that are automatically created are prefixed with your input. This prefix can help you identify which zones were automatically created by IBM Spectrum Control when reviewing your fabric.
- **Make changes to the active zone set:** Whether changes are made to the active zone set, or to a new inactive active zone set. If this check box is selected, changes are made to the active zone set. If the check box is cleared, changes are made to the inactive zone set. The new inactive set will contain only the new zones.

3. Click **Save**. A window is displayed to confirm the result of the save action (Figure 6-38).




Figure 6-38 Changes successfully saved

For more information about zoning and storage provisioning tasks, see the IBM Spectrum Control documentation in the IBM Knowledge Center at the following link:

<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>

Additionally, you can see the User and Administrator guides, which can be found at the following link:

http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.9/com.ibm.spectrum.sc.doc/fqz0_r_printable_pdf_files.html

6.9 Alerting

You can use IBM Spectrum Control to monitor your b-type switches and to automatically notify you whenever a switch resource changes its condition or violates a performance threshold. The conditions that generate alerts are detected during the data collection and event processing. You can customize which alert will trigger a notification, and in some cases like the performance thresholds violations, you will be required to enter values for triggering such alerts.

When an event occurs and triggers an alert, the alert is written to the IBM Spectrum Control log. You can also select one or more other ways to be notified of the event. These alerts notifications include SNMP traps, IBM Tivoli Enterprise Console® events, entries in Windows event log or UNIX syslog, and emails.

6.9.1 Prerequisites for using alerts

To successfully use the IBM Spectrum Control alerting feature, the following conditions must be met:

- ▶ Data collection schedules must be configured and scheduled to run regularly. Additionally and to detect performance thresholds violations, the performance monitors must be running to collect performance data from your switches.
- ▶ If you want to be notified about an alert in some way other than an entry in the IBM Spectrum Control log and display in the Web GUI, such as using SNMP traps, Tivoli Enterprise Console events, or email notifications, you must configure those alert destinations beforehand.
- ▶ If an alert is triggered based on an SNMP trap from the monitored switch, you must properly configure the SNMP server of the monitor resource to enable IBM Spectrum Control to listen to SNMP traps. The default port number is 162, and the default community is public.

6.9.2 Setting up the alert notification settings

To configure the IBM Spectrum Control alert notification settings that will be used by default by the monitored resources, complete the following steps:

1. Select **Settings** on the IBM Spectrum Control top menu, and then select **Alert Notifications**. The alert notifications settings window will be displayed, as shown in Figure 6-39.

The screenshot shows the 'Alert Notifications' window in IBM Spectrum Control. The left sidebar has 'Email' selected. The main panel shows the 'Email' configuration. There are 'Save' and 'Cancel' buttons at the top. Below them is the 'Email server for sending alert notifications' section with input fields for 'Reply to address' (ITSO_Tito@us.ibm.com), 'Mail server' (10.18.228.179), and 'Port' (25). There are 'Test' and 'Remove' buttons below these fields. At the bottom, under 'Global email notification settings', there are two sections: 'Configuration Alerts' and 'Performance Threshold Alerts'. Each has a checked 'Email' checkbox and the email address ITSO_Tito@us.ibm.com.

Figure 6-39 Alert notification settings window

2. Enter the email server values to be used for sending alert notifications:
 - Reply to address: If a user replies to an email that was triggered by an alert, the reply is sent to this email address. This email address also receives any undeliverable email messages for alerts that are configured with incorrect or invalid email address.
 - Mail server: The name of the mail server to use. You can specify a host name, an IPv4 address, or an IPv6 address.
 - Port: The port number for the outgoing SMTP (email) server. The default SMTP server listening port is 25.
3. When all of the email server values are entered, test the communication by clicking **Test**. You will need to provide an email address to send the email test to.

4. Set the global email notification settings. The email addresses that you define to receive alert notifications are applied globally to all alert definitions and all resources (like switches), unless they are specifically overridden. You can later configure a different notification method for a specific resource or set of resources by overriding these global settings when editing the resource alerts definitions. For more information, see “Override the default notification settings” on page 181.

6.9.3 Enabling the default alert definitions

Enabling the default alert definitions for your b-type switches allows you to be automatically notified when certain predefined conditions are detected.

To enable these alerts, complete the following steps:

1. Select **Network** on the IBM Spectrum Control top menu, and then select **Switches**.
2. In the Switches view, right-click the switch that you want, and then select **Edit Alerts Definitions**, as shown in Figure 6-40.

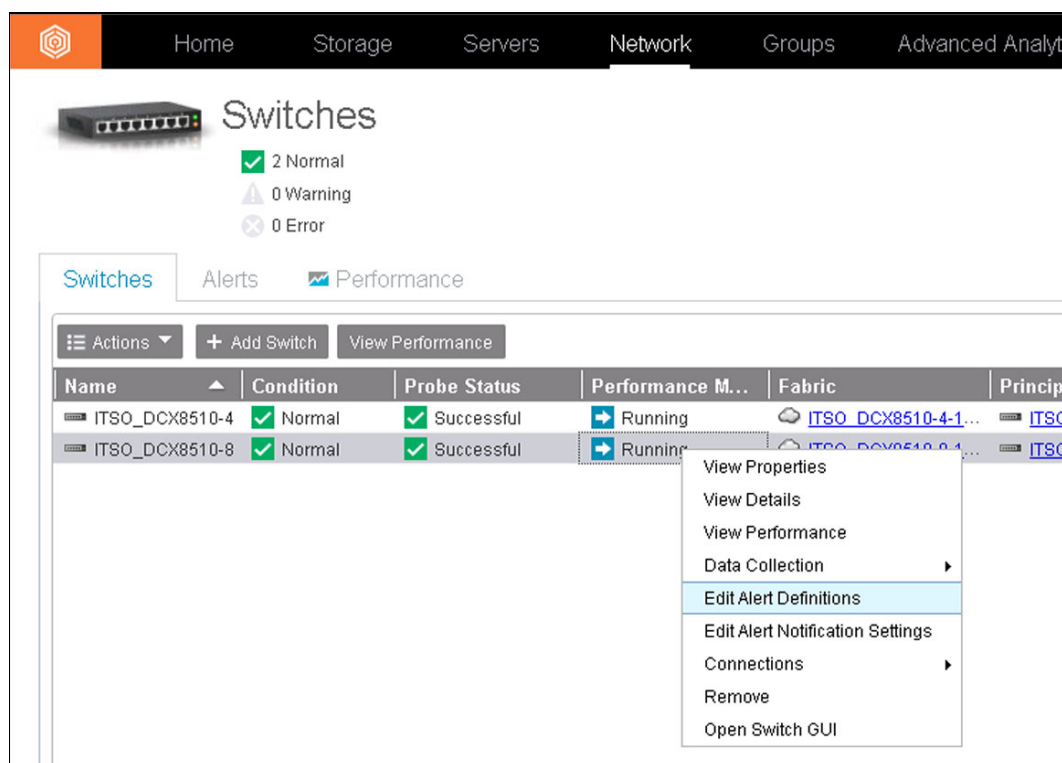


Figure 6-40 Editing alert definitions

As you can see in Figure 6-41, there are three categories of alerts: Switches, Inter-Switch Connections, and Ports. Each one of these categories has a set of predefined alerts that can be enabled for quick and easy use of switch alerting capabilities. You can also customize these alerts definitions for alignment with your alerting needs (see 6.9.4, “Custom alert definitions” on page 179).

Edit Alert Definitions

Switches 1/4 Inter-Switch Connections 0/0 Ports 9/18

General 1/4 Performance 0/0

<input checked="" type="checkbox"/>	Status	error	⚠	✉	🚫	+
<input type="checkbox"/>	Performance Monitor Status		⚠	✉	🚫	+
<input type="checkbox"/>	Last Successful Monitor		⚠	✉		+
<input type="checkbox"/>	Firmware		⚠	✉	🚫	+

Save Cancel

Figure 6-41 Editing alert definitions: General condition alerts

3. Click **Cancel** to close the dialog. The default alerts for the switch are automatically enabled after you access the Edit Alert Definitions window.

6.9.4 Custom alert definitions

You can customize your alerting environment by enabling all of those additional alerts that are not enabled by the default alert definitions. To enable these alerts, complete the following steps:

1. Open the Edit Alert Definitions window as shown in Figure 6-41.

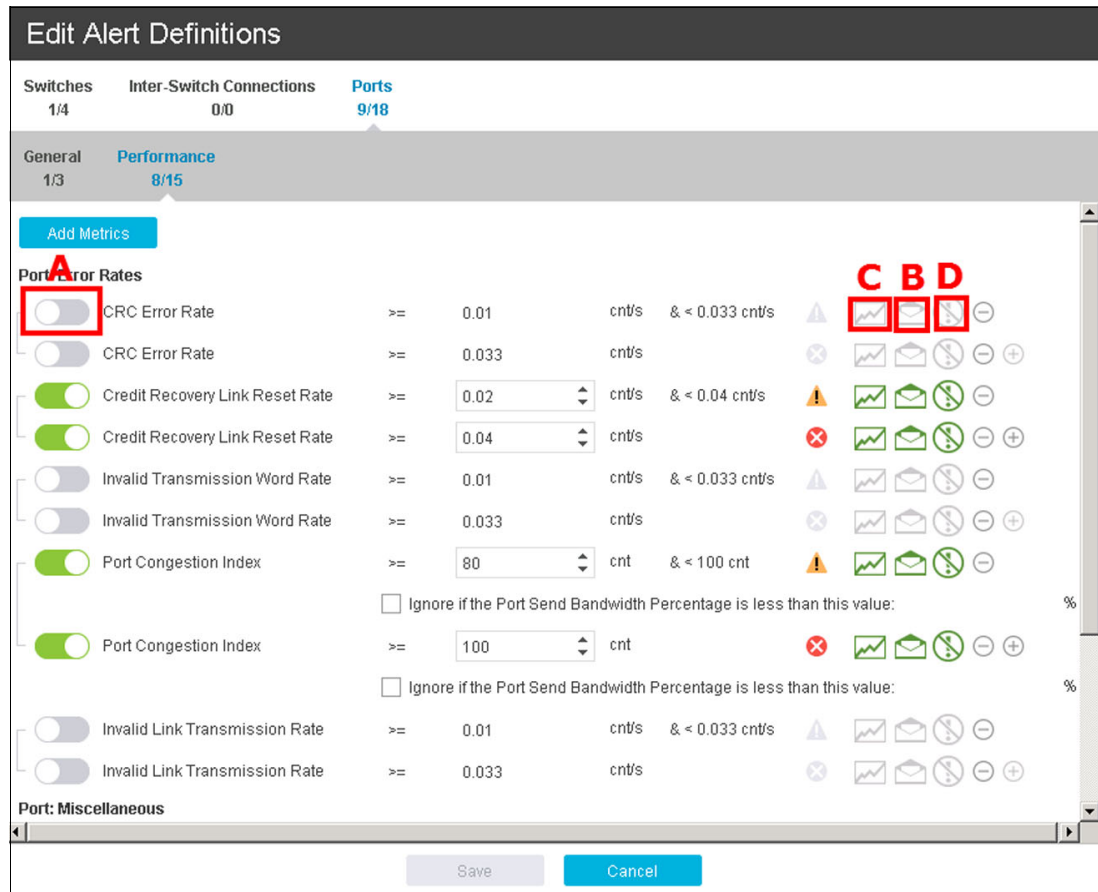


Figure 6-42 Custom alert definitions

2. Modify (enable or disable) the wanted alert by switching the green switch A (see Figure 6-42) on or off.

3. If the alert was enabled, configure the threshold values by which this alert will be triggered. Click the C button (see Figure 6-42 on page 179) and the alert threshold configuration pane will be presented, as shown in Figure 6-43.

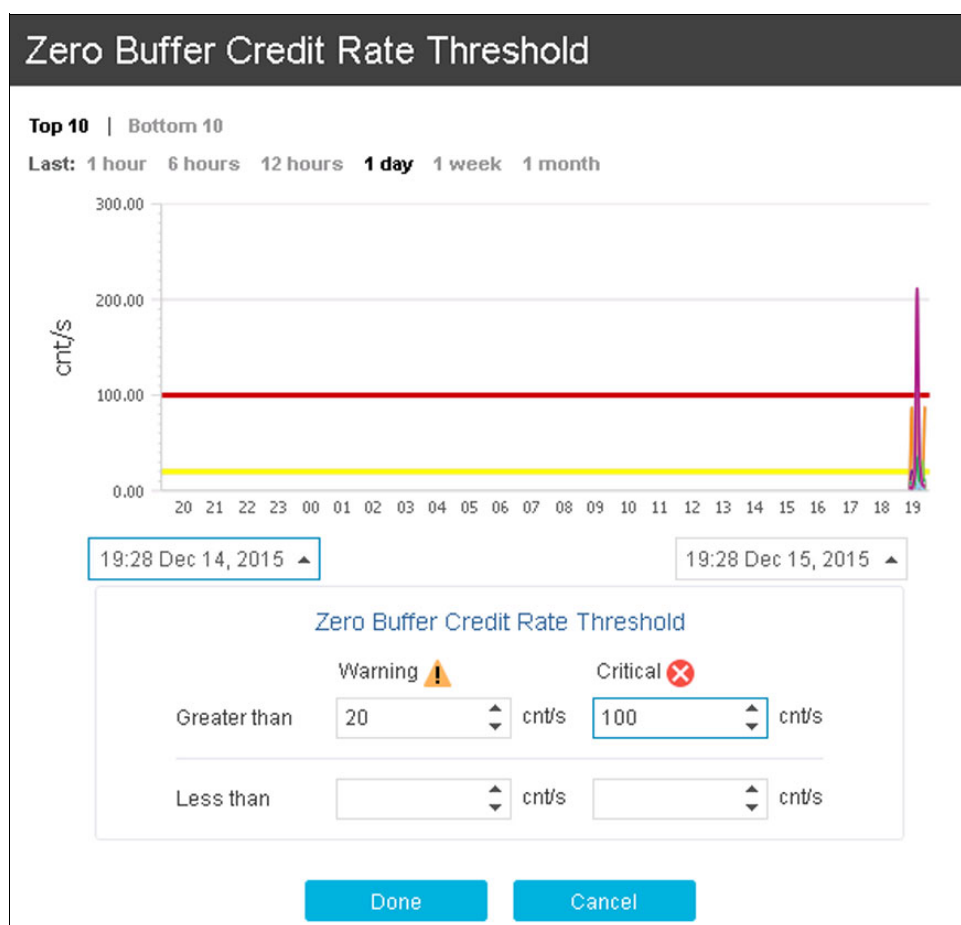


Figure 6-43 Editing the custom alerts thresholds

4. Input the wanted threshold values by which you want IBM Spectrum Control to trigger both the Warning and Critical alerts. If any information is already available for the alert that you are configuring, that information will be charted above your threshold fields. Consider this effect when you define your threshold values.
5. When you are satisfied with your input, click **Done** to go back to the main Edit Alert definitions page.
6. Repeat steps 2 - 5 for each alert that you want to customize. When you are done, click **Save** to commit and enable your changes to the alerting definitions.

7. You can also modify the severity of each specific alert, allowing you to align them according to your own organization's needs. In order to do so, click the **Severity** icon, as shown in Figure 6-44.

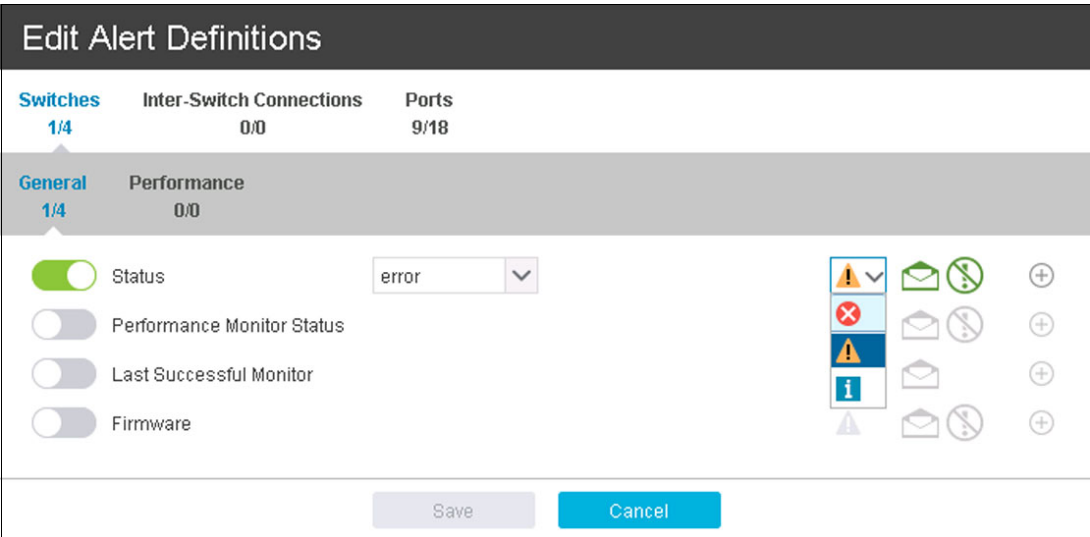


Figure 6-44 Changing the alert severity

Select the alert severity that you want based on the guidelines that are shown in Figure 6-45 and click **Save**.

Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes. For example, assign an informational severity to alerts that are generated when a new pool is added to a storage system.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the status of a data collection job is not normal.
	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the amount of available space on a file system falls below a specified threshold.

Figure 6-45 Alerts severity guidelines

Override the default notification settings

In some situations, you might want to have specific notification settings for a particular switch alert instead of using the global notification settings (see 6.9.2, “Setting up the alert notification settings” on page 176).

To override the default notification settings, complete the following steps:

1. While editing the alerts definitions, click the **B** button (see Figure 6-42 on page 179) and the alert notification settings for that particular alert and switch will be presented, as shown in Figure 6-46.

CRC Error Rate Threshold Setti...

Specify actions to be taken when this alert is generated

☐ Run script

☒ Override notification settings ?

☒ Email ITSO_Tito@us.ibm.com Customize

☐ Netcool / OMNIbus ?

☐ SNMP

☐ Windows log

Done Cancel

Figure 6-46 Overriding the default notification settings for an alert

2. Select **Override notification settings** to enable the additional configuration fields.
3. Modify the notification settings as desired. If you are setting an email address, you can also click **Customize** to customize the email's subject. If you select the IBM Netcool® option, ensure that the Netcool EIF probe's IP address is already configured under IBM Spectrum Control settings.
4. When you are satisfied with your input, click **Done** to go back to the main Edit Alert definitions page.
5. Repeat steps 1 - 4 for each alert you want to customize and, when you are done, click **Save** to commit and enable your changes to the alerting definitions.

Suppressing alerts

In some scenarios, you might want to have some control over the conditions during which the alert is triggered, beyond the thresholds definitions. In order to do so, complete the following steps:

1. While editing the alerts definitions, click the **D** button (see Figure 6-42 on page 179) and the alert custom suppression settings for that particular alert and switch will be presented, as shown in Figure 6-47.

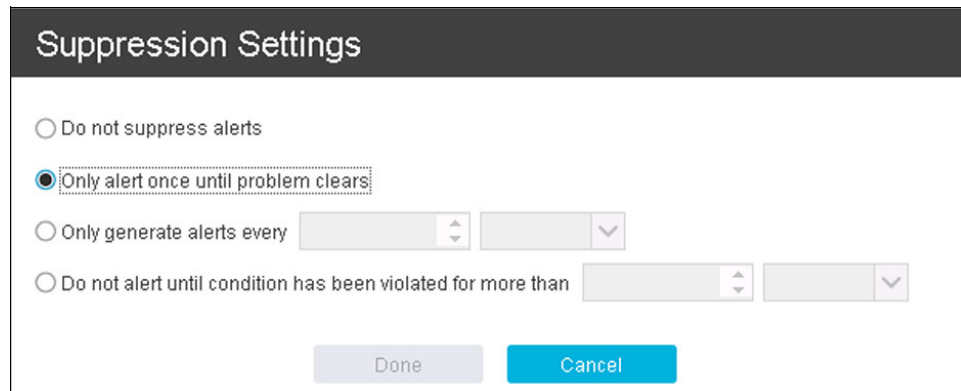


Figure 6-47 Setting alerts suppression conditions

2. Modify the wanted notification settings. The following options are available:
 - If you want to receive alert notifications on every detected alert violation, select **Do not suppress alerts**.
 - **Only alert once until problem clears** will notify you once whenever a new alert violation is detected and will discard future alerts about the same condition.
 - If you want to receive an alert notification only once and discard the following violations for a certain period, select **Only generate alerts every**. A new alert notification will be triggered when the defined time period ends, considering the time of the first alert occurrence as the start time.
 - Use the **Do not alert until condition has been violated for more than** option to define a time period for which you will allow this violation to occur until an alert is sent. For example, you might want to be notified about a specific resource high utilization condition, but only when that condition has exceeded a certain amount of time.
3. When you are satisfied with your input, click **Done** to go back to the main Edit Alert definitions page.
4. Repeat steps 1 - 3 for each alert that you want to customize. When you are done, click **Save** to commit and enable your changes to the alerting definitions.

6.9.5 Managing and acknowledging alerts

Some alerts are triggered by conditions that commonly occur and can be ignored. The context of the environment (applications, servers, business, users, and so on) can determine that an alert that might be meaningful for some is not for another. In such cases, you can acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

To manage and acknowledge alerts, complete the following steps:

1. Select **Network** on the IBM Spectrum Control top menu, and then select **Switches**.
2. In the Switches view, click the Alerts tab. The Alerts main window is displayed as shown in Figure 6-48.

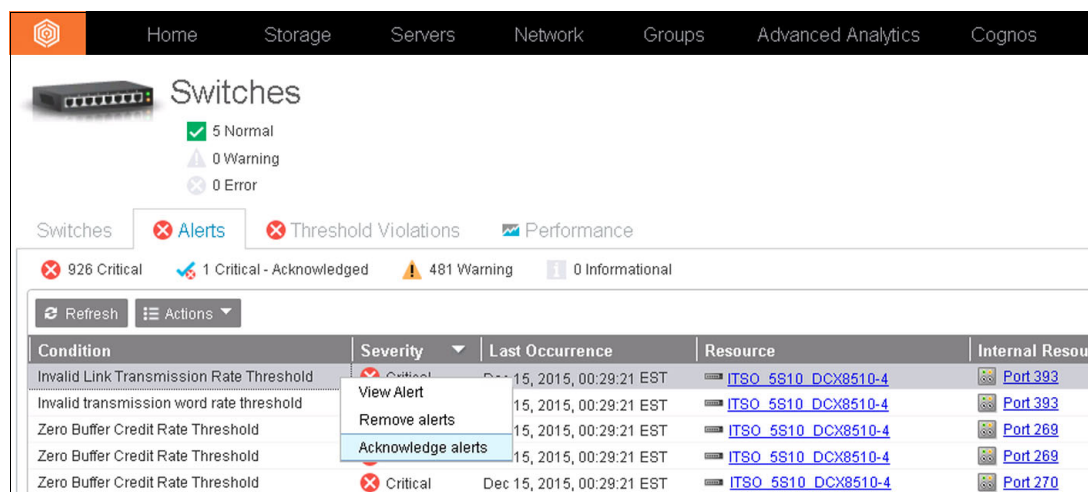


Figure 6-48 Managing alerts

3. To acknowledge an alert, right-click the alert and select **Acknowledge alerts**. The alert will be removed from the list.

For more information about alerting, see the IBM Spectrum Control documentation in the IBM Knowledge Center at the following link:

<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>

Additionally, you can refer to the User and Administrator guides that at the following link:

http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.9/com.ibm.spectrum.sc.doc/fqz0_r_printable_pdf_files.html

6.10 Performance monitoring

IBM Spectrum Control can collect information about the performance of your switches. This information includes key performance metrics and notifications of threshold violations that can help you identify and troubleshoot issues in your fabrics.

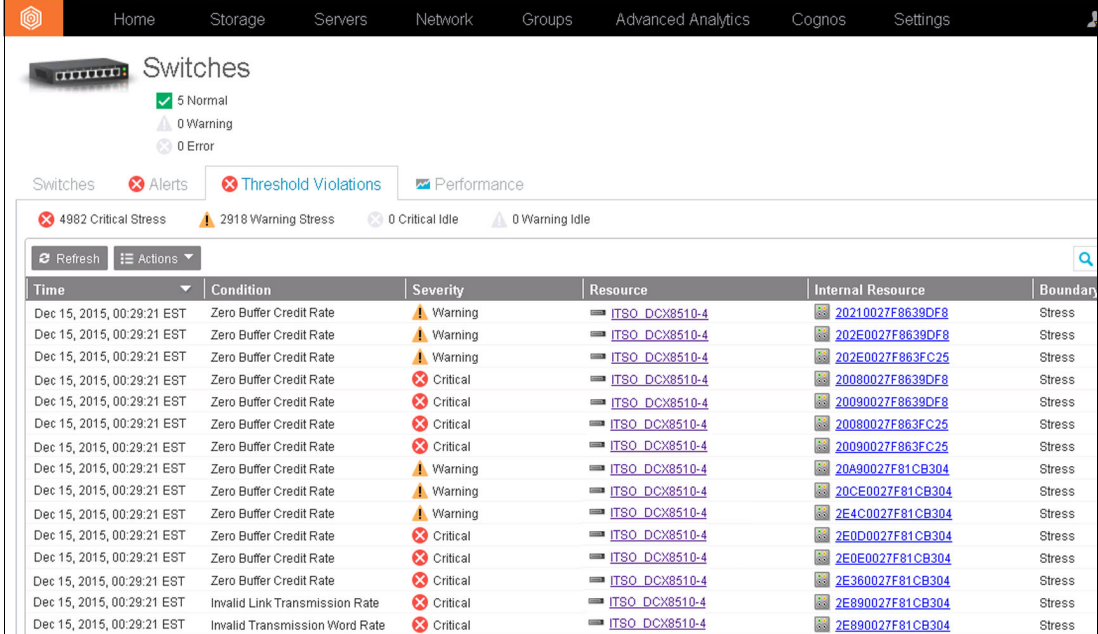
Before you review the performance of your devices, you must complete the following tasks:

- Add your switches to IBM Spectrum Control to enable the switch monitoring and schedule the data collection task (see 6.4, “Adding switches and fabrics” on page 148)
- Ensure that performance data is being collected (see 6.6, “Enabling the switch performance monitoring” on page 160)
- Optionally, configure the performance alerting by defining the corresponding alerts thresholds (see 6.9, “Alerting” on page 175)

6.10.1 Performance thresholds violations review

Similarly to the way the Alerts are displayed on a centralized view (see Figure 6-48 on page 184), IBM Spectrum Control allows you to review and manage all the performance thresholds violations by using the Thresholds Violations centric view. This view enables you to quickly identify potential issues on your fabric and giving you a start point for your more complex troubleshooting investigations. In order to access this view, complete the following steps:

1. Select **Network** on the IBM Spectrum Control top menu, and then select **Switches**.
2. In the Switches view, click the Threshold Violations tab, as shown in Figure 6-49.



Time	Condition	Severity	Resource	Internal Resource	Boundary
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Warning	ITSO_DCX8510-4	20210027F8639DF8	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Warning	ITSO_DCX8510-4	202E0027F8639DF8	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Warning	ITSO_DCX8510-4	202E0027F863FC25	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Critical	ITSO_DCX8510-4	20080027F8639DF8	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Critical	ITSO_DCX8510-4	20090027F8639DF8	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Critical	ITSO_DCX8510-4	20080027F863FC25	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Critical	ITSO_DCX8510-4	20090027F863FC25	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Warning	ITSO_DCX8510-4	20A90027F81CB304	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Warning	ITSO_DCX8510-4	20CE0027F81CB304	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Warning	ITSO_DCX8510-4	2E4C0027F81CB304	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Critical	ITSO_DCX8510-4	2E0D0027F81CB304	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Critical	ITSO_DCX8510-4	2E0E0027F81CB304	Stress
Dec 15, 2015, 00:29:21 EST	Zero Buffer Credit Rate	Critical	ITSO_DCX8510-4	2E360027F81CB304	Stress
Dec 15, 2015, 00:29:21 EST	Invalid Link Transmission Rate	Critical	ITSO_DCX8510-4	2E890027F81CB304	Stress
Dec 15, 2015, 00:29:21 EST	Invalid Transmission Word Rate	Critical	ITSO_DCX8510-4	2E890027F81CB304	Stress

Figure 6-49 Managing Threshold Violations

6.10.2 Performance review

In order to review and troubleshoot the performance of your switches, complete the following steps to access the IBM Spectrum Control Performance view:

1. Select **Network** on the IBM Spectrum Control top menu, and then select **Switches**.
2. In the Switches view, click the Performance tab. The Performance view is displayed as shown in Figure 6-50.

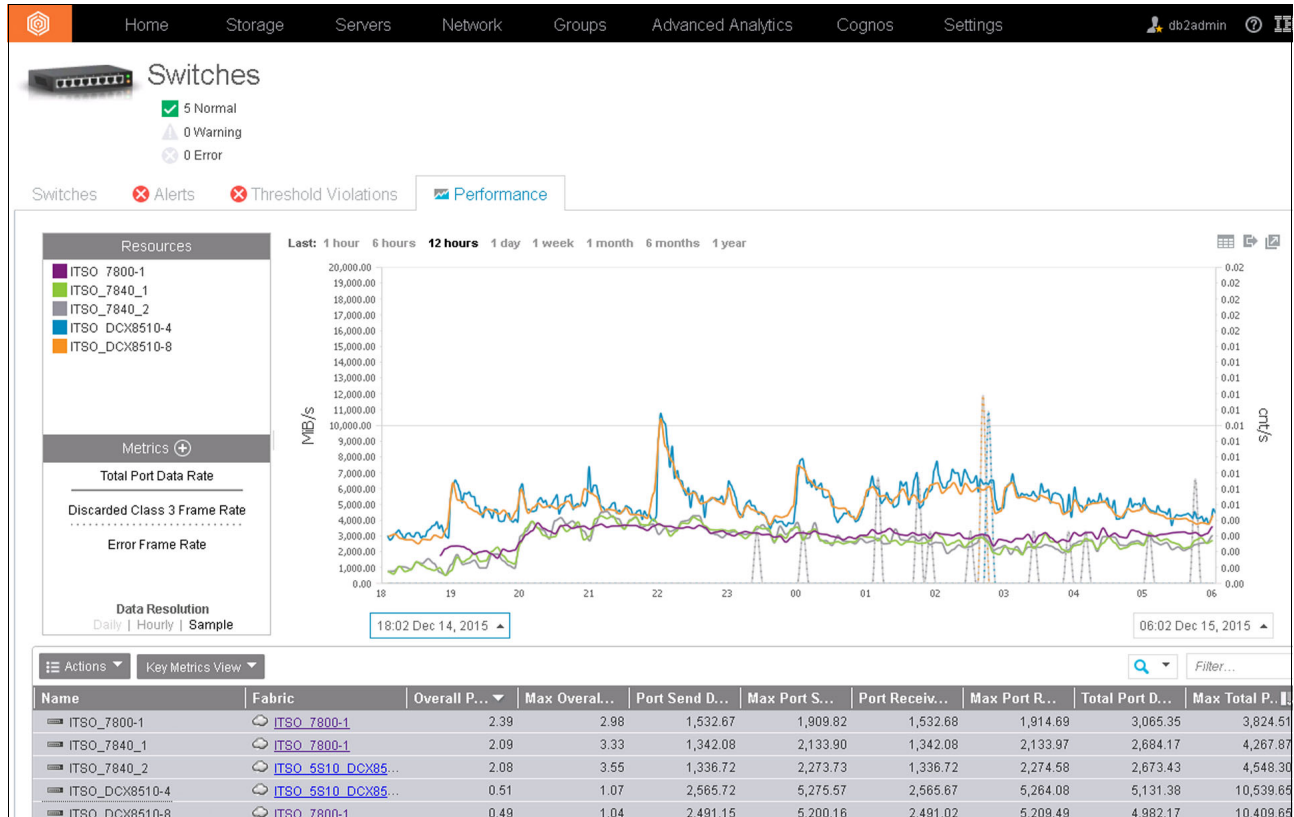


Figure 6-50 Performance view

Note: You can also access the performance view by selecting a switch on the main Switches view, right-clicking the switch that you are interested in, and then selecting **View Performance**. With this method, you only work with the performance of that selected switch.

3. By default, the main window displays charts the default switch metrics (Total Data Rate) for all of the switches. From here, you can start your troubleshooting by changing the display parameters according to your needs:
 - a. Switch chart performance metrics: Select the metrics to be displayed on the chart by clicking the **Metrics** icon (the plus sign). Figure 6-51 shows the available metrics for selection. When you are done, click **OK**.

Select Chart Metrics 3 Selected

Port Metrics (3)

Data Rate (MiB/s)	<input type="checkbox"/> Send	<input type="checkbox"/> Receive	<input checked="" type="checkbox"/> Total
Bandwidth (%)	<input type="checkbox"/> Send	<input type="checkbox"/> Receive	<input type="checkbox"/> Overall
Other (cnt/s)	<input type="checkbox"/> Total Port Error ...		
More			
I/O Rates			
Frame Rate (frms/s)	<input type="checkbox"/> Send	<input type="checkbox"/> Receive	<input type="checkbox"/> Total
Error Rates			
Frame Errors (cnt/s)	<input checked="" type="checkbox"/> Error Frames	<input type="checkbox"/> CRC Errors	<input type="checkbox"/> Short Frames
	<input type="checkbox"/> Long Frames	<input type="checkbox"/> F-BSY Frames	<input type="checkbox"/> F-RJT Frames
	<input checked="" type="checkbox"/> Discarded Class 3...	<input type="checkbox"/> Bad EOF CRC Error...	
Port Protocol Errors	<input type="checkbox"/> Discarded Frames ...	<input type="checkbox"/> Link Reset Transm...	<input type="checkbox"/> Link Reset Receiv...
	<input type="checkbox"/> Zero Buffer Credi...	<input type="checkbox"/> Class 3 Send Time...	<input type="checkbox"/> Class 3 Receive T...
	<input type="checkbox"/> Credit Recovery L...	<input type="checkbox"/> RDY Priority Over...	<input type="checkbox"/> Port Congestion I...
Link Errors (cnt/s)	<input type="checkbox"/> Link Failures	<input type="checkbox"/> Signal Loss	<input type="checkbox"/> Sync Loss

OK **Cancel**

Figure 6-51 Selecting performance chart metrics

- b. Resources: By default, the performance view is loaded with all of the resources (switches) preselected, meaning all of them will be charted simultaneously. You can select a single switch under the Resources pane to chart only the performance metrics for that particular switch, as shown in Figure 6-52. To select multiple switches, press and hold the **Ctrl** key while you are selecting the resources that you want.



Figure 6-52 Single resource performance chart

- c. Time range: You can select the time range for the metrics to be charted by clicking either of the two time stamps on the bottom sides of the chart.
- d. Data resolution: By default, the charts are generated with a data resolution value of Sample, which means each value that is charted corresponds to a performance data collection sample for that switch device. The granularity of the samples is defined at the time that the Performance Data Collection is configured for each switch device. By default, IBM Spectrum Control collects 1-minute samples. However, this value can be changed, which affects the representation of the samples on the performance charts. Other possible values for Data Resolution are **Hourly** and **Daily**, in which cases they represent samples average values for each hour and day.
4. By manipulating the three icons in the upper right corner of the chart, you can do the following tasks:
- Toggle the metrics display between chart (default) and table.
 - Export the information to a CSV file.
 - Open the Performance view tab on a separate window (detach).

For more information about performance monitoring, see the IBM Spectrum Control documentation in the IBM Knowledge Center at the following link:

<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>

Additionally, you can refer to the User and Administrator guides that can be found at the following link:

ftp://public.dhe.ibm.com/software/tivoli/tpc/v527/Tivoli_Storage_Productivity_Center/English/



Fabric Vision

Fabric Vision functionality is part of the Gen 5 Fibre Channel technology. Fabric Vision provides rich set of tools that provides extra management, monitoring, and diagnostic tools.

This chapter includes the following sections:

- ▶ Fabric Vision
- ▶ ClearLink Diagnostics Port
- ▶ Bottleneck detection
- ▶ Buffer credit depletion and recovery
- ▶ Fabric Performance Impact monitoring
- ▶ Managing Forward Error Correction
- ▶ Monitoring, Alerting, and Performance Suite

7.1 Fabric Vision

IBM b-type Fabric Vision technology is an advanced hardware and software architecture. It combines capabilities from Fabric Operating System (FOS), b-type Gen 5 devices, and IBM Network Advisor to help administrators address problems before operations are affected and accelerate new application deployments.

The Fabric Vision license includes support for Flow Vision, Monitoring Alerting Policy Suite (MAPS), Clear Link Diagnostic Port, Fabric Watch, and Advanced Performance Monitoring. Switches with Fabric Watch (FW) and Advanced Performance Monitoring (APM) licenses automatically obtain FOS V7.2.x Fabric Vision license features simply by upgrading to FOS V7.2.x. Switches with only FW or APM can upgrade to Fabric Vision by purchasing other licenses.

Fabric Vision technology includes the following features:

- ▶ ClearLink diagnostic tests: Ensure optical and signal integrity for Gen 5 Fibre Channel optics and cables.
- ▶ Credit Loss Recovery: Helps overcome performance degradation and congestion due to buffer credit loss.
- ▶ Bottleneck Detection: Identifies and alerts administrators to device or ISL congestion, and abnormal levels of latency in the fabric. This feature works with Brocade Network Advisor to automatically monitor and detect network congestion and latency in the fabric.
- ▶ Monitoring, Alerting, and Performance Suite (MAPS): Provides a new, easy-to-use solution for policy-based threshold monitoring and alerting. MAPS proactively monitors the health and performance of the SAN infrastructure to ensure application uptime and availability. By utilizing both pre-built rules and policy-based templates, MAPS simplifies fabric-wide threshold configuration, monitoring, and alerting. Administrators can configure the entire fabric (or multiple fabrics) at one time by using common rules and policies, or customize policies for specific ports or switch elements through Brocade Network Advisor.
- ▶ Fabric Performance Impact (FPI) monitoring: Provides the ability to automatically mitigate the effects of slow drain devices or even resolve the slow drain behavior at the source.

Note: FPI now provides improved latency monitoring and detection. It is the preferred tool to replace the credit loss and bottleneck detection tools in FOS 7.3.x and later.

- ▶ Forward Error Correction (FEC): Enables recovery from bit errors in inter-switch links (ISLs), enhancing transmission reliability and performance.

7.1.1 Flow Vision

Flow Vision enables administrators to identify, monitor, and analyze specific application flows to preform troubleshooting, maximize performance, avoid congestion, and optimize resources. Flow Monitor is integrated with Brocade MAPS to enable threshold-based monitoring and alerting of flows. Flow Vision includes these features:

- ▶ Comprehensive visibility into application flows in the fabric, including the ability to learn (discover) flows automatically.
- ▶ Monitoring of application flows within a fabric at a specific port.
- ▶ Pre-defined flow to discover all application flows going through all device ports on a switch for network provisioning and planning.

- ▶ Statistics that are associated with the specified flows to gain insights into application performance, including the following statistics:
 - Transmit frame count
 - Receive frame count
 - Transmit throughput
 - Receive throughput
 - SCSI Read frame count
 - SCSI Write frame count
 - Number of SCSI Reads and Writes per second (IOPS)
- ▶ When NPIV is used on the host, users can monitor Virtual Machine (VM)-to-LUN-level performance.
- ▶ Monitoring of various frame types at a switch port to provide insights into the storage I/O access pattern at the LUN level, reservation conflicts, and I/O errors.
- ▶ MAPS to enable threshold-based monitoring and alerting of flows.
- ▶ Flow Generator, which is a built-in traffic generator for pre-testing and validating the data center infrastructure including route verification and integrity of optics, cables, ports, back-end connections, and ISLs before deployment.

Note: Flow Vision is covered in greater detail in 12.9, “Flow Vision” on page 322. It offers many performance-related tools that work together with Fabric Vision and more specifically MAPS to provide a rich set of performance monitoring and troubleshooting tools.

7.2 ClearLink Diagnostics Port

ClearLink Diagnostic Port (D_Port) mode allows you to convert a Fibre Channel port into a diagnostic port for testing link traffic and running electrical loopback and optical loopback tests. The test results can be useful in diagnosing various port and link problems.

D_Port functionality is supported only on 16 Gbps-capable platforms, running Fabric OS 7.0.0 or later. The ports must use 10 Gbps or 16 Gbps Brocade-branded small form-factor pluggable (SFP) transceivers. It is also supported on 8-Gbps LWL SFP, 8-Gbps ELWL SFP, ICL ports, QSFP, and QSFP+ ports.

A Fabric Vision license is required to use Clear Link Diagnostics for third-party adapters and controllers.

Note: Starting with Fabric OS 7.4.0, D_Port tests can be performed if the Fabric Watch and Performance Monitor feature licenses are present.

Table 7-1 shows switches and FOS levels that support ClearLink Diagnostics.

Table 7-1 Switches and minimum FOS levels that support ClearLink Diagnostics

IBM Model	Machine type	Brocade Model	Minimum FOS
SAN768B-2	2499-816	DCX 8510-8	v7.0.0
SAN384B-2	2499-416	DCX 8510-4	v7.0.0
SAN24B-5	2498-F24	Brocade 6505	v7.0.1
SAN48B-5	2498-F48	Brocade 6510	v7.0.0
SAN96B-5	2498-F96	Brocade 6520	v7.1.0
SAN42B-R	2498-R42	Brocade 7840	v7.3.0

7.2.1 Enabling D_port Diagnostics by using IBM Network Advisor

Enabling D_port diagnostics by using IBM Network Advisor is accomplished by selecting **Monitor** → **Troubleshooting** → **FC** → **Diagnostic Port Test**.

Figure 7-1 shows D_Port diagnostics being enabled.

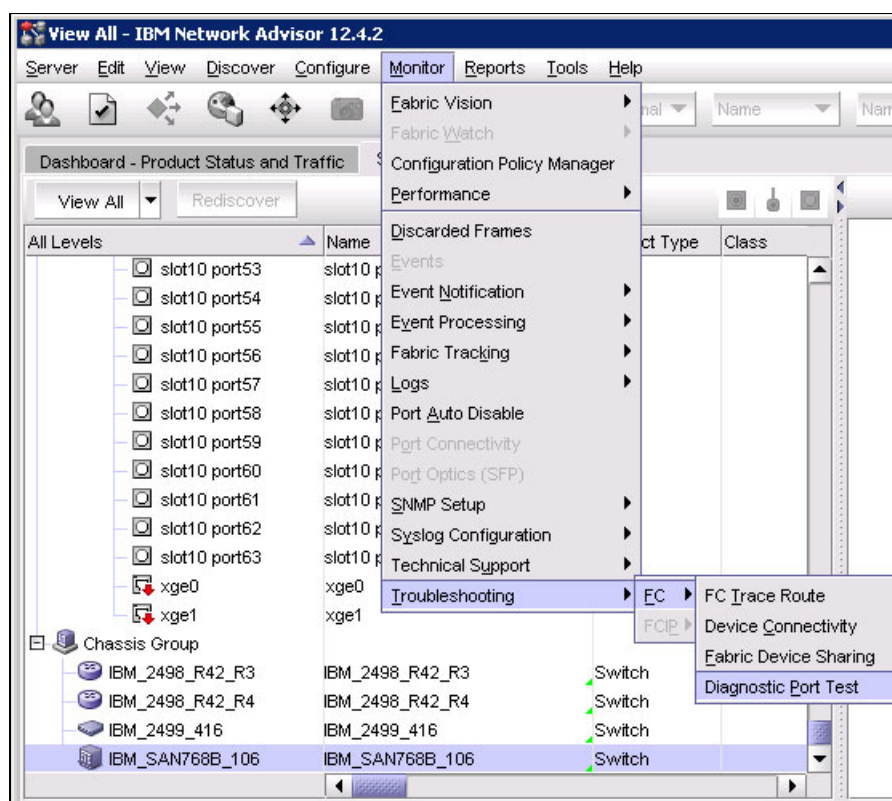


Figure 7-1 D_Port diagnostics being enabled

7.2.2 Selecting and running D_port diagnostics

Select D_port diagnostics testing and the Diagnostic Port Test panel is displayed. All ports eligible for D_Port testing are presented in the left pane. Individual ports are selected from that pane by completing these steps:

1. Highlight the port and click the **Left** arrow. Conversely, to remove a port, highlight the port in the Selected Ports pane and click the **Right** arrow.
2. To start the test select the **Start** button. The results are displayed in the lower right pane.
3. When all tests are complete, remove ports from the selected ports pane and click **Close** to close the panel.

