IBM

# Implementing an Open IBM SAN

- Discover the latest additions to the IBM SAN family
- Enhance your skills while using an easy-to-follow format
- Grow with the new technology

Jon Tate
Khalid Ansari
Pauli Rämö
Glen Routley

Redbooks

ibm.com/redbooks

International Technical Support Organization

# Implementing an Open IBM SAN

October 2003

SG24-6116-03

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

**Fourth Edition (October 2003)**

This edition applies to the SAN hardware and software products described herein.

# Contents



 

# Figures

# Tables

# Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates 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. 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 IBM's application programming interfaces.

# Trademarks

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

AIX®
BladeCenter™
DB2®
Enterprise Storage Server®
ESCON®
@server™
FICON™
Illustra™
IBM®
ibm.com®
Netfinity®
NUMA-Q®
OS/390®
PowerPC®
pSeries®
Rational®
Redbooks™
Redbooks (logo) ™
RS/6000®
S/390®
Storage Tank™
System/390®
SANergy™
SP2®
Tivoli®
TotalStorage®
Wave®
xSeries®
zSeries®

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, 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 Sun Microsystems, Inc. in the United States, other countries, or both.

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

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

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

# Preface

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

In this IBM® Redbook, which is an update and major revision of the previous version, we have tried to consolidate as much of the critical information as possible while covering procedures and tasks that are likely to be encountered on a daily basis.

Each of the products described has much, much more functionality than we could ever hope to cover in just one redbook. The IBM SAN portfolio is rich in quality products that bring a vast amount of technicality and vitality to the SAN world. Their inclusion and selection is based on a thorough understanding of the storage networking environment that positions IBM, and therefore its customers and partners, in an ideal position to take advantage by their deployment.

We cover the latest additions to the IBM SAN family, which includes products from companies such as Brocade, Cisco, CNT, and McDATA. We show how they can be implemented in an open systems environment, and we focus on the Fibre Channel protocol (FCP) environment in particular. We address some of the key concepts that they bring to the market, and in each case, we give an overview of those functions that are essential to building a robust SAN environment.

In other redbooks we explore in greater depth the IBM SAN product family, Fibre Channel basics, and SAN design concepts. More information can be found in the IBM Redbooks™:

- *Introduction to Storage Area Networks*, SG24-5470
- *IBM SAN Survival Guide*, SG24-6143
- *Designing and Optimizing an IBM SAN*, SG24-6419

## The team that wrote this redbook

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


*L-R Glen, Jon, Khalid, and Pauli*

**Jon Tate** is a Project Manager for IBM TotalStorage® SAN Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2 support for IBM storage products. Jon has 17 years of experience in storage software and management, services and support, and is both an IBM Certified IT Specialist, and an IBM SAN Certified Specialist.

**Khalid Ansari** is a SAN Solutions Technical Support Specialist based in Research Triangle Park, North Carolina, US. He has more than 5 years of experience in TCP/IP Networking and Storage Networking Technologies. His areas of proficiency include Multi Protocol (ATM, Ethernet, Token Ring) Switches and Routers, FC based SAN Fabric (Cisco, McDATA, Brocade, and CNT) and Fibre Channel/TCP-IP Trace Analysis. He holds SNIA level 3 SAN specialist and Cisco Certified Network Professional certifications.

**Pauli Rämö** is an advisory IT Specialist in IBM Finland. He has 11 years of experience in RS/6000®, IBM @server™ pSeries®, AIX®, and Linux. Additionally, his other areas of expertise include HACMP, open systems storage solutions, and SAP R/3 Basis. He is an IBM Certified Advanced Technical Expert - RS/6000 AIX V4 and SAP Certified Technical Consultant for UNIX/Oracle and UNIX/DB2® UDB.

**Glen Routley** is an SNIA level 3 SAN certified specialist based in Melbourne, Australia. He has 19 years of experience in the S/390® field. His areas of expertise include providing technical support to the Australia / New Zealand geography for the SAN and NAS product lines, and he also specializes in large systems networking and VTS tape subsystems.

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

Rufus Credle
International Technical Support Organization, Raleigh Center

Tom Cady
Yvonne Lyon
Deanna Polm
Sokkieng Wang
International Technical Support Organization, San Jose Center

George DeBiasi
Brian Cartwright
Sven Eichelbaum
Uwe Hofmann
Thomas Jahn
Glen Routley
Eric Wong
The previous authors of this redbook

William Champion
Cedric Cook
Scott Drummond
Parker Grannis
Chuck Grimm
Andy Huryn
Edith Kropf
Tien Nguyen
Victoria Perris
Michael Starling
Peter Thurston
Diana Tseng
Karen Ward
Michelle Wright
Ruoyi Zhou
IBM Storage Systems Group

Pat Bodin
Skyline Computer

Jim Baldyga
James Carpignano
Henry Fong
Al Hicks
Daniel Krueger
Brian Steffler
Brocade Communications Systems

Dave Burchwell
CNT Technologies Corporation

Barbara Sewell
Larry Shane
McDATA Corporation

Tom and Jenny Chang
Garden Inn Hotel, Los Gatos, California

## Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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 Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

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

  **ibm.com**/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. QXXE  Building 80-E2
  650 Harry Road
  San Jose, California 95120-6099

1

# Implementing an IBM TotalStorage SAN Switch

In this chapter we introduce the IBM TotalStorage SAN Switch (2109 and 3534) products. We perform the steps required to install and configure a fabric using IBM TotalStorage SAN Switches; and then we perform some management functions, including upgrading firmware, implementing a secure fabric, and monitoring performance within the fabric.

## 1.1 Introducing the IBM TotalStorage SAN Switch

The IBM TotalStorage SAN Switch (3534 and 2109) family of products provide 2 Gb/s port-to-port non-blocking throughput with auto-sensing capability for connecting to older 1 Gb/s host servers, storage, and switches. They come in eight, sixteen, thirty-two, and up to 2 x 64 port models. Unlike hub-based Fibre Channel Arbitrated Loop (FC-AL) solutions, which reduce performance as devices are added by sharing the bandwidth, an IBM TotalStorage SAN Switch Fabric throughput continues to increase as additional ports are interconnected.

All of these models are fully interoperable with the previous IBM TotalStorage SAN Switches (Models S08 and S16), and can be added to existing fabrics, enabling transition from existing Fibre Channel storage networks to the faster technology.

### Features

In the following paragraphs, we describe some of the capabilities available on all the IBM TotalStorage SAN Switches.

#### *Auto-sensing speed negotiation*

The IBM TotalStorage SAN Switch uses internal Application Specific Integrated Circuits (ASICs) supporting link operation at either 2 Gb/s or 1 Gb/s. As a device is connected to a port, the link speed is negotiated to the highest speed that is supported by the device. This speed selection is auto-negotiated by the ASIC driver on a per-port basis.

If multiple devices are connected to a port (for example, on an FL_Port), the driver auto-negotiates for the highest common speed and sets the transmitter and receiver accordingly. Auto-sensing negotiation allows easy configuration.

#### *Frame filtering*

Zoning is a fabric management service that can be used to create logical subsets of devices within a SAN and enable partitioning of resources for management and access control purposes. Frame filtering enables the switch to provide zoning functions with finer granularity. Frame filtering can be used to set up port level zoning, world wide name zoning, device level zoning, protocol level zoning, and LUN level zoning. After the filter is set up, the complicated function of zoning and filtering can be achieved at wire speed.

#### *Performance Monitoring*

Performance Monitoring is a licensed feature that provides error and performance information to manage your storage environment. There are three types of monitoring:

- **Arbitrated Loop Physical Address (AL_PA) monitoring:** This provides information regarding the number of CRC errors.
- **End-to-end monitoring:** This provides information regarding a configured source identifier (SID) to destination identifier (DID) pair. Information includes the number of CRC errors for frames with the SID-DID pair, Fibre Channel words transmitted from the port for the SID-DID pair, and Fibre Channel words received for the port for the SID-DID pair.
- **Filter-based monitoring:** This provides error information with a customer-determined threshold.

#### *Trunking*

Trunking is a feature that enables traffic to be balanced across available inter-switch links (ISLs) while still preserving in-order delivery. On some Fibre Channel protocol devices, frame traffic between a source device and destination device must be delivered in-order within an exchange.

The requirement for in-order delivery in conjunction with the FSPF forces current devices to fix a routing path within a fabric. Consequently, certain traffic patterns in a fabric can cause all active traffic to be allocated to a single available path and leave other paths unused.

A trunking group (a set of available paths linking two adjacent switches) is created at the ASIC level within a switch, and therefore all paths in a trunking group must be within the same ASIC to form the trunk.

Ports in the trunking group are called trunking ports. One trunking port is designated as the trunking master port and is used to set up all routing paths for the entire trunking group. The trunk provides up to an 8 Gb/s single-aggregate ISL pipe between switches using four ISL paths.

### 1.1.1 Software specifications

In this section we describe the software for the IBM TotalStorage SAN Switches.

#### Fabric Operating System

The Fabric Operating System (FOS) manages the operation of the switch and delivers the same, and compatible, functionality to the different models of switches. The switch firmware is designed to make the switches easy to install and use while retaining the flexibility needed to accommodate user requirements. A fabric constructed with cascaded 2109 switches automatically assigns individual switch addresses, establishes frame routes, configures the internal name server, and so on.

Users can access internal management functions using standard host-based Simple Network Management Protocol (SNMP) software or Web browsers. They can access these functions using network connectivity through the Ethernet port or using Internet Protocol (IP) over the Fibre Channel ports. SCSI Enclosure Services (SES) is also supported as a management method. The management functions of the switch allow a user to monitor frame throughput, error statistics, fabric topology, fans, cooling, media type, port status, IDs, and other information to aid in system debugging and performance analysis.

## Fabric OS version

The FOS includes all the basic switch and fabric support software as well as optionally licensed software that is enabled using license keys. It is comprised of two major software components: firmware that initializes and manages the switch hardware, and diagnostics.

The F32 and M12 use Fabric OS (FOS) Version 4.x and is a Linux Based OS, while the F08 and F16 use FOS Version 3.x and the earlier S08 and S16 use version 2.x, which were based on a VxWorks operating system. We show the models and required Firmware versions in Table 1-1.

*Table 1-1 Firmware versions*

| Model | FOS | Secure FOS |
|---|---|---|
| 3534-1RU<br>2109-S08, S16 | v2.x.x | v2.6.1 |
| 3534-F08<br>2109-F16 | v3.0.x | v3.1.x |
| 2109-F32, M12 | v4.0.x | v4.1.x |

In the last column of the table, we show the levels of FOS required to implement a secure fabric which we describe in 1.12, “Advanced Security” on page 211. These Secure FOS levels may also be used for fabrics without implementing fabric security.

## Initialization

When the switch is powered on or restarted, the following operations are performed:

1. Early power-on self test (POST) diagnostics are run. POST is run before the FOS is started.
2. The FOS is initialized.

3. The hardware is initialized. The switch is reset, the internal addresses are assigned, the Ethernet port is initialized, the serial port is initialized, and the front panel is initialized.
4. A full POST is run.
5. The links are initialized. Receiver and transmitter negotiation is run to bring the connected ports online.
6. During the Fabric Login (FLOGI), link parameters are exchanged. This determines whether any ports are connected to other switches. If so, it negotiates who becomes the principal switch.
7. Domain addresses are assigned. After the principal switch is identified, port addresses are assigned. Each switch tries to keep the same domain ID that it used previously. Previous IDs are stored in the configuration Flash memory.
8. The routing table is constructed. After the addresses are assigned, the unicast routing tables are constructed.
9. Normal Nx_Port operation is enabled.

## Routing

The switches control processor maintains two routing tables, one for unicast and one for multicast. The unicast routing tables are constructed during fabric initialization. The multicast tables are initially empty, except for broadcast addresses. After the tables have been constructed they are loaded into each ASIC.

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

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

## Service functions

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

### *SNMP*

Simple Network Management Protocol (SNMP) allows network devices to be monitored, controlled, and configured remotely from a network management station running a network manager program.

SNMP agent code in the network device allows management by transferring data that is specified by a Management Information Base (MIB).

The switch agent supports the following features:

- SNMPv1 manager
- Command-line utilities to provide access to and command the agent
- MIB-II system group, interface group, and SNMP group
- Fabric-element MIB
- IBM-specific MIBs
- Standard generic traps
- IBM-specific traps

### *Diagnostics*

The switch supports a set of power-on self tests (POSTs), as well as tests that can be invoked using a command line interface. These diagnostics are used during the manufacturing process as well as for fault isolation of the product in customer installations.

### *Diagnostic environment*

Most diagnostics are written to run in the FOS environment. However, as the FOS does not run without a working SDRAM, a SDRAM/boot EEPROM test is run as part of the pre-FOS startup code to verify that the basic processor connected memories are functioning properly.

### *Hardware support*

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

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

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

***Diagnostic coverage***
The POST and diagnostic commands concentrate on the Fibre Channel ports and verify switch functionality of the switch.

## 1.2 IBM TotalStorage SAN Switch F08

The IBM TotalStorage SAN Switch F08, known as a 3534-F08, is an 8-port entry level switch for customers with a small number of servers requiring the ability to share SAN attached storage devices. The base model, which can be connected to only one other switch, allows for growth and investment protection. It uses a single auto-ranging power supply.

In Figure 1-1 we show the 3534-F08 view from the front.

*Figure 1-1 3534-F08 switch*

### 1.2.1 IBM TotalStorage SAN Switch F08 product overview

The IBM 3534-F08 is equivalent to a Brocade SilkWorm 3200, and provides the following features:

- Eight Fibre Channel ports, each port capable of full-duplex throughput at either 2 gigabits per second (200MB/s).
- Auto-sensing ports that self-negotiate to the highest speed supported by the attached server, storage, or switch.
- A single E-port for connecting to another switch.
- Each port supports the new Small Form-Factor Pluggable (SFP) media with options for either shortwave optical connection for distances up to 300 meters, or longwave optical connections for distances up to 10 kilometers.
- A 10/100 ethernet port used for out of band management of the switch.
- A serial interface used for initial configuration and diagnostics.
- A small 1U package that can be either rack-mounted or used in a table-top configuration.
- Distributed fabric services such as name serving, zoning, routing, and microcode upgrade.

- WEB TOOLS, which provides a comprehensive set of management tools that support a Web browser interface for flexible, easy-to-use operations.

Optional features activated by License key include:

- Extended fabric, Remote Switch, QuickLoop, and Security.
- Performance Bundle feature, which provides Inter-Switch Link (ISL) Trunking and Advanced Performance Monitoring.
- Full Fabric Activation, providing WEB TOOLS, Advanced zoning, and FabricWatch.

## 1.2.2 Upgrading Entry Level to Full Fabric

The IBM TotalStorage SAN Switch F08 Entry Level can be concurrently upgraded to Full Fabric capabilities by purchasing a Full-fabric license key.

The upgrade provides a cost effective and scalable approach to growing fabric based SANs.

The IBM TotalStorage SAN Switch Full-fabric activation feature offers:

- Upgrade to full-fabric switch capability on all eight ports, (F_Port, FL_Port or E_Port)
- Support for up to eight interswitch links, or E_Ports
- Advanced Zoning
- Fabric Watch

Figure 1-2 shows the layout of the different elements in the F08 faceplate.

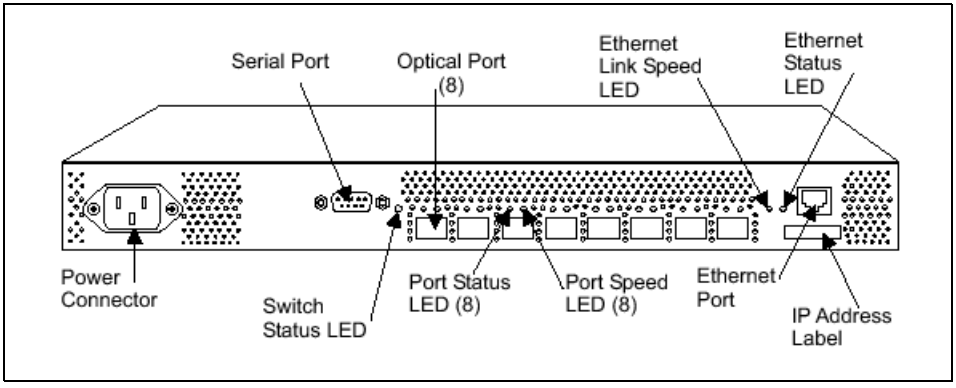


*Figure 1-2 IBM TotalStorage SAN Switch F08 faceplate*

Figure 1-3 shows the IBM TotalStorage SAN Switch F08 rear view.

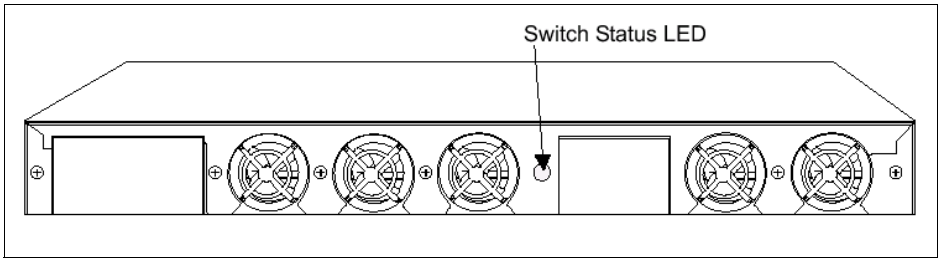


*Figure 1-3 IBM TotalStorage SAN Switch F08 back panel*

## 1.3 IBM TotalStorage SAN Switch F16

The IBM TotalStorage SAN Switch F16, also known as a 2109-F16, is a Full Fabric, sixteen 2 Gb/s Fibre Channel port switch, with an optional second concurrently replaceable power supply, and replaceable fan assembly. It provides a cost affective high availability solution for small and medium size SANs.

We show a picture of the 2109-F16 in Figure 1-4.

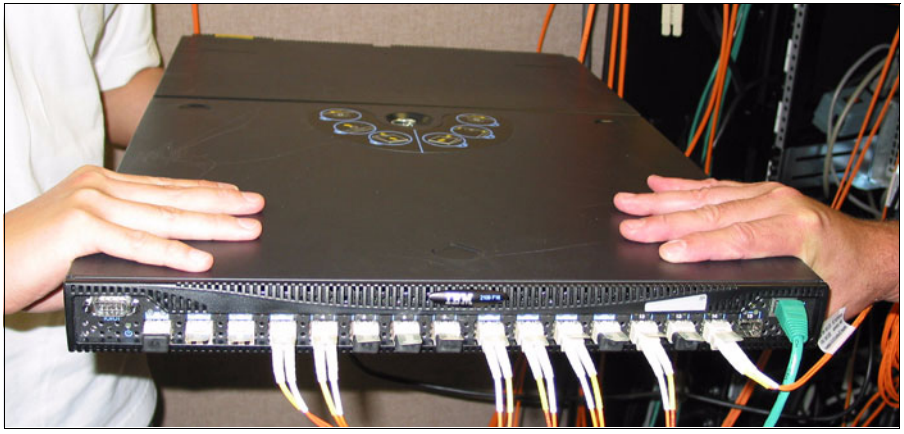

*Figure 1-4 2109-F16 switch*

### 1.3.1 IBM TotalStorage SAN Switch F16 product overview

IIBM TotalStorage SAN Switch F16 is equivalent to a Brocade SilkWorm 3800, and provides the following features:

- Sixteen Fibre Channel ports, each capable of full-duplex throughput at either 1 or 2 gigabits per second.
- All ports are auto-sensing ports that self-negotiate to the highest speed supported by the attached server, storage, or switch.
- They are all Universal ports that self-configure as F_Ports, FL_Ports, or E_Ports.
- Each Fibre Channel port uses Small Form-Factor Pluggable (SFP) media with options for either shortwave optical connection for distances up to 300 meters, or longwave optical connections for distances up to 10 kilometers.
- A 1U package that can be either rack-mounted or used in a table-top configuration, with the option of a redundant power supply, providing a highly available switch.
- Hardware zoning controlled at the port level, and at the worldwide name level.
- Cascading support for flexibility in creating scalable fabric topologies.
- Distributed fabric services such as name serving, routing, Advanced Zoning, Fabric Watch, and microcode upgrade.
- WEB TOOLS, which provides a comprehensive set of management tools that support a Web browser interface for flexible, easy-to-use operations.

Optional features activated by License key include:

- Extended fabric, Remote Switch, QuickLoop, Security.
- Performance Bundle, which provides Inter-Switch Link (ISL) Trunking and Advanced Performance Monitoring.

Figure 1-5 shows the layout of the different elements in the F16 faceplate.

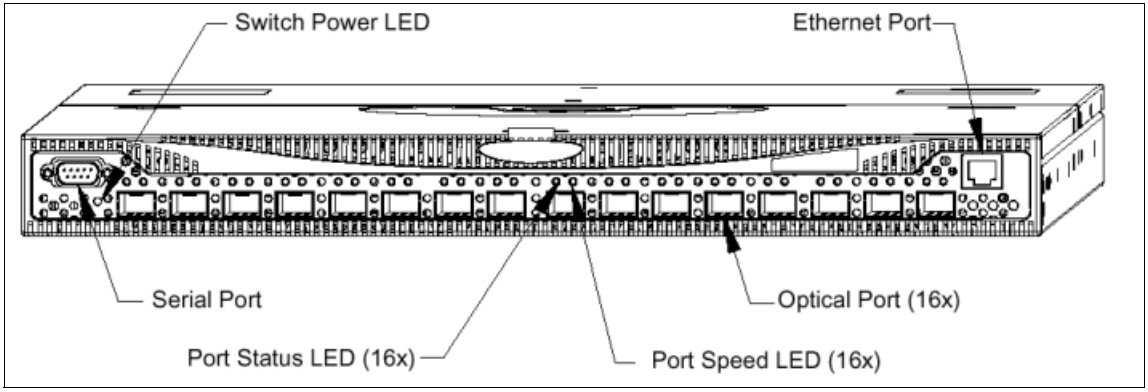


*Figure 1-5 IBM TotalStorage SAN Switch F16 faceplate*

Figure 1-6 shows the IBM TotalStorage SAN Switch F16 rear view.

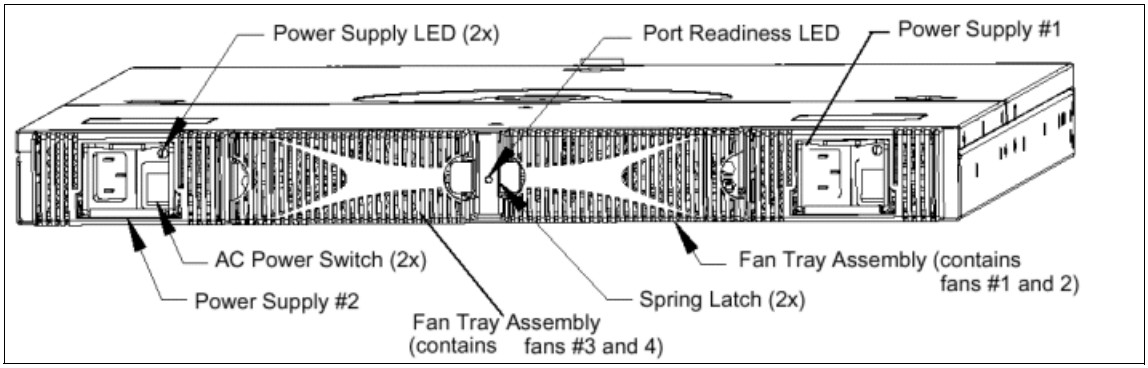


*Figure 1-6 IBM TotalStorage SAN Switch F16 back panel*

# 1.4 IBM TotalStorage SAN switch F32

The IBM TotalStorage SAN switch F32, also known as IBM 2109-F32, is a Thirty-two port, 2 Gigabits per second, Full Fabric, highly available Fibre Channel switch. With it's concurrently replaceable redundant power supplies, and concurrent firmware upgrade ability, it is a good solution for customers with growing SAN environments and limited down time opportunities.

In Figure 1-7 we show a picture of the IBM TotalStorage SAN switch F32:

*Figure 1-7 2109-F32 switch*

## 1.4.1 IBM TotalStorage SAN switch F32 product overview

IBM TotalStorage SAN switch F32 is equivalent to a Brocade SilkWorm 3900 and includes the following features:

- Thirty-two ports, each port capable of full-duplex throughput at either 1 or 2 gigabits per second.
- Auto-sensing ports that self-negotiate to the highest speed supported by the attached server, storage, or switch.

- Universal ports that self-configure as F_Ports, FL_Ports, or E_Ports.
- Each port supports the new Small Form-Factor Pluggable (SFP) media with options for either shortwave optical connection for distances up to 300 meters, or longwave optical connections for distances up to 10 kilometers.
- A 1.5U package that can be either rack-mounted or used in a table-top configuration, with a redundant power supply, providing a highly available switch.
- Hardware zoning controlled at the port level, and at the worldwide name level.
- Performance Bundle feature, which provides Inter-Switch Link (ISL) Trunking and Advanced Performance Monitoring.
- Cascading support for flexibility in creating scalable fabric topologies.
- Distributed fabric services such as name serving, Advanced Zoning, routing.
- WEB TOOLS, which provides a comprehensive set of management tools that support a Web browser interface for flexible, easy-to-use operations.
- Concurrent code activation, allowing switch firmware upgrades without the need to remove the switch from the fabric.

Optional features activated by License key include:

- Extended fabric, Remote Switch, Security.

**Attention:** QuickLoop support will not be provided with the F32 Switch. IBM TotalStorage SAN Switches with QuickLoop capability may be used to attach private loop devices in a core/edge fabric.

In Figure 1-8 we identify the various indicators and ports on the front panel of the IBM TotalStorage SAN switch F32.

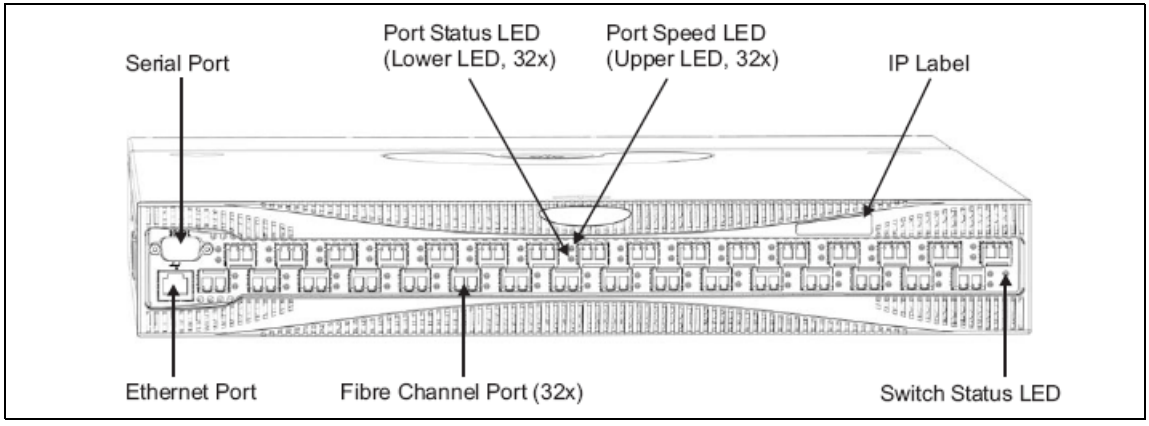


*Figure 1-8 IBM TotalStorage SAN switch F32 faceplate*

Figure 1-9 shows the rear components of the IBM TotalStorage SAN switch F32.

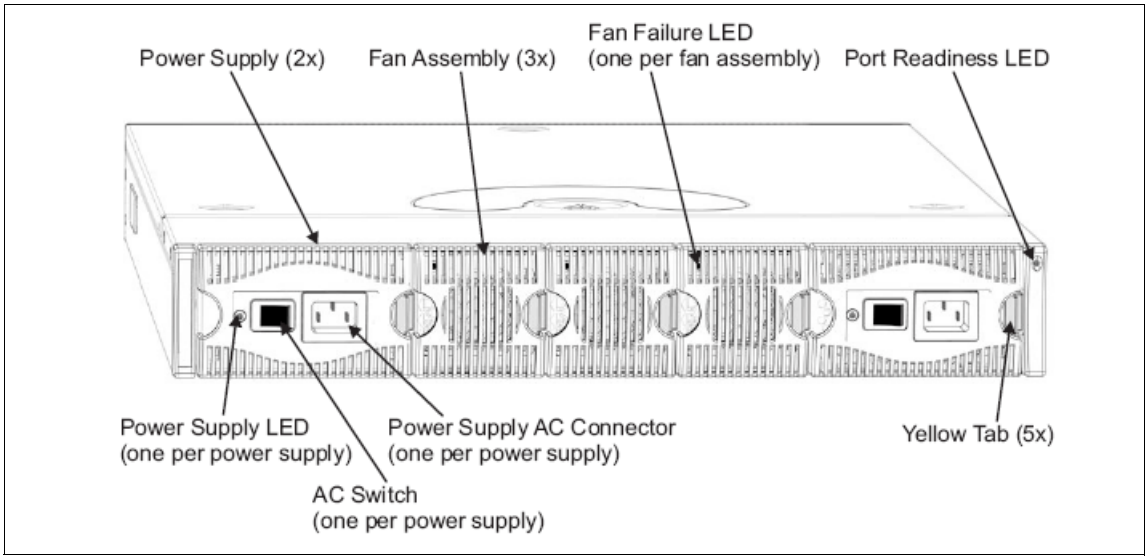


*Figure 1-9 Rear components of the IBM TotalStorage SAN switch F32*

## 1.5 IBM TotalStorage SAN Switch 2109-M12

The IBM TotalStorage SAN Switch M12 is designed to be used as the core of a fabric, with its high availability and port count eliminating the need for multiple small switches. It can be used to upgrade from smaller existing 2109 core switches, or to implement a new large Fabric, and may be used as one or two separate switches.

It is scalable, in 16 port increments, from a single 32 port switch up to 2 x 64 port switches, providing 128 ports in one chassis, utilizing autosensing 1 or 2 Gb/s full duplex SFP ports. Also providing redundant Control Processors (CP) for non-disruptive fail-over, concurrent code upgrade activation, and multiple redundant power supplies and fans, it has been designed to minimize any outage to SAN operation.

We show a picture of the 2109-M12 in Figure 1-10.


*Figure 1-10 2109-M12 switch*

## 1.5.1 IBM TotalStorage SAN Switch M12 product overview

The director class IBM TotalStorage SAN Switch M12 is based on the same ASIC switching technology used in the IBM TotalStorage SAN Switches F08, F16, and F32. The M12 core switch supports both 1 Gb/s and 2 Gb/s auto-sensing ports as well as advanced fabric services that can simplify the design, administration, and management of enterprise SANs. The M12 provides investment protection, since it is fully compatible with existing IBM SAN Switches S08, S16, F08, F16, and F32, and it will provide a high performance, scalable, flexible, function rich, and reliable core to a large SAN fabric solution.

The IBM TotalStorage SAN Switch M12 is equivalent to a Brocade silkworm 12000, and provides the following features:

- High availability with redundant / hot swappable components, FSPF rerouting around failed links, and non-disruptive firmware upgrades.
- Up to one hundred and twenty-eight non-blocking ports, each with full-duplex throughput at either 2 Gb/s per second or 1 Gb/s per second.
- The dual switch capability allows either one or two 64-port switches per chassis. The switches may be interconnected to create a high port count solution, or they can be used to create two independent fabrics.
- Auto-sensing ports that self-negotiate to the highest speed supported by the attached switch, server, or storage.
- Universal ports that self-configure as E_Ports, F_Ports, or FL_Ports.
- Each port supports the Small Form-Factor Pluggable (SFP) media with options for either shortwave optical connection for distances up to 300 meters, or longwave optical connections for distances up to 10 kilometers.
- The 14U package can be mounted in a standard 19 inch rack or purchased mounted in the IBM 2109-C36 SAN rack.

Licenses for the M12 chassis apply to both logical switches, and include:

- WEB TOOLS
- Full Fabric
- Advanced Zoning
- FabricWatch
- ISL-Trunking
- Advanced Performance Monitoring

Optional licenses:

- Fabric Manager
- Extended Fabric
- Remote Switch
- Security

**Attention:** QuickLoop support will not be provided with the M12 Switch. IBM TotalStorage SAN Switches with QuickLoop capability may be used to attach private loop devices in a core/edge fabric.

All of the licensed features mentioned above are described in detail throughout this chapter.

## 1.5.2 Hardware components

In this section we describe some of the hardware components of the IBM TotalStorage SAN Switch M12.

### Backplane

At the heart of the M12 is a backplane, which has been designed to allow for future enhancements, including 10 Gb/s throughput, and 128 port single switch. It easily provides continuous 100% utilization and non-blocking throughput on all ports at the same time. The design also allows for the hot plugging of the blade assemblies.

### CP blade assembly

The M12 contains two CP blade assemblies for redundancy. Referring to Figure 1-11, we show the CP blades installed in slots 5 and 6, providing a physical divider between the two logical switches of the chassis. The active CP provides control and management functions, including these:

- System initialization
- Switch drivers
- High availability drivers
- Name server
- System management
- Fabric OS
- Fabric Access
- Extended Fabrics
- Fabric Watch
- Remote Switch
- WEB TOOLS
- Zoning

Each CP blade has a PowerPC® 405GP 200-MHz microprocessor (PPC405) on the assembly. It contains a high-performance reduced instruction set computer (RISC) core, synchronous dynamic random access memory (SDRAM) controller, PCI bus interface, direct memory access (DMA) engine, serial ports, IC interface, read-only memory, and general purpose I/O. In addition, the CP blade assembly offers the following features:

- Hot-plugged interface circuitry to support reliability, availability, serviceability, and failover; if one CP stops functioning, the other CP automatically takes its place.
- An amber LED to indicate error status for the CP.
- A green LED to indicate the proper operation for the CP power.

### *Ethernet ports*

Each CP blade has its own 10/100 BASE-T ethernet port, giving the ability to connect remotely to the switch through your ethernet network. Each port has LEDs to indicate speed and status, and they also are assigned separate IP addresses, allowing full redundancy, and LAN management access to the offline CP.

### *Serial ports*

Each CP blade also contains 2 serial ports. The top serial port is for connection of a modem, which allows remote dial-in support to the switch. The lower serial port is for connection of an RS-232 compatible terminal port for command line interface (CLI) communication.

## Switch blade assembly

Each switch blade assembly has 16 external Fibre Channel ports that run at an auto-negotiated rate of 1 Gb/s or 2 Gb/s. They support trunking, and are universal (E_Port, F_Port, and FL_Port). Port speed can be managed through the management interface. A full chassis, shown in Figure 1-11, may consist of up to eight switch blade assemblies, providing a maximum of 128 ports; while the minimum configuration is two switch blade assemblies, providing 32 ports.

It is responsible for Fibre Channel switching circuitry, and houses the ASIC, backplane serial-deserializer (SERDES), external SERDES, and status LEDs for external SERDES such as port speed and port state, as well as the SFP fiber optic media. Each switch blade assembly is hot-pluggable, allowing installation of new blades while the switch is running to increase the port count of the switch or to replace a failed blade. This is accomplished by the high performance connectors to the backplane using long pins, short pins, or both to assure proper ground-voltage-signal sequencing. Field effect transistor (FET) switches, such as QuickSwitches, are used to isolate the PCI interfaces.

When a switch blade assembly is inserted, the power regulation circuitry inhibits the in-blade DC converter (DCC) and keeps the switch blade assembly turned off. The CP, under software control, enables the DCC, and thus turns on the switch blade assembly. When the switch blade assembly is ready, it interrupts the CP for initialization.

Each switch blade assembly has an on-board serial EEPROM that is only accessible through the IC bus interface. This serial EEPROM can be accessed by a CP to determine information, including:

- OEM serial number
- IBM serial number
- Manufacturing date
- Manufacturing location

- Part number
- Revision
- Error logs

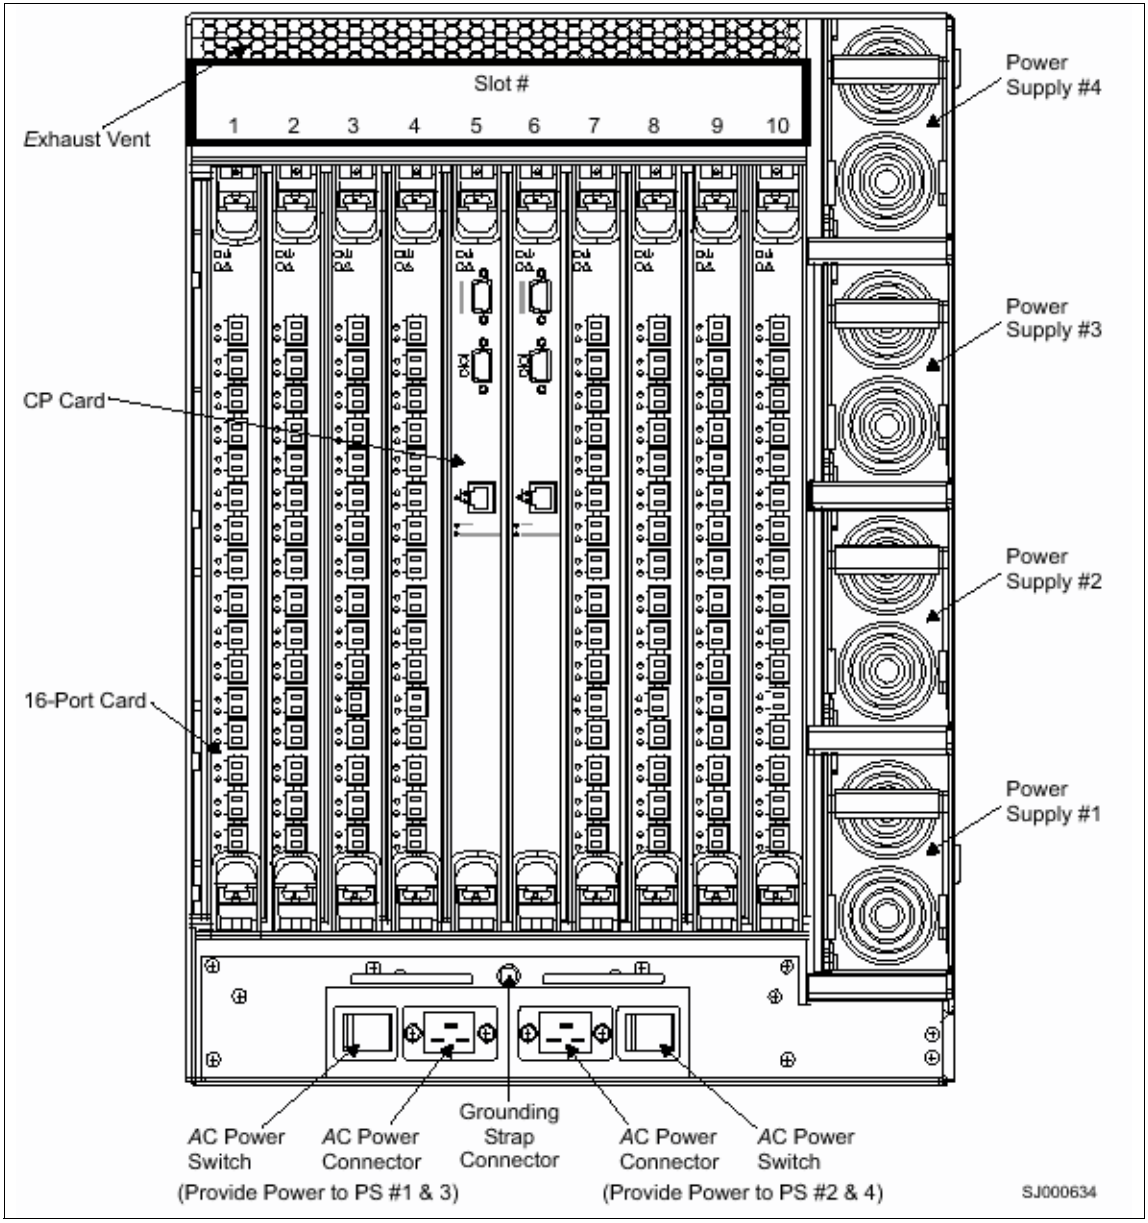


*Figure 1-11 Port side view of the M12*

## Power supplies

Looking at the port side of the M12, we can see four power supplies on the right hand side; we show this in Figure 1-11. These power supplies are split across the two AC inputs.

That is to say, power supplies 1 and 3 are fed from the left-hand AC input, and power supplies 2 and 4 are fed from the right-hand AC input.

With this power distribution configuration, a fully configured chassis is capable of continuous operation during periods of losing any two power supplies.

The power supplies are hot-pluggable for non-disruptive repair actions.

## Blower assembly side

The WWN bezel is found at the top of the blower side of the chassis. The WWN card contains status LEDs, chassis serial number, the IP addresses assigned to the CP cards, and the logical switch names, IP addresses, and WWNs for the two logical switches (switch 0 assumes the WWN of the chassis, and switch 1 is the chassis WWN + 1). The LEDs show the OK/Attention status of the port side blades and power supplies, as shown in Figure 1-12.

The WWN card is concurrently replaceable after v4.1 of firmware. At levels prior to v4.1, it is not replaceable without disrupting the chassis operation.

### *Blowers*

There are three blower assemblies in the M12; these provide the cooling to the chassis components. If a blower fails, the remaining two blowers increase their speed to continue adequate cooling. The blower assemblies are hot-pluggable for non-disruptive replacement.

Should a second blower fail, the M12 has been designed to sequence down the switch blades, so as not to overheat and damage any components. This sequence is predefined, but may also be modified to suit your configuration, and can maintain a degraded system configuration during such circumstances.

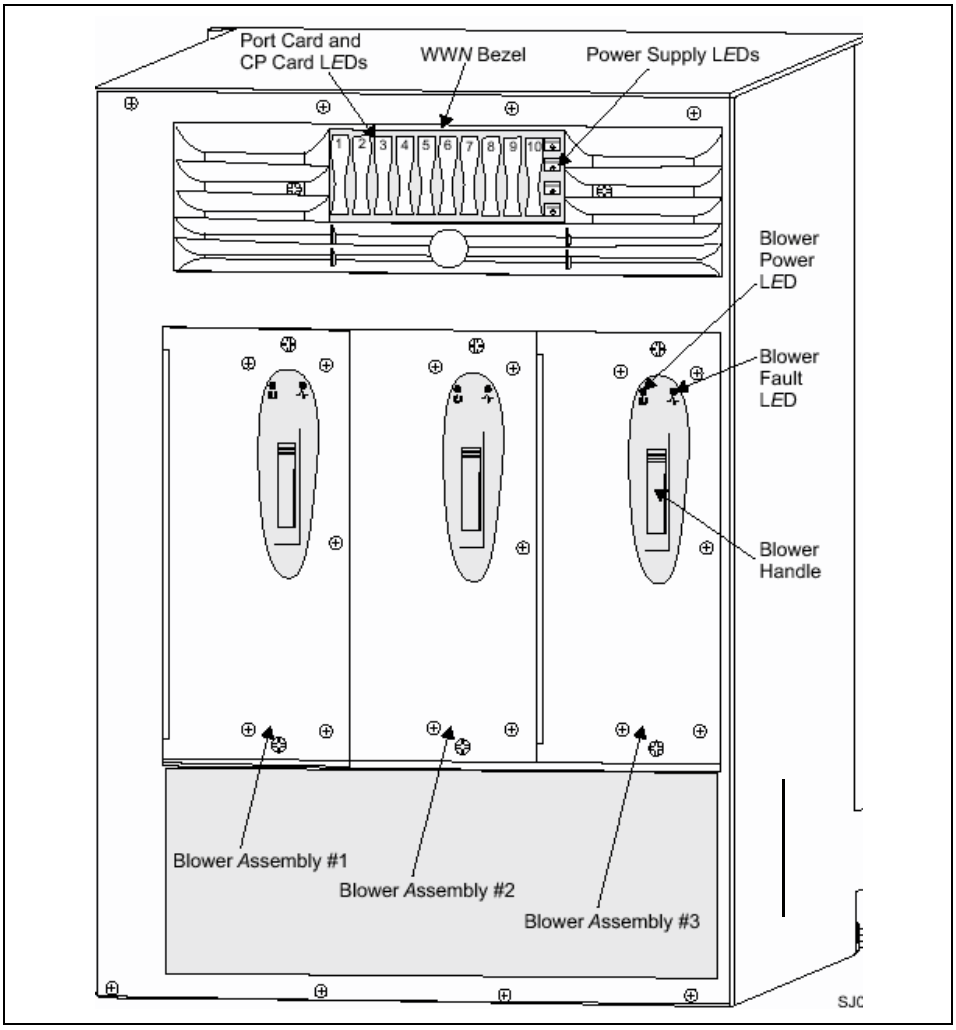


*Figure 1-12 Blower side view*

## Locations

The M12 uses a numbering scheme to reference components and ports.

Component numbering starts from left to right, and bottom to top, as shown in Figure 1-11 on page 18. Power supplies number from 1 at the bottom through 4 at the top. Slots are numbered from 1 on the left through 10 on the right.

Port blades installed in slots 1 through 4 are part of logical switch 0, while blades installed in slots 7 through 10 are part of logical switch 1. We show this in Figure 1-13.

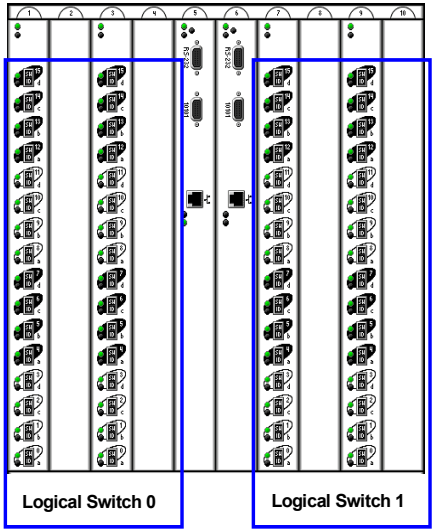


*Figure 1-13 Logical Switch layout*

Physical port numbering for each port card begins with port 0 at the bottom, and port 15 at the top of the card, as shown in Figure 1-14.

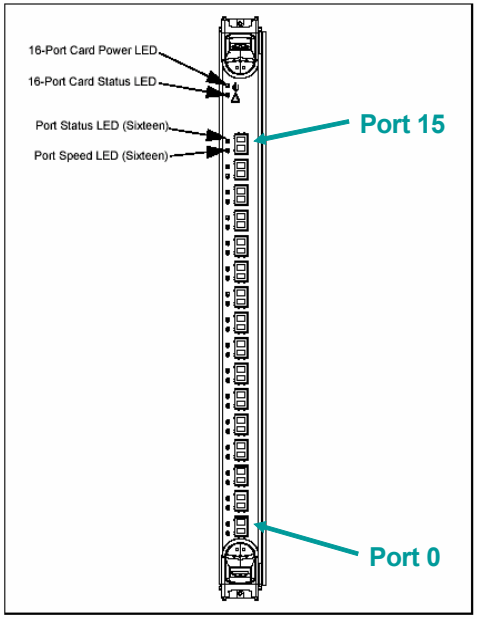


*Figure 1-14 Physical port numbering*

As each logical switch can have up to 64 ports, there is a need to number these ports from a switch perspective, not just at a blade level. This switch port numbering is known as *port area* (sometimes referred to as the absolute port number), numbering the ports for each logical switch from area 0 through 63.

Using the Command Line interface, zoning commands use the port area numbering, while other commands require the slot/port method. Refer to section “Selecting ports on the M12” on page 55 for more information on this.

At a telnet command line session, we can display the slot / port numbering in relation to the area numbering by using the **switchShow** command.

We show the output of this for our logical switch 0:

```
sw96:admin> switchshow
switchName:     sw96
switchType:     10.1
switchState:    Online
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:00:60:69:80:06:7a
switchBeacon:   OFF
blade1: Beacon: OFF
blade3: Beacon: OFF

Area Slot Port Gbic Speed State
===================================
  0    1    0   id    N2   No_Light  LE
  1    1    1   id    N2   No_Light
  2    1    2   id    N2   No_Light
  3    1    3   id    N2   No_Light
  4    1    4   id    N2   No_Light
  5    1    5   id    N2   No_Light
  6    1    6   id    N2   No_Light
  7    1    7   id    N2   No_Light
  8    1    8   id    N2   No_Light
  9    1    9   id    N2   No_Light
 10    1   10   id    N2   No_Light
 11    1   11   id    N2   No_Light
 12    1   12   id    N2   No_Light
 13    1   13   id    N2   No_Light
 14    1   14   id    N2   No_Light
 15    1   15   id    N2   No_Light
 32    3    0   id    N2   No_Light
 33    3    1   id    N2   No_Light
 34    3    2   id    N2   No_Light
 35    3    3   id    N2   No_Light
 36    3    4   id    N2   No_Light
```

```
37    3    5   id    N2   No_Light
38    3    6   id    N2   No_Light
39    3    7   id    N2   No_Light
40    3    8   id    N2   No_Light
41    3    9   id    N2   No_Light
42    3   10   id    N2   No_Light
43    3   11   id    N2   No_Light
44    3   12   id    N2   No_Light
45    3   13   id    N2   No_Light
46    3   14   id    N2   No_Light
47    3   15   id    N2   No_Light
```

The table presented in Figure 1-15 shows the *area* number for each physical port location. The Physical Slot number refers to logical switch 0 / logical switch 1 slot position.

The Logical Slot numbering is only used to help define the FC address.

| Logical Slot 0 | | Logical Slot 1 | | Logical Slot 2 | | Logical Slot 3 | |
|---|---|---|---|---|---|---|---|
| Physical Slot 1/7 | | Physical Slot 2/8 | | Physical Slot 3/9 | | Physical Slot 4/10 | |
| Area # | Physical Port | Area # | Physical Port | Area # | Physical Port | Area # | Physical Port |
| 15 | 15 | 31 | 15 | 47 | 15 | 63 | 15 |
| 14 | 14 | 30 | 14 | 46 | 14 | 62 | 14 |
| 13 | 13 | 29 | 13 | 45 | 13 | 61 | 13 |
| 12 | 12 | 28 | 12 | 44 | 12 | 60 | 12 |
| 11 | 11 | 27 | 11 | 43 | 11 | 59 | 11 |
| 10 | 10 | 26 | 10 | 42 | 10 | 58 | 10 |
| 9 | 9 | 25 | 9 | 41 | 9 | 57 | 9 |
| 8 | 8 | 24 | 8 | 40 | 8 | 56 | 8 |
| 7 | 7 | 23 | 7 | 39 | 7 | 55 | 7 |
| 6 | 6 | 22 | 6 | 38 | 6 | 54 | 6 |
| 5 | 5 | 21 | 5 | 37 | 5 | 53 | 5 |
| 4 | 4 | 20 | 4 | 36 | 4 | 52 | 4 |
| 3 | 3 | 19 | 3 | 35 | 3 | 51 | 3 |
| 2 | 2 | 18 | 2 | 34 | 2 | 50 | 2 |
| 1 | 1 | 17 | 1 | 33 | 1 | 49 | 1 |
| 0 | 0 | 16 | 0 | 32 | 0 | 48 | 0 |

*Figure 1-15 Physical port location to area numbering cross reference*

## 1.6 Installing the IBM TotalStorage SAN Switch

The first step to install an IBM TotalStorage SAN Switch is the physical mounting and connection to electrical outlets. This is the customer's responsibility, although the IBM TotalStorage SAN Switch 2109-M12 can be purchased separately or pre-installed in the 2109-C36 SAN rack, and it is the responsibility of the IBM service representative to physically install the chassis or rack in the location you have planned.

Once the switch is installed and powered on, it will require some initial configuration parameters to be set. A service may be purchased from IBM to perform these steps:

- **IP addresses:** To access the management interfaces of a switch from a remote workstation on a network, we need to set the IP address, subnetmask, and gateway address for the switch, or for each of the logical switches in an M12. These settings can be modified using the **`ipAddrSet`** command.

  We show the steps to perform this in "Setting the IP address using the serial port" on page 25.

  The default IP address and subnet mask for the F08, F16 and F32 switches is as follows:

  – 10.77.77.77 255.255.255.0

  The default logical IP addresses, subnet mask and switch names for an M12 are as follows. These IP addresses correspond to "sw0", (slots 1-4), and "sw1", (slots 7-10):

  – 10.77.77.77 255.255.255.0 sw0
  – 10.77.77.76 255.255.255.0 sw1

  An M12 also has native IP addresses to access each CP card. The default native IP addresses, subnet masks, and hostnames are as follows:

  – 10.77.77.75 255.255.255.0 CP0 (the CP Card in slot 5 at the time of configuration)

  – 10.77.77.74 255.255.255.0 CP1 (the CP Card in slot 6 at the time of configuration)

- **Domain ID:** For switches to be connected together within a fabric, each switch must have different domain ID's. The default domain ID for a switch is 1. In an M12 the domain ID for the logical switch in slots 1 through 4, and the logical switch in slots 7 through 10, are both 1 by default. If two switches are connected via ISL after initialization is complete, they will segment due to both switches having the same domain ID. Domain IDs can be modified using the **`configure`** command. We show an example of how to do this in "Connecting to the switch" on page 35.

- **Switch names:** Setting a switch name to identify different switches within a site is recommended. This is very helpful in identifying a switch easily that you are connected to. By using the **`switchname`** command you can assign your own switch names that can be up to 15 characters long, must begin with an alpha character, and can include alpha, numeric, and underscore characters.

Following are the steps we took to configure the above settings and connect our switch for use in a network and fabric. We also include the extra steps required to configure a 2109-M12.

The time required to accomplish this is approximately 15 minutes. The items required are:

- 2109 physically installed and connected to a power source
- Workstation that has a terminal emulator application (we used HyperTerminal)
- Serial cable provided with the switch, for connecting the switch to the workstation
- An unused IP address (2109-M12 requires four IP addresses)
- Ethernet cable for connecting the switch to the workstation or to a network containing the workstation
- SWL or LWL SFPs and fiber optic cables as required

## 1.6.1 Setting the IP address using the serial port

As IBM TotalStorage SAN Switches ship with a default IP address, it is possible to perform initial configuration using a telnet connection and the 10.77.77.77 address. However, we recommend not to connect the switch to your LAN until IP settings are properly configured and will not conflict with any other devices in your network.

Following are the steps we used to set the IP address using the serial port on an IBM TotalStorage SAN Switch F16. The procedure is the same for the IBM TotalStorage SAN Switch F08 and IBM TotalStorage SAN switch F32, although for the IBM TotalStorage SAN Switch M12, the procedure is different. We show the steps for an M12 in “M12 configuration procedure” on page 29.

1. Remove the shipping plug from the serial port and insert the serial cable provided with the switch.
2. Connect the other end of the serial cable to an RS-232 serial port on the workstation.

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

3. Verify that the switch is on and initialization has completed; refer to “Initialization” on page 4.
4. Disable any serial communication programs running on the workstation, such as PDA synchronization.
5. Open a terminal emulator application (such as HyperTerminal on a PC, or TERM in a UNIX® environment), and configure as follows:
   a. In a Microsoft® Windows® environment, adjust the following parameters and values if necessary:
      - Bits per second: 9600
      - Databits: 8
      - Parity: None
      - Stop bits: 1
      - Flow control: None

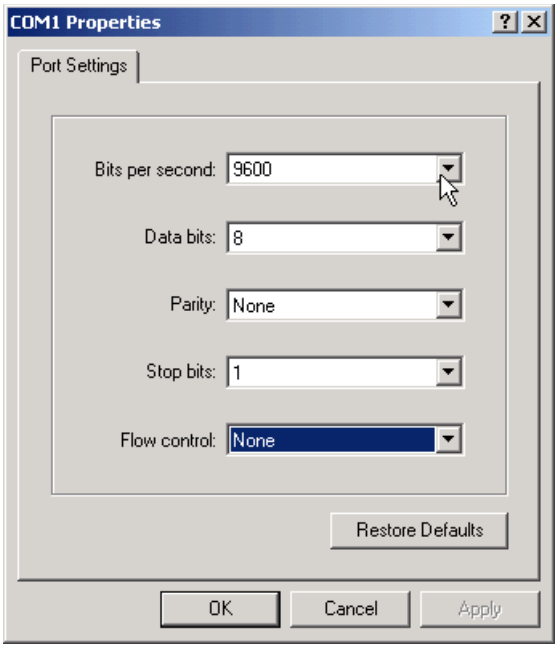


*Figure 1-16 HyperTerm COM1 properties window*

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

   **tip /dev/ttyb -9600**

6. From the terminal emulator application, log on to the switch through the serial connection. The default administrative logon is `admin` and the default password is `password`.
7. Enter the following at the prompt:

   `ipAddrSet`

8. Enter the following information at the corresponding prompts, listed below:

   **Ethernet IP Address [10.77.77.77]:** *Enter new ethernet IP address*
   **Ethernet Subnetmask [255.255.254.0]:** *Enter new ethernet subnetmask*
   **Fibre Channel IP Address [none]:** *Enter new Fibre Channel IP address if desired*
   **Fibre Channel Subnet Mask [none]:** *Enter new Fibre Channel subnet mask if desired*
   **Gateway Address [none]:** *Enter new gateway address*
   **Set IP address now? [y = set now, n = next reboot]:** *Enter "y" to set now*

9. We can verify that the address was correctly set by entering the following: **ipAddrShow**
10. Once the IP address is verified as correct, remove the serial cable, and replace the shipping plug in the serial port.

    **Note:** The serial port is intended only for use during the initial setting of the IP address and for service purposes. Using the serial port during normal switch operation or for regular maintenance is not recommended.

11. Record the IP address for future reference.

In Figure 1-17, we show how the foregoing steps are performed.

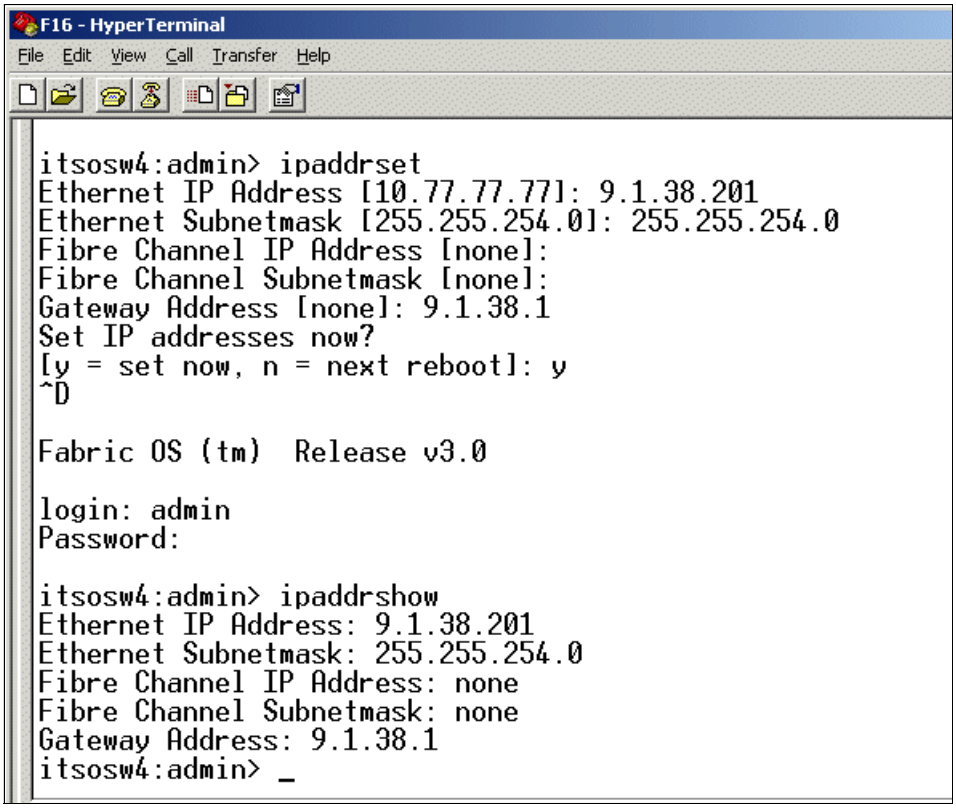


*Figure 1-17 Setting the Ethernet IP address for the 2109-F16*

Once the IP address is set, we are able to connect the switch to the workstation computer by ethernet cable (this can be a direct cross-over connection or through a network) by following these steps:

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

   **Note:** The switch can now be accessed remotely, through Telnet or WEB TOOLS. As a result, it is important to ensure that the switch is not being modified simultaneously from any other connections during the remaining steps.

4. Continue with “Connecting to the switch” on page 35.

## M12 configuration procedure

After verifying that the M12 is on and POST tests have completed, we log on to the CP Card installed in slot 5 by establishing a serial connection from our workstation that has a terminal emulator application (such as HyperTerminal for Windows, or TERM in a UNIX environment).

**Tip:** Disable any serial communication programs running on the workstation (such as synchronization programs for a PDA).

1. Remove the protective shipping cap from the terminal serial port on the CP Card in slot 5, and insert the serial cable. The terminal serial port is the second serial port from the top of the CP Card, shown in Figure 1-11 on page 18.
2. Open your terminal emulator application and configure as described below.

   For Windows, we must set our terminal emulator for the following parameters:

   ```
   Bits per second:            9600
   Databits:                   8
   Parity:                     None
   Stop bits:                  1
   Flow control:               None
   ```

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

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

   When the terminal emulator application stops reporting information, press Enter to get the following prompt:

   At the following prompt, we enter *0* to log in to switch 0:

   ```
   Enter switch number to login <0 or 1>:  0
   ```

3. We enter the administrative logon *admin* and the password of *password*.

   ```
   CP0 Console Login:
   ```

   At the initial login we are prompted to enter new Admin and User passwords. The same administrative account applies to both logical switches. If the passwords are changed on switch 0, they are automatically changed on switch 1.

   We bypassed modifying the password, by pressing CTRL-C.

4. When we arrive at the Command prompt we need to determine which CP Card is active by using the **haShow** command:

   ```
   switch:admin> haShow
   Local CP (Slot 5, CP0): Active
   Remote CP (Slot 6, CP1): Standby
   HA Enabled, Heartbeat Up
   ```

We can see that CP0 is active and must use this CP to perform the IP configuration for both CP cards.

**Tip:** If the CP Card in slot 5 is not the active CP Card, disconnect the serial cable from the CP Card, connect it to the CP Card in slot 6, and log on again.

5. First we set the IP addresses for the CP cards:
   a. Using the **ipAddrSet** command at the prompt, and entering *2* for the CP Card in slot 5, or *3* for the active CP Card in slot 6. For example:

      ```
      switch:admin> ipAddrSet 2
      ```

      Entering **ipAddrSet** alone will prompt for the switch or CP number:

      ```
      Switch number [0 for switch0, 1 for switch1, 2 for CP0, 3 for CP1]: 2
      ```

   b. Enter the requested information at the prompts, as shown below (the current information is shown in square brackets):

      ```
      Ethernet IP Address [10.77.77.75]: 9.43.236.107
      Ethernet Subnetmask [255.0.0.0]: 255.255.254.0
      Host Name [switch0]:
      Gateway Address [0.0.0.0]: 9.43.236.1
      ```

**Important:** The same gateway address must be used for both CP Cards (these gateway addresses are referenced for the logical IP addresses).

      The native IP address of the active CP Card is updated immediately. The native IP address of the standby CP Card is updated at the next reboot.

   c. Repeat Step 5a-b for the other CP Card by issuing **ipAddrSet 3** command.
6. Next we configure the two logical switch IP addresses:
   a. To configure the first logical switch IP address, we enter **ipAddrSet 0** at the prompt and enter the requested information:

      ```
      switch:admin> ipAddrSet 0
      Ethernet IP Address [10.77.77.77]: 9.43.236.96
      Ethernet Subnetmask [0.0.0.0]: 255.255.254.0
      Fibre Channel IP Address [none]:
      Fibre Channel Subnet Mask [none]:
      ```

      The logical IP address is updated immediately.

   b. Next we configure the second logical switch IP address, by entering **ipAddrSet 1** at the prompt:

      Enter the requested information for this IP address at the prompts, as described in Step 6a.
7. Reboot the switches by entering **reboot** at the prompt:

```
switch:admin> reboot
```

**Important:** If the reboot command is issued from the active CP, it will reboot the entire cabinet (both switches). If reboot is issued while connected to the standby CP, only the standby CP will reboot.

After monitoring the messages during reboot for any errors, we remove the serial cable and replace the protective dust cap on the port.

Now connection of the ethernet interfaces to our corporate LAN enables us to manage the switches by using Fabric Manager, WEB TOOLS, or telnet. To change any CP setting, we are only able to perform this by using the Serial connection, or Telnet to the Active CPs assigned IP address.

**Tip:** The M12 supports a maximum of two telnet sessions with administrative privileges at the same time.

8. Now by using telnet to a logical switch IP address, we are able to change the switch name by using the **switchName** command, this will cause a domain address format RSCN to be issued to the fabric. We recommend not changing the switch name unless necessary for your installation; we chose to change ours to *sw96* as shown:

   ```
   switch:admin> switchName "sw96"
   ```

9. As we changed the slot1-4 switch name (sw0), we also need to change the name of the logical switch in slots 7-10 (sw1). To do this we need to telnet to the address of that switch and use the **switchName** command again to change the name to *sw97* in our case, as shown in Figure 1-18.

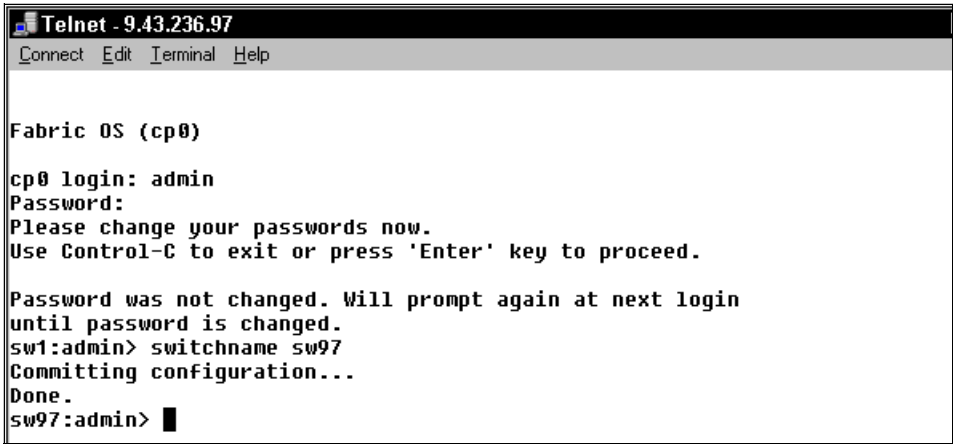


*Figure 1-18 Telnet login to Logical switch 1 (slots 7-10)*

**Attention:** Telnet to the active CP, the logical switch in slots 1-4 (sw0), or the logical switch in slots 7-10 (sw1), all look the same at initial logon. Telnet to sw1 will be different with the command prompt showing its switch name, although telnet to the active CP or sw0 will display the same logical switch name at the prompt. Care must be taken as some commands will not be available between CP sessions and logical switch sessions, and some commands may have very different results from what was expected.

10. Next we can change the default domain ID for both switches if required. Both logical switches are set to ID "1" from the factory, so to prevent a domain ID conflict, we make the domain IDs unique before connecting the switches to a fabric. A list of current domain IDs is available by using the **`fabricShow`** command. To change the domain ID from telnet, we perform the following steps, and this is also shown in Figure 1-19.
    a. Disable the switch using the **`switchDisable`** command.
    b. Enter the **`configure`** command.
    c. Reply **`y`** to the configure Fabric parameters prompt.
    d. Enter the new ID **`2`** at the Domain prompt.
    e. Complete the remaining prompts (or press CTRL+d to accept the other settings and exit).
    f. Enter **`switchEnable`** to re-enable the switch.
    g. Perform steps 10a through 10f for the other switch if required.

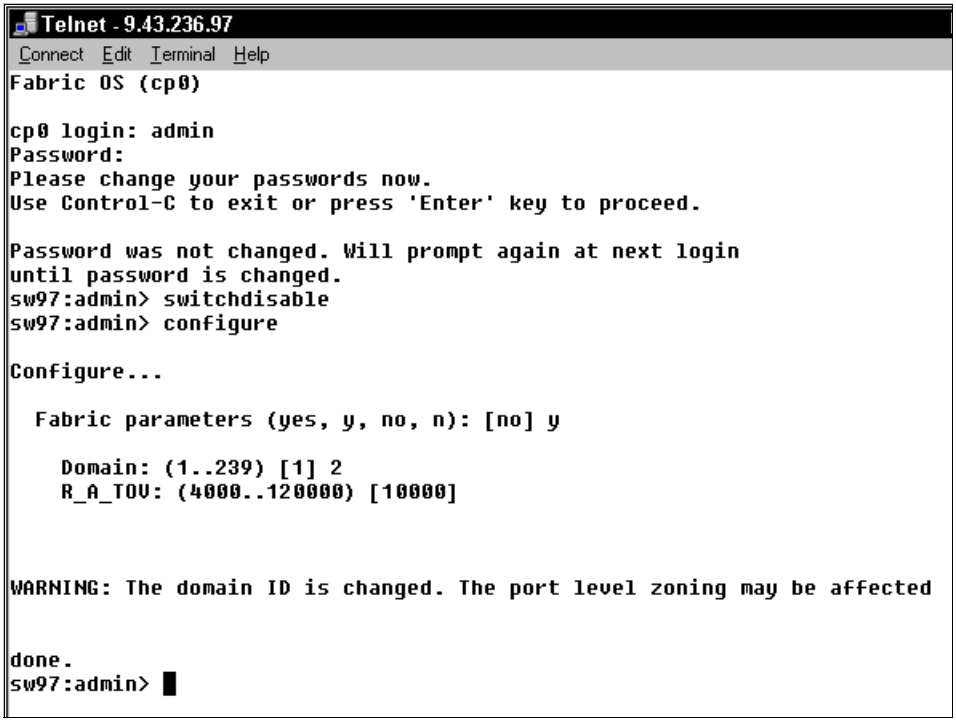


*Figure 1-19 Configuring Domain ID from Telnet*

## Optional modem setup

The M12 has the ability for modems to be connected to each CP card, to allow a redundant remote dial up facility into the switches for remote support. The modems are not supplied with the switch, and Hayes compatible modems must be purchased separately if you wish to use this facility.

As the modems can only be detected at power on of the chassis, or CP failover, we recommend connecting them now.

One modem is attached to each CP card. The active CP will answer an incoming call after one ring. If for some reason the Active CP is unable to accept the call, the Standby CP will answer the call after 7 rings.

For this to occur, both modems need to share the same telephone line, as this sharing also allows a single number to be used to dial into the chassis no matter which CP is active. See Figure 1-20.

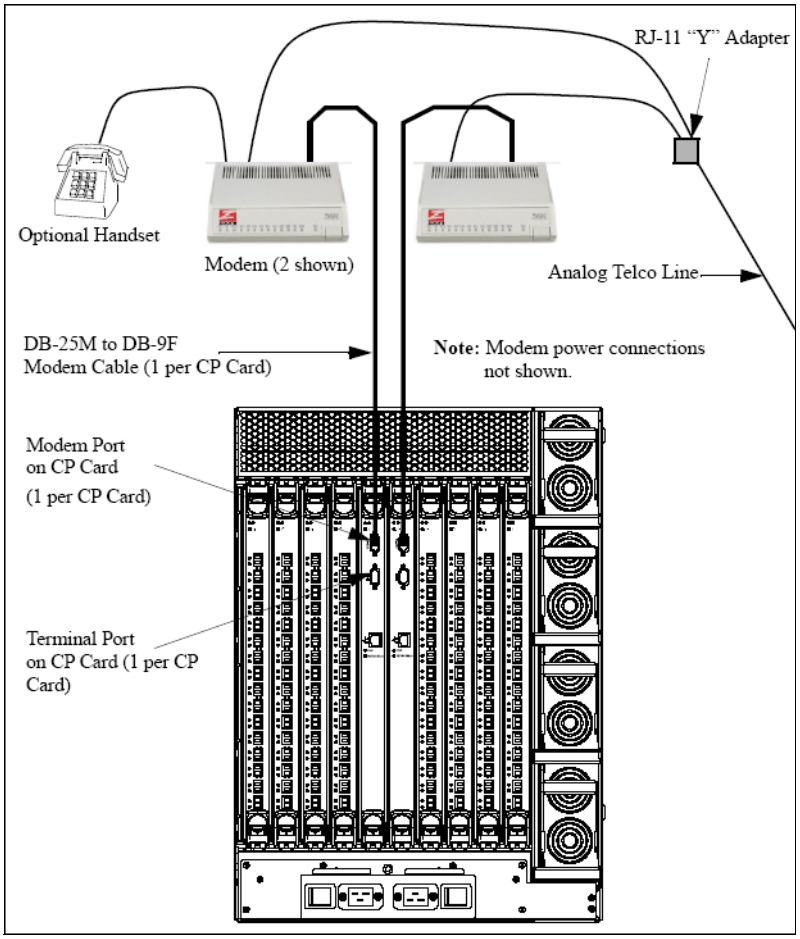


*Figure 1-20 Optional modem line and data connections*

## Remote connection settings

For Windows, we must set our terminal emulator for the following parameters:

| Parameter | Value |
|---|---|
| Port Speed (bits per second) | 115200 |
| Data Protocol | Standard EC |
| Compression | Enabled |
| Flow control | Hardware |
| Data bits | 8 |
| Parity | None |
| Stop bits | 1 |
| Modulation | Standard |

Now, entering the telephone number into your Emulator program should dial and successfully connect to the active CP allowing the command line interface to be used.

## 1.6.2 Connecting to the switch

Once the switch IP address is set, we need to set the operating parameters (for example, domain ID) and insert SFPs before connecting it to the fabric.

We perform the following steps to prepare for fabric connection:

1. Log on to the switch by using Telnet. The default administrative logon is *admin* and the default password is *password*.
2. Modify the domain IDs if desired.

   **Note:** The default domain ID is 1. If both switches are powered on, and the domain ID is already in use when the switch is connected to the fabric, the fabric will segment. If the new switch is connected and then powered on, the domain ID for the new switch is automatically negotiated to a unique value. The domain IDs currently in use can be determined by issuing the command **fabricShow.**

   a. Disable the switch by entering:

      ```
      switchDisable
      ```

   b. Enter the following command:

      ```
      configure
      ```

   c. Enter “y” after the “Fabric parameters” prompt:

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

   d. Enter a unique domain ID (such as the domain ID used by the previous switch, if still available).

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

   e. Complete the remaining prompts (or press CTRL+D to accept the remaining settings without completing all the prompts).

   f. Re-enable the switch by entering:

      ```
      switchEnable
      ```

3. We set the switchname to itsosw4 by entering:

   ```
   switchname “itsosw4”
   ```

   Entering switchname without any parameter will display the current name.

4. An optional step is to specify any custom status policies for the fabric:
   a. Enter the following at the prompt:

      ```
      switchStatusPolicySet
      ```

   b. Specify the desired status policies. To completely deactivate the alarm for a particular condition, enter “0” at the prompt for that condition.
5. Add SFPs and fiber optic cables to the ports as required.

   **Note:** The ports and cables used in trunking groups must meet specific requirements.

   a. Remove the shipping plug from the ports to be used.
   b. Position the SFP so that the key (the tab near the cable-end of the SFP) is on top, and insert the SFP into the port until it is firmly seated and the latching mechanism makes a clicking sound. For specific instructions, refer to the SFP manufacturer’s documentation.

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

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

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

6. Verify the correct operation of the switch.
   a. Enter the following at the Telnet prompt:

      ```
      switchShow
      ```

   **Note:** This command provides information about the status of the switch and the ports. Backing up the configuration after any initial configuration changes, and periodically thereafter, is strongly recommended. This ensures that a complete configuration is available if ever required for uploading to a replacement switch.

## 1.6.3 Setting Core PID format

The Core PID format parameter is a fabric wide parameter that needs to be set in switches with 8 or 16 ports for port addressing compatibility with the 2109-F32 and 2109-M12 switches. As the change to set this parameter is disruptive to switch and fabric operation, we recommend setting this parameter during fabric install, to make adding an F32 or M12, at a later date, be of minimal impact.

**Important:** The Core PID format must be set on ALL switches with 16 ports or less in a fabric if your SAN includes or will include a 2109-F32 or 2109-M12. By setting it without an F32 or M12 present, we are preparing our fabric for a future capacity upgrade with minimal disruption. The Core PID option requires a minimum firmware level of v2.6.0c (for 1RU, S08, and S16 switches) and v3.0.2c (for F08 and F16).

To set Core PID format, open a telnet session to the switch.

First we will check if the Core PID parameter has been set by issuing the **configShow "fabric"** command:

```
itsosw4:admin> configshow "fabric"
fabric.domain:  4
fabric.ops.BBCredit:    16
fabric.ops.E_D_TOV:     2000
fabric.ops.R_A_TOV:     10000
fabric.ops.dataFieldSize:       2112
fabric.ops.max_hops:    7
fabric.ops.mode.SeqSwitching:   0
fabric.ops.mode.fcpProbeDisable:        0
fabric.ops.mode.isolate:        0
fabric.ops.mode.longDistance:   0
fabric.ops.mode.noClassF:       0
fabric.ops.mode.pidFormat:      0
fabric.ops.mode.sync:   0
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
lines deleted for clarity
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

Type <CR> to continue, Q<CR> to stop:
```

To set the Core PID follow these steps:

1. Disable the switch with the switchDisable command:
   **switchdisable**
2. Run the `configure` command:
   **configure**

3. The command prompts you to set Fabric Parameters. Type **y**:

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

4. Press Enter to use default parameters for settings until you are prompted for the Core PID format setting. Set the parameter to **1**.

   ```
   Core Switch PID Format: (0..1) [0] 1
   ```

5. Continue to press Enter to skip other settings. You should get the following message:

   ```
   Committing configuration...done.
   ```

6. Enable the switch:

   ```
   switchenable
   ```

7. Fastboot the switch:

   ```
   fastboot
   ```

## 1.6.4 Setting the date

Now is also a good opportunity to set the date and time in the switch. Although a switch with incorrect date and time will function properly, it is used for time stamping during logging of events. We suggest setting these parameters prior to any further operations.

We do this by using the **`date "MMDDhhmmYY"`** command, where MM = Month, DD = Day, hh = hour, mm = minutes, YY = Year. Setting the Time and Date in one switch will also set the other switch. An example of this is shown in Figure 1-21.


*Figure 1-21 Setting the time and date with telnet*

The firmware is year 2000 compliant. Year values greater than 69 are interpreted as 1970 - 1999; year values less than 70 are interpreted as 2000 - 2069.

We have now completed the steps to install, although we would recommend upgrading to the latest level of firmware available at this time before making the switch available for use.

Refer to 1.9, “Upgrading switch firmware” on page 182 to perform this step. Also, if the switch just installed it to become part of another fabric, we cover the topics of “Merging SAN fabrics” on page 176, and “Migrating the M12 into a core fabric” on page 209.

# 1.7 Management

The WEB TOOLS interface is a Web browser based interface that assists with the installation, setup, and management of the IBM TotalStorage SAN Switch.

## 1.7.1 Launching WEB TOOLS

Access to the WEB TOOLS interface is provided by using a Java™-enabled Web browser. The following are the minimum levels required for some popular browsers:

- For Microsoft Windows:
  - Internet Explorer 5.5(SP2®) or above
- For UNIX:
  - Netscape 4.77 or above

In addition to the foregoing, Java Plug-In 1.3.1_04 is recommended.

***To launch***

1. Start the Web browser, if it is not already active.
2. Enter the switch name or IP address in the Location/Address field.

**Tip:** When managing a multi switch fabric, it is recommended to enter the switch name or IP address of the switch with the largest port count, and the highest firmware level.

3. A *Fabric View* appears in the left column, displaying all compatible switches in the fabric. Also a *Switch View* and detail of the switch we targeted with the IP address displays in the larger area on the right side of the browser.

Figure 1-22 shows the WEB TOOLS view window for a single switch fabric.

**Note:** The WEB TOOLS display is very different after v3.1 or v4.1 firmware. We show the earlier view later on, in 1.9, “Upgrading switch firmware” on page 182.

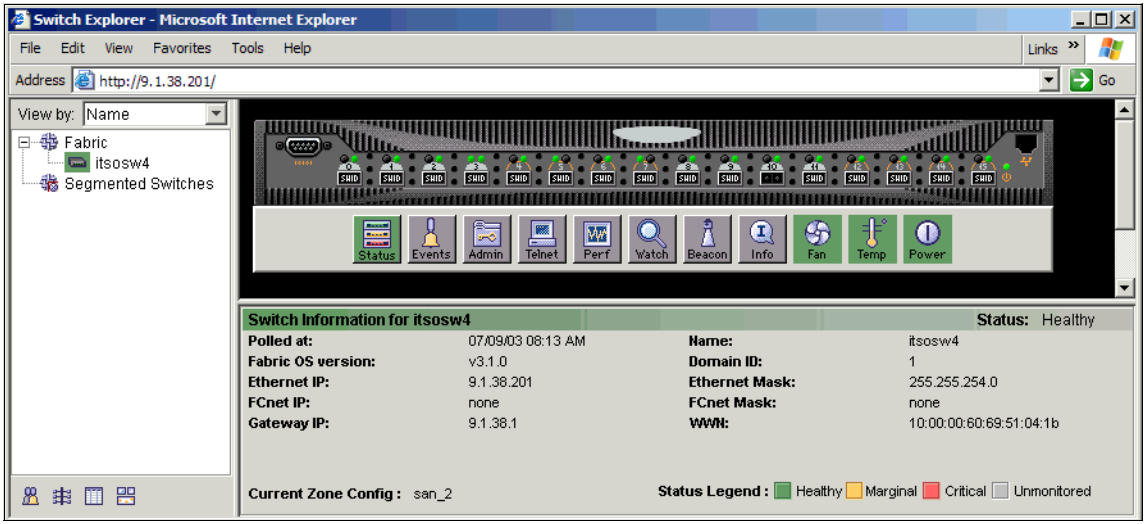


*Figure 1-22 WEB TOOLS — Single Switch Fabric View*

In Figure 1-23 we show a multi-switch fabric where we have selected to view the 2109-F32 switch view from the fabric view list in the left column. In this figure we have not defined anything to see more switches; it is a feature of WEB TOOLS that it will display all interconnected switches within a fabric.

It can also be noted that while we have pointed our browser to the IP address of the M12 in the fabric, the switch view is the one we have selected from the fabric view, and not the one to which the browser is pointed.

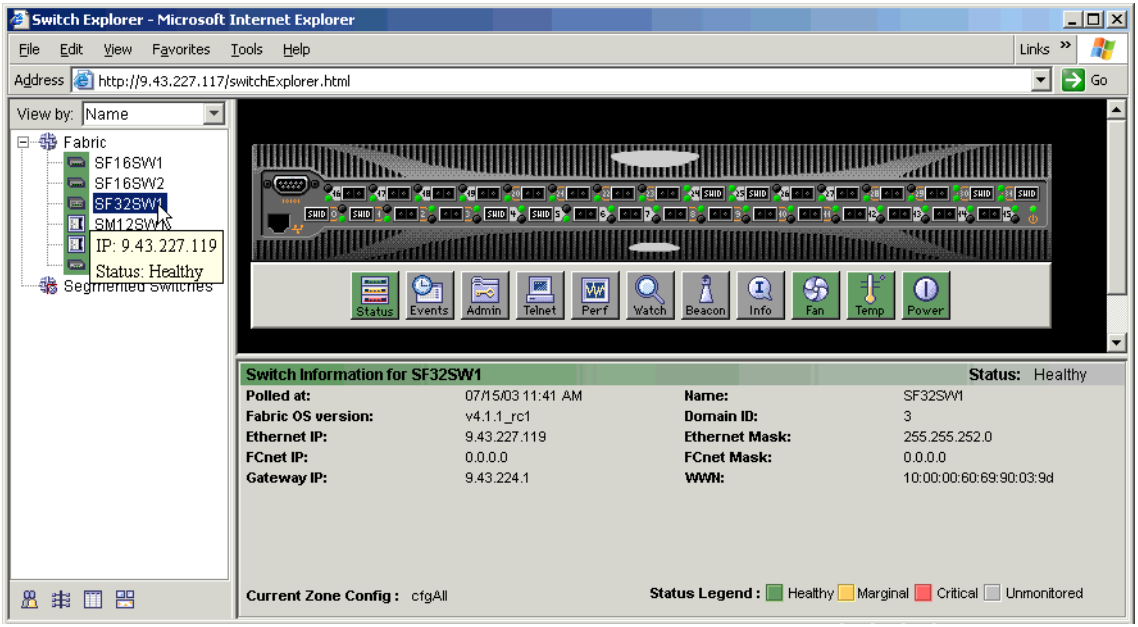


*Figure 1-23 2109-F32 Selected from fabric view*

There are three main components (frames) of the Fabric View window. On the left-hand side is the fabric management frame, also shown in Figure 1-24, which includes a list of all the switches in the fabric. At the bottom of the frame are buttons for opening separate fabric events, topology, nameserver, and zoning windows.

The larger two frames display the switch view and information view of the switch IP address we pointed our Web browser to. After the initial browser connection to a switch within the fabric, we can select other switch views by clicking the desired switch within the fabric frame.

In our example shown in Figure 1-23 we selected the switch named *SF32SW1* in the left column, our right upper frame shows the switch view for the selected switch, and the lower frame displays its details such as switch name, domain ID, IP addressing information, firmware level, switch WWN, and also the name of the fabrics active zoning configuration. A legend is also included at the bottom to explain the colors used to indicate the status of different components within the fabric. In our view all switches are shown as green and therefore all are in a Healthy state.

In the following sections we describe the WEB TOOLS features in more detail.

## 1.7.2 WEB TOOLS Fabric View

The fabric frame shown in Figure 1-24 lists all the switches that are connected to form the SAN fabric. This list will show all the switches that are managed by the WEB TOOLS interface and can include switches from back-level versions of code, managed hubs, and switches, with or without licenses for Fabric Watch. Displayed in our fabric are six switches, and we have chosen to view them by name, although we also have the option to view by WWN or IP address.

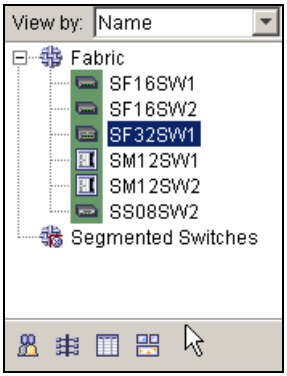


*Figure 1-24 Fabric View*

Apart from the thumbnail of each switch in the fabric, also displayed at the bottom are four command buttons. These command buttons launch fabric wide services which we describe in the following sections. Switch specific functions launched from within the Switch View frame will be covered in 1.7.5, “WEB TOOLS Switch View” on page 70.

## Fabric Events

Fabric Events is a log of all the events that have occurred across the fabric. The Fabric Watch conditions will be logged as well as other Fabric-wide events.

From the bottom of the Fabric View frame, we can select the button as shown in Figure 1-25.

*Figure 1-25 Fabric Events Button*

This will display the Fabric Event log as shown in Figure 1-26.

Fabric Events

| Switch | Number | Time | Count | Level | Message |
|---|---|---|---|---|---|
| SM12SW1 | 80 | Dec 4 17:06:23 | 1 | 4 | SULIB-FWDL_END FirmwareDownload has completed successfully. |
| SM12SW1 | 79 | Dec 4 17:01:56 | 1 | 4 | SULIB-CP_REBOOT_OK Standby CP rebooted successfully. |
| SM12SW2 | 47 | Dec 4 16:58:52 | 1 | 2 | EM-CP_ERR CP in slot 6 set to faulty because of CP ERROR |
| SM12SW1 | 77 | Dec 4 16:58:52 | 1 | 4 | HAMKERNEL-NON_REDUNDANT (session=2) Reset Standby CP due to |
| SM12SW1 | 78 | Dec 4 16:58:52 | 1 | 2 | EM-CP_ERR CP in slot 6 set to faulty because of CP ERROR |
| SM12SW1 | 76 | Dec 4 16:58:33 | 1 | 4 | SULIB-CP_REBOOT Standby CP reboots. |
| SM12SW1 | 75 | Dec 4 16:54:51 | 1 | 4 | SULIB-ACTIVE_FAILOVER Active CP forced failover succeeded. This CP is |
| SM12SW2 | 46 | Dec 4 16:52:00 | 1 | 4 | EM-BOOT Restart reason: Failover |
| SM12SW1 | 74 | Dec 4 16:52:00 | 1 | 4 | EM-BOOT Restart reason: Failover |
| SM12SW1 | 73 | Dec 4 16:51:57 | 1 | 4 | HAMKERNEL-TAKEOVER (session=1) Now this CP should be active |
| SM12SW2 | 45 | Nov 20 17:01:14 | 1 | 2 | EM-HIL_FAIL HIL Error: hilPwrOnOffSlot failed, rc=-4 for Slot 10 |
| SM12SW2 | 44 | Nov 20 17:01:10 | 1 | 2 | EM-FRU_FAULTY Slot 10 set to faulty, rc=0x30024 |
| SM12SW2 | 42 | Nov 14 14:45:40 | 1 | 3 | DIAG-CLEARERR Pt8/15(60) Ch7/4 Diagnostics Error Cleared Err# 0120 |
| SM12SW2 | 43 | Nov 14 14:45:40 | 1 | 3 | DIAG-CLEARERR Pt8/14(61) Ch7/5 Diagnostics Error Cleared Err# 0120 |
| SM12SW2 | 41 | Nov 14 14:44:48 | 1 | 1 | DIAG-PORTDIED spinsilkPort Loopback, pass 2, Pt8/15(60) Ch7/4 2Gbps |
| SM12SW2 | 40 | Nov 14 14:44:34 | 1 | 1 | DIAG-PORTDIED spinsilkPort Loopback, pass 2, Pt8/14(61) Ch7/5 2Gbps |
| SM12SW2 | 38 | Nov 7 11:22:46 | 1 | 2 | EM-SLOT_NOT_SEATED Slot 7 ejector not closed |
| SM12SW2 | 39 | Nov 7 11:22:46 | 1 | 4 | EM-BOOT Restart reason: Power-on |
| SM12SW1 | 72 | Nov 7 11:22:46 | 1 | 4 | EM-BOOT Restart reason: Power-on |
| SM12SW2 | 37 | Oct 30 16:52:22 | 1 | 2 | EM-CP_ERR CP in slot 6 set to faulty because of CP ERROR |
| SM12SW1 | 70 | Oct 30 16:52:22 | 1 | 4 | HAMKERNEL-NON_REDUNDANT (session=1) Reset Standby CP due to |
| SM12SW1 | 71 | Oct 30 16:52:22 | 1 | 2 | EM-CP_ERR CP in slot 6 set to faulty because of CP ERROR |

*Figure 1-26 Fabric Event Log*

We can sort the columns into ascending or descending order by clicking the column headings; in our example we have sorted by time, indicated by the small arrow head in the *Time* column heading. We can also re-arrange the columns to suit your requirements by dragging and dropping them as required. To exit from the log, just close the window.

## Topology

The Topology is the physical configuration of the fabric, including active domains and paths. The topology report is as viewed from the local domain (the local domain is the switch that was selected in the fabric view frame).

To view the **Fabric Topology** report, we click the button shown in Figure 1-27.

*Figure 1-27 Fabric Topology button*

This takes us to the Fabric Topology report, as shown in Figure 1-28 and continued in Figure 1-29.

**View Fabric Topology from Switch SM12SW1:** Print

**There are total of 6 domains in the fabric**

| |
|---|
| Local Domain ID: 1 (Switch Name: SM12SW1) |
| Domain ID: 1 (Switch Name: SM12SW1) |
| Domain ID: 2 (Switch Name: SM12SW2) |
| Domain ID: 3 (Switch Name: SF32SW1) |
| Domain ID: 4 (Switch Name: SS08SW2) |
| Domain ID: 21 (Switch Name: SF16SW2) |
| Domain ID: 22 (Switch Name: SF16SW1) |

**Active Paths:**

| |
|---|
| **Destination Domain ID:2** (Switch Name: SM12SW2) |
| Destination's World Wide Name: 10:00:00:60:69:80:06:7b |
| Number of Path(s) to Domain 1: 1 |
| Metric to reach the domain: 500 |

**Path Number 1:**

| Output Port | Input Ports | Total Bandwidth | Bandwidth Demand | Hop Count | Flag |
|---|---|---|---|---|---|
| 4/0 | 1/0 , 3/5 | 4 | 200 | 1 | D |

*Figure 1-28 Fabric Topology report*

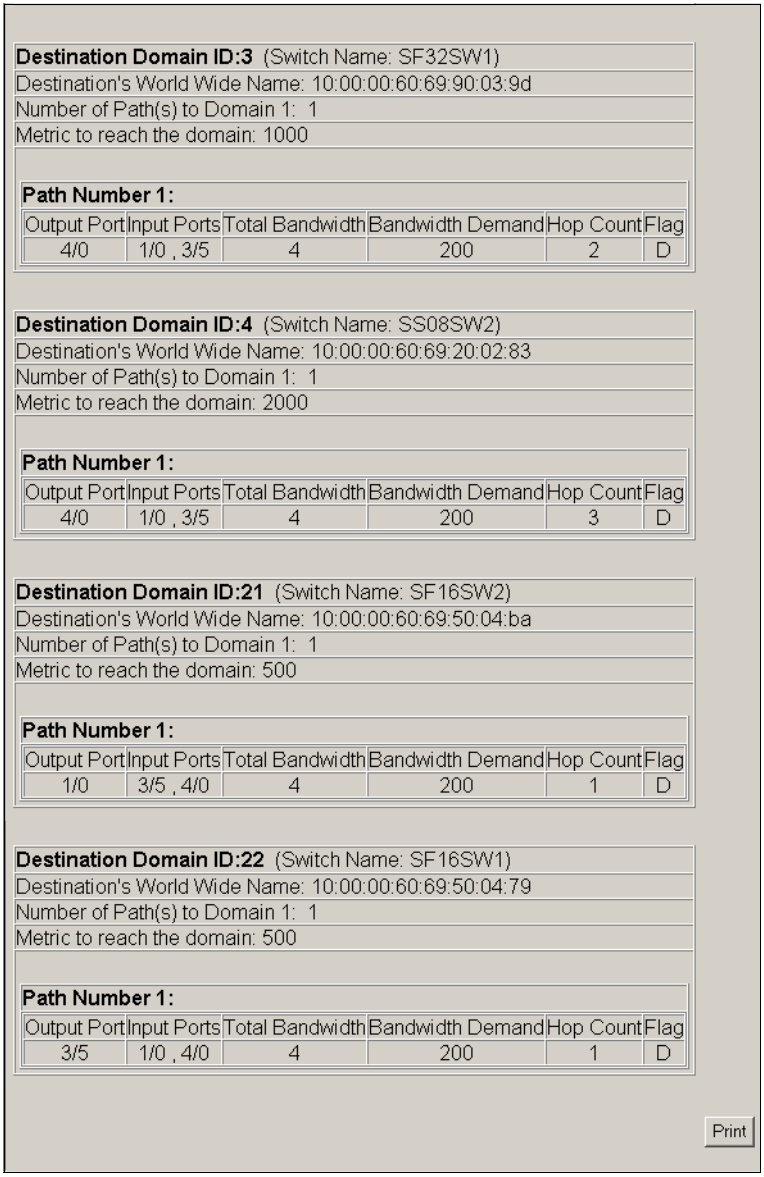

**Destination Domain ID:3** (Switch Name: SF32SW1)
Destination's World Wide Name: 10:00:00:60:69:90:03:9d
Number of Path(s) to Domain 1: 1
Metric to reach the domain: 1000

**Path Number 1:**

| Output Port | Input Ports | Total Bandwidth | Bandwidth Demand | Hop Count | Flag |
|---|---|---|---|---|---|
| 4/0 | 1/0 , 3/5 | 4 | 200 | 2 | D |

**Destination Domain ID:4** (Switch Name: SS08SW2)
Destination's World Wide Name: 10:00:00:60:69:20:02:83
Number of Path(s) to Domain 1: 1
Metric to reach the domain: 2000

**Path Number 1:**

| Output Port | Input Ports | Total Bandwidth | Bandwidth Demand | Hop Count | Flag |
|---|---|---|---|---|---|
| 4/0 | 1/0 , 3/5 | 4 | 200 | 3 | D |

**Destination Domain ID:21** (Switch Name: SF16SW2)
Destination's World Wide Name: 10:00:00:60:69:50:04:ba
Number of Path(s) to Domain 1: 1
Metric to reach the domain: 500

**Path Number 1:**

| Output Port | Input Ports | Total Bandwidth | Bandwidth Demand | Hop Count | Flag |
|---|---|---|---|---|---|
| 1/0 | 3/5 , 4/0 | 4 | 200 | 1 | D |

**Destination Domain ID:22** (Switch Name: SF16SW1)
Destination's World Wide Name: 10:00:00:60:69:50:04:79
Number of Path(s) to Domain 1: 1
Metric to reach the domain: 500

**Path Number 1:**

| Output Port | Input Ports | Total Bandwidth | Bandwidth Demand | Hop Count | Flag |
|---|---|---|---|---|---|
| 3/5 | 1/0 , 4/0 | 4 | 200 | 1 | D |

Print

*Figure 1-29 Fabric Topology report - continued*

The Fabric Topology report lists the domain IDs and switch names for all the active domains in the fabric.

For each switch in the fabric, the window displays the active paths to the local domain (these are the Inter-Switch Links (ISLs). Also shown are the output port numbers (ISL ports), input port numbers and the hop count.

## Name Server

The Name Server table provides the name server entries listed in the name server database. This includes all name server entries for the fabric, not only those local to the host domain. Each row in the table represents a different device which has logged into the fabric.

To access the **Name Server** table, click the button as shown in Figure 1-30.

*Figure 1-30 Name Server button*

The Name Server table, as shown in Figure 1-31 and Figure 1-32, shows the list of all devices that are attached to our SAN fabric. This provides a good cross reference of WWPN / WWN and the port position on the switch. It also lists the zones that the port is a member of, and therefore can be a very useful problem determination tool.

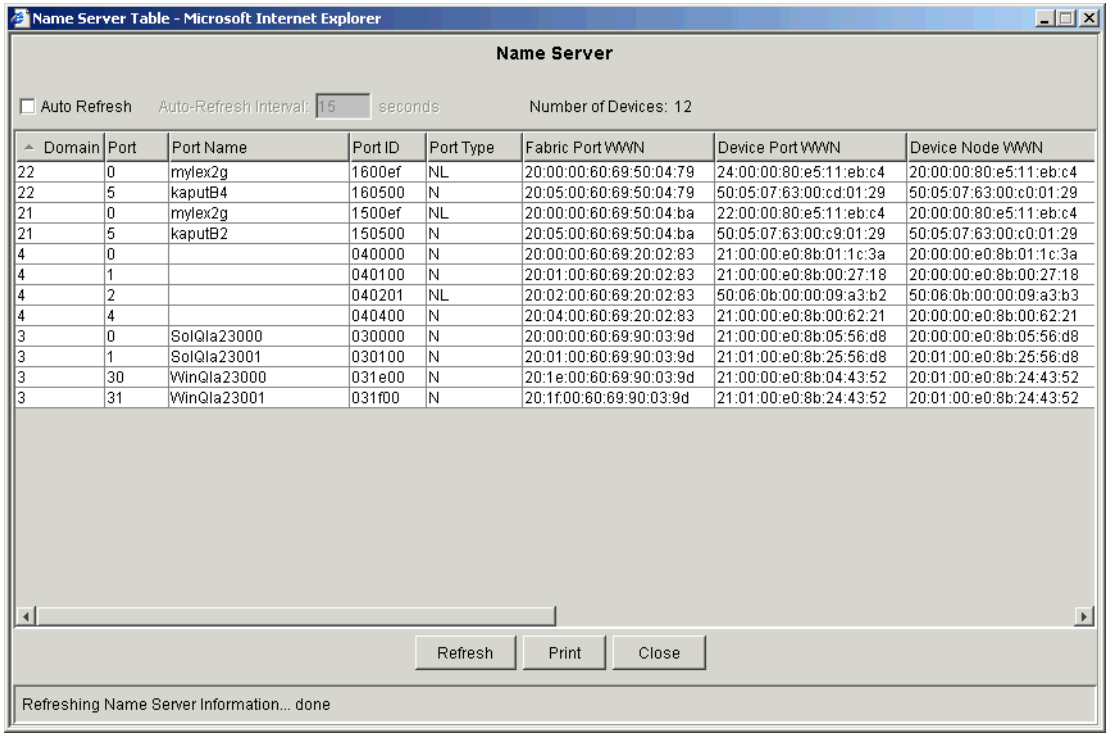


| Domain | Port | Port Name | Port ID | Port Type | Fabric Port WWN | Device Port WWN | Device Node WWN |
|---|---|---|---|---|---|---|---|
| 22 | 0 | mylex2g | 1600ef | NL | 20:00:00:60:69:50:04:79 | 24:00:00:80:e5:11:eb:c4 | 20:00:00:80:e5:11:eb:c4 |
| 22 | 5 | kaputB4 | 160500 | N | 20:05:00:60:69:50:04:79 | 50:05:07:63:00:cd:01:29 | 50:05:07:63:00:c0:01:29 |
| 21 | 0 | mylex2g | 1500ef | NL | 20:00:00:60:69:50:04:ba | 22:00:00:80:e5:11:eb:c4 | 20:00:00:80:e5:11:eb:c4 |
| 21 | 5 | kaputB2 | 150500 | N | 20:05:00:60:69:50:04:ba | 50:05:07:63:00:c9:01:29 | 50:05:07:63:00:c0:01:29 |
| 4 | 0 | | 040000 | N | 20:00:00:60:69:20:02:83 | 21:00:00:e0:8b:01:1c:3a | 20:00:00:e0:8b:01:1c:3a |
| 4 | 1 | | 040100 | N | 20:01:00:60:69:20:02:83 | 21:00:00:e0:8b:00:27:18 | 20:00:00:e0:8b:00:27:18 |
| 4 | 2 | | 040201 | NL | 20:02:00:60:69:20:02:83 | 50:06:0b:00:00:09:a3:b2 | 50:06:0b:00:00:09:a3:b3 |
| 4 | 4 | | 040400 | N | 20:04:00:60:69:20:02:83 | 21:00:00:e0:8b:00:62:21 | 20:00:00:e0:8b:00:62:21 |
| 3 | 0 | SolQla23000 | 030000 | N | 20:00:00:60:69:90:03:9d | 21:00:00:e0:8b:05:56:d8 | 20:00:00:e0:8b:05:56:d8 |
| 3 | 1 | SolQla23001 | 030100 | N | 20:01:00:60:69:90:03:9d | 21:01:00:e0:8b:25:56:d8 | 20:01:00:e0:8b:25:56:d8 |
| 3 | 30 | WinQla23000 | 031e00 | N | 20:1e:00:60:69:90:03:9d | 21:00:00:e0:8b:04:43:52 | 20:01:00:e0:8b:24:43:52 |
| 3 | 31 | WinQla23001 | 031f00 | N | 20:1f:00:60:69:90:03:9d | 21:01:00:e0:8b:24:43:52 | 20:01:00:e0:8b:24:43:52 |

*Figure 1-31 Name Server table*

The Name Server table contains the following parameters:

**Domain** Domain ID of the switch to which the device is connected

**Port** Port number of the switch to which the device is connected

**Port Name** Symbolic name as defined in Port Settings

**Port ID** The Fibre Channel Port address of the device (basically a 24 bit hexadecimal number)

**Port Type** Shows whether the port is a public loop port (NL) or whether it is a normal switch fabric port (N)

**Fabric Port WWN** World-Wide Name of the switch port

**Device Port WWN** World-wide name for the device port (WWPN)

**Device Node WWN** World-wide name of the device node (WWNN)

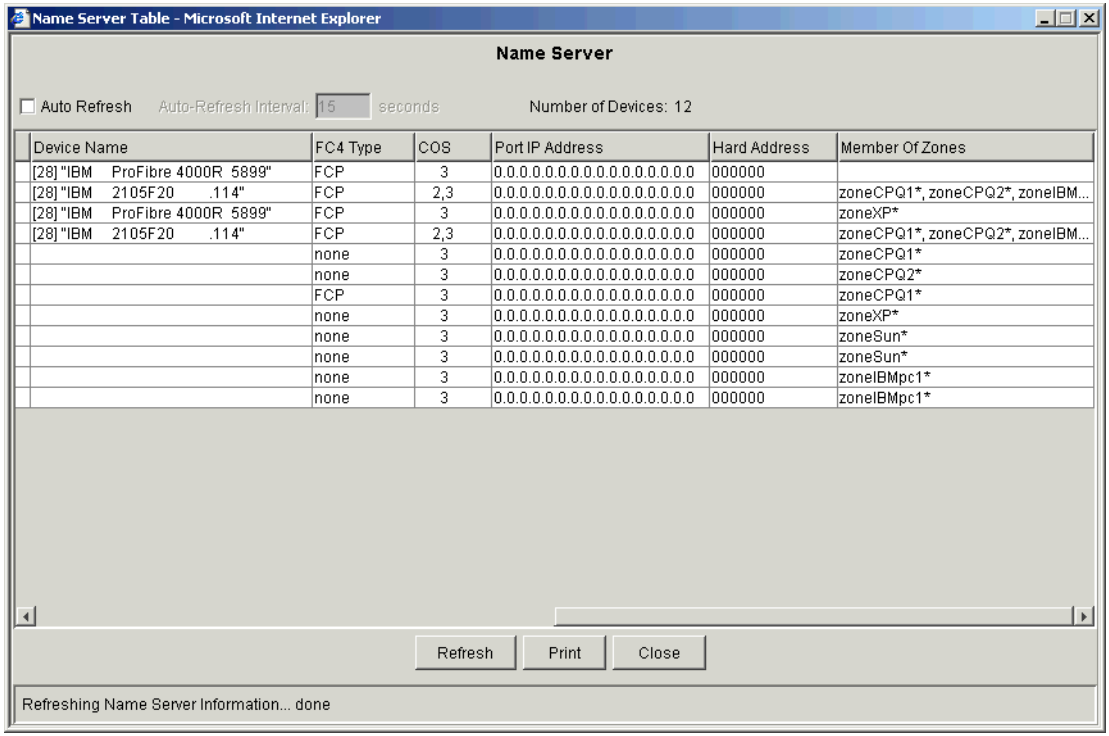


*Figure 1-32 Name Server table - continued*

| | |
|---|---|
| **Device name** | Name of the device according to the SCSI INQUIRY command |
| **FC4 type** | Fibre Channel FC4 layer types supported by the device, such as FCP or IP. |
| **COS** | Class of Service that the device supports |
| **Port IP address** | IP address of the fabric port (may be zeros) |
| **Hard Address** | Hard address assigned to the fabric port |
| **Member of zones** | List of zones to which the port belongs |

## 1.7.3 Zone Admin

The Zone Admin function is used to set up, maintain and activate the zones across the fabric. From here we can also define aliases for members in a zone and can create the zones that will form the active configuration across the fabric.

A zoning license and administrative privileges are required to access this function.

**Note:** All 2109 switches are delivered with the zoning license pre-installed. The 3534 requires an optional Advanced zoning license.

When administering zoning on an IBM TotalStorage SAN Switch, the following steps are recommended:

- Define zone aliases to establish groupings.
- Add zone members.
- Place zones into one or more zone configurations.
- Enable one of the zone configurations (only one can be enabled at a time).

There are three separate methods for adding members to a zone. Each method corresponds to a zoning “mode”, and the combination of the methods corresponds to an additional mode. Once a mode is selected, all operations on zones must use the zoning object selected. Zoning operations must correspond to that mode, and any zones, aliases, and configuration files which do not, cannot be selected.

### Zone administration

To access zone administration, we click the icon shown in Figure 1-33.

*Figure 1-33 Zone Admin button*

After we click the **Zone Admin** button, a prompt displays requesting User Name and Password, as shown in Figure 1-34.

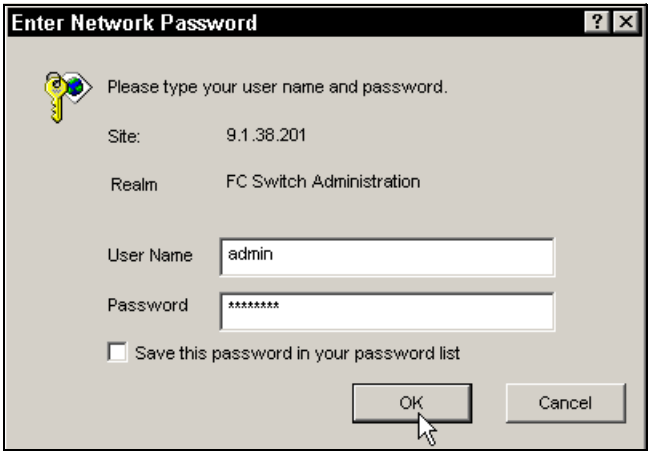


*Figure 1-34 User Authentication*

Enter the user name and the password and click **OK**. (The defaults are *admin* / *password*).

By clicking the **View** menu we can select a Zoning Scheme we wish to use from the pulldown menu, as shown in Figure 1-35.

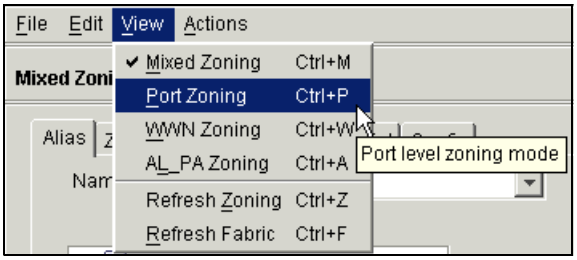


*Figure 1-35 Zoning scheme selection*

We describe the zoning schemes in the following sections.

### *Mixed zoning*

In this scheme, all objects are displayed in the *Member Selection List*. Any object, being a WWN, switch, port, AL_PA, or alias, can be selected to be managed in the *Members* list. When the Zoning management function is opened, this is the default scheme.

Working in the mixed zoning scheme allows us to define a WWN and a physical port to be within the same configuration. If we have mixed members in a zone, the zoning uses session based hard zoning.

### *Port zoning*

This zoning scheme only offers physical switches and ports to be selected and defined as members for alias, zoning, QuickLoop, Fabric Assist, and configuration groups. Aliases, zones, and configuration groups which have objects other than physical ports will not be displayed in this scheme.

If we work only in this Port zoning scheme, our configuration will be hardware enforced by the switch ASICs (hard zoned).

### *WWN zoning*

This scheme only allows aliases, zoning and configuration file operations on WWNs. Aliases, zones, and configuration files which have objects other than WWNs will not be displayed within this scheme.

If we work only in this WWN zoning scheme, our configuration will be hardware enforced by the switch ASICs (hard zoned).

### *AL_PA zoning*

This scheme allows only aliases, zoning and configuration file operations on AL_PA's in a QuickLoop. Aliases, zones and configuration files which have objects other than AL_PA's in a QuickLoop will not be displayed.

If we work only in this AL_PA zoning scheme, our configuration will be hardware enforced by the switch ASICs (hard zoned).

## 1.7.4 Implementing zoning

In the following example we show the windows in which we apply zoning concepts that have previously been discussed. For our purpose we have chosen the Port zoning scheme, although the procedure is the same for WWN, AL_PA, and Mixed zoning schemes.

Upon selecting *Port Zoning* from the View menu, the Port Zoning view appears, with the *Alias* tab displayed. This is shown in Figure 1-36.

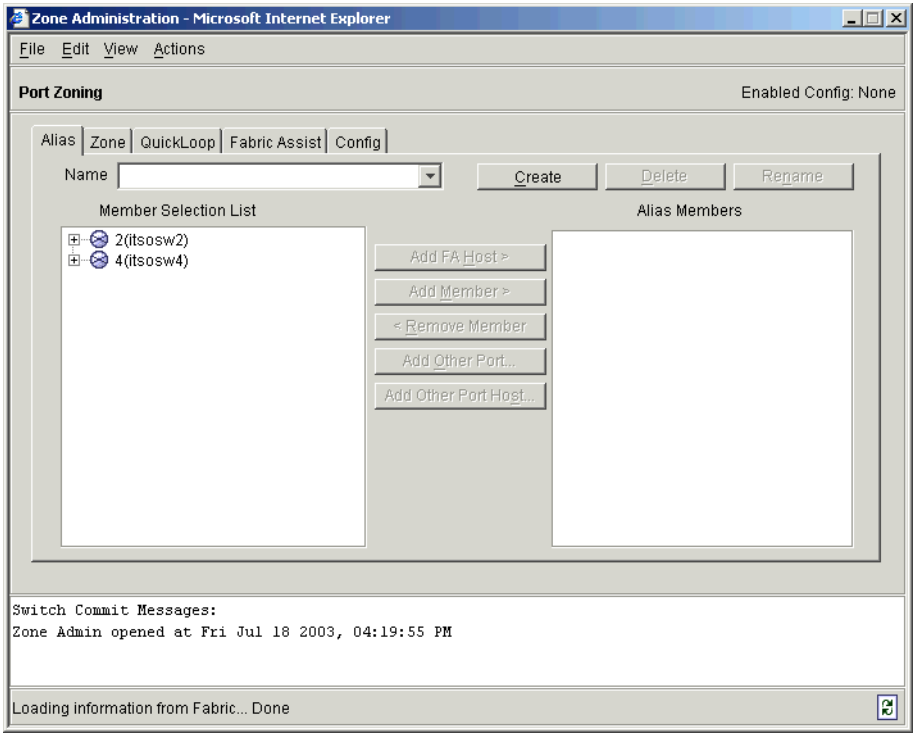


*Figure 1-36 Port Zoning initial view*

There are five tabs provided in the Zoning view:

- Alias
- Zone
- QuickLoop
- Fabric Assist
- Config

(The *AL_PA* scheme will not have the Fabric Assist tab.)

## Alias tab

By defining an alias to a port(s) or WWN(s), we simplify understanding what the device is that we are working with in other tabs. We recommend assigning aliases and ensuring they are maintained to correctly identify SAN components.

The **Alias** tab lists the fabric Domains/Ports/Nodes in the Member Selection List. This tab is used to define alias names to port(s) or node(s) so that we easily recognize them.

To create an alias, we first click the **Create** button, and enter a meaningful name for the Host / Storage attached. The Create New Alias window is shown in Figure 1-37.

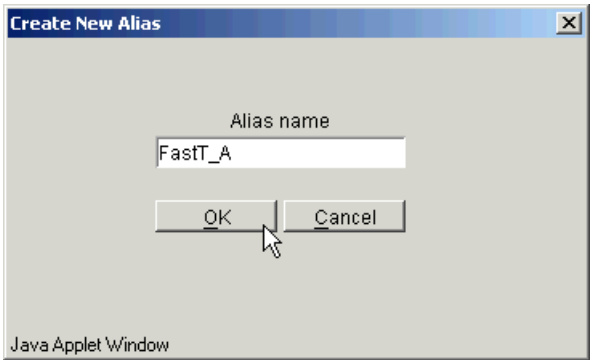


*Figure 1-37 Create New Alias*

After clicking **OK**, we see the name displayed in the Name field, shown in Figure 1-38. We can now select a member, or multiple members, from the *Member Selection List* on the Left.

We select port 5 on switch domain 4, and then click the **Add Member** button to add it to the *Alias Members* list in the right panel.

If a host or device has multiple HBAs, we may wish to add more members to our alias. As we are defining an alias for one controller on a FAStT, we wish to only define this single port, as shown. We have successfully identified port 5 on switch domain 4 to have an alias of FAStT_A.

To allow easier identification of hosts and devices listed in the selection list, we can click the plus (+) beside the port to open further detail of the WWNN, WWPN, and VPD descriptor information. In Figure 1-38 we can see that both ports 5 and 7 on switch domain 4 are attached to a WWNN `20:02:00:a0:b8:0c:bc:e7` with different WWPNs each being identified by the VPD descriptor of `IBM 1742,` which is the machine type for a FAStT 700.

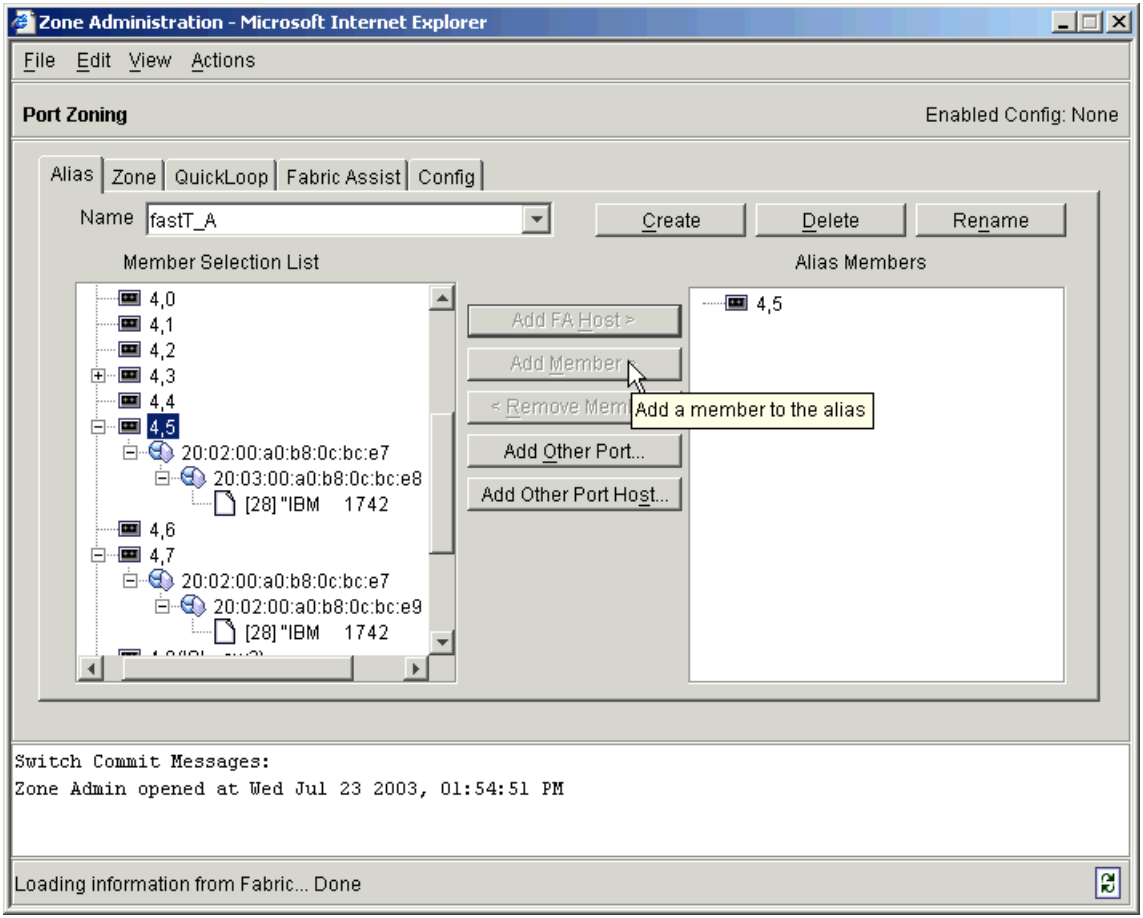


*Figure 1-38 Alias administration*

We would follow the same procedure for all our hosts, and storage, before adding them to zones.

Table 1-2 describes the fields and buttons on the **Alias** tab.

*Table 1-2 Alias tab description:*

| Button | Function |
|---|---|
| Name | Select an existing alias name to be modified. |
| Create | Select to create a new alias. A new alias dialog displays. Enter a new alias name that is unique. The new alias name cannot contain spaces. |
| Delete | Select to delete the alias selected in the Name field. Deleting an alias automatically removes it from all zones. |
| Rename | Select to rename the alias selected in the Name field. A dialog displays in which you can edit the alias name. Renaming an alias automatically renames it in all zones. |
| Member Selection List | This field contains a list of potential alias members, including switches, ports, Nodes, WWNs, and QuickLoop AL_PAs. |
| Add FA Host > | Use this button to add a Fabric Assist Host to the member list. |
| Add Member > | Select to add the item selected in the Member Selection List to the Alias Members list. You can add individual ports or an entire switch. If a switch is added, all ports on the switch are added. To add a device WWN, select either a node WWN (folder icon) or port WWN (blue circle icon) from the WWN sub-tree. |
| < Remove Member | Select to remove the member selected from the Alias Name Members Selection list. |
| Add Other Port | Select to add a switch/port combination that currently is not part of the fabric. |
| Add Other Port Host | Select to add a switch/port combination of a host that currently is not part of the fabric. |

## Selecting ports on the M12

Some consideration must be taken to understand the port addressing when zoning an IBM TotalStorage SAN Switch M12. In previous versions of the Fabric OS (version 2.0 and version 3.0), the primary method for identifying a port within the fabric was the “domain, port” combination.

For example, to add port 1 on domain 5 to a zone:

```
sw96:admin>zoneadd "bluezone","5,1"
```

The “domain, port” method of selecting ports cannot be used in the M12 because of the addition of slots and the high port count of the switch. This method was replaced in Fabric OS version 4.0 by two methods to specify a particular port: the slot/port method and the port area number method.

### *Slot/port method*

To select a specific port, you must identify both the slot number and port number that you are working with.

When specifying a particular slot and port for a command, the slot number operand must be followed by the slash (/) and then a value for the port number. For example, to enable port 63, we specify:

```
portEnable 10/15
```

The M12 has a total of 10 slots, counted 1 - 10. Slot number 5 and slot number 6 are CP cards, and slots 1 - 4 and 7 - 10 are switch cards. On each switch card, there are 16 ports counted from the bottom 0 - 15. A particular port must be represented by both slot number (1 - 10) and port number (0 - 15).

**Restriction:** No spaces are allowed between the slot number, the slash (/), and the port number.

### *Port area number method*

Some commands, such as zoning commands, require you to specify ports using the port area number method. In the Fabric OS version 4.0 each port on a particular domain is given a unique area ID.

The chassis contains two logical switches. The Area IDs for both logical 64-port switches range from 0 - 63. Both logical switch 0 and 1 have a port that is referenced with Area ID 0.

An area ID for each port is unique inside each logical switch (that is, each assigned domain ID). These are two of the three parts of a 24-bit Fibre Channel address ID: 8-bit domain ID, 8-bit area ID, and 8-bit port ID.

Use the **switchShow** command to display all ports on the current (logical) switch and their corresponding area IDs, as follows:

```
Area Slot Port Media Speed State
====================================
  0    7    0   --    N2   No_Module
  1    7    1   --    N2   No_Module
  2    7    2   --    N2   No_Module
  3    7    3   --    N2   No_Module
  4    7    4   id    N2   Online
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

   lines removed for clarity
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

 60   10   12   --    N2   No_Module
 61   10   13   --    N2   No_Module
 62   10   14   --    N2   No_Module
 63   10   15   id    N2   No_Light
```

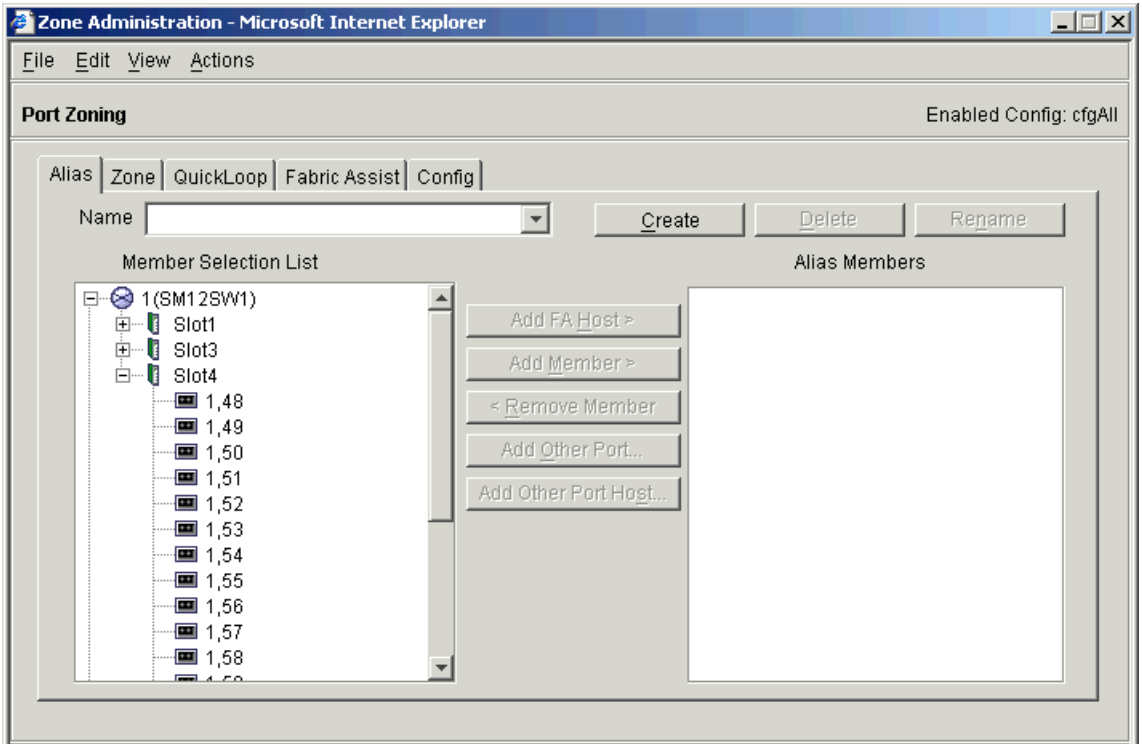


*Figure 1-39 M12 Zoning — Slot / port area number*

Figure 1-39 shows how the WEB TOOLS interface for the M12 Zoning view displays the slot and associated ports for a domain (switch).

## Zone tab

We use the **Zone** tab to specify which switch ports are to be in the selected zone and to create and manage zones. A zone can have one or multiple members, and can include switches, ports, WWNs, aliases, and QuickLoop AL_PAs.

**Important:** We recommend creating individual zones of each host to the disk storage subsystems. Also, hosts should have a separate HBA for Tape communication, and again be in another individual Host / Tape zone.

This small granularity of zoning removes unnecessary PLOGI activity from host to host, as well as removing the risk of problems caused by a faulty HBA affecting others.

The **Zone** tab is shown in Figure 1-40.

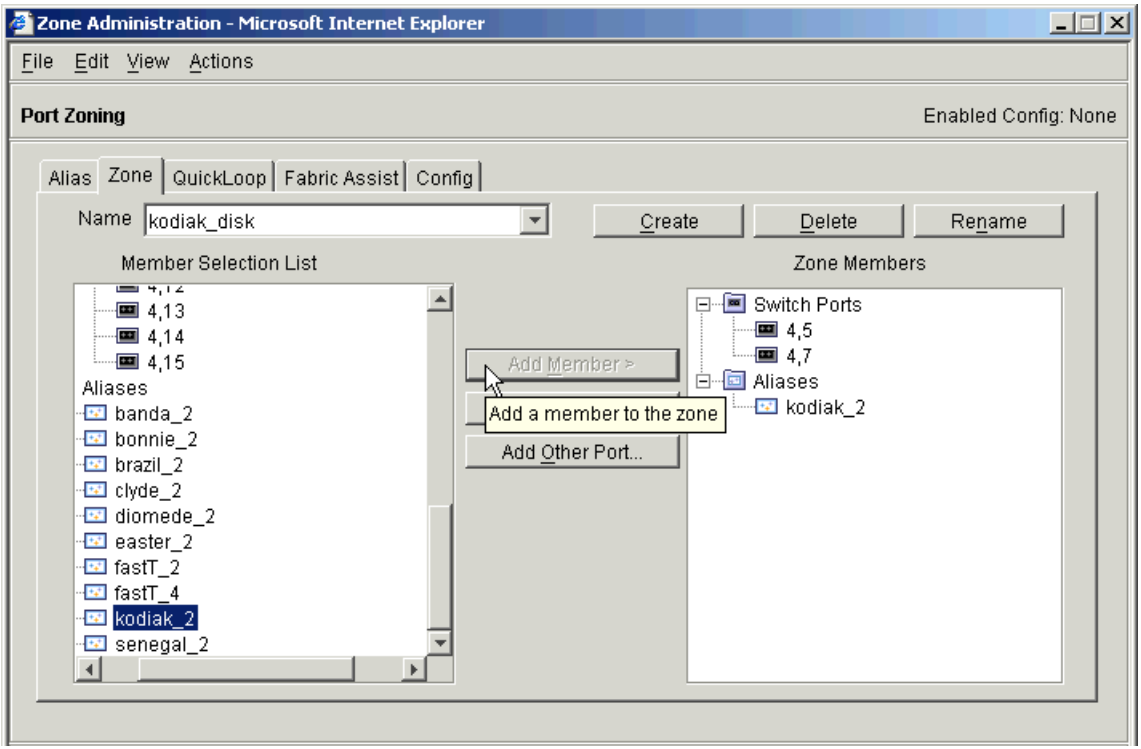


*Figure 1-40 Zone creation*

In this example we have created a zone name of *kodiak_disk*.

Then we added the switch ports which have our FAStT attached, and also added the alias *kodiak_2*, representing our host which we defined in the previous topic.

As mentioned in the previous recommendation, we could add another HBA installed in our kodiak server to this zone, but we do not recommend adding other hosts. We choose to define a separate zone for each host.

Table 1-3 describes the fields and buttons on the **Zone** tab.

*Table 1-3 Zone tab description*

| Button | Function |
|---|---|
| Name | Select an existing alias name to be modified. |
| Create | Select to create a new alias. A new alias dialog displays. Enter a new alias name that is unique. The new alias name cannot contain spaces. |
| Delete | Select to delete the alias selected in the Name field. Deleting an alias automatically removes it from all zones. |
| Rename | Select to rename the alias selected in the Name field. A dialog displays in which you can edit the alias name. Renaming an alias automatically renames it in all zones. |
| Member Selection List | This field contains a list of potential alias members, including switches, ports, Nodes, WWNs, and QuickLoop AL_PAs. |
| Add Member > | Select to add the item selected in the Member Selection List to the Alias Members list. You can add individual ports or an entire switch. If a switch is added, all ports on the switch are added. To add a device WWN, select either a node WWN (folder icon) or port WWN (blue circle icon) from the WWN sub-tree. |
| < Remove Member | Select to remove the member selected from the Alias Name Members Selection list. |
| Add Other Port | Select to add a switch/port combination that currently is not part of the fabric. |

## QuickLoop tab

A QuickLoop license is required to use this tab. You can use the **QuickLoop** tab to create and manage QuickLoops if used in conjunction with zoning.

The **QuickLoop** tab is shown in Figure 1-41.

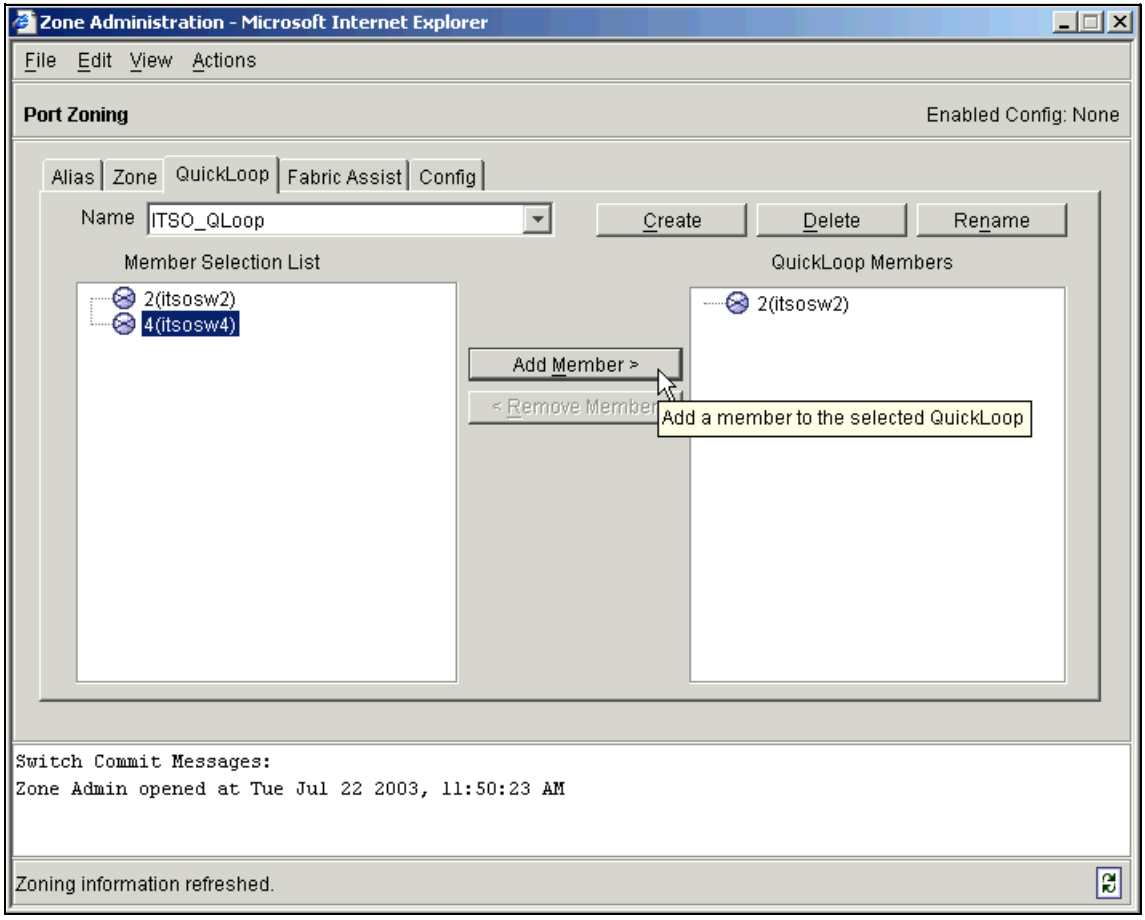


*Figure 1-41 QuickLoop zoning tab*

In this example we have created a QuickLoop Name called *ITSO_QLoop* which will have two switch members. We have already added sw2 to the Members list and have selected sw4 so that we can add it also to the *ITSO_QLoop* Member list. For more information regarding QuickLoop, refer to “QuickLoop” on page 108.

Table 1-4 describes the fields and buttons on the **QuickLoop** tab.

*Table 1-4 QuickLoop tab description*

| Field | Function |
|---|---|
| Name | To modify an existing QuickLoop, select a QuickLoop name. |
| Create | Click to create a new QuickLoop. A dialog displays in which you can enter the name of the new QuickLoop. All names must be unique and contain no spaces. |
| Delete | Click to delete the QuickLoop selected in the QuickLoop Name field. Deleting a QuickLoop automatically removes it from all aliases, zones, and zone configurations, including the associated AL_PAs. |
| Rename | Click to edit the name of the QuickLoop selected in the QuickLoop Name field. A dialog displays in which you can edit the name of the QuickLoop. |
| Member Selection List | A list of valid members available to add to a QuickLoop. |
| Add Member > | Click to add the switch selected in the Switch Selection List to the QuickLoop Members list. |
| < Remove Member | Click to remove the selected member from the QuickLoop Member list. |

## Fabric Assist Tab

Fabric Assist (FA) is a feature that allows Private Loop hosts on QuickLoop enabled ports, to communicate with fabric storage devices on F_Ports. We use the **Fabric Assist** tab to create and manage Fabric Assists. A QuickLoop license is required to use this tab.

The **Fabric Assist** tab is shown in Figure 1-42.

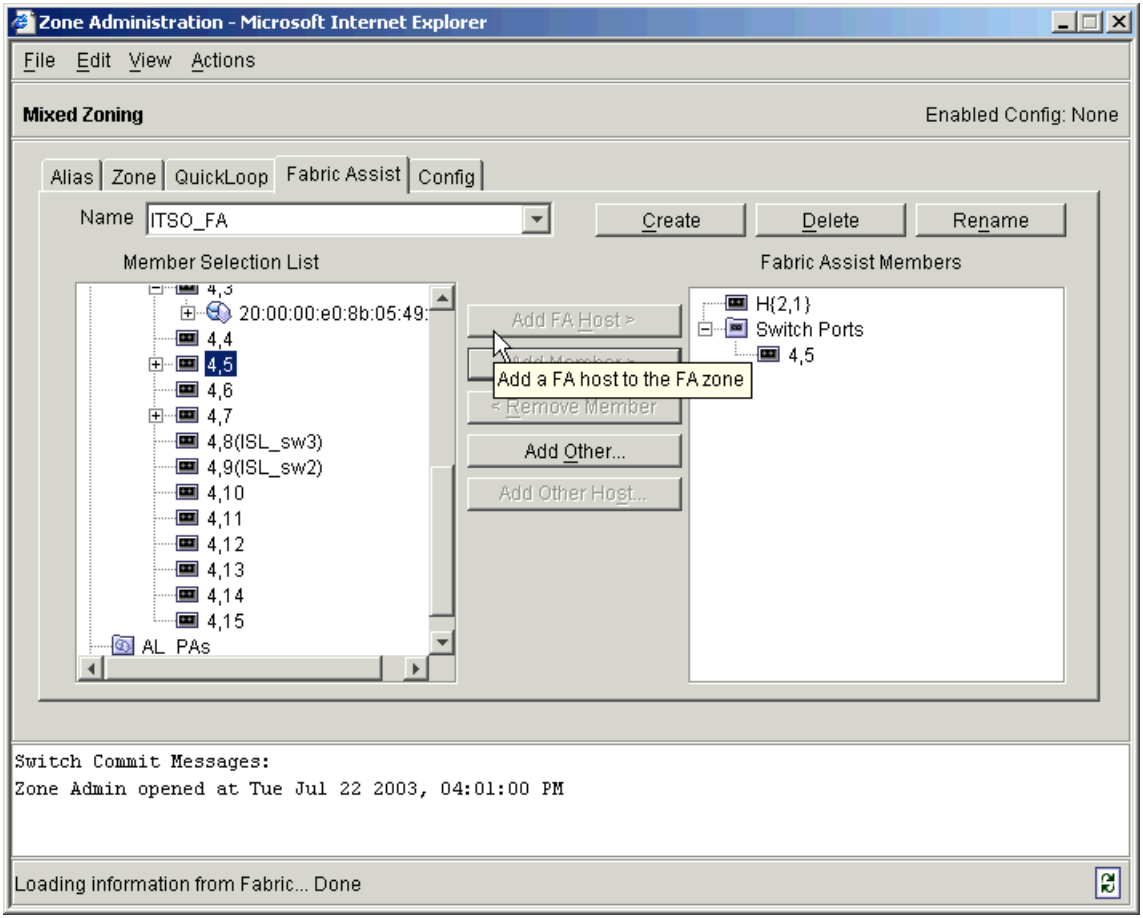


*Figure 1-42 Fabric Assist zoning tab*

In this example we created an ISTO_FA group name, where we then added Host port 2,1 by using the **Add FA Host** button. We have then added the target storage, on port 4,5 by using the **Add Member** button.

Table 1-5 describes the fields and buttons in the **Fabric Assist** tab.

*Table 1-5 Fabric Assist tab description*

| Field | Function |
|---|---|
| Name | Select an existing Fabric Assist name to be modified or viewed. |
| Create | Click to create a new Port Fabric Assist name. A dialog displays. Enter the name of the new Port Fabric Assist. All names must be unique and contain no spaces. |
| Delete | Click to delete the Port Fabric Assist selected in the FA Name field. Deleting a Port Fabric Assist automatically removes it from all aliases, zones, and zone configurations, including the associated AL_PAs. |
| Rename | Select to edit the name of the Port Fabric Assist selected in the FA Name field. |
| Member Selection List | This field displays a list of members available to add to the Port Fabric Assist list. |
| Add FA Host > | Click to add the selected item in the Member Selection List as a host to the Fabric Assist name list. Only one domain port or a WWN can be added as a host. |
| Add Member > | Click to add the member selected in the Member Selection List to the Fabric Assist name list. |
| < Remove Member | Click to remove the selected member from the Fabric Assist name list. |
| Add Other... | Click to add a switch/port combination that is not currently part of the fabric. |
| Add Other Host... | Click to add a switch/port/host combination that currently is not part of the fabric. |
| Fabric Assist Members | This field displays a list of the members that belong to the Fabric Assist group currently selected in the Name field. |

## Config tab

We now use the **Config** tab to create a zone configuration. Zone configurations are used to enable or disable a group of zones at the same time. The **Config** tab is shown in Figure 1-43.

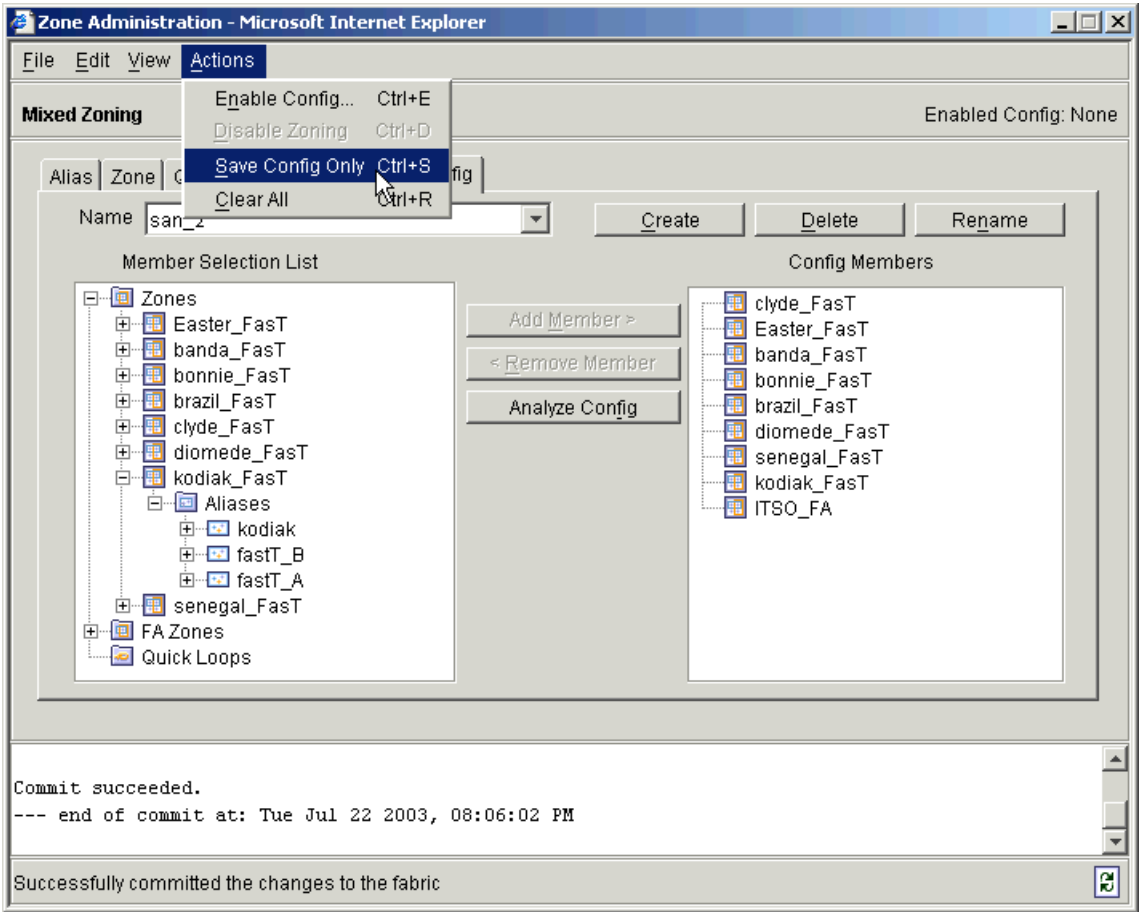


*Figure 1-43 Zoning Config tab*

In this example we have created a config name called *san_2*. We then used the **Add Member >** button to move the zones we created in previous steps, listed in the left column, to the Config Members list on the right. This process creates a configuration containing all the desired zones we wish to activate.

**Tip:** We are able to quickly confirm the members of each zone in the selection list by clicking the small plus (**+**) beside each listed zone, to view the individual definitions.

Table 1-6 contains a description of the fields and buttons that appear on the **Config** tab.

*Table 1-6 Config tab description*

| Button | Function |
|---|---|
| Name | Select an existing configuration to modify. |
| Create | Click to create a new configuration. A dialog displays. Enter the name of the new configuration. All names must be unique and contain no spaces. |
| Delete | Click to delete the configuration selected in the Cfg Name field. |
| Rename | Click to edit the name of the configuration selected in the Cfg Name field. |
| Member Selection List | This field provides a list of the zones and QuickLoops available to add to the configuration. |
| Add Member > | Click to add the switch selected in the Zone/QLoop Selection List to the Config Members list. |
| < Remove Member | Click to remove the selected member from the Config Members list. |
| Analyze Config | Analyzes the configuration that is selected along with it's member zones and aliases. A zoning configuration error window appears in the event of a conflict. |

Before we activate our zone config, we save our it to the switch, using the *Save Config Only* function from the *Actions* pulldown menu. This only saves the config to nonvolatile storage, it does not bring the config active.

After our config is saved, we click the **Analyze Config** button. This checks the validity of our zoning configuration, and alerts us to ports and WWNs that we have not included. Before running, we are prompted to refresh the current configuration from the switch; this is so that the config Analyze can be checked against the most recent information from the switch. The output from the Analyze run against our config is shown in Figure 1-44.

The Zoning Configuration Analyze window displays a summary of the saved configuration and attempts to point out some of the zoning conflicts before applying the changes to the switch. Some of the potential errors it might catch are:

- Ports/WWNs/Devices that are part of the selected configuration, but not part of the fabric.
- Zones with only a single member.

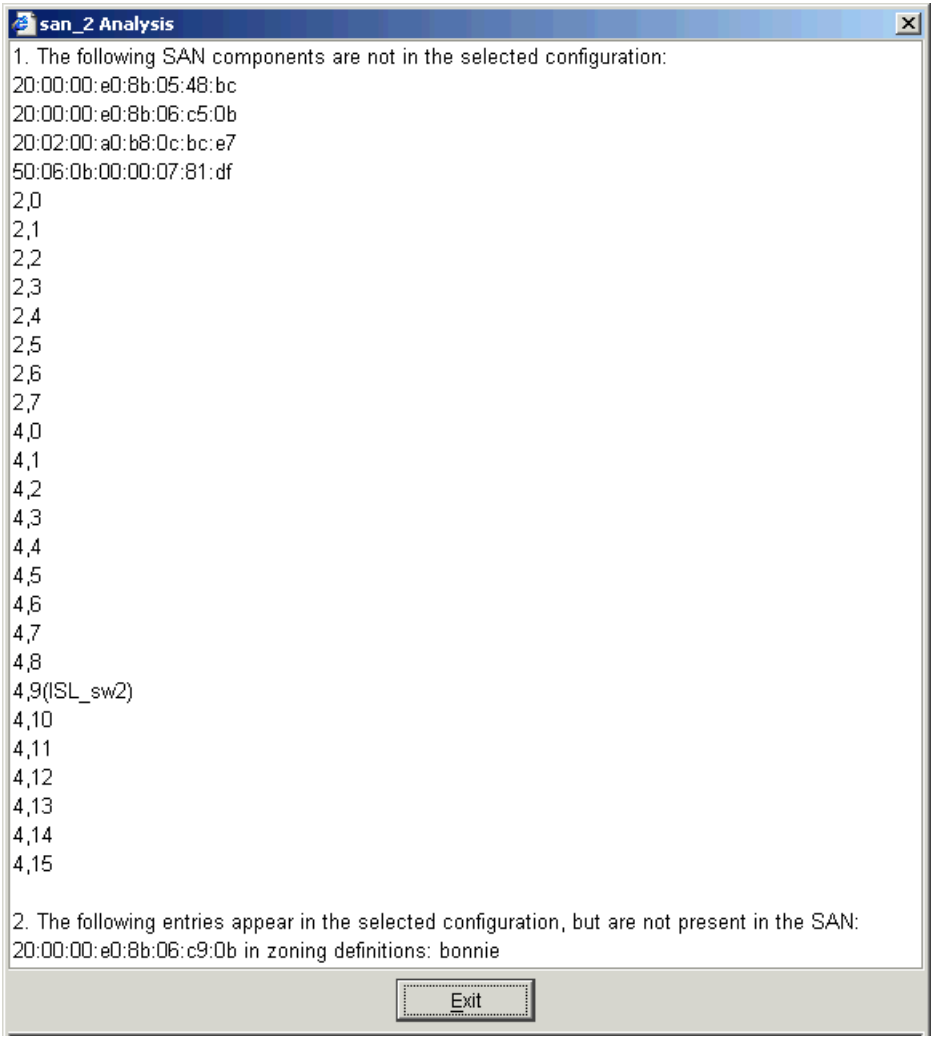


*Figure 1-44 Analyze config output*

After reviewing the Analyze output and ensuring that any of the reported errors do not concern our required configuration, we can now activate the configuration.

## Activating a zoning configuration

To make our zoning definitions active, we need to enable the configuration that we have built. We do this by using **Enable Config...** from the *Actions* pulldown menu, shown in Figure 1-45. We then select the correct config as detailed below.

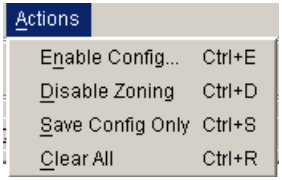


*Figure 1-45 Actions pulldown menu*

We are then prompted to select which config we would like to enable, shown in Figure 1-46.

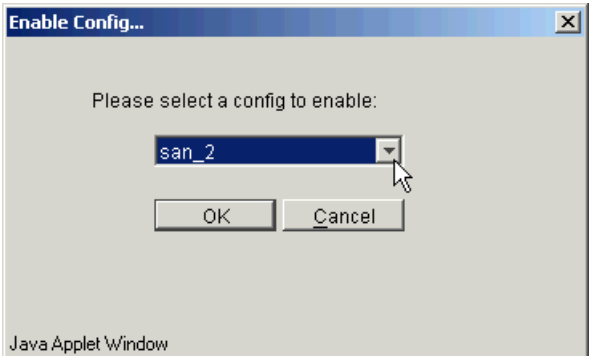


*Figure 1-46 Select config prompt*

We are then warned with a dialogue shown in Figure 1-47, to ensure we are making the correct decision in enabling this configuration.

**Attention:** Care must be taken when enabling zone configs. Adding new zones will not impact any currently running definitions, although removing a zone may have a large impact to the current environment.

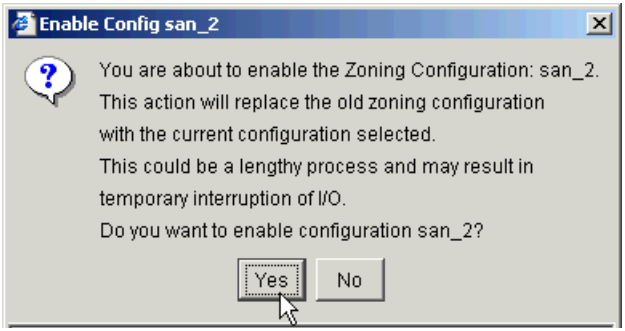


*Figure 1-47 Config enable warning*

At this point the new zone config definitions take place on the SAN fabric.

## Modifying an existing configuration

When adding a new host or a new device into the fabric, changes to the zoning will be necessary. For example, we add a new host, define a *newhost* alias, create a *newhost_FAStT* zone. Using the procedures previously described in this topic, we then add the newhost_FAStT zone to our config.

We then have two choices, immediate implementation, or we can save our updates and perform the activate at a later time:

- Choose **Enable Config...** from the *Actions* pulldown menu, the changes are saved and take effect immediately.
- Choose **Save Config only** from the *Actions* pulldown menu. The changes are saved, but will not take effect immediately. For the changes to take effect, we have to select the configuration in the names list, and then select **Enable Config...** from the *Actions* pulldown menu.

## Zoning and E_Ports

When creating a zone, we only work with device ports or host ports (F_Ports, FL_Ports, L_Ports). ISL Ports (E_Port) should not be included in zone definitions.

If we take the example presented in Figure 1-48.

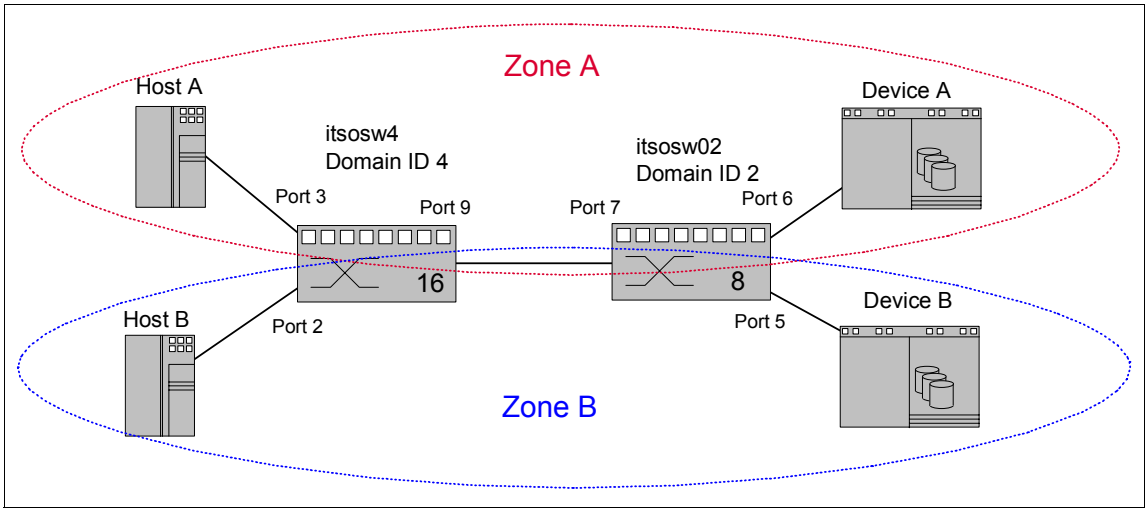


*Figure 1-48 Zoning implementation — E_Ports and Zoning*

To create Zone A, we include:

- Domain ID 4, Port 3 (4,3)
- Domain ID 2, Port 6 (2,6)

But *do not* include any ISL ports, that is to say:

- Domain ID 4, Port 2 (4,2)
- Domain ID 4, Port 9 (4,9)
- Domain ID 2, Port 5 (2,5)
- Domain ID 2, Port 7 (2,7)

Similarly, to create Zone B, we only include:

- Domain ID 4, Port 2 (4,2)
- Domain ID 2, Port 5 (2,5)

Zones do not affect data traffic across ISLs in cascaded switch configurations. Because Hard Zoning enforcement is performed at the destination, an ISL can carry data traffic from all zones.

Therefore, when dealing with zoning, the fabric should be seen as a “cloud” to which are attached devices and hosts. That is, we define the end-to-end destinations, and do not include the path to get there.

## 1.7.5 WEB TOOLS Switch View

From the Switch View of WEB TOOLS, we are able to view a summary of the state of the individual switch, quickly understanding firmware version, IP addresses, port state, and if there is any out-of-line status.

In this section we will use an F16 to describe the interface, although the functions are identical on any of the IBM TotalStorage SAN Switch family. The IBM TotalStorage SAN Switch M12 has some extra functions, and we will point these out where required.

The Switch View presents a picture of the switch as shown in Figure 1-49.

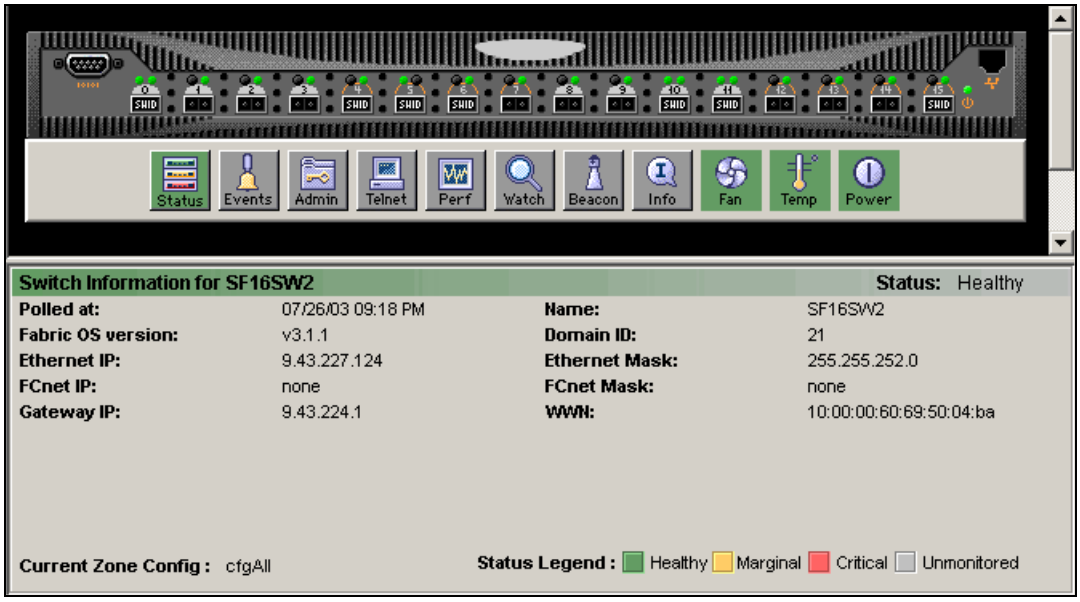


*Figure 1-49 Switch View*

From the Switch View, we have an overview of the actual switch front panel and monitor LEDs.

There are buttons that allow us to drill down further into the switch — we can select to view individual switch settings, performance, status, event logs, port information, FabricWatch events, or open a telnet session. These options are covered in more detail in the following topics.

We can also click any of the ports and be presented with more detailed information on the individual ports.

Next we point our browser to the IP address of logical switch 0 in an IBM TotalStorage SAN Switch M12 as shown in Figure 1-50. Here we can see its detailed information, which would be similar for the other models.

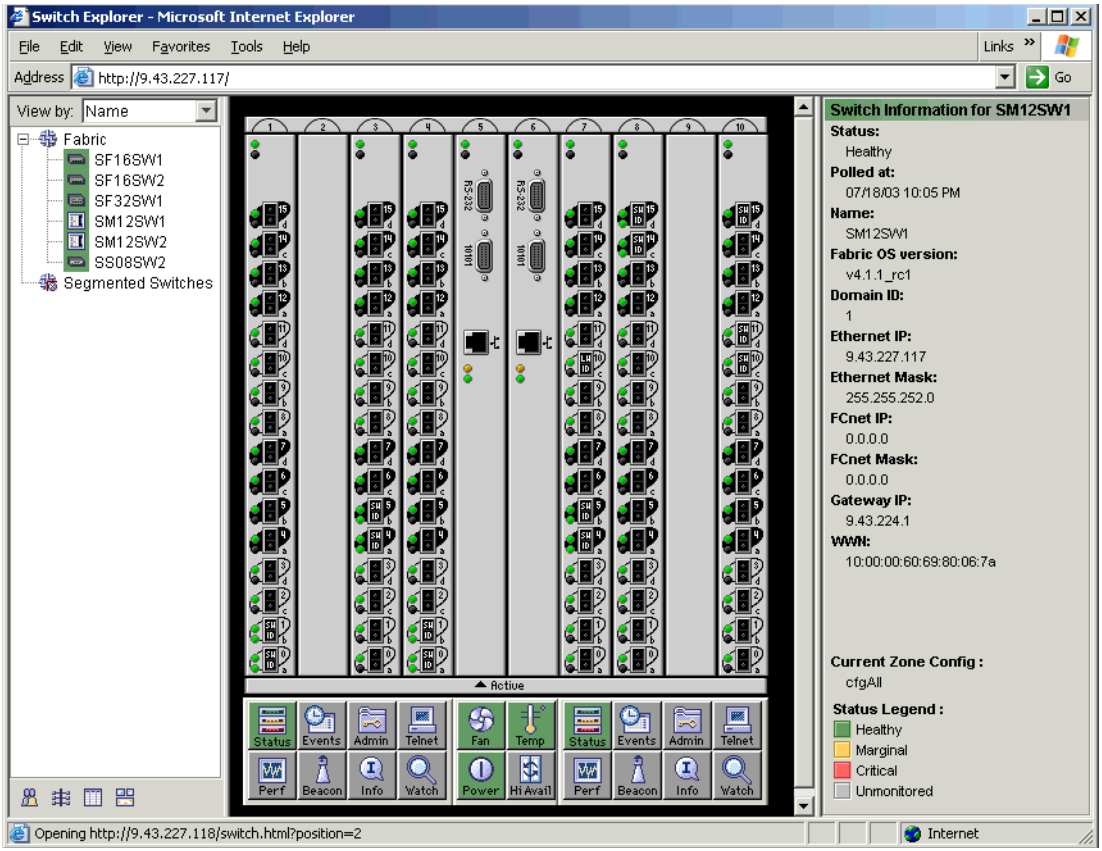


*Figure 1-50 2109-M12* WEB TOOLS *view*

From the M12 Switch view, we can also look at temperature, fan speeds, CP status, and power supply status for the overall chassis.

## Port Information

To access the detailed port information, click the port as shown in Figure 1-51.


*Figure 1-51 Go to Port Information*

The port information will be displayed for the switch, as shown in Figure 1-52.

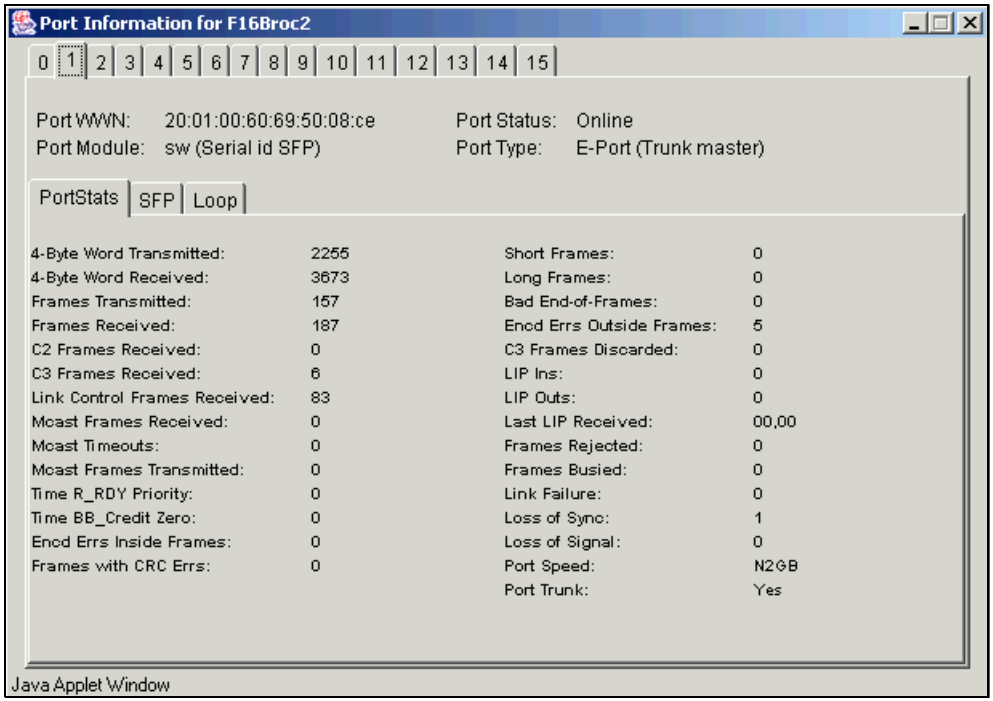


*Figure 1-52 Port Information*

From this window, we can select any of the switch ports. If an SFP is installed, then additional information on the SFP itself can be selected by accessing the **SFP** tab. The **Loop** tab contains information about the loop on a port, including QuickLoop statistics if a QuickLoop license is available.

## Port Information for the M12

The graphical representation of the physical M12 chassis, in the middle frame, includes both logical switches, as shown in Figure 1-53.

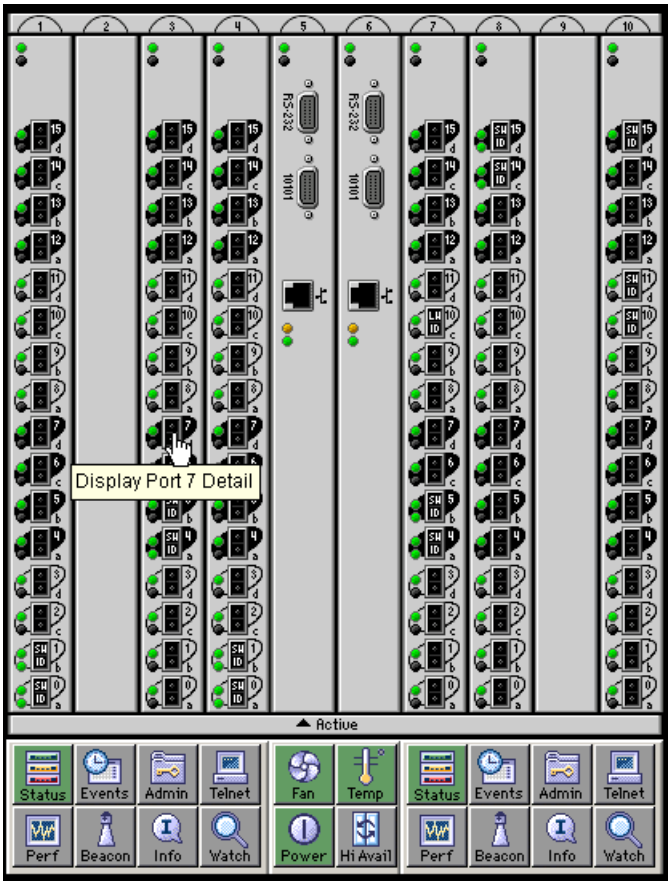


*Figure 1-53 2109-M12 Switch view*

This view shows only the physically installed port blades for each switch. We have port blade slots 1, 3, and 4 for switch 0; and port blades 7, 9, and 10 for switch 1. The Active CP is also indicated by the arrow below it.

Double-clicking a particular port gives us a view of the detailed information for that port, as seen in Figure 1-54.

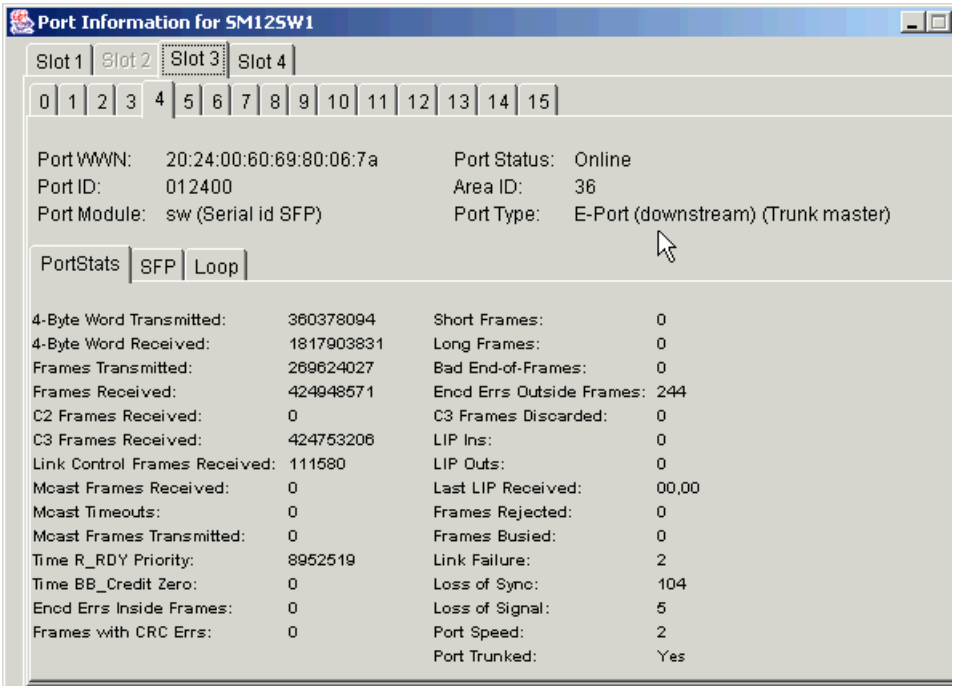


*Figure 1-54 Port detail view*

While the actual port information view is similar to other 2109 models, this display also has tabs along the top to select the particular port card slot installed in the logical switch. In our example, we chose logical switch 0, and then we double-clicked port 4 in slot 3. While this takes us directly to statistics for the port we selected, we can also now click any port or slot tab to view statistics for the other ports in this logical switch.

Other information from the switch view is available by clicking the appropriate button at the bottom of the view. In Figure 1-55, we can see that the buttons for an M12 are logically divided into three groups, as per the divider marks between the buttons. The first group of eight buttons on the left are for logical switch 0 (slots 1-4), the next group of four buttons in the center are for overall chassis functions, and the last group of eight buttons on the right are for logical switch 1 (slots 7-10).

While most of these buttons perform the same function on all switches, we chose to explain the Status button using the M12, as it also allows us to explain the buttons that only the M12 has.


*Figure 1-55 M12 Switch view buttons with failure*

In the following example we have caused a fan failure. This affects the chassis status and both of the individual switch statuses, indicated by both Status buttons and the Fan button being yellow.

## Status button

The Status button is available on all IBM TotalStorage SAN Switch models, and once for each logical switch on an M12. It displays the overall health status of each individual switch. The button will appear yellow if there is a marginal condition with the switch, or red if the switch is considered “down”. Clicking the button shows a text message box describing the current condition, as shown in Figure 1-56.

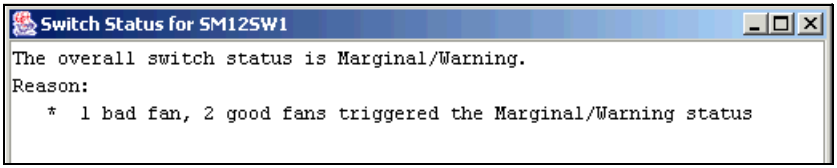


*Figure 1-56 Switch status display*

At a telnet command prompt, the same information could be displayed by entering **switchstatusshow**:

```
SM12SW1:admin> switchstatusshow
The overall switch status is Marginal/Warning
  Contributing factors:
   *  1 bad fan, 2 good fans triggered the Marginal/Warning status
```

## Fan button

The Fan button is an alerting icon on all models except the M12. If all conditions are normal according to the switch policy settings, the icon should be green. On the M12, it is a chassis wide status button.

Clicking the Fan button displays an informational window describing the state of each fan, as shown in Figure 1-57. In this example, as the chassis is designed to be able to run on two fans without any disruption, a single fan failure will show as a yellow marginal condition.

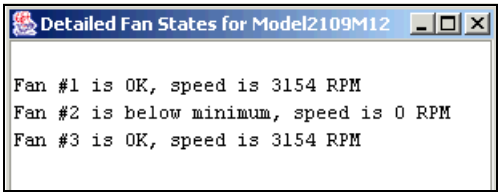


*Figure 1-57 M12 Fan detailed status*

It is possible to gather the same information from a telnet command line by typing **`fanshow`** if preferred:

```
SM12SW1:admin> fanshow
Fan #1 is OK, speed is 3154 RPM
Fan #2 is below minimum, speed is 0 RPM
Fan #3 is OK, speed is 3040 RPM
```

## Temp button

The Temp button is an alerting icon on all switch models except the M12. It will change color from green to show that all temperatures are within the defined limits, and to yellow or red depending on the policy thresholds. In our example we have a failed fan, therefore it is a wise idea to monitor the temperatures of the switch closer. On the M12, clicking the Temp button will display detailed temperature information for each slot in the chassis, and is shown in Figure 1-58.

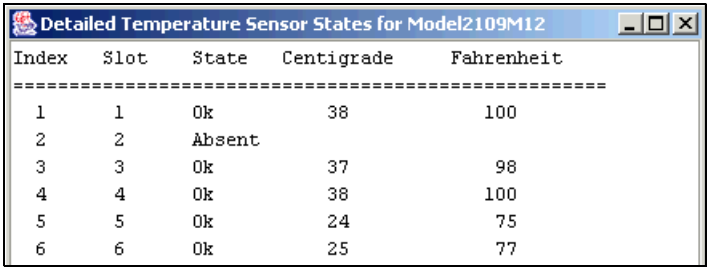


*Figure 1-58 Temperature detail display*

To display similar information at a telnet command line, we need to log into each logical switch separately and perform a **`tempShow`**:

### Logical switch 0 (SM12SW1)

```
SM12SW1:admin> tempShow

Index   Slot    State   Centigrade     Fahrenheit
=================================================
  1      1      Ok          38            100
  2      2      Absent
```

```
  3       3       Ok          38              100
  4       4       Ok          38              100
  5       5       Ok          24               75
  6       6       Ok          25               77
```

**Logical switch 1 (SM12SW2)**

```
SM12SW2:admin> tempShow

Index   Slot    State   Centigrade     Fahrenheit
==================================================
  1       5     Ok          24              75
  2       6     Ok          25              77
  3       7     Ok          36              96
  4       8     Ok          35              95
  5       9     Absent
  6      10     Ok          39             102
```

## 1.7.6 Admin button

In 1.6, “Installing the IBM TotalStorage SAN Switch” on page 24, we showed how to configure many settings using the Command Line Interface. Most of these settings may also be configured using the WEB TOOLS Administration Tools interface.

To perform administration and setup functions on a single switch, we select the appropriate switch from the fabric view, then from the switch view frame we click the Admin button as shown in Figure 1-59.

With an M12, choosing the Admin Button from the Left or the Right group of buttons works on only that logical switch.

**Tip:** We recommend checking the Name of the switch, found in the Admin view, to ensure that you are working on the correct switch.


*Figure 1-59 Select Admin*

### *Administration Tools window layout*

Once the administration window has opened, we can see it is composed of five areas, as shown in Figure 1-60.

**Tip:** By hovering the mouse over buttons and other areas of the window, we can find out their function.

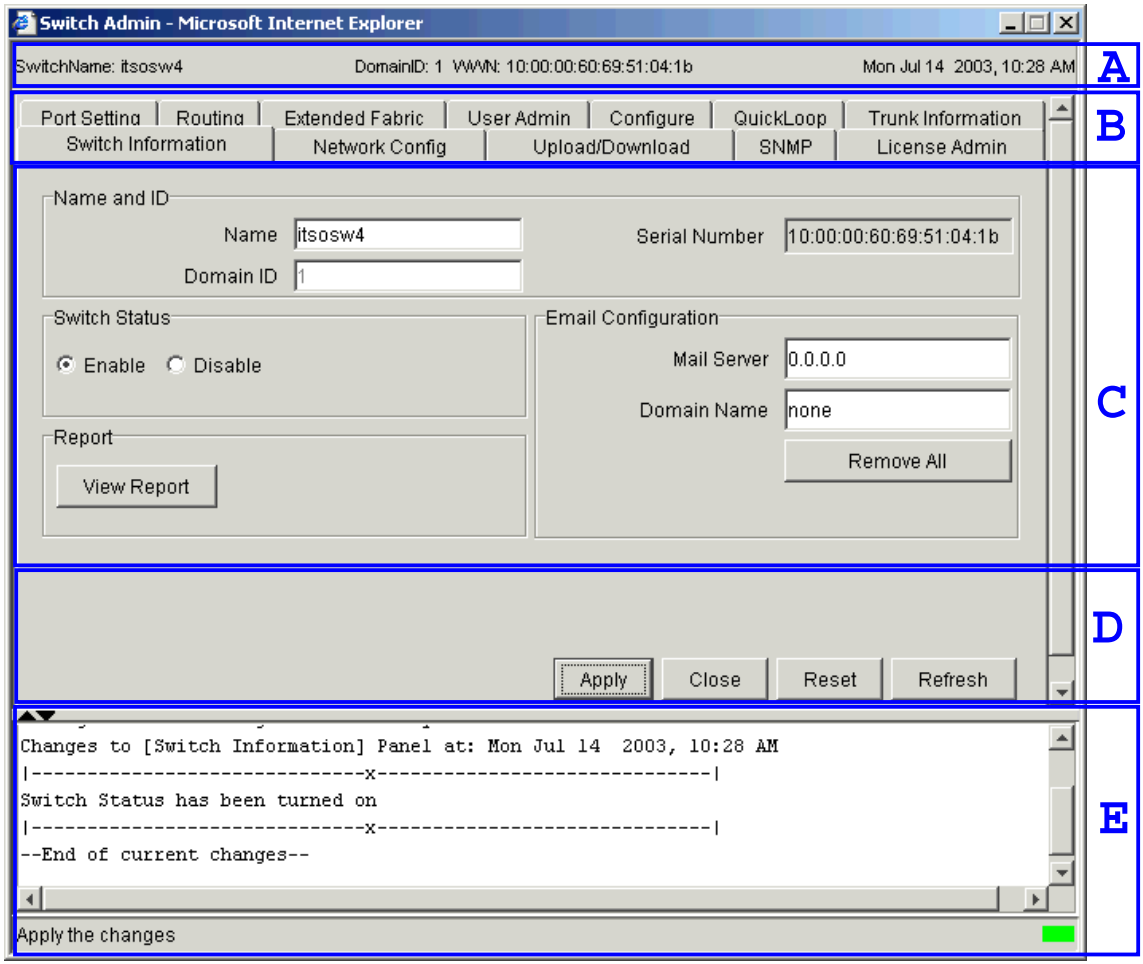


*Figure 1-60 Administration window layout*

- **Area A:** Displays summary information about the current switch.
- **Area B:** Allows navigation through the different management panels. The content of this area depends on the licenses installed on the switch.

- **Area C:** Contains parameters to be set in the current panel.
- **Area D:** Contains the buttons bar.
- **Area E:** Contains the report window that allows viewing of the switch report upon operation completion.

## Switch Information

When the administration window is first opened, the Switch Information tab is displayed by default. We show this in Figure 1-61.

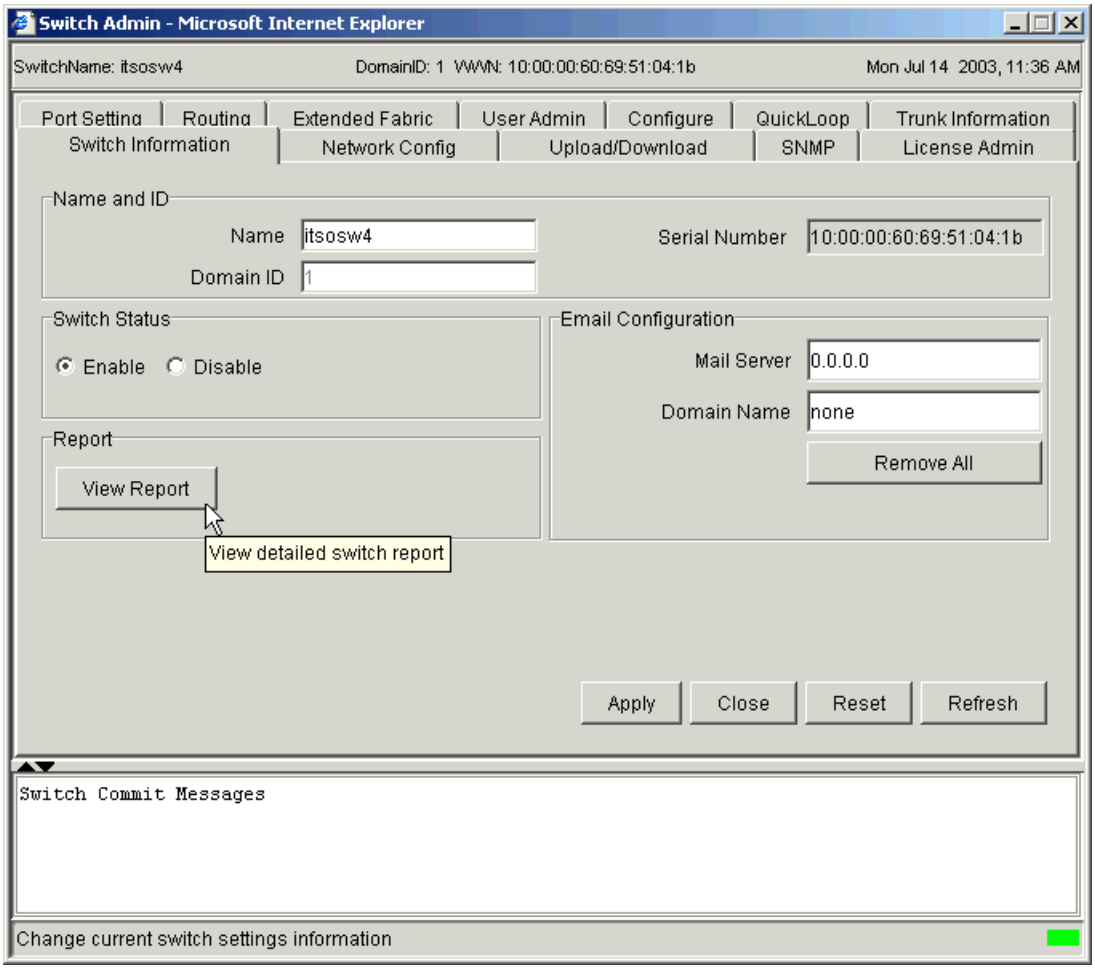


*Figure 1-61 Settings View*

On this first tab we can define the switch name and the domain ID, set the base e-mail configuration, enable or disable the entire switch, and view a detailed report of the switch.

Table 1-7 describes the fields on the **Switch Information** tab.

*Table 1-7 Switch Information tab*

| Field | Description |
|---|---|
| Name | Enter data for the switch name. Enter a new name to change a name in this field. |
| Domain ID | Displays or sets switch domain ID. Domain IDs must be unique within a fabric. To change domain ID, enter new domain ID in this field. Use a number from 1 to 239 for normal operating mode (FCSW compatible) and a number from 0 to 31 for VC encoded address format mode (backward compatible to SilkWorm 1000 series). |
| Serial Number | Displays the WWN of the switch. |
| (Status) Enable | Click the radio button to enable the switch. |
| (Status) Disable | Click the radio button to disable the switch. |
| Apply | Click to save any changes made to this tab and remain in the current tab. Additional changes can be made and the **Apply** button clicked when making changes incrementally. |
| Close | Click to exit the Switch Admin view. If changes have been made and not committed by clicking the **Apply** button, a dialog box is presented. It allows the changes to be committed or deleted. |
| Reset | Click to reset the tab to the last set of saved changes. |
| Refresh | Click to retrieve current values from the switch. |

### *View Report*

Clicking the **View Report** button will display the window shown in Figure 1-62. We do not show all of the detail here, although the detailed report on the Switch configuration includes a list all the types of switches connected to our local switch, the inter-switch links, list of ports, the Name Server information, details on the configured zones and SFP serial ID information.

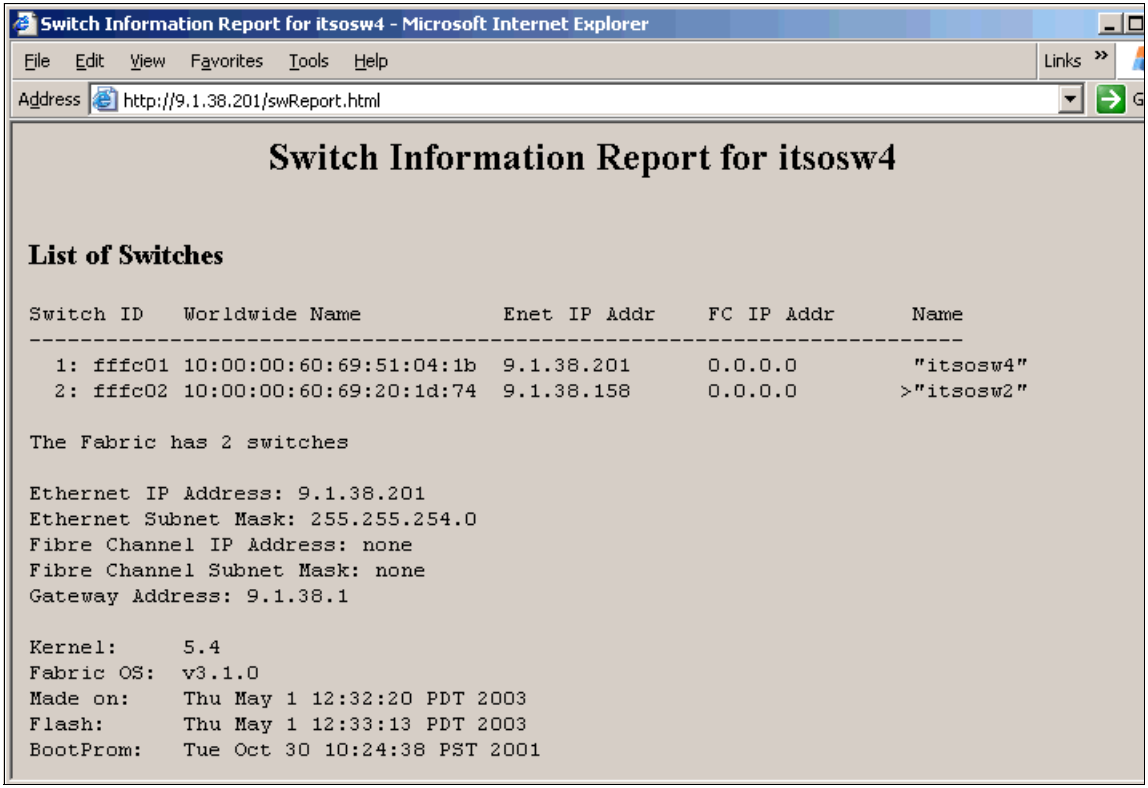


*Figure 1-62 Switch information report*

## Network Config

Use the **Network Config** tab to modify the IP settings of the switch as shown in Figure 1-63.

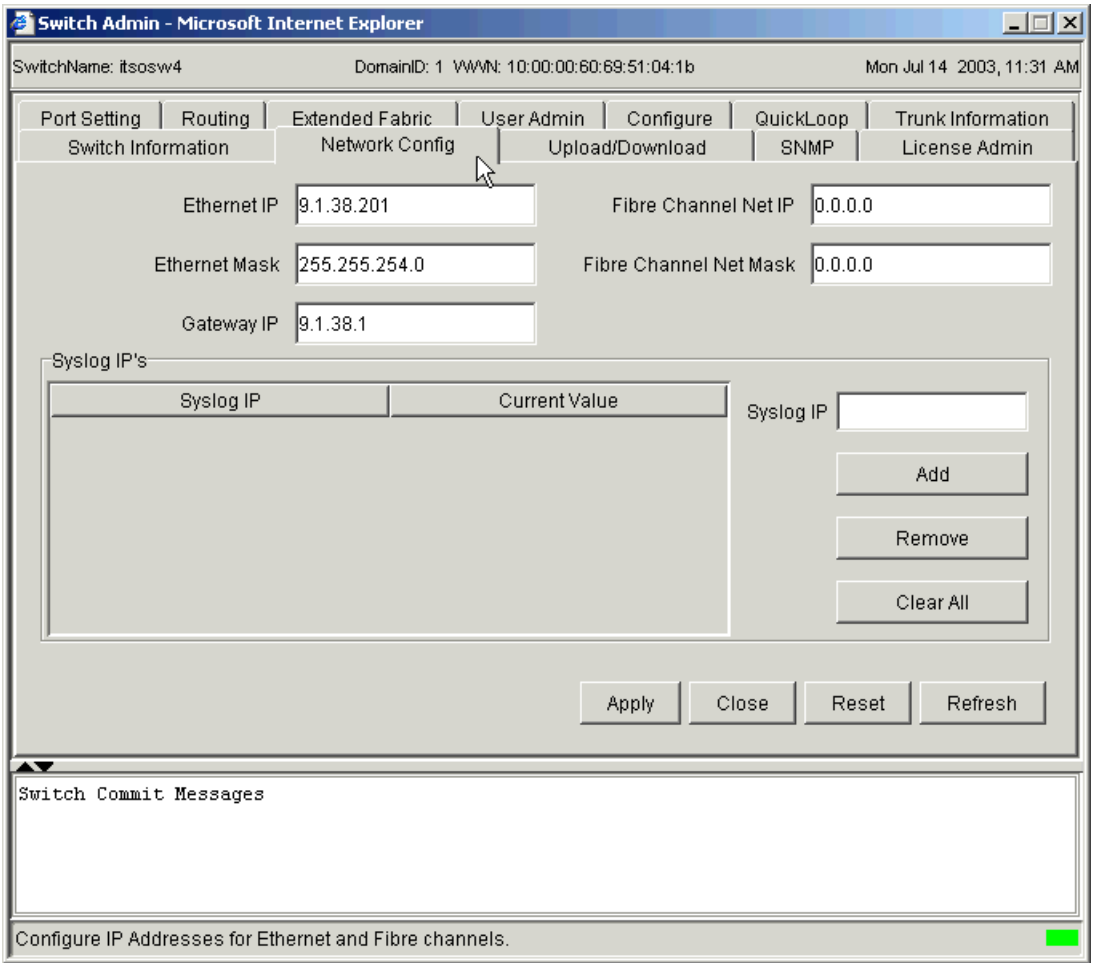


*Figure 1-63 Network configuration panel*

The lower section of the window is for configuring the Syslog daemon.

Table 1-8 describes the fields on the **Network Config** tab.

*Table 1-8 Network config tab*

| Field | Description |
|---|---|
| Ethernet IP | Displays or sets the Ethernet IP address |
| Ethernet Mask | Displays or sets the Ethernet IP Subnet Mask. |
| Gateway IP | Displays or sets the Gateway IP address. |
| Fibre Channel Net IP | Displays or sets the Fibre Channel IP address. |
| Fibre Channel Net Mask | Displays the Fibre Channel SubnetMask address. |
| Syslog IPs | Displays the six Syslog IP address for a user to configure. |
| Add | Add syslog IP address entered in field. |
| Remove | Remove syslog IP address in field. |
| Clear All | Remove all previous syslog IP entries. |
| Apply | Click to save the changes made to this tab and to stay in the current tab. Additional changes can be made and the **Apply** button clicked when making changes incrementally. |
| Close | Click to exit Admin window. If changes have been made but not committed by clicking the **Apply** button, a dialog box displays. |
| Reset | Click to reset the options to the last set of committed changes. If the **Apply** button has not been clicked on this tab, the parameters are returned to the original values the tab had when it was initially displayed. |
| Refresh | Click to retrieve current values from the switch. |

### *Overview of syslogd*

The Fabric OS maintains an internal log of all error messages, but the internal log buffers are limited in capacity; when the internal buffers are full, new messages overwrite old messages.

The IBM TotalStorage SAN Switch can be configured to send error log messages to a UNIX host system that supports **syslogd**. This host system can be configured to receive error/event messages from the switch and store them in its file system, overcoming the size limitations of the internal log buffers on the switch.

The host system can be running UNIX, Linux, or any other operating system as long as it supports standard **syslogd** functionality. The IBM TotalStorage SAN Switch by itself does not assume any particular operating system to be running on the host system.

To configure the syslog function, we simply put the IP address of the host running the **syslogd** in the **Syslog IP** field, and click **Add**. After adding all logging host IP addresses to the list, we must click **Apply** to save the changes.

### *Network Config on an M12*

When configuring the network settings on an M12 using this tab, extra care should be taken that we have opened the Admin function for the correct logical switch, as the settings only apply to that logical switch. There is also an extra button to allow setting the IP address and subnet mask for each CP, as shown in Figure 1-64.

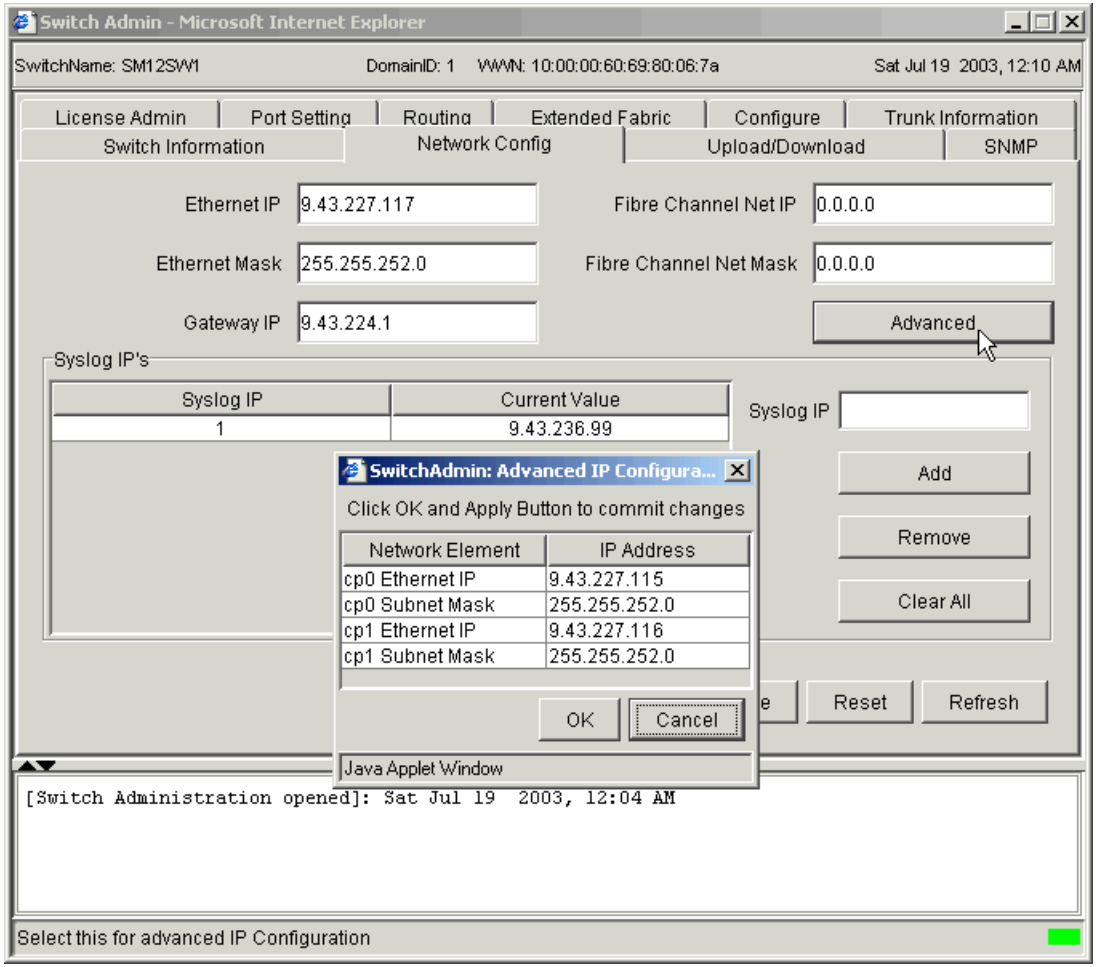


*Figure 1-64 Admin View — Network Config*

Selecting the Advanced button takes us to a window where we are able to set the ethernet IP addresses for both of the CP cards.

These same settings were configured earlier by using the command line install procedure, detailed in “M12 configuration procedure” on page 29.

## Upload / download

As shown in Figure 1-65, we can use this tab in order to:

- Upgrade the firmware of the switch.
- Save the switch configuration file to a host running FTP server for future restore to the switch.
- Download a previously saved configuration file from a host running FTP or RSHD.
- Set the switch configuration back to default.
- Reboot or Fastboot the switch.

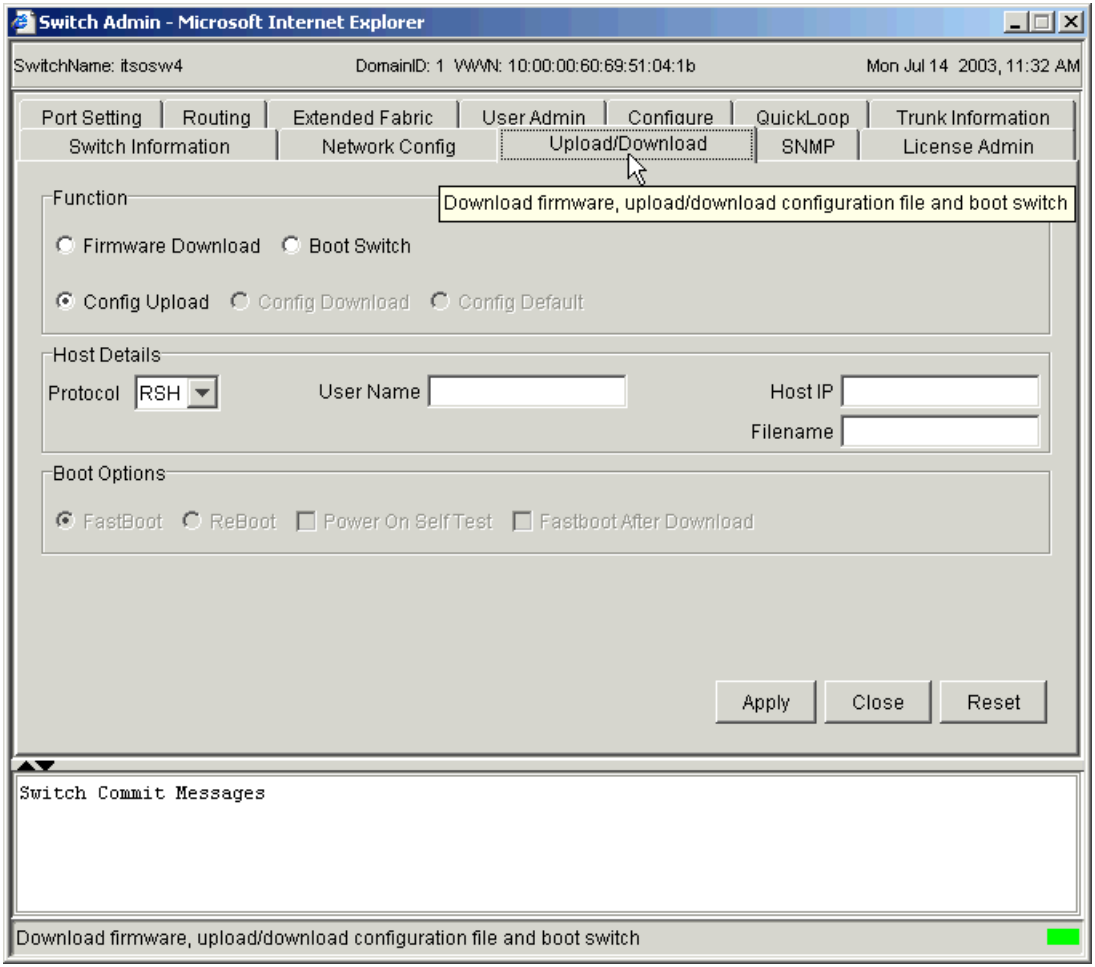


*Figure 1-65 Upload / Download panel*

When uploading the configuration file to an FTP server, make sure that the user has read/write access for the specified file path. The saved file includes current operating parameters, SNMP settings, as well as zoning configuration.

POST is the Power On Self Test and takes about two minutes to complete. A series of commands can be executed to test the switch. The Fabric OS POST includes the following tests:

- ramTest
- portRegTest
- centralMemoryTest
- cmiTest
- camTest
- portLoopbackTest

In Table 1-9 we describe the fields on the **Upload / Download** tab.

*Table 1-9 Upload / Download tab*

| **Available in the Function section:** | |
|---|---|
| Firmware Download | Select the radio button to download firmware. |
| Bootswitch | Select the radio button to boot the switch. |
| Config Upload | Select the radio button to upload the configuration file to the specified host. This allows for saving of the configuration file to the switch using the specified filename (full path). The User and Password must be valid for the specified host, and the file path must be read-write capable by the user. |
| ConfigDownload | Select the radio button to restore a previously saved configuration file (switch must be disabled). |
| Config Default | Select the radio button to reset the configuration to the default value (switch must be disabled). |
| **Available in the Host Details section:** | |
| Protocol | Select the pull-down menu to chose the downloading protocol (for F32 and M12 there is FTP only). |
| User name | Enter the rsh or ftp User Name. |
| Host IP | Enter the IP address of the rshd or ftp server. |
| Filename | Enter the filename to be downloaded including the full path. |

| Available in the Boot Options section: | |
|---|---|
| FastBoot | Select the radio button to do a Fastboot on the switch. This will cause the switch to skip over the POST test. |
| Reboot | Select the radio button to reboot the switch. |
| Power on Self-Test | Check the box to have the switch do a POST test when it reboots. |
| Fastboot After Download | Check the box to have the switch fastboot automatically after the firmware download completes. |
| Apply | Select to execute the choices made within this tab. |
| Close | Select to exit the Admin window. |
| Reset | Select to reset the tab to the last set of applied changes. |

## SNMP

Use the **SNMP** tab for administration of the SNMP Subsystem. From the **SNMP** tab we can specify the switch community string, location, trap level, and trap recipients.

We can also set SNMP parameters with Telnet using the **agtcfgSet** command and the **agtcfgShow** command to display the current SNMP settings.

**Note:** In order for the switches to send SNMP traps, we must first enter the Telnet command **snmpMibCapSet**. This enables the MIBs on all switches to be monitored.

```
itsosw4:admin> snmpMibCapSet
The SNMP Mib/Trap Capability has been set to support
 FE-MIB SW-MIB
FA-MIB (yes, y, no, n): [no] y
SW-TRAP (yes, y, no, n): [no] y
FA-TRAP (yes, y, no, n): [no] y
SW-EXTTRAP (yes, y, no, n): [no] y
Committing configuration...done.
```

The **SNMP** tab is shown in Figure 1-66.

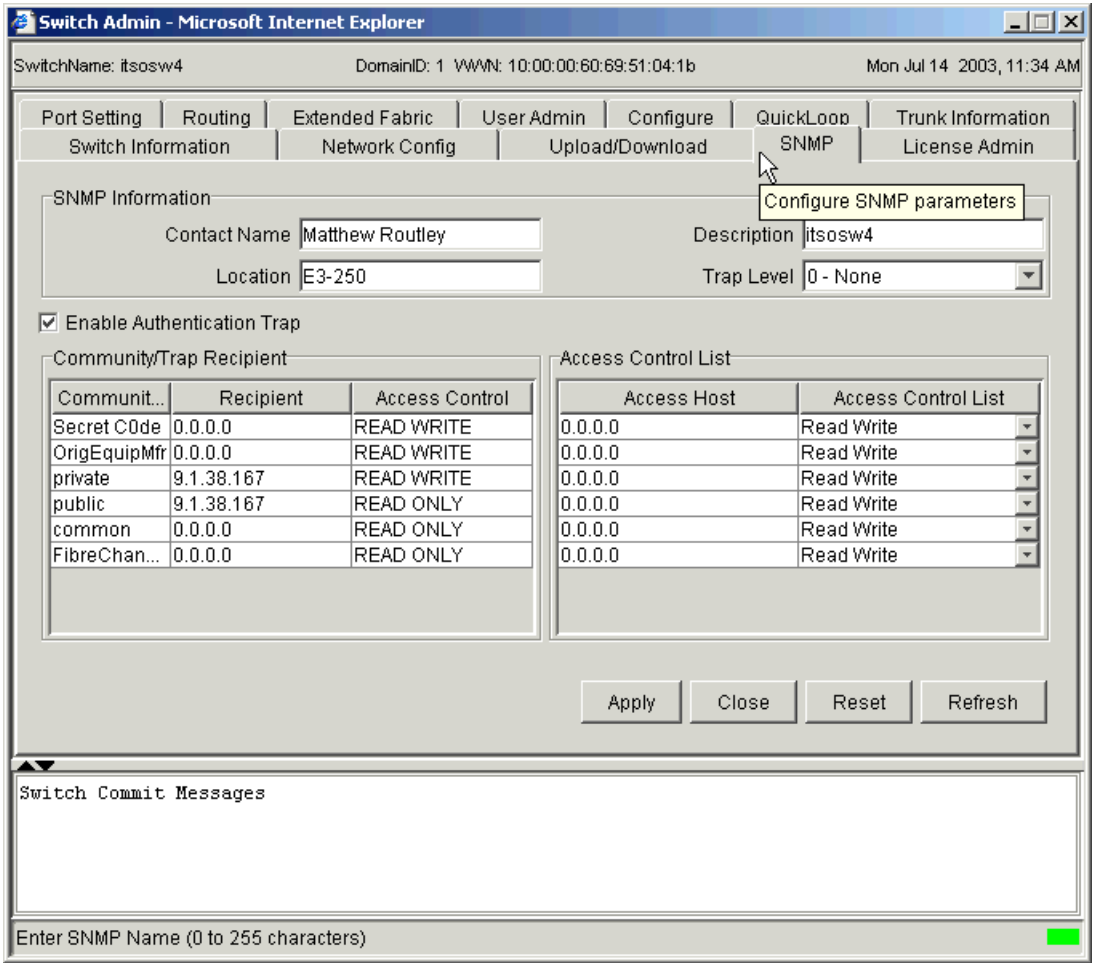


*Figure 1-66 SNMP configuration window*

In Table 1-10 we describe the fields on the **SNMP** tab.

*Table 1-10 SNMP tab*

| **Basic information:** | |
|---|---|
| Name | Displays or sets contact information for switch. Default is Field Support. |
| Location | Displays or sets the location of switch. Default is End User Premise. |
| Description | Displays or sets system description. Default is Fibre Channel Switch. |

| | |
|---|---|
| Trap Level | Sets severity level of switch events that prompt SNMP traps. Default is 0. |
| Enable Authentication Trap | Check to enable authentication traps; uncheck to disable (recommended). |
| **Community and Trap Recipient Configuration:** | |
| Community String | Displays the community strings that are available to use. A community refers to a relationship between a group of SNMP managers and an SNMP agent, in which authentication, access control, and proxy characteristics are defined. A maximum of six community strings can be saved to the switch. |
| Recipient | Displays the IP address of the Trap Recipient. A trap recipient receives the message sent by an SNMP agent to inform the SNMP management station of a critical error. |
| Access Control | Displays the Read/Write access of a particular community string. Read only access means that a member of a community string has the right to view, but cannot be changed. Read/Write access means that a member of a community string can be both viewed and changed. |
| **Access Control List Configuration:** | |
| Access Host | Displays the IP address of the host of the access list. |
| Access Control List | Displays the Read/Write access of a particular access list. Read only access means that a member of an access list has the right to view, but cannot make changes. Read/Write access means that a member of an access list can both view and make changes. |
| Apply | Click to save the changes made to this tab. Additional changes can be made and the **Apply** button clicked when making changes incrementally. |
| Close | Click to exit the Admin Window. If changes have been made but not committed by clicking the **Apply** button, a dialog box displays. |
| Reset | Click to reset the tab to the last set of committed changes. If the **Apply** button has not been clicked on this tab, the parameters are returned to the original values the tab had when it was initially displayed. |
| Refresh | Click to retrieve current values from the switch. |

## License admin

We use the **License Admin** tab to install license keys that have been purchased. License keys are used to enable additional features on a switch. We can also use the table within the **License Admin** tab to remove a listed license from the switch. The **License Admin** tab is shown in Figure 1-67.

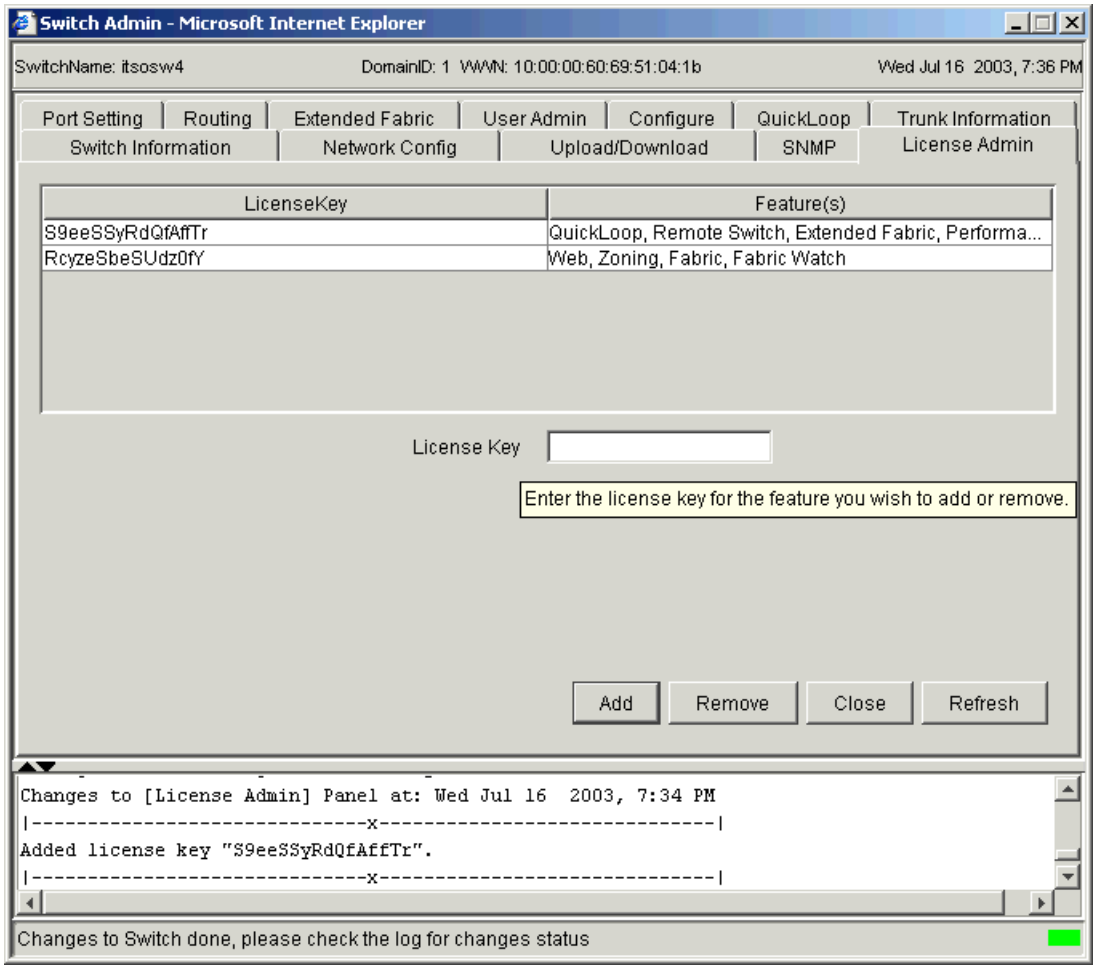


*Figure 1-67 License Admin window*

In Table 1-11 we describe the fields on the **License Administration** tab.

*Table 1-11 License admin tab*

| Field | Description |
|---|---|
| Feature(s) | A list of the licenses installed on the switch. |
| License Key | Enter license key to be added or removed. |
| Add | Select to add the specified license. |
| Remove | Select to remove the specified license. |
| Close | Select to exit the Admin Window. |
| Refresh | Click to retrieve current values from the switch. |

## Installing a license key through CLI

To install a license key feature using the CLI, perform the following steps:

1. From a command prompt, use the Telnet command to log onto the switch using an account that has administrative privileges. For example:

   **C:\telnet** *address*

   Here, *address* is replaced with switch IP address.

2. To determine which licenses are already installed on the switch, type **licenseShow** at the command line.

   A list displays of all the licenses currently installed on the switch. For example:

   ```
   admin> licenseShow
   1A1AaAaaaAAAA1a:
   Release v3.0
   Web license
   Zoning license
   SES license
   Trunking license
   ```

3. To install a license key enter the following on the command line:

   **licenseAdd "**key**"**

   Here, "key" is the license key provided to you, enclosed in double quotes. The license key is case sensitive and must be entered exactly as given.

4. Verify the license was added by entering the following on the command line:

   **licenseShow**

   If the license is listed, the feature is installed and immediately available.

   If the license is not listed, repeat step 3.

## Port Setting

Clicking the **Port Setting** tab will display the panel shown in Figure 1-68.

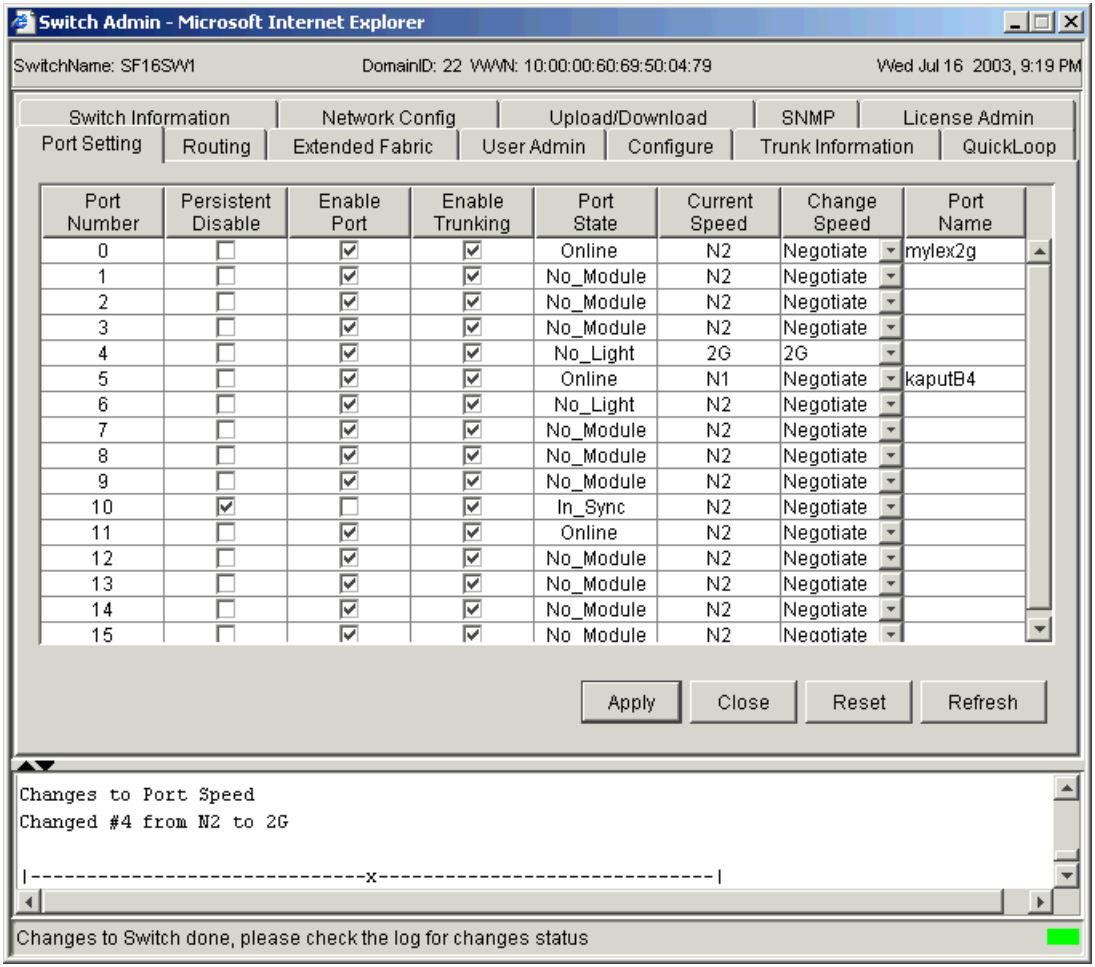


*Figure 1-68 Port Setting panel*

In this panel you can:

- Set or reset a persistent Disable per port
- Disable or enable a specific port
- Disable or enable trunking for a specific port (default value is enabled)
- View the current port state
- View the current speed for the switch ports
- Manually set the speed for a specific port
- Define a symbolic name to identify what is attached to the port

**Note:** On 2109-F16 and 3534-F08, the trunking column will only be present if a trunking license has been purchased and applied to the switch. 2109-F32 and M12 have the trunking feature installed by default.

Table 1-12 describes the fields on the **Port Settings** tab.

*Table 1-12 Port settings tab*

| Field | Description |
|---|---|
| Port Number | The port number. |
| Persistent Disable | Check to disable port, remains disabled thru switch reboots and power cycles. UnCheck to enable the port. |
| Enable Port | Check to disable the port, uncheck to enable. At power on or reboot the port will be enabled. |
| Enable Trunking | Check to enable the port trunking. Four trunk ports form a group, with one of them in the role of master port. |
| Port State | Displays the current state of each port. |
| Current Speed | Displays the Speed of the port connection. N1, N2 indicated autonegotiated to 1 or 2 Gb. 1G or 2G indicates manually set speed. |
| Change Speed | Click to select port speed can be fixed to 1G, 2G, or auto-negotiate. If the speed is set to negotiate, the speed will depend on the negotiated result. |
| Port Name | Double Click in this field to enter a symbolic name fore the port. |
| Apply | Click to save the changes made to this tab. Additional changes can be made and the **Apply** button clicked again to make changes incrementally. |
| Close | Click to exit the Admin window. If changes have been made but not committed by clicking the **Apply** button, a confirmation box displays. |
| Reset | Click to reset the tab to the last set of committed changes. If the **Apply** button has not been clicked on this tab, the parameters are returned to the original values the tab had when it was initially displayed. |
| Refresh | Click to retrieve current values from the switch. |

## Routing

Clicking the **Routing** tab displays the panel shown in Figure 1-69.

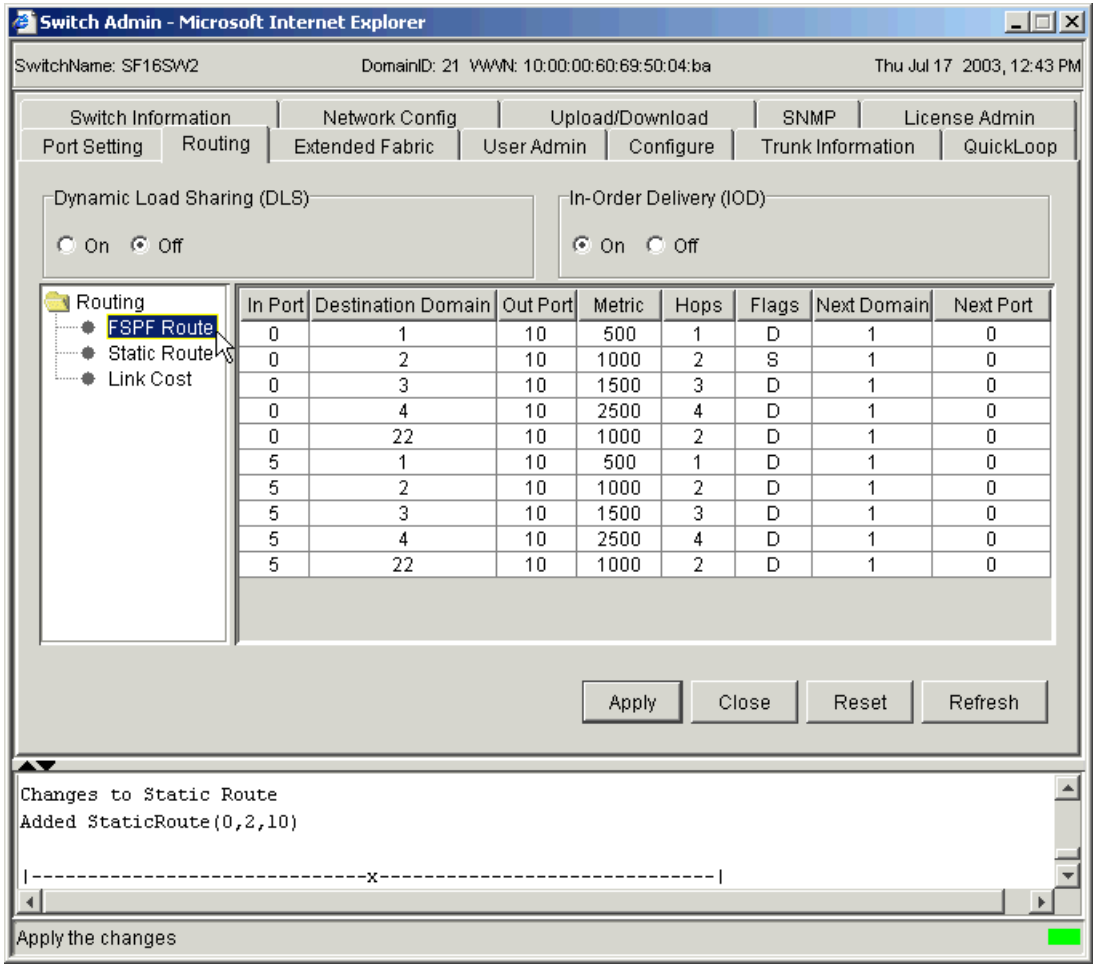


*Figure 1-69 Routing panel with FSPF selected*

In the following topics we discuss the various settings available in Figure 1-69.

### *Dynamic Load Sharing (DLS)*

Routing is generally based on the incoming port and the destination domain. This means that all the traffic coming in from a port (either E_Port or Fx_Port) directed to the same remote domain is routed through the same output E_Port.

To optimize fabric routing, when there are multiple equivalent paths to a remote switch, traffic is shared among all the paths. Load sharing is recomputed when a switch is booted up or every time a change in the fabric occurs. A change in the fabric is defined as an E_Port going up or down, or an Fx_Port going up or down.

In an IBM fabric, if DLS is turned off, load sharing is performed only at boot time or when an Fx_Port comes up. Optimal load sharing is rarely achieved with DLS disabled.

If DLS is turned on, routing changes can affect working ports. For example, if an Fx_Port goes down, another Fx_Port may be rerouted from one E_Port to a different E_Port. The switch minimizes the number of routing changes, but some are necessary in order to achieve optimal load sharing.

Turning on DLS can affect performances when using it in conjunction with the In-Order Delivery option.

### *In-Order Delivery (IOD)*

Use the IOD option to enforce in-order delivery of frames during a fabric topology change.

In a stable fabric, frames are always delivered in-order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for instance, a link goes down), traffic is rerouted around the failure. When topology changes occur, generally, some frames are delivered out-of-order. This option ensures that frames are not delivered out-of-order, even during fabric topology changes.

In an IBM fabric, the IOD option is to be set on.

This option should be used with care, because it can cause a delay in the establishment of a new path when a topology change occurs. Only if there are devices connected to the fabric that do not tolerate occasional out-of-order delivery of frames, should this command be used.

### *FSPF Route*

As shown in Figure 1-69, the FSPF Route option is selected (highlighted) under the *Routing* tree. The main area of the window then displays the FSPF routing table, including the destination domain and port, hop count, and the metric being the cost assigned to that link. We define the different columns in Table 1-13.

*Table 1-13 FSPF Route Field Descriptions*

| Field | Description |
|---|---|
| In Port | Displays the Port number where the frames enter the switch. |
| Destination Domain | Displays the destination domain ID for the participating static routes for a particular In Port. The destination domain is the target of the out port. |
| Out Port | Displays the Out port. It should be within the range of ports that are available for static routes for the current domain. More than one out port can be used for any In port with a different domain id. Each domain id requires an out port. |
| Metric | Displays the calculated cost of reaching the destination domain. |
| Hops | Displays the number of hops in the “shortest path” route. |
| Flags | Displays whether the route is Static (*S*) or Dynamic (*D*). |
| Next Domain | Displays the next domain ID in the routing path. The Next Domain is the switch that the “Out Port” is connected to. |
| Next Port | Displays the next Port in the routing path. The Next Port is the port number that the “Out Port” is physically connected to. |

### *Static Route*

This section can be used to define static routes. A static route is a route that is defining a specific path, and will not change when a topology changes occur, unless the path defined by the route becomes unavailable.

In Figure 1-70 we are defining a static route so that all frames received on port *0* with a destination domain of *2* will be transmitted through port *10*. Clicking **OK** will add our definition to the list. We then need to click **apply** to bring this definition active; the active definition can be seen in the FSPF routing table in Figure 1-69 identified by the *S* flag. To remove a static route, we need to select the specific definition in the static routs list and then click **Delete**.

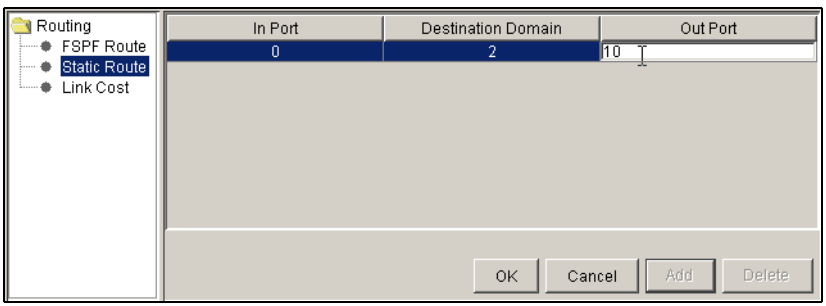


*Figure 1-70 Routing - Static Route*

### *Link Cost*

By selecting the next option under the *Routing* tree, we can view the link cost for a specific link. By double-clicking in the *Cost* field for the specific port, we are able to modify the cost; we show this in Figure 1-71. This setting will have an effect on the cost value the local switch has for this link. It will use this value to calculate the lowest cost path to a destination on other switch(es) within the fabric. For a 1 Gb/s per second ISL, the default cost is 1000. For a 2 Gb/s ISL, the default cost is 500. Valid values for link cost are from 1 to 9999.

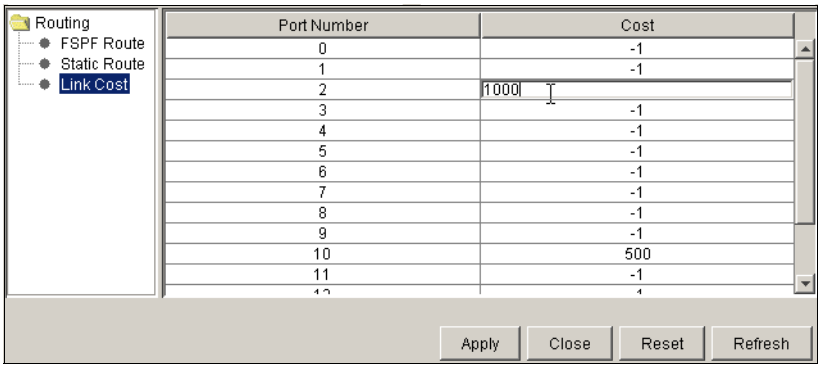


*Figure 1-71 Routing - Link Cost*

## Extended Fabric

This Tab is only present if an optional Extended Fabric License has been purchased and the key installed. Clicking the **Extended Fabric** tag displays the panel shown in Figure 1-72.

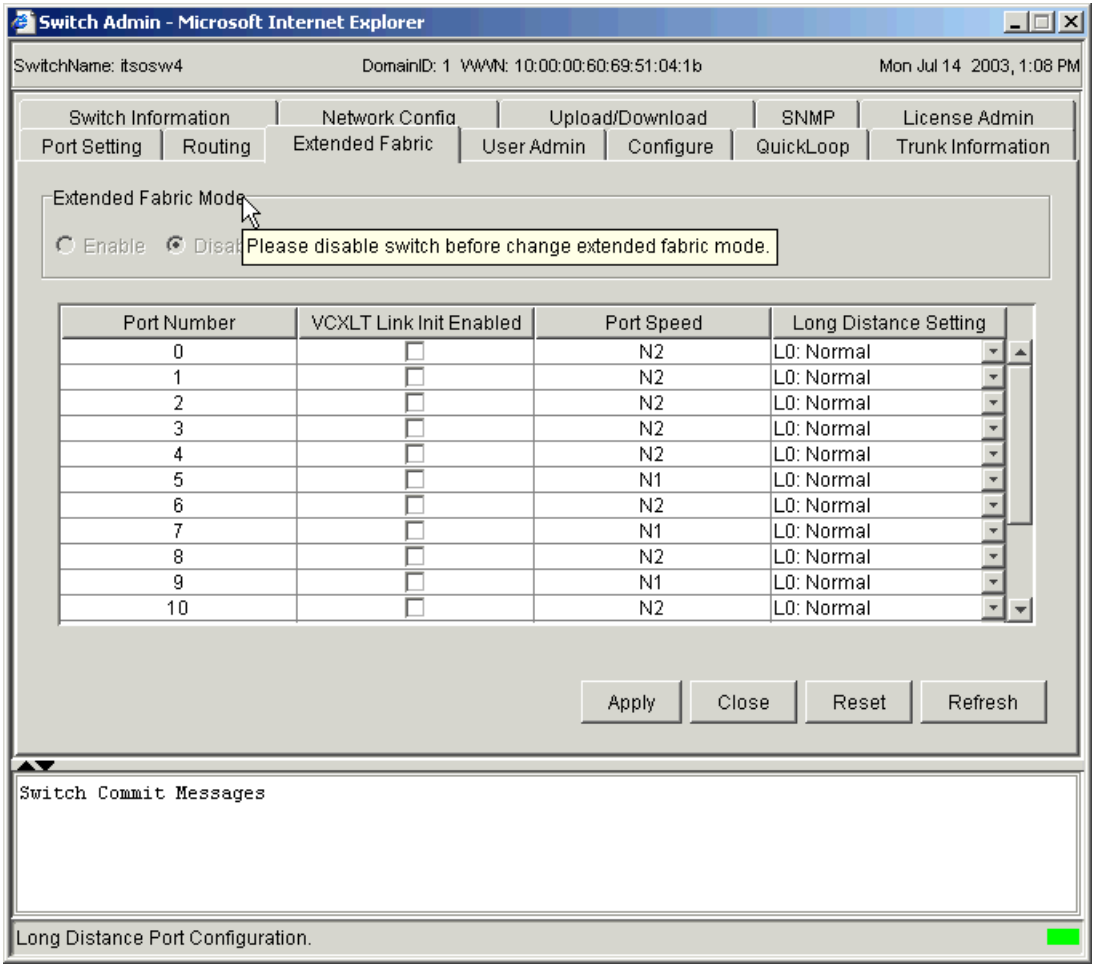


*Figure 1-72 Extended Fabric panel*

Use the **Extended Fabric** tab to manage the Extended Fabric feature. From the Extended Fabric tab, we can specify which ports are to be configured for distance and at what level. The **Extended Fabric** feature can only be enabled or disabled when the switch is in the Disabled condition. For ports that are disabled, the rows appear grayed-out in the table within the Extended Fabric tab. We discuss the Extended Fabric feature in greater detail in “Extended Fabrics” on page 206.

Table 1-14 describes the fields on the **Extended Fabric** tab.

*Table 1-14 Extended fabric tab*

| Field | Description |
|---|---|

| | |
|---|---|
| Extended Fabric Mode | Click the box to allow ports to be configured for long distance, or uncheck to turn the option off. |
| Port Number | Port Number being used for the Extended Fabric. |
| VCXLT Link init enabled | The option “VC Translation Link Init” is to turn on the long distance link initialization sequence. This option defaults to off. |
| Long Distance Setting | Click to view Long Distance settings. |
| Apply | Click to save the changes made to this tab and to stay in the current tab. Additional changes can be made and the **Apply** button clicked when making changes incrementally. |
| Close | Click to exit the Admin Window. If changes have been made but not committed by clicking the **Apply** button, a dialog box displays. |
| Reset | Click to reset the tab to the last set of committed changes. If the **Apply** button has not been clicked on this tab, the parameters are returned to the original values the tab had when it was initially displayed. |
| Refresh | Click to retrieve current values from the switch. |

## User Admin

To go to the User Admin window, click the **User Admin** tab as shown in Figure 1-73.

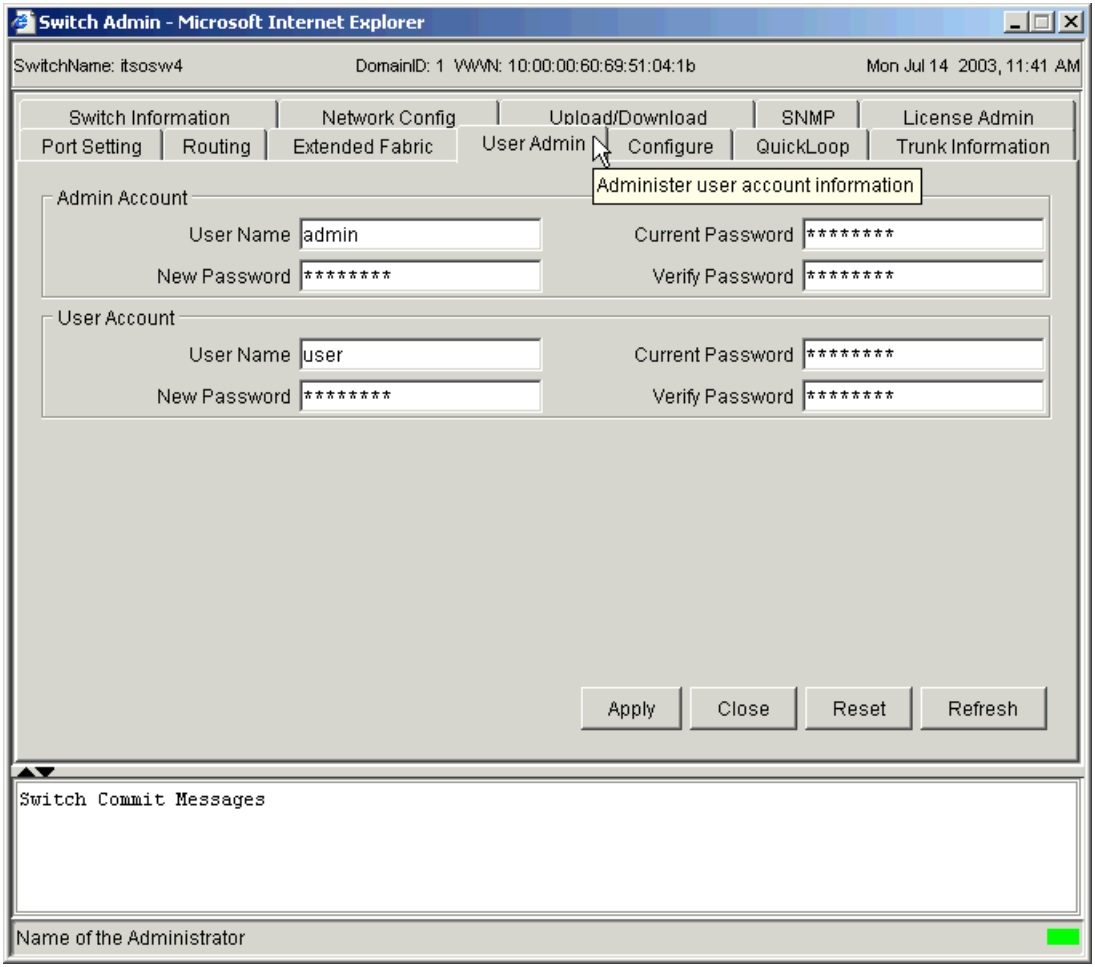


*Figure 1-73 User Admin*

From this window, we can change the user names and passwords that allow access to the switches from the TotalStorage Switch Specialist.

To change a User Name, we overtype the entry in the User Name field and specify the current admin password and the new password in both the *New Password* and *Verify Password* fields.

**Restriction:** Changing the User Name does $NOT$ create additional users, it is only changing the existing ID to a new name.

If we only wish to change the password of the admin user, we enter our current password, and the new password into both the *New Password* and *Verify Password* fields. Clicking **Apply** validates the changes, and they are reported to the report window as shown in Figure 1-74.

After the changes are successfully committed to the switch, we are required to log back in using the new user and / or password.

### *Admin access level*

This access level allows change and view access to all functions. From telnet access, the Admin level allows use of all commands within the Help Menu. Typically, most switch administration is performed at this level.

### *User access level*

This access level provides view access only. Users are not able to make zoning changes or any switch configuration changes. This level is recommended for monitoring switch activity.

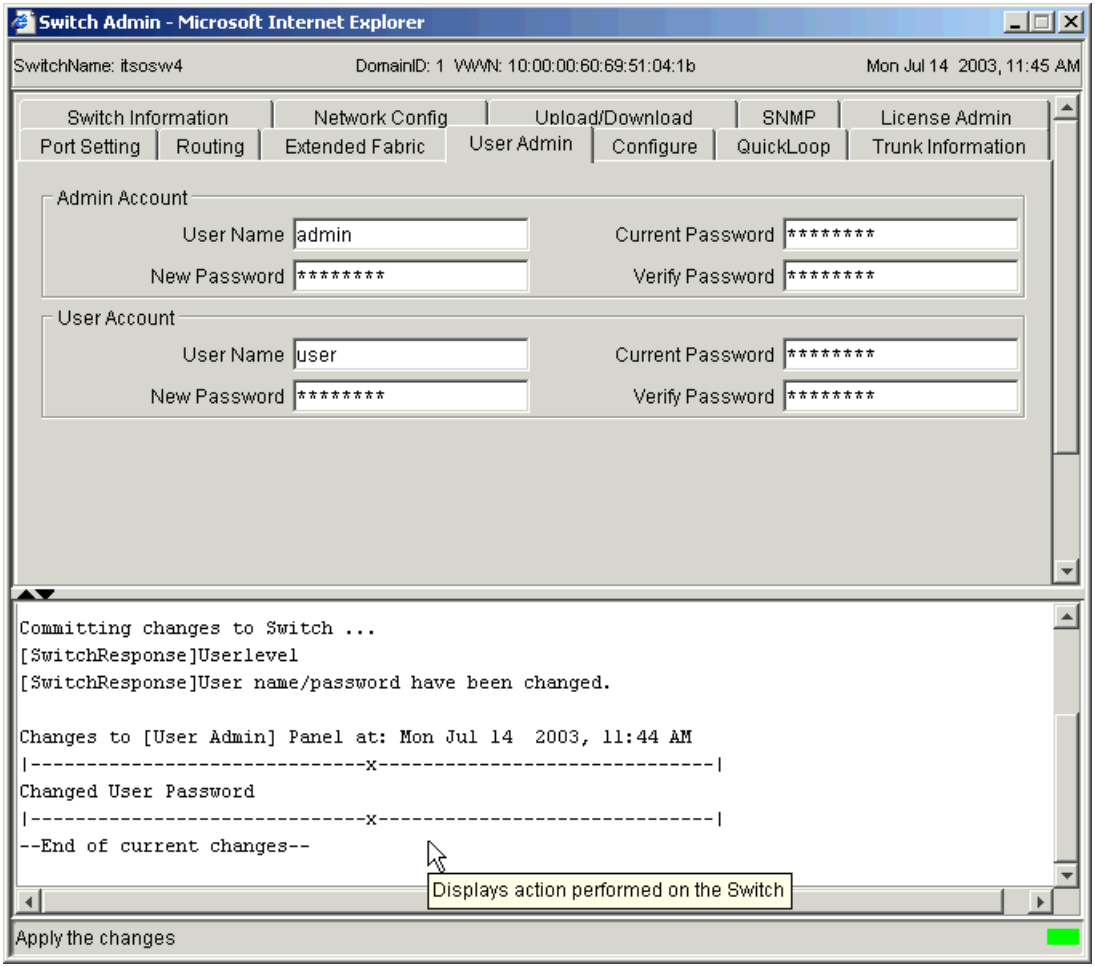


*Figure 1-74 Users information changes*

Table 1-15 describes the fields on the **User Admin** tab.

*Table 1-15 User admin tab*

| Field | Description |
|---|---|
| User Name | Enter new user name or modify the existing user name. |
| Current Password | Enter the current password. |
| New Password | Enter a new password. |
| Verify Password | Re-enter new password to verify. |
| Apply | Click to save the changes made to this tab. Additional changes can be made and the **Apply** button clicked when making changes incrementally. |
| Close | Click to exit the Switch Administration view. If changes have been made but not committed by clicking the **Apply** button, a dialog box displays to ask the user if they want to save the changes before exiting the view. |
| Reset | Click to reset the tab to the last set of committed changes. If the **Apply** button has not been clicked on this tab, the parameters are returned to the original values the tab had when it was initially displayed. |
| Refresh | Click to retrieve current values from the switch. |

## Configure

Clicking the **Configure** tab displays the panel shown in Figure 1-75. We are unable to make any changes to the settings on this tab if the switch is enabled.

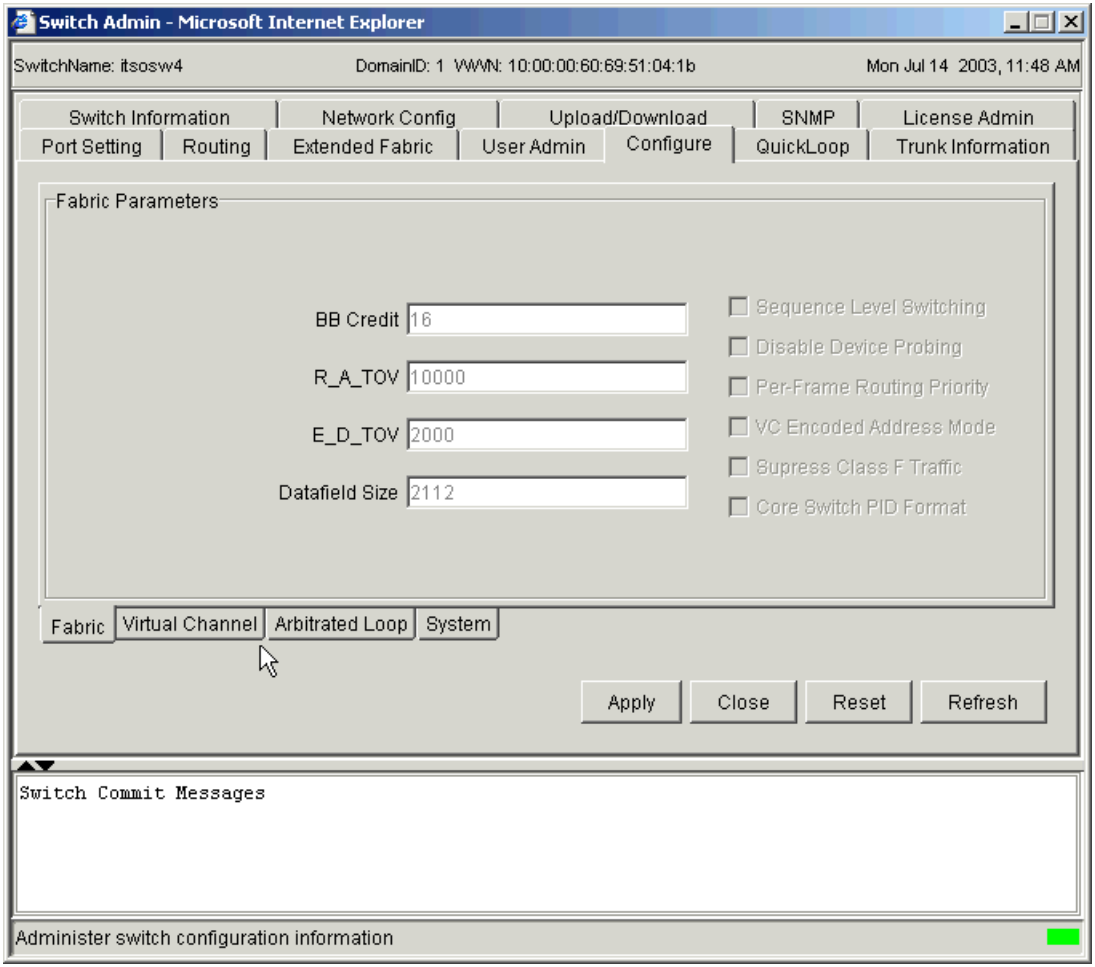


*Figure 1-75 Configure panel*

The following paragraphs describe the different parameters found on the sub-tabs shown in Figure 1-75.

### *Fabric parameters*

The Fabric parameters available are:

- `BB Credit:` The buffer-to-buffer (BB) credit represents the number of buffers available to attached devices for frame receipt. This value ranges from 1 to 27. Default value is 16.

- `R_A_TOV`: The Resource Allocation Time Out Value (R_A_TOV) is displayed in milliseconds. Allocated circuit resources with detected errors are not released until this time value has expired. If the condition is resolved prior to the time out, the internal time out clock resets and waits for the next error condition.
- `E_D_TOV`: Error Detect Time Out Value (E_D_TOV) is displayed in milliseconds. This timer is used to flag a potential error condition when an expected response is not received (an acknowledgment or reply in response to packet receipt, for example) within the set time limit. If the time for an expected response exceeds the set value, then an error condition occurs.
- `Datafield Size`: The largest data field size in bytes.
- `Sequence Level Switching`: When Sequence Level Switching is enabled, frames of the same sequence from a particular source are transmitted together as a group. When this feature disabled, frames are transmitted interleaved among multiple sequences. Under normal conditions, Sequence Level Switching should be disabled for better performance.
- `Disable Device Probing`: When Disable Device Probing is enabled, devices that do not register with the Name Server are not present in the Name Server data base. Set this mode only if the switch N_Port discovery process (PLOGI, PRLI, INQUIRY) causes an attached device to fail.
- `Per-Frame Routing Priority`: In addition to the eight virtual channels used in frame routing priority, support is also available for per-frame based prioritization when this value is set. When Per-frame Route Priority is enabled, the virtual channel ID is used in conjunction with a frame header to form the final virtual channel ID.
- `VC Encoded address Mode`: When enabled, frame source and destination address utilize an address format compatible with older switches. Also known as Silkworm 1000 compatibility mode, this mode is required for Numa-Q.
- `Suppress Class F Traffic`: When enabled, all class F interswitch frames are transmitted as class 2 frames. This is to support remote fabrics which involve ATM gateways which don't support class F traffic.
- `Core Switch PID Format`: When enabled, allows full 256 port addressing that is used for core switches. This parameter must be set the same on all switches in the fabric, for more information refer to “Setting Core PID format” on page 37.

### *Virtual Channels parameters*

This feature enables fine tuning of Inter Switch Links by configuring parameters for the eight virtual channels. These parameters are used for congestion control. We recommend to leave the default values for these parameters alone unless expert advice is available.

### *Arbitrated Loop parameters*

These are the Arbitrated Loop Parameters:

- `Send Fan Frames`: Specifies that fabric address notification (FAN) frames be sent to public loop devices to notify them of their node ID and address. When enabled, frames are sent; when disabled, frames are not sent.
- `Always send RSCN`: Following the completion of loop initialization, a remote state change notification (RSCN) is issued when FL_Ports detect the presence of new devices or the absence of pre-existing devices. When this mode is enabled, a RSCN is issued upon completion of loop initialization, regardless of the presence or absence of new or preexisting devices.
- `Do Not Allow AL_PA 0x00`: This option disallows AL_PA values from being 0.

### *System Services parameter*

The System Services parameter lets you set activity monitoring on the switch:

- `rstatd`: Allows you to dynamically enable or disable a server that returns details about system operation information through remote procedure calls (RPCs). Note that only Ethernet statistics and system up time are supported. The retrieval of this information is supported by a number of operating systems. For example, most UNIX-based systems use the **`rup`** or **`rsysinfo`** command to retrieve the information. Figure 1-76 shows an example using the **`rup`** command on an AIX host.

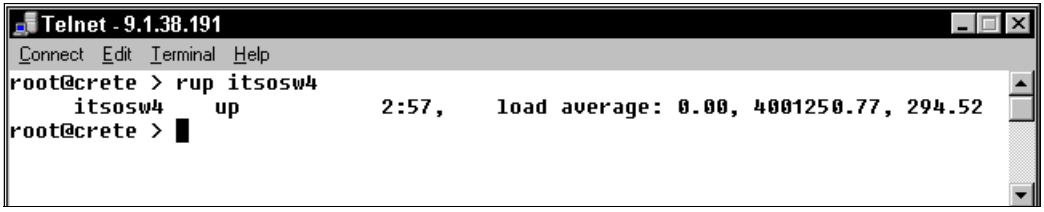


*Figure 1-76 Retrieving switch information using rup and rstatd*

- `rapid`: Allows you to dynamically enable or disable a service that handles RPC requests for the API server.
- `rusersd`: Allows you to dynamically enable or disable a server that returns information about the user logged into the system through RPC. The retrieval of this information is supported by a number of operating systems. For example, most UNIX-based systems use the **`rusers`** command to retrieve the information. Figure 1-77 shows an example using the **`rusers`** command on an AIX host.

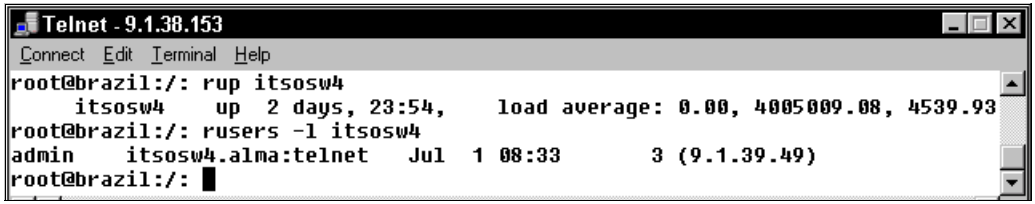


*Figure 1-77 Retrieving logged user information using rusers and rusersd*

**Tip:** Use the **rusersd** daemon to determine the active telnet connection to the switch.

- `Disable RLS probing`: Allows you to disable Read Link Error Status of the AL_PAs.

For details about the parameters available in this panel, refer to the **configure** command description in the *Command Reference for OS 3.0 and OS 4.0,* GC26-7469.

## QuickLoop

QuickLoop is an additional feature enabled by license key, which may be installed on all IBM TotalStorage SAN Switch models except the 2109-F32 and 2109-M12 Core switch. The tab displayed in Figure 1-78 only appears after the license key has been installed.

QuickLoop is a unique method used to enable Private arbitrated loop devices to connect to a fabric, and complies with FC-AL standards. Because this allows private loops to be attached to fabrics, it can best be described as a Private Loop Fabric Attach (PLFA), as compared to a Private Loop Direct Attach (PLDA).

In the following sections, we discuss the different sections of the QuickLoop tab displayed.

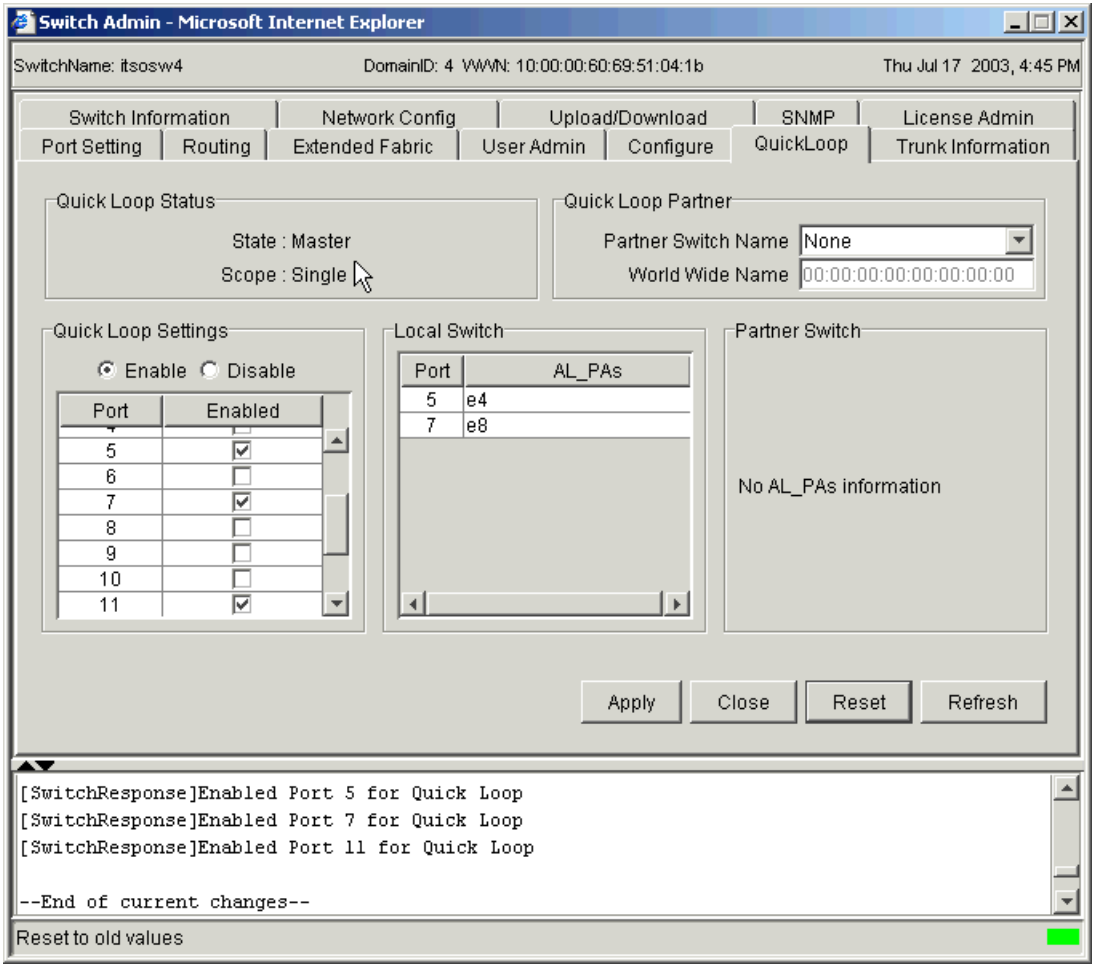


*Figure 1-78 QuickLoop tab*

### *QuickLoop status*

The QuickLoop status area displays the current state of the QuickLoop feature. In our example, we can see that our switch is the `Master`, if there are no QuickLoops in operation, that `state` should be `online`. Also displayed in this area is the `scope` of the QuickLoop partnership; ours is in `single` switch mode.

### *QuickLoop partner*

A QuickLoop is either a “single switch”, where all looplets are located on a single switch; or a “dual switch”, where looplets are located on either of the two partner switches.

In this area we can select a partner switch by scrolling down the `Partner Switch Name` box. The switch name and World Wide Name are then displayed. The same would be required to be performed on the partner switch also.

**Restriction:** A switch can only be in one QuickLoop partnership.

### *QuickLoop settings*

In this area we can set QuickLoop at the switch level or at the individual port level:

- `Enable`: Allows you to enable QuickLoop on all ports on the switch (except E_Ports).
- `Disable`: Allows you to disable QuickLoop on all ports by checking an individual box.
- `Enabled`: Allows individual ports to use the QuickLoop feature.

### *Local switch*

This section allows us to view all AL_PAs for Loop devices connected to the local switch ports.

### *Partner switch*

This section allows us to view AL_PAs for Loop devices connected to the partner switch ports if a partner switch is defined.

## Looplets

A QuickLoop consists of multiple private arbitrated looplets (a set of devices connected to a single port) that are connected by a fabric. All devices in a QuickLoop share a single AL_PA space and behave as if they are in one loop. This allows private devices to communicate with other devices over the fabric, provided they are in the same QuickLoop.

QuickLoop has the following characteristics:

- A QuickLoop can include up to two switches and support up to 126 devices.
- Each individual switch can only be included in one QuickLoop.
- A QuickLoop can include all or a subset of ports on an individual switch.
- Multiple QuickLoops can exist in a fabric of multiple switches.
- QuickLoop enabled switches can exist in the same fabric as non-QuickLoop enabled switches.
- A device attached to a QuickLoop can communicate with all other devices attached to the same QuickLoop.

- A private device in a QuickLoop can communicate with devices in the same QuickLoop only. Existing PLDA capable host drivers need no modification to perform I/O operations with storage devices.
- Public devices that are arbitrated loop capable are treated as private devices when connected to QuickLoop ports (their fabric login, or “FLOGI,” is rejected).
- QuickLoop supports the use of legacy devices, allowing them to be attached to a fabric and operate as if in a Private Loop Direct Attach (PLDA) environment.
- QuickLoop functionality can be enabled or disabled for either the entire switch or for individual ports. When QuickLoop is disabled on an individual port, that port returns to fabric mode.
- Each looplet in a QuickLoop has its own unshared bandwidth and can support transfer rates up to 200 MB/sec.
- Multiple devices can communicate simultaneously and at full bandwidth within multiple looplets located in the same QuickLoop.
- If a looplet error is detected, QuickLoop automatically takes the looplet out of service. If the error condition is cleared, the looplet is automatically reinstated.

## Private loop migration

QuickLoop provides a potential migration path from deploying a single private loop to deploying a fabric-based Storage Area Network (SAN). Initially, QuickLoop-enabled switches can be used to replace hubs when the SAN is first deployed and only has private devices attached. Then, as the SAN grows, fabric switches can be added without any detrimental effect to the QuickLoop enabled switches.

### *Address translation*

QuickLoop address translation is transparent and requires no actions on the part of the user. It is achieved through hardware translative mode (also known as phantom mode), in which a device not physically located in a looplet is made addressable by a unique AL_PA in that looplet. There are two hardware translative modes available to a QuickLoop enabled switch:

- **Standard translative mode:** This allows public hosts to communicate with private target devices across the fabric.
- **QuickLoop mode:** This allows private hosts to communicate with private target devices across the fabric.

The switch automatically determines and sets the appropriate mode.

## QuickLoop and zoning

QuickLoop can be used in conjunction with zoning. Using the products together provides the following additional features:

- AL_PAs from multiple QuickLoops can be used to add members to a zone. This is due to the Zoning ability to name QuickLoops and therefore identify the QuickLoop to which the AL_PA belongs.
- Additional control over access to QuickLoop devices is possible. Fabric devices in a zoned fabric can only access the QuickLoop (and fabric) devices that are in the same zone.

Zones can be created within QuickLoops. Zoning can be used to partition QuickLoops. This creates “QuickLoop zones” (as opposed to fabric zones), which support identification by either physical port number or AL_PA. For more information on QuickLoop zoning, refer to “QuickLoop tab” on page 59.

### *Managing QuickLoop*

We can enable QuickLoop for each port or for the whole switch by using the WEB TOOLS as detailed in “QuickLoop” on page 108, or by using the telnet commands, **`qlEnable`**, **`qlDisable`**, **`qlPortEnable`**, and **`qlPortDisable`**.

## Trunk Information

This panel is used to view information related to trunk groups as shown in Figure 1-79.

This panel is for viewing only. Disabling or enabling trunking is done through the Port Setting panel as explained in “Port Setting” on page 93.

We describe the Trunking feature in detail in below.

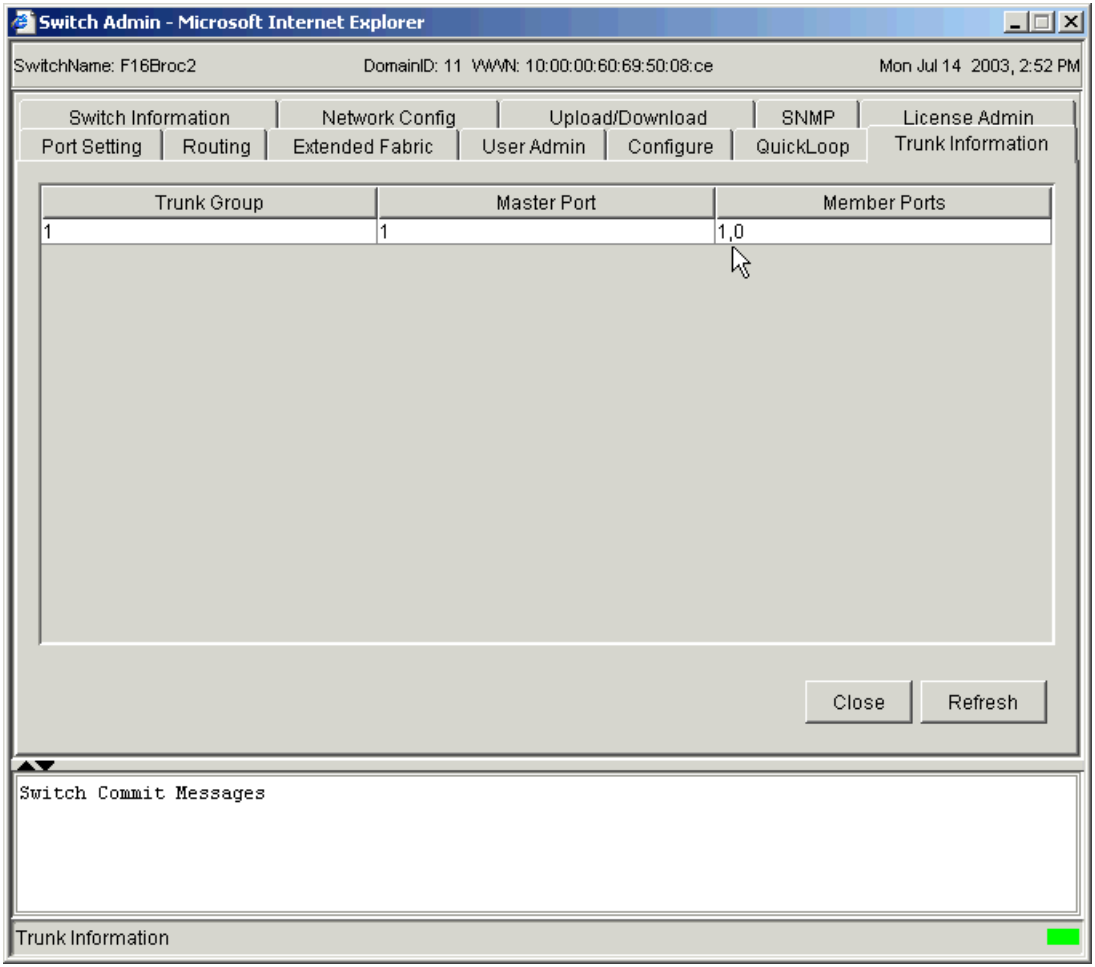


*Figure 1-79 Trunking Information panel*

ISL Trunking is shipped as standard with the 2109-F32 and M12 switches, with the 3434-F08 or 2109-F16 it requires a separate *Performance Bundle* License key to be purchased and installed.

The ISL Trunking feature allows up to four Interswitch Links (ISLs) to merge logically into a single link. An ISL is a connection between two switches through an Expansion Port (E_Port).

When using ISL Trunking to aggregate bandwidth of up to four ports, the speed of the ISLs between switches in a fabric is quadrupled. For example, at 2 Gb/s speeds, trunking delivers ISL throughput of 4, 6, and up to 8 Gb/s.

ISL Trunking supports high-bandwidth, large-scale SANs which include core switches. The primary task of ISL Trunking is to provide high bandwidth path between switches in a fabric, while balancing the traffic across the individual links and maintaining In-Order Delivery of data packets to their destination.

**Attention:** In-Order Delivery is the default setting in an IBM fabric, this setting may be changed by the user.

ISL Trunking may be managed using Telnet commands or the WEB TOOLS interface.

### *Advantages of ISL Trunking*

The ISL Trunking feature has many advantages; for example, it ensures optimal ISL bandwidth use across trunked links, while preserving in-order delivery (see previous Attention box). ISL Trunking uses frame-level load balancing, as opposed to Fibre Channel Shortest Path First (FSPF), to achieve faster fabric convergence, as well as higher availability in the fabric.

### *Routing without the ISL Trunking feature*

Prior to the implementation of the ISL Trunking feature, device-level load sharing was done through Fibre Channel networks that created ISLs and operated using the FSPF routing protocol. The FSPF routing protocol established and communicated the shortest paths for data to be carried from source to destination.

Although FSPF compliant switches ensure fixed routing paths, and guarantee that all frames are delivered in-order, congestion occurs if the aggregation of the stream exceeds the capacity of one of the ISLs in the path. For example, four untrunked ISLs have a maximum capacity of 2 Gb/s, which provides for a maximum throughput of 8 Gb/s. Due to traffic that is not trunked, the throughput of the four ISLs is determined as follows:

2 Gb/s + 1.5 Gb/s +.5 Gb/s + 1 Gb/s, which gives a 5 Gb/s total.

This is because two 2 Gb/s data streams are competing for the same path.

### *Routing with the ISL Trunking feature*

With ISL Trunking four ISLs provide 8 Gb/s of total throughput. With the implementation of ISL Trunking, bandwidth is shared across the trunked ISLs, permitting a total throughput of:

2 Gb/s + 1.5 Gb/s + 0.5 Gb/s + 1 Gb/s + 2 Gb/s, for a total 7 Gb/s in this case.

Because the trunk aggregates the four individual paths into one and preserves in-order deliver of frames, the total throughput is increased compared to a non-trunked group of ISLs.

## Trunking groups, ports, and masters

ISL Trunking dynamically performs load balancing, at the frame level, across a set of available links between two adjacent switches to establish a trunking group. Ports that belong to a trunking group are called trunking ports. One port is used to assign traffic for the group, and is referred to as the trunking master.

### *Trunking groups*

A trunking group is identified by the trunking master that represents the entire group. The rest of the group members are referred to as slave links that help the trunking master direct traffic across ISLs, allowing efficient and balanced in-order communication.

### *Trunking ports*

Trunking ports in a trunking group should meet the following criteria:

- Port must be configured as E_Ports.
- Ports must reside in the same contiguous four-port groups (quad). Each switch has the four-port quads identified on the port panel with alternating colors:
  - Group 1: port 0 to port 3
  - Group 2: port 4 to port 7
  - Group 3: port 8 to port 11
  - Group 4:port 12 to port 15
  - etc...
- Ports must run at the 2 Gb/s speed.
- The cable difference between all ports in a trunking group must be less than 500 meters.

### *Trunking masters*

The trunking master implicitly defines the trunking group. All ports with the same master are considered to be part of the same group. Each trunking group includes a single trunking master and several trunking slave links. The first ISL found in any trunking group is assigned to be the trunking master, also known as the principal ISL. After the trunking group is fully established, all data packets intended for transmission across the trunk are dynamically distributed at frame level across the ISLs in the trunking group, while preserving in-order delivery.

## Installing ISL Trunking

The ISL Trunking feature requires a 2109-M12, 2109-F32, 2109-F16 or a 3534-F08 switch. The M12 and F32 ship with the feature already installed. The F16 and F08 require that a *Performance Bundle* license be installed to enable trunking using either Telnet or the Web interface.

Both switches at either end of an ISL Trunk require an active license for trunking to work. A license may have been installed in the switch at the factory. If not, contact your switch supplier to obtain a license key.

## Administering ISL Trunking

The ISL Trunking feature is managed by performing some administration tasks. These tasks include:

- Enabling or disabling the trunking
- Enabling and disabling ports of a switch
- Setting the speed of a port
- Debugging a trunking link failure

The ISL Trunking feature is administered using Telnet commands.

### *ISL Trunking Telnet commands*

Table 1-16 describes the Telnet commands used to manage the ISL Trunking feature.

*Table 1-16 ISL Telnet commands*

| Command | Description | Example |
|---|---|---|
| `portCfgTrunkport` | Use this command to configure a port to be enabled or disabled for trunking. | To enable port 5 for ISL TRUNKING, enter:<br>`portCfgTrunkport 5, 1`<br>To disable port 5 for ISL TRUNKING, enter:<br>`portCfgTrunkport 5, 0` |
| `switchCfgTrunk` | Use this command to enable or disable trunking on all ports of a switch. | To enable trunking on all ports of a switch, enter:<br>`switchCfgTrunk 1`<br>To disable ISL Trunking on all ports of a switch, enter:<br>`switchCfgTrunk 0` |
| `trunkDebug` | Use this command to debug a trunk link failure. | To debug ports 1 and 2, enter:<br>`trunkDebug 1, 2` |
| `trunkshow` | Use this command to display ISL Trunking membership information. | To display ISL Trunking membership information about users, enter:<br>`trunkshow` |

## 1.7.7 The Telnet interface

The IBM TotalStorage SAN Switch also has a Telnet interface that can be accessed by clicking the picture of the monitor, as shown in Figure 1-80.


*Figure 1-80 Go to Telnet session*

In Figure 1-81 we show the Telnet window that is presented. From this window, the login and password are required.

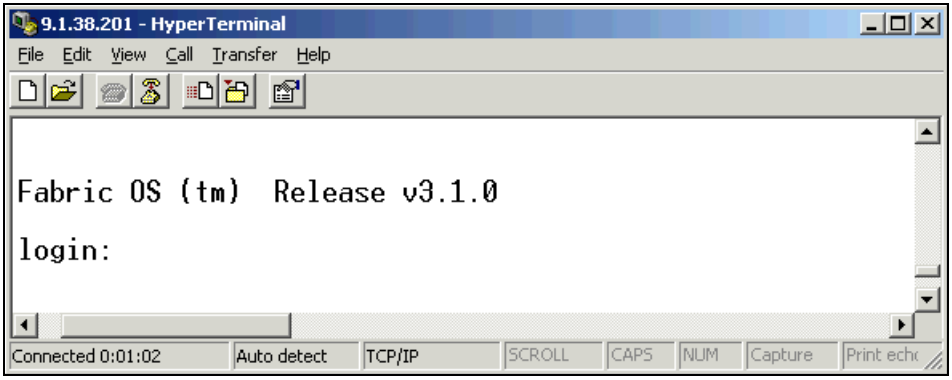


*Figure 1-81 Telnet session*

Only one Telnet session per F08 or F16 switch can be active at a time. Trying to open a second Telnet session will display the window shown in Figure 1-82.

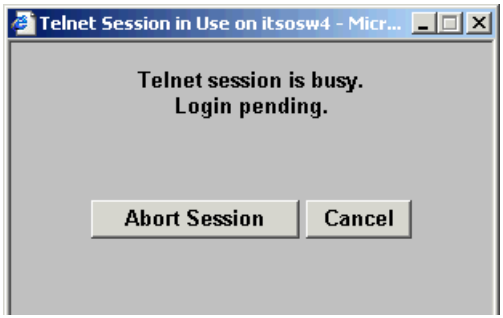


*Figure 1-82 Abort Telnet Session*

From this window we can abort the active session and start a new session, or by cancelling the action, keep the existing session open.

## 1.7.8 Telnet Commands: Overview

In Table 1-17 we provide an overview of the commands for managing and monitoring the 2109 Model F16 and 3534 Model F08 switches with Telnet. You can use these commands and settings to configure and operate the switch through the Telnet interface.

*Table 1-17 Telnet commands: Overview*

| Command | Description | Synopsis |
|---|---|---|
| **General purpose** | | |
| agtcfgSet | Modifies the SNMP agent configuration. | **agtcfgSet** |
| agtcfgShow | Prints SNMP agent configuration. | **agtcfgShow** |
| agtcfgDefault | Resets SNMP agent configuration to default values. | **agtcfgDefault** |
| aliasShow | Prints Alias information server. This command is not related to zoning. | **aliasShow** |
| backSpace | Sets or clears the alternate backspace character (0 for BACKSPACE, 1 for DEL). | **backSpace [0 \| 1]** |
| bsn | Displays the IBM serial number of a switch. | **bsn** |
| configure | Changes the system configuration settings. | **configure** |
| configShow | Displays the system configuration settings. | **configShow ["textfilter"]** |
| configDefault | Restores the system configuration to the default. | **configDefault** |
| configDownload | Downloads the switch configuration from a host file. | **configDownload ["host","user","file"[,"passwd"]]** |
| configUpload | Backs up the switch configuration to an ASCII file on a host workstation. | **configUpload ["host","user","file"[,"passwd"]]** |

| Command | Description | Synopsis |
|---|---|---|
| date | Displays or sets the system date and time. | **date ["newDate"]** |
| errDump | Displays the error log without page breaks. | **errDump** |
| errShow | Scrolls through the error log. | **errShow** |
| fabricShow | Displays the fabric membership information. | **fabricShow** |
| faShow | Displays Fabric Assist information. | **faShow** |
| faStatsShow | Displays Fabric Assist statistics. | **faStatsShow** |
| fanShow | Displays the fan status. | **fanShow** |
| fastboot | Restarts the switch, bypassing POST. | **fastboot** |
| firmwareDownload | Downloads a switch firmware file from a host. | **firmwareDownload ["host","user","file" [,"passwd"]]** |
| fpgaDownload | Downloads FPGA netlist into switch. | **fpgaDownload** |
| sfpShow | Displays serial ID SFP information. | **sfpShow [portnumber]** |
| h | Displays the shell history. | **h** |
| help | Displays help information for commands. | **help [command]** |
| i | Displays the task summary. | **i [taskId]** |
| ifModeSet | Sets the link operating mode for a network interface. | **ifModeSet ["interface"]** |
| ifModeShow | Displays the link operating mode for a network interface. | **ifModeShow ["interface"]** |
| ifShow | Displays the network interface information. | **ifShow ["ifName"]** |
| interopMode | Enables or disables switch interoperability with switches from other manufacturers. | **interopMode [0\|1]** |

| Command | Description | Synopsis |
|---|---|---|
| ipAddrSet | Sets the Ethernet and FC IP addresses. | **ipAddrSet** |
| ipAddrShow | Displays ISL switch group topology and status. | **ipAddrShow** |
| islShow | Displays ISL information. | **islShow** |
| islTopoCheck | Displays ISL sgroup connections for a switch. | **islTopoShow [sgroup]** |
| islTopoShow | Displays ISL switch group topology and status. | **islTopoShow** |
| login | Login as a new user. | **login** |
| logout | Logout from a telnet, rlogin, or serial port session. | **logout** |
| msConfigure | Configures the Management Server. | **msConfigure** |
| msPlatShow | Displays the Management Server Platform Database. | **msPlatShow** |
| msPlCapabilityShow | Displays the platform database management capability. | **msPlCapabilityShow** |
| msPlClearDB | Clears the management server platform database on all switches in the fabric. | **msPlClearDB** |
| msPlMgmtActivate | Displays the network interface information and activates platform database management on all switches in the fabric. | **msPlMgmtActivate** |
| msPlMgmtDeactivate | Deactivates platform database management on all switches in the fabric. | **msPlMgmtDeactivate** |
| msTdDisable | Disables the management server Topology Discovery Management service. | **msTdDisable ["ALL"]** |
| msTdEnable | Enables the management server Topology Discovery Management service. | **msTdEnable ["ALL"]** |

| Command | Description | Synopsis |
|---|---|---|
| nsAllShow | Displays global name server information. | **nsAllShow [type]** |
| nsShow | Displays the local name server information. | **nsShow** |
| passwd | Changes system login name and password. | **passwd ["user"]** |
| portCfgDefault | Sets port configurations to default. | **portCfgDefault port_number** |
| portCfgEport | Enables or disables a port from becoming an E_Port. | **portCfgEport [port_number, mode]** |
| portCfgGport | Designates a port as a locked G_Port. | **portCfgGport port_number, mode** |
| portCfgLport | Locks a port as an L_PORT. | **portCfgLport port_number, mode [,mode1]** |
| portCfgLongDistance | Configures a port to support long distance links. | **portCfgLongDistance port_number [,long_distance_level]** |
| portCfgShow | Displays port configuration settings. | **portCfgShow** |
| portCfgSpeed | Configures the port speed level. | **portCfgSpeed port[, speed_level]** |
| portCfgTrunkport | Configures a port to be enabled or disabled for trunking. | **portCfgTrunkport port, 1\|0** |
| portDisable | Disables a switch port. | **portDisable port** |
| portEnable | Enables a switch port. | **portEnable port** |
| portErrShow | Displays port error summary. | **portErrShow** |
| portLogClear | Clears the port log. | **portLogClear** |
| portLogDump | Displays ports log without page breaks. | **portLogDump [count[, saved]]** |
| portLogDumpPort | Displays the port log of a specified port without page breaks. | **portLogDumpPort port** |
| portLogShow | Displays the port log. | **portLogShow [count[, saved]]** |
| portLogShowPort | Displays the port log of a specified port. | **portLogShowPort port** |

| Command | Description | Synopsis |
|---|---|---|
| portPerfShow | Displays the port throughput performance in bytes, kilobytes, or megabytes. | **portPerfShow [interval]** |
| portShow | Displays the port status. | **portShow port** |
| portStatsShow | Displays port hardware statistics. | **portStatsShow port** |
| psShow | Displays the power supply status. | **psShow** |
| quietMode | Sets or clears the shell quiet mode. | **quietMode [0\|1]** |
| reboot | Restarts the switch. | **reboot** |
| snmpMibCapSet | Views and modifies options for configuring SNMP MIB and trap capability. | **snmpMibCapSet** |
| ssn | Displays and sets soft serial number (switch WWN) | **ssn** |
| switchBeacon | Sets switch beaconing mode on or off. | **switchBeacon 0\|1** |
| switchCfgSpeed | Configures all ports of the switch to a particular speed level. | **switchCfgSpeed speed_level** |
| switchCfgTrunk | Enables or disables trunking on all the ports of a switch. | **switchCfgTrunk 0\|1** |
| switchDisable | Disables the switch. | **switchDisable** |
| switchEnable | Enables the switch. | **switchEnable** |
| switchName | Displays or sets switch name. | **switchName ["newName"]** |
| switchShow | Displays switch and port status. | **switchShow** |
| switchStatusShow | Displays the overall status of the switch. | **switchStatusShow** |
| switchStatusPolicyShow | Displays the policy parameters that determine the overall switch status. | **switchStatusPolicyShow** |
| switchStatusPolicSet | Sets the policy parameters that determine the overall switch status. | **switchStatusPolicSet** |

| Command | Description | Synopsis |
|---|---|---|
| syslogdIpAdd | Adds the IP address of a syslog daemon. | **syslogdIpAdd IPaddress** |
| syslogdIpRemove | Removes the IP address of a syslog daemon. | **syslogdIpRemove IPaddress** |
| syslogdIpShow | Displays all syslog daemon IP addresses. | **syslogdIpShow** |
| tempShow | Displays temperature readings. | **tempShow** |
| timeOut | Used to set or clear idle telnet connection timeout value. | **timeOut [0 \| minutes]** |
| trunkDebug | Debugs a trunk link failure. | **trunkDebug port1, port2** |
| trunkShow | Displays trunking information. | **trunkShow** |
| uptime | Displays length of time the system has been operational. | **uptime** |
| version | Displays firmware version information. | **version** |
| **License** | | |
| licenseAdd | Adds a license key to this switch. | **licenseAdd license** |
| licenseRemove | Removes a license key from this switch. | **licenseRemove license** |
| licenseShow | Shows current license key. | **licenseShow** |
| **Performance** | | |
| perfCfgSave | Saves Performance configuration. | **perfCfgSave** |
| perfCfgRestore | Restores Performance configuration. | **perfCfgRestore** |
| perfCfgClear | Clears Performance settings from flash. | **perfCfgClear** |
| perfClrAlpaCrc | Clears AL_PA device's CRC count. | **perfClrAlpaCrc port[, ALPA]** |
| perfShowAlpaCrc | Gets AL_PA CRC count by port and AL_PA. | **perfShowAlpaCrc port[, ALPA]** |

| Command | Description | Synopsis |
|---|---|---|
| perfAddEEMonitor | Adds end-to-end monitor to a port. | **perfAddEEMonitor port, "SourceID", "DestID"** |
| perfDelEEMonitor | Deletes an end-to-end monitor on port. | **perfDelEEMonitor port[, monitor]** |
| perfShowEEMonitor | Shows user-defined end-to-end monitors. | **perfShowEEMonitor port[, interval]** |
| perfSetPortEEMask | Sets overall mask for E-to-E monitors. | **perfSetPortEEMask port,"TxSIDMsk","TxDIDMsk","RxSIDMsk", "RxDIDMsk"** |
| perfShowPortEEMask | Shows the current end-to-end mask. | **perfShowPortEEMask port** |
| perfAddUserMonitor | Adds filter-based monitor. | **perfAddUserMonitor port, "grouplist" [, "alias"]** |
| perfAddReadMonitor | Adds filter-based monitor - SCSI Read. | **perfAddReadMonitor port[, "alias"]** |
| perfAddWriteMonitor | Adds filter-based monitor - SCSI Write. | **perfAddWriteMonitor port [, "alias"]** |
| perfAddRWMonitor | Adds monitor - SCSI Read and Write. | **perfAddRWMonitor port[, "alias"]** |
| perfAddSCSIMonitor | Adds monitor for SCSI frame count. | **perfAddSCSIMonitor port[, "alias"]** |
| perfAddIPMonitor | Adds monitor for IP traffic frame count. | **perfAddIPMonitor port[, "alias"]** |
| perfDelFilterMonitor | Deletes filter-based monitor. | **perfDelFilterMonitor port[, monitor]** |
| perfShowFilterMonitor | Shows filter-based monitors. | **perfShowFilterMonitor port[, interval]** |
| **QuickLoop** | | |
| qlDisable | Disables quick loop mode. | **qlDisable** |
| qlEnable | Enables quick loop mode. | **qlEnable** |
| qlPortDisable | Sets port in non-quick loop mode. | **qlPortDisable port** |
| qlPortEnable | Sets port in quick loop mode. | **qlPortEnable port** |
| qlLIFA | Enables/disables fill LIFA with 0xFF. | **qlLIFA [0\|1]** |

| Command | Description | Synopsis |
|---|---|---|
| qlPartner | Prints/sets quick loop partner. | **qlPartner [0\|wwn]** |
| qlStatsShow | Prints quick loop statistics. | **qlStatsShow** |
| qlShow | Prints quick loop information | **qlShow** |
| qlPortShowAll | Prints quick loop port information. | **qlPortShowAll** |
| **Route** | | |
| bcastShow | Prints broadcast tree information. | **bcastShow** |
| dlsReset | Turns off Dynamic Load Sharing. | **dlsReset** |
| dlsSet | Turns on Dynamic Load Sharing. | **dlsSet** |
| dlsShow | Prints state of Dynamic Load Sharing. | **dlsShow** |
| fspfShow | Prints FSPF global prints. | **fspfShow** |
| interfaceShow | Prints FSPF interface information. | **interfaceShow [port]** |
| iodReset | Turns off In-Order Delivery. | **iodReset** |
| iodSet | Turns on In-Order Delivery. | **iodSet** |
| iodShow | Prints state of In-Order Delivery. | **iodShow** |
| linkCost | Sets or prints the FSPF cost of a link. | **linkCost** |
| LSDbShow | Prints Link State Database entry. | **LSDbShow [domain]** |
| mcastShow | Prints multicast tree information. | **mcastShow [group_ID]** |
| nbrStateShow | Prints neighbor's summary information. | **nbrStateShow [port]** |
| nbrStatsClear | Resets FSPF neighbor's counters. | **nbrStatsClear [port]** |
| topologyShow | Prints paths to domain(s). | **topologyShow [domainnumber]** |
| uRouteConfig | Configures static unicast route. | **uRouteConfig port, domain, outputport** |
| uRouteRemove | Removes static unicast route. | **uRouteRemove port, domainnumber** |

| Command | Description | Synopsis |
|---|---|---|
| uRouteShow | Prints Print port's unicast routing information. | **uRouteShow [port],[domainnumber]** |
| **Zoning** | | |
| aliAdd | Adds a member to a zone alias. | **aliAdd "aliName", "member; member"** |
| aliCreate | Creates a zone alias. | **aliCreate "aliName", "member; member"** |
| aliDelete | Deletes a zone alias. | **aliDelete "aliName"** |
| aliRemove | Removes a member from a zone alias. | **aliRemove "aliName", "member; member"** |
| aliShow | Prints zone alias information. | **aliShow ["pattern"]** |
| cfgAdd | Adds a member to a configuration. | **cfgAdd "cfgName", "member; member"** |
| cfgCreate | Creates a zone configuration. | **cfgCreate "cfgName", "member; member"** |
| cfgDelete | Deletes a zone configuration. | **cfgDelete "cfgName"** |
| cfgRemove | Removes a member from a configuration. | **cfgRemove "cfgName", "member; member"** |
| cfgShow | Prints zone configuration information. | **cfgShow ["pattern"]** |
| qloopAdd | Adds a member to a qloop. | **qloopAdd "qloopname", "member;member"** |
| qloopCreate | Creates a qloop. | **qloopCreate "qloopname", "member;member"** |
| qloopDelete | Deletes a qloop. | **qloopDelete "qloopName"** |
| qloopRemove | Removes a member from a qloop. | **qloopRemove "qloopName", "member;member"** |
| qloopShow | Prints qloop information. | **qloopShow [pattern]** |
| zoneAdd | Adds a member to a zone. | **zoneAdd "zoneName", "member;member"** |
| zoneCreate | Creates a zone. | **zoneCreate "zoneName", "member;member"** |
| zoneDelete | Deletes a zone. | **zoneDelete "zoneName"** |
| zoneRemove | Removes a member from a zone. | **zoneRemove "zoneName", "member;member"** |
| zoneShow | Prints zone information. | **zoneShow [pattern]** |

| Command | Description | Synopsis |
|---|---|---|
| fazoneAdd | Adds a member to a fabric assist zone. | **fazoneAdd "zoneName", "member;member"** |
| fazoneCreate | Creates a fabric assist zone. | **fazoneAdd "zoneName", "member;member"** |
| fazoneDelete | Deletes a fabric assist zone. | **fazoneDelete fazoneName** |
| fazoneRemove | Removes a member from a fabric assist zone. | **fazoneRemove "zoneName", "member;member"** |
| fazoneShow | Prints Fabric Assist Zone information. | **fazoneShow [pattern [, transflag]]** |
| cfgClear | Clears all zone configurations. | **cfgClear** |
| cfgDisable | Disables a zone configuration. | **cfgDisable** |
| cfgEnable | Enables a zone configuration. | **cfgEnable "cfgName"** |
| cfgSave | Saves zone configurations in flash. | **cfgSave** |
| cfgTransAbort | Aborts zone configuration transaction. | **cfgTransAbort** |

## 1.7.9 Performance Monitor

The Performance Monitor performs the following functions:

- It graphically displays throughput (megabytes per second) for each port and for the entire switch. Port throughput is the number of bytes that are received at a port plus the number of bytes that are transmitted. Switch throughput is the sum of the throughput for all the ports. The Performance Monitor also allows the graphing of traffic based on the Source ID and the Destination ID hardware-filtering mechanism.
- It provides the ability to change the configuration of a switch or port visually by using the graphics.

To access the Performance Monitor, we click the Perf button from the switch view in WEB TOOLS as shown in Figure 1-83.