Figure 7-2 shows the D_Port testing panel with port added and testing results in the lower pane.

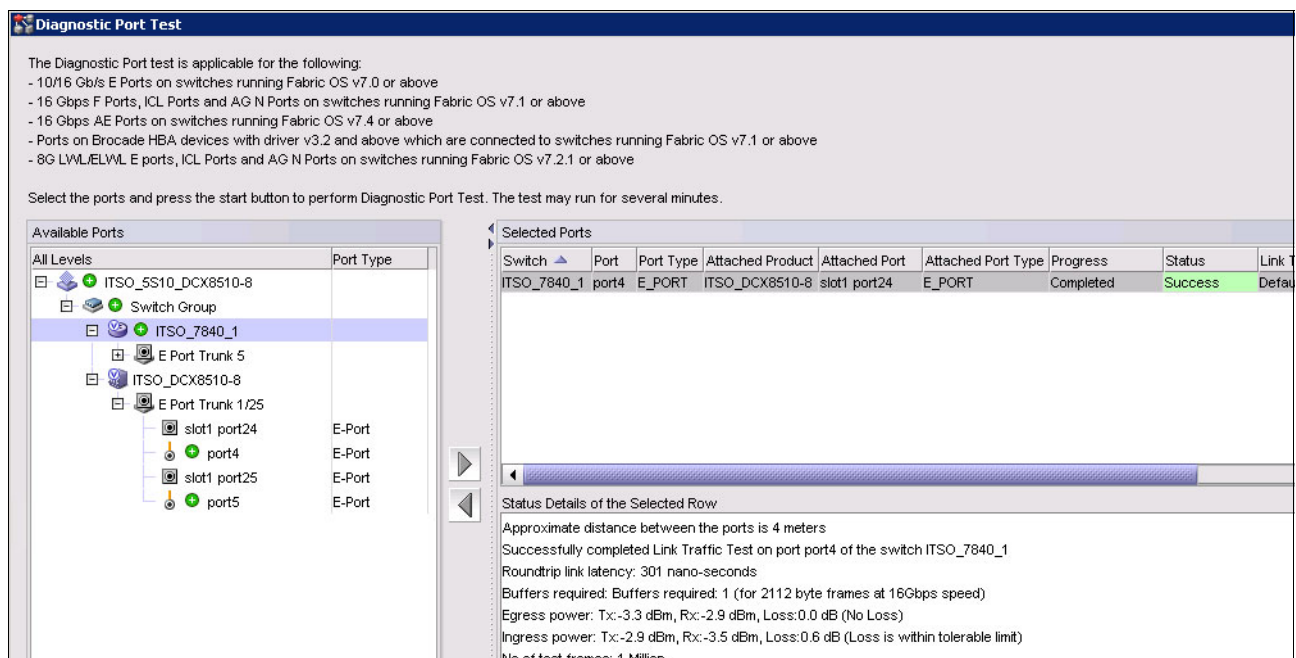


Figure 7-2 Shows the D_Port testing panel with switch Port added and testing results

Note: For ClearLink use and detailed configuration information that includes sample test scenarios, see the FOS Administration Guide at:

<http://www.brocade.com/content/html/en/administration-guide/fos-741-adminguide/GUID-2A84645C-8F20-44B9-AA4F-BC20D438EFF2.html>

7.3 Bottleneck detection

Bottleneck Detection identifies and alerts administrators to device or ISL congestion, and abnormal levels of latency in the fabric. When it is applied to F_Ports, Bottleneck Detection can continuously monitor for medium or high levels of latency on a device port and provide notification about the nature and duration of the latency. Bottleneck Detection can also serve as a confirmation about host information when storage latencies are suspected as the cause of poor host performance.

The reverse (eliminating the storage as the source of poor performance) is also true. When applied to E_Ports, Bottleneck Detection can alert administrators when it detects high levels of latency on an ISL. High levels are often the result of congestion or latency from elsewhere in the fabric, but also can be a condition that can occur as a result of device latencies from multiple flows.

Network Advisor works with Bottleneck Detection to automatically monitor and detect network congestion and latency in the fabric, providing visualization of bottlenecks in a connectivity map and product tree. Network Advisor can also show exactly which devices and hosts are affected by a bottlenecked port.

FOS v6.3 marked the initial release of Bottleneck Detection and later enhanced in v6.3.1b to include latency detection for F_Ports. Alerting is accomplished through RASlog entries only by producing a RASlog message (AN-1003) when a latency threshold is exceeded. The message has a severity level of WARNING.

7.3.1 Enabling, displaying, and disabling bottleneck detection

Enabling, displaying, and viewing bottleneck detection in the command line interface is accomplished by using the commands below. These commands set the alerts with default parameters that can be altered as needed based on the environment and level of monitoring wanted.

The **bottleneckmon --status** command returns information about whether the monitor is enabled and what parameters are set as shown in Example 7-1.

Example 7-1 Results of the bottleneckmon --status

```
switch:admin> bottleneckmon --status
Bottleneck detection - Enabled
=====

Switch-wide sub-second latency bottleneck criterion:
=====
Time threshold           - 0.800
Severity threshold       - 50.000

Switch-wide alerting parameters:
=====
Alerts                   - Yes
Latency threshold for alert - 0.100
Congestion threshold for alert - 0.800
Averaging time for alert   - 300 seconds
Quiet time for alert       - 300 seconds
```

The **bottleneckmon [--enable |--disable] -alert** command enables or disables the monitor with the default settings. If no response is returned after running these commands, the CLI will return only a prompt. See Example 7-2.

Example 7-2 Bottleneck monitor being enabled and disabled for all ports on a switch

```
switch:admin>bottleneckmon --enable -alert
switch:admin>
switch:admin>bottleneckmon --disable -alert
switch:admin>
```

When viewing the statistics, the **--show** command reports the number of bottlenecked ports or will report that the service is not enabled if the tools are disabled or FOS v7.4 or greater is installed. See Example 7-3.

Example 7-3 *bottleneckmon --show* command output

```
switch:admin> bottleneckmon --show
```

```
=====
                        Tue Nov 10 20:27:42 UTC 2015
=====

List of bottlenecked ports in most recent interval:
None
=====
```

From	To	Number of bottlenecked ports
Nov 10 20:27:32	Nov 10 20:27:42	0
Nov 10 20:27:22	Nov 10 20:27:32	0
Nov 10 20:27:12	Nov 10 20:27:22	0

7.3.2 Enabling bottleneck monitor in IBM Network Advisor

To enable the bottleneck monitor with IBM Network Advisor, select **Monitor** → **Performance** → **Bottlenecks**.

Figure 7-3 shows the Bottleneck monitors being enabled.

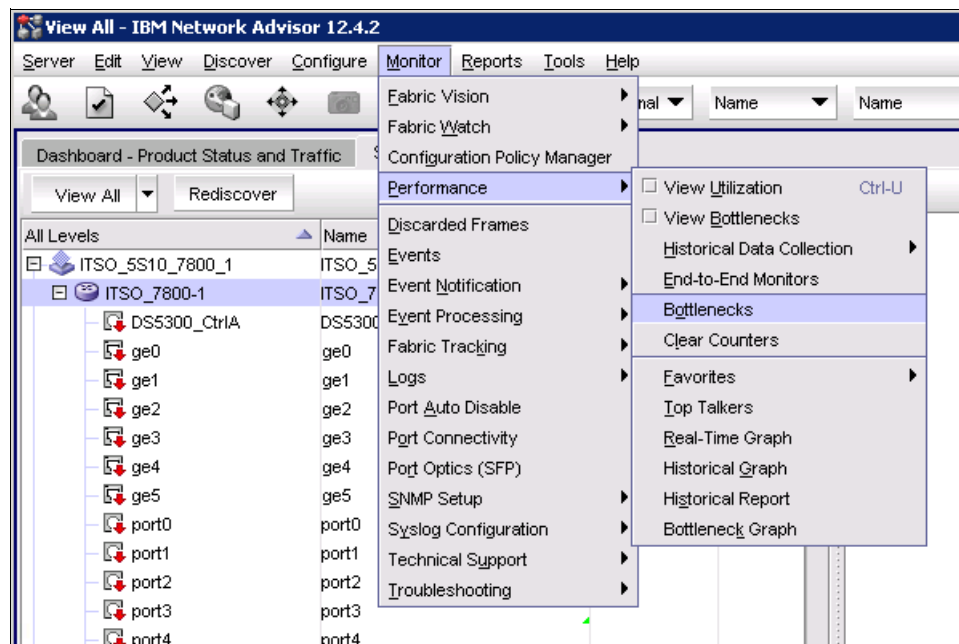


Figure 7-3 Enabling the Bottleneck monitor

7.3.3 Configuring Bottleneck monitors

After you select the Bottlenecks option, a Bottleneck monitor window is displayed. This window allows for configuration of notification thresholds on a whole switch or by ports basis.

Figure 7-4 shows the Bottlenecks configuration window.

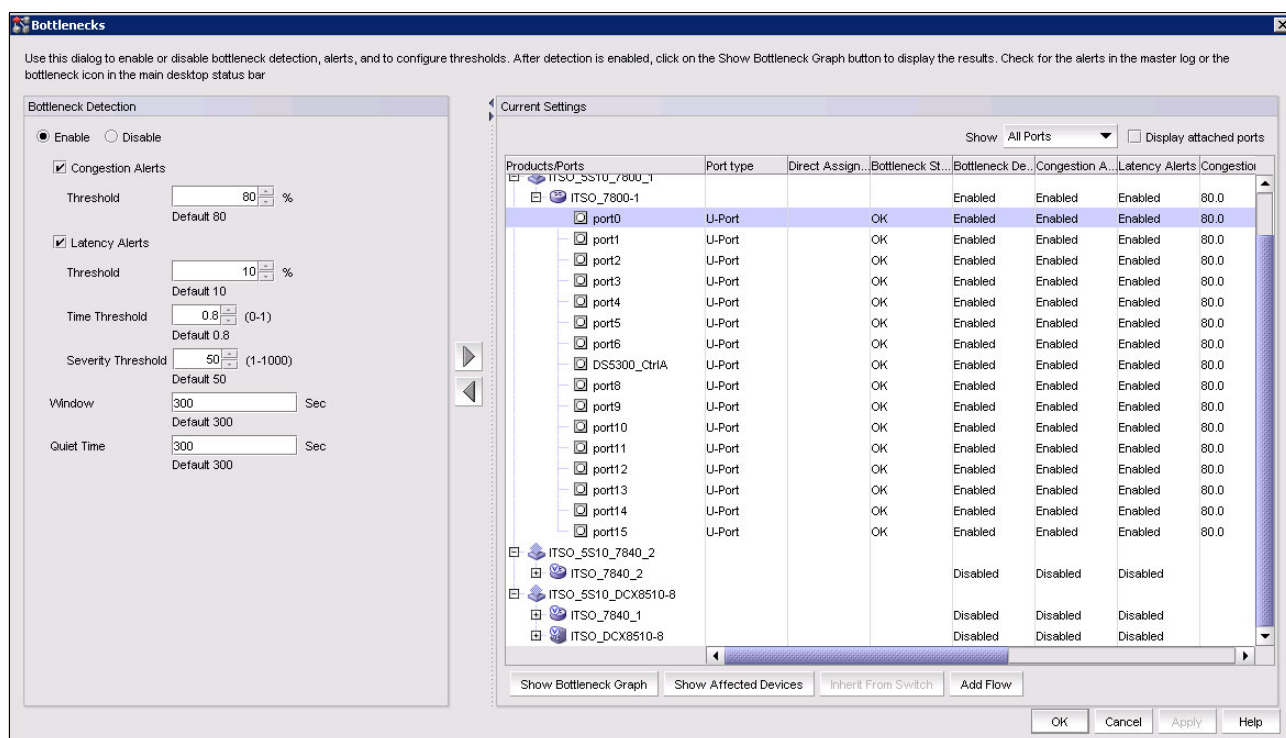


Figure 7-4 Bottlenecks configuration window

Note: The *FOS Administration Guide* contains detailed configuration instructions and recommendations. It is available on the following website:

<http://www.brocade.com/content/html/en/administration-guide/fos-741-adminguide/GUID-D41B5965-6282-417B-B8FC-9B8492578676.html>

7.3.4 Bottleneck Monitor Suggested initial settings

Field experience shows that the original strategy of enabling Bottleneck Detection with conservative values for latency thresholds almost always yields no results. There was a concern that aggressive values would result in Bottleneck Detection alert storms, but this has not been the case. Even the most aggressive values result in relatively few alerts being generated. As a result, it is now a preferred practice that the most aggressive settings are tried first and then backed off gradually if too many alerts are seen.

The following is a suggested set of graphical user interface (GUI)-based settings for all F-ports being monitored:

- ▶ Congestion 50%
- ▶ Latency 20%
- ▶ Window 60 seconds
- ▶ Quiet Time 60 seconds

The window time represents the size of the time window to look at when determining whether to alert.

Quiet time is used to throttle the frequency of the alerts between consecutive alerts.

Table 7-2 provides some suggested starting values for various levels of bottleneck monitoring.

Table 7-2 Suggested starting values for three levels of monitoring

Parameter	Conservative Setting	Normal Setting	Aggressive Setting
Congestion	0.8	0.5	0.1
Latency	0.3	0.2	0.1
Window	300	60	5
Quiet time	300	60	1

7.3.5 Displaying bottleneck statistics

Bottlenecks are reported through alerts in the Master Log. A bottleneck cleared alert is sent when the bottleneck is cleared.

Bottlenecks can be highlighted in the Connectivity Map and Product List. Select **Monitor** → **Performance** → **View Bottlenecks**. If a port is experiencing a bottleneck, a **Bottleneck** icon is displayed in the Connectivity Map for the switch and fabric, and in the Product List for the port, switch, and fabric, as shown in Figure 7-5. In example shown, port15 and port22 are bottlenecked.

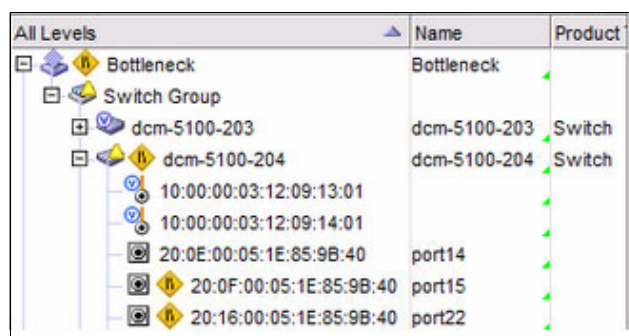


Figure 7-5 Ports icons displaying bottlenecks

Note: For detailed instructions on displaying bottleneck output, see the Brocade FOS Administration Guide at:

http://www.brocade.com/content/html/en/user-guide/networkadvisor-1242-san-manual/wwhelp/wwhimpl/common/html/wwhelp.htm#href=Ch_Performance.31.07.html&single=true

7.4 Buffer credit depletion and recovery

FOS V7.1 and later supports buffer credit recovery. This feature allows links to recover after buffer credits are depleted when the logic is enabled. The buffer credit recovery feature also maintains performance. If a credit is lost, a recovery attempt is initiated. During link reset, the frame and credit loss counters are reset without performance degradation.

Credit recovery is supported on E_Ports, F_Ports, and EX_Ports. Buffer credit recovery is enabled automatically across any long-distance connection for which the E_Port, F_Port, or EX_Port buffer credit recovery mechanism is supported.

To enable backend port credit loss recovery with the link reset only option and to display the configuration, issue these commands:

```
switch:admin> creditrecovmode --cfg onLrOnly
switch:admin> creditrecovmode --show
Internal port credit recovery is Enabled with LrOnly
C2 FE Complete Credit Loss Detection is Enabled
```

To enable back-end port credit loss recovery with the link reset threshold option and to display the configuration, issue these commands:

```
switch:admin> creditrecovmode --cfg onLrThresh
switch:admin> creditrecovmode --show
Internal port credit loss recovery is Enabled with LrThresh
C2 FE Complete Credit Loss Detection is Enabled
```

To disable back-end port credit loss recovery and to display the configuration, issue these commands:

```
switch:admin> creditrecovmode --cfg off
switch:admin> creditrecovmode --show
Internal port credit loss recovery is Disabled
C2 FE Complete Credit Loss Detection is Enabled
```

Note: In FOS 7.2 and lower, use the **bottleneckmon --cfgcredittools** command to enable buffer credit depletion and detection tools. For syntax and use, see the Fabric OS Command Reference Supporting Fabric OS v7.2.x and earlier at MyBrocade:

<https://my.brocade.com/>

7.5 Fabric Performance Impact monitoring

The FPI tool monitors the latency on E_Ports and F_Ports over different time periods and uses that to determine the performance impact to the fabric and network.

FPI monitoring is automatically enabled for new b-type switches already running Fabric OS 7.3.0 or later, or if the legacy bottleneck monitoring feature was not enabled before the switch firmware was upgraded to Fabric OS 7.3.x or 7.4.0.

Note: To use FPI monitoring, the legacy bottleneck monitoring feature cannot be enabled. You can use the **bottleneckmon --status** command to verify the bottleneck monitoring status. To disable the legacy bottleneck monitoring, use the **bottleneckmon --disable** command.

To enable FPI, if previously disabled, complete the following steps:

1. Log in to IBM Network Advisor using an ID with administrator privileges.
2. Select **Monitor** → **Fabric Vision** → **MAPS** → **Configure** and the MAPS Configuration window is displayed (Figure 7-6).

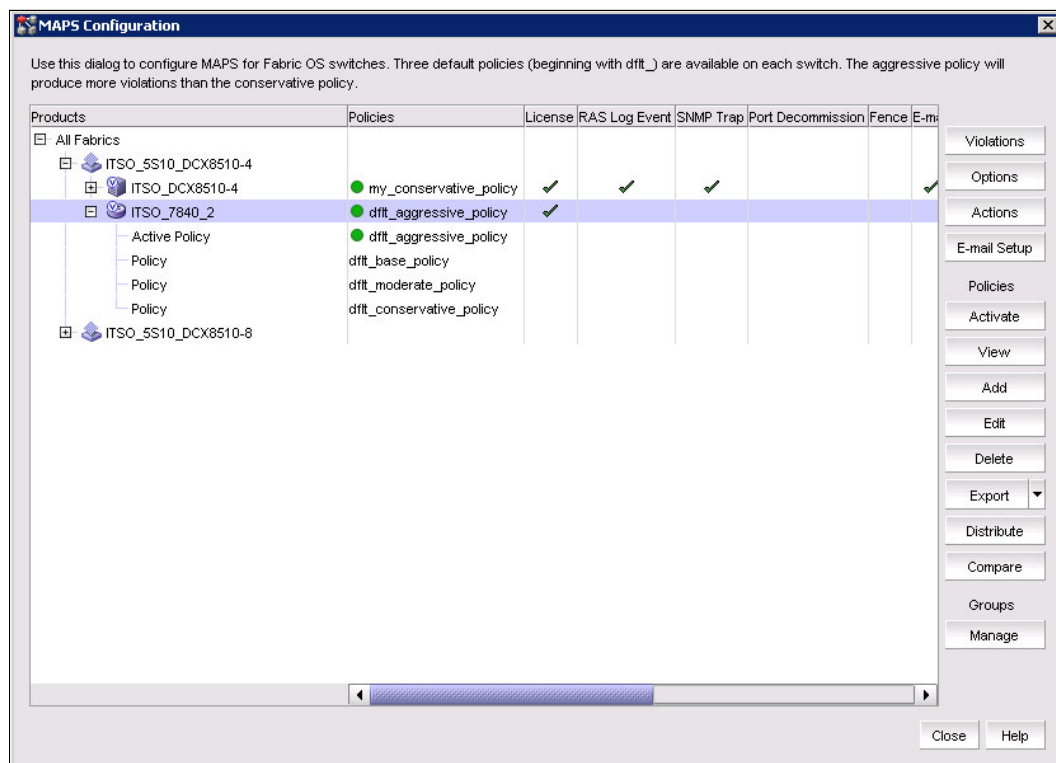


Figure 7-6 Shows the MAPS Configuration window

3. Select the **Options** button to display the MAPS Options window (Figure 7-7). Select **FPI (Fabric Performance Impact) Monitoring**.

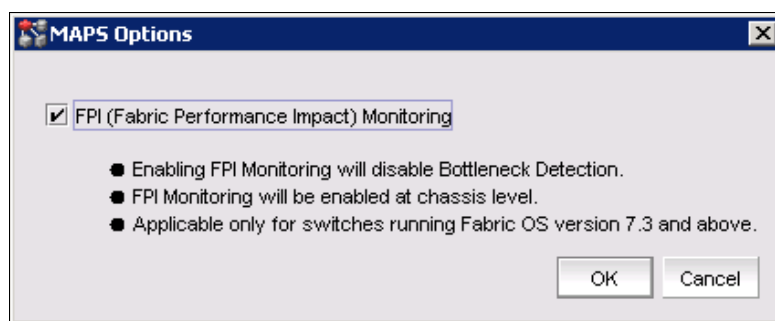


Figure 7-7 Shows the MAPS Options panel with FPI check box set

4. Select **OK** to return to the MAPS Configuration window.
5. Select **Close** to return to the main IBM Network Advisor window.

For more information about CLI commands to enable and display FPI settings, see the Fabric OS Command Reference available on the following website:

<http://brocade.com>

7.6 Managing Forward Error Correction

FEC is described in “Forward Error Correction” on page 16.

7.6.1 Enabling, disabling, and viewing FEC status

Important: Enabling FEC is disruptive to traffic.

To enable FEC on a single port and to display the configuration, issue these commands:

```
switch:admin> portcfgfec --enable -FEC 5/28
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
switch:admin> portcfgfec --show 5/28
Port: 412
FEC Capable: YES
FEC Configured: ON
FEC via TTS Configured: OFF
FEC State: active
```

To disable the FEC feature on a port range, issue these commands:

```
switch:admin> portcfgfec --disable -FEC 0-8
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
```

To disable the FEC feature on a port range, issue these commands:

```
switch:admin> portcfgfec --disable -FEC 0-8
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
```

Note: Additional use and syntax can be found in the Fabric OS Command Reference v7.4.1 Guide at:

<https://my.brocade.com/>

7.7 Monitoring, Alerting, and Performance Suite

MAPS provides an easy-to-use solution for policy-based threshold monitoring and alerting. MAPS proactively monitors the health and performance of the SAN infrastructure to ensure application uptime and availability. By utilizing both pre-built rules and policy-based templates, MAPS simplifies fabric-wide threshold configuration, monitoring, and alerting. Administrators can configure the entire fabric (or multiple fabrics) at one time by using common rules and policies, or customize policies for specific ports or switch elements through Brocade Network Advisor.

MAPS replaces FW as the preferred method of fabric monitoring, and cannot coexist with FW. The FW configuration must be migrated to a MAPS compatible format.

In Fabric OS 7.4.0, MAPS is no longer optional. It replaces Fabric Watch, and provides a set of monitors that work even if the MAPS license is not activated. MAPS monitors field-replaceable units (FRUs), environmental and switch resources (such as memory), and so on.

Important: *MAPS activation is a non-reversible process.* After it is enabled, it is enabled for any version of Fabric OS 7.2.0 or later. It is not required to be enabled at FOS 7.4.0 or later. Fabric Watch is no longer an option. If you want to use a monitoring and alerting tool, you *must* enable MAPS.

The Fabric Watch feature is not part of Fabric OS 7.4.0, and only becomes active if you downgrade to a version of Fabric OS before 7.2.0. This process enables Fabric Watch with its last configured settings. These settings will be the defaults if no custom settings are available.

For detailed migration instructions, see the *Monitoring and Alerting Policy Suite Administrator's Guide*, which is available on the following website:

<http://www.brocade.com/content/dam/common/documents/content-types/administration-guide/fos-740-maps.pdf>

7.7.1 Enabling MAPS

If Fabric Watch is not in use, or a fresh configuration that ignores previous Fabric Watch policies is wanted, MAPS can quickly start monitoring a switch using MAPS with one of the predefined policies that are available in MAPS.

If Fabric Watch is in use, the policies can be converted MAPS policies and will use the same thresholds.

Note: *Fabric Watch policies must be converted into MAPS policies before you install Fabric OS 7.4.0.* MAPS will then automatically use the policy named fw_active_policy to provide the same monitoring functionality as Fabric Watch is using.

For instructions on how to convert Fabric Watch Policies into MAPS Policy, see “Converting Fabric Watch policies to MAPS policies” on page 204.

To enable MAPS in IBM Network Advisor with a default policy, select the menu options **Monitor → Fabric Vision → MAPS → Enable**.

Figure 7-8 shows MAPS being enabled by using IBM Network Advisor.

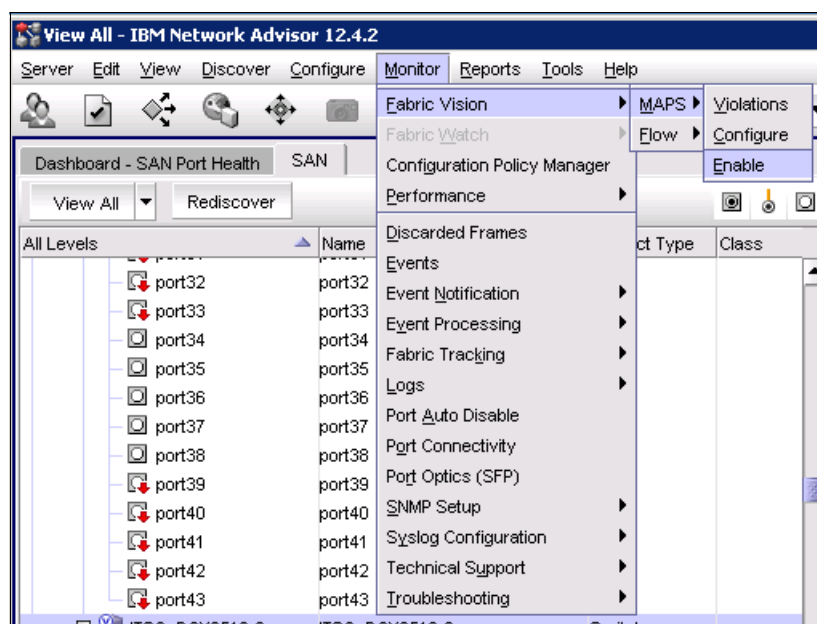


Figure 7-8 Enabling MAPS by using IBM network Advisor

Maps can also be enabled by using the command-line interface (CLI). Detailed instructions and examples are available in the *Monitoring and Alerting Policy Suite Administrator's Guide* for Fabric OS v7.4.0, which is available on the following website:

<https://my.brocade.com/>

Converting Fabric Watch policies to MAPS policies

Note: To retain the Fabric Watch policies and thresholds, this conversion must be completed before you install Fabric OS 7.4.0.

To convert an existing Fabric Watch policy to a MAPS policy, complete the following steps:

1. Start a CLI session to the switch using an ID with administrative privileges.
2. Enter the command `mapsconfig --fwconvert`.

The command creates a MAPS policy named `fw_active_policy` that can then be applied when MAPS is enabled as described in 7.7.2, "Configuring MAPS with Fabric Watch Rules" on page 204.

Note: This command converts the Fabric Watch policy only. MAPS must then be enabled and have the converted policy applied.

7.7.2 Configuring MAPS with Fabric Watch Rules

To continue using the same Fabric Watch thresholds in MAPS, the FW rules must be converted into MAPS policies before you install FOS V7.4.0. MAPS will then use the policy named `fw_active_policy` to provide the same monitoring rules.

To enable MAPS and apply the converted policy, complete the following steps:

1. Log in to IBM Network Adviser with an ID with administrative privileges.
2. Select **Monitor** → **Flow Vision** → **MAPS** → **Enable** to display the Enable MAPS window as shown in Figure 7-9.

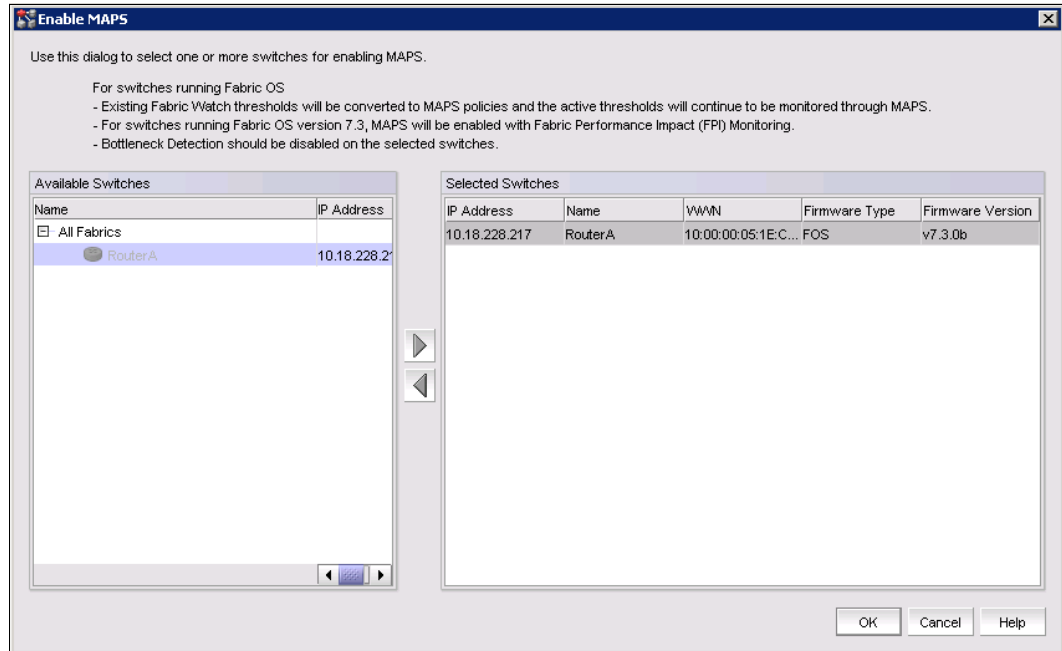


Figure 7-9 Enable MAPS window

3. Select the switch or switches from the Available Switch pane and click the **Right Arrow** to move it to the selected switches pane.
4. Click **OK** to start the MAPS enable process. A status panel is displayed when it is complete as shown in Figure 7-10.

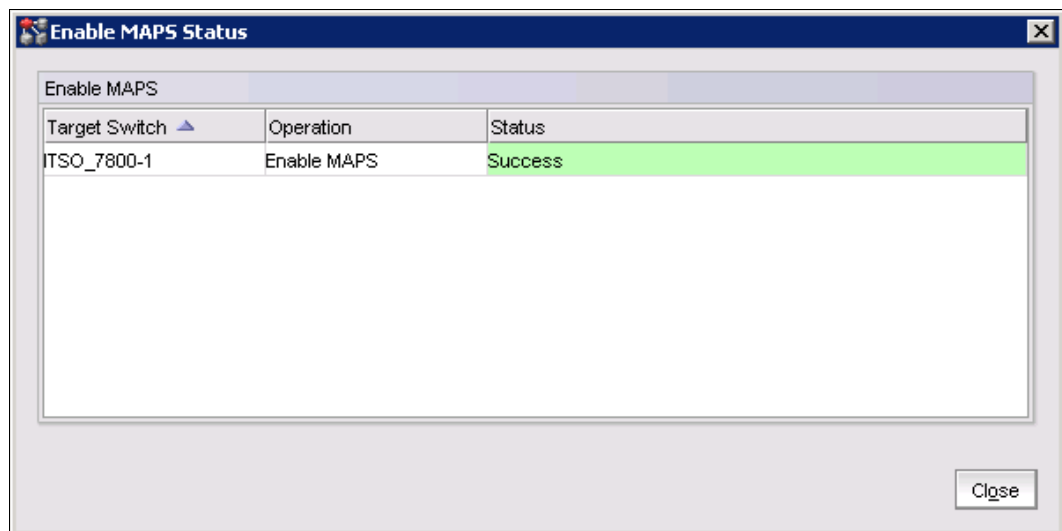


Figure 7-10 Enable Maps Status window

Note: If the switch is not at or above Fabric OS v7.2.0, IBM Network Advisor displays a message that indicates that the switch is not available for conversion.

5. Click **Close** return to the main IBM Network Advisor window.

When MAPS is enabled with an existing converted Fabric Watch policy, the policy is automatically applied and enabled on the switch that has been migrated to MAPS. To confirm that the policy has been enabled, select **Monitor** → **Fabric Vision** → **MAPS** → **Configure**. The MAPS Configuration window is displayed as shown in Figure 7-11.

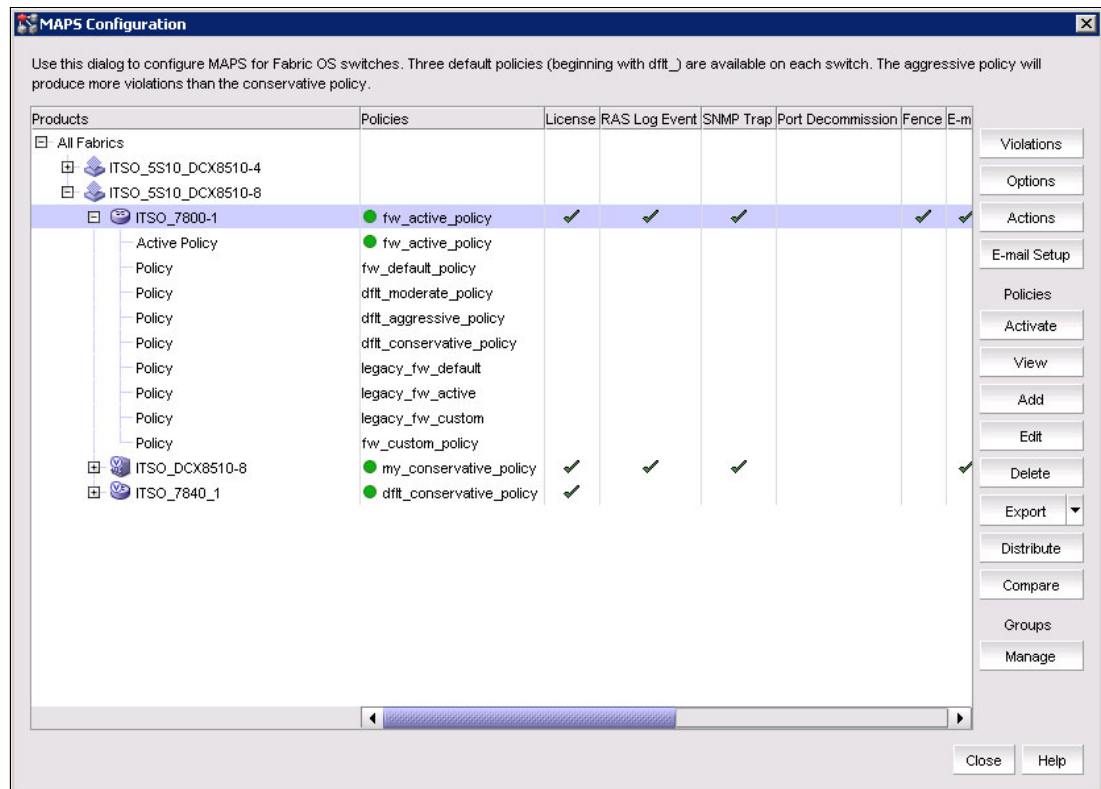


Figure 7-11 MAPS Configuration window with the "fw_active_policy" applied

If the `fw_active_policy` is present for the switch, it is automatically applied to the switch when MAPS is enabled.

7.7.3 MAPS Slow Drain Device Quarantine

In a fabric, many flows share a link or virtual circuit (VC). However, the credits that are used to send traffic or packets across the link are common to all the flows that are using the same link. Because of this structure, a slow-draining device might slow down the return of credits and have a negative effect on the healthy flows through the link.

Figure 7-12 illustrates how this process works. In the figure, the inability of the slow-draining device (1) to clear frames quickly enough is causing backpressure not just to the edge device (2), but also to the core (3) and the appliance (4) that is trying to reach the slow-draining device.




Figure 7-12 Slow drain device

To remedy this effect, Slow Drain Device Quarantine (SDDQ) was created. This feature, with Quality of Service (QoS) monitoring, allows MAPS to identify a slow-draining device and quarantine it by automatically moving all traffic destined to the F_Port that is connected to the slow-draining device to a low-priority VC so that the traffic in the original VC does not experience backpressure.

Enabling and disabling actions at a global level allows you to configure rules with stricter actions, such as port fencing, but disable the action globally until you can test the configured thresholds. After you validate the thresholds, you can enable port fencing globally without having to change all of the rules.

Note: For SDDQ to take effect, the Fabric Vision license must be installed on the switch where the slow draining device is detected, and on the switch where the quarantine action is to occur. Intermediate switches do not need the Fabric Vision license for this feature to work, but they must have QoS enabled on all ISLs.

To enable SDDQ, complete the following steps:

1. Log in to IBM Network Advisor by using an account with administrative privileges.
2. Select **Monitor** → **Flow Vision** → **MAPS** → **Configure** to display the MAPS Configuration window.

3. Select **All Fabrics** in the Products window and click **Actions**. The MAPS Policy Actions window is displayed as shown in Figure 7-13.

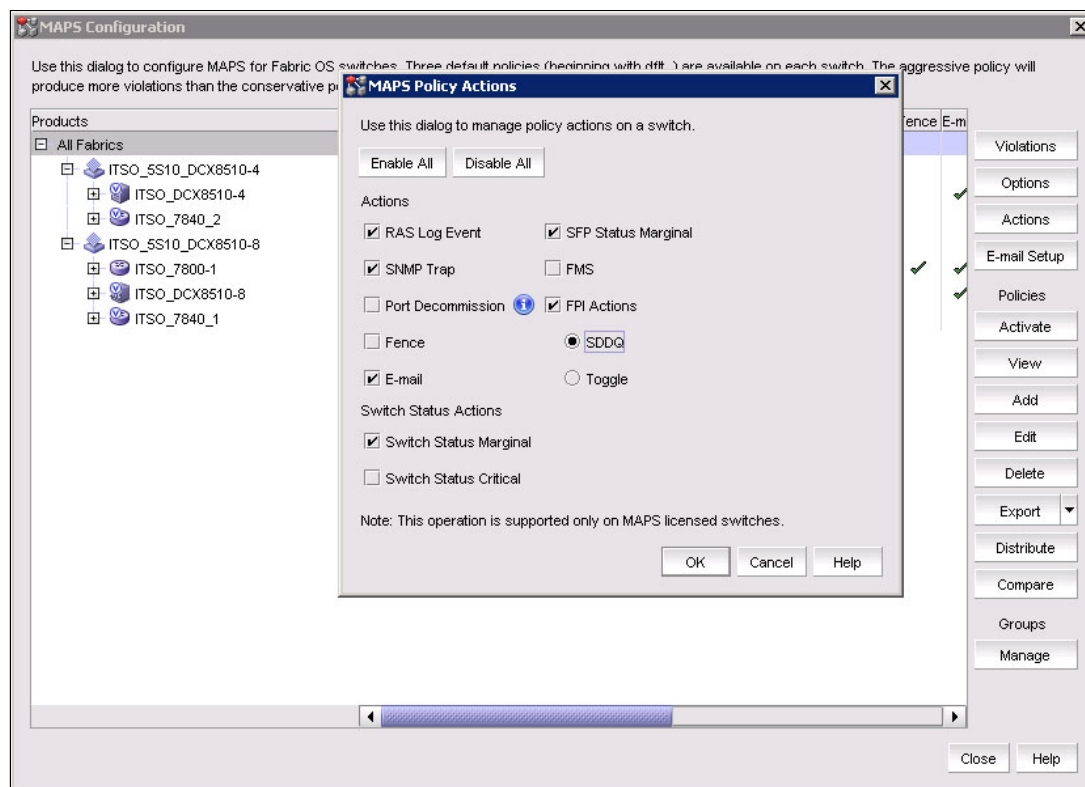


Figure 7-13 MAPS Policy Action panel

4. Select **FPI Actions**, and then select the **SDDQ** radio button.
5. Click **OK** on the MAPS Policy Actions window.
6. Click **Close** on the MAPS Configuration window.

SDDQ can be disabled in the same manner by clearing the **FPI Actions** check box.

7.7.4 MAPS port fencing

MAPS supports port fencing and port decommissioning for both E_Ports and F_Ports. These actions automatically take ports offline when configured thresholds in a specific rule are exceeded. Port fencing immediately takes ports offline, which might cause loss of traffic. Port decommissioning takes a port offline more slowly, but without loss of traffic. Both are disabled by default. Port decommissioning and port fencing can only be configured for the port health monitoring system rules, which affect the monitoring systems for which decommissioning is supported.

Port fencing is enabled by using the MAPS Policy Actions window. For instructions on how to access this panel, see 7.7.3, “MAPS Slow Drain Device Quarantine” on page 206. When the panel is open, select **Fence** under the Actions options. Figure 7-14 shows the **Fence** option selected.

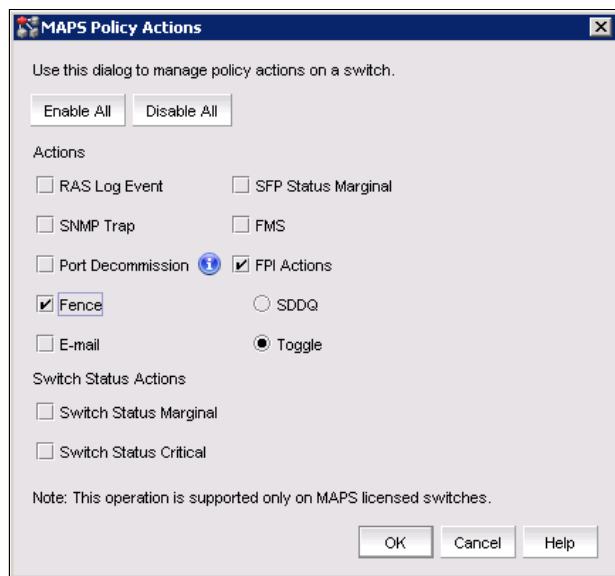


Figure 7-14 MAPS Policy Actions panel with the fenced option selected

7.7.5 MAPS Toggling and decommissioning

MAPS supports port fencing and port decommissioning for both E_Ports and F_Ports. These actions automatically take ports offline when configured thresholds in a rule are exceeded. Port fencing immediately takes ports offline, which might cause loss of traffic. Port decommissioning takes a port offline more slowly, but without loss of traffic. Both are disabled by default. Port decommissioning and port fencing can only be configured for the port health monitoring system rules, which affect the monitoring systems for which decommissioning is supported.

MAPS supports port toggling, which allows Fabric OS to recover a port that has been bottlenecked by a target device. Although there are many reasons why the target device could be bottlenecked, one of the most common is a temporary glitch in an adapter or its software.

Port toggling in MAPS temporarily disables a port and then enables it again, allowing the port to reset and recover from the issue. If the port does not recover, Fabric OS suspends the port, forcing the port traffic to switch over to a different path if one is available.

To enable recovering ports by using port toggling, MAPS assumes that there is a redundant path to the target device. It does not check to see whether there is one, nor can it check to see whether traffic to or from the target device has been switched over to a redundant path. MAPS also assumes that while the port is being toggled, the operational state of the port will not be changed by any other mechanism, such as an administrator disabling or moving the port, or a port fencing operation.

Port Toggling and decommissioning are enabled by using the MAPS Policy Actions window. For instructions on how to access this window, see 7.7.3, “MAPS Slow Drain Device Quarantine” on page 206. When the panel is open, select the **Toggle** check box under the FPI options or the **Port Decommission** check box under the Actions options. Figure 7-14 on page 209 shows the **Fence** option selected. Figure 7-15 shows the MAPS Policy Actions window with the toggle and decommission options selected.