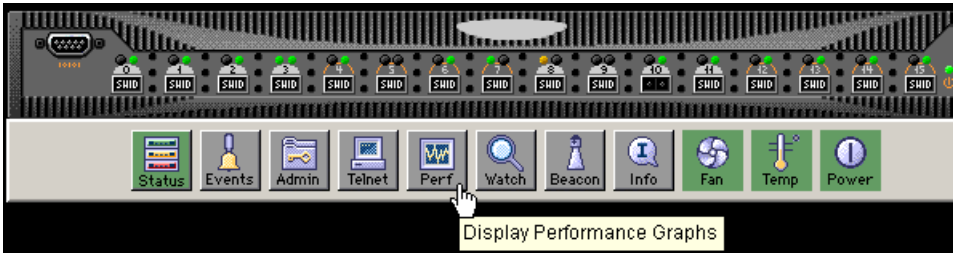


*Figure 1-83 Switch management window*

The Performance Monitor contains a collection of graphs on the display panel, or canvas. The graphs are sized based on the number of graphs loaded on the canvas. Double-clicking a graph expands the graph to the size of the display.

## Features

These are some of the features available in the Performance Monitor:

- An existing report can be selected from a list of reports that are predefined. In some cases, you can supply the object to be monitored and graphed (such as port number, SID/DID pair, AL_PA, or switch domain number).
- Graphs are displayed on a canvas, which can hold a maximum of eight graphs simultaneously. An individual graph can be maximized to occupy the entire canvas. The size of the graphs on the canvas is determined by the number of graphs being displayed. The window does not need to be scrolled to view all the selected graphs.
- The collection of graphs in the canvas can be stored for later retrieval on the switch. Up to 20 individual canvases can be saved. Each canvas is saved with its name, a brief description, and the graphs that comprise the canvas.
- Any graph can be magnified and detached from the main canvas or removed from the main canvas using a pop-up menu. You can display the pop-up menu by pointing the mouse at any graph on the main canvas and clicking the right mouse button. To reattach the detached (Zoomed Out) graph back to the main canvas, you can point the mouse to the detached graph, click the right button and select **Zoom In**.
- Each graph can be printed.

After clicking the Perf button from the Switch View, we see the default performance graph as shown in Figure 1-84.

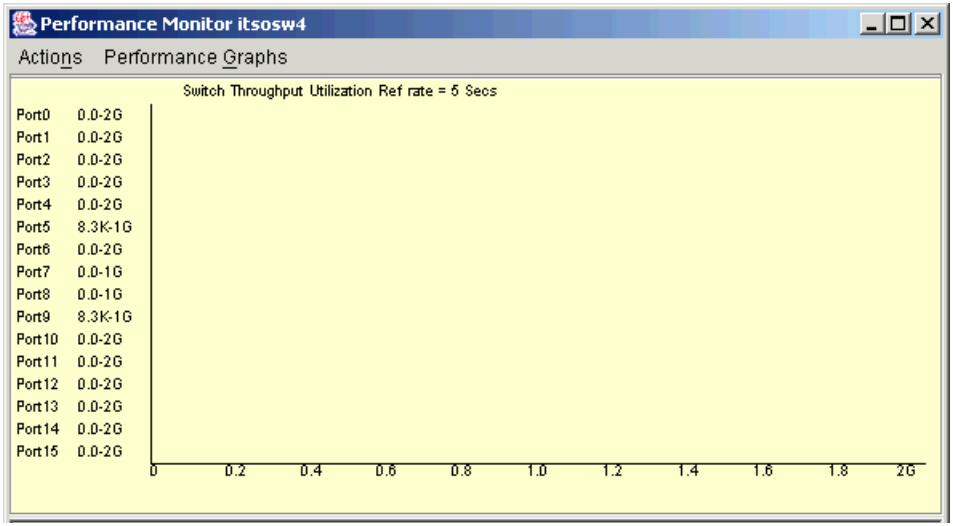


*Figure 1-84 Performance Monitoring — Default graph*

All graphs are real-time. Depending on the graph chosen, it is updated either every 5 or 15 seconds.

## Performance Monitor menus

The Performance Monitor is made up of two main menus:

- Actions menu
- Performance graphs menu

### *Actions menu*

The Actions menu of the Performance Monitor feature, shown in Figure 1-85, is made up of the following sub-menus:

- Display canvas configurations
- Save current canvas configuration
- Resource usage display
- Print all graphs

A canvas is a collection of predefined graphs.

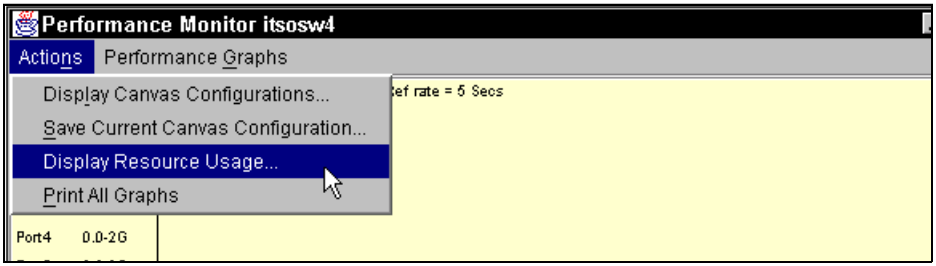


*Figure 1-85 Actions menu displaying choices*

### *Display canvas configurations*

Use this item to display and edit the various canvas configurations previously saved, as shown in Figure 1-86.

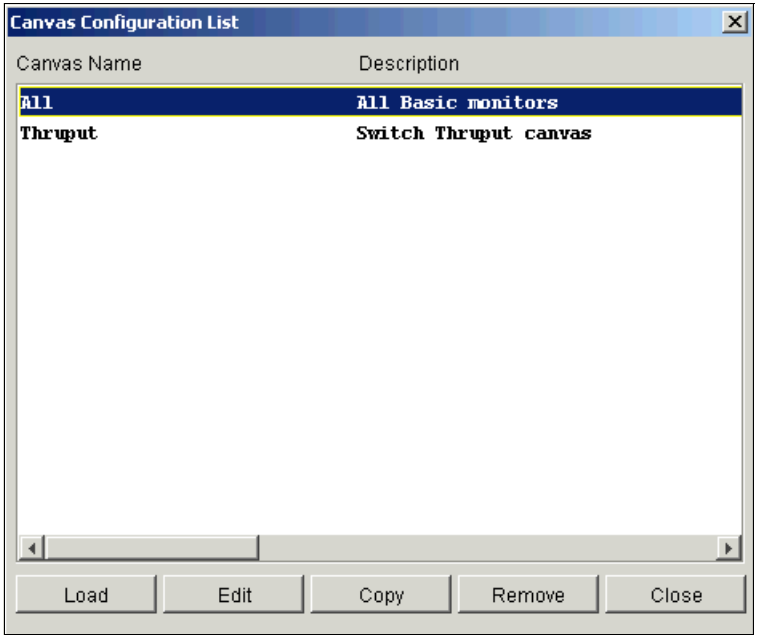


*Figure 1-86 Canvas Configuration List window*

Table 1-18 describes the fields on the Canvas Configuration List window.

*Table 1-18 Canvas Configuration List window — fields*

| | |
|---|---|
| **Available in Canvas Configuration List** | |
| Load | Select to load a canvas of 1 to 8 graphs onto the Performance Monitor facility by choosing the highlighted canvas name. |
| Edit | Select to make changes to a canvas or change configurations. A list of graphs which comprise the highlighted canvas will appear. |
| Copy | Select to copy the highlighted canvas configuration from the list to the switch flash. You will be prompted to type in the name and description of the canvas to which you want to copy your chosen graph. |
| Remove | Select to remove a highlighted canvas from the list and the switch flash. You will be prompted with a warning that you are going to delete the selected canvas. |
| Close | Select to close the canvas configuration list. |
| **Available in Edit Canvas Window** | |
| Save | Select to save an edited canvas. |
| Edit | Select to make changes to a graph on a canvas. A data entry frame will appear. |
| Add | Select to add a graph to a canvas. A pop-up menu of available graphs will display. Use this option to select the type of graph to add. For more information, refer to the *Basic Monitoring* and *Advanced Monitoring* sections of this chapter. |
| Remove | Select to delete a graph. The graph currently highlighted will be removed. |
| Cancel | Select to exit the window without making any changes. |
| **Available in Copy Canvas List** | |
| Name | Type in the name of the canvas to which you want to copy the graph. |
| Description | Type in a description of the graph to be copied. |
| Copy Canvas | Select to copy the selected graph to another canvas. |
| Cancel | Select to exit the window without making a copy. |

### *Save Current Canvas Configuration*

The Save Current Canvas Configuration menu saves the currently configured canvas to the switch. It uses a canvas name and a brief description to save the canvas, as shown in Figure 1-87.

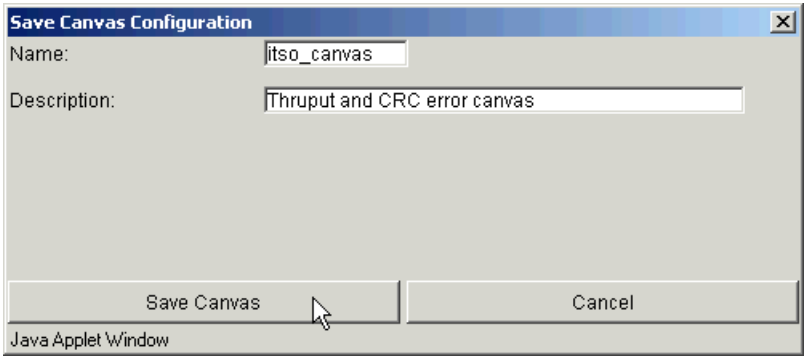


*Figure 1-87 Save Canvas Configuration window*

If the canvas already exists, the Confirm Override Canvas confirmation window pops up. Use the override option when you need to update an existing canvas.

### *Display Resource Usage*

The Resource Usage Display window allows you to view the resources that are allocated for end-to-end use, as well as providing filter-based monitoring for each port, as shown in Figure 1-88.

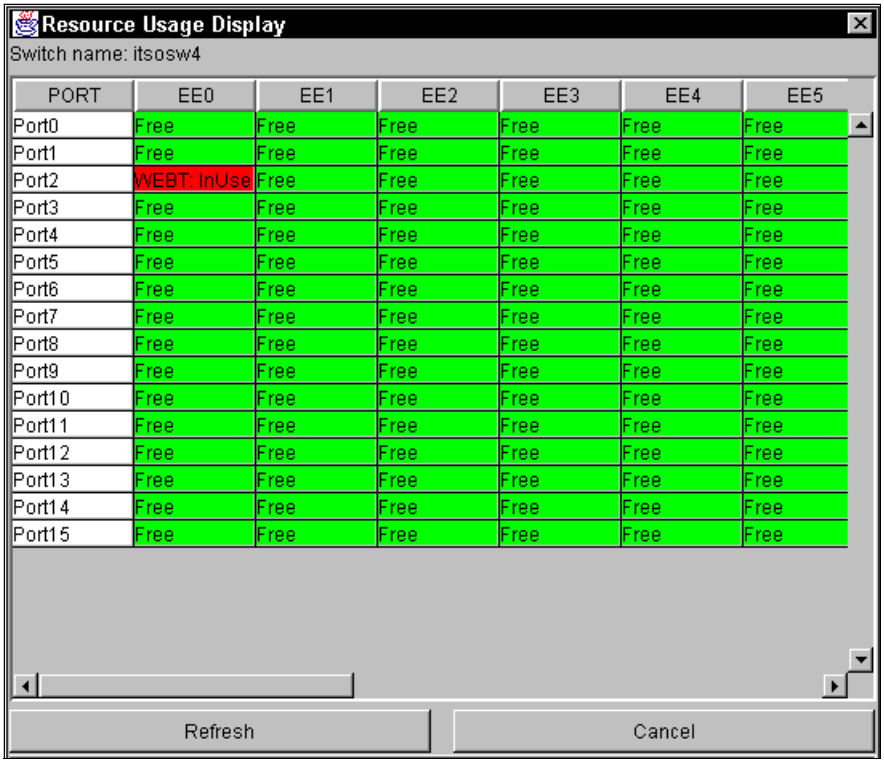


*Figure 1-88 Resource Usage Display window*

These are the fields available in the Resource Usage Display window:

- **Refresh**: Select to refresh the window immediately.
- **Cancel**: Select to close the window.

In Figure 1-88, we can see that the End-to-End resource EE0, corresponding to a certain filter is used on port 2 by the WEB TOOLS.

### *Print all graphs*

Use this item to print all the graphs on the selected canvas.

## Performance Graphs menu

We show the Performance Graphs menu in Figure 1-89.

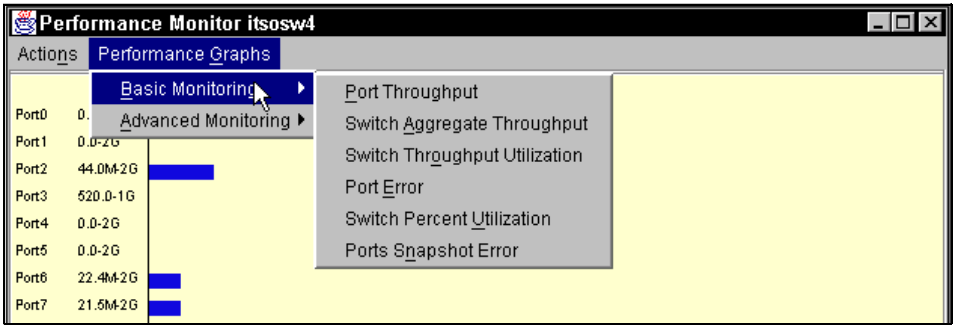


*Figure 1-89 Performance graphs menu, showing choices*

The Performance Graphs menu gives access to two sets of performance graphs:

- Basic Monitoring
- Advanced Monitoring (requires an additional license key)

### *Basic Monitoring*

We have selected all the options available in basic monitoring and have created a canvas that includes them. This is shown in Figure 1-90.

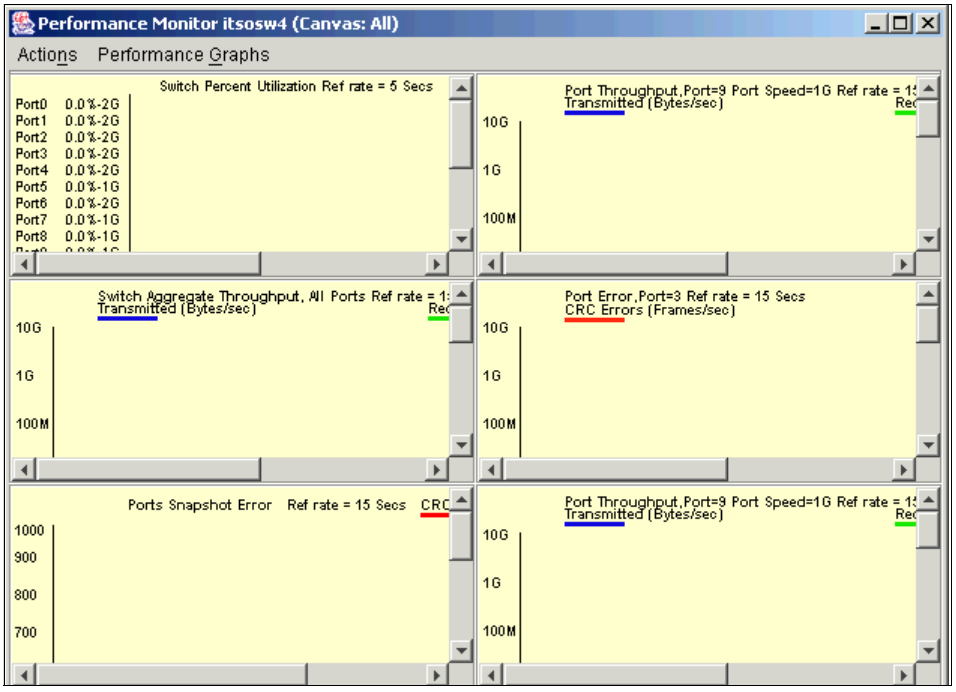


*Figure 1-90 Basic monitoring full functions*

The graphs available on this canvas are described in Table 1-19.

*Table 1-19 Graphs available in Basic Monitor*

| Graph Name | Type | Description |
|---|---|---|
| Port Throughput Graph | Line | Displays the performance of a port based on four-byte frames received and transmitted. |
| Switch Aggregate Throughput Graph | Line | Displays the aggregate performance of all ports of a switch. S |
| Switch Throughput Utilization Graph | Horizontal Bar | Displays the port throughput at the time the sample is taken. |
| Port Error Graph | Line | Displays a line of CRC errors for a given port. |
| Switch Percent Utilization Graph | Horizontal Bar | Displays the percentage of usage of a chosen switch at the time the sample is taken. |
| Ports SnapShot Error Graph | Vertical Bar | Displays the CRC error count between sampling periods for all the ports on a switch. |

For each graph, additional options are available by right-clicking the graph as shown in Figure 1-91.

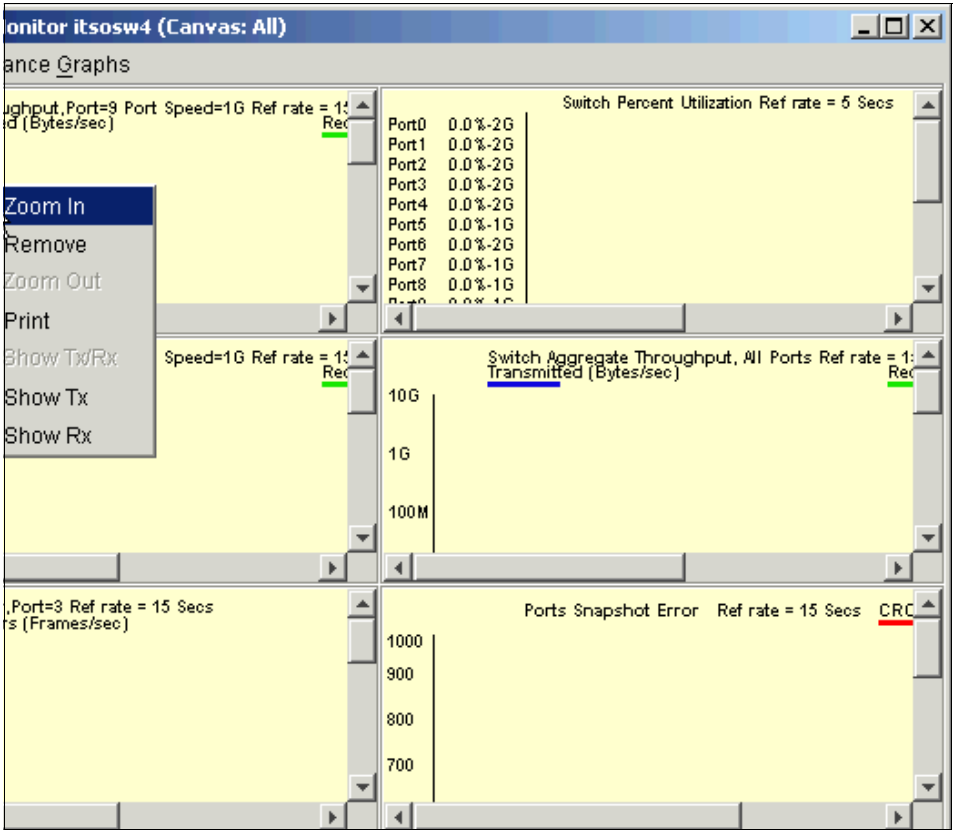


*Figure 1-91 Graphs additional options*

These are the options:

- `Zoom In`: Detach the graph from the canvas in a larger window. We then have the option to `Zoom Out` to place the graph back on the canvas.
- `Remove`: Remove the graph from the canvas
- `Print`: Print the graph
- `Show Tx/Rx`: Display both transmitted and received bytes
- `Show Tx`: Display only transmitted bytes
- `Show Rx`: Display only received bytes

### *Example: Port throughput graph*

To view the throughput on port 3 of the current switch, select **Performance Graphs —> Basic Monitoring —> Port Throughput**. The Port Throughput Setup is then displayed, shown in Figure 1-92.

**Note:** To expand the Domain folder, we need to double-click it to open the port tree.

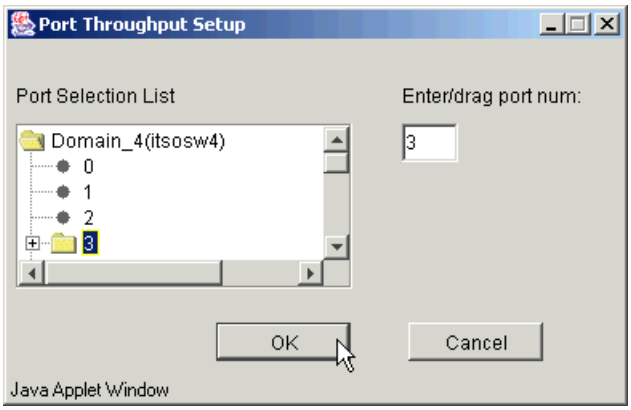


*Figure 1-92 Port throughput graph setup*

We enter the number of the port we want to monitor. A new graph is then added to the canvas. If we zoom in, we get the window shown in Figure 1-93.

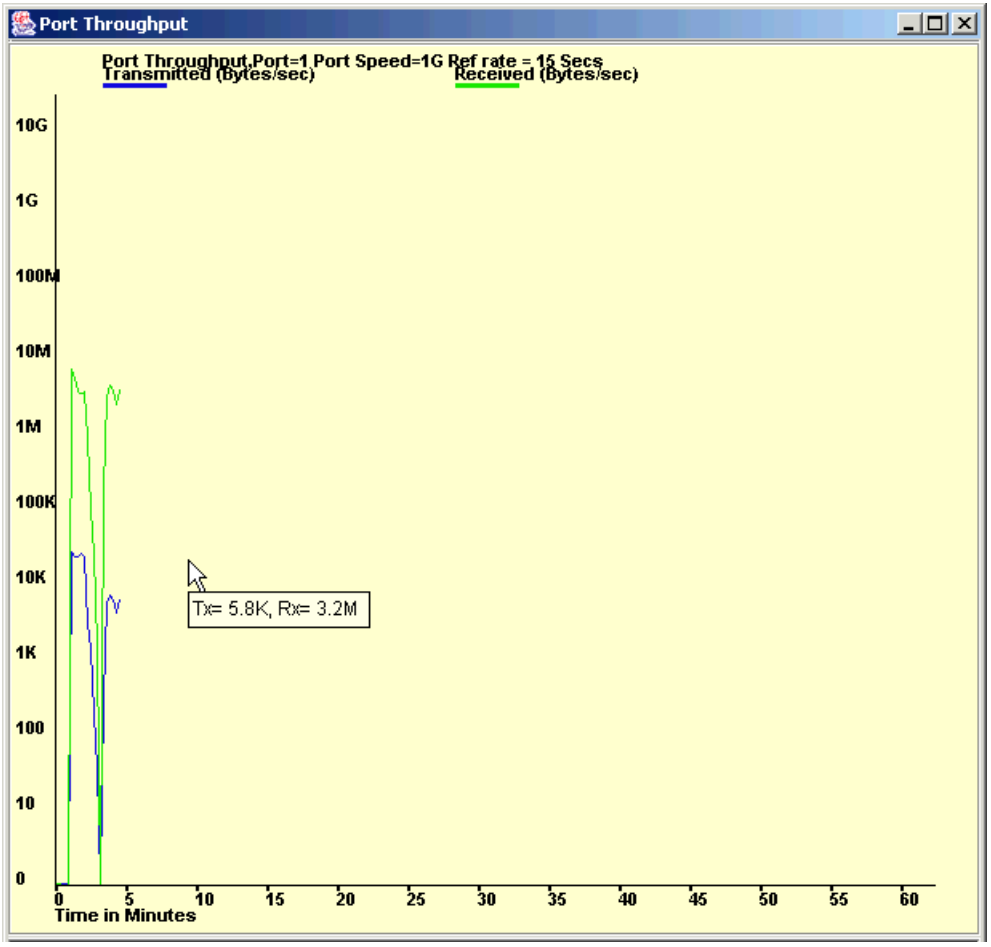


*Figure 1-93 Port Throughput graph*

**Tip:** We can get more detailed information by dragging the mouse pointer over a graph.

## 1.7.10 Advanced Performance Monitoring

Performance Monitoring is an optionally licensed product that runs on the F16 and F08. It provides SAN performance management through an end-to-end monitoring system that enables you to:

- Increase end-to-end visibility into the fabric

- Enable more accurate reporting for service level agreements and charged access applications
- Improve performance tuning and resource optimization
- Shorten troubleshooting time
- Promote better capacity planning
- Simplify administration and setup
- Increase productivity with pre formatted and customizable windows and reports

The Performance Monitoring product:

- Monitors transaction performance from its source to its destination
- Provides device performance measurements by port, AL_PA, and LUN
- Reports CRC error measurement statistics
- Measures trunking performance
- Compares IP versus SCSI traffic on each port
- Includes a wide range of predefined reports
- Allows you to create customized user-defined reports

You can administer Performance Monitoring through either Telnet commands or WEB TOOLS. If you use WEB TOOLS, a WEB TOOLS license must also be installed on the switch.

## 1.7.11 Performance Monitoring with Telnet commands

Three different types of Performance Monitoring can be done using Telnet commands:

- AL_PA monitoring
- End-to-end monitoring
- Filter-based monitoring

### AL_PA monitoring

AL_PA monitoring provides information about the number of CRC errors occurring in Fibre Channel frames in a loop configuration. AL_PA monitoring collects CRC error counts for each AL_PA that is attached to a specific port.

### End-to-end monitoring

End-to-end monitoring provides information about transaction performance between the transactions source (SID) and destination (DID) on a fabric or a loop. Up to 16 SID-DID pairs per port can be specified. For each of the SID-DID pairs, the following information is available:

- CRC error count on the frames for the SID-DID pair
- Fibre Channel words transmitted from the port for the SID-DID pair
- Fibre Channel words received by the port for the SID-DID pair

### Filter-based monitoring

Filter-based monitoring provides information about a filter's hit count. Any parameter in the first 64 bytes of the Fibre Channel frame can be measured. The counter increases each time a frame is filtered through the corresponding port. Examples of port filter statistics that can be measured are:

- SCSI read, write, or read/write commands
- CRC error statistics (port and AL_PA)
- IP versus SCSI traffic comparison

## 1.7.12 Performance Monitoring with WEB TOOLS

You can monitor performance using the WEB TOOLS if a WEB TOOLS license is also installed. The enhanced Performance Monitoring features in WEB TOOLS provide:

- Predefined performance graphs for AL_PA, end-to-end, and filter-based
- User-defined graphs
- Performance canvas for application-level or fabric-level views
- Configuration editor (save, copy, edit, and remove multiple configurations)
- Persistent graphs across restarts (saves parameter data across restarts)
- Print capabilities

### Predefined performance graphs

Predefined graphs are provided to simplify performance monitoring. A wide range of end-to-end fabric, LUN, device, and port metrics are included.

Figure 1-94 shows the predefined performance graphs available.

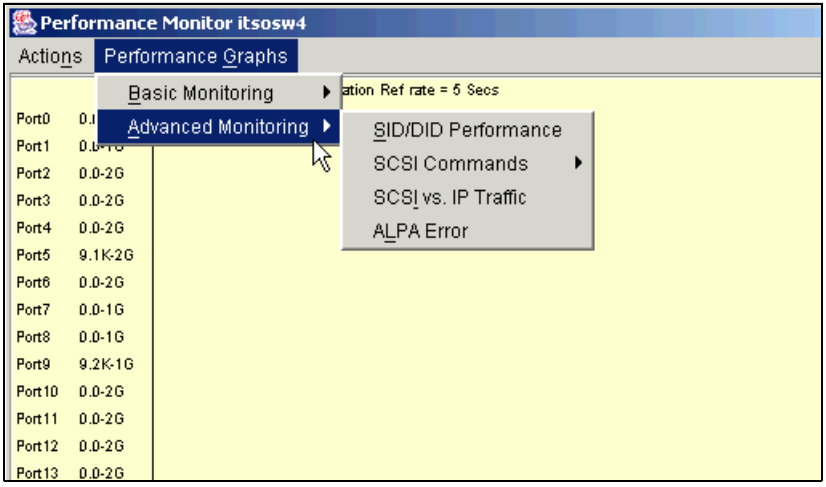


*Figure 1-94 Advanced monitoring options*

## Installing Performance Monitoring

To enable Performance Monitoring, you must install a license on each switch that will use this feature. Contact your switch supplier to obtain a license key.

**Note:** A license might have already been installed on the switch at the factory.

You can install a Performance Monitoring license through Telnet commands or using WEB TOOLS. as discussed in “License admin” on page 91.

## Using Advanced Performance Monitoring with WEB TOOLS

In the sample windows that follow, we show the functions that are available using the TotalStorage Specialist.

**Attention:** As the monitoring of any switch is subjective by nature, we just show the windows to give the reader some familiarity with features that can be monitored.

Figure 1-95 shows all the options that are available.

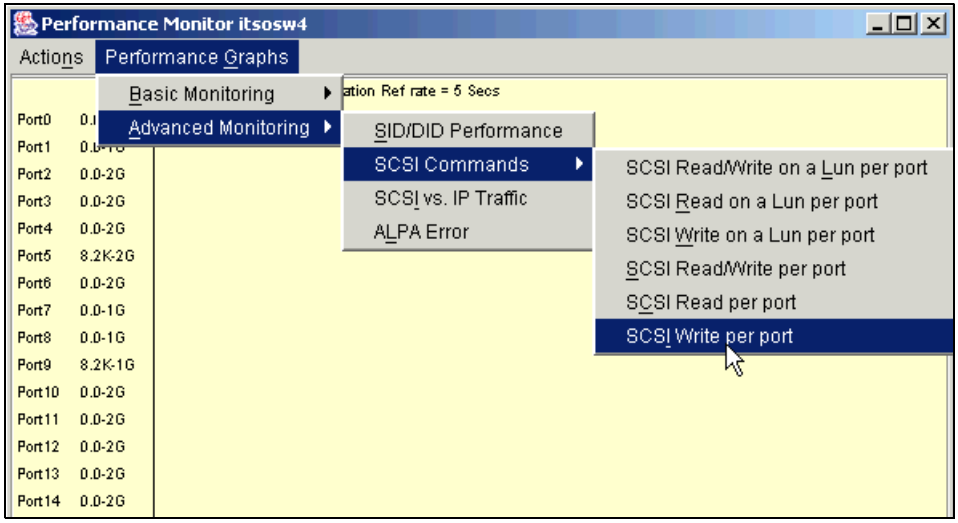


*Figure 1-95 Advanced monitoring range of options*

Table 1-20 describes the types of graphs available in the Advanced Monitoring menu.

*Table 1-20 Graphs available in Advanced Monitoring feature*

| Graph Name | Type | Description |
|---|---|---|
| SID/DID Performance Graph | Line | This graph charts the traffic between a SID (or WWN) and a DID (or WWN) pair on the switch being managed. |
| SCSI Commands Graph | Line | The total number of Read/Write commands on a given port to a specific LUN.<br>Provides the following choices:<br>SCSI Read/Write on a LUN per port.<br>SCSI Read on a LUN per port.<br>SCSI Write on a LUN per port.<br>SCSI Read/Write per port.<br>SCSI Read per port.<br>SCSI Write per port. |
| SCSI vs IP Graph | Vertical Bar | Shows percentage of SCSI versus IP frame traffic on each individual port. |
| AL_PA Error Graph | Line | Displays CRC errors for a given port and a given AL_PA. |

### *SID/DID Performance Graph*

Go to **Performance Graphs —> Advanced Monitoring —> SID/DID Performance**. To set up our parameters for SID/DID performance monitoring, we then use the window shown in Figure 1-96.

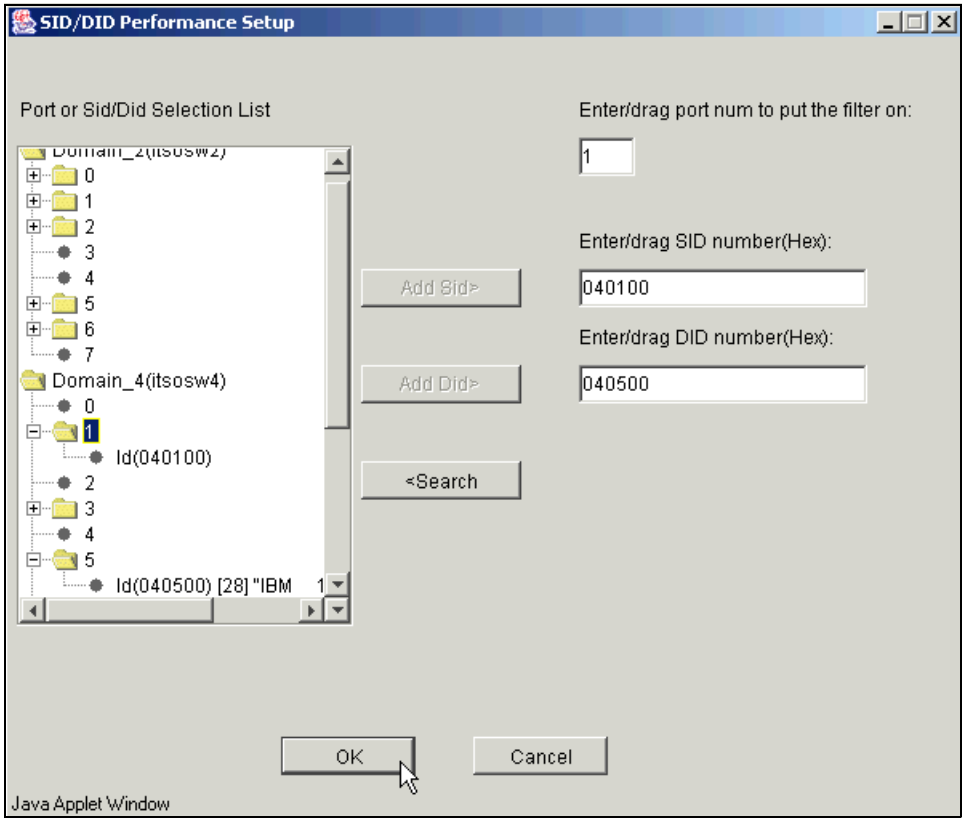


*Figure 1-96 SID/DID performance setup*

To choose the port and SID/DID that you want to graph:

1. Double-click a folder in the `Port Selection List` window. A drop-down list of ports will appear.
2. Select the port that you want to monitor or change by using one of the following methods:
   a. Type the port number in the `Enter /Drag Port Numbers` window.
   b. Drag the port “folder” from the Port Selection window to the `Enter/Drag Port Number` window.
3. Select the port “folder”, or the small icon that appears next to it. A drop-down list of SID/DID files will appear.

4. Select the SID/DID numbers that you want to graph by using one of the following methods:
   a. Type the SID number in the `Enter /drag SID Numbers` window. Repeat for the DID number.
   a. Drag the SID “file” from the Port Selection window to the `Enter/drag SID Number` window. Repeat for the DID number.
5. Select **OK**.

An example of an SID/DID graph, displaying the traffic between a SID and a DID pair, is shown in Figure 1-97.

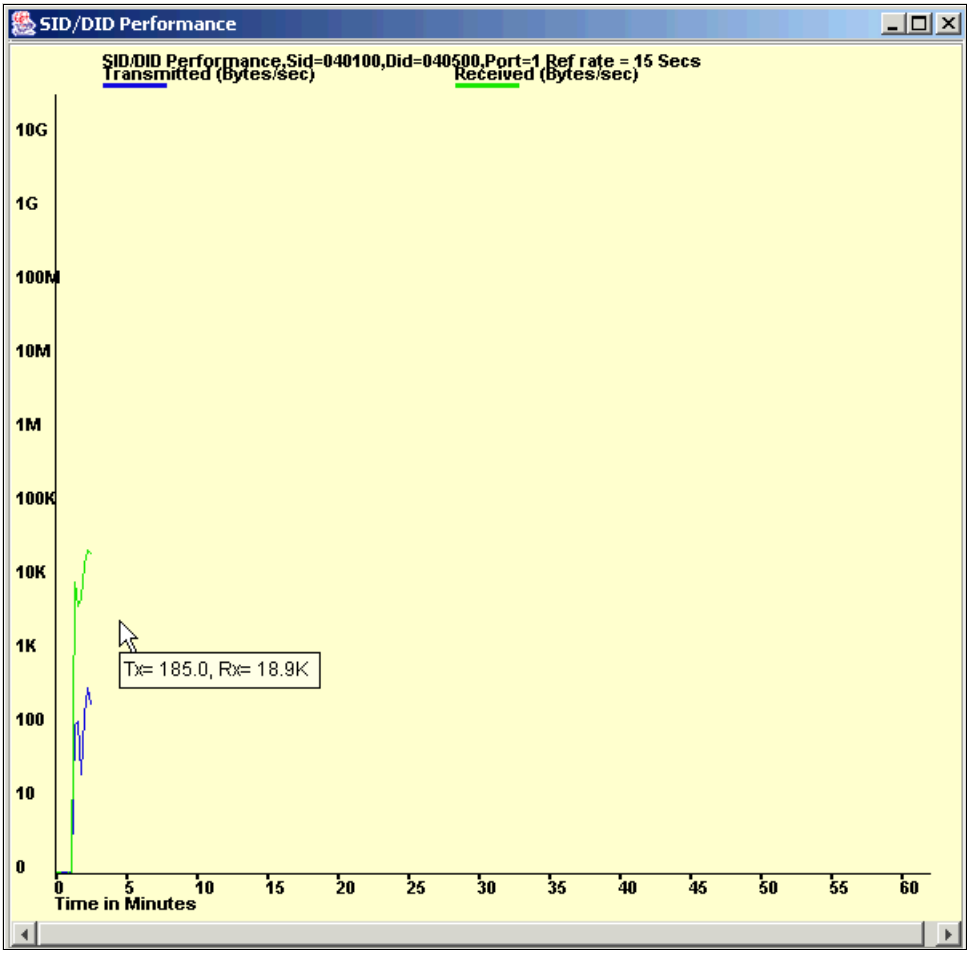


*Figure 1-97 SID/DID graph example*

Note that SID/DID monitoring monitors traffic on the port logically closest to the SID on the current switch.

Figure 1-98 shows several switches and the proper ports on which to add performance monitors for a specified SID/DID pair.

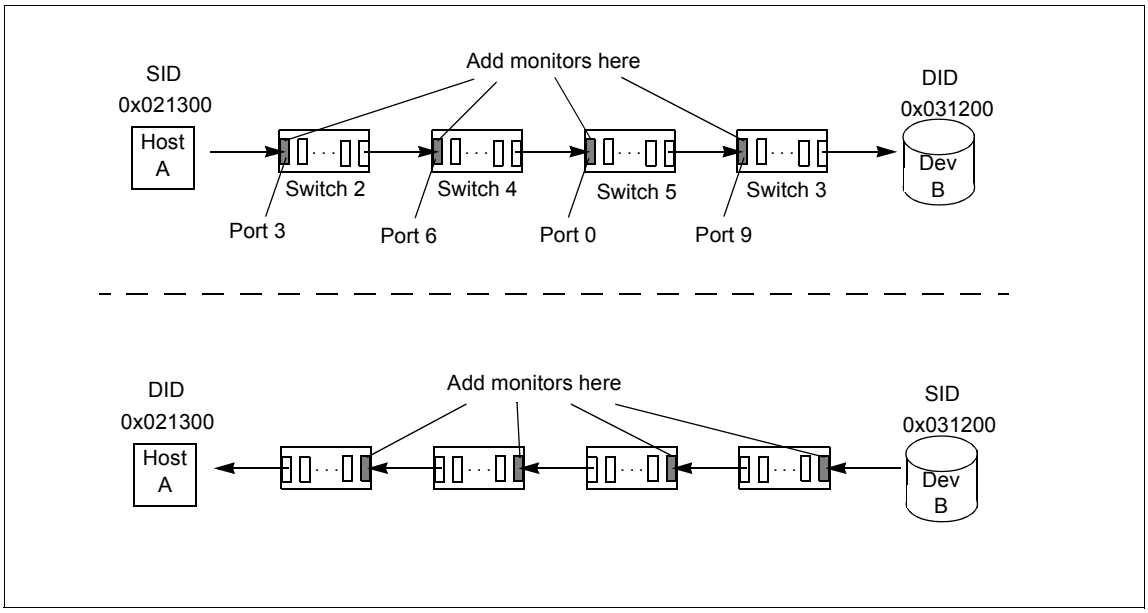


*Figure 1-98 Proper placement of SID/DID performance monitors*

In Figure 1-98, monitoring Port 6 on Switch 4 specifying Host A as the SID and Dev B as the DID is correct.

But monitoring Port 6 on Switch 4 specifying Dev B as the SID and Host A as the DID will not display a valid graph, as traffic will be shown as null.

### *SCSI command graph*

When you select the SCSI graph in **Performance Graphs —> Advanced Monitoring —> SCSI Commands**, the following options will be displayed in a pull-down menu:

- SCSI Read/Write on a LUN per port
- SCSI Read on a LUN per port
- SCSI Write on a LUN per port
- SCSI Read/Write per port
- SCSI Read per port
- SCSI Write per port

Each graph will prompt you with a data entry window to select the port and LUN to be monitored, as shown in Figure 1-99. In this example, we want to monitor SCSI Read and Writes command on LUN 1 going through port 2 of the current switch.

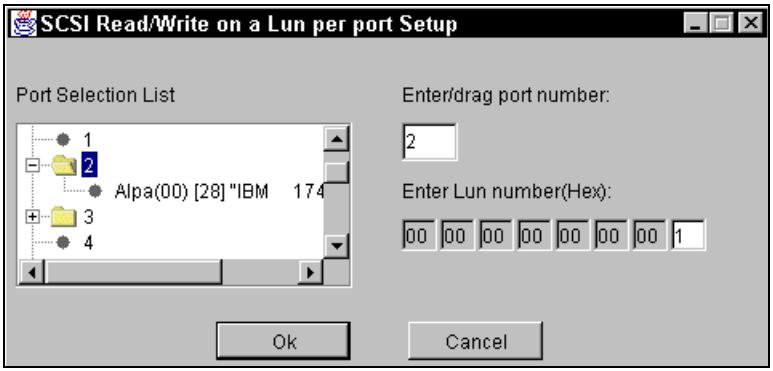


*Figure 1-99 SCSI read/write LUN per port setup*

To select the port and LUN to monitor:

1. Double-click the folder in the `Port Selection List` window. A drop-down list of ports will appear.
2. Select the port that you want to monitor or change by using one of the following methods:
   a. Type the port number in the `Enter/Drag Port Numbers` window.
   b. Drag the port "file" from the `Port Selection` window to the `Enter/Drag Port Number` window.
3. Enter a LUN number in the `Enter LUN Number (Hex)` window.

   You can enter only four LUN numbers at a time.
4. Select **OK**.

A graph displaying the total number of Read and/or Write commands on a given port to a specific LUN will be displayed.

An example of a SCSI graph, using the Write on a LUN per port option, is shown in Figure 1-100.

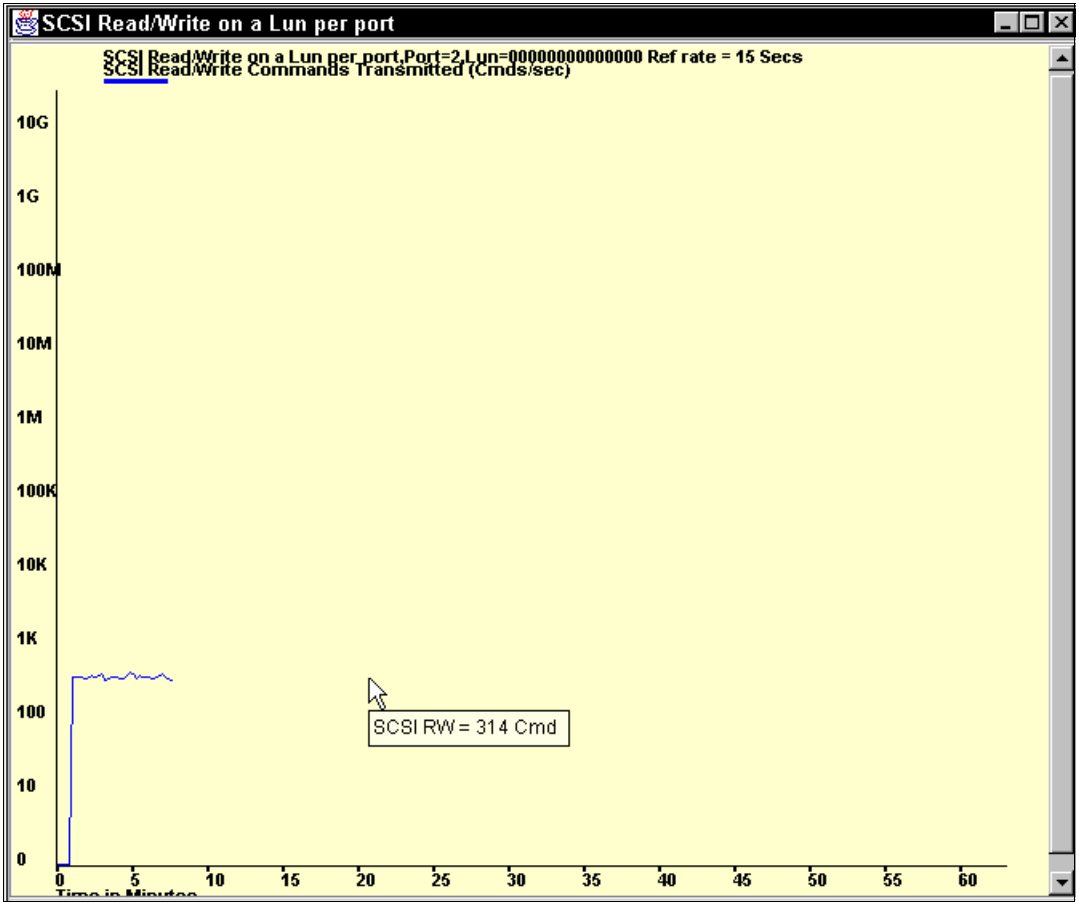


*Figure 1-100 SCSI Read/Write on a LUN per port graph*

### *SCSI versus IP Traffic Graph*

The SCSI versus IP Traffic graph is accessible via **Performance Graphs —> Advanced Monitoring —>SCSI versus IP Traffic.**

An example of this graph, displaying the percentage of SCSI versus IP frame traffic, is shown in Figure 1-101.

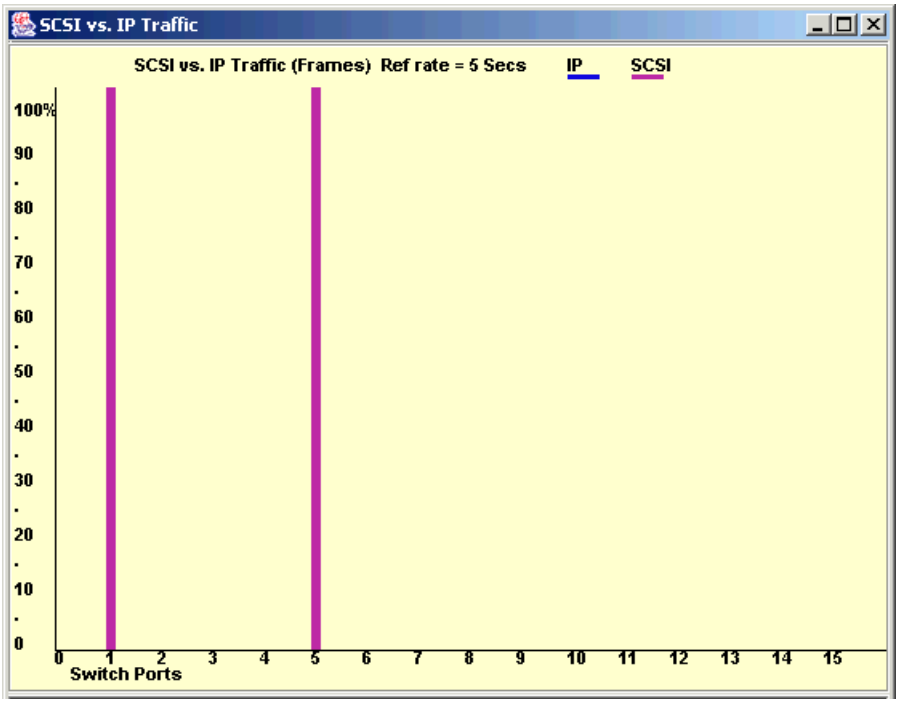


*Figure 1-101 SCSI versus IP traffic graph*

This graph gives us the percentage of IP and SCSI traffic on the current switch on a port basis.

### *AL_PA Error graph*

When you select an AL_PA Error graph via **Performance Graphs —> Advanced Monitoring —>AL_PA Error**, you will be prompted to choose the port that you want to monitor for various errors.

Figure 1-102 is an example of the data entry window that you will see when you choose to create an AL_PA Error Graph.

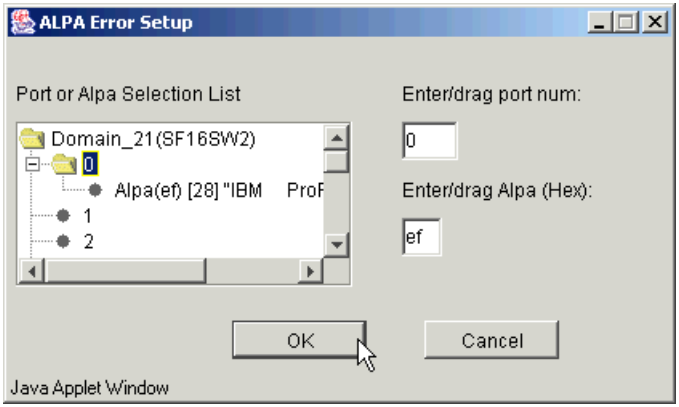


*Figure 1-102 AL_PA error graph setup window*

To choose the port and AL_PA that we wish to graph:

1. Double-click the **Domain** folder in the `Port or AL_PA Selection List` window. A drop-down list of ports will appear.
2. We select the port that we wish to monitor or change by using one of the following methods:
   a. Type the port number in the `Enter/Drag Port Numbers` window.
   b. Drag the port “folder” from the `Port Selection` window to the Enter/Drag Port Number window.
3. We select the small plus that appears next the port “folder”. A drop-down list of AL_PAs on that port will appear.
4. We select the AL_PA number that we wish to graph by using one of the following methods:
   a. Type the AL_PA number in the `Enter/drag SID Numbers` window.
   b. Drag the AL_PA “file” from the `Port Selection` window to the `Enter/drag ALPA Number` window.
5. Select **OK**.

An AL_PA Error graph will be displayed, as shown in Figure 1-103.

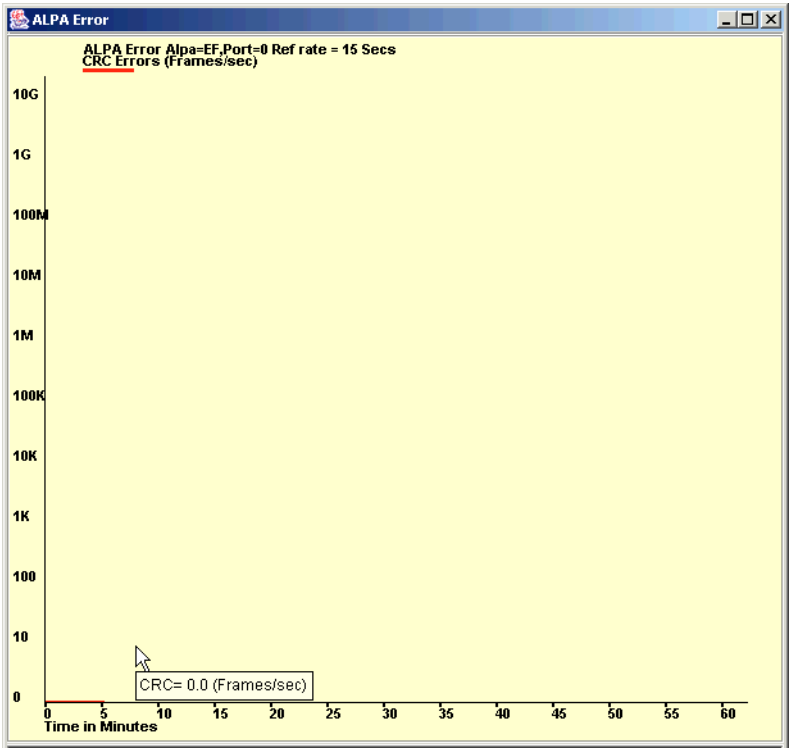


*Figure 1-103 AL_PA error graph*

## Using Advanced Performance Monitoring with Telnet

Three different types of Performance Monitoring can be done using Telnet commands:

- AL_PA monitoring
- End-to-end monitoring
- Filter-based monitoring

### *AL_PA monitoring*

AL_PA monitoring provides information about the number of CRC errors occurring in Fibre Channel frames in a loop configuration. AL_PA monitoring collects CRC error counts for each AL_PA that is attached to a specific port.

AL_PA-based performance monitoring does not require explicit configuration. The switch hardware and firmware automatically monitors CRC errors for all valid AL_PAs.

### *Displaying the CRC Error Count*

Use the **perfShowAlpaCrc** command to display the CRC error count for all AL_PA devices or a single AL_PA on a specific port. The port must be an active L_Port.

Figure 1-104 shows the CRC error count for all AL_PA devices on port 0.


*Figure 1-104 AL_PA CRC error count display*

We can display the CRC error count for one AL_PA by specifying this AL_PA as:

```
perfShowAlpaCrc 0, 0xef
```

### *Clearing the CRC Error Count*

Use the **perfClrAlpaCrc** command to clear the CRC error count for AL_PA devices on a specific port. We can clear the error counts for all the AL_PA devices on a port as shown in Figure 1-105.

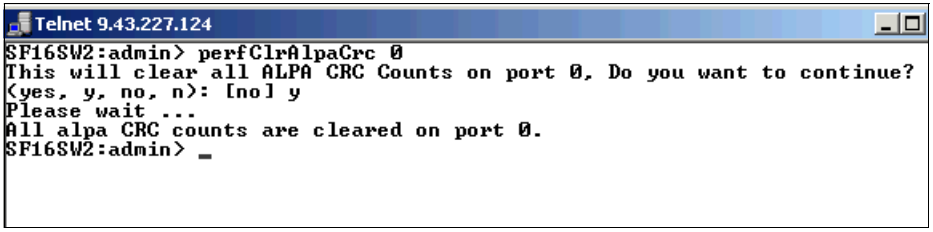


*Figure 1-105 Clear AL_PA CRC error count*

We clear the CRC error count for a specific AL_PA by specifying this AL_PA:

```
perfClrAlpaCrc 0, 0xef
```

## End-to-end monitoring

End-to-end monitoring provides information about transaction performance between the transactions source (SID) and destination (DID) on a fabric or a loop. Up to 16 SID-DID pairs per port can be specified. For each of the SID-DID pairs, the following information is available:

- CRC error count on the frames for the SID-DID pair
- Fibre Channel words transmitted from the port for the SID-DID pair
- Fibre Channel words received by the port for the SID-DID pair

To enable end-to-end performance monitoring, you must configure an end-to-end monitor on a port, specifying the SID-DID pair. The monitor counts only those frames with matching SID and DID.

Each SID or DID has three fields, listed in the following order:

- Domain ID (DD)
- Port ID (AA)
- AL_PA (PP)

For example, the SID `0x118a0f` has domain ID 0x11, Port ID 0x8a, and AL_PA 0x0f (the prefix “0x” denotes a hexadecimal number).

### *Adding End-to-end Monitors*

Use the **`perfAddEEMonitor`** command to add an end-to-end monitor to a port. With this command we specify the port, the SID, and the DID that we want to monitor. Depending on the application, we can select any port along the routing path for monitoring.

Figure 1-106 shows two devices: Host A, which is connected to port 3 on switch 2; and Dev B, which is connected to port 2 on switch 3.

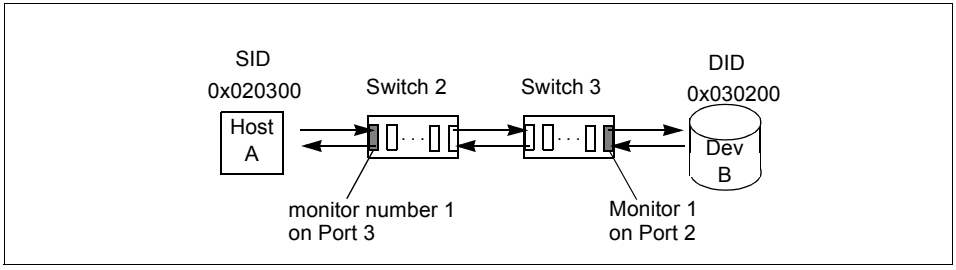


*Figure 1-106 Setting end-to-end monitor on a port*

To monitor the traffic from Host A to Dev B, work on Switch 2 and add a monitor to port 3, specifying 0x020300 as the SID and 0x030200 as the DID. To monitor the traffic from Dev B to Host A, work on Switch 3 and add a monitor to port 2, specifying 0x030200 as the SID and 0x020300 as the DID.

We use **`perfAddEEMonitor`** as shown in Figure 1-107.


*Figure 1-107 Add an end-to-end monitor to switch2 port 3*

As shown in Figure 1-107, monitor number 0 counts the frames that have an SID of 0x020300 and a DID of 0x030200. For monitor number 0, RX_COUNT is the number of words from Host A to Dev B, CRC_COUNT is the number of frames from Host A to Dev B with CRC errors, and TX_COUNT is the number of words from Dev B to Host A.

Note that the monitor must be properly placed as explained in “SID/DID Performance Graph” on page 142.

In Figure 1-106, if we add a monitor to switch2, port 3 specifying Dev B as the SID and Host A as the DID, no counters are incremented:

- Valid: `perfAddEEMonitor 3,"0x020300","0x030200"`
- Not valid: `perfAddEEMonitor 3,"0x030200","0x020300"`

### *Setting a Mask for End-to-End Monitors*

End-to-End monitors count the number of words in Fibre Channel frames that match a specific SID/DID pair. If we want to match only part of the SID or DID, we can set a mask on the port to compare only certain parts of the SID or DID. With no mask set, the frame must match the entire SID and DID to trigger the monitor. By setting a mask, we can choose to have the frame match only one or two of the three fields (Domain ID, Area ID, AL_PA) to trigger the monitor.

**Note:** We can set only one mask per port. The mask is applied to all of the end-to-end monitors on a port. If we subsequently create new monitors on the port, the mask is applied to these new monitors as well. All of the counters are reset when we set the mask.

The mask is specified in the form “dd:aa:pp” where dd is the domain ID mask, aa is the Port ID mask, and pp is the AL_PA mask. The values for dd, aa, and pp are either:

- `ff` (the field must match)
- `00` (the field is ignored).

Use the **`perfSetPortEEMask`** to set a mask for end-to-end monitors. The command sets the mask for all end-to-end monitors of a port.

The **perfSetPortEEMask** command sets a mask for the domain ID, Port ID, and AL_PA of the SIDs and DIDs for frames transmitted from and received by the port. Figure 1-108 shows the mask positions in the command.

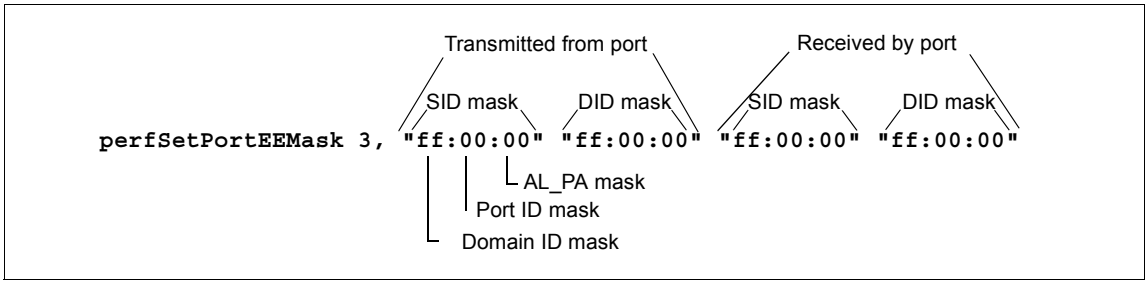


*Figure 1-108 Mask positions for end-to-end monitors*

In Figure 1-108, a mask (“ff”) is set on port 3 to compare the domain ID fields on the SID and DID in all frames (transmitted and received) on port 3. The AL_PA and Port ID fields in all frames are ignored, as no mask is set on these fields.

If we set the following monitor on port 3:

```
perfAddEEMonitor 3,"0x020300","0x030200"
```

Then, without any mask, then the SID must be `0x020300` and the DID must be `0x030200` to trigger the monitor.

If you set the mask shown in Figure 1-108, then the frame SID and DID must match only the domain ID portion of the specified SID-DID pair. That is, frames with SID of “0x02nnnn” and DID of “0x03nnnn” trigger the monitor, where nnnn is any number.

Each port can have only one EE mask. The mask is applied to all end-to-end monitors on the port. You cannot specify individual masks for each monitor on the port. If you define a new end-to-end monitor on a port after you have created a mask for that port, the mask is automatically applied to the new monitor.

The default EE mask value upon power-on is “ff:ff:ff” for everything—SID and DID on all transmitted and received frames.

In Figure 1-109, we use the **perfSetPortEEMask** command to set a mask on the SID and DID domain ID of frames transmitted from switch 2, port 3. After the mask is set, the monitor number created previously in Figure 1-107 on page 153 counts the number of words in incoming Fibre Channel frames that have an SID of 0x02nnnn and a DID of 0x03nnnn, where nnnn is any number.

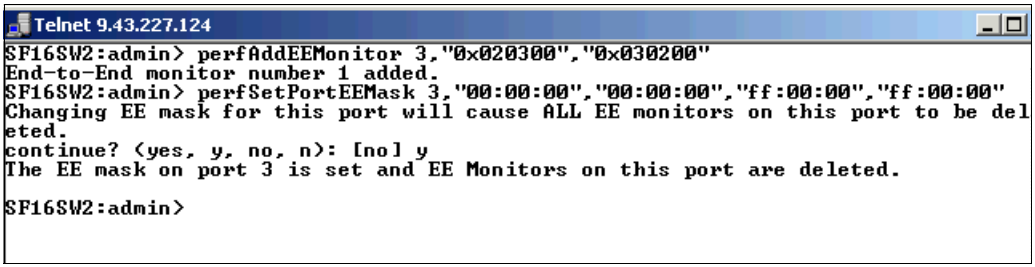


*Figure 1-109 Set a mask on switch2, port 3*

### *Displaying the end-to-end mask of a port*

You can use the **perfShowPortEEMask** command to display the current end-to-end mask of a port as shown in Figure 1-110.

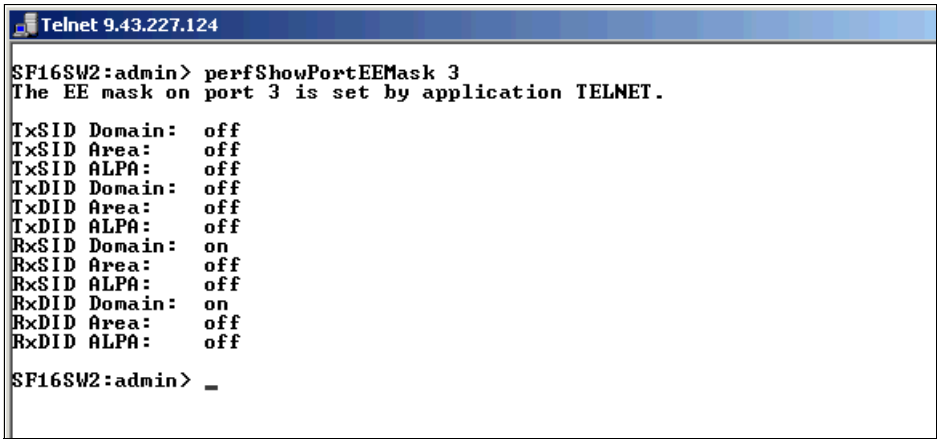


*Figure 1-110 Displaying the end-to-end mask of a port*

The end-to-end mask has 12 fields, with each having a value of `on` or `off`.

### *Displaying the end-to-end monitors*

We use the **perfShowEEMonitor** command to display the end-to-end monitors defined on the port. We can display cumulative counters as shown in Figure 1-111.

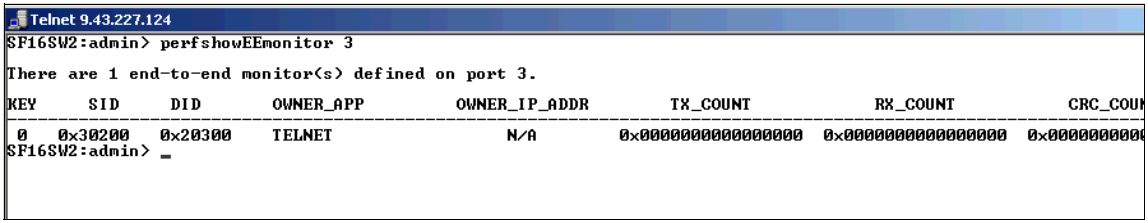


*Figure 1-111 Displaying end-to-end monitor using perfShowEEMonitor*

This command displays:

- `Key`: Monitor number
- `SID`: Source ID
- `DID`: Destination ID
- `OWNER_APP`: TELNET or WEB_TOOLS
- `OWNER_IP_ADDR`: IP address of the owner of the filter monitor
- `TX_COUNT`: Transmitting frame count
- `RX_COUNT`: Receiving frame count
- `CRC_COUNT`: CRC error count

The cumulative counters are 64-bit values in hexadecimal format.

If we specify an interval number in the **`perfShowEEMonitor`** command, the command displays a rolling table of CRC error, Tx, and Rx counters on a per-interval basis for all the valid monitors on the port as shown in Figure 1-112. The counter values are the number of bytes, in decimal format.

```
switch2:admin> perfShowEEMonitor 3,6

perfShowEEMonitor 3, 6: Tx/Rx are # of bytes and crc is # of crc errors

       0              1
 -------------  -------------
 crc    Tx    Rx  crc    Tx    Rx
 =============  =============
   0     0     0    0     0     0
   0   178m   87m    0   178m   87m
   0   155m   89m    0   155m   89m
   0   174m   85m    0   174m   85m
   0   168m   89m    0   168m   89m
   0   205m   85m    0   205m   85m
   0   178m   88m    0   178m   88m
   0   163m   87m    0   163m   87m
   0   186m   86m    0   186m   86m

switch2:admin>
```

*Figure 1-112 Displaying end-to-end monitor with a interval*

The counter values in Figure 1-112 are the number of bytes in decimal format. The “m” stands for megabytes. You may also see “g” which stands for gigabytes, or “k” which stands for kilobytes.

**Note:** The minimum interval value that can be specified is 5 seconds.

### *Deleting end-to-end monitors*

Use the **perfDelEEMonitor** command to delete an end-to-end monitor on a port as shown in Figure 1-113. Indicate which monitor to delete by specifying the monitor number that was returned by a previous **perfAddEEMonitor** command.

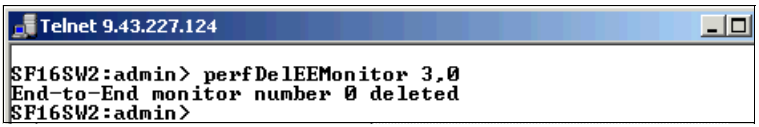


*Figure 1-113 Deleting end-to-end monitors*

The following command deletes all of the end-to-end monitors on port 2:

```
sw1:admin> perfDelEEMonitor 2
This will remove ALL EE monitors on port 2, continue? [y|n]y
```

### *Clearing end-to-end monitor counters*

To clear all of the end-to-end monitor counters on a port, use the **perfSetPortEEMask** command to reset all of the end-to-end monitor counters on that port.

The **perfSetPortEEMask** command also sets the end-to-end mask, so if you do not want to change the mask, you must re-specify the current mask settings. You can view the current mask settings using the **perfShowPortEEMask** command.

To clear the counters for a single end-to-end monitor, delete the monitor using the **perfDelEEMonitor** command, and then add the monitor again, using the **perfAddEEMonitor** command.

## Filter-based monitoring

Filter-based monitoring provides information about a filter’s hit count. Any parameter in the first 64 bytes of the Fibre Channel frame can be measured. The counter increases each time a frame is filtered through the corresponding port. Examples of port filter statistics that can be measured are:

- SCSI read, write, or read/write commands
- CRC error statistics (port and AL_PA)
- IP versus SCSI traffic comparison

The filter can be a standard filter (for example, a read command filter that counts the number of read commands that have been received by the port) or a user-defined filter that you customize for your particular use.

The maximum number of filters is eight per port, in any combination of standard filters and user-defined filters.

### *Adding standard filter-based monitors*

This section describes how to add standard filter-based monitors to a port. Use the telnet commands listed in Figure 1-21to define filter-based monitors on a port.

*Table 1-21 Add Filter based monitor commands*

| Command | Description |
|---|---|
| perfAddReadMonitor | Count the number of SCSI Read commands |
| perfAddWriteMonitor | Count the number of SCSI Write commands |
| perfAddRWMonitor | Count the number of SCSI Read and Write commands |
| perfAddSCSIMonitor | Count the number of SCSI traffic frames |
| perfAddIPMonitor | Count the number of IP traffic frames |

In Figure 1-114 we add several filter monitors to switch2, port 3.

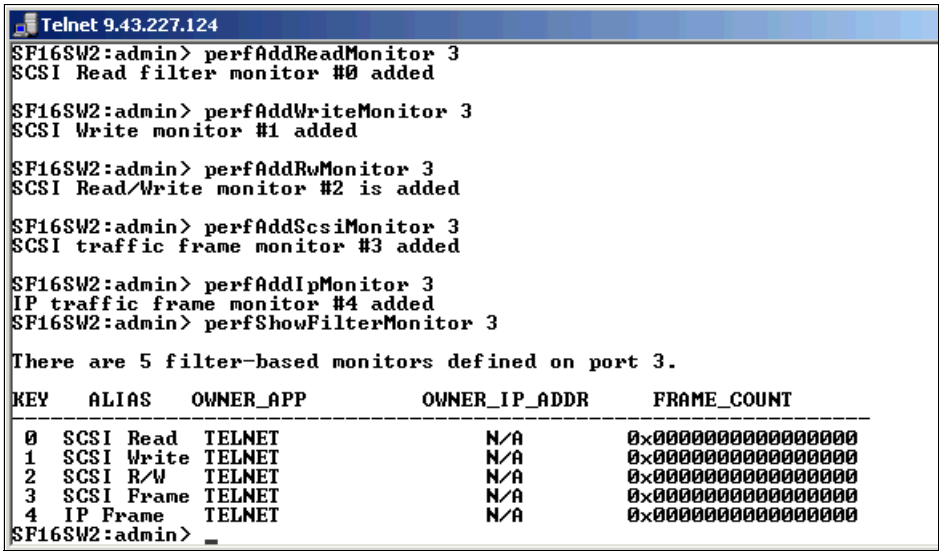


*Figure 1-114 Adding filter monitors to a port*

### *Adding user-defined filter-based monitors*

In addition to the standard filters (read, write, read/write, and frame count), you can create custom filters to qualify frames for statistics gathering to fit your own special needs.

To define a custom filter, use the **perfAddUserMonitor** telnet command. With this command, you must specify a series of offsets, masks, values and an alias for the monitor. The following actions are performed. For all incoming frames, the switch:

1. Locates the byte found in the frame at the specified offset
2. Applies the mask to the byte found in the frame
3. Compares the value with the given values in the **perfAddUserMonitor** command
4. Increments the filter counter if a match is found

You can specify up to six different offsets for each port, and up to four values to compare against each offset.

If more than one offset is required to properly define a filter, the bytes found at each offset must match one of the given values for the filter to increment its counter. If one or more of the given offsets does not match any of the given values, the counter does not increment.

The value of the offset must be between 0 and 63, in decimal format. Byte 0 indicates the first byte of the Start of Frame (SOF), byte 4 is the first byte of the frame header, and byte 28 is the first byte of the payload. Thus only the SOF, frame header, and first 36 bytes of payload may be selected as part of a filter definition.

### *Displaying filter-based monitors*

Use the **perfShowFilterMonitor** command to display all the filter-based monitors of a port. You can display a cumulative count of the traffic detected by the monitors, or you can display a snapshot of the traffic at specified intervals.

**Note:** Intervals must be specified in multiples of 5 seconds, for example, 5, 10, 15, 20, 25, etc., because registers are scanned every 5 seconds.

This command displays all the filter-based monitors defined on the specified port. It displays all the valid monitor numbers and user-defined aliases on the specified port.

Figure 1-115 shows the traffic at a specified interval of six seconds on port 0.

```
switch2:admin> perfshowFilterMonitor 0,6

perfShowFilterMonitor 0, 6

    0         1         2         3         4
  #CMDs     #CMDs     #CMDs    #Frames   #Frames
-------------------------------------------------
   7         1         8         15        0
 937       429       1.3k      3.7k      0
 1.0k      444       1.4k      3.9k      0
 851       440       1.2k      3.7k      0
 885       444       1.3k      3.8k      0
 981       435       1.4k      3.8k      0
 863       432       1.2k      3.6k      0
 886       433       1.3k      3.6k      0
 982       449       1.4k      3.8k      0
 862       426       1.2k      3.6k      0
 895       440       1.3k      3.7k      0
 757       429       1.1k      3.5k      0
 877       454       1.3k      3.8k      0

switch2:admin>
```

*Figure 1-115 Displaying filter monitor*

**Note:** A defined filter will only increment if set on receiving ports.

### *Deleting filter-based monitors*

To delete a filter-based monitor, first list the valid monitor numbers using the **perfShowFilterMonitor** command, then use the **perfDelFilterMonitor** command to delete a specific monitor. If you do not specify which monitor number to delete, you will be asked if you want to delete all entries.

## 1.7.13 Fabric Watch

Fabric Watch monitors key fabric and switch elements, making it easy to quickly identify and escalate potential problems. It monitors each element for out-of-boundary values or counters and provides notification when any exceed the defined boundaries. Fabric Watch can configure elements, such as error status, and performance counters within a switch, and how they are monitored. If an element exceeds the specified threshold or trigger value, Fabric Watch will issue an alert. This can be in the form of writing to the event log, logging to the port log, issuing an SNMP trap, or sending an e-mail (or a combination of any of these).

The Fabric Watch feature monitors the performance and status IBM TotalStorage SAN Switch, and can alert SAN managers when problems arise. The real-time alerts from Fabric Watch software help SAN managers solve problems before they become costly failures. SAN managers can configure Fabric Watch software to monitor any of the following occurrences:

- Fabric events (such as topology re-configurations and zone changes)
- Physical switch conditions (such as fans, power supplies, and temperature)
- Port behavior (such as state changes, errors, and performance)
- SFPs (for switches equipped with SMART SFPs)

### Range monitoring

With Fabric Watch, each switch continuously monitors error and performance counters against a set of defined ranges. This and other information specific to each monitored element is made available by Fabric Watch for viewing and, in some cases, modification. This set of information about each element is called a *threshold*, and the upper and lower limits of the defined ranges are called *boundaries*.

If conditions break out of acceptable ranges, an *event* is considered to have occurred, and one or more *alarms* (reporting mechanisms) are generated if configured for the relevant threshold. There are three types of alarms:

- SNMP trap
- Entry in the switch event log
- Locking of the port log to preserve the relevant information

For more information, refer to the following Web site:

http://www.storage.ibm.com/ibmsan/index.htm

## Element categories

Fabric Watch elements include any component of the fabric or switch that Fabric Watch software monitors. To monitor elements, Fabric Watch software categorizes them into areas, and groups these areas into classes.

### *Classes*

Classes (also known as agents) are high-level categories of elements. Fabric Watch software monitors elements that compose the following classes:

- Fabric
- Environment
- Port (includes E_Port, Optical F/FL_Port, Copper F/FL_Port)
- SPF
- Performance Monitor (AL_PA, End-to-End, Filter)

### *Areas*

Areas are the behaviors that Fabric Watch software monitors. Table 1-22 lists all Fabric Watch classes, the areas within those classes, and a description of each area.

*Table 1-22 Fabric Watch Classes and Area*

| Class | Area | Area Description |
|---|---|---|
| Fabric | E_Ports downs | Monitors E_Port status. |
| | Fabric Reconfigure | Monitors changes to the fabric configuration. |
| | Domain ID Changes | Monitors forcible domain ID changes. |
| | Segmentation Changes | Monitors segmentation changes. |
| | Zone Changes | Monitors changes to currently enabled zoning configurations. |
| | Fabric <-> QL | Monitors changes to QuickLoop |
| | Fabric logins | Monitors the number of host device fabric logins (FLOGI). |
| | SFP State Change | Monitors insertion/removal of smart SFP. |
| Environmental | Temperature | Monitors switch temperature in degrees Celsius. |
| | Fan | Monitors switch fan speed in RPMs. |

| Class | Area | Area Description |
|---|---|---|
| Port | Link Loss | Monitors the link failure rate of each port. Tracks the number of link failures per configured time interval. |
| | Sync Loss | Monitors the number of synchronization loss errors per configured time interval. |
| | Signal Loss | Monitors the number of signal loss errors per configured time interval. |
| | Protocol Error | Monitors the number of protocol errors per configured time interval. |
| | Invalid Words | Monitors the number of invalid words transmitted (from a device to a port) per configured time interval. |
| | Invalid CRCs | Monitors the number of CRC errors per configured time interval. |
| | Rx Performance | Monitors receive rate in KB/sec. |
| | Tx Performance | Monitors transmit rate in KB/sec. |
| | State Changes | Monitors state changes. |
| SFP | Temperature | Monitors SFP temperature in degrees Celsius. |
| | Rx Power | Monitors SFP receiver power in uWatts. |
| | Tx Power | Monitors SFP transmitter power in uWatts |
| | Current | Monitors SFP current in mAmps. |
| | Voltage | Monitors SFP power in mVolts. |
| Performance Monitor | CRC Errors | Monitors the number of CRC errors that occur (for AL_PA or for a SiD-DiD pair) per configured time interval (in seconds). |
| | FCW Received | Monitors receive rate of a SiD-DiD pair in KB per second. |
| | FCW Transmitted | Monitors transmit rate of a SiD-DiD pair in KB per second. |
| | Custom Filter Counter | Monitors the filter-based counter that the user defines. |

## Accessing Fabric Watch

To access the Fabric Watch function, click the “magnifying glass” button (labeled **Watch**) from the Switch View, as shown in Figure 1-116.

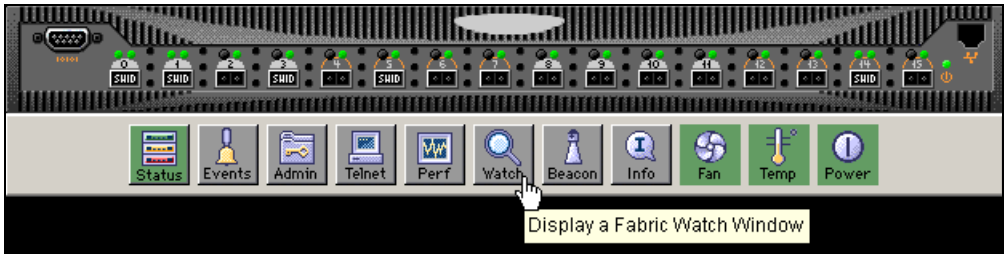


*Figure 1-116 Go to Fabric Watch*

The Fabric Watch window, as shown in Figure 1-117, is then displayed.

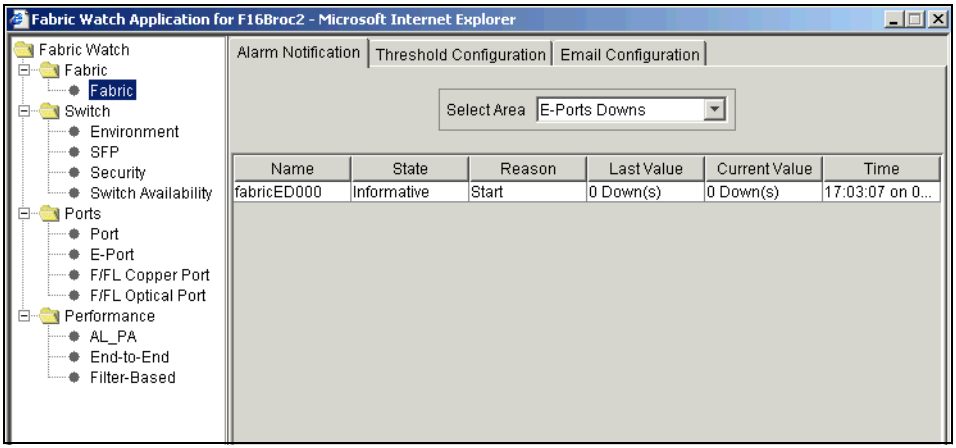


*Figure 1-117 Fabric Watch panel*

The window is divided into two sections. The left-hand side has a tree structure that lists the *Classes* that can be monitored using Fabric Watch. If you expand the *Classes*, all the *Areas* that are associated with a particular *Class* are displayed.

The main part of the window on the right-hand side has a display with three tabs:

- Alarm Notifications tab
- Thresholds Configuration tab
- Email Configuration tab

## Alarm Notifications

Use the **Alarm Notifications** tab to view the information for all elements of the Fabric Watch, Fabric, or Performance Monitor classes. The information displayed includes:

- The name of the threshold
- The current value
- The last event type
- The last event time
- The last event value
- The last event state

The Alarm Notifications will refresh the displayed information according to the threshold configuration.

The **Alarm Notifications** tab is shown in Figure 1-118.

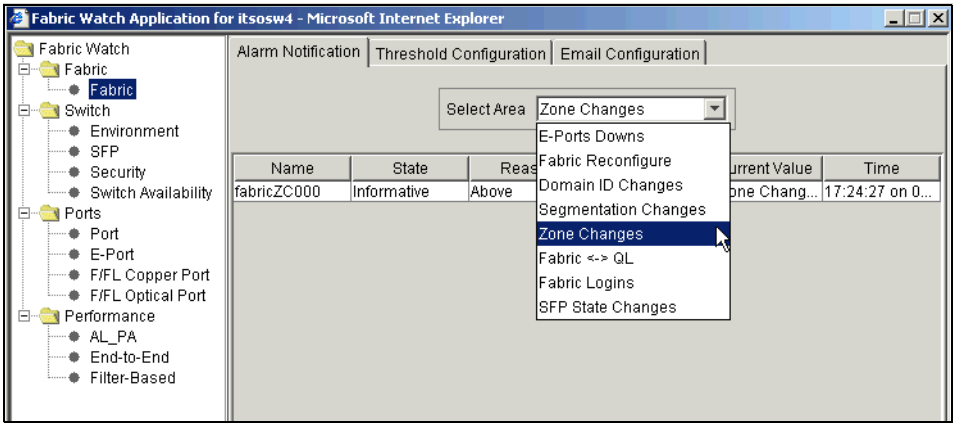


*Figure 1-118 Alarm Notifications*

## Configure Thresholds

Use the **Configure Thresholds** tab to view and configure Fabric Watch thresholds for the Fabric Watch class currently selected in the organizational tree on the left side of the window. The **Configure Thresholds** tab is shown in Figure 1-119.

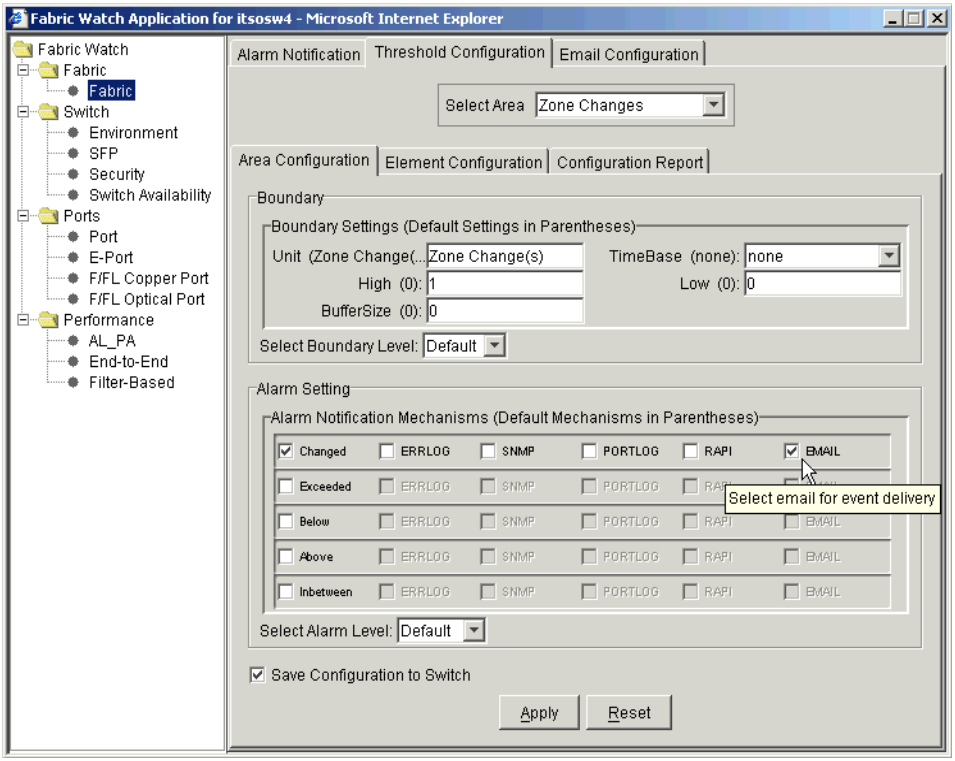


*Figure 1-119 Configure Thresholds*

The Configure Threshold display changes according to the Class and Area selected in the organizational tree. However, the **Configure Thresholds** tab always contains the same buttons described below.

### *Default*

Click to return settings to default values.

### *Apply*

Click to apply the values specified in the current display. When **Apply** is clicked after a change, the Alarm Mechanism dialog displays as shown in Figure . This dialog box is used to specify alarm delivery mechanisms by checking the appropriate boxes.

When e-mail is chosen as one of the alarm delivery mechanisms, the text fields are enabled in the “Mail Server, “Domain Name”, “Mail From” and “Mail To” information to configure the E-mail Alert by Fabric Watch Daemon function. If E-mail is not chosen as one of the delivery mechanisms, the text fields are grayed out.

When RAN is chosen, changes apply to the API application, but not regular usage.

To continue, select the type of alarm you want and click **Apply**.

Note that changes made within this window will apply to all areas within the class you are working.

### *Reset*

Undo the changes.

## Thresholds for the Environmental classes

The Environmental classes are displayed by selecting the **Environmental** file from the separate column left of the **Configure Thresholds** tab as shown in Figure 1-120.

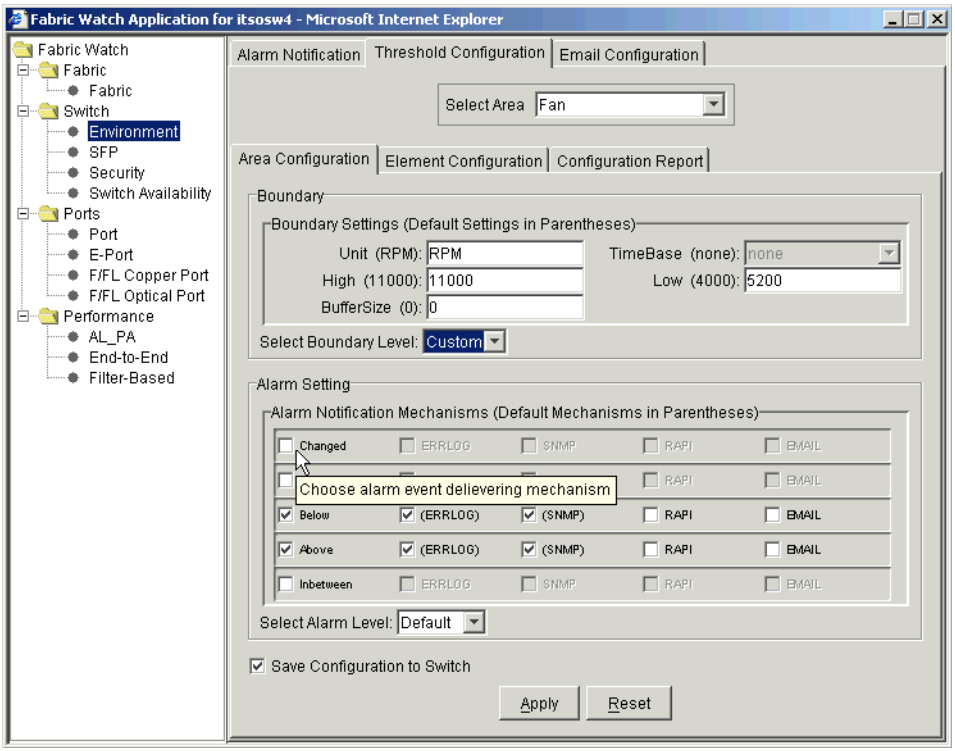


*Figure 1-120 Environmental Thresholds*

The panel contains three columns into which you browse by clicking one of these check boxes: **Temperature** or **Fan**. Only one area can be active at a time. The following sections discuss the parameters available in each area.

### *Temperature and Fan*

These parameters are available:

- **Threshold Type drop-down list:** Select the threshold type (exceeded, above, in-between, below, changed).
- **High/Low drop-down list:** Select to enter the high and low settings for the threshold type selected in the Threshold Type drop-down list (not available for all areas).
- **Thresh element drop-down list:** Check or uncheck to specify if you want this element to be monitored.

## Thresholds for the SFP Classes

The SFP classes are displayed by selecting the **SFP** file from the separate column left of the **Configure Thresholds** tab as shown in Figure 1-121.

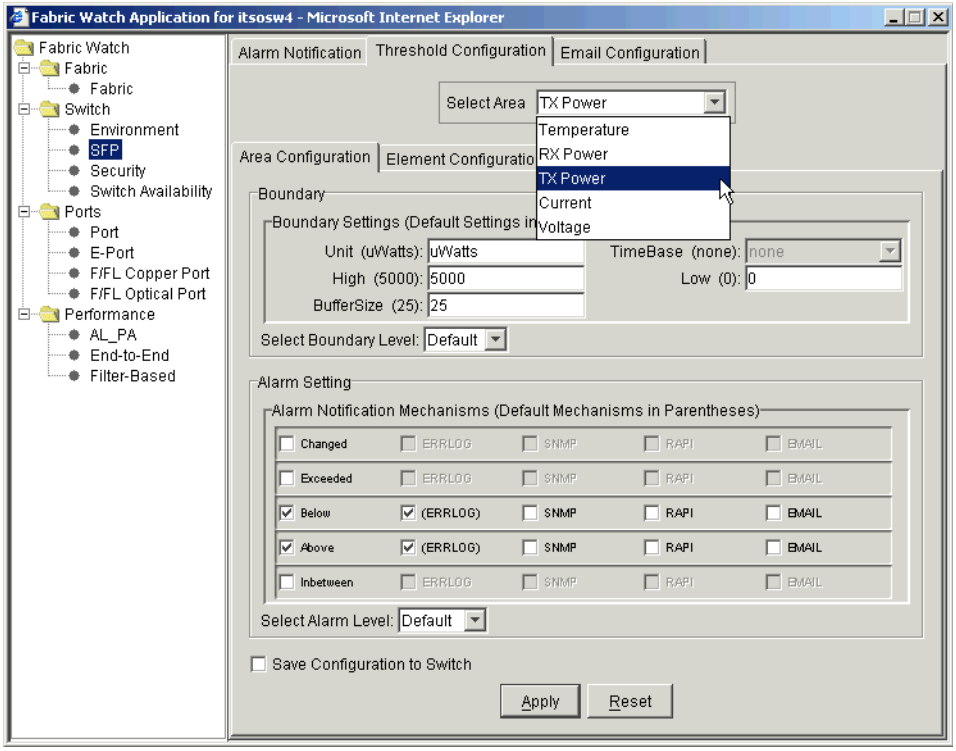


*Figure 1-121 SFP thresholds*

The columns for the Temperature, RX Power, TX Power, and Current areas contain the following items:

- **Threshold Type drop-down list:** Select the threshold type (exceeded, above, in-between, below, changed).
- **High/Low drop-down list:** Select to enter the high and low settings for the threshold type selected in the Threshold Type drop-down list (not available for all areas).
- **Scale:** Specify the point at which you want to set the specified threshold.
- **Area Select checkbox:** Select a Fabric Watch area to configure. Only one area can be selected at a time.

### Thresholds for the remaining classes

The Port, E_Port, F/FL Copper Port, F/FL Optical Port classes display the following fields for each area (Link Loss, Sync Loss, Signal Loss, Protocol Error, Invalid Words, Invalid CRCs, State Changes, RX Performance, TX Performance):

- **Low text box:** Enter the low threshold boundary.
- **High text box:** Enter the high threshold boundary.
- **Threshold Type drop-down list:** Select the type of threshold.
- **Time period drop-down list:** Select the time period for which you want a time-based count to be measured.

The thresholds for the Port class are displayed as shown in Figure 1-122.

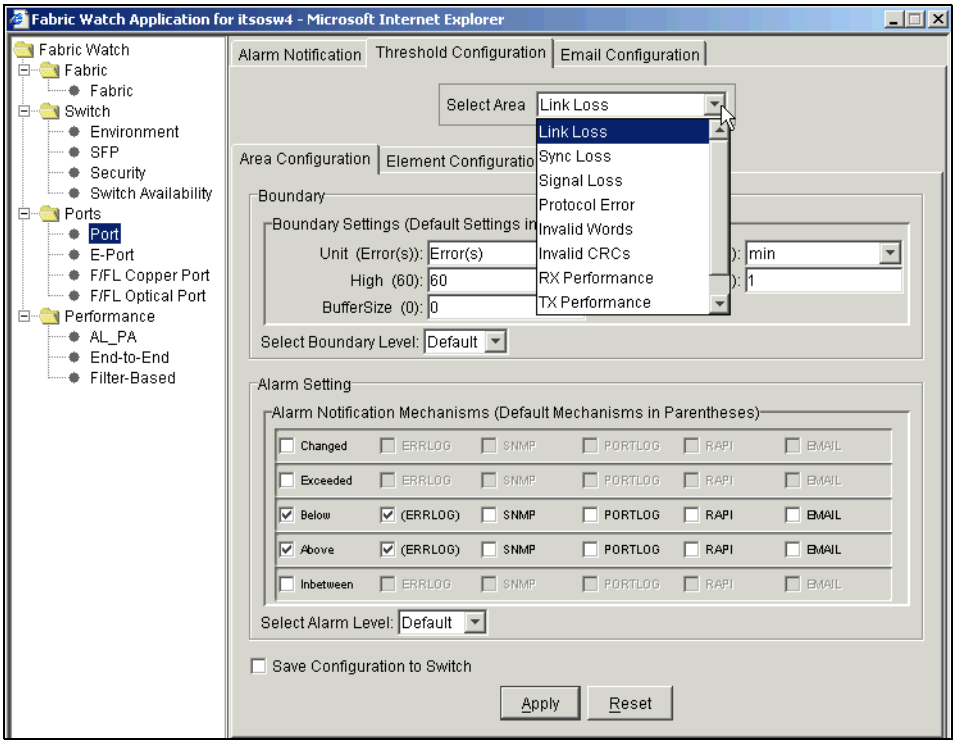


*Figure 1-122 Port Thresholds*

Use the **Configure Thresholds** tab to view and configure End-to-End thresholds for the Performance class currently selected in the organizational tree on the left side of the window.

Note that you must define the SID/DID pair through the Performance Monitor before you can monitor the threshold in the End-to-End class. For more information on the Performance Monitor, refer to “Performance Monitor” on page 127.

The **Configure Thresholds** tab for the End-to-end Thresholds is shown in Figure 1-123.

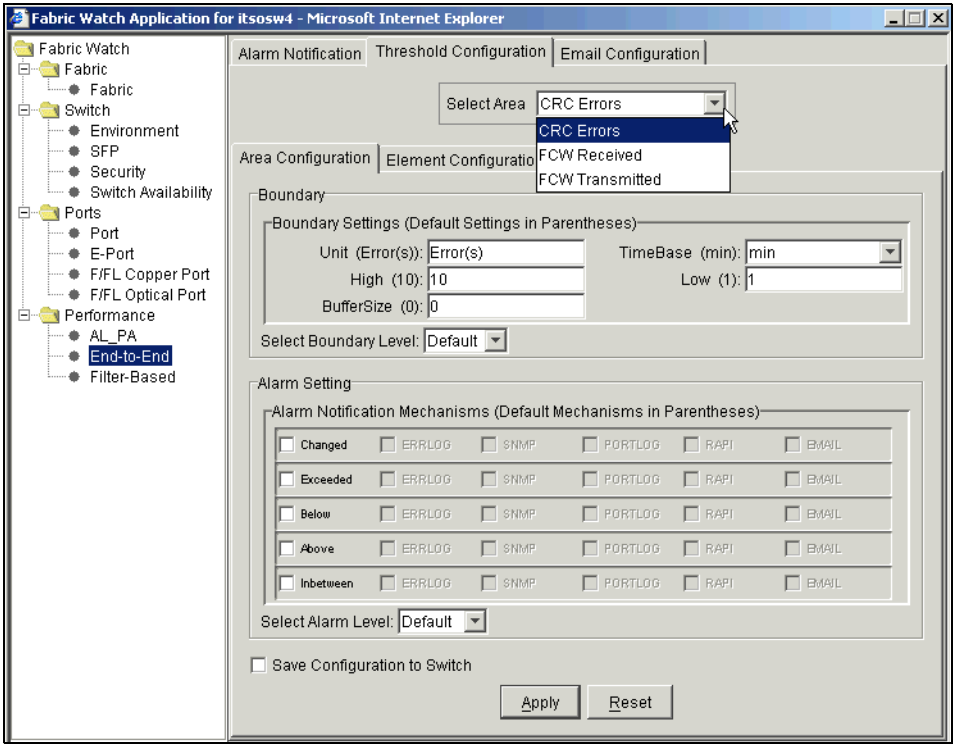


*Figure 1-123 Thresholds Tab with End-to-End class selected in Performance View*

Use the **Configure Thresholds** tab to view and configure Filter-based thresholds for the Performance class currently selected in the organizational tree on the left side of the window.

Note that the filter type must be predefined in the Performance Monitor before you can use the Filter-Based thresholds. For more information on the Performance Monitor, refer to “Performance Monitor” on page 127.

The **Configure Thresholds** tab is shown in Figure 1-124.

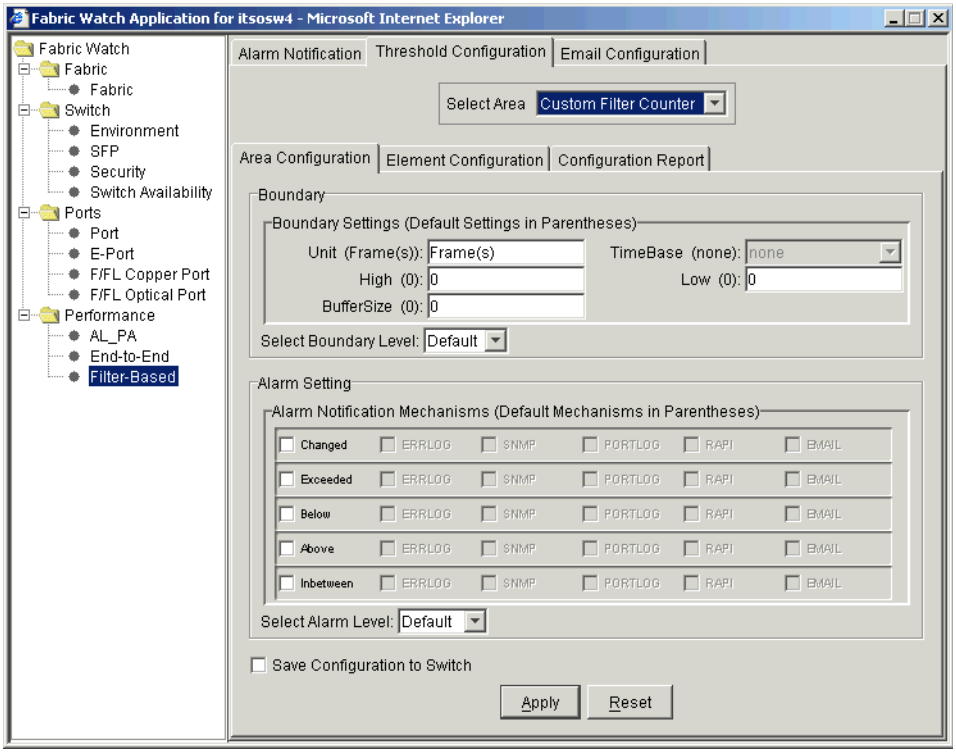


*Figure 1-124 Thresholds Tab with Filter Based class selected in Performance View*

## Current Settings tab

Use the **Current Settings** tab to view the current Fabric Watch threshold parameters for the area selected in the Fabric Watch tree.

The **Current Settings** tab is shown in Figure 1-125.

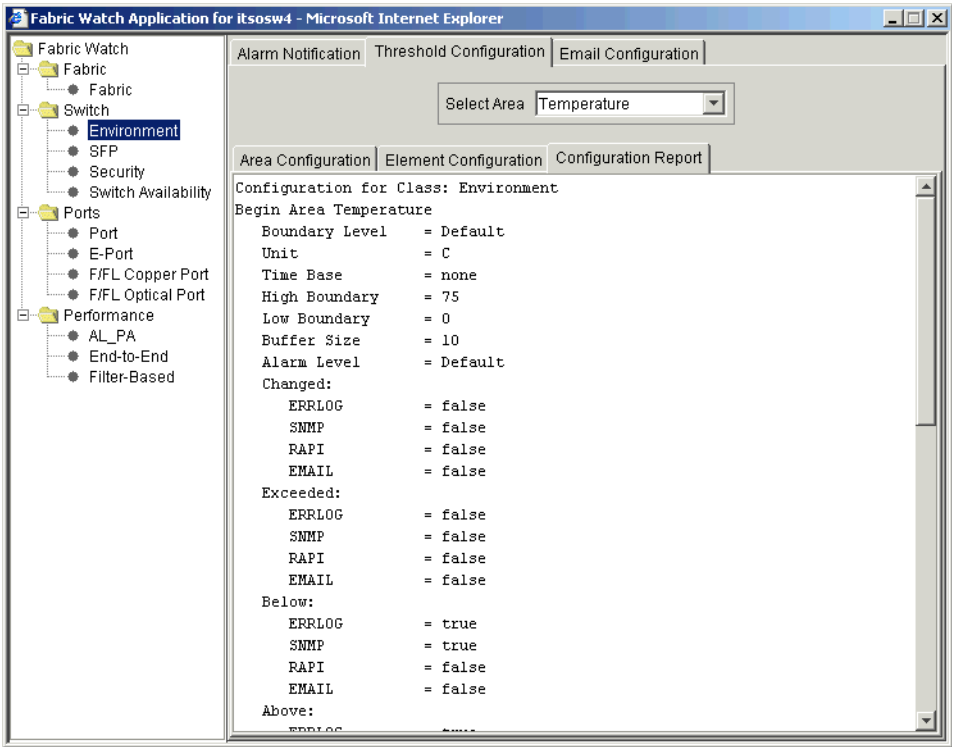


*Figure 1-125 Current Settings Tab in the Fabric Watch View*

## Modifying settings for switches with one power supply

The IBM default settings for Fabric Watch will cause a switch with a single power supply to appear yellow in the WEB TOOLS, indicating a *MARGINAL* status. The status can also be clicking the Status button in the switch view, this opens a window describing the cause of our marginal state as shown in Figure 1-126.

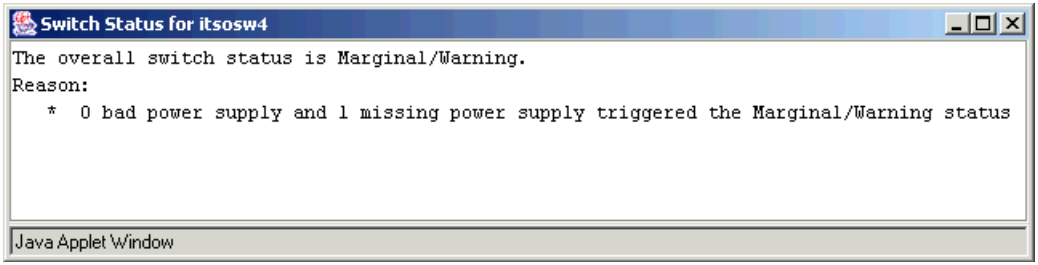


*Figure 1-126 Checking the switch status*

The switch status can be changed to HEALTHY using a Telnet connection. Figure 1-127 shows the command that we issued to change the status of the switch to ensure that a switch with only one power supply is shown with a *HEALTHY* status.

```
itsosw4:admin> switchstatuspolicyset

The current overall switch status policy parameters:
                  Down    Marginal
 ---------------------------------
       FaultyPorts  2         1
       MissingSFPs  0         0
     PowerSupplies  2         1
      Temperatures  2         1
             Fans  2         1
        PortStatus  0         0
         ISLStatus  2         1

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


The minimum number of
  FaultyPorts contributing to DOWN status: (0..16) [2]
  FaultyPorts contributing to MARGINAL status: (0..16) [1]
  MissingSFPs contributing to DOWN status: (0..16) [0]
  MissingSFPs contributing to MARGINAL status: (0..16) [0]
  Bad PowerSupplies contributing to DOWN status: (0..2) [2] 0
  Bad PowerSupplies contributing to MARGINAL status: (0..2) [1] 0
  Bad Temperatures contributing to DOWN status: (0..3) [2]
  Bad Temperatures contributing to MARGINAL status: (0..3) [1]
  Bad Fans contributing to DOWN status: (0..4) [2]
  Bad Fans contributing to MARGINAL status: (0..4) [1]
  Down PortStatus contributing to DOWN status: (0..16) [0]
  Down PortStatus contributing to MARGINAL status: (0..16) [0]
  down ISLStatus contributing to DOWN status: (0..16) [2]
  down ISLStatus contributing to MARGINAL status: (0..16) [1]

Policy parameter set has been changed
Committing configuration...0x1026ec50 (tThad): Jul 18 11:04:21
    WARNING FW-STATUS_SWITCH, 3, Switch status changed from Marginal/Warning toK

done.
itsosw4:admin>
```

*Figure 1-127 Changing the default setting*

To change the default settings, we issue the command: **switchstatuspolicyset**.

The first section of response to the command is the same as if we had issued the **switchstatuspolicyshow** command and displays a list of the current settings. Here we can see that the *PowerSupplies* line is defined to be Marginal if the switch is powered by one power supply. These default settings assume that the switch has two power supplies and that one has failed. Obviously, for a switch purchased with a single power supply, this is not valid.

We are then prompted to enter the new values for each setting, starting with the *DOWN* value for the Faulty Ports, then the *MARGINAL* value for Faulty Ports. We press Enter to use default values; we are prompted for the next setting, and eventually, for the Power supply *DOWN* and *MARGINAL* values. We enter zero for the number of *bad power supplies contributing to the DOWN status* and zero for the number of *bad power supplies contributing to the MARGINAL status*. Indeed, as we are working with only one power supply, if this power supply goes down, then the whole switch goes down. There is no marginal status.

At the bottom of the Telnet display in Figure 1-127, after our change to the policy parameter takes affect, Fabric Watch (FW) issues a message indicating that the status of the switch has changed from *MARGINAL* to *HEALTHY*.

## Email Configuration

Use the **Email Configuration** tab to configure the destination e-mail ID to receive any alerts selected in the threshold configuration to deliver to e-mail as shown in Figure 1-128. Also on this tab, we are able to generally enable or disable the e-mail function for fabric Watch alerts, and send a test e-mail to ensure that the function is working.

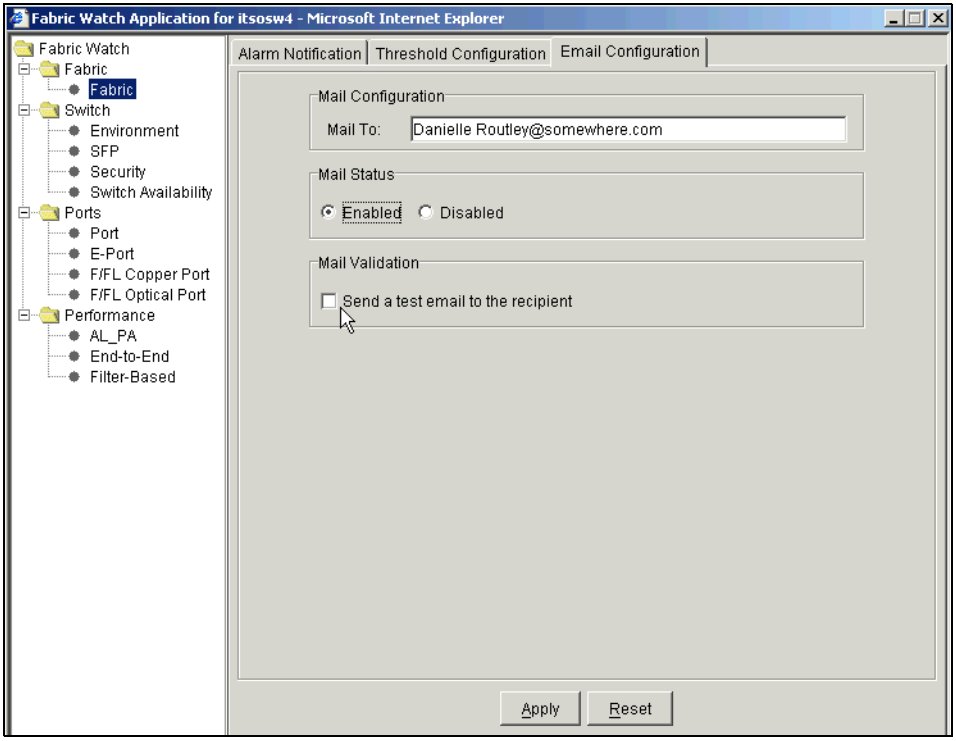


*Figure 1-128 Email Configuration tab*

### 1.7.14 Beaconing

The Beaconing function will locate a switch by sending a signal to the specified switch, which causes an LED yellow light pattern to flash from side to side of the switch. This makes the switch very easy to find.

To activate Beaconing, click the lighthouse icon on the Switch View as shown in Figure 1-129.


*Figure 1-129 Start Beaconing*

This function can be toggled on and off once the switch is identified.

## 1.8 Merging SAN fabrics

Merging a SAN fabric occurs where two or more separate fabrics are combined. An example of this is shown in Figure 1-130.

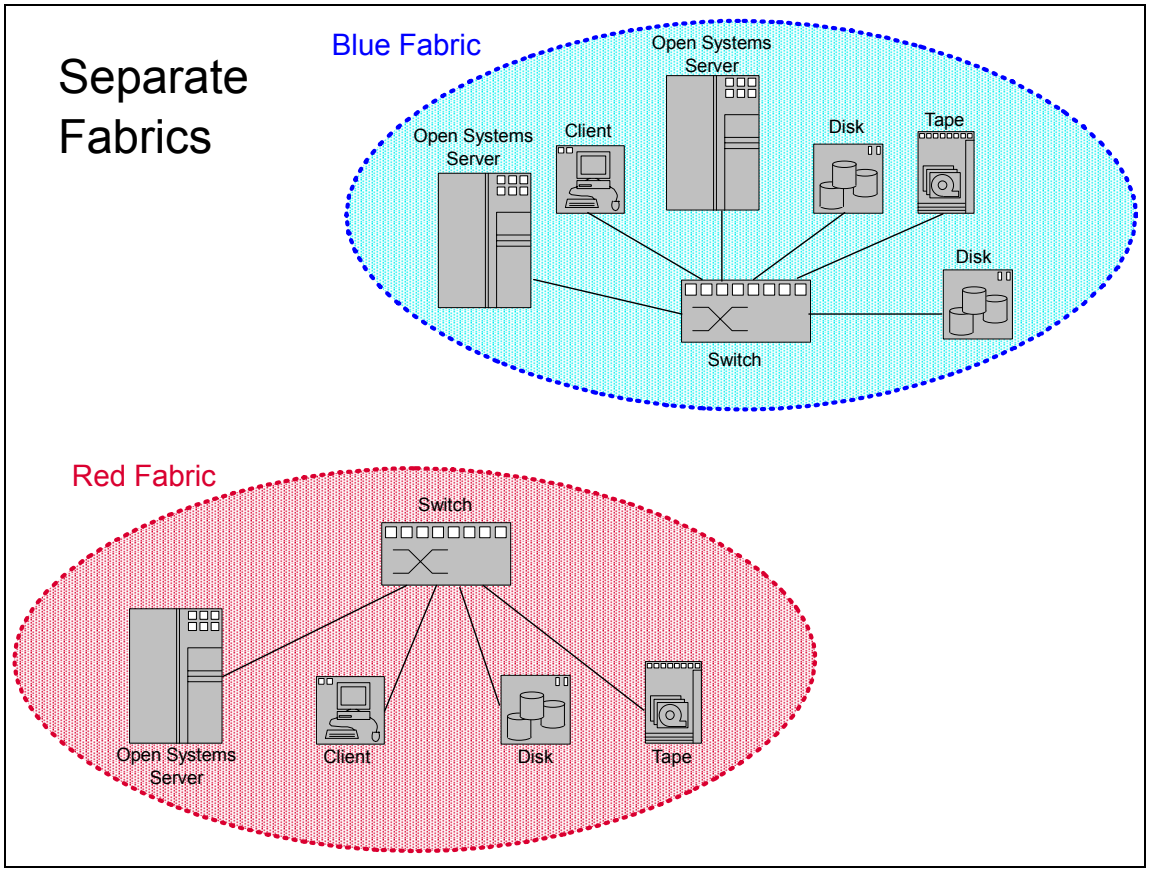


*Figure 1-130 Two separate SAN fabrics*

These separate SAN fabrics can be merged to form a larger SAN fabric by connecting the switches using an Inter-Switch Link (ISL) as shown in Figure 1-131.

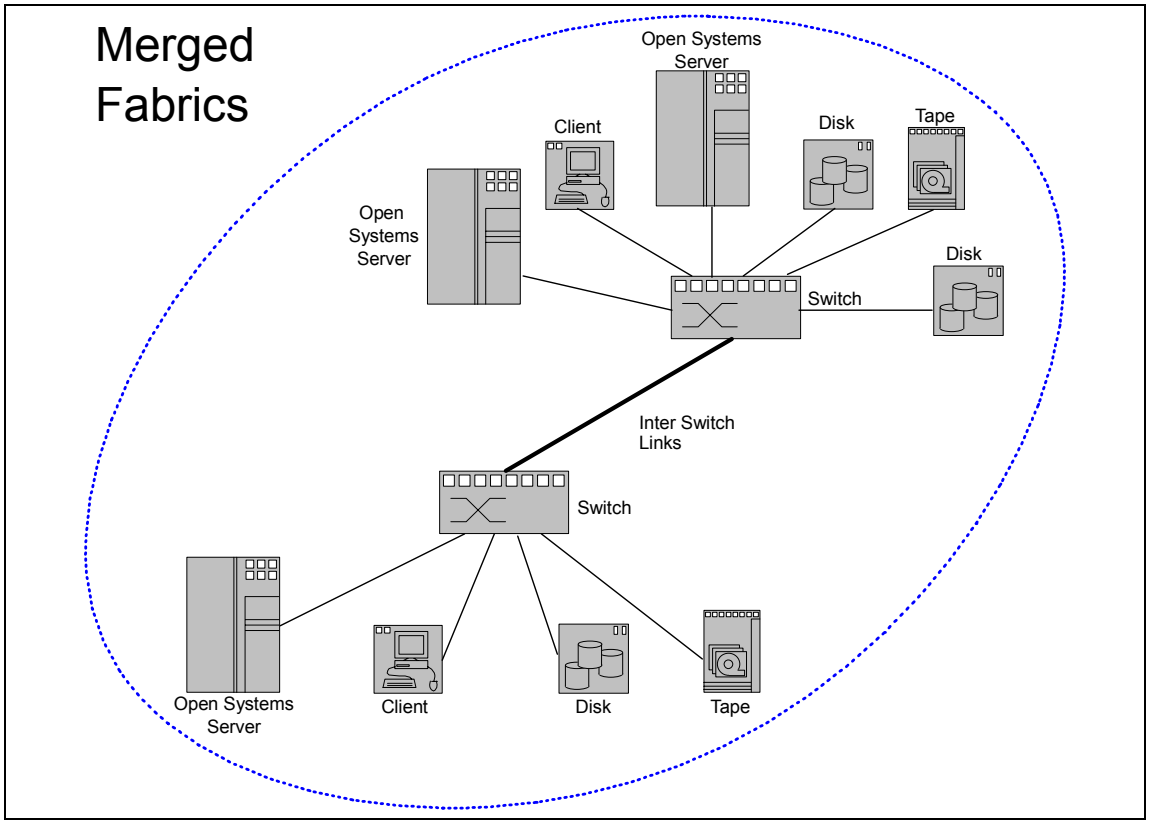


*Figure 1-131 A merged fabric*

The zoning information for each fabric is retained as are the domain IDs for the switches, assuming that there are no conflicting definitions.

This could happen when an organization acquires another company or when two business units within one company merge. The result is that a SAN fabric is extended through the addition of another complete fabric.

**Important:** You should always disable a switch before adding it to an existing fabric.

Some conflicts may occur as two fabrics are merged. Some of the most common sources of conflict are:

- Duplicate domain ID
- Zoning configuration conflicts
- Operating parameters inconsistency (for example, Core PID format)

When this occurs, part of the SAN fabric is said to be *segmented*. You can identify a segmentation from the slow flashing orange LED on the ISL port.

The following section describes these three conflicts and their possible solution.

## 1.8.1 Duplicate domain IDs

Domain IDs are used to uniquely identify a switch within a fabric. Therefore, each switch within the same fabric must have a unique domain ID. Duplicate domains causes the ISL between the two switches to be segmented as shown in Figure 1-132.

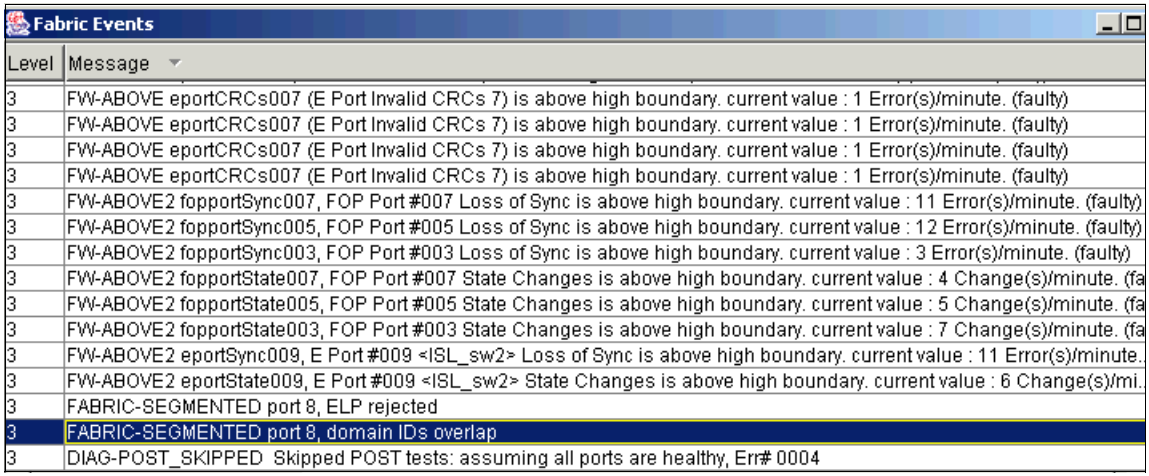
Fabric Events

| Level | Message |
|---|---|
| 3 | FW-ABOVE eportCRCs007 (E Port Invalid CRCs 7) is above high boundary. current value : 1 Error(s)/minute. (faulty) |
| 3 | FW-ABOVE eportCRCs007 (E Port Invalid CRCs 7) is above high boundary. current value : 1 Error(s)/minute. (faulty) |
| 3 | FW-ABOVE eportCRCs007 (E Port Invalid CRCs 7) is above high boundary. current value : 1 Error(s)/minute. (faulty) |
| 3 | FW-ABOVE eportCRCs007 (E Port Invalid CRCs 7) is above high boundary. current value : 1 Error(s)/minute. (faulty) |
| 3 | FW-ABOVE2 fopportSync007, FOP Port #007 Loss of Sync is above high boundary. current value : 11 Error(s)/minute. (faulty) |
| 3 | FW-ABOVE2 fopportSync005, FOP Port #005 Loss of Sync is above high boundary. current value : 12 Error(s)/minute. (faulty) |
| 3 | FW-ABOVE2 fopportSync003, FOP Port #003 Loss of Sync is above high boundary. current value : 3 Error(s)/minute. (faulty) |
| 3 | FW-ABOVE2 fopportState007, FOP Port #007 State Changes is above high boundary. current value : 4 Change(s)/minute. (fa |
| 3 | FW-ABOVE2 fopportState005, FOP Port #005 State Changes is above high boundary. current value : 5 Change(s)/minute. (fa |
| 3 | FW-ABOVE2 fopportState003, FOP Port #003 State Changes is above high boundary. current value : 7 Change(s)/minute. (fa |
| 3 | FW-ABOVE2 eportSync009, E Port #009 <ISL_sw2> Loss of Sync is above high boundary. current value : 11 Error(s)/minute. |
| 3 | FW-ABOVE2 eportState009, E Port #009 <ISL_sw2> State Changes is above high boundary. current value : 6 Change(s)/mi. |
| 3 | FABRIC-SEGMENTED port 8, ELP rejected |
| 3 | FABRIC-SEGMENTED port 8, domain IDs overlap |
| 3 | DIAG-POST_SKIPPED Skipped POST tests: assuming all ports are healthy, Err# 0004 |

*Figure 1-132 Domain ID segmentation error log*

To solve this overlap, change the domain ID of one of the switches participating in the ISL. This can be done using the WEB TOOLS GUI in the **Switch Settings** tab or using the **`configure`** telnet command as shown in 1.6.2, “Connecting to the switch” on page 35.

Domain ID overlap can be easily avoided by disabling the switches first using the `switchDisable` command. When bringing back the switches online automatically, the domain ID is negotiated set to a valid value.

## 1.8.2 Zoning configuration conflicts

When merging two fabrics, zoning information from the two previously separate fabrics is merged as much as possible into the new fabric.

Sometimes, zoning inconsistency can occur and zoning information cannot be merged.

An example of segmentation due to zoning is shown in Figure 1-133.

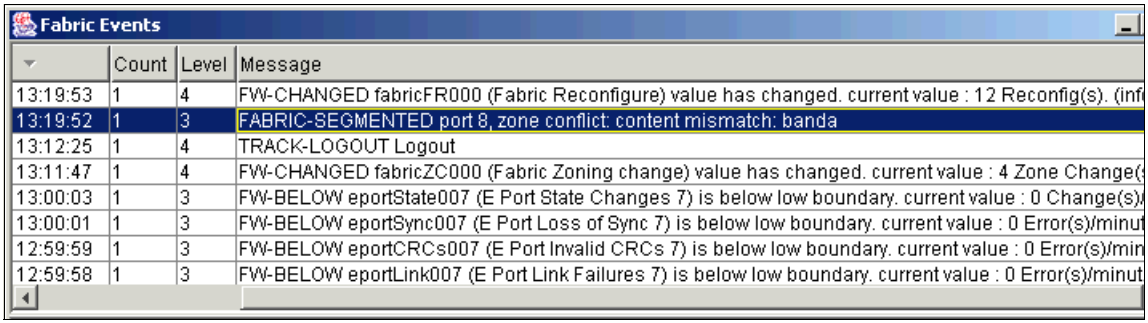


*Figure 1-133 Zone conflict error log*

In the example above, we have a different active configuration enabled on each of the two fabrics, and each of the configurations we have an alias defined for *banda*, each alias definition pointing to a different switch/port.

One of the solutions is to make sure, before attempting the merge, that zoning information on both fabrics does not have any duplicate name definitions.

The other solution is to make sure that the switch we are adding to the fabric is cleared of any zoning information. This can be done by following this process:

1. Disable the active configuration using **`cfgdisable`**.
2. Issue the **`cfgclear`** command to clear all zoning information.
3. Issue the **`cfgsave`** command to save the changes.
4. Issue **`switchenable`** to enable the switch.

Figure 1-134 shows an example command flow of this process.

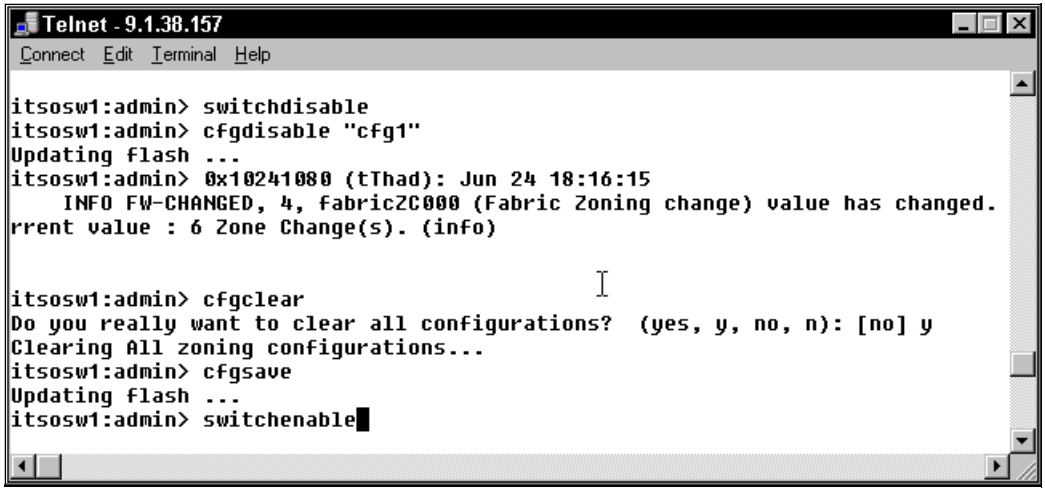


*Figure 1-134 Clearing all zoning information*

## 1.8.3 Operating parameters conflicts

Conflicts due to fabric wide operating parameters are less common since default values for these settings suit most needs. They can occur when dealing with multi vendor environment or distance solution installations, for example.

Error log messages vary a lot depending on the source of the problem. An example is shown in Figure 1-135.

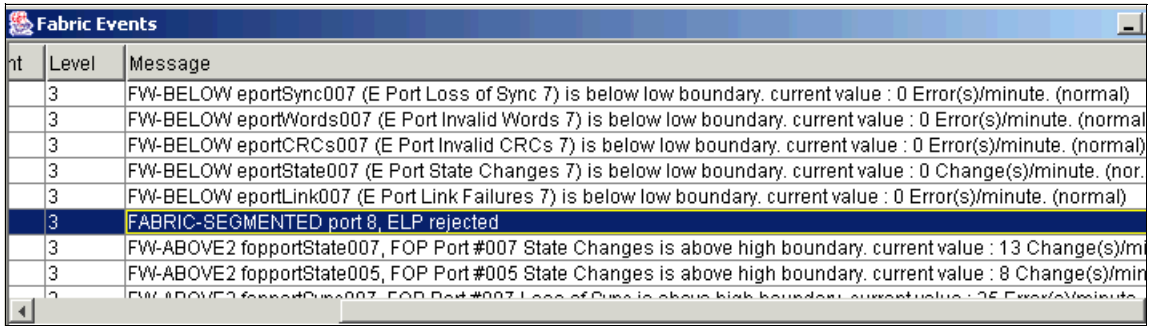


*Figure 1-135 Fabric parameter segmentation error log*

In the example above, we have core PID set on in one fabric and not in the other which caused the segmentation.

One solution to this problem is to make sure the fabric wide operating parameters are consistent across all participating switches.

If default values are used, then follow these steps to reset the settings:

1. Telnet into the switch that you are adding, for example, **`telnet 9.1.38.1.157`**, and press Enter.
2. Login, enter the switch userid and password
3. Disable the switch with **`switchdisable`**
4. Reset parameters using **`configdefault`**
5. Set IBM fabric parameters **`iodset`** and **`dlsreset`**
6. Use **`configure`** to set required domain ID and other specific parameters, ensuring all except the domain ID are identical.
7. Reboot the switch using the **`reboot`** or **`fastboot`** commands (the switch will be enabled after the boot completes).

## 1.9 Upgrading switch firmware

From time to time new versions of firmware will be released, in the following example we have documented the steps to upgrade a switch from v3.0.2q to the new v3.1.0 Secure FOS of code. This can be performed using Telnet or by using the WEB TOOLS interface. We will perform both methods.

The latest microcode levels can be downloaded from these sites:

- For the 2109-M12:

  `http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/2109m12`
- For the 2109-F32:

  `http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/2109f32`
- For the 2109-F16:

  `http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/2109f16`
- For the 3534-F08:

  `http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/3534f08`

**Note:** As new firmware levels are introduced regularly, the process we document here will apply to subsequent firmware releases. It is likely that by the time you read this redbook, later firmware will be available.

In this example we have chosen the link for the 2109-F16.

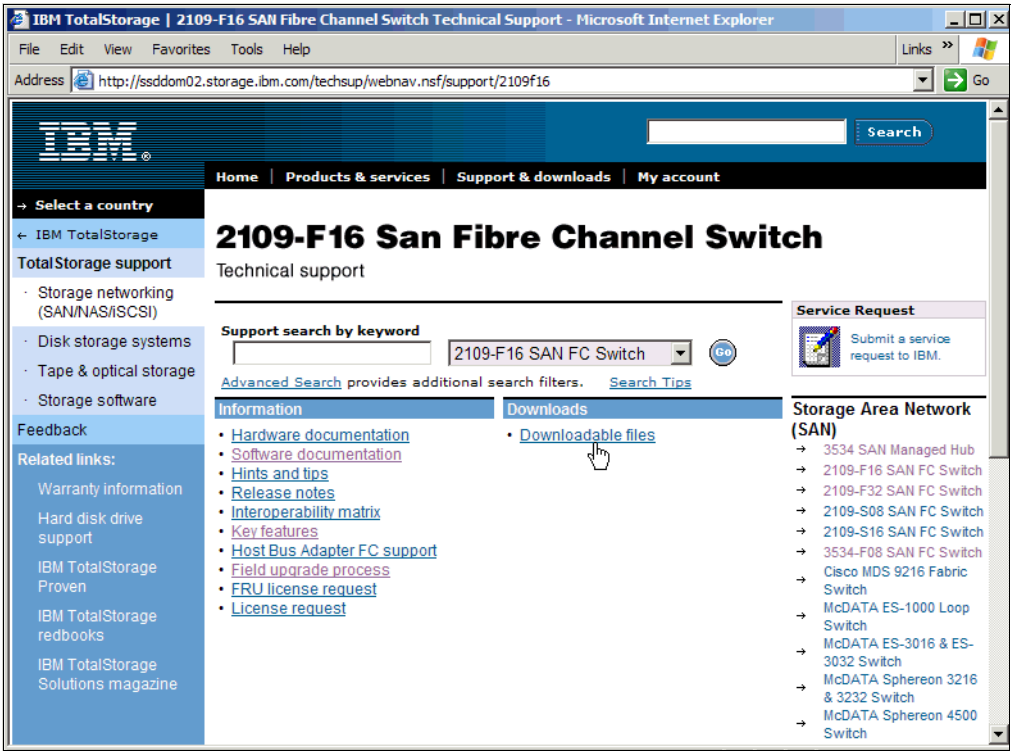


*Figure 1-136 IBM product support Web page*

By clicking the **`Downloadable files`** link on the Web page shown in Figure 1-136, we are redirected to a Brocade download site which allows us to download firmware and documentation for all of the IBM TotalStorage SAN Switch products. A pop-up window appears warning us of the redirection off the IBM hosted Web site, shown in Figure 1-137.

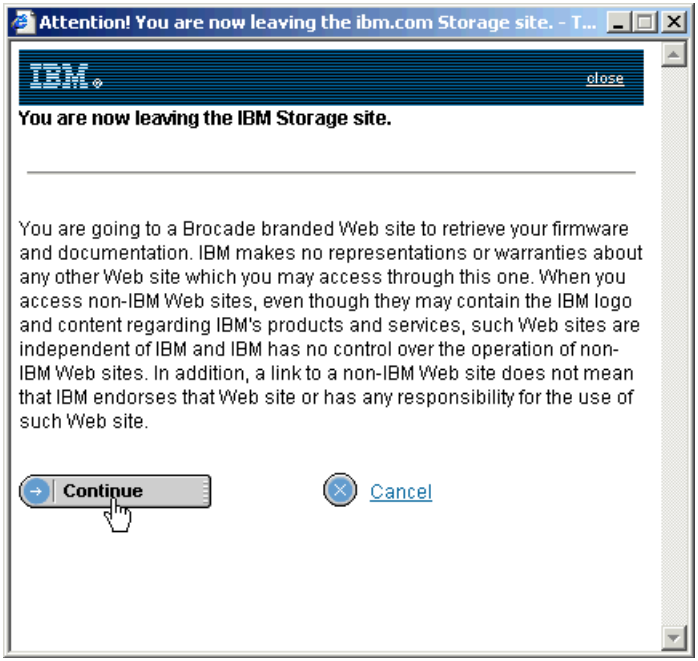


*Figure 1-137 Redirect to Brocade confirmation*

As the Brocade download site lists all versions of firmware for all the IBM TotalStorage SAN Switch products, shown in Figure 1-138, we need to understand how to determine the latest level of firmware, and which release level is appropriate for our switch. In Table 1-1 on page 4 we show which levels of firmware relates to the model of switch you have.

We determine the base level for our switch as follows:

- v2.x.x is for 3534-1RU and 2109-Sxx models
- v3.x.x is for 3534-F08 and 2109-F16
- v4.x.x is for 2109-F32 and 2109-M12

Once the base version is identified for our switch, the next digit is a firmware feature support identifier, for example v3.1.x includes support for Advanced Security, while v3.0.x does not. The last digit is a fix level identifier; we select the highest number and suffix. At the time of writing, the most recent version is v3.1.0 for our 2109-F16; or without Advanced Security support (Brocade's Secure Fabric OS), the latest level would be v3.0.2q.

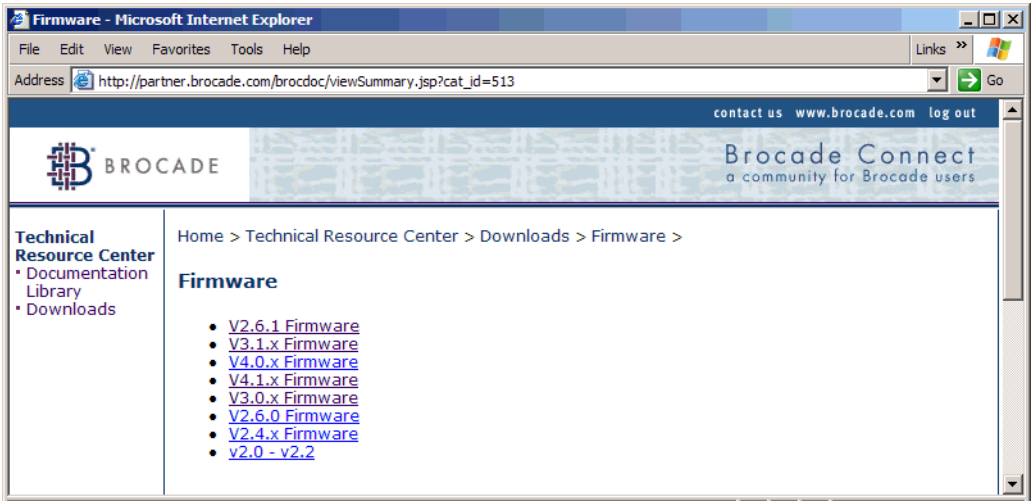


*Figure 1-138 Firmware levels download list*

From the list shown in Figure 1-138, we select the *V3.1.x Firmware* link and save the most recent level to our c:\2109\V3 directory.

The SNMP MIBs and the RSH utility (rsh.zip file) may be downloaded from the above site also by clicking the **Downloads** link on the left side of the window.

### *Using the RSH daemon (RSHD)*

In the examples that follow, we use the FTP method to perform firmware upgrades and configuration back-ups. However, RSHD may be used instead, by performing the following to run the daemon, and selecting **rshd** instead of **ftp** when prompted for protocol in the following procedures.

**Restriction:** RSHD is *not* supported with the 2109-F32 or 2109-M12.

Using the RSH daemon provided on the Brocade download site is a simple method to set up a server for transferring files to and from a switch. To do this, we need to start a DOS session to run the RSHD program. This RSH daemon validates the user and delivers the files to the switch where it is stored in flash memory. We show the running the RSH daemon in Figure 1-139.

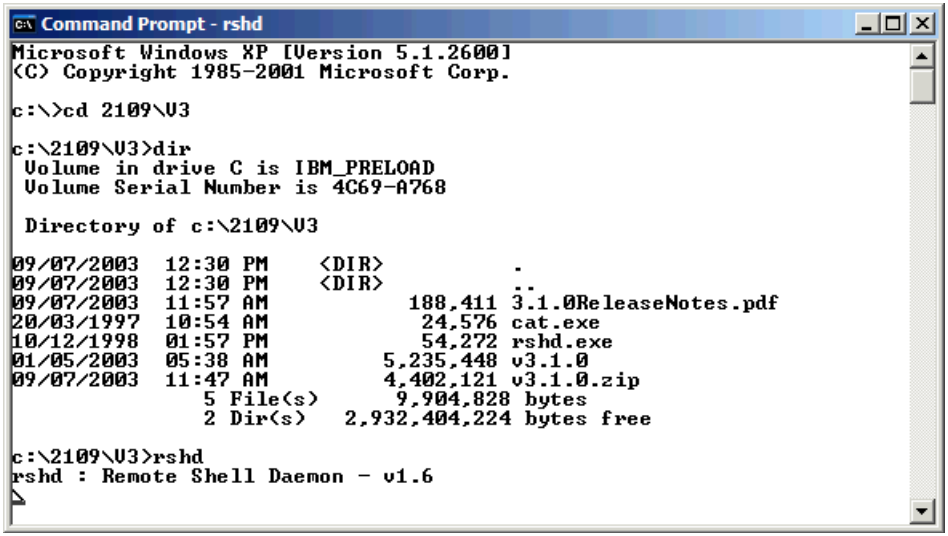


*Figure 1-139 Running the Remote Shell Daemon (RSH)*

In this example, the directory `c:\2109\V3` is used to store the switch firmware as well as the *rshd* and *cat* utilities.

**Tip:** The RSH daemon must be run from the same directory that the firmware is in.

The symbol underneath the *rshd : Remote Shell Daemon - v1.6* heading will rotate to indicate that the process is alive and well.

After a successful transfer the daemon should display a message indicating that it is processing the command to transfer the firmware to the switch. We can now close this process.

## Upgrading the Firmware using Telnet

To transfer the new level of the firmware to the switch, we need to know the IP address or hostname of the server where we stored the firmware. If these are not known already, open a DOS window on the server and issue the **`ipconfig`** command as shown in Figure 1-140.

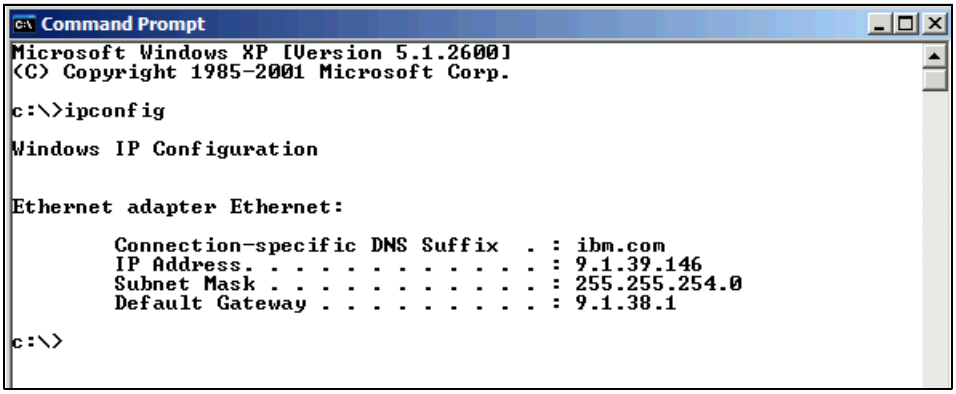


*Figure 1-140 Displaying Hostname and IP configuration details*

In our case, on the Windows server, we used the Ethernet adapter to connect to the switch. We make a note of the IP Address details, as this will be required when transmitting the firmware to the switch.

To download the firmware to the switch, we need to open a Telnet session to the switch. There are many Telnet clients available; for this example; we just type **`telnet`** followed by the IP address or hostname of the switch at a DOS command prompt.

**Tip:** While performing the firmware upgrade, we recommend that you take advantage of your scheduled fabric outage and enable the core PID setting. If this is not already set, refer to “Setting Core PID format” on page 37, and “Configure” on page 104 to enable using WEB TOOLS.

Before we upgrade to a new level of firmware we recommend performing a backup of the current configuration by entering the **`configupload`** command and following the prompts. This will also back up the security policies, if Advanced Security (Secure Fabric OS) was already in use on the switch.

When we are logged into the switch, we issue the **`firmwareDownload`** command to download the firmware as shown in Figure 1-141.

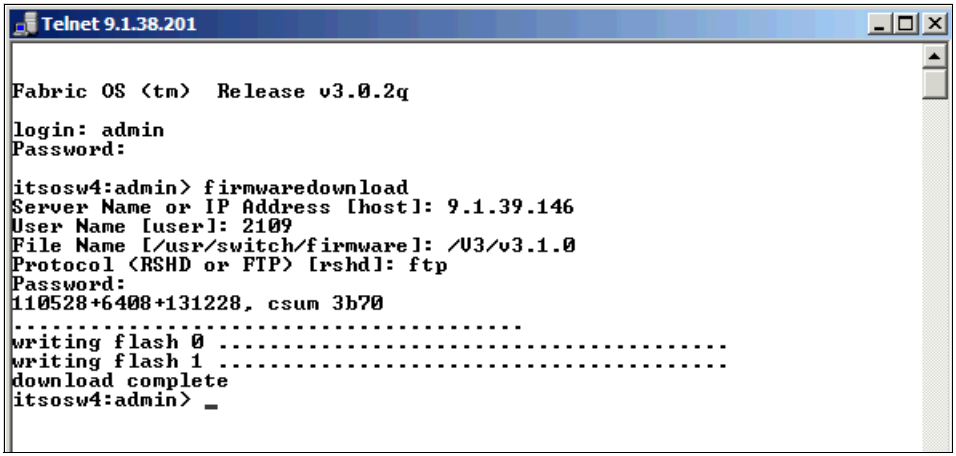


*Figure 1-141 Upgrading the switch firmware with FTP*

The switch prompts us for:

1. `Server Name or IP Address [host]:`

   The host IP address on which the FTPD program is running.

2. `User Name [user]:`

   We entered 2109 as the user defined in our FTP server.

3. `File Name [/usr/switch/firmware]:`

   We enter the full path to the firmware file.

**Tip:** We defined a userid of 2109, to our FTP server, defining it to have a base read only directory where the firmware resides c:\2109 in our case, to simplify the File Name path statement. Also the syntax always uses a forward slash (/) no matter which operating system is the server.

4. `Protocol (RSHD or FTP) [rshd]:`

   We entered FTP.

5. `Password:`

   We enter the password as defined on the FTP server for the 2109 userid.

Alternatively, you may specify all the variables with the **`firmwareDownload`** command in the following format:

**`firmwareDownload [“host”,“user”,“file”[,"passwd"]]`**

In this case, *host* is the IP address of the ftp server that the firmware is on, *user* is normally ‘user’ for **rshd** or a valid defined user name in the FTP server *trusers* file, *file* is the full path and filename of the firmware, and *passwd* is only required if you are using an FTP server to transfer the firmware to the switch.

To achieve the same as our previous FTP example above, we could enter:

```
firmwaredownload "9.1.39.146","2109","/V3/v3.1.0","2109"
```

On completion of the firmware upgrade, we need to reboot the switch. This can be done anytime, but the new firmware changes will only take effect after a reboot is successfully completed.

This can be performed simply by issuing the **reboot** or **fastboot** command. The fastboot does not perform the full POST diagnostics, therefore it can be used to re-initialize the switch with the new firmware faster.

It will take approximately 2 minutes for the switch to reboot, but this does depend on a number of factors. Some of those factors are the number of switches that are in the fabric and the amount of zoning and name server information that has to be propagated to each switch in the fabric.

It is also possible to reboot the switch from within the IBM StorWatch™ Switch Specialist as described in, “Upgrading the firmware using the WEB TOOLS” on page 192.

After the reboot is complete, we use the **version** command to confirm that we have successfully upgraded.

```
itsosw4:admin> version
Kernel:     5.4
Fabric OS:  v3.1.0
Made on:    Thu May 1 12:32:20 PDT 2003
Flash:      Thu May 1 12:33:13 PDT 2003
BootProm:   Tue Oct 30 10:24:38 PST 2001
itsosw4:admin>
```

### *2109-M12 firmware download using telnet*

To perform a firmware upgrade on a 2109-M12 using Telnet, we must have an FTP server running which has the unzipped or un-tarred firmware package within an accessible directory for the FTP login.

To upgrade the firmware, we performed the following steps:

1. Telnet into the M12 and execute the **haShow** command to verify that both CPs are available.

```
sw76:admin> haShow
Local CP (Slot 5, CP0): Active
Remote CP (Slot 6, CP1): Standby
HA Enabled, Heartbeat Up, HA State synchronized
```

2. Use the **firmwareDownload** command to download the new version of the firmware.

> **Note:** The *File Name* field requires forward slashes (/) be used in the directory path, and we must specify `release.plist` as the file name.

```
swd76:admin> firmwaredownload
This command will upgrade both CPs in the switch. If you
what to upgrade a single CP only, please use -s option.

You can run firmwareDownloadStatus from a telnet session
to get the status of this command.

This command will cause the active CP to reset. This will
cause disruption to devices attached to both switch 0 and
switch 1 momentarily and will require that existing telnet
sessions be restarted.

Do you want to continue [Y]: y
Server Name or IP Address: 9.42.164.42
User Name: root
File Name: /tmp/v4.1.0/release.plist
Password:
FirmwareDownload has started on Standby CP. It may take up to 10 minutes.
FirmwareDownload has completed successfully on Standby CP.
Standby CP reboots.
Standby CP booted up.
Standby CP booted up with new firmware.
```

> **Note:** At this point the telnet session disconnects due to the CP reboot. We have logged back in to the switch with a new telnet session.

3. As we have just logged back into the switch, we issue the **firmwareDownloadStatus** command to find out the current status of the upgrade. We know that it has completed by the *Firmwaredownload command has completed successfully* statement.

```
swd76:admin> firmwareDownloadStatus
[0]: Fri Jul 11 13:44:31 2003
cp0: Firmwaredownload has started on Standby CP. It may take up to 10
minutes.

[1]: Fri Jul 11 13:50:27 2003
cp0: Firmwaredownload has completed successfully on Standby CP.

[2]: Fri Jul 11 13:50:30 2003
cp0: Standby CP reboots.

[3]: Fri Jul 11 13:53:11 2003
cp0: Standby CP booted up.

[4]: Fri Jul 11 13:53:13 2003
cp0: Standby CP booted up with new firmware.

[5]: Fri Jul 11 13:55:58 2003
cp1: Active CP forced failover succeeded. Now this CP becomes Active.

[6]: Fri Jul 11 13:56:01 2003
cp1: Firmwaredownload has started on Standby CP. It may take up to 10
minutes.

[7]: Fri Jul 11 14:02:12 2003
cp1: Firmwaredownload has completed successfully on Standby CP.

[8]: Fri Jul 11 14:02:16 2003
cp1: Standby CP reboots.

[9]: Fri Jul 11 14:05:24 2003
cp1: Standby CP booted up with new firmware.

[10]: Fri Jul 11 14:05:28 2003
cp1: Firmwarecommit has started on both Active and Standby CPs.

[11]: Fri Jul 11 14:11:04 2003
cp1: Firmwarecommit has completed successfully on Active CP.

[12]: Fri Jul 11 14:11:06 2003
cp1: Firmwaredownload command has completed successfully.

swd76:admin> exit
```

4. We now issue the **haShow** command to view the status of the Standby CP.

```
sw76:admin> haShow
Local CP (Slot 5, CP0): Active
Remote CP (Slot 6, CP1): Non-Redundant
```

```
sw76:admin> haShow
Local CP (Slot 5, CP0): Active
Remote CP (Slot 6, CP1): Standby
HA Enabled, Heartbeat Up, HA State synchronized
```

**Note:** In the first example above of the **haShow** command, the standby CP is in the process of rebooting. In the second example it is ready to become the active CP.

Both switches are now controlled by the newly upgraded CP (CP1) firmware.

5. Issue the **firmwareshow** command and confirm that all the primary and secondary partitions on both CP cards firmware levels all match.

```
sw76:admin> firmwareshow
Local CP (Slot 5, CP0): Standby
        Primary partition:      v4.1.0
        Secondary Partition:    v4.1.0
Remote CP (Slot 6, CP1): Active
        Primary partition:      v4.1.0
        Secondary Partition:    v4.1.0

Note: If Local CP and Remote CP have different versions
of firmware, please retry firmwaredownload command.
```

**Note:** For more information on any of the commands used in this procedure, refer to the *Brocade Fabric OS Reference Version 4.1.0 — 53-0000519-02*, which contains detailed information on commands, and command syntax.

This completes the command line method of firmware upgrade.

## Upgrading the firmware using the WEB TOOLS

As with upgrading the firmware using Telnet, we need to make sure that our FTP server is running, and we also know the server IP address. Refer to “Upgrading the Firmware using Telnet” on page 186 for these details.

To upgrade the firmware using the WEB TOOLS, we point our Web browser to the IP address of the SAN switch. We see a Fabric View as shown in Figure 1-142.

**Tip:** While performing the firmware upgrade, we recommend that you take advantage of your scheduled fabric outage and enable the core PID setting. If this is not already set, please refer to “Setting Core PID format” on page 37, and “Configure” on page 104 to enable using WEB TOOLS.

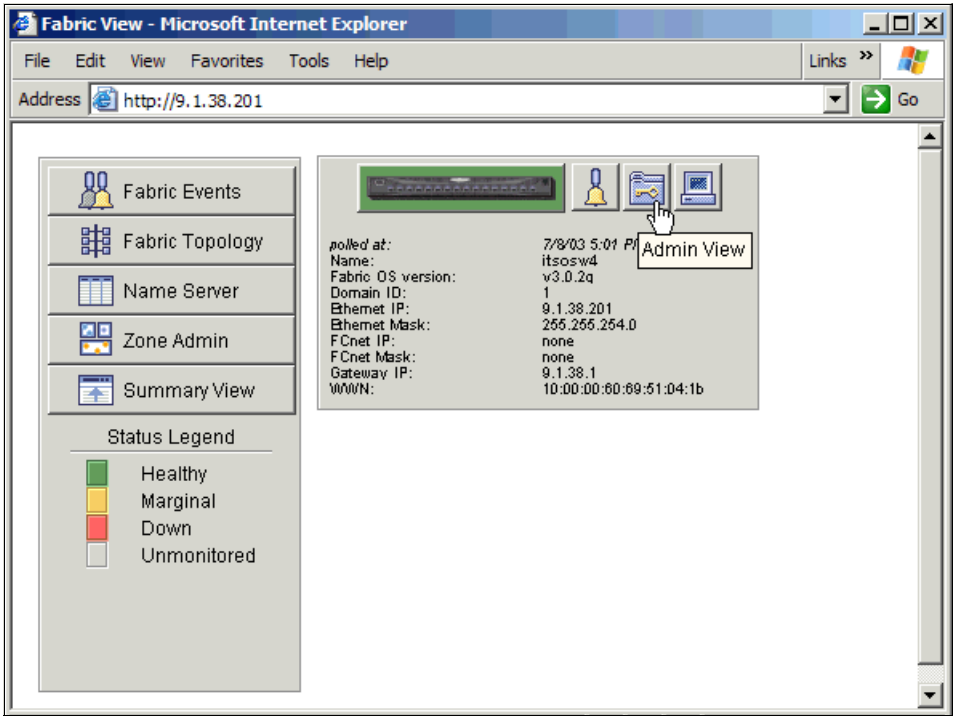


*Figure 1-142 Go to Switch Admin window*

From the Fabric View we locate the Switch that is to receive the new Firmware, and then click the **Admin View** button.

This will display the Switch Admin window, as shown in Figure 1-143.

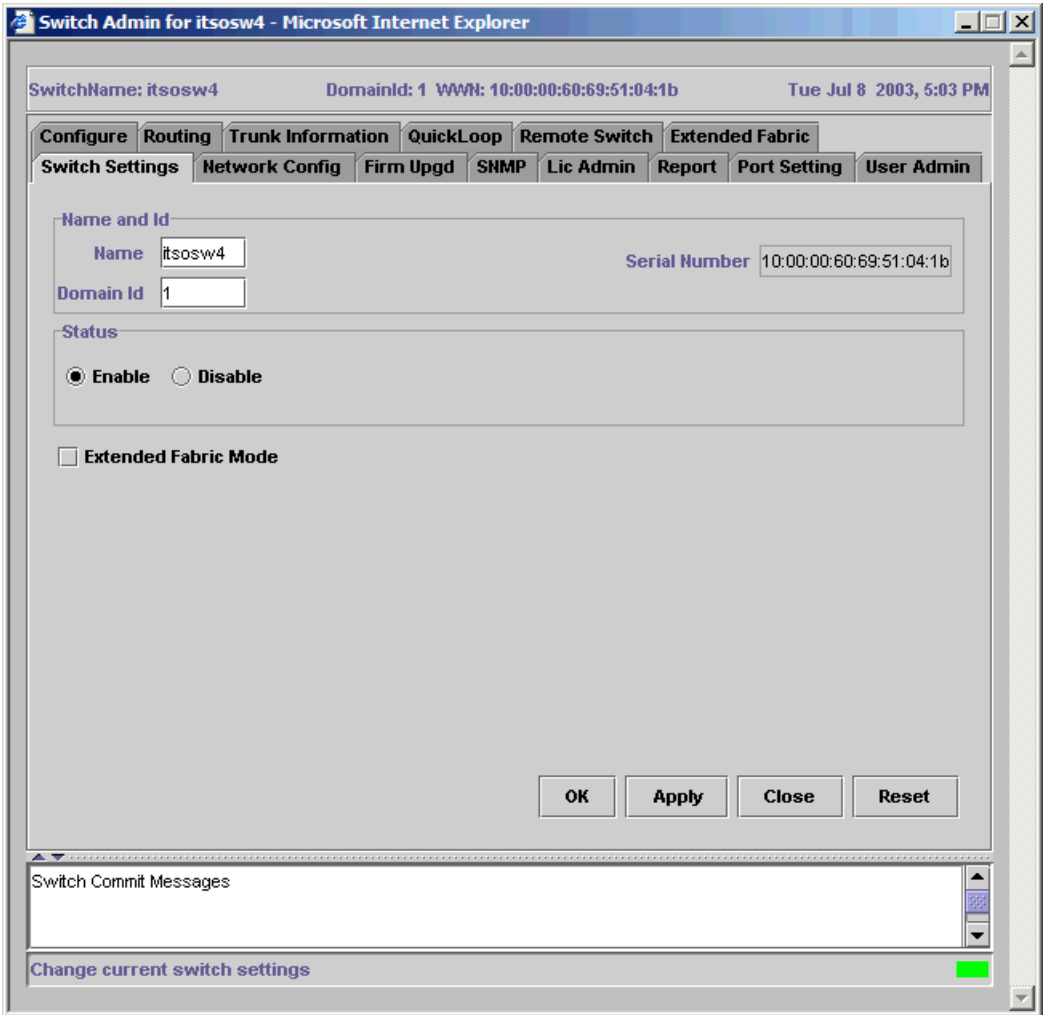


*Figure 1-143 Switch Admin window*

We then select the **Firm Upgd** tab and see the window shown in Figure 1-144.

**Attention:** After v3.1.0 firmware, the **Firm Upgd** tab is known as **Upload/Download**.

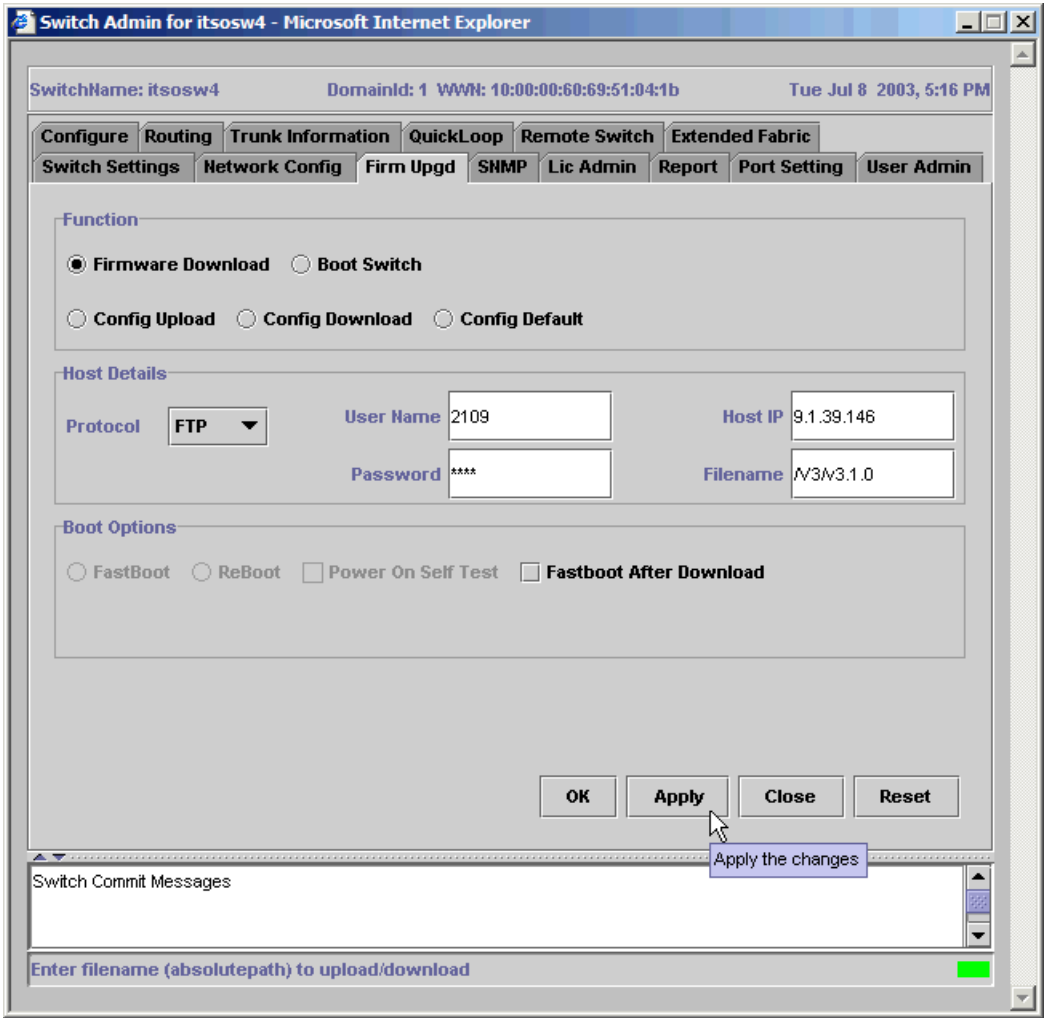


*Figure 1-144* Firmware Upgrade window using FTP

We click the **Firmware Download** radio button, select the **FTP** protocol option, then specify the FTP userid, Host IP address, FTP password, and the directory and firmware file name.

Note that the FTP server must be running before clicking **Apply** to initiate the download process.

**Tip:** By not selecting **Fastboot After Download**, we are able to verify the success of the download in the Message window.

After initiating the firmware download by clicking **Apply**, the process will take two to three minutes to complete. Upon completion, the message window provides us with a report on the actions performed on the switch as shown in Figure 1-145.

**Tip:** We expand and contract the message box by clicking the small up / down arrow heads on the top left of the message area.

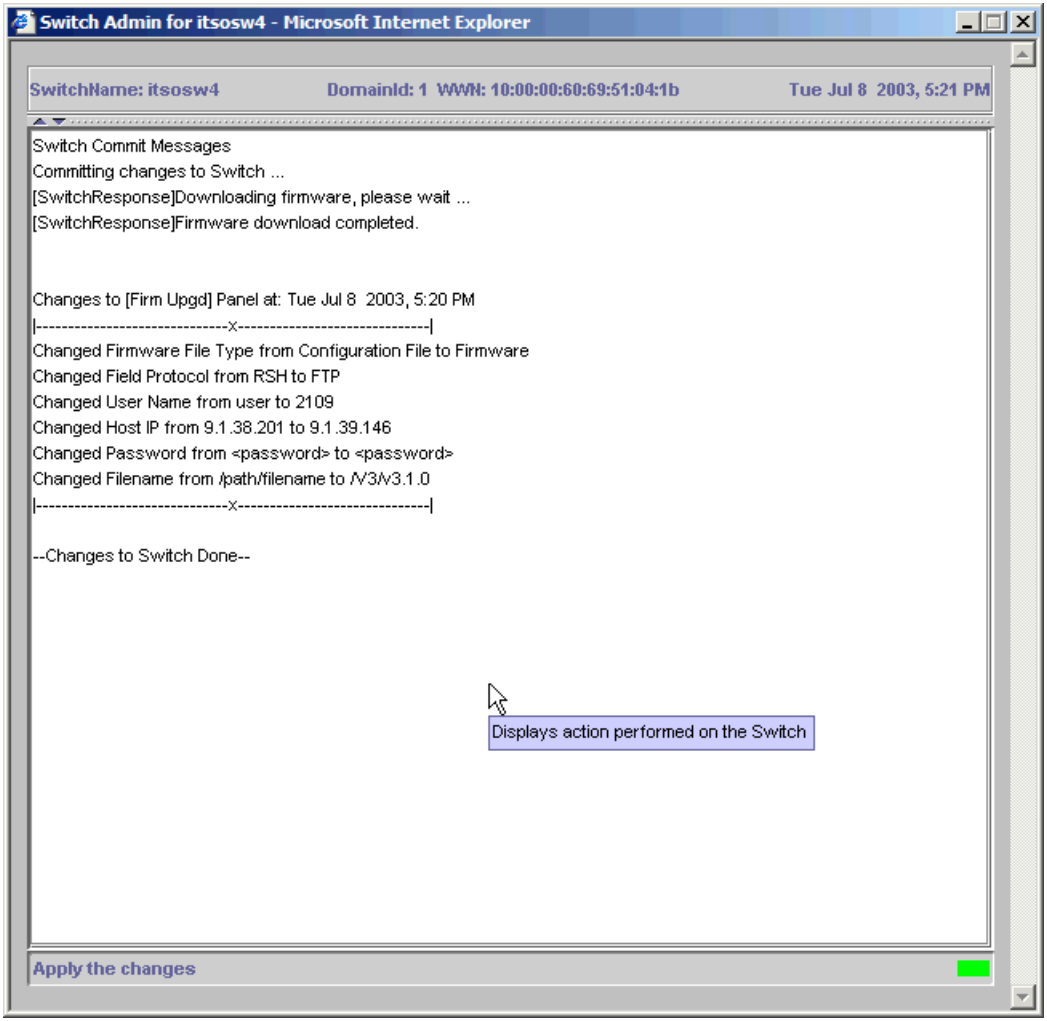


*Figure 1-145 2109-F16 firmware download report*

In order for the new firmware to take effect, we have to reboot the switch. This can be done in the same tab (**Firm Upgd** tab). We click **Boot Switch** in the **Function** area. This enables the **Boot Options** area. We then click **Fastboot** and **OK** as shown in Figure 1-146.

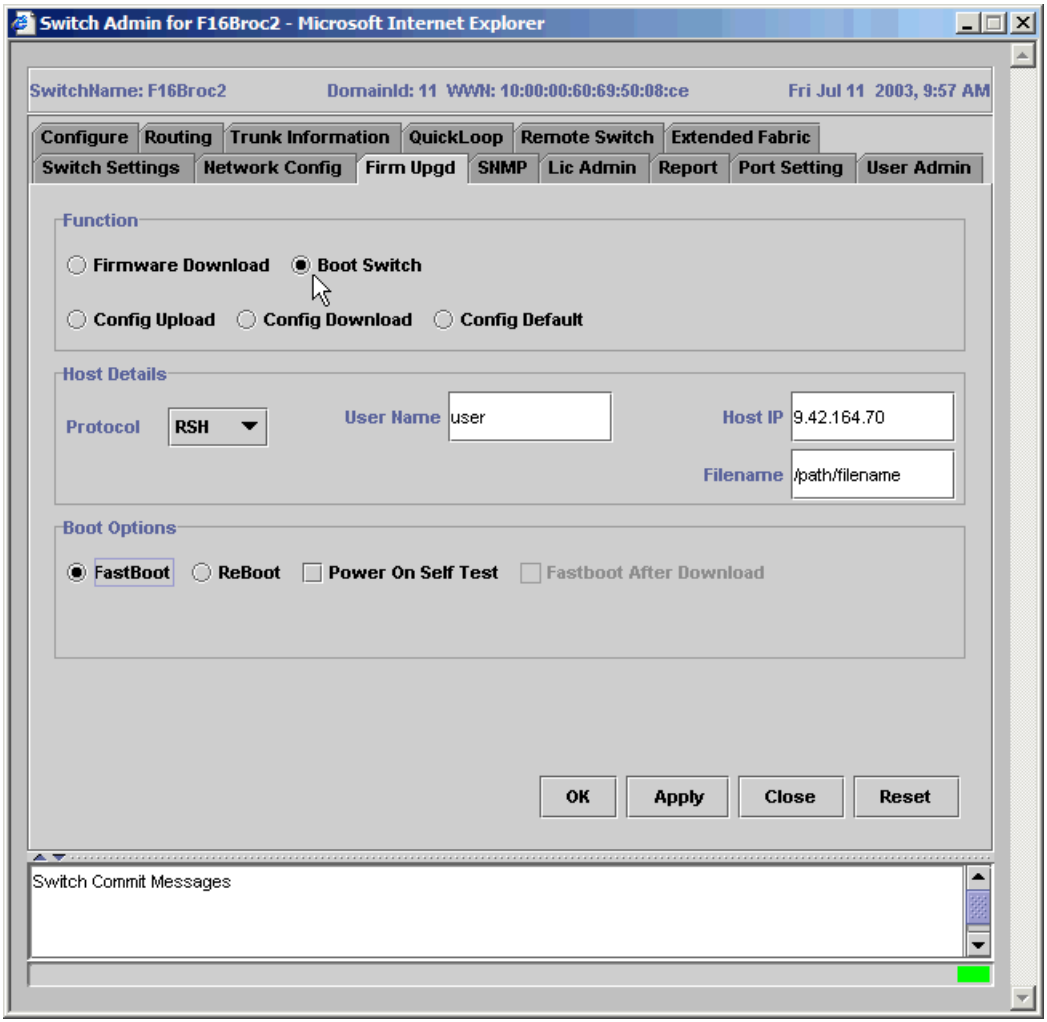


*Figure 1-146 Rebooting the 2109-F16*

Clicking **OK** brings up a window warning that the switch is about to reboot. We click **Yes**.

**Note: Fastboot** compared to **Reboot** reduces boot time significantly, as it bypasses Power On Self Test (POST). **Fastboot** should *not* be used if a hardware problem is suspected on the switch.

After performing the fastboot on our switch, we received the message displayed in Figure 1-147; this is due to a change in the firmware upgrade handling between v3.0.x and v3.1.x. This message, in this case, can be expected and does not indicate the firmware upgrade failed.


*Figure 1-147 Firmware upgrade Exception message*

Closing the Admin view and refreshing the Fabric View main window, we can verify that the firmware level has been upgraded by checking the switch details as shown in Figure 1-148.

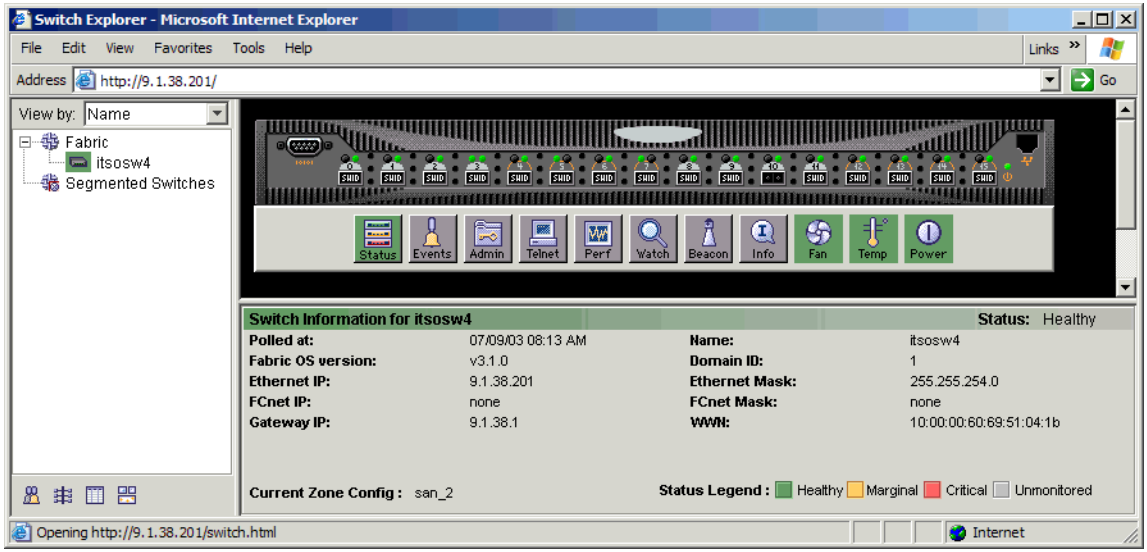


*Figure 1-148 Switch View showing new firmware level*

In the upgrade we have just performed, there is a significant difference in the WEB TOOLS interface. It is not always expected that significant visual changes like this will be the result of a firmware upgrade. We installed the upgrade to v3.1.0, which allows our 2109-F16 to support Secure FOS. We detail these enhancements in 1.12, “Advanced Security” on page 211.

We have now successfully implemented the changes to the switch.

## Version 4.x.x firmware upgrade

As version 4 of the Fabric OS is Linux based, there are some differences in the upgrade process, compared to version 2 or version 3 firmware. We show the steps by performing a firmware upgrade on a 2109-M12, taking it from v4.0.2c to v4.1.0. The same procedure may be used on the 2109-F32 also. The file transfer method for 2109-F32 and 2109-M12 is different from the method used for previous 2109 models, in that only FTP is supported.

**Attention:** In this procedure we will be upgrading the switch to SFOS v4.1. Upgrading a switch to this level will cause a 20 to 30 second disruption to fabric services, and therefore is considered disruptive.

For an M12 having dual CPs, the Upload/Download tab downloads firmware onto both CPs in an automated sequence. During this sequence a monitoring process is initiated to check the progress by requesting the status from the CP and reporting the status or error message in the Status and Error Window at the bottom of the view. All other switch administration tabs and buttons are disabled during the firmware download monitoring process.

First we need to download the latest level of firmware from the Web by selecting the *downloads* link from the following URL:

Then unzip the downloaded file to a directory on your workstation. The unzip utility will extract the code to a new sub-directory for the firmware.

For example: `C:\2109\V4\v4.1.0`, where v4.1.0 was created by the unzip.

Make sure that an FTP server is running in your workstation and that read access for the directory above is given to the userid you will use.

**Tip:** We defined a special `2109` userid to our FTP server, which we gave read access to the `c:\2109` directory in the *trusers* file. Having a specific ID defined makes the directory path simpler to specify.

As shown in Figure 1-149, we fill in the **User Name**, **Host IP, Password,** and **Filename** according to what we have defined in our FTP server. We also must select the **Firmware Download** radio button. The **ftp** option in the ***protocol*** pulldown is the only option available.

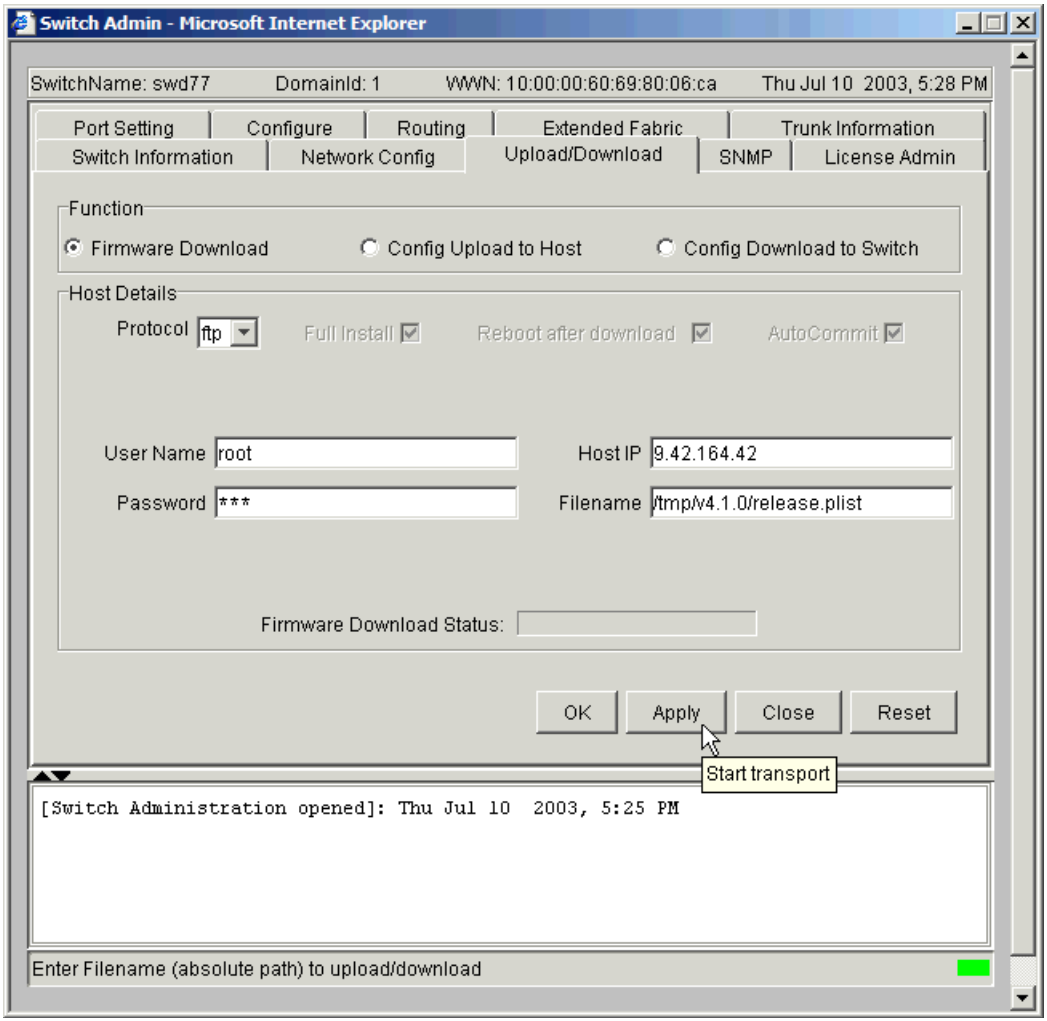


*Figure 1-149 Firmware download setup*

We must specify *release.plist* for the filename. If we do not specify this, or if our path to the file is incorrect, we will receive the failure message shown in Figure 1-150.

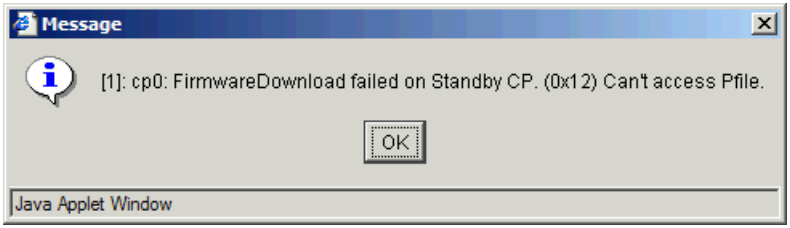


*Figure 1-150 Failure message*

6. Clicking the **Apply** button presents us with a confirmation button and warning message, as shown in Figure 1-151. This message explains that the procedure will take some time, and will cause the loss of ethernet communication briefly as the CPs failover.

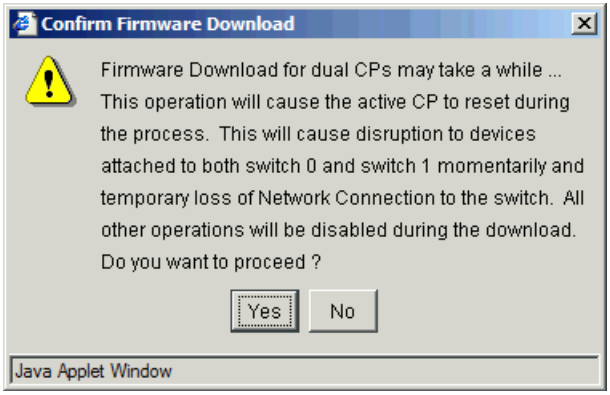


*Figure 1-151 firmware download confirm*

7. After clicking the **Yes** button, the automated firmware download begins. The Firmware Download Status bar progresses to lets us know that something is still happening; this blue progress bar will start over a few times throughout the procedure.

   The information area at the bottom of the view will keeps us aware of what step is in progress, or what error has occurred. The progress can be seen in Figure 1-152.

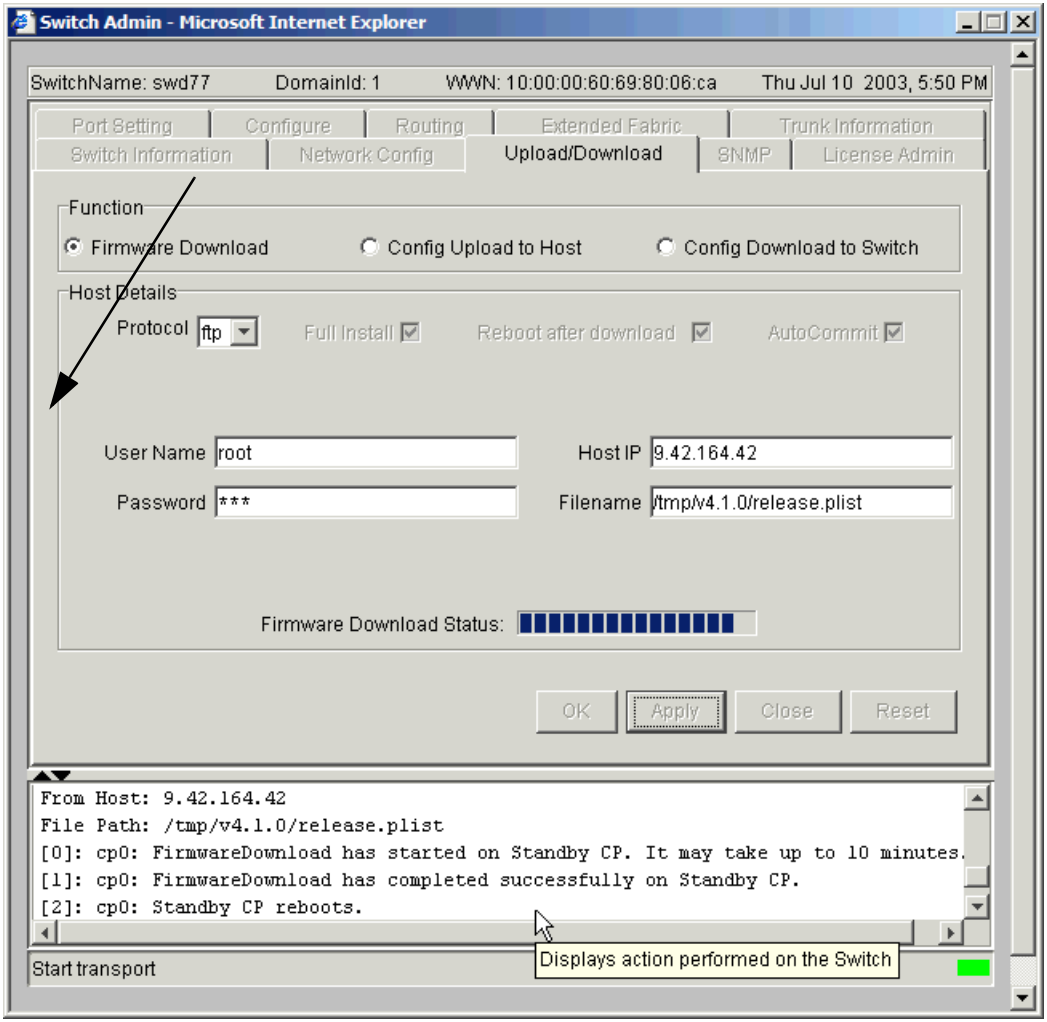


*Figure 1-152 Firmware download progress*

The procedure will perform the following steps:

1. Download firmware to the standby CP and reboot it.
2. Perform a failover of the CPs, swapping the active and standby CPs.
3. Download firmware to the new standby CP and reboot it.
4. Perform a firmware commit to both CPs.

This can be seen in the message window at completion of the end of the Firmware Download procedure, as we show by expanding the message box in Figure 1-153.

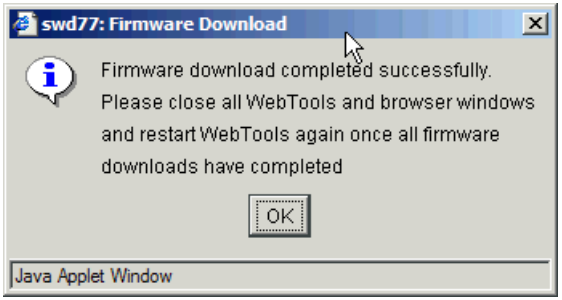


*Figure 1-153 Firmware download complete*

# 1.10 Distributed fabrics

There are three features available on the IBM TotalStorage SAN Switch that allow for remote distribution of the fabric:

- ISL R_RDY mode
- Remote switch
- Extended fabrics

We discuss these features in the topics that follow.

## 1.10.1 ISL R_RDY Mode

ISL R_RDY Mode is a new standard feature with v3.1 and v4.1 of FOS. It is designed as a replacement of the Remote Switch feature as it is more flexible and is supported by many gateway manufacturers.

When first establishing a connection to another switch or Node, switch ports initialize using Exchange Link Parameters (ELP) Mode 1. Gateways however, expect an initialization that uses ELP mode 2. Setting a port ISL R_RDY mode prepares the port for Gateway connections by causing the port initialization to use the expected method (ELP mode 2). Therefore, the WAN gateway does not need to support a special mode for these switches.

To enable R_RDY on port 9, we use the **portcfgislmode** command:

```
itsosw4:admin> portcfgislmode 9, 1
Committing configuration...done.
ISL R_RDY Mode is enabled for port 9. Please make sure the PID
formats are consistent across the entire fabric.
itsosw4:admin>
```

After ensuring that the above steps have been performed on the other remote switch, and all parameters, including core PID, match — our remote switch link is now operational.

## 1.10.2 Remote Switch

The Remote Switch feature is an optionally licensed feature on the 3534 and 2109 series switches with Fabric OS version 3.0, 4.0 or higher.

Remote Switch enables us to connect two remote IBM TotalStorage SAN Switch fabrics over an IP network, enabling communication of IP or ATM protocols as well as the normal Fibre Channel traffic.

The Remote Switch feature functions with the aid of a “bridging device” or network bridge. The network-bridge must support both a Fibre Channel physical interface and a secondary non-Fibre Channel physical interface such as IP or ATM. With Remote Switch on both sides of a fabric, the network-bridge accepts Fibre Channel frames from one side of a fabric, tunnels them across the network, and then passes them to the other side of the fabric.

The two switches are cascaded together to form a fabric that, from the viewpoint of the connected hosts and storage devices, interact the same as locally connected switches. The performance limitations depend only on the type of connection that is used.

The Remote Switch feature supports a maximum of two switches in a fabric, and provides these benefits:

- **Coordinated fabric services:** The Remote Switch fabric configuration fully supports all fabric services, including Distributed Name Services, Registered State Change Notifications, and Alias Services.
- **Distributed management:** Access to the management facilities (WEB TOOLS, Telnet, SNMP, and SES) is available from either the local or the remote switch. Interconnect for switch management is routed through the Fibre Channel connection; no additional network connection is required between sites.
- **Ability to support multiple interswitch links (ISLs):** Sites requiring redundant configurations can connect multiple E_Ports to remote sites by using multiple gateways. Standard Fabric OS routing facilities automatically maximize throughput by using the E_Ports to load share traffic during normal operation, with automatic failover and failback during interruption on the Wide Area Network (WAN) connection.

## 1.10.3 Using Remote Switch

To transfer frames across a WAN using ATM protocol, the Fibre Channel frames (from 256 to 2112 bytes) must be broken into smaller pieces (53 byte ATM cells) at the local end of the ATM network. After the frames are broken into smaller pieces, they are tunnelled inside ATM cells to be transmitted across the ATM network. At the remote end of the ATM network, these pieces are reassembled back into complete Fibre Channel frames and are transmitted through the remote Fibre Channel interface.

To accomplish this, the gateway provides an E_Port interface that links to the IBM TotalStorage SAN Switch E_Port. After the link between the two E_Ports is negotiated, the gateway E_Port moves to pass-through mode and passes Fibre Channel traffic from the IBM TotalStorage SAN Switch E_Port to the ATM network.

## 1.10.4 Configuring a Remote Switch fabric

A Remote Switch fabric requires two 3534 or 2109 series switches with identical configurations. A separate Extended Fabric license is not required to operate the switch at distances greater than 100 km. This is achieved when the switch operates over the gateway network. Performance is limited to the link used.

In addition to normal switch configuration options, the following parameters must be configured for the remote switch environment:

- **Time-out values:** The Resource Allocation Time-out Value (R_A_TOV), and Error Detect Time-out Value (E_D_TOV) must be increased, as appropriate, for all switches participating in the Remote Switch fabric. This provides for the possible increase in transit time caused by the introduction of WAN links into the fabric.
- **Data field size:** All switches participating in the Remote Switch fabric must have the data field size configured to the maximum of 2048 bytes to accommodate the maximum field size that is supported by ATM gateways. Data field sizes smaller than 2048 bytes can be set, but they might cause significant performance degradation.
- **Class F frame suppression:** All switches participating in the Remote Switch fabric must have the Class F frame suppression flag set. Class F frames are automatically converted to Class 2 frames.
- **BB credit:** The setting for BB credit must be the same on both switches. Switches with a different value will segment.

### Setting parameter values through Telnet

Using the telnet interface, we will use the **configure** command to set the following parameter values:

- BB credit
- R_A_TOV and E_D_TOV
- Data field size
- Class F frame suppression flag

Using the following commands, we changed the parameter values:

```
itsosw2:admin>switchDisable
itsosw2:admin> configure

Configure...

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

    Domain: (1..239) [4]
    BB credit: (1..27) [16]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000] 5000
    WAN_TOV: (1000..120000) [0]
    Data field size: (256..2112) [2112] 2048
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0] 1
    SYNC IO mode: (0..1) [0]
    VC Encoded Address Mode: (0..1) [0]
    Core Switch PID Format: (0..1) [1] ^D

Committing configuration...done.
itsosw4:admin>itsosw2:admin>switch:admin>switchEnable
```

## 1.10.5 Extended Fabrics

Extended Fabrics is an optionally licensed product that runs on 3534 or 2109 series switches with Fabric OS version 3.0, 4.0 or higher.

The Extended Fabrics feature creates an interconnected fabric at distances of up to 100 km using Fibre Channel technology. Extended Fabrics can increase the allowable distance between two switches.

Extended Fabrics optimizes the internal buffering algorithm for IBM TotalStorage SAN Switches. It provides maximum buffering between E_Ports that are connected over an extended distance through buffer re-configuration. This results in line speed performance of close to full Fibre Channel speed for switches that are interconnected at 100 km, thus providing the highest possible performance for transfers between switches.

The Fibre Channel connection extensions can be provided by extended distance SFPs, Fibre Channel repeaters, or wave division multiplexing (WDM) devices.

**Note:** Performance can vary depending on the condition of the fiber optic connections between the switches. Losses due to splicing, connectors, tight bends, and other degradation can affect the performance over the link and the maximum distance possible.

To enable Extended Fabrics, an Extended Fabrics license must be installed. If a fabric is created with a 2109 Model F16 switch, the long distance extended fabric configuration needs to be set only once for each fabric at the edge port connector switch. The edge port connector switch automatically works with the rest of the switches in the fabric.

**Note:** To enable Extended Fabrics in a fabric created with 3534 switches, each switch in the fabric must be configured individually.

## 1.10.6 Using Extended Fabrics

We can configure ports to support long distance links through the Telnet or WEB TOOLS interfaces.

### Supported configurations

An Extended Fabric can be created with either 3534 or 2109 series switches respectively, that are running Fabric OS v3.0 or v4.0 at a minimum. An Extended Fabric can consist of:

- 3534 switches only
- 2109 series switches only
- A combination of 3534 and 2109 series switches

**Note:** In a combination (3534 and 2109 series) configuration, the long-distance ISL that connects the fabrics must be installed between edge-port switches of same series. An Extended Fabric does not work if the long distance ISL is installed between non-matching edge port switches.

## 1.10.7 Configuring Extended Fabrics

In order to run Extended Fabrics, the following two parameters need to be set:

- Switch configuration to enable long distance
- Port configuration to select the long distance mode

In the 3534 switches, each switch within the fabric must have the switch configuration turned on. In the 2109 series switches, only the edge-port switches need to have the switch configuration turned on.

Perform the following steps to set the long distance fabric mode bit:

1. Login to the switch through Telnet.
2. At the command line, type the following command:

   **`switchDisable`**
3. At the command line, type the following command:

   **`configure`**
4. Type **`Y`** at the Fabric parameters prompt.
5. Type **`1`** at the following prompt:

   `Long Distance Fabric [0]:`

There are three possible long distance levels for a port:

- **Level 0** — Re-configures the port as a regular switch port. The number of buffers reserved for the port supports up to 10 km links.
- **Level 1** — Distances up to 50 km will support 1 Gb/s and 2 Gb/s switches (3534 and 2109 series).
- **Level 2** — Distances up to 100 km will support 1 Gb/s and 2 Gb/s switches (3534 and 2109 series).

Ports are grouped into quads, each of which consists of four adjacent ports that share a common pool of frame buffers. The possible quad groupings are:

- Ports 0 -3
- Ports 4 -7
- Ports 8 -11
- Ports 12 - 15

Certain buffers are dedicated for each port, and others are shared among the ports. In Extended Fabric mode, one port is given an increase of dedicated buffers from this pool.

The total number of frame buffers in a quad is limited, and the Extended Fabric port matrix introduces a combination of long distance ports that are available.

This is shown in Table 1-23.

*Table 1-23 Combination of long distance ports that are available*

| Port 0 | Port 1 | Port 2 | Port 3 |
|---|---|---|---|
| L1 | F or E | F or E | F or E |
| L1 | L1 | F or E | F or E |
| L1 | L1 | L1 | F or E |
| L1 | L1 | L1 | L1 |
| L2 | F | F | F |
| L2 | E | F | |
| L2 | E | | |
| L2 | L1 | F | |
| L2 | L1 | | |

Where:

- L0 represents an Extended Fabric mode of 10 km
- L1 represents an Extended Fabric mode of 50 km
- L2 represents an Extended Fabric mode of 100 km
- F represents the F_Port that is used when connected to devices
- E represents the E_Port that is used for interswitch connectivity

### Setting the port configuration

We can configure a port to support long distance links by using the Telnet command **`portCfgLongDistance`** or by using the WEB TOOLS.

## 1.11 Migrating the M12 into a core fabric

In this section we describe the requirements for migrating the M12 into a core to edge fabric.

### 1.11.1 Prerequisites

The fabric we are going to migrate the M12 into must meet the following requirements before any physical fabric connections are made to the M12:

- All IBM 2109 S-series switches in your fabric must run Fabric OS version 2.6 or higher.
- All IBM 3534-F08 and 2109-F16 series switches in your fabric must run Fabric OS version 3.0.2c or higher.
- All edge switches must have the Core PID format set; refer to 1.6.3, “Setting Core PID format” on page 37.
- The four IP addresses that you reserve for the two logical switches and the two CPs in each 2109-M12 must appear in the same subnet.
- You must remove any hosts that connect to the core and connect them to the edge. You may need to update port-based zoning for any devices attached directly to the cores, as their port addresses may change.
- All setup has been performed on the 2109-M12 according to the “M12 configuration procedure” on page 29.

## Final preparation of the M12

To ensure that the M12 is ready to be migrated into the fabric, we perform the following steps before connecting to the fabric.

**Attention:** Connecting the M12 to the fabric, and then performing the following steps will delete all defined zoning information in our fabric.

1. With the M12 not connected to any fabric, we issue **`cfgDisable`** to disable any active configuration.
2. Enter the **`cfgClear`** command to remove all configuration data and zoning information from the switch.
3. To ensure that trunking is enabled, we type **`switchCfgTrunkEnable 1`**
4. Disable both logical switches by using **`switchDisable`**

   Disabling the switches stops constant fabric re-configurations occurring while we connect the ISLs.

## Verify current fabric

In this step we will take a snapshot of our current fabric configuration. This will help in verification of the fabric after the M12 has been migrated in.

1. Use the **`switchShow`**, **`nsShow`**, **`fabricShow`**, and **`nsAllShow`** to record the total number of devices on each switch and the number of ISLs that connect to each switch.
2. On every host system, ensure that your multipathing is operational.

### Replace the core switches with the M12

With our current fabric devices all at the correct firmware levels and the core PID format set, the M12 installed, configured, cleared, and disabled, and our current fabric configuration saved and documented, we are now ready to perform the swap of our core switches.

With correct redundancy built into the fabric, this can be accomplished by performing one core switch at a time, or in a fully redundant Fabric configuration, the entire core of a fabric can be performed, with minimal impact to system availability.

**Attention:** As switches are disabled or enabled, fabric re-configuration will occur. This may cause pauses of I/O activity for short periods of time. The length of the pause depends on the host configuration.

1. Perform a **`switchDisable`** on one or all fabric core switches. Transfer the cables from the disabled core switch(es) to the M12.

   Connecting the cables in the M12 in a predefined layout to maximize redundancy is recommended.
2. On the M12, use the **`switchEnable`** command for the switch that has just been cabled.
3. Verify the Fabric according to the information gathered in “Verify current fabric” on page 210.
4. Repeat steps 1 through 4 for the other core switch, or the other Fabric.

### Cleaning the old core switches

Before re-using the old core switches, they should be cleaned. that is to say they have had all configuration and zoning information removed from them. We perform this by:

1. Issuing **`cfgDisable`** to disable any active configuration.
2. Entering **`cfgClear`** command to remove all configuration data and zoning information from the switch

We now have completed the upgrade to a M12 core switch in an existing fabric.

## 1.12 Advanced Security

To implement a secure fabric on an IBM TotalStorage SAN Switch, we require two things: an optional Advanced Security (AS) license key, and a firmware version supporting Secure Fabric OS (SFOS). When installed and configured, it

provides a comprehensive SAN security solution for IBM 2109 and 3534 switches and the devices that are attached to them. All IBM 2109 and 3534 switch models are supported, and may be used in a mixed environment.

**Note:** IBM has OEM'd Brocade's Secure Fabric OS, and the IBM name for this product is *Advanced Security.* At some stages throughout this topic, we will interchange the nomenclature.

### *Features*

Advanced Security provides the ability to:

- Secure the SAN infrastructure from unauthorized management and device access.
- Share resources within the same fabric by tightly controlling where devices (servers / hosts) can attach.
- Provide a secure means for distributing fabric wide security and zoning information (trusted switch).
- Create a "trusted SAN infrastructure".

### *Control*

The security level for the fabric is defined by a Fabric Management Policy Set (FMPS) that consists of:

- Fabric Configuration Server (FCS) policy
- Management Access Control (MAC) policies
- Device Connection Control (DCC) policies
- Switch Connection Control (SCC) policy
- Options policy (prevents Node WWN usage)

### *Management*

To manage an Advanced Security environment, we can use Telnet, Fabric Manager, or API integration into SAN Management software, such as Tivoli® SAN Manager.

## Planning

Before we leap ahead and enable security on our fabric, we need to do some planning to minimize any disruption to our SAN services:

- Document the switch name, WWN, and IP address of every switch in the fabric(s).
- Identify which switches will be the Fabric Configuration Server (FCS), and also identify at least one to be the backup FCS.
- Determine the policy requirements for each device and host.

- Identify management workstations to install secure Telnet or SSH client on.
- All switches must have minimum firmware levels to support SFOS as listed in Table 1-1 on page 4.
- All switches in the fabric must have a zoning and security license.
- Digital certificates must be installed on each switch in the fabric before enabling security.

**Note:** Only switches *upgraded* to v2.6.1, v3.1 and v4.1 firmware will require digital certificates to be added. All new switches shipped with these levels of firmware pre-installed will already have the digital certificates loaded.

## 1.12.1 Implementing Advanced Security

We will now perform the steps to implement security on our fabric, assuming that we have completed upgrading firmware to the required levels by following the procedure in 1.9, “Upgrading switch firmware” on page 182. We also assume that the security license key has been purchased and installed on all switches in the fabric.

The first step we perform is to back up the configuration of all the switches in our fabric. This is an important step that allows us to be able to restore the switch to its current condition if anything should go wrong during our implementation process. To do this, we follow the procedures outlined in “Upload / download” on page 86 for each switch, ensuring that we select the *Config Upload* option. This may also be accomplished using the **configUpload** command in a telnet session.

Our next step is to determine if digital certificates are installed on our switches in the fabric. We perform this on all switches by using the **configshow “pki”** command on switches at v2.6.1+ or v3.1+ as follows:

```
SF16SW1:admin> configshow "pki"
pki.CSR:        Exist
pki.Certificate:        Empty
pki.Passphrase: Exist
pki.Private_Key:        Exist
pki.Root_CA_Cert:       Exist
SF16SW1:admin>
```

For switches using firmware v4.1+ we use the **pkishow** command as follows:

```
SM12SW1:admin> pkishow
Passphrase      : Exist
Private Key     : Exist
CSR             : Exist
Certificate     : Empty
```

```
Root Certificate: Exist
SM12SW1:admin>
```

We can see in both cases that the Certificate shows as *Empty,* therefore we need to install the certificates. We will perform this for an F16 although the procedure is the same on all switch models.

We visit the IBM TotalStorage SAN Switch Web site at:

`http://www.storage.ibm.com/ibmsan/products/2109/san_switch_solu.html`

From this Web site, we select the model of the switch we are working with; in our case we select the SAN switch F16. From the displayed Web page, we now select the **Feature Keys** tab, which displays the Web page in Figure 1-154.

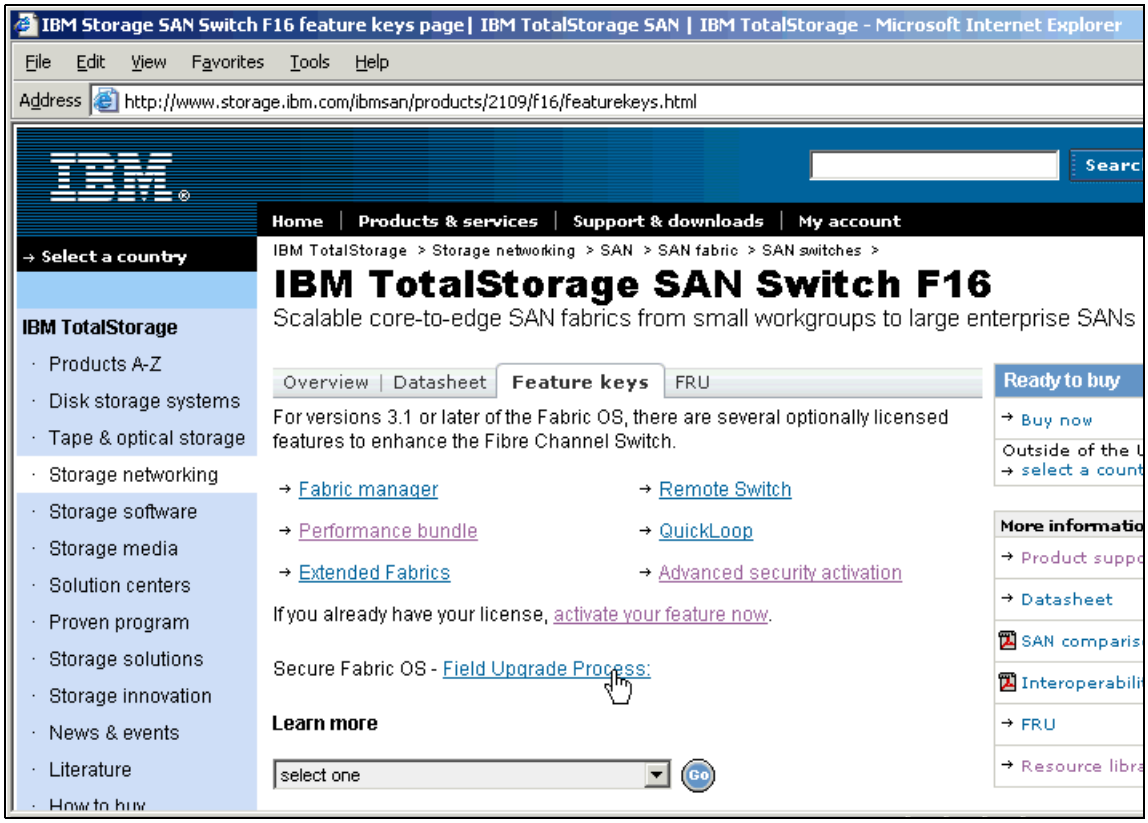


*Figure 1-154 Feature Keys Web page*

From the Feature Keys tab we select the **Secure Fabric OS Field upgrade Process:** link, which takes us to the Web page shown in Figure 1-155.


*Figure 1-155 Field Upgrade Process Web page*

From the Field Upgrade Process site, we are given links to download the Secure Fabric OS users guides, and instructions on how to implement the Advanced Security. We will be following the steps as outlined on this Web site.

We have already completed steps 1, 2, and 3 according to the instructions, and therefore click step 4, *Obtain PKICert,* this downloads a file called pki_v1.0.5.zip.

We extract the zip file to a temporary directory, where we can then run the Setup.exe to install the utility on our workstation. During the install process, we select all the default options. When the install completes, we run **c:\nt_pki\pkicert.exe**. After this opens, we press Enter to accept the default log file, and are then presented with the menu shown Figure 1-156.

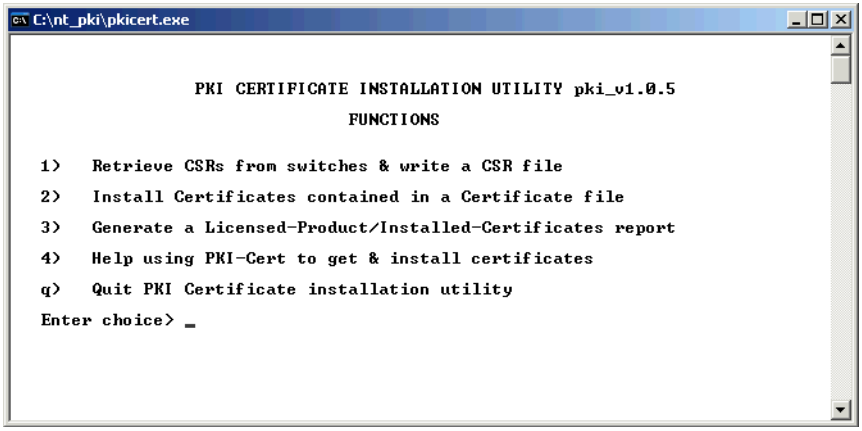


*Figure 1-156 PKI Cert Utility menu*

## Obtain CSRs

From the menu we take option **1**, to retrieve CSRs from switches and write a CSR file. This takes us to another menu where we are given the following options:

```
1)   Manually enter fabric address
2)   Read addresses from a file (name to be given)
r)   Return to Main menu
```

We take option **1** to allow us to manually enter our fabric(s) address. From the next window we only need to enter an IP address of one switch within a fabric, we can enter multiple fabrics if we wish, and by just hitting enter without entering an address on a line continue to the next window.

At this point the PKI Cert utility connects to the fabric, and prompts us for the userid and password (we are given 5 attempts). The next window prompts us for a file name as shown in Figure 1-157, where we enter a fully qualified file name and path where we would like to store the CSR information from the fabric switches.

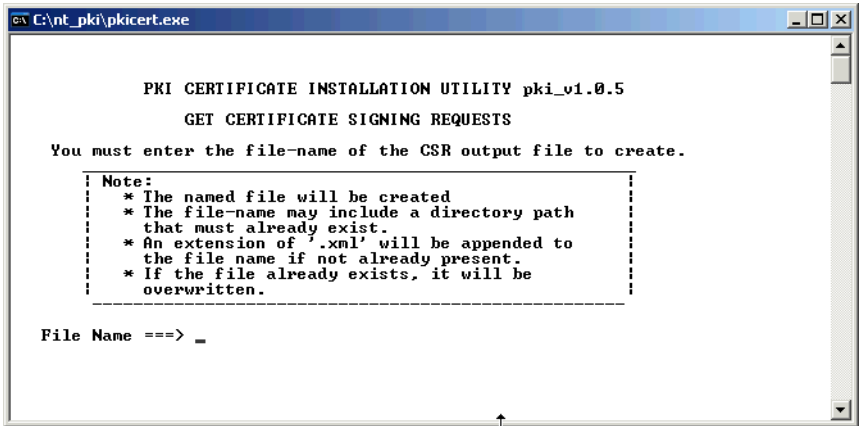


*Figure 1-157 PKI CSR file name*

After entering the file name we are asked if we would like to Include (optional) licensed product data; we replied **yes** to save the optional data. We are then asked if we want to get CSRs from switches that already have certificates. As our aim here is to install certificates on switches without them currently, we answer **No** to this question.

Next we are asked which fabric we wish to retrieve from; we selected all. Now the utility retrieves the CSRs from each switch, giving us its progress as shown in Figure 1-158.

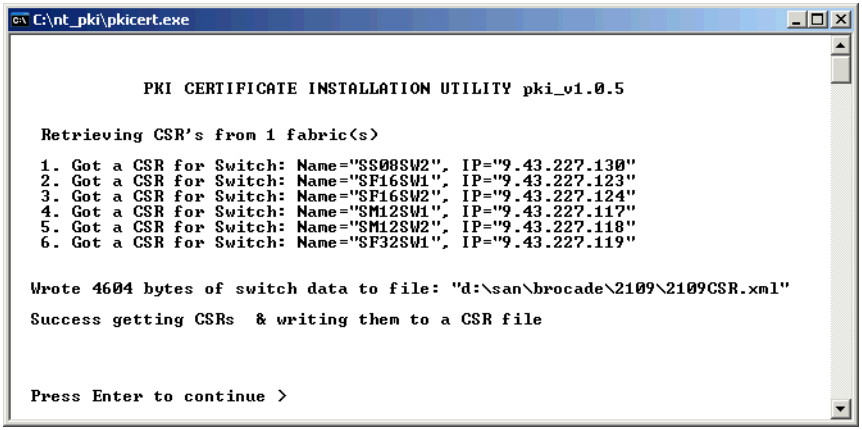


*Figure 1-158 PKI Certificate retrieval status*

Once this completes, we press Enter to continue. This returns us to the first menu, where we select **q** to quit.

## Request Certificates

Now that we have saved the CSR file on our workstation, we return to step 6 on the Field Upgrade process Web page, as shown in Figure 1-155 on page 215.

We click the Request Certificates link at step 6, and are taken to the Brocade switch key activation site. After agreeing to the licensing, and filling out our details, we point the browser to the CSR file we saved from the switches in the previous steps, and click the **Submit** button. We verify our information and click **Submit** again.

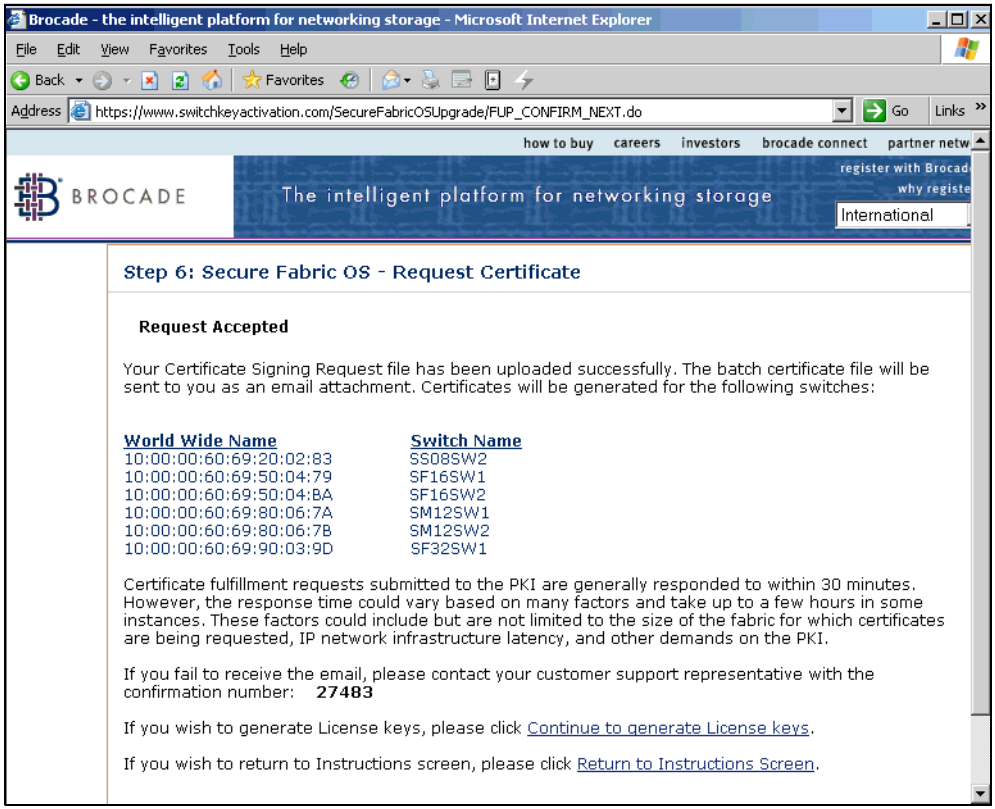


*Figure 1-159 Brocade request Certificate confirmation*

After we have submitted our collected file, an automated machine will process it, shortly after we have received the digital certificates at the e-mail address we provided in the submit form. We detach the certificates file to a temporary directory, and execute the **`c:\nt_pki\pkicert.exe`** utility again.

**Note:** If the CSR collected includes a switch without a Security license, the submitted CSR file will not be processed.

## Install the certificates

This time, from the PKICert utility menu shown in Figure 1-156 on page 216, we select option **2** to Install Certificates contained in the Certificate file we received. We then select option **1** to Manually enter the fabric IP address. We show the IP address entry in Figure 1-160, where pressing Enter on the second line (instead of supplying another IP address) advances us to the next window.

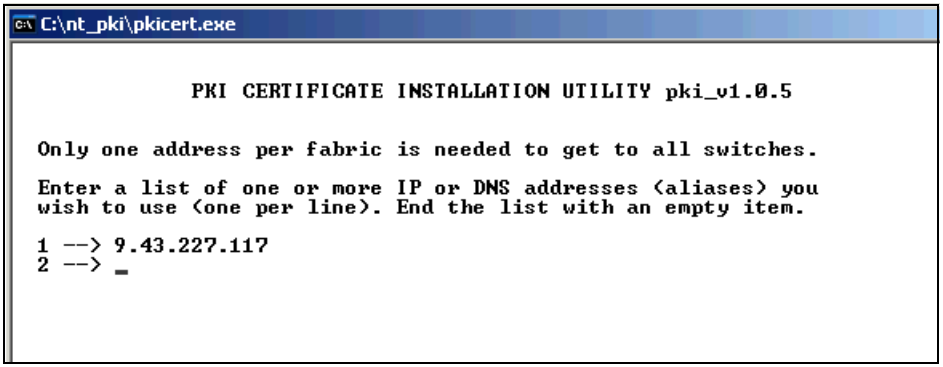


*Figure 1-160 IP address input*

At this point we are asked to provide the login *user* and *password* for PKICert to connect to the fabric. Once PKICert successfully connects to the fabric, we are prompted for the full path and file name of the Certificate file we received in the e-mail earlier.

Next we select the target fabric as shown in Figure 1-161.

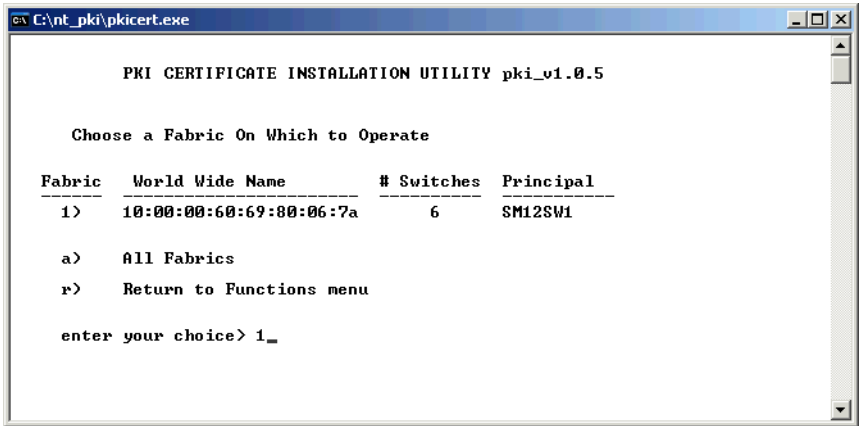


*Figure 1-161 Target fabric selection*

If we had entered multiple fabric IP addresses earlier we could now select an individual fabric or all the fabrics listed. In our case, we have only entered a single fabric.

The utility now installs the certificates on each switch in the fabric, confirming the success or failure as displayed in Figure 1-162.

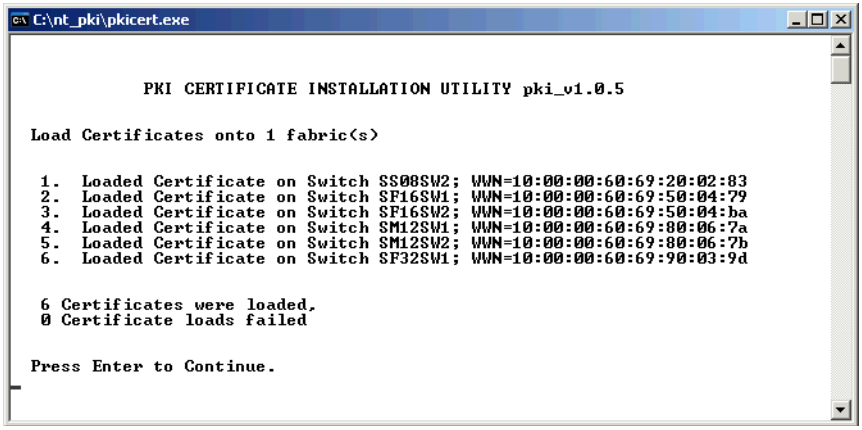


*Figure 1-162 Certificate installation success*

After pressing Enter to continue, we select **q** to quit the PKICert Utility.

We now confirm that we have successfully installed the digital certificates by issuing the **configshow "pki"** command for v2.6.1 and v3.1:

```
SF16SW1:admin> configshow "pki"
pki.CSR:        Exist
pki.Certificate:        Exist
pki.Passphrase: Exist
pki.Private_Key:        Exist
pki.Root_CA_Cert:       Exist
SF16SW1:admin>
```

or **pkishow** command for v4.1

```
SM12SW1:admin> pkishow
Passphrase      : Exist
Private Key     : Exist
CSR             : Exist
Certificate     : Exist
Root Certificate: Exist
SM12SW1:admin>
```

## How to telnet to a switch securely

Now that we have successfully installed the digital certificates on all our switches, we need to prepare our workstation to be able to securely communicate with the FCS switches in the fabric once we enable security, as normal telnet will not be allowed to connect.

From step 8 in the Web page shown in Figure 1-155 on page 215, we click the *Obtain Secure Telnet Client* link, and to download the client, we are taken to another Web page where we may select a Windows or Solaris client. We selected the Windows download link and saved *ntsectelnet.zip* to our workstation.

We then unzip the file, making sure we maintain the directory structure (if the directory structure is not maintained, the install will fail).

From our temporary unzip location, we then execute **setup.exe**.


*Figure 1-163 Secure Telnet Install*

Figure 1-163 shows the Install shield splash window for the Brocade Secure Telnet client installer, we use the Next button to install the client with all default values and complete the install process. This puts a Secure Telnet Icon on our desktop, we double click this icon to open the window shown in Figure 1-164.

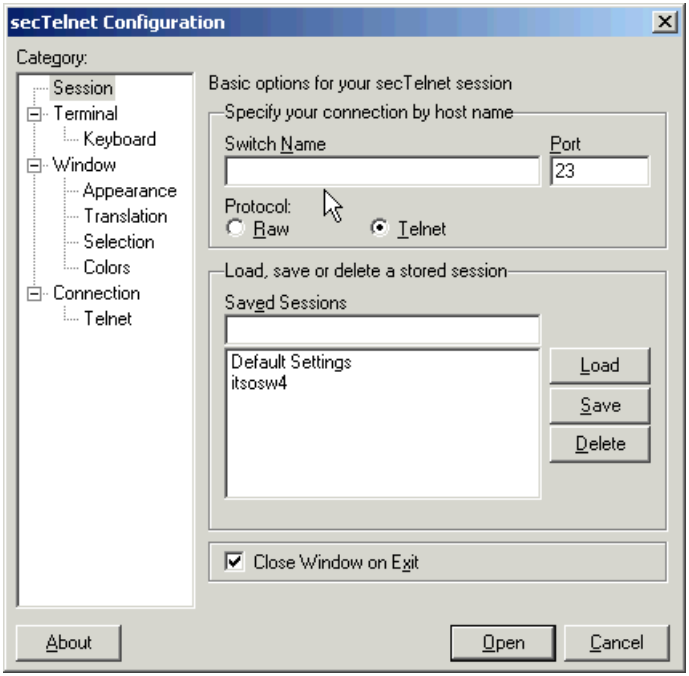


*Figure 1-164 Secure Telnet client configuration*

In this secTelnet Configuration window, we enter the IP address of the FCS switch we want to connect to in the Switch Name field, and then click the **Open** button. We also have an option of saving the connection definition, by entering a name in the Saved Sessions field and clicking the **Save** button. In our example we have saved a session for the itsosw4 switch. Now, by double-clicking the name, we launch a secure Telnet session to that switch, as shown in Figure 1-165.


*Figure 1-165 Secure Telnet session*

As the secure Telnet session uses the digital certificates that we have previously installed on the switch, establishing a connection verifies that we are ready to begin enabling Advanced Security.

**Tip:** Before enabling Advanced Security on the fabric, we recommend performing the secure Telnet session establishment to each switch in the fabric to verify that the certificates are working properly before we lock the fabric with security policies.

## 1.12.2 Enabling Advanced Security

Before continuing, we recommend performing a backup of the configuration of all the switches in our fabric again. This lets us restore the switch to this checkpoint in the procedure, if all is well currently. To do this, we follow the procedures outlined in “Upload / download” on page 86, ensuring that we select the *Config Upload* option. This may also be accomplished using the **`configUpload`** command in a telnet session. If a restore of these saved configuration is required, this may be accomplished using the **`configDownload`** command.

**Tip:** Using different configUpload save names will ensure that we have two different restore points.

We have now prepared our fabric for Advanced Security; also, during our planning step, we have identified which switches we will make the Primary and Backup FCSs. To continue, we need to schedule a fabric outage, as enabling Advanced Security is a fabric-wide setting, and will cause all switches in the fabric to reboot.

Enabling secure mode:

- Creates a default Fabric Management Policy Set (FMPS) using the FCS policy containing the WWNs that are specified in the list.
- Distributes the FMPS to all switches in the fabric
- Activates the FMPS
- Reboots all switches

The Primary FCS switch:

- Distributes the default policy sets to all switches in the fabric
- Activates the zoning configurations and any future zone management
- Applies the FMPS policy set

Using the secTelnet client we installed earlier, we now connect to the switch we have identified as being our Primary FCS. After logging in to the switch, we use the **`secModeEnable`** command as shown in Figure 1-166, where we must read and agree to the End User License Agreement.

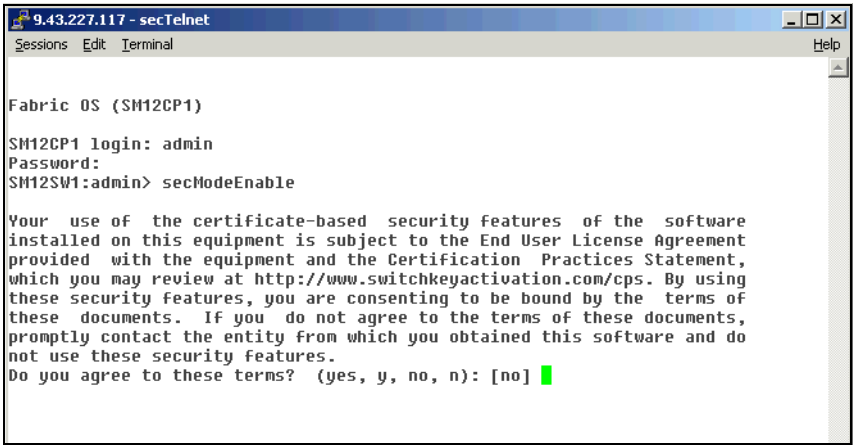


*Figure 1-166 The secModeEnable command*

We enter **y** to agree to the terms. Next we are asked to define the FCS list; at a minimum, we recommend defining two separate switches as FCS. One switch will operate as the primary Fabric Configuration Server and the other as backup, in case the primary were ever to fail. More FCS switches may be defined, although we do recommend that these switches also be located in a physically secure environment.

The following sample coding shows how we defined an M12 and the F32 in our fabric as FCS switches:

```
This command requires Switch Certificate,  Security license and Zoning
license to be installed on  every switch in the fabric.

PLEASE NOTE: On successful completion of this command, all login sessions
will be closed and all switches will go through a reboot to form a secure
fabric.

This is an interactive session to create a FCS list.

The new FCS list is empty.

Enter WWN, Domain, or switch name(Leave blank when done): SM12SW1
Switch WWN is 10:00:00:60:69:80:06:7a.

The new FCS list:
  10:00:00:60:69:80:06:7a

Enter WWN, Domain, or switch name(Leave blank when done): SF32SW1
Switch WWN is 10:00:00:60:69:90:03:9d.
```

```
The new FCS list:
  10:00:00:60:69:80:06:7a
  10:00:00:60:69:90:03:9d

Enter WWN, Domain, or switch name(Leave blank when done):
Are you done?  (yes, y, no, n): [no] y
Is the new FCS list correct?  (yes, y, no, n): [no] y
```

In our example we defined the FCS switches by entering their switch names; we could also define them by entering their domain ID, or WWN.

The process continues by prompting us to change the current passwords, which include:

- Root password for the FCS switch
- Factory password for the FCS switch
- Admin password for the FCS switch
- User password for the fabric
- Admin password for the non-FCS switches

The following coding shows the prompts to define each of these passwords. Also shown is the case where we entered a password that was too short; passwords must be between 8 and 40 characters in length:

```
Please enter current admin account password:
Changing password for root
New FCS switch root password:
Password must be between 8 and 40 characters long.
New FCS switch root password:
Re-type new password:
Changing password for factory
New FCS switch factory password:
Re-type new password:
Changing password for admin
New FCS switch admin password:
You cannot reuse the old password.
New FCS switch admin password:
Re-type new password:
Changing password for user
New fabric wide user password:
Re-type new password:
Changing password for admin
New Non FCS switch admin password:
Re-type new password:
```

After entering the last password verification, we received the following messages as all switches in the fabric reboot:

```
Broadcast message from root Tue Jul 29 11:36:46 2003...

Security Policy or Password Change: admin user will be logged out on switch 1

Broadcast message from root Tue Jul 29 11:36:48 2003...

Security Policy or Password Change: root factory will be logged out on switch 1

Broadcast message from root Tue Jul 29 11:36:52 2003...

Security Policy or Password Change: root factory admin user will be logged out
on switch 0
```

After all switch reboots are complete, the fabric is now secured using default policies.

With the secure fabric now enabled, we are only able to manage the fabric from the FCS switches.

If we are running FCS switches that have v4.1 or higher firmware, we can secure our fabric further by disabling the telnet daemon to our FCS switches, only allowing SSH sessions to be established. To disable the telnet interface, we use secTelnet to our FCS switch and run the **configure** command.

**Note:** The **configure** command on a secure FCS switch does not require the switch to be disabled as it normally is in a non-secure or non-FCS switch, and only presents specific options which may be changed concurrently.

```
SM12SW1:admin> configure

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

Configure...

  System services (yes, y, no, n): [no] y

    rstatd (on, off): [off]
    rusersd (on, off): [off]
    telnetd (on, off): [on] off
Broadcast message from root (pts/1) Tue Jul 29 15:24:29 2003...

Security policy change: TTY pts on switch instance 0 will be logged out.
```

As we have now disabled the telnetd daemon completely, we are only able to use an SSH client to connect to the switch. An example of an SSH client is PuTTY, which may be freely downloaded from the Internet.

Some other useful commands to view and manage the security policies are:

- **secPolicyFcsRemove**: Used to change the position of a switch in the FCS list.
- **secFcsFailover**: Used to cause the primary FCS switch to failover to the next FCS switch in the list.
- **secPolicyAdd**: Used to add members to a specified policy.
- **secPolicyRemove**: Used to remove a member from a specified policy.
- **secPolicyShow**: Displays a list of current FCS switches and identifies the primary. The output of secPolicyShow for our fabric is shown in Figure 1-167.

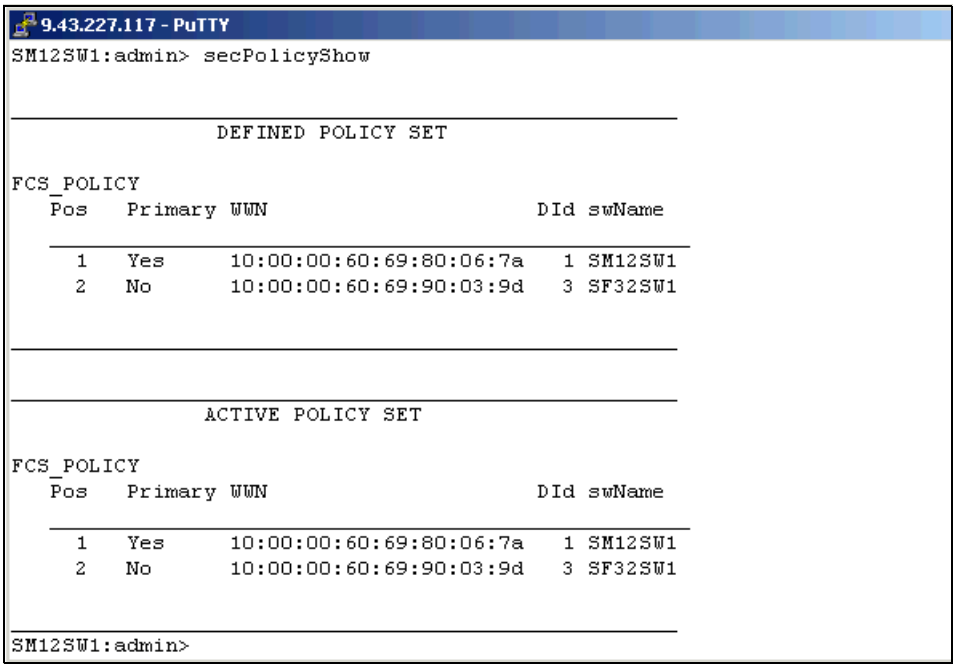


*Figure 1-167 The secPolicyShow output*

For further details on configuring security policies, refer to *Brocade Secure Fabric User's Guide*, 53-0000526.

## 1.13 Fabric Manager

Fabric Manager is an application which provides a graphical interface allowing us to monitor and manage multiple fabrics from a standard workstation. Fabric Manager can be used to manage fabric wide settings such as zoning and also manage settings at an individual switch level.

Fabric Manager provides high-level summary information about all switches in a fabric, automatically launching the WEB TOOLS interface when more detailed information is required. The launching of WEB TOOLS is transparent, providing a seamless user interface. In addition to the ability to view switches as a groups, Fabric Manager provides improved performance over WEB TOOLS alone.

Fabric Manager installs on a workstation, and can be used to manage IBM TotalStorage SAN Switches that have Fabric OS version 2.2 or later and the WEB TOOLS license installed. All the switches in the fabric are represented in the main window of Fabric Manager, but only those with a WEB TOOLS license can be managed through Fabric Manager.

## Advantages

Fabric Manager is a complete SAN management tool, and provides the following advantages:

- Provides a highly scalable Java-based application that manages multiple switches and multiple fabrics in real-time.
- Assists you with configuring, monitoring, dynamic provisioning, and daily management of SANs.
- Lowers the cost of SAN ownership by intuitively facilitating SAN management tasks.
- Saves time by enabling the global integration and running of processes across multiple fabrics through its single-point SAN management platform.
- Allows more effective management by providing rapid access to critical SAN information across both Fabric OS SANs and enhanced Fabric OS SANs.

## Capabilities

With WEB TOOLS, Fabric Manager provides the following information and capabilities:

- Configures and manages the fabric on multiple efficient levels.
- Intelligently groups multiple SAN objects and SAN management functions to provide ease and time-efficiency in administering tasks.
- Identifies, isolates, and manages SAN events across multiple switches and fabrics.
- Provides drill-down capability to individual SAN components through tightly coupled WEB TOOLS and Fabric Watch integration.
- Discovers all SAN components and views the real-time state of all fabrics.
- Provides multi-fabric administration of secure Fabric OS SANs through a single encrypted console.

- Implements scalable SAN management tasks through functionality and tools that intelligently span eight fabrics and 200 switches.
- Monitors ISLs.
- Manages switch licenses.
- Performs fabric stamping.

## Concepts

The following is a description of the concepts that are supported by Fabric Manager.

### *Logical groups*

We can create logical groups to monitor the status of their component switches and propagate actions over the chosen group of switches. We can also use this feature to quickly determine the status of a large number of switches without looking through each one. A logical group differs from a physical group in that it does not necessarily represent a physically grouped set of switches.

### *Local files*

Fabric Manager saves groups and other information to local files. Fabric Manager stores these files in our home directory. Log files are under the following directory:

```
user home/Fabric Manager/log
```

### *Import/export*

Logical groups and other configuration information can be saved to local files and shared between hosts through the Import and Export options. Additionally, configuration information can be imported from files.

### *ISL checking*

ISL checking is done by stamping or taking a snapshot of a topology. When we turn on ISL checking for a fabric, a stamp is taken of the topology of the ISLs. Then when a change occurs in these ISLs, the status of the switch changes and the detailed information is shown on the Events page.

### *Security*

**Note:** This feature is not available without Advanced Security.

Security is implemented on a policy basis. Advanced Security enables sensitive operations to be restricted to a few trusted switches. It allows us to designate a small number of switches (known as Fabric Configuration servers) for fabric-wide management operations. Individual switches will still be accessed for local

configuration. It is possible to configure Advanced Security in such a way that Fabric Manager is unable to access most of the switches. In this case Fabric Manager can only be used in a reduced mode without most monitoring features and lacking many of the administration launch points.

## 1.13.1 Requirements for Fabric Manager

Following is a description of some of the requirements for Fabric Manager.

### Switch requirements

Fabric Manager can be used to manage IBM TotalStorage SAN Switches that meet the following requirements:

- WEB TOOLS license installed.
- Fabric OS v2.2 or greater required. Fabric Manager can be used to manage switches with earlier versions of Fabric OS, but status and event information will not be available.

### Workstation requirements

The following items are required for the correct installation and operation of Fabric Manager on the computer workstation:

- One of the following operating systems:
  - `Fabric Manager Server`: Windows 2000 pro or server
  - `Fabric Manager Client`: Windows NT® 4.0, Windows 2000, Solaris 2.7 or Solaris 2.8
- Adequate RAM:
  - 128 MB for fabrics of 21 switches or less
  - 256 MB for fabrics containing more than 21 switches
- 10 MB of free disk space
- One of the following Web browsers:
  - Netscape Communicator 4.7x or 6.2.
  - Internet Explorer 5.5 or 6.x.

## 1.13.2 Installing Fabric Manager on Windows

To install Fabric Manager 4.0:

1. Insert the Fabric Manager 4.0 CD-ROM in the CD drive of the computer workstation. Fabric Manager Install will automatically start if the CD autoplay feature is enabled, otherwise run **`install.exe`**.

**Note:** If Fabric Manager is already installed on the computer, a window will display at this point to indicate this fact. If this window displays, it is necessary to exit the installer and uninstall the existing version.

2. As Fabric Manager is a Licensed software application, we are given the choice of installing the *Full version* by entering our purchased Serial number and License key, or the *Evaluation Version* of the software. We chose the Evaluation button and clicked **Next**.
3. We now must read and agree to the terms of the License agreement. After agreeing, we select **I accept the terms of the License Agreement** and click **Next**.
4. We are then asked to choose if we want to install the Client, or Server, or both Client and Server. We select Server and Client as shown in Figure 1-168, and then click **Next**.

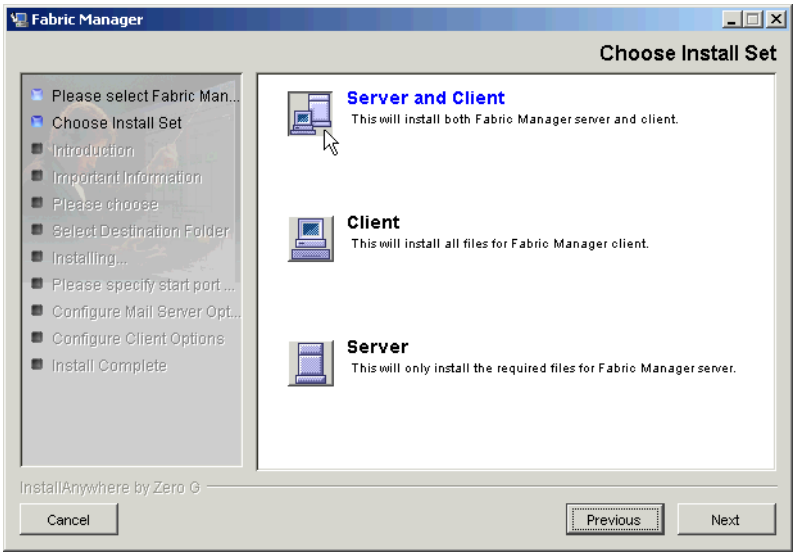


*Figure 1-168 Choose Install Set*

5. We are then asked to confirm our install choice by clicking **Next**.
6. Determine whether a compatible browser is installed. If a compatible browser is not installed, exit the Fabric Manager installation, install or upgrade the browser, and then restart the Fabric Manager installation.

**Note:** The browser should be installed before Fabric Manager is installed so that the pathname for the browser can be written to the Fabric Manager properties file.

7. Click **Next** to continue:
   a. A window that allows selection of the installation path for the client displays.
   b. Select a destination Folder using the **Choose** button or accept the default by clicking **Next**.
8. As we chose to also install the Server, we get a window asking for the Server installation path. Click **Install** to accept the default installation location, or browse for a custom location and then click **Install**.

   Once **Install** is clicked, a window showing the progress of the installation displays, with the name of the file currently being installed in the lower portion of the window.

   The install program searches the registry for the Web browser and adds the complete pathname to the FabricManager.properties file. If the installer is unable to locate a Web browser, a window displays warning that no browser was found. If this window displays, exit the Fabric Manager installation, install the browser, and then relaunch the Fabric Manager installation.
9. The Installation now configures Fabric Manager and requests us to specify a TCP/IP port range as Fabric Manager requires seven consecutive ports. We accept the default beginning from port 24600 by clicking **Next**.
10. Version 4.0 Fabric manager supports an e-mail alerting function. We are now requested to provide our SMTP mail server address, and also an e-mail ID that is assigned to Fabric Manager. We enter this information and click **Next**.

**Restriction:** These e-mail settings cannot be changed after installation, so we recommend finding the correct information before continuing with the installation.

11. The next window, shown in Figure 1-169, asks for the Windows Domain Name. This is important, as Fabric Manager authenticates logins with the controller specified by this Domain. That is, the windows domain login is used when logging into the Fabric Manager application.

**Tip:** To find out our windows Domain, we type `set` at a command prompt and locate the statement *USERDOMAIN=* .

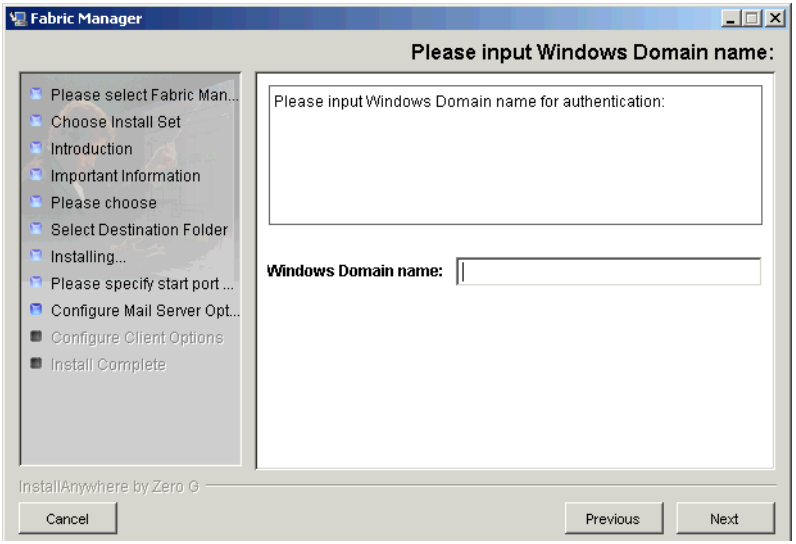


*Figure 1-169 Domain name entry*

12. After ensuring that we do not have any *services* windows open, we click **Next**.
13. The Server IP address is now requested. As we are installing the Fabric Manager Server on the same workstation, we enter `localhost` and click **Next**.

    Once the installation of Fabric Manager is complete, the window shown in Figure 1-170 displays.


*Figure 1-170 Install progress window*

14. Check the checkbox if you want to view the ReadMe file, then click **Done** to close the installer.

## 1.13.3 Installing Fabric Manager on Solaris

Preparing to use Fabric Manager to manage fabric(s) in a Solaris environment requires the following steps:

1. Installing the appropriate browser, if not already installed on the workstation.
2. Installing the required Java plug-in on the workstation, if not already installed.
3. Checking that Java 1.2 is the default Java VM installed on your system.
4. Adding Solaris Patches for the Java plug-in for systems running Solaris 7. Solaris 8 does not require patches.

### Installing the browser for Solaris

1. Find out which version of Netscape you're running by executing the following from a terminal window:

   ```
   #netscape -version
   ```

   Fabric Manager requires Netscape Communicator 4.51 or later.

2. If your version of Netscape is older than 4.51, a new version can be downloaded from the Netscape Web site or installed from CD.

On systems running Solaris 8, Netscape v4.51 is the default browser, so on Solaris 8, you should not have to make any browser adjustments. However, on previous releases, the default browser is HotJava. Fabric Manager does not support HotJava.

When Netscape is installed in the directory you specified (typically the following: /usr/dt/appconfig/netscape directory on a Solaris 8 system) you must create a symbolic link to it from the /usr/dt/bin directory.

Execute the following command from a terminal window.

```
#cd /usr/dt/bin
#ln -s ../appconfig/netscape/netscape netscape
```

Make sure /usr/dt/bin is in your path; if not, edit Users home directory *.profile* or *.cshrc* file to add it, depending upon which shell you use.

```
#vi /.profile
PATH=$PATH:/usr/dt/bin
export PATH
```

## Installing the Java Plug-in for Solaris

1. Find out if you have the correct Java plug-in installed by going to the help menu on your browser's toolbar and selecting the About Plug-ins feature.

   The Java plug-in 1.2.2_02 should display.

2. If your Java Plug-in is older than 1.2.2_02, download it from

   http://www.sun.com/software/solaris/netscape/jpis/

   It downloads in compressed tar file format into the directory you specified for download:

   – Unarchive Java Plug-in by using the following commands:

   ```
   #uncompress plugin-12-sparc.tar.Z
   #tar -xf plugin-12-sparc.tar
   ```

   – Install the Java Plug-in by using the following command from a terminal window:

   ```
   #pkgadd -d . SUNWj2pi
   ```

   – It will be installed in /opt/NSCPcom/ by default.

The browser won't load the plug-in unless it resides in your Netscape base directory. If that is /opt/NSCPcom, you don't need to make any changes; otherwise, execute the following commands:

```
#cd /opt/NSCPcom/plugins
#mv javaplugin12.so /usr/dt/appconfig/netscape/plugins
#cd ..
#mv j2pi /usr/dt/appconfig/netscape
```

– Now set your environment variable NPX_PLUGIN_PATH to point to the new plugin location by using the following commands:

```
#vi /.profile
NPX_PLUGIN_PATH=/usr/dt/appconfig/netscape/plugins
export NPX_PLUGIN_PATH
```

– Logout of your current session and log back in and this environment variable will be set.

More information about Java plug-ins is available from:

http://www.sun.com/solaris/netscape/jpis/userguide-java_plugin.html

## Checking the Java default for Solaris

Next you need to ensure that Java VM™ 1.2.2_07 is the default for Solaris:

1. Find out which Java VM is the default on your system by executing the following command:

   ```
   #java -version
   ```

The recommended VM for Solaris Fabric Manager is 1.2.2_07.

2. Downloaded the recommended VM from:

   http://www.sun.com/solaris/netscape/jpis/userguide-java_plugin.html

3. Download the self-extracting binary file into your /usr directory; then execute the following:

   ```
   chmod +x Solaris_JDK_1.2.2_07_sparc.bin./Solaris_JDK_1.2.2_07_sparc.bin
   ```

4. Make the recommended VM your default by putting the directory in your path before /usr/java.

   For example, vi.profi 6:

   ```
   PATH= /usr/Solaris_JDK_1.2.2_07/bin:$PATH
   Export PATH
   ```

5. Patches may be required to be installed for certain releases of Solaris to use this VM. If required, these may be downloaded from the same source.

6. Logout from your session, and the path will be set when you log back in. Java 1.2 VM should now be the default.

## Installing Solaris patches

Download the patches (if needed) from:

http://sunsolve.sun.com/pub-cgi/show.pl?target=patches/patch-access

## Configuring Solaris

The Solaris system should be configured to have the “number of open file descriptors” of at least 256.

1. Check the current configuration by using the **`ulimit -a`** command. The following is sample output:

   ```
   time(seconds) unlimited
   file(blocks) unlimited
   data(kbytes) unlimited
   stack(kbytes) 8192
   coredump(blocks) unlimited
   nofiles(descriptors) 64
   vmemory(kbytes) unlimited
   ```

2. Check the item `nofiles(descriptors)`. It should be at least 256. The default is 256 for Solaris 8 systems.

If it is less than 256, login as root and use the command **`ulimit -n 256`** to set it to 256, or refer to the Solaris System Administrator’s manual to configure this limit.

## 1.13.4 Launching Fabric Manager

Following we show how to launch Fabric Manager from Solaris and Windows.

### Launching in Solaris

You can launch Fabric Manager once Fabric Manager and the Java Plug-in are both installed on the workstation, and a WEB TOOLS license is installed on the switch.

To launch Fabric Manager, change to the directory in which you installed Fabric Manager, and execute the following command:

```
#./startFabricManager
```

The Fabric Manager View window displays.

### Launching in Windows

We can launch Fabric Manager once Fabric Manager and the Java Plug-in are both installed on the workstation, and a WEB TOOLS license is installed on the switch.

To launch Fabric Manager:

**Select Start —>Programs —> Fabric Manager —> Fabric Manager**

We first get a logon window where we use our Windows domain userid and password. Once authenticated, the Fabric Manager View window displays.

## 1.13.5 Implementing Fabric Manager

In the following paragraphs, we go through some of the more useful functions. For more functions and a detailed description of Fabric Manager, refer to the *Brocade Fabric Manager User's Guide,* 53_0000823.

### Fabric Manager view

The Fabric Manager detail view is the first view that displays when we launch Fabric Manager. It provides access to specific information about the fabric and switches through a panel that represents each switch. Every switch in the fabric, including any unlicensed switches, is represented by a switch panel in Fabric Manager view. However, only switches with a WEB TOOLS license can be managed from Fabric Manager. To add a license for an unlicensed switch, click the corresponding switch icon in Fabric Manager view, and a license window automatically displays.

The initial Fabric Manager view opens as shown in Figure 1-171.

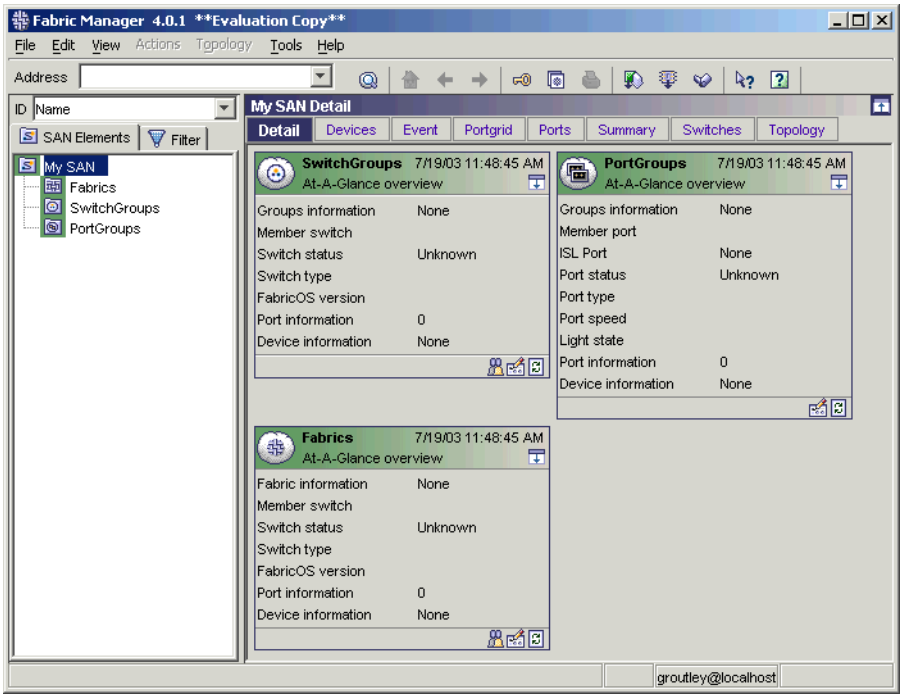


*Figure 1-171 Fabric Manager address window*

1. Type the switch name or IP address in the Address field.

**Note:** When working in a multiswitch environment, we recommend you enter the IP address of the switch with the highest port count and highest level of firmware. If an M12 is installed, then use that IP address.

2. Press Enter to submit the address.

After a correct address is entered, Fabric Manager displays the Fabric view shown in Figure 1-172.

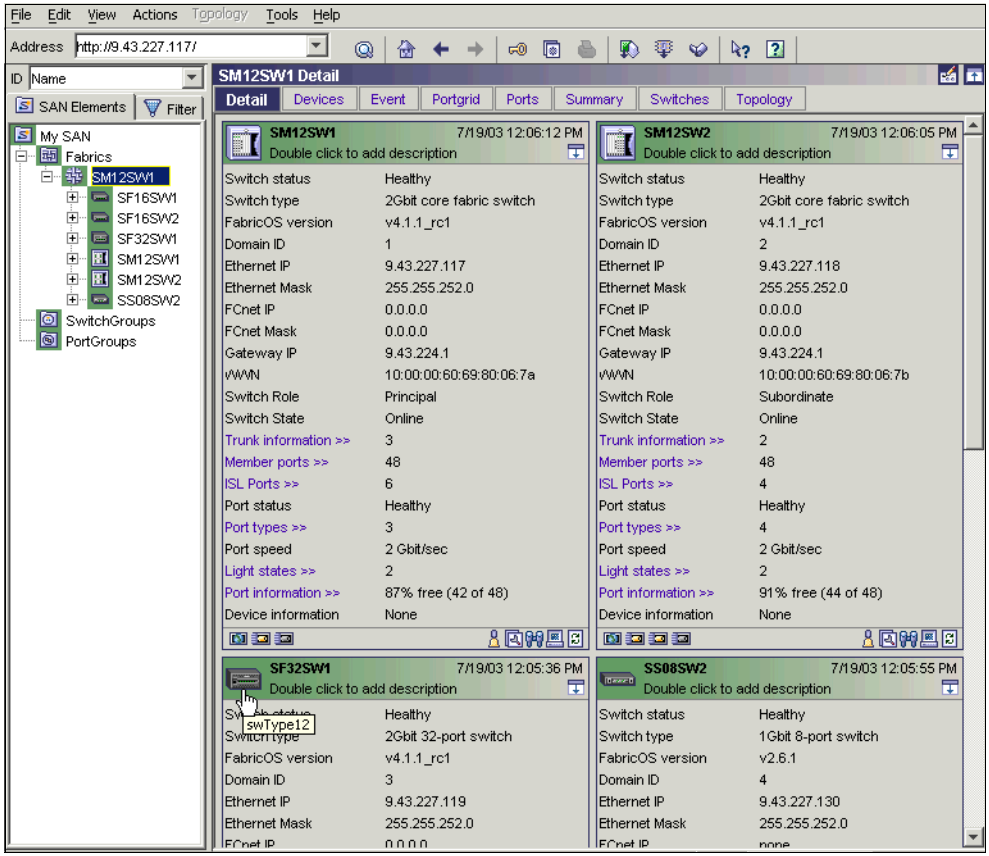


*Figure 1-172 Default View window*

The left-hand side is the SAN Elements panel. It is comprised of a pull-down menu where we can select to display by “Name”, “IP”, “Domain ID” , “WWN”, the Navigation Tree control, and two tabs (“SAN elements” and “Filter”).

The Navigation Tree control of the SAN Elements panel displays various nodes, such as Fabrics, Groups, Reboot Groups, Devices, Switches, Ports, and so on.

By selecting one of the options from the pulldown menu, we can modify the display of the SAN elements on the SAN Elements panel:

- **Name:** Displays the defined switch name.
- **IP:** Displays the switch IP address.
- **WWN:** Displays the switch WWN.
- **Domain ID:** Displays each switch’s domain ID.

The Filter panel allows us to filter the browser display and show only switches matching one of the following criteria:

- IP
- Name
- Type
- Version
- WWN
- Domain ID

To filter the display, choose one criteria in the list box, type the desired value in the edit box, and press Enter. This displays a window similar to Figure 1-173.

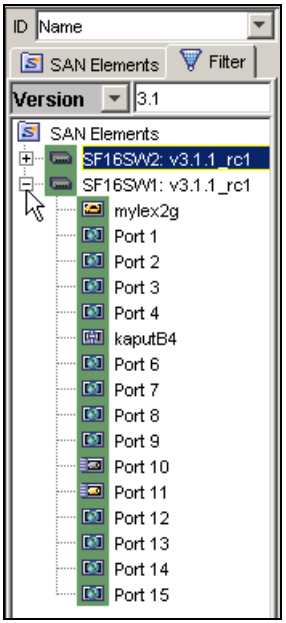


*Figure 1-173 Applying filter to SAN elements display*

In Figure 1-173, we want to restrict WWN display to devices running firmware version v3.1.

The right-hand side of the Fabric View window is the Switch View portion of the Fabric View. We can use it to manage individual switches.

From this view, we can access switch specific operations such as:

- Switch events
- Switch settings
- Telnet window
- Switch front panel view

Launching the Switch View in Fabric Manager actually launches the WEB TOOLS interface for that switch.

Depending on our selection in the navigation tree, the Switch View will display either a fabric icon or individual switch icons.

Figure 1-174 shows the window display at a fabric level.

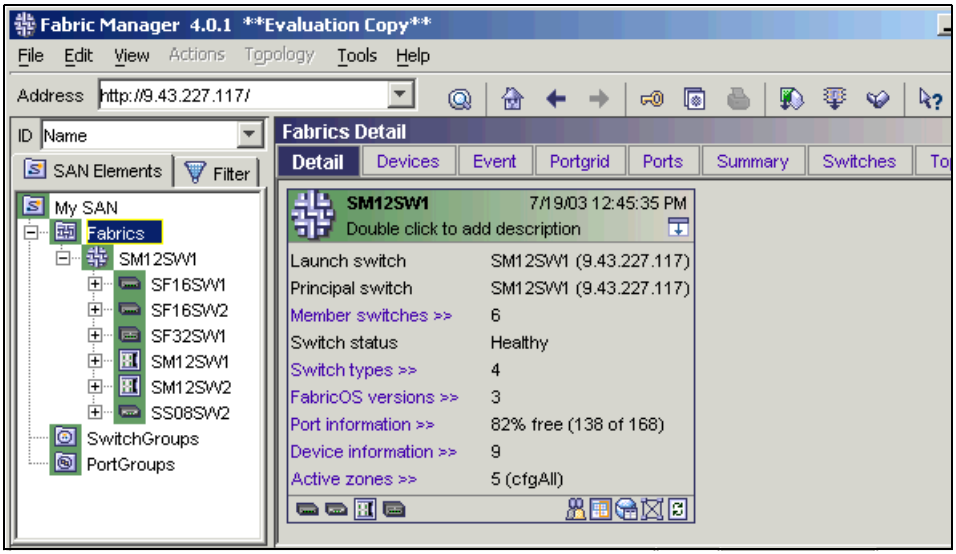


*Figure 1-174 Fabric Detail*

From the icons on the right hand side of this window, we can access fabric-wide operations such as:

- Fabric events
- Zone administration
- Name server
- Fabric topology

## Setting the File Transfer options

In order to get certain information from fabric switches, Fabric Manager needs to be able to connect to an FTP server. This FTP connection would be used, for example, to retrieve all configuration information.

To set the File Transfer options, go to **File —> Options**. This displays the window shown in Figure 1-175, where we select *File Transfer* from the *Configurations* tree menu on the left.

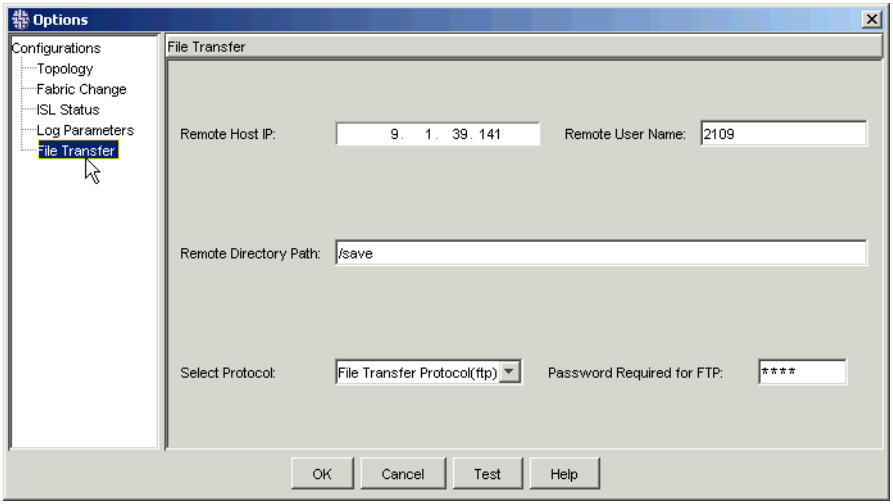


*Figure 1-175 File Transfer options*

In this window we set the IP address of the remote FTP server as well as all necessary valid user information. We can then test the connection using the “Test” button. Fabric Manager will attempt a connection and return the result.

## Creating logical groups

Logical groups allow us to operate on a set of switches that are not necessarily physically connected or part of the same fabric. For example, we could create logical groups according to the switch model.

To create a logical group, right-click the **Group** node in the Navigation Tree and choose **Edit** from the context menu, as shown in Figure 1-176.

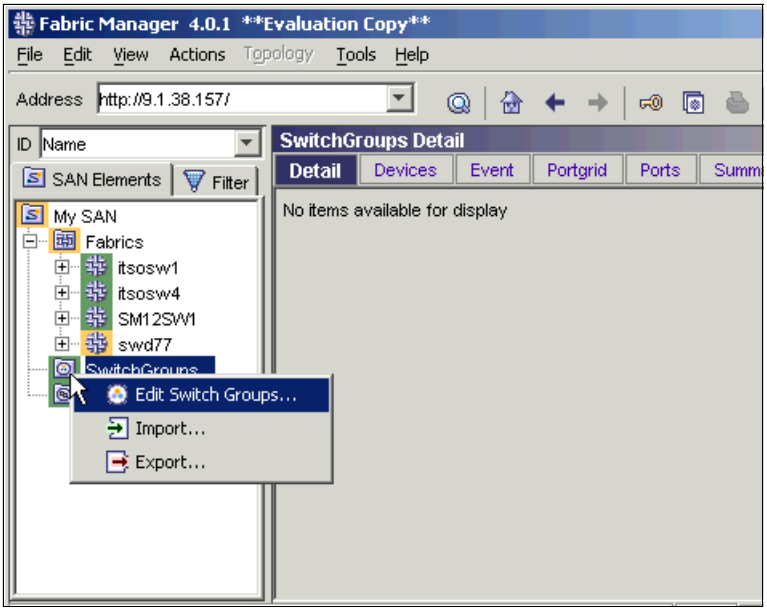


*Figure 1-176 Access the edit group window*

The window shown in Figure 1-177 is then displayed.

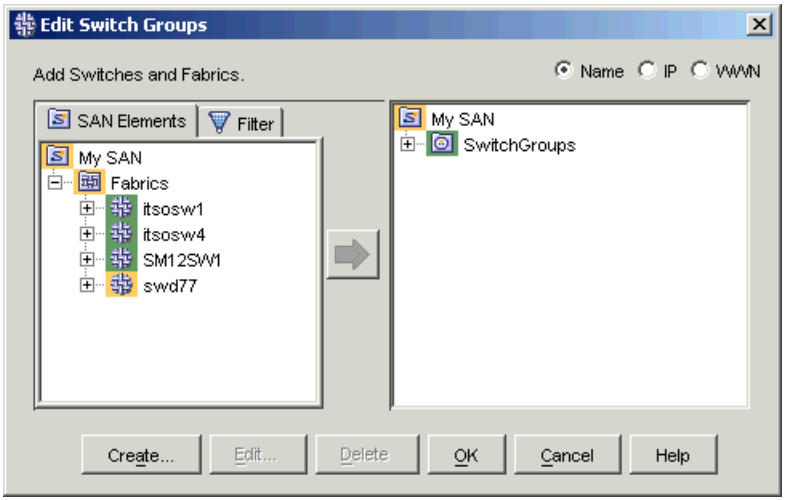


*Figure 1-177 Edit group window*

On the left-hand side of this window is the list of all the switches in the registered fabrics. We can filter the display as well by choosing the filter tab, and we can select to view the fabrics and switches by Name, IP address, or WWN.

On the right-hand side is the list of existing groups. In our example, we currently have no defined groups.

To define a new group, click the **Create...** button. This brings up the window shown in Figure 1-178.

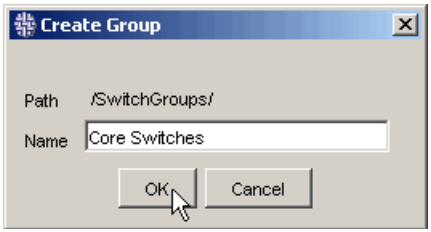


*Figure 1-178 Enter the group name*

Enter the group name and click **OK**.

The group name now appears in the left-hand side list. To add a switch to this group:

- Select the group.
- Select the switches.

**Tip:** Hold the Ctrl key to select more than one switch.

- Click the right-arrow or drag and drop the selections in the navigation tree to the group.

Figure 1-179 shows an example of this.

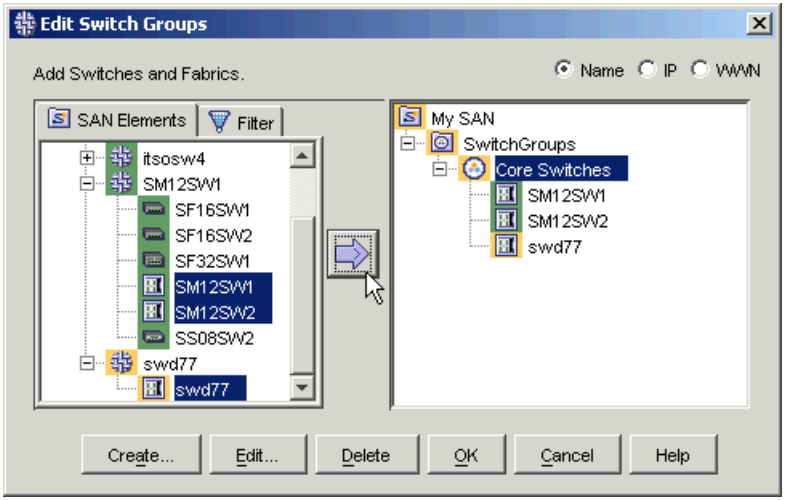


*Figure 1-179 Add switches to the group*

Click **OK** to close this window. The group is now visible in the *SwitchGroups* View in the navigation tree as shown in Figure 1-180. We have also chosen to view our group with the **Switches** tab in this figure.

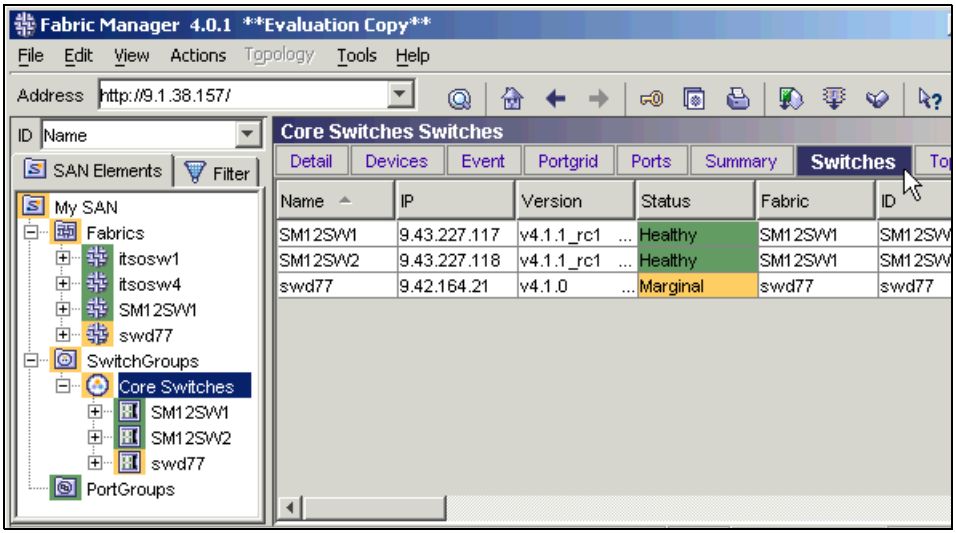


*Figure 1-180 Groups in the Fabric View*

### Sharing logical groups definitions

We can export logical group definitions in order to back up our configuration or to share this definition with another host.

To share logical groups definitions, perform the following steps:

1. Select **File —> Groups**.
2. Select **Export**.
3. Use the Browse button to select a file to Export a Group to.
4. Type a name for your “group” file.
5. Highlight the name of the group(s) to be exported from the navigation-tree.
6. Add the group to be exported by clicking the arrow button, or by dragging and dropping selections from the navigation-tree to the table.
7. Select **Save**.

We can now import our group to a separate Fabric Manager machine.

1. Select **File —> Groups**.
2. Select **Import**.
3. Browse to select the file you previously exported to.

## 1.13.6 Fabric Login

In order to be able to operate on the switches in the fabric, we need to perform a “Fabric Login”. Fabric Login is necessary, for example, to perform firmware upgrades or a switch reboot.

To define the Fabric Login procedure, click the key icon in the Fabric View as shown in Figure 1-181, which will launch the process.

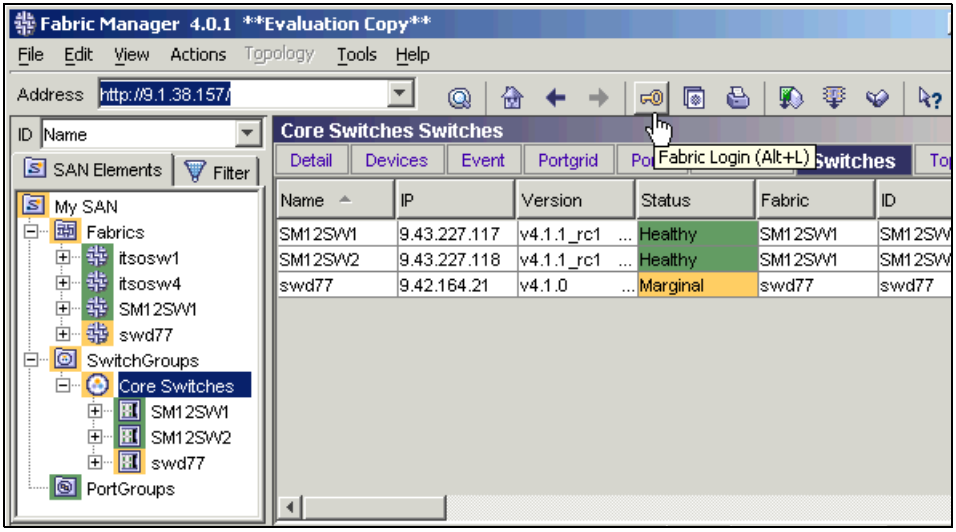


*Figure 1-181 Launch Fabric Login*

To login to multiple switches:

- From the left-hand side navigation tree, highlight the switches or groups of switches to be selected. (We can select multiple items by holding down the Ctrl key while clicking).
- Use the Add/Delete arrows in the middle column to select the switches.
- The selected switches will be applied in a table with all their details.
- Enter the User Name and Password that apply to the switches you selected. This User Name is the same as the one you would use to log into the switch using a Telnet command.
- Choose the **Apply** button to test and apply the login.

Figure 1-182 shows an example of the Fabric Login window.

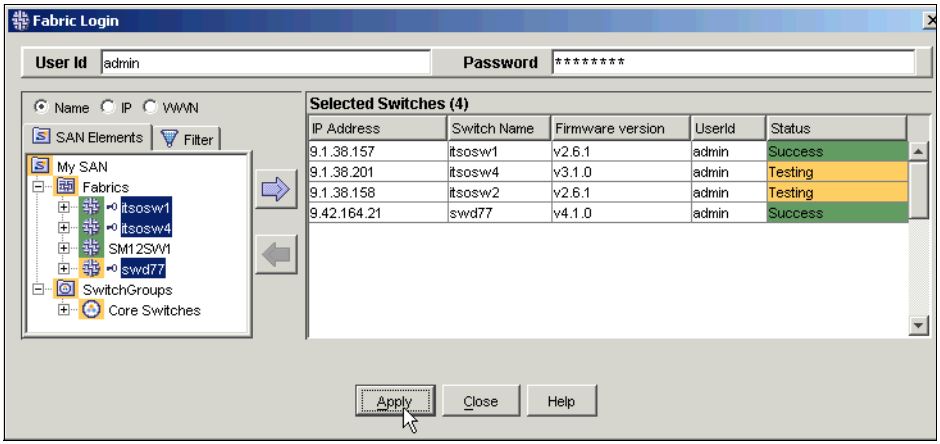


*Figure 1-182 Test and apply login information*

The Status field in this window gives us information about the login process.

Referring back to the Fabric View window, all switches for which the login test has been successful are identified by a key icon in the navigation tree, as shown in Figure 1-183.

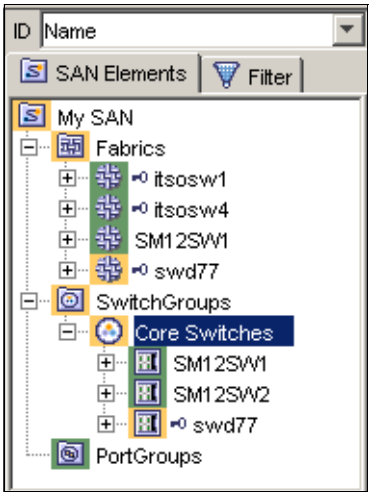


*Figure 1-183 Navigation tree after successful login test*

## Downloading firmware to multiple switches

Fabric Manager allows us to upgrade firmware on multiple switches without having to log into every single device and run the firmware download process.

Prior to downloading firmware to multiple switches, you should make sure that you are logged into the switches you want to upgrade.

We access the firmware download by clicking the **Download Firmware to switches** icon as shown in Figure 1-184.

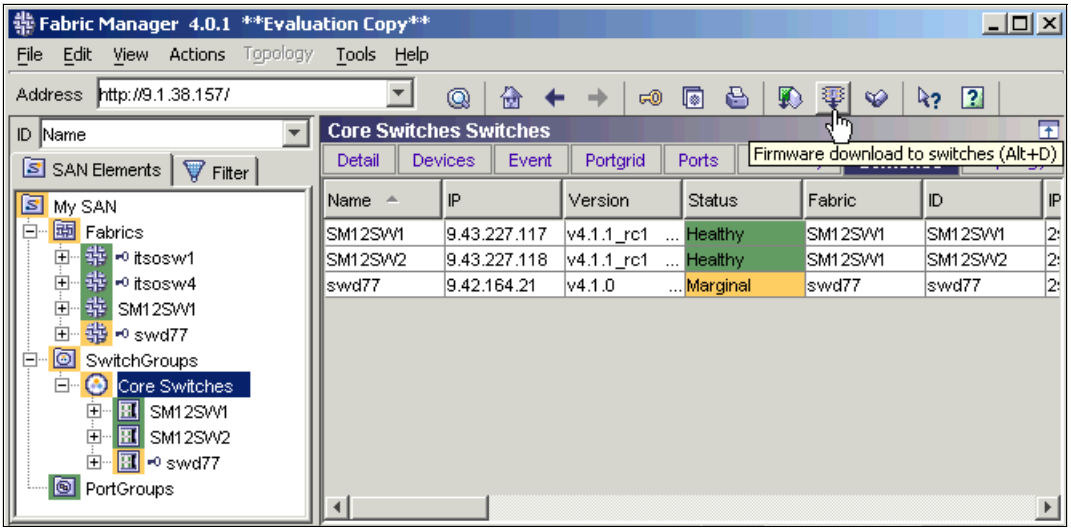


*Figure 1-184 Firmware download icon*

The firmware download window is then displayed as shown in Figure 1-185.

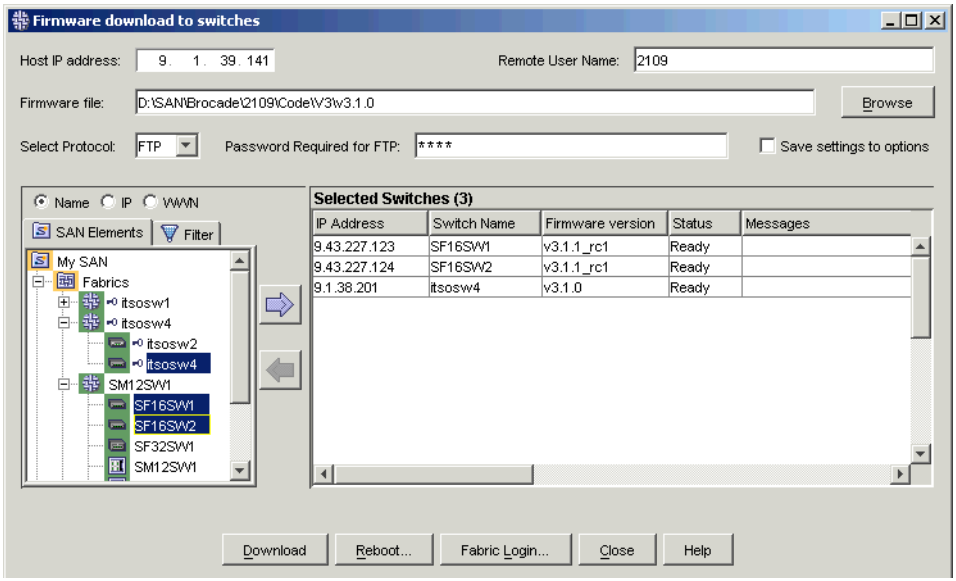


*Figure 1-185 Firmware download window*

To use the Download Firmware window to upgrade the firmware of multiple switches:

- Highlight switches or groups of switches to be targeted for firmware upgrade.
- Use the Select/Deselect arrows in the middle column to move the switches or drag and drop from the navigation window to the table.
- The selected switches will be applied in a table with all their details.
- Enter the Host Name or Host IP address.
- Enter the Remote User Name.
- Use the **Browse** button to select a firmware file from the local host.
- Select download protocol (RSHD or FTP).
- If FTP is the chosen protocol, enter the FTP password.
- Choose the **Download** button to begin firmware download.

Once the download process is begun, you can check the process status in the status field, as shown in Figure 1-186.

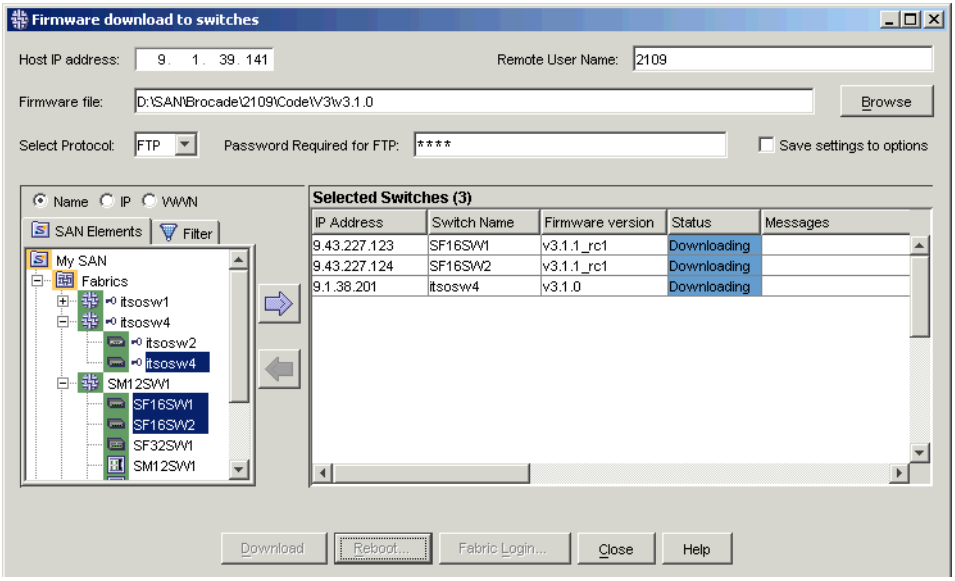


*Figure 1-186 Firmware download status*

As soon as the firmware download is completed successfully, the Status field will turn green. Note that for the new firmware to take effect, we need to reboot the switches. This can be done by clicking the **Reboot** button and following the steps described in the next section.

## 1.13.7 Rebooting switches

Fabric Manager allows us to manage switch reboots and operate on multiple switches at a time.

## Create a Reboot Group

The first step is to create Reboot Groups. To do so, select **Tools —> Reboot —> Create Reboot Sequence**. This displays the window shown in Figure 1-187.

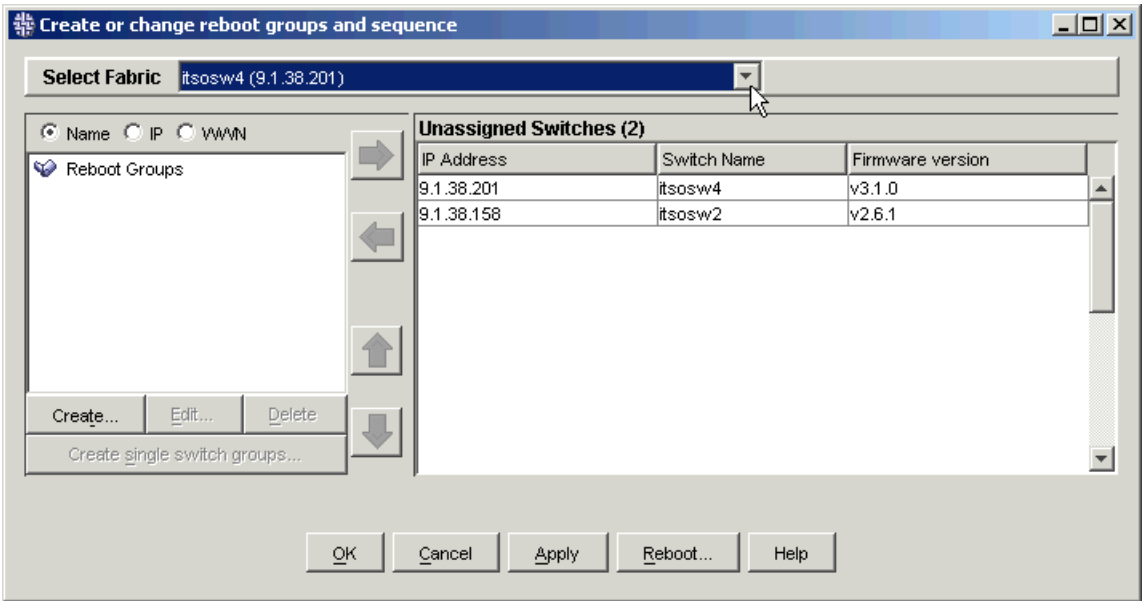


*Figure 1-187 Create reboot groups*

On the left hand side of the window are the created groups and on the right hand side are the switches available in the fabric which we have chosen from the *Select Fabric* pulldown list.

To create a group, click the **Create** button. This displays a new window. We enter a group name in the *"Name of the reboot group"* field. Then click **OK** to save and return to the main window. We can then add switches following these steps:

1. Highlight the group on the left side list.
2. Highlight the switches to add on the right side list.
3. Click the left *Assign Switches to Reboot Group* arrow.

The switches then appear under the Group name as shown in Figure 1-188.

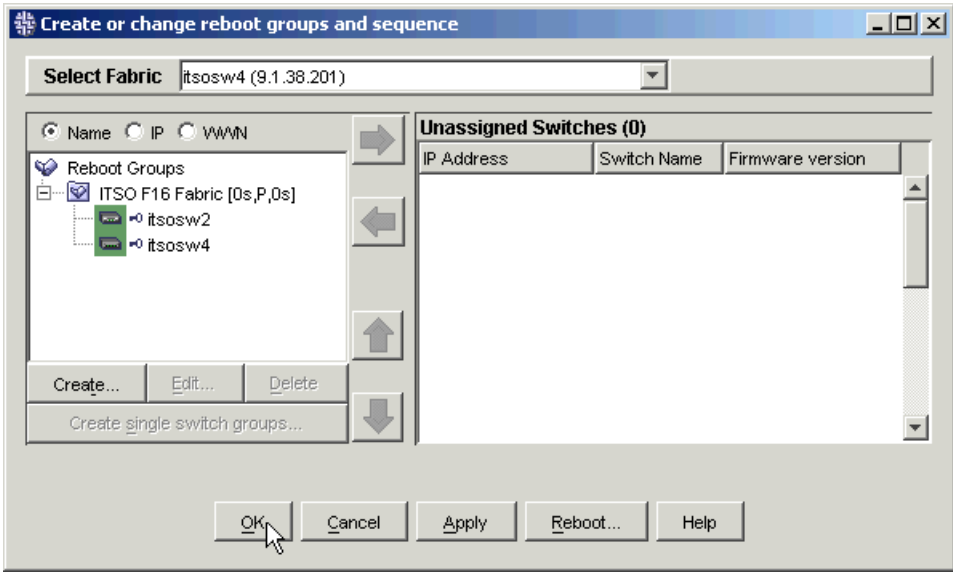


*Figure 1-188 Add switches to a reboot group*

We then click **Apply** to save or click **OK** to save and exit.

## Rebooting the switches

To reboot switches, either select **Tools—> Reboot—> Sequence Reboot** or click the **Sequenced reboot** button shown in Figure 1-189.

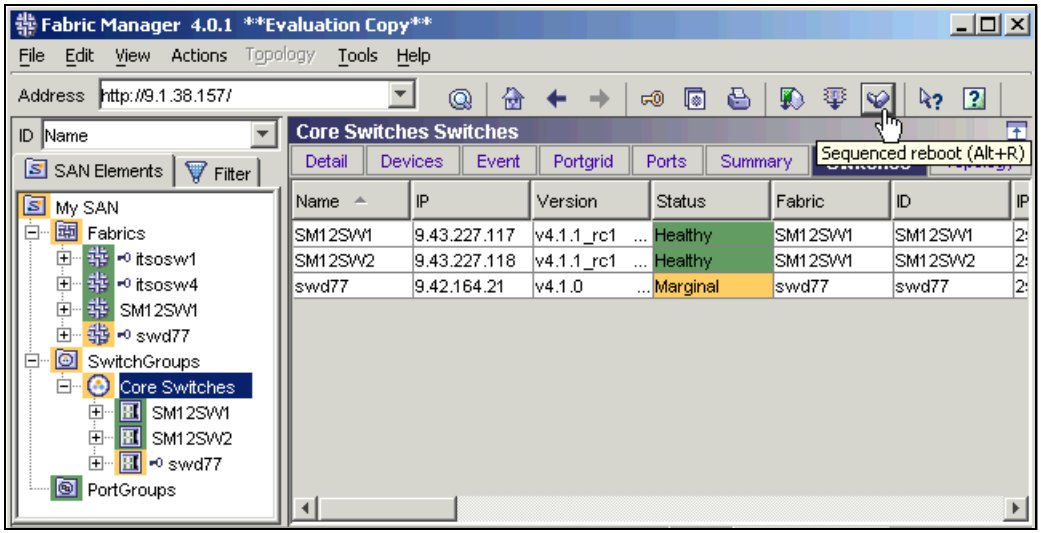


*Figure 1-189 Sequenced reboot window*

Once the Sequenced reboot window is open, the list on the left side displays the group we defined for our fabric in “Create a Reboot Group” on page 253 and the list on the right displays the switch(es) selected for reboot.

To add switches to reboot, highlight a switch, or a group, then click the right *Select Switches* arrow.

Then we select either the **Fastboot** or **Reboot** button to perform on the selected switches. We can see the switch status of the reboot process as shown in Figure 1-190. As we have chosen to reboot both switches in the fabric, the fabric in the navigation tree also turns red, indicating that the fabric is down.

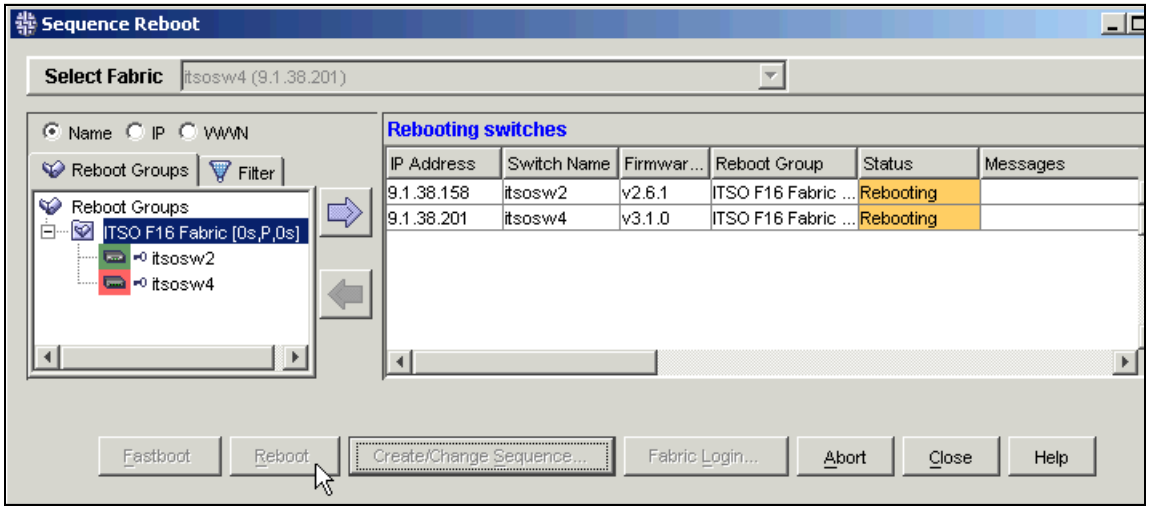


*Figure 1-190 Switches rebooting*

Once the reboot is finished, we receive an Information window notifying us that the reboot sequence is complete; also the “Status” field will display *Done* in green. We can then click **Close** to exit the window.

## 1.13.8 Fabric Merge

When merging two different fabrics, conflicts related to zoning, domain ID or operating parameters can occur, causing the new fabric to be segmented.

The Fabric Merge function allows you to check the compatibility of two fabrics before actually merging them.

You can launch “Fabric Merge” by going to **Tools —>Fabric Merge** as shown in Figure 1-191.

For example, in this section, we will work with two fabrics:

- Fabric A with one hub
- Fabric B with two switches

Each of these fabrics has its own set of domain IDs, zoning configurations and operating parameters.