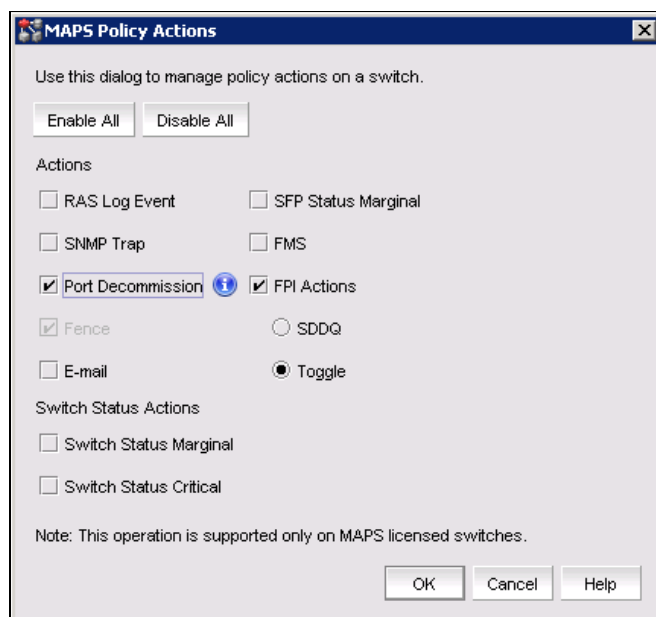


Figure 7-15 MAPS Policy Actions window with Toggle and Port Decommission options selected

7.7.6 MAPS Dashboard

The MAPS dashboard provides a summary view of the switch health status that allows administrators to easily determine whether everything is working according to policy or whether further investigation is required.

IBM Network Advisor provides a MAPS Widget on the Dashboard view of the main window. See Chapter 4, “IBM Network Advisor” on page 49 for more information about IBM Network Advisor views and widgets.

The widget displays the number of MAPS threshold violations for all network objects (such as ports, trunks, switches, and circuits) for all MAPS-capable devices. In addition, the MAPS dashboard widgets include the Fabric Watch threshold violations for devices with the Fabric Watch license or FC devices running Fabric OS 7.2.0 or later with the Fabric Vision license but not migrated to MAPS.

By default, the MAPS Dashboard widget that is shown in Figure 7-16 refreshes every minute. If any violations occur on fabrics in your area of responsibility (AOR) during the minute refresh time frame, the widget refreshes every 10 seconds. If you delete, discover, or unmonitor a device, the widget refreshes. MAPS violation data is stored in the database for 30 days. The system purges old data (over 30 days) every night at 12 midnight. The system also purges violations from deleted or unmonitored devices.

Category	Violation Count	Network Obj...
Fabric Health	0	0 Switches
FCIP Health	0	Circuits / Tunne
Virtual Machine Vi...	0	Virtual Machine
Switch Resources	0	0 Switches
FRU Health	10	1 Switches
Traffic Performance	0	0 Ports / Flows
Backend Port Health	0	Ports / Switches
Port Health	0	Ports / Switches
Fabric Performanc...	0	Ports / Switches
Switch Status Policy	0	0 Switches
Security Violations	0	0 Switches

Figure 7-16 Out of Range Violations Widget

The Out of Range Violations widget includes the following fields and components. It always displays whether or not there is an associated violation.

- ▶ Fabric Health
- ▶ FCIP Health
- ▶ FRU Health
- ▶ Port Health
- ▶ Backend Port Health
- ▶ Security Violations
- ▶ Switch Resources
- ▶ Switch Status Policy
- ▶ Traffic Performance
- ▶ Virtual Machine Violations

The widget offers color coding to highlight the worst severities and device counts of the violation category. The widget can also be customized to display the categories by user preference.

To access additional data from the widget, right-click any row and select **Violations** to navigate to the Violations window.

NoteD_Port: Detailed information and instructions about configuration and displaying the MAPS violations is available in the *Monitoring and Alerting Policy Suite Administrator's Guide*, which is available on the following website:

<http://www.brocade.com>



Virtual Fabrics

This chapter describes Virtual Fabrics, provides examples of how to implement this feature, and highlights a sample environment where Virtual Fabrics are deployed.

Virtual Fabrics is an architecture to virtualize hardware boundaries. The Virtual Fabrics feature allows storage area network (SAN) design and management to be done at the granularity of a port. It is done by allowing a customer to partition a single physical switch into multiple Logical Switches.

This chapter includes the following sections:

- ▶ IBM b-type Virtual Fabric
- ▶ Virtual Fabric features
- ▶ Configuring Virtual Fabrics

8.1 IBM b-type Virtual Fabric

This section explains what the IBM b-type Virtual Fabrics feature is, and how it is configured to an operational state.

8.1.1 Virtual Fabrics introduction

IBM b-type Virtual Fabrics allows IT organizations to manage IT assets by corporate function, use 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.

Physical switches can be partitioned into independently managed logical switches, each with their own data, control, and management paths.

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 nondisruptive in existing environments.

8.1.2 Logical switches and logical fabrics

This section describes some of the capabilities of logical switches and logical fabrics.

Figure 8-1 introduces logical switches and logical fabrics.

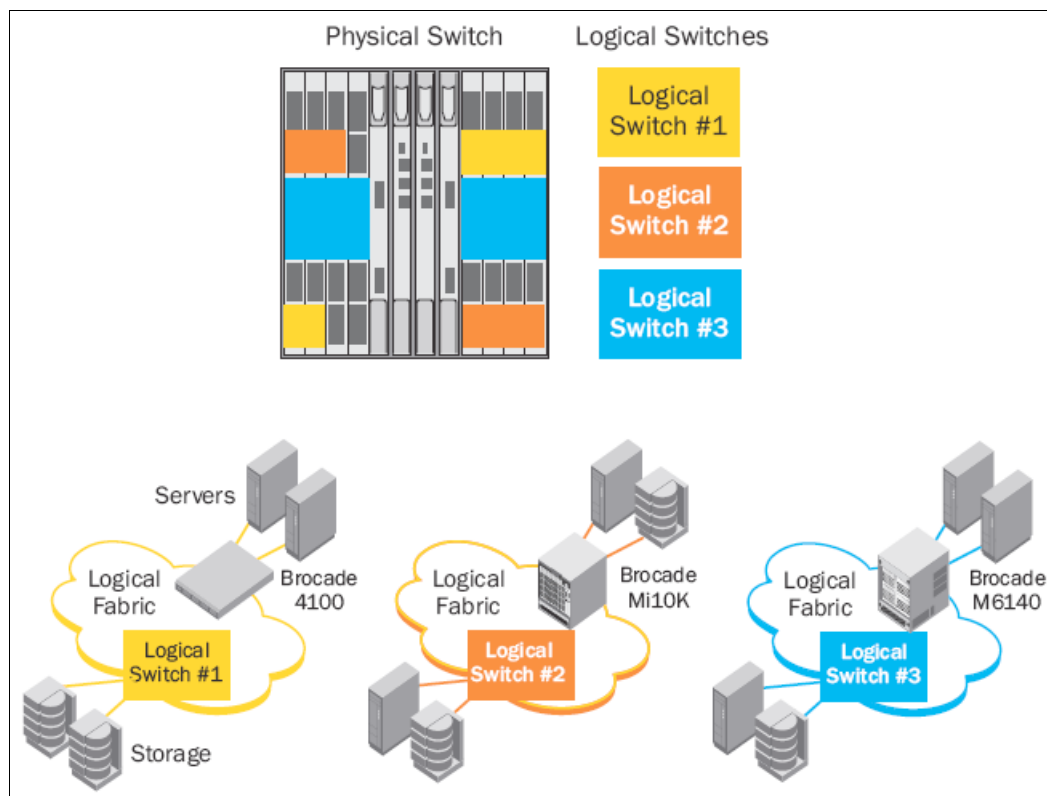


Figure 8-1 Logical switches and logical fabrics

Layer-2 traffic isolation is available with a special extended inter-switch link (XISL) shared by multiple logical fabrics, or with dedicated inter-switch link (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 devices in any edge fabric.

Virtual Fabrics was introduced with FOS 6.2.0. It is available on IBM b-type Gen 5 16 Gbps products that are *Virtual Fabrics-capable* such as the IBM SAN768B, IBM SAN384B, IBM SAN80B, the IBM SAN40B switches. For investment protection, products that are not Virtual Fabrics capable (such as earlier 2 and 4 Gbps Fabric Operating System (FOS) products) can seamlessly connect to Virtual Fabrics-capable products without requiring a reconfiguration of the existing switches. See Table 8-1.

Table 8-1 Supported logical switch creation limits

Platform	Maximum Logical Switches ^a /Chassis
SAN768B	8
SAN384B	8
SAN80B-4	4
SAN48B-5	3

a. Numbers include the Default Switch and Base Switch.

Important: On the SAN80B-4 and the SAN40B-4 switches, the Default Switch and Base Switch can be the same.

To simplify Virtual Fabrics management, organizations can use IBM Network Advisor. After they are created, logical switches and fabrics are managed the same as their physical counterparts. Alternatively, organizations can use the standard FOS command-line interface (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. These feature include frame trunking, DPS, Fibre Channel Routing, Adaptive Networking, Access Gateway, Access Gateway trunking, and Fibre Channel over IP (FCIP) for extension.

8.2 Virtual Fabric features

This section describes different features of Virtual Fabrics.

The Virtual Fabrics suite consists of the following specific features:

- ▶ Logical switch
- ▶ Logical fabric
- ▶ Device sharing

For more information, see the *Fabric OS version 7.4 Administrators Guide* at the following website:

<https://my.brocade.com>

8.2.1 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 that is inserted into the chassis initially belong to the default switch. Ports that are required by user-defined logical switches are dynamically allocated from the default switch by the chassis administrator. When 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

A *base switch* is a logical switch that is used to communicate among different logical switches. The legacy EX_port is connected to the base logical switch. Also, ISLs connected to the Base Switch are used to communicate among different fabrics. The default logical switch supports E_ and EX_ports.

Logical switch

A *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.

8.2.2 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, ICL connections, or both), supporting frame trunking and DPS. Similar to 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.

8.2.3 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. The physical XISL connection between two base switches automatically forms an LISL “tunnel” dedicated to the traffic to and from logical switches.

8.2.4 User accounts

Table 8-2 lists the predefined user accounts in Fabric OS that are available in the local switch user database.

Table 8-2 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.

8.3 Configuring Virtual Fabrics

This section describes a limited set of instructions and commands for configuring and managing Virtual Fabrics. For a complete list, see the *Fabric OS version 6.2.0 Administrators Guide and Fabric OS Command Reference Manual* available only through the Partner Network website (go to Product Documentation and register or log in):

<http://www.brocade.com/data-center-best-practices/resource-center/index.page>

Virtual Fabrics can be managed with Data Center Fabric Manager (IBM Network Advisor). This section demonstrates how to configure VF by using the standard Fabric OS v6.4.+ CLI and IBM Network Advisor.

8.3.1 Changing the context to a different logical switch

When Virtual Fabric is enabled, you want to move between the defined virtual switches. This process can be done with either Web Tools or the **setcontext** command from the CLI.

Here is one method to change the context:

1. Connect to the physical chassis and log in with an account that is assigned to the admin role.
2. Set the context to the logical switch that you want to manage (if you are not already in that context). Example 8-1 shows how to switch to FID 20.

Example 8-1 Changing the logical switch context to FID 20

```
IBM_SAN384B_27:admin> setcontext 20
```

8.3.2 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 Control processor blades (CPs) are rebooted and all EX_Ports are disabled after the reboot.

Using IBM Network Advisor

Complete the following steps to enable Virtual Fabrics:

1. Right-click the switch where VF will be configured to display the drop-down menus for the switch and select **Enable Virtual Fabric** as shown in Figure 8-2.

Requirement: SNMP V3 must be enabled and configured for management of Virtual Fabrics.

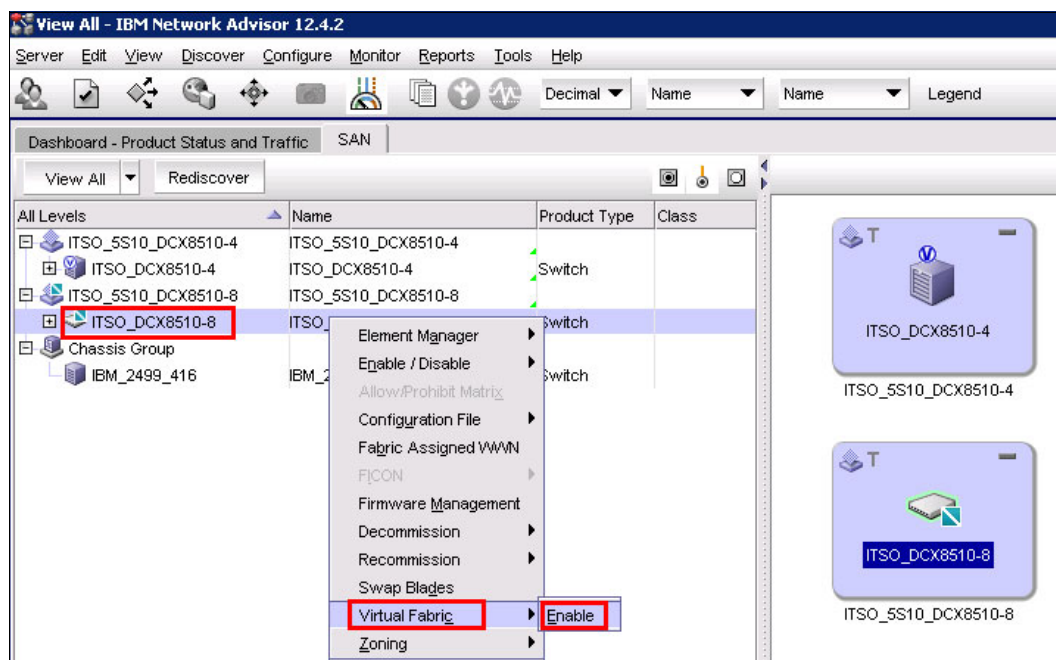


Figure 8-2 Enable Virtual Fabric

A warning message is displayed as shown in Figure 8-3.

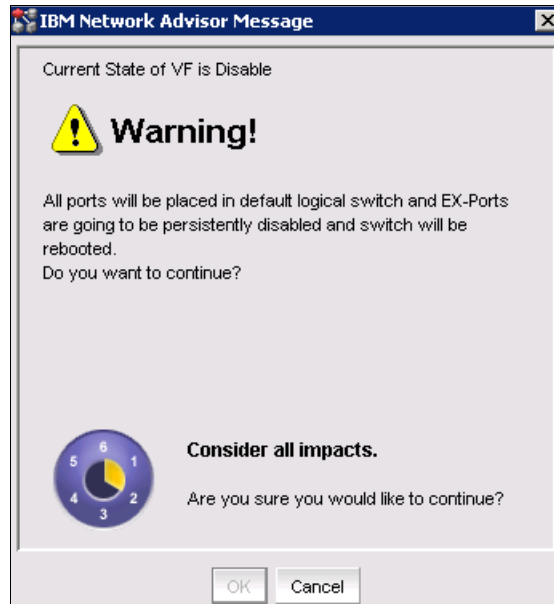


Figure 8-3 VF warning message

2. Read the warning message and select the **OK** button to continue.

When this operation is completed and the reboot is done, you will see a **V** symbol above the VF enabled switch and a chassis group appears in the product list, as shown in Figure 8-4.

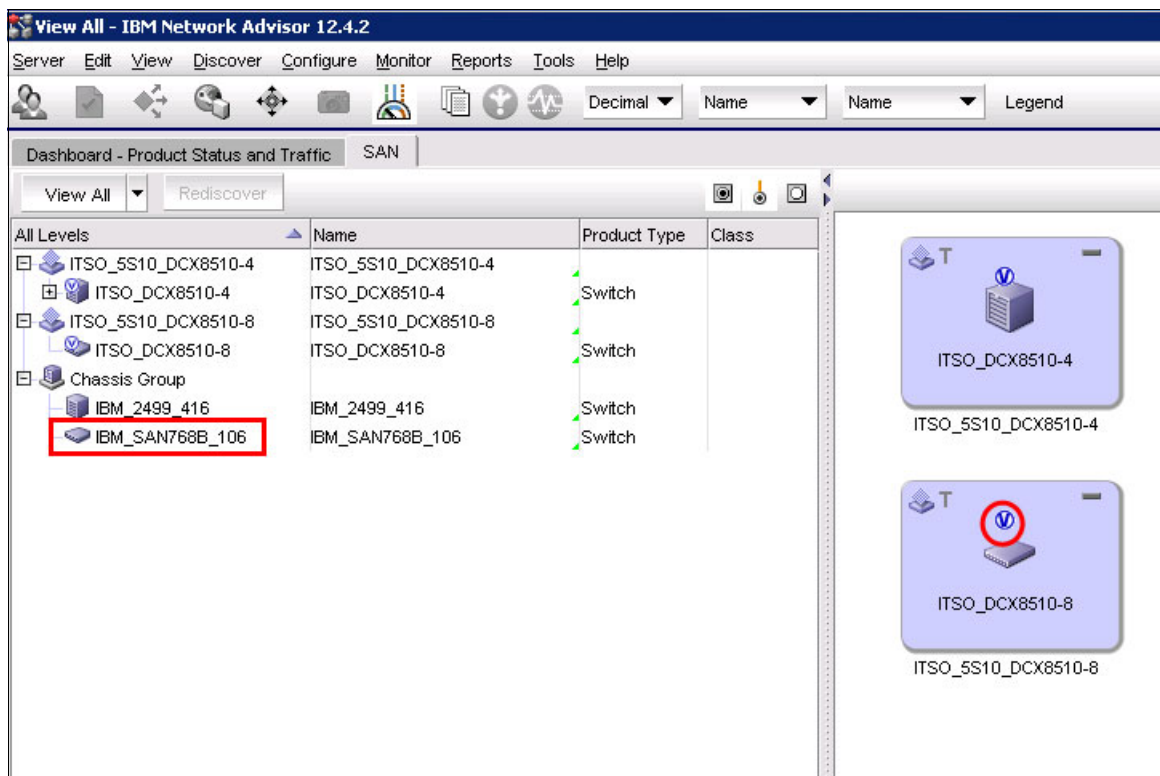


Figure 8-4 VF enabled switch

Using the command-line interface

To manage virtual fabrics, you must have admin privileges on the switch chassis.

Example 8-2 shows how to check whether Virtual Fabrics is enabled or disabled and then enable it.

Example 8-2 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
```

8.3.3 Disabling Virtual Fabrics

This section describes how to disable Virtual Fabrics.

Using IBM Network Advisor

Complete the following steps to disable Virtual Fabrics in IBM Network Advisor:

1. Select the switch in the chassis group that is displayed in the product list, right-click to open the drop-down menu options, and select the option to disable Virtual Fabrics, as shown in Figure 8-5.

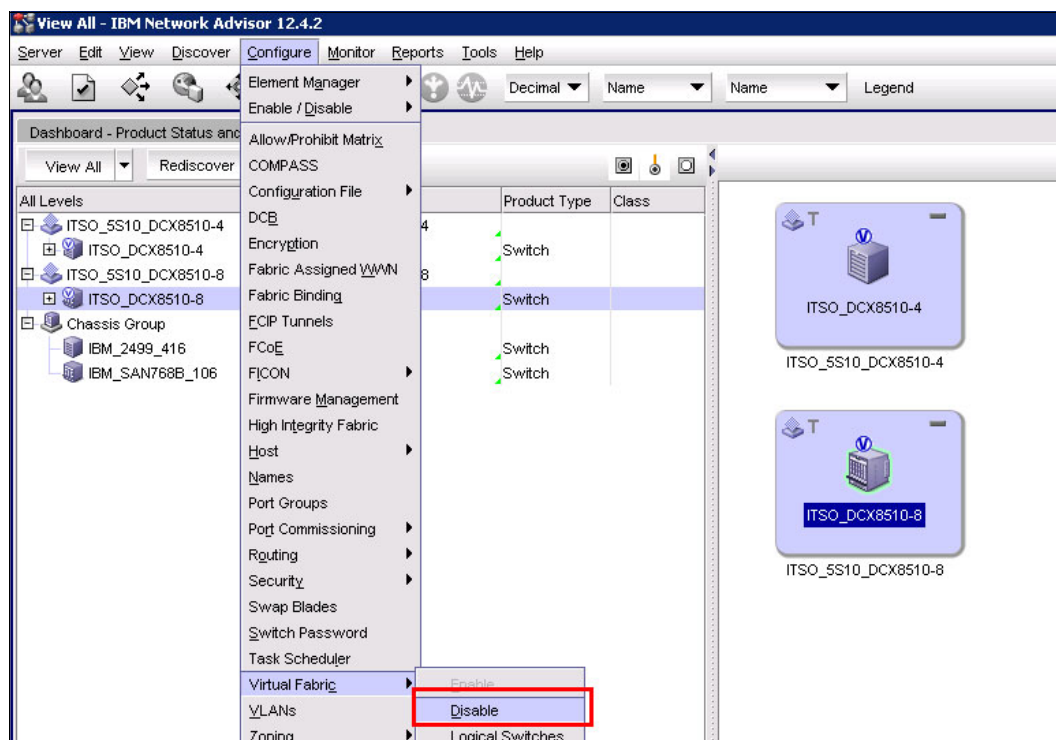


Figure 8-5 Disable Virtual Fabric

A warning message appears. This is the same warning message that is shown for enabling the VF, as shown in Figure 8-3 on page 219.

2. Read the warning and select **OK** to close the window.

Using the command-line interface

Example 8-3 shows how to use the CLI to check whether Virtual Fabrics is enabled or disabled, and then disable it.

Example 8-3 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.

8.3.4 Logical switch management

IBM Network Advisor is used to manage logical switches after Virtual Fabrics is enabled. From the IBM Network Advisor **Configure** drop-down menu, right-click the switch and select **Logical Switches**, as shown in Figure 8-6.

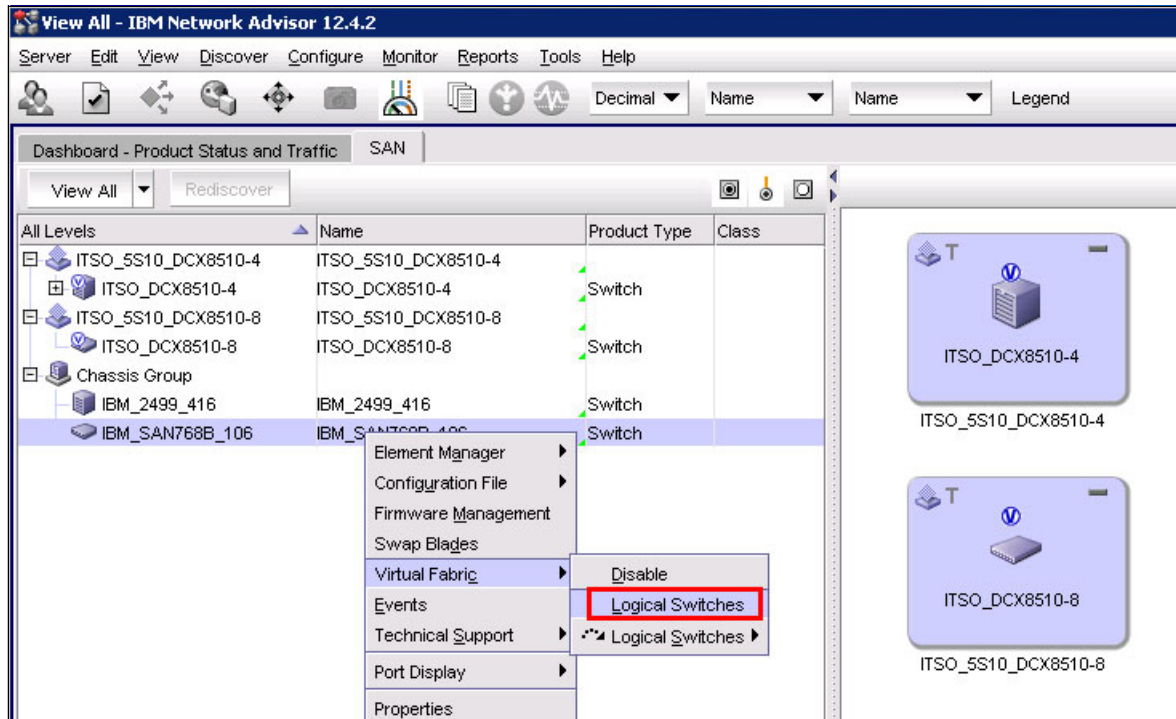


Figure 8-6 Logical switches IBM Network Advisor

The Logical Switches management window is displayed as shown in Figure 8-7. 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.

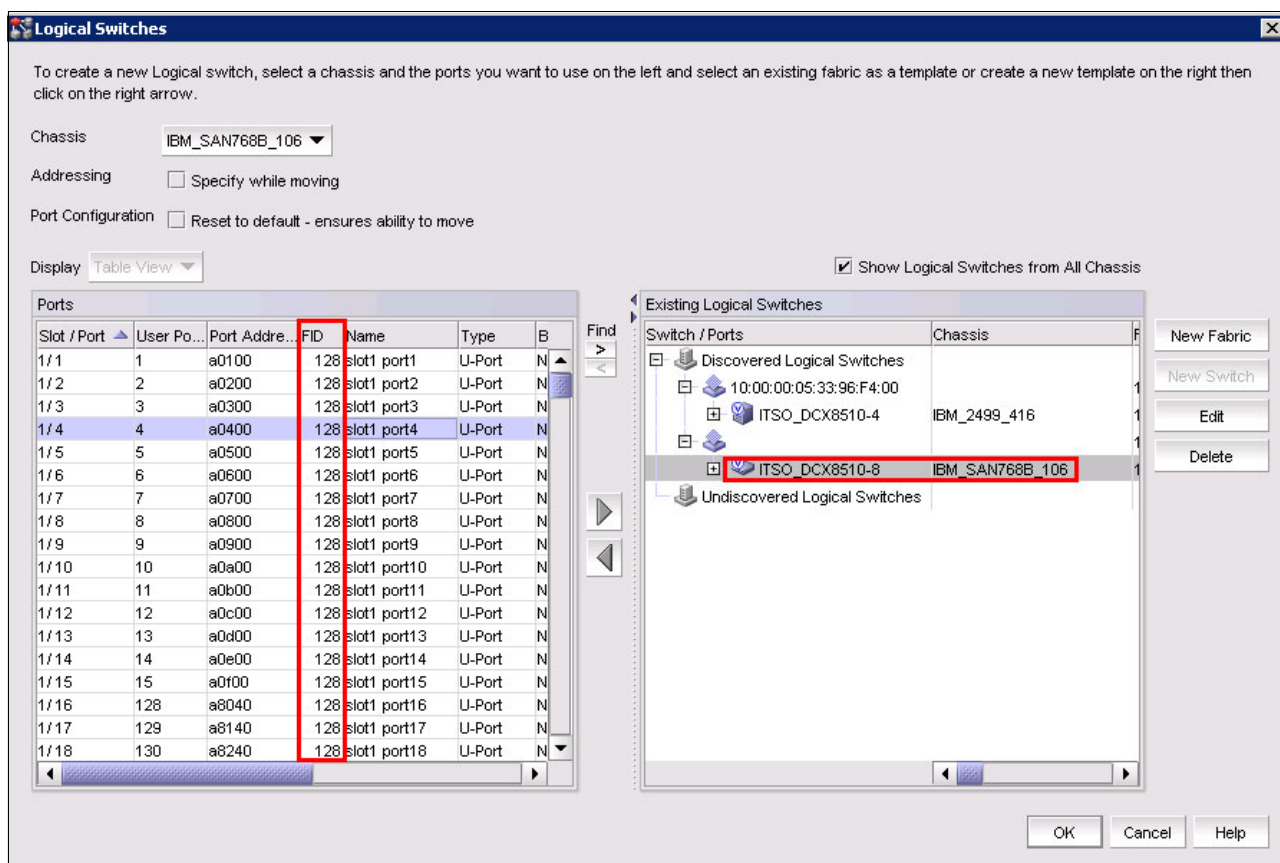


Figure 8-7 Logical Switch management

8.3.5 Modifying the base switch

To modify the base switch, select the base switch from the Logical Switches window and select the **Edit** button. You can modify all base switch parameters, as shown in Figure 8-8.

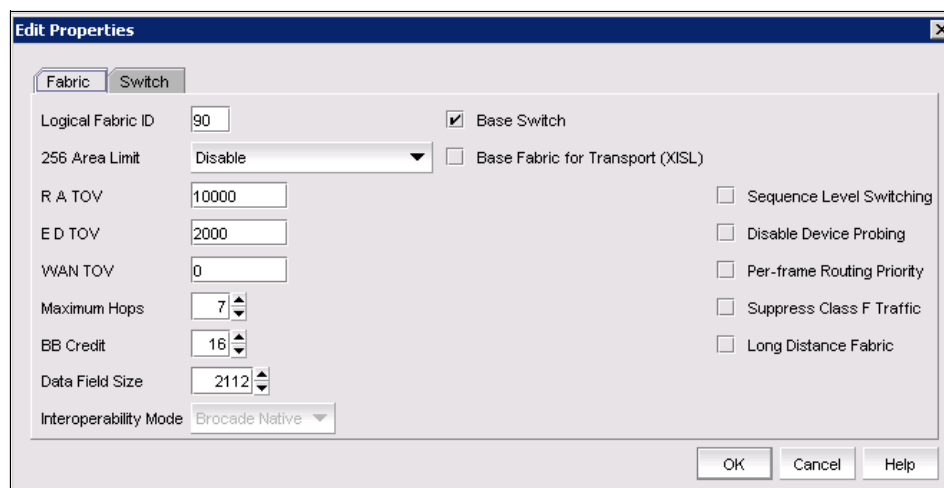


Figure 8-8 Edit Properties

After the configuration edit is complete, Click **OK** in the Edit window, and then click **OK** in the Logical Switches management window.

This action opens a confirmation window. Read the message on the window and select **OK**. The system then performs a configuration operation and displays the progress of the command under the status field, as shown in Figure 8-9.

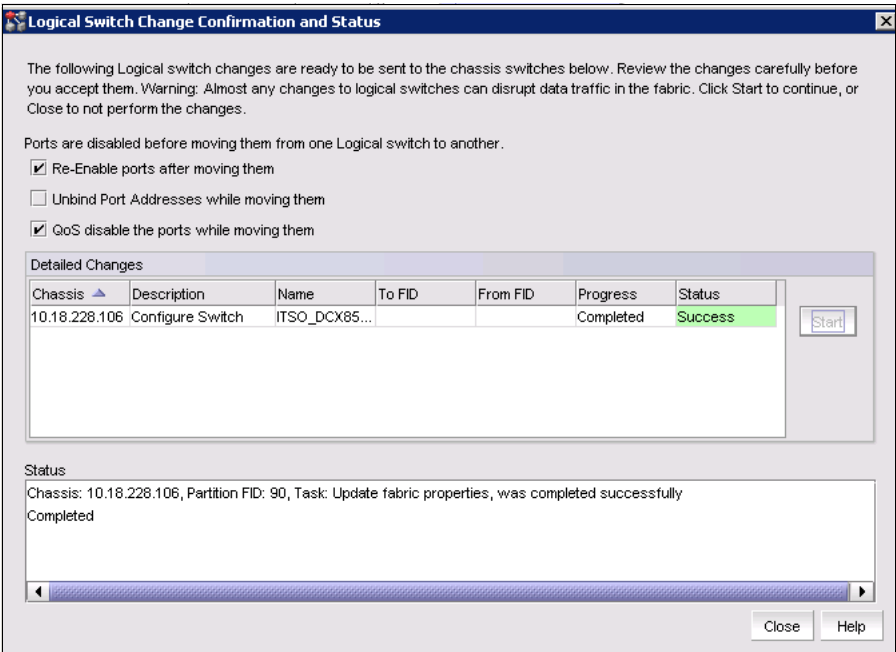


Figure 8-9 Confirmation window

8.3.6 Creating a logical switch

When the logical switch is created, it is automatically enabled and it has no ports assigned. Complete the following steps to create a logical switch:

1. Open the Logical Switches view and select the **New Fabric** option. This action opens the New Logical Fabric template.
2. Select the options that are required for the new fabric and click **OK** (Figure 8-10).

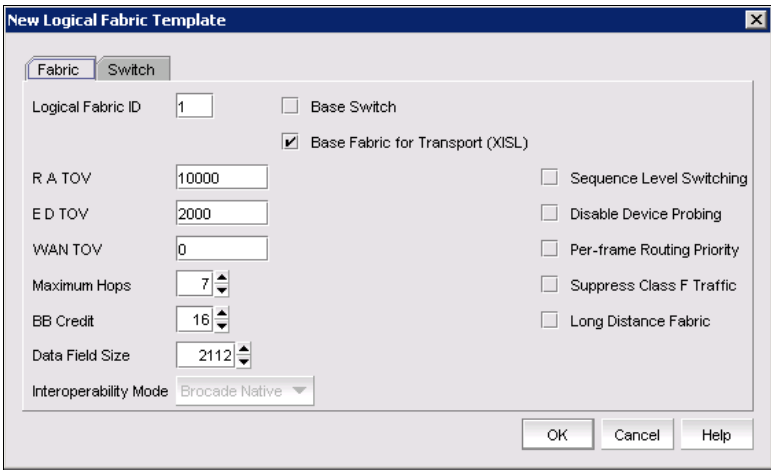


Figure 8-10 New Logical Fabric Template

The new logical fabric is displayed in the Logical Switches window.

3. Select the new fabric and then click **New Switch**, as shown in Figure 8-11.

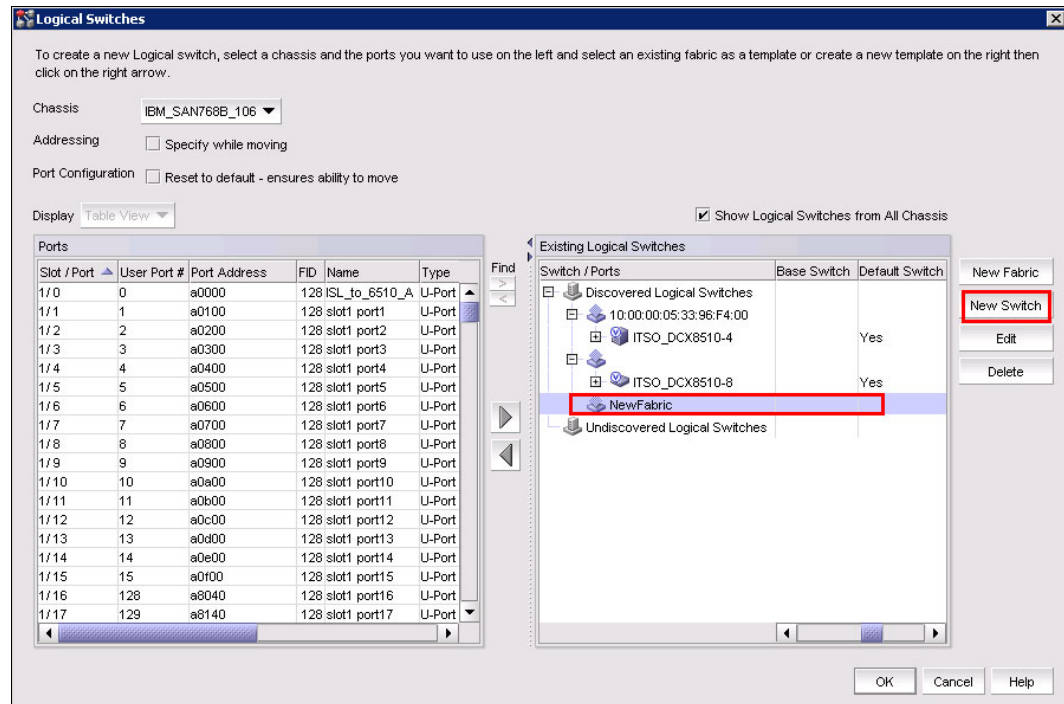


Figure 8-11 Adding a switch

The new logical switch window opens.

4. Configure the new logical switch as required by modifying the fields as shown in Figure 8-12.




Figure 8-12 New logical switch fabric parameters

Under the Switch option, you can change the switch name and domain ID, as shown in Figure 8-13.

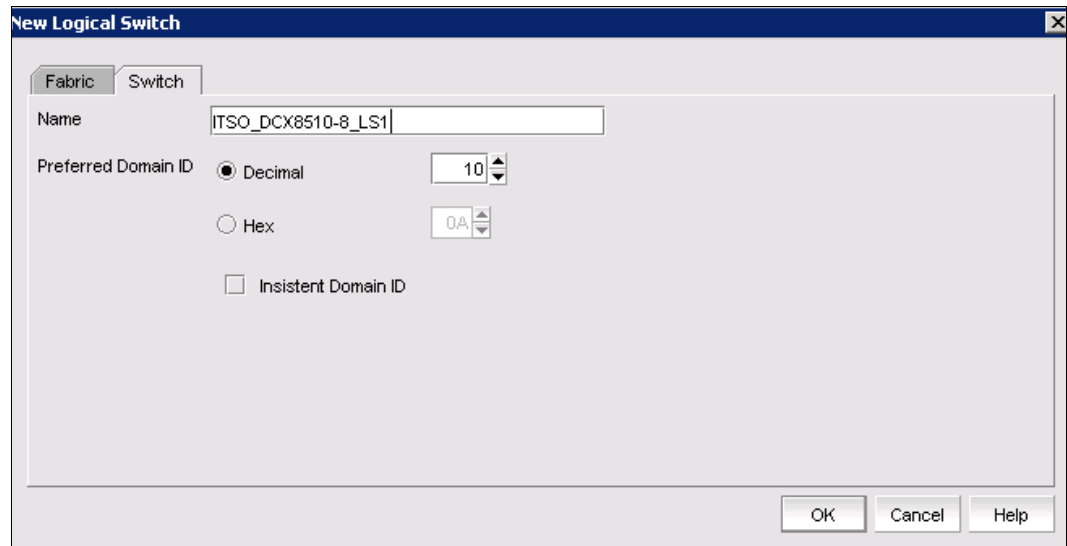


Figure 8-13 New logical switch

5. Click **OK** to add the switch.
6. 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 8-14. This process can be used at any time to add or delete ports from the logical switch.

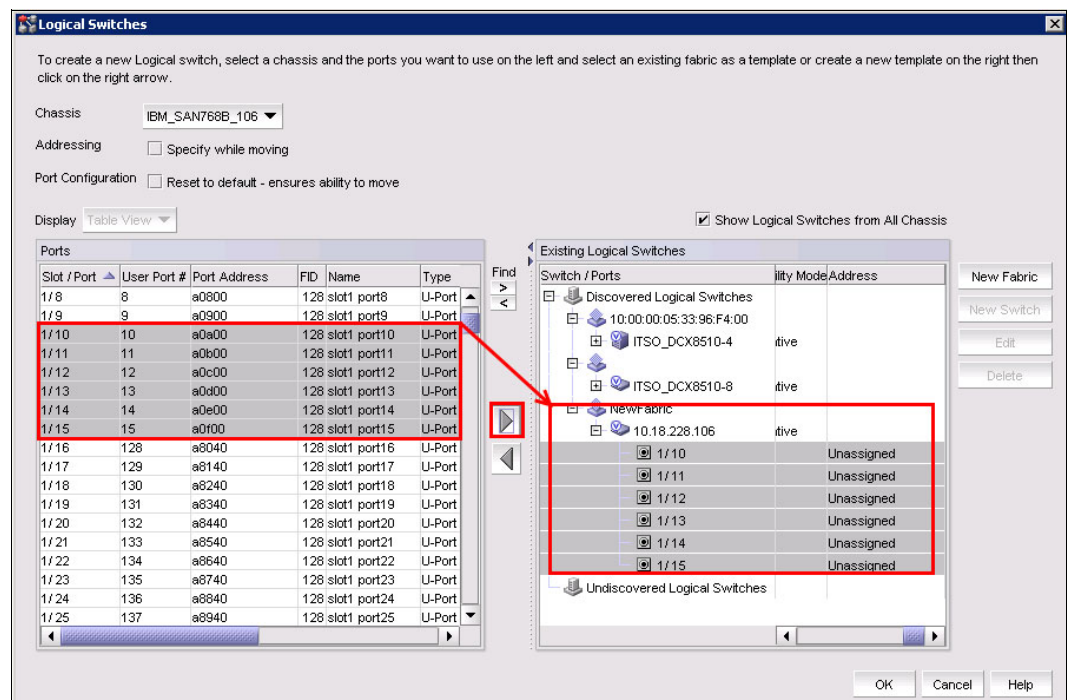


Figure 8-14 Add ports to logical switch

7. Click **OK** to process the new configuration.
The Logical Switch Change Confirmation and Status window is displayed.

Read the information in the Logical Switch Change Confirmation and Status window and then click **Start** to complete the addition of the logical switch, as shown in Figure 8-15.

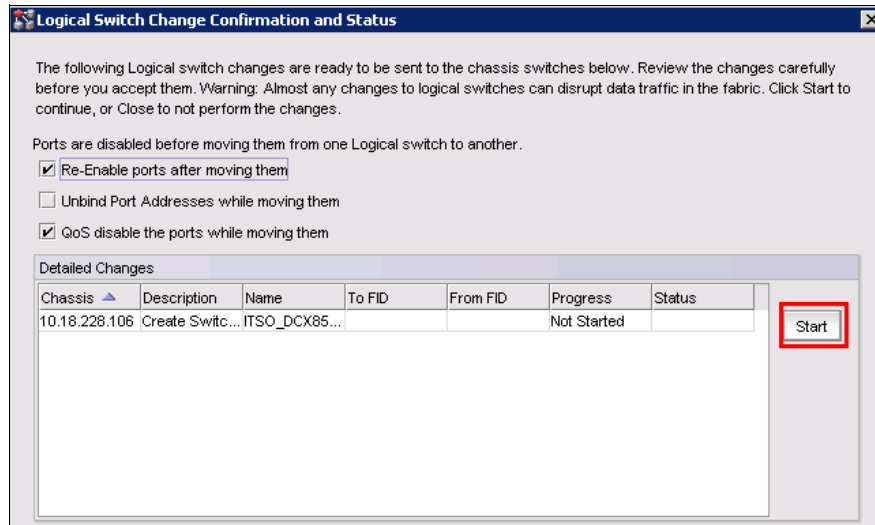


Figure 8-15 Logical switch change confirmation and status

The status bar displays the status of the activation. It will change to *Success* when completed, and the newly created fabric and switch will be displayed in IBM Network Advisor, as shown in Figure 8-16.

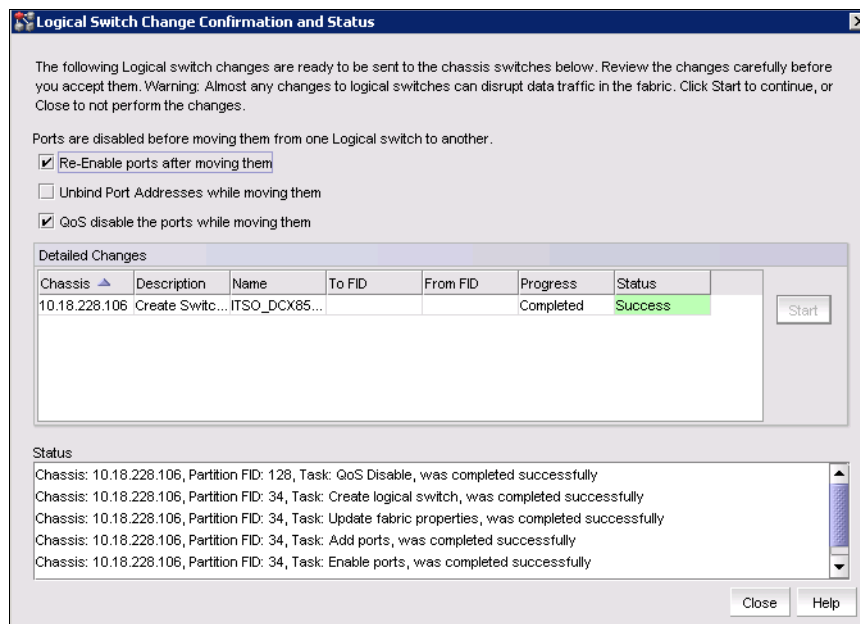


Figure 8-16 Logical Switch Status

Figure 8-17 shows the logical switch.

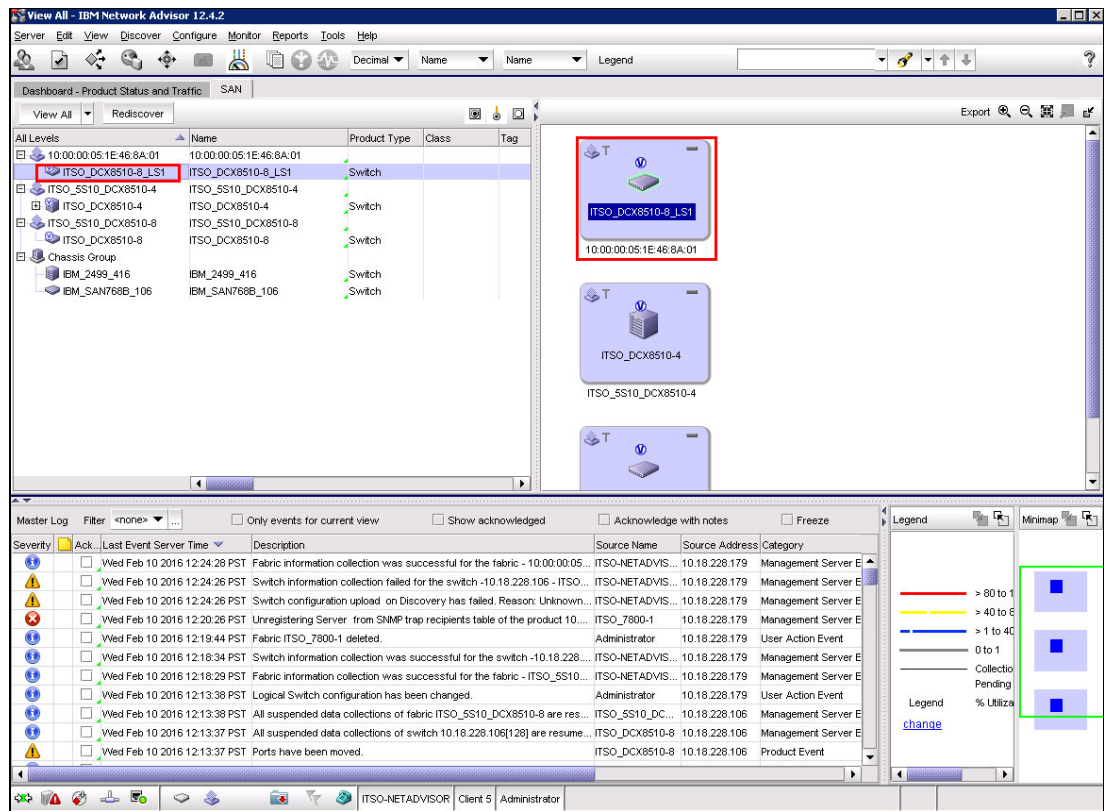


Figure 8-17 IBM Network Advisor logical switch

8.3.7 Deleting a logical switch

Complete the following steps to delete a logical switch:

1. Open the Logical Switches configuration window, select the switch that you want to delete, and click **Delete**.
2. A warning message appears. Read the warning and click **Yes**, as shown in Figure 8-18.

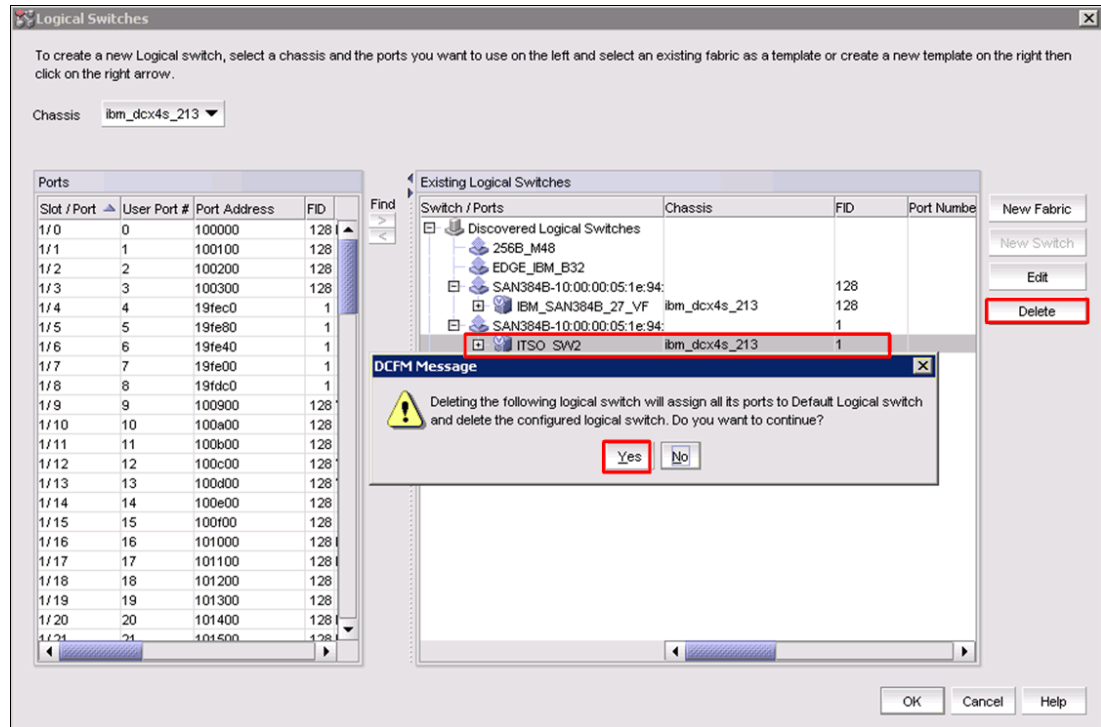


Figure 8-18 Delete switch

3. When the delete is completed, click **OK** to activate the new configuration.

Important: The default logical switch cannot be deleted.

8.3.8 Displaying the logical switch configuration

Example 8-4 shows the configuration that is created by using the CLI.

Example 8-4 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

8.3.9 Changing the fabric ID of a logical switch

Complete the following steps to change the fabric ID of an existing logical switch:

1. Select the Logical Switches window and click **Edit**. This step opens the Edit Properties window, where you can change the logical fabric ID, as shown in Figure 8-19.

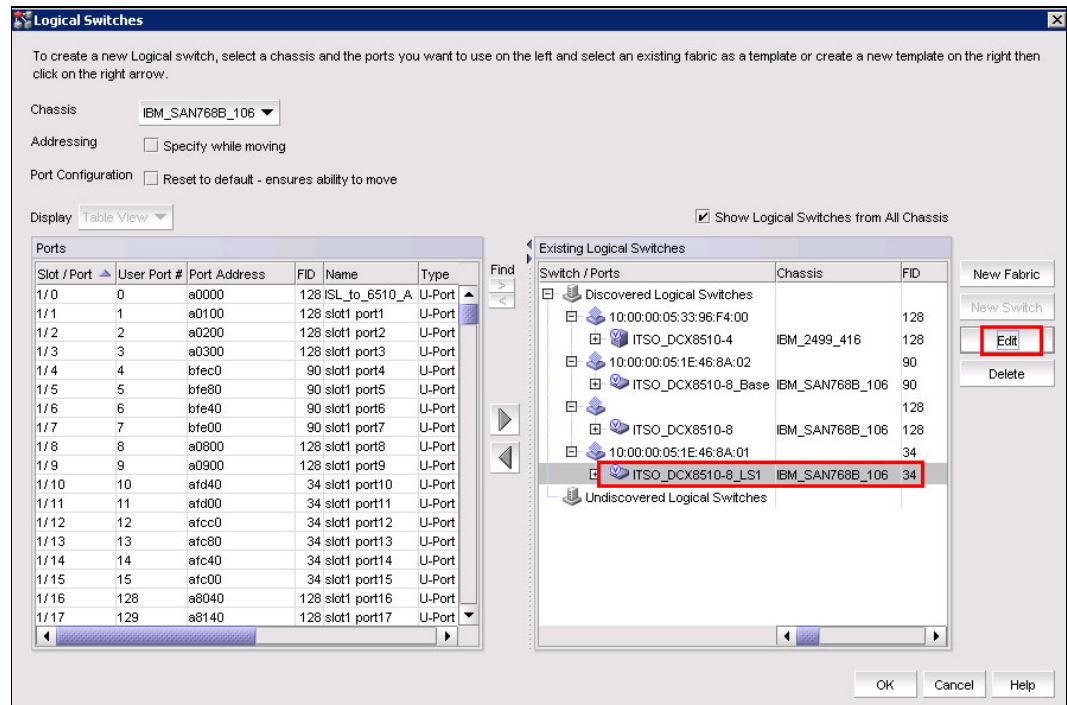


Figure 8-19 Change Logical Fabric ID

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.

In the Logical Switches window, the switch will display under the new fabric ID, as shown in Figure 8-20.

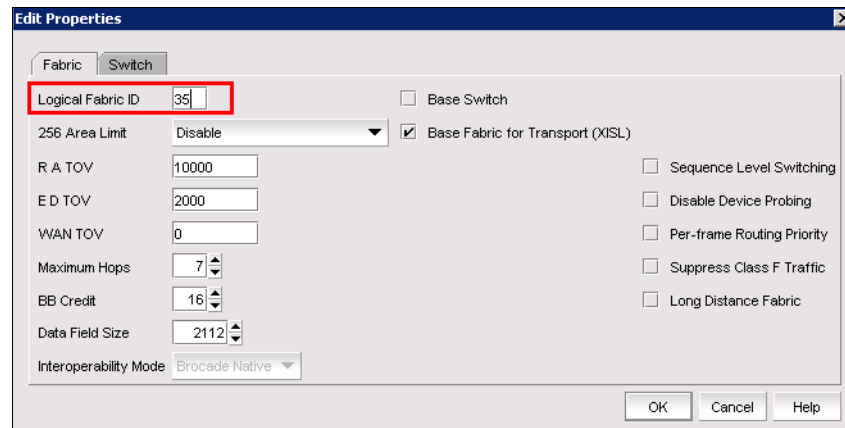


Figure 8-20 Logical Switch view with changed ID

2. To activate the change, click **OK** on the Logical Switches window, then read the confirmation message and click **Start** to complete the operation.

The updated logical fabric ID is displayed in the logical switch view as shown in Figure 8-21.

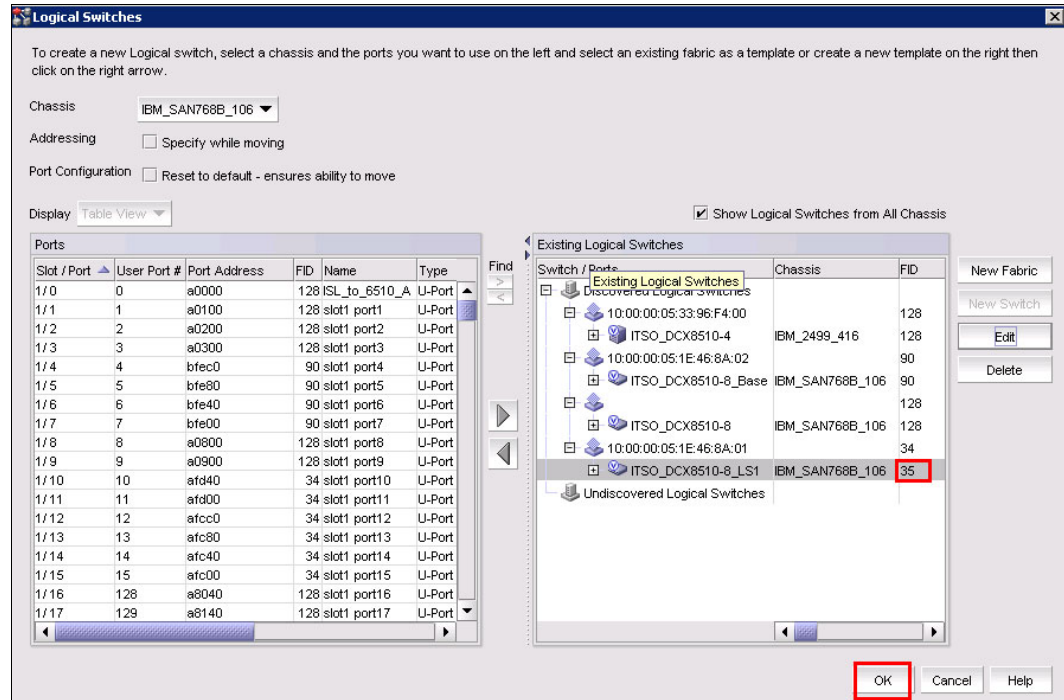


Figure 8-21 Updated logical fabric ID in the logical switch view

8.3.10 Changing a logical switch to a base switch

Only the base switch can be used for ISLs. If there is no base switch, you might want to change one of the logical switches to a base switch.

Complete the following steps to change a logical switch to a base switch:

1. Select the logical switch in the Logical Switch View window and click **Edit**.

The Edit Properties window displays.

2. Select **Base Switch** and click **OK**, as shown in Figure 8-22.

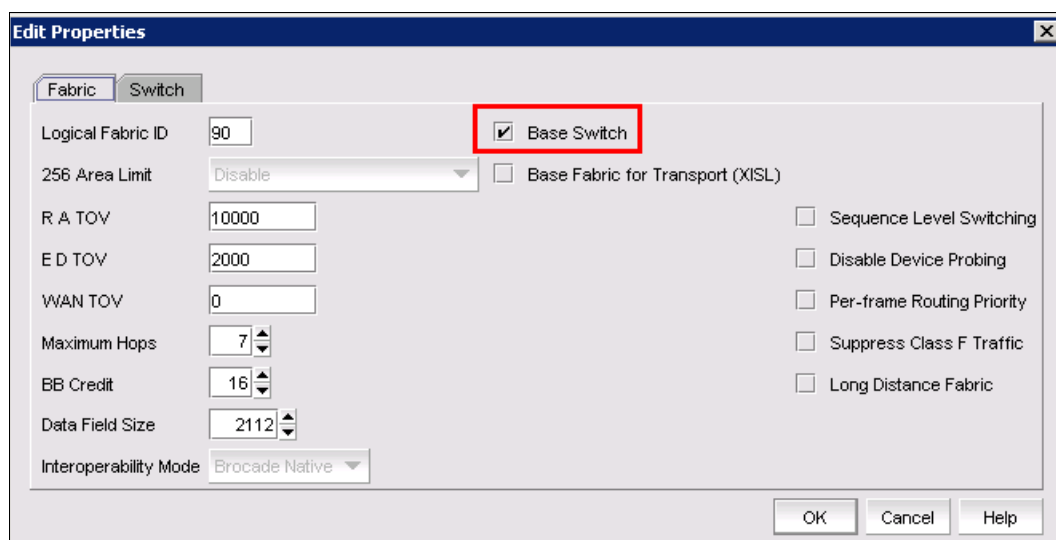


Figure 8-22 Edit Properties base unit

3. To activate the change, click **OK** in the Logical Switches window, read the confirmation message, and click **Start** 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 run this command successfully.

8.3.11 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. Run the **switchshow** command to check whether the switch is enabled for XISL use as shown in Example 8-5.

Example 8-5 Check XISL with the switchshow command