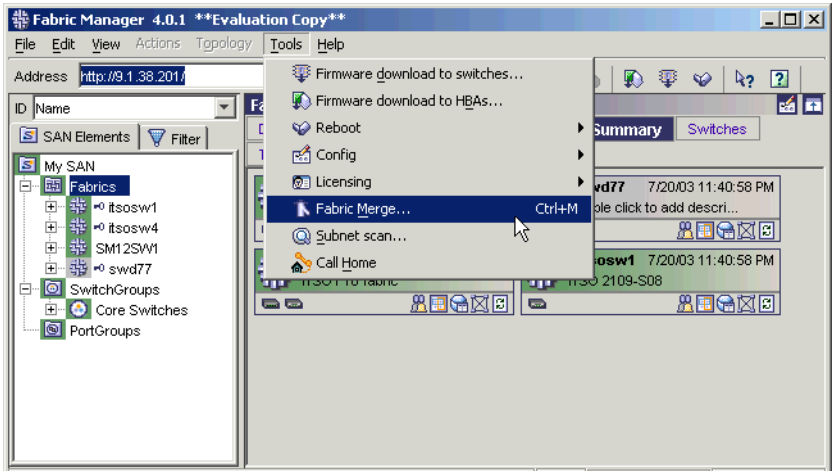


*Figure 1-191 Launch the Fabric Merge window*

The first step is to choose the two fabrics to merge, as shown in Figure 1-192.

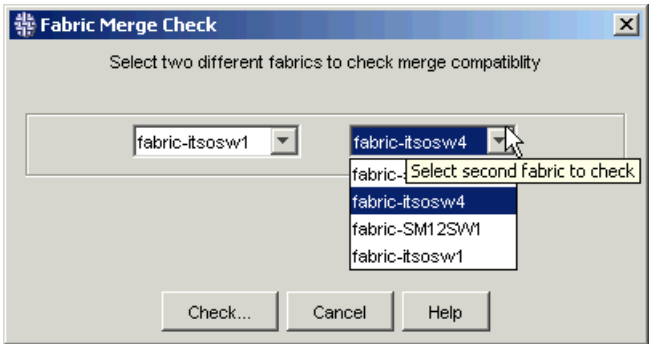


*Figure 1-192 Choose two fabric to merge*

For the two fabrics specified here, Fabric Manager downloads the configuration file and checks for any inconsistencies with respect to zoning, domain IDs, and various operating parameters.

Once you have clicked the **Check** button, Fabric Manager attempts to connect to each of the fabrics and download their configuration files to the FTP server defined in “Setting the File Transfer options” on page 242.

Once the Fabric Manager gets the configuration files, it compares them. In Figure 1-193 we show an example of the parameters not matching, due to core PID not matching.

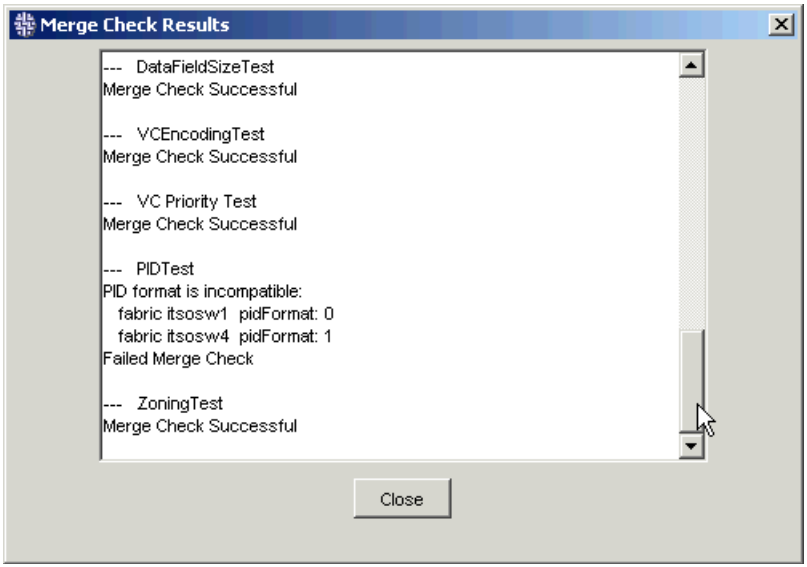


*Figure 1-193 Merge check failure*

At this point, we would now close the Merge manager, and manually configure our core PID to match in both fabrics.

If all fabric parameter settings pass the checking, we are then prompted to run the zone merge manager as shown in Figure 1-194.


*Figure 1-194 Zone merge manager prompt*

By clicking **OK** we let Fabric Manager help us to resolve conflicts. Fabric manager displays a window shown in Figure 1-195 with each fabric's configuration listed.

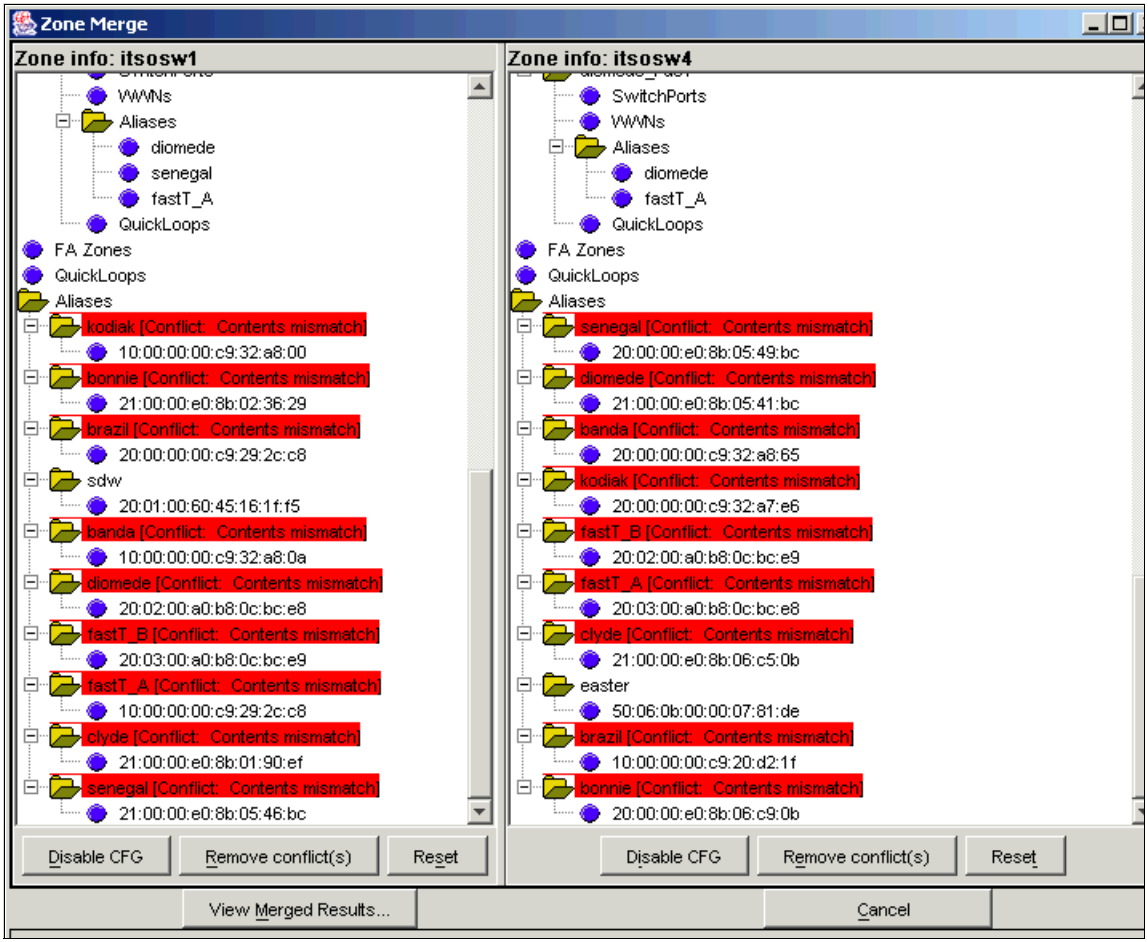


*Figure 1-195 Zone Merge window*

The conflicts are highlighted in red in each config tree. In our example, we have conflicts because the configurations both have duplicate alias names.

We can remove the conflicts in one of the fabrics by selecting the conflicts and clicking the **Remove conflict(s)** button. After removing a conflict, we could restore it by clicking the **Reset** button.

In our example, this will remove all the aliases for second HBA in each host. This would not be a desirable result, so we cancel the Merge Manager, and alter our aliases on one fabric. Then, when rerunning the Merge Manager, our configs do not have any conflicts, although the config names are highlighted in red, as shown in Figure 1-196.

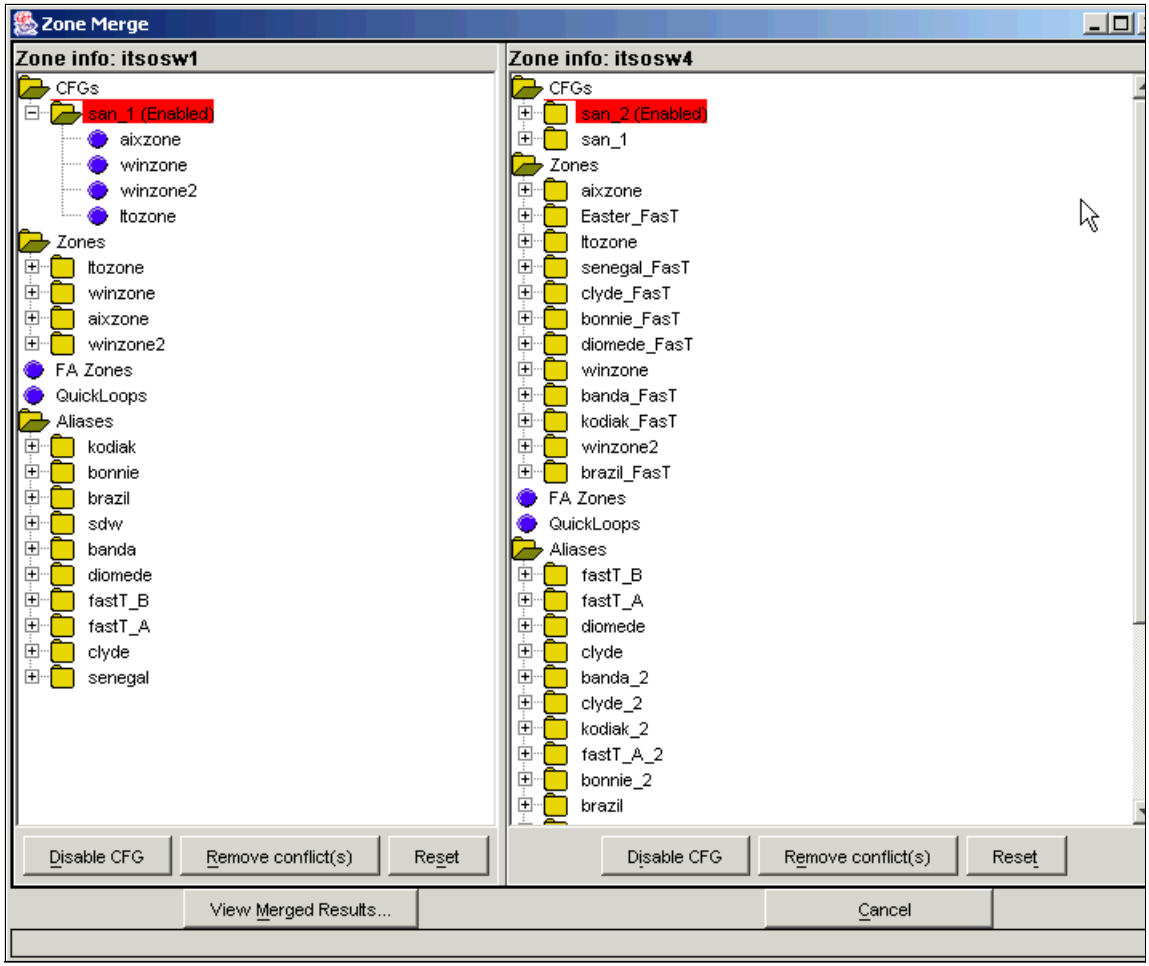


*Figure 1-196 Zone merge conflict removed*

We need to disable one of the fabrics configs, so that the merge can occur. We use the appropriate **Disable CFG** button to do this.

Now we can click **View Merged Results** to display the final zoning information as shown in Figure 1-197.

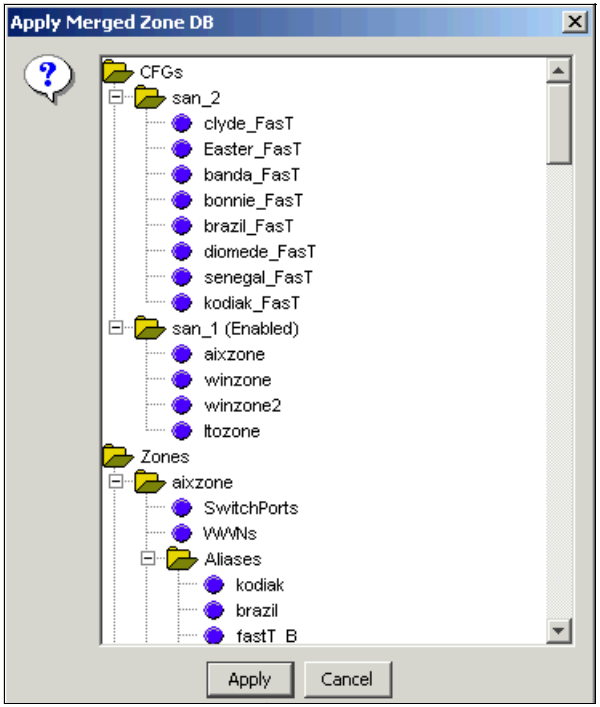


*Figure 1-197 Merged zone window*

From this window we can apply the displayed zoning configuration or cancel to return to the previous window.

**Attention:** Clicking **Apply** will modify the zoning configuration in both fabrics according to the display shown in Figure 1-197, even if the merge is not completed. In our example, the previously active configuration “SAN_2” in Fabric itsosw4 was disabled.

Once these steps have completed, without errors, the two fabrics are ready for merging by connecting a physical ISL between them.

**Tip:** We can use Fabric Manager’s ability to load configuration parameters to multiple switches to configure a whole fabric without having to logon to every single switch.

## 1.13.9 Loading switch configuration

Fabric Manager allows us to download switch configuration parameters to a file and upload this configuration or part of it to multiple switches.

This can be used, for example, to set SNMP information or fabric operating parameters to multiple switches without having to set these values on each individual device.

The first step is to save an existing configuration from a switch. This can be done by accessing the switch configuration menu **Tools —> Config —> Save Baseline** in the Fabric View. This brings up the window shown in Figure 1-198.

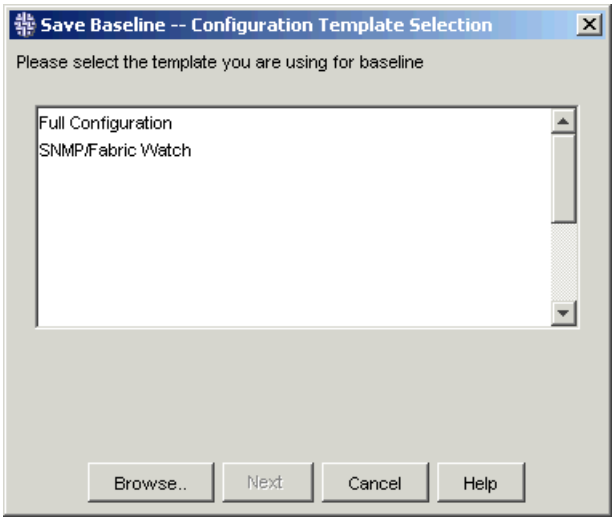


*Figure 1-198 Save Baseline selection window*

In this window you can select the way in which Fabric Manager will present the configuration parameters:

- **Full Configuration:** This lets you choose from among all the parameters.
- **SNMP/Fabric Watch:** This restricts the selection to SNMP and Fabric Watch parameters only.

In our example, we will choose **Full Configuration**.

Selecting one of the above templates will enable the **Next** button.

The next step is to choose the switch from which you wish to download the configuration, as shown in Figure 1-199.

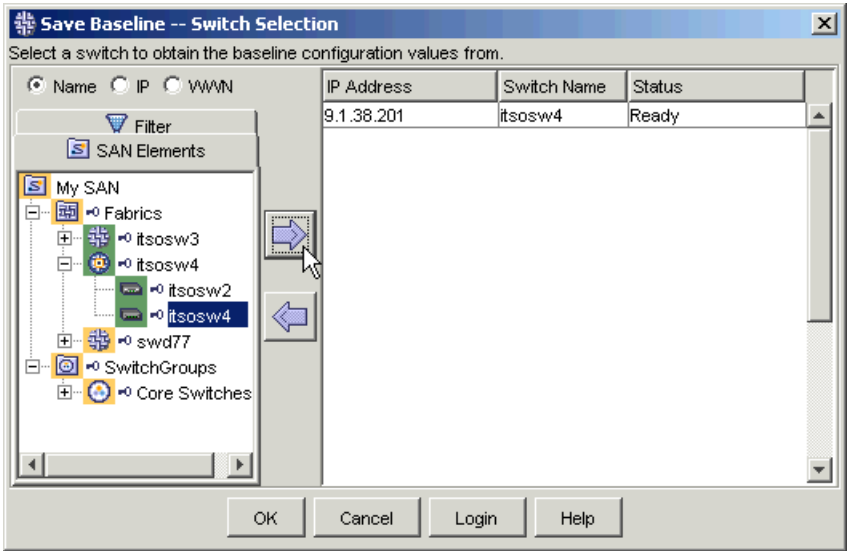


*Figure 1-199 Save Baseline — Switch selection*

Select the switch from the left-hand list and click the right facing arrow. This adds the switch to the left-hand list. You can download the configuration from only one switch at a time.

You can use the **Login** button to define the log into the switch if it is not already done.

At this time, you should make sure that the FTP server specified in the options is running. Clicking **OK** will start the download of the switch configuration file for file manager internal process. The window shown in Figure 1-200 is displayed.

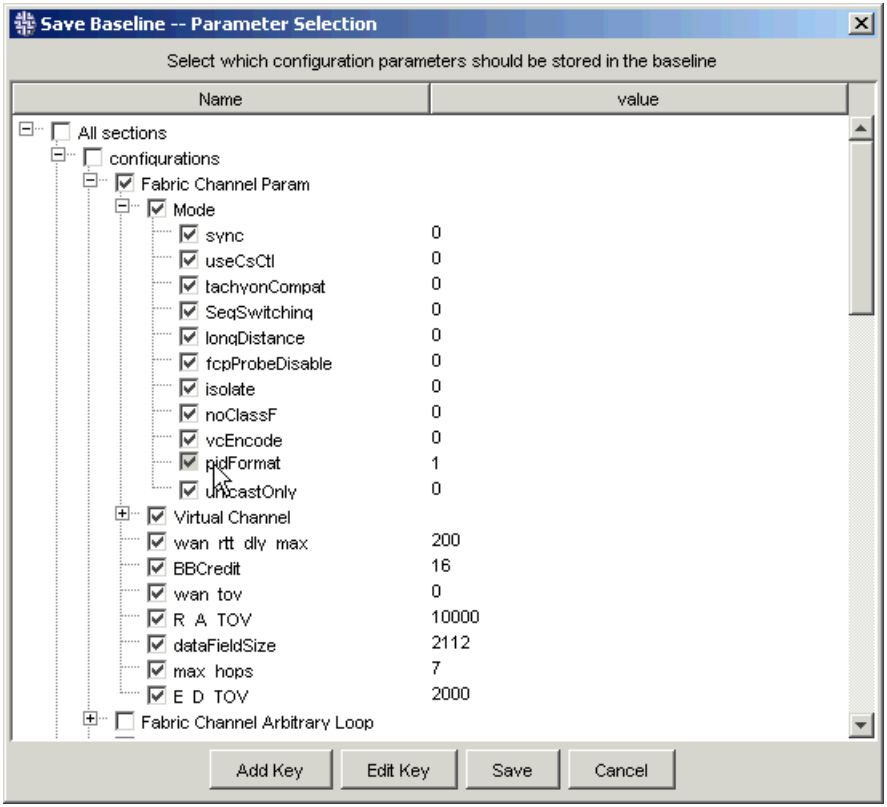


*Figure 1-200 Save Baseline — Parameter Selection*

From this window, we can choose which parameters or set of parameters we would like to save by checking the corresponding check boxes. In this example, we choose to save only information related to Fabric Parameters. If we would like to change a parameter before saving this Baseline, we can select the *key*, we chose **pidFormat** (the checkbox is slightly greyed), and then clicked the **Edit Key** button, giving us the window shown in Figure 1-201.

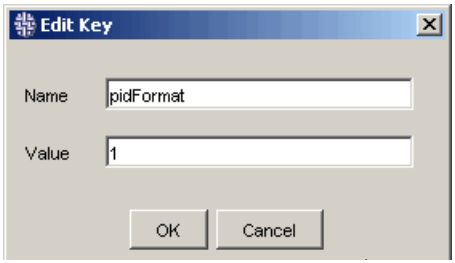


*Figure 1-201 Edit parameter key*

From the Edit Key window we can change the *Value* field to what we desire to be set as our Baseline save.

Once we have chosen the parameters to be saved we click **Save**. This will open a file browsing window where we are able to specify a location for the configuration file as shown in Figure 1-202.

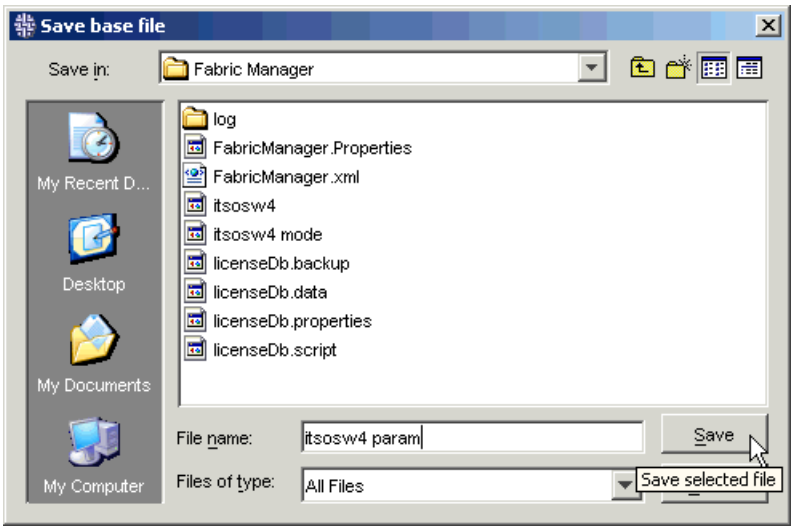


*Figure 1-202 Choose a location for configuration file*

The saved file can now be used to upload the parameters to another switch later on, or can be kept as a backup.

## Compare and download file from a file

We can use the file saved in the preceding paragraph to propagate the saved parameters to multiple switches. This can be useful for SNMP information or fabric wide parameters, for example.

Go to **Tools —> Config —>Compare/Download from File**.

The first step is to choose the file in which configuration parameters are stored. We are prompted to choose a configuration file as shown in Figure 1-203.

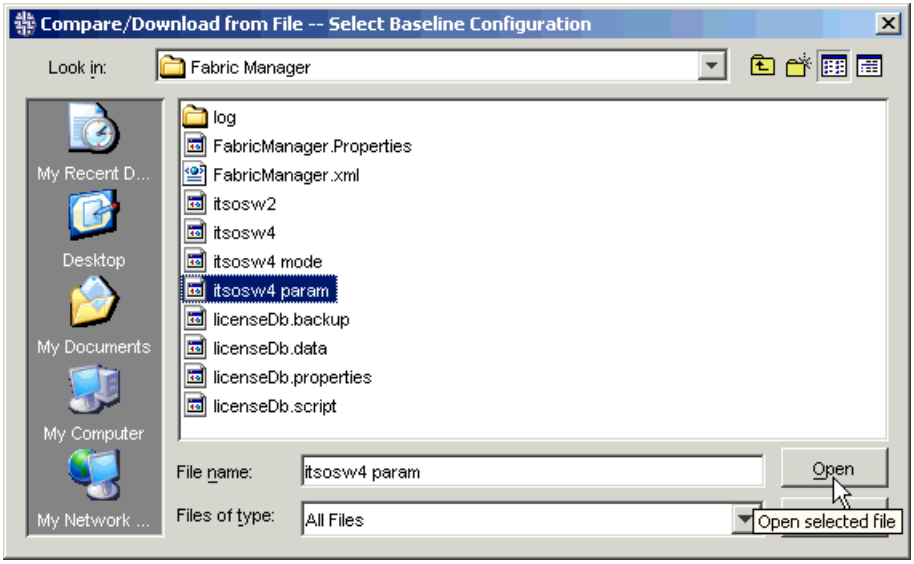


*Figure 1-203 Select configuration file to compare/download*

Next, you have to choose the target switches, that is to say, the switches to which you want to apply the configuration. This is shown in Figure 1-204.

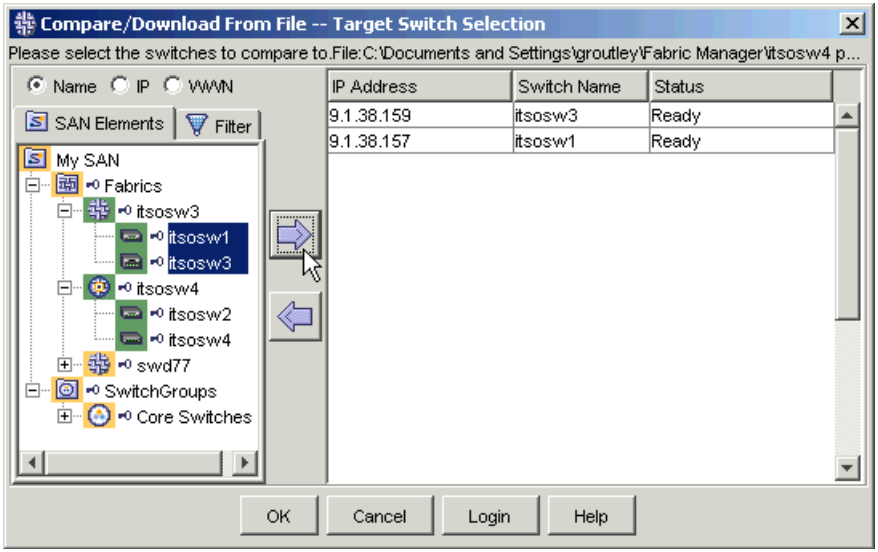


*Figure 1-204 Compare download from file — Target Switch Selection*

From the left-hand side list we can select multiple switches. Then click the right facing arrow or drag and drop the selection to the right-hand side list.

Clicking **OK** will start the configuration download from the target switches. Fabric Manager then compares the parameters available in the baseline file to the ones set in the target switch and displays the window shown in Figure 1-205.

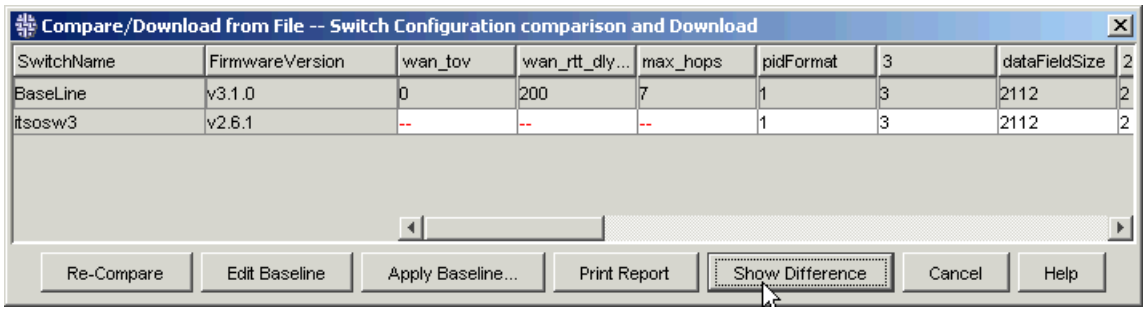


*Figure 1-205 Compare/Download from file — Comparison*

This window displays in red the differences between the baseline file and the current switches settings. Clicking the **Show Difference** button will show *only* the differences. Then we have the choice to print the comparison report, cancel the operation, edit or apply the baseline, or perform the compare again.

We chose to apply the baseline, so the window in Figure 1-206 is displayed.

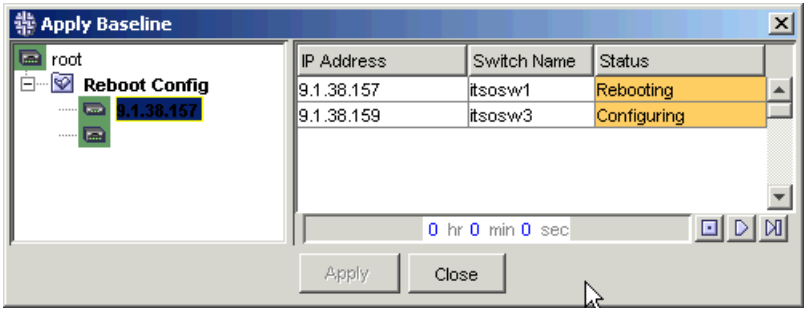


*Figure 1-206 Apply baseline to the switches*

Fabric Manager will upload the parameters to each switch, one at a time and reboot it. As one switch is done (configured and rebooted), it will have a strike-through in the switch list in the left-hand side of the window. Note that you can check the status of the switch being updated in the Status field.

Once the baseline is applied to all switches, you can click **Close** to return to the Fabric View.

## 1.13.10 Managing licenses

Fabric Manager lets you manage licenses on switches across the fabric. You can:

- View licensing information on each individual switch
- Save licensing information from a switch to a local file for backup, for example
- Download a license file to a switch for upgrade

To manage licensing, go to **Tools —>Licensing —> Load from switch**. This displays a switch selection window. Select one or more switches in the left-hand side list and click the right arrow. Validate with “OK”.

Note that you have to be logged into the switch. If not, Fabric Manager will display the fabric login window and let you enter login information.

The License Administration window is shown in Figure 1-207.

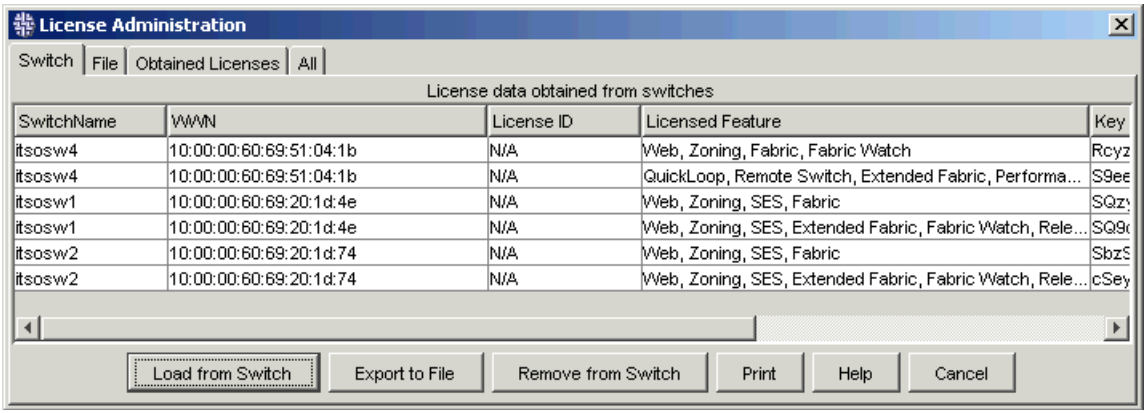


*Figure 1-207 License administration — Switch tab*

Four tabs are available in this window:

- **Switch:**
  - Allows us to view licenses currently installed on the selected switches.
  - Loads licensing information from switches by clicking the “Load from switch” button.
  - Saves the selected license information to an XML file by highlighting the appropriate line(s) and clicking “Export to file”.
  - Allows a specific license from the display and removed, using the “Remove from switch” button.
- **File:**
  - Allows us to load licensing information from a saved XML file for display

- Lets us select a displayed license and install it to the corresponding switch

- **Obtained Licenses:**
  - Allows the management and installation of electronically purchased Licenses.
- **All:**
  - Allows you to have a consolidated view of all licenses displayed on the other three tabs.

**Tip:** Do not remove the Web license, as it is required to use Fabric Manager on a switch!

The **File** tab is shown in Figure 1-208.

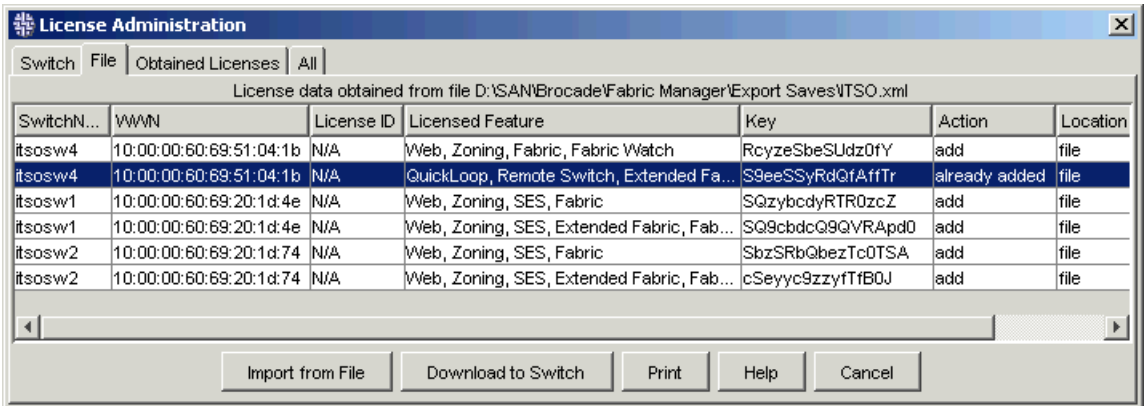


*Figure 1-208 License Administration — File tab*

## ISL Checking

Use the ISL option of the Actions menu to record and monitor the ISL configuration for a fabric. There are two separate actions that can be taken when using the ISL option:

- ISL Checking
- Restamp (available only when ISL Checking is enabled)

### *Enabling ISL checking*

When ISL Checking is enabled, a snapshot (or stamp) is taken of the topology. When a change occurs in the ISLs, the detailed information will be shown on the Status Reason section of the Events.

To enable ISL Checking:

- Highlight the Fabric.
- Select `Actions`.
- Select `ISL`.
- Select `ISL Checking`.

A check mark should appear, showing that ISL checking is enabled and the node icon will change.

When ISL Checking is turned on, the following changes will take place:

- If an event occurs, the Status Legend color will change to represent the appropriate Event.

  In our example we have lost the ISL on ITSOsw4 fabric. The navigation frame now shows itsosw2 and itsosw4 as two separate fabrics, with itsosw4 in red indicating a down condition.

- Any events will be shown in the Fabric Events window.

  We received an event stating that the connection between switch 4 port 9 and switch 2 port 7 has been removed, resulting an a Down fabric status.

Figure 1-209 shows the event entry for our example of ISL loss.

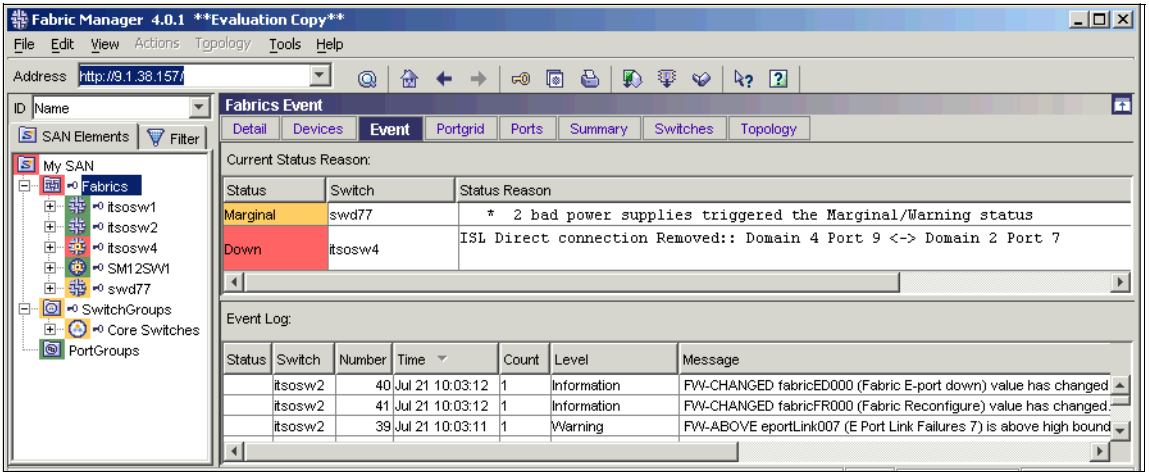


*Figure 1-209 ISL Checking event entry*

### *Using Restamp*

As enabling ISL checking takes a persistent snapshot of the fabric topology, we need to refresh this snapshot when adding or removing ISLs.

To restamp the fabric topology, and have the most current topology information noted in Fabric Events, use the Restamp option. This option is only available if ISL Checking is already enabled. To restamp the fabric:

- Select **Actions —> ISL**.
- Select **Restamp**. A snapshot is taken of the fabric
- Select **Actions —> Fabric Events**.

You can now view the latest changes within the fabric in the top of the Fabric Events window.

## Security

After enabling an Advanced Security fabric as discussed in 1.12, “Advanced Security” on page 211, we are able to manage the security policies from Fabric Manager.

By right-clicking our fabric icon we launch a menu as shown in Figure 1-210 where we select the **Security...** option.

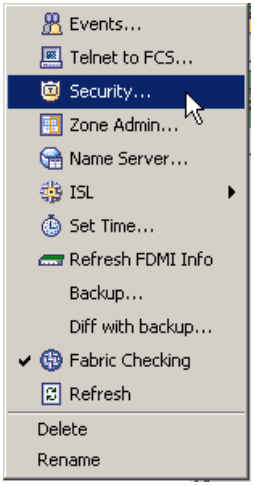


*Figure 1-210 Selecting Security management*

When we do this we receive a message shown in Figure 1-211 indicating that passwords have not been learned, although Fabric Manager previously had been defined with passwords for this fabric, during the enabling of Advanced Security we were forced to change all the passwords.

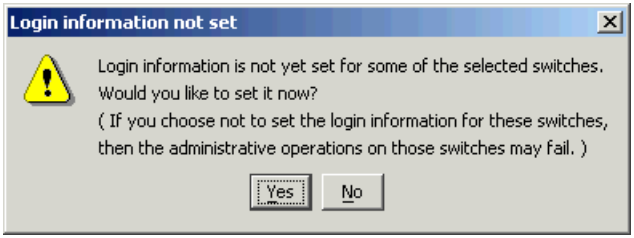


*Figure 1-211 Password error message*

We answer **Yes** to the message and re-define the passwords as defined in our enabling Security section. Once the passwords have been successfully learned, the Security Administration window opens, as shown in Figure 1-212.

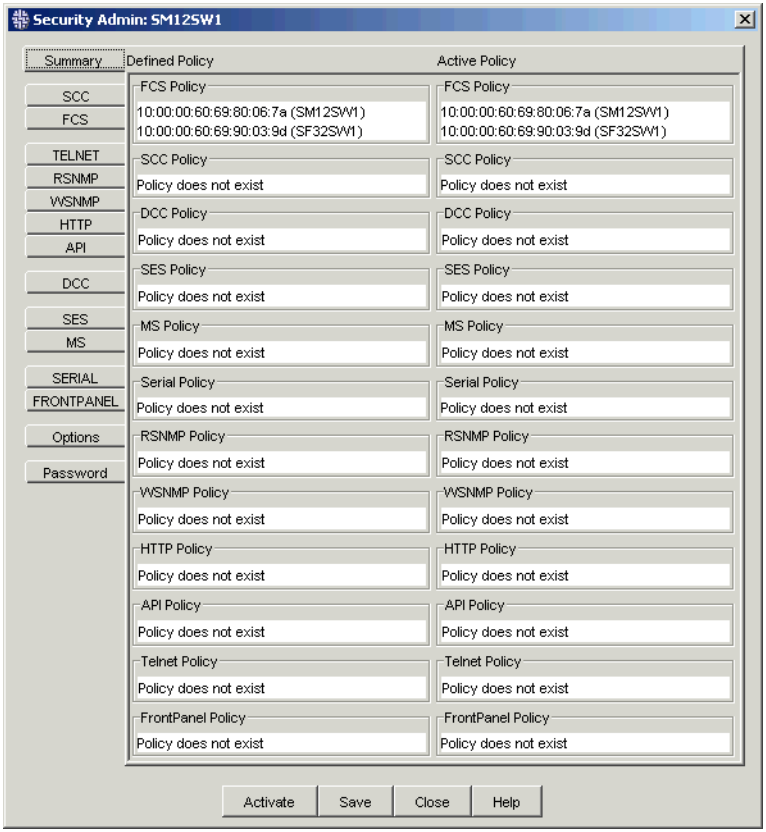


*Figure 1-212 Security Policy management*

From this window we can view the various security policies, and define them by clicking the appropriate tab on the left side of the window.

## 1.14 Interoperability

In the topics that follow, we describe:

- Switch interoperability
- Host operating system support
- HBA interoperability

### *Switch interoperability*

The IBM TotalStorage SAN Switches all support both 1 Gb/s and 2 Gb/s transmit and receive rates with auto-negotiation. The actual data signaling rate that is used on a port is automatically sensed and is set to the rate that is supported by a device or devices that are attached to the port. All IBM TotalStorage SAN Switches have been tested and are compliant with the current FC standards. They are compatible with most current-generation switches N_Ports, NL_Ports, and E_Ports, as well as host adapters, Redundant Array of Independent Disks (RAID) storage devices, hubs, and Fibre-SCSI bridge devices.

Fabric OS V4.x includes interoperability with the IBM SAN Fibre Channel Managed Hub, 35341RU and 2109 S series switches with a minimum firmware V2.6.0, and also the 3534-F08 and 2109 F16 series switches running level 3.0.1 or above firmware.

McDATA Directors are compatible with interoperability mode after release 3.2 of the McDATA EOS.

### *Implementation in existing environments*

Because the IBM TotalStorage SAN Switch F08, IBM TotalStorage SAN Switch F16, IBM TotalStorage SAN switch F32, and IBM TotalStorage SAN Switch M12 have a compatible 1 Gb/s auto-negotiated signaling rate on each port, they can be used as a replacement for the 3534-1RU and 2109-S08/16 series switches. As newer technology is added to existing systems that support 2 Gb/s signaling, the ports can accept these devices and interoperate with existing 1 Gb/s devices. If one of the models is connected to a third-party device but is unable to negotiate the signaling rate, it is also an option to manually set the speed of each port through the management interfaces.

### *Heterogeneous inter-switch operations*

IBM TotalStorage SAN Switch Fabric OS supports interoperability for the following functions:

- Basic switch functions:
  - Link initialization
  - Principal switch selection
  - Routing (FSPF)

- Basic services:
  - Simple name service
  - State change notification
  - WWN zoning (name server zoning)

The following facilities are switch-based facilities and will continue to function on any IBM TotalStorage SAN Switch:

- SNMP facilities
- Translative mode (private target support on fabrics)
- Enhanced performance metrics

The following facilities are IBM value-added facilities that are *not* supported in a multi-vendor fabric. Use of these facilities causes the Fabric to segment:

- QuickLoop zones
- QuickLoop Fabric assist mode
- Port, protocol, or LUN zoning
- Trunking

IBM is not aware of any areas of non-compliance with any ratified standards at this time.

### *Host Operating system support*

All versions of IBM TotalStorage SAN Switch Fabric OS have no specific Host OS dependencies. The Fabric OS in the switches allows for any Fibre Channel compliant device to attach to the switches as long as it conforms to the standards for device login, name service, and related Fibre Channel features. Regardless of the operating environment, a proper interface to the fabric requires a Fibre Channel HBA with a standards-compliant driver.

IBM provides support for environments using HBA, host operating system, and driver combinations which have been fully tested by IBM labs. For a complete list of the supported host operating system levels, refer to the current version of the Interoperability guide found on the IBM TotalStorage SAN Switch site:

- For the IBM TotalStorage SAN Switch M12:

  ftp://service.boulder.ibm.com/storage/san/2109m12/SM2109M12.pdf

- For the IBM TotalStorage SAN switch F32:

  ftp://service.boulder.ibm.com/storage/san/2109f32/SM2109F32.pdf

- For the IBM TotalStorage SAN Switch F16:

  ftp://service.boulder.ibm.com/storage/san/2109f16/SM2109F16.pdf

- For the IBM TotalStorage SAN Switch F08:

  ftp://service.boulder.ibm.com/storage/san/3534f08/SM3534F08.pdf

### *HBA Interoperability*

To check the HBA and driver versions that are supported by IBM on various Operating systems, please refer to:

http://knowledge.storage.ibm.com/HBA/HBASearchTool

## 1.15 Interoperability with IBM BladeCenter

In this section we describe the settings and limitations of adding an IBM BladeCenter™ to an IBM TotalStorage SAN Switch fabric. We will show the steps we used to add the BladeCenter to our single 2109-F16 fabric, although the steps are the same for any IBM TotalStorage SAN Switch fabric.

### Initial configuration on IBM BladeCenter Switch

For detail on how to connect to the IBM BladeCenter Switch, refer to 5.3, “BladeCenter management” on page 591.

First we verify the required configuration options listed in Table 1-24 to interconnect IBM BladeCenter with IBM TotalStorage SAN Switches.

For successful integration of the IBM BladeCenter with IBM 2109 fabric, all the features listed in Table 1-24 must be applied on both IBM BladeCenter and the IBM TotalStorage SAN Switch.

**Attention:** Changing the fabric configuration from native mode to an interop mode is a disruptive process and requires careful planning and implementation. The domain IDs allowed range on the BM BladeCenter switch and IBM 2109 series is between the range of 97 - 127. If any IBM TotalStorage SAN Switch is configured with a domain ID of 1 in native mode, then in interop mode the domain ID 1 becomes 97 automatically.

The fabric addresses of the end device will change due to the domain ID change of the switch. In an environment with AIX or HP hosts installed, the host adapter (FC devices) should be deleted and configured to re-discover the end devices (targets). Another implication of fabric migration from native to interop mode is that the multipathing and persistent binding functions also require the old definitions to be deleted and re-configured using the new Fibre Channel address.

*Table 1-24 Configuration options required for IBM BladeCenter and IBM 2109*

| Configuration features | IBM BladeCenter | IBM 2109 Series |
|---|---|---|
| Firmware level | 1.4.0.49 or later | F08/F16 - 3.0.2g or later<br>F32 - 4.0.2b or later<br>M12 - 4.0.0e or later |
| Domain ID lock | True | Static (Manual Configuration) |
| Domain IDs Allowed Range | 97 - 127 | 97-127 |
| E_D_TOV | 2000 milli seconds | 2000 milli seconds |
| R_A_TOV | 10000 milli seconds | 10000 milli seconds |
| IO Stream Guard | disabled on the ISL | Not Applicable |
| Core PID | Not Applicable | 1 |
| Principal Switch Priority | 254 | Not Applicable |
| Interop Mode | Not Applicable | Interop mode = 1 |
| Zoning | WWPN based | WWPN based |
| Default Zone | Disabled | Not Applicable |
| Switch Port Mode | GL Port < default> | E Port |
| Platform Management Services | Not Applicable | Disabled |

## Configuration limitations

Here we list the various configuration limitations:

- When merging IBM BladeCenter and 2109-XXX fabrics, a maximum of 31 interconnected switches per fabric are supported.
- Only WWPN based zoning is supported. Port based or WWNN based zoning is not supported in an interop fabric
- The domain ID allowed on IBM BladeCenter and the 2109 series switches is between 97 and 127.
- The following 2109 features are not supported in the interop mode:
  - QuickLoop
  - QuickLoop Fabric Assist
  - Remote Switch
  - Extended Fabric

- Trunking
- Advanced Performance Monitor
- Advanced Security
- Fabric Services
- Alias Server
- Management Server
- Platform Support
- Virtual Channels
- Broadcast Zones

## Important guidelines

These are some important guidelines:

- It is not recommended to use IBM BladeCenter SAN Utility and 2109 WEB TOOLS for zone configuration. Zone configuration may be performed from either the IBM BladeCenter SAN Utility or 2109 WEB TOOLS, but *not* both.
- We strongly recommended to upload the switch configuration before applying any configuration changes on the IBM 2109 and BladeCenter.

## Verify BladeCenter configuration

The output of the **show switch** command lists the active Switch WWN, Firmware Version, Domain ID, Operation State, Principal Switch Role, and Diagnostic Status, as shown in Example 1-1.

*Example 1-1 show switch command output*

```
show switch
  Switch Information
  ------------------
  SymbolicName                     FCSM
  SwitchWWN                        10:00:00:c0:dd:01:c3:0c
  SwitchType                       SwitchBlade
  PROMVersion                      V1.4.0.2-0 (Tue Jun  3 16:52:53 2003)
  CreditPool                       0
  DomainID                         125 (0x7d)
  FirstPortAddress                 7d0000
  FlashSize - MBytes               128
  LogLevel                         Critical
  MaxPorts                         16
  NumberOfResets                   39
  ReasonForLastReset               NormalReset
  SWImageVersion (1) - build date  V1.4.0.42-0 (Mon Feb  3 22:33:00 2003)
  SWImageVersion (2) - build date  V1.4.0.49-0 (Tue Jun  3 16:52:53 2003)
  ActiveConfiguration              default
  ActiveSWImage                    2
  AdminState                       Online
  AdminModeActive                  False
```

```
  BeaconOnStatus                    False
  OperationalState                  Online
  PrincipalSwitchRole               False
  BoardTemp (1) - Degrees Celsius   45
  BoardTemp (2) - Degrees Celsius   45
  SwitchDiagnosticsStatus           Passed
  SwitchTemperatureStatus           Normall
```

To verify the zone configuration, we use the **show config zoning** command as shown in Example 1-2.

*Example 1-2 zone configuration*

```
FCSM: admin> show config zoning
  Configuration Name: default
  Zoning Configuration Information
  AutoSave              True
  Default               All
```

Similarly, we use the **Show Config Switch** command to verify the domain ID Lock, TimeOut values (E_D_TOV, R_A_TOV) and Principal Switch Priority parameters as shown in Example 1-3.

*Example 1-3 BladeCenter switch configuration*

```
FCSM: USERID > show config switch
  Configuration Name: default
  Switch Configuration Information
  -------------------------------
  AdminState            Online
  BroadcastEnabled      True
  InbandEnabled         False
  DomainID              1 (0x1)
  DomainIDLock          False
  SymbolicName          FCSM
  R_T_TOV               100
  R_A_TOV               10000
  E_D_TOV               2000
  FS_TOV                5000
  DS_TOV                5000
  PrincipalPriority     254
  ConfigDescription     IBM BladeCenter(TM) 2-port Fibre Channel Switch Module
  ConfigLastSavedBy     Initial
  ConfigLastSavedOn     Initial
```

The I/O StreamGuard function must be disabled on the E_Port.

The I/O StreamGuard parameter, E_Port administrative / operational state, port speed, and port type can be verified by using the **show config port 15** command as shown in Example 1-4.

*Example 1-4 displaying E_Port configuration*

```
FCSM: admin> show config port 15
Configuration Name: default

  Port Number: 15
  ------------
  AdminState       Online
  LinkSpeed        Auto
  PortType         GL
  TLPortMode       TLTargetMode
  ISLSecurity      Any
  SymbolicName     Port15
  ALFairness       False
  ARB_FF           False
  InteropCredit    0
  ExtCredit        0
  FANEnable        True
  LCFEnable        False
  MFSEnable        True
  MFS_TOV          10
  MSEnable         True
  NoClose          False
  IOStreamGuard    Disabled
  VIEnable         False
  CheckAlps        False

FCSM: admin>
```

If any of the configuration options in the preceding examples did not match the required options listed in Table 1-24, then we use the **set config switch** command to configure the required settings, as shown in Example 1-5.

*Example 1-5 Set Config Switch Configuration menu*

```
FCSM: USERID> admin start
FCSM (admin): USERID> config edit
FCSM (admin-config): USERID> set config switch

A list of attributes with formatting and current values will follow. Enter a
new value or simply press the ENTER key to accept the current value. If you
```

```
wish to terminate this process before reaching the end of the list press 'q' or
'Q' and the ENTER key to do so.
AdminState         (1=Online, 2=Offline, 3=Diagnostics) [Online ]
  BroadcastEnabled   (True / False)  [True ]
  InbandEnabled      (True / False)     [True ]
  DefaultDomainID    (decimal value, 1-239) [1 ] 125
  DomainIDLock       (True / False)   [False ]  True
  SymbolicName       (string, max=32 chars)  [FCSM ]
  R_T_TOV            (decimal value, 1-1000 msec)  [100 ]
  R_A_TOV            (decimal value, 100-100000 msec) [10000 ]
  E_D_TOV            (decimal value, 10-20000 msec) [2000 ]
  FS_TOV             (decimal value, 100-100000 msec)  [5000 ]
  DS_TOV             (decimal value, 100-100000 msec)  [5000 ]
  PrincipalPriority  (decimal value, 1-255)  [254 ]  254
  ConfigDescription  (string, max=64 chars)  [IBM BladeCenter(TM) 2-port Fibre
Channel Switch  Module]
Finished configuring attributes.
  This configuration must be saved (see config save command) and activated (see
config activate command) before it can take effect.
  To discard this configuration use the config cancel command.
FCSM2 (admin-config): USERID> config save
FCSM (admin): USERID> config activate
The configuration will be activated.  Please confirm (y/n): [n] y
```

Once all the steps are complete on the IBM BladeCenter switch, we then proceed to configure our IBM 2109 F16 switch.

**Attention:** Changing the IBM TotalStorage SAN Switch series from native to interop mode is a disruptive process and requires reboot / fastboot after setting the domain ID, Timeout Values, and the Interop mode. It is strongly recommended to schedule a maintenance window in case of reverting to the old configuration, due to unexpected results.

We use an IBM 2109-F16 switch to demonstrate the procedure of implementing the Interop fabric with IBM BladeCenter. The same procedure also applies for IBM TotalStorage SAN Switch series.

## Configuring the IBM 2109-F16 using the CLI

Before making any configurations, we verify the current configuration, in case we should need to revert to the original configuration if the fabric migration from native to interop mode is unsuccessful.

We verify and configure the following parameters:

1. Firmware is at a correct level to interoperate with BladeCenter.
2. Configuration Parameters (Domain ID, E_D_TOV, R_A_TOV, core_PID)
3. InteropMode = 1
4. The Active/Defined Zone Configuration does not conflict with IBM BladeCenter definitions.
5. Platform Management services are disabled.

### *Step # 1*

We verify the firmware version using the **version** command as shown in Example 1-6.

*Example 1-6 displaying the switch firmware version*

```
IBM_2109_FC_Switch:admin> version
Kernel:     5.3.1
Fabric OS:  v3.0.2f
Made on:    Mon Jun 3 14:46:11 PDT 2002
Flash:      Mon Jun 3 14:48:15 PDT 2002
BootProm:   Tue Oct 30 10:24:38 PST 2001
IBM_2109_FC_Switch:admin>
```

As our firmware needs to be upgraded to 3.0.2g or later we perform the upgrade as described in “Upgrading switch firmware” on page 182.

### *Step # 2*

We verify the Active Zone Configuration using the **cfgshow** command as shown in Example 1-7.

*Example 1-7 Displaying zone configuration*

```
IBM_2109_FC_Switch:admin> cfgshow
Defined configuration:
cfg: BC_INTEROP_FABRIC
Win2K_HBA1_FastT200; Win2K_HBA2_FastT200
 zone:Win2K_HBA1_FastT200
21:01:00:e0:8b:28:37:3d; 20:07:00:a0:b8:07:54:49;
20:06:00:a0:b8:07:54:49
 zone: Win2K_HBA2_FastT200
21:00:00:e0:8b:08:37:3d; 20:07:00:a0:b8:07:54:49;
20:06:00:a0:b8:07:54:49

Effective configuration:
 no configuration in effect
IBM_2109_FC_Switch:admin>
```

We compare our output to the Bladecenter output in Example 1-2 to ensure there are no overlapping definitions.

The Zone configuration can be defined and activated using 2109 WEB TOOLS or CLI. For zoning procedures for the IBM TotalStorage SAN Switch, refer to 1.7.4, “Implementing zoning” on page 51.

### *Step # 3*

Configure the Interop mode by using the **interopmode 1** command as shown in Example 1-8.

*Example 1-8 Changing the switch from native to interop mode.*

```
IBM_2109_FC_Switch:admin> interopmode 1
Committing configuration...done.
interopMode is 1
NOTE: It is recommended that you boot this switch to make this change take
effect
IBM_2109_FC_Switch:admin> fastboot
Rebooting...
```

### *Step # 4*

After verifying that the correct firmware is active, we proceed to configure the switch into interopmode, and assign a unique domain ID between 97-127.

We need to disable the switch first to change the configuration parameters (domain ID and TimeOut Values). The command used to change the switch configuration parameters is **configure** as shown in Example 1-9.

*Example 1-9 switch configuration*

```
IBM_2109_FC_Switch:admin> switchdisable
IBM_2109_FC_Switch:admin> configure
Configure...

  Fabric parameters (yes, y, no, n): [no] y
    Domain: (1..239) [97] 98
    BB credit: (1..27) [16]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    WAN_TOV: (1000..120000) [0]
    WAN_RTT_DLY_MAX: (100..5000) [200]
    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]
    SYNC IO mode: (0..1) [0]
    VC Encoded Address Mode: (0..1) [0]
```

```
    Core Switch PID Format: (0..1) [1]
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
  Virtual Channel parameters (yes, y, no, n): [no]
  Zoning Operation parameters (yes, y, no, n): [no]
  RSCN Transmission Mode (yes, y, no, n): [no]
  NS Operation Parameters (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]
Committing configuration...done.
```

### *Step # 5*

We disable the Platform Management Services by using the **msplmgmtdeactivate** command, as shown in Example 1-10.

*Example 1-10 disabling platform management services*

```
IBM_2109_FC_Switch:admin> msplmgmtdeactivate
This will erase all Platform entries. Are you sure?  (yes, y, no, n): [no] y
Committing configuration...done.
Request Fabric to Deactivate Platform Management services....
Done.
```

### *Step # 6*

After the switch completes the reboot and comes online, we verify the fabric merge by using the **fabricshow** command on the 2109, as shown in Example 1-11.

*Example 1-11 Verifying the fabric merge from the 2109*

```
   IBM_2109_FC_Switch:admin> fabricshow
   Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
   -------------------------------------------------------------------------
    97: fffc62 10:00:00:60:69:51:3c:9d  9.42.164.20     0.0.0.0
   >"IBM_2109_F C_Switch"
   125: fffc7d 10:00:00:c0:dd:01:c3:0c  192.168.70.131  0.0.0.0         "FCSM"

   The Fabric has 2 switches
```

We also verify the fabric merge using the **show domains** command from the BladeCenter CLI, as shown in Example 1-12.

*Example 1-12 Verifying the fabric merge from the BladeCenter*

```
FCSM: admin> show domains

  Principal switch is (remote): 10:00:00:60:69:51:3c:9d
  Upstream Principal ISL is   : 15

Domain ID List:
        Domain 98 (0x62)  WWN = 10:00:00:60:69:51:3c:9d
        Domain 125 (0x7d)  WWN = 10:00:00:c0:dd:01:c3:0c

FCSM: admin>
```

Both these examples show that the fabric merge was successful, with the 2109 elected as principal switch and the IBM BladeCenter is the subordinate switch.

At this point the implementation of IBM BladeCenter and IBM TotalStorage SAN Switch interop fabric is complete and successful.

2

# Implementing a SAN with Cisco

In this chapter we introduce the new Cisco MDS 9000 family of Fibre Channel switches and directors. We describe the initial setup required to activate the Cisco Fabric Manager GUI, and describe how to configure the Cisco SAN with the GUI. We also introduce the Cisco Command Line Interface (CLI), and discuss the implications of Interoperability Mode on the Cisco fabric.

**Note:** We used Cisco Multilayer intelligent SAN operating system (SAN-OS) version 1.1(1a) for all our testing. If your SAN-OS level is different, some of the panels may not look the same. However, the concepts introduced here should still apply.

# 2.1 Introducing the Cisco MDS 9000 products

The Cisco MDS 9000 family is a growing Fibre Channel switch family, with several unique, advanced features, such as virtual storage area network (VSAN) support. The MDS 9000 family supports Fibre Channel connectivity at 1 Gb/s and 2 Gb/s, and can also support FCIP and iSCSI over gigabit ethernet.

The high availability features available on all models include redundant power supplies, redundant system clocks, non-disruptive restart of failed service processes, PortChannel for Fibre Channel, and both Port Channel and virtual routing redundancy protocol (VRRP) for IP connectivity. Some members of the MDS 9000 family also support advanced high availability features, such as dual supervisor modules and non-disruptive software upgrade capability.

## 2.1.1 MDS 9000 models

The Cisco MDS 9000 family currently consists of three members:

- Cisco MDS 9216 Multilayer Fabric Switch (IBM 2062-D01)
- Cisco MDS 9506 Multilayer Director (IBM 2062-D04 and 2064-T04)
- Cisco MDS 9509 Multilayer Director (IBM 2062-D07 and 2064-T07)

All of the switches use the same firmware and management tools. In addition, the line cards are designed to be interchangeable between the switches. The different models are described in the following sections.

### Cisco MDS 9216 Multilayer Fabric Switch (MDS 9216)

The MDS 9216 is an entry level member of the Cisco MDS 9000 family, providing a cost-effective SAN solution for small environments and remote locations not requiring the high-availability features of the larger models.

The MDS 9216 has the following features:

- 16 built-in Fibre Channel ports in the supervisor card
- One free slot for an additional line card
- 16-48 Fibre Channel ports
- Dual redundant power supplies
- Rack height 3U

The MDS 9216 is shown in Figure 2-1.

*Figure 2-1 Cisco MDS 9216 Multilayer Fabric Switch*

### *Power supplies for the MDS 9216*

The MDS 9216 has two 845W auto-ranging AC power supplies to the rear. The power input connectors are within the power supplies. The power supplies can accept any input voltage between 100 and 240 volts.

Some of the early MDS 9216 switches were delivered with a single power supply, however, that configuration is not available anymore.

## Cisco MDS 9506 Multilayer Director (MDS 9506)

The MDS 9506 is an enterprise level member of the Cisco MDS 9000 family, and has the following features:

- Four slots for line cards (at least one line card required)
- 16-128 Fibre Channel ports
- Dual redundant power supplies
- Dual redundant supervisor modules
- Support for non disruptive firmware upgrades
- Available with AC or DC input power supplies:
  - 110-220VAC: 2062-D04
  - 48-60VDC: 2062-T04
- Rack height 7U

**Note:** If you purchase the MDS 9506 through IBM, it always contains redundant power supplies and supervisor modules. This may not always be the case if you purchase the switch from another vendor. We recommend that you ensure that your configuration includes these redundant features.

The MDS 9506 is shown in Figure 2-2.

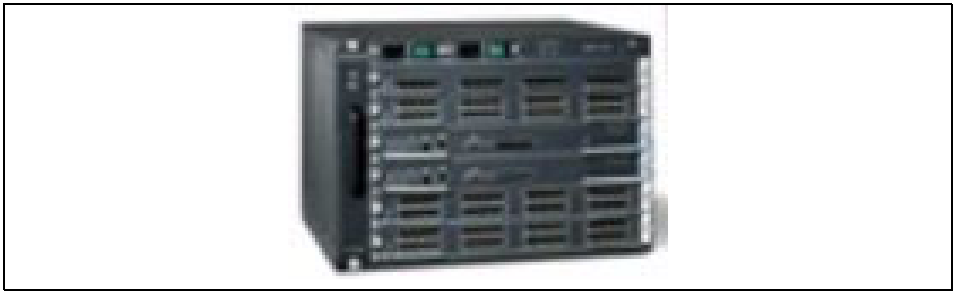

*Figure 2-2 Cisco MDS 9506 Multilayer Director*

### *Power supplies for the MDS 9506*

The MDS 9506 has two 1900W AC or DC power supplies to the rear. The power input connectors are in the Power Entry Module (PEM) to the front. The AC power supplies are capable of working with either 110V or 220V, but the usable power depends on the input voltage. With 110V AC input voltage, each power supply can only provide 1100W.

The power supplies can be run in two modes; redundant mode or combined mode. In redundant mode the usable power is the output power of the smallest installed power supply. In combined mode, the usable power is twice the output power of the smaller installed power supply. Since combined mode reduces system availability, we recommend only using redundant mode.

**Note:** When used with 110V AC input power, the AC power supplies of the MDS 9506 can only supply 1100W of power each. In the case of power supply or power circuit failure, the remaining power supply is not adequate enough to run a fully occupied MDS 9506. Therefore we strongly recommend that you use 220V power circuits for the MDS 9506.

## Cisco MDS 9509 Multilayer Director (MDS 9509)

The MDS 9509 is an enterprise level member of the Cisco MDS 9000 family for applications requiring a higher port count that can be achieved with the MDS 9506, and has the following features:

- Seven slots for line cards, at least two line cards required
- 32-224 Fibre Channel ports
- Dual redundant power supplies
- Dual redundant supervisor modules
- Support for non-disruptive firmware upgrades
- Available with AC or DC input power supplies
  - 110-220VAC: 2062-D07
  - 48-60VDC: 2062-T07
- Rack height 14U

**Note:** If you purchase the MDS 9509 through IBM, it always contains redundant power supplies and supervisor modules. This may not always be the case if you purchase the switch from another vendor. We recommend that you ensure that your configuration includes these redundant features.

The MDS 9509 is shown in Figure 2-3.

*Figure 2-3 Cisco MDS 9509 Multilayer Director*

### *Power supplies for the MDS 9509*

The MDS 9509, as purchased from IBM, has two 2500W AC or DC power supplies to the front. The power input connectors are directly in the power supplies. The AC power supplies are capable of working with either 110V or 220V, but the usable power depends on the input voltage. With 110VAC input voltage, each power supply can only provide 1300W.

The power supplies can be run in two modes; redundant mode or combined mode. In redundant mode, the usable power is the output power of the smallest installed power supply. In combined mode, the usable power is twice the output power of the smallest installed power supply. Since combined mode reduces system availability, we recommend only using redundant mode.

**Note:** When used with 110V AC input power, the AC power supplies of the MDS 9509 can only supply 1300W of power each. In the case of power supply or power circuit failure, the remaining power supply cannot supply adequate power for a fully occupied MDS 9509. Therefore we strongly recommend that you use 220V power circuits for the MDS 9509.

Cisco also sells 4000W AC power supplies for the MDS 9509. The 4000W AC power supplies support 220V AC input only.

## 2.1.2 Supervisor modules

The MDS 9216 contains a single, built-in supervisor module, with 16 integrated Fibre Channel ports. It also contains an interface module for local and remote management interfaces.

The MDS 9506 and MDS 9509 contain one or two supervisor modules. Each supervisor module contains the local and remote management interfaces. IBM always delivers these storage switches with two supervisor modules.

### The built-in supervisor and interface modules of MDS 9216

The supervisor module of the MDS 9216 contains 16 Fibre Channel ports, a crossbar switch, and the management functions of the switch, and can be only installed into slot number one of the MDS 9216.

Since the MDS 9216 supports only one supervisor module, supervisor failover is not supported. There is only a limited non-disruptive software upgrade capability.

The interface module, located above the supervisor module in slot number one, contains the following interfaces:

- Console for serial connection via RJ-45 port
- 10/100 Mgmt for ethernet management port (mgmt0)
- Auxiliary COM2 port for connecting an external serial communications device

### The supervisor modules of MDS 9506 and MDS 9509

The supervisor modules of the MDS 9506 and MDS 9509 provide the management interfaces for the switch. Each module also contains a 720 Gb/s capable crossbar switch, that is used to route the traffic between the ports in the switch. The supervisor modules are installed in slots five and six of the switch.

When both supervisor modules are installed, one of them is active and the other works as a hot-standby. Synchronization is managed between the active and standby supervisor, and in the case of supervisor failover, the traffic across the switch is unaffected.

In addition to the internal Flash memory, the supervisor modules also support an optional CompactFlash card. This card can be used to store additional software images, configuration files, logs, and core dumps.

**Note:** If you purchase the MDS 9506 or MDS 9509 through IBM, it always contains two supervisor modules. This may not always be the case if you purchase the switch from another vendor. We recommend that you ensure that your configuration includes two supervisor modules.

Each supervisor module contains the following interfaces:

- Console for serial connection via RJ-45 port
- 10/100 Mgmt for ethernet management port (mgmt0)
- Auxiliary COM2 port for connecting an external serial communications device
- CompactFlash slot for the optional CompactFlash card

The supervisor module is shown in Figure 2-4.

*Figure 2-4 Supervisor module*

## 2.1.3 Line cards and transceivers

The Cisco MDS 9000 family currently has three different line cards:

- 16-port Fibre Channel Line Card (feature code (f/c) 2116)
- 32-port Fibre Channel Line Card (f/c 2132)
- 8-port IP Line Card for iSCSI and FCIP (f/c 2208)

By combining the different line cards in a single, modular chassis, you can configure an optimized SAN for devices with different protocols and performance requirements.

The line card features do not include any optical SFP transceivers. An adequate number of SFP transceivers must be included in the configuration. The configuration must include one SFP transceiver for each Fibre Channel port, and at least one SFP transceiver for each IP line card.

The different line cards and transceivers are described below.

### 16-port Fibre Channel Line Card (f/c 2116)

The 16-port Fibre Channel Line Card supports a fully-sustained, full-duplex 2 Gb/s data rate for each port, with an aggregate total bandwidth of 64 Gb/s. It also provides up to 255 buffer credits for each port. Therefore it is best suited for applications requiring very high bandwidth, such as ISL connections, high-end servers, and fast storage devices.

The 16-port Fibre Channel Line Card is shown in Figure 2-5.

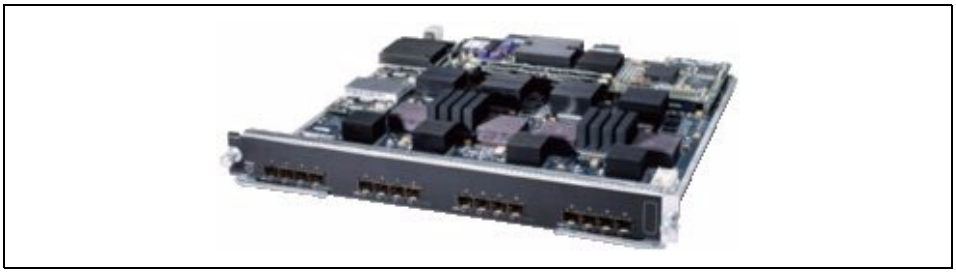

*Figure 2-5 16-port Fibre Channel Line Card*

## 32-port Fibre Channel Line Card (f/c 2132)

The 32-port Fibre Channel Line Card is internally divided into eight four-port groups. The four ports in each group share access to 2.65 Gb/s of internal bandwidth. Each port is capable of working at 1 Gb/s or 2 Gb/s data rate, and has 12 buffer credits.

Since very few open systems servers are able to continuously utilize the bandwidth of a Fibre Channel interface, the 32-port line card provides a cost-effective solution to attaching a large number of servers to the Fibre Channel fabric. You can even consider the cards to be edge switches in a core-edge fabric contained in a single box — without the need to allocate Fibre Channel ports for ISLs or manage multiple, separate switches.

Since ISLs require continuous bandwidth, only the first port of each group of four ports can be configured as an ISL. This also causes the other three ports in the same group to become disabled. Therefore we do not recommend using the ports in the 32-port line card for ISLs.

The 32-port Fibre Channel Line Card is shown in Figure 2-6.

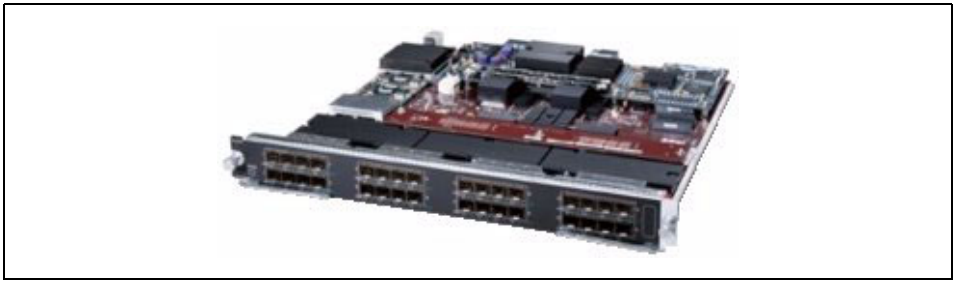

*Figure 2-6 32-port Fibre Channel Line Card*

## 8-port IP Line Card (f/c 2208)

The 8-port IP Line Card provides eight gigabit ethernet ports that can support iSCSI and FCIP protocols simultaneously. Since the bit rate of gigabit ethernet is different from the bit rate of Fibre Channel, the card only supports the new Tri-rate SFPs.

The 8-port IP Line card is shown in Figure 2-7.

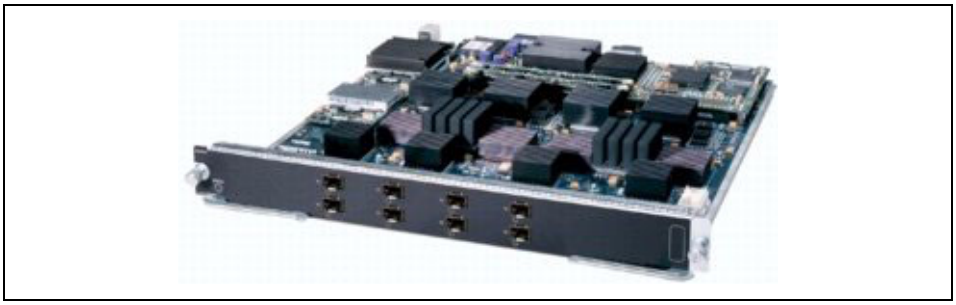

*Figure 2-7 8-port IP Line Card*

### *Ports configured to run FCIP*

The ports configured for FCIP can support up to three virtual ISL connections (FCIP tunnels). This way you can transport the Fibre Channel traffic transparently over an IP network between two FCIP capable switches. Each virtual ISL connection acts as a normal Fibre Channel ISL or EISL.

To use FCIP, you need to purchase the “FCIP Activation for 8-port IP Services Line Card” feature for every 8-port IP Line Card that needs to support FCIP. The feature codes are f/c 2209 for the MDS 9216, and f/c 2210 for the MDS 9506 and MDS 9509.

### *Ports configured to run iSCSI*

Ports configured to run iSCSI work as a gateway between iSCSI hosts and Fibre Channel attached targets. The module terminates iSCSI commands and issues new Fibre Channel commands to the targets.

## SFP transceivers

There are currently four SFP transceivers available for the MDS 9000 family:

- Tri-rate SW SFP transceiver (f/c 5210)
- Tri-rate LW SFP transceiver (f/c 5220)
- FC port SW SFP transceiver (f/c 5230)
- FC port LW SFP transceiver (f/c 5240)

The Tri-rate transceivers can support the bit rates for both Fibre Channel and gigabit ethernet, and are therefore the only transceivers supported by the 8-port IP Line Card. The older FC port transceivers only support the bit rates for Fibre Channel.

In addition, special transceivers are available to support a coarse wavelength-division multiplexing (CWDM) extended distance solution (2062-CW1), that allows the transmission of up to eight Fibre Channel or IP traffic streams over a single physical fiber, without using a separate dense wavelength-division multiplexing (DWDM) system.

## 2.1.4 Diagnostic tools

The Cisco MDS 9000 series switches allow non-disruptive monitoring of the Fibre Channel traffic with the Switched Port Analyzer (SPAN) feature. This feature can be used to copy all Fibre Channel frames from one or more Fibre Channel ports to the Span Destination (SD) port.

The Cisco MDS 9000 Port Analyzer Adapter supplements this feature by encapsulating the Fibre Channel frames to standard ethernet frames over a 1000base-T ethernet interface. The frames can then be analyzed using free network protocol analyzer software, such as Ethereal. This allows cost-effective monitoring of the Fibre Channel traffic.

The Cisco MDS 9000 Port Analyzer Adapter is shown in Figure 2-8.


*Figure 2-8 Cisco MDS 9000 Port Analyzer Adapter*

Since the actual data within the Fibre Channel frame is often not relevant for protocol analysis, the Cisco MDS 9000 Port Analyzer Adapter can also truncate the Fibre Channel frames. The truncate modes available are described in Table 2-1. The modes can be configured with a four-position DIP switch located in the rear of the adapter.

*Table 2-1 Cisco MDS 9000 Port Analyzer Adapter truncate modes*

| Mode | Description |
|---|---|
| No truncate | Fibre Channel frames are passed without any modification to the payload. |
| Ethernet truncate | Fibre Channel frames are truncated to 1496 bytes to fall within the maximum ethernet frame size. |
| Shallow truncate | Fibre Channel frames are truncated, if the payload of the frame is more than 256 bytes. |
| Deep truncate | Fibre Channel frames are truncated, if the payload of the frame is more than 64 bytes. |
| Management | Only fixed 288 byte ethernet frames that contain internal debug information are transmitted. |

## 2.1.5 Port addressing and port modes

The Fibre Channel ports in the Cisco MDS 9000 family are addressed with addresses in the form of f/c<slot>/<port>, where <slot> is the slot number of the line card (1-9), and <port> is the port number on the line card (1-32). For example, the first port of the line card in slot 1 is f/c1/1, and the seventh port of the line card in slot 3 is f/c3/7.

### Fibre Channel IDs and Persistent FC IDs

Contrary to other switch manufacturers, there is no direct correlation between physical Fibre Channel ports and Fibre Channel IDs (FC ID). This is necessary to allow intermixing line cards with different number of ports, while being able to utilize all port addresses, to allow both fabric and loop devices to coexist, and also to allow switches larger than 256 ports in the future.

The following rules apply to the FC ID assignment for any VSAN:

- When an N_Port or NL_Port logs into the switch it is assigned a FC ID.
- N_Ports receive the same FC ID if disconnected and reconnected to any port within the same switch.
- NL_Ports receive the same FC ID only if reconnected to the same port within the same switch, where the port was originally connected.

If the Persistent FoIds feature is not enabled for a VSAN, these rules apply:

- The WWN of the N_Port or NL_Port and the assigned FC ID are stored in a volatile cache, and are not saved across switch reboots.
- The switch preserves the binding of FC ID to WWN on a best-effort basis.

- The volatile cache has room for a maximum of 4000 entries and if the cache gets full, oldest entries are overwritten.

If the Persistent FoIds feature is enabled for a VSAN, the following rules apply:

- The FC ID to WWN mapping of the WWNs currently in use is stored to a nonvolatile database, and is saved across reboots.
- The FCID to WWN mapping of any new device connected to the switch is automatically stored into the non-volatile database.
- You can also manually configure the FCID to WWN mappings if necessary.

**Note:** If you attach AIX or HP-UX hosts to a VSAN, you need to have persistent FC IDs enabled for that VSAN. This is because these operating systems use the FC IDs in device addressing. If the FC ID of a device changes, the operating system considers it to be a new device, and gives it a new name.

## Port modes

The Fibre Channel ports in the Cisco MDS 9000 family can operate in several modes. The operational modes are described in Table 2-2.

*Table 2-2 Fibre Channel port operational modes*

| Mode | Description |
|---|---|
| E_Port | An expansion port (E_Port) interconnects two Fibre Channel switches, forming an ISL between an E_Port in each switch. The ISL belongs to a single VSAN, and can also be connected to third-party switches. |
| F_Port | A fabric port (F_Port) connects the switch to a N_Port in a host or storage device using a point-to-point link. Only one N_Port can connect to the F_Port. |
| FL_Port | A fabric loop port (FL_Port) connects the switch to a public FC-AL loop. Only one FL_Port can be operational in a single FC-AL loop at any given time. |
| TE_Port | A trunking E_Port (TE_Port) interconnects two Fibre Channel switches, forming an extended ISL (EISL) between a TE_Port in each switch. The EISL can multiplex the traffic of several VSANs. However, the EISL is currently only available in the Cisco MDS 9000 family of switches. |
| TL_Port | A translative loop port (TL_Port) connects the switch to a private FC-AL loop. |

| Mode | Description |
|---|---|
| SD_Port | A SPAN destination port (SD_Port) acts as a snooper port, allowing the monitoring of the switch traffic with a standard Fibre Channel analyzer. |
| B_Port | A bridge port (B_Port) is used to connect some SAN extender devices to the switch, instead of E_Port. |
| Fx_Port | A Fx_Port can operate as either F_Port or FL_Port, depending on the device connected to it. The port mode is determined during interface initialization. |
| Auto | A port configured as auto can operate as E_Port, F_Port, FL_Port, or TE_Port, depending on the device connected to it. The port mode is determined during interface initialization. |

## 2.1.6 Zoning

The Cisco MDS 9000 family zoning can be administrated from any switch in the fabric, and all changes are automatically distributed to all of the switches.

The Cisco MDS 9000 family supports zoning by the following criteria:

- Port world wide name (pWWN) - the WWN of the Nx_Port (device) attached to the switch
- Fabric pWWN (fWWN) - the WWN of the fabric port (port-based zoning)
- FC ID — the FC ID of the N_Port attached to the switch

To make management of zoning easier, the Cisco MDS 9000 family supports alias names for all of the elements above.

The Cisco MDS 9000 family supports a default zone. All ports and WWNs not assigned to any zone belong to the default zone. If zoning is not activated, all devices belong to the default zone. You can control access between default zone members by default zone policy.

The Cisco MDS 9000 family supports both soft and hard zoning.

### Soft zoning

In soft zoning, zoning restrictions are applied during the interaction between the name server and the end device. If an end device somehow manages to know or guess the FC ID of another end device, it can access that device.

### Hard zoning

In hard zoning, the zoning is enforced for each frame sent by an Nx_Port as the frame enters the switch. This prevents any unauthorized access at all times. The enforcement is done by the switch hardware at wire speed.

## 2.1.7 VSAN

VSAN is a unique feature of Cisco MDS 9000 series that enables dividing the physical Fibre Channel fabric to virtual SAN fabrics. Each VSAN is a completely separate SAN fabric, with its own set of domain IDs, fabric services, zones, namespace, and interoperability mode.

Each port in the switch fabric belongs to exactly one of the VSANs at any given time, with the exception of trunking E_Ports (TE_Ports) that can multiplex the traffic of several VSANs over a single physical link.

Up to 256 VSANs can be configured in a single switch. The VSAN numbers can range from 1 to 4094. VSAN number 1 is called the default VSAN, and is the VSAN that initially contains all of the ports in the switch. If you do not need to divide the fabric into VSANs, you can leave all ports in the default VSAN.

The VSAN number 4094 is called the isolated VSAN, and any port configured into that VSAN is isolated from all other ports. If you delete a VSAN, all ports in it are moved to the isolated VSAN to avoid implicit transfer of the ports to the default VSAN.

## 2.1.8 Trunking and PortChannel

In Cisco terminology, the term trunking is used to describe a single trunking E_Port (TE_Port) with the remit to multiplex the traffic of more than one VSAN. This is in contrast to other Fibre Channel switch manufacturers who use that term (trunking) to describe the aggregation of several physical interfaces into a single logical interface. Cisco calls this feature PortChannel. We are using the Cisco terminology to describe the features.

Both trunking PortChannel features are available for both Fibre Channel and gigabit ethernet interfaces on the Cisco MDS 9000 family. Since the configuration rules for these features are different, we describe both of them separately.

### FC Trunking

Trunking, also known as VSAN trunking, enables interconnect ports to transmit and receive frames in more than one VSAN over the same physical link. In this case the link is configured as an extended ISL (EISL) link using the EISL frame format.

Trunking is only applicable to E_Ports and used for inter-switch connections. Trunking is normally enabled for all ports in the switch but can be disabled on a port-by-port basis. If the port becomes operational as a trunking E_Port, it is referred to as a TE_Port. If a port, with trunking enabled, is connected to a third-party switch, it works as a normal E_Port.

### FC PortChannel

The PortChannel feature can be used to aggregate up to 16 ISL or EISL links into a single logical link. The Fibre Channel ports can be any Fibre Channel ports in any 16-port Fibre Channel line card.

The PortChannel feature increases the available aggregate bandwidth of the logical link since the traffic is distributed among all functional links in the channel. It also provides high availability, since the channel remains active as long as at least one of the links forming it remains active, and the traffic is transparently distributed over the remaining links.

Since PortChannel can be built on EISL links, both trunking and PortChannel are supported simultaneously.

### Ethernet Trunking

The gigabit ethernet ports support VLANs, and can be configured as trunking ports according to the 802.1Q standard. Each physical gigabit ethernet port can have multiple logical VLAN interfaces, as long as all of the logical interfaces are in separate subnets.

### Ethernet PortChannel

The ethernet PortChannel feature can be used to aggregate a pair of gigabit ethernet ports on a 8-port IP line card into a single logical interface. The allowed combinations of ports for PortChannel are 1-2, 3-4, 5-6, and 7-8.

## 2.1.9 iSCSI and FCIP support

The Cisco MDS 9000 series supports both iSCSI and FCIP on the 8-port IP line card simultaneously.

### iSCSI

The iSCSI support is used to connect iSCSI capable hosts to Fibre Channel storage devices. Support for iSCSI is included in the price of the 8-port IP line card.

### FCIP

The FCIP support is used to connect separate SAN islands over an IP network. Each defined connection is a virtual E_Port (VE_Port), and can work as an E_Port or a TE_Port. Each gigabit ethernet interface can support up to three FCIP connections.

To use FCIP, you need to purchase the “FCIP Activation for 8-port IP Services Line Card feature for every installed 8-port IP line card in the switch. The feature codes are f/c 2209 for the MDS 9216, and f/c 2210 for the MDS 9506 and MDS 9509.

## 2.1.10 Features and ordering

The Cisco MDS 9000 family devices can be configured with the normal IBM econfig tool. The terminology used in `econfig` reports is used here.

In addition to the features listed below, fiber optic cables can also be included into the configuration. However, since we do not consider them to be part of the switch, we do not list those features.

### Cisco MDS 9216 Multilayer Fabric Switch

These are the feature codes available for the Cisco MDS 9216 Multilayer Fabric Switch:

- Base switch, with 16 Fibre Channel ports:
  - 2062-D01 CISCO STORAGE SWITCH MDS 9216
- Optionally, one of the following line cards:
  - f/c 2116 16 PORT FIBRE CHAN LINE CARD
  - f/c 2132 32 PORT FIBRE CHAN LINE CARD
  - f/c 2208 8-Port IP Line Card for iSCSI and FCIP
- Optionally, FCIP activation feature for the 8 port IP line card:
  - f/c 2209 FCIP activation for 8-port IP Services Line Card
- As many SFP transceivers, as there are installed ports (16-48):
  - f/c 5210 Tri-rate SW SFP Transceiver
  - f/c 5220 Tri-rate LW SFP Transceiver
  - f/c 5230 FC PORT SW SFP TRANSCEIVER
  - f/c 5240 FC PORT LW SFP TRANSCEIVER
- One of the following power cord features, depending on your geography:
  - f/c 9110 AC Power Cords, North America
  - f/c 9111 AC Power Cords, Australia
  - f/c 9112 AC Power Cords, Europe

- f/c 9113 AC Power Cords, Italy
- f/c 9114 AC Power Cords, United Kingdom
- f/c 9115 AC Power Cords, Argentina

### Cisco MDS 9506 Multilayer Director

These are the feature codes available for the Cisco MDS 9506 Multilayer Director:

- Base switch, one of the following, depending on power supply:
  - 2062-D04 CISCO STORAGE SWITCH MDS 9506
  - 2062-T04 CISCO STORAGE SWITCH MDS 9506
- Exactly two supervisor cards:
  - 2010 Supervisor Line Card
- Up to four line cards from the following list (at least one Fibre Channel line card):
  - f/c 2116 16 PORT FIBRE CHAN LINE CARD
  - f/c 2132 32 PORT FIBRE CHAN LINE CARD
  - f/c 2208 8-Port IP Line Card for iSCSI and FCIP
- Optionally, FCIP activation features for the 8 port IP line cards:
  - f/c 2210 FCIP activation for 8-port IP Services Line Card
- Optionally, up to two flash memory cards:
  - f/c 5810 Cisco Flash Memory Card
- As many SFP transceivers, as there are installed ports (32-128):
  - f/c 5210 Tri-rate SW SFP Transceiver
  - f/c 5220 Tri-rate LW SFP Transceiver
  - f/c 5230 FC PORT SW SFP TRANSCEIVER
  - f/c 5240 FC PORT LW SFP TRANSCEIVER
- For the 2062-D04 only, one of the following power cord features, depending on your geography and environment:
  - f/c 9221 Power Cords, US, 250Vac 20A, Right Angle C19, NEMA 6-20 Plug
  - f/c 9222 Power Cords, US, 250Vac 20A, Right Angle C19, NEMA L6-20 Plug
  - f/c 9223 Power Cords, EU, 250VAC 16A, Right Angle C19, CEE 7/7 Plug
  - f/c 9224 Power Cords, Intl, 250Vac 20A, Right Angle C19, IEC 309 Plug

### Cisco MDS 9509

These are the feature codes available for the Cisco MDS 9509 Multilayer Director:

- Base switch, one of the following, depending on power supply:
  - 2062-D07 CISCO STORAGE SWITCH MDS 9509
  - 2062-T07 CISCO STORAGE SWITCH MDS 9509
- Up to seven line cards from the following list (at least two Fibre Channel line cards):
  - f/c 2116 16 PORT FIBRE CHAN LINE CARD
  - f/c 2132 32 PORT FIBRE CHAN LINE CARD
  - f/c 2208 8-Port IP Line Card for iSCSI and FCIP
- Optionally, FCIP activation features for the 8 port IP line cards:
  - f/c 2210 FCIP activation for 8-port IP Services Line Card
- Optionally, up to two flash memory cards:
  - f/c 5810 Cisco Flash Memory Card
- As many SFP transceivers, as there are installed ports (32-224):
  - f/c 5210 Tri-rate SW SFP Transceiver
  - f/c 5220 Tri-rate LW SFP Transceiver
  - f/c 5230 FC PORT SW SFP TRANSCEIVER
  - f/c 5240 FC PORT LW SFP TRANSCEIVER
- For the 2062-D07 only, one of the following power cord features, depending on your geography and environment:
  - f/c 9210 Power Cords, US, 250Vac 16A NEMA 6-20
  - f/c 9211 Power Cords, Europe, 16A
  - f/c 9212 Power Cords, International, 250Vac 16A
  - f/c 9213 Power Cords, US, 250Vac 16A Twist NEMA L6-20

## 2.2 Initial setup of the Cisco MDS 9000 products

Before you can manage the Cisco MDS 9000 series switch through the network, you have to set up the TCP/IP parameters for the switch.

The first time the switch is powered on, it automatically runs the setup program, and prompts you for the IP address and other configuration information necessary to communicate over the management ethernet interface. You can also start the setup program with the **`setup`** command later if necessary.

### 2.2.1 Preparing to configure the switch

Before you configure the switch for the first time, you need to gather the following information:

- New administrator password
- Switch name
- IP address for the management ethernet
- Subnet mask for the management ethernet
- Default gateway IP address (optional)
- DNS server IP address (optional)
- NTP server IP address (optional)
- SNMP v3 secret key (optional)

### 2.2.2 Connecting to the switch via the serial port

1. Connect the serial cable provided with the switch to the RJ-45 socket in the switch, using the console port in these modules:
   - Interface module in MDS 9216
   - Supervisor module in slot 5 in MDS 9506 and MDS 9509
2. Connect the other end of the serial cable to an RS-232 serial port on the workstation.
3. Disable any serial communication programs running on the workstation.
4. Open a terminal emulation application (such as HyperTerminal on a PC), and configure it as follows:
   - Bits per second: 9600
   - Data bits: 8
   - Parity: none
   - Stop bits: 1
   - Flow control: none

   An example of the HyperTerminal serial port properties window is shown in Figure 2-9.

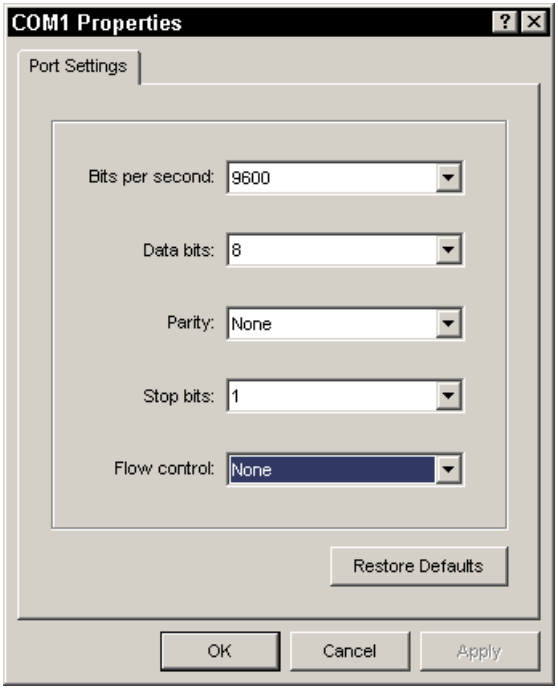


*Figure 2-9 HyperTerminal serial port properties window*

## 2.2.3 Setting up the initial parameters with the setup program

We will assume that you are already connected to the console serial port of the switch, but that the switch is still powered off.

Note that the steps you have to take may be different depending on which features you want to activate. However, the prompts of the setup program should be self-explanatory.

1. Power on the switch:

```
...
---- Basic System Configuration Dialog ----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime
to skip away remaining dialogs.
```

2. Enter **yes** to enter setup mode:

```
Would you like to enter the basic configuration dialog (yes/no): yes
```

3. Enter the new password for the administrator (user admin):

```
Enter the password for "admin" : newpass
```

4. Enter **no** to not create additional accounts at this time:

```
Create another login account (yes/no) [n]: no
```

5. Enter **yes** to create a SNMPv3 account for Fabric Manager:

```
Configure SNMPv3 Management parameters (yes/no) [y]: yes
```

   a. Enter the user name:

```
SNMPv3 user name [admin]: snmpuser
```

   b. Enter the SNMPv3 password (eight characters minimum):

```
SNMPv3 user authentication password : snmppass
The same password will be used for SNMPv3 privacy as well.
```

6. Enter **no**, to not create read-only SNMP community string:

```
Configure read-only SNMP community string (yes/no) [n]: no
```

7. Enter **no**, to not create SNMP community string:

```
Configure read-write SNMP community string (yes/no) [n]: no
```

8. Enter a name for the switch:

```
Enter the switch name : 9216
```

9. Enter **yes** to configure out-of-band management (the ethernet management interface mgmt0):

```
Continue with Out-of-band (mgmt0) management configuration? (yes/no)
[y]: yes
```

   a. Enter the mgmt0 IP address:

```
Mgmt0 IP address : 9.42.164.78
```

   b. Enter the mgmt0 subnet mask:

```
Mgmt0 IP netmask : 255.255.255.0
```

10. Enter **no** to not configure in-band management at this time:

```
Continue with In-band (vsan1) management configuration? (yes/no) [n]: no
```

11. Enter **yes** to enable the IP routing and default gateway:

```
Enable the ip routing capabilities? (yes/no) [y]: yes
```

   a. Enter **no** to not configure a static route:

```
Configure static route? (yes/no) [y]:no
```

   b. Enter **no** to not configure the default network:

```
Configure the default network? (yes/no) [y]: no
```

c. Enter **yes**, if you want to configure the default gateway:

```
Configure the default gateway? (yes/no) [y]: yes
```

i. Enter the default gateway IP address:

```
IP address of the default gateway : 9.42.164.1
```

12. Enter **no**, to not configure a DNS server:

```
Configure the DNS IP address? (yes/no) [n]: no
```

13. Enter **no**, to not configure the default domain name (for DNS):

```
Configure the default domain name? (yes/no) [n]: no
```

14. Enter **yes** to enable the telnet service:

```
Enable the telnet service? (yes/no) [y]: yes
```

15. Enter **yes** to enable the SSH service:

```
Enable the ssh service? (yes/no) [n]: yes
```

a. Enter the SSH key type (dsa, rsa or rsa1):

```
Type of ssh key you would like to generate (dsa/rsa/rsa1) : dsa
```

b. Enter the number of bits for the SSH key (512-2048):

```
Number of  key bits <768-2048> : 1024
```

16. Enter **no**, to not configure a NTP server (time server):

```
Configure the ntp server? (yes/no) [n]: no
```

17. Enter **shut** to configure the default switchport interface to the shut state:

```
Configure default switchport interface state (shut/noshut) [shut]: shut
```

18. Enter **on** to configure the default switchport trunk mode:

```
Configure default switchport trunk mode (on/off/auto) [on]: on
```

19. Enter **deny**, to deny traffic across the default zone:

```
Configure default zone policy (permit/deny) [deny]: deny
```

20. Review the configuration that you have just entered:

```
The following configuration will be applied:
  username admin password admin role network-admin
  snmp-server user Falkon network-admin auth md5 password priv password
  switchname 9216
  interface mgmt0
    ip address 9.42.164.78 255.255.255.0
    no shutdown
  ip default-gateway 9.42.164.1
  telnet server enable
  ssh key dsa 1024 force
  ssh server enable
```

```
system default switchport shutdown
system default switchport trunk mode on
no zone default-zone permit vsan 1-4093
```

21. Enter **no**, if you are pleased with the configuration; (otherwise, enter **yes** and go back to step 2):

```
Would you like to edit the configuration? (yes/no) [n]: no
```

22. Enter **yes** to save the configuration:

```
Use this configuration and save it? (yes/no) [y]: yes
```

**Note:** If you do not save the configuration at this point, none of your changes are updated the next time the switch is rebooted. Ensure that you type **yes** here to save the new configuration.

23. Wait until the configuration has been saved, and you get the command prompt:

```
####################################] 100%



Exiting the basic config setup.


9216#
```

24. Your basic configuration is now finished, and you can proceed to install the Cisco Fabric Manager and Device Manager.

## 2.2.4 Installing the Cisco Fabric Manager and Device Manager

To install the Cisco Fabric Manager, you should already have Java Runtime Environment (JRE) 1.4, or later, and Java Web Start installed.

1. Start your Web browser and open the Web page from the IP address of your switch. You should see the Cisco Fabric Manager installation page, as shown in Figure 2-10.


*Figure 2-10 The install page of Cisco software*

If you get any error messages during this phase, you may not have the prerequisite software installed. For example, if you don't have the Java Web Start installed, you will get an error message similar to that shown in Figure 2-11.

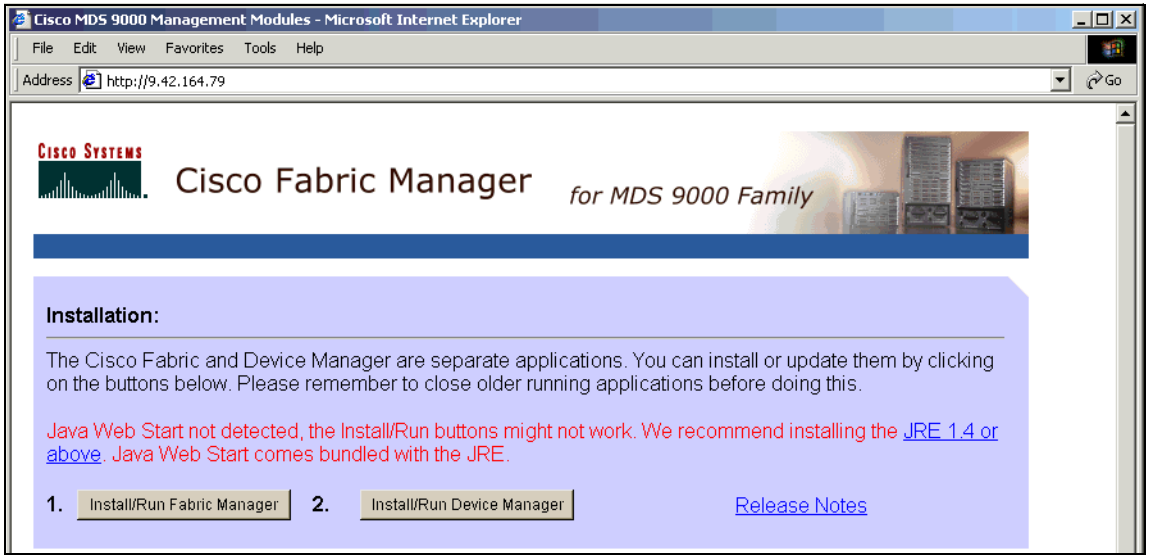


*Figure 2-11 The install page of Cisco software - no Java Web Start installed*

We recommend that you install the JRE and Java Web Start using the link provided by the installation page.

2. Start the installation of Fabric Manager from the Web browser by clicking the **Install/Run Fabric Manager** button.

   If this is a first time install of the Cisco applications, you may get the security warning shown in Figure 2-12. Click the **Start** button to proceed.

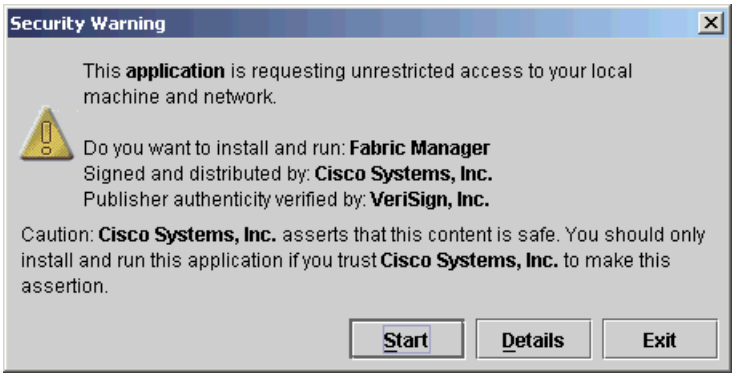


*Figure 2-12 Fabric Manager security warning*

3. When the installation is complete, you should see the Fabric Manager login window as shown in Figure 2-13. The Fabric Manager is now ready for use.

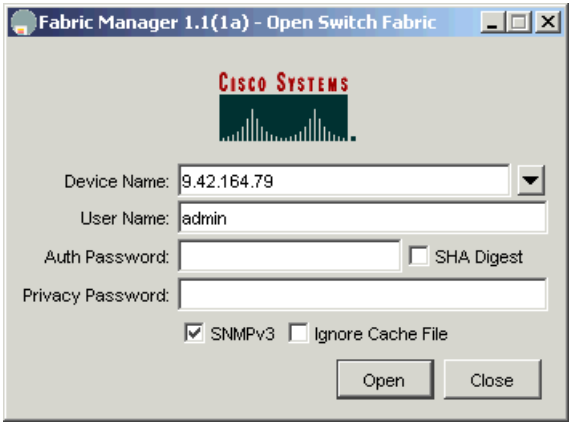


*Figure 2-13 Fabric Manager initial login window*

4. Start the installation of Device Manager from the Web browser by clicking the **Install/Run Device Manager** button.
5. When the installation is complete, you see the Device Manager login window, as shown in Figure 2-14. The Device Manager is now ready for use.


*Figure 2-14 Device Manager initial login window*

## 2.3 Managing the Cisco switch with the Device Manager

The Cisco Device Manager provides a GUI for managing the parameters of a single Cisco MDS 9000 series switch.

While features affecting multiple switches, such as PortChannel, can be configured using the Device Manager, we recommend using Fabric Manager to configure them. This way, the changes are coordinated across multiple switches. VSANs, however, are considered to affect only a single switch, and can be configured with either application.

### 2.3.1 Getting started

You can start the Device Manager from the icon from your desktop or Windows Start menu. Enter the IP address or host name of your switch, the SNMPv3 user name and password, and the privacy password, if you have one set up.

If you have used the Device Manager before, you can choose one of the devices you have used previously from the pull-down menu, as shown in Figure 2-15.


*Figure 2-15 Device Manager device names pull-down menu*

If the connection cannot be established, you will get an error message similar to that shown in Figure 2-16.

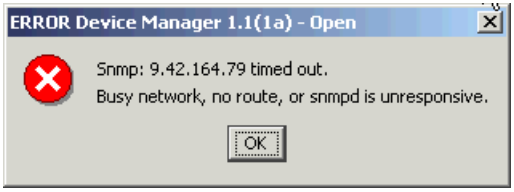


*Figure 2-16 Device Manager login error message*

## 2.3.2 User interface

The main user interface of the Device Manager represents the front panel of the switch, and is different for different members of the MDS 9000 family. The Device Manager views of the MDS 9216 and MDS 9509 are shown in Figure 2-17 and Figure 2-18.

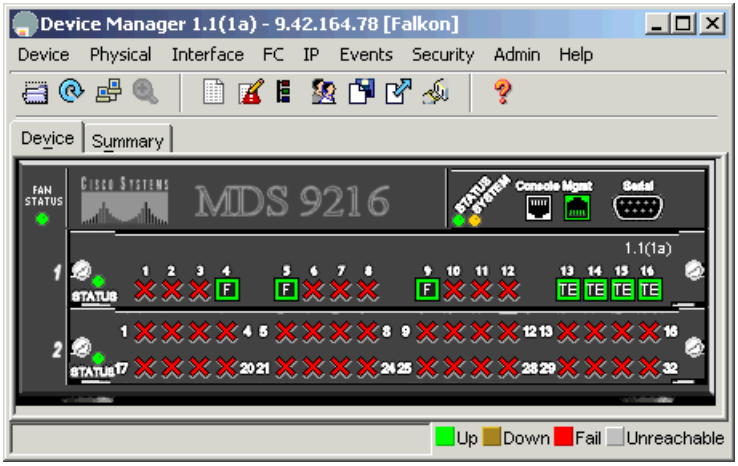


*Figure 2-17 Device Manager view of MDS 9216*

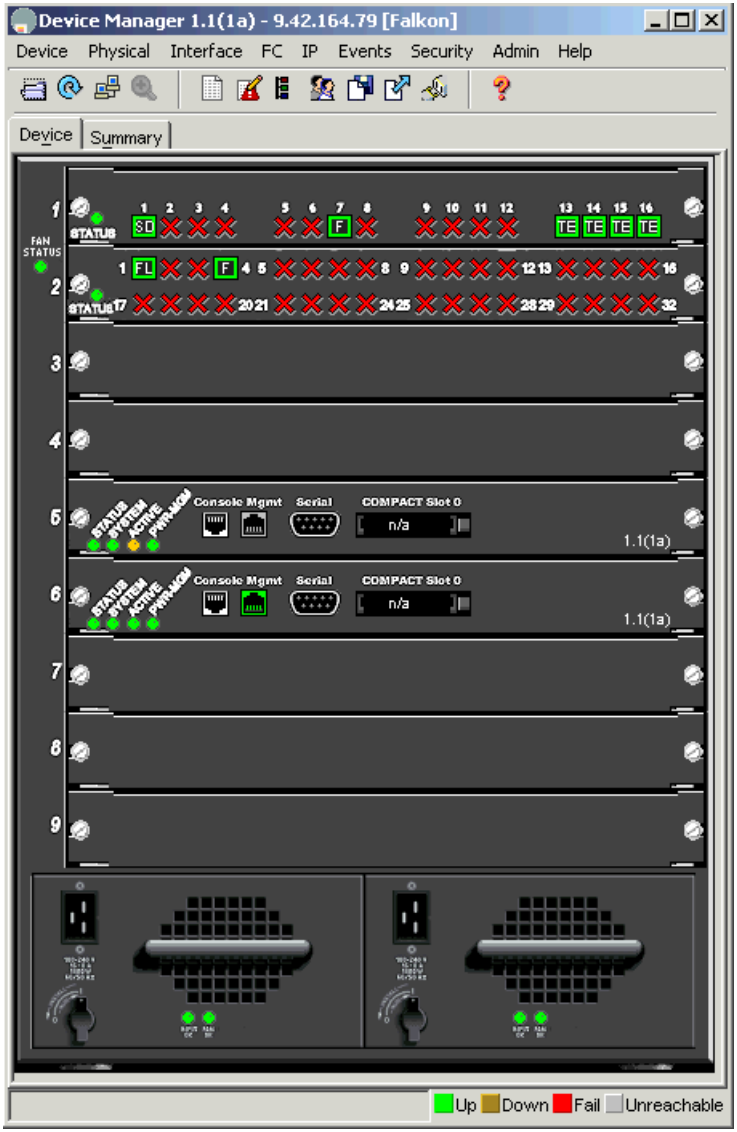


*Figure 2-18 Device Manager view of MDS 9509*

In the views described above, active ports are shown as green squares, with the current port mode shown inside the square. Enabled ports that are not active, as well as failed ports, are shown as red X. Ports that have been disabled by the administrator are shown as brown, empty squares.

By clicking the **Summary** tab, you will get a summary of the devices currently connected to the ports in the switch, as well as the CPU load, memory usage, and flash disk usage, as shown in Figure 2-19.

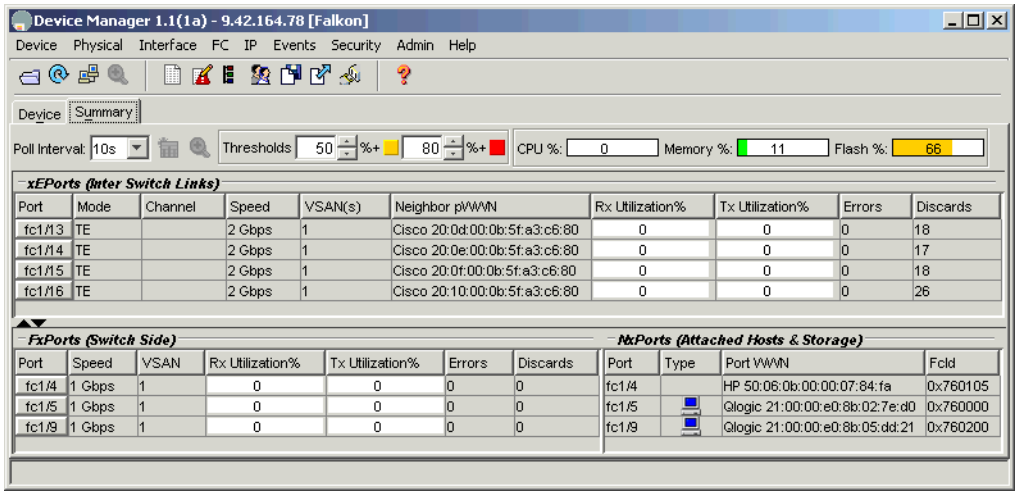


*Figure 2-19 Device Manager summary of connected devices*

## SNMP timeouts

The Device Manager uses the SNMP protocol to communicate with the switch. SNMP is a stateless protocol. When you apply changes to the switch, the Device Manager sends a request packet with the changes to the switch and waits for a response packet.

Depending on your network, either the request packet or the response packet may end up being dropped. This results in an SNMP timeout message as shown in Figure 2-20.

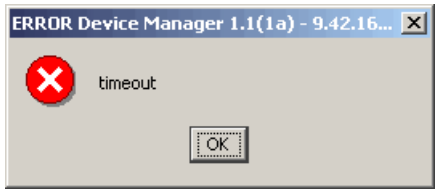


*Figure 2-20 Device Manager SNMP timeout message*

If you get this message, you do not know which one of the packets was dropped. This means that you do not know if your changes are applied to the switch or not. We recommend that you click the **Refresh Display** button to ensure that the information in the Device Manager is up to date before making any further changes.

### Stopping the Device Manager

If you have made changes to the Cisco running configuration that have not yet been copied to the startup configuration, you will get a message as shown in Figure 2-21.

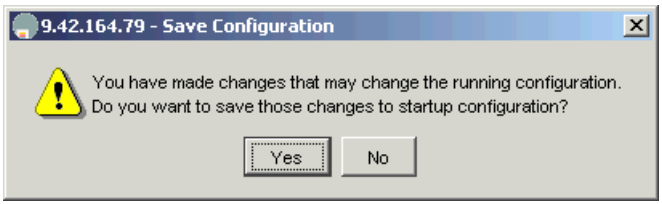


*Figure 2-21 Unsaved running configuration warning*

You can click **Yes** to copy the running configuration to the startup configuration. The Device Manager can be closed when the copy process is finished.

## 2.3.3 Context-sensitive menus

The interfaces, the management ethernet, modules, power supplies, and the system all have context sensitive menus. You can activate a context sensitive menu for an element by clicking it with the right mouse button.

### Fibre Channel interface menu

The menu for a Fibre Channel interface is shown in Figure 2-22.


*Figure 2-22 Fibre Channel interface menu*

The Configure option can be used to configure various parameters of the interface and is described in detail below.

You can also use the menu options to monitor the traffic on the interface, enable or disable the interface, or turn on or off beaconing (flashing LED) for the interface for identification purposes. Beaconing has no effect on the switch operation.

### *Configure General tab*

The **General** tab can be used to view and configure most of the settings for the port, such as the VSAN the port belongs to, the port mode, the trunking mode, and port speed. It can also be used to give the port a description, or enable or disable the port. The **General** tab is shown in Figure 2-23.

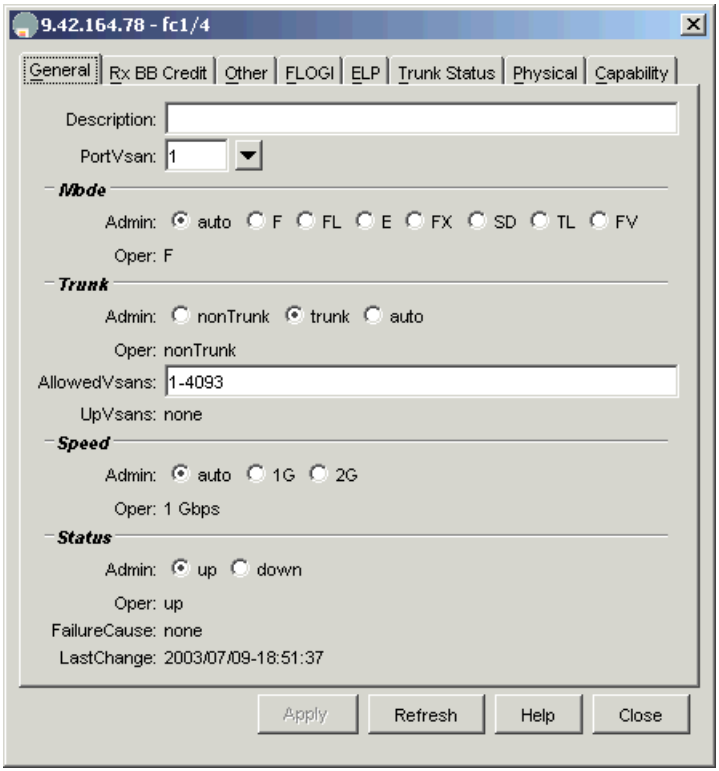


*Figure 2-23 Configure General tab*

### *Configure Rx BB Credit tab*

The **Rx BB Credit** tab can be used to configure the BB credits for the port. The allowed values for ports in a 16 port line card are 1-255. The value for the ports in a 32 port line card is fixed to 12, and any attempt to change it results in an error. The **Rx BB Credit** tab is shown in Figure 2-24.

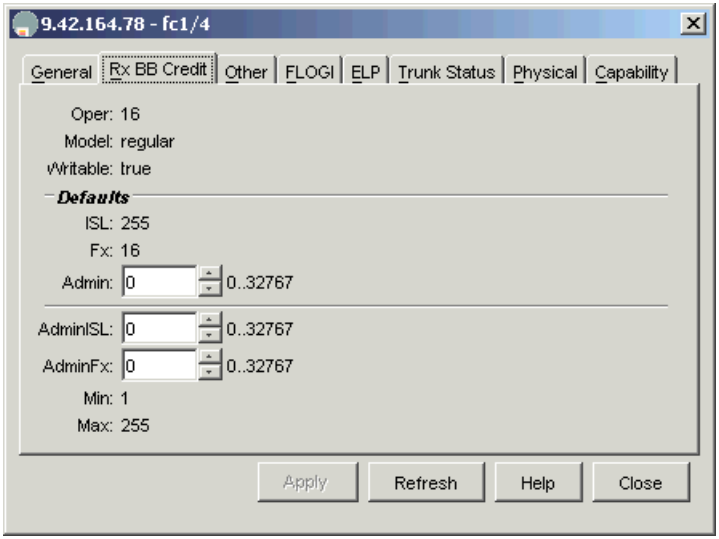


*Figure 2-24 Configure Rx BB Credit tab*

### *Configure Other tab*

The **Other** tab can be used to configure the maximum data field size of the interface. This setting does not usually need to be changed. The **Other** tab is shown in Figure 2-25.

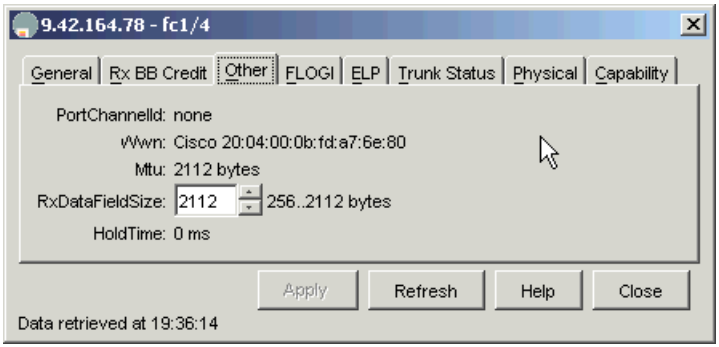


*Figure 2-25 Configure Other tab*

### *Configure FLOGI tab*

The **FLOGI** tab can be used to view the device ports (N_Ports or NL_Ports) connected to the switch port. The **FLOGI** tab is shown in Figure 2-26.

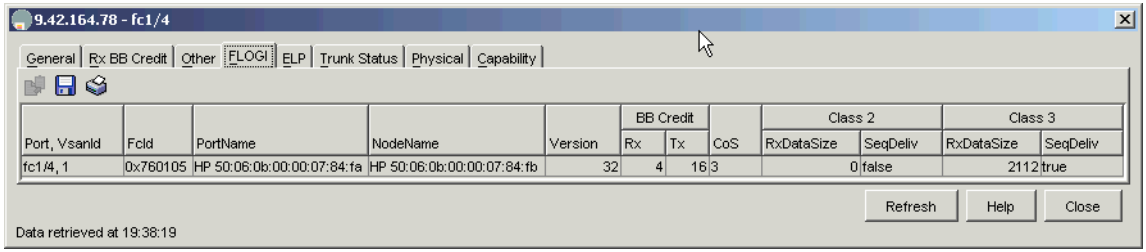


*Figure 2-26 Configure FLOGI tab*

### *Configure ELP tab*

The **ELP** tab can be used to view the exchange link parameter (ELP) attributes, and contains valid information only for E_Ports and TE_Ports. The settings cannot be changed. The **ELP** tab is shown in Figure 2-27.

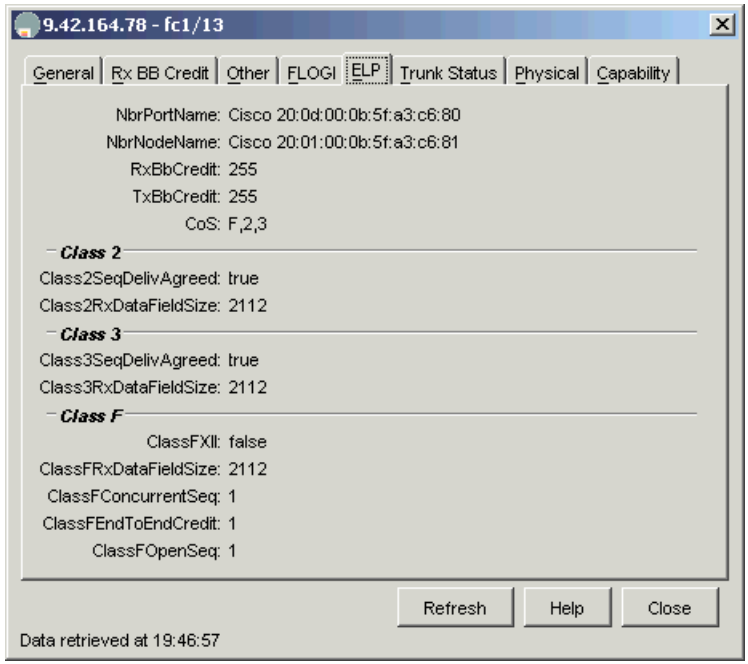


*Figure 2-27 Configure ELP tab*

### *Configure Trunk Status tab*

The **Trunk Status** tab can be used to view the status of different VSANs on a TL_Port. The **Trunk Status** tab is shown in Figure 2-28.

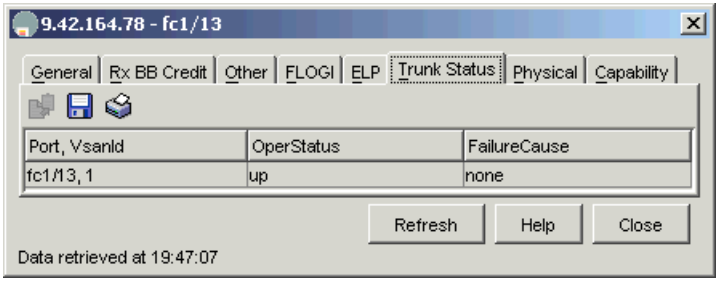


*Figure 2-28 Configure Trunk Status tab*

### *Configure Physical tab*

The **Physical** tab can be used to view information about the SFP transceiver installed in the port. It can also be used to turn on beaconing for the port for identification purposes. Beaconing has no effect on the switch operation. The **Physical** tab is shown in Figure 2-29.

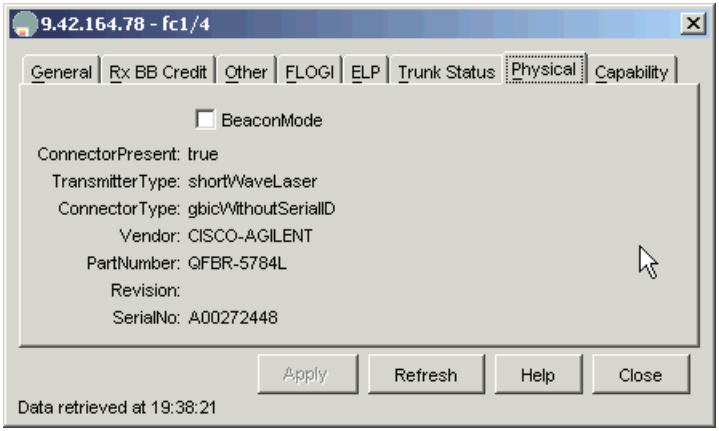


*Figure 2-29 Configure Physical tab*

### *Configure Capability tab*

The **Capability** tab can be used to view some general port capabilities. The **Capability** tab is shown in Figure 2-30.

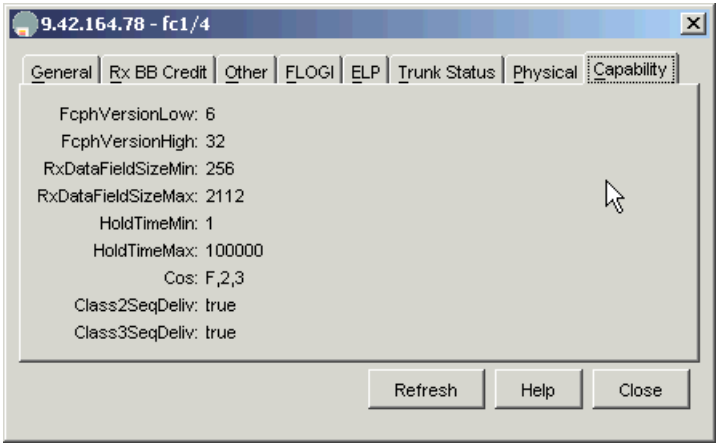


*Figure 2-30 Configure Capability tab*

## Gigabit ethernet interface menu

The menu for the gigabit ethernet interfaces is shown in Figure 2-31.


*Figure 2-31 Gigabit ethernet interface menu*

The Configure option can be used to configure various parameters of the interface. The Monitor option can be used to monitor the traffic on the interface. Both menus are described in detail below.

### *Configure General tab*

The **General** tab of the Configure window can be used to view and configure the basic settings of the port, such as the IP address and maximum transmission unit (MTU).

Configuring the MTU to a higher value than the ethernet default of 1500 bytes allows a single ethernet packet to contain a maximum length Fibre Channel frame. However, if you want to use this functionality, ensure that your ethernet hardware supports these size frames.

The **General** tab is shown in Figure 2-32.

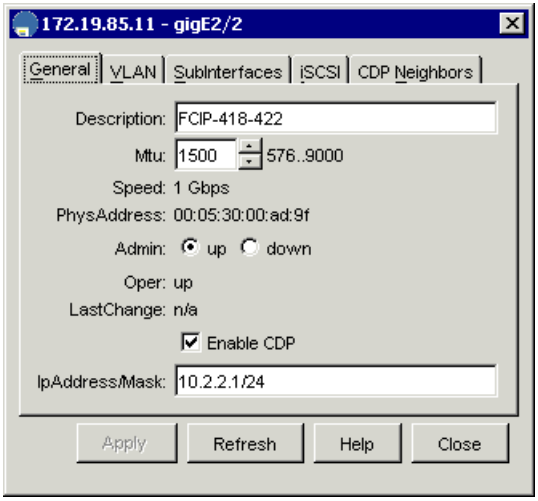


*Figure 2-32 Configure General tab*

### *Configure VLAN tab*

The **VLAN** tab of the Configure window can be used to define any virtual LANs (VLANs) for the ethernet interface, if the interface is connected to an 802.1Q compliant switch. You can list a comma separated list of VLAN numbers here. Each VLAN has a separate subinterface. The **VLAN** tab is shown in Figure 2-33.

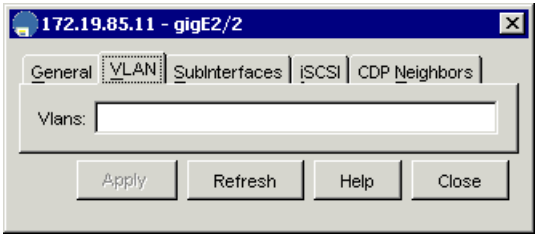


*Figure 2-33 Configure VLAN tab*

### *Configure SubInterfaces tab*

The **SubInterfaces** tab of the Configure window allows you to define the parameters of each VLAN subinterface defined for the physical ethernet interface. The parameters are the same as the parameters defined for the physical ethernet interface in the **General** tab. The **SubInterfaces** tab is shown in Figure 2-34.


*Figure 2-34 Configure SubInterfaces tab*

### *Configure iSCSI tab*

The **iSCSI** tab of the Configure window allows you to turn on or off the iSCSI functionality for the physical port. The **iSCSI** tab is shown in Figure 2-35.

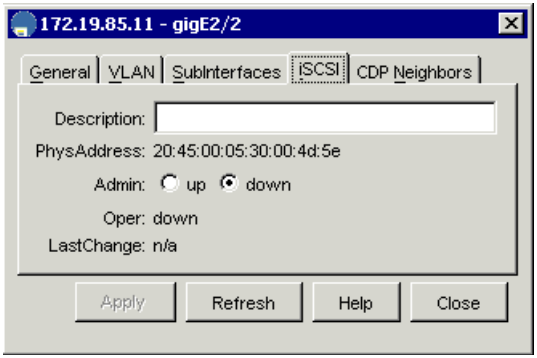


*Figure 2-35 Configure iSCSI tab*

### *Configure CDP Neighbors tab*

The **CDP Neighbors** tab of the Configure window shows a list of the other Cisco switches discovered with the Cisco Discovery Protocol (CDP) through this interface. The **CDP Neighbors** tab is shown in Figure 2-36.

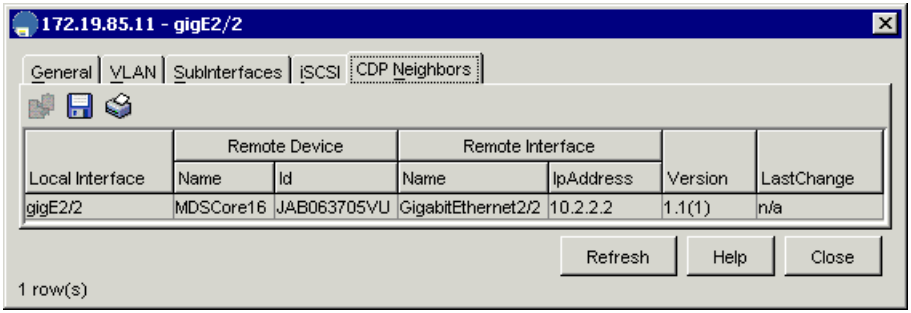


*Figure 2-36 Configure CDP Neighbors tab*

### *Monitor window*

The Monitor window allows you to monitor the IP traffic of the gigabit ethernet interface. The Monitor window is shown in Figure 2-37.

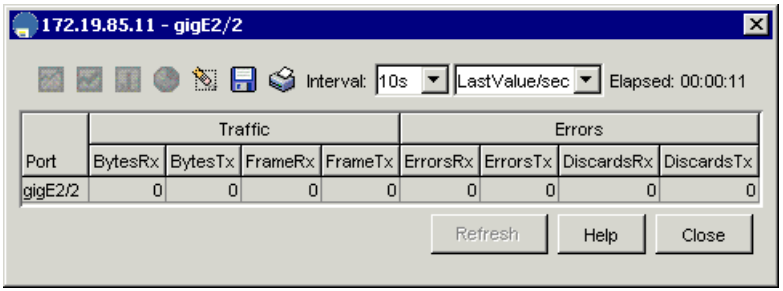


*Figure 2-37 Monitor window*

## Management ethernet menu

The menu for the management ethernet (mgmt0) is shown in Figure 2-38.


*Figure 2-38 Management ethernet menu*

The only option in the menu is Configure. The management ethernet configuration window is shown in Figure 2-39.

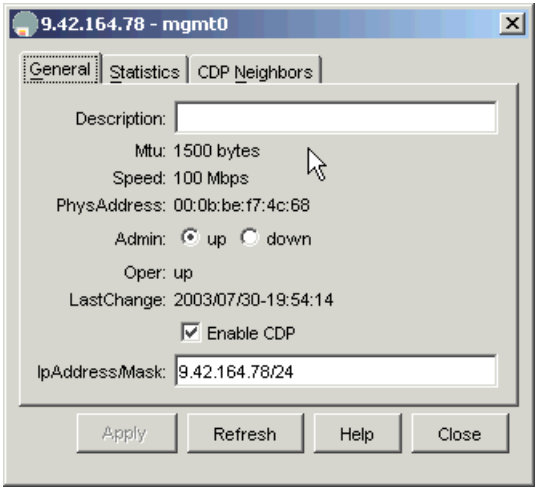


*Figure 2-39 Management ethernet configuration window*

Note that while you can change the parameters for the management ethernet interface here, many of the changes cause you to lose connection to the switch. Therefore, we do not recommend changing the network settings here.

## Module menu

The menu for a module is shown in Figure 2-40.

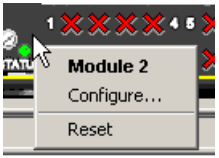


*Figure 2-40 Module menu*

You can only view the configuration of the module or reset it. The module configuration window is shown in Figure 2-41.

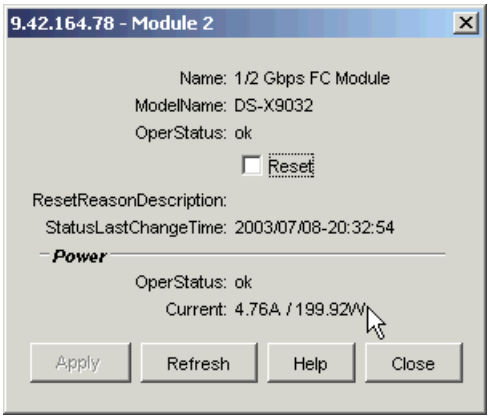


*Figure 2-41 Module Configure window*

## Power Supplies menu

You can use the Power Supplies menu to view or change the configuration of power supplies in certain members of the MDS 9000 family. The Power Supplies menu is shown in Figure 2-42.


*Figure 2-42 Power Supplies menu*

The only option in the menu is Configure. The power supply configuration window is shown in Figure 2-43.

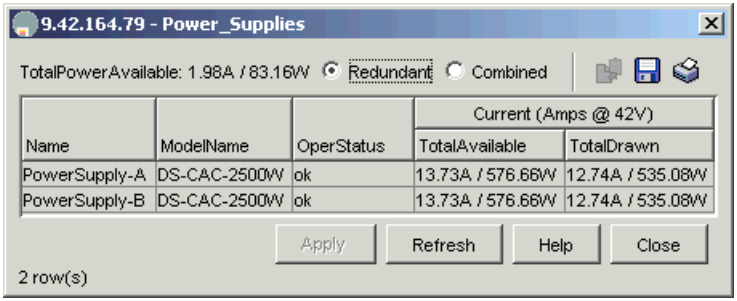


*Figure 2-43 Power supply configuration window*

You can see the total power and available power for the power supplies in the switch in the window. You can also change the power supply mode between Redundant and Combined.

## System menu

The System menu is shown in Figure 2-44.

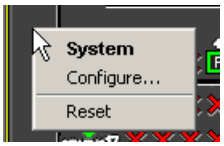


*Figure 2-44 System menu*

You can use the System menu to configure certain system-level parameters, or reset the switch. The Configure window is shown in Figure 2-45.

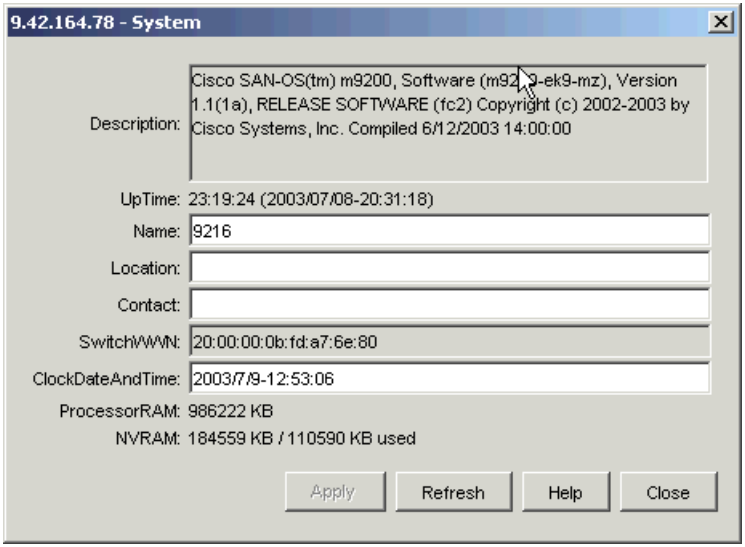


*Figure 2-45 System configuration window*

You can change the switch name, and the time and date of the switch. You can also enter the switch location and contact information.

## 2.4 Managing the Cisco SAN with the Fabric Manager

The Fabric Manager is a centralized tool used to manage the Cisco SAN fabric and the devices connected to it.

### 2.4.1 Getting started

You can start the Fabric Manager from the icon from your desktop or Windows Start menu. Enter the IP address or host name of your switch, the SNMPv3 username and password, and the privacy password, if you have one set up.

If you have used the Fabric Manager before, you can choose one of the devices you have used previously from the pull-down menu, as shown in Figure 2-46.

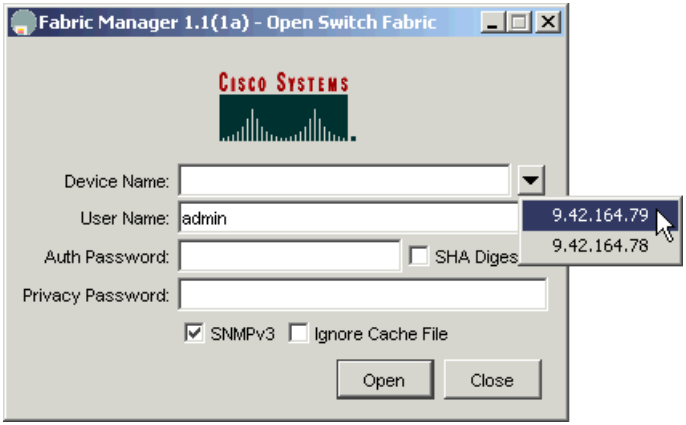


*Figure 2-46 Fabric Manager device names pull-down menu*

If the connection cannot be established, you get an error message, as in Figure 2-47.

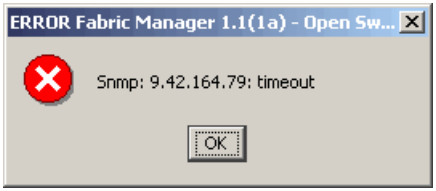


*Figure 2-47 Fabric Manager login error message*

## 2.4.2 User interface

When you start the Fabric Manager, you will see the logical view of your switch fabric, as shown in Figure 2-48.

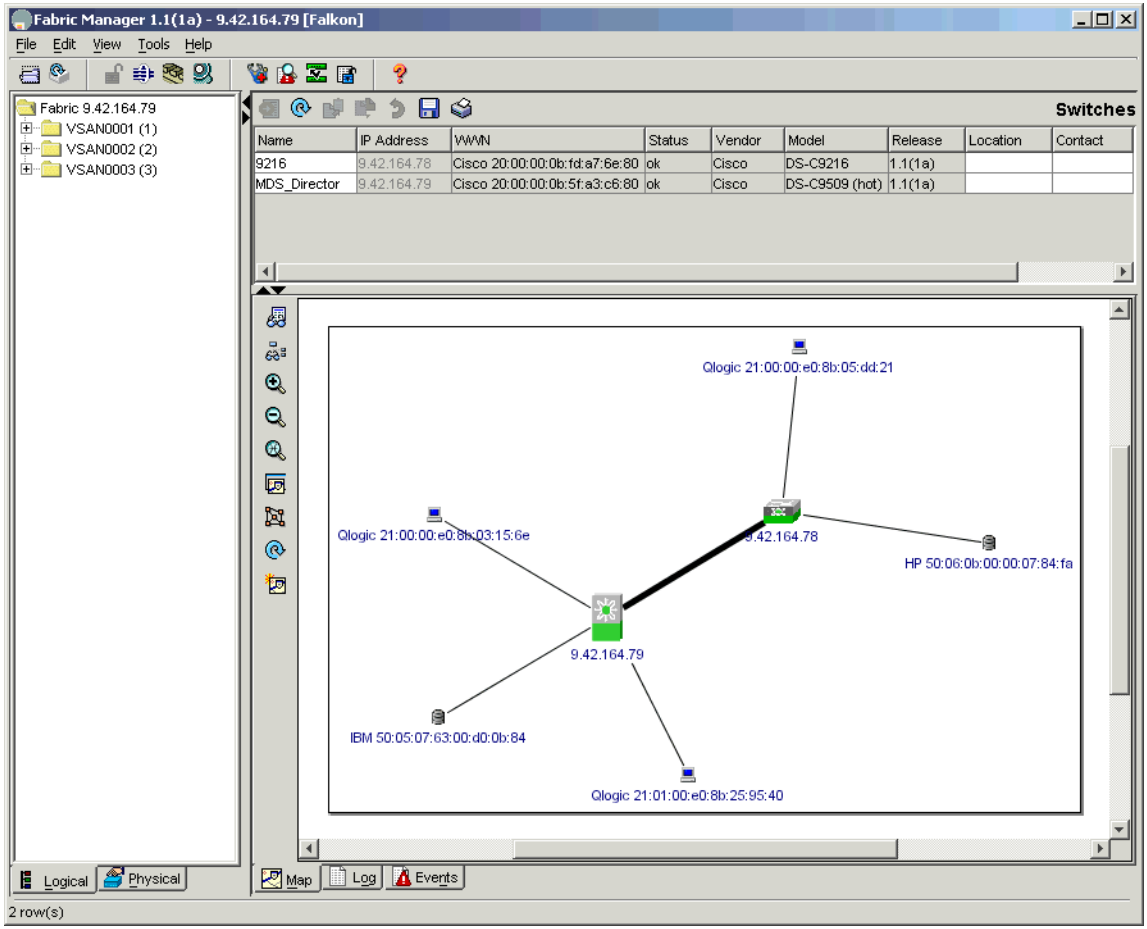


*Figure 2-48 Fabric Manager logical view*

The window contains the graphical representation of your switch fabric on the bottom right, an information area on the top right, and a navigation menu on the left. Any of the areas can be hidden to give more space to the other windows. The content of the information area changes automatically to represent the selection chosen in the navigation menu, and the current selection is shown on top of the information area.

There are two navigation menus available, and the menus can be selected by the tabs below the menu area. The logical menu is a representation of the VSANs defined in the network and the zone sets, zones and zone members within each VSAN. The physical menu is a representation of all the physical assets in the SAN, and can also be used to configure most operating parameters of all of the switches in the SAN.

## SNMP timeouts

The Fabric Manager uses the SNMP protocol to communicate with the switch. SNMP is a stateless protocol, and when you apply changes to the switch, the Fabric Manager sends a request packet with the changes to the switch and waits for a response packet.

Depending on your network, either the request packet or the response packet may end up being dropped. This results in a SNMP timeout message, as shown in Figure 2-49.

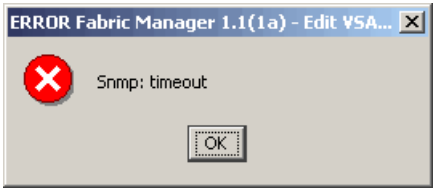


*Figure 2-49 Fabric Manager SNMP timeout message*

If you get this message, you do not know which of the packets was dropped. This means that you do not know if your changes are applied to the switch or not. We recommend that you click the **Refresh Values** button to ensure that the information in the Fabric Manager is up to date before making any further changes.

## Stopping the Fabric Manager

If you have made changes to the Cisco running configuration that have not yet been copied to the startup configuration you will get a message as shown in Figure 2-50.

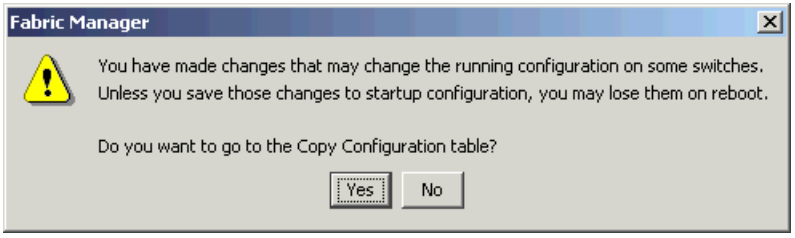


*Figure 2-50 Unsaved running configuration warning*

You can click **Yes** to go to the Copy Configuration window, and then click **Apply Changes** to do the actual copy, and wait for the copy processes to finish. After all of the copy processes are finished you can close the Fabric Manager.

The Fabric Manager also saves information about your switch fabric into a local database in your workstation. If you have changes that have not been saved, you will get a message as shown in Figure 2-51.

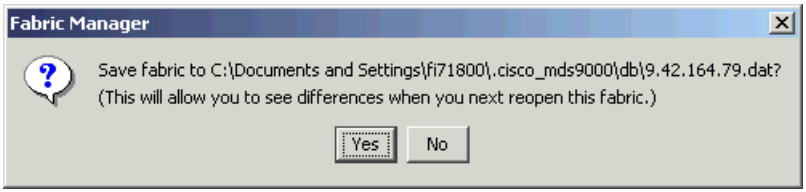


*Figure 2-51 Unsaved local fabric database warning*

Since having the local database up to date helps you to see any changes to the fabric, when you open the Fabric Manager again, it is a good idea to click **Yes** here.