```
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. Run the **switchdisable** command to disable the switch. No output will be returned.
3. Use the **configure** command to configure the switch as shown in Example 8-6.

Example 8-6 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. Run the **switchenable** command to re-enable the switch. No output will be returned.

8.3.12 Creating a logical fabric that uses XISLs

The following procedure describes the flow of creating a logical fabric that uses XISLs, but does not provide every required detail.

Complete these steps to create a logical fabric that uses multiple chassis and XISLs:

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 feature 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 these 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 process 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 these 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 that are 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.



Implementation and migration strategies

This chapter introduces designing a SAN, redundancy and resiliency, migration assessment, migration strategy, and creating a migration plan to replace or expand an existing SAN. These topics can also be applied to planning of a new SAN because most of the requirements and concerns about planning a new SAN represent a subset of the requirements for replacing an existing SAN.

A licensing overview and information about specific usage are provided to explain the available advanced functionalities and their licensing requirements.

This chapter includes the following sections:

- ▶ Designing a storage area network
- ▶ Migration assessment
- ▶ Migration strategy
- ▶ Developing a migration plan
- ▶ Preparing to migrate
- ▶ Performing the migration and validation
- ▶ Completing the migration
- ▶ Licensing

9.1 Designing a storage area network

The storage area network (SAN) planning process is similar to any type of project planning and includes the following phases:

- ▶ Phase I: Gathering requirements
- ▶ Phase II: Developing technical specifications
- ▶ Phase III: Estimating project costs
- ▶ Phase IV: Analyzing return on investment (ROI) or total cost of ownership (TCO) (if necessary)
- ▶ Phase V: Creating a detailed SAN design and implementation plan

When you select which criteria to meet, the subject matter experts (SMEs) from the different areas (server, storage, and networking) should be engaged to understand the role of the fabric. Because most SANs tend to operate for a long time before they are renewed, consider future growth because SANs are difficult to redesign. Deploying new SANs or expanding existing ones to meet additional workloads in the fabrics requires critical assessment of business and technology requirements. Proper focus on planning will ensure that the SAN, when it is deployed, meets all current and future business objectives, including availability, deployment simplicity, performance, future business growth, and cost.

A critical aspect for successful implementation that is often overlooked is the ongoing management of the fabric. Identifying systems-level individual SMEs for all the components that make up the SAN, and adequate and up-to-date training on those components, is critical for efficient design and operational management of the fabric. When you are designing a new SAN or expanding an existing SAN, take the following parameters into account:

- ▶ Application virtualization
 - Which applications will run in a virtual machine (VM) environment?
 - How many VMs per server?
 - Under what conditions will migration of VMs take place, such business or non-business hours, and is more CPU or memory needed to maintain response times?
 - Is there a need for Flash systems to improve read response times?
- ▶ Homogeneous/heterogeneous server and storage platforms
 - Will the system be using blade servers or rack servers?
 - Is auto-tiering in place?
 - Fabric OS compatibility and feature support?
 - What is the refresh cycle of servers and storage platforms?
- ▶ Scalability
 - How many user ports are needed now?
 - How many inter-switch links (ISLs)/UltraScale inter-chassis links (ICLs) are required for minimizing congestion?
 - Will you scale out at the edge or the core?
- ▶ Backup and disaster tolerance
 - Is there a centralized backup? This will determine the number of ISLs needed to minimize congestion at peak loads.
 - What is the impact of backup on latency-sensitive applications?

- Is the disaster solution based on a metro Fibre Channel (FC) or Fibre Channel over IP (FCIP) solution?
- ▶ Diagnostics and manageability
 - What is the primary management interface to the SAN (command-line interface, IBM Network Advisor, or third-party tool)?
 - How often will the Fabric Operating System (FOS) be updated?
 - How will you validate cable and optics integrity?
- ▶ Investment protection
 - Support for future FC technology and interoperability
 - Support for alternative technologies such as Fibre Channel over Ethernet (FCoE)

9.1.1 Redundancy and resiliency

An important aspect of SAN topology is the resiliency and redundancy of the fabric. The main objective is to remove any single point of failure. Resiliency is the ability of the network to continue to function, recover from a failure, or both. Redundancy describes duplication of components, even an entire fabric, to eliminate a single point of failure in the network. IBM b-type fabrics have resiliency built into Fabric OS, which can quickly “repair” the network to overcome most failures. For example, when a link between switches fails, Fibre Channel shortest path first (FSPF) quickly recalculates all traffic flows if a second route is available, which is when redundancy in the fabric becomes important.

The key to high availability and enterprise-class installation is redundancy. By eliminating a single point of failure, business continuance can be provided through most foreseeable and even unforeseeable events. At the highest level of fabric design, the complete network should be redundant, with two completely separate fabrics that do not share any network equipment (routers or switches). Servers and storage devices should be connected to both networks by using some form of Multi-Path I/O (MPIO) solution.

MPIO allows data to flow across both networks seamlessly in either an active/active or active/passive mode. It ensures that if one path fails, an alternative is readily available. Ideally, the networks would be identical, but at a minimum they should be based on the same switch architecture. In some cases, these networks are in the same location. However, to provide for Disaster Recovery (DR), two separate locations are often used, either for each complete network or for sections of each network. Regardless of the physical location, there are two separate networks for complete redundancy.

9.1.2 Switch interconnections

As mentioned in the previous section, there should be at least two of every element in the SAN to provide redundancy and improve resiliency. The number of available ports and device locality (server/storage tiered design) determines the number of ISLs needed to meet performance requirements. This requirement means that there should be a minimum of two trunks, with at least two ISLs per trunk. Each source switch should be connected to at least two other switches, and so on.

In addition to redundant fabrics, redundant links should be placed on different blades, different ASICs, or at least different port groups whenever possible. Whatever method is used, it is important to be consistent across the fabric.

9.1.3 UltraScale ICL connectivity for Gen 5 directors

The SAN768B-2 and SAN384B-2 platforms use second-generation UltraScale ICL technology from Brocade with optical QSFP. The SAN768B-2 allows up to 32 QSFP ports, and the SAN384B-2 allows up to 16 QSFP ports to help preserve switch ports for end devices. Each QSFP port has four independent 16 Gbps links, each of which terminates on a different ASIC within the core blade. Each core blade has four ASICs. A pair of connections between two QSFP ports can create 32 Gbps of bandwidth.

9.1.4 SAN768B-2 and SAN384B-2 UltraScale ICL connection preferred practices

Each core blade in a chassis must be connected to each of the two core blades in the destination chassis to achieve full redundancy. For redundancy, use at least one pair of links between two core blades.

Follow these guidelines when designing switch ISL and UltraScale ICL connectivity:

- ▶ There should be at least two core switches.
- ▶ Every edge switch should have at least two trunks to each core switch.
- ▶ Select small trunk groups (keep trunks to two ISLs) unless you anticipate very high traffic volumes. This configuration ensures that you can lose a trunk member without losing ISL connectivity.
- ▶ Place redundant links on separate blades.
- ▶ Trunks should be in a port group (ports within an ASIC boundary).
- ▶ Allow no more than 30 m in cable difference for optimal performance for ISL trunks.
- ▶ Use the same cable length for all UltraScale ICL connections.
- ▶ Avoid using ISLs to the same domain if there are UltraScale ICL connections.
- ▶ Use the same type of optics on both sides of the trunks: Short Wavelength (SWL), Long Wavelength (LWL), or Extended Long Wavelength (ELWL).

9.1.5 Device placement

Device placement is a balance between traffic isolation, scalability, manageability, and serviceability.

With the growth of virtualization and multinode clustering on the UNIX platform, frame congestion can become a serious concern in the fabric if there are interoperability issues with the end devices.

Traffic locality

Designing device connectivity depends on the expected data flow between devices. For simplicity, communicating hosts and targets can be attached to the same switch. However, this approach does not scale well. Given the high-speed, low-latency nature of Fibre Channel, attaching these host-target pairs on different switches does not mean that performance is adversely affected. Although traffic congestion is possible, it can be mitigated with proper provisioning of ISLs/UltraScale ICLs. With current generation switches, locality is not required for performance or to reduce latencies. For mission-critical applications, architects might want to localize the traffic when using Flash systems or in exceptional cases, particularly if the number of ISLs available is restricted or there is a concern for resiliency in a multi-hop environment.

One common scheme for scaling a core-edge topology is dividing the edge switches into a storage tier and a host/initiator tier. This approach lends itself to ease of management and ease of expansion. In addition, host and storage devices generally have different performance requirements, cost structures, and other factors that can be readily accommodated by placing initiators and targets in different tiers.

Fan-in ratios and oversubscription

Another aspect of data flow is the “fan-in-ratio” or “oversubscription”, in terms of source ports to target ports and device to ISLs. This is also referred to as the “fan-out-ratio” if viewed from the storage array perspective. The ratio is the number of device ports that share a single port, whether ISL, UltraScale ICL, or target. This number is always expressed from the single entity point of view, such as 7:1 for seven hosts that are using a single ISL or storage port.

Note: One approach to ISL oversubscription is to use the base values 3:1 for very high-performance applications, 7:1 when intermediate performance is required, and 15:1 when performance is not the key consideration. When you are working with environments where high performance is required and you do not want to sacrifice a high port count for ISL connectivity, consider the use of the b-type enterprise ICL technology to interconnect your fabric backbone.

What is the optimum number of hosts that should connect per to a storage port? This seems like a fairly simple question. However, when you consider clustered hosts, VMs, and the number of logical unit numbers (LUNs) (storage) per server, the situation can quickly become much more complex. Determining how many hosts to connect to a particular storage port can be narrowed down to three considerations:

- ▶ Port queue depth
- ▶ I/O per second (IOPS)
- ▶ Throughput

Of these three, throughput is the only network component. Thus, a simple calculation is to add up the expected bandwidth usage for each host that is accessing the storage port. The total should not exceed the supported bandwidth of the target port.

However, in practice it is highly unlikely that all hosts perform at their maximum level at any one time. With the traditional application-per-server deployment, the host bus adapter (HBA) bandwidth is overprovisioned. However, with virtual servers (KVM, Xen, Hyper-V, proprietary UNIX OSs, and VMware) the situation can change radically. Network oversubscription is built into the virtual server concept. To the extent that servers use virtualization technologies, the network-based oversubscription should be reduced proportionally. Therefore, it might be prudent to oversubscribe ports to ensure a balance between cost and performance.

Another method is to assign host ports to storage ports based on capacity (density). The intended result is a few high-capacity hosts and a larger number of low-capacity servers assigned to each storage port, thus distributing the load across multiple storage ports.

Regardless of the method that is used to determine the fan-in/fan-out ratios, port monitoring should be used to determine actual utilization and what adjustments, if any, should be made. In addition, ongoing monitoring provides useful heuristic data for effective expansion and efficient assignment of existing storage ports. To determine the device-to-ISL fan-in ratio, a simple calculation method works best. The storage port should not be oversubscribed into the core. For example, an 8 Gbps storage port should have an 8 Gbps pipe into the core.

9.2 Migration assessment

It is important to understand the current application environment and the new SAN requirements before attempting a migration. There is more than one way to proceed with the migration process, depending on the current SAN architecture, fabric topology, size, and number of active devices attached. A SAN fabric migration can be done both offline or online, depending on the application or project requirements. An offline migration is the simpler of the two approaches, although careful planning is required. However, in many environments where planned downtime is not possible, the migration must be performed online. An online migration in a single or redundant fabric requires careful evaluation of the application availability and currently deployed topology to plan for a methodical migration path.

Consider these factors, regardless of the migration approach:

- ▶ Assessing the existing fabric topology
- ▶ Assessing the new fabric topology
- ▶ Logistic planning for hardware installation
- ▶ Topology and zone planning
- ▶ Preliminary migration planning

9.2.1 Assessing the existing fabric topology

Determine whether the current environment is a single fabric or a redundant fabric. If the current environment uses a redundant fabric, a rolling migration might be an option, where one fabric is active and the other fabric is migrated offline. Similarly, device paths in a single resilient fabric can be failed over as devices are moved to the new fabric. Both methods minimize fabric downtime and I/O interruptions, if multipathing software is in use. Migrating a single non-resilient fabric is more complex and requires application interruption or an outage if the host must be rebooted.

Consider these elements when assessing the migration activity:

- ▶ Application failover considerations: If multipathing software such as Microsoft MPIO, IBM AIX® MPIO, Hitachi HiCommand Dynamic Link Manager, or EMC PowerPath is in use, collect metrics to determine how long it takes to fail over and fail back in the existing SAN.
- ▶ Storage failover considerations: Move all the LUNs to a single controller if not dual-pathed. Verify that the number of LUNs from a single port does not exceed the vendor recommendation.
- ▶ Topology change at the time of migration: Migrating to a new fabric is a good opportunity to address any performance bottlenecks, server and storage scalability, and general maintenance of the fabric, such as structured cable management. High-density directors with ICLs offer an opportunity to simplify traditional SAN designs.
- ▶ Zone configuration export/modify strategy: If some or all of the devices in the old fabric are being migrated to the new fabric, the existing zone database can be exported and then imported into the IBM b-type SAN to minimize the migration time frame.
- ▶ Server and storage device placement: Although hop count is no longer an issue, keeping the number of hops between server and storage to no more than two can minimize possible congestion issues as the SAN expands. Whatever method is used for device placement, be consistent across switches and fabrics.

9.2.2 Assessing the new fabric

Consider these points when assessing the new fabric:

- ▶ Fabric OS upgrade requirements: Before connecting any devices, verify that the switches are running the correct version of Fabric OS.
- ▶ Capture configuration parameters of the existing switch: Capture the configuration of the existing switches and compare that configuration with the new ones.
- ▶ Analyze the existing zoning: Assess the existing zone database. Clean the zone database by removing any zone members that are no longer part of the fabric.
- ▶ Trunking setup considerations: ISL Trunking is a hardware-based stripping mechanism with predictable latencies for traffic flows. In a multi-switch environment, multiple trunks should exist such that during an entire trunk failure, the remaining trunks are not congested.
- ▶ Future server or storage expansion: Planning for the future is key to ensuring that the architecture that is put in place for the new SAN will meet long-term requirements.

9.2.3 Logistic planning of hardware installation

When planning a migration, consider the following concerns about facilities and logistics:

- ▶ Rack space requirements: IBM b-type Gen 5 directors use front-to-back airflow, which allows a narrow rack and the implementation of hot/cold aisle cooling.
- ▶ Power requirements: SAN768B-2 and SAN384B-2 have a power consumption of 1952 W and 1064 W when fully loaded.
- ▶ Cable requirements: Confirm that the cable plant is within the required specifications and uses structured cabling, when possible, to minimize device placement errors during the migration.

9.2.4 Preliminary migration planning

When developing a preliminary migration plan, consider the following topics:

- ▶ To complete a successful migration, identify the personnel that are needed during the key phases of the project: Facilities management, network administration, SAN administration/engineering, server administration with knowledge of the dual-pathing and failover software, storage administration with knowledge of redundant paths, and project management.
- ▶ Identify and analyze key implemented features, and define equivalent solutions for IBM b-type SAN infrastructure.
- ▶ Identify and analyze advanced features that might need to be considered, such as FCIP, Encryption, or FICON for the new SAN.
- ▶ Define move groups based on applications, storage ports, and zoned hosts.
- ▶ Identify and resolve any service level agreement (SLA) conflicts within move groups.
- ▶ Create port maps for host/storage on the migrated SAN.
- ▶ Review the migration plan with the user or business group, and revise as needed.
- ▶ Complete the final migration plan.

Establish the following post-migration verification criteria:

- ▶ Total number of devices in the fabric.
- ▶ Baseline performance metrics for ISL and ICLs.
- ▶ Baseline latency measurements for server and storage ports.
- ▶ Hosts are dual-pathed to the fabric so that the use of failover mechanisms minimizes the disruption to production I/O.
- ▶ The IBM b-type devices basic setup and configuration has been performed in advance.
- ▶ All required switch licenses have been acquired and installed.
- ▶ If IBM Network Advisor is being used, it has already been set up and is able to discover the IBM b-type fabric.
- ▶ Regardless of the type of fabric, perform the migration during non-peak business hours.

9.2.5 Gather infrastructure information

Detailed information is required for the following items:

- ▶ Individual Fabric Details
- ▶ Device Details
- ▶ Device Mapping Details
- ▶ Application-Specific Details

9.3 Migration strategy

The migration process can be simplified by preparing a migration plan in advance. Besides cabling, rack space, and power requirements, other factors such as scheduling downtime, personnel security, and application change windows as well as host and storage failover can significantly affect the SAN operations. The current configuration and operational requirements of a target SAN might impose additional constraints. The key to a successful migration is to minimize fabric interruption or to completely eliminate downtime, whenever possible, by identifying issues in advance.

Effective planning provides the preliminary groundwork for the evaluation phase and sets the foundation for the migration process. After reviewing the requirements that apply to a specific situation, the migration process will fall into one of the following categories:

- ▶ **Online Redundant Fabric Migration**

A redundant fabric provides the flexibility to upgrade one fabric by bringing it offline while redirecting active I/O to the other fabric. Current I/O operations are not affected by the migration activity. With this strategy, the hosts are operating in a degraded mode with no data path protection. Any failure on an active path completely ceases I/O.

With proper planning, any downtime or outage is minimized. When the fabric upgrade has been completed and verified, it can be brought back online by restoring the I/O paths. The migration process is repeated for the second fabric after all I/O paths are successfully restored on the first fabric.

- ▶ **Offline fabric migration**

An offline fabric migration assumes that a fabric can be brought offline to perform the migration and that I/O is stopped during the downtime. This method is the safest and most convenient method for migration.

9.3.1 Migration methods

Infrastructure resiliency or redundancy of the fabric determines the primary migration strategy. While you are preparing for the development of a migration plan, the strategy to be used must be identified. The migration has the following options:

- ▶ Port-to-Port migration: This method is a straightforward port-port migration from one fabric to another. This method requires all logically grouped initiator/target pairs to be moved during a single migration activity. This strategy is called migrating by “move groups.” For example, when a storage port is moved, all associated HBAs that are accessing LUNs through this port must also be moved.
- ▶ Application migration: This method is possible if the physical infrastructure is not shared across application tiers. If the application happens to run on a new server and storage infrastructure, a validation is required to check whether all the data has been migrated before the cutover. SANs tend to be logically identified as database, web services, backup, and so on.
- ▶ Device migration: This method is a logical approach to offline migration because customers physically isolate servers and storage devices in racks or sections of the data center. Migrating devices by using this method provides a clear high-level accounting, especially for the racks that are relocated as part of the migration.

9.4 Developing a migration plan

A plan is the foundation for a successful project or, in this case, a SAN migration. The best practices for such a plan can be derived from any number of formal methodologies. However, any good plan should include at least the following steps:

- ▶ Project scope and success criteria
- ▶ Phases, tasks, and subtasks
- ▶ Resource definitions
- ▶ Timelines
- ▶ Task dependencies
- ▶ Tracking criteria
- ▶ Checkpoints
- ▶ Deliverables for procedures, designs, and configurations
- ▶ Fallback plans
- ▶ Signoff criteria

9.5 Preparing to migrate

Performing the following steps ahead of time helps to minimize the time that is required for migration. Keep checklists to track switch configurations, zone information, and port mappings.

Migration preparation falls into the following categories:

- ▶ Build the new SAN infrastructure
- ▶ Configure the SAN
- ▶ Validate the new SAN

9.6 Performing the migration and validation

The migration includes the following tasks:

- ▶ Run the migration plan
- ▶ Validate migration per phase
- ▶ Validate application operations per phase
- ▶ Sign off on the SAN migration phase

When the migration assessment, qualification, and preparation are complete, the SAN can be migrated. Based on the criteria that are listed in the previous sections, select the primary migration strategy from these options:

- ▶ Offline migration
- ▶ Redundant and single fabric online migration

9.6.1 Offline migration

Although this method requires the fabric to be offline, it is also the safest option for migration.

9.6.2 Redundant and single fabric online migration

This strategy involves migrating devices by keeping the applications online. It is challenging to facilitate and requires a great deal of planning. However, if planned properly, and if the key applications that need to remain online are designed to support high availability and redundancy, this method can allow migration with no interruption of service. The key to this approach is setting the correct expectations in advance.

9.7 Completing the migration

When the migration activity is complete, it is critical to use a post-migration plan. There are several steps to ensuring that all the work that you completed is protected and validated. The following are some of the post-migration activities:

- ▶ Run IBM SAN Health
- ▶ Validate new SAN configurations
- ▶ Validate application operations
- ▶ Back up new SAN configurations
- ▶ Sign off on SAN migration
- ▶ Decommission the old SAN infrastructure

9.8 Licensing

The following license types are supported in Fabric OS:

- ▶ Permanent license: A permanent license enables a license-controlled feature to run on the switch indefinitely.
- ▶ Temporary license: A temporary license enables a license-controlled feature to run on the switch on a temporary basis. A temporary license enables demonstration and evaluation of a licensed feature, and can be valid for up to 45 days.

- ▶ Universal temporary license: A universal temporary license can only be installed once on a switch, but can be applied to as many switches as required. Temporary use duration (the length of time the feature will be enabled on a switch) is provided with the license keys.
- ▶ Slot-based licensing: A slot-based license allows you to select the slots that the license will enable up to the capacity purchased. You can increase the capacity without disrupting slots that already have licensed features running. Each licensed feature that is supported on the blade has a separate slot-based license key.

9.8.1 Available Fabric OS licenses

This subsection provides a description and usage of the available licenses.

10 Gigabit FCIP/Fibre Channel (10G license)

The 10G license has these characteristics:

- ▶ Allows 10 Gbps operation of FC ports on the Brocade SAN48B-5 or SAN96B-5 switches or the FC ports of FC16-32 or FC16-48 port blades installed on a SAN384B-2 or SAN768B-2 Backbone.
- ▶ Enables the two 10-GbE ports on the FX8-24 extension blade when installed on the SAN384B, SAN768B, SAN384B-2, or SAN768B-2 Backbone.
- ▶ Allows selection of the following operational modes on the FX8-24 blade:
 - 10 1-GbE ports and 1 10-GbE port
 - 2 10-GbE ports
- ▶ License is slot-based when applied to a Brocade Backbone. It is chassis-based when applied to a SAN48B-5 or SAN96B-5 switch.

This license allows the establishment of a 10G ISL in metro optical connectivity or with 10G dense wavelength division multiplexing (DWDM) devices. It also allows a 10G FCIP connection.

SAN06B-R Upgrade

The SAN06B-R Upgrade has these characteristics:

- ▶ Enables full hardware capabilities on the Brocade SAN06B-R base switch, increasing the number of Fibre Channel ports from four to sixteen and the number of GbE ports from two to six.
- ▶ Supports up to eight FCIP tunnels instead of two.
- ▶ Supports advanced capabilities such as Open Systems tape read/write pipelining.

Note: The SAN06B-R switch must have the SAN06B-R Upgrade license to add FICON Management Server (CUP) or Advanced FICON Acceleration licenses.

Adaptive Networking with QoS

This license provides a rich framework of capability, allowing a user to ensure that high-priority connections obtain the bandwidth necessary for optimum performance, even in congested environments. The QoS SID/DID Prioritization and Ingress Rate Limiting features are included in this license, and are fully available on all 8G and 16G platforms.

Notes:

- ▶ The SAN96B-5 does not require an Adaptive Networking with QoS license to enable the capabilities that are associated with this license. These capabilities are included by default on the SAN96B-5.
- ▶ This license is automatically enabled for new switches that operate with only Fabric OS 7.2.0 or later, and for existing switches that are upgraded to Fabric OS 7.2.0 or later.

Advanced Extension

The Advance Extension license provides these features:

- ▶ Enables two advanced extension features: FCIP Trunking and Adaptive Rate Limiting.
- ▶ FCIP Trunking feature allows all of the following configurations:
 - Multiple (up to 4) IP source and destination address pairs (defined as FCIP Circuits) using multiple (up to 4) 1-GbE or 10-GbE interfaces to provide a high-bandwidth FCIP tunnel and failover resiliency.
 - Support for up to 4 of the following quality of service (QoS) classes: Class-F, high, medium, and low priority, each as a TCP connection.
- ▶ The Adaptive Rate Limiting feature provides a minimum bandwidth guarantee for each tunnel with full usage of available network bandwidth without any negative impact to throughput performance under high traffic load.
- ▶ Available on the SAN06B-R switch, SAN42B-R, and on the SAN768B, SAN384B, SAN768B-2 and SAN384B-2 platforms for the FX8-24 on an individual slot basis.

Advanced Acceleration for FICON

This license allows use of specialized data management techniques and automated intelligence to accelerate FICON tape read and write and IBM Global Mirror data replication operations over distance. It maintains the integrity of command and acknowledgment sequences.

This license is available on the SAN06B-R and SAN42B-R switches. It is also available on the SAN384B, SAN768B, SAN384B-2, and SAN768B-2 directors for the FX8-24 blade on an individual slot basis.

Advanced Performance Monitoring

The Advanced Performance Monitoring (APM) license offers these features:

- ▶ Enables performance monitoring of networked storage resources
- ▶ Includes the Top Talkers feature
- ▶ Helps to identify end-to-end bandwidth usage by host/target pairs and is designed to provide information for capacity planning

The Fabric Vision license is equivalent to the combination of both the APM and Fabric Watch (FW) licenses. If you have both the APM and the FW licenses installed, you do not need the Fabric Vision license.

Fabric Watch

The FW license constantly monitors mission-critical switch operations for potential faults and automatically alerts administrators about problems before they become costly failures. Fabric Watch includes Port Fencing capabilities.

The Fabric Vision license is equivalent to the combination of both the APM and FW licenses. If you have both the APM and FW licenses installed, you do not need the Fabric Vision license.

Extended Fabrics

This license provides greater than 10 km of switched fabric connectivity at full bandwidth over long distances (depending on the platform, this distance can be up to 3000 km).

Fabric interconnectivity over Fibre Channel at longer distances

ISLs can use long-distance dark fiber connections to transfer data. Wavelength-division multiplexing, such as DWDM, coarse wavelength division multiplexing (CWDM), and time-division multiplexing (TDM), can be used to increase the capacity of the links. As Fibre Channel speeds increase, the maximum distance decreases for each switch.

The Extended Fabrics feature extends the distance the ISLs can reach over an extended fiber. This extension is accomplished by providing enough buffer credits on each side of the link to compensate for latency that is caused by the extended distance.

Simplified management over distance

Each device that is attached to the SAN appears as a local device, which simplifies deployment and administration.

Optimized switch buffering

When Extended Fabrics is installed on gateway switches (with E_Port connectivity from one switch to another), the ISLs (E_Ports) are configured with a large pool of buffer credits. The enhanced switch buffers help ensure that data transfer can occur at near-full bandwidth to use the connection over the extended links efficiently. This efficiency ensures the highest possible performance on ISLs.

Note: This license is not required for long-distance connectivity that uses licensed 10G ports.

ISL Trunking

The ISL Trunking license provides these features:

- ▶ Provides the ability to aggregate multiple physical links into one logical link for enhanced network performance and fault tolerance.
- ▶ Includes Access Gateway ISL Trunking on those products that support Access Gateway deployment.

Ports on Demand

The Ports on Demand license allows you to instantly scale the fabric by provisioning additional ports by using license key upgrades.

Note: This license applies to the SAN24B-4, SAN40B-4, SAN80B-4, SAN24B-5, SAN48B-5, SAN96B-5, and Flex System FC5022 16 Gb SAN Scalable switches.

DataFort Compatibility

This license provides ability to read, write, decrypt, and encrypt the NetApp DataFort encrypted Disk LUNs and tapes to the following switches:

- ▶ Encryption Switch (SAN32B-E4)
- ▶ Enterprise b-type directors with FS8-18 blade

The DataFort Compatability license includes metadata, encryption, and compression algorithms.

Encryption Performance Upgrade

This license provides additional encryption bandwidth on encryption platforms. For the Brocade Encryption Switch, two Encryption Performance Upgrade licenses can be installed to enable the full available bandwidth. On a Brocade enterprise platform, a single Performance License can be installed to enable full bandwidth on all FS8-18 blades that are installed in the chassis.

Enhanced Group Management

The Enhanced Group Management license enables full management of the device in a data center fabric with deeper element management functionality and greater management task aggregation throughout the environment. This license is used with Brocade Network Advisor application software. This license is applicable to all of the IBM b-type 8G and 16G FC platforms.

Note: This license is enabled by default on all 16G FC platforms, and on SAN768B and SAN384B platforms that are running Fabric OS v7.0.0 or later. This license is not included by default on 8G FC fixed port switches (SAN24B-4, SAN40B-4, SAN80B-4, and 8G FC blade server SAN I/O modules).

Fabric Vision

The Fabric Vision (FV) license allows you to activate the following features:

- ▶ Monitoring and Alerting Policy Suite (MAPS)
- ▶ Flow Vision
- ▶ Run D_Port tests between a switch and non-Brocade HBAs

This license replaces the Advanced Performance Monitor (APM) and Fabric Watch (FW) licenses. If you have the Fabric Vision license, you can use Advanced Performance Monitoring and Fabric Watch features without the APM and FW licenses.

Integrated Routing

Integrated Routing (IR) allows any ports on the SAN40B-4, SAN80B-4, SAN48B-5, SAN96B-5, SAN42B-R, SAN32B-E4, Flex System FC5022 switches, or on the SAN768B, SAN384B, SAN384B-2 and SAN768B-2 platforms to be configured as an EX_Port supporting FC-FC routing. This configuration provides improved scalability and fault isolation, along with multivendor interoperability.

FICON CUP

FICON with control unit port (CUP) Activation provides in-band management of the supported SAN b-type switch and director products by system automation for z/OS. To enable in-band management on multiple switches and directors, each chassis must be configured with the appropriate FICON CUP feature. This support provides a single point of control for managing connectivity in active FICON I/O configurations. System automation for OS/390 or z/OS can now concurrently manage IBM 9032 ESCON directors, in addition to supported SAN b-type switch and director products with FICON.