## 2.4.3 Managing zones and zone sets

In the switches in the Cisco MDS 9000 family, each VSAN has its own zones and zone sets. Only one zone set can be active in a VSAN at any given time. The zone set can contain multiple zones, and a zone can belong to multiple zone sets.

Each switch has a local zone database for each VSAN, that can be used to create zoning configurations. The zoning information for the active zone set is propagated to all of the switches when a zone set defined in the local zone database is activated. However, the local zone database is not replicated to the other switches. We recommend that you choose one of the switches in the fabric, and use it to maintain your zone database.

You can open the local zone database for a VSAN by opening the logical menu, clicking the name of the VSAN with the right mouse button, and choosing the **Edit Local Zone Database** option. An example of the local zone database window is shown in Figure 2-52.

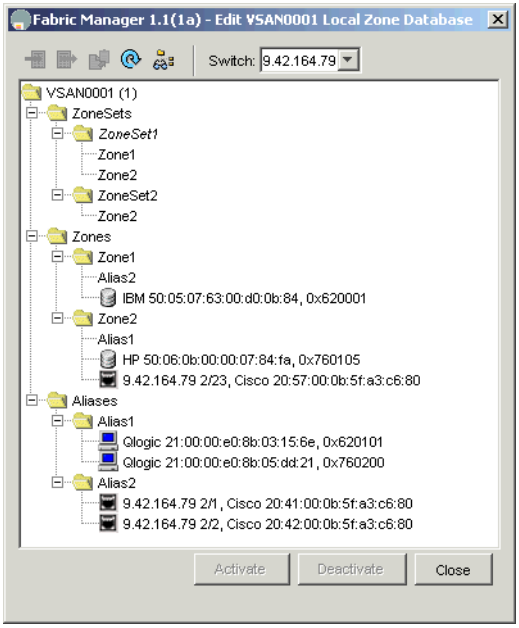


*Figure 2-52 Local zone database*

In this example, ZoneSet1 consists of both Zone1 and Zone2, while ZoneSet2 consists of only Zone1. ZoneSet1 is currently the active zone set, since its name is shown in italics. Zone1 consists of Alias2 and one port WWN (pWWN), Zone2 consists of Alias1, one pWWN, and one fabric WWN (fWWN). Alias1 consists of two pWWNs, and Alias2 consists of two fWWNs.

**Note:** You cannot use the Fabric Manager to add members to an alias or a zone by FC ID. If you want to add the members by FC ID, you have to use the CLI. Also, the members added by FC ID look exactly the same as the members added by pWWN, when viewing the local zone database. They can only be distinguished with the CLI. We recommend that you use only one method to manage zoning - either GUI or CLI, but not both.

## Creating aliases, zones, and zone sets

When creating the initial zoning for the fabric, we recommend starting from the bottom and working your way up the tree. Start by first creating aliases for your devices, then zones, and finally the zone set, and finish by activating your zone set.

We will start from an empty local zone database. In this example, we use the buttons in the top of the window, but you can also find the same commands by clicking an object with the right mouse button.

### *Creating an alias*

To create a new alias, activate the Aliases line and click the **Insert** button. You can now give the name for the new alias in a window as shown in Figure 2-53.

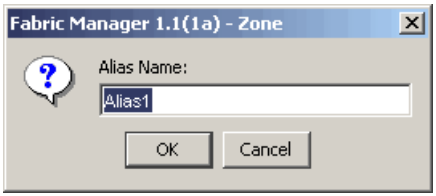


*Figure 2-53 Creating an alias*

When you click **OK**, the alias is created, but it does not have any members yet. To add pWWNs or fWWNs as members to the alias, activate the newly created alias and click the **Insert** button. You will see a window similar to that shown in Figure 2-54.

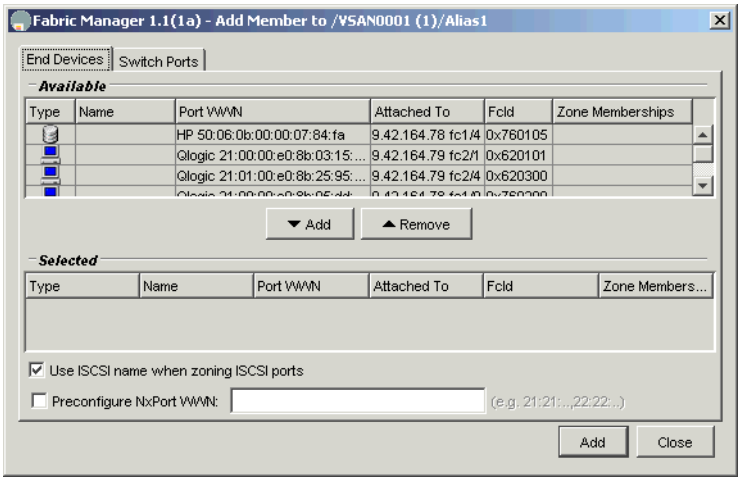


*Figure 2-54 Adding members to the alias*

You will see a list of all active end devices (pWWNs) in the switch fabric. If you want to add members based on the end devices (pWWN), you can choose one or more of the devices in the Available list, and copy them to the Selected list by using the **Add** button between the lists, as shown in Figure 2-55.

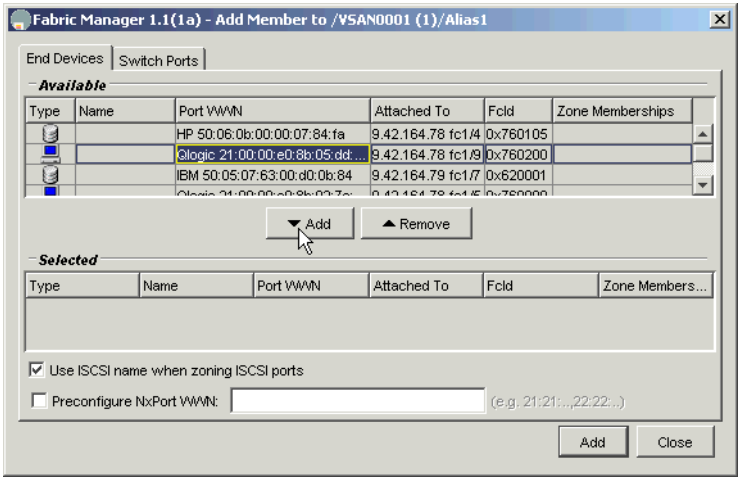


*Figure 2-55 Adding members by pWWN 1*

When you have all the pWWNs you want to add in the Selected list, click the **Add** button in the lower right corner of the window to add them into the alias as shown in Figure 2-56.

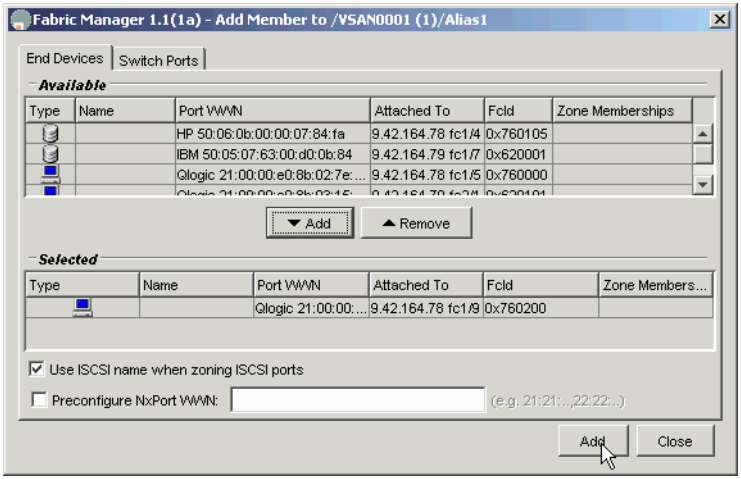


*Figure 2-56 Adding members by pWWN 2*

The Add Member window remains open, and you can add more members to the alias. When you have finished adding members to the alias, click **Close** to close the window.

You can also add members to the alias by their switch ports (fWWN) by clicking the **Switch Ports** tab. You will see a window as shown in Figure 2-57.

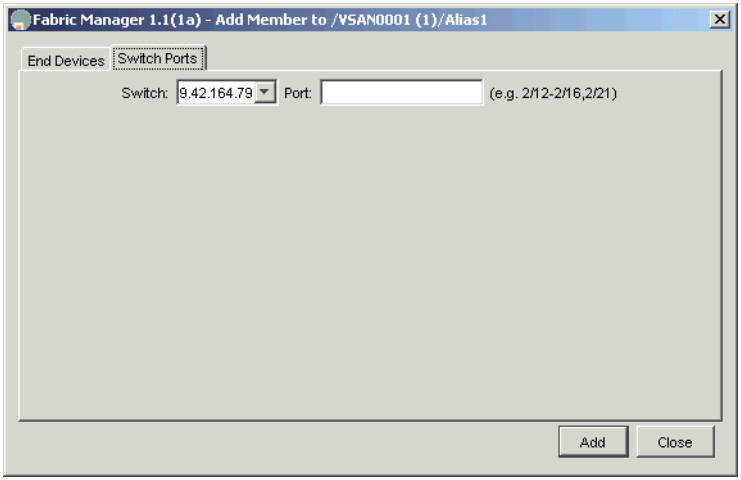


*Figure 2-57 Adding members by fWWN*

Enter the port names you want to add to the alias on the line. Note that you can also choose the ports from another switch in the fabric by choosing the IP address of the switch in the left field. Click **Add** to add the fWWNs to the alias.

The Add Member window remains open, and you can add more members to the alias. When you have finished adding members to the alias, click **Close** to close the window.

If you try to add a port that is not in the same VSAN you are working with to your alias, you will get a warning similar to that shown in Figure 2-58.

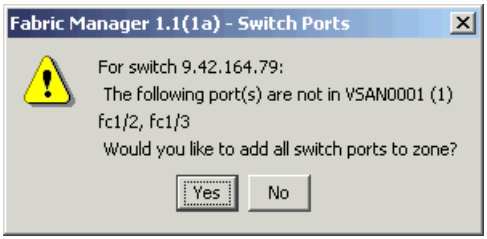


*Figure 2-58 Different VSAN warning*

While you can add any ports in your switches to the zone, only the ports in the same VSAN can communicate with each other. Adding ports not in the same VSAN has no effect on the zone operation.

### *Creating a zone*

To create a new zone activate the Zones line and click the **Insert** button. You can now enter the name for the new zone in the window as shown in Figure 2-59.

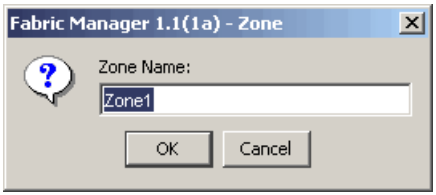


*Figure 2-59 Creating a zone*

When you click **OK**, the zone is created, but it does not have any members yet. To add members to the zone, activate the newly created zone and click the **Insert** button. You will see a window similar to that shown in Figure 2-60.

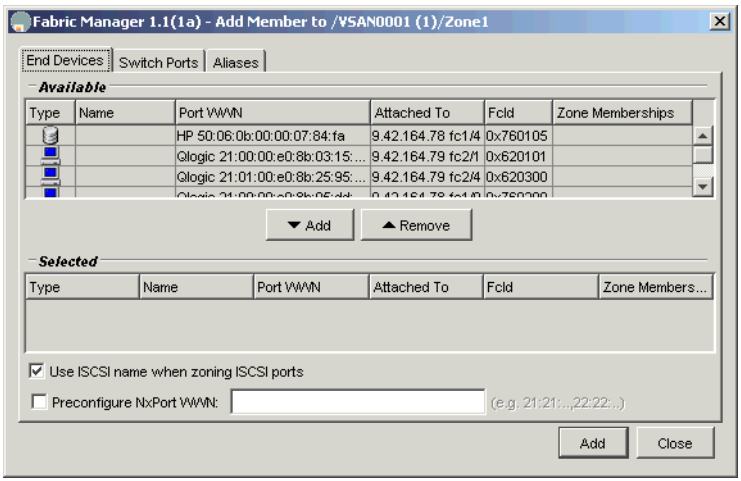


*Figure 2-60 Adding members to the zone*

Adding members by end device (pWWN) and switch port (fWWN) works exactly the same as adding members in an alias, and is not described here. However, there is also an option to add aliases to the zone. To add aliases to the zone, click the **Aliases** tab, and you will see a window similar to that shown in Figure 2-61.

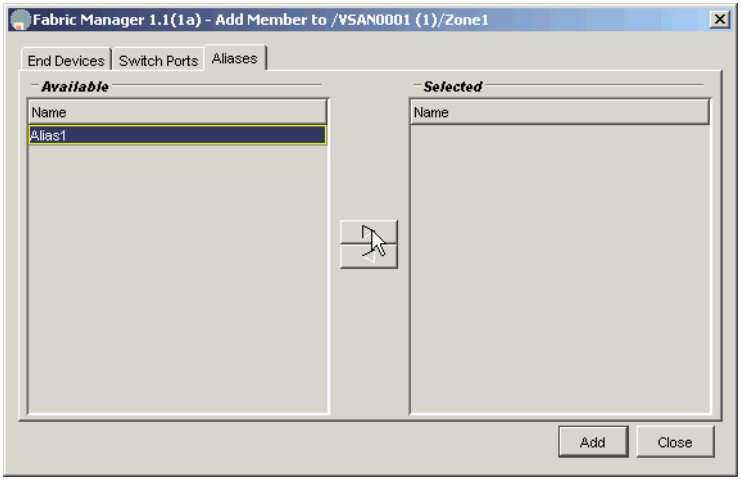


*Figure 2-61 Adding members by alias 1*

Choose the aliases you want to add from the Available list and move them to the Selected list by clicking the right arrow button in the middle. When you have selected all the aliases you want, click the **Add** button as shown in Figure 2-62.

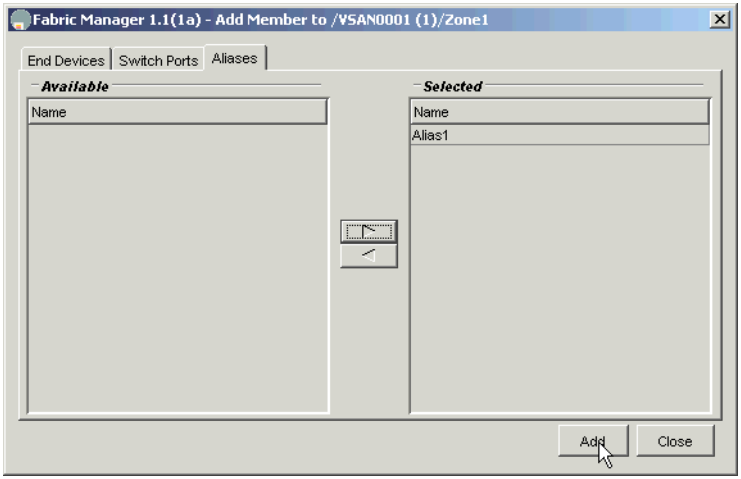


*Figure 2-62 Adding members by alias 2*

The Add Member window remains open, and you can add more members to the alias. When you have finished adding members to the alias, click **Close** to close the window.

### *Creating a zone set*

To create a new zone set, activate the ZoneSets line and click the **Insert** button. You can now give the name for the new zone set in a window as shown in Figure 2-63.

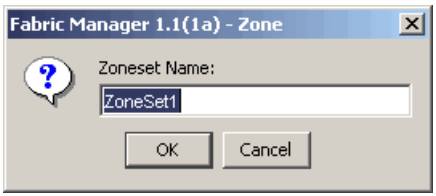


*Figure 2-63 Creating a zone set*

When you click **OK**, the zone set is created, but it does not have any zones yet. To add zones to the zone set, activate the newly created zone set and click the **Insert** button. You will see a window similar to that shown in Figure 2-64.

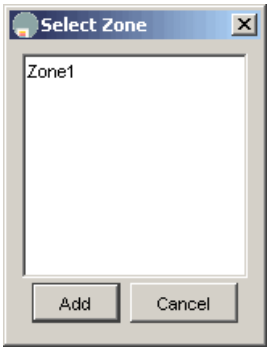


*Figure 2-64 Adding zones to the zone set*

Choose the zones you want to add to the zone set, and click **Add** to add them.

## Cloning aliases, zones and zone sets

You can make an identical copy of any alias, zone or zone set by selecting it and clicking the **Clone** button. This is especially useful if you need to make temporary changes to your zone set, while retaining the original definitions. You can make a clone of your zone set, edit it, and then activate the cloned and edited zone set.

Cloning works in the same way for all the elements. For example, to clone a zone set, select it and click the **Clone** button. You will see a window similar to that shown in Figure 2-65.

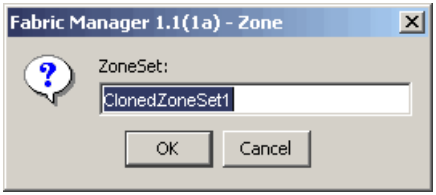


*Figure 2-65 Cloning a zone set*

Enter a name for your cloned zone set, and click **OK** to create it.

## Deleting zoning elements

You can delete any element from the local zone database, except the top level ZoneSets, Zones and Aliases entries, by choosing the element and clicking the **Delete** button. You will get a confirmation window as shown in Figure 2-66.

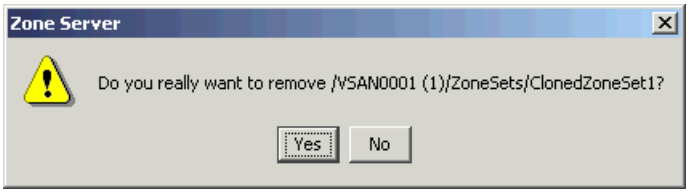


*Figure 2-66 Delete confirmation window*

Click **Yes** to delete the element, or **No** to not delete it.

To maintain consistency of the local zoning database, the following dependencies apply:

- If you delete an alias from a zone, or a zone containing an alias, the alias definition is retained.
- If you delete an alias definition from the Aliases tree, the alias is also deleted from any zones.
- If you delete a zone from a zone set, or a zone set containing a zone, the zone definition is retained.
- If you delete a zone definition from the Zones tree, the zone is also deleted from any zone sets.

## Activating and deactivating a zone set

To activate a zone set, choose the zone set from the local zone database and click the **Activate** button at the bottom of the window. You will see a window as shown in Figure 2-67.

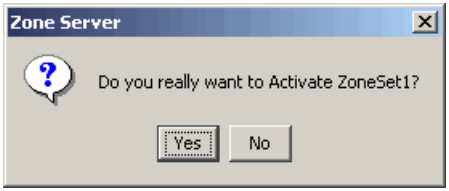


*Figure 2-67 Activating a zone set*

Click **Yes** to activate the zone set. If you had another zone set active previously, it is automatically deactivated.

**Note:** Activating a zone set also saves the complete current running configuration to the startup configuration. This permanently saves all changes made to the running configuration, not just zoning changes.

Usually there is very rarely a need to deactivate the active zone set. If you need to deactivate the active zone set, choose it from the local zone database and click the **Deactivate** button. You will see a window as shown in Figure 2-68.

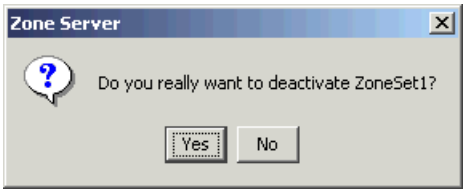


*Figure 2-68 Deactivating the zone set*

Click **Yes** to activate the zone set. After the zone set is deactivated, all devices in the fabric belong to the default zone, and the default zone policy (permit or deny) applies to all communication between them.

## 2.4.4 Managing VSANs

You can open the Fabric Manager VSAN list by choosing the path **Switches—>FC—>VSANs** from the physical menu. All of the VSANs in all of the switches in the fabric are listed in tabular form as shown in Figure 2-69.

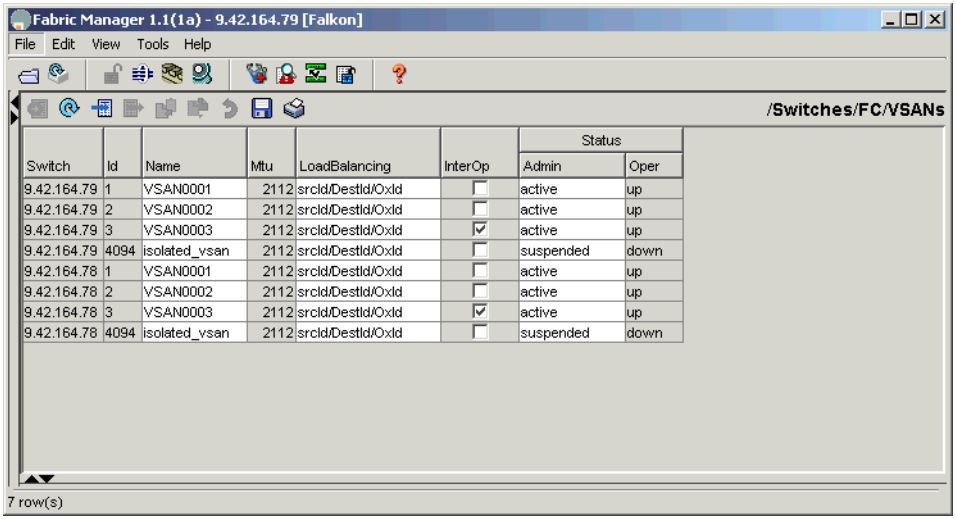


*Figure 2-69 VSAN list*

Note that different switches may have different sets of VSANs. However, there are two VSANs that exist in every switch, numbers 1 and 4094. You cannot delete either of these, but you can suspend VSAN 1 if necessary.

## Creating a new VSAN

If you want to create a new VSAN, click the **Create Row...** button. You will see a window as shown in Figure 2-70.

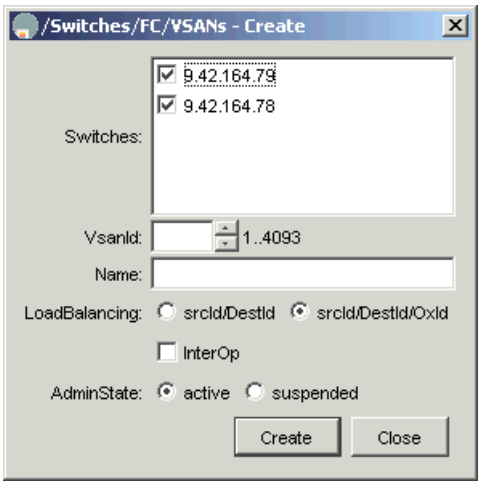


*Figure 2-70 VSANs - Create dialog*

You can choose the switches where you want to have the VSAN. By default, all switches are chosen.

You have to give the VSAN a number. A VSAN with that number cannot exist in any of the switches where you are creating the new VSAN.

By default, the VSAN is given a name in the form VSAN<number>, where <number> is the number of the VSAN zero-padded to four digits. You can overwrite this name with any name you choose.

You can also specify the load balancing method used in this VSAN as well as if you want to use interoperability mode in this VSAN. The interoperability mode is discussed in more detail in 2.6, “Interoperability mode implications” on page 397.

Our completed VSAN dialog is shown in Figure 2-71.

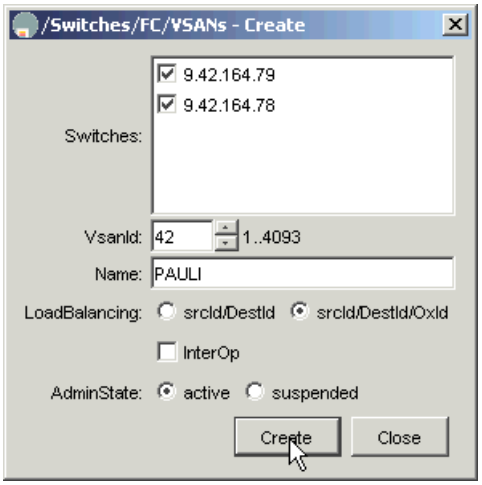


*Figure 2-71 VSANs - Create dialog completed*

Click **Create** to create the new VSAN and the VSAN is immediately created. The VSAN creation window remains open and changes to the next higher VSAN number automatically, allowing you to quickly create more than one similar VSAN. When you are finished creating VSANs, click **Close** to close the window.

The new VSAN can now be seen in the VSAN list, as shown in Figure 2-72.

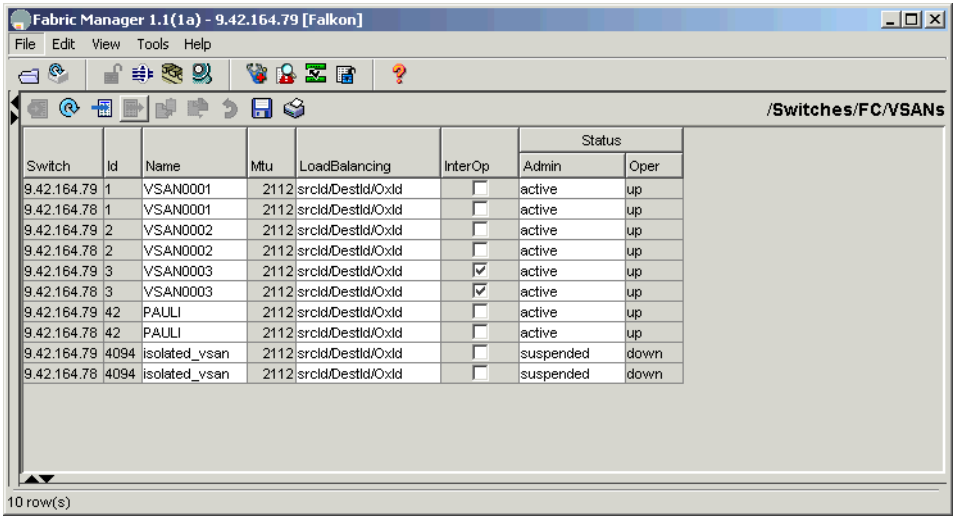


*Figure 2-72 VSAN list with the new VSAN*

You can now assign ports to the new VSAN.

### Deleting a VSAN

You can delete VSANs by choosing one or more lines in the VSAN list and clicking the **Delete Row** button. You will get a confirmation dialog as shown in Figure 2-73.

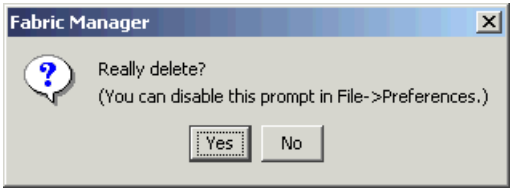


*Figure 2-73 Delete VSAN confirmation*

If you delete a VSAN that has ports assigned to it, the ports are moved to VSAN 4094, the isolated_vsan. This causes the ports to be isolated from any other ports, including each other.

## 2.4.5 Managing administrator access

The Cisco MDS 9000 family of switches use two different methods for access control.

The Fabric Manager and Device Manager communicate with the switch using the SNMP protocol, and uses the access control methods provided with SNMP. Both SNMPv2 and SNMPv3 are supported. We do not recommend using SNMPv2, due to the lack of security mechanisms in it.

You can connect to the command line interface (CLI) by using telnet or ssh protocols, or by using the console serial port of the switch. The switch can have its own usernames and passwords for the command line, or you can use an external RADIUS authentication server.

Both of the methods (SNMP and CLI) have their own set of usernames and passwords. You can use the Fabric Manager to manage the SNMP user accounts, but the CLI user accounts can only be managed by the CLI.

You can get a list of the currently defined SNMPv3 users by choosing the path **Switches—>Security—>SNMP** from the physical menu. All of the VSANs in all of the switches in the fabric are shown in tabular form as shown in Figure 2-74.

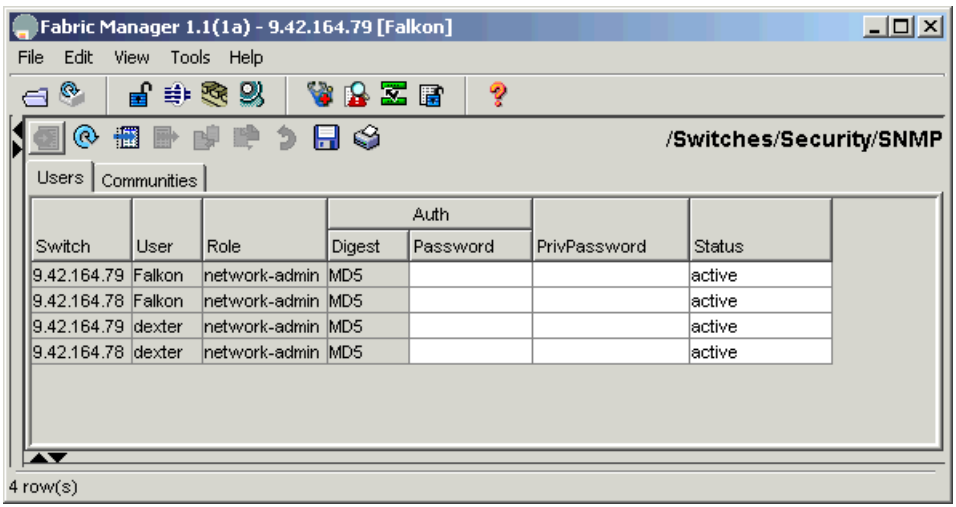


*Figure 2-74 List of SNMP users*

## Adding a new SNMP user

You can add a new user by clicking the **Create Row...** button. You can enter the details for the new user in the window as shown in Figure 2-75.

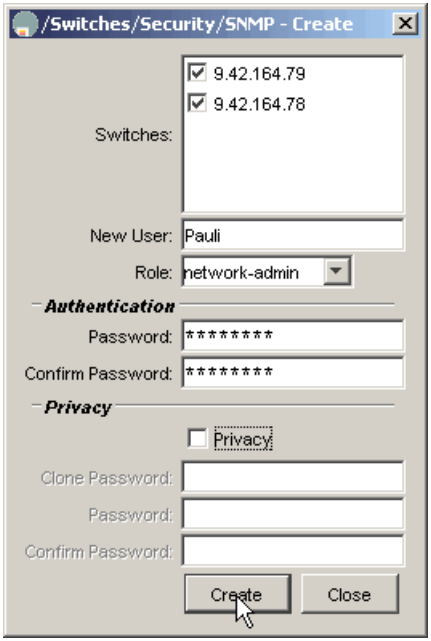


*Figure 2-75 Adding a SNMP user*

Enter the user name, role and password in the fields in the window. You can also choose to create the user in only a subset of the switches here. The authority of the user depends on the role chosen. There are two roles available: network-admin and network-operator.

If you want to encrypt all the management traffic for this user you can enter a privacy password as well. The user has to know the privacy password to log in.

Click **Create** to create the user when you are finished. The window stays open allowing you to create additional users. When you are finished creating the users, click **Close**.

## Deleting a SNMP user

You can delete a SNMP user by choosing the user from the table and clicking the **Delete Row** button. You see a confirmation window, as shown in Figure 2-76.

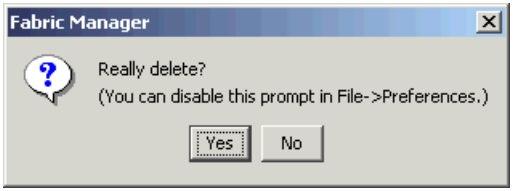


*Figure 2-76 Deleting a SNMP user*

Click **Yes**, if you want to delete the user.

**Note:** If you delete all the SNMP users, you will have to use the CLI to create new users before you can use Fabric Manager or Device Manager.

## 2.4.6 Managing software and configuration files

All Cisco MDS 9000 family members contain an internal bootflash memory that resides in the supervisor module. The MDS 9506 and MDS 9509 also have a slot for an additional CompactFlash card.

### Upgrading the switch software

The current version of Fabric Manager has limited support for upgrading the software in the switches. However, we still recommend that you use the CLI to upgrade the software. The upgrade procedure is described in 2.5.4, “Upgrading the switch software with the CLI” on page 390.

### Saving the running configurations

Each switch has two configurations:

- Running configuration
- Startup configuration

The running configuration is the currently active configuration. The startup configuration is the configuration that is activated the next time the switch is started. In the case of software upgrade, the startup configuration is automatically converted to the format of the new software, when the new software is started.

If you have made changes to the running configuration, and want them to be active across switch restarts, you have to save the running configuration to the startup configuration.

To copy the running configuration of one or more switches to the startup configuration, choose the path **Switches—>Copy Configuration** from the physical menu, click the check boxes for the switches where you want to run the copy, and click the **Apply Changes** button, as shown in Figure 2-77.

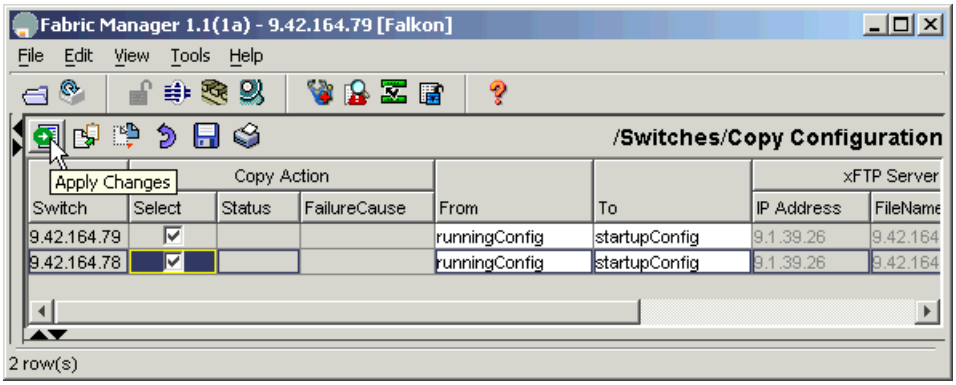


*Figure 2-77 Copying running configuration to startup configuration*

When the copy is running, the status fields show inProgress. When the copy is finished, the status fields change to Success.

## Copying configuration files to an FTP server

You can also copy the running configuration or startup configuration to an external FTP or TFTP server. In this case you have to change the To field to serverFile, and fill in the details on your FTP server as shown in Figure 2-78.

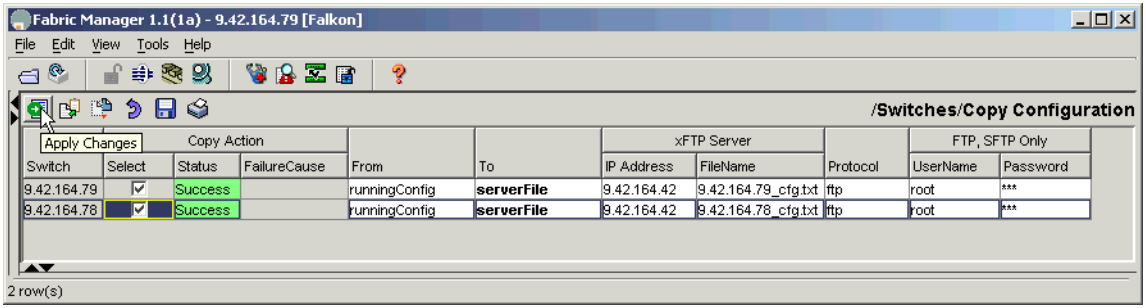


*Figure 2-78 Copying configuration to a FTP server*

You can use FTP, SFTP or TFTP as the protocol depending on the capabilities of your remote server. If you are using FTP or SFTP you also have to fill in the username and password fields. The default IP address is the address you are logging in from. When you have finished entering the values, click the **Apply Changes** button.

When the copy is running, the status fields show inProgress. When the copy is finished, the status fields change to Success. If the copy fails for any reason, the status fields show Failed, and the cause of the failure is shown in the failure cause field.

## 2.4.7 Managing interfaces

In this section we discuss general interface attributes and PortChannels.

### Managing general interface attributes

You can change the attributes of Fibre Channel interfaces by choosing the path **Switches—>FC—>Physical Interfaces** from the physical menu. All of the interfaces in the fabric are shown in tabular form as shown in Figure 2-79.

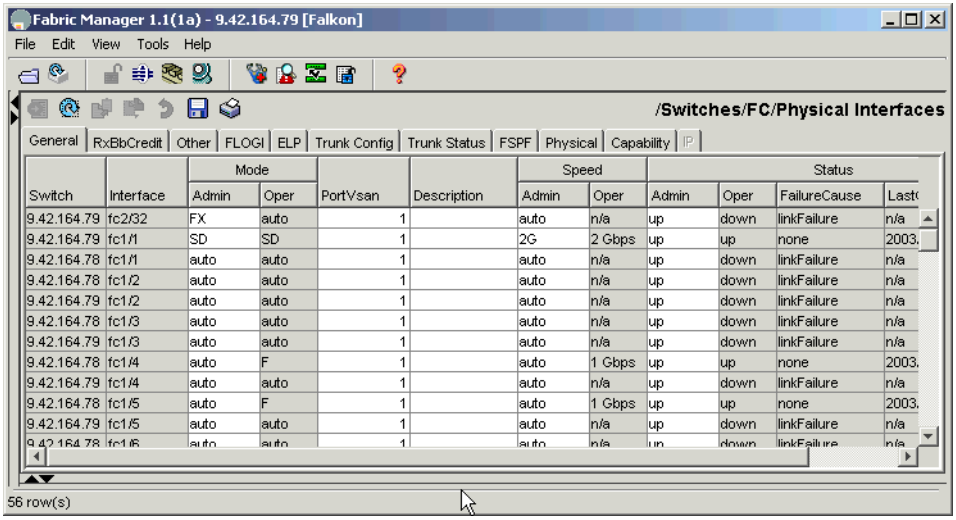


*Figure 2-79 Physical Interfaces view*

To make it easier to find the interfaces you are interested in, the interfaces can be sorted by the values of any column by clicking the column header.

You can change the fields with white background. The fields with gray background cannot be changed. For some fields, such as Mode, Speed and Status, you get a pull-down menu of possible values when you click the field. For others, such as PortVsan and Description, you may type the value in yourself.

Changes are not immediately applied to the switch. Instead, they are shown in bold, as shown in Figure 2-80. After you have made all the changes, you can apply them to the switches by clicking the **Apply Changes** button. If you decide that you do not want to save the changes, you can click the **Refresh Values** button instead to reset the fields to their original values.

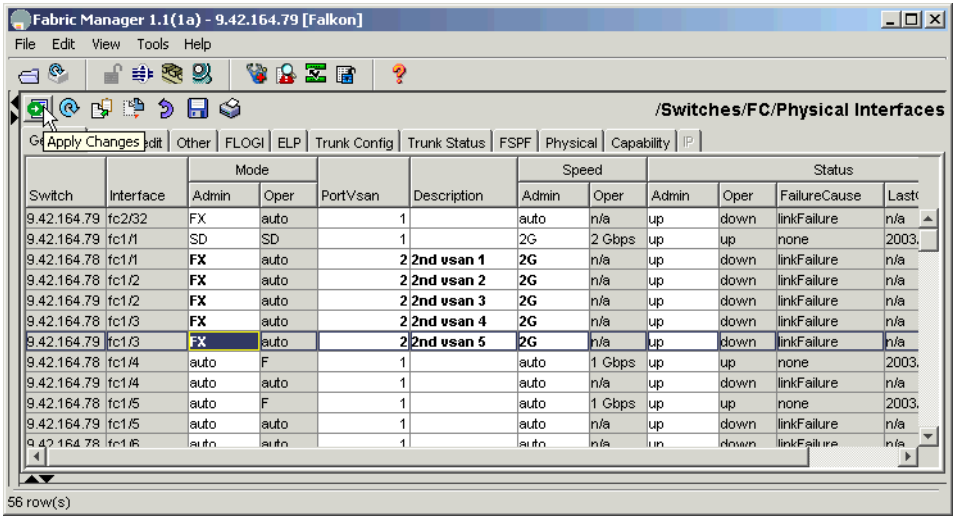


*Figure 2-80 Physical Interfaces view with changes*

For some fields, only some values are valid. For example, the PortVsan field has to be set to the VSAN number of a VSAN defined in the switch where the port is located. If you give the field an invalid value, you will get an error message as shown in Figure 2-81, when you click the **Apply Changes** button.

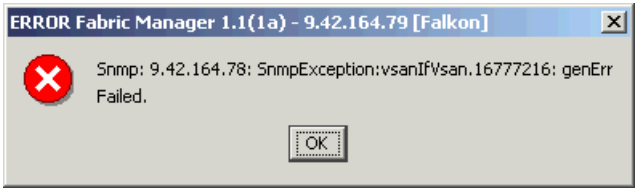


*Figure 2-81 Physical Interfaces error message*

## Creating a new PortChannel

We will start with the configuration presented in Figure 2-82, with four separate EISL links between the switches shared by several VSANs. Our goal is to build a PortChannel using two of the ISL links, and share the PortChannel for all of the VSANs.

**Note:** While it is possible to use all four links to create the PortChannel at once, this causes disruption to the traffic. Therefore we create the PortChannel first with two links, and add the remaining two links to the PortChannel when it is operational.

1. Click one of the existing (E)ISLs with the right mouse button, and choose the option **Add to Port Channel** from the menu that opens.

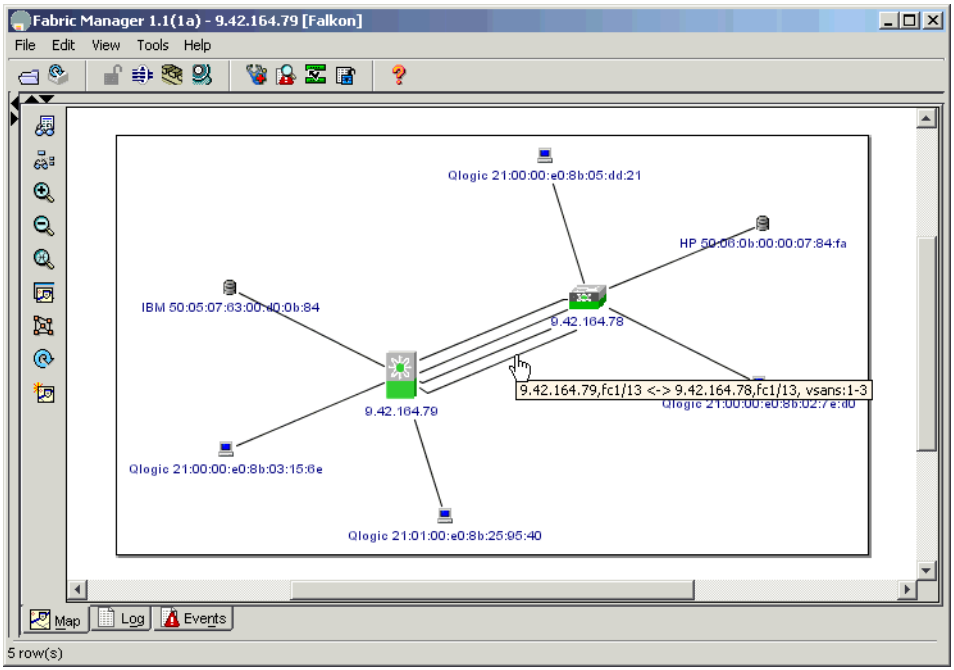


*Figure 2-82 Creating a PortChannel 1*

2. Select the (E)ISLs you want to use for the PortChannel, as shown in Figure 2-83, and click **Next**.

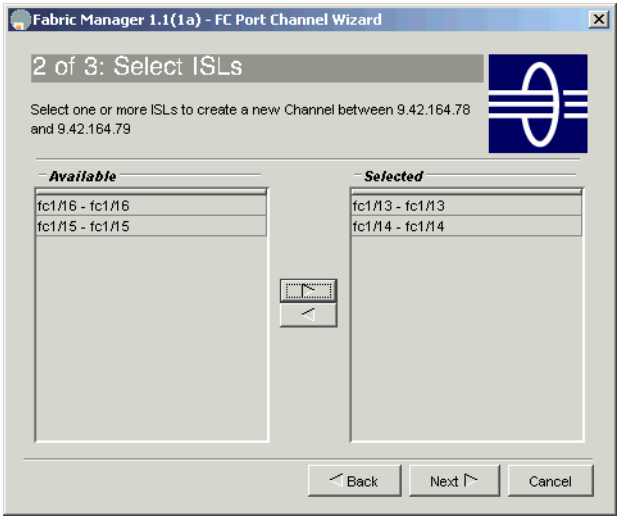


*Figure 2-83 Creating a PortChannel 2*

3. Choose the parameters for the new PortChannel, as shown in Figure 2-84. Note that you have to choose the correct Trunk Mode parameter here, if you want to use the PortChannel for multiple VSANs. Click **Finish**, when finished.

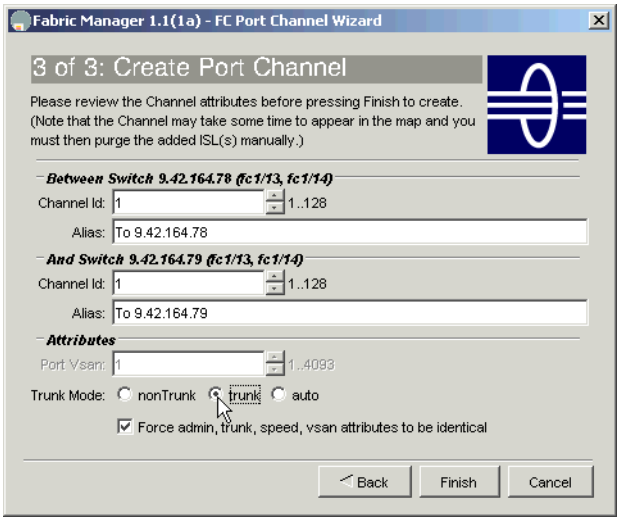


*Figure 2-84 Creating a PortChannel 3*

4. The PortChannel is now being created. However, it takes some time before the map display reflects the change. Click the **Refresh Map** button, until you see the PortChannel in the map display, as shown in Figure 2-85. The PortChannel is shown as a slightly thicker line than a normal link.

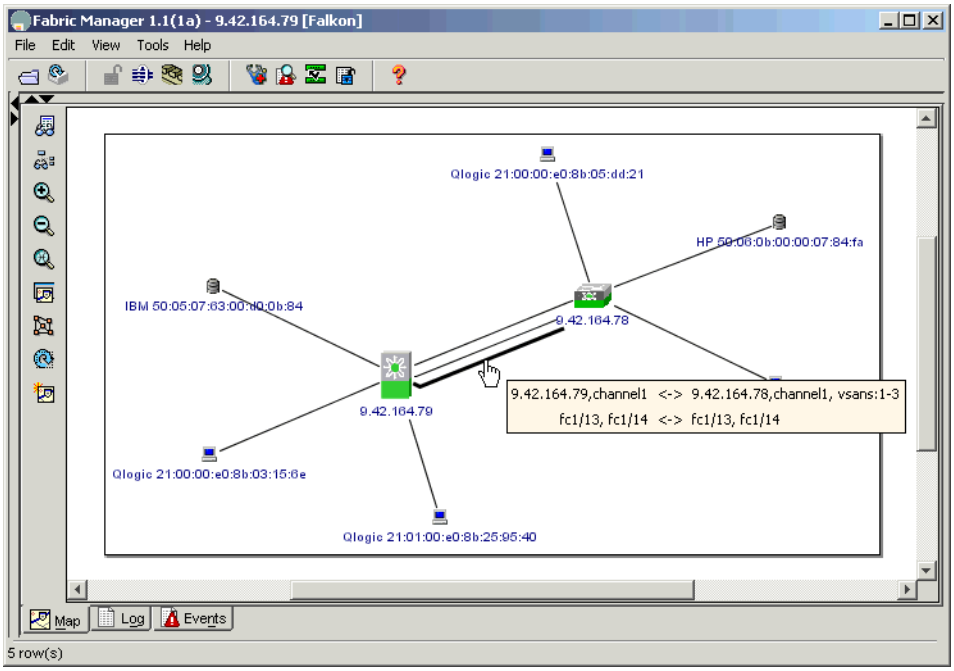


*Figure 2-85 PortChannel with two ports*

## Adding new ports to the PortChannel

Now that the PortChannel is operational, we want to add the remaining two EISL links to it.

1. Click PortChannel with the right mouse button, and choose the option **Edit...** from the menu that opens, as shown in Figure 2-86.

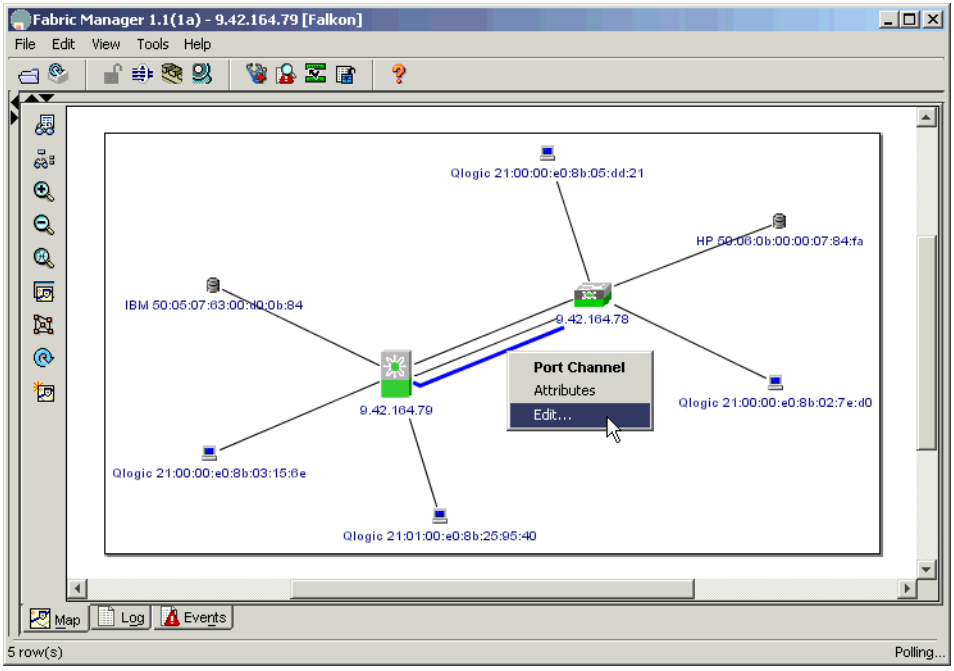


*Figure 2-86 Adding ports to PortChannel 1*

2. Add the remaining two ports to the window, as shown in Figure 2-87. Click **Finish**, when finished.

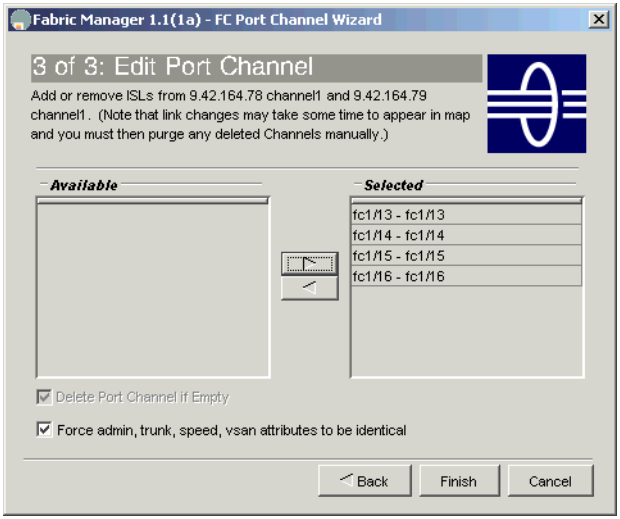


*Figure 2-87 Adding ports to PortChannel 2*

3. The ports are now added to the PortChannel. However, it takes some time before the map display reflects the change. Click the **Refresh Map** button, until you see only the PortChannel in the map display, as shown in Figure 2-88.

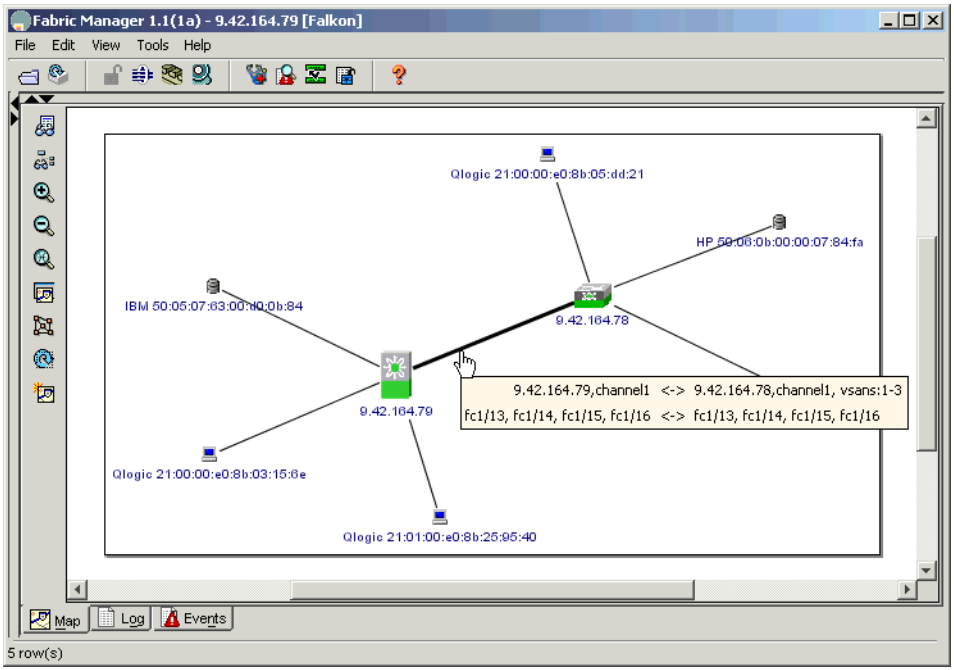


*Figure 2-88 PortChannel with four ports*

### Removing links from a PortChannel

You can also remove links from a PortChannel with the same procedure that you used to add links to it. When the last link is removed, the PortChannel is removed as well.

## 2.4.8 Managing events and alarms

The Cisco MDS 9000 series switches can use a wide range of methods for reporting and responding to network events. The methods available include:

- SNMP events
- RMON alarms
- Call Home
- Syslog

## SNMP events

The SNMP events are preconfigured notifications, including SNMPv2 traps and SNMPv3 informs. You can use the SNMP facility to send events to an external network management console.

To edit the SNMP configuration, choose this path from the physical menu: **Switches—>Events—>Notifications/Traps**. You will see a window listing the currently active SNMP destinations, as shown in Figure 2-89.

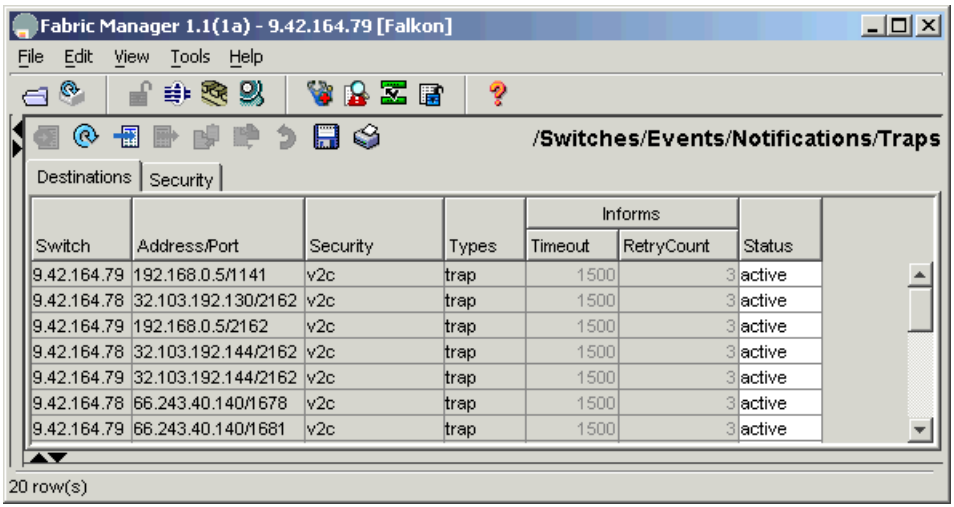


*Figure 2-89 SNMP destinations*

To add a new SNMP destination, click the **Create Row...** button. You will see a window as shown in Figure 2-90.

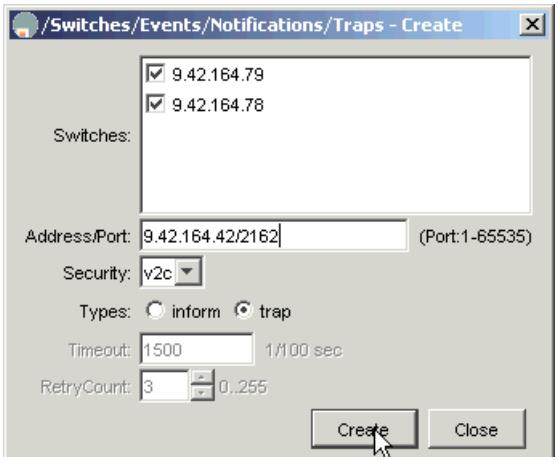


*Figure 2-90 Add SNMP destination*

Enter the parameters of your SNMP server here, and click the **Create** button to create the SNMP destination.

There are three predefined security settings: v1, v2c, and v3. You can also configure event security by clicking the **Security** tab. However, we recommend that you only do this if you have experience with SNMPv3.

You can delete an SNMP server by clicking the **Delete Row** button.

Event filters allow you to configure what SNMP messages you want to send. To configure event filters, choose the path **Switches—>Events—>Filters**. You can use the **FC** tab to configure filters for the Fibre Channel messages as shown in Figure 2-91, and the **Other** tab to configure filters for other messages as shown in Figure 2-92.

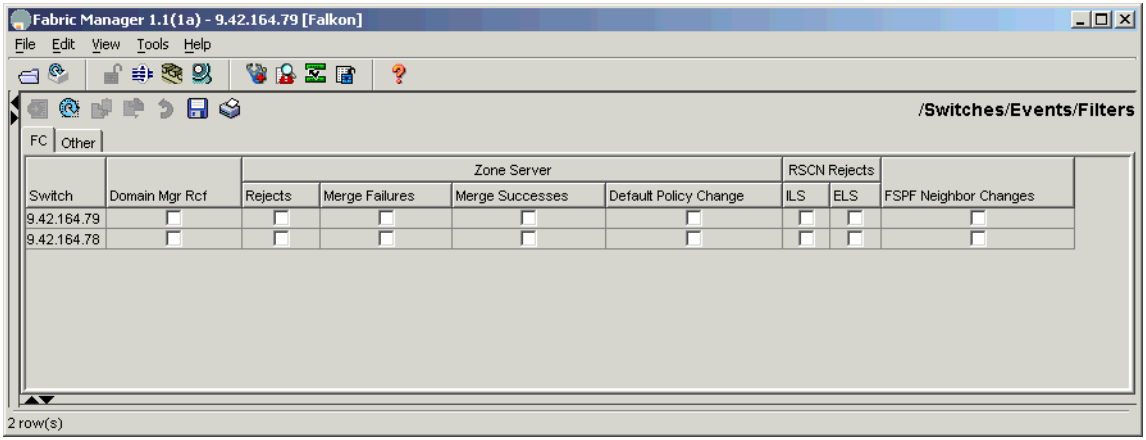


*Figure 2-91 Filters for FC messages*

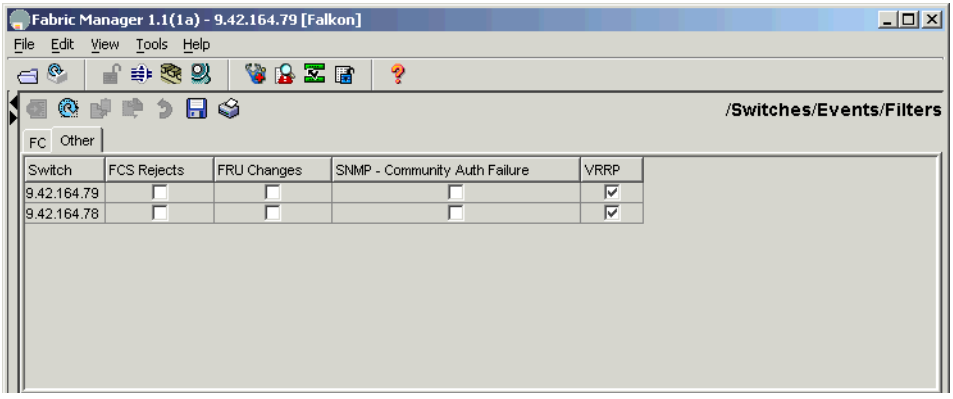


*Figure 2-92 Filters for other messages*

## RMON alarms

The RMON alarms are configurable notifications that you can set based on thresholds for various network events.

You have to enable RMON alarms by port from the Device Manager using the following procedure:

1. Choose **Events—>Threshold Manager** from the main menu and click the **Ports** tab. You will see a window as shown in Figure 2-93.

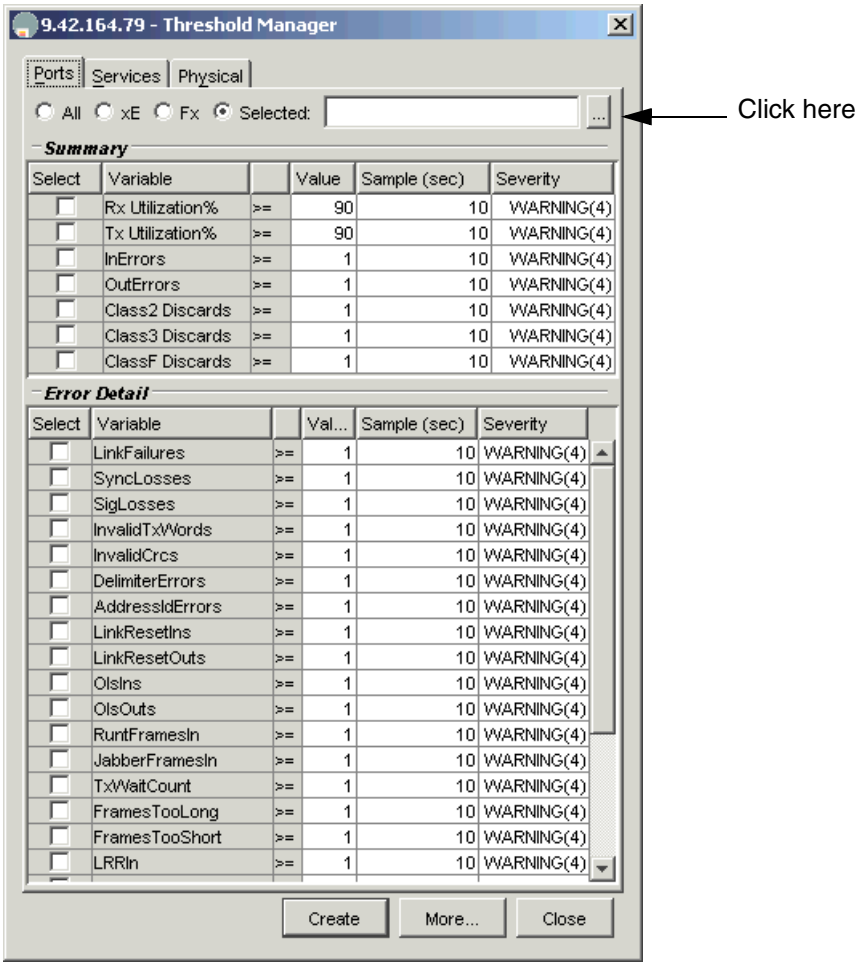


*Figure 2-93 Threshold manager Ports tab*

2. Select the ports you want to monitor by choosing the **Selected:** radio button, and clicking the **...** button on the right side. When you click the button, you will see a selection window as shown in Figure 2-94.

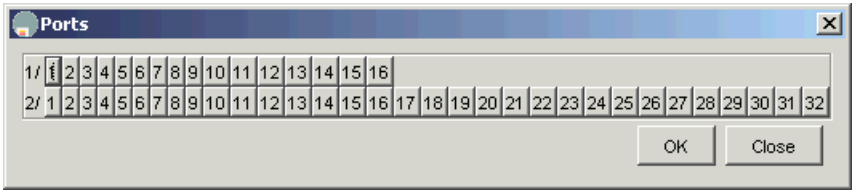


*Figure 2-94 Threshold manager port selection*

Select the ports you want to monitor, and press **OK**.

Alternatively, you can also choose one of the other radio buttons in the Threshold manager window. This way you can easily choose all ports, all Fx_Ports, or all xE_Ports.

3. Click the check box of each variable you want to monitor.
4. Enter the threshold value for each variable in the Value column.
5. Enter the sampling period for each variable in seconds.
6. Select one of the following severity levels for each alarm:
   - Fatal
   - Warning
   - Critical
   - Error
   - Information
7. Click **Create**.

   You see a confirmation window as shown in Figure 2-95.

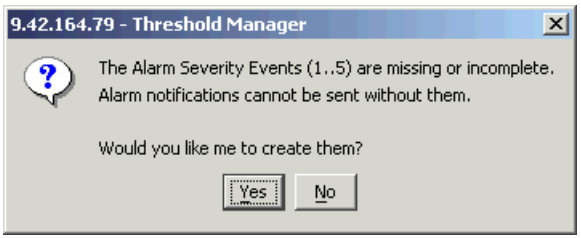


*Figure 2-95 Threshold manager confirmation window*

8. Click **Yes** to create the alarm.

   If you click **No**, the system defines only a log event, and does not send alarms.

You can also create alarms for VSAN events in the same way by clicking the **Services** tab, and for physical components by clicking the **Physical** tab. The **Services** tab is shown in Figure 2-96, and the **Physical** tab in Figure 2-97.

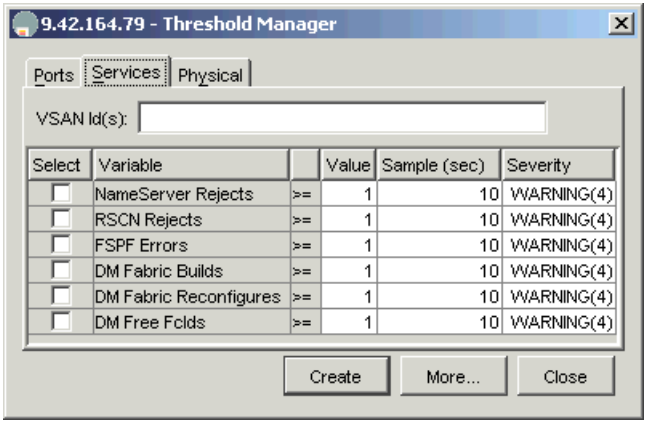


*Figure 2-96 Threshold manager Services tab*

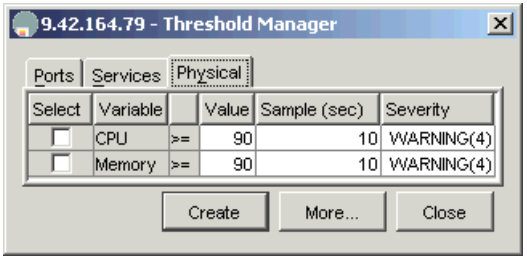


*Figure 2-97 Threshold manager Physical tab*

You can change the default controls for the RMON alarms by clicking the **More** button in the Threshold manager window. You will see a window as shown in Figure 2-98.

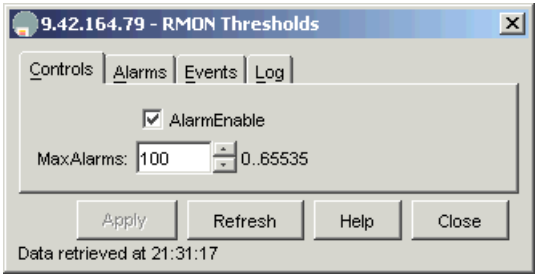


*Figure 2-98 RMON thresholds Control window*

By clicking the **Alarms** tab you can view the defined alarms as shown in Figure 2-99.

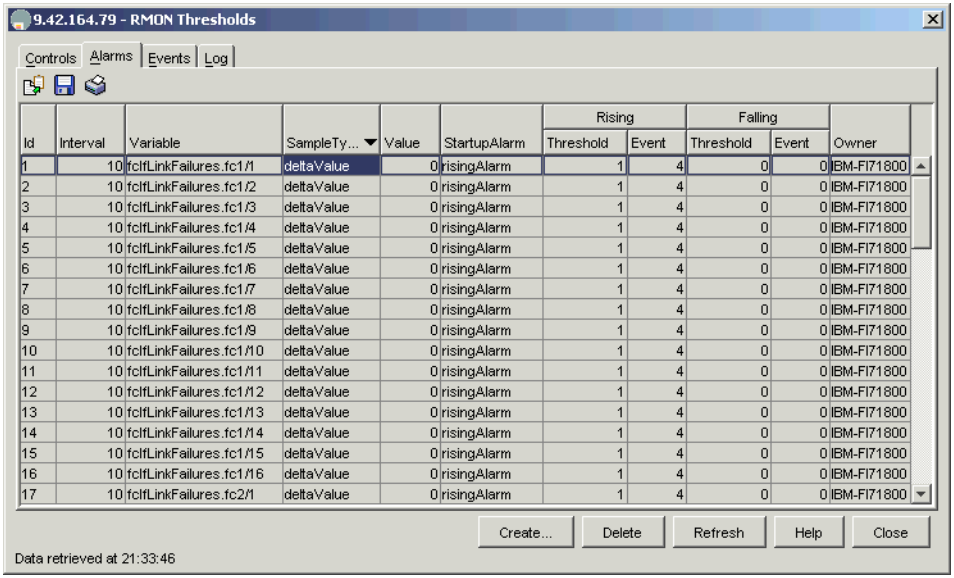


*Figure 2-99 RMON thresholds Alarm window*

## Call Home

The Call Home feature allows you to configure automatically generated e-mail messages or other responses to specific events. You can use the Call Home for direct paging of a network support engineer, e-mail notification to a network operations center, and utilize the Cisco AutoNotify services to automatically create a case with the Technical Assistance Center (TAC).

The Call Home feature supports the following message formats:

- Short Text
- Plain Text
- XML

The Short Text message format is suitable for pagers and printed reports. The Plain Text message is the best format for any other messages that are read by humans, such as e-mail.

The XML message is a machine readable format, and is used for communication between applications, as well as for communication to the Cisco TAC.

The system allows up to 50 concurrent e-mail destinations for each message format. The message categories supported include system, environment, switching module hardware, services module hardware, supervisor module hardware, inventory, and test.

If you want to configure the settings for the Call Home feature, choose the path **Switches—>Events—>Call Home** in the physical menu, and click the **General** tab. You will see a window as shown in Figure 2-100.

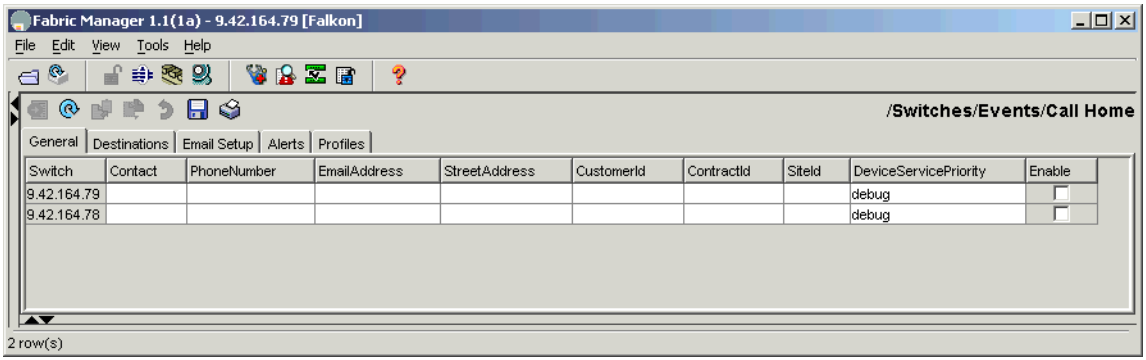


*Figure 2-100 Call Home General tab*

In the **General** tab you can fill in the contact information for all your switches.

To view the list of defined Call Home destinations, click the **Destinations** tab. You will see a window as shown in Figure 2-101.

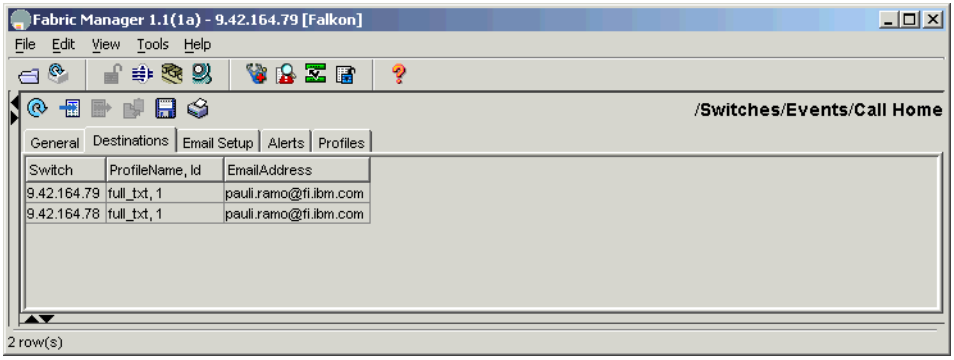


*Figure 2-101 Call Home Destinations tab*

If you want to add a new destination click the **Create Row...** button. You will see a window as shown in Figure 2-102.

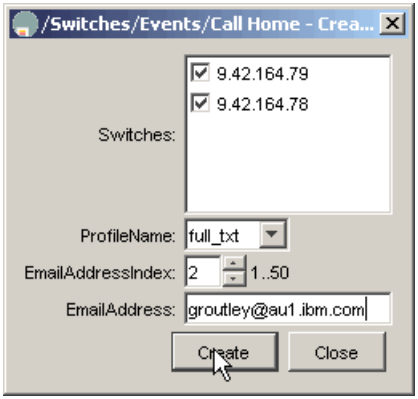


*Figure 2-102 Adding a call home destination*

You can choose the message format by choosing the corresponding profile name. You have to give the e-mail address a unique index in the range 1-50.

Click **Create** to create the new destination. The Create Destination window remains open allowing you to create additional destinations. Click Close when you have finished creating the destinations.

You also need to configure the e-mail setup parameters in the **Email Setup** tab, as shown in Figure 2-103.

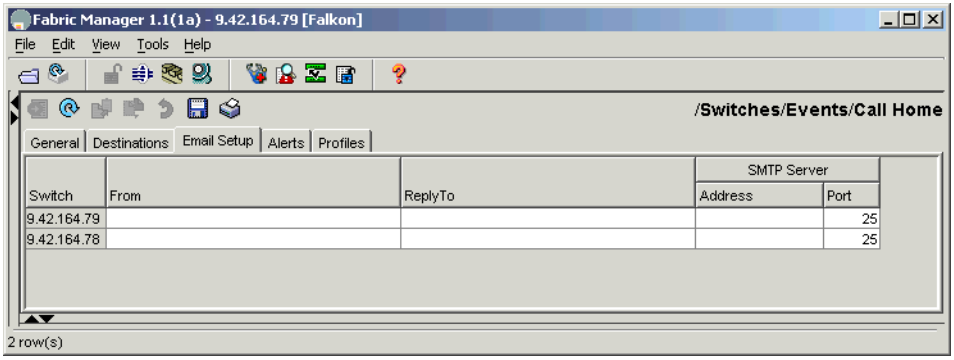


*Figure 2-103 Call Home Email Setup tab*

The From and Reply To fields are included in the headers of all e-mails sent. You also need to give the address of an e-mail server, and the port the server is listening to, if it is not the standard SMTP port (25).

You can send a test e-mail with the **Alerts** tab as shown in Figure 2-104.

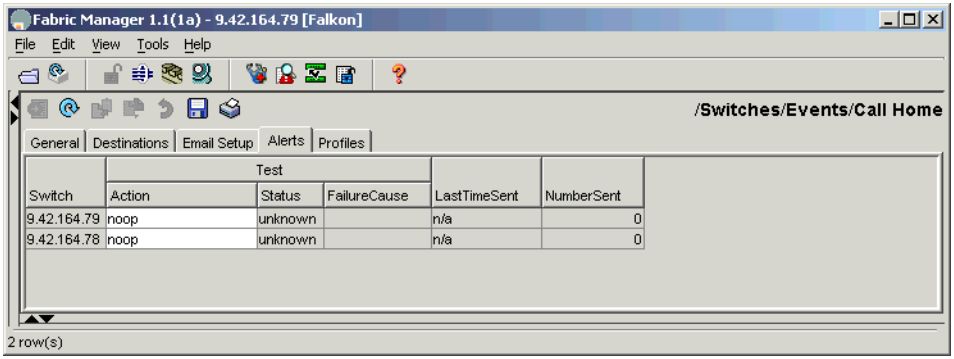


*Figure 2-104 Call Home Alerts tab*

Choose the type of the test message to send in the Action column, and click the **Apply Changes** button.

Finally, you can set the maximum message size for various message formats in the **Profiles** tab as shown in Figure 2-105.

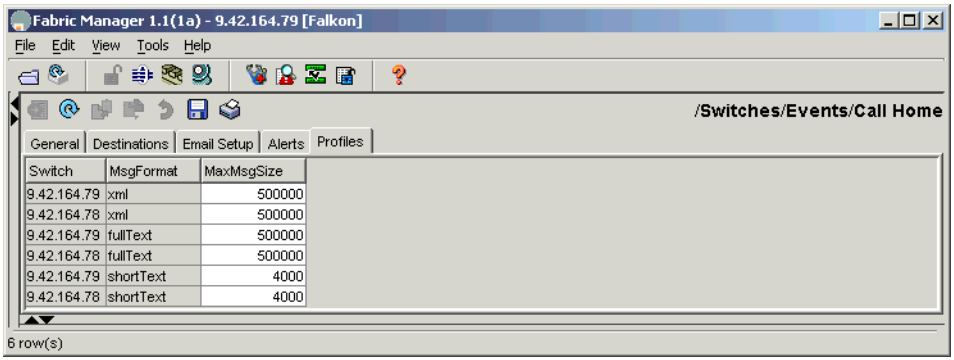


*Figure 2-105 Call Home Profiles tab*

## Syslog

The syslog facility is used for internal error reporting in the switch, and can also be used to report errors to external syslog servers.

To edit the syslog attributes, choose the path **Switches—>Events—>SysLog** in the physical menu, and click the **General** tab. You will see the general syslog settings window, as shown in Figure 2-106.

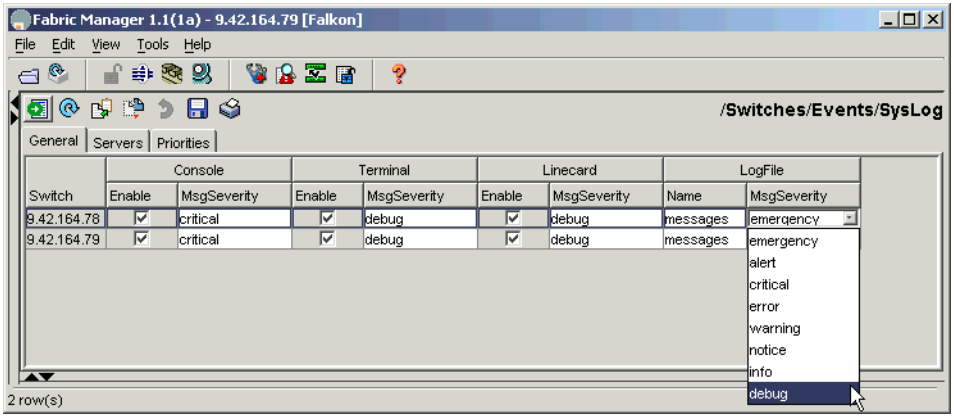


*Figure 2-106 Syslog general settings*

In the **General** tab, you can change the parameters for internal error reporting destinations. You can enable or disable each of them, and select the minimum severity level of messages for each destination. The severity levels range from emergency (highest) to debug (lowest).

You can configure up to three external servers to receive syslog messages from the switch. To configure external servers, use the **Servers** tab as shown in Figure 2-107.

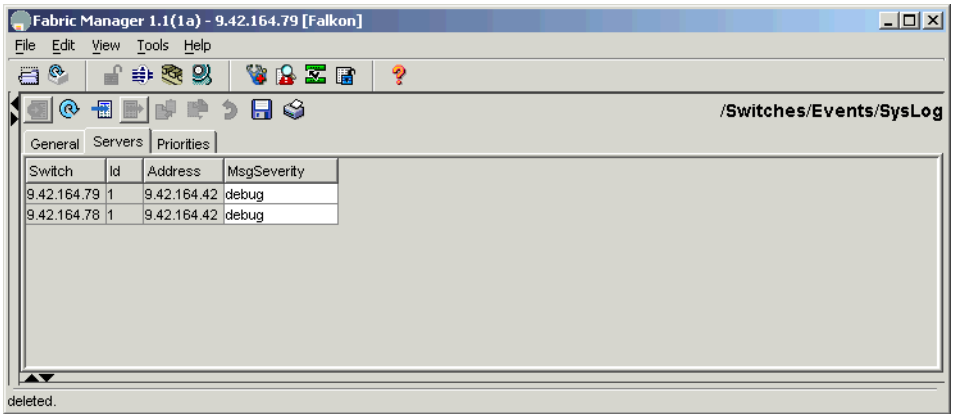


*Figure 2-107 External syslog servers*

If you want to add a new syslog server, click the **Create Row...** button. You will see a window similar to the one shown in Figure 2-108.

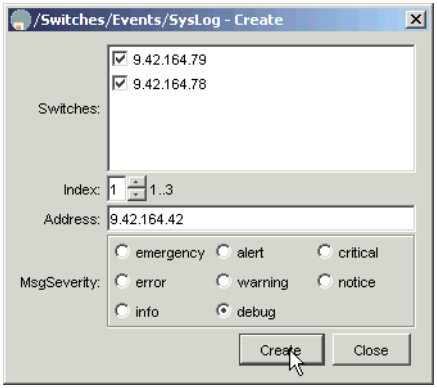


*Figure 2-108 Adding a syslog server*

Note that you have to give each syslog server a numeric index value, and two syslog servers cannot have the same index value. You can also specify the minimum severity level that is reported to a given syslog destination.

You can also delete a syslog server by clicking the **Delete Row** button.

You can edit the syslog message severity levels for the various facilities in the switch with the **Priorities** tab, as shown in Figure 2-109.

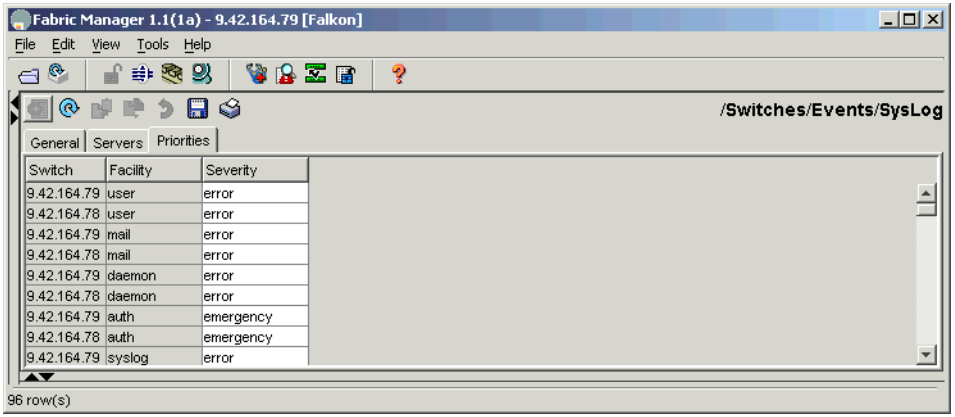


*Figure 2-109 Syslog severity levels*

## 2.4.9 Managing the system and components

You can see the system attributes of the switches in your fabric by choosing the path **Switches** in the physical menu. The attributes are listed in a window as shown in Figure 2-110.

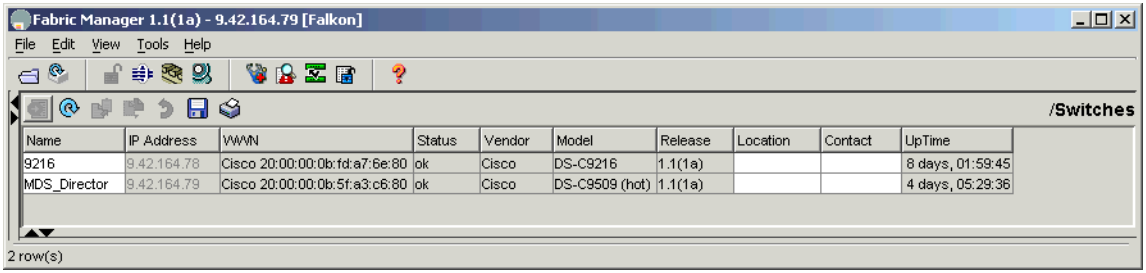


*Figure 2-110 Switches*

You can see the information and change some of the settings of the different modules in the switches in your fabric by choosing this path from the physical menu: **Switches—>Modules**. If you click the **Inventory** tab, you will see a list of all the modules in the SAN fabric, as in Figure 2-111.

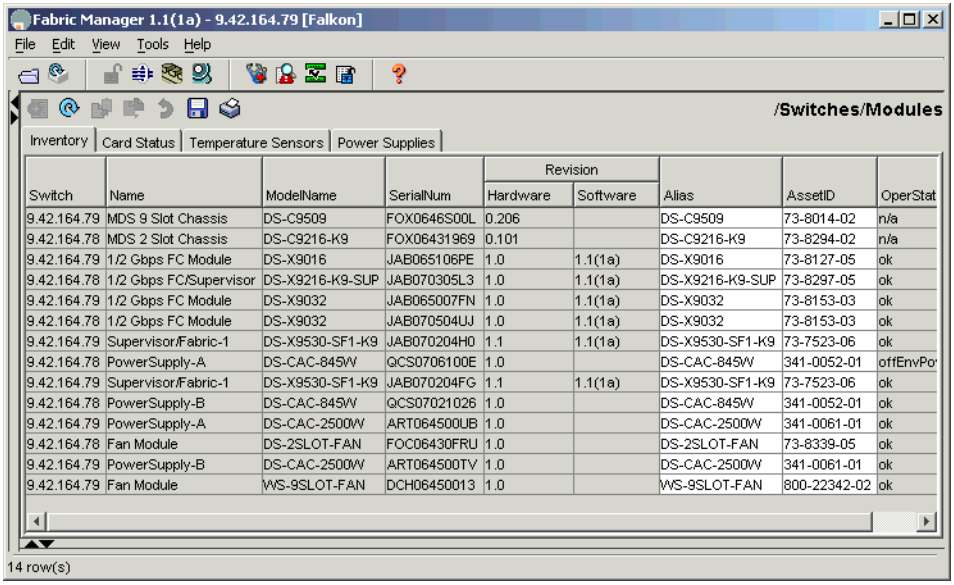


*Figure 2-111 Modules Inventory tab*

If you want to see the status of the line cards and supervisors in the fabric, click the **Card Status** tab. You will see a window as shown in Figure 2-112.

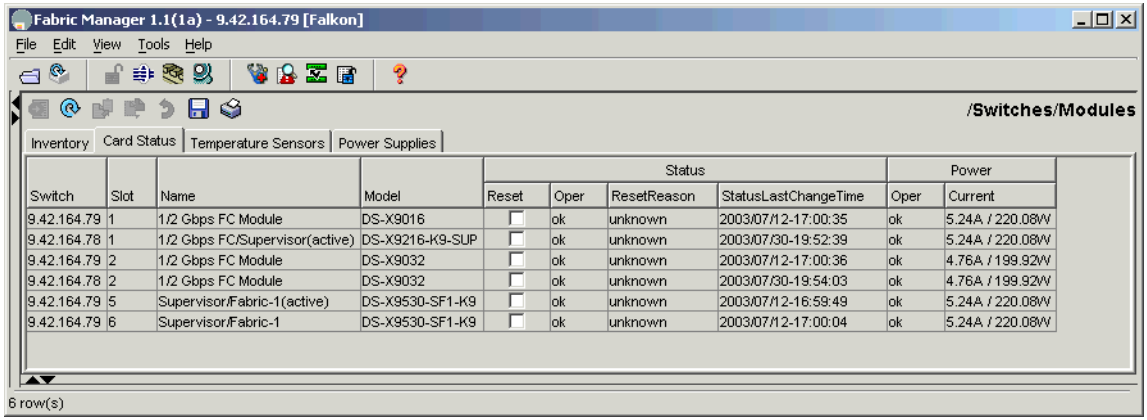


*Figure 2-112 Modules Card Status tab*

Note that the power usage of each line card and supervisor module is shown here. The power requirements of the fan modules are not listed here.

You can monitor the temperature sensors of the line cards and supervisors with the **Temperature Sensors** tab, as shown in Figure 2-113.

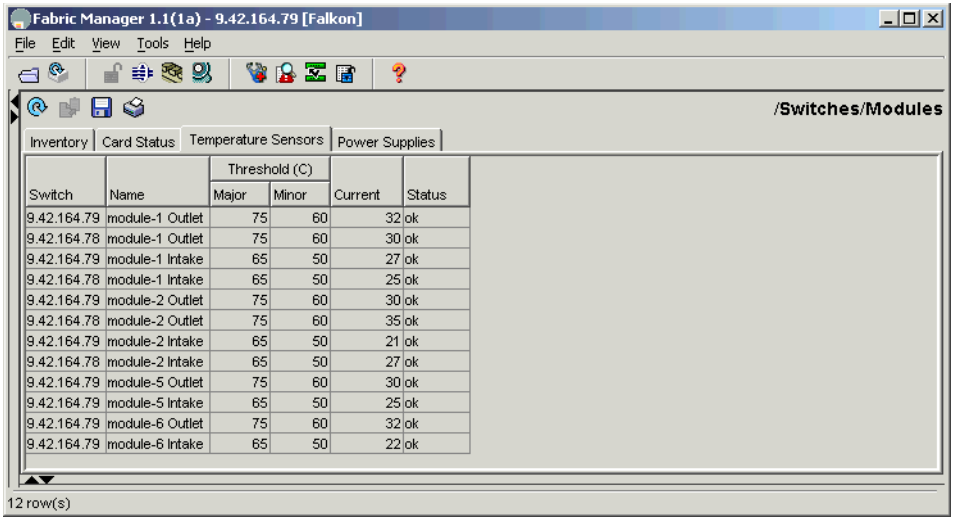


*Figure 2-113 Modules Temperature Sensors tab*

You can see the status of all power supplies of all switches in your SAN in the **Power Supplies** tab as shown in Figure 2-114.

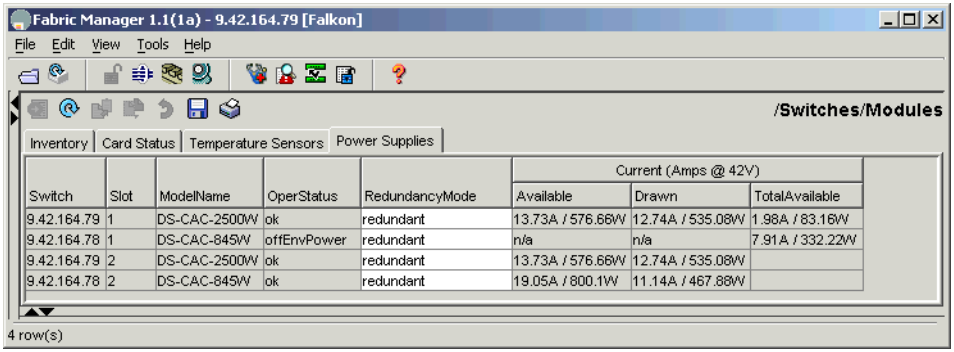


*Figure 2-114 Modules Power Supplies tab*

The power supply 1 of the switch 9.42.164.78 does not have input power in this example. The power supply 1 line for each switch contains the power available for any new modules. Note that the switch 9.42.164.79 does not have sufficient power available to accommodate new line cards.

If you want to change the power supply redundancy mode, you can do it here. However, we recommend that you always use redundant mode for maximum switch availability.

## 2.4.10 Managing IP storage services

The Cisco MDS 9000 series supports two different types of ethernet interfaces:

- Management ethernet (10/100 Mb/s, mgmt0)
- Gigabit ethernet ports on the 8-port IP line card

The management ethernet can only be used for management traffic. If you want to use the IP storage services, you have to have the 8-port IP line card installed in the switch.

The gigabit ethernet interfaces support virtual LANs (VLANs) using the IEEE 802.1Q standard for VLAN encapsulation. They also support combining two interfaces together into a single logical interface with the PortChannel feature.

**Note:** If you are connecting the gigabit ethernet port to a Cisco ethernet switch, and want to use VLANs, verify that the following requirements are met:

- The ethernet switch port is configured as a trunking port.
- The encapsulation is set to 802.1Q, and not ISL, which is the default.

Configuring the management interface to the same subnet with any of the gigabit ethernet interfaces is not supported. There are also other limitations on configuring the VLAN interfaces and the major interfaces into the same subnet. For more information, refer to the *Cisco MDS 9000 Family Configuration Guide*, DOC-7814893.

In the following sections we discuss the following topics:

- Configuring the ethernet interfaces
- Configuring VLANs on ethernet interfaces
- Configuring ethernet PortChannel
- Configuring VLANs on PortChannel interfaces
- Configuring FCIP

While the 8-port line card also supports the iSCSI protocol, we do not discuss the configuration of iSCSI here. The reason is that the iSCSI protocol is still very much work in progress, and configuring a gateway between iSCSI and Fibre Channel is complex enough, due to the differences in the security and device addressing models between Fibre Channel and iSCSI, to need a separate redbook.

## Configuring the ethernet interfaces

You can configure the IP addresses of the physical ethernet interfaces by choosing the path **Switches—>IP—>Physical Interfaces** from the physical menu. You can configure parameters, such as the IP address and maximum transmission unit (MTU).

Configuring the MTU to a higher value than the ethernet default of 1500 bytes allows a single ethernet packet to contain a maximum length Fibre Channel frame. However, if you want to use this functionality, ensure that your ethernet hardware supports these size frames.

The **General** tab of the physical IP interface window is shown in Figure 2-115.

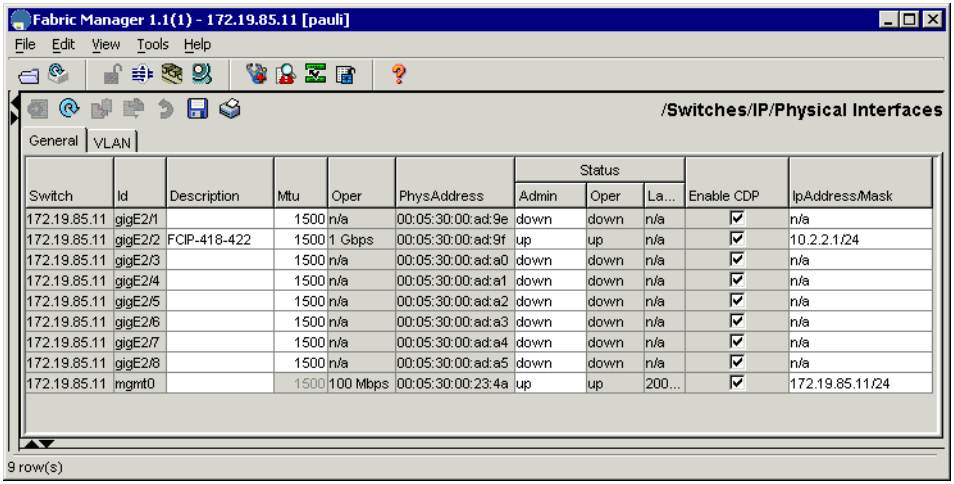


*Figure 2-115 IP interfaces - General tab*

## Configuring VLANs on ethernet interfaces

In the **VLAN** tab you can define any virtual LANs (VLANs) for the ethernet interfaces, if the interface is connected to an 802.1Q compliant switch. You can list a comma separated list of VLAN numbers for each physical interface here. Each VLAN has a separate subinterface. The **VLAN** tab is shown in Figure 2-116.

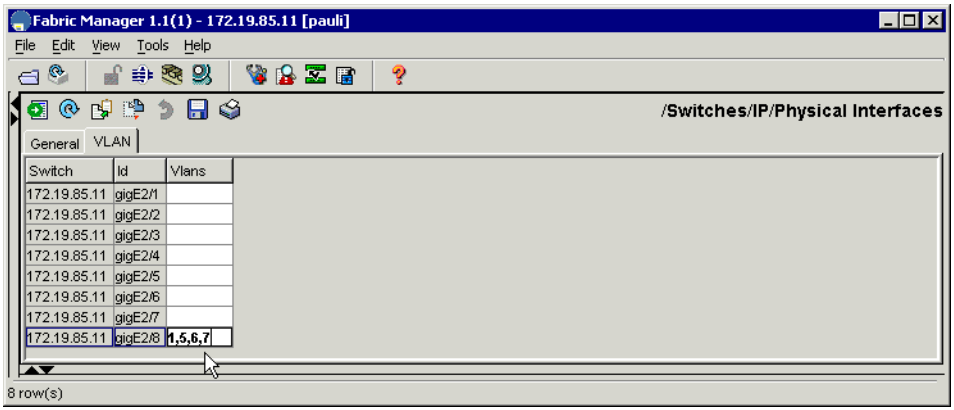


*Figure 2-116 IP interfaces - VLAN tab*

When you have the VLAN numbers defined for one or more physical interfaces, the VLAN subinterfaces are shown in the **General** tab, along with the physical gigabit ethernet interfaces, as shown in Figure 2-117.

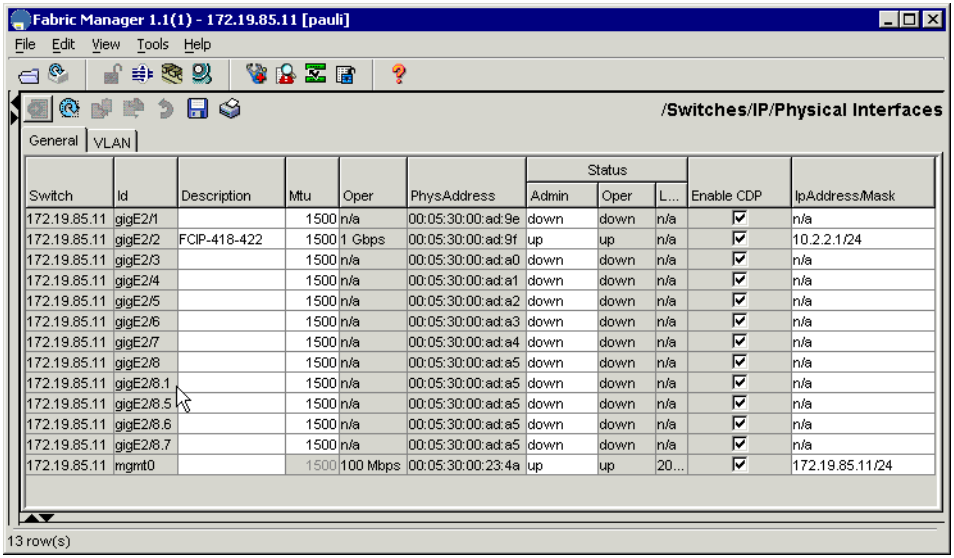


*Figure 2-117 IP interfaces - General tab with VLANs defined*

## Configuring ethernet PortChannel

With the PortChannel feature, you can combine two adjacent gigabit ethernet interfaces into a single logical interface. This requires that the interfaces are connected to a switch supporting the Cisco EtherChannel. The allowed combinations of ports for a PortChannel are 1-2, 3-4, 5-6, and 7-8.

To configure an ethernet PortChannel, choose this path from the physical menu: **Switches—>PortChannels**, and click the **Create Row...** button. You see a window, as in Figure 2-118.

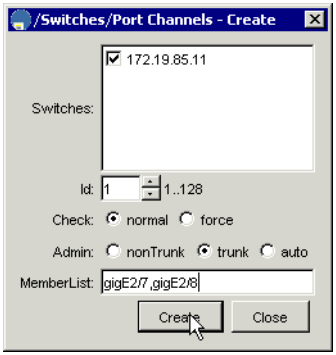


*Figure 2-118 Creating an ethernet PortChannel*

Give the PortChannel an unused number and list the physical interfaces that you want to include in the PortChannel, and click **Create** to create the PortChannel. The PortChannel is now listed in the list of PortChannels, as shown in Figure 2-119.

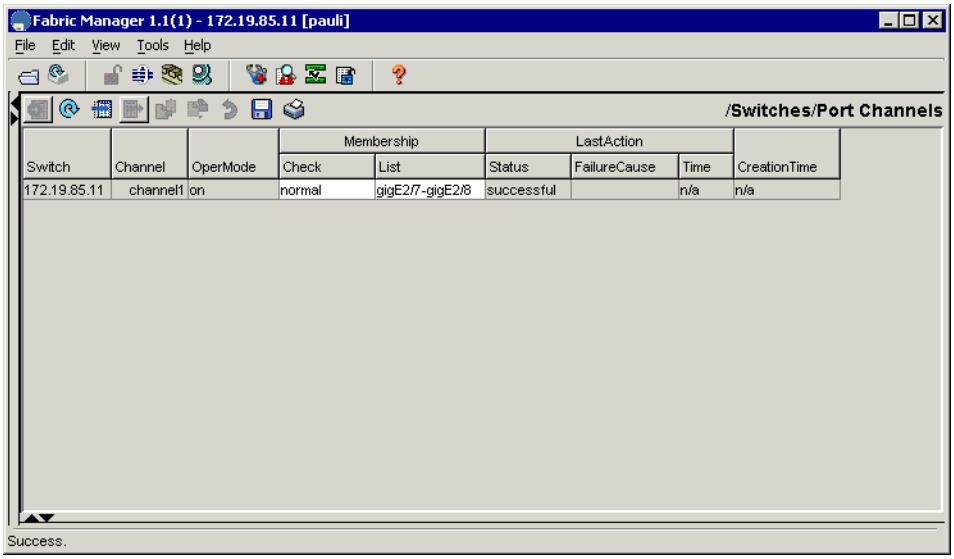


*Figure 2-119 List of PortChannels with an ethernet PortChannel defined*

When you want to configure the parameters for the PortChannel choose the path **Switches—>IP—Logical Interfaces** from the physical menu, as shown in Figure 2-120.

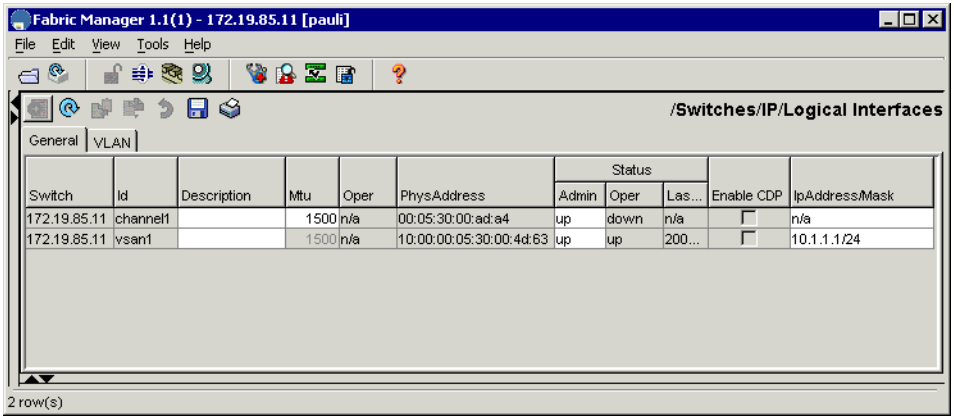


*Figure 2-120 List of logical IP interfaces with an ethernet PortChannel defined*

## Configuring VLANs on PortChannel interfaces

You can also configure VLANs for the logical PortChannel interface with the **VLAN** tab as shown in Figure 2-121.

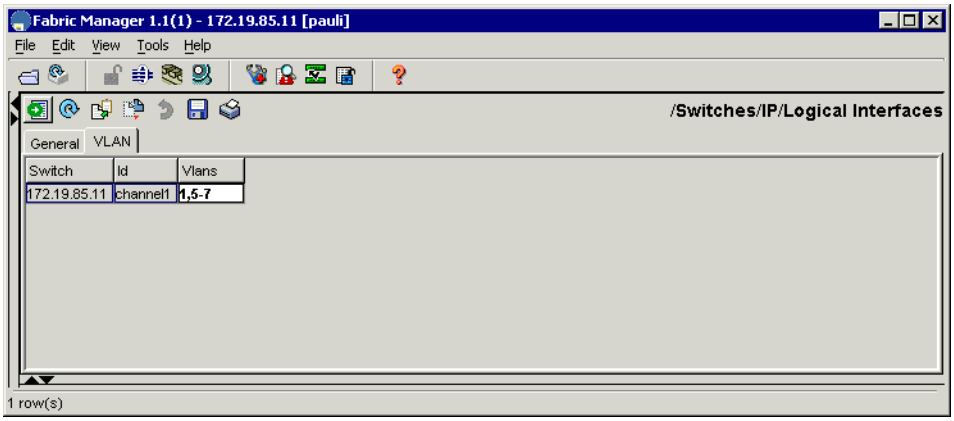


*Figure 2-121 Configuring VLANs on ethernet PortChannel interfaces*

List the VLANs for each PortChannel interface, and click the **Apply Changes** button. The logical subinterfaces for the VLANs are shown in the table of logical IP interfaces as shown in Figure 2-122.

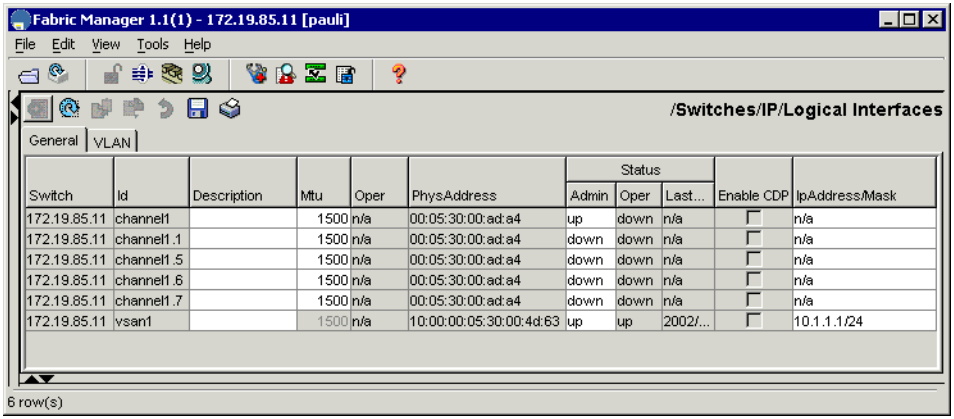


*Figure 2-122 List of logical IP interfaces*

## Configuring FCIP

Fibre Channel over TCP/IP (FCIP) allows Fibre Channel SAN islands to be interconnected over IP networks to form a single Fibre Channel fabric. These connections are called FCIP tunnels. Each gigabit ethernet interface in the Cisco MDS 9000 series can support up to three active FCIP tunnels at one time.

All of the procedures described below assume that the IP connectivity already exists between the gigabit ethernet ports you are going to use.

### *Creating FCIP tunnels with the FCIP wizard*

The easiest way to create a FCIP tunnel between two Cisco MDS 9000 switches with Fabric Manager is the FCIP wizard.

1. Start the FCIP wizard by choosing the path **Edit—>FCIP Tunnel** from the Fabric Manager main menu. You will see the wizard window as shown in Figure 2-123.

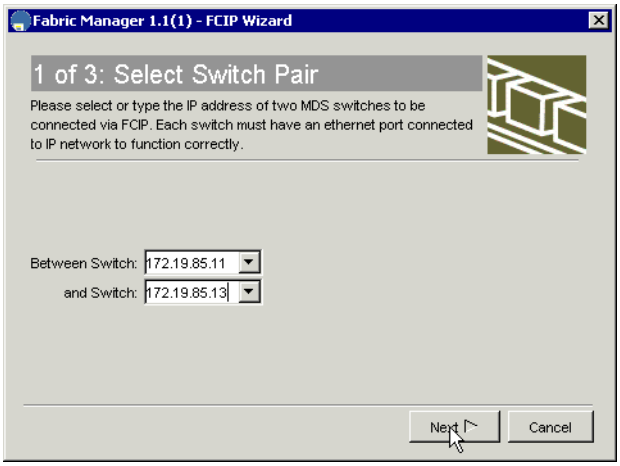


*Figure 2-123 FCIP tunnel wizard 1*

2. Enter the IP addresses of the switches in each end of the IP link and click **Next**. You see the second window of the wizard, as shown in Figure 2-124.

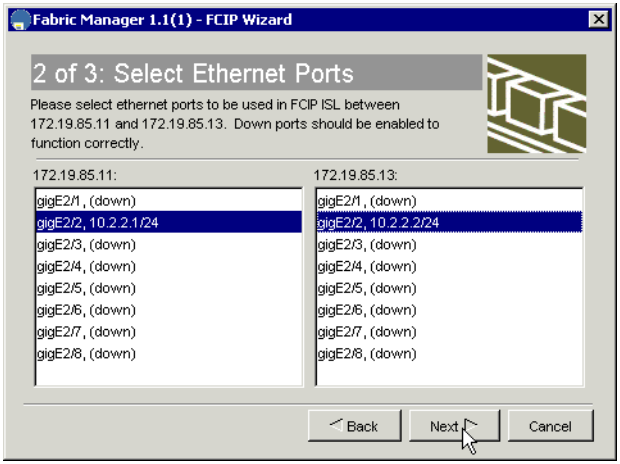


*Figure 2-124 FCIP tunnel wizard 2*

3. Choose the IP addresses for the gigabit ethernet interfaces, and click **Next**. You will see the third window of the wizard as shown in Figure 2-125.

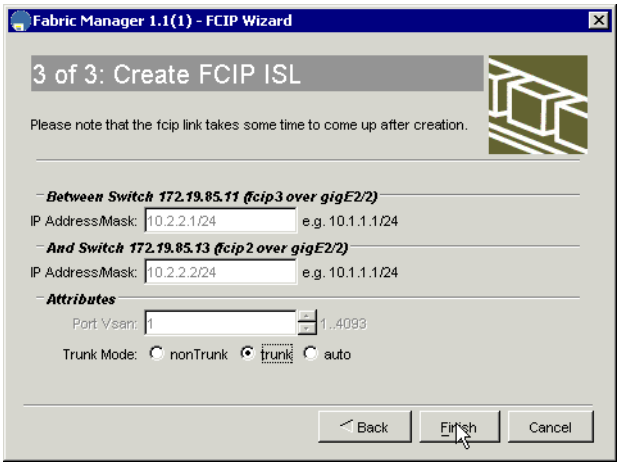


*Figure 2-125 FCIP tunnel wizard 3*

4. Specify the link attributes for the new link and click **Finish** to create the FCIP tunnel.

## 2.4.11 Managing advanced features

The Cisco MDS 9000 series switches implement a number of advanced features. We describe some of the most important of them here.

### Domain parameters

In the Cisco MDS 9000 family, each VSAN has its own set of domain parameters.

If you want to view the currently active domain parameters, choose the path **Switches—>FC—>Domain Manager** from the physical menu, and click the **General** tab. You will see a window as shown in Figure 2-126.

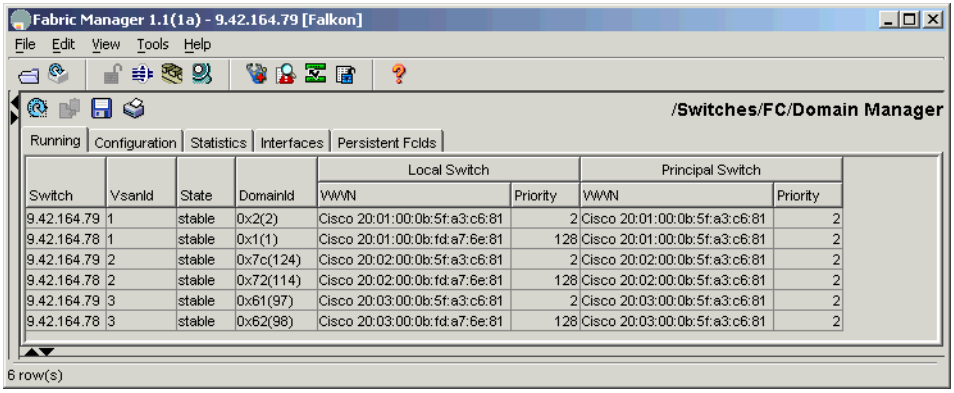


*Figure 2-126 Active domain parameters*

Note that each VSAN has a separate domain ID value for each switch.

If you want to change the parameters, click the **Configuration** tab. You will see a window as shown in Figure 2-127.

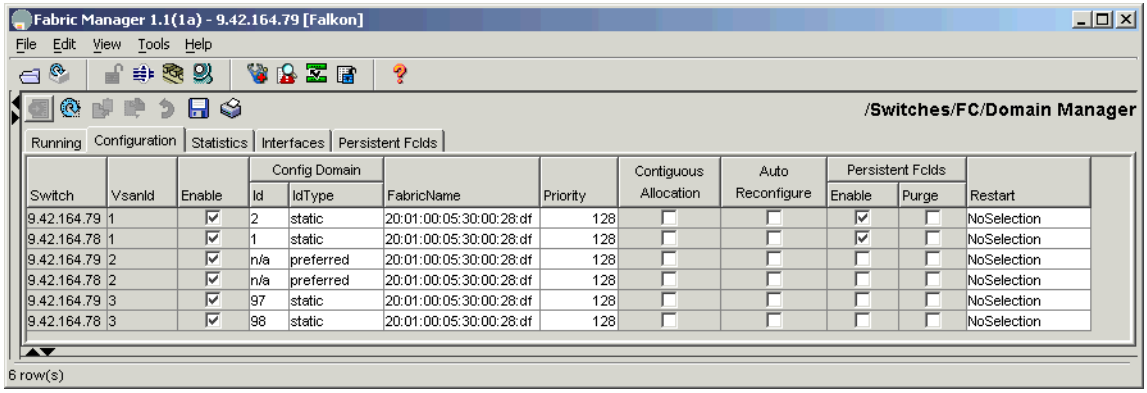


*Figure 2-127 Domain configuration*

Note that many of the options here, such as the domain IDs, cannot be changed non-disruptively.

You can activate the Persistent FcIds feature here. However, before you can activate the feature, your configuration has to fulfill the following pre-requisites:

- The switches in the VSAN must have static domain IDs defined.
- The defined static domain IDs must match the domain IDs in the **Running** tab.

**Note:** If you attach AIX or HP-UX hosts to a VSAN, you need to have persistent FC IDs enabled for that VSAN. This is because these operating systems use the FC IDs in device addressing. If the FC ID of a device changes, the operating system considers it to be a new device, and gives it a new name.

If you need to purge any non-active FC ID from the database, you can do it by clicking the check box in the Purge column and clicking the **Apply Changes** button.

If you define static domain IDs to a VSAN that is configured to run in interoperability mode, you have to make sure that you are only using domain IDs in the permitted range (97-127). Fabric Manager allows you to configure any static domain ID to the switches even in interoperability mode, but using a domain ID outside the permitted range causes the switch to become isolated from the fabric. However, this only affects the VSAN with the incorrect configuration.

To view the domain statistics for a domain, click the **Statistics** tab and you will see a window as shown in Figure 2-128.

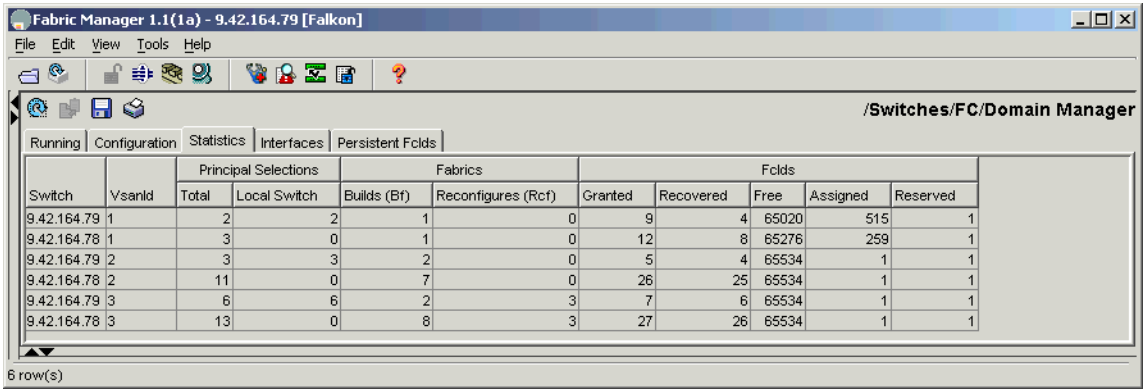


*Figure 2-128 Domain statistics*

The **Domain Interfaces** tab can be used to view the E_Ports between the switches in each VSAN as shown in Figure 2-129.

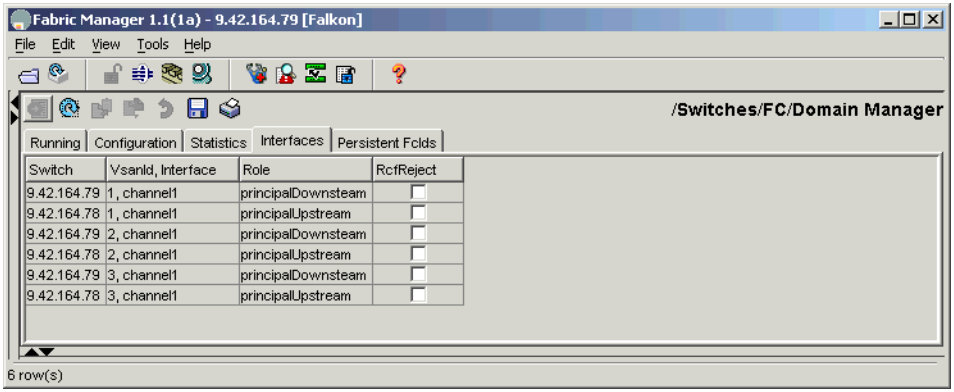


*Figure 2-129 Domain interfaces*

If you have the Persistent FcIds feature turned on for a VSAN, you can see the WWN to FC ID mapping information in the **Persistent FcIds** tab as shown in Figure 2-130.

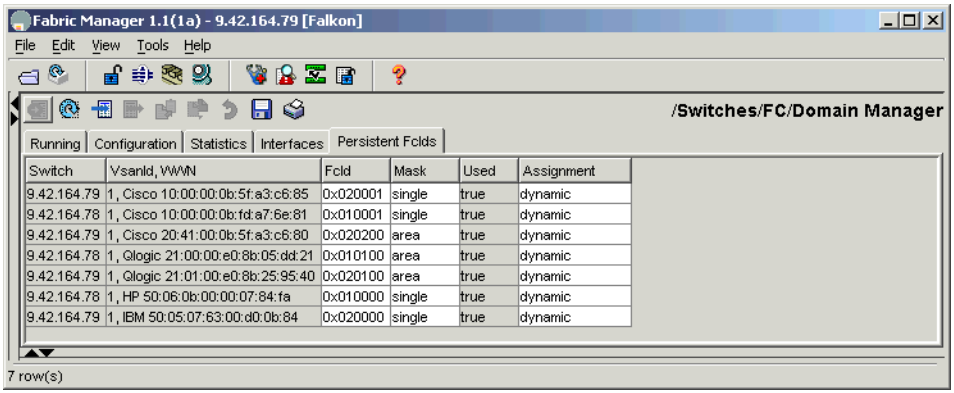


*Figure 2-130 List of persistent FcIds*

You can also define the mappings manually in this window. However, due to the added management complexity, we recommend using the automatic mapping.

## RSCN

Registered State Change Notification (RSCN) is a Fibre Channel service that is used to inform hosts about changes in the fabric. To receive these notifications, hosts must register with the fabric controller to receive them.

The following events can cause an RSCN to be sent:

- Disks joining or leaving the fabric
- A name server registration change
- A new zone enforcement
- IP address change
- Any other similar event, that affects the operation of the host

In addition to sending these events to registered hosts, a switch RSCN (SW-RSCN) is sent to all reachable switches in the fabric.

After receiving the notification it is up to the nodes to query the Name Server to obtain the new information. No details about the changed information are delivered by the switch in the RSCN. It is also up to the nodes to decide how the changes affect their operation, and how they implement the changes.

You can view the list of nodes currently registered to receive RSCN by choosing the path **Switches—>FC—>RSCN** in the physical menu and clicking the **Registrations** tab as shown in Figure 2-131.

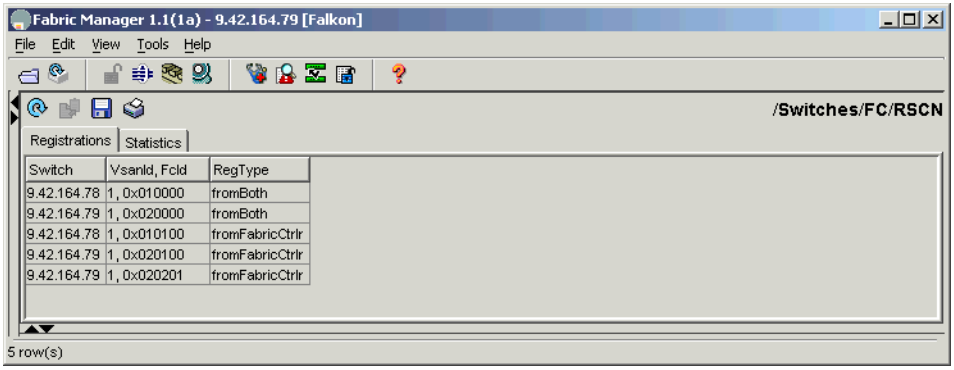


*Figure 2-131 RSCN Registrations*

You can also view the RSCN statistics for different VSANs by clicking the **Statistics** tab as shown in Figure 2-132.

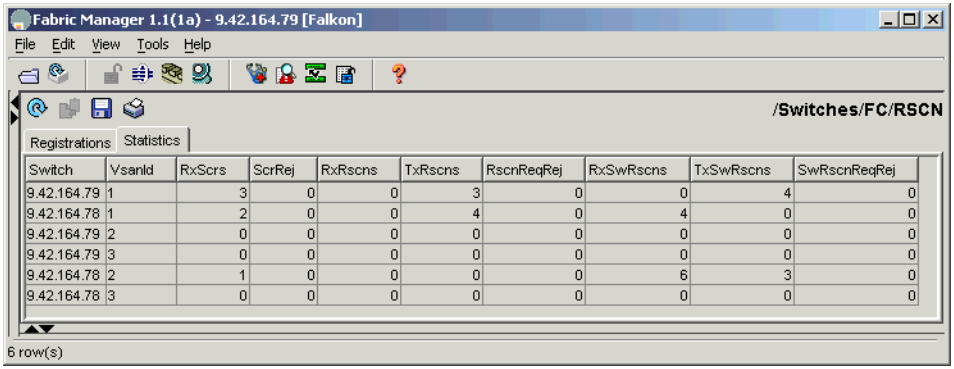


*Figure 2-132 RSCN Statistics*

## VRRP

The Virtual Routing Redundancy Protocol (VRRP) is a restartable application that can be used to provide redundant paths to a gateway switch. It can be used with all the different IP interfaces supported by the Cisco MDS 9000 family.

In our example we will use VRRP to implement a highly available management path to the switches in the SAN fabric. Our target environment is shown in Figure 2-133.

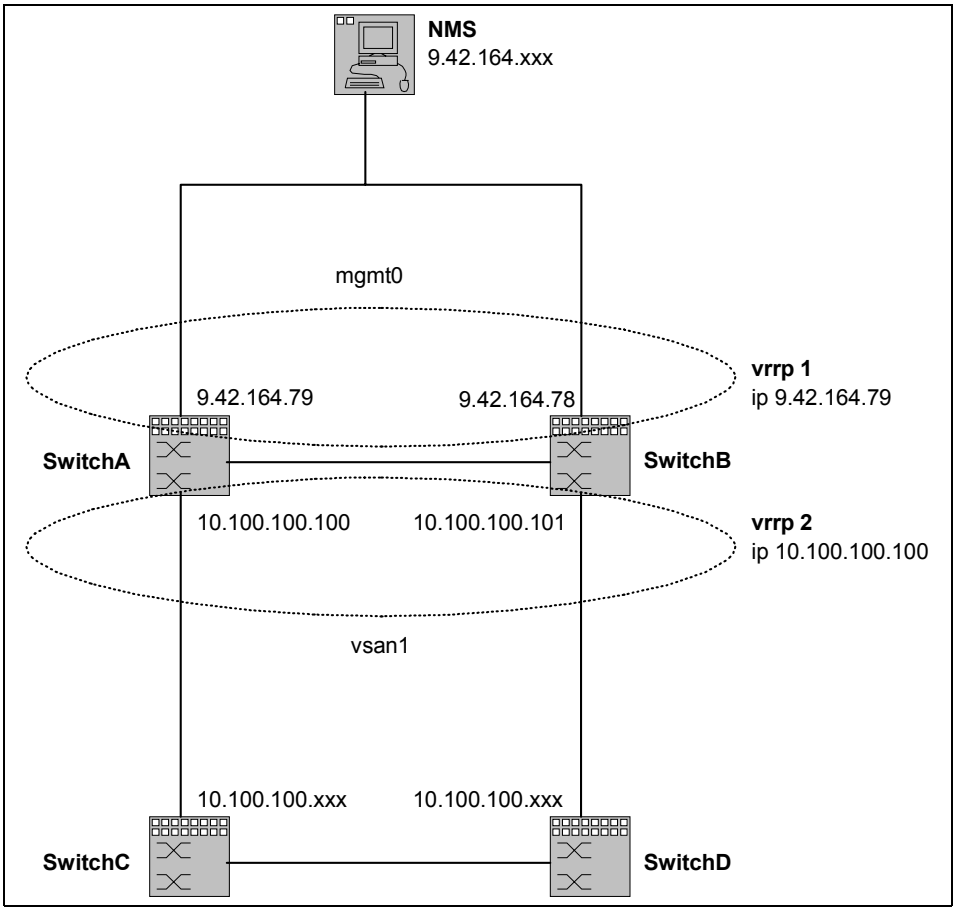


*Figure 2-133 VRRP target environment*

In our target environment the switch SwitchA is the VRRP master, and is normally working as a router and routing the management traffic to the virtual network, vsan1, running over the Fibre Channel fabric. SwitchB works as the VRRP backup switch.

The SwitchA address in the vsan1 network (10.100.100.100) is set as the default gateway address for SwitchC and SwitchD. If SwitchA fails, the other switches do not have to change their gateway settings to be able to communicate with the external network and the network management station (NMS).

In the same way, the gateway to the vsan1 network has to be set to the mgmt0 interface of SwitchA (9.42.164.79) for the NMS.

Since a single VRRP group can only contain interfaces of one type, we have to define two VRRP groups; one for the management ethernet interfaces (mgmt0), and one for the virtual SAN IP interfaces (vsan1). Note that only the vsan1 interfaces of SwitchA and SwitchB are included in the VRRP group.

We use the following procedure to implement our environment:

1. Create the VRRP groups by choosing the path **Switches—>IP—>VRRP** from the physical menu, and clicking the **Add Row...** button.
2. Fill in the interface and VRRP number for the first VRRP as shown in Figure 2-134.

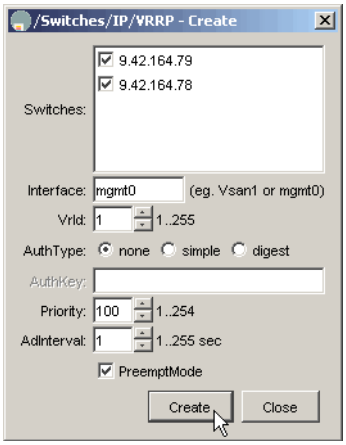


*Figure 2-134 Creating VRRP group for mgmt0*

3. Click **Create** to create the first VRRP group.
4. Change the interface and VRRP number to the values for the second VRRP, as shown in Figure 2-135.

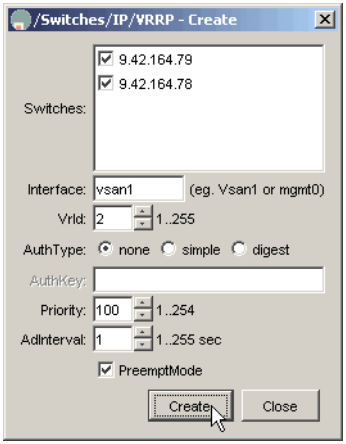


*Figure 2-135 Creating VRRP group for vsan1*

5. Click **Create** to create the second VRRP group.
6. Click **Close** to stop creating the VRRP groups.
7. You should see a list of the VRRP groups that you defined as shown in Figure 2-136.

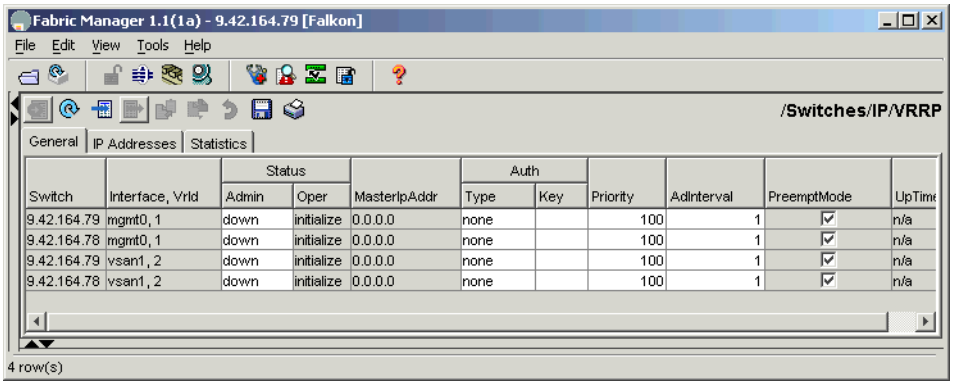


*Figure 2-136 List of VRRP groups*

8. Click the **IP Addresses** tab and click the **Add Row...** button to start defining the IP addresses for the VRRP groups.
9. Fill in the interface, VRRP number, and IP address for the first VRRP group, as shown in Figure 2-137.

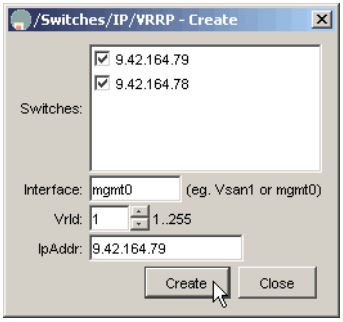


*Figure 2-137 Creating IP address for VRRP for mgmt0*

10. Click **Create** to create the IP address for the first VRRP group.
11. Change the interface, VRRP number, and IP address to the values for the second VRRP group as shown in Figure 2-138.

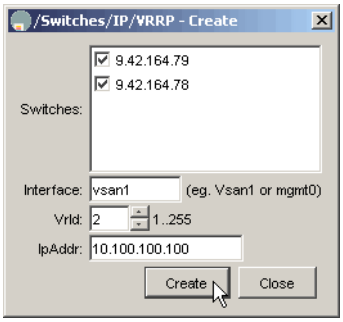


*Figure 2-138 Creating IP address for VRRP for vsan1*

12. Click **Create** to create the IP address for the second VRRP group.
13. Click **Close** to stop creating IP addresses.
14. You should see a list of the IP addresses of the VRRP groups you defined as shown in Figure 2-139.

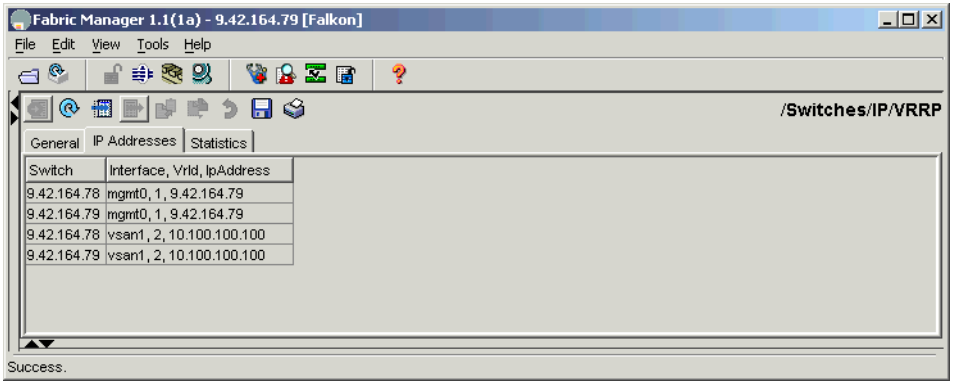


*Figure 2-139 List of defined IP addresses for VRRP*

15. Click the **General** tab to start your VRRP groups.
16. Change the admin status of all your VRRP groups to **up** as shown in Figure 2-140.

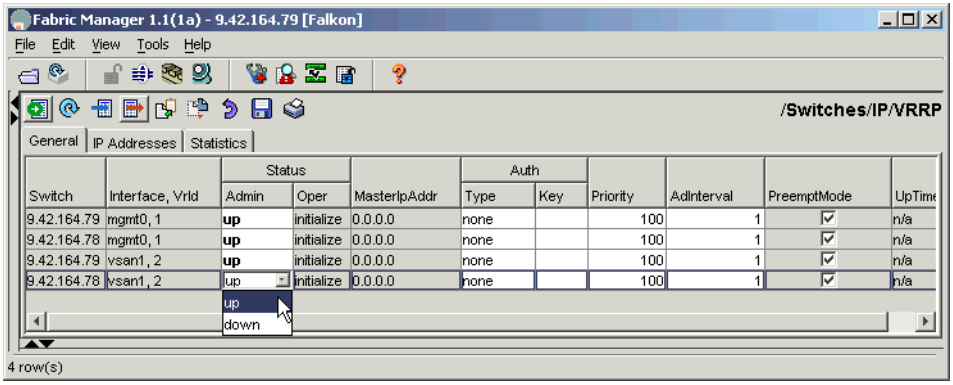


*Figure 2-140 Starting VRRP*

17. Click the **Refresh Values** button until you can confirm that your VRRP groups are active as shown in Figure 2-141.

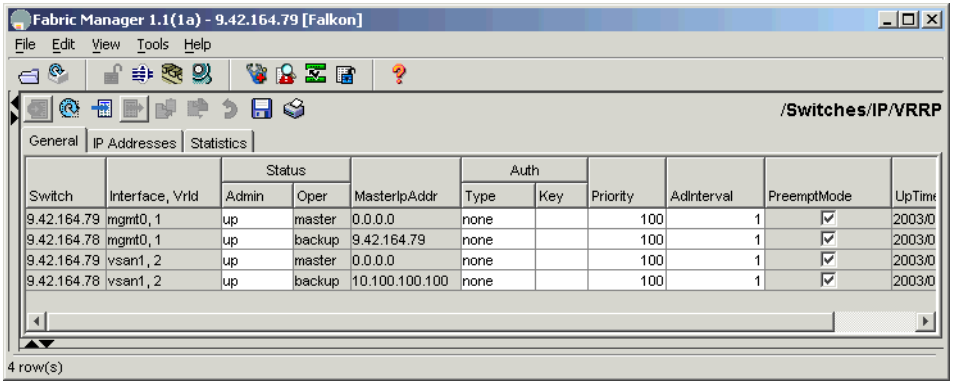


*Figure 2-141 VRRP active*

18. The VRRP definition is now complete.

## 2.5 Managing the Cisco SAN with the CLI

The Cisco MDS 9000 family has a very rich command line interface that should look familiar to anyone who has experience with any other Cisco product. The command line interface can be used to configure any parameter of the switch and fabric.

### 2.5.1 Getting started

You can login to the switch in several ways:

- Serial connection to the console port
- Telnet connection to the management ethernet interface
- SSH connection to the management ethernet interface

For network connections, we recommend using the SSH connection whenever possible because all traffic in a SSH connection is encrypted.

The commands can be abbreviated, as long as the command can still be distinguished from all other commands. For example, the **`config`** command can be abbreviated as **`con`**, but not as **`co`**, since that would confuse it with the **`copy`** command.

In addition, you can press the Tab key to complete a command.

## 2.5.2 CLI command modes

The Cisco 9000 family of switches have two main command modes:

- **EXEC mode:** EXEC mode enables you to temporarily configure terminal settings, perform basic tests, and display system information. Changes made in this mode are generally not saved across system resets. You are in EXEC mode immediately after logging to the switch.

  The command prompt for EXEC mode is **`switch#`**.

- **Configuration mode:** Configuration mode enables you to configure features that affect the whole system. Changes made in this mode are saved across system resets, if you choose to save them. You can get to configuration mode from EXEC mode with the **`config terminal`** command, and back to the EXEC mode with the **`end`** command or the **`exit`** command.

  The command prompt for configuration mode is **`switch(config)#`**.

The Configuration mode is further divided into several submodes that can be used to configure every parameter of the switch and fabric.

## 2.5.3 Overview of CLI commands

The CLI commands available in EXEC mode are described in Table 2-3.

*Table 2-3 EXEC mode commands*

| Command | Description |
|---|---|
| attach | Connect to a specific line card |
| callhome | Call home commands |
| cd | Change current directory |
| clear | Reset functions |
| clock | Manage the system clock |
| config | Enter configuration mode |
| copy | Copy from one file to another |
| debug | Debugging functions |
| delete | Remove files |
| dir | Directory listing for files |
| discover | Discover information |
| exit | Exit from the EXEC |

| Command | Description |
|---|---|
| fcping | Ping an N-port |
| fctrace | Trace the route for an N-port |
| find | Find a file below the current directory |
| format | Format disks |
| install | Upgrade software |
| mkdir | Create a new directory |
| move | Move files |
| no | Disable debugging functions |
| ping | Send echo messages |
| purge | Deletes unused data |
| pwd | View the current directory |
| reload | Reboot the entire box |
| rmdir | Remove existing directory |
| run-script | Run shell scripts |
| send | Send message to all open sessions |
| setup | Run the basic SETUP command facility |
| show | Show running system information |
| sleep | Sleep for the specified number of seconds |
| system | System management commands |
| tail | Display the last part of a file |
| telnet | Telnet to another system |
| terminal | Set terminal line parameters |
| test | Test command |
| traceroute | Trace route to destination |
| undebug | Disable debugging functions (see also debug) |
| write | Write the current configuration |
| zone | Execute Zone Server commands |

The CLI commands available in Configuration mode are described in Table 2-4.

*Table 2-4 Configuration mode commands*

| Command | Description |
|---|---|
| aaa | Configure authentication, authorization and accounting (AAA) parameters |
| arp | Remove an entry from the ARP cache (with the no prefix) |
| boot | Configure boot variables (kickstart and system) |
| callhome | Configure the callhome parameters |
| cdp | Configure the Cisco Discovery Protocol (CDP) parameters |
| clock | Configure timezone and daylight savings time parameters, or set the current time |
| do | Run a single EXEC mode command, without exiting from Configuration mode |
| end | Exit from configuration mode |
| exit | Exit from configuration mode |
| fcalias | Configure a Fibre Channel alias |
| fcanalyzer | Configure the Cisco Fabric Analyzer |
| fcc | Configure Fibre Channel Congestion Control |
| fcdomain | Configure Fibre Channel domain parameters |
| fcdroplatency | Configure network or switch Fibre Channel drop latency |
| fcflow | Configure fcflow statistics |
| fcinterop | Configure interoperability mode parameters |
| fcns | Configure Fibre Channel name server proxy |
| fcroute | Configure Fibre Channel routes |
| fcs | Configure fabric configuration server |
| fctimer | Configure Fibre Channel timers |
| fspf | Configure fabric shortest path first (fspf) |
| in-order-guarantee | Set in-order delivery guarantee |
| interface | Configure an interface |

| Command | Description |
|---|---|
| ip | Configure Internet protocol (IP) parameters |
| iscsi | Configure iSCSI parameters |
| kernel | Kernel options |
| line | Configure a terminal line |
| logging | Configure logging (syslog) parameters |
| no | Negate a command or set its defaults |
| ntp | Configure network time protocol (NTP) parameters |
| power | Configure power supply |
| poweroff | Power off a module in the switch |
| qos | Configure priority of Fibre Channel control frames |
| radius-server | Configure RADIUS authentication related parameters |
| role | Configure roles |
| rscn | Configure registered state change notification (RSCN) parameters |
| snmp-server | Configure simple network management protocol (snmp) parameters |
| span | Configure a switched port analyzer (SPAN) session |
| ssh | Configure SSH parameters |
| switchname | Configure name of the switch |
| system | Configure system parameters |
| telnet | Enable or disable telnet server |
| trunk | Enable or disable trunk protocol for new ISL connections |
| username | Configure user information |
| vsan | Configure the virtual SAN (VSAN) database |
| wwn | Allocate secondary MAC addresses |
| zone | Configure zones |
| zoneset | Configure zone sets |

## 2.5.4 Upgrading the switch software with the CLI

The upgrade procedure described here requires that the software version in the switches before upgrade is at least 1.0(3). If you need to upgrade a switch with an older software version, you have to use the manual method described in *Cisco MDS 9000 Family Configuration Guide*, DOC-7814893.

The procedure described here is the easiest way to upgrade the software. The **`install all`** command takes care of both downloading the new code to the switch and actually installing the code, with the smallest possible impact on the switch operation.

The software in the Cisco MDS 9000 family of switches consists of two packages: the kickstart package and the system package. These packages are delivered as separate files and they have to be compatible with each other.

You can use the following protocols to download the new software to the Cisco switch:

- SCP
- FTP
- SFTP
- TFTP

### Upgrading a switch with one supervisor

A switch with only one supervisor can be non-disruptively upgraded only if the kickstart image does not change, and if the old and new software are compatible. If either the kickstart image changes or the software versions are incompatible, downtime needs to be scheduled for the upgrade.

The major steps in the upgrade require you to:

- Download the new code to the supervisor
- Check for compatibility
- Install the Loader, if required
- Install the BIOS, if required
- Update boot variables
- Save configuration
- Reload the supervisor

The actual upgrade procedure is shown in Example 2-1. Note that the complete **`install all`** command has to be entered as a single line with all the required parameters.

*Example 2-1 Upgrading a switch with one supervisor*

```
9216# install all system ftp://9.42.164.42/tmp/cisco/m9200-ek9-mz.1.1.1a.bin
kickstart ftp://9.42.1564.42/tmp/cisco/m9200-ek9-kickstart-mz.1.1.1a.bin


Copying ftp://9.42.164.42/tmp/cisco/m9200-ek9-mz.1.1.1a.bin to
bootflash:/m9200-ek9-mz.1.1.1a.bin
Enter username:root
Password:


Copying ftp://9.42.164.42/tmp/cisco/m9200-ek9-kickstart-mz.1.1.1a.bin to
bootflash:/m9200-ek9-kickstart-mz.1.1.1a.bin
Enter username:root
Password:



Image verification is in progress, please wait.
This command is going to install system image m9200-ek9-mz.1.1.1a.bin
and kickstart image m9200-ek9-kickstart-mz.1.1.1a.bin
on this system
The command will:
 - Install the Loader, if required
 - Install the BIOS, if required
 - Update boot variables
 - Save configuration
 - Reload the supervisor


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


Installing Loader, please wait.
Installing Loader on module 1 ... not required (same version)

Installing BIOS, please wait.
Installing BIOS on module 1 ...  successful
Installing BIOS on module 2 ...  successful

Updating boot variables .. successful
Saving configuration, please wait.
```

After this message, the supervisor reboots and our login session is terminated. The reboot usually takes a few minutes. After the reboot is finished we can login to the switch and check the version as shown in Example 2-2.