Full Fabric

The Full Fabric license enables a switch to connect to a multi-switch fabric through E_Ports, forming ISL connections.

Note: This license is only required on select blade server SAN I/O models and the SAN24B-4, and does not apply to other fixed-port switches or chassis-based platforms.

ICL 8-Link

ICL 8-Link activates all eight links on ICL ports on a Brocade SAN384B or half of the ICL bandwidth for each ICL port on the SAN768B platform by enabling only eight links out of the sixteen links available. This license allows you to purchase half the bandwidth of SAN768B ports initially and upgrade with an additional ICL 8-Link license to use the full ICL bandwidth later.

This license is also useful for environments that need to create ICL connections between a SAN768B and a SAN384B. The latter cannot support more than eight links on an ICL port.

It is available on the SAN768B and SAN384B backbones only.

ICL 16-Link

This license activates all 16 links on ICL ports on a SAN768B chassis. Each chassis must have the ICL 16-Link license installed in order to enable the full 16-link ICL connections.

This license is available only on the SAN768B.

Inter-Chassis Link (1st POD)

This license activates half of the ICL bandwidth on a SAN768B-2, or all the ICL bandwidth on a SAN384B-2. It allows you to enable only the bandwidth that is needed and upgrade to more bandwidth later. This license is also useful for environments that need to create ICL connections between a SAN768B-2 and a SAN384B-2. The latter platform supports only half the number of ICL links that the former platform supports.

This license is available only on the Gen5 Backbones.

Inter-Chassis Link (2nd POD)

Activates the remaining ICL bandwidth on the Brocade SAN768B-2 chassis. Each chassis must have this ICL license installed to enable all available ICL connections.

This license is available only on the Gen5 backbones.

Enterprise ICL

This license allows you to connect four or more chassis to a SAN768B-2 or SAN384B-2 chassis by using ICLs. For each Gen5 backbone, you can connect up to three Gen5 Backbones with ICLs without this license. This license is required only on the Gen5 Backbone that is connected to four or more Gen5 Backbone chassis.

This license requirement does not depend on the total number of Gen5 Backbone chassis that exist in a fabric. Rather, it depends only on the number of chassis connected directly to a Gen5 Backbone with ICLs.

You must also have an ICL POD license on each Gen5 Backbone to activate the ICL ports.

The Enterprise ICL license allows only connection of more than four chassis by using ICLs. It does not enable the ICL ports on a chassis.

Note: This license applies only to the Gen5 Backbone family.

Server Application Optimization

When deployed with Brocade server adapters, the Server Application Optimization (SAO) license optimizes overall application performance for physical servers with 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.

Notes:

- ▶ The Brocade SAN96B-5 does not require an SAO license to enable the capabilities that are associated with this license. These capabilities are included by default on the SAN96B-5.
- ▶ This license is automatically enabled for new switches that operate with only Fabric OS 7.2.0 or later and for existing switches that are upgraded to Fabric OS 7.2.0 or later.

WAN Rate Upgrade 1

This license provides additional WAN transmission throughput up to 10 Gbps on a SAN42B-R. Without WAN Rate Upgrade 1 license, the SAN42B-R provides WAN throughput of 5 Gbps. Upgrade licenses do not impose restriction on the number of physical ports that are used when the aggregate bandwidth of all configured FCIP tunnels does not exceed the licensed limit.

WAN Rate Upgrade 2

Provides unlimited WAN transmission throughput (other than the physical port limit) and enables two 40 GbE ports on a SAN42B-R switch. The 40 GbE ports cannot be used without the WAN Rate Upgrade 2 license.

The WAN Rate Upgrade 1 license must be installed before you install and activate the WAN Rate Upgrade 2 license. The WAN Rate Upgrade 1 license cannot be removed until you remove the WAN Rate Upgrade 2 license.

Enterprise software bundle for SAN768B-2 and SAN384B-2

The enterprise software bundle is a bundle of FOS features on top of the base FOS functionality that is included in the hardware base for both the SAN768B-2 and the SAN384B-2. It includes the following features (licenses):

- ▶ Adaptive Networking
- ▶ Advanced Performance Monitoring
- ▶ Extended Fabrics
- ▶ Fabric Watch
- ▶ ISL Trunking
- ▶ Server Application Optimization
- ▶ Fabric Vision
- ▶ FICON with CUP activation
- ▶ Integrated routing
- ▶ Inter-Chassis License with eight 16 Gbps 2 km quad small form factor pluggables (QSFPs)
- ▶ Inter-Chassis License with sixteen 16 Gbps 100m QSFPs
- ▶ 16 Gbps 2 km QSFP
- ▶ Inter-chassis QSFP bundle

- Inter-Chassis License conversion
- Enterprise ICL license

9.8.2 License administration

When you receive a transaction key, you need to retrieve the license ID (LID) of the b-type switch by using the `licenseidshow` command:

```
switch:admin> licenseidshow
a4:f8:69:33:22:00:ea:18
```

Go to the IBM Storage License Keys portal:

<http://www.ibm.com/storage/key>

Select the license activation type **Generate SAN b-type switch feature license key**, and follow the instructions to generate a SAN b-type switch feature license key. See Figure 9-1.

Switch feature activation

Activate a switch feature(s)

The fields indicated with an asterisk (*) are required to complete this transaction; other fields are optional. If you do not want to provide us with the required information, please use the Back button on your browser or close the window or browser session that is displaying this page to return to the previous page.

Your feature activation key(s) will be forwarded to the email address you provide below.

Email notification

* Email address

* Verify email address

* Installation site information

World Wide Name/License ID	Transaction key
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>

Figure 9-1 Switch feature activation

Use your web browser to connect to Web Tools (or **Element Manager** → **Hardware** through IBM Network Advisor) to find the WWN of the switch and manage licensing. After logging in to Web Tools using admin credentials, click **Configure** → **Switch Admin** to open the Switch Administration window.

In the Switch Administration window, the WWN of the switch is displayed, along with a License tab that you can use to manage licensing.

For an in-depth licensing overview, see the *Fabric OS Software Licensing Guide* that is available at the following website:

<http://my.brocade.com>



Fabric administration

This chapter covers zoning, Fibre Channel Routing (FCR), the use of syslog, and Network Time Protocol (NTP) servers. The chapter also covers some typical SAN administration practices and techniques.

This chapter includes the following sections:

- ▶ Administration practices
- ▶ Initial setup
- ▶ Audit and syslog configuration
- ▶ Network Time Protocol
- ▶ Zoning
- ▶ Trunking
- ▶ Fibre Channel over distance
- ▶ Fibre Channel over IP
- ▶ FC-FC routing overview
- ▶ FCIP and FCR
- ▶ Access Gateway and N_Port ID Virtualization
- ▶ Inter-chassis links
- ▶ Fabric OS management
- ▶ Upgrading firmware or rolling back to an earlier version

10.1 Administration practices

Many sources of documentation are available to assist with administration of a fabric for IBM b-type switches and fabrics. A defined management strategy within your organization will improve usability, speed time to implementation of new resources and changes, and reduce time to resolution of unplanned events.

The best resources for the creation and investigation of administration policies, tasks, and problem determination are this publication and others on the IBM Redbooks web site, the Fabric OS administration guides, and various topic-specific guides offered at the following websites:

<http://www.redbooks.ibm.com/>
<https://my.brocade.com/>

A simple way to create a SAN diagram and inventory of your organization's storage area network is to use the b-type SAN Health Assessment tool. For more information, see 12.6, "SAN health" on page 311.

Setting up items such as audit log capture (see 10.3, "Audit and syslog configuration" on page 255), network time protocol server connections (see 10.4, "Network Time Protocol" on page 259), and user IDs (see Chapter 4, "IBM Network Advisor" on page 49) are a few of the techniques that will maximize the SAN resources and provide valuable information. These techniques illustrate a few of the good administration policies that are contained in the Fabric OS administration guides and IBM Redbooks publications.

For more IBM Network Advisor topics such as installation, configuration of backups, and user management, see Chapter 4, "IBM Network Advisor" on page 49.

10.2 Initial setup

Zoning enables you to partition a storage area network (SAN) into logical groups of devices that can access each other.

Before you configure any IBM SAN switch, it must be physically assembled, racked, and connected to the appropriate electrical outlet and network. The hardware requirements and specifications can be found in the specific b-type hardware installation guide that comes with the product.

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

For information about initial installation and configuration, see the installation, service, and user guides for your switch at the IBM product support portal:

<https://www.ibm.com/support/entry/portal/support>

Note: EZSwitchSetup is an easy-to-use graphical user interface application for setting up and managing single switch fabrics.

For full compatibility, see the *EZSwitchSetup Administrator's Guide*, which you can download at the following website:

<http://my.brocade.com>

10.3 Audit and syslog configuration

The b-type audit log and syslog provide valuable information about what has occurred on a switch and in the fabric.

The last 1024 messages are persistently saved in the audit log, but all audit events are sent to the system message log, which (assuming there are no bottlenecks) will be forwarded to your syslog server.

Audit logging assumes that your syslog is operational and running. Before configuring an audit log, ensure that the host syslog is operational and configured to preserve audit log entries past the 1024 message limit. See 10.3.2, “Syslog” on page 256 for more information.

10.3.1 Audit log

The audit log is a collection of information that is created when specific events are identified on an IBM b-type platform. The log can be dumped by running **auditdump** command-line interface (CLI) command, and audit data can also be forwarded to a syslog server for centralized collection.

Information is collected about many different events that are associated with zoning, security, trunking, Fibre Channel over IP (FCIP), FICON, and so on. Each release of the Fabric Operating System (FOS) provides more audit information.

By default, all event classes are configured for audit. To create an audit event log for specific events, you must explicitly set a filter. See the Fabric OS administration and the Fabric OS command guides for specific commands to set filters at. These guides are available at the following website:

<https://my.brocade.com/>

Information is related to event classes tracked and made available. For example, you can track changes from an external source by the user name, IP address, or type of management interface that is used to access the switch.

To view the audit log in IBM Network Advisor, select **Monitor** → **Logs** → **Audit** to display the Audit log window.

Figure 10-1 shows the menu selection to open the Audit log and the opened Audit log.

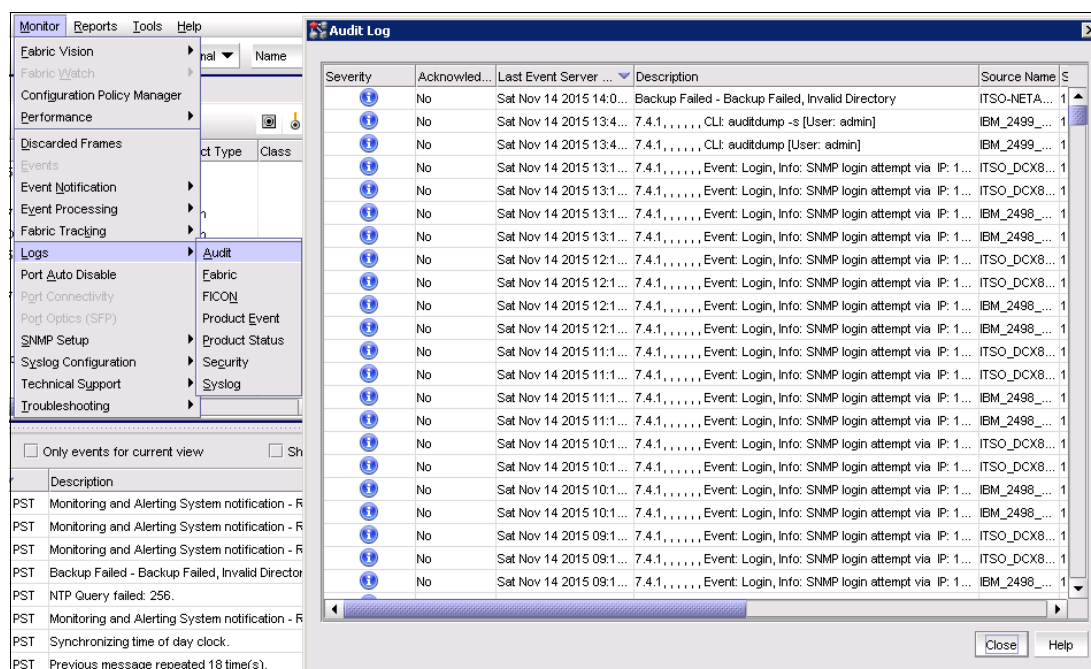


Figure 10-1 Menu section to open Audit log and open audit log pane

10.3.2 Syslog

Fabric OS 7.4.0 supports these functions:

- ▶ Configuring a switch to forward all error log entries to a remote syslog server
- ▶ Setting the syslog facility to a specified log file
- ▶ Removing a syslog server,
- ▶ Displaying the list of configured syslog servers

Brocade switches use the syslog daemon. Up to six servers are supported.

By default, the switch uses the User Datagram Protocol (UDP) to send the error log messages to the syslog server. The default UDP port is 514. However, you can configure the switch to send the error log messages securely by using the Transport Layer Security (TLS) protocol.

For more information, see Chapter 4, “IBM Network Advisor” on page 49.

Setting the syslog recipient

To automatically register the IBM Network Advisor Management application server as the syslog recipient on all managed SAN products, complete the following steps:

1. Start IBM Network Advisor.
 1. Select **Server** → **Options**.
 2. In the Options window, select **Syslog Registration** from the Category pane.
 3. Select **Auto register server as syslog recipient**.

Figure 10-2 shows the Options window with the Syslog Registration category selected.

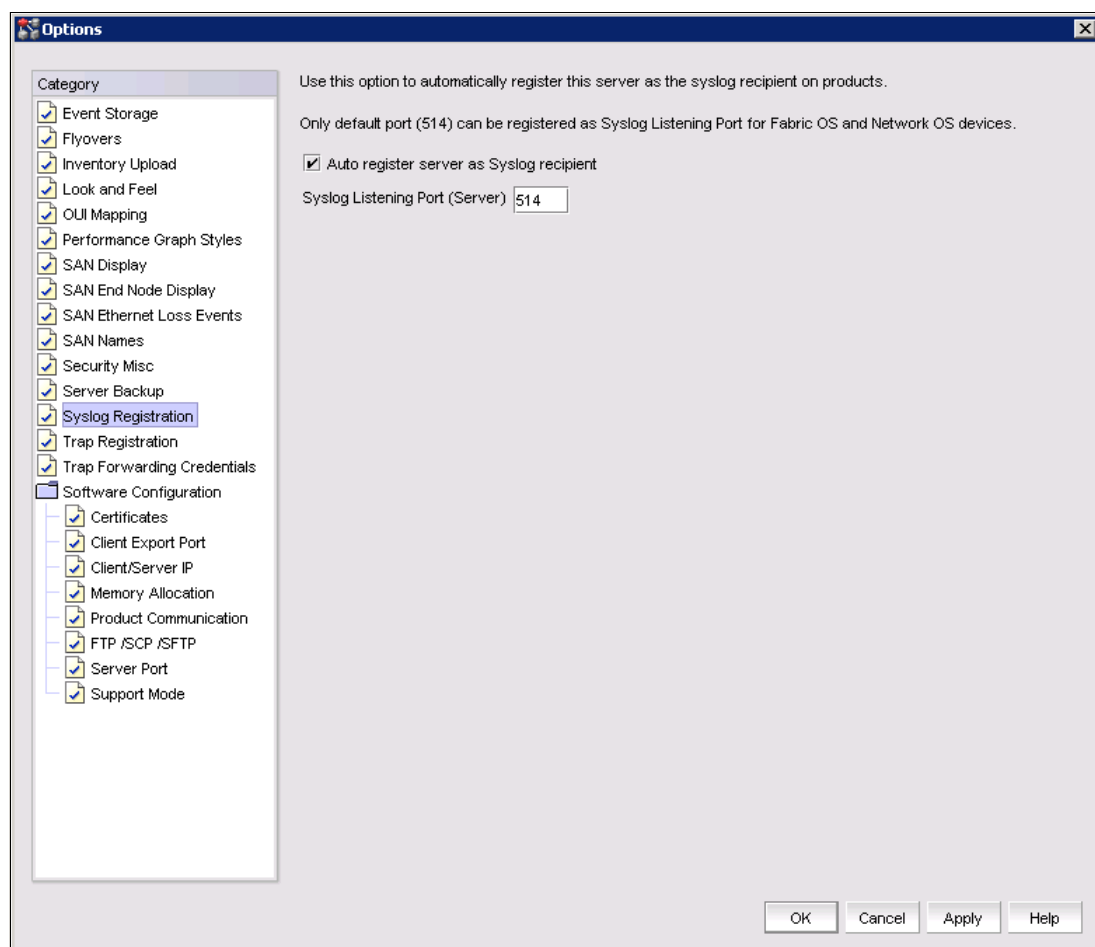


Figure 10-2 Syslog Registration panel

Note: The syslog listening port number is 514 in the default settings. If you change the port number to something else, auto-registration is disabled.

Syslog forwarding

This section describes the process of configuring the syslog server connection by using IBM Network Advisor Management server. The procedure can also be completed by using the CLI. For CLI instructions and examples, see the *Fabric OS Administrators Guide* and *Fabric OS Command Reference* at the following website:

<https://my.brocade.com/>

Complete the following steps in IBM Network Advisor to perform syslog forwarding:

1. Select **Monitor** → **Syslog Configuration** → **Syslog Forwarding**. The Syslog Forwarding window is displayed.
2. Select **Enable Syslog forwarding** as shown in Figure 10-4 on page 259.
3. Click **Add**. The Add Syslog Destination window is displayed as shown in Figure 10-3 on page 258.
4. In the Add Syslog Destination window, enter a name and the IP for the syslog server.

5. Select **Syslog Repeater** if you want to forward all syslogs, whether the source is managed or unmanaged. If the **Syslog Repeater** check box is not selected, syslogs from the managed products are sent to the server. If no filter is selected, then syslogs from all products are sent.

Figure 10-3 shows the Add Syslog Destination window.

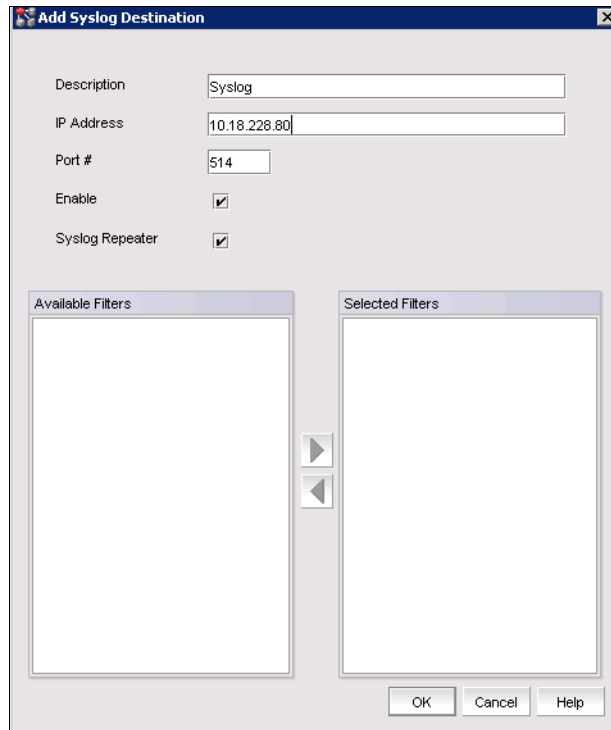


Figure 10-3 Add Syslog Destination window

6. Click **OK** to add the syslog destination.

While enabling secure syslog mode, you must specify a port that is configured to receive the log messages from the switch.

You are returned to the Syslog Forwarding window.

7. Choose not to select a filter (zero) or you can select up to five filters from the **Available Filters** window. Click the right arrow button to move them to the Selected Filters list. This selection is enabled only when **Syslog Repeater** is not selected.

Note: When configuring filters, define them before you add the syslog server to make them available during configuration. For information about defining filters, see the *Network Advisor SAN User Manual* at the following website:

<https://my.brocade.com>

8. Select **OK** to close and apply the settings.

Figure 10-4 shows the Syslog Forwarding window with the syslog server added.

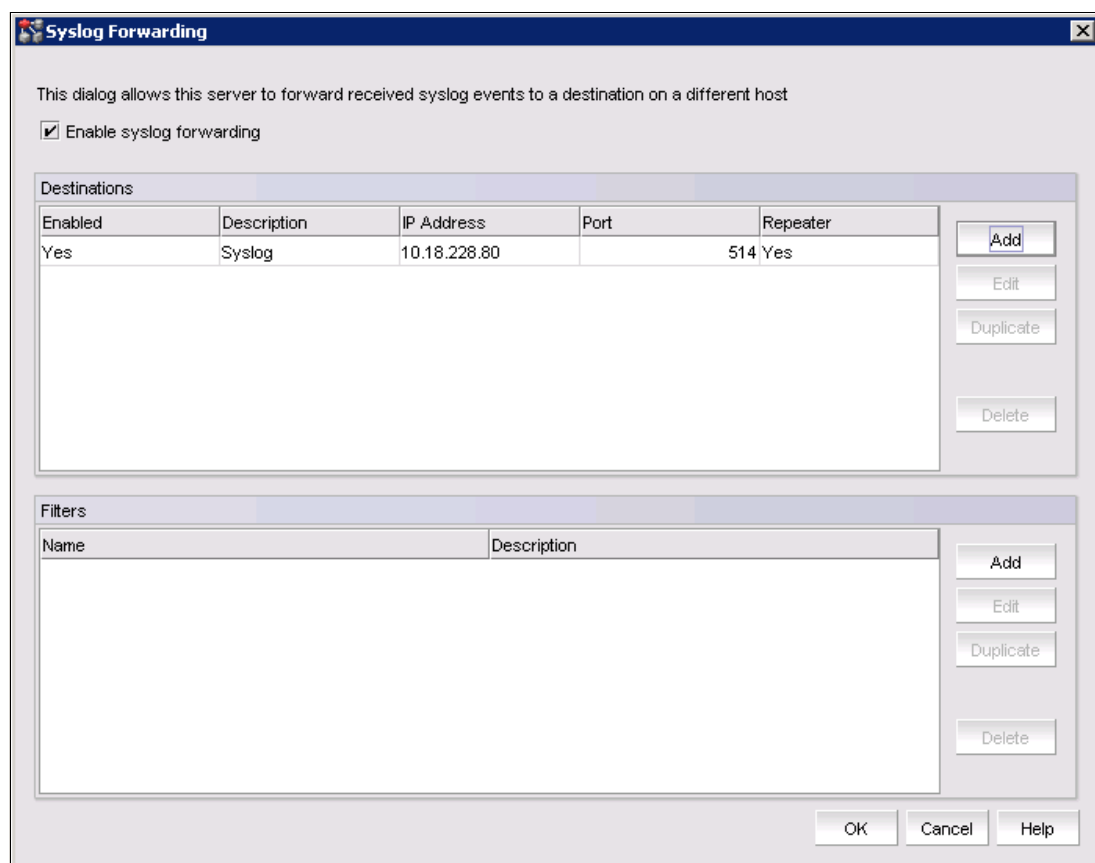


Figure 10-4 Syslog Forwarding window

10.4 Network Time Protocol

Switches maintain the current date and time internally and receive the time from the fabric's principal switch.

Switches with incorrect date and time values work. However, because the date and time are used for logging, error detection, and troubleshooting, synchronizing the local time of the principal or primary FCS switch with at least one external Network Time Protocol (NTP) server is preferred.

If the Virtual Fabrics feature is enabled, the switch behavior has these characteristics:

- ▶ All switches in a chassis must be configured for the same set of NTP servers. This configuration ensures that time does not go out of sync in the chassis. It is not recommended to configure the local server (LOCL) in the NTP server list.
- ▶ Default switches in the fabric can query the NTP server. If Virtual Fabrics is not enabled, only the principal or primary FCS switch can query the NTP server.
- ▶ Logical switches in a chassis receive clock information from the default logical switch, and not from the principal or primary Fabric Configuration Server (FCS) switch.

Complete the following steps to synchronize the local time with NTP:

1. Log in to the switch by using the CLI.
2. Enter the **tsClockServer** command:

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

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

```
switch:admin> tsclockserver  
LOCL  
switch:admin> tsclockserver "132.163.135.131"  
switch:admin>  
switch:admin> tsclockserver  
132.163.135.131  
switch:admin>
```

Note: NTP configuration is also available in IBM Network Advisor by using Configuration and Operational Monitoring Policy Automation Services Suite (COMPASS), which is supported in Professional Plus and Enterprise editions only. For more information about COMPASS, see the *Brocade Network Advisor SAN User Manual* at:

<http://my.brocade.com>

10.5 Zoning

Zoning enables you to partition a SAN into logical groups of devices that can access each other.

A device in a zone can communicate only with other devices that are connected to the fabric within the same zone. A device that is not included in the zone is not available to members of that zone. When zoning is enabled, devices that are not included in any zone configuration are inaccessible to all other devices in the fabric.

Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device that is attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should be utilized only in a small and tightly controlled environment.

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 FOS level to perform zoning tasks.

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 worldwide name (WWN)), zone boundaries are enforced at the hardware level. In these cases, 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. This type of zone is referred to as a soft zone.

10.5.1 Zoning preferred practices

Consider the following preferred practices when you use zoning:

- ▶ Zone using the core switch in preference to using an edge switch.
- ▶ Zone using a backbone instead of a switch. A backbone has more resources to handle zoning changes and implementations.
- ▶ Zone by single host bus adapter (HBA) where practical (see Figure 10-5) unless Peer zoning is in use (see 10.5.2, “Peer Zoning” on page 262). Each zone that is created has only one HBA (initiator) in the zone, and each target device is added to the zone. Typically, a zone is created for the HBA and the disk storage ports are added. In this manner, zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change.
 - If the HBA also accesses tape devices, a second zone is created with the HBA and associated tape devices in it.
 - For clustered systems, it might be appropriate to have an HBA from each of the cluster members included in the zone.
- ▶ When you add a switch to an existing fabric, before joining the fabric, set the default zone (defzone) policy of the switch being added as follows:
 - If the joining switch has locally attached devices that are online, the defzone policy of the switch being added should be set to No Access.
 - If the joining switch has no online locally attached devices, the defzone policy of the switch being added can be set to All Access.

Important: This setting is required to prevent a transitional state where the All Access policy might allow excessive registered state change notification (RSCN) activity. Extreme cases might create the potential for more adverse effects. This setting is for fabrics that have a very high device count.

Figure 10-5 shows examples of preferred zoning practices.




Figure 10-5 Preferred zoning practices

10.5.2 Peer Zoning

Peer Zoning allows a “principal” device to communicate with the rest of the devices in the zone. The principal device manages a Peer Zone. Other “non-principal” devices in the zone can communicate with the principal device only. They cannot communicate with each other.

Before the introduction of Peer Zoning, Single-Initiator Zoning was considered a more efficient zoning method in terms of hardware resources and RSCN volume. However, the added storage requirements of defining a unique zone for each host and target rapidly exceeds zone database size limits. As the number of zones increase, it becomes more difficult to configure and maintain the zones.

In a Peer Zone setup, principal to non-principal device communication is allowed, but non-principal to non-principal device and principal to principal device communication are not allowed. This approach establishes zoning connections that provide the efficiency of Single-Initiator Zoning with the simplicity and lower memory characteristics of One-to-Many Zoning. Figure 10-6 shows a comparison of traditional and peer zoning.




Figure 10-6 Comparison of traditional zoning (left) and peer zoning (right)

Peer Zone connectivity rules

Peer Zoning adheres to the following connectivity rules:

- ▶ Non-principal devices can communicate only with the principal device.
- ▶ Non-principal devices cannot communicate with other non-principal devices, unless allowed by some other zone in the active zone set.
- ▶ Principal devices cannot communicate with other principal devices, unless allowed by some other zone in the active zone set.
- ▶ The maximum number of Peer Zones is determined by the zone database size. The supported maximum zone database size is 2 MB for systems that are running only IBM System Storage SAN768B-2 and SAN384B-2 platforms, and their predecessors. The presence of any other platform reduces the maximum zone database size to 1 MB.

Peer Zone configuration

In the Peer Zone, devices must be identified as either WWN or D,I devices. Mixing WWN and D,I device identification within a single Peer Zone is not allowed. Aliases are not supported as devices for a Peer Zone.

Note: To create a peer zoning by using the CLI see the *Fabric OS Administrators Guide* for version 7.4 and later at the following website:

<https://my.brocade.com/>

To create a Peer Zone by using IBM Network Advisor, complete the following steps:

1. In the IBM Network Advisor window, select **Configure** → **Zoning** → **Fabric** to display the Zoning window.
2. Select the correct zoning scope from **Zoning Scope** menu. (For more information about selecting the most appropriate scope, see “Zoning preferred practices” on page 261.
3. Select **New Peer Zone** from the **New Zone** menu. Figure 10-7 shows the New Peer Zone option.

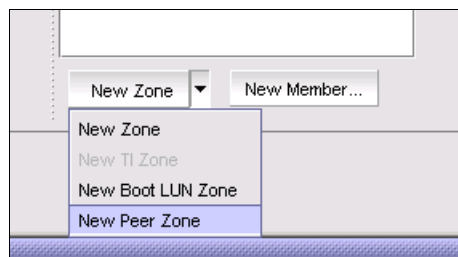


Figure 10-7 New Peer Zone pull-down option

4. Provide a meaningful name for the new Peer zone in the **Peer Zone Name** field.
5. Select a principal device or devices from the Potential Members pane and click the upper right facing arrow (see Figure 10-8 on page 264) to add it to the Principal Members pane.
6. Select an additional device or devices in the Potential Members pane and click the lower right facing arrow to add them to the Peer Members pane (see Figure 10-8 on page 264).
7. Click **OK** to create the peer zone and return to the Zoning panel.

The peer zone now appears in the list of available zones that can be added to a zoning configuration and activated in the Zoning panel.

Figure 10-8 shows the Add Peer Zone window.

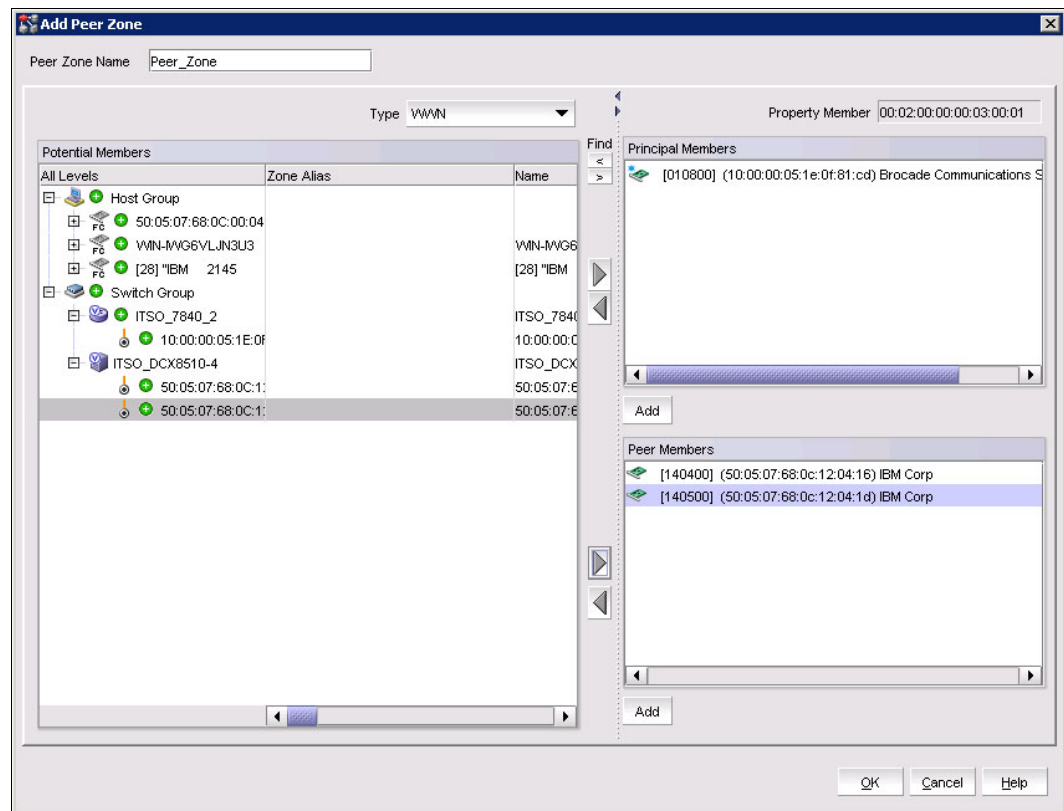


Figure 10-8 Add Peer Zone window

10.6 Trunking

Trunking optimizes the use of bandwidth by allowing a group of links to merge into a single logical link, called a trunk group. Traffic is distributed dynamically and in order over this trunk group, achieving greater performance with fewer links. Within the trunk group, multiple physical ports appear as a single port, thus simplifying management. Trunking also improves system reliability by maintaining in-order delivery of data and avoiding I/O retries if one link within the trunk group fails.

Trunking is frame-based instead of exchange-based. Because a frame is much smaller than an exchange, frame-based trunks are more granular and better balanced than exchange-based trunks and provide maximum utilization of links.

For information about trunking and trunking requirements, see Chapter 4, “IBM Network Advisor” on page 49.

10.6.1 Configuring trunk groups

After the Trunking license is installed, the ports that are to be used in trunk groups must be reinitialized (cycle the port through the offline state). This procedure needs to be performed only once, and is required for all types of trunking. Alternatively, the switch can be disabled and enabled.

Displaying trunking

Ports that are involved in trunking configurations can be displayed by using both IBM Network Advisor and the CLI.

To view trunk ports in CLI, complete the following steps:

1. Connect to the switch and log in using an account that is assigned to the admin role.
2. Enter the **trunkShow** command (see Example 10-1).

Example 10-1 shows trunking groups 1, 2, and 3; ports 4, 13, and 14 are masters.

```
switch:admin> trunkshow
1: 6-> 4 10:00:00:60:69:51:43:04 99 deskew 15 MASTER
2: 15-> 13 10:00:00:60:69:51:43:04 99 deskew 16 MASTER
      12-> 12 10:00:00:60:69:51:43:04 99 deskew 15
      14-> 14 10:00:00:60:69:51:43:04 99 deskew 17
      13-> 15 10:00:00:60:69:51:43:04 99 deskew 16
3: 24-> 14 10:00:00:60:69:51:42:dd 2 deskew 15 MASTER
```

Ports that are involved in trunking configurations show in a number of areas of IBM Network Advisor. The easiest place to view trunking configurations is in the product view of the main IBM Network Advisor window.

To view the ports that are used in a trunk configuration, complete the following steps:

1. Open IBM Network Advisor and log in using an account assigned to the admin role.
2. Click the SAN tab to display the product list pane.
3. In the product pane list, click the + symbol to expand the port view for the selected switch.
4. Ports that are part of a trunk group will display as a trunk port, as shown in Figure 10-9.




Figure 10-9 Product list pane

10.6.2 Enabling Trunking

Trunking can be enabled for a single port or for an entire switch. Trunking is automatically enabled when you install the Trunking license. The procedure is only required if trunking has been disabled on a port or switch. Enabling trunking disables and reenables the affected ports. As a result, traffic through these ports might be temporarily disrupted.

Complete the following steps to enable trunking with the CLI:

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **portCfgTrunkPort** command to enable trunking on a port. Enter the **switchCfgTrunk** command to enable trunking on all ports on the switch.

```
portcfgtrunkport[slot/]port mode  
switchcfgtrunk mode
```

Mode 1 enables trunking. Example 10-2 shows trunking being enabled.

Example 10-2 Trunking is being enabled on slot 1, port 3.

```
switch:admin> portcfgtrunkport 1/3 1
```

To enable port trunking with IBM Network Advisor, complete the following steps:

1. Open IBM Network Advisor and log in using an account that is assigned to the admin role.
2. Select the SAN tab to display the Product list pane.
3. Right-click the switch that you want to enable trunking for, and select **Element Manger** → **Ports**. The Web Tools window is displayed with the FC Ports tab selected as shown in Figure 10-10 on page 267.
4. Change the window view to Advanced by selecting **View** and then **Advanced** to enable the required menu options for the following steps.
5. Highlight the port for which trunking is to be enabled in the FC Ports Explorer pane.

- Click **Actions** and select **Trunking** → **Enable**. Ensure that trunking is enabled for that port.

Figure 10-10 shows the enable trunking option for a port.

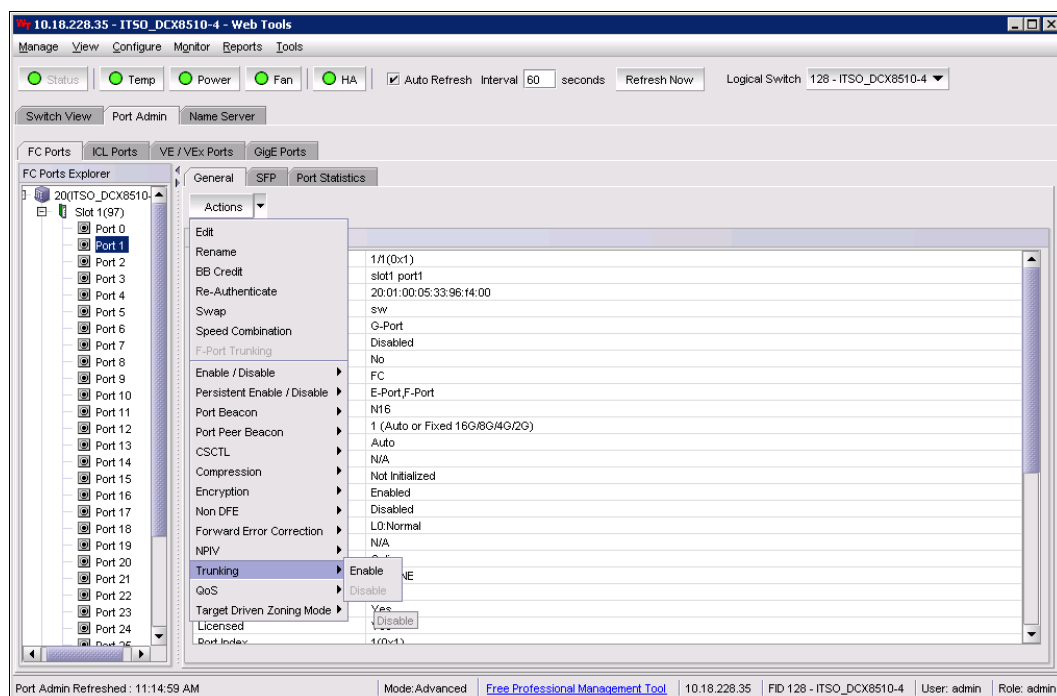


Figure 10-10 Enable the trunking option in Web Tools

10.6.3 Types of trunking

Trunking can be between two switches, between a switch and an Access Gateway module, or between a switch and a compatible adapter. The following configurations are possible:

- ▶ ISL trunking, or E_Port trunking, is configured on an inter-switch link (ISL) between two Fabric OS switches and is applicable only to E_Ports.
- ▶ ICL trunking is configured on an inter-chassis link (ICL) between two IBM System Networking SAN384B-2 (2499-416) and IBM System Networking SAN768B-2 (2499-816) and is applicable only to ports on the core blades.
- ▶ EX_Port trunking is configured on an inter-fabric link (IFL) between an FC router (EX_Port) and an edge fabric (E_Port). The trunk ports are EX_Ports connected to E_Ports.
- ▶ F_Port trunking is configured on a link between a switch and either an Access Gateway module or a Brocade adapter. The trunk ports are F_Ports (on the switch) connected to N_Ports (on the Access Gateway or adapter).
- ▶ N_Port trunking is configured on a link between a switch and either an Access Gateway module or a compatible adapter. It is similar to F_Port trunking. The trunk ports are N_Ports (on the Access Gateway or adapter) connected to F_Ports (on the switch).

ISL trunking is the typical use case. Detailed information is available in the *Fabric OS Administrators Guide*, which is available at the following website:

<https://my.brocade.com>

10.7 Fibre Channel over distance

When you need fabric connectivity over longer distances, SANs are typically connected over metro or long-distance networks. In both cases, path latency is critical for mirroring and replication solutions. For native Fibre Channel links, the amount of time that a frame spends on the cable between two ports is negligible because that aspect of the connection speed is limited only by the speed of light. The speed of light in optics amounts to approximately 5 microseconds per kilometer, which is negligible compared to the typical disk latency of 5 - 10 milliseconds. The Extended Fabrics feature enables full-bandwidth performance across distances spanning up to hundreds of kilometers. It extends the distance ISLs can reach over an extended fiber by providing enough buffer credits on each side of the link to compensate for latency that is introduced by the extended distance.

10.7.1 Buffer credits

Buffer credits are a measure of frame counts, and are not dependent on the data size (a 64-byte and a 2-KB frame both consume a single buffer). Consider the following parameters when allocating buffers for long-distance links that are connected through dark fiber or through a D/CWDM in a pass-through mode:

- ▶ Round Trip Time (RTT) (that is, the distance)
- ▶ Frame processing time
- ▶ Frame transmission time

Below are some suggested guidelines for calculating the number of required buffer credits:

- ▶ Number of credits = $6 + ((\text{link speed Gbps} * \text{Distance in KM}) / \text{frame size in KB})$.
Example: 100 KM @2k frame size = $6 + ((8 \text{ Gbps} * 100) / 2) = 406$
- ▶ A buffer model should be based on the average frame size.
- ▶ If compression is used, the number of buffer credits that is needed is 2x the number of credits without compression.

On the IBM b-type 16 Gbps backbones platform, 4 K buffers are available per ASIC to drive the 16 Gbps line rate to 500 KM at a 2 KB frame size. Additional control is available when LD or LS links are configured. Using these links allows the number of required buffer credits to be specified based on the anticipated or real average frame size for a long-distance port. Using the frame size option, the number of buffer credits that are required for a port is automatically calculated. These options give extra flexibility to optimize performance on long-distance links.

In addition, FOS 7.1 and later provides commands to gain better insight into long-distance link traffic patterns by displaying the average buffer usage and average frame size through the CLI.

The **portBufferCalc** command can automatically calculate the number of buffers that are required per port given the distance, speed, and frame size. The number of buffers that is calculated by this command can be used when configuring long-distance ports.

Note: If no options are specified, then the current port's configuration is used to calculate the number of buffers that are required.

10.7.2 Fabric interconnectivity over Fibre Channel at longer distances

SANs spanning data centers in different physical locations can be connected through dark fiber connections by using Extended Fabrics with wave division multiplexing, such as dense wavelength division multiplexing (DWDM), coarse division multiplexing (CWDM), and time-division multiplexing (TDM). Extended Fabrics is a FOS optionally licensed feature. This situation is similar to connecting switches in the data center with one exception: Additional buffers are allocated to E_Ports connecting over distance. The Extended Fabrics feature extends the distance that the ISLs can reach over an extended fiber. This task is accomplished by providing enough buffer credits on each side of the link to compensate for the latency that is introduced by the extended distance.

Any of the first eight ports on the 16 Gbps port blade can be set to 10 Gbps FC for connecting to a 10 Gbps line card D/CWDM without needing a specialty line card. If you connect to DWDMs in a pass-through mode where the switch is providing all the buffering, a 16 Gbps line rate can be used for higher performance.

Extended Fabrics device limitations

It is suggested that FC8-64 and FC16-64 port blades not be used for long distance because of their limited buffers. These blades do not support long-wavelength (LWL) fiber optics and therefore only support limited distance. However, the ports can still be configured to reserve frame buffers for the ports that are intended to be used in long-distance mode through DWDM.

Note: A limited number of reserved buffers can be used for long distance for each blade. If some ports are configured in long-distance mode and have buffers reserved for them, insufficient buffers might remain for the other ports. In this case, some of the remaining ports might come up in degraded mode.

10.8 Fibre Channel over IP

FCIP enables you to use the existing IP wide area network (WAN) infrastructure to connect Fibre Channel SANs.

The TCP connections ensure in-order delivery of Fibre Channel (FC) frames and lossless transmission. The Fibre Channel fabric and all Fibre Channel targets and initiators are unaware of the presence of the IP WAN.

The advantage of this configuration is that the existing, less expensive IP WAN network can be used to encapsulate and transmit the Fibre Channel frame over long distances while still ensuring the high integrity required for storage I/O.

Figure 10-11 shows the Fibre Channel frame being encapsulated and decapsulated for transport across the WAN.

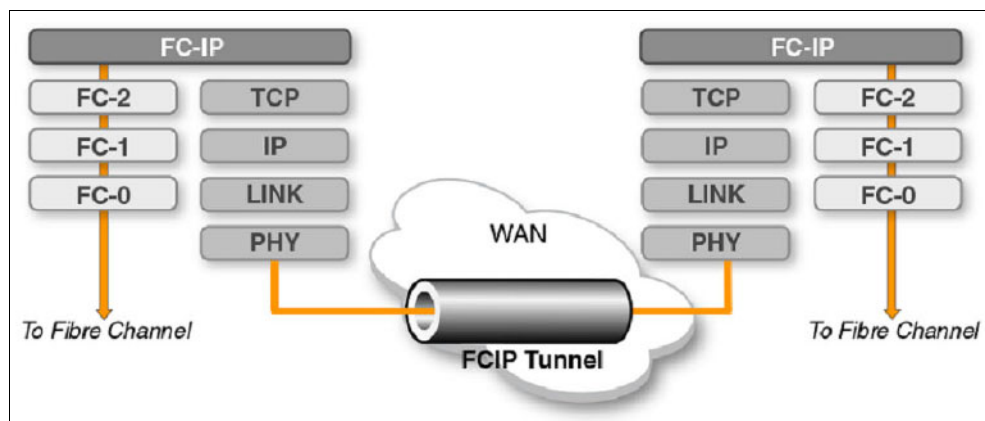


Figure 10-11 FCIP encapsulation

The extension switch or blade accomplishes the FCIP connection by establishing a tunnel that emulates FC ports on each end of the tunnel.

When the tunnel is configured and the connection is established between the devices, a circuit is formed. Either an ISL or an IFL can be created depending on the configuration that you want.

When the circuit is formed and functioning, the virtual FC that is created can be configured as a Virtual E port (VE-Port) or an FCR connection, a Virtual Extended port (VEX-port) to link the two ends.

Figure 10-12 shows the tunnel and circuit concept.




Figure 10-12 Tunnel and circuits concept

10.8.1 Configuring an FCIP tunnel

The detailed steps for FCIP tunnel configuration are provided in the *Fabric OS FCIP Administrator's Guide* that you can download at the following website:

<http://my.brocade.com>

The following is the general outline of those steps:

1. Persistently disable VE_Ports.
2. If required, configure VEX_Ports. For more information, see 10.10, "FCIP and FCR" on page 274.
3. Set the media type or operating mode.
4. Create an IP interface (IPIF) for each circuit that you want on a port by assigning an IP address, netmask, and an IP maximum transfer unit (MTU) size to an Ethernet port.

5. Create one or more IP routes to a port.
6. Test the IP connection.
7. Create the FCIP tunnel or tunnels.

Note: Configuring a tunnel automatically configures circuit 0 for the tunnel.

8. Persistently enable the VE_Ports.

10.9 FC-FC routing overview

Fibre Channel routing allows two or more FC devices to communicate across fabrics without requiring them to merge.

For routing to take place, there must be either a router or routing blade, or the switch must have the capability to perform integrated routing.

The integrated routing license allows ports in supported switches to be configured as EX_Ports or VEX_Ports supporting Fibre Channel Routing. This configuration eliminates the need to add a routing blade or the need for an external router for Fibre Channel Routing (FCR) purposes.

Routing is accomplished by creating virtual domains in the remote switch to facilitate traffic flow into and out of the individual fabrics. The virtual domains are known as translate (XD) or front (FD) domains, depending on where they are:

- ▶ The XD domain is in the switch or blade that is responsible for routing between fabrics.
- ▶ The FD domain is projected into the remote fabric.

See Figure 10-13 for an example.




Figure 10-13 Routing concept

Using FC-FC routing, you can share resources across multiple fabrics without the administrative problems, such as change management, network management, scalability, reliability, availability, and serviceability, that might result from merging the fabrics.

A Fibre Channel router (FC router) is a switch that is running the FC-FC routing service. The FC-FC routing service can be simultaneously used as an FC router and as a SAN extension over WANs by using FCIP.

10.9.1 Setting up FC to FC routing

Complete the following tasks to set up and configure FC to FC routing:

1. Verify that you have the proper setup for FC to FC routing.
2. Assign backbone fabric IDs.
3. Configure FCIP tunnels if you are connecting Fibre Channel SANs over IP-based networks.
4. Configure IFLs for edge and backbone fabric connection.
5. Modify port cost for EX_Ports, if you want to change from the default settings.
6. Enable shortest IFL mode if you want to choose a lowest cost IFL path in the backbone fabric.
7. Configure trunking on EX_Ports that are connected to the same edge fabric.
8. Configure Logical SAN (LSAN) zones to enable communication between devices in different fabrics.

The detailed commands for each step are documented in the *Fabric OS Administrator's Guide* at the following link:

<http://my.brocade.com>

10.9.2 Logical SAN Zones

An LSAN consists of zones in two or more edge or backbone fabrics that contain the same devices. LSANs essentially provide selective device connectivity between fabrics without requiring the fabrics to merge.

Note: A backbone fabric consists of one or more FC switches with configured EX_Ports. These EX_Ports in the backbone connect to edge fabric switches through E_Ports. This type of EX_Port-to-E_Port connectivity is called an IFL.

To enable device sharing across multiple fabrics, LSAN zones must be created on the edge fabrics, and optionally on the backbone fabric as well. Use normal zoning operations to create zones (see 10.5, “Zoning” on page 260) with names that begin with the special prefix “LSAN_”, and adding host and target port WWNs from both local and remote fabrics to each local zone as wanted. Zones on the backbone and on multiple edge fabrics that share a common set of devices will be recognized as constituting a single multi-fabric LSAN zone. The devices that they have in common will then be able to communicate with each other across fabric boundaries.

SAN zone members in all fabrics must be identified by their WWN. You cannot use the port IDs that are supported only in Fabric OS fabrics.

Note: The name of an LSAN zone begins with the prefix “LSAN_”. The LSAN name is case-insensitive, so `lsan_` is equivalent to `LSAN_`, `Lsan_`, and so on.

Peer LSAN zone support

Starting with Fabric OS 7.4.0, the FC router supports the peer LSAN zones if configured in the edge fabric. Peer zoning rules are applied by the edge fabric switch. The FC router treats peer LSAN zones as normal LSAN zones and imports the devices in the edge fabric as per a pair-matching algorithm. For more information, see 10.5.2, “Peer Zoning” on page 262.

10.9.3 Fibre Channel routing and virtual fabrics

If the Virtual Fabrics feature is enabled (see Chapter 8, “Virtual Fabrics” on page 213), then in the FC-FC routing context, a base switch is like a backbone switch and a base fabric is like a backbone fabric.

If Virtual Fabrics is enabled, the following rules apply:

- ▶ EX_Ports and VEX_Ports can be configured only on the base switch.
- ▶ When Virtual Fabrics are enabled, the chassis is automatically rebooted. When the switch comes up, only one default logical switch is present, with the default fabric ID (FID) of 128. All previously configured EX_Ports and VEX_Ports are persistently disabled with the reason that ExPort is a non-base switch. A base switch must be explicitly created and the EX_Ports and VEX_Ports moved to that base switch, before you can enable the ports.
- ▶ If EX_Ports or VEX_Ports are moved to any logical switch other than the base switch, these ports are automatically disabled.
- ▶ EX_Ports can connect to a logical switch that is in the same chassis or in a different chassis. However, the following configuration rules apply:
 - If the logical switch is on the same chassis, the EX_Port FID must be set to a different value than the FID of the logical switch to which it is connecting.
 - If the logical switch is on a different chassis, no FID for any logical switch in the FC router backbone fabric can be the same as the FID of the logical switch to which the EX_Port is connecting.
- ▶ If an EX_Port or VEX_Port is connected to an edge fabric, ensure that there are no logical switches with extended ISL (XISL) use enabled in that edge fabric. If any logical switch in the edge fabric allows XISL use, then the EX_Port or VEX_Port is disabled.
- ▶ Backbone-to-edge routing is not supported in the base switch.
- ▶ All FC router commands can be executed only in the base switch context.
- ▶ The **fcrConfigure** command is not allowed when Virtual Fabrics is enabled. Instead, use the **lsCfg** command to configure the FID.
- ▶ Although the IBM System Networking SAN48B-5 and IBM System Networking SAN96B-5 support up to four logical switches, if you are using FC-FC routing, they can have a maximum of only three logical switches.

10.10 FCIP and FCR

The FCIP tunnel traditionally traverses a WAN or IP cloud, which can have characteristics that adversely impact a Fibre Channel network. The FCIP link across a WAN is essentially an FC ISL over an IP link. In any design, it should be considered an FC ISL.

Repeated flapping of a WAN connection can cause disruption in directly connected fabrics. This disruption might come about from many fabric services trying to reconverge repeatedly. This situation causes the processor on the switch or director to go to full capacity.

If the processor can no longer process the various tasks that are required to operate a fabric, an outage might occur. If you limit the fabric services to within the local fabric itself and do not allow them to span across the WAN, you can prevent this situation from occurring.

FCR provides a termination point for fabric services, referred to as a *demarcation point*. EX_Ports and VEX_Ports are demarcation points in which fabric services are terminated, forming the “edge” of the fabric. A fabric that is isolated in such a way is referred to as an *edge fabric*. There is a special case in which the edge fabric includes the WAN link because a VEX_Port was used. This type of edge fabric is referred to as a *remote edge fabric*.

FCR does not need to be used unless a production fabric must be isolated from WAN outages. When connecting to array ports directly for Remote Data Replication (RDR), FCR provides no benefit. Mainframe environments are precluded from using FCR, as it is not supported by FICON.

When a mainframe host writes to a volume on the direct access storage device (DASD), and that DASD performs RDR to another DASD, then DASD to DASD traffic is not using FICON. It is using an open systems RDR application such as IBM Metro Mirror or Global Mirror. These open-system RDR applications can use FCR, even though the volumes they are replicating are written by the FICON host.

The following are some basic FCR architectures:

- ▶ No FCR or one large fabric: This type of architecture is used with mainframes and when the channel extenders are directly connected to the storage arrays.
- ▶ Edge-backbone-edge: Edge fabrics bookend a transit backbone between them.
- ▶ VEX_Port: When a VEX_Port is used, the resulting architecture can be either backbone-remote edge or edge-backbone-remote edge, depending on whether devices are connected directly to the backbone or an edge fabric hangs from the backbone. Both are possible.

10.10.1 Using EX_Ports and VEX_Ports

If an FCR architecture is indicated, an “X” port is needed. An “X” port is a generic reference for an EX_Port or a VEX_Port. The only difference between an EX_Port and a VEX_Port is that the “V” indicates that it is FCIP-facing. The same holds true for E_Ports and VE_Ports; VE_Ports are E_Ports that are FCIP-facing.

The preferred practice in an FC routed environment is to build an edge fabric to backbone to edge fabric (EBE) topology. This topology provides isolation of fabric services in both edge fabrics. This topology requires an EX_Port from the backbone to connect to an E_Port in the edge fabric, as shown in Figure 10-14. The backbone fabric continues to be exposed to faults in the WAN connections. However, because its scope is limited by the VE_Ports in each edge fabric, and because edge fabric services are not exposed to the backbone, it does not pose any risk of disruption to the edge fabrics in terms of overrunning the processors or causing a fabric service to become unavailable. The edge fabric services do not span the backbone. See Figure 10-14.




Figure 10-14 Edge-backbone-edge FCR architecture

There might be cases in which an EBE topology cannot be accommodated. Alternatively, the main production fabric can be isolated from aberrant WAN behavior while allowing the backup site to remain exposed. This topology provides a greater degree of availability and less risk compared to not using FCR at all. This topology uses VEX_Ports that connect to a remote edge fabric. The point is that the remote edge fabric continues to be connected to the WAN, and the fabric services span the WAN all the way to the EX_Port demarcation point. The fabric services that span the WAN are subject to disruption and repeated reconvergence, which can result in an outage within the remote edge fabric. This situation might not be of great concern if the remote edge fabric is not being used for production (merely for backup) because such WAN fluctuations are not ongoing.

You can build two topologies from remote edge fabrics. In the first, shown in Figure 10-15, production devices are attached directly to the backbone. In the second, shown in Figure 10-16 on page 276, the backbone connects to a local edge fabric. In both cases, the other side is connected to a remote edge fabric through a VEX_Port. Also in both cases, the production fabrics are isolated from the WAN. Between the two architectures, the second architecture with the edge fabric is preferred for higher scalability. The scalability of connecting devices directly to the backbone is relatively limited.

Figure 10-15 shows an example of Backbone-remote edge architecture

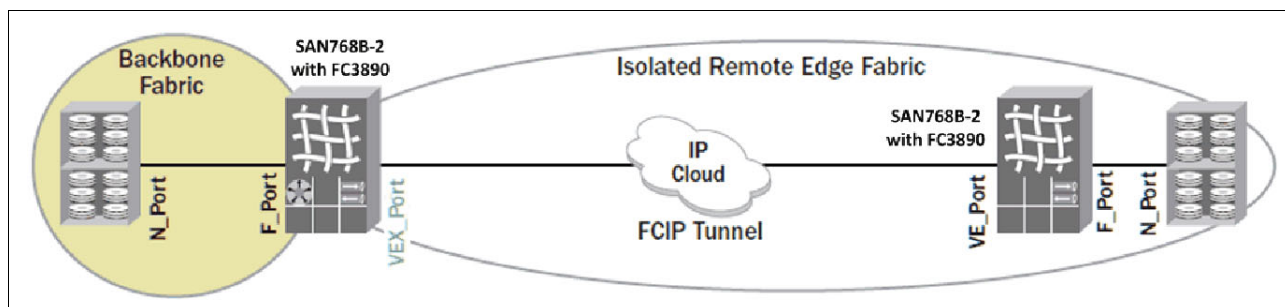


Figure 10-15 Backbone-remote edge architecture

Figure 10-16 shows an example of Edge-remote edge architecture.

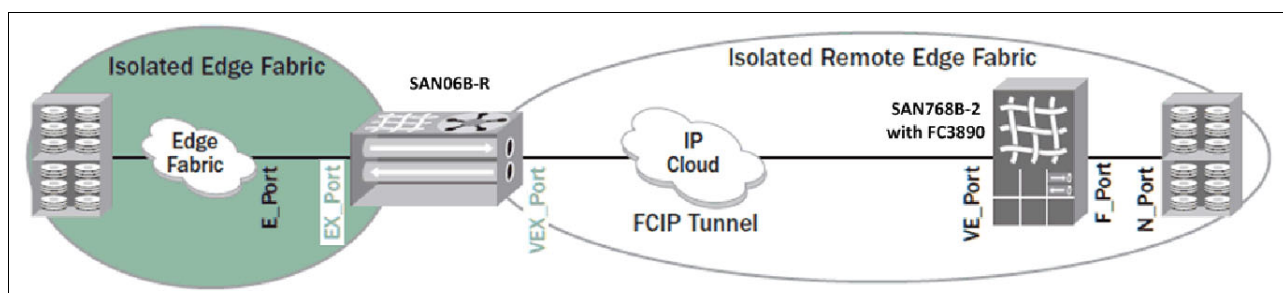


Figure 10-16 Edge-remote edge architecture

There are more considerations with “X” ports. When FC routing is configured, the path from initiator to target might pass through only one “X” port. Paths with more than one “X” port are not supported.

Figure 10-17 shows supported and unsupported paths through FC routing configurations.

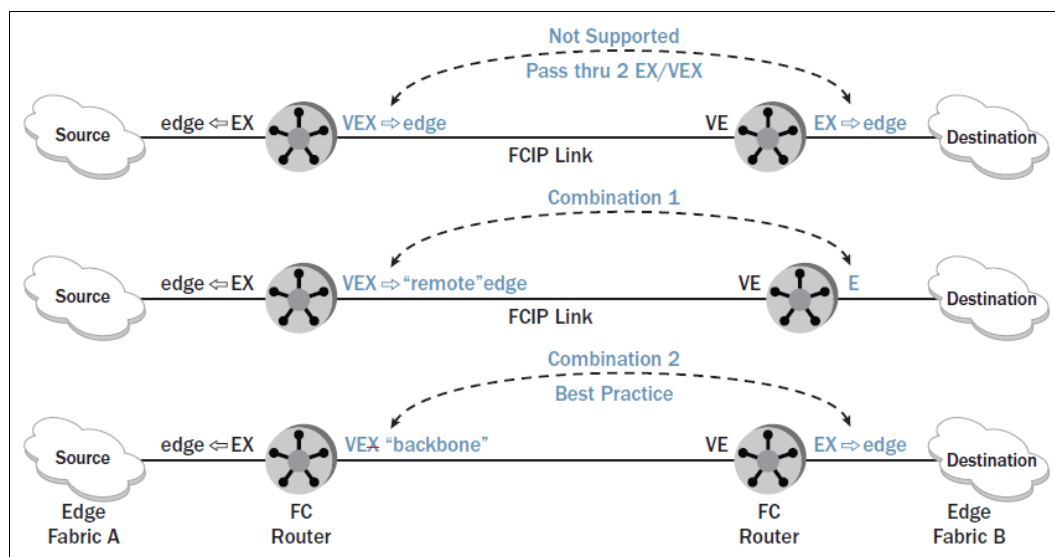


Figure 10-17 Supported and unsupported paths through FC routing configurations

The Integrated Routing (IR) license, which enables FCR on IBM b-type switches and directors, is needed only on the switches or directors that implement the “X” ports. Any switches or directors that connect to “X” ports and have no “X” ports of their own do not need the IR license. The IR license is not needed on the E_Port/VE_Port side to connect to the EX_Port/VEX_Port side.

Tip: VEX ports are not supported on the SAN42B-R. An alternative approach is to use VE to VE FCIP connections.

10.11 Access Gateway and N_Port ID Virtualization

One of the main limits to Fibre Channel scalability is the maximum number of domains (individual physical or virtual switches) in a fabric. Keeping the number of domains low reduces much of the impact that is typically attributed to SAN fabrics. Small-domain-count fabrics are more reliable, perform better, and are easier to manage. When configuring the edge switches in Access Gateway (NPV for Cisco or Transparent for Qlogic) mode, the mode eliminates the usage of a domain ID.

In an environment with multivendor edge switches, the N_Port ID Virtualization (NPIV) feature allows those switches to be presented in a transparent mode to the fabric, which reduces interoperability issues.

Figure 10-18 illustrates the difference between traditional E port connectivity and a switch that is configured to function as an access gateway.

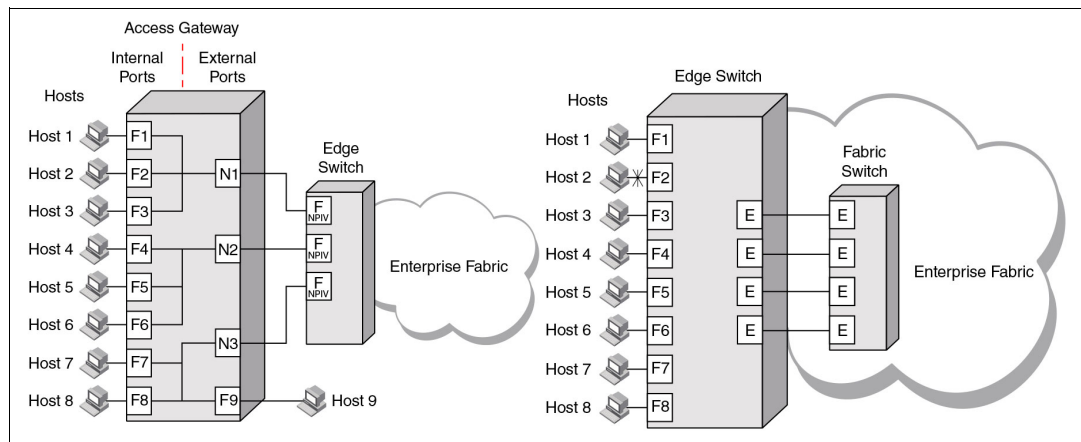


Figure 10-18 Differences between E port connectivity (right) and Access Gateway connectivity (left)

NPIV enables a single Fibre Channel protocol port to appear as multiple, distinct ports, providing separate port identification within the fabric for each operating system image behind the port. This configuration appears as though each operating system image had its own unique physical port. NPIV assigns a different virtual port ID to each Fibre Channel protocol device. NPIV enables you to allocate virtual addresses without affecting your existing hardware implementation. The virtual port has the same properties as an N_Port, and can register with all services of the fabric.

Each NPIV device has a unique device port ID (PID), Port WWN, and Node WWN, and behaves the same as all other physical devices in the fabric. Multiple virtual devices that are emulated by NPIV appear no different from regular devices that are connected to a non-NPIV port.

The same zoning rules apply to NPIV devices as non-NPIV devices. Zones can be defined by domain, port notation, by WWN zoning, or a combination of these. However, to perform zoning to the granularity of the virtual N_Port IDs, you must use WWN-based zoning.

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

Note: For more information, see the *Fabric OS Administrator's Guide* or the *Access Gateway Administrator's Guide*, which you can find at the following website:

<http://brocade.com>

10.12 Inter-chassis links

ICLs are high-performance ports for interconnecting multiple backbones, enabling industry-leading scalability while preserving ports for server and storage connections. Optical UltraScale ICLs, based on Quad Small Form Factor Pluggable (QSFP) technology, connect the core routing blades of two backbone chassis. Each QSFP-based ICL port combines four 16 Gbps links, providing up to 64 Gbps of throughput within a single cable. It offers up to 32 QSFP ports in a Brocade DCX 8510-8 chassis or 16 QSFP ports in a Brocade DCX 8510-4 chassis, with up to 2 Tbps ICL bandwidth and support for up to 100 meters on universal optical cables. The optical form factor of the QSFP-based ICL technology offers several advantages over the copper-based ICL design in the original platforms.

The second generation increased the supported ICL cable distance from 2 meters to 50 meters. The 100-meter ICL is supported beginning in Fabric OS 7.1.0, when using 100-meter-capable QSFPs over OM4 cable only, providing greater architectural design flexibility.

Note: Before Fabric OS 7.3.0, all the FE ports and ICL ports used the same buffer credit model. In Fabric OS 7.3.0 and later, ICL ports support a 2 km distance. To support this distance, you must use specific QSFPs and allocate a greater number of buffer credits per port.

The combination of four cables into a single QSFP provides incredible flexibility for deploying various different topologies, including a 9-chassis full-mesh design with only a single hop between any two points within the fabric.

In addition to these significant advances in ICL technology, the ICL capability still provides dramatic reduction in the number of ISL cables that are required, a four to one reduction compared to traditional ISLs with the same amount of interconnect bandwidth. Because the QSFP-based ICL connections are on the core routing blades instead of consuming traditional ports on the port blades, up to 33% more FC ports are available for server and storage connectivity.

ICL Ports on Demand are licensed in increments of 16 ICL ports. Connecting five or more chassis through ICLs requires an Enterprise ICL license.

10.12.1 Supported topologies

Two network topologies are supported by SAN768B-2 and SAN384B-2 platforms and UltraScale ICLs: Core/edge and mesh. Both topologies deliver unprecedented scalability while reducing ISL cables.

10.12.2 QSFP-based ICL connection requirements

To connect multiple b-type chassis through ICLs, a minimum of four ICL ports (two on each core blade) must be connected between each chassis pair. With 32 ICL ports available on the SAN768B-2 (with both ICL POD licenses installed), this configuration supports ICL connectivity with up to eight other chassis and at least 256 Gbps of bandwidth to each connected 16 Gbps b-type backbones.

Figure 10-19 shows a diagram of the minimum connectivity between a pair of SAN768B-2 chassis. The physical location of ICL connections might be different from what is shown here, but there should be at least two connections per core blade.




Figure 10-19 Minimum connections that are needed between a pair of SAN768B-2 chassis

The dual connections on each core blade must be within the same ICL trunk boundary on the core blades. ICL trunk boundaries are described in detail in 10.12.3, “ICL trunking and trunk groups” on page 280. If more than four ICL connections are required between a pair of SAN768B-2/SAN384B-2 chassis, add ICL connections in pairs (one on each core blade).

ICL connection preferred practice: Each core blade in a chassis must be connected to each of the two core blades in the destination chassis to achieve full redundancy. (For redundancy, use at least one pair of links between two core blades.)

A maximum of 16 ICL connections or ICL trunk groups between any pair of SAN768B-2/SAN384B-2 chassis is supported, unless they are deployed by using Virtual Fabrics, where a maximum of 16 ICL connections or trunks can be assigned to a single Logical Switch. This limitation is because of the maximum supported number of connections for FSPF routing. Effectively, there should never be more than 16 ICL connections or trunks between a pair of SAN768B-2/SAN384B-2 chassis, unless Virtual Fabrics is enabled and the ICLs are assigned to two or more Logical Switches. The exception to this rule is if eight port trunks are created between a pair of SAN768B-2/SAN384B-2 chassis, as described in 10.12.3, “ICL trunking and trunk groups”.

QSFP-based ICLs and traditional ISLs are not concurrently supported between a single pair of SAN768B-2/SAN384B-2 chassis. All inter-chassis connectivity between any pair of SAN768B-2/SAN384B-2 chassis must be done by using either ISLs or ICLs. The final layout and design of ICL interconnectivity is determined by the customer's unique requirements and needs, which dictate the ideal number and placement of ICL connections between SAN768B-2/SAN384B-2 chassis.

10.12.3 ICL trunking and trunk groups

Trunking involves taking multiple physical connections between a chassis or switch pair and forming a single “virtual” connection, aggregating the bandwidth for traffic to traverse across. This section describes the trunking capability that is used with the QSFP-based ICL ports on the IBM b-type 16 Gbps chassis platforms. Trunking is enabled automatically for ICL ports, and cannot be disabled by the user.

Each QSFP-based ICL port has four independent 16 Gbps links, each of which terminates on one of four ASICs on each SAN768B-2 core blade, or two ASICs on each SAN384B-2 core blade. Trunk groups can be formed by using any of the ports that make up contiguous groups of eight links on each ASIC.

Figure 10-20 shows that each core blade has groups of eight ICL ports (indicated by the blue box around the groups of ports) that connect to common ASICs in such a way that their four links can participate in common trunk groups with links from the other ports in the group.




Figure 10-20 Core blade trunk groups

Each SAN384B-2 core blade has one group of eight ICL ports, and each SAN768B-2 core blade has two groups of eight ICL ports.

Because there are four separate links for each QSFP-based ICL connection, each of these ICL port groups can create up to four trunks, with up to eight links in each trunk.

A trunk can never be formed by links within the same QSFP ICL port because each of the four links within the ICL port terminates on a different ASIC for the SAN768B-2 core blade, or on either different ASICs or different trunk groups within the same ASIC for the SAN384B-2 core blade. Thus, each of the four links from an individual ICL is always part of independent trunk groups.

When connecting ICLs between a SAN768B-2 and a SAN384B-2, the maximum number of links in a single trunk group is four. This limitation is because the different number of ASICs on each product's core blades, and the mapping of the ICL links to the ASIC trunk groups. To form trunks with up to eight links, ICL ports must be deployed within the trunk group boundaries that are shown in Figure 10-20. They can be created only when deploying ICLs between a pair of SAN768B-2 chassis or SAN384B-2 chassis. It is not possible to create trunks with more than four links when connecting ICLs between a SAN768B-2 and SAN384B-2 chassis.

As a preferred practice, deploy trunk groups in groups of up to four links by ensuring that the ICL ports that are intended to form trunks all are within the groups that are indicated by the red boxes in Figure 10-21.




Figure 10-21 Core blade suggested trunk groups

Following this preferred practice, trunks can be easily formed by using ICL ports, whether two SAN768B-2 chassis are being connected, two SAN384B-2 chassis, or a SAN768B-2 and a SAN384B-2.

Any time that more ICL connections are added to a chassis, they should be added in pairs by including at least one additional ICL on each core blade. It is also a preferred practice that trunks on a core blade are always composed of equal numbers of links, and that connections be deployed in an identical fashion on both core blades within a chassis. As an example, if two ICLs are deployed within the group of four ICL ports in trunk group A in Figure 10-21, a

single additional ICL can be added to trunk group A, or a pair of ICLs to any of the other trunk groups on the core blade. This configuration ensures that no trunks are formed that have a different total bandwidth from other trunks on the same blade. Deploying a single additional ICL to trunk group B might result in four trunks with 32 Gbps of capacity (those created from the ICLs in trunk group A) and four trunks with only 16 Gbps (those from the single ICL in group B).

The port mapping information that is shown in Figure 10-22 on page 282 and Figure 10-23 on page 283 also indicates the preferred ICL trunk groups by showing ports in the same preferred trunk group with the same color.

Core blade (CR16-8) port numbering layout

Figure 10-22 shows the layout of ports 0 - 15 on the SAN768B-2 CR16-8 line card. You can also see what the **switchshow** output would be if you ran a **switchshow** command within FOS by using the CLI.

External ICL Port #	Switchshow Port #	External ICL Port #	Switchshow ICL Port #
7	28-31	15	60-63
6	24-27	14	56-59
5	20-23	13	52-55
4	16-19	12	48-51
3	12-15	11	44-47
2	8-11	10	40-43
1	4-7	9	36-39
0	0-3	8	32-35

Figure 10-22 SAN768B-2 CR16-8 core blade - external ICL port numbering to “switchshow” (internal) port numbering

The colored groups of external ICL ports indicate those ports that belong to common preferred trunk groups. For example, ports 0 - 3 (shown in blue in Figure 10-22) form four trunk groups, with one link being added to each trunk group from each of the four external ICL ports. For the SAN768B-2, you can create up to 16 trunk groups on each of the two core blades.

The first ICL POD license enables ICL ports 0 - 7. Adding a second ICL POD license enables the remaining eight ICL ports, ports 8 - 15. This change applies to ports on both core blades.

Note: To disable ICL port 0, you must run **portdisable** on all four “internal” ports that are associated with that ICL port.

Core blade (CR16-4) port numbering layout

Figure 10-23 shows the layout of ports 0 - 7 on the SAN384B-2 CR16-4 line card. You can also see what the **switchshow** output would be if you ran **switchshow** within FOS by using the CLI.

External ICL Port #	Switchshow Port #
7	28-31
6	24-27
5	20-23
4	16-19
3	12-15
2	8-11
1	4-7
0	0-3

Figure 10-23 SAN384B-2 core blade - external ICL port numbering to “switchshow” (internal) port numbering

The colored groups of external ICL ports indicate those ports that belong to a common preferred trunk group. For example, ports 0 - 3 (shown in blue in Figure 10-23) form four trunk groups, with one link being added to each trunk group from each of the four external ICL ports. For the SAN384B-2, you can create up to eight trunk groups on each of the two core blades.

A single ICL POD license enables all eight ICL ports on the SAN384B-2 core blades. This change applies to ports on both core blades.

Note: To disable ICL port 0, you must run **portdisable** on all four “internal” ports that are associated with that ICL port.

10.12.4 ICL diagnostic tests

FOS V7.4.1 provides Diagnostic Port (D_Port) support for ICLs, helping administrators quickly identify and isolate ICL optics and cable problems. The D_Port on ICLs measures link distance and performs link traffic tests. It skips the electrical loopback and optical loopback tests because the QSFP does not support those functions. In addition, FOS V7.1 offers D_Port test CLI enhancements for increased flexibility and control.

10.12.5 Summary

The QSFP-based optical ICLs enable simpler, flatter, low-latency chassis topologies, spanning up to a 100-meter distance with standard cables. These ICLs reduce inter-switch cabling requirements and provide up to 33% more front-end ports for servers and storage, providing more usable ports in a smaller footprint with no loss in connectivity.

10.13 Fabric OS management

Fabric OS can be downloaded to a Backbone, which is a chassis, and to a non-chassis-based system, also referred to as a fixed-port switch. The difference in the download process is that Backbones have two Control Processors (CPs) and fixed-port switches have one CP. Both IBM Network Advisor and THE CLI can be used to upgrade or downgrade Fabric OS from either an FTP or SSH server by using FTP, SFTP, or SCP to the switch. Or, you can use a Brocade-branded USB device. IBM Network Advisor also offers a built-in FTP server.

Note: For detailed Firmware Management options including CLI-based instructions, see the *Fabric OS Administrators Guide* available at the following website:

<http://brocade.com>

Not every release of FOS is certified for use on IBM branded switches. It is important to ensure that the level is listed in the IBM Support Portal for the switch platform.

All upgrades require, at minimum, that the release notes for the Fabric OS level be reviewed. However, you might also need to the *Fabric OS Administrator Guide* and the *IBM Hardware Installation, Service and User Guide* for the platform that is being upgraded or rolled back depending on your familiarity with the hardware. All these items are available through the IBM support portal link.

All code is obtained by using the IBM support portal. However, this portal redirects to the Brocade IBM assist website to download the code.

The IBM support portal is available at the following link:

<https://www.ibm.com/support/entry/portal/support>

New firmware consists of multiple files in the form of RPM packages listed in a .plist file. The .plist file contains specific firmware information (time stamp, platform code, version, and so on) and the names of packages of the firmware to be downloaded. These packages are made available periodically to add features or to remedy defects. Contact your switch support provider to obtain information about available firmware versions.

All systems maintain two partitions (a primary and a secondary) of nonvolatile storage areas to store firmware images. The firmware download process always loads the new image into the secondary partition. It then swaps the secondary partition to be the primary and high availability (HA) reboots the system (which is nondisruptive). After the system boots up, the new firmware is activated. The firmware download process then copies the new image from the primary partition to the secondary partition.

In dual-CP systems, the firmware download process, by default, sequentially upgrades the firmware image on both CPs using HA failover to prevent disruption to traffic flowing through the Backbone. This operation depends on the HA status on the Backbone.

For an IBM System Storage SAN384B-2 and SAN768B-2 Backbone family product, with one or more AP blades, Fabric OS automatically detects mismatches between the active CP firmware and the blade's firmware, and triggers the auto leveling process. This auto-leveling process automatically updates the blade firmware to match the active CP. At the end of the auto-leveling process, the active CP and the blade run the same version of the firmware.

Since version 6.0.0, a one level nondisruptive upgrade is supported. Therefore, to go from 7.0.x to 7.4.x, the path would require multiple updates, if the intent is to do so non-disruptively. The suggested practice is to use the last release in each major release level. The example above would require an update path consisting of 7.1.x → 7.2.x → 7.3.x → 7.4.x where “x” is a sub release within the major release of that level.

10.13.1 Firmware download enhancements

FOS v7.4 introduces the following enhancements for firmware download:

- ▶ Staged Firmware Download

FOS v7.4 supports staged firmware download so that users can download firmware package to a switch first and choose to install and activate the downloaded firmware later.

- ▶ Firmware Clean Installation

FOS v7.4 supports firmware clean installation, which installs a firmware package without retaining the existing configuration or maintaining HA. With this feature, customers receiving a new switch from the factory can install firmware in a single step to the wanted version that their networks are running, without going through multiple steps of nondisruptive firmware download.

- ▶ Firmware Auto Sync Enhancement

FOS v7.4 enhances the firmware auto sync feature to support automatic synchronization of firmware versions on a standby CP with a version different from the active CP. If an active CP runs FOS v7.4 or higher, automatic upgrades and downgrades can be done in these circumstances:

- A standby CP with firmware version as early as FOS v6.4 can be upgraded automatically.
- A standby CP with firmware version later than FOS v7.4 can be downgraded automatically.

Note: For each switch in the fabric to be updated, complete all firmware download changes on the current switch before starting the upgrade or downgrade on the next switch. This process avoids disruption of traffic between switches in your fabric.

10.14 Upgrading firmware or rolling back to an earlier version

In most cases, firmware will be upgraded, which is installing a newer firmware version than the one currently running. However, some circumstances might require installing an older version, which is rolling back the firmware. The procedures in this section assume that firmware is being upgraded, but they also work for rolling back to an earlier version when the old and new firmware versions are compatible. Always reference the latest release notes for updates that might exist regarding rollbacks under particular circumstances.

Considerations for FICON Control Unit Port (CUP) environments: To prevent channel errors during nondisruptive firmware installation, the switch CUP port must be taken offline from all host systems.

10.14.1 Preparing for upgrades

Before you run a firmware upgrade, complete the tasks in this section. In the unlikely event of a failure or timeout, these preparatory tasks enable you to provide your switch support provider the information that is required to troubleshoot the firmware download.

1. Download the wanted Fabric OS package and the corresponding md5 sum file from the support portal at the following website:
<https://www.ibm.com/support/entry/portal/support>
2. Read the release notes for the new firmware to find out whether there are any updates related to the firmware download process.
3. Connect to the switch and log in using an account with admin permissions and verify the current version of Fabric OS and HA status. To confirm Fabric OS level and confirm HA status in IBM Network Advisor, complete the following steps:
 - a. In IBM Network Advisor, right-click the switch in the product or topology panes of the SAN tab and select **Element Manager** → **Hardware** to display the Hardware window in Web Tools (see Figure 10-24).

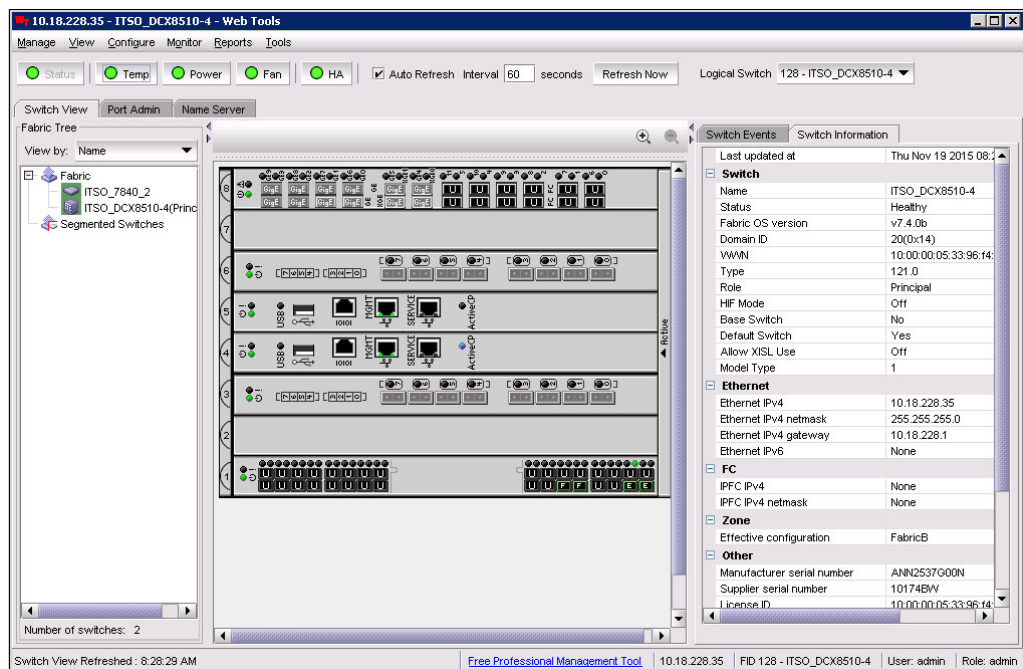


Figure 10-24 Web Tools Hardware window

- b. Select the **Switch information** tab in the right pane to display the switch hardware information, including the Fabric OS version.
- c. Click the **HA** button at the top of the panel to display the High Availability status window. If the status is not HA enabled, Heartbeat Up, HA State synchronized, you will need to troubleshoot the problem before the firmware upgrade can proceed.

Figure 10-25 shows the High Availability window.

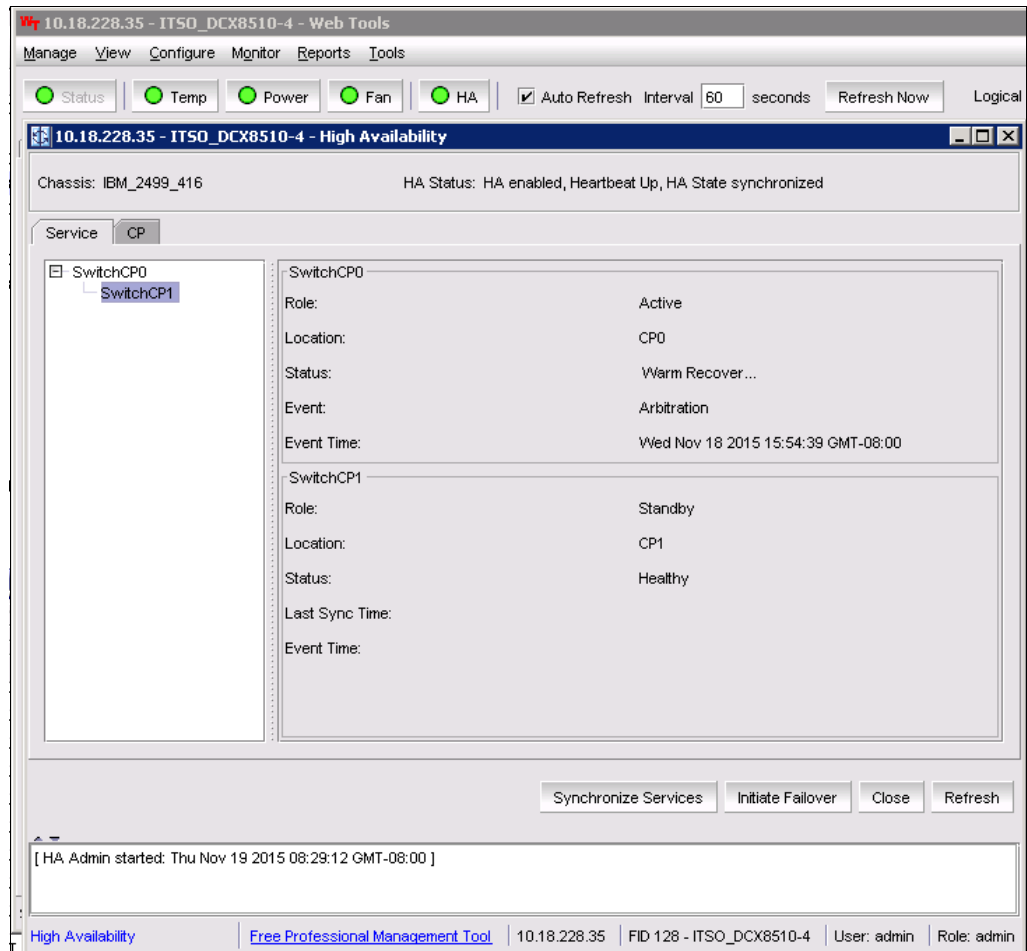


Figure 10-25 Web Tools High Availability window

4. Collect and save a switch configuration file on your FTP or SSH server, or USB memory device on supported platforms.

For information about collecting the switch configuration, see 4.12, “Scheduling daily or weekly backups for the fabric configuration” on page 101.

5. Connect to the switch and log in using an account with admin permissions. Collect and save a current supportsave before you run the Fabric OS upgrade. This information helps to troubleshoot the firmware download process if a problem is encountered.

For instructions for collecting a supportsave, see 12.7, “Collecting support data” on page 314.

10.14.2 Staging the Fabric OS package for download to the switch

After the Fabric OS version package has been downloaded locally to the Host running the IBM Network Advisor server, it must be added to the Firmware repository.

The firmware repository is used by the internal FTP, SCP, or SFTP server that is delivered with the Management application software. It can be used by an external FTP server if it is installed on the same platform as the Management application software. The repository is not available to FTP servers on external platforms.

To add a Fabric OS package to the firmware repository, complete the following steps:

1. Log in to IBM Network Advisor using an ID with admin privileges.
2. Select **Configure** → **Firmware Management** and the Firmware Management window will be displayed. Figure 10-26 shows the Firmware Management window with the Repository tab selected.