*Example 2-2 Checking the software version of the upgraded switch*

```
9216# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.


Software
  BIOS:      version 1.0.7
  loader:    version 1.0(3a)
  kickstart: version 1.1(1a)
  system:    version 1.1(1a)

  BIOS compile time:       03/20/03
  kickstart image file is: bootflash:/m9200-ek9-kickstart-mz.1.1.1a.bin
  kickstart compile time:  6/12/2003 14:00:00
  system image file is:    bootflash:/m9200-ek9-mz.1.1.1a.bin
  system compile time:     6/12/2003 14:00:00


Hardware
  RAM 963108 kB

  bootflash: 500736 blocks (block size 512b)
  slot0:          0 blocks (block size 512b)

  9216 uptime is 0 days 0 hour 2 minute(s) 32 second(s)

  Last reset at 740018 usecs after Sat Jan 12 15:52:50 1980
    Reason: Reset Requested by CLI command reload
    System version: 1.0(4)


9216#
```

## Upgrading a switch with two supervisors

A switch with two supervisors can, in most cases, be upgraded without impact on the switch operation. The only exception is if the old and new software are incompatible.

The major steps in the upgrade require you to:

- Download the new code to the supervisors
- Check for compatibility

- Install the Loader, if required
- Install the BIOS, if required
- Update boot variables
- Save the configuration
- Reload the standby supervisor
- Perform an HA switch-over (causing reload of the old active supervisor)
- Perform hitless upgrade of the line cards

Note that since only one HA switch-over is done the supervisor module, that was the standby supervisor before the upgrade, is the active supervisor after the upgrade.

The status of the modules in our switch before upgrade is shown in Example 2-3. Note that the supervisor module in slot 5 is the active supervisor module.

*Example 2-3 The module status before upgrade*

```
MDS_Director# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    32     1/2 Gbps FC Module               DS-X9032           ok
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
1    1.0(4)       1.0     20:01:00:0b:5f:a3:c6:80 to 20:10:00:0b:5f:a3:c6:80
2    1.0(4)       1.0     20:41:00:0b:5f:a3:c6:80 to 20:60:00:0b:5f:a3:c6:80
5    1.0(4)       1.1     --
6    1.0(4)       1.1     --

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
1    00-0b-46-a1-a6-18 to 00-0b-46-a1-a6-1c  JAB065106PE
2    00-0b-46-a1-a4-80 to 00-0b-46-a1-a4-84  JAB065007FN
5    00-0b-be-f7-45-a8 to 00-0b-be-f7-45-ac  JAB070204H0
6    00-0b-be-f7-45-cc to 00-0b-be-f7-45-d0  JAB070204FG

* this terminal session
MDS_Director#
```

The actual upgrade procedure is shown in Example 2-4. Note that the complete **`install all`** command has to be entered as a single line, with all the required parameters.

*Example 2-4 Upgrading a switch with two supervisors*

```
MDS_Director# install all system
ftp://9.42.164.42/tmp/cisco/m9500-sf1ek9-mz.1.1 .1a.bin kickstart
ftp://9.42.164.42/tmp/cisco/m9500-sf1ek9-kickstart-mz.1.1.1a.bin


Copying ftp://9.42.164.42/tmp/cisco/m9500-sf1ek9-mz.1.1.1a.bin to
bootflash:/m9500-sf1ek9-mz.1.1.1a.bin
Enter username:root
Password:


Copying ftp://9.42.164.42/tmp/cisco/m9500-sf1ek9-kickstart-mz.1.1.1a.bin to
bootflash:/m9500-sf1ek9-kickstart-mz.1.1.1a.bin
Enter username:root
Password:



Image verification is in progress, please wait.
This command is going to install system image m9500-sf1ek9-mz.1.1.1a.bin
and kickstart image m9500-sf1ek9-kickstart-mz.1.1.1a.bin
on this system
The command will:
 - Install the Loader, if required
 - Install the BIOS, if required
 - Update boot variables
 - Save configuration
 - Reload the standby supervisor
 - Perform a HA Switchover
 - Perform a hitless upgrade of module 1, 2

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


Image synchronization is in progress, please wait.

Installing Loader, please wait.
Installing Loader on module 5 ... not required (same version)
Installing Loader on module 6 ... not required (same version)

Installing BIOS, please wait.
Installing BIOS on module 1 ...  successful
Installing BIOS on module 2 ...  successful
Installing BIOS on module 5 ...  successful
Installing BIOS on module 6 ...  successful

Updating boot variables .. successful
```

```
Saving configuration, please wait.
Reload of the standby supervisor is in progress, please wait
Success, the standby supervisor is online and ready to takeover
```

After this message, the standby supervisor takes over and our login session is terminated. After the takeover, the previously active supervisor also reboots automatically, and finally the rest of the modules are upgraded. You can login to the switch and monitor the progress of the upgrade, as shown in Example 2-5, Example 2-6, and Example 2-7. Note also that the supervisor module in slot 6 is now the active supervisor module.

*Example 2-5 Upgrading the first module*

```
MDS_Director# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           upgrading
2    32     1/2 Gbps FC Module               DS-X9032           ok
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
1    1.0(4)       1.0     20:01:00:0b:5f:a3:c6:80 to 20:10:00:0b:5f:a3:c6:80
2    1.0(4)       1.0     20:41:00:0b:5f:a3:c6:80 to 20:60:00:0b:5f:a3:c6:80
5    1.1(1a)      1.1     --
6    1.1(1a)      1.1     --

Mod  MAC-Address(es)                          Serial-Num
---  --------------------------------------   ----------
1    00-0b-46-a1-a6-18 to 00-0b-46-a1-a6-1c   JAB065106PE
2    00-0b-46-a1-a4-80 to 00-0b-46-a1-a4-84   JAB065007FN
5    00-0b-be-f7-45-a8 to 00-0b-be-f7-45-ac   JAB070204H0
6    00-0b-be-f7-45-cc to 00-0b-be-f7-45-d0   JAB070204FG

* this terminal session
MDS_Director#
```

*Example 2-6 Upgrading the second module*

```
MDS_Director# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    32     1/2 Gbps FC Module               DS-X9032           upgrading
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
```

```
Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
1    1.1(1a)      1.0     20:01:00:0b:5f:a3:c6:80 to 20:10:00:0b:5f:a3:c6:80
2    1.0(4)       1.0     20:41:00:0b:5f:a3:c6:80 to 20:60:00:0b:5f:a3:c6:80
5    1.1(1a)      1.1     --
6    1.1(1a)      1.1     --

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
1    00-0b-46-a1-a6-18 to 00-0b-46-a1-a6-1c  JAB065106PE
2    00-0b-46-a1-a4-80 to 00-0b-46-a1-a4-84  JAB065007FN
5    00-0b-be-f7-45-a8 to 00-0b-be-f7-45-ac  JAB070204H0
6    00-0b-be-f7-45-cc to 00-0b-be-f7-45-d0  JAB070204FG

* this terminal session
MDS_Director#
```

*Example 2-7 All modules upgraded*

```
MDS_Director# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    32     1/2 Gbps FC Module               DS-X9032           ok
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
1    1.1(1a)      1.0     20:01:00:0b:5f:a3:c6:80 to 20:10:00:0b:5f:a3:c6:80
2    1.1(1a)      1.0     20:41:00:0b:5f:a3:c6:80 to 20:60:00:0b:5f:a3:c6:80
5    1.1(1a)      1.1     --
6    1.1(1a)      1.1     --

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
1    00-0b-46-a1-a6-18 to 00-0b-46-a1-a6-1c  JAB065106PE
2    00-0b-46-a1-a4-80 to 00-0b-46-a1-a4-84  JAB065007FN
5    00-0b-be-f7-45-a8 to 00-0b-be-f7-45-ac  JAB070204H0
6    00-0b-be-f7-45-cc to 00-0b-be-f7-45-d0  JAB070204FG

* this terminal session
MDS_Director#
```

After the upgrade is finished, we can also check the software version as shown in Example 2-8.

*Example 2-8 Checking the software version on a switch with two supervisors*

```
MDS_Director# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.


Software
  BIOS:      version 1.0.7
  loader:    version 1.0(3a)
  kickstart: version 1.1(1a)
  system:    version 1.1(1a)

  BIOS compile time:       03/20/03
  kickstart image file is: bootflash:/m9500-sf1ek9-kickstart-mz.1.1.1a.bin
  kickstart compile time:  6/12/2003 13:00:00
  system image file is:    bootflash:/m9500-sf1ek9-mz.1.1.1a.bin
  system compile time:     6/12/2003 13:00:00


Hardware
  RAM 1027620 kB

  bootflash: 500736 blocks (block size 512b)
  slot0:          0 blocks (block size 512b)

  MDS_Director uptime is 0 days 0 hour 6 minute(s) 57 second(s)

  Last reset
    Reason: Watchdog Timeout/External Reset
    System version: 1.0(4)


MDS_Director#
```

## 2.6 Interoperability mode implications

The switch interoperability mode enables the switches from multiple vendors to connect to each other. This includes attaching systems with built-in switch modules, such as the IBM BladeCenter, to an existing fabric.

Switch interoperability mode turns off a number of advanced or proprietary features of the switch forming a common subset of features that the switches from all vendors can implement.

### 2.6.1 General implications

The general implications of using the interoperability mode with the Cisco MDS 9000 family are described in Table 2-5.

*Table 2-5 Interoperability mode in Cisco MDS 9000 family*

| Switch feature | Changes in interoperability mode |
|---|---|
| Domain IDs | Domain IDs are restricted to the range 97-127 (0x61-0x7f), since some vendors cannot use the full range of 239 domains within a fabric |
| Timers | All Fibre Channel timers (F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV) must be the same on all switches |
| Trunking | VSAN trunking is not supported between switches from different vendors, but can still be used between the Cisco MDS 9000 series switches |
| Default zone | The default zone behavior of permit or deny may change |
| Zone attributes | Only WWPN zoning is allowed, port based zoning or FC ID based zoning are not supported |
| Zone propagation | Some vendors pass only the active zoneset to other switches, instead of the full zone configuration |
| VSAN | Switch interoperability mode is an attribute of a VSAN, and turning it on for a VSAN does not affect other VSANs |
| TE_Ports | TE_Ports cannot be used to connect to switches from another vendor, but can still be used to connect to other Cisco MDS 9000 series switches, and VSANs that have the interoperability mode turned on can use them |
| PortChannels | PortChannels cannot be used to connect to switches from another vendor, but can still be used to connect to other Cisco MDS 9000 series switches, and VSANs that have the interoperability mode turned on can use them |

**Note:** By default, Brocade switches use proprietary frames to exchange platform information, and the Cisco MDS 9000 series switches do not understand those frames. This causes the E_Ports between the switches to become isolated. To avoid this, you have to run the command **`msplmgmtdeactivate`** in the Brocade switch before connecting the switches together.

## 2.6.2 Changing an existing fabric to interoperability mode

In the Cisco fabric, interoperability mode is a VSAN level configuration parameter. If you change one VSAN to interoperability mode, the other VSANs are not affected.

However, we recommend that the VSAN to be changed should be suspended for the change. Also, the domain IDs of the switches in the VSAN may have to be changed to conform to the interoperability mode addressing limitations.

**Note:** The domain ID of a switch is part of the FC ID of any device connected to the switch. AIX and HP-UX hosts use the FC IDs of devices in their internal device addressing. If the FC ID of a device changes, these operating systems see the device as a new device. In this case, you have to take any steps required by your operating system and multipathing software to correct the device configuration.

## 2.6.3 Settings required for IBM BladeCenter attachment

In this example we will attach an IBM BladeCenter server to an existing Cisco MDS 9509 director. We will describe the settings required to successfully implement the attachment and show a step-by-step procedure for the attachment using the Cisco CLI.

You can also implement the attachment using the Fabric Manager GUI. Fabric Manager is described in 2.4, “Managing the Cisco SAN with the Fabric Manager” on page 326.

We recommend that you set one of the Cisco switches to be the primary switch by setting the Cisco switch to the highest possible priority (1) and setting the switch within the BladeCenter to the lowest possible priority (254), and that you manage the zoning with the Cisco tools only.

The settings required for both the IBM BladeCenter and the Cisco MDS 9000 are described in Table 2-6.

*Table 2-6 Settings used for IBM BladeCenter attachment*

| Configuration Features | IBM BladeCenter | Cisco MDS 9000 |
|---|---|---|
| Firmware level | V1.4.0.49-0 | 1.1(1a) |
| Domain ID | 125<br>(allowed range 97-127) | 99<br>(allowed range 97-127) |
| Domain ID lock | True | Static |
| E_D_TOV | 2000 ms<br>(has to be same for both) | 2000 ms<br>(has to be same for both) |
| R_A_TOV | 10000 ms<br>(has to be same for both) | 10000 ms<br>(has to be same for both) |
| IO Stream Guard | disabled on the ISL | Not Applicable |
| Principal Switch Priority | 254 | 1 |
| VSAN | Not Applicable | VSAN 9<br>(can be any VSAN) |
| Interop Mode | Not Applicable | VSAN in interop mode |
| Zoning | WWPN based | WWPN based |
| Default Zone | disabled | disabled |
| Switch Port Mode | GL Port < default> | E Port |

## Step by step procedure using the CLI

**Attention:** The example below is a sequential process and skipping *any* step listed below may lead the user into a different menu from where the following steps are not allowed.

This procedure is applicable to all the members of the Cisco 9000 family.

1. Configure a new VSAN in Interop mode or change the existing VSAN from Cisco Native to Interop mode as shown below.

```
MDS_Director login: admin
Password:
MDS_Director# config t
Enter configuration commands, one per line.  End with CNTL/Z.
MDS_Director(config)# vsan database
MDS_Director(config-vsan-db)# vsan 9 name Interop_VSAN interop
MDS_Director(config-vsan-db)#
```

2. Assign Fibre Channel Interfaces into the Interop VSAN. In our example below, five interfaces, fc 1/5 to 1/9, are assigned to VSAN 9.

```
MDS_Director(config-vsan-db)# vsan 9 interface fc 1/5
MDS_Director(config-vsan-db)# vsan 9 interface fc 1/6
MDS_Director(config-vsan-db)# vsan 9 interface fc 1/7
MDS_Director(config-vsan-db)# vsan 9 interface fc 1/8
MDS_Director(config-vsan-db)# vsan 9 interface fc 1/9
MDS_Director(config-vsan-db)# exit
MDS_Director(config)#
```

3. Configure the domain ID for VSAN 9 from the Domain Manager config prompt. The domain IDs must fall in the range of 97 -127 for domains that belong to the Interop VSAN, and we recommend that you use the static option to ensure that the domain ID does not change. In the example below static domain ID 99 is configured for VSAN 9.

```
MDS_Director(config)# fcdomain domain 99 static vsan 9
MDS_Director(config)#
```

4. Configure the Switch Priority value to 1 (highest priority) so that the MDS 9509 is elected as the Principal switch as shown below.

```
MDS_Director(config)# fcdomain priority 1 vsan 9
MDS_Director(config)#
```

**Note:** We recommend that the IBM BladeCenter FC switch always be the subordinate by lowering its principal switch priority value to 254.

5. Once the domain ID and switch priority values are set, perform a disruptive restart of the Domain Manager for the VSAN, as shown below.

```
MDS_Director(config)# fcdomain restart disruptive vsan 9
MDS_Director(config)#
```

6. Configure E_Port mode for the port that is going to be the ISL between the Cisco switch and the BladeCenter, and exit from the Configuration mode.

```
MDS_Director(config)# interface fc 1/5
MDS_Director(config-if)# switchport mode e
MDS_Director(config-if)# exit
MDS_Director(config)# exit
MDS_Director#
```

7. Verify that all of the switches in the VSAN can be seen as shown below.

```
MDS_Director# show fcdomain domain-list vsan 9
Number of domains: 2
Domain ID              WWN
---------    -----------------------
0x63(99)    20:09:00:0b:5f:a3:c6:81 [Local] [Principal]
0x7d(125)    10:00:00:c0:dd:01:c3:0c
```

```
MDS_Director#
```

8. Create a zone set to hold the zones for the VSAN.

```
MDS_Director# conf t
MDS_Director(config)# zoneset name BladeCenter_MDS9509 vsan 9
MDS_Director(config-zoneset)#
```

9. Create zones for the zone set. In our case we want to include port 1 of the BladeCenter and an ESS to the zone, and the name we choose reflects that.

```
MDS_Director(config-zoneset)# zone name BladeSrv_P1_Shark
MDS_Director(config-zoneset-zone)#
```

10. Add members into the zone and return to Config mode, as shown below. Note that you should only select the zone members by pwwn.

```
MDS_Director(config-zoneset-zone)# member pwwn 21:00:00:09:6b:36:00:66
MDS_Director(config-zoneset-zone)# member pwwn 50:05:07:63:00:c7:0b:e8
MDS_Director(config-zoneset-zone)# exit
MDS_Director(config-zoneset)# exit
MDS_Director(config)#
```

11. Activate Zone Set for VSAN 9 and return to EXEC mode as shown below.

```
MDS_Director(config)# zoneset activate name BladeCenter_MDS9509 vsan 9
Zoneset activation initiated. check zone status
MDS_Director(config)# exit
MDS_Director#
```

12. Verify the Zone merge and activation as shown below.

```
MDS_Director# show zoneset active vsan 9
zoneset name BladeCenter_MDS9509 vsan 9
  zone name BladeSrv_P1_Shark vsan 9
  * fcid 0x630002 [pwwn 50:05:07:63:00:d0:0b:84]
  * fcid 0x7d0100 [pwwn 21:00:00:09:6b:36:00:66]

  zone name BladeSrv_P2_Shark vsan 9
  * fcid 0x630002 [pwwn 50:05:07:63:00:d0:0b:84]
  * fcid 0x630100 [pwwn 21:00:00:e0:8b:05:dd:21]


MDS_Director#
```

13. The implementation is now complete. We can see the internal Fibre Channel switch of the BladeCenter in the Cisco Fabric Manager as a QLogic switch, and the WWPNs of any blades that are installed in the BladeCenter behind that switch, as shown in Figure 2-142. The red diagonal line over the QLogic switch means that Fabric Manager does not have IP connectivity to the QLogic switch.

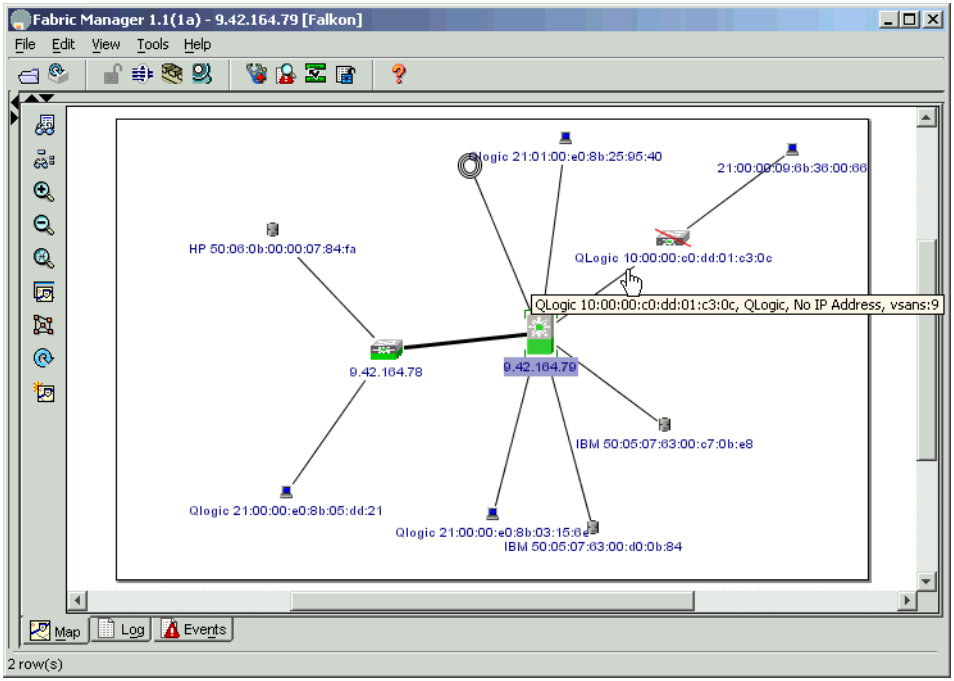


*Figure 2-142 Cisco fabric with BladeCenter installed*

3

# Implementing a SAN with CNT

The IBM machine type 2042, CNT FC/9000 Fibre Channel Director, is the core product of a reseller agreement between IBM and CNT Technologies, which adds the CNT FC/9000 Fibre Channel Director to IBM's growing list of enterprise-class SAN fabric offerings.

# 3.1 Introducing the CNT FC/9000 products

The CNT FC/9000 family is a family of expandable Fibre Channel directors that supports FCP and FICON™.

## 3.1.1 Director models

There are currently three models of directors available in the CNT FC/9000 family:

- CNT FC/9000 Fibre Channel Director 64 (FC/9000-64)
- CNT FC/9000 Fibre Channel Director 128 (FC/9000-128)
- CNT FC/9000 Fibre Channel Director 256 (FC/9000-256)

The differences between the models are described in Table 3-1.

*Table 3-1 FC/9000 models*

| | **FC/9000-64** | **FC/9000-128** | **FC/9000-256** |
|---|---|---|---|
| IBM part number | 2042-001 | 2042-128 | 2042-256 |
| XFIO2 modules included in base configuration | 2 | 3 | 4 |
| Additional XFIO2 or FIO modules | 6 | 13 | 28 |
| Ports included | 16 | 24 | 32 |
| Ports maximum | 64 | 128 | 256 |
| Rack included | No | Yes | Yes |

All of the models share the following common characteristics:

- Support for both Fibre Channel and FICON, including FICON cascading
- Support for 1 Gb/s and 2 Gb/s Fibre Channel
- Management with a graphical user interface
- Optionally, fully redundant architecture

In addition, two FC/9000-64 directors can be connected to form a FC/9000-128, or two FC/9000-128 four FC/9000-64 directors can be connected to form a FC/9000-256 using the interconnection kits available.

You can see the CNT FC/9000 family of Fibre Channel directors in Figure 3-1. The models from left to right are FC/9000-256, FC/9000-128, and FC/9000-64.

*Figure 3-1 The CNT FC/9000 family*

## 3.1.2 Modules and transceivers

The CNT FC/9000 series currently supports two modules:

- Additional 2 Gb XFIO2 Module (8 ports) (feature code (f/c) 5210)
- Additional 1 Gb FIO Module (8 ports) (f/c 5010)

We recommend that you use the new XFIO2 module whenever possible. However, since the XFIO2 module does not support Fibre Channel Arbitrated Loop (FC-AL), some peripheral devices are only supported by the FIO module.

### XFIO2 module

The XFIO2 module has eight Fibre Channel ports that are capable of working at both 1 Gb/s and 2 Gb/s, and two mirror ports that can be used to monitor the Fibre Channel traffic. It has 128 buffer credits per port to support ISL distances of up to 100 km, and also supports FICON addressing. The ports use SFP transceivers and eight SFP transceivers must be ordered for each XFIO2 module. There are three different SFP transceivers currently available:

- Shortwave SFP Transceiver (f/c 2210)
- Longwave SFP Transceiver (f/c 2220)
- Longwave 35km SFP Transceiver (f/c 2235)

All of the SFP transceivers use LC duplex connectors.

### FIO module

The FIO module has eight Fibre Channel ports that are capable of working only at 1 Gb/s. The ports use Gigabit Interface Converters (GBICs), and eight GBICs

must be ordered for each FIO module. There are three different GBICs currently available:

- Shortwave GBIC (multimode fiber) (f/c 2010)
- Longwave GBIC (singlemode fiber) (f/c 2020)
- Extended Longwave GBIC (80 km) (f/c 2030)

All of the GBICs use SC duplex connectors.

### 3.1.3 ISL modes

There are two ISL modes that can currently be used to connect the CNT FC/9000 family of directors:

- T_Port mode
- E_Port mode

The T_Port mode is a CNT proprietary ISL mode and was the only ISL mode supported with firmware levels below 3.0. The E_Port mode is the ISL mode defined by the FC-SW-2 standard. Currently both modes are supported, and we recommend that you only use the E_Port mode for new installations.

T_Port mode is discussed in Appendix A, “CNT FC/9000 T_Port mode” on page 701.

### 3.1.4 Port modes

The different Fibre Channel port modes supported by the CNT FC/9000 directors are described in Table 3-2.

*Table 3-2 Fibre Channel port modes*

| Mode | Description |
|---|---|
| E_Port | ISL port in the E_Port mode |
| F_Port | Fabric port, can connect to one N_Port |
| FL_Port | Public Loop port, can connect to one or more NL_Ports in the loop |
| T_Port | ISL port in the legacy T_Port mode |
| TL_Port | Translative loop port for connecting private loop devices to the fabric, only available in the legacy T_Port mode |

In E_Port mode, the required port mode is detected automatically by the director and cannot be set manually. The only setting that can be changed is whether loop devices are allowed to connect to a particular port, or not.

## 3.1.5 Zoning

In E_Port mode zones can be specified based on device WWN or Host WWN. Zones are grouped in zone sets. One zone set can be active in the fabric at any given time.

If E_Port mode is not specified, in Appendix A, “CNT FC/9000 T_Port mode” on page 701 we discuss other CNT zoning options.

## 3.1.6 Management capabilities

FC/9000 fabrics can be managed using different communication methods:

- IN-VSN-Manager (IP based client/server management software
- Simple network management protocol (SNMP)
- Serial interface (RS232, dedicated for CNT / IBM engineers)
- Call Home (modem connection for notification purposes)
- Trivial file transfer protocol (TFTP) to load microcode IP-settings

The most commonly used interface with the CNT FC/9000 director is the IN-VSN software tool. In the following topics our focus is on how to use the IN-VSN tool.

## 3.1.7 Supported protocols

In most environments a homogeneous landscape of servers and storage is hard to find. For example, most tape-drives solely support FC_AL (known as arbitrated loop), whereas modern disk systems are widely used with point-to-point protocol (sometimes called P2P).

CNT supports these protocols in E_Port mode:

- Open Systems:
  - Arbitrated Loop
  - FC-SW (Fibre Channel Switched Fabric)
- S/390 systems (zSeries®):
  - FICON

All these can attach to a single CNT director at the same time.

### Support of cascading

To create even larger fabrics FC/9000 directors can be cascaded. By doing this you can create fabrics with more than a thousand external ports. Today, IBM supports fabrics with up to 8 cascaded CNT FC/9000 directors.

Cascading is also supported when using FICON attachments.

### 3.1.8 Supported device attachment

For the latest support matrix of the CNT FC/9000, refer to:

http://www.storage.ibm.com/ibmsan/products/directors/prod_data/supserver-042.html

## 3.2 Getting started

Most of the management activities for the CNT directors can be performed from the IN-VSN management console.

### 3.2.1 Initial setup of CNT FC/9000 IP settings

The CNT Director is delivered with the current supported level of firmware and default IP addresses 10.1.1.51 and 10.1.1.52, and a subnet mask of 255.255.255.0. There is also a default chassis ID, switch ID, and fabric ID. Without doing any kind of cascading, there is no need to change these IDs.

It is the responsibility of an IBM Customer Engineer (CE) to reset all default addresses to reflect the environment in which it is being installed.

For added security, the TCP/IP address can only be set or reset by the CE, using an RS232 connection and entering the CE user id and password. The new TCP/IP address is displayed in an LED panel that can be located on the FCM module.

The PC that is used for the IN-VSN server needs to be connected to the FCM modules using the Ethernet ports.

You should also ascertain whether the applicable microcode level is installed or needs to be installed. The same is true for the IN-VSN management software:

- The director microcode level that we will use is: 3.0.1.1.
- The IN-VSN management tool release we will use is: 8.0.0.6.

Since these codes are subject to change as new functions are added and improvements made, ask your IBM or CNT technical contact what the currently supported levels are.

### 3.2.2 Establishing network connection

As mentioned previously, the CNT director and its management PC are delivered with pre-installed IP settings. They are adequate to set up a small private network with just directors and the management PC as members.

You can leave the initial IP setup as it was delivered for use as, or similar to, a private network. Consequently, only local users can attach to the directors and its IN-VSN management tool.

You can also adopt your corporate network settings to enable remote IN-VSN access.

## Leaving all IP settings as a private network

For using the IN-VSN software from this local private network, the supplied hub is sufficient. To enable all IP based components to communicate, plug in the Management PC and both CNT IP ports to this Hub. So initially, our network setup looks like that shown in Figure 3-2.

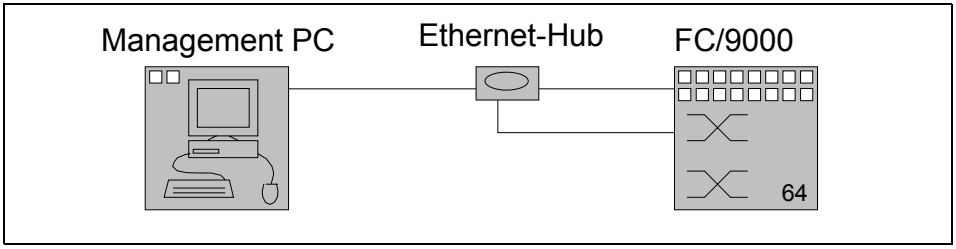


*Figure 3-2 Private IP network for initial IN-VSN management ability*

## Enabling IN-VSN access from a corporate network

To exploit the remote management capabilities of the IN-VSN management software, we recommend that you connect this network as a subnet to the corporate LAN.

Before actually changing any parameters, it is important to obtain all the information you need in advance, such as:

- Available IP addresses
- Valid subnet mask
- Default gateway

To connect the CNT subnet to your corporate LAN follow these steps:

1. Change the IP setting of the primary FCM blade using the RS232 Interface (this should be done by CNT or IBM Customer Engineers).
2. Change the IP setting of the secondary FCM blade using the RS232 interface (this should be done by CNT or IBM Customer Engineers).
3. Change the IP settings of the management PC using the common Windows Control Panel tools, and reboot the Windows server.

4. From the management PC, **`ping`** both IP addresses of the director to ensure that everything is set properly (provided that all components are still connected to the 3Com IP-Hub).
5. Attach all required ports to a switch or hub that is part of the corporate network, or connect the supplied hub to the corporate network.

After attaching our CNT setup to the corporate network, as shown in Figure 3-3, we are now able to access the IN-VSN software from wherever we are in the corporate network.

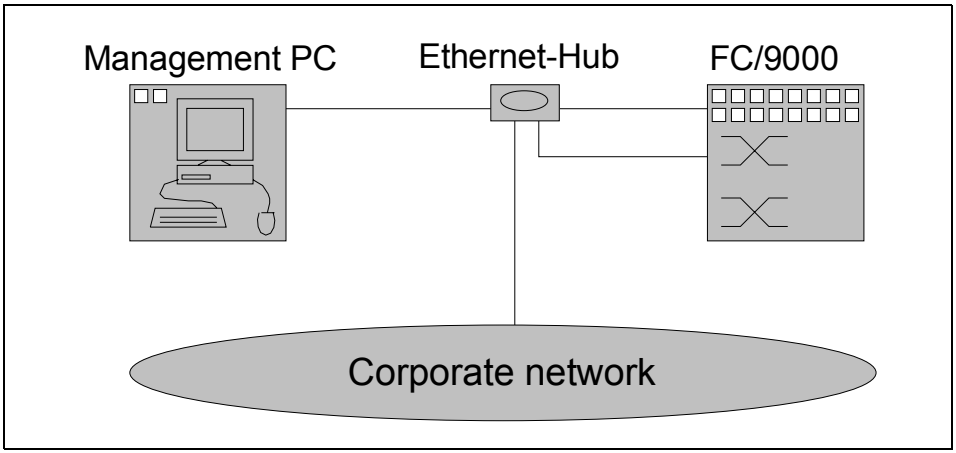


*Figure 3-3 CNT setup attached to a corporate network*

To gain actual access to the IN-VSN management tool, you need to install an IN-VSN client that can then communicate with the IN-VSN server over the corporate network.

Instructions on how to install both the server and the client part of IN-VSN software are provided in 3.2.3, “Installing the IN-VSN Enterprise Manager software” on page 414.

## Setting up high security network access

In the previous setup we described a network layout in which the Management PC as well as the fabric components were connected to the corporate network.

However, if this corporate network itself cannot be considered as secure enough, we recommend that you separate the fabric components from the corporate network. The only bridge between such a separated fabric management network and the corporate network would be the Management PC.

Referring to Figure 3-4, we see that now only the Management PC can access the directors's IP ports.

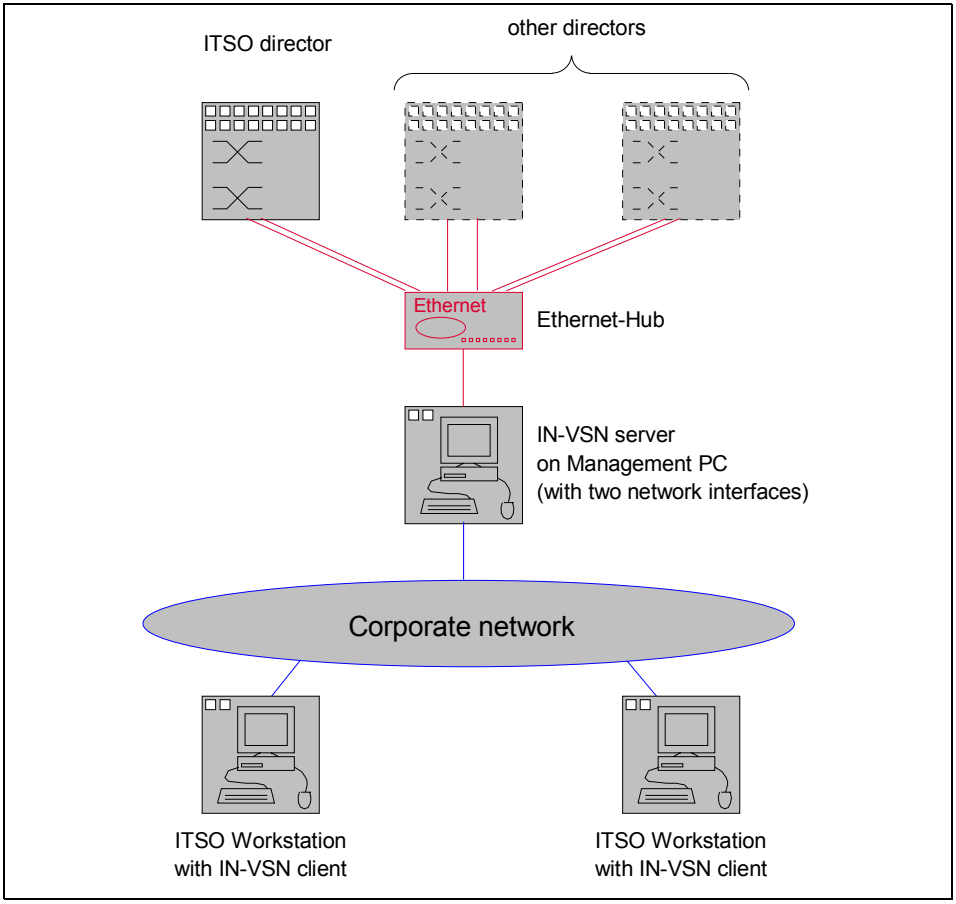


*Figure 3-4 CNT setup with secure director access*

Direct IP access from the corporate network to directors is now impossible. The only way to gain access is using the IN-VSN server.

We consider this as the most secure network setup for remote IN-VSN access. However, you need two network interfaces in the Management PC.

### 3.2.3 Installing the IN-VSN Enterprise Manager software

There are a number of different ways in which the IN-VSN software can be delivered.

If software feature code (f/c) 7600 is configured, the workstation pack includes a PC, monitor, Ethernet Hub, and Modems. The management software comes pre-installed on the PC.

If you require more than one version of the IN-VSN client code, this is achieved by using f/c 7201.

If no additional software feature codes are configured when ordering the CNT Director, you receive a CD that contains the IN-VSN server and client code. This CD is only licensed for one copy of the server and one copy of the client.

#### Prerequisites for installing the IN-VSN Enterprise Manager

The IN-VSN server and client can be installed into a PC running Windows 2000 service pack 2 or higher.

You need to have Java Runtime Environment (JRE) version 1.4.1 or later installed in your machine before installing either the IN-VSN server or the IN-VSN client. If you don't already have the JRE installed, you can install it with the **`j2re-1_4_1_01-windows-i586-i.exe`** installer that you can find on your CD.

#### Running the setup procedures of IN-VSN server and client

Both the IN-VSN server and client are installed with the same setup program **`setup.exe`**, that you can find on your installation CD.

To install the IN-VSN software, follow these steps:

1. From a Windows browser or a command prompt, start **`D:\setup.exe`**.
2. The install wizard appears. Click **Next** to continue the installation as shown in Figure 3-5.

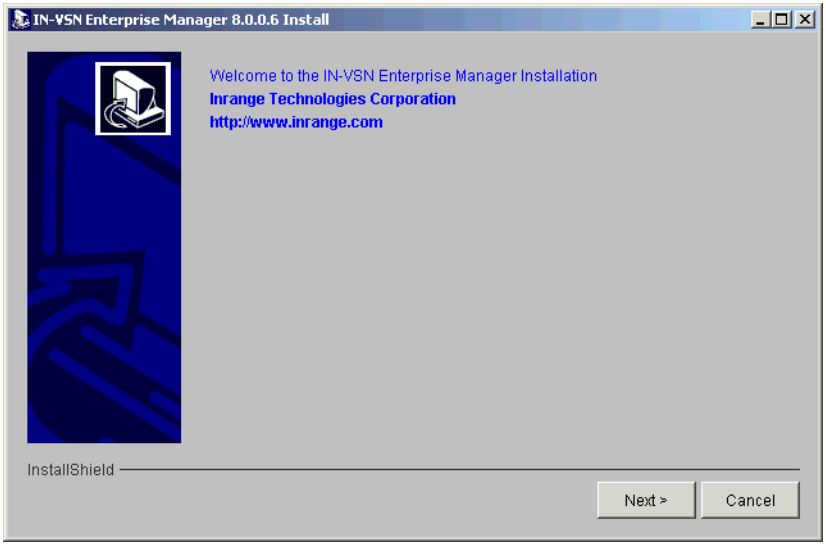


*Figure 3-5 IN-VSN Setup window*

3. Click **I Agree** and **Next** in the License Agreement Window as shown in Figure 3-6.

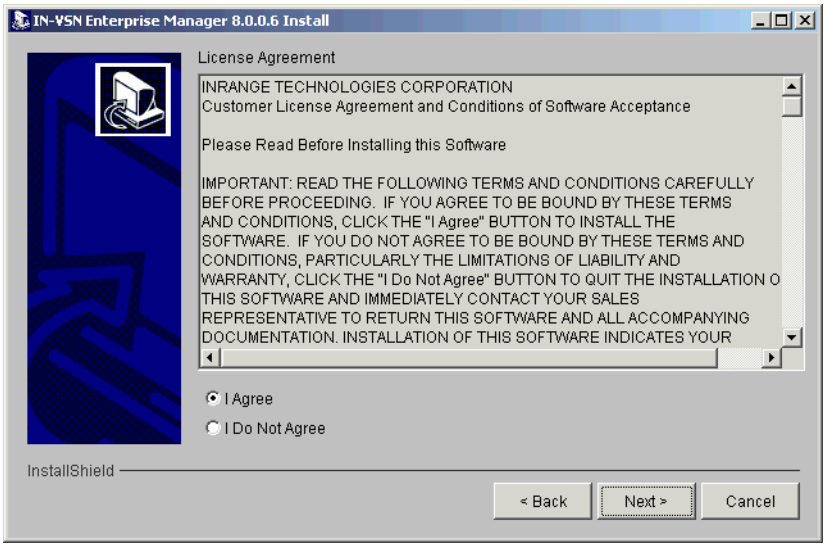


*Figure 3-6 IN-VSN License Agreement*

4. Choose which part of the IN-VSN software should be installed, server and/or client, as shown in Figure 3-7.

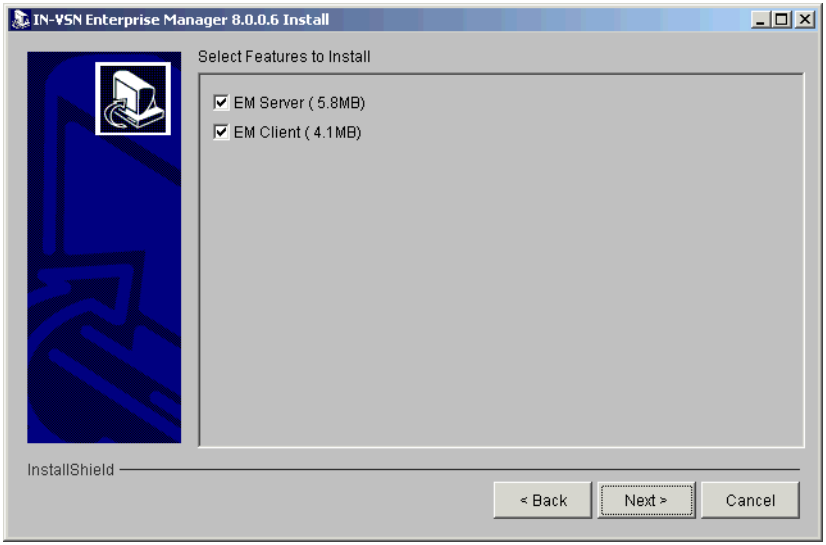


*Figure 3-7 Choosing the IN-VSN parts*

5. After selecting the features that you wish to install, you may either accept the default installation path or enter a new one, as shown in Figure 3-8. The default installation path for both client and server is **`C:\INRANGE\Fc9000`**.

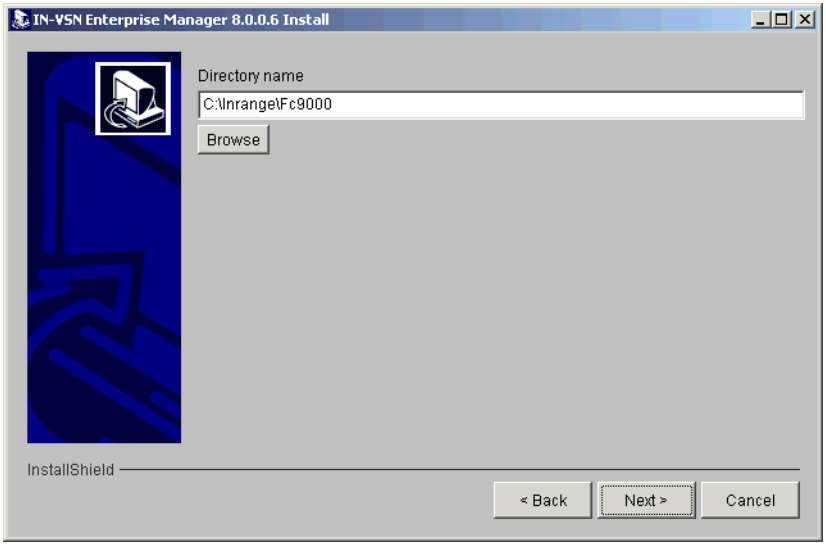


*Figure 3-8 Choosing the path for IN-VSN*

6. If the directory does not exist, the setup program asks your permission to create it as shown in Figure 3-9.

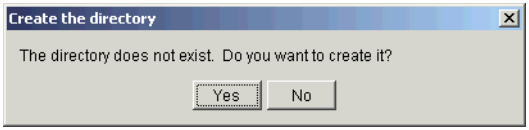


*Figure 3-9 Confirm the creation of the new directory*

7. Before actually starting the installation you see an overview of the choices you made, and the total disk space requirements, as shown in Figure 3-10. Click **Next** to accept these settings.

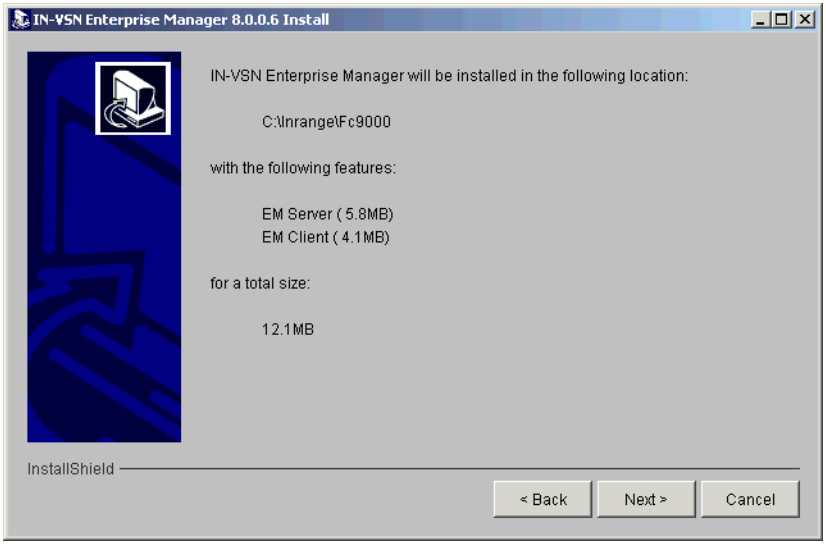


*Figure 3-10 Overview of IN-VSN installation settings*

8. The installation proceeds with no need for additional input. After installing these IN-VSN packages you are given an installation summary which tells you that all components were installed successfully, as shown in Figure 3-11.


*Figure 3-11 IN-VSN installation complete*

9. Click **Finish** to close the installer window. You can now start the IN-VSN server and log in with the IN-VSN client.

## 3.3 Managing the fabric with IN-VSN

To be able to use IN-VSN, you have to have the IN-VSN server running in your management PC, and a network connection between your client and the management PC.

### 3.3.1 Starting the IN-VSN server

If you do not have the server running, start it on your management PC by double-clicking the **IN-VSN Enterprise Mgr** icon, or use the Windows start menu. You will see the IN-VSN server status window as shown in Figure 3-12.


*Figure 3-12 IN-VSN Server*

Another indicator of the running status is the permanent change of colors of the text in the middle of the IN-VSN server window. Once the IN-VSN server is running, you can then start multiple IN-VSN client sessions pointing to the servers IP address or DNS name.

## 3.3.2 Logging into the IN-VSN

The following prerequisites must be fulfilled before being able to start an IN-VSN client session:

- Have the IN-VSN client code installed.
- On the workstation running the IN-VSN client, have the appropriate Java Virtual Machine installed.
- Have an IN-VSN server running locally or somewhere else in the network.
- Have an IP connection from IN-VSN client to IN-VSN server.

To start the IN-VSN client, just double-click the **IN-VSN client** icon on your desktop, or use the Windows Start menu. You will see a login window as shown in Figure 3-13.


*Figure 3-13 IN-VSN login window*

Enter user name (default: *admin*), password (default: *admin*), and the address of your IN-VSN server. If you are starting the client in the same machine where the server is running, you can just use *localhost* as the server name.

When logged into the IN-VSN server for the first time, all the settings are still the default values, and no specific fabric or director information is available. This is shown in Figure 3-14.

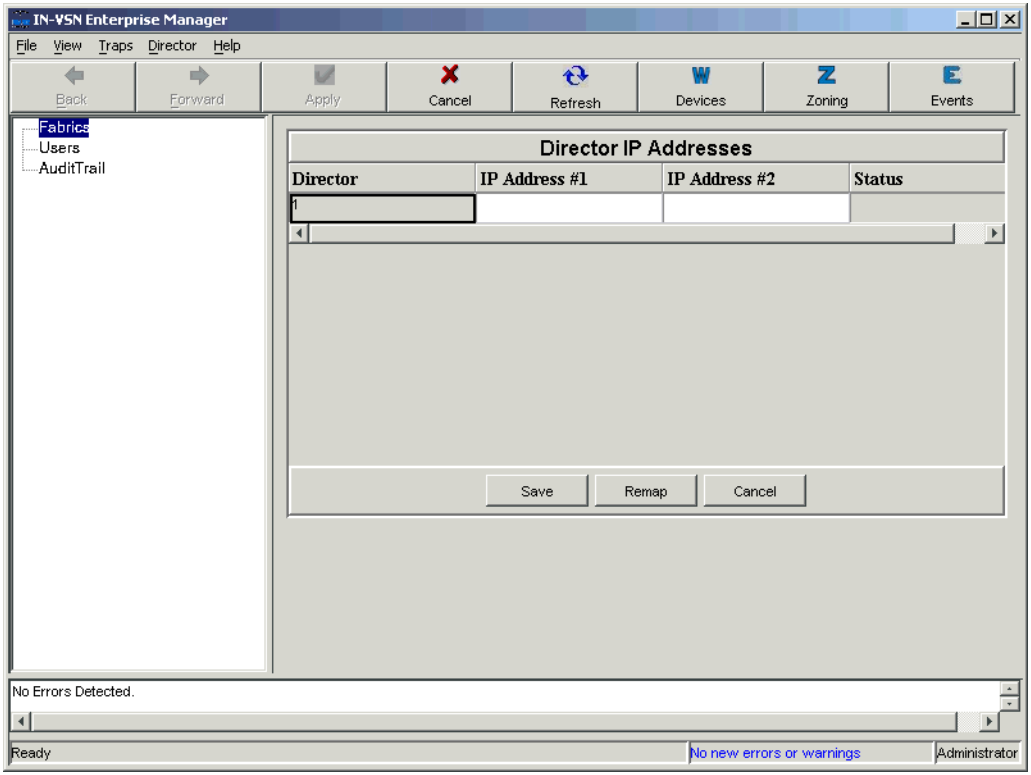


*Figure 3-14 Initial IN-VSN view*

## 3.3.3 Defining users

Different user-levels are provided by CNT. You can have multiple users with different levels defined at a time. Default user names are the same as their level and are as shown in Table 3-3.

*Table 3-3 User levels and default users*

| Default user | Default password | Explanation |
|---|---|---|
| viewer | viewer | View and monitoring |
| oper | oper | Operations like hard zoning, naming or port definitions |
| admin | admin | Any zoning, naming, definitions, user-management |
| maint | Restricted distribution | Special user for customer engineer (CE) access |

**Note:** The maint user is not shown in IN-VSN user management, and the maint role cannot be granted to any other user.

These default users and passwords are the same for all IN-VSN servers. We strongly advise you to change them as soon as you start your fabric definition to avoid any unauthorized and unwanted access.

To manage the users, click **Users** on the menu on the left side of the window. You need to enter your password again as an additional security measure. You then see the list of the defined users as shown in Figure 3-15.

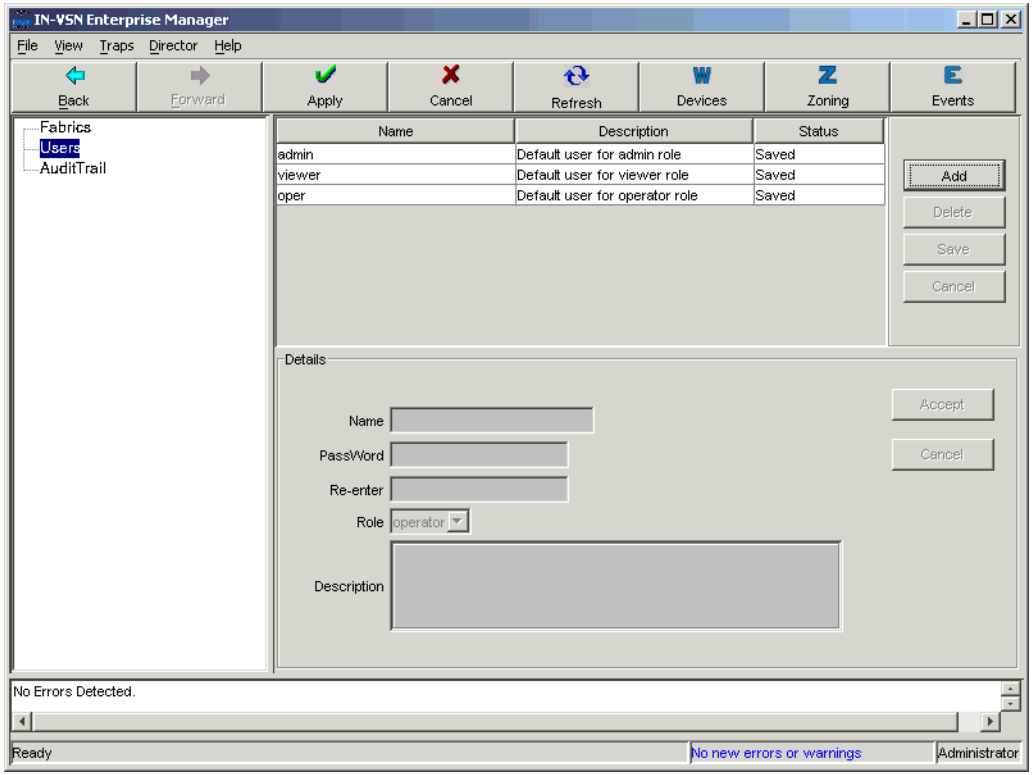


*Figure 3-15 List of defined users*

## Adding a new user

To add a new user, click the **Add** button on the right side of the window. Enter the information about the user into the bottom part of the window as shown in Figure 3-16.

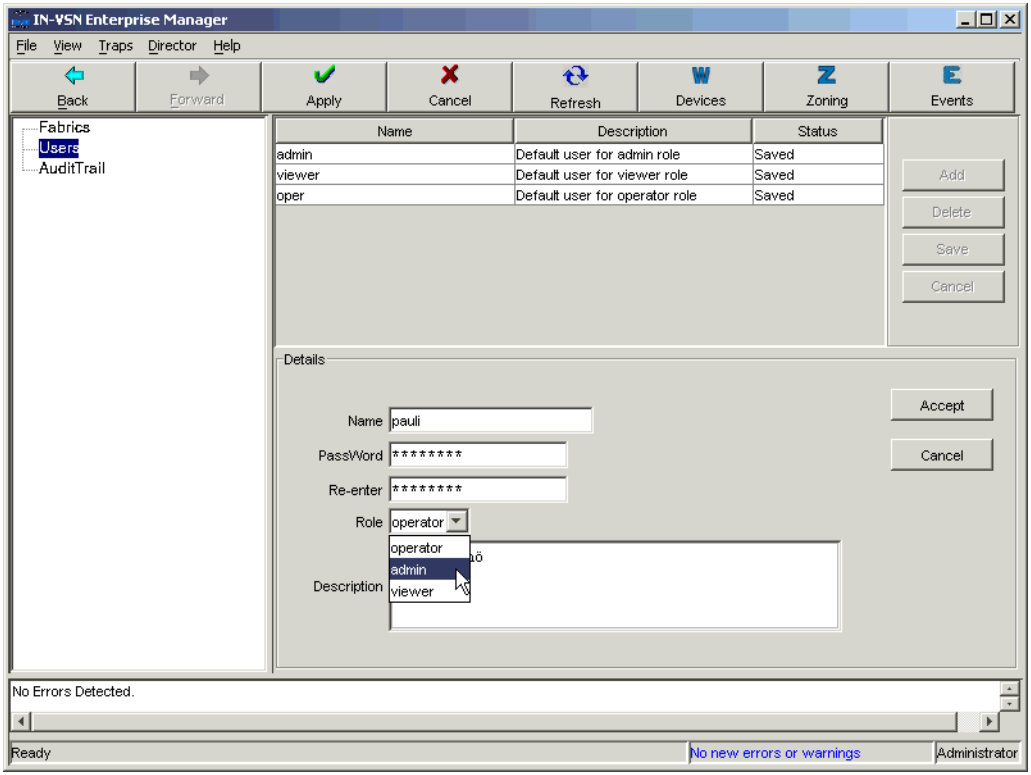


*Figure 3-16 Adding a user*

Note that you can choose one of the three roles for the new user: Admin, Operator, or Viewer. When you have finished filling in the details, click **Accept** to add this newly created user into the list of known users.

The new user you created is now shown in the list of defined users as shown in Figure 3-17. However, the user is not saved and activated until you click the **Save** button.

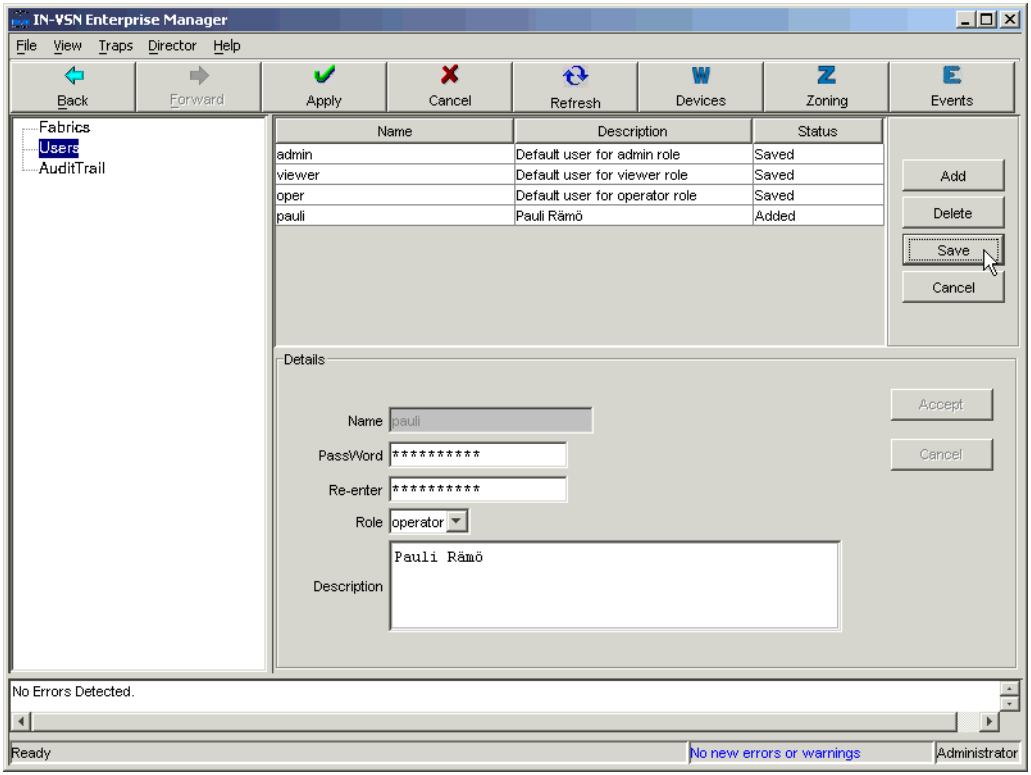


*Figure 3-17 Saving new user configuration*

Note that the default users are not deleted automatically, but they can be deleted with the IN-VSN like any other user. To maintain maximum security, we recommend that you delete the default users as soon as you have your own administrative user IDs defined and tested for successful access.

## Deleting users

Deleting users is done using the same panel as we used for adding users. Note that you cannot delete the user that you are logged in with. This ensures that you always have at least one user with Admin role defined.

To delete any user select the victim in the list of already defined and saved users. Click **Delete**, and click **Yes** on the confirmation dialog box to confirm the deletion.

Once you have confirmed this deletion that particular user is deleted in the IN-VSN database of the IN-VSN server. It is also removed immediately from the user list of the IN-VSN client.

## Changing user definitions

Once you have added users, you can change their attributes. All attributes, except the name, can be changed. These are:

- Password
- User rights
- Description

To change a user, first select it in the **Users** list. Once you have selected a user, you are automatically in change mode. Change the desired settings and click **Accept** as shown in Figure 3-18.

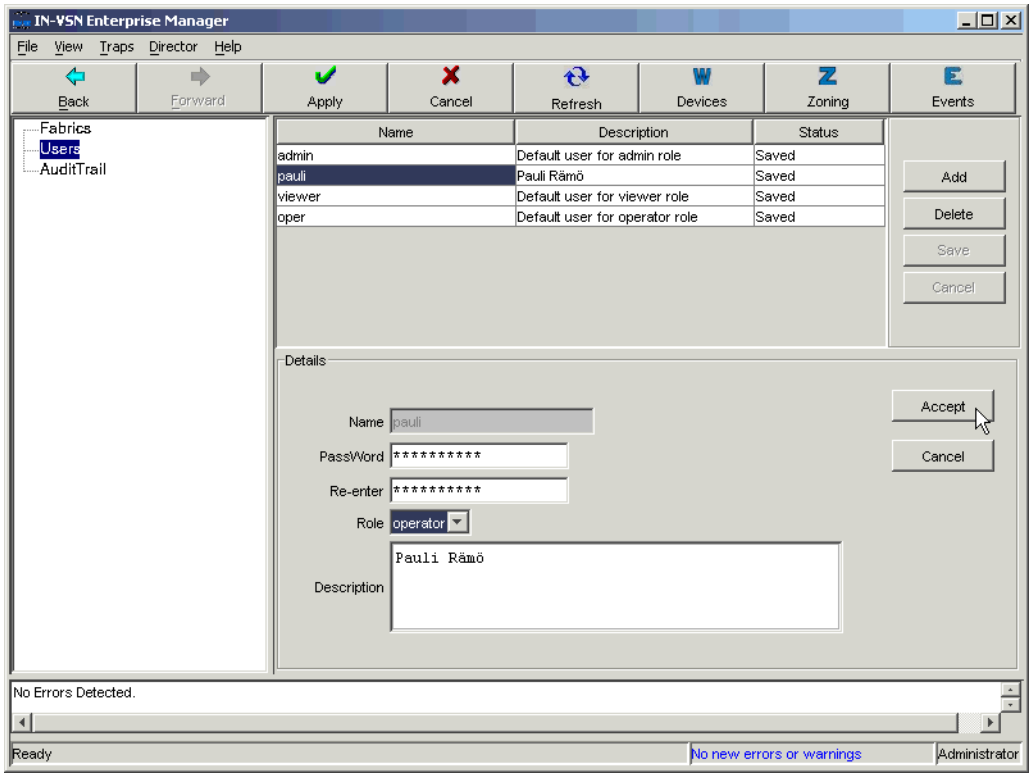


*Figure 3-18 Changing a user*

We have changed the rights of user *pauli* from *admin* to *operator*.

After accepting the changes, this modified user appears with updated status. You can modify multiple users before actually making these changes effective by clicking **Save**.

If you are logged in with admin rights, you can change all attributes of all users. If you are logged in with only operator rights, you can only change your own password and description. An operator is not allowed to change its own rights (operator to admin, for example).

## 3.3.4 Connecting IN-VSN to a CNT fabric

A fabric is made up of one or multiple CNT directors and or switches which are linked via Inter Switch Links (ISLs). Directors that are not linked via ISL are considered to be different fabrics.

Initially, the IN-VSN software is not aware of any fabric components at all as shown in Figure 3-19.

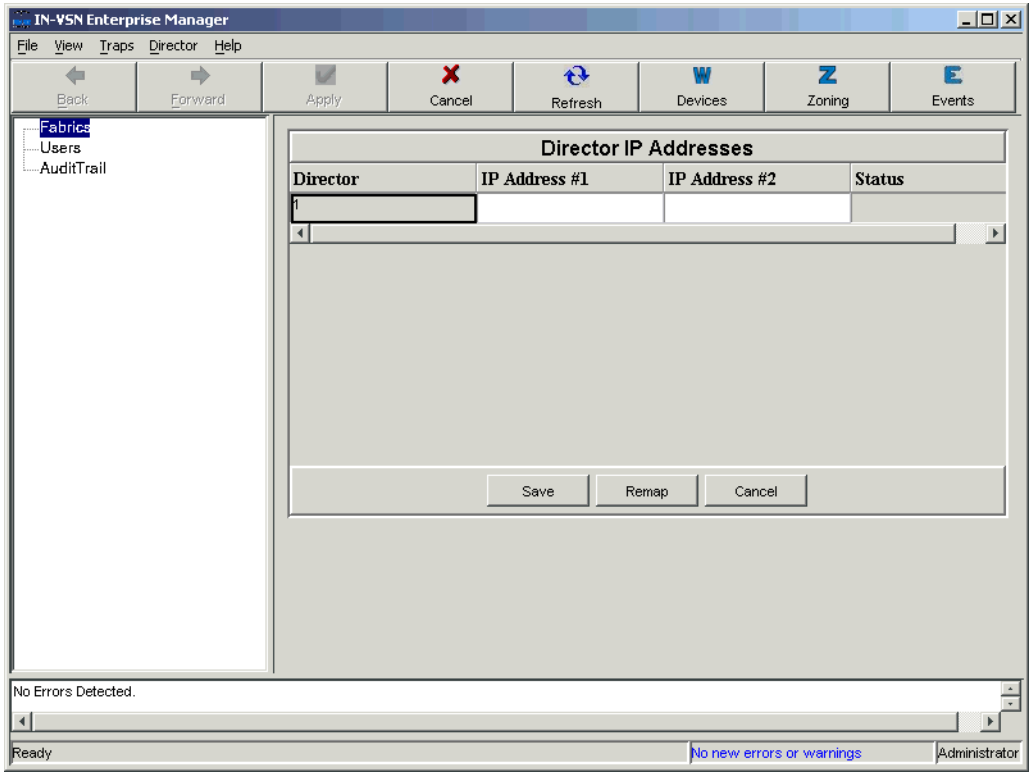


*Figure 3-19 IN-VSN with no fabrics defined*

To make a specific fabric known to the IN-VSN software, enter the IP addresses defined for the switch in the IP Address #1 and IP Address #2 fields.

For each director in the fabric you must use a dedicated row in the Director IP Addresses field. In our case we have one director and therefore only one row. After entering all fabric information, click **Save** as shown in Figure 3-20.

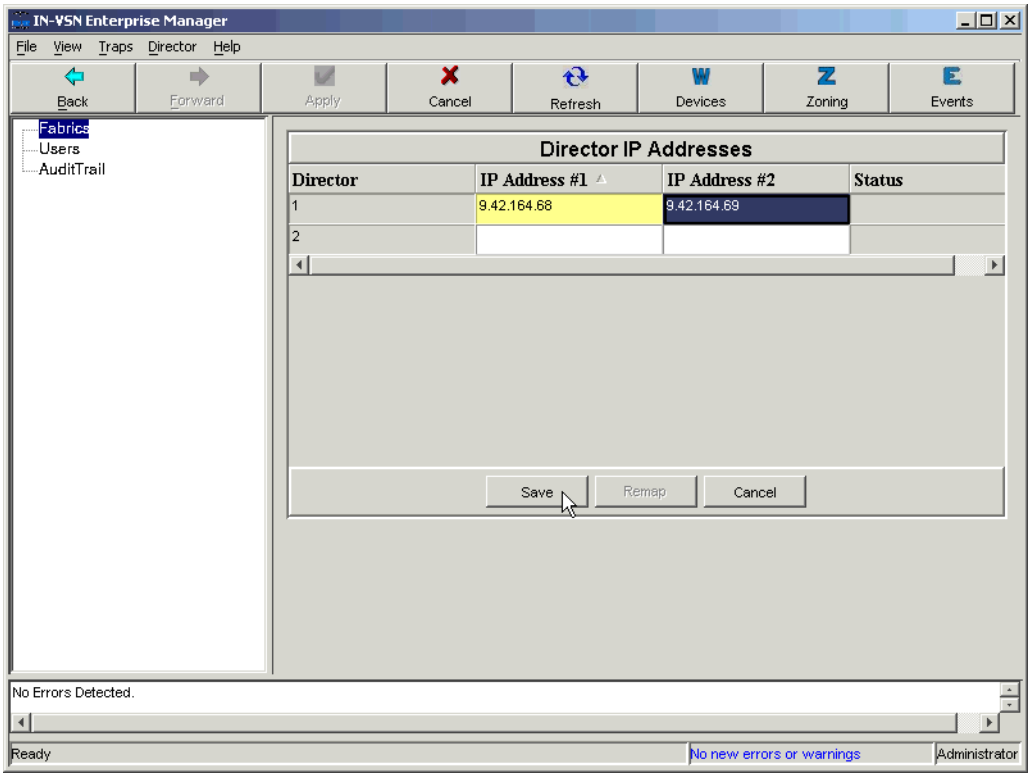


*Figure 3-20 Connecting to a new fabric*

Using the entered IP addresses the Enterprise Manager (EM) auto-discovers the fabric. Each discovered fabric is auto-named based on the principal switch for an E_Port fabric or based on the fabric ID for a T_Port fabric.

Once the fabric is discovered, you can see it by clicking the **Fabrics** icon as shown in Figure 3-21.

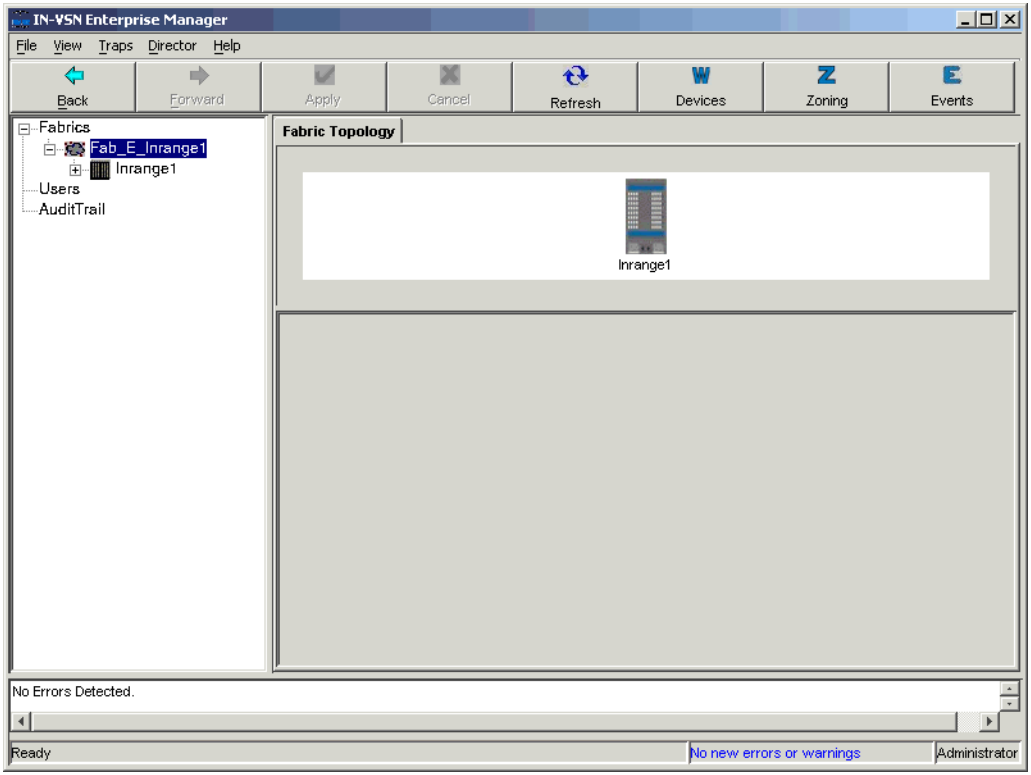


*Figure 3-21 Initial fabric view*

The name of this director is *Inrange1*.

If you had multiple directors in one fabric, they would appear as multiple symbols in the fabric view.

You can use the **Remap** button to rediscover the fabric. Any director to which EM has lost connection is deleted.

## 3.3.5 Setting the director clock

All directors are delivered with preset time and date settings. However, in most cases, these clock settings do not match with the local times.

These clock settings do not affect the fabric functionality at all. However, it is important to set them because this makes reading and understanding the time-stamped logs much easier.

To set the director clock, just click the specific director in the navigation tree and then choose the path **Director—>Set Director Clock** from the main menu.

Enter your desired time settings and apply this by clicking **OK** as shown in Figure 3-22.

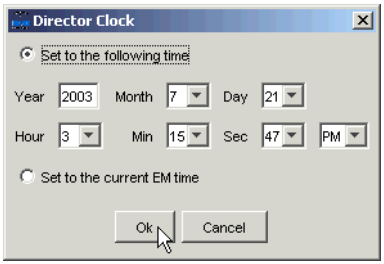


*Figure 3-22 Setting the director clock*

## 3.3.6 Assigning names and aliases

To make it easier to manage large SANs, you can assign names to directors and individual ports. While this step is not mandatory, we recommend giving the directors and ports meaningful names. This makes the management of the SAN much easier.

**Note:** The names you give to the director ports belong to the physical ports of the director, and not to the attached devices. When changing the cabling of the director, we recommend that you also change the names to reflect the new environment.

### Assigning the director's name

To assign a new name to the director, click the specific director in the IN-VSN tool and click the **General** tab as shown in Figure 3-23. You can now change the name and description of the director and click the **Apply** button to confirm the change.

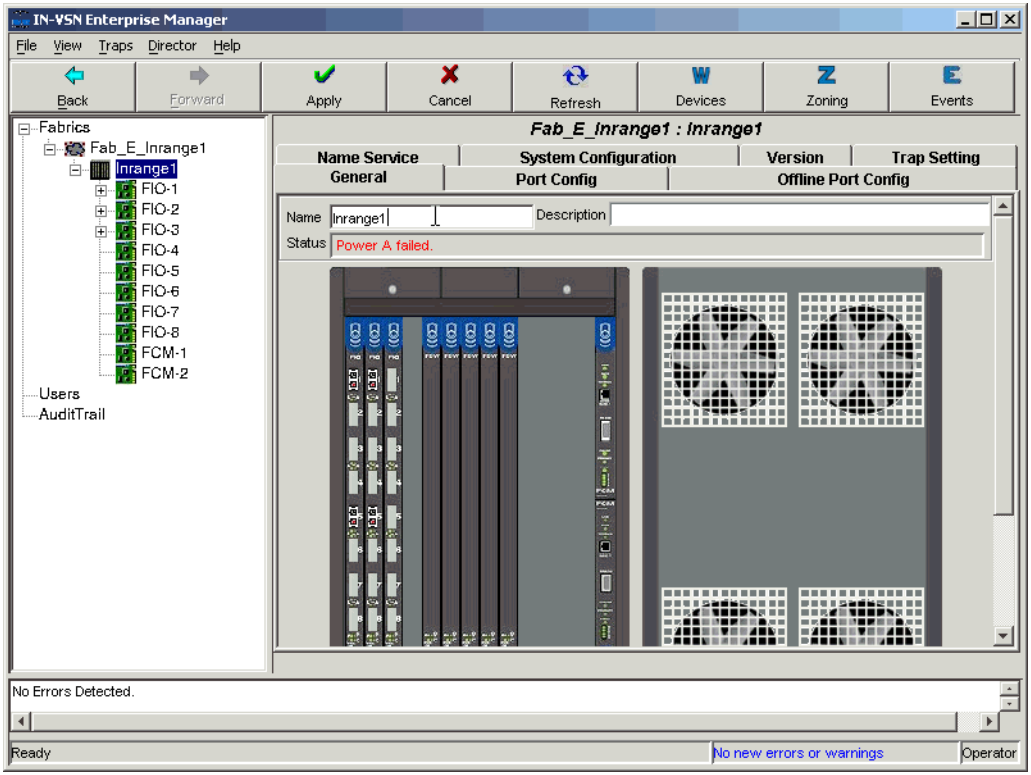


*Figure 3-23 Changing the name of the director*

## Assigning names to ports

To assign names to individual FC ports of a director, click the director in the IN-VSN tool, and click the **Port Config** tab, as shown in Figure 3-24. You can now change the names of any ports you wish to change, and click the **Apply** button to confirm the change.

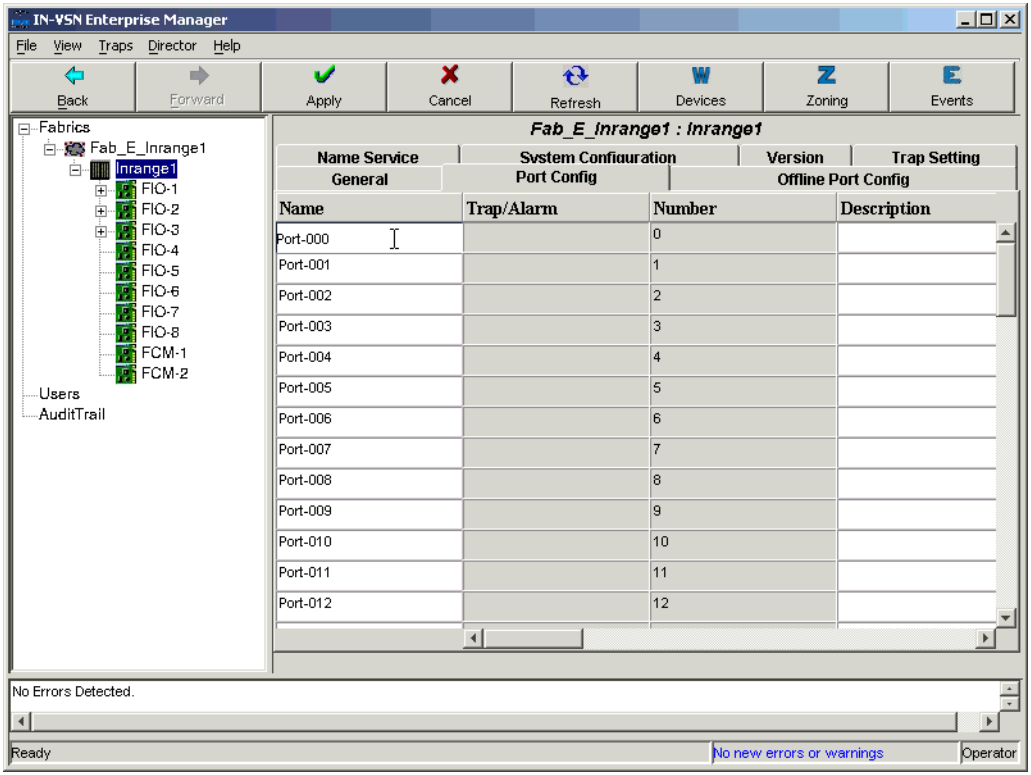


*Figure 3-24 Changing the port names*

Once the new names are effective, they are used throughout the whole IN-VSN fabric management, including zoning, FIO-blade monitoring and name services.

## Assigning names to devices

You can also assign nicknames to the device WWNs that the name server knows about. These nicknames can then be used to define zones.

**Note:** We recommend that you plan the naming convention of the hosts carefully to gain maximum benefit from it. An example of a good naming convention is <hostname>_<interfacename>.

You can define the nicknames by clicking the **Devices** button as shown in Figure 3-25.

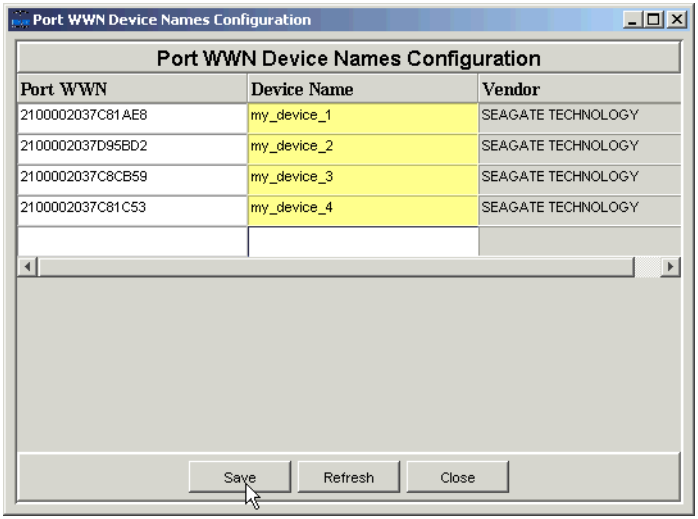


*Figure 3-25 Defining nicknames for devices*

In the window, you see a list of all connected Port WWNs. You can set a nickname for each of these WWNs. Click **Save** to save your changes. Click **Close** to close the window, after you have finished making changes.

## 3.3.7 Attaching loop ports

Today, the storage and server industry is moving rapidly towards switched fabrics. However, there are still a lot of systems that use Loop protocol. For instance, most tape devices, as well as lots of legacy FC host adapters, use FC-AL.

**Note:** To read the following topics, it is useful to understand the differences between terms like director port, loop port, loop devices, initiators, targets:

- A director port is an actual physical port of the FC/9000 director.
- Loop ports are the external ports attached to a director port. Loop ports use loop protocols like private loop or public loop. They are sometimes referred as loop devices or loop nodes.
- Loop devices running public loop are referred as NL_Ports.

To make loop node attachment possible, you have to enable the director ports to autosense loop devices. You can enable autosense for a single port at a time, or all ports at once. Note that the ports in the new XFIO2 module do not support loop devices.

## Enabling loop attachments for a specific director port

To enable loop attachments for a single port, choose the port from the device tree in the left side of the IN-VSN Enterprise Manager window. You will see the properties of that port as shown in Figure 3-26. Activate the Auto Sense Arbitrated Loop Enabled check box and click the **Apply** button.

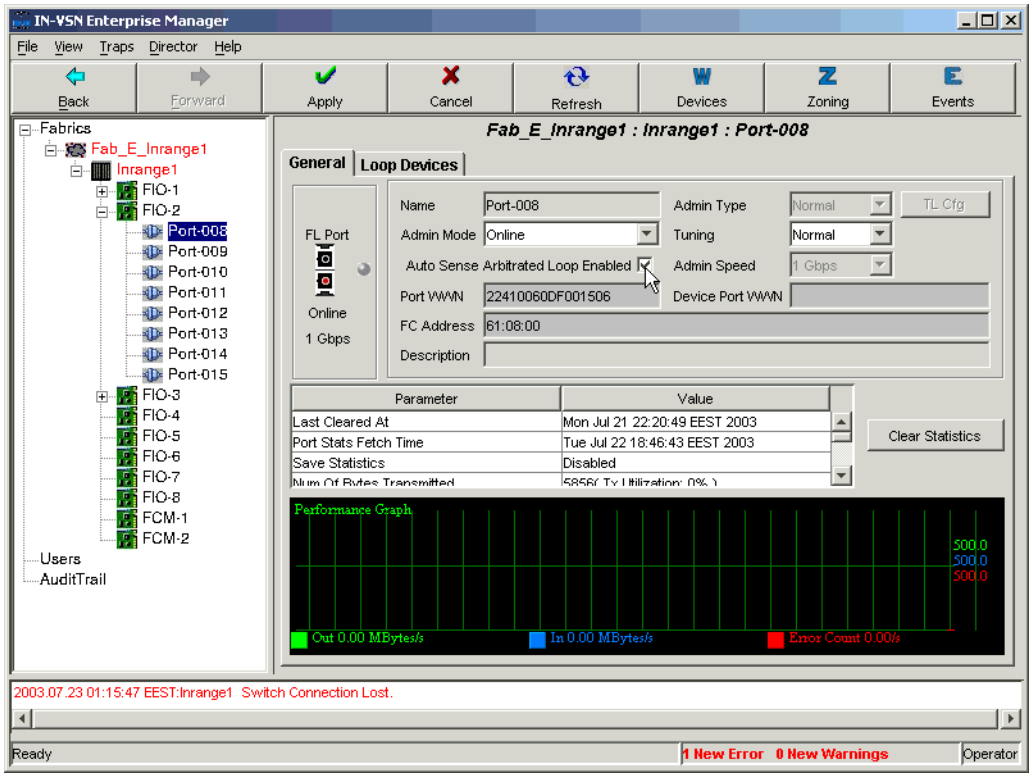


*Figure 3-26 Enabling the loop attachment for a single port*

You will have to confirm the changes, and after that, you are able to use loop attachments on that port.

## Enabling loop attachment for the entire director

To enable loop attachment for all ports in the director, choose the director in the device tree and choose the menu path **Director—>Auto Sense Arbitrated Loop Enable**.

You will have to confirm the changes, and after that, you are able to use loop attachments on all of the ports in the director.

## Loop ports in a name server table

Once you have a port enabled to sense Arbitrated Loop ports and there is actually a loop port attached, you can verify that it has been recognized correctly by reading the name server table. You can do this by choosing the director from the device menu and clicking the **Name Service** tab*- as shown in Figure 3-27.

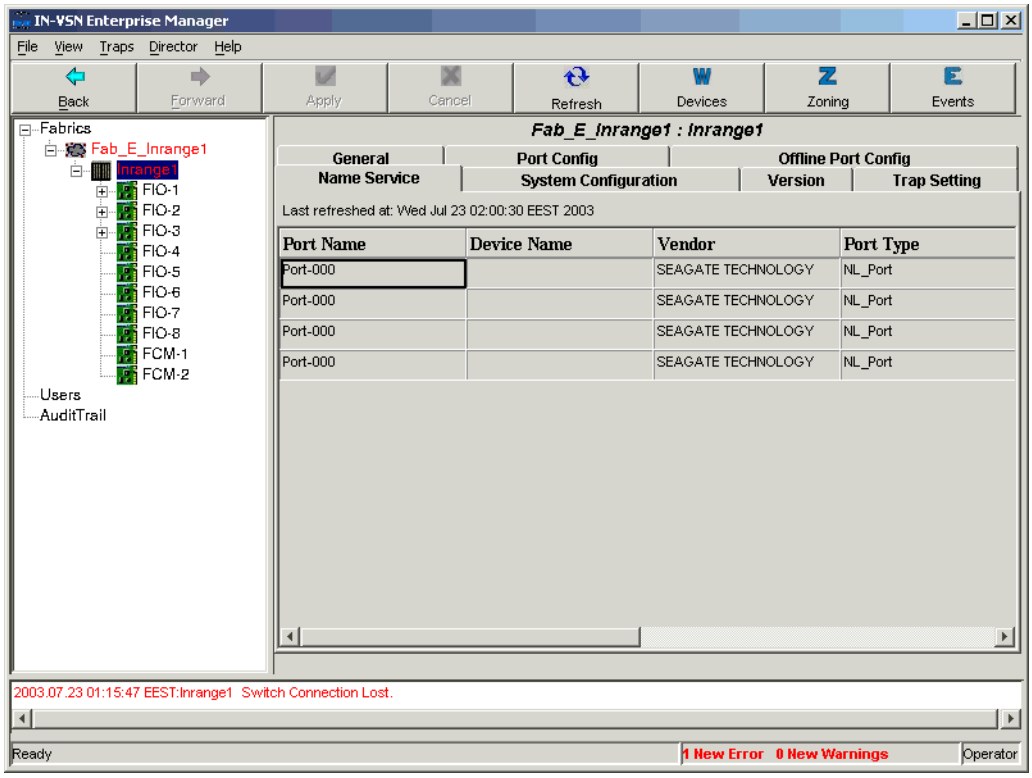


*Figure 3-27 Loop ports in name server table*

Public loop ports are displayed as NL_Ports (Node Loop).

The port *Port-000* was set to enable AutoSense AL and we actually attached this port physically. After logging in, this port is displayed as an NL_Port, since this is a public loop port. There are also four devices connected to the loop in that port.

## Bypassing loop devices

In cases where you have multiple loop devices attached to one director port, you can specify which devices should be actually used in the fabric.

In our example we have four loop devices attached to one director port. This is shown in Figure 3-28.

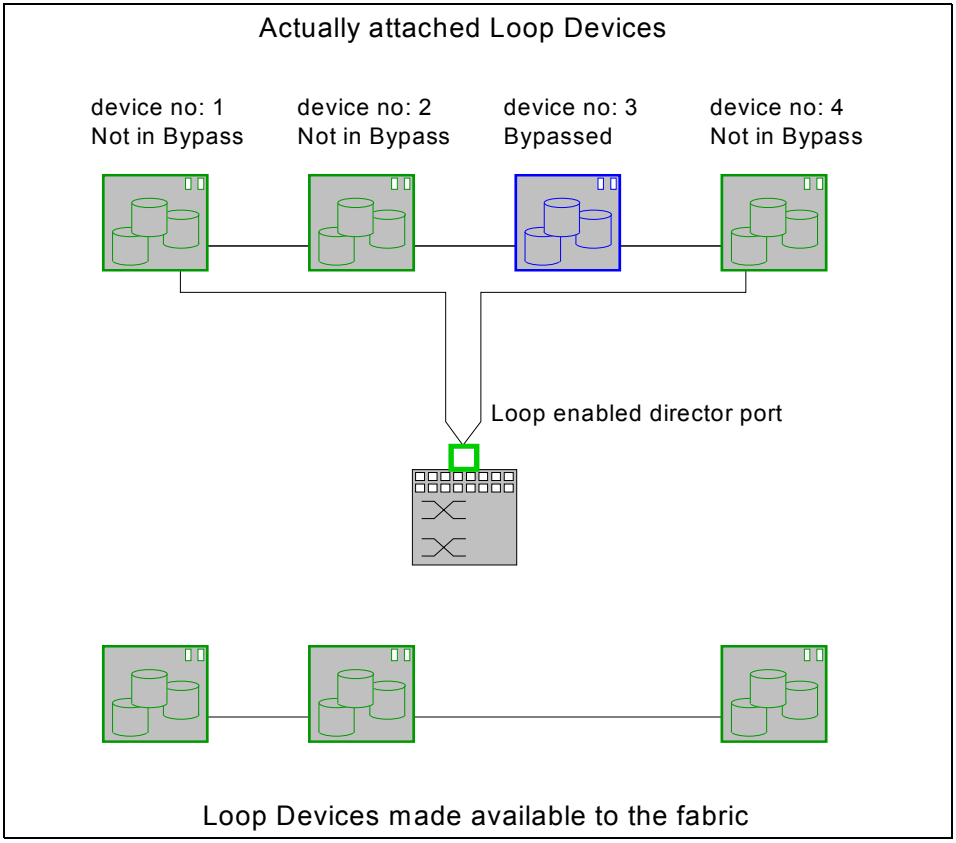


*Figure 3-28 IN-VSN: Bypassing loop devices*

We have decided that device 3 should not be seen in the fabric. To achieve this type of filtering, device 3 must be disabled and therefore set to *Bypassed*. In this case it means that a bypassed device is disabled from being used in the fabric.

This does not change the availability of the external Arbitrated Loop itself. For instance, even if a bypassed device fails, it might affect the other external loop members as well. So, take into account that this kind of bypassing is just a way of filtering, but is not necessarily an improvement of the availability.

To disable a particular device, select it in the **Loop Devices** tab of the specific director port. Click **Disable** to change its state to *Bypassed* as shown in Figure 3-29.

To enable a particular device, use the same menu and click **Enable**.

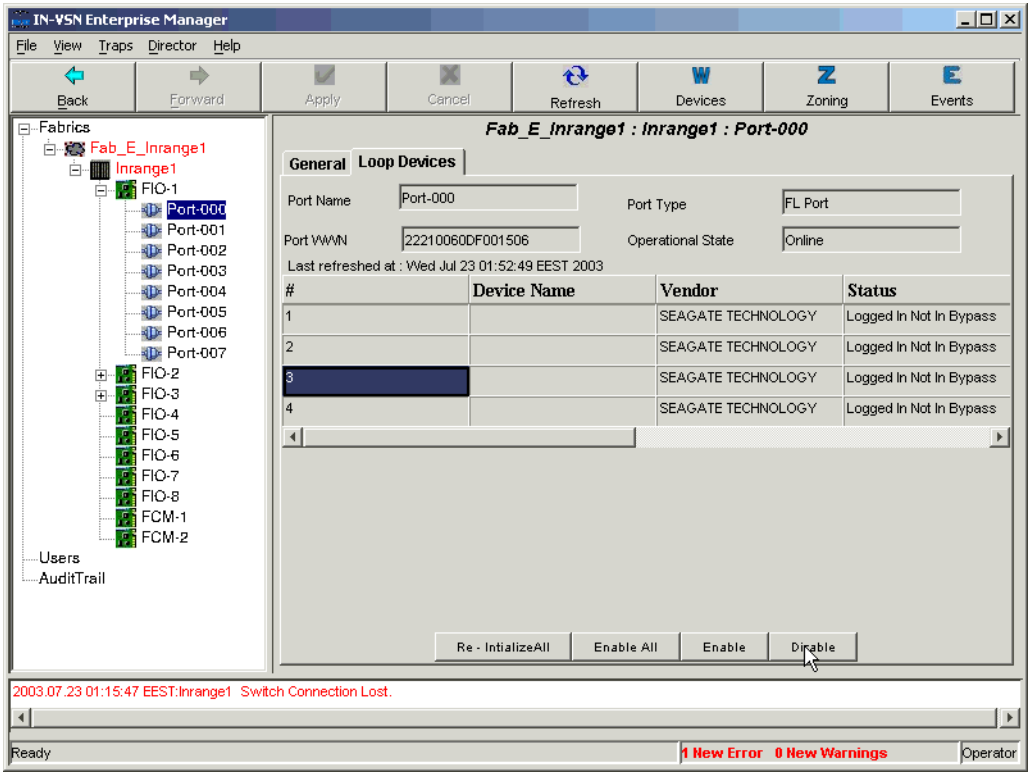


*Figure 3-29 IN-VSN: Enabling and Disabling Loop devices*

You should have at least one device per port left that is *Not in Bypass*. By default, all attached loop devices are set to *Not in Bypass*.

**Note:** Not all loop devices support Bypass. In such cases, even after clicking the **Disable** button, they remain in *Not in Bypass* mode. That means such devices are always enabled.

## Usage of public loop

You can attach public loop initiators or targets without any additional settings provided that the director port is set to enable AL.

However, attaching a public loop port does not automatically mean that this port can talk to any other ports in the fabric. This port is recognized as an NL_Port and can be zoned in the same way as normal N_Ports can be zoned.

### Impact of LIP in the fabric

Ports in loop networks use a process called *Loop Initialization Primitive* (LIP) sequence to establish their port addresses. All members of that loop are involved in a LIP.

Loop initialization occurs whenever there is a change in the layout of a loop, such as adding a new node, a node leaving, or breaks in service in the loop.

The start of a LIP causes data transfers in progress to stop momentarily thereby severely affecting the performance and availability of Arbitrated Loops.

These LIPs are not propagated to other fabric members. This is true even if multiple loop ports are zoned together. Therefore, all LIP impact is limited only to the external physical loop (for example, an FC_AL hub).

## 3.3.8 Implementing zoning

One of the basic purposes of SAN fabric products is to enable or disable communication between the different ports (devices) attached to them.

In most cases it is helpful to limit the potential access of ports. Zoning provides an effective tool to limit and control the communication between fabric ports.

There are multiple reasons to want to limit access:

- We may want to avoid Windows servers seeing all disks in a fabric. Otherwise there would be a high risk of getting signatures written on all disks which would then mean these disks are unusable by other operating systems.
- For security reasons we may want to limit the access to disk with confidential data to only selected servers.
- We would like to get control of the amount of paths a FC host adapter has to a specific disk. This is because not all environments are flexible in their usage of multipathing software.

### Understanding WWN zoning

WWN zoning allows you to designate devices using their WWPN. This means you can group devices by WWNs with WWN zoning. These zones can then be grouped into *zonesets.* All zones within a zoneset are in effect at the same time and only one zoneset is active in the fabric at any given time.

Enterprise Manager lets you manage several zonesets across the fabric with only one of them active at a time. Note that the active zoneset cannot be modified.

One zoneset may be comprised of 1 to 256 zones. That zoneset may accommodate up to 3500 member devices. Note that all zones are present in the fabric and available to any user which may have access to the fabric.

WWN gives greater flexibility to manage a fabric as devices can be moved across the fabric without having to change zoning configurations.

## 3.3.9 Defining WWN zones

To access WWN zoning, click the **Zoning** button and you will see the E_Port zoning window, as shown in Figure 3-30.

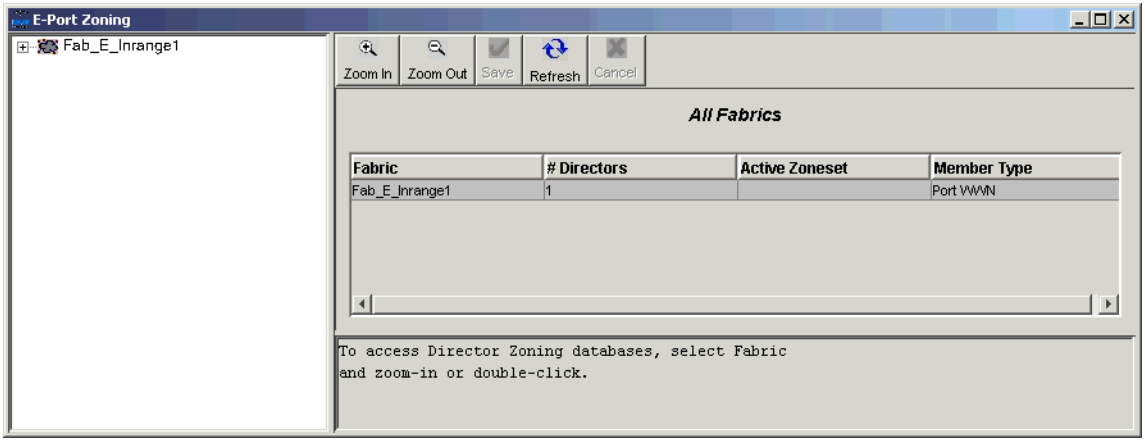


*Figure 3-30 WWN Zoning window*

Note that WWN zoning is referred to as E_Port zoning in the IN-VSN management application. The word “E_Port” refers to the switch mode and does not mean that you have to include ISL ports when building the zones.

The navigation tree in the left hand side of the window displays all ports and attached devices. For each attached device, the tree displays the corresponding WWN and Port ID (Fibre Channel address) under the port node.

**Tip:** Follow the instructions in the message area located at the bottom of the window to browse through WWN zoning administration.

To access zoning information, double-click the fabric name in the list displayed on the right-hand side. Another option is to highlight the fabric and click the **Zoom In** button.

Next, choose the director on which you wish to create the zoneset and click the **Zoom In** button.

**Note:** Although zones and zonesets are defined at the director level, they are valid fabric wide. Zone and zoneset information is propagated automatically to other directors and switches within the same fabric.

As shown in Figure 3-31 we are now viewing the Zoneset screen. All of the zonesets which have been created on this director are listed here along with the active zoneset which has been propagated in the fabric.

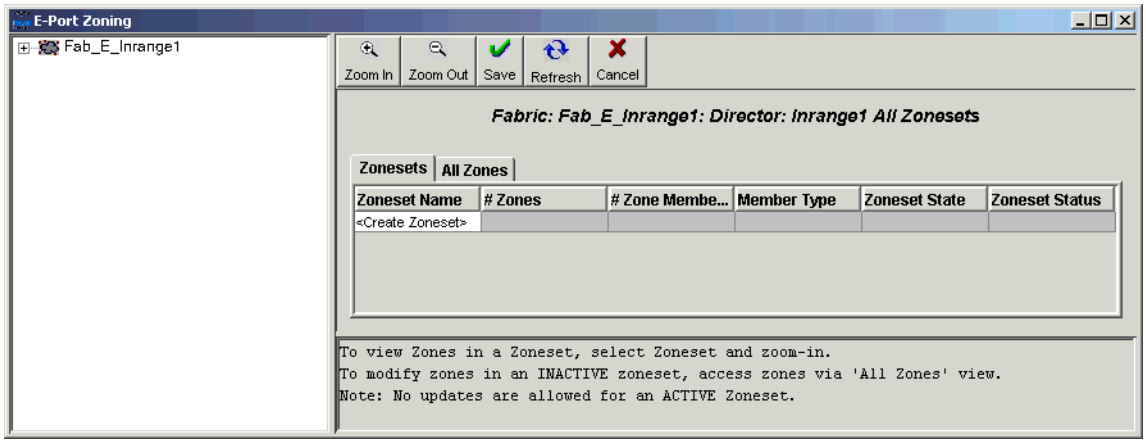


*Figure 3-31 WWN Zoning window - Zonesets tab*

The first step is to create the zones that you use to populate the zoneset. You can access the zones by clicking the **All Zones** tab as shown in Figure 3-32.

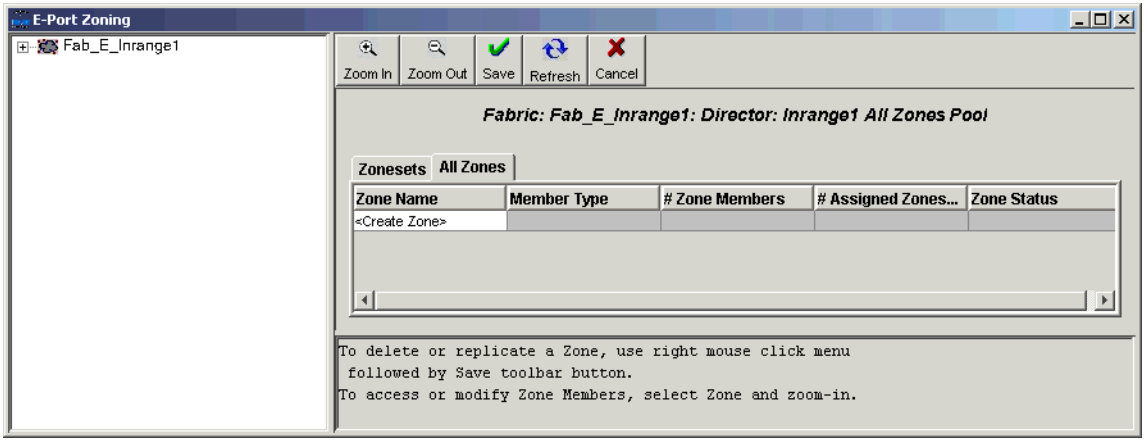


*Figure 3-32 WWN Zoning window - All Zones tab*

Highlight the <Create Zone> field, type in the name of the first zone, and press the Enter key. Your new zone is now created but does not contain any members yet.

Next, choose the newly created zone, and click the **Zoom In** button. You should see an empty list of zone members as shown in Figure 3-33.

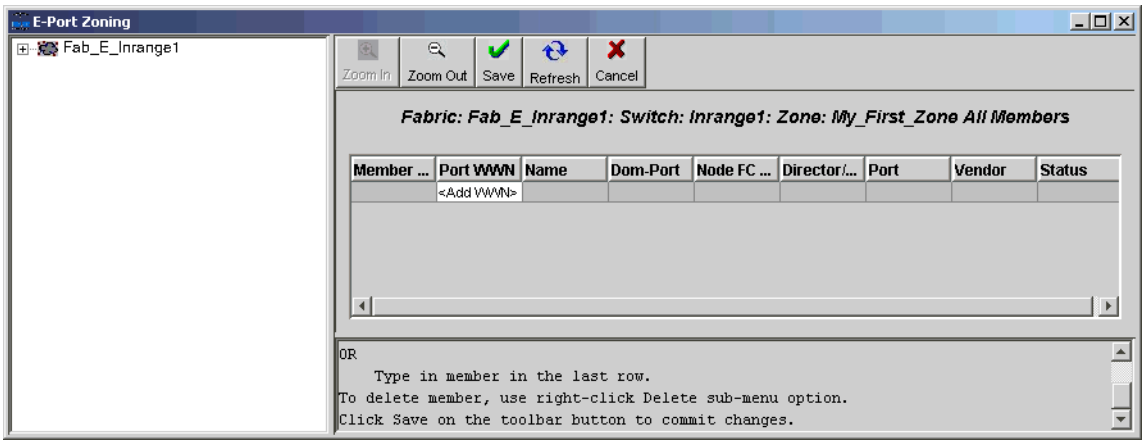


*Figure 3-33 WWN Zoning window - empty member list*

There are a couple of ways that you can add device WWNs at this screen. You can simply type them in, or you can drag and drop them from the tree to the left. You may also right-click over **<Add WWN>** and click **Add zone Members**. A list of known WWNs and their nicknames pops up.

We will drag and drop them from the tree to the left. Note that instead of WWNs we can choose the devices by the nicknames we have given them. The new member list is shown in Figure 3-34.

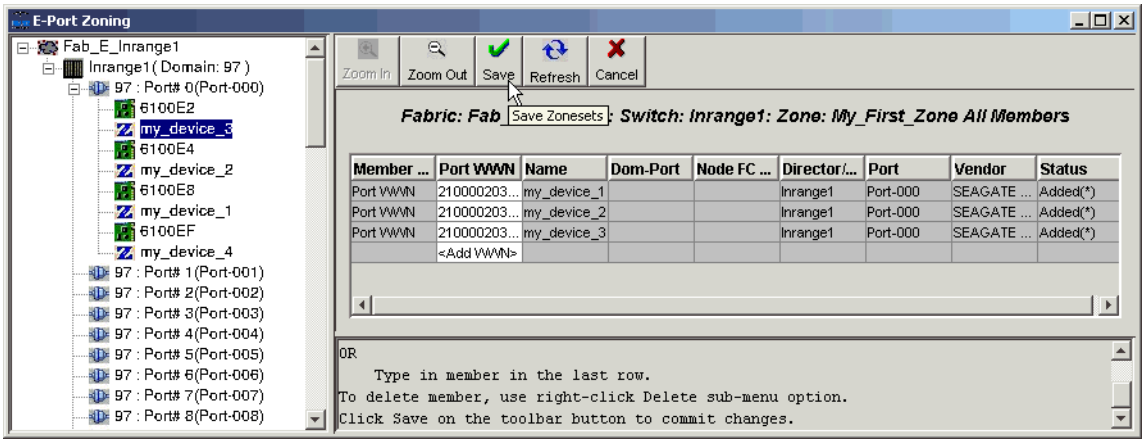


*Figure 3-34 WWN Zoning window - filled member list*

Note that the port names are displayed for information purposes only.

Once the required WWNs have been added, click the **Save** button. To create additional zones, click **Zoom Out** from the previous panel and repeat the same process that you followed to create the first zone.

The next step is to create the *z*oneset. Click the **Zoom Out** button to go back to the list of zones, click the **Zonesets** tab, and highlight <Create Zoneset>. Type in the name of the zoneset and press Enter.

**Note:** The zoneset name *must* begin with an alphabetic character (A-Z, a-z). Use alphanumeric ASCII characters to create the rest of the zoneset name. You can also use the following three characters in a zoneset name: “$” or “_” or “-”.

To add zones to this zoneset, right click over the zoneset and choose **Add Zones** in the contextual menu as shown in Figure 3-35.

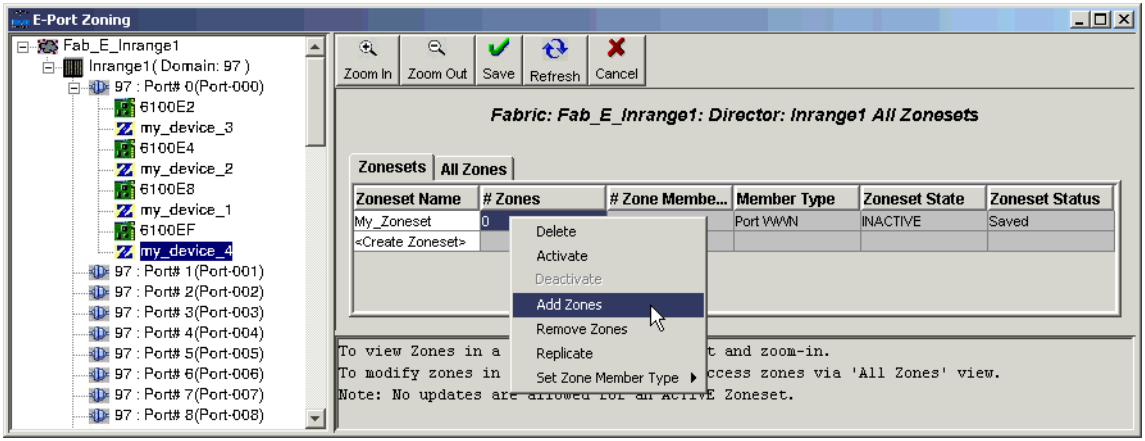


*Figure 3-35 WWN Zoning window - adding zones to zoneset*

The window listing all the available zones is displayed as shown in Figure 3-36.

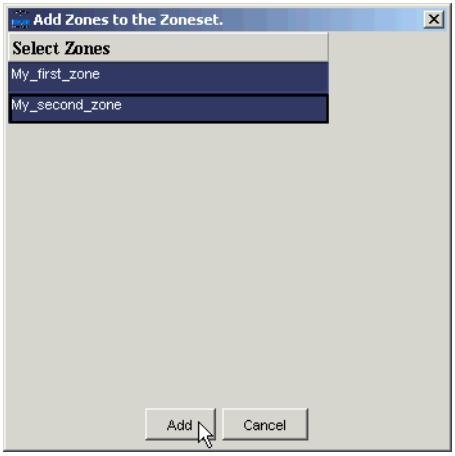


*Figure 3-36 Selecting zones for zoneset*

Select the zones you want to add. You can select multiple entries by holding down the Ctrl key while selecting the zones. When you have all selected all the zones you want to add to the zoneset, click **Add**.

Go back to the main window and click **Save** to save the zoneset configuration.

The zoneset is now defined but you still have to activate it. To activate the zoneset, right-click it and choose **Activate** in the contextual menu as shown in Figure 3-37.

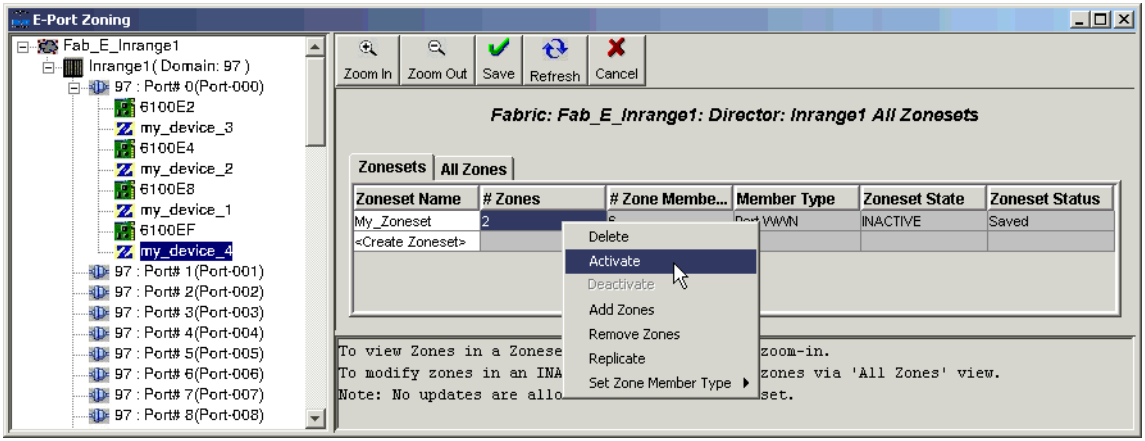


*Figure 3-37 E_Port mode zoning - Activate the zoneset*

Confirm that you want to activate the zoneset at the confirmation screen by clicking **Yes**. The result is shown in Figure 3-38.

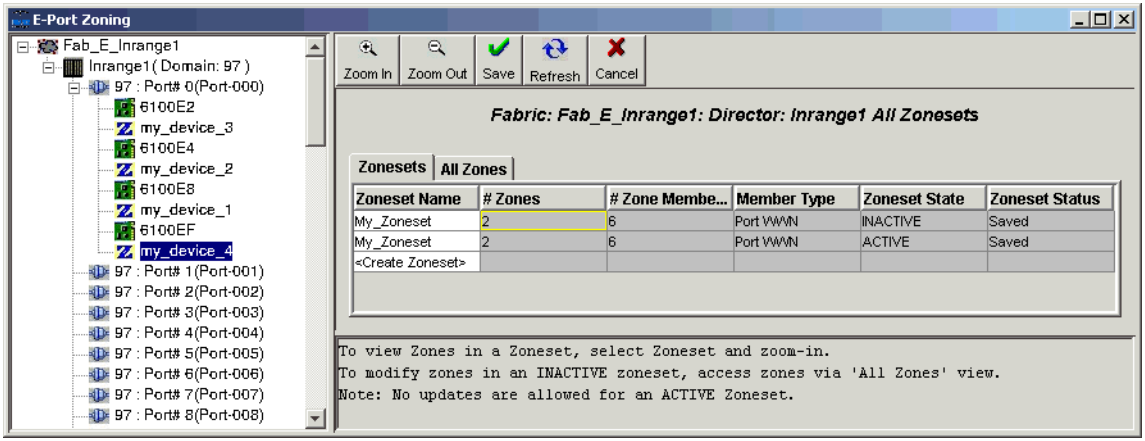


*Figure 3-38 E_Port mode zoning - Zoneset active*

Note that when you activated the zoneset an inactive duplicate of the zoneset was created. This allows you to modify a copy of the activated zoneset since we cannot modify the active zoneset itself. You could then activate the copy with any modifications you might want to make.

## 3.3.10 Implementing multi-switch fabrics

Since the CNT directors support large port counts within a single domain, from 24 all the way to 256 ports, the motivation to use ISLs with the CNT director is more likely to be for implementations involving distance such as connecting remote data centers in a campus or metro environment. Fabrics of up to eight directors are supported.

For software release 3.0 and later, the CNT directors implement the industry standard E_Port (as defined in ANSI FC-SW-2) to provide the interconnect protocol for ISLs. The advantage of implementing the FC-SW-2 methods for E_Port is to allow the CNT director to interoperate with a large number of Fibre Channel switches, directors, routers, and bridges, that are also designed to conform to the FC-SW-2 standard.

The CNT directors support ISLs for creating cascade, mesh, core, and multi-stage topologies for open systems (FCP) as well as support for FICON cascaded environments. WWN zoning, as defined in FC-SW-2, can be used for the grouping of ports and devices in an extended fabric topology the same way it is implemented in a single domain fabric.

## 3.3.11 Managing the IN-VSN server

In this topic we describe the parameter setup which can be performed without having any directors known to the IN-VSN software:

- Database backup
- Call Home and Page Home settings (CE's responsibility)

### Database backup

CNT provides the ability to backup all director related data to an external disk. This is called database backup. You can use this backup to restore your configuration data. Database backup settings are performed using the IN-VSN server.

You can run a manual backup of the IN-VSN database by choosing **File—>Backup** from the menu of the IN-VSN server. You can then specify the path for the backup file as shown in Figure 3-39.

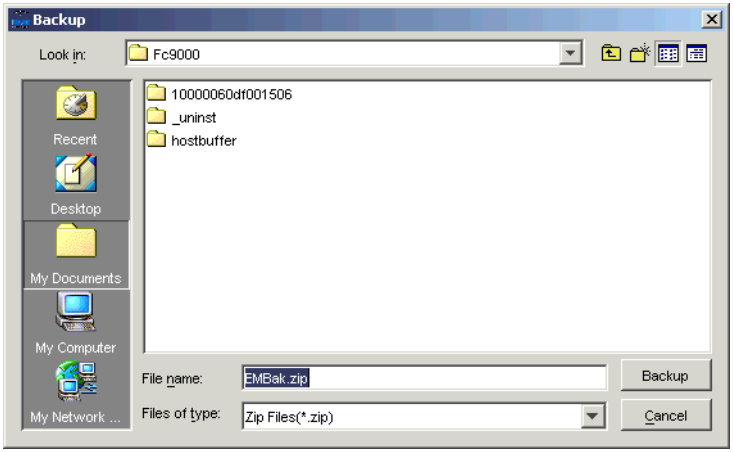


*Figure 3-39 IN-VSN database backup*

To achieve a higher level of disaster recoverability we recommend that you place this database file on a different drive than the one you are using for the database. You can also use a network mapped drive, copy the backup file to a removable media for safekeeping, or copy it to an external backup system, such as Tivoli Storage Manager.

The database file created by this manual backup could be used for IN-VSN database restore. The IN-VSN server also supports automatic database backups. We recommend that you use this option to ensure that you always have a current backup of the database.

To set up the automatic backup option you have to log into the IN-VSN server. You can use the same user names and passwords as with the IN-VSN client. You can login to this by choosing the option **File—>Logon** from the IN-VSN server menu. You will see a logon window as shown in Figure 3-40.


*Figure 3-40 IN-VSN server logon window*

Enter your username and password and click **OK** to log into the server.

After you have logged into the server you can set up the automatic backup by choosing the path **Configuration—>AutoBackUp** from the menu. You will see a window as shown in Figure 3-41.

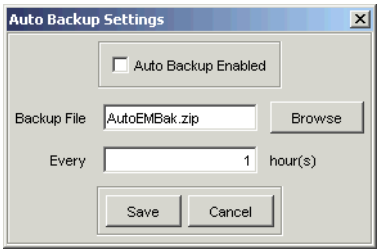


*Figure 3-41 Defining automatic backup*

Here you can choose where the backup is written, and how often a new backup copy is made. Remember to activate the Auto Backup Enabled check box as well. If possible you should use a network drive for the backup, or copy it to an external backup server, such as Tivoli Storage Manager.

In our case we have chosen to write the backup to the network path **`H:\backup.zip`** every four hours. Our settings window is shown in Figure 3-42.

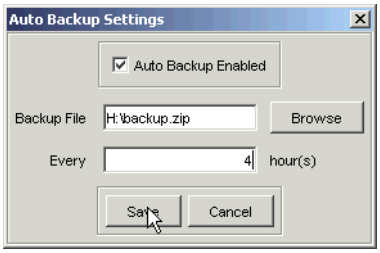


*Figure 3-42 Automated backup settings*

### Setting up page home/call home

In some environments it may be useful to have CNT or IBM automatically contacted via telephone or pager by the IN-VSN server. This is called Page Home or Inventory Update Call.

However, these parameters should be set up by IBM or CNT Customer Engineers only, and cannot be changed with any other user id than maint.

## 3.3.12 Security considerations

As a SAN may store a great deal of sensitive and often confidential corporate information, it might be interesting to consider some security issues such as:

- Physical security
- User management

We describe these as they relate to the usage of the CNT FC/9000 director.

### Physical security

To limit the physical access to the SAN, consider these possibilities:

- Use the CNT cabinet with locked doors to prevent access to the cabling.
- Limit the access to any patch panels used in your SAN infrastructure.
- Use access control methods for the datacenter facility.

### User management

As described in 3.3.3, “Defining users” on page 421, CNT provides different types of users.

User settings are saved in the IN-VSN servers database, but not on the directors themselves. To ensure the ability to recover from an IN-VSN server failure, we advise you to use the Auto-Backup function. This backup setting is done using the IN-VSN server interface as described in 3.3.11, “Managing the IN-VSN server” on page 444.

You can have multiple users defined and also logged into the IN-VSN software at the same time. Be aware of these important issues:

- Different users with the same level (for example, admin) can have access to the fabric at the same time. This could result in chaos since both active users are allowed to make changes. Consequently, we recommend that only a limited number of people have admin rights to the same fabric.
- The IN-VSN default users (admin, oper, view) are always the same globally, including their passwords. Consequently, you should not consider them as secure enough. Therefore, we advise you to delete them or change their passwords after you have added your own users.

## 3.4 Monitoring and maintenance

This topic briefly covers these topics:

- Communication protocols to manage an CNT SAN
- Microcode load information
- Monitoring user activities
- Using the CNT event log

### 3.4.1 Management communication protocols

Management communication can be done using different physical interfaces:

- Director:
  - RS-232 Serial Interface
  - Ethernet 10/100 Base T Connector
- Management PC:
  - Modem connectivity
  - RS-232 Serial Interface
  - Ethernet 10/100 Base T Connector to directors
  - Industry standard network interfaces to the corporate network

Using these physical interfaces the following management services can be used:

#### IN-VSN management software

Client/Server based management tool dedicated for users.

#### Maintenance interface using the RS-232 serial interface

Terminal based interface to change basic settings like director ID's or IP settings. This is dedicated for IBM or CNT Customer Engineers.

#### Simple network management protocol (SNMP)

An interface to integrate CNT directors into the corporate systems management network.

#### Trivial file transfer protocol (TFTP)

TFTP is used to load new firmware to directors and to retrieve and/or change director management configuration settings (for example, IP, SNMP). This service is dedicated to IBM or CNT Customer Engineers.

#### Call home functionality

Call home enables the IN-VSN software to propagate event information and configurations to CNT or IBM systems. This allows better response times in case of component failures. Call home is disabled by default but can be enabled by Customer Engineers. A modem connection for the Management PC is required.

#### Changing the IP address

The IP settings of CNT directors should be done by IBM or CNT Customer Engineers.

The change of the IP settings of a CNT FC/9000 director is performed using a Terminal session via the RS-232 serial interface.

Each FCM blade has its own IP settings and also its own serial interface.

You should always keep both IP settings valid to ensure management access in the event of FCM failures.

To change both addresses, you have to login to both FCM blades one after the other. First you login to one FCM to set its IP settings, and then you login to the other FCM to change its IP settings.

The change process for IP settings involves a booting of the FCM. However, this local FCM boot process does not affect other director parts. Neither the alternate FCM blade nor the director functions are disturbed.

This results in a non-disruptive IP change.

### 3.4.2 Microcode-loads

Microcode-loads can be done non-disruptively using the IN-VSN Enterprise Manager in conjunction with the TFTP interface. However, this action should be performed by IBM or CNT customer engineers.

### 3.4.3 Monitoring user activities

IN-VSN management software allows you to monitor user activities using an Audit Trail.

All user levels (*admin*, *oper*, *viewer*) can access the audit trail.

All activities are logged and categorized into different types. These are some examples of these operation types:

- User login
- IP address of user login
- Fabric definition
- User definitions
- Name server zoning
- Switch name changed

To look at the audit trail, click **AuditTrail** in the navigation tree and you will see the audit trail for the system as shown in Figure 3-43.

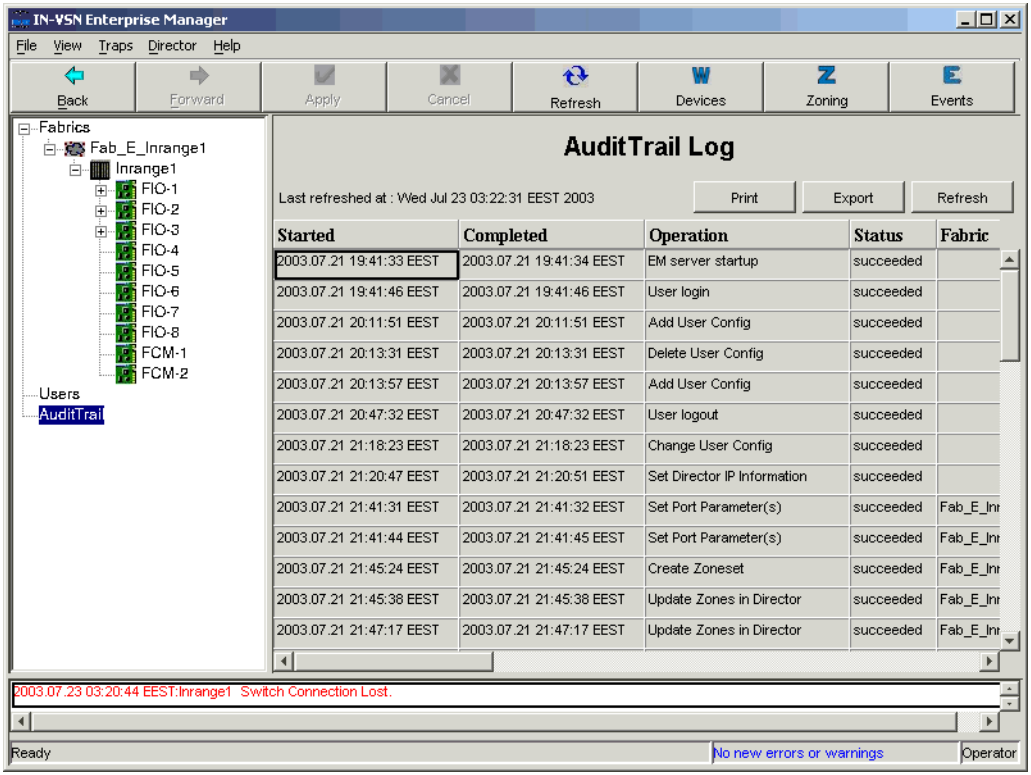


*Figure 3-43 Using the CNT Audit-Trail to monitor user activities*

## 3.4.4 Using the IN-VSN event log

The event log contains all important events that have occurred. This includes events triggered by users and events caused by other external or internal influences such as FRU failures or losing power.

To read the event log, click **EventLog** button and you will see the event log as shown in Figure 3-44.

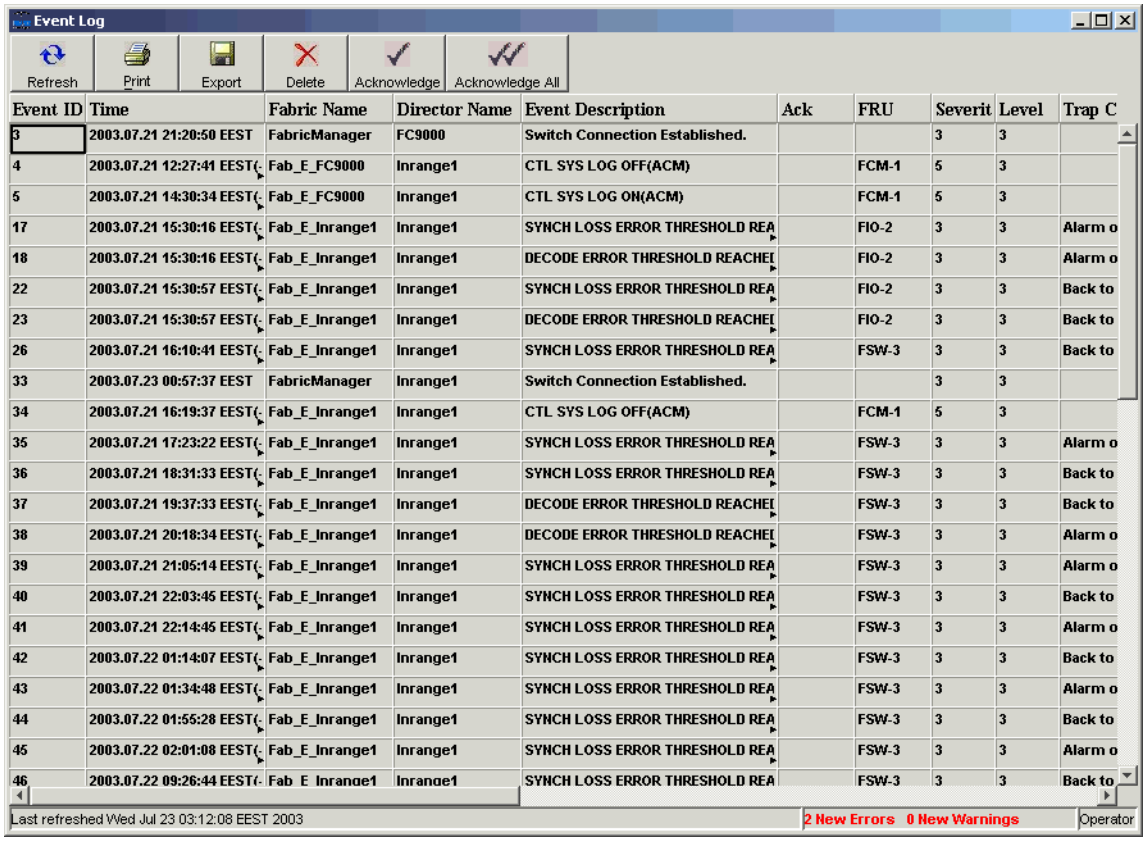


*Figure 3-44 IN-VSN event log*

This event log contains information for all fabrics managed by this IN-VSN server.

To file this log for future usage you are able to export it by clicking the **Export** button. Two file types can be used to export the log:

- Comma Separated Value Files (*.csv)
- Text Files (*.txt)

## 3.5 Interoperability mode implications

When the CNT FC/9000 director is configured into E_Port mode, it is also effectively running in interoperability mode. There are no settings that have to be changed in the director before connecting it to switches from other vendors.

### 3.5.1 BladeCenter attachment

The settings required for both the IBM BladeCenter and the CNT are described in Table 3-4. We recommend that you manage zoning with the CNT tools only.

*Table 3-4 Configuration settings for interoperability*

| **Configuration Options** | **IBM BladeCenter** | **CNT FC9000** |
|---|---|---|
| **Firmware Version** | **V1.4.0.42-49** | **V3.0.3.2** |
| Domain ID Allowed Range | 97 - 127 | 97 - 127 |
| Domain ID Configured | 125 is required | 99 |
| Domain ID Lock | True | Yes |
| Switch Port Mode | G_Port or GL_Port | E_Port |
| Switch Port Speed | 1 Gb/s | 1 Gb/s |
| Principal Switch Priority | 254 | Not applicable |
| R_A_TOV | 10000 (milliseconds) | 10000 (milliseconds) |
| E_D_TOV | 2000 (milliseconds) | 2000 (milliseconds) |
| Default Zone State | Disabled | Not applicable |
| Interop Mode | Not applicable | Not applicable |
| Zoning | WWPN based | WWPN based |
| IO Stream Guard | Disabled on ISL (E Port) | Not Allowed |

Before proceeding to configure the BladeCenter and the CNT, identify the common and product specific configurations and also understand the implications of migrating to interoperability mode.

**Attention:** Changing the fabric configuration from native mode to an interoperability mode is a disruptive process and requires planning before implementation. The domain ID values supported to merge the IBM BladeCenter and CNT director is from 97 — 127. In an environment with AIX and HP hosts installed, the HBAs (FC devices) should be deleted and reconfigured for new device discovery. Another implication of fabric migration from native to interoperability mode is that the multipathing and persistent binding functions requires re-configuration using the new fabric address.

### 3.5.2 BladeCenter initial configuration

The following are essential for the BladeCenter:

- Firmware version V1.4.0.42-49
- Domain ID set to 125
- Principal switch priority set to 254
- Zoning is based on FC-SW2 compliant
- I/O StreamGuard is enabled on the E_Port

Optionally, it is preferred to configure the E_Port speed to 1 Gb/s.

If any of these settings need to be changed, use the **set config switch** command as shown in Example 3-1.

*Example 3-1 The set config switch menu*

```
Login: USERID
Password: xxxxxxxx
FCSM: USERID> admin start
FCSM (admin): USERID> config edit
FCSM (admin-config): USERID> set config switch

A list of attributes with formatting and current values will follow. Enter a
new value or simply press the ENTER key to accept the current value. If you
wish to terminate this process before reaching the end of the list press 'q' or
'Q' and the ENTER key to do so.
AdminState         (1=Online, 2=Offline, 3=Diagnostics) [Online ]
  BroadcastEnabled   (True / False)  [True ]
  InbandEnabled      (True / False)     [True ]
  DefaultDomainID    (decimal value, 1-239) [1 ] 125
  DomainIDLock       (True / False)   [False ]  True
  SymbolicName       (string, max=32 chars)  [FCSM ]  FCSM
  R_T_TOV            (decimal value, 1-1000 msec)  [100 ]
  R_A_TOV            (decimal value, 100-100000 msec) [10000 ]
  E_D_TOV            (decimal value, 10-20000 msec) [2000 ]
  FS_TOV             (decimal value, 100-100000 msec)  [5000 ]
  DS_TOV             (decimal value, 100-100000 msec)  [5000 ]
```

```
  PrincipalPriority  (decimal value, 1-255)  [254 ]
  ConfigDescription  (string, max=64 chars)  [IBM BladeCenter(TM) 2-port Fibre
Channel Switch  Module]
Finished configuring attributes.
  This configuration must be saved (see config save command) and activated (see
config activate command) before it can take effect.
  To discard this configuration use the config cancel command.
FCSM2 (admin-config): USERID> config save
FCSM (admin): USERID> config activate
The configuration will be activated.  Please confirm (y/n): [n] y
```

### Limitations on merging BladeCenter and CNT

When merging, the maximum number of switches (domain IDs) that can be configured depends upon the CNT switch model:

- On 8/16 port directors, 16 domains are allowed.
- On 64 port directors, 56 domains are allowed.
- On 128 port directors, 48 domains are allowed.
- On 256 port directors, 32 domains are allowed.
- Only WWPN based zoning is supported.

The IN-VSN Manager or BladeCenter SAN Utility can be used to configure zoning — but do *not* use both management utilities simultaneously.

## 3.5.3 CNT configuration

The IN-VSN manager 8.1.0 application is used to configure the director for BladeCenter attachment as shown in Figure 3-45.

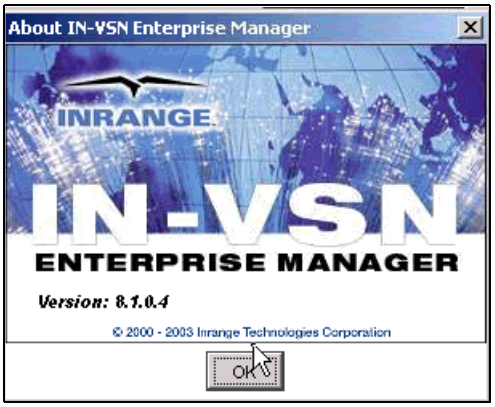


*Figure 3-45 IN-VSN application*

The system configuration menu is available from the IN_VSN client application by selecting the director icon.

Click the **Version** tab to make sure that the firmware level is FC 3.0.3.2 or greater.

From the IN-VSN Enterprise Manager menu, select the **System Configuration** tab, as shown in Figure 3-46.

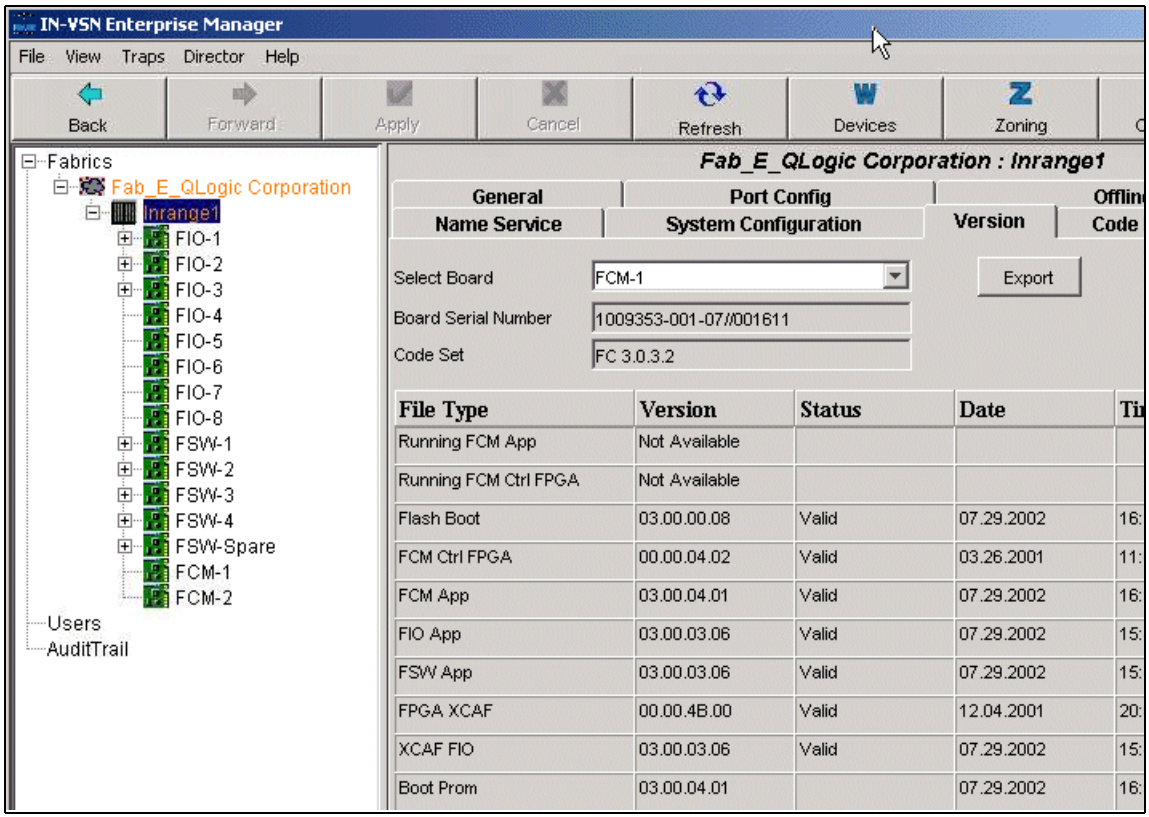


*Figure 3-46 Firmware version check*

The **System Configuration** menu will display. We set the director system configuration options to those shown in Figure 3-47.

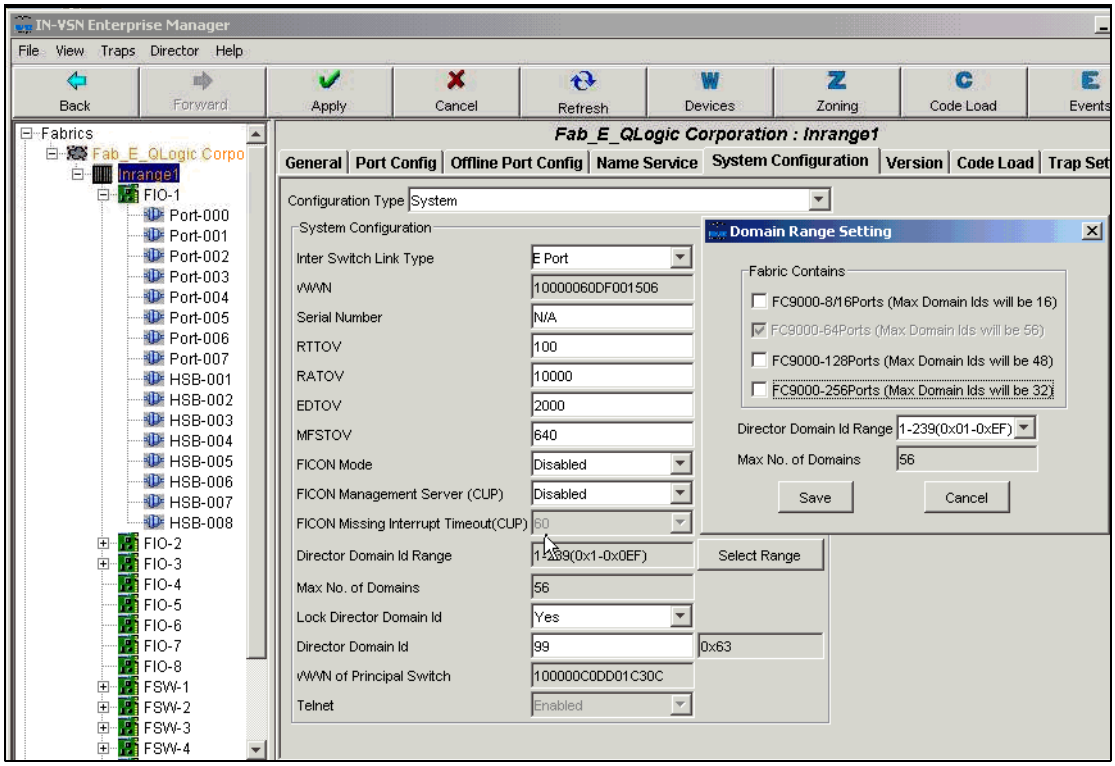


*Figure 3-47 System configuration menu*

Once all the configuration steps are complete and have been applied, verify if the fabric has merged by checking the E_Port status on both the director and the BladeCenter fabric view.

As shown in Figure 3-48, the E_Port status on the director as is Online at a speed of 1 Gb/s.

From the IN-VSN manager menu, select the port on the FIO card. In the following example the BladeCenter is attached to the third port (Port-018) on the third FIO card (FIO-3) on the director as shown in Figure 3-48.

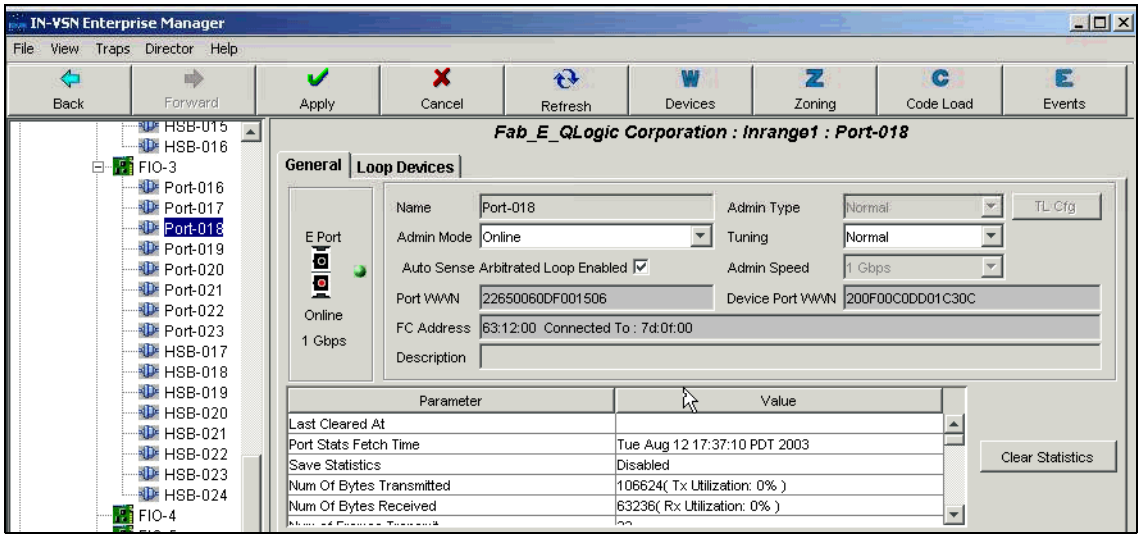


*Figure 3-48 E_Port Online status*

## Zoning

Zoning the CNT director is described in 3.1.5, “Zoning” on page 409 and zoning from the BladeCenter is described in Chapter 5.4, “Zoning” on page 608.

To verify that the **Active Zoneset** name and **Member Type** matches the fabric you expect it to be in or join, use the menu as shown in Figure 3-49.


*Figure 3-49 Active Zoneset view*

The BladeCenter is now successfully interoperating with the director.

4

# Implementing a SAN with McDATA

The McDATA family of products range from the director class products ED-6140 and ED-6064, to the entry level switches Sphereon ES-4500 and ES-3232. All are designed to meet the continuously growing demands for zero downtime along with high speed (2 Gb/s) connectivity to adapt to the fast emerging technologies such as on-demand computing, storage virtualization, and digital media. With dedicated bandwidth available via point-to-point connectivity between high end servers (open systems) and the enterprise class storage devices over the Fibre Channel protocol (FCP), the McDATA directors also provide support for OS/390® mainframes with FICON technology. Directors are usually installed as the core device in the backbone of an enterprise SAN.

The new Sphereon ES-4500 and ES-3232 switches with 2 Gb/s speed address the demands for high speed SAN connectivity for midsize to small enterprise businesses. The ES-4500 and ES-3232 switches are deployed as edge switches in a core-to-edge topology that allows seamless scalability using the flex port feature.

The firmware release EOS 5.01 provides common features such as hard zoning, OpenTrunking, SANtegrity, and hot code activation (HotCat). The same device firmware is installed across the entire McDATA product line.