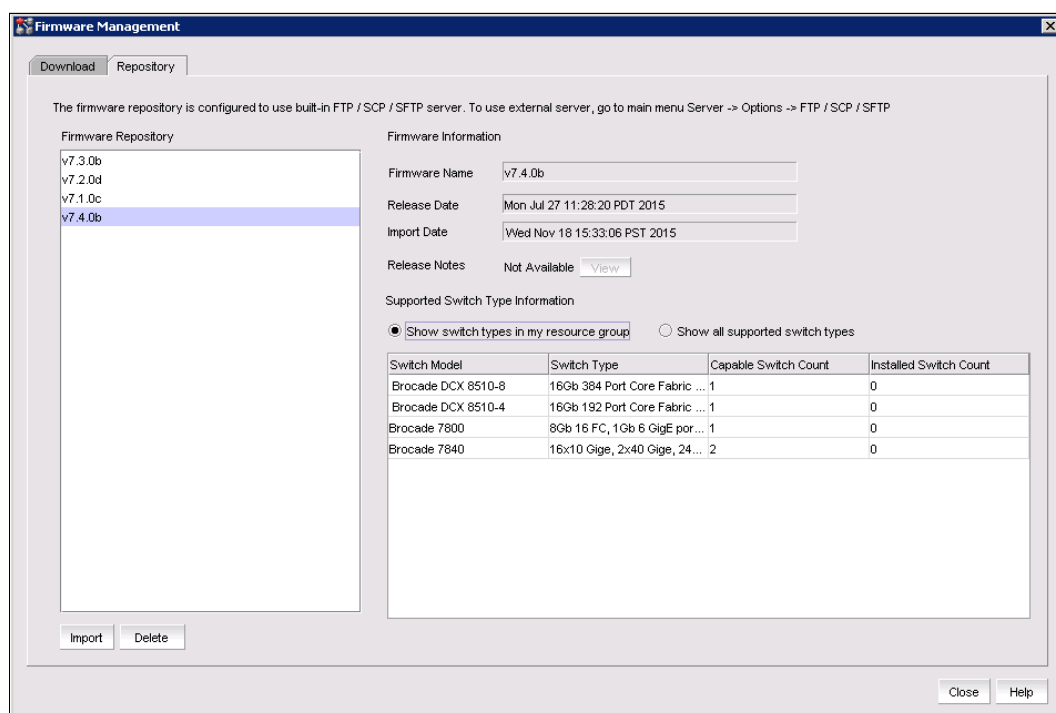


Figure 10-26 Firmware Management panel with the Repository tab selected

3. Select the Repository tab.

The right pane displays the current Fabric OS packages that are available in the repository. The left pane displays supported switches.

Note: In the Supported Switches pane, the **show supported switches in my resource group** or **show supported switches** radio buttons can be selected. The **show supported switches in my resource group** button shows only switches in the inventory that is currently discovered that support the fabric OS package that is selected in the Firmware Repository pane. The **show supported switches** option will show all supported products that are available from IBM.

4. To import a downloaded Fabric OS package into the repository, click **Import** to display the Import Firmware File window.

Packages can also be deleted by selecting a package in the repository pane and clicking **Delete** (see Figure 10-27).

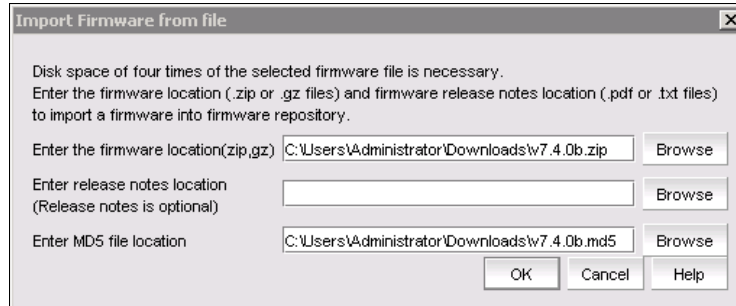


Figure 10-27 Import Firmware from File panel

5. Click **Browse** to browse to the downloaded Fabric OS package location on the local server and repeat the step for the md5 sum file.
6. Click **OK** to add the package to the repository. IBM Network Advisor confirms the file's integrity by using the md5 file, and then extract and stage the file in the repository.
An Import completed successfully window will be displayed when complete.
7. Click **OK** to close the notification and the Firmware Repository window is displayed.
8. Click **Close** to close the Firmware Management window.

10.14.3 Upgrading firmware

To complete a Fabric OS upgrade with IBM Network Advisor, complete the following steps:

1. Log in to IBM Network Advisor using an ID with Admin privileges.
2. Select **Configure** → **Firmware Management** to display the Firmware Management window.

3. Select the Download tab.

The left pane displays switches that are available for firmware update, the right pane displays selected switches, and the lower right pane displays firmware update status details when a firmware update is underway (see Figure 10-28).

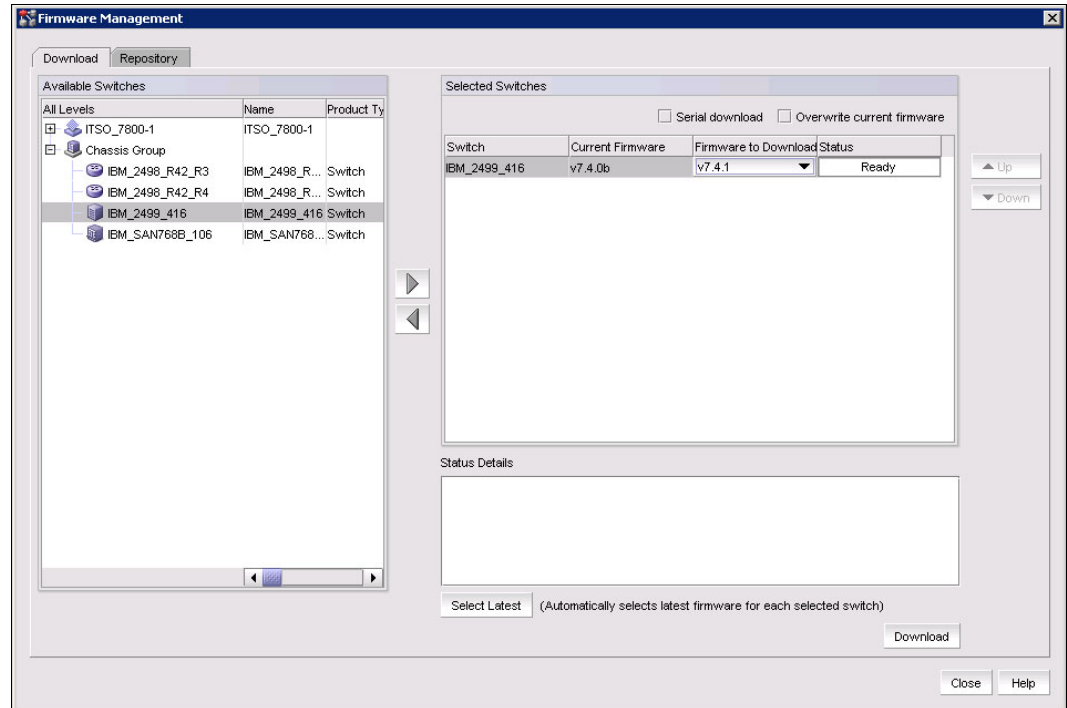


Figure 10-28 Firmware Download window with the Download tab selected

4. Select one or more switches from the Available Switches pane.
5. Click the right arrow to move the switches to the **Selected Switches** pane.
6. Select a specific version from the Firmware to Download column, or use **Select Latest** to automatically select the latest version.
7. To download the firmware to the selected switches one at a time, select **Serial download**.
 - Use the **Up** and **Down** buttons to determine the order in which the firmware is downloaded to the switches. If firmware download fails on one switch, all other switches in the queue will be skipped.
 - If the **Serial download** check box is cleared, the download occurs in parallel on the switches (up to 20 at a time).
8. To overwrite the current firmware, even if the selected version is the same as the version currently running on the switch, select **Overwrite current firmware**.
 - While the firmware is downloaded to the device, the Status column displays the current download status. After the firmware download is complete, the Message column displays whether the download was a success or failure.



Security

Many components in SAN security relate to SAN design, and the decision to use these components depends on installation requirements rather than network functioning or performance. One clear exception is the zoning feature that is used to control device communication. The proper use of zoning is key to fabric functioning, performance, and stability, especially in larger networks. For more information, see 10.5, “Zoning” on page 260. Other security-related features are largely mechanisms for limiting access and preventing attacks on the network (and are mandated by regulatory requirements). They are not required for normal fabric operation.

This chapter includes the following sections:

- ▶ Role-Based access controls
- ▶ Default accounts
- ▶ User accounts
- ▶ Security protocols
- ▶ Access control lists
- ▶ Policy Database Distribution
- ▶ In-flight encryption and compression: b-type (16 Gbps) platforms only
- ▶ In-flight encryption and compression guidelines

11.1 Role-Based access controls

One way to provide limited accessibility to the fabric is through user roles. FOS has predefined user roles, each of which has access to a subset of the CLI commands. These roles are known as role-based access controls (RBACs), and they are associated with the user login credentials. When you log in to a switch, your user account is associated with a predefined role or a user-defined role. The role that your account is associated with determines the level of access you have on that switch and in the fabric. For more information about RBAC, see the “Managing User Accounts” chapter in the *Fabric OS Administrator's Guide*, which you can find at the following website:

<http://my.brocade.com/>

11.2 Default accounts

FOS offers four predefined accounts: Admin, factory, root, and user. Although the root and factory accounts are reserved for development and manufacturing, the password for all default accounts should be changed and secured during the initial installation and configuration of each switch. Recovering passwords might require significant effort and fabric downtime.

11.3 User accounts

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

Each user account is associated with the following things:

- ▶ Admin Domain list: Specifies the Administrative Domains to which a user account is allowed to log in.
- ▶ Home Admin Domain: Specifies the Admin Domain that the user is logged in to by default. The home Admin Domain must be a member of the user's Admin Domain list.
- ▶ Permissions: Associate roles with each user account to determine the functional access levels within the bounds of the user's current Admin Domain.
- ▶ Virtual Fabric list: Specifies the Virtual Fabric that a user account is allowed to log in to.
- ▶ Home Virtual Fabric: Specifies the Virtual Fabric that the user is logged in to, if available. The home Virtual Fabric must be a member of the user's Virtual Fabric list. If the fabric ID is not available, the next-lower valid fabric ID is used.
- ▶ LF Permission List: Determines functional access levels within the bounds of the user's Virtual Fabrics.
- ▶ Chassis role. Similar to switch-level roles, but applies to a different subset of commands.

11.4 Security protocols

Security protocols provide endpoint authentication and communications privacy by using cryptography. Typically, the user is authenticated to the switch while the switch remains unauthenticated. This authentication means that you can be sure that you know 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.

Table 11-1 shows the security protocols that are supported by Fabric OS.

Table 11-1 Security protocols that are supported by Fabric OS

Protocol	Description
CHAP	Challenge Handshake Authentication Protocol (CHAP) uses shared secrets to authenticate switches.
HTTPS	HTTPS is a Uniform Resource Identifier scheme that is used to indicate a secure HTTP connection. Web Tools supports the use of Hypertext Transfer Protocol over SSL (HTTPS).
IPSec	Internet Protocol Security (IPSec) is a framework of open standards for providing confidentiality, authentication, and integrity for IP data transmitted over untrusted links or networks.
LDAP	Lightweight Directory Access Protocol (LDAP) with Transport Layer Security (TLS) uses a certificate authority (CA). By default, LDAP traffic is transmitted unsecured. With the import of signed certificates, you can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL) / TLS technology with LDAP.
SCP	Secure Copy Protocol (SCP) is a means of securely transferring computer files between a local and a remote host or between two remote hosts that use the Secure Shell (SSH) protocol. Configuration upload and download support the use of SCP.
Secure Syslog	Secure syslog requires importing syslog CA certificates by using the secCerUtil command.
SFTP	Secure File Transfer Protocol (SFTP) is a network protocol for securely transferring files on a network. Configuration upload and download support the use of SFTP.
SNMP	Simple Network Management Protocol (SNMP) is used in network management systems to monitor network-attached devices for conditions that warrant administrative attention. Supports SNMPv1 and v3.
SSH	SSH is a network protocol that allows data to be exchanged over a secure channel between two computers. Encryption provides confidentiality and integrity of data. SSH uses public-key cryptography to authenticate the remote computer and allow the remote computer to authenticate the user, if necessary.
SSL	Fabric OS uses SSL to support HTTPS. A certificate must be generated and installed on each switch to enable SSL. This configuration supports SSLv3, 128-bit encryption by default. It also supports TLSv1.0, TLSv1.1, and TLSv1.2.

Note: Some of the security protocols that are listed in Table 11-1 require additional software or certificates that you must obtain to deploy secure protocols. For more information and configuration instructions, see the *Fabric OS Administrator's Guide* at the following website:

<http://my.brocade.com>

11.5 Access control lists

Access control lists (ACLs) are used to provide network security through policy sets. FOS provides several ACL policies, including a Switch Connection Control (SCC) policy, a Device Connection Control (DCC) policy, a Fabric Configuration Server (FCS) policy, an IP Filter, and others. The following sections briefly describe each policy and provide basic guidelines. A more in-depth description of ACLs can be found in the *Fabric OS Administrator's Guide*, which you can find at the following website:

<http://my.brocade.com/>

11.5.1 SCC policy

The SCC policy restricts the fabric elements (FC switches) that can join the fabric. Only switches that are specified in the policy are allowed to join the fabric. All other switches fail authentication if they attempt to connect to the fabric, resulting in the respective E_Ports being segmented because of the security violation.

Use the SCC policy in environments where you need strict control of fabric members. Because the SCC policy can prevent switches from participating in a fabric, it is important to regularly review and properly maintain the SCC ACL.

11.5.2 DCC policy

The DCC policy restricts the devices that can attach to a single Fibre Channel (FC) port. The policy specifies the FC port and one or more worldwide names (WWNs) that are allowed to connect to the port. The DCC policy set comprises all of the DCC policies that are defined for individual FC ports. Not every FC port must have a DCC policy, and only ports with a DCC policy in the active policy set enforce access controls. A port that is present in the active DCC policy set allows only WWNs in its respective DCC policy to connect and join the fabric. All other devices fail authentication when they attempt to connect to the fabric, resulting in the respective F_Ports being disabled because of the security violation.

Use the DCC policy in environments where you need strict control of fabric members. Because the DCC policy can prevent devices from participating in a fabric, it is important to regularly review and properly maintain the DCC policy set.

11.5.3 FCS policy

Use the FCS policy to restrict the source of fabric-wide settings to one FC switch. The policy contains the WWN of one or more switches, and the first WWN (that is online) in the list is the primary FCS. If the FCS policy is active, then only the primary FCS is allowed to make or propagate fabric-wide parameters. These parameters include zoning, security (ACL) policies databases, and other settings.

Use the FCS policy in environments where you need strict control of fabric settings. As with other ACL policies, it is important to regularly review and properly maintain the FCS policy.

11.5.4 IP Filter

The IP Filter policy is a set of rules that are 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.

As a preferred practice, non-secure IP protocols that are used for switch management such as Telnet and HTTP should be blocked. SSH is enabled on the default IP Filter policy and SSL should be configured to use HTTPS for web access.

The IP Filter policy should be used in environments where you need strict control of fabric access. As with other ACL policies, it is important to regularly review and properly maintain the IP Filter policy.

11.5.5 Authentication protocols

Fabric Operating System (FOS) supports both Fibre Channel Authentication Protocols (FCAPs) and Diffie-Hellman CHAPs (DH-CHAPs) on E_Ports and F_Ports. Authentication protocols provide extra security during link initialization by ensuring that only the wanted device/device type is connecting to a given port.

11.6 Policy Database Distribution

Security Policy Database Distribution provides a mechanism for controlling the distribution of each policy on a per-switch basis. Switches can individually configure policies to either accept or reject a policy distribution from another switch in the fabric. In addition, a fabric-wide distribution policy can be defined for the SCC and DCC policies with support for strict, tolerant, and absent modes. This setting can be used to enforce whether the SCC or DCC policy must be consistent throughout the fabric.

The Policy Database Distribution has three modes:

- ▶ **Strict mode:** All updated and new policies of the type specified (SCC, DCC, or both) must be distributed to all switches in the fabric, and all switches must accept the policy distribution.
- ▶ **Tolerant mode:** All updated and new policies of the type specified (SCC, DCC, or both) are distributed to all switches (FOS V6.2.0 or later) in the fabric. However, the policy does not need to be accepted.
- ▶ **Absent mode:** Updated and new policies of the type specified (SCC, DCC, or both) are not automatically distributed to other switches in the fabric. Policies can still be manually distributed.

Together, the policy distribution and fabric-wide consistency settings provide a range of control on the security policies from little or no control to strict control.

11.7 In-flight encryption and compression: b-type (16 Gbps) platforms only

IBM b-type Fibre Channel (16 Gbps) platforms support both in-flight compression and encryption at a port level for both local and long-distance inter-switch links (ISLs). In-flight data compression is a useful tool for saving money when either bandwidth caps or bandwidth usage charges are in place for transferring data between fabrics. Similarly, in-flight encryption enables a further layer of security with no key management impact when transferring data between local and long-distance data centers besides the initial setup.

Enabling in-flight ISL data compression or encryption increases the latency as the ASIC processes the frame compression or encryption. The approximate latency at each stage (encryption and compression) is 6.2 microseconds. For example (see Figure 11-1), compressing and then encrypting a 2 KB frame incurs approximately 6.2 microseconds of latency on the sending Condor3-based switch, and approximately 6.2 microseconds of latency at the receiving Condor3-based switch to decrypt and decompress the frame. This combination results in a total latency time of 12.4 microseconds, again not counting the link transit time.

Figure 11-1 shows the total accumulated latency when encryption and compression are in use.

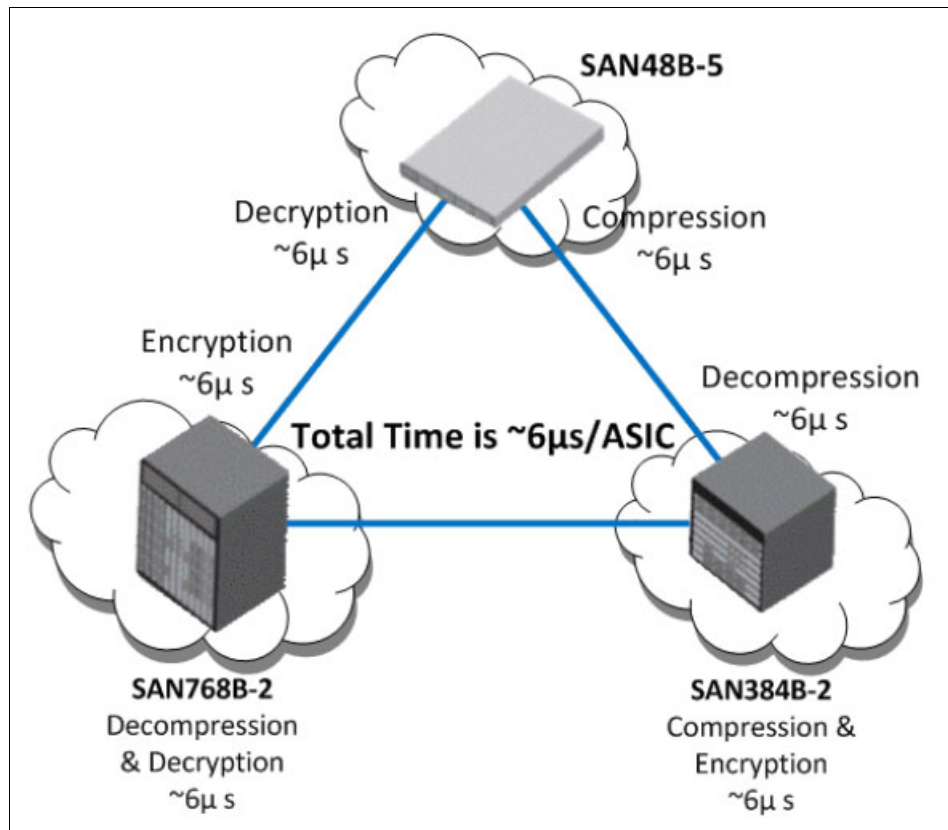


Figure 11-1 Latency for encryption and compression

Virtual Fabric considerations (encryption and compression): The E_Ports in the user-created Logical Switch, Base Switch, or default switch can support encryption and compression. Both encryption and compression are supported on XISL ports, but not on LISL ports. If encryption or compression is enabled and ports are being moved from one LS to another, it must be disabled before the ports are moved.

11.8 In-flight encryption and compression guidelines

Consider these in-flight encryption and compression guidelines:

- ▶ Encryption and compression are supported on E_Ports and EX_Ports.
- ▶ ISL ports must be set to Long-Distance (LD) mode when compression is used.
- ▶ Twice the number of buffers should be allocated if compression is enabled for long distance because frame sizes might be half the size.
- ▶ If both compression and encryption are used, enable compression first.

When you implement ISL encryption, using multiple ISLs between the same switch pair requires that all ISLs be configured for encryption, or none at all.

No more than two ports on one ASIC can be configured with encryption, compression, or both when running at 16 Gbps speed. With FOS V7.1, additional ports can be used for data encryption, data compression, or both if the system is running at lower than 16 Gbps speeds.

Encryption is not compliant with Federal Information Processing Standards (FIPS).

Note: For more information about securing a SAN network and associated management and IP WAN connectivity, see the *Fabric OS Administration Guide* that is available at the following website:

<http://my.brocade.com>



Troubleshooting

This chapter describes the steps that you can take to ascertain the health of the storage area network (SAN) fabric and perform problem determination when unplanned events occur. Although it provides many troubleshooting tips and techniques, it does not teach troubleshooting methodology.

This chapter provides the following information:

- ▶ General problem determination
- ▶ Errors and symptoms
- ▶ Switch and port status
- ▶ Port errors
- ▶ System messages and RAS logs
- ▶ SAN health
- ▶ Collecting support data
- ▶ Using MAPS for problem determination
- ▶ Flow Vision

12.1 General problem determination

Troubleshooting should begin at the center of the SAN, which is the fabric. Because switches are located between the hosts and storage devices and have visibility into both sides of the storage network, starting with the fabric or fabrics can help narrow the search path. After eliminating the possibility of a fault within the fabric, determine whether the problem is on the storage side or the host side, and continue a more detailed diagnosis from there. Using this approach can quickly pinpoint and isolate problems.

For example, if a host cannot detect a storage device, determine whether the storage device is logically connected to the switch by using IBM Network Advisor to view the physical port status at the switch and the name server to determine whether the device is logically logged in to the fabric. If not, focus first on the switch directly connecting to storage. Use your vendor-supplied storage diagnostic tools to better understand why it is not visible to the switch. If the storage can be detected by the switch, and the host still cannot detect the storage device, then the problem is between the host and the switch.

12.2 Errors and symptoms

For problem determination of common problems and symptoms, the *Fabric OS Troubleshooting and Diagnostics Guide* provides information about identifying and working through them. It includes information about the following topics:

- ▶ LED statuses
- ▶ Connectivity
- ▶ Performance
- ▶ Fabric Merge and ISL
- ▶ FCIP and FCR problem determination
- ▶ Hardware problem determination
- ▶ Firmware upgrade problems

Each of the sections provides information about how to review and determine corrective action for some of the problems that might be encountered in the fabric.

The *Fabric OS Troubleshooting and Diagnostics Guide* is available for each major release of Fabric OS and can be obtained at the following website:

<https://my.brocade.com>

12.3 Switch and port status










It is important to understand the overall switch status for all switches in the fabric when you investigate problems within the storage area network. After you determine that the switches in the fabric are healthy and operational, you can do further investigation of port status to ensure that all relevant ports are healthy.

12.3.1 Displaying the switch status

Switch status is available in several locations in IBM Network Advisor. However, the product list provides a quick overview of all switches in a fabric and their present state by displaying icons.

Table 12-1 shows the status icons and their meanings.

Table 12-1 Status icons that are displayed in the product list of IBM Network Advisor

Icon	Status
No Icon	Healthy and Operational
	Attention
	Degraded or Marginal
	Device Added
	Device Removed or Missing
	Down or Failed
	Routed In
	Routed Out
	Unknown or Link Down
	Unreachable

To display the product list, complete the following steps:

1. Log in to IBM Network Advisor with an ID with administrator privileges.
2. Click the SAN tab.

The product list is displayed in the right pane (see Figure 12-1).

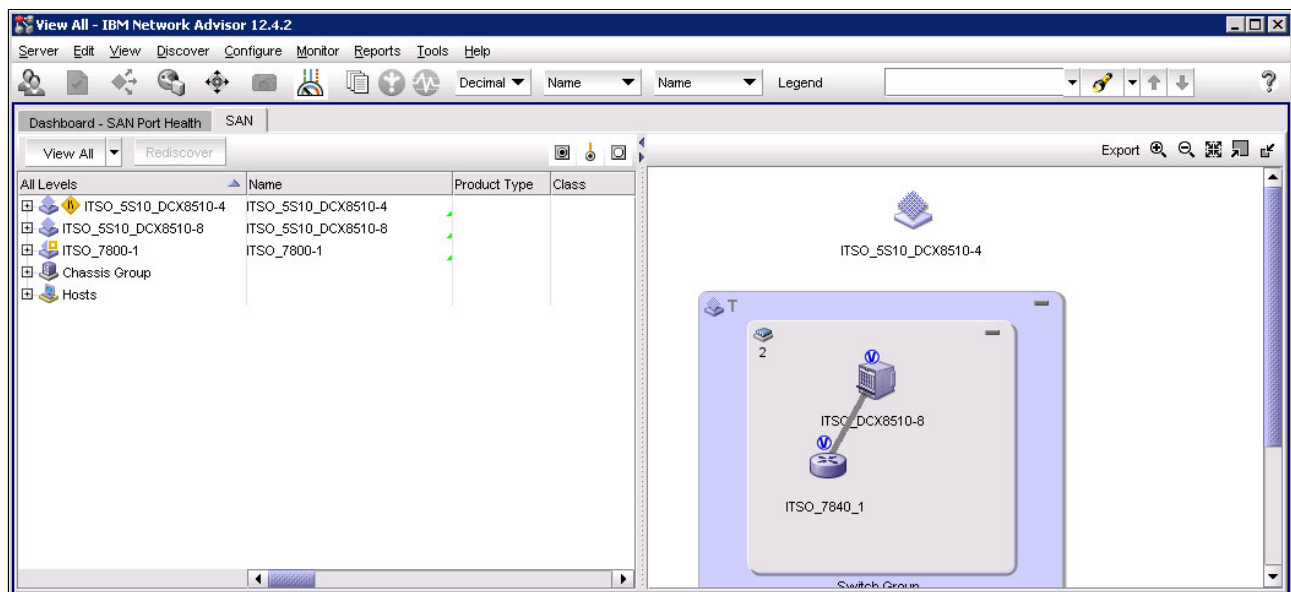


Figure 12-1 IBM Network Advisor Product list

- To display more details about a specific switch, right-click the switch in the right topology pane and select **Properties** from the menu.

The switch Properties window is displayed as shown in Figure 12-2.

Note: Use the **switchStatusShow** command in the CLI to display the overall status of the switch, including its power supplies, fans, and temperature. If the status of any one of these components is either marginal or down, the overall status of the switch is also displayed as marginal or down. If all components have a healthy status, the switch displays a healthy status. For more information about the **switchStatusShow** command, see the *Fabric OS Troubleshooting and Diagnostics Guide*, *Fabric OS Administrator's Guide* and *Fabric OS Command Reference* that are available at the following website:

<https://my.brocade.com>

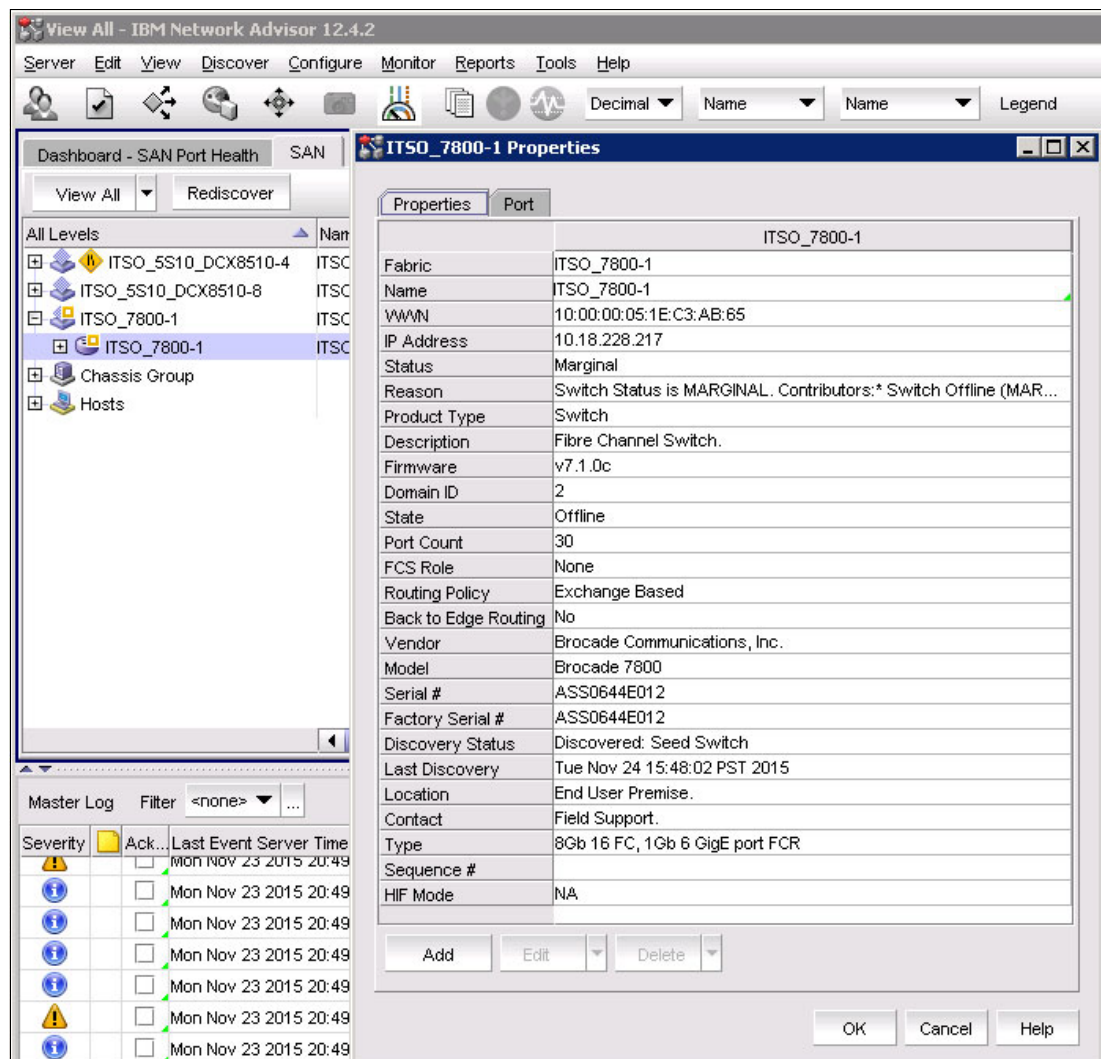


Figure 12-2 Switch Properties window

12.3.2 Port status















This section explains how to view port status.

Viewing the status of a port in IBM Network Advisor

Port status can also quickly be determined by expanding the product list. Each port has a corresponding icon to indicate its status. After a port has been identified as needing further investigation, you can obtain more information by using the displayed list.

The icons in Table 12-2 are displayed to indicate port status.

Table 12-2 Port status icons

Icon	Status
	Occupied FC Port
	Unoccupied FC Port
	Attached FC Port
	Trunk (port group)
	IP and 10 GE Port
	Attached IP and 10 GE Port
	Attached-to-Cloud 10 GE Port
	Virtual Port
	Virtual FCoE Port
	Attached FCoE Port
	Pre-boot Virtual Port
	Virtual Attached Port
	Mirror Port
	Bottleneck Port

To display the port statuses in the Product List in IBM Network Advisor, complete the following steps:

1. Open IBM Network Advisor and use an ID with admin privileges.
2. Click the SAN tab.
The product list is displayed in the left pane.
3. Click the + icon to expand the list of the fabric that contains the switch or switches.
4. Click the + icon next to the switch group that contains the switches for which the port statuses are to be displayed.
5. Click the + icon next to the switch or switches for which the port statuses are to be determined. A list of ports is displayed with the statuses.

Figure 12-3 shows the Product list expanded to view port statuses in IBM Network Advisor.

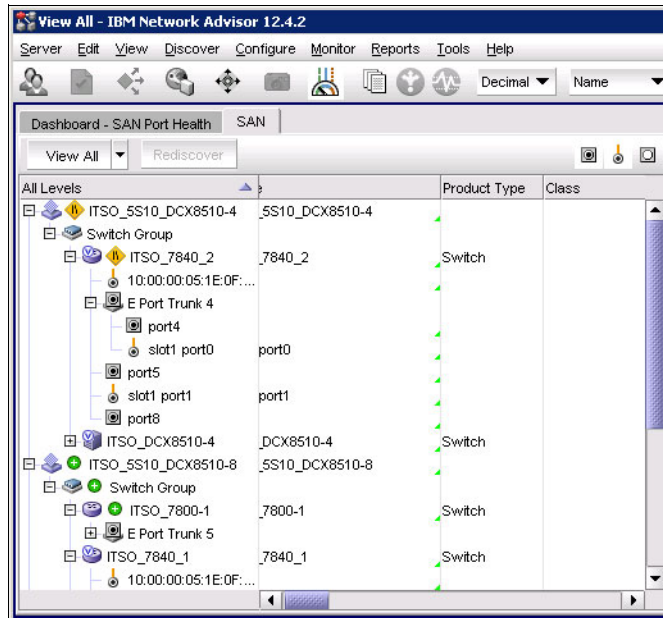


Figure 12-3 Port status

- For more information about a specific port in the list, right-click the port and select **Properties** from the menu. The port Properties window is displayed (Figure 12-4).

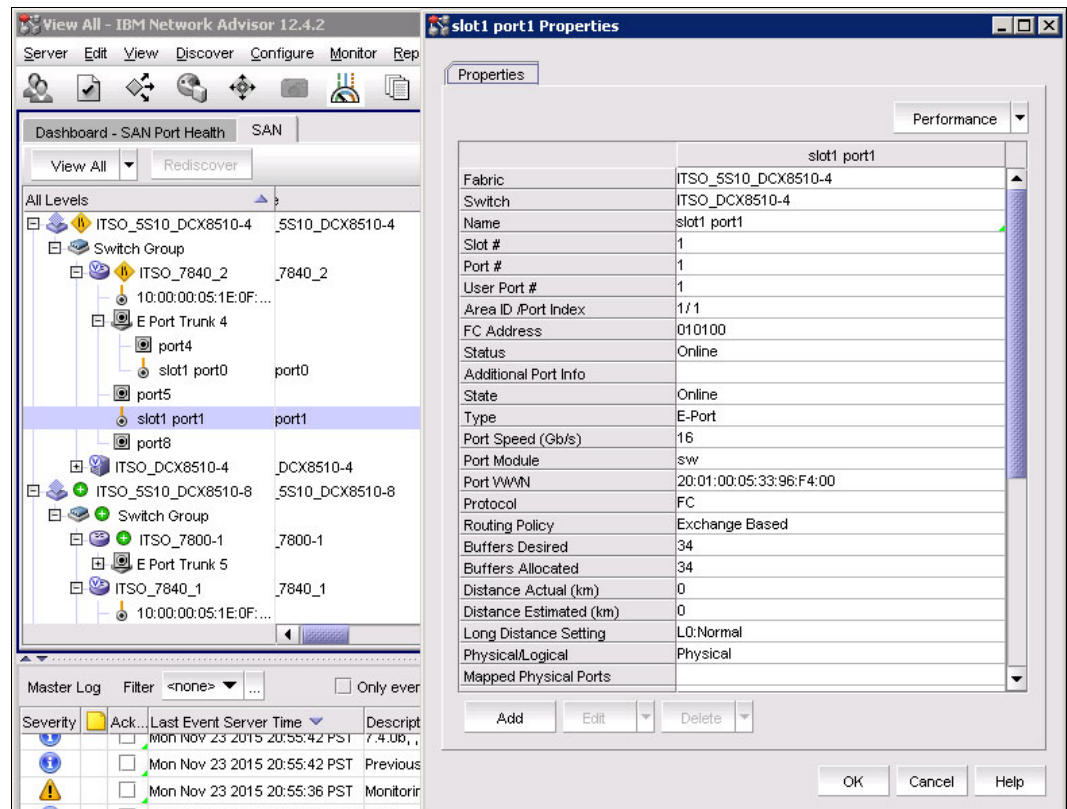


Figure 12-4 Port Properties panel

Viewing the status of a port in the CLI

To view the status of a port, run the **portShow [slot/] port** command with admin permissions and specify the number that corresponds to the port you are troubleshooting. Example 12-1 shows the output of an example **portShow** command.

Example 12-1 Port status command in the CLI

```
ITS0_DCX8510-4:FID128:admin> portshow 1/4
portIndex: 4
portName: slot1 port4
portHealth: OFFLINE

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x1 PRESENT U_PORT
LocalSwcFlags: 0x0
portType: 24.0
portState: 2 Offline
Protocol: FC
portPhys: 4 No_Light portScn: 2 Offline
port generation number: 8
state transition count: 1

portId: 010400
portIfId: 4312001b
portWwn: 20:04:00:05:33:96:f4:00
portWwn of device(s) connected:

Distance: normal
portSpeed: N16Gbps

FEC: Inactive
Credit Recovery: Inactive
LE domain: 0
Peer beacon: Off
FC Fastwrite: OFF
Interrupts: 0 Link_failure: 0 Frjt: 0
Unknown: 0 Loss_of_sync: 0 Fbsy: 0
Lli: 5 Loss_of_sig: 1
Proc_rqrd: 33 Protocol_err: 0
Timed_out: 0 Invalid_word: 371778
Rx_flushed: 0 Invalid_crc: 0
Tx_unavail: 0 Delim_err: 0
Free_buffer: 0 Address_err: 0
Overrun: 0 Lr_in: 0
Suspended: 0 Lr_out: 0
Parity_err: 0 Ols_in: 0
2_parity_err: 0 Ols_out: 0
CMI_bus_err: 0

ITS0_DCX8510-4:FID128:admin>
```

12.4 Port errors

After you understand the status of the port, the port error statistics also provide information about any physical or logical errors that are occurring on the link or links that are involved.

Port error statistics can be obtained in IBM Network Advisor or the command-line interface (CLI).

12.4.1 Viewing port statistics with IBM Network Advisor

Port error statistics are viewed with the Element manager in IBM Network Advisor. To display them, complete the following steps:

1. Log in to IBM Network Advisor with an ID that has admin privileges.
2. Click the SAN tab.
The product list and topology panes are displayed.
3. Right-click the switch for which port statistics are required in either the product list pane or the Topology pane, and select **Element Manager** → **Ports**. The port admin tab from Web Tools is displayed.
4. Select **View** → **Advanced** to change the panel options to advance mode and enable the port statistics tab.
5. Select the FC Ports tab to display the list of FC ports on that switch.
6. Select the port from the FC Ports Explorer pane and the general port properties tab is displayed in the right pane.
7. Select the Port Statistics tab and the port statistics are displayed for the selected port in the lower Port Statistics pane. Selecting either the Advanced tab or the Error Details tab in the Port Statistics pane provides error statistics on the ports (see Figure 12-5).

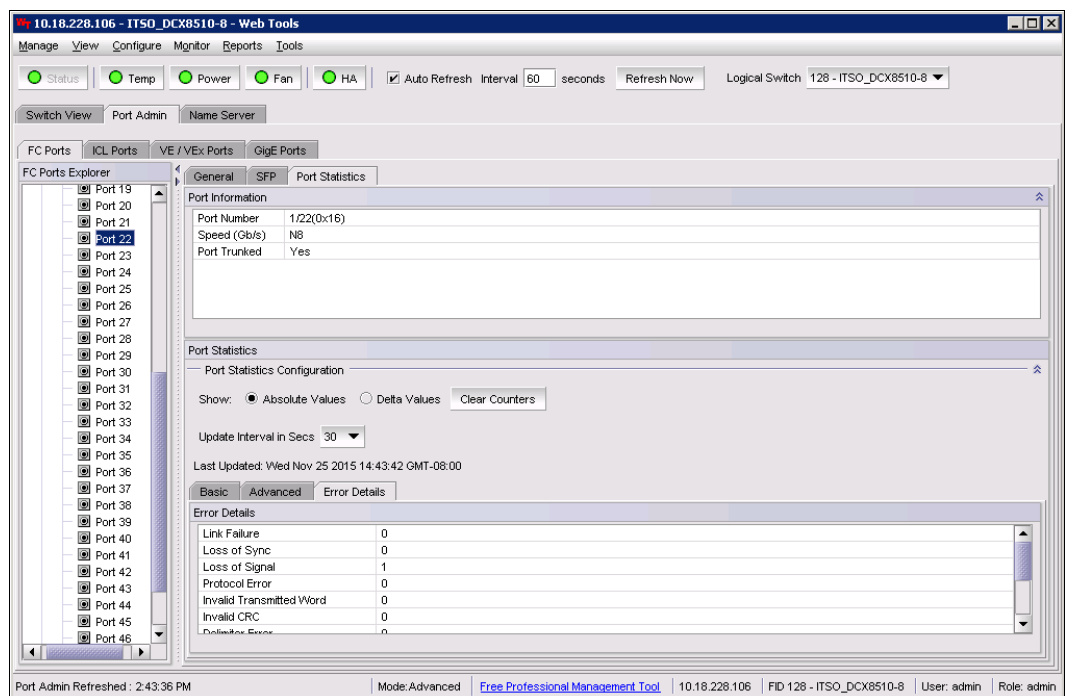


Figure 12-5 Port error statistics

12.4.2 Viewing the port statistics in the CLI

The **portShow [slot/] port**, as shown in Example 12-1 on page 305, displays port statistics in the lower half of the output.

The **portstatsshow [slot/]port** command provides more detailed output on the port statistics as shown in Example 12-2.

Example 12-2 Output of the portstatsshow [slot/]port command

```
ITS0_DCX8510-4:FID128:admin> portstatsshow 1/4
stat_wtx          2761009470  4-byte words transmitted
stat_wrx          1219327107  4-byte words received
stat_ftx          32748199   Frames transmitted
stat_frx          29806760   Frames received
stat_c2_frx       0          Class 2 frames received
stat_c3_frx       29806760   Class 3 frames received
stat_lc_rx        0          Link control frames received
stat_mc_rx        0          Multicast frames received
stat_mc_to        0          Multicast timeouts
stat_mc_tx        0          Multicast frames transmitted
tim_rdy_pri       0          Time R_RDY high priority
tim_txcrd_z       0          Time TX Credit Zero (2.5Us ticks)
tim_txcrd_z_vc 0- 3: 0      0      0      0
tim_txcrd_z_vc 4- 7: 0      0      0      0
tim_txcrd_z_vc 8-11: 0      0      0      0
tim_txcrd_z_vc 12-15: 0     0      0      0
tim_latency_vc 0- 3: 1      1      1      1
tim_latency_vc 4- 7: 1      1      1      1
tim_latency_vc 8-11: 1      1      1      1
tim_latency_vc 12-15: 1     1      1      1
fec_cor_detected  0          Count of blocks that were corrected by FEC
fec_uncor_detected 0          Count of blocks that were left uncorrected by
FEC
er_enc_in         0          Encoding errors inside of frames
er_crc            0          Frames with CRC errors
er_trunc          0          Frames shorter than minimum
er_toolong        0          Frames longer than maximum
er_bad_eof        0          Frames with bad end-of-frame
er_enc_out        371778     Encoding error outside of frames
er_bad_os         343338     Invalid ordered set
er_pcs_blk        0          PCS block errors
er_rx_c3_timeout  0          Class 3 receive frames discarded due to
timeout
er_tx_c3_timeout  0          Class 3 transmit frames discarded due to
timeout
er_unroutable     0          Frames that are unroutable
er_unreachable    0          Frames with unreachable destination
er_other_discard  0          Other discards
er_type1_miss     0          frames with FTB type 1 miss
er_type2_miss     0          frames with FTB type 2 miss
er_type6_miss     0          frames with FTB type 6 miss
er_zone_miss      0          frames with hard zoning miss
er_lun_zone_miss  0          frames with LUN zoning miss
er_crc_good_eof   0          Crc error with good eof
er_inv_arb        0          Invalid ARB
```

```

er_single_credit_loss    0          Single vcrdy/frame loss on link
er_multi_credit_loss     0          Multiple vcrdy/frame loss on link
phy_stats_clear_ts       0          Timestamp of phy_port stats clear
lgc_stats_clear_ts       0          Timestamp of lgc_port stats clear
ITS0_DCX8510-4:FID128:admin>

```

The **porterrshow** CLI command provides an overview of port statistics for the entire switch.

12.4.3 Resetting the port error statistic counters

When errors are reported on a link, it is important to understand that the error count is cumulative. There is no indication in these counters as to when the errors occurred.

In order to prove that the errors are relevant and still occurring, clear the error statistic counters and recheck them after some time passes. That is, a single snapshot does not provide enough information to determine whether the error is still occurring. For example, check at 5- and 60-minute intervals.

To reset all port statistics counters to zero on a selected device or fabric, complete the following steps:

1. Right-click a device or a fabric on the Connectivity Map or Product List and select **Monitor** → **Performance** → **Clear Counters**. An attention message is displayed.
2. Click **Yes** on the message.

All the port statistics counters and port logs will be cleared on all reachable switches in that switch group. The audit events log generated by the switches is displayed in the Master Log.

12.4.4 Understanding error counters

Based on the errors that occur, some generalizations can be made for further investigation. Table 12-3 lists port error counters, what the error is counting, and suggestions for further investigation.

Table 12-3 Error summary description

Error type	Description
frames tx	Frames transmitted.
frames rx	Frames received.
enc in	Encoding errors inside frames.
crc err	Frames with CRC errors.
crc g_eof	CRC errors that occur on frames with good end-of-frame delimiters.
too shrt	Frames shorter than minimum.
too long	Frames longer than maximum.
bad eof	Frames with bad end-of-frame delimiters.
enc out	Encoding error outside of frames.

Error type	Description
disc c3	Number of Class 3 frames discarded (Rx). This counter includes the sum of the following Class 3 discard counters that are reported by the portStatsShow command: er_rx_c3_timeout, er_tx_c2_timeout, er_c2_dest_unreach, and er_other_disc. For a description of these counters, run portStatsShow help .
link fail	Link failures (LF1 or LF2 states).
loss sync	Loss of synchronization.
loss sig	Loss of signal.
frjt	Frames rejected with F_RJT.
fbsy	Frames busied with F_BSY.

Here are tips for further investigation of errors:

- ▶ crc_err and enc_out errors together imply a small form-factor pluggable (SFP) issue.
- ▶ enc_out errors on their own imply a cable/connector issue. This error can cause a performance problem because of buffer recovery.
- ▶ too_long or too_short errors indicate an unreliable link.
- ▶ disc_c3 relates to port congestion.
- ▶ loss_sig can indicate incompatible speeds between two points. These error messages can also be caused by severe physical layer errors or by devices that are being reset, such as server reboots.

12.4.5 SFP and optic levels

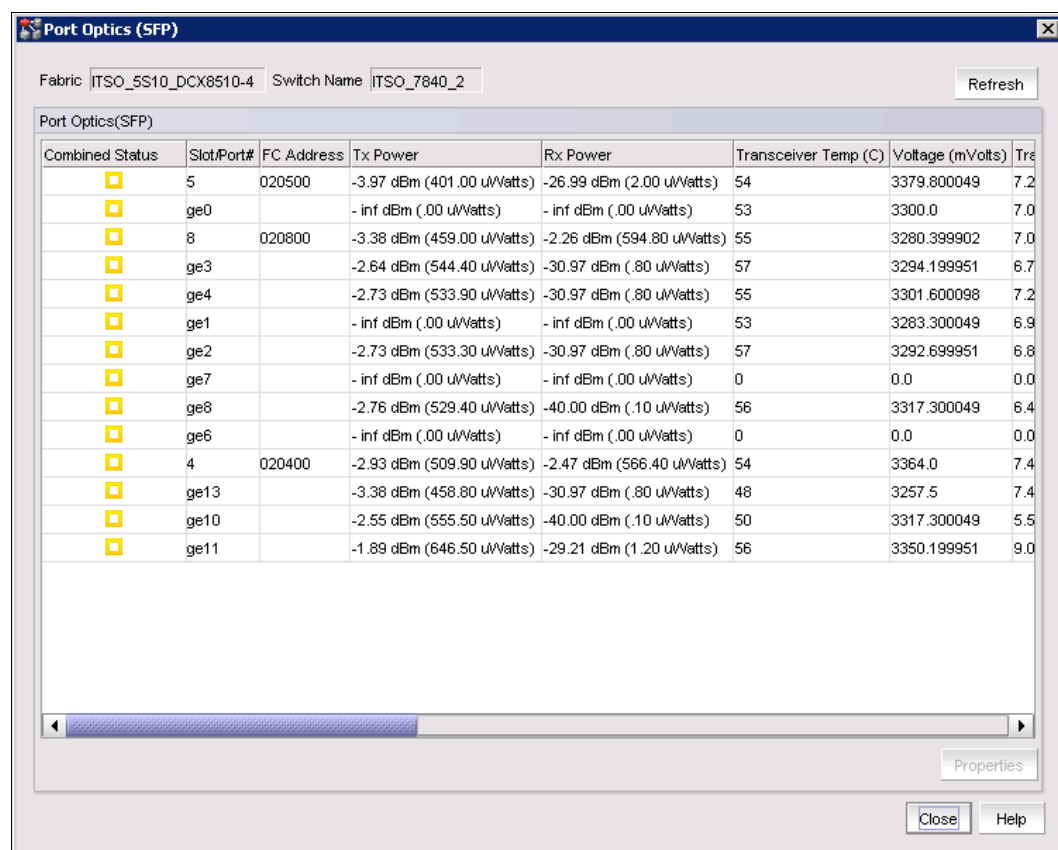
To complete the review of port health, complete the following steps to check the SFP status and optic levels:

1. Log in to IBM Network Advisor with an account that has admin privileges.
2. Select the SAN tab.

The Product List and Topologies panes are displayed.

3. In the Product List, highlight the switch for which SFP and optic levels are required and select **Monitor** → **Port Optics (SFP)**.

The Port Optics (SFP) window is displayed as shown in Figure 12-6.



The screenshot shows the 'Port Optics (SFP)' window. At the top, there are fields for 'Fabric' (ITSO_SS10_DCX8510-4) and 'Switch Name' (ITSO_7840_2), along with a 'Refresh' button. Below this is a table titled 'Port Optics(SFP)' with columns: Combined Status, Slot/Port#, FC Address, Tx Power, Rx Power, Transceiver Temp (C), Voltage (mVolts), and Tr. The table lists data for multiple ports, including ge0, ge3, ge4, ge1, ge2, ge7, ge8, ge6, ge13, ge10, and ge11. Each row shows the port's status (indicated by a yellow square icon), its slot/port number, FC address, and various power and temperature readings.

Combined Status	Slot/Port#	FC Address	Tx Power	Rx Power	Transceiver Temp (C)	Voltage (mVolts)	Tr
	5	020500	-3.97 dBm (401.00 uWatts)	-26.99 dBm (2.00 uWatts)	54	3379.800049	7.2
	ge0		- inf dBm (.00 uWatts)	- inf dBm (.00 uWatts)	53	3300.0	7.0
	8	020800	-3.38 dBm (459.00 uWatts)	-2.26 dBm (594.80 uWatts)	55	3280.399902	7.0
	ge3		-2.64 dBm (544.40 uWatts)	-30.97 dBm (.80 uWatts)	57	3294.199951	6.7
	ge4		-2.73 dBm (533.90 uWatts)	-30.97 dBm (.80 uWatts)	55	3301.600098	7.2
	ge1		- inf dBm (.00 uWatts)	- inf dBm (.00 uWatts)	53	3283.300049	6.9
	ge2		-2.73 dBm (533.30 uWatts)	-30.97 dBm (.80 uWatts)	57	3292.699951	6.8
	ge7		- inf dBm (.00 uWatts)	- inf dBm (.00 uWatts)	0	0.0	0.0
	ge8		-2.76 dBm (529.40 uWatts)	-40.00 dBm (.10 uWatts)	56	3317.300049	6.4
	ge6		- inf dBm (.00 uWatts)	- inf dBm (.00 uWatts)	0	0.0	0.0
	4	020400	-2.93 dBm (509.90 uWatts)	-2.47 dBm (566.40 uWatts)	54	3364.0	7.4
	ge13		-3.38 dBm (458.80 uWatts)	-30.97 dBm (.80 uWatts)	48	3257.5	7.4
	ge10		-2.55 dBm (555.50 uWatts)	-40.00 dBm (.10 uWatts)	50	3317.300049	5.5
	ge11		-1.89 dBm (646.50 uWatts)	-29.21 dBm (1.20 uWatts)	56	3350.199951	9.0

Figure 12-6 Port Optics (SFP) window

The Port Optics (SFP) window contains the SFP information and optic readings for all switches in the port.

12.5 System messages and RAS logs

This section describes the types of system messages and what to do with them.

12.5.1 System message types

FOS supports three types of system messages. A system message can be of one or more of the following types:

- ▶ RASLog messages
- ▶ Audit log messages
- ▶ First Failure Data Capture messages

12.5.2 RASLog messages

RASLog messages report significant system events (failure, error, or critical conditions) or information, and are also used to show the status of the high-level user-initiated actions. RASLog messages are forwarded to the console, to the configured syslog servers, and to the

Simple Network Management Protocol (SNMP) management station through SNMP traps or informs.

The **errDump** command shows the error log without pagination, and the **errShow** command shows the error log messages with pagination.

The system messages are documented in the *Fabric OS Message Reference* guide, which helps you diagnose and fix problems. You can find this guide at the following website:

<http://my.brocade.com/>

The messages are organized alphabetically by module name. A module is a subsystem in the Fabric Operating System (FOS). Each module generates a set of numbered messages. For each message, the guide provides message text, probable cause, recommended action, and severity level. There might be more than one cause and more than one recommended action for any specific message, but the guide describes the most probable cause and typical action that is recommended.

12.5.3 Audit log messages

Event auditing is designed to support post-event audits and problem determination. It is based on high-frequency events of certain types, such as security violations, zoning configuration changes, firmware downloads, and certain types of fabric events. Audit messages that are flagged as AUDIT are not saved in the switch error logs. The switch can be configured to stream audit messages to the switch console and to forward the messages to specified syslog servers. The audit log messages are not forwarded to an SNMP management station. There is no limit to the number of audit events.

12.5.4 First-Failure data capture messages

First-Failure data capture (FFDC) is used to capture failure-specific data when a problem or failure is noted for the first time and before the switch reboots or trace and log buffers are wrapped. All subsequent iterations of the same error are ignored. This critical debug information is saved in nonvolatile storage and can be retrieved by running **supportSave**. The FFDC data is used for debugging or analyzing the problem. FFDC is intended for use by Brocade technical support.

12.6 SAN health

Brocade SAN Health is a no-charge software utility that is designed to securely audit and analyze your SAN environment. To help optimize your SAN's performance, SAN Health automatically discovers critical fabric characteristics and reports their details in easy-to-understand Excel and Visio formats. In addition, SAN Health performs critical tasks, such as the following ones:

- ▶ Taking inventory of devices, switches, firmware versions, and fabrics
- ▶ Capturing and displaying historical performance data
- ▶ Comparing zoning and switch configurations against preferred practices
- ▶ Assessing performance statistics and error conditions
- ▶ Producing detailed graphical reports and diagrams

SAN Health gives you a powerful tool that helps you focus on optimizing your SAN rather than manually tracking its components. In fact, a wide variety of useful features make it easier for you to collect data, identify potential issues, and check your results over time.

To provide a comprehensive report about your SAN environment, SAN Health uses two main components:

- ▶ Data capture application
- ▶ Back-end report processing engine

After SAN Health finishes the capture of switch diagnostic data, the back-end reporting process automatically generates a Visio topology diagram and a detailed snapshot report of your SAN configuration. This summary report contains information about the entire SAN and specific details about fabrics, switches, and individual ports. Other useful items in the report include alerts, historical performance graphs, and preferred practices. SAN Health delivers topology diagrams, comprehensive reports, detailed explanations, and more.

It is important to note that the SAN health report is a snapshot, so reflects only the condition of the fabric at that specific point. In complex or busy environments, statuses can change quickly, so the SAN health report should be used as a reference and confirmed with real-time investigation.

12.6.1 Installing Brocade SAN Health

To install Brocade SAN Health, complete the following steps:

1. Go to the following link and download SAN Health Diagnostics Capture:

<http://www.brocade.com/en/support/support-tools/support-download-san-health-diagnostics-capture.html>

2. Download, extract, and install the application with the default settings.

12.6.2 Using SAN Health Diagnostics Capture

To run the application to collect and upload the fabric information for report generation, complete the following steps:

1. Launch the SAN Health collection tool from the desktop icon or **Start** menu option on MS Windows based systems. The SAN Health application window will be displayed (see Figure 12-7 on page 313).
2. Click the **New** button at upper left of the SAN Health application window.
3. Select the Site Details tab and complete the fields. These details are used on the title page of the reports. The report is processed at Brocade and is returned by using a secure single sign-on web page. An ID is automatically created from the site details if one does not already exist.
4. Click the Add Switches tab and add or confirm switch IPs and log-in credentials.
5. Click the Fabric tab, select the fabric or fabrics in the tree view to provide a name for each fabric, and then select the performance capture duration from the pull-down menu.
6. Click **Test Fabric Connectivity Get Switch Details** to ensure that all switches can communicate with the SAN Health Application.
7. Click the **Save** button on the top task bar to save the configuration for future use.

8. Click the Capture tab and click **Start Audit**. Before the audit starts, a set of “pre-flight” checks are run to ensure that the data values have been entered correctly. Any items reported as incorrectly entered must be corrected before the application allows an audit to start.

When the audit begins, processing time depends on the capture performance data interval that is set on the Fabric tab. The progress of the tool as it completes the audit is displayed in left pane of the window.

Figure 12-7 shows the SAN Health application window with the Site Details tab selected.

The screenshot shows the SAN Health Version 4.0.6 application window. The 'Site Details' tab is selected. The form contains the following fields:

- Name this Report: ITSD Rebooks
- User Details: Salutation (Mr/Ms), First Name (Jane), Last Name (Doe), Job Title (Fabric Administrator), Phone (123-123-1234)
- Company Details: Name (IBM Ltd), Address1 (3600 Steeles Ave. E), Address2, City (Markham), Zip/Postal Code (P7G 1R5), State/Province (Ontario), Country (Canada)
- Report Return: Email (ITSD@ibm.com), Retype Email (ITSD@ibm.com)
- Optional: Send a copy of the report to the following people? (A Brocade staff member you are working with, Brocade Support (Case number required), Another company that you are working with (Company Name: IBM Ltd, Email Address: ITSD@ca.ibm.com))

On the right, there are instructions for completing a SAN Health audit:

- 1) Name the report, enter site details and report return options. These details are used on the title page of the report and to ensure correct report return. Report processing is centralized at Brocade and reports are returned via a secure single sign on web page. If you don't already have a sign on, one is automatically created from the site details you enter.
- 2) Enter a Switch IP Address and login credentials to add a fabric. Ensure that switch login credentials are correct and optionally set the Visio diagram position for each switch.
- 3) Click on the fabric(s) in the tree view and complete the fabric details. Name the fabric, enter the performance capture duration and specify the support provider for each fabric.
- 4) Test Connectivity to the fabric members. Click on the "Test Connectivity" button at the fabric or individual switch levels to test login credentials. Each switch fabric membership, type, model number and communication capabilities are determined.
- 5) Save the audit set file you have just created. Save the SET file so that you don't have to enter all of these details again for future audits.
- 6) Click on "Start Audit"

A legend on the right explains various icons: Incomplete data, Complete data, Audit completed, Warning, Audit aborted, Searching for open SSH or Telnet port, Session refused, Bad login credentials, Capturing throughput data, Session timed out, Exchanging login credentials, Gathering the output from CLI diagnostic commands, Brocade device discovered, Cisco device discovered, Unknown device, and Logical Switches.

At the bottom, an 'ACTIVITY LOG' shows the following entries:

- 10:03:48 HTTPS Activity - Sending data to the server
- 10:03:48 HTTPS Activity - Waiting for response from the server
- 10:03:48 HTTPS Activity - Returned from the specific page request OK
- 10:03:48 HTTPS Activity - Returned from data transmission OK
- 10:03:48 HTTPS Activity - Connected (idle)
- 10:03:48 File Upload Completed Successfully
- 10:03:48 HTTPS Activity Completed OK
- 10:04:10 (1Fabric, 2 Switches, 2 Selected, 0 Completed, 0 Failed)

Figure 12-7 SAN Health application with the Site Details tab selected

9. To receive the report, the encrypted SAN Health file (.BSH) file must be sent to the Brocade report generator. You have three options to complete this task:
 - Click **Send the diagnostics data file via HTTPS**.
 - Upload the file at <https://my.brocade.com/upload/ReportGeneration.jsp>.
 - Send the file as an email attachment to <mailto:SHUpload@brocade.com>.
10. A report generation notification email is sent from the Brocade SAN Health Administrator when the report is available for download at the MyBrocade portal. The report contains a spreadsheet, a Visio file, and SHData files. This last file can be loaded by SAN Health Professional to do advanced analysis.

12.6.3 SAN Health Professional

Brocade SAN Health Professional provides an easy-to-understand framework for analyzing SAN components and configuration data that is captured by the SAN Health Diagnostics Capture utility. It provides a straightforward, easy-to-navigate user interface for auditing SAN Health data captures, making it a valuable tool for SAN inventory tracking and change management activities. You can import up to two SAN Health Diagnostic Capture captures to SAN Health Professional for immediate, detailed analysis about any SAN component.

In addition to its standard data analysis and search capabilities, the SAN Health Professional framework supports optional add-on modules.

SAN Health Professional provides a straightforward, easy-to-navigate user interface for auditing SAN Health data captures, making it a valuable tool for inventory tracking and change management activities. Organizations can import up to two SAN Health captures to SAN Health Professional for immediate, detailed analysis about any component.

To enable the highest level of flexibility, SAN Health Professional provides extensive searching and filtering capabilities. Searches can be broad (for example, a single search string such as “HBA” or “CHIPID”) or precise. Precision searching narrows the search to any combination of attribute names, devices, ports, switches, directors, fabrics, aliases, zones, or configurations.

For more information about SAN Health Professional, including package download and installation instructions, see the Brocade SAN Health website:

<http://www.brocade.com>

12.7 Collecting support data

The following sections detail some of the diagnostic features that can gather relevant support information.

The **supportShow** CLI command can be run on the switch to dump important diagnostic and status information to the session window for review and capture. Most Telnet and SSH clients offer the ability to capture session data before opening the session. IBM Network Advisor and CLI (use the **supportsave** CLI command) allow administrators to collect a supportsave. The **supportsave** command collects a number of important outputs in a series of files and compresses them into a package. This package can then be sent to the service provider for support and problem determination. The **supportsave** data package is preferred by most support providers.

12.7.1 Saving comprehensive diagnostic files to the server

To save comprehensive diagnostic files to the server, connect to the switch, log in as the admin user, run **supportSave -c**, and respond to the prompts. The **-c** flag uses the FTP, SCP, or SFTP parameters that are saved by the **supportFtp** command. If this flag is omitted, the FTP, SCP, or SFTP parameters must be specified through command-line options or interactively. To display the current **supportFtp** parameters, run **supportFtp**. On a dual-CP system, run **supportFtp** on the active CP.

12.7.2 Scheduling technical support and event information collection by using IBM Network Advisor

Technical support and event information can be collected for up to 50 devices. Technical SupportSave uses the built-in FTP, SCP, or SFTP server that is configured on the Management server to save data.

To capture technical support and event information at a predetermined date and time, complete the following steps:

1. Click **Monitor** → **Technical Support** → **Product/Host SupportSave**. The Technical SupportSave window opens.
2. Click the Schedule tab.
3. Select **Enable scheduled Technical Support Data**.
4. Select how often the scheduled collection will occur from the **Frequency** list.
5. Select the start date for the scheduled collection from the **Start Date** list. This list is only available when **Weekly** or **Monthly** is selected from the **Frequency** list.
6. Select the time that the scheduled collection will begin from the **Start Time Hour** and **Minute** lists.
7. Click the SAN Products tab, if necessary, and complete the following steps:
 - a. Right-click in the **Available SAN Products** table and select **Expand All**.
 - b. Select the switches to collect data for in the **Available SAN Products** table, and click the right arrow to move them to the **Selected Products and Hosts** table.

The **Available SAN Products** table displays the following information:

- All Levels: All discovered devices and ports as both text and icons.
- Name: The name of the available switch.
- Product Type: The type of product.
- Tag: The tag number of the device.
- Serial #: The serial number of the device.
- WWN: The switch port's worldwide name.
- IP Address: The switch port's IP address.
- Domain ID: The switch port's top-level addressing hierarchy of the domain.
- Vendor: The hardware vendor's name.
- Model: The name and model number of the hardware.
- Port Count: The total number of ports.
- Firmware: The firmware version.
- Location: The customer site location.
- Contact: The primary contact at the customer site.
- Description: A description of the customer site.
- State: The switch state, for example, online or offline.
- Status: The operational status of the switch, for example, unknown or marginal.

8. Click the Hosts tab and complete the following steps:
 - a. Right-click in the **Available SAN Products** table and select **Expand All**.
 - b. Select the products to collect data for in the **Available Hosts** table and click the right arrow to move them to the **Selected Products and Hosts** table. The **Selected Products and Hosts** table displays the following information:
 - IP Address: The IP address of the selected product or host.
 - Name: The name of the selected product or host.
 - WWN: The worldwide name of the selected product or host.

- **Firmware Type:** The type of firmware: FOS (Fabric OS).
- **Firmware version:** The firmware version of the selected product or host.
- **Support Save Credentials:** Whether the product or host has SupportSave credentials or not.

The **Available Hosts** table displays the following information:

- **Name:** The name of the available host.
- **IP Address:** The host port's IP address.
- **Network Address:** The network address of the host.
- **Fabrics:** The fabric of the host.

9. Select how often to purge the support data from the **Purge Support Data** list.
10. Click **OK** on the Technical SupportSave window.

12.7.3 Starting immediate technical support information collection

To capture technical support and event information for specified devices, complete the following steps:

1. Click **Monitor** → **Technical Support** → **Product/Host SupportSave**. The Technical SupportSave window is displayed.
2. Click the **Generate Now** tab, if necessary.
3. Click the **SAN Products** tab, if necessary, and complete the following steps:
 - a. Right-click in the **Available SAN Products** table and select **Expand All**.
 - b. Select the switches to collect data from in the **Available SAN Products** table and click the right arrow to move them to the **Selected Products and Hosts** table. Technical SupportSave data for Fabric OS devices is saved to the following directory:
`Install_Home\data\ftproot\technicalsupport\`
4. Click the **Hosts** tab, if necessary, and complete the following steps:
 - a. Right-click in the **Available Hosts** table and select **Expand All**.
 - b. Select the hosts to collect data from in the **Available Hosts** table and click the right arrow to move them to the **Selected Products and Hosts** table.
5. Click **OK** on the Technical SupportSave window. The Technical SupportSave Status window opens with the following details:
 - **Name:** The name of the product.
 - **IP Address:** The product's IP address.
 - **Firmware Type:** The type of product.
 - **Progress:** The status of the supportsave collection. On products that are running FOS V7.0 or later, this field shows the percentage complete and is updated every minute. For Host products, and FOS products that are running Version 6.4 or earlier, this field cannot display the percentage (it only displays whether it is "In Progress" or "Completed").
 - **Status:** The status of the supportsave process, for example, Success or Failure.
6. Click **Close** on the Technical SupportSave Status window.

12.8 Using MAPS for problem determination

Monitoring and Alerting Policy Suite (MAPS) as described in “Flow Vision” on page 192 provides a number of tools that can be used for problem determination and monitoring.

Because MAPS can be used to monitor and report on port health, back-end health, and field-replaceable unit (FRU) health, it is an excellent tool to review data if problems occur in the fabric.

The following sections describe some of these features from a problem determination perspective.

More documentation:

The *Monitoring and Alerting Policy Suite Administrator's Guide* provides detailed instructions for configuring and monitoring all aspects of MAPS by using the CLI.

Detailed instructions for using MAPS with the graphical user interface are provided in the *Brocade Network Advisor SAN User Manual*.

You can download these publications at the following website:

<https://my.brocade.com>

12.8.1 Port health and cyclic redundancy checks monitoring

The Port Health category monitors port statistics and takes action based on the configured thresholds and actions. Thresholds can be configured per port type and optionally applied to all ports of the specified type. Ports whose thresholds can be monitored include physical ports, D_Ports, E_Ports, F_Ports, and Virtual E_Ports. The Port Health category also monitors the physical aspects of an SFP transceiver, such as voltage, current, receive power (RXP), transmit power (TXP), and state changes in physical ports, D_Ports, E_Ports, and F_Ports.

Note: For detailed information and instructions on configuring, creating, and editing thresholds and actions, see the *Monitoring and Alerting Policy Suite Administrator's Guide*, which is available at:

<https://my.brocade.com>

Cyclic redundancy checks (CRCs) serve to validate the integrity of a frame. When validation fails, the frame that arrived at the receiving port is corrupted. Frame corruption can be caused by device disconnection, an optical transceiver failure at the device, fiber optic cabling, or a poor connection.

Two types of CRC errors can be logged on a switch. Taken together, they can help determine which link introduced the error into the fabric. The two types are plain CRCs, which have bad end-of-frame (EOF) markers, and CRCs with good EOF (crc g_eof) markers. When a crc g_eof error is detected on a port, it indicates that the transmitter or path from the sending side might be a possible source. When a complete frame that contains a CRC error is first detected, the error is logged, and the good EOF (EOFn) is replaced with a bad EOF marker (EOFni). Because Brocade switches forward all packets to their endpoints, changing the EOF marker allows the packet to continue but not be counted.

For MAPS threshold and fencing purposes, only frames with CRC errors and good end-of-frame markers are counted. This process enables you to know exactly how many errors were originated in a specific link.

You can use the MAPS dashboard to display an overview of port health and then investigate it further in more detail.

To display port health, complete the following steps:

1. Open IBM Network Advisor with an ID that has admin privileges.
2. Click the Dashboard tab to display the dashboard as shown in Figure 12-8.

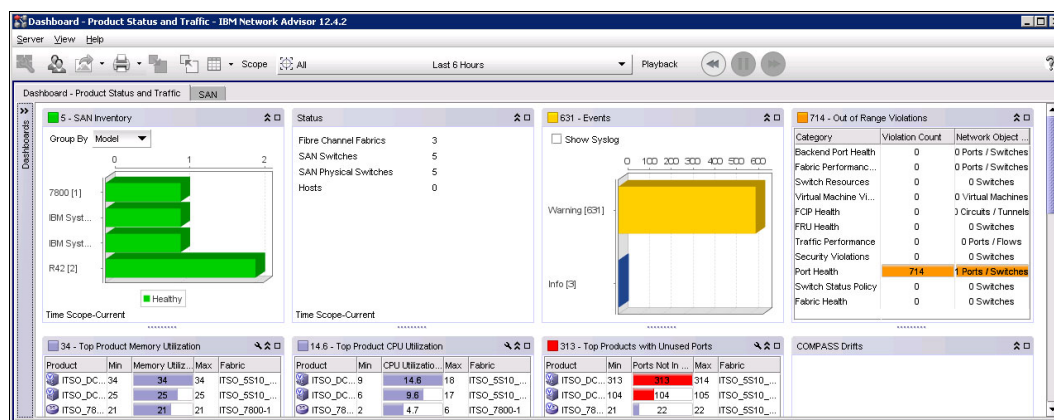


Figure 12-8 MAPS Dashboard

3. To display the port health metrics, select the **Dashboard** pull-out and the available dashboards are displayed.
4. Click the + to expand the default dashboards available and select **SAN Port Health** (see Figure 12-9 on page 318).
5. Select the **Scope** pull-down menu to alter the time span over which the default metrics are reported from in the display in the lower right pane.

Figure 12-9 shows MAPS Dashboard with SAN Port health selected.

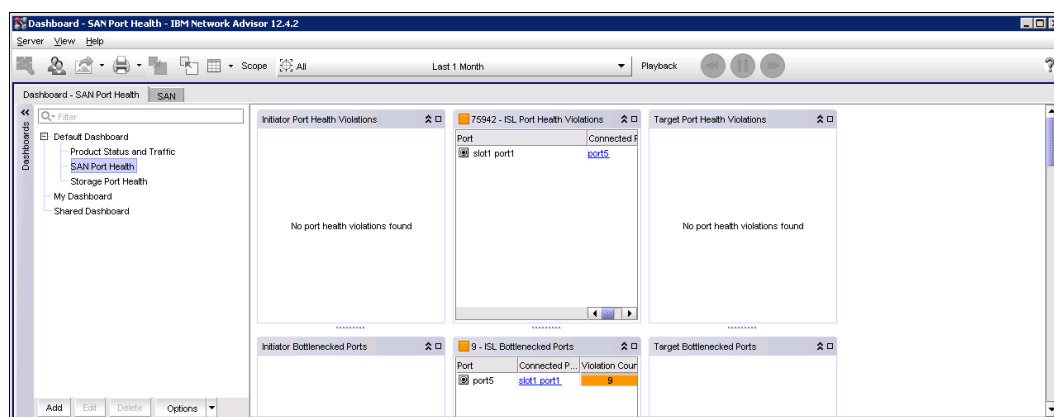


Figure 12-9 MAPS Dashboard with SAN Port Health selected

6. Scroll to view the CRCs reported for both inter-switch link (ISL) and fabric ports for the time frame that is selected.

7. Use the Top Initiator Ports CRC Errors, Top ISL Ports CRC Errors, and Top Target Ports CRC errors widgets to view all relevant CRC errors.

Figure 12-10 shows an example of the Top ISL Ports CRC Errors widget.

Port	Connected P...	CRC Errors	CRC Errors/s...	Link Failures	Seq
slot1 port1 port5		18	0	0	0.0

Refreshed- 11:46 AM

Figure 12-10 TOP ISL Ports CRC Errors widget

When you use this method, ports that encounter CRC errors can quickly be discovered and corrected.

12.8.2 Back-end port monitoring

Back-end ports can be connected to ports within a fixed-port switch or to other blades within the switch chassis. Therefore, their functionality is different from front-end ports, which connect to devices outside of the switch. The primary task of back-end ports is to route packets that pass through a switch's ASICs. Switch (and consequently fabric) performance degrades when there are errors in back-end ports. MAPS error notification allows for earlier corrective action.

MAPS monitors the back-end port counter statistics for back-end ports through the group ALL_BE_PORTS, which identifies each port by using a slot-port combination such as 3/3/1. For fixed-port switches, the slot number is 0. Worldwide name (WWN) IDs and Port Names are not supported. The predefined groups for front-end ports do not apply to back-end ports. History data for back-end ports is collected for seven days.

The Back-end health category in the MAPS widget "Out of Range Violations" enables monitoring of the health of the back-end switch ports for these errors:

- ▶ CRC and Link reset error rates
- ▶ Invalid transmission words
- ▶ BAD_OS
- ▶ Frame length (either too long or truncated)

Figure 12-11 shows the Back-end port monitor in the MAPS Out of Range Violations widget that is available in the IBM Network Advisor dashboard.

620 - Out of Range Violations		
Category	Violation Count	Network Object Count
Backend Port Health	0	0 Ports / Switches
Fabric Performanc...	214	7 Ports / Switches
Switch Resources	0	0 Switches
Virtual Machine Vi...	0	0 Virtual Machines
FCIP Health	0	0 Circuits / Tunnels
FRU Health	0	0 Switches
Traffic Performance	0	0 Ports / Flows
Security Violations	0	0 Switches
Port Health	406	2 Ports / Switches
Switch Status Policy	0	0 Switches
Fabric Health	0	0 Switches

Figure 12-11 Back-end Port monitor

To display the Back-end Port Health monitor, complete the following steps:

1. Open IBM Network Advisor with an ID that has admin privileges.
2. Click the Dashboard tab and the dashboard is displayed.
3. To display the back-end port health metrics, select the **Dashboard** pull-out and the available dashboards are displayed.
4. Click the **+** to expand the default dashboards available and select **Product Status and Traffic** (see Figure 12-9 on page 318).
5. Select the **Scope** pull-down menu to alter the time span over which the default metrics are reported from in the display in the lower right pane.
6. Scroll to view the Out of Range Violations widget.
7. If any violations are reported, right-click the row and select **Violations** to display the Back-end Port Violation window. More information about the parameter that is out of range is displayed in the lower pane.

The window displays a number of valuable information columns, including the device that reported the violation, the time the violation occurred, the metric that was out of range, and the recommended action.

12.8.3 FRU Health

The FRU Health category enables rules to be defined for FRUs. The following table lists the monitored parameters in this category. Possible states for all FRU measures are faulty, inserted, on, off, and out.

Table 12-4 shows a list of supported FRU Health category parameters.

Table 12-4 FRU Health category parameters

Monitored parameter	Description
Power supplies (PS_STATE)	State of a power supply has changed.
Fans (FAN_STATE)	State of a fan has changed.
Blades (BLADE_STATE)	State of a slot has changed.
SFPs (SFP_STATE)	State of the SFP transceiver has changed.
WWN (WWN_STATE)	State of a WWN card has changed.

The FRU Health category is reported in the MAPS Out of Range Violations widget available in the IBM Network Advisor dashboard.

To open the widget, complete the following steps:

1. Open IBM Network Advisor with an ID with that has admin privileges.
2. Click the **Dashboard** tab to display the dashboard.
3. To display the FRU Health metrics, select the **Dashboard** pull-out, which displays the available dashboards.
4. Click the **+** to expand the default dashboards that are available and select the **Product Status and Traffic** menu option (see Figure 12-9 on page 318).
5. Select the **Scope** pull-down menu to alter the time span over which the default metrics are reported from in the display in the lower right pane.
6. Scroll to view the Out of Range Violations widget.
7. If any violations are reported, right-click the row and select **Violations** to display the FRU Health Violation panel.

More information about the parameter that is out of range is displayed in the lower pane (see Figure 12-12).

Time	Product	Objec...	Port ...	Rule Con...	Meas...	Unit	Margi...	Critical	RAS ...	SNMP	Port ...	Fence	E-mail	Recommended Action
Thu Nov 1 ...	ITSO_DCX8510-8	U-Port	ALL_PO...	IN					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-8	U-Port	ALL_PO...	OUT					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-8	U-Port	ALL_PO...	IN					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-8	U-Port	ALL_PO...	IN					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-8	U-Port	ALL_PO...	OUT					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-8	U-Port	ALL_PO...	IN					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-8	U-Port	ALL_PO...	OUT					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-4	U-Port	ALL_PO...	OUT					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-4	U-Port	ALL_PO...	IN					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-4	U-Port	ALL_PO...	IN					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-4	U-Port	ALL_PO...	IN					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-4	U-Port	ALL_PO...	OUT					Disable	Disable				Disable SFP state changes occur when the SFP is inser
Thu Nov 1 ...	ITSO_DCX8510-4	U-Port	ALL_PO...	IN					Disable	Disable				Disable SFP state changes occur when the SFP is inser

Figure 12-12 FRU Health Violations panel

12.8.4 Fibre Channel over IP (FCIP) Health

The FCIP Health category enables definition of rules for FCIP health, including circuit state changes, circuit state utilization, and packet loss. The FCIP monitors support Minute, Hour, Day, and Week time bases for these parameters:

- ▶ Tunnel state change
- ▶ Tunnel throughput
- ▶ Tunnel QoS throughput
- ▶ Tunnel QoS Packet loss
- ▶ FCIP Circuit State Changes
- ▶ FCIP Circuit Utilization
- ▶ FCIP Packet loss
- ▶ FCIP Circuit Round Trip Time
- ▶ FCIP connection variance

The three default monitoring policies (Aggressive, Moderate, and Conservative) have preconfigured thresholds for FCIP criteria for MAPS monitoring. Each default policy has actions that are triggered when the reported value is greater than the threshold value.

Note: Additional FCIP parameters are monitored on the IBM SAN42B-R extension switch. In addition to the monitors that are available in the default policies, the following things are also monitored:

- ▶ FCIP tunnel state change (STATE_CHG): The number of FCIP tunnel state changes. This count applies to the tunnel group only.
- ▶ FCIP tunnel or tunnel QoS utilization (UTIL): The percentage of FCIP utilization. This this monitoring applies to both the tunnel and the tunnel QoS groups.

For more information about FCIP categories and monitoring, see the *Monitoring and Alerting Policy Suite Administrator's Guide* that is available at the following website:

<http://my.brocade.com>

To view FCIP Health, complete the following steps:

1. Open IBM Network Advisor with an ID that has admin privileges.
2. Click the Dashboard tab to display the dashboard.
3. To display the FCIP Health metrics, select the **Dashboard** pull-out and the available dashboards are displayed.
4. Click the **+** icon to expand the default dashboards available and select **Product Status and Traffic** (see Figure 12-9 on page 318).
5. Select the **Scope** pull-down menu to change the time span over which the default metrics are reported from in the display in the lower right pane.
6. Scroll to view the **Out of Range Violations** widget.
7. If any violations are reported, right-click the row and select **Violations** to display the FCIP Health Violation window. More information about the parameter that is out of range is displayed in the lower pane.

12.9 Flow Vision

Flow Vision (described in Chapter 7, “Fabric Vision” on page 191) provides a number of tools that can be employed for troubleshooting both before implementation or during unplanned events. Flow vision is divided into three categories: Flow Monitor, Flow Generator, and Flow Mirror. These categories will be reviewed in the following sections as they pertain to troubleshooting.

The statistics that are generated by using Flow Vision can also be monitored with the Monitoring and Alerting Policy Suite (MAPS) threshold service that makes the two features a powerful diagnostics and monitoring tool when they are used together.

Note: For more information about the Flow Vision tool, see the *Flow Vision Administrator's Guide*, which is available at the following website:

<https://my.brocade.com>

12.9.1 Flow Monitor

A common use of flow monitors is to monitor traffic that is flowing from a particular ingress port to a specified target egress port. Monitoring of various frame types at a switch port can provide insights into storage I/O access patterns at a LUN, reservation conflicts, and I/O errors. Examples of the frame types that can be monitored include SCSI Aborts, SCSI Read, SCSI Write, SCSI Reserve, and all rejected frames.

When a problem is discovered in a flow either by MAPS threshold monitoring of preconfigured flows or a Flow has been defined for the path that is to be investigated, the Flow Monitoring window can be displayed to view detailed information about that flow. See “Flow Vision” on page 19 for more information about configuring flows.

To view the summary data for a Flow Monitor flow, complete the following steps:

1. Select the device on which a defined flow has been configured and select **Monitor** → **Fabric Vision** → **Flow Monitor**. The Flow Vision window is displayed pre-populated with a list of all defined flows in the **Flow Definitions** table (see Figure 12-13).

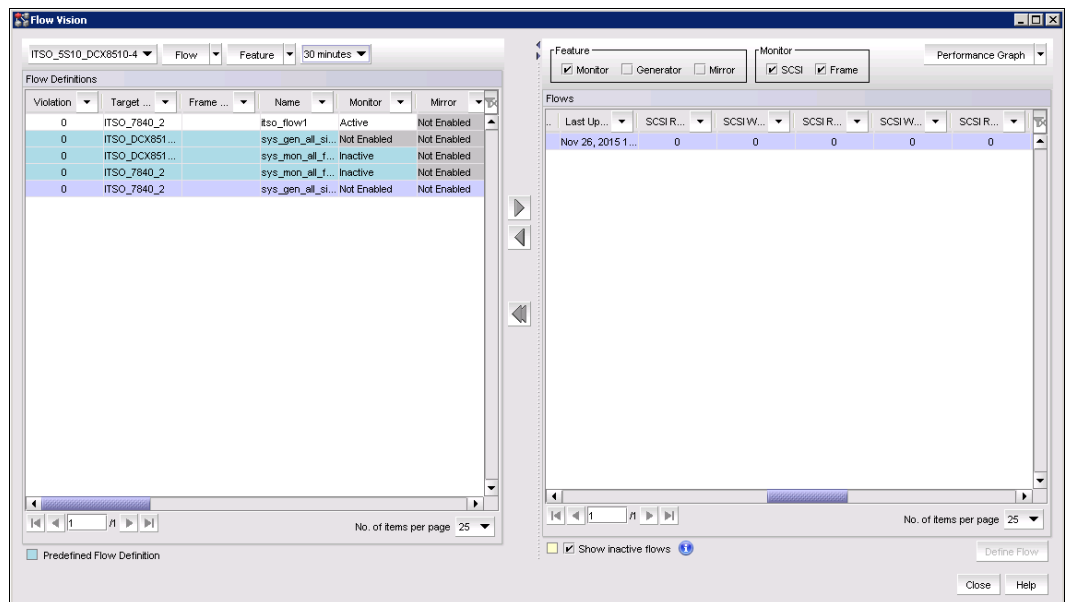


Figure 12-13 Flow Vision monitor panel

2. Select a time interval for monitoring the flow in the Time duration list. Possible values are 30 minutes, 1 hour, 6 hours, 12 hours, 1 day, 3 days, 1 week, and 1 month.
3. Select the flow to be monitor in the **Flow Definitions** table.
4. Click the **right arrow** button to display the selected flow in the **Flows** table.
5. Select or clear the **SCSI** check box to display or hide SCSI-related measures.
SCSI-related measures include SCSI read count, write count, read rate, write rate, read data, write data, and read and write frame data.
6. Select or clear the **Frame** check box to display or hide frame-related measures.
Frame-related measures include transmit (Tx) and receive (Rx) frame count, Transmit frame and receive frame rate, Transmit and receive word count, and transmit and receive throughput.

7. The subflow data for the selected Flow Monitor flow is displayed for review.

Data updates dynamically every 5 minutes.

Note: Many metrics are available. However, all of them might not display, depending on the configured flow. For descriptions of Flow Monitor review and the metrics that are available, see the *Flow Vision Administrator's Guide* that is available at the following website:

<https://my.brocade.com>

12.9.2 Flow Generator

Flow Generator is a test traffic generator for pre-testing the SAN infrastructure (including internal connections) for robustness before deploying it.

The configuration of Flow Generator and flow monitor is described in the *Flow Vision Administrator's Guide* which is available at the following website:

<https://my.brocade.com>

Flow Generator flows can be monitored by using Flow Monitor. For example, a combination of Flow Generator flows and Flow Monitor flows can be used to verify per-flow throughput at an ingress or egress port. This feature can be useful when more than one Flow Generator flow share an ingress or egress port. To do this, a flow must be created using both the Flow Generator and Flow Monitor features that share the ingress or egress port. In this manner, a link can be tested before implementation.

12.9.3 Flow Mirroring

Flow Mirror duplicates the specified frames in a user-defined flow and sends them to the switch CPU or sends them to a mirror port.

Note: Only 64 bytes of the FC frame is captured when mirrored to the switch CPU.

In this manner, Flow Mirror can assist with diagnosing a number of symptoms and implementing an analyzer without interruption to the link.

Here are examples of problems that can be investigated with Flow Mirror:

- ▶ Diagnosing excessive SCSI reserve and release activity
- ▶ Diagnosing a slow-draining F_Port
- ▶ Tracking SCSI commands
- ▶ Tracking latency between a host and all connected targets
- ▶ Troubleshooting protocol errors

For more information about how to create a Flow Mirror port, and detailed use cases and examples, see the *SAN User Manual and Flow Vision Administrator's Guide* that is available at the following website:

<https://my.brocade.com>

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

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *IBM b-type Gen 5 16 Gbps Switches and Network Advisor*, SG24-8186
- ▶ *IBM Network Advisor*, TIPS1124
- ▶ *IBM System Networking SAN24B-5 Switch*, TIPS1128
- ▶ *IBM System Networking SAN96B-5*, TIPS1103
- ▶ *IBM System Storage SAN06B-R Extension Switch*, TIPS1126
- ▶ *IBM System Storage SAN42B-R Extension Switch*, TIPS1209
- ▶ *IBM System Storage SAN48B-5*, TIPS1125
- ▶ *IBM System Storage SAN768B-2 and SAN384B-2 Fabric Backbones*, TIPS1127

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- ▶ Data and storage management
<http://www.ibm.com/software/tivoli/csi/cloud-storage/>
- ▶ IBM System Storage SAN b-type family
<http://www.ibm.com/systems/storage/san/b-type/>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

Implementing or Migrating to an IBM Gen 5 b-type SAN

SG24-8331-00

ISBN 0738441821



(0.5" spine)

0.475" <-> 0.873"

250 <-> 459 pages



SG24-8331-00

ISBN 0738441821

Printed in U.S.A.

Get connected