## 4.1 Introducing the McDATA Directors

The ED-6140 and ED-6064 with UPM cards are almost identical except for the port density. The ED-6140 has 140 ports and the ED-6064 has 64 ports and both support 2 Gb/s speed. Also, a difference between the ED-6140 and ED-6064 directors is the operating voltage. The ED-6140 operates at 240 volts versus the ED-6064, which can auto-sense and operate at 110 or 220 volts. The ED-6064 with FPM cards only supports 1 Gb/s speed and can be upgraded to support 2 Gb/s port speed.

These are the key features available with the ED-6140 and ED-6064:

- **High Availability:** In order to satisfy customer demands for high availability, the directors are shipped with redundant configuration of critical components dual CTP cards, S-BARs, power supplies, cooling fans, hot swappable GBICs and the port cards. If one component fails due to a hardware defect, then the second component still keeps the director running and the nodes connected to it are still operational. Hot Code activation (HotCat) is a software feature that allows for dynamic firmware upgrades without having to schedule downtime.
- **Scalability:** With a high port count of 140 ports per single ED-6140 and 64 ports on the ED-6064, the port count can grow from a stand-alone unit to a multi-switch fabric of up to 31 switches using the E_Port. The core-to-edge design allows for the install of a new switch into the existing SAN without causing any major disruption to Fibre Channel link.
- **Performance:** The McDATA director architecture includes high speed ASICs that support 2 Gb/s full duplex port speed, 60 buffer credits per port to handle throughput traffic for distances up to 120 km. The virtual output queuing (non-blocking) architecture eliminates the head of line blocking, a known problem with the old switching technology and manages link level and end to end congestion points by adequately allocating buffering and queuing resources. The OpenTrunking technology, a software feature that provides load balancing by continuously monitoring the inter switch links for congestion points and automatically distributes the load over the under utilized links.
- **Class Of Service Support:** The Fibre Channel signaling protocol provides several classes of transmission service that support framing protocol and flow control between ports. Directors support class 2, class 3, and class F traffic between the end devices and the switches over E_Ports.
- **Multiple Topology Support:** The directors support point to point topology between N_Port devices, switched topology through inter switch links using class F traffic and indirectly provide support for arbitrated loop topology using Sphereon ES-4500 switch through an ISL inter connecting to the director in the fabric.

- **Manageability:** The McDATA products are managed by the Enterprise Fabric Connectivity Server that runs server and client services for the EFC Fabric Manager and EFC Product Manager applications on a McDATA supplied PC. The EFCM server with a central point of control can manage up to 48 devices in a single or multiple fabrics.

  The SANpilot is a limited, cost effective Web based management interface used by pointing the Web-browser to the Embedded Web Server (EWS) in the director or switch. The EWS gives a single switch management view, not a fabric view.

  These products also provide the Command Line Interface (CLI) to perform configuration and debugging via an out-of-band telnet connection. More information on the Command Line Interface may be found in the *McDATA Command Line Interface user manual* P/N 620-00013.

The ED-6064 with 1 Gb/s port speed can be field upgraded to provide 2 Gb/s port speed. The upgrade requires careful planning to minimize the impact to your operation. The upgrade does require that the existing CTP cards be replaced with CTP2 cards. While these cards may be upgraded concurrently, all the existing FPM port cards also need to be replaced. This can be managed with redundancy in your fabric design, or by scheduling a maintenance window.

**Important:** Both 1 Gb/s (FPM) and 2 Gb/s (UPM) port cards can be intermixed in the McDATA ED-6064 Enterprise Fibre Channel Director. However, to enable 2 Gb/s operations, all port cards must be UPM cards. If there are any 1 Gb/s FPM 4-port modules installed, the entire Director will operate at 1 Gb/s.

For in-depth information on the ED-6064, also see the McDATA ED-6064 manuals provided with the director:

- *McDATA ED-6064 Director Planning Manual,* 620-000106
- *McDATA ED-6064 Director User Manual,* 620-000107
- *McDATA ED-6064 Director Installation and Service Manual,* 620-000108
- *McDATA ED-6064 Director Product Manager User Manual,* 620-000120

### 4.1.1 ES-3232

The ES-3232 switch is designed to suit the needs of departmental SANs or SANs that do not need the high availability of directors. The ES-3232 switch comes with redundant hot swappable fans, power supplies, and HotCAT (hot code load and activation). It also serves as an edge switch connected to a director at the core of the SAN backbone. The ports are G_Ports that can configure as an F_Ports when connected to N port device or as E_Ports if connected to another switch.

The ES-3232 is shown in Figure 4-1.

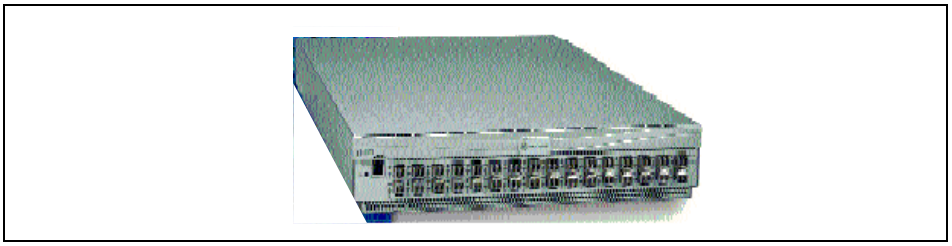

*Figure 4-1 The Sphereon ES-3232 Fabric Switch*

The switch can be used as a desk top unit, or it can reside in the Fabricenter Equipment Cabinet or a standard 19 inch rack. The switch can be managed by the EFC Manager, or via the SANpilot with limited functionality.

The switch offers high performance with 2 Gb/s non-blocking bandwidth for every port. The ES-3232 also offers the FlexPort Technology feature to scale from 16 to 32 ports non-disruptively.

Further information about the switches can be found in these McDATA publications:

- *McDATA Products in a SAN Environment Planning Manual,* 620-000124
- *McDATA ES-3232 Switch Product Manager User Manual,* 620-000137
- *McDATA ES-3232 Switch Installation and Service Manual,* 620-000142

## 4.1.2 ES-4500

The ES-4500 switch consists of the FlexPort Technology in 8-port increments up to a maximum of 24-ports. It is the only McDATA switch that directly supports Fibre Channel loop topology. The ES-4500 replaces the ES-1000 hub as it provides support for OpenTrunking, Hot Code Activation (HotCAT), SANtegrity and SANpilot. The various port configuration features on ES-4500 are GX_Port, FX_Port, G Port, E_Port, and F_Port, and it also has an embedded NL_Port.

The ES-4500 is shown in Figure 4-2.

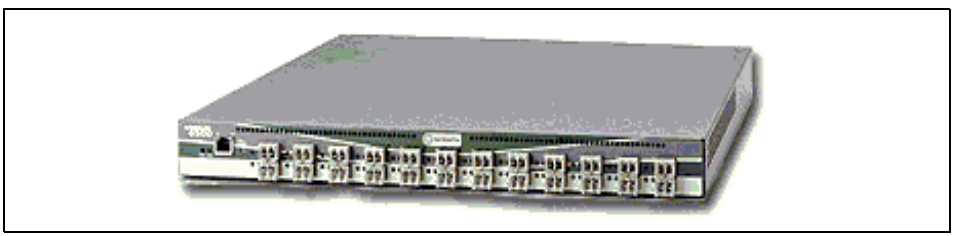

*Figure 4-2 The Sphereon ES-4500 FlexPort Switch*

The ES-4500 has shared BB_Credits. The first four ports (numbered 0 through 3) are pre-configured to a BB_Credit value of 12 to support extended data transmission distance of up to 20 km at 1 Gb/s, and up to 10 km at 2 Gb/s. The remaining ports (numbered 4 to 23) are pre-configured to a BB_Credit value of 5, and do not support extended distance operation.

For more information refer to the ES-4500 Service and Install Manual and Sphereon ES-4500 User Manual found at Web site:

http://www.mcdata.com/filecenter/template?page=docs.search

### 4.1.3 The Fabricenter cabinet

The Fabricenter Management configuration (Part number FC-512M) includes:

- EFC Manager software providing a fabric-wide view of the entire storage network allowing IT administrators to monitor and control all switched enterprise components from a single remote or local console.
- EFC Manager laptop server mounted on slide out tray
- External zip drive
- Modem for call-home support
- Ethernet hub, 24-port, 10/100 Base-T
- Power distribution package (24 single connections or 12 with dual independent power distribution)
- One Fabricenter Management product can manage up to three expansion cabinets of equipment.
- The Fabricenter Expansion configuration (Part number FC-512E) includes:
- Ethernet hub, 24-port, 10/100 Base-T
- Power distribution package (24 single connections or 12 with dual independent power distribution)

Figure 4-3 shows the Fabricenter.

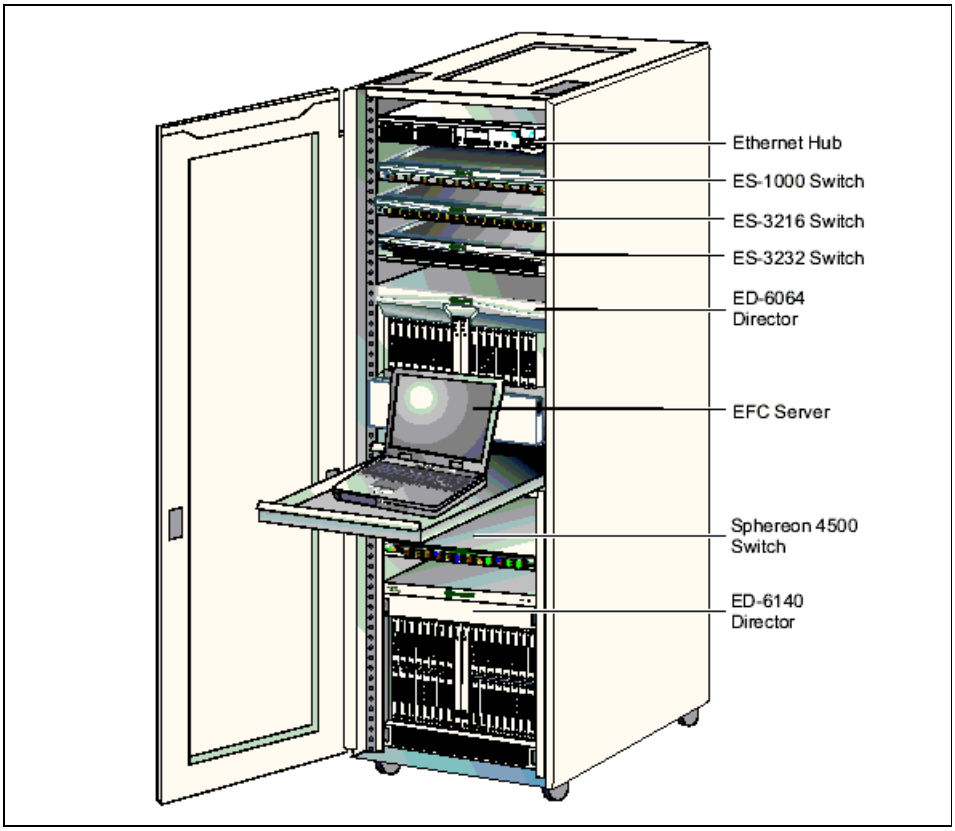


*Figure 4-3 The Fabricenter*

## 4.2 Setting up the network environment

Before we proceed with configuring the McDATA products (directors, switches and EFCM) for out-of-band management via Ethernet over TCP/IP interface, we first need to focus on designing the LAN architecture to maintain high availability, security, and optimal throughput.

**Tip:** We recommend connecting the primary ethernet interface on the EFC server to the corporate network to allow remote management of the Fabric via the EFCM, SANpilot and telnet applications. The secondary ethernet interface on the EFC server should be connected to the private network to protect from known vulnerabilities such as broadcast storms, viruses, and hacking. The broadcast storms are a common phenomenon, especially in ethernet over TCP/IP environments. Any malfunctioning device can bring the entire LAN to a halt by continuously transmitting broadcast packets, causing the Fabric devices to interrupt the IO activity and handle the TCP/IP broadcast packets. As a result the Fibre Channel Read/Write operations get timed out.

### 4.2.1 McDATA SAN on a dedicated TCP/IP ethernet LAN

For security and high availability reasons, we will connect the McDATA products to a private LAN. We use a second ethernet switch to establish a private ethernet connection from the local EFC server and the SAN. This is illustrated in Figure 4-4.

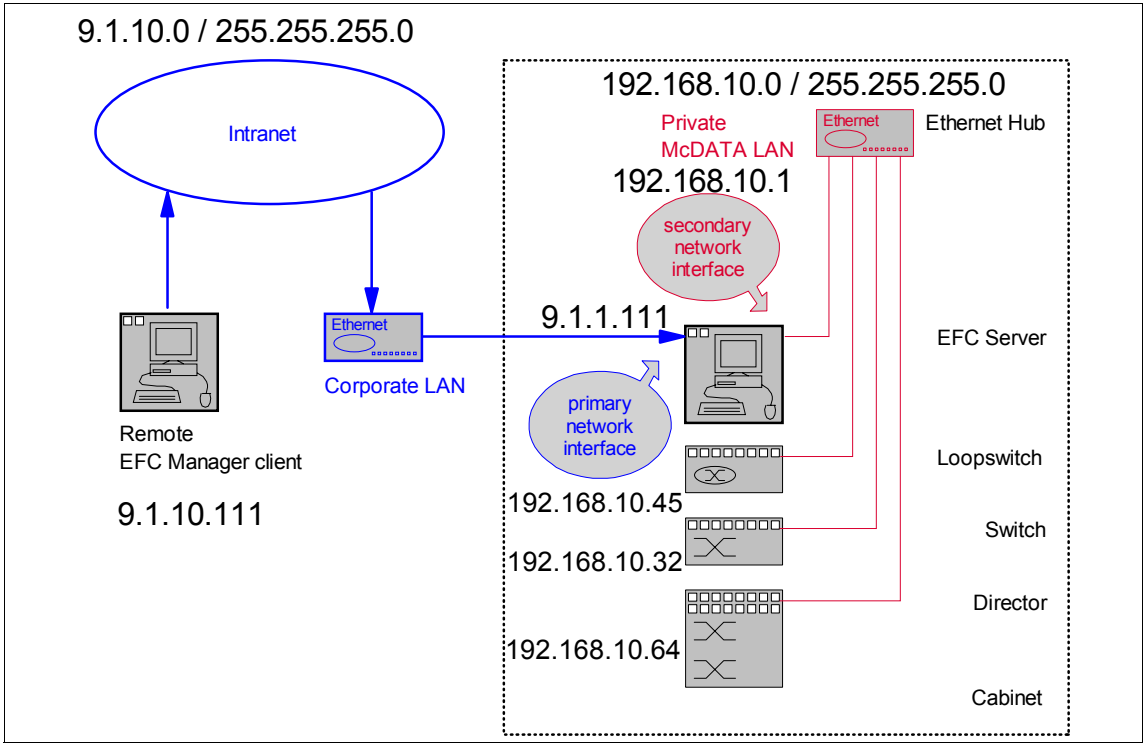


*Figure 4-4 Suggested McDATA network setup*

To simplify our implementation we assign the IP address range of 9.1.1.0/255.255.0.0 for the corporate LAN and use the 192.168.10.0 / 255.255.255.0 range for the private LAN. In the figure below we configure IP address 9.1.1.111 on the primary Ethernet interface for the EFC server and assign 192.168.10.1 on the secondary interface of the EFC server.

The arrows indicate the path from the remote EFC Manager client to the EFC Server. As illustrated the McDATA SAN is segregated from the corporate public network. We strongly recommended this LAN architecture to maintain high availability, manageability, fabric integrity and optimal performance.

The primary ethernet interface of the EFC Server connecting to the corporate LAN can be manually configured with a valid and unique IP address or it can be configured to obtain the IP address automatically from a DHCP server. The secondary Ethernet interface must be hard configured with a IP address since we do not use a DHCP server on the private LAN.

For more information on configuring the network environment, refer to the SAN Planning documentation found at Web site:

http://www.mcdata.com/knowcenter/techpubs/index.html

## 4.3 Product management

The McDATA products can be managed using the out-of-band and in-band product management interfaces.

The following out-of-band management access methods are currently available:

- Management through the EFC Manager and Product Manager applications.
- From the Internet using the SANpilot interface installed on the product. This interface supports configuration, statistics monitoring, and basic operation of the product, but does not offer all the capabilities of the corresponding EFC Product Manager application.
- From a Telnet session using the command line interface. Any platform that supports Telnet client software can be used.
- From a serial connection to the RS-232 port also known as the maintenance port using the null modem cable. The default baud speed configured on com1 port is 57600 with flow control set to none.
- Management using simple network management protocol (SNMP).

### 4.3.1 SANpilot: the Web based interface

With product firmware Version 1.2 (or later) installed, administrators or operators with a browser-capable PC and an Internet connection can monitor and manage the director or switch through the SANpilot interface. The interface provides a GUI similar to the Product Manager application, and supports product configuration, statistics monitoring, and basic operation.

The SANpilot interface does not replace nor offer the management capability of the EFC Manager and Product Manager applications. For example, the SANpilot interface does not support all product maintenance functions. In addition, the SANpilot interface manages only a single product but it has a hyperlink access to other switches in a fabric.

Some of the major limitations with SANpilot are:

- It does not provide the zoning and firmware libraries to store more than one zone set and firmware releases.
- The option to upload the data collection zip file under the maintenance menu is not available. This is a critical tool used to perform problem determination and configuration analysis.

Figure 4-5 shows the SANpilot interface initial panel.

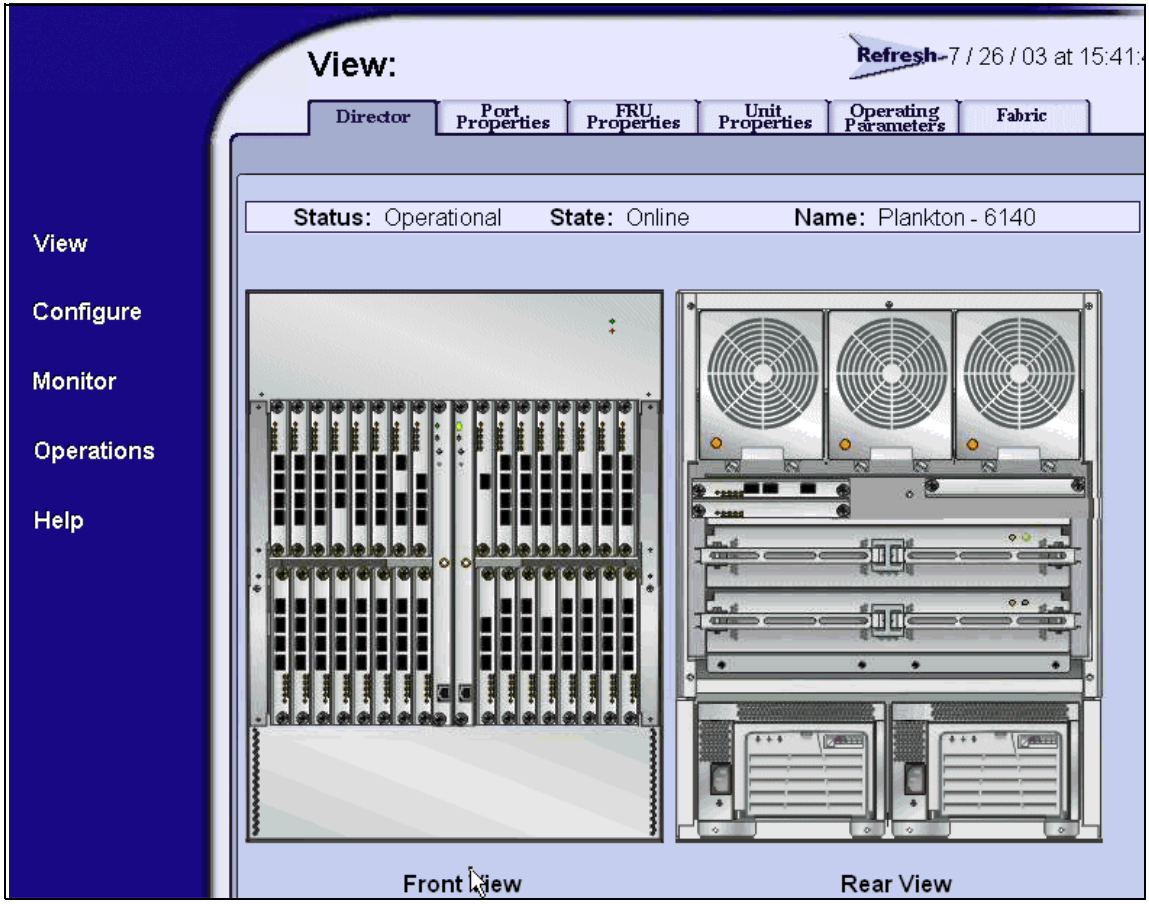


*Figure 4-5 ED-6140 hardware view from the SANPilot Web interface*

**Important:** With E/OS 05.01.00, SANpilot now provides support for the configuration of OpenTrunking and Port Binding. The zoning interface has also been enhanced to allow the user to easily add members to a zone that are not directly attached to the locally managed switch. The user is presented with a list of all devices logged in to the fabric and can conveniently select the additional members for addition to a zone.

For more details about the various configuration and management features with SANpilot, refer to the SAN planning documentation found at this Web site:

`http://www.mcdata.com/knowcenter/techpubs/index.html`

### 4.3.2 Introducing the EFC Manager

The EFC Manager is a centralized management tool for the McDATA SAN. It is a Java based GUI that provides a graphical view of all the fabric devices and functions as a central control point for all common management and monitoring tasks for up to 48 directors or switches.

To manage the McDATA directors and switches running firmware release 5.01, the EFCM must be running version 07.01.009. The devices are configured using specific Product Manager views, and the Fabric Manager component is used for fabric wide configuration, such as zoning and fabric operating parameters.

**Note:** The EFC Server and EFC Manager application provide a GUI to monitor and manage McDATA products, and are a dedicated hardware and software solution that should not be used for other tasks. McDATA tests the EFC Manager application installed on the EFC Server, but does not test the compatibility of any third-party software. Modifications to the EFC Server hardware or installation of additional software, including patches or service packs may interfere with normal operation.

For detailed information about the EFC Manager and how to use it, refer to the *McDATA Enterprise Fabric Connectivity Manager User Manual,* 620-005001.

The EFC Manager may be used locally from the EFC Server laptop, but it is also possible to use it from a remote workstation. To do so we need to download the code from the EFC server, and install it on our workstation. This workstation can be running Windows, AIX, Solaris, HP-UX or LINUX. In our case we will be using a PC running Windows NT.

## Why EFC manager is required for enterprise SANs

One of the major factors for enterprise businesses to use the EFC server to manage their SAN is the capability to backup and restore device and fabric configuration for all the products managed by the local or remote EFC server. It enables the enterprise SAN to become disaster proof.

The EFC Server uses the product manager application to backup and restore the configuration data stored in the nonvolatile random-access memory (NV-RAM) on a director or switch CTP card on the EFC Manager data directory. The other feature backs up (to the Zip drive) or restores the entire EFC Manager data directory.

The NV-RAM data includes:

- Product identification data, port configuration data, and link incident (LIN) alerts
- Operating parameters such as flow control values, preferred domain ID, Active zoning configuration and SNMP configuration

The EFC Manager data directory includes:

- All EFC Manager configuration data such as product definitions, user definitions session options and remote event notifications
- All log files, such as EFC Manager logs and individual director or switch Product Manager logs
- Zoning library includes all configured zone sets and zone definitions
- Firmware library
- Call-home settings such as phone numbers and dialing options
- Configuration data for each managed product, stored on the EFC Server and in NV-RAM on each director or switch

### 4.3.3 Accessing the EFC Manager client installation software

The software is downloaded from the EFC server. The download and installation of the EFC Manager is done using a Web and Java based installation procedure. All we need is a Web browser and we will use Microsoft Internet Explorer. In the Address (URL) field of the Explorer we point to the IP address of the EFC server to access the initial page.

This takes us to the start page for the remote EFC Manager client installation, as shown in Figure 4-6 and Figure 4-7, where we can choose the correct client for the operating system of our remote workstation.

From here, we can also download the SNMP MIB files later on, if required.

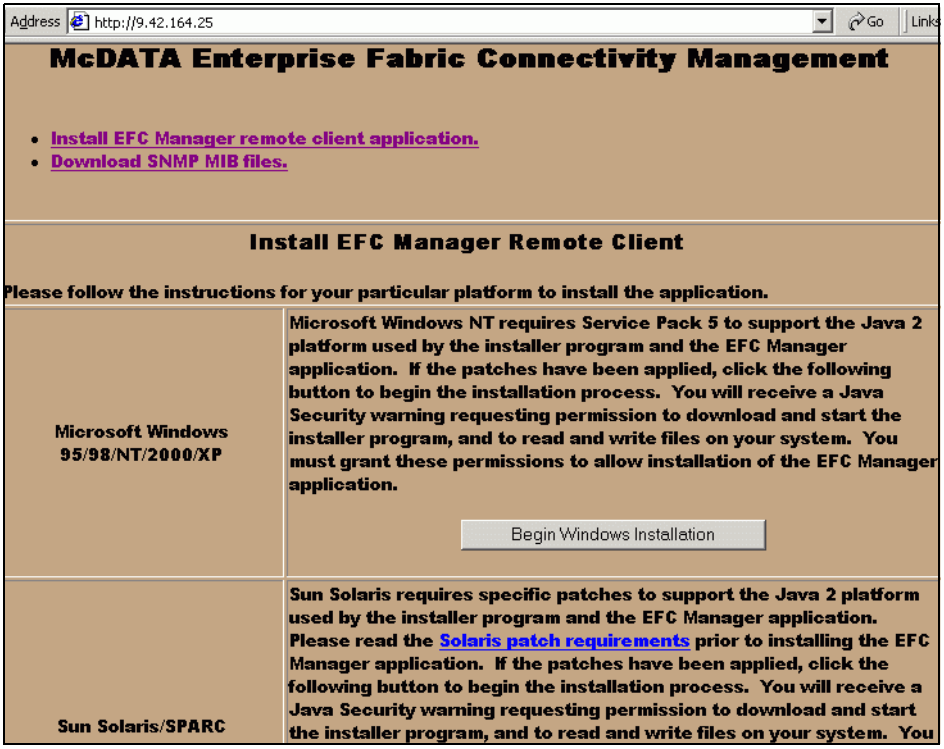


*Figure 4-6 Start page for remote EFC Manager client installation*

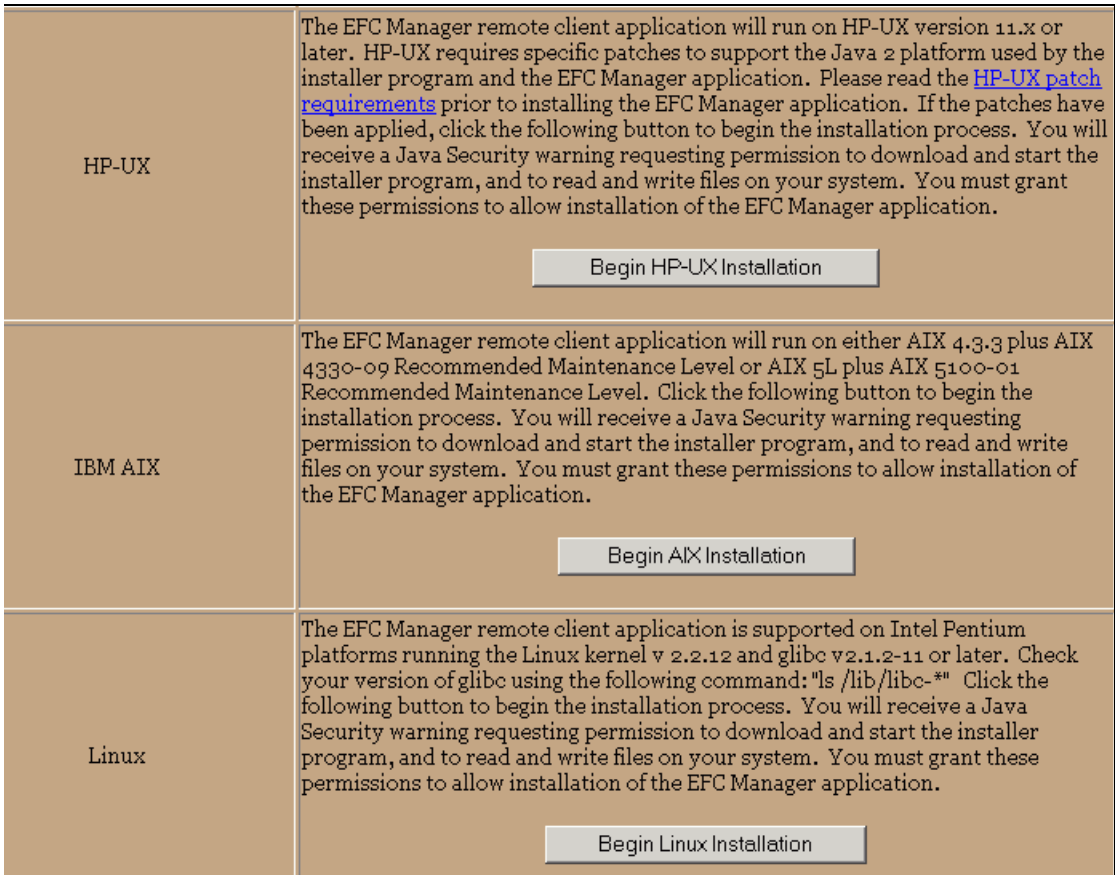


*Figure 4-7 Start page for remote EFC Manager client installation continued*

## 4.3.4 Downloading and installing the EFC Manager client

We will be installing the EFC Manager client software on a Microsoft Windows system, so we select that option. After doing so, we are taken to the Web page from which we can begin the download and installation procedure. We are prompted to grant additional privileges to the Java based installation software. First, we have to grant the right to start programs, and then grant the right to read, modify, or delete files.

This is shown in Figure 4-8.

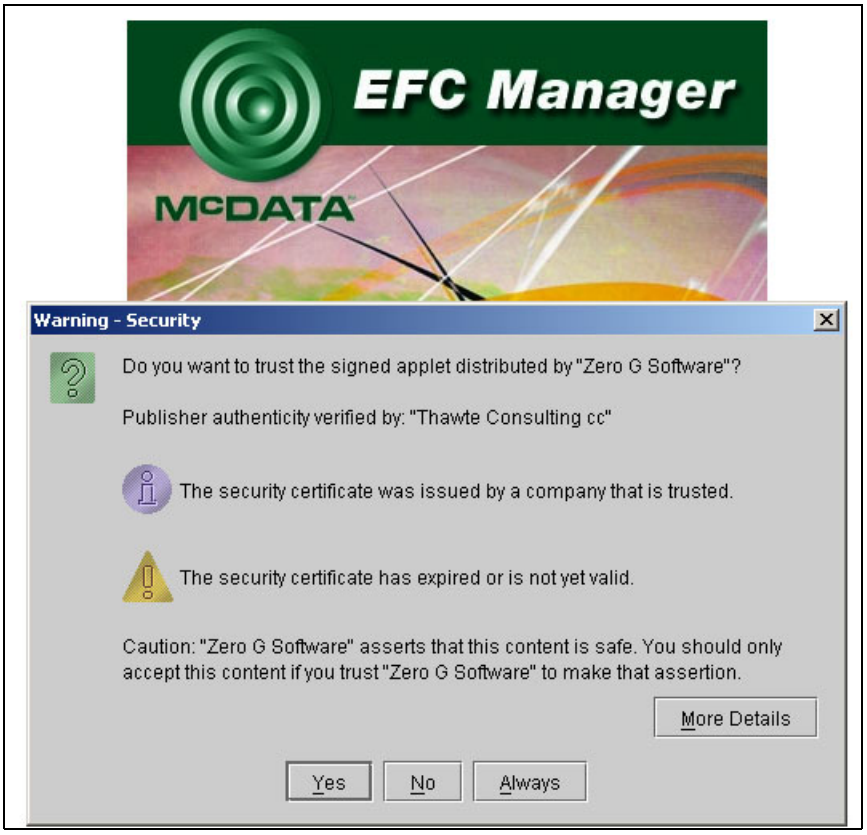


*Figure 4-8 Security prompt for Java application*

This warning message appears because, for security reasons, a Java applet is usually not allowed to perform the tasks mentioned. You should only allow Java programs from trusted sources to perform such actions.

If you do not grant the additional privileges, you can download the installer to your local workstation by scrolling lower in the browser window, after Denying the security warning, and execute it locally.

After granting the rights, the button to select the **Start Installer for Windows...** appears, and we are able to start the installation process, as shown in Figure 4-9.

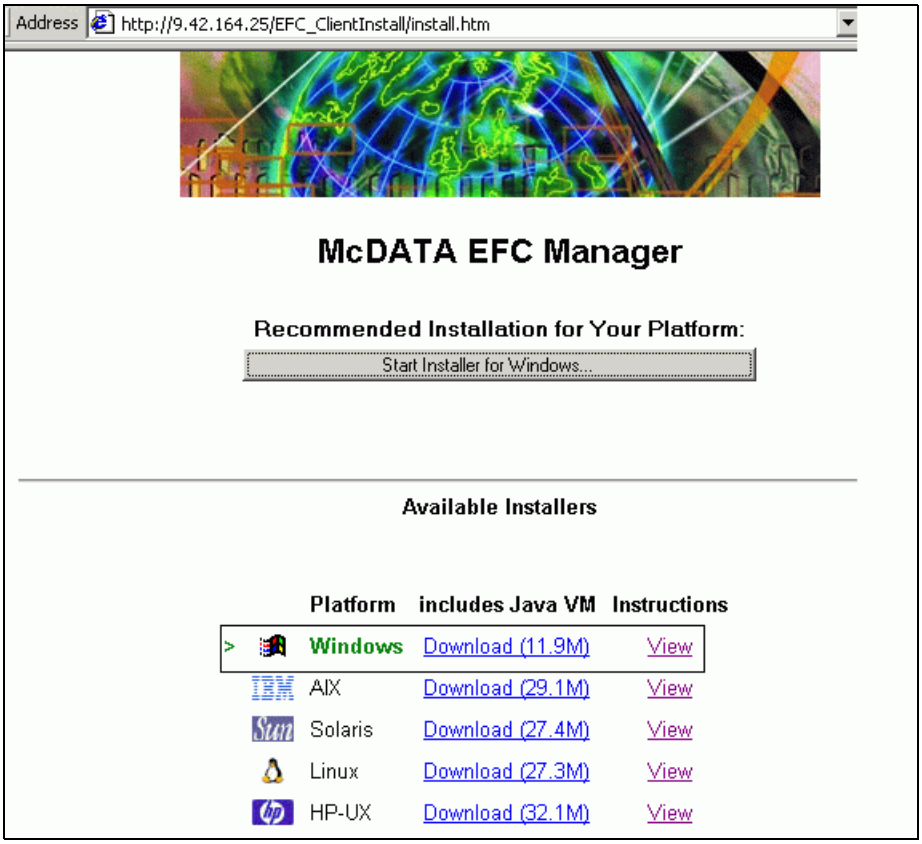


*Figure 4-9 Start download prompt*

After clicking the button to start the installation process, the software starts downloading to our local machine and begins the installation. The same installation would begin after executing the installer, **`mcdataClientInstall.exe`**, which could be manually downloaded from a list of available installers. The instructions are found by scrolling down on the window shown in Figure 4-10.

We now follow the prompts to install the EFC Manager client. After confirming the License agreement, we get information about which version we are going to install, as shown in Figure 4-10.

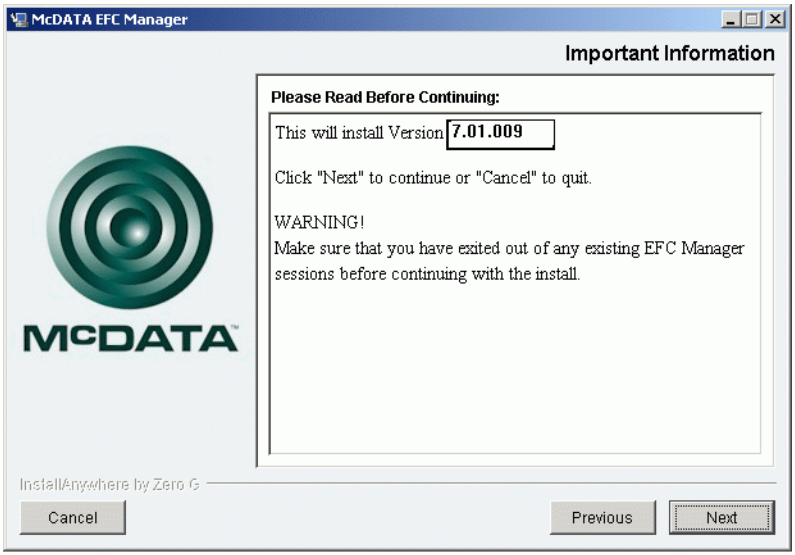


*Figure 4-10 EFC Manager client installation*

The next steps will ask us where we want to save the program files, if we want to remove a previous installation, and where we would like the program icons to be saved. Before it begins to install, the pre-installation summary menu provides one last chance to verify and confirm the installation path. If all is OK, then select **Install** as shown in Figure 4-11.

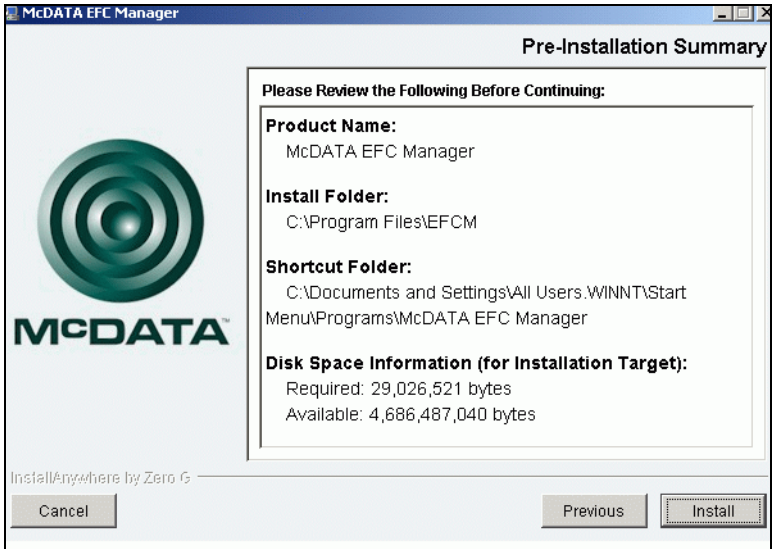


*Figure 4-11 The pre-Installation Summary menu*

Once the installation is complete, we are asked if we would like to start the EFCM client.

**Important:** If you plan on accessing the EFC Server across the firewall, some manual editing of the EFC Server configuration must be made.

### 4.3.5 Configuring EFCM access through a firewall

EFCM clients use random TCP ports to connect to EFC Servers. This allows a client machine to run multiple instances of the EFCM client to the same or different EFC Servers. Because of the use of random TCP ports, pass-through ports cannot be configured on network firewalls.

To get around this, manual configuration is required on the server, on the client, and on the firewall itself as follows:

1. On the EFCM Server, edit the file named:

   *C:\Program Files\EFCM\Config.properties*

   Add the following line to the file: **`RmiExportPort = 1098`**

   This will hard-code the server to listen on TCP port 1098 instead of on random ports. This value can be any decimal TCP port desired. In addition to this port, the server also listens for connections on the well-known RMI port 1099, which cannot be changed.

2. On the EFCM client, create a new file in the EFCM installation directory (default on Windows platforms is "`C:\Program Files\EFCM`", the default on UNIX platforms is `EFCM` in your home directory) named `Config.properties` and add the following line to the file: **`RmiExportPort = 1098`**

   This will hard-code the client to listen for connections back from the server on port 1098 instead of on random ports. This value can be any decimal TCP port desired, and it does not need to match the value used on the server side.

3. On the network firewall between the EFCM client and the EFCM server, configure the firewall to allow the configured ports through. EFCM also uses FTP between the client and the server, so the firewall administrator must allow the well-known FTP port 21 through the firewall as well.

### 4.3.6 Configuring the IP address for out-of-band management

The default IP address will be changed to allow access from local and remote hosts using SANpilot, EFC Manager and remote client in our environment. We will use the hyper terminal application to access the serial port using the null modem cable. The COM1 port properties are shown in Figure 4-12.

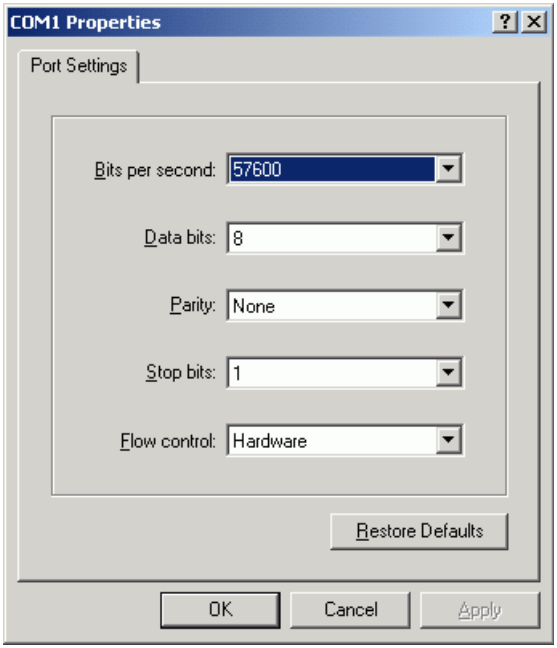


*Figure 4-12 COM1 properties*

Once you login to the switch, run the **ipconfig** command to verify the current IP configuration and configure the new IP address as shown in Figure 4-13.

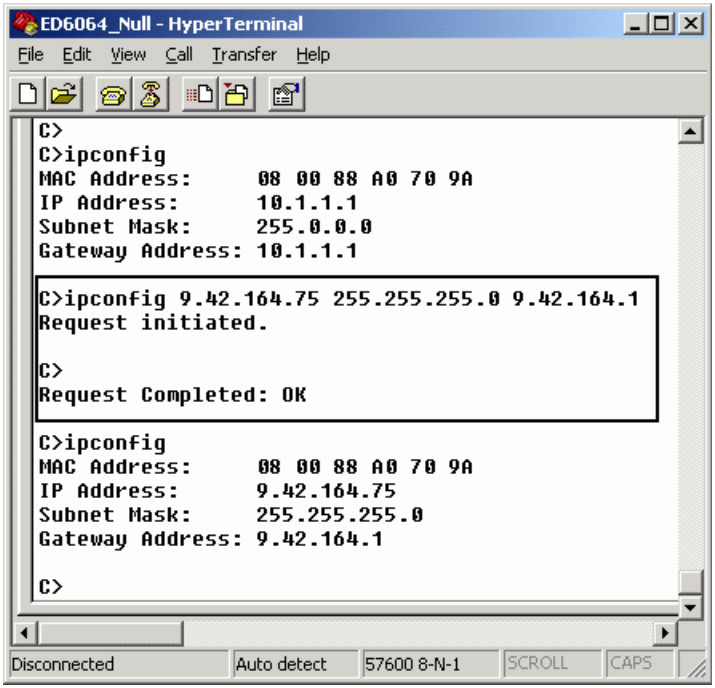


*Figure 4-13 IP address configuration procedure*

We can now manage the SAN using the EFC Manager, EFC Client, SANpilot and from a Telnet session.

## 4.4 Managing the environment using the EFC Manager

In our environment, we have an ED-6064 and ES-4500 building a core-to-edge topology. The EFC server is operational and has access to the ED-6064 and ES-4500 switch.

The EFC Manager is used for fabric specific administration of the McDATA SAN. The EFC Manager also serves as the entry point to the applications for managing the McDATA devices as well as the fabric. The application used for these tasks are the Product Manager and Fabric Manager.

In the following topics we will access the EFC Manager and perform some of the administration tasks that have to be accomplished using the EFC Manager, before we move on to the management of the devices and the fabric.

## 4.4.1 Logging in to the EFC Manager

The following administration and configuration steps can be done locally from the EFC Server or remotely using the EFCM client

### Logging in to the EFC Manager on the EFC Server

To start the EFC Manager, we logon with the same user ID and password we used when logging on to Windows 2000, `Administrator` and `password`. We are working locally on the EFC server and therefore we specify in the EFC Host Server entry field `localhost`, which is shown in Figure 4-14.

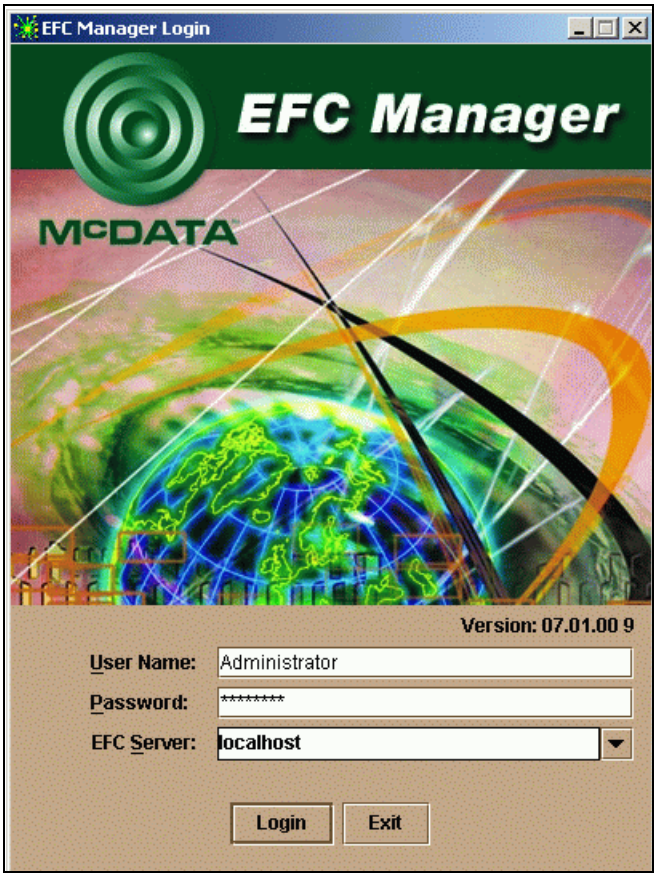


*Figure 4-14 Logging in to the EFC Manager on the EFC Server*

## Remote login to the EFC Manager

After finishing the installation of our client workstation, there will be a shortcut to the EFC Manager on the desktop as shown in Figure 4-15.


*Figure 4-15 EFC Manager workstation icon*

By double-clicking on the icon, we get to the EFC Manager login window. We now login to the EFC Manager using the default username and password and the IP address of the EFC server to login, as shown in Figure 4-16.

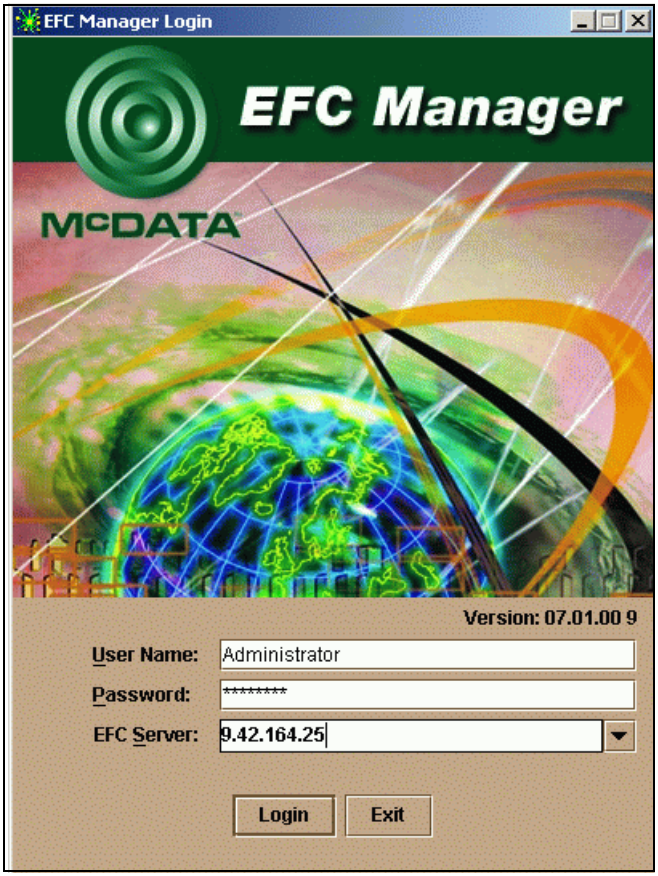


*Figure 4-16 Remote login in to the EFC Manager*

## 4.4.2 Administering the SAN using the EFC Manager

After a successful login, we can move on to administer the McDATA SAN products. Following are examples of the options that we can configure using the EFC Manager.

After logging on to the EFC Manager, it opens with the Product View shown in Figure 4-17.

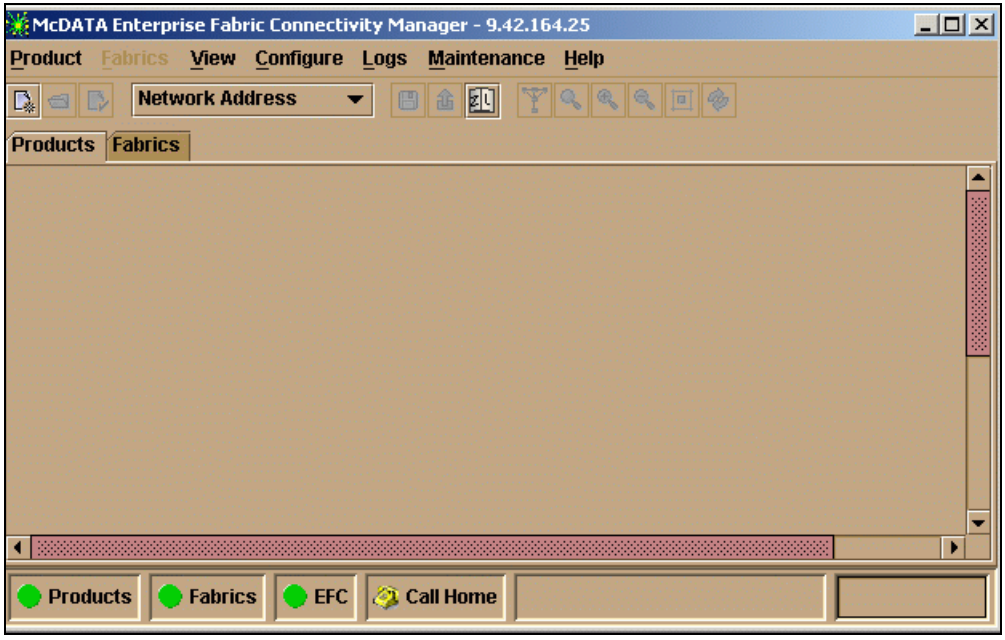


*Figure 4-17 EFC Manager, Product View, no switches defined*

At the top of the window, there is a pull-down menu bar that allows us to perform various configuration and monitoring tasks. The main window is empty because there are no devices configured on the EFC Manager yet. The same applies to the Fabric View of the EFC Manager. We can switch to it by clicking the **Fabrics** tab as shown in Figure 4-17.

At the bottom of the window are icons which indicate the health of individual products, the fabric, and the EFC Manager. These icons provide a very quick method of alerting us to any potential problems in the McDATA SAN. As we can see, all three icons are green circles, which means that we have no major issues.

### 4.4.3 Defining users on the EFC Manager

First, we want to define users on the EFC Manager because we do not want the Administrator user ID to be used remotely, so we will create a new user and use that for remote access.

We can define up to 16 users for the EFC Manager but only a maximum of four can log on concurrently. With the user of the EFC Manager running locally on the EFC server there can be five sessions open concurrently.

From the pull-down menu, we select **Configure —> Users...** as shown in Figure 4-18.


*Figure 4-18 EFC Manager, Product View*

We are presented with a list of defined users and the options to add users, modify existing users, view the rights of a user, and delete users. We will add another user by clicking the **New** button and then specifying the name, password, and description of the new user. Also, this window is used to specify the rights that the new user should have. This is shown in Figure 4-19.

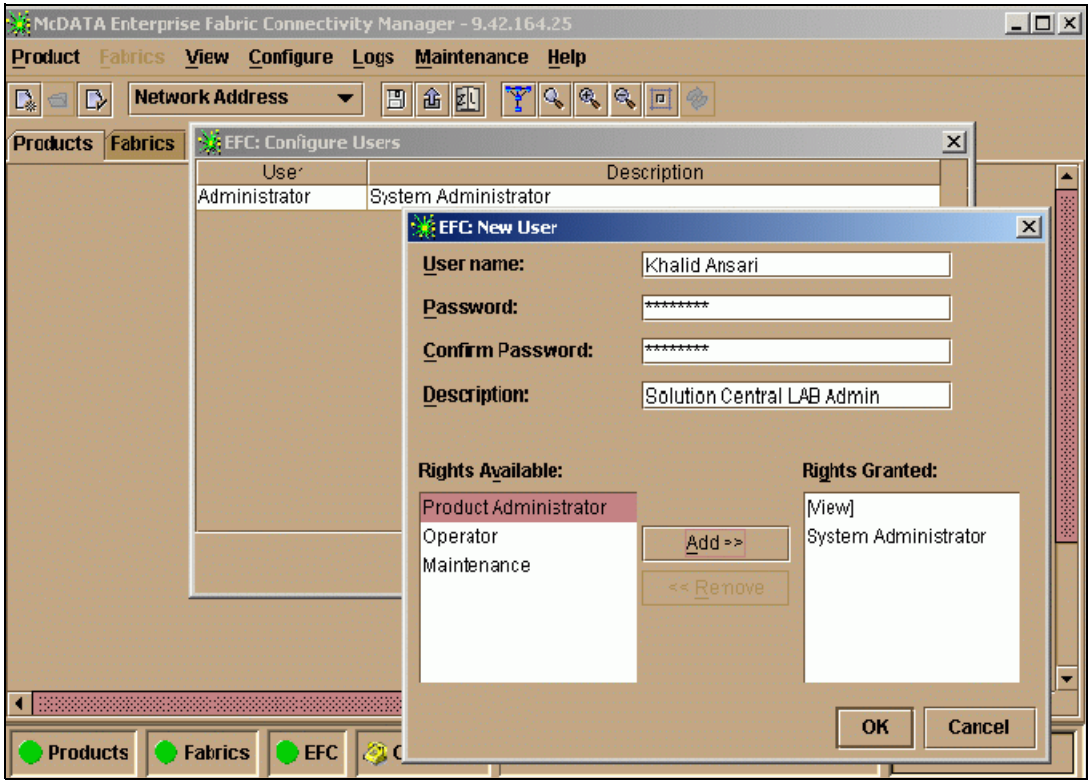


*Figure 4-19 EFC Manager, Configure Users, New User*

To assign rights to the user, we click one of the **Rights Available** choices and click the **Add** button. The rights are:

- System Administrator
- Product Administrator
- Operator
- Maintenance
- View

The *System Administrator* right grants access to every control and configuration task that needs to be performed from within the EFC Manager and can be viewed as the highest level of authority. It only has “view” rights while operating in a product manager application. Here we need the *Product Administrator* right to perform changes.

All new users initially have view rights and this cannot be removed. For a table of user rights of product manager functions, refer to the *McDATA Enterprise Fabric Connectivity Manager User Manual*, P/N 620-005001.

To change the settings for a user, for instance, to change the password, we go to **Configure —> Users**. Using the **Modify** button, we are presented with a window similar to the New window where we can change our password and the user rights. This is shown in Figure 4-20.

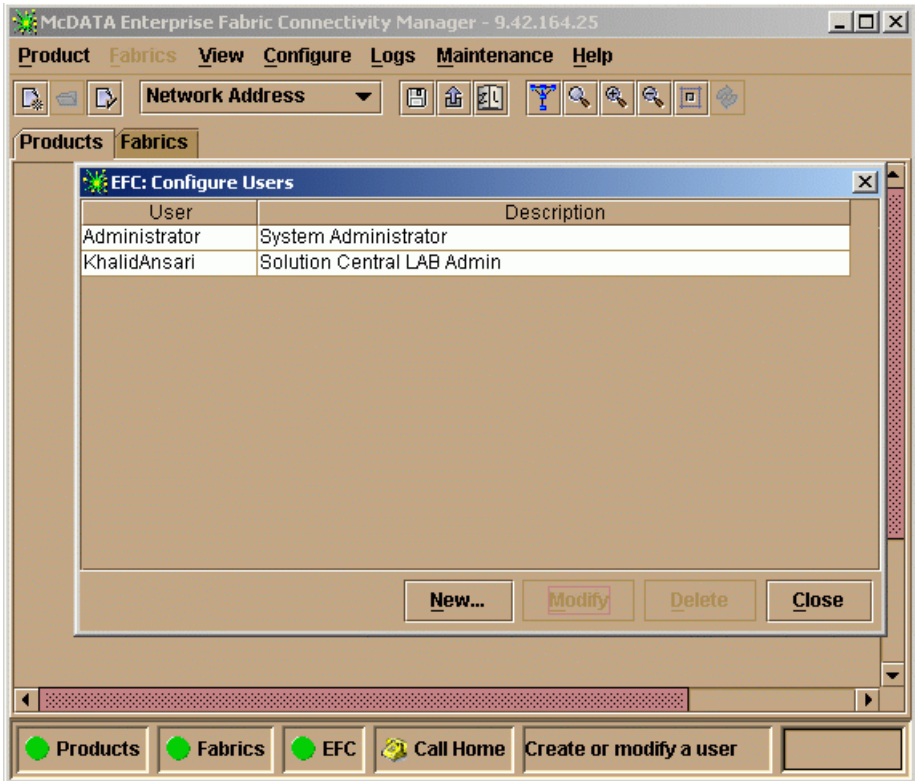


*Figure 4-20 EFC Manager, Configuring Users, Modify User*

Once the new user has been defined, we can login to the EFC server with the newly created user ID and password.

### 4.4.4 Identifying devices to the EFC Manager

We have to identify the devices that are going to be configured and monitored through this EFC Manager. Those devices, then, cannot be managed by another active EFC Server.

After logging on to the EFC server the Product View opens with no devices installed as shown in Figure 4-21.

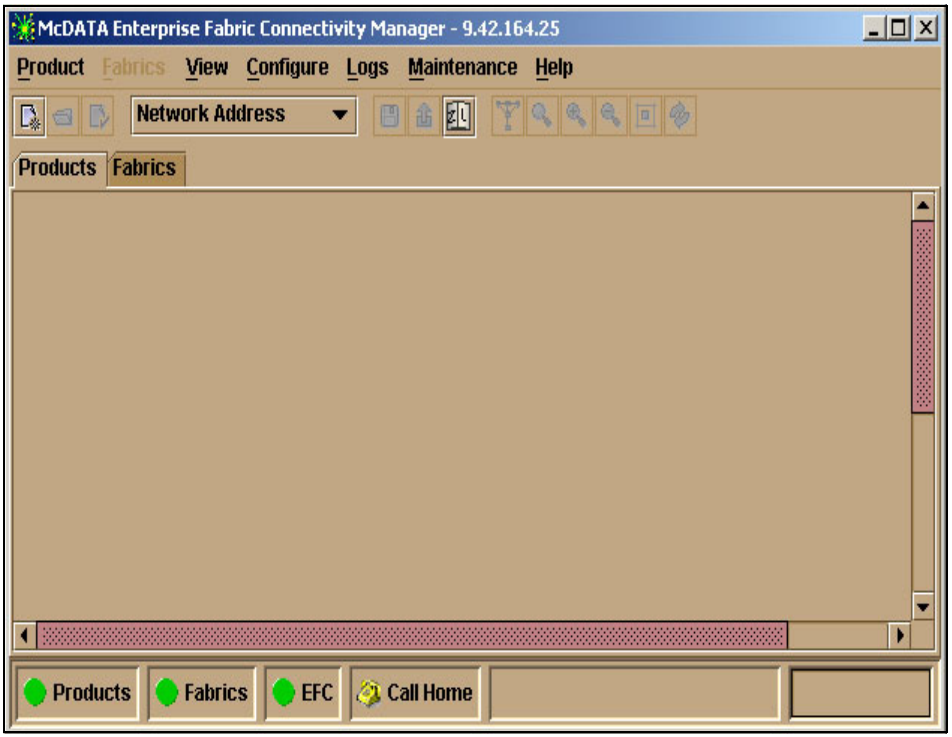


*Figure 4-21 EFC Manager, Product View, no switches defined*

To identify the ED-6064 to the EFC Manager we need to tell the EFC Manager the IP address of the ED-6064. This is accomplished by selecting **Product -> New...**, as shown in Figure 4-22.

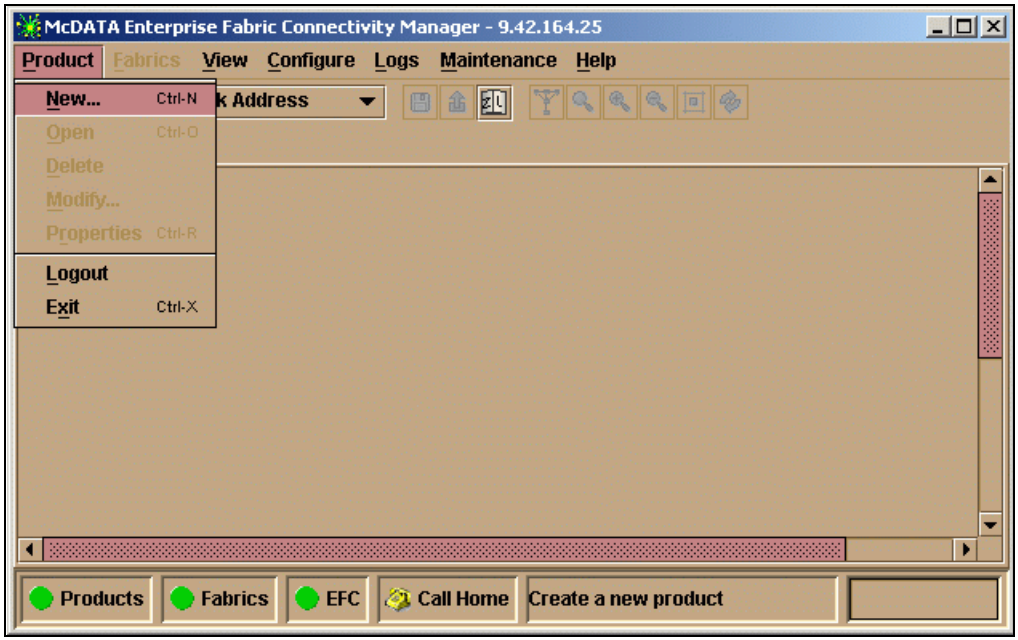


*Figure 4-22 EFC Manager, New Product...*

Selecting this takes us to the New Product entry field, where we have to fill in the IP address of the director that we want to add. This is shown in Figure 4-23.

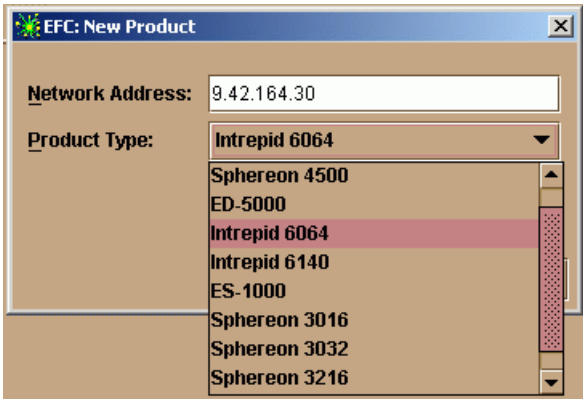


*Figure 4-23 Defining new ED-6064 with its IP address*

The ED-6064 was correctly installed in the network previously, and now the EFC server can communicate with it. Therefore, the newly added director appears as an icon in the Product View window, as shown in Figure 4-24.

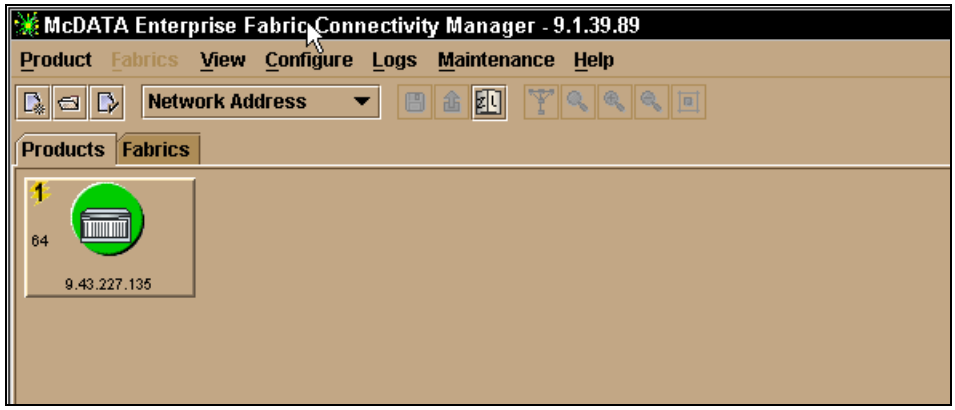


*Figure 4-24 EFC Manager, Product View, new ED-6064 icon*

We can also see the IP address of the ED-6064. A green circle indicates that the switch is up and running with no known problems. We show in Figure 4-25 what each of the different areas of an individual product icon mean.

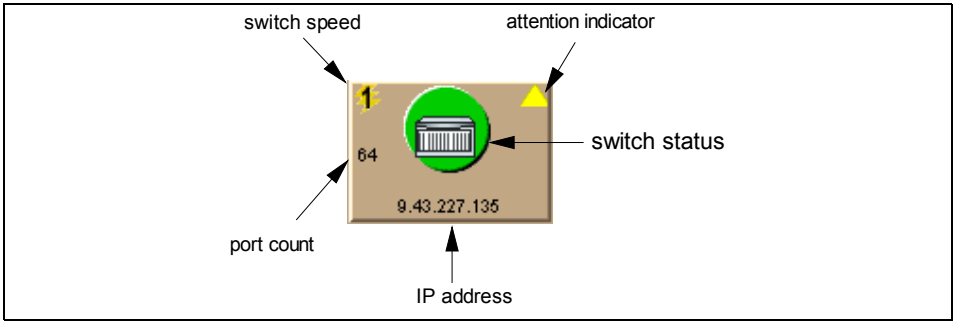


*Figure 4-25 Product View, icon description*

## 4.4.5 Assigning nicknames to World Wide Port Names

As with IP addresses and the DNS, managing the SAN can be made easier by defining nicknames for WWNs. The names can be the DNS host name in the case of only one adapter in one host. If there is more than one adapter in one device, we recommend that the nickname should consist of the hostname and some extension to distinguish between adapters. This will help later when we have to identify devices, for instance while configuring zoning.

This is especially true because the ESS has Emulex adapters built in. We have other Emulex adapters in our SAN so it would be useful to distinguish between the ones in workstations and the ones found in the ESS. For our ESS we chose, as an example, ESS Bay1-Host2, and Bay3-Host2.

We use the information to include specific adapters in a zone. However, this is not our only means of restricting access to volumes on the ESS.

On the ESS, every ESS logical volume is normally accessible through every FC port. However, we can restrict the host FC ports that are configured in the ESS to only see ESS logical volumes through specific ESS FC ports. This means that those ports can only access the volumes assigned to it through those ESS FC ports.

Including ESS FC adapters within specific zones might be useful when we want to influence the bandwidth that a specific group of host FC ports (members of a zone) can get, and through which bay we want the data to go.

In our example, we configure nicknames for two ESS host ports, two FastT600 controllers, IBM 3584 Tape device and the AIX host. This is done by selecting **Configure —> Nicknames...**, which opens the window without any nicknames configured. We use the **New...** button to add some nicknames, as shown in Figure 4-26.

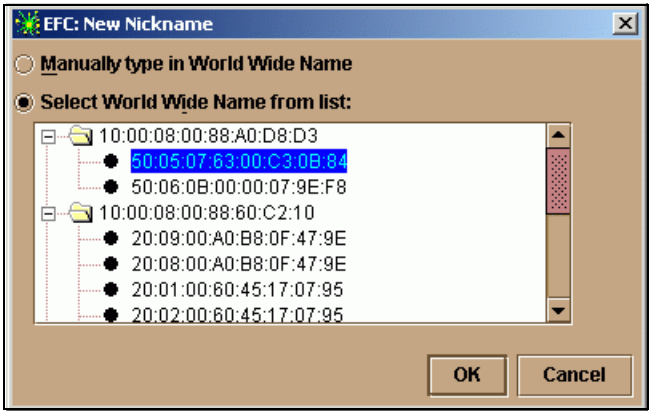


*Figure 4-26 EFC Manager, Configure Nicknames, Add Nickname*

After assigning nicknames to the identified WWPNs, the window looks like that shown in Figure 4-27.

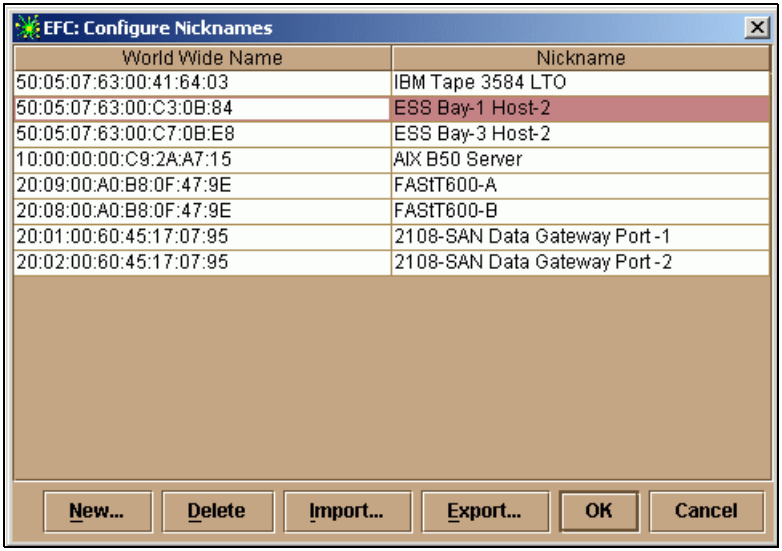


*Figure 4-27 EFC Manager, Configure Nicknames, nicknames assigned*

Doing this is similar to using a **`hosts`** file on an IP host.

In our simple case it might be tempting to work with the WWN and to skip or ignore the task of assigning nicknames. However, as more devices attach, maintaining the fabric with names is more convenient and will be easier than figuring out which WWN belongs to which machine at a later date.

After assigning nicknames, the Node List View of the Product Manager shows the names of those that are currently attached. With a growing SAN, it becomes more and more important to be able to distinguish between the node ports.

## 4.5 Managing the devices using the Product Manager

Now we are in a position to configure the devices. The Product Manager provides different options for every device type to be configured which reflect the specific hardware and the configuration options that the devices provide. For instance, the ES-3232 and ES-4500 switches do not feature all of the high availability features available with ED-6140 and ED-6064, and the ES-4500 is the only switch from the McDATA product line currently marketed that supports arbitrated loop topology. We only configure options that are necessary for our SAN, for instance, the operating parameters. We do not cover administration tasks such as configuring SNMP.

Also, the Product Manager presents different front and rear views of the devices in the Hardware View. These are interactive views which display the status of monitored units, for instance, if they have failed. Additionally, clicking a unit opens a window with more information about the unit.

## 4.5.1 Managing the ES-3232

The Hardware View of the ES-3232 looks like that shown in Figure 4-28.

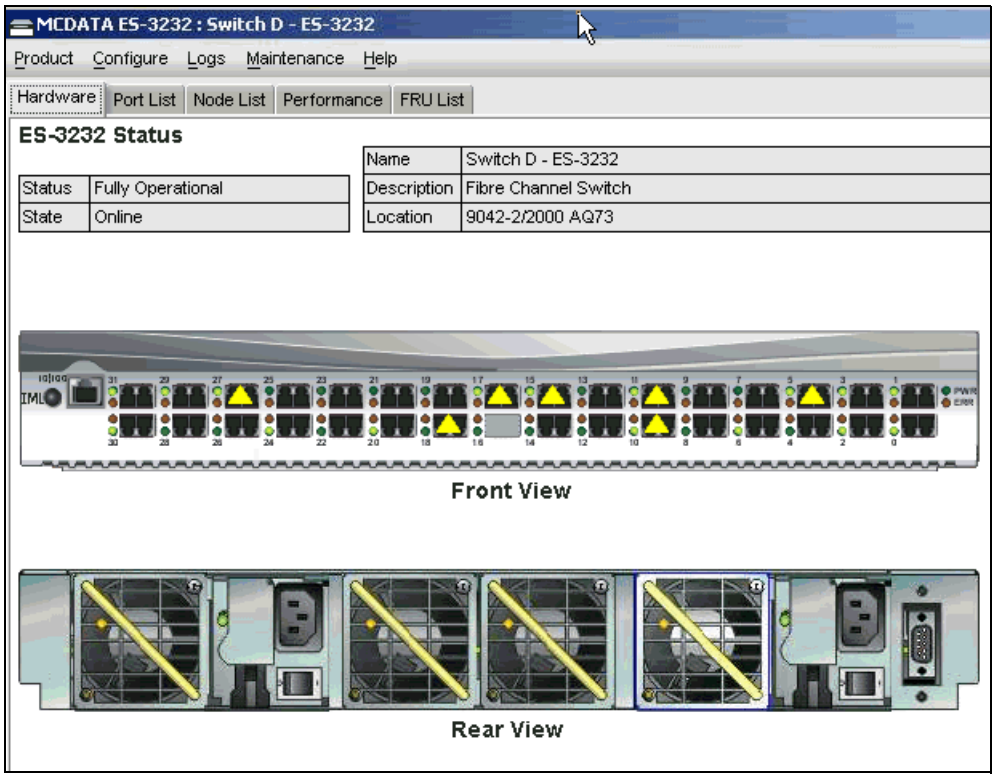


*Figure 4-28 Product Manager ES-3232 Hardware View*

The other options available to configure, for example the operating mode, the operating parameters and the ports, are the same as with the ED-6064, As that is the case we will use the ED-6064 to manage our McDATA SAN but it is safe to assume that the operations are similar for the other models.

## 4.5.2 Configuring the ED-6064 using EFC Product Manager

We double-click the **ED-6064** icon for the director to open the Product Manager.

This opens the Product Manager in its own window with the Hardware View, as shown in Figure 4-29. This illustrates the different types of monitored interactive parts as they show up when we move the mouse cursor over them.

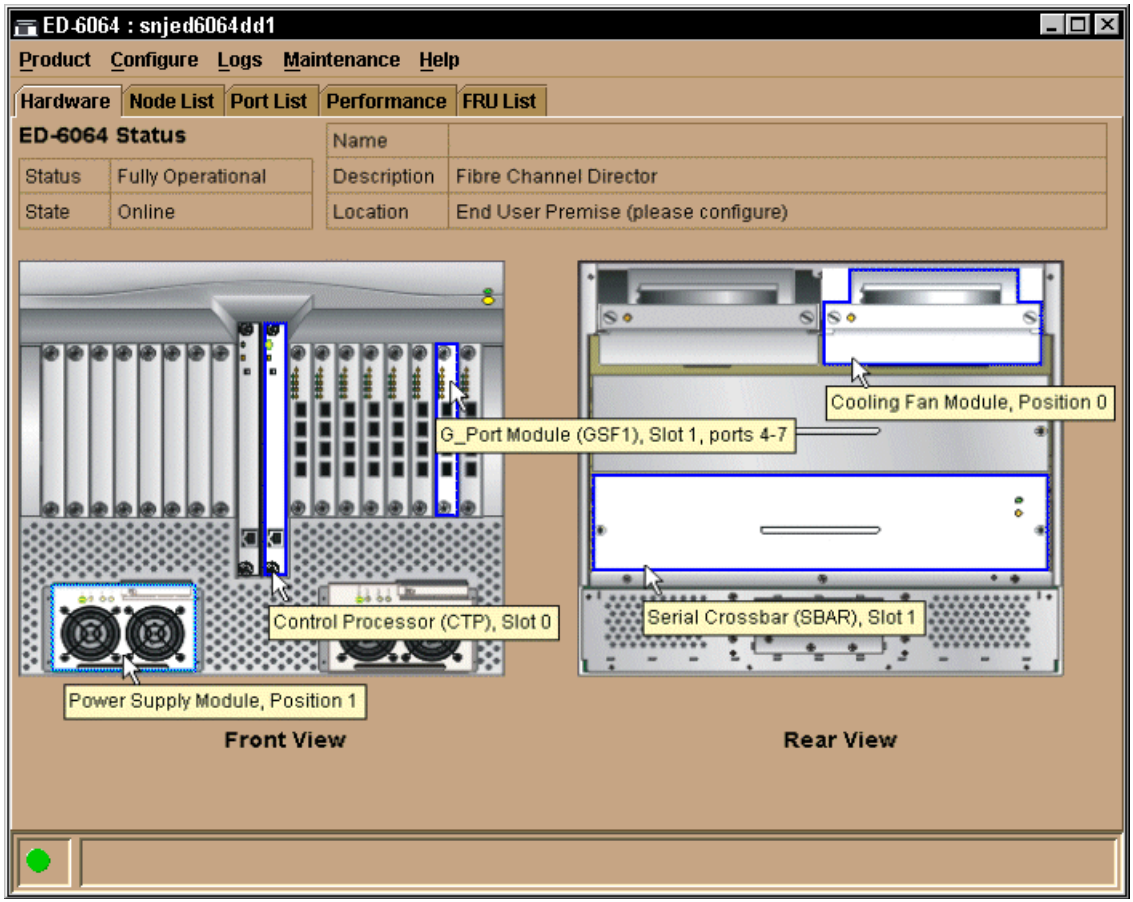


*Figure 4-29 Product Manager ED-6064: Hardware View*

At the top of the window we see the operational status of the switch itself. The switch is fully operational and online. The Front View and the Rear View of the unit show the installed components. The graphics representing the components of the switch are interactive, which means by selecting them, we are able to view more information, or to perform configuration or maintenance actions. They are also monitored so that we have a graphical representation of the failed part in the front and rear view.

## Using the interactive port card view

As we can see, there are 8 port cards installed in our switch, which makes a total of 32 ports. Double-clicking one of the port cards opens the Port Card View as shown in Figure 4-30.

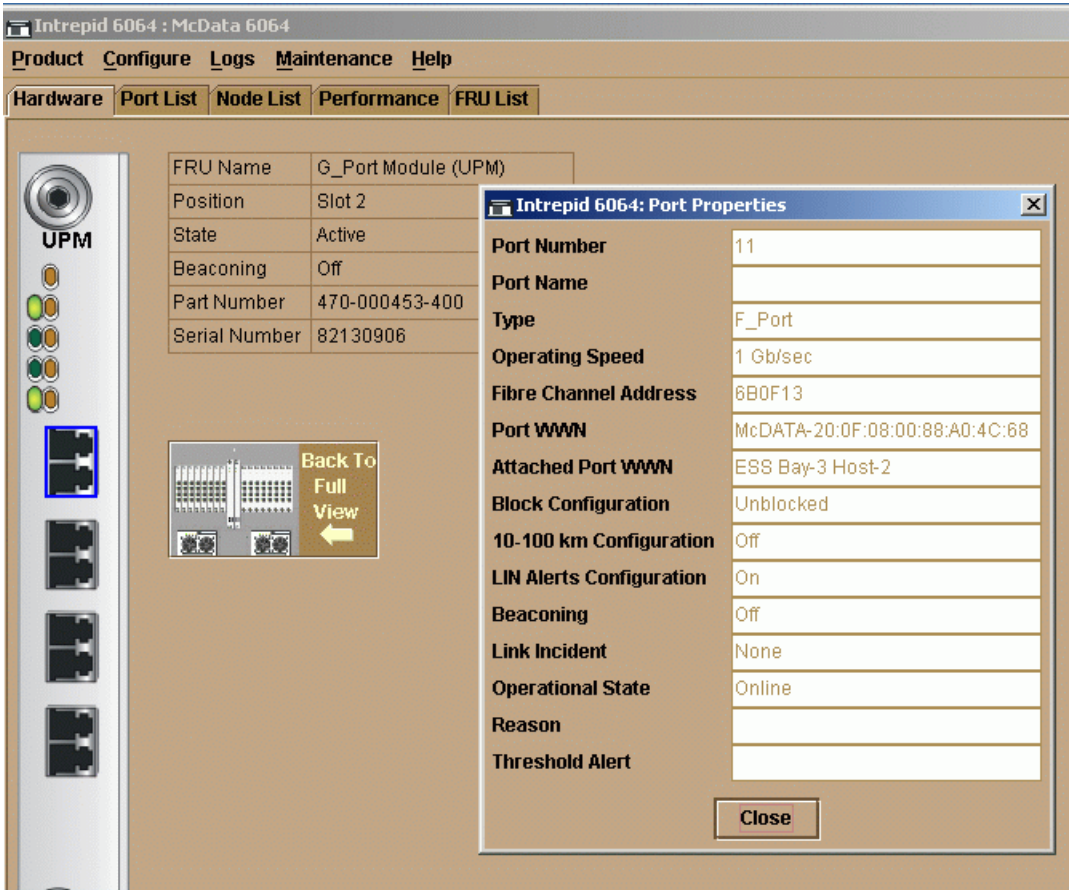


*Figure 4-30 Product Manager ED-6064: Port card view and properties*

By moving the mouse over a specific port, we see its port number. By double-clicking it, we get detailed port information and this is also shown in Figure 4-30. The port we selected is Blocked and shows as a G_Port. Also, we see the parameters that are currently defined for this port and that the port is online. Right-clicking a port gives us a context menu as shown in Figure 4-31.

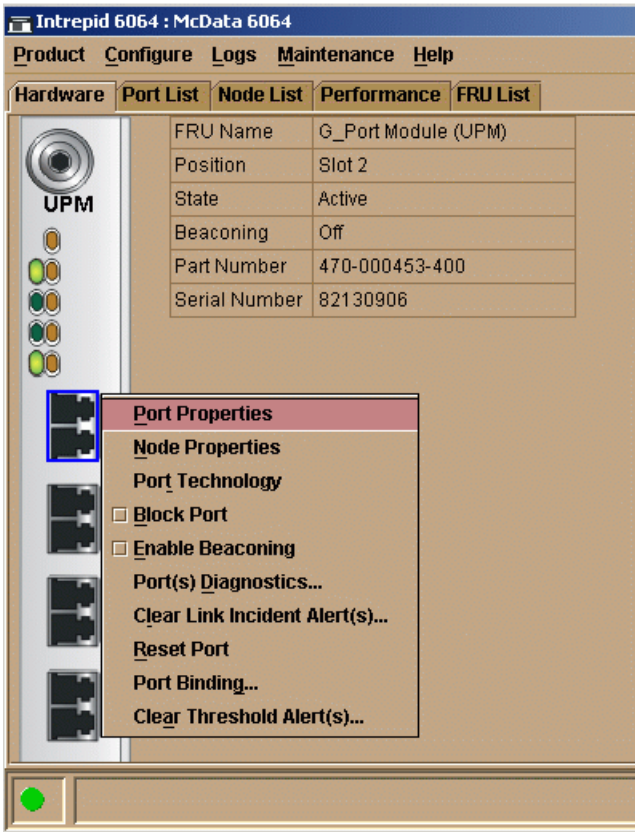


*Figure 4-31 ED-6064 port card viewing and configuration options*

From here, we can perform actions on the port such as resetting the port or performing diagnosis. To go back to the full Hardware View, we click once in the the field shown in Figure 4-32.

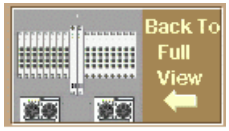


*Figure 4-32 Product Manager ED-6064: Back to Hardware View*

## Configuring the director identification

There are fields for the name, description, location and a contact point for the director in the main window. This is useful to distinguish among a number of installed directors.

To configure this information, we select **Configure —> Identification...**, and are presented with a dialog window with data entry fields, as shown in Figure 4-33.

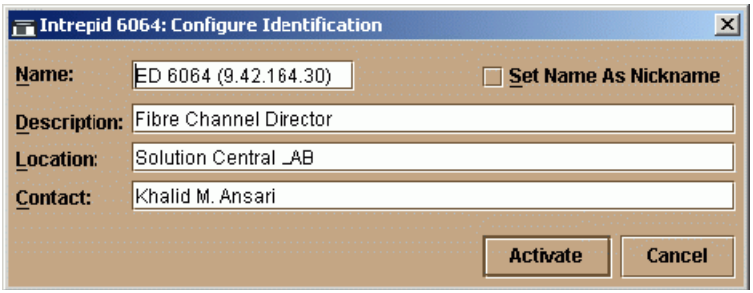


*Figure 4-33 Product Manager ED-6064: Configure Identification*

After activation, the display of the main window changes and places the name of the director in the title bar, and the name, description and location displayed in the window, as shown in Figure 4-34. This information is used in various locations of the Product Manager to identify the selected director.

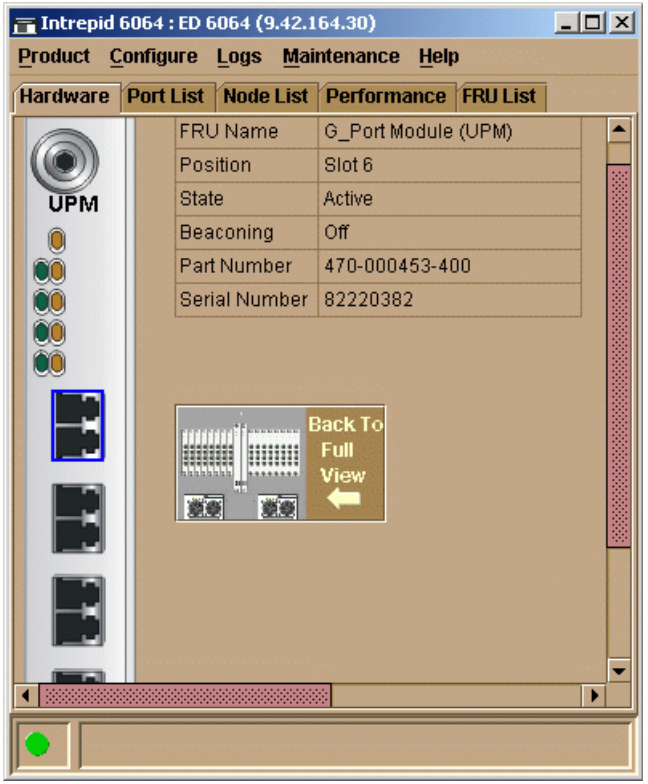


*Figure 4-34 ED-6064 Hardware View changed director information*

## Configuring the Management Style

The ED-6064 features the capability to change the operating mode. To configure it, we select **Configure —> Management Style** and get the following option menu shown in Figure 4-35.

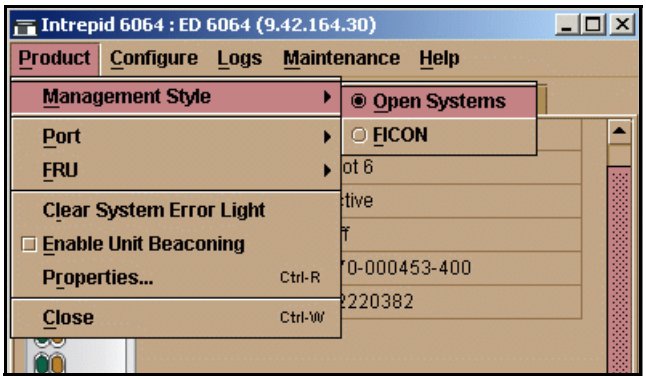


*Figure 4-35 Product Manager ED-6064: Configure Management Style*

This provides a secondary menu with options for Open Systems and FICON management styles. These options change some Product Manager menu options to allow management of the director in open systems or FICON environments.

**Open Systems**. Select **Management Style --> Open Systems** button for (non-FICON) FCP environments.

**FICON**. select **Management Style -->FICON** button when attaching an IBM S/390 Parallel Enterprise or zSeries server to the director and implementing in-band director management through a Fibre Connection (FICON) channel.

For more details, refer to the ED-6140 / ED-6064 user's manual second edition at the following Web site:

`http://www.mcdata.com/downloads/pslgn/techdocs/ed6064/6064_6140_User2nd_ed.pdf`

## Configure Open Fabric 1.0 Mode

McDATA provides with their switches and directors the Open Fabric 1.0 option that allows the McDATA switch or director to interconnect with multi-vendor fabrics such as Cisco, Brocade, BladeCenter, and CNT.

We select the following menu options: **Configure -->Operating Parameters-->Fabric Parameters**, as shown in Figure 4-36.

**Restriction:** The port based zoning is not available with McDATA switch and director in Open Fabric 1.0 mode. Contact McDATA support for details on compatible firmware levels and tested configurations in Open Fabric mode.

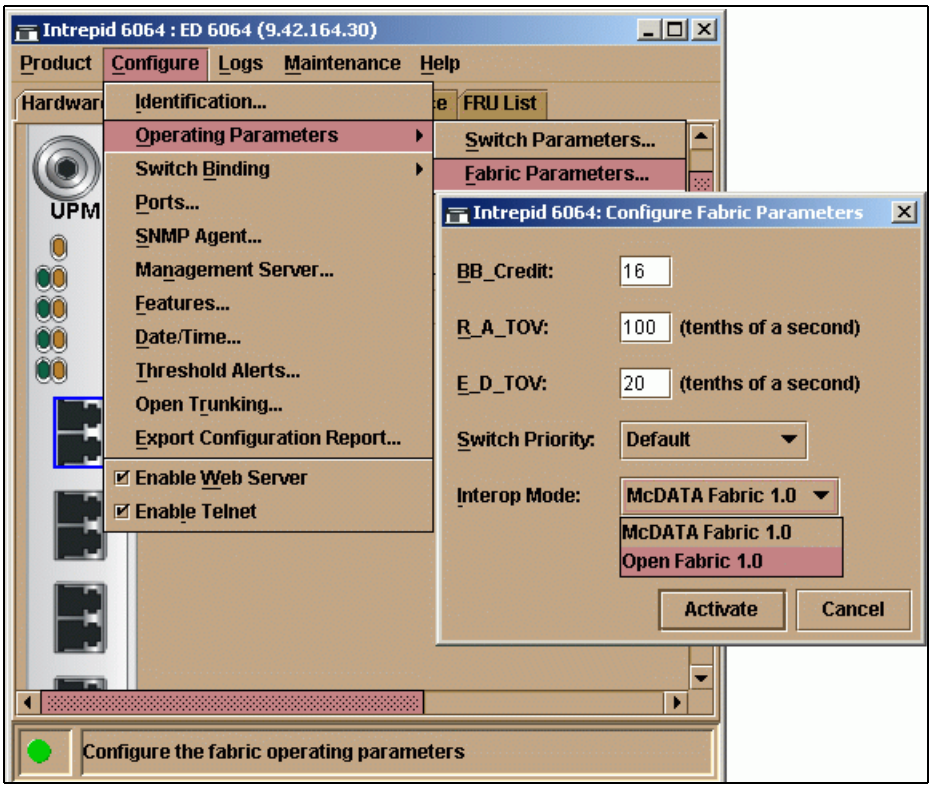


*Figure 4-36 Product Manager ED-6064: Configure Operating Mode Open Fabric*

If the Interop Mode is configured for Open-Fabric 1.0, so any open fabric compliant switch can be connected. Those are then visible in the EFC Manager and the Fabric can be zoned from within the Fabric Manager. However, these switches cannot be managed through the Product Manager.

In McDATA Fabric 1.0 mode, the connectivity is restricted to McDATA switches. If connected to a non-McDATA switch it will be marked as 'Invalid Attachment'.

## Configuring the FC ports

To configure the options relating to each port, we select **Configure —> Ports...** and we are now presented with the Configuration Ports window, which is shown in Figure 4-37.

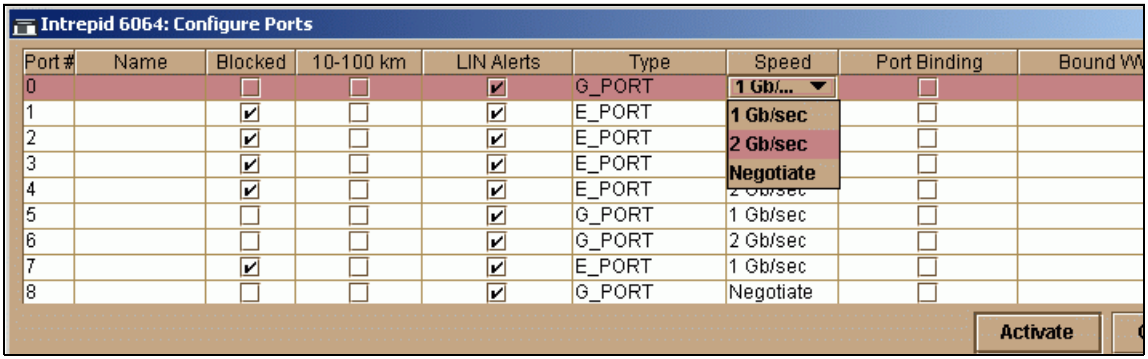


*Figure 4-37 Product Manager ED-6064: Configure Ports*

The port number is automatically assigned and cannot be changed. We can specify a port name here, but this is only useful if the cabling on the port does not change often. The port name then appears in the Product Manager to identify the port, for example in the Port Properties dialog box.

Of more interest here is the ability to block specific ports, or to use extended distance buffering when connecting remote sites with channel extenders, and to define the port type. The ports are by default G_Ports, which means they can auto configure by sensing the type of node ports attached on the port. For example, a G_Port will act as an E_Port if connected to another switch port. By left-clicking in a ports **Type** field we can lock the port type to be only used as E_Ports for ISLs, or as F_Ports for connectivity to node ports. Alternatively, we can right-click any **Type** field as shown in Figure 4-38 to reset any port to G_Port.

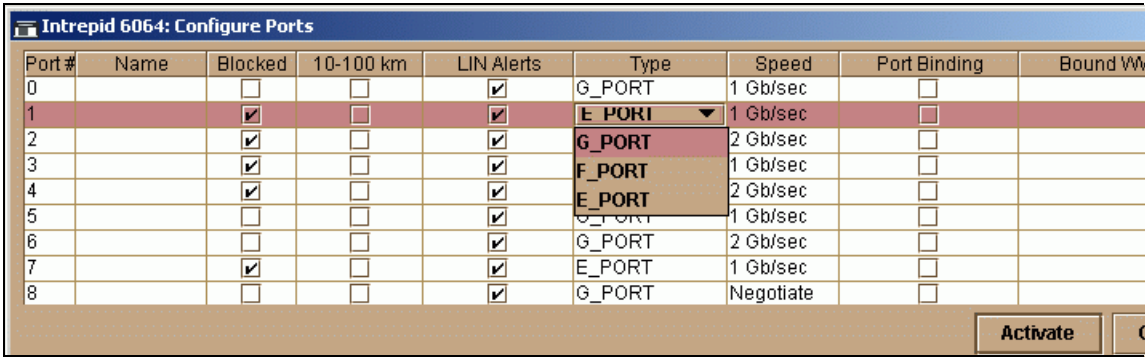


*Figure 4-38 Product Manager ED-6064: Configure Ports port type*

The Port speed can be set in a similar method to the port type, the **Speed** may be set to 2 Gb/s, 1 Gb/s, or Auto-negotiate. Also,.link incident (LIN) alerts can be disabled here. A link incident is a problem on a link which is visible in the Link Incident Log. It is indicated with a small yellow triangle next to the port.

To view the LIN log, go to **Logs -> Link Incident Log...**, as shown in Figure 4-39.

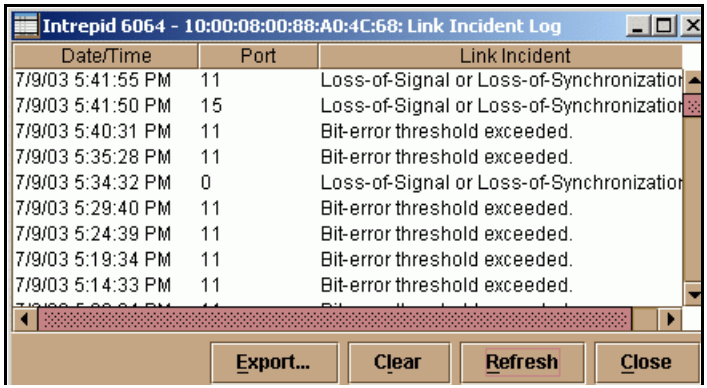


*Figure 4-39 Product Manager ED-6064: Link Incident Log*

## Using the Port List View

To view the status of all installed ports in a tabular view and see the changes that have been made, we select the Port List Tab, as shown in Figure 4-40.

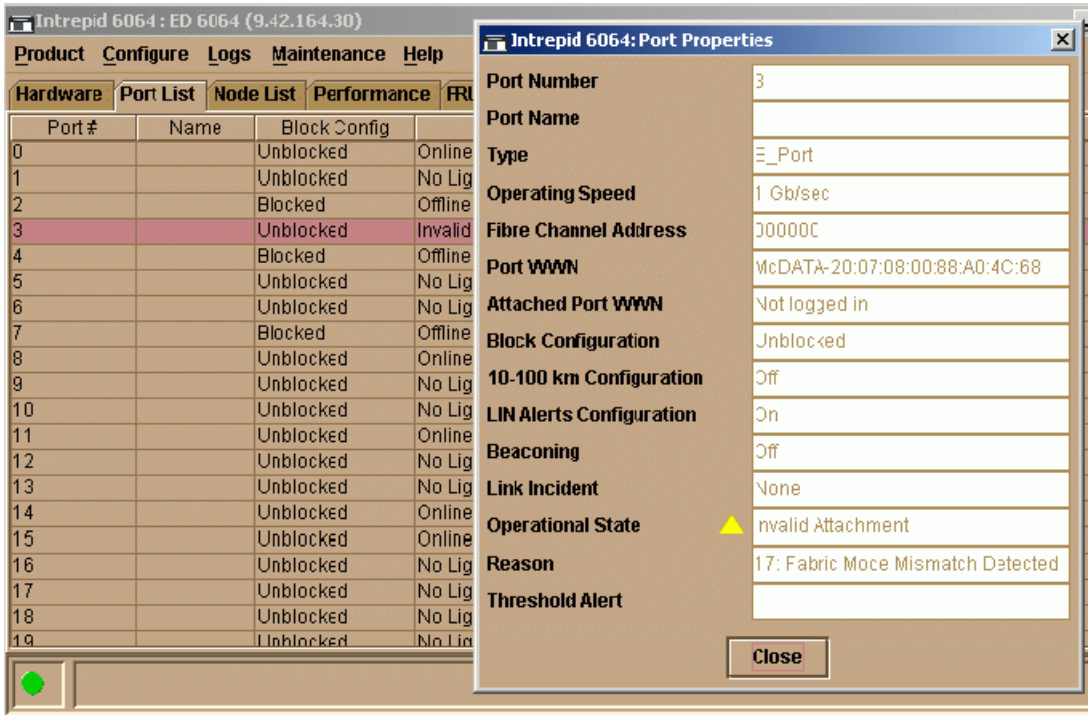


*Figure 4-40 Product Manager ED-6064: Port List View Port Properties*

We can see that Port 3 has a problem. Double-clicking the Port line opens the Port Properties window indicating that the port is not operational due to an invalid attachment. Additionally there is a reason for this indicated as a fabric mode mismatch detected.

## Configuring the FC operating parameters

To change the operating parameters, we first have to set the ED-6064 offline. To set the director offline, which will terminate all FC operations, we select **Maintenance —> Set Online State...** which is shown in Figure 4-41.

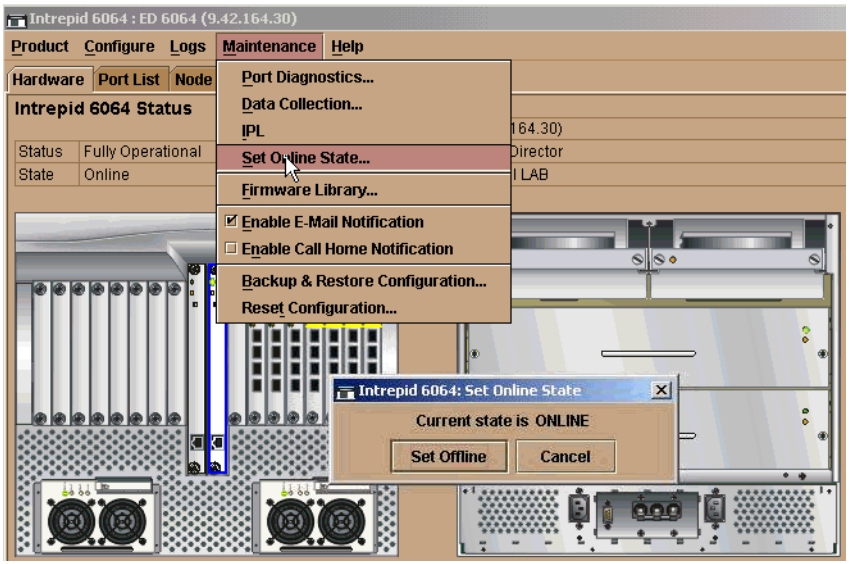


*Figure 4-41 Product Manager ED-6064: Set Online State*

Now we can go to the Configure Operating Parameters window by selecting **Configure —> Operating Parameters.---> Fabric Parameters..**. Here we can change some of the Fibre Channel parameters for the director, for example, the switch priority and flow control values. This is shown in Figure 4-42.

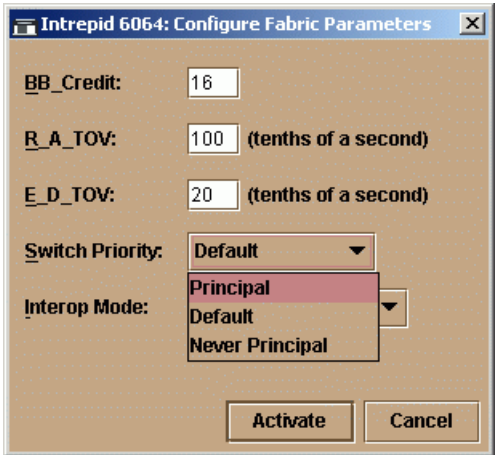


*Figure 4-42 Product Manager ED-6064: Configure Operating Parameters*

The BB_Credit of 16 is the default credit value. The R_A_TOV is a time-out value for operations that depend on the maximum time that frames can be delayed and still be delivered. The E_D_TOV defines the time that the director waits for a response before declaring an error condition. The flow control values R_A_TOV,E_D_TOV values must be the identical on all switches in order to build a multi switch fabric.

The switch priority is used to define the principal switch in a multi switch fabric. We may want this director to always be the Principal switch, or we may select Never Principal on an edge switch, for example an ES-4500. By using Default as our switch priority, the Principal is automatically negotiated. We can also set the Interop Mode form here.

We now move on to configure the switch parameters by selecting **Configure ---> Operating Parameters ---> Switch Parameters.** An important configuration parameter here is the Preferred Domain ID. The director speed if set to 2 Gb/s allows the director to support 1 and 2 Gb/s on all the ports. The Preferred Domain ID and Director Speed configuration options are shown in Figure 4-43.

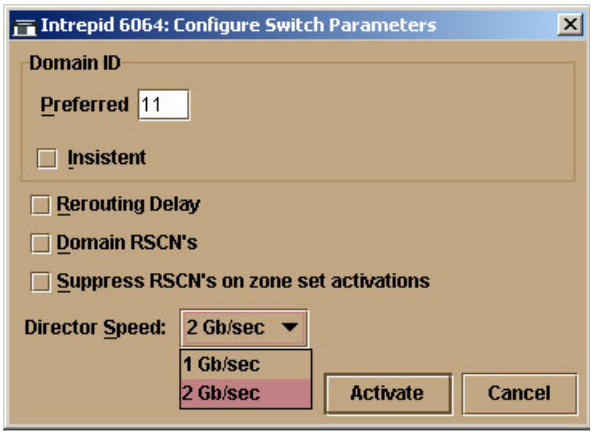


*Figure 4-43 Configure Preferred Domain ID and Speed*

The domain ID has to be unique for each switch or director within a multi switch fabric. If we add a switch or director to our SAN, its domain ID is allocated by the principal switch. If the preferred domain ID is already in use by another switch, then the switch will get another unused domain ID. However, when two fabrics join and they recognize a domain ID conflict, the fabric will become segmented.

If other switches or directors join the fabric, there will be a rerouting delay. This is to ensure that frames are delivered in the correct order in a multi-switch fabric. Also, the routes through the fabric will be recalculated to make sure that the shortest path is taken first.

Now the director is ready for production use in the Fibre Channel network. It can be connected to devices, such as other switches, storage, or hosts.

### 4.5.3 Configuring ES-4500 switch for arbitrated loop

The ES-4500 replaces the ES-1000 switch and it is an enhancement to ES-1000 hub. It provides arbitrated loop topology support directly by configuring its ports as FL_Port.

We are using the ES-4500 switch configuration using EFCM to demonstrate the arbitrated loop topology support by connecting an IBM 3584 LTO device.

We will identify the ES-4500 Sphereon switch from the EFC Manager by selecting **Product---->New** and select the ES-4500 Sphereon from the drop down menu and assign a valid IP address, as shown in Figure 4-44.

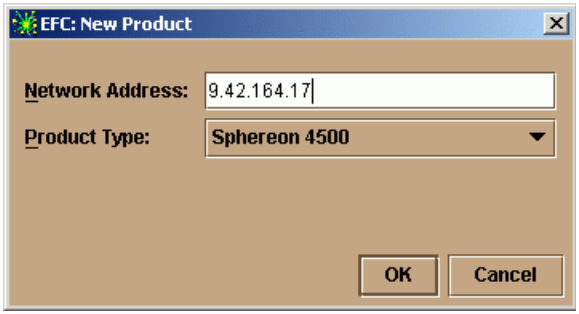


*Figure 4-44 Configure ES-4500 Identification from EFC Product Manager*

Adding the new device will create an icon in the product manager window, as shown in Figure 4-45.

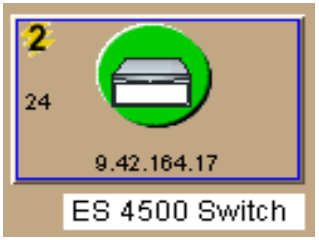


*Figure 4-45 ES-4500 Sphereon Switch icon in the EFC Product manager*

Select the ES-4500 icon from the product manager window to perform configuration and management as shown in Figure 4-46, and to show the switch's front and rear view.

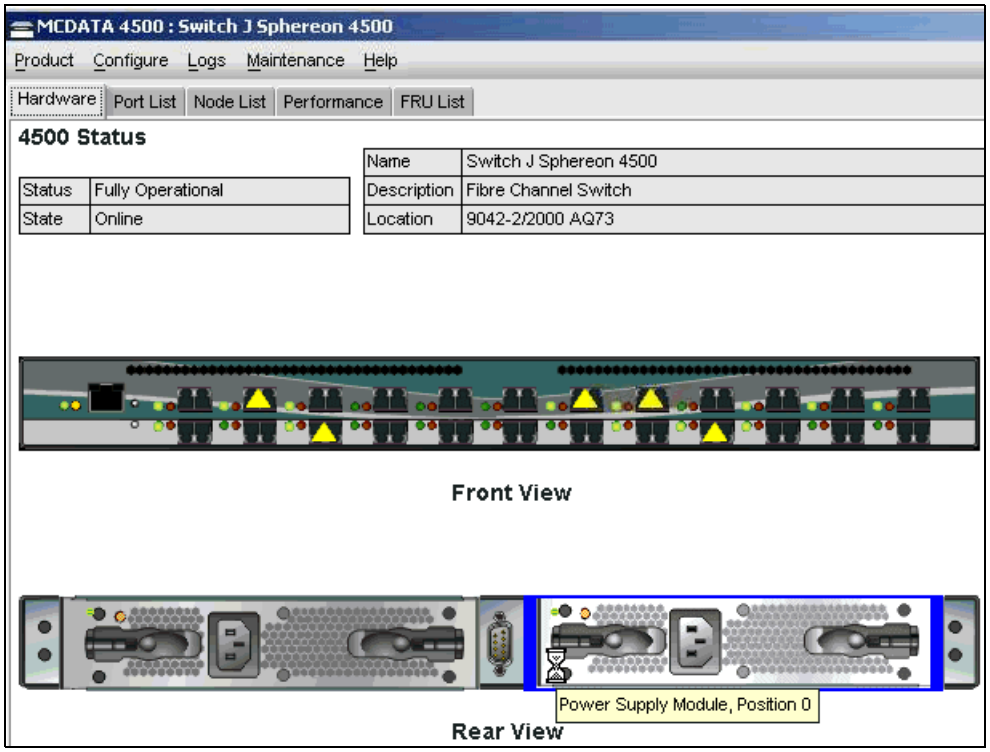


*Figure 4-46 ES-4500 switch front and rear view*

From the ES-4500 product manager we will configure the operating parameters. To do this, we select **Configure—>Operating Parameters** as shown in Figure 4-47.

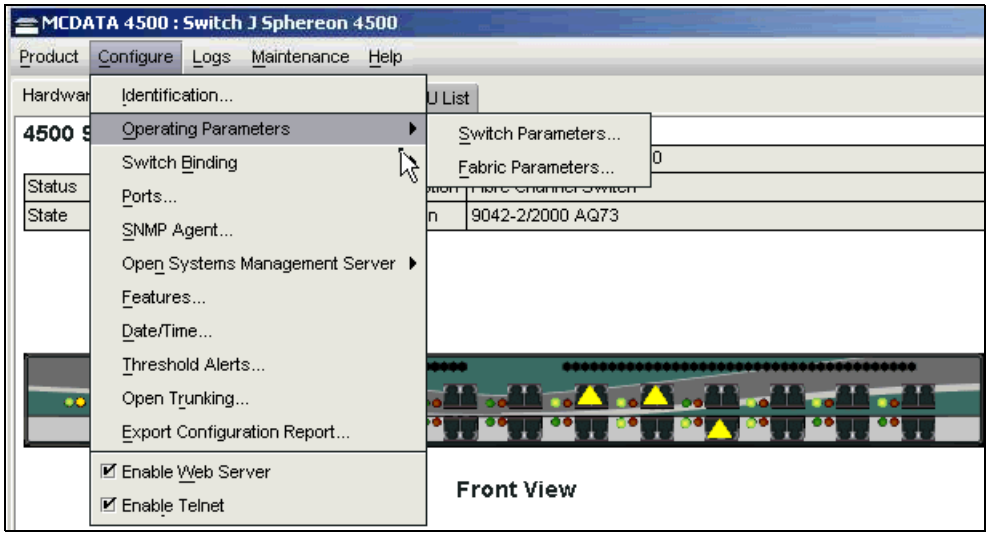


*Figure 4-47 ES-4500 Operating Parameters menu*

## Configuring the ES-4500 operating parameters

As the ES-4500 is deployed in a core-to-edge topology, we will select the ES-4500 to be **Never Principal** so that the ED-6064 is always the principal switch in the fabric. Notice in the Configure Fabric Parameters menu the BB_Credit option is not listed. This is because of the ES-4500 shared memory architecture and it has a fixed allocation of BB_Credits across the 24 ports. We will configure the ES-4500 to be **Never Principal**, as we want the director to be the **Principal**.

To configure this, select **Configure----> Operating parameters ----> Fabric Parameters** from the ES-4500 product manager and set the switch priority to be **Never Principal**. The R_A_TOV and E_D_TOV values will not be changed from the default, as shown in Figure 4-48.

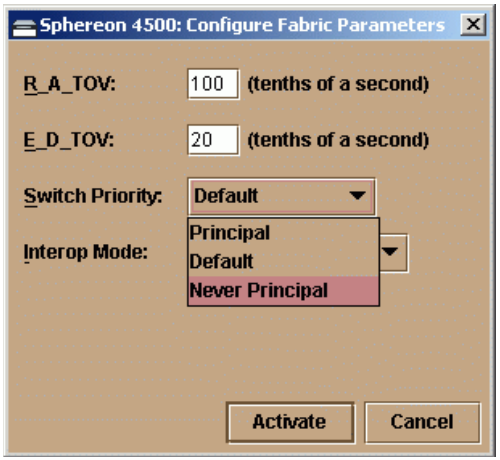


*Figure 4-48 Configure Fabric Parameters menu*

To configure the preferred Domain ID for ES-4500 switch select **Configure -----> Operating Parameters -----> Switch Parameters**, the domain ID has to be unique throughout the fabric. We will assign 5 as the domain ID value on the ES-4500 switch, as shown in Figure 4-49.

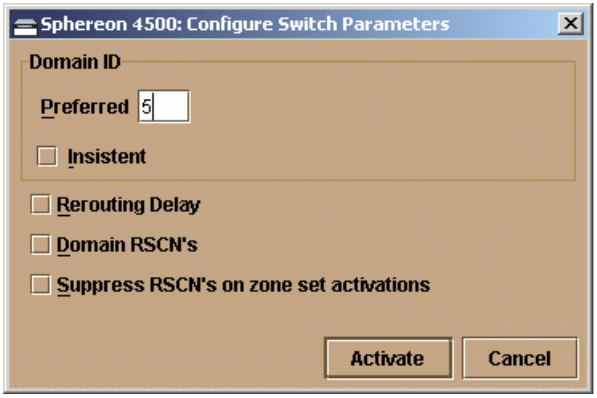


*Figure 4-49 Configure Switch Parameters menu*

The other parameters from the Configure Switch Parameters menu will be set to default and remain unselected as shown above.

## 4.5.4 ES-4500 port configuration options

The ES-4500 port configuration features are unique from the other IBM/McDATA products. It provides GX, FX, G, F and E_Port options; while the ED-6140, ED-6064, and ES-3232 provide F_Port, E_Port, and G Port options.

### GX_Port

The GX_Port is the default option and it can auto configure to F_Port, FL_Port or E_Port. The GX_Port should always be the preferred port setting in order to connect an ISL or N port (fabric node) or an FL_Port (arbitrated loop public or private device). A private device can only be attached to a GX_Port.

### FX_Port

The FX_Port option will lock the port to auto-configure as either an F_Port or an FL_Port. The FX_Port does not allow an ISL to another switch.

### G_Port

The G_Port option will allow the port to auto-configure as an F_Port or an E_Port.

### F_Port

If the port is chosen as an F_Port, then it disables the E_Port and FL_Port function on that port.

### E_Port

If selected, then only inter switch links (E_Port) are allowed on that port.

## 4.5.5 ES-4500 switch port configuration

The ES-4500 port configuration menu is displayed. Port #'s 8 and 10 are used as inter switch links (ISLs) to the ED-6064, and on Port # 15, a Windows host is attached. The GX_Port can auto-sense the connecting device and configure itself accordingly. For example, if Port # 8 is set as a **GX**_Port, it should automatically detect the ED-6064 and automatically configure itself as an E_Port. Port # 15 is also set as a **GX_Port** and can automatically detect the Windows host attachment and automatically configure itself as an F_Port.

Similarly, Port # 5 is configured as an FX_Port so that it can auto-sense the connecting arbitrated loop device and automatically configure itself as an FL_Port.

Figure 4-50 shows that Ports # 8 and # 15 are set as GX_Ports, Port # 10 is set as an E_Port, and Port # 5 is configured as an FX_Port.

Click the **Activate button** after making the changes to the port types as shown in Figure 4-50.

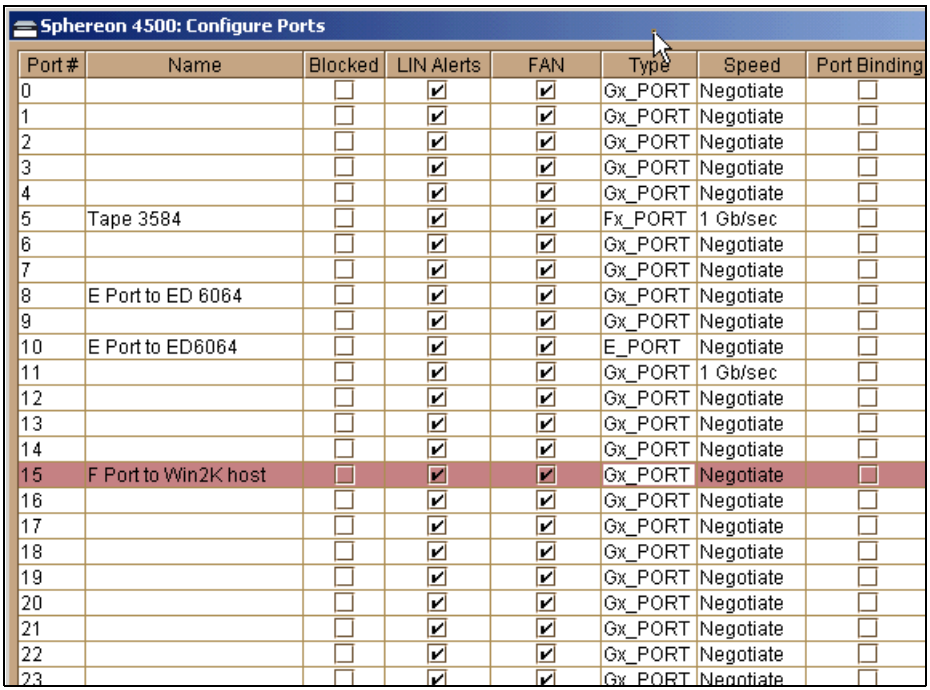


| Port # | Name | Blocked | LIN Alerts | FAN | Type | Speed | Port Binding |
|---|---|---|---|---|---|---|---|
| 0 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 1 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 2 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 3 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 4 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 5 | Tape 3584 | ☐ | ☑ | ☑ | Fx_PORT | 1 Gb/sec | ☐ |
| 6 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 7 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 8 | E Port to ED 6064 | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 9 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 10 | E Port to ED6064 | ☐ | ☑ | ☑ | E_PORT | Negotiate | ☐ |
| 11 | | ☐ | ☑ | ☑ | Gx_PORT | 1 Gb/sec | ☐ |
| 12 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 13 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 14 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 15 | F Port to Win2K host | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 16 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 17 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 18 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 19 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 20 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 21 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 22 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |
| 23 | | ☐ | ☑ | ☑ | Gx_PORT | Negotiate | ☐ |

*Figure 4-50 ES-4500 port configuration options*

## ES-4500 port list view

To view the port status and verify if all ports have automatically configured themselves correctly, click the **Port List** tab from ES-4500 product manager view as shown in Figure 4-51.

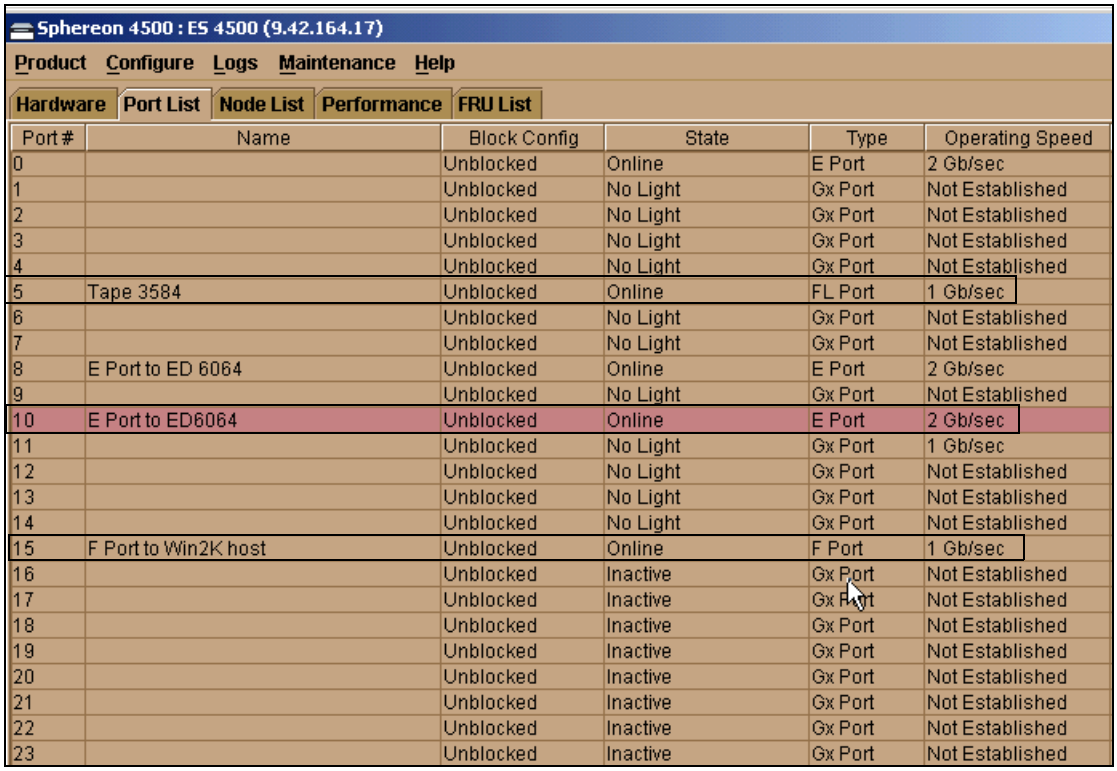


| Port # | Name | Block Config | State | Type | Operating Speed |
|---|---|---|---|---|---|
| 0 | | Unblocked | Online | E Port | 2 Gb/sec |
| 1 | | Unblocked | No Light | Gx Port | Not Established |
| 2 | | Unblocked | No Light | Gx Port | Not Established |
| 3 | | Unblocked | No Light | Gx Port | Not Established |
| 4 | | Unblocked | No Light | Gx Port | Not Established |
| 5 | Tape 3584 | Unblocked | Online | FL Port | 1 Gb/sec |
| 6 | | Unblocked | No Light | Gx Port | Not Established |
| 7 | | Unblocked | No Light | Gx Port | Not Established |
| 8 | E Port to ED 6064 | Unblocked | Online | E Port | 2 Gb/sec |
| 9 | | Unblocked | No Light | Gx Port | Not Established |
| 10 | E Port to ED6064 | Unblocked | Online | E Port | 2 Gb/sec |
| 11 | | Unblocked | No Light | Gx Port | 1 Gb/sec |
| 12 | | Unblocked | No Light | Gx Port | Not Established |
| 13 | | Unblocked | No Light | Gx Port | Not Established |
| 14 | | Unblocked | No Light | Gx Port | Not Established |
| 15 | F Port to Win2K host | Unblocked | Online | F Port | 1 Gb/sec |
| 16 | | Unblocked | Inactive | Gx Port | Not Established |
| 17 | | Unblocked | Inactive | Gx Port | Not Established |
| 18 | | Unblocked | Inactive | Gx Port | Not Established |
| 19 | | Unblocked | Inactive | Gx Port | Not Established |
| 20 | | Unblocked | Inactive | Gx Port | Not Established |
| 21 | | Unblocked | Inactive | Gx Port | Not Established |
| 22 | | Unblocked | Inactive | Gx Port | Not Established |
| 23 | | Unblocked | Inactive | Gx Port | Not Established |

*Figure 4-51 Port list menu*

It can be confirmed that:

- Port # 5 has automatically configured itself as an FL_Port because an arbitrated loop device is connected.
- Port # 8 has automatically configured itself as an E_Port as it has an ISL to ED-6064.
- Port # 10 has automatically configured itself as an E_Port as it has an ISL to ED-6064.
- Port # 15 has automatically configured itself as an F_Port as a Windows device is attached on that port.

In all cases the operating speed has also been automatically negotiated.

## Tape attachment to the ES-4500

The ES-4500 switch with its **GX** or **FX_Port** capability provides connectivity to Fibre Channel arbitrated loop capable devices. The port type, if set to the default **GX_Port** or if configured as an **FX_Port**, will automatically detect an attached FC-AL device and configure the port as an FL_Port.

**Tip:** The latency involved during the link initialization process can be reduced by manually configuring the port type to FX_Port and changing the speed mode from auto-negotiate to 1 or 2 Gb/s. If the port type and speed are set to default values (GX_Port and auto-negotiate) then the port has to go through the various stages of link initialization such as speed negotiation and port configuration before bringing the port to an online state.

The IBM 3584 Tape Library is attached on Port # 5 as an FL_Port with 1 Gb/s speed. If you notice in Figure 4-53, Port # 5 configured type is set to default **GX_Port** but it can also be locked as an **FX_Port** so that it only allows an F_Port connection or an FL_Port connection.

The tape device attached on Port # 5 is in an Online state with the port type displayed as FL_Port and 1 Gb/s operating port speed as shown in Figure 4-52.

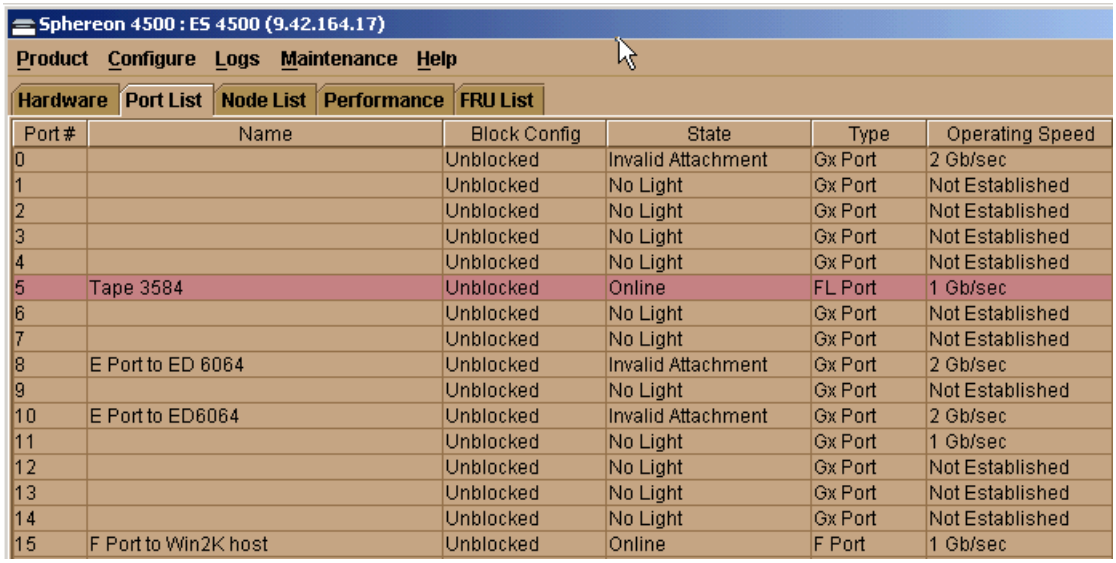


| Port # | Name | Block Config | State | Type | Operating Speed |
|---|---|---|---|---|---|
| 0 | | Unblocked | Invalid Attachment | Gx Port | 2 Gb/sec |
| 1 | | Unblocked | No Light | Gx Port | Not Established |
| 2 | | Unblocked | No Light | Gx Port | Not Established |
| 3 | | Unblocked | No Light | Gx Port | Not Established |
| 4 | | Unblocked | No Light | Gx Port | Not Established |
| 5 | Tape 3584 | Unblocked | Online | FL Port | 1 Gb/sec |
| 6 | | Unblocked | No Light | Gx Port | Not Established |
| 7 | | Unblocked | No Light | Gx Port | Not Established |
| 8 | E Port to ED 6064 | Unblocked | Invalid Attachment | Gx Port | 2 Gb/sec |
| 9 | | Unblocked | No Light | Gx Port | Not Established |
| 10 | E Port to ED6064 | Unblocked | Invalid Attachment | Gx Port | 2 Gb/sec |
| 11 | | Unblocked | No Light | Gx Port | 1 Gb/sec |
| 12 | | Unblocked | No Light | Gx Port | Not Established |
| 13 | | Unblocked | No Light | Gx Port | Not Established |
| 14 | | Unblocked | No Light | Gx Port | Not Established |
| 15 | F Port to Win2K host | Unblocked | Online | F Port | 1 Gb/sec |

*Figure 4-52 Port # 5 is Online as an FL_Port type.*

If we click the **Node List** tab, we see that the IBM 3584 Tape is connected on Port # 5 and the fabric address is 6509**CA**, where **CA** is the arbitrated loop physical address (AL_PA) of the tape device. This is shown in Figure 4-53.

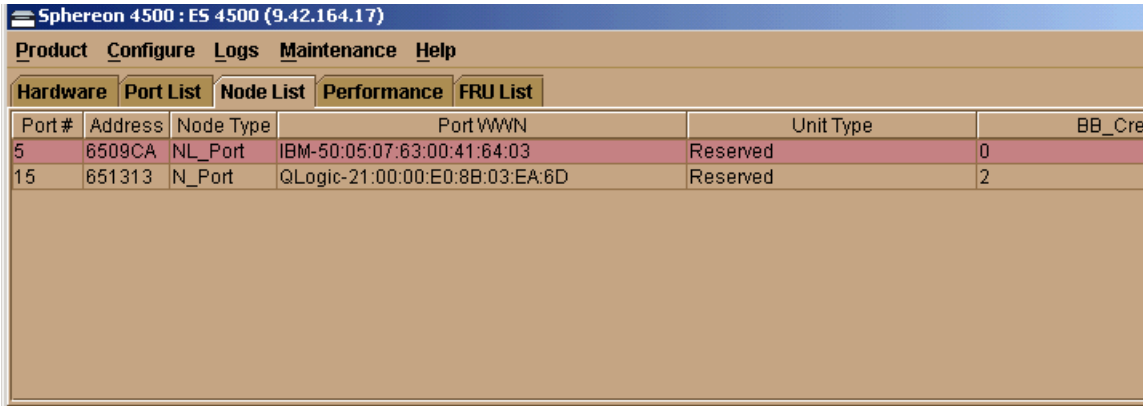


*Figure 4-53 Node List display of tape device*

# 4.6 Troubleshooting the McDATA SAN

In the sections that follow we will show some of the ways in which you can troubleshoot the SAN.

## 4.6.1 Identifying and resolving hardware symptoms

In this section, we will identify products that have their attention indicator on (indicating a problem) and then show the steps taken to identify and resolve the cause.

In Figure 4-54 we can see that our ED-6064 director and ES-3016 require attention.

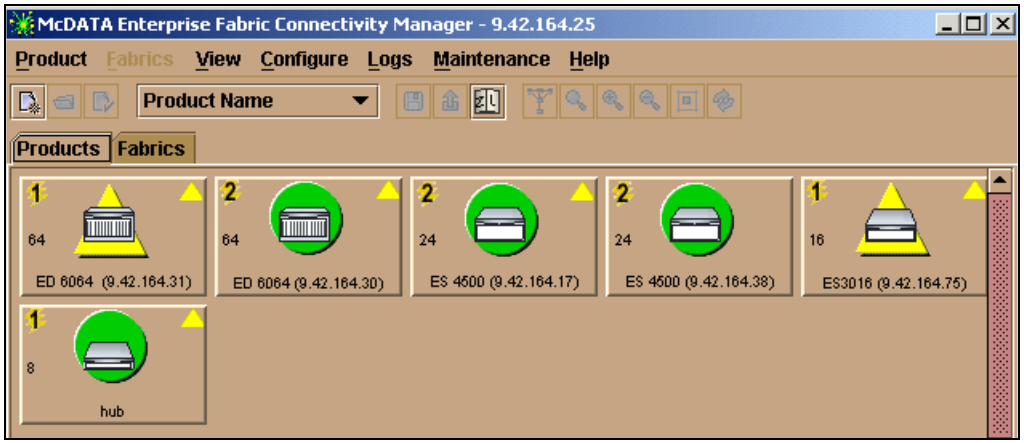


*Figure 4-54 EFCM indicating attention required*

By double-clicking the ED-6064 icon (with the IP address of 9.42.164.31) the product menu window is opened as shown in Figure 4-55.

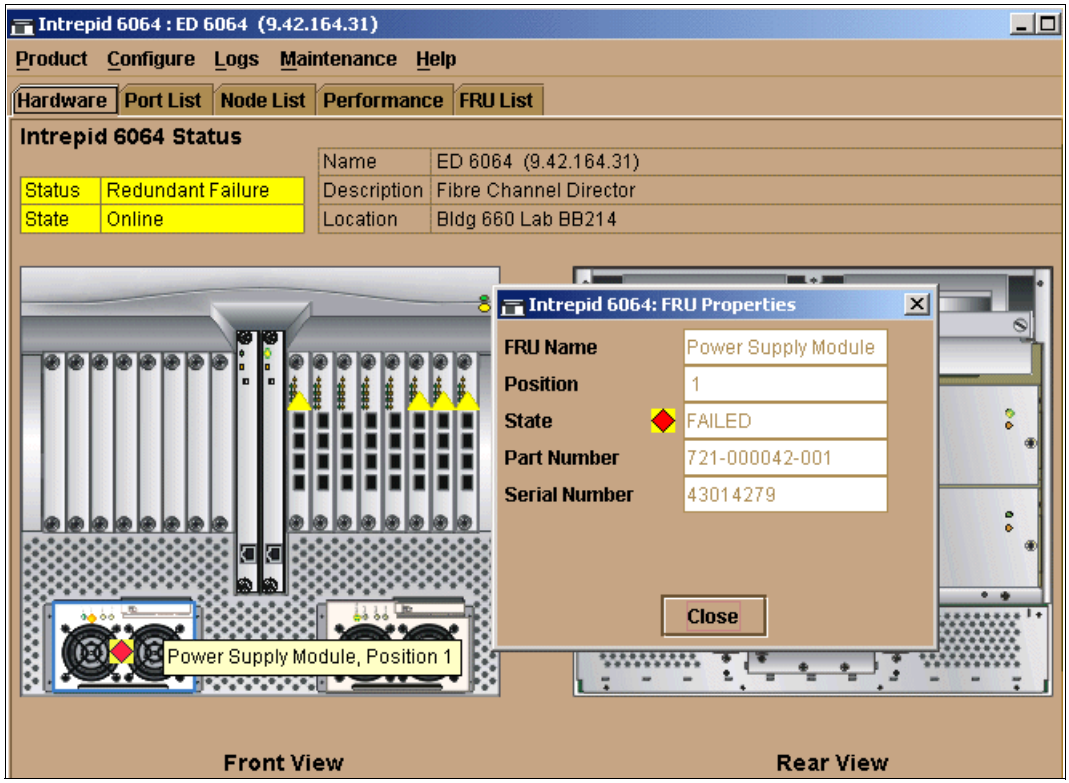


*Figure 4-55 Attention indicators show a failed power supply module*

We notice that the attention indicator is blinking on the ED-6064 power supply # 1, and by double-clicking the blinking icon, the new pop-up window lists the details of the FRU and its state. We can see that the power supply module is in a failed state and is the cause for the attention indicator.

To fix the problem and clear the attention indicator, a service call has to be placed. To open a defect call, you need to gather the device type and serial number of the ED-6064 and then initiate a call to replace the failed power supply.

The S/N can be obtained from the ED-6064 product manager by selecting **Product--->Properties** from the EFC product manager menu as shown in Figure 4-55.

You can also view the ED-6064 event log to retrieve the problem description, its severity, and FRU-position, as shown in Figure 4-56.

Intrepid 6064 - 10:00:08:00:88:A0:D8:D3: Event Log

| Date/Time | Event | Description | Severity | FRU-Position |
|---|---|---|---|---|
| 7/12/03 2:28:41 PM | 081 | Port set to invalid attachment state | Informational | |
| 7/12/03 2:26:38 PM | 081 | Port set to invalid attachment state | Informational | |
| 7/12/03 1:52:48 PM | 081 | Port set to invalid attachment state | Informational | |
| 7/12/03 1:52:47 PM | 201 | Power supply DC voltage failure. | Major | PWR-1 |
| 7/12/03 1:52:46 PM | 070 | E_Port has become segmented. | Informational | |
| 7/12/03 1:52:43 PM | 410 | CTP card reset. | Informational | CTP-0 |
| 7/12/03 1:50:55 PM | 001 | System power-down. | Informational | |

*Figure 4-56 Event log indicates problem*

After installing the new power supply, the attention indicator will disappear and the power redundancy in ED-6064 is restored as shown in Figure 4-57 for the director at IP address 9.42.164.31.

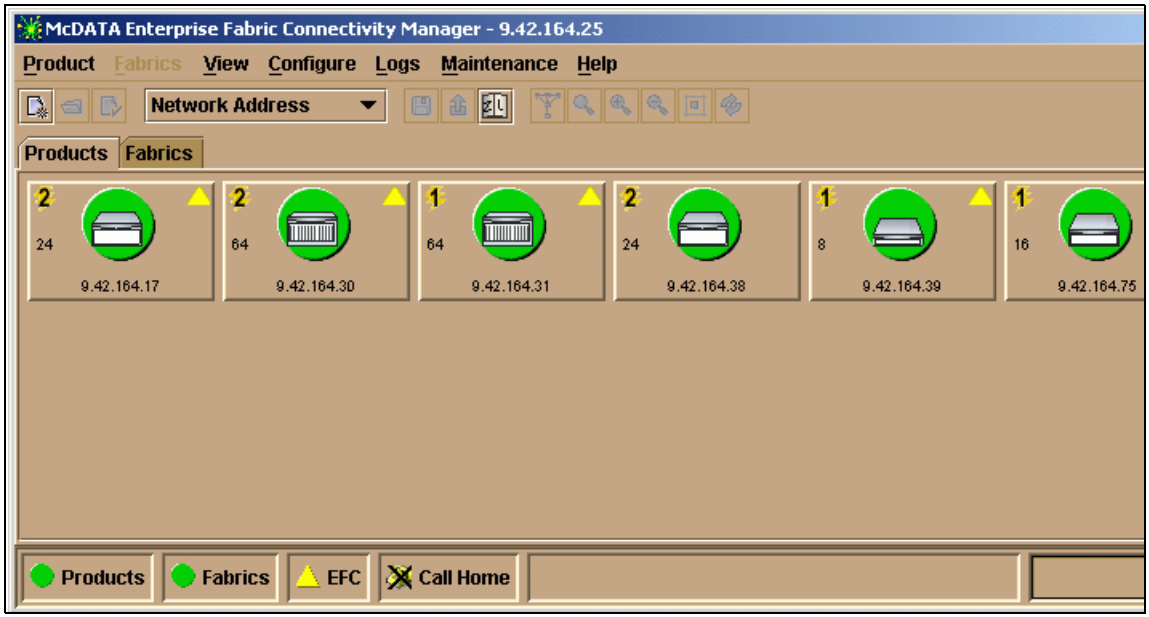


*Figure 4-57 Product icon changed to normal state*

Similarly, the bad power supply and fan units on the ES-3016 were also replaced to restore the switch status from degraded to normal operation.

## 4.6.2 Identifying and resolving fabric segmentation

We will discuss various scenarios that can cause fabric segmentation. In Figure 4-58 we can see that the two ED-6064s are segmented due to incompatible zone configurations. The link between ED-6064 (9.42.164.30) and ED-6064 (9.42.164.31) is segmented as shown in Figure 4-60.

### Segmentation due to incompatible zone configurations

The product event log and the port properties menu are used to quickly identify the operational state of the E_Port and understand the cause of segmentation.

From the port properties menu retrieved by selecting the segmented port # 2 from the port list menu, the fabric has segmented due to incompatible zoning configurations, as shown in Figure 4-58.

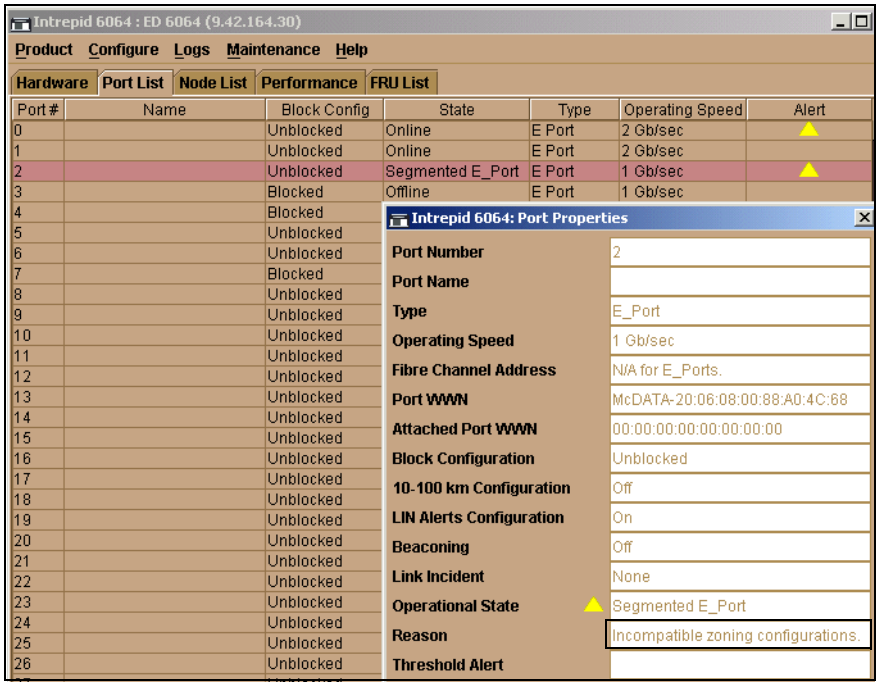


*Figure 4-58 Port List menu shows that the E_Port has segmented*

The EFCM fabric manager Event log is another place to verify the details of the segmentation. It lists Event ID: 070 (E_Port has become segmented) and Event ID:150 (Zone Merge Failure) as shown in Figure 4-59.

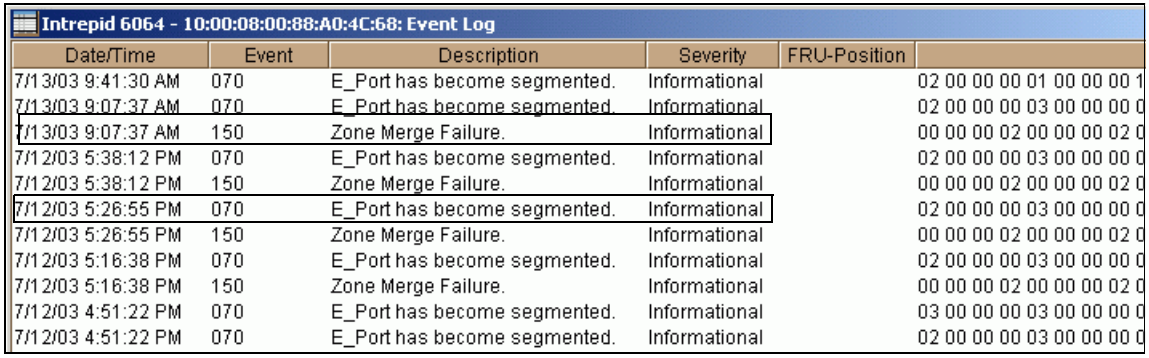


*Figure 4-59 Event log entries indicating segmentation*

To debug the incompatible zone set configuration, we need to identify and select the ED-6064 from the EFC fabric manager, **Fabrics** menu, as shown in Figure 4-60.

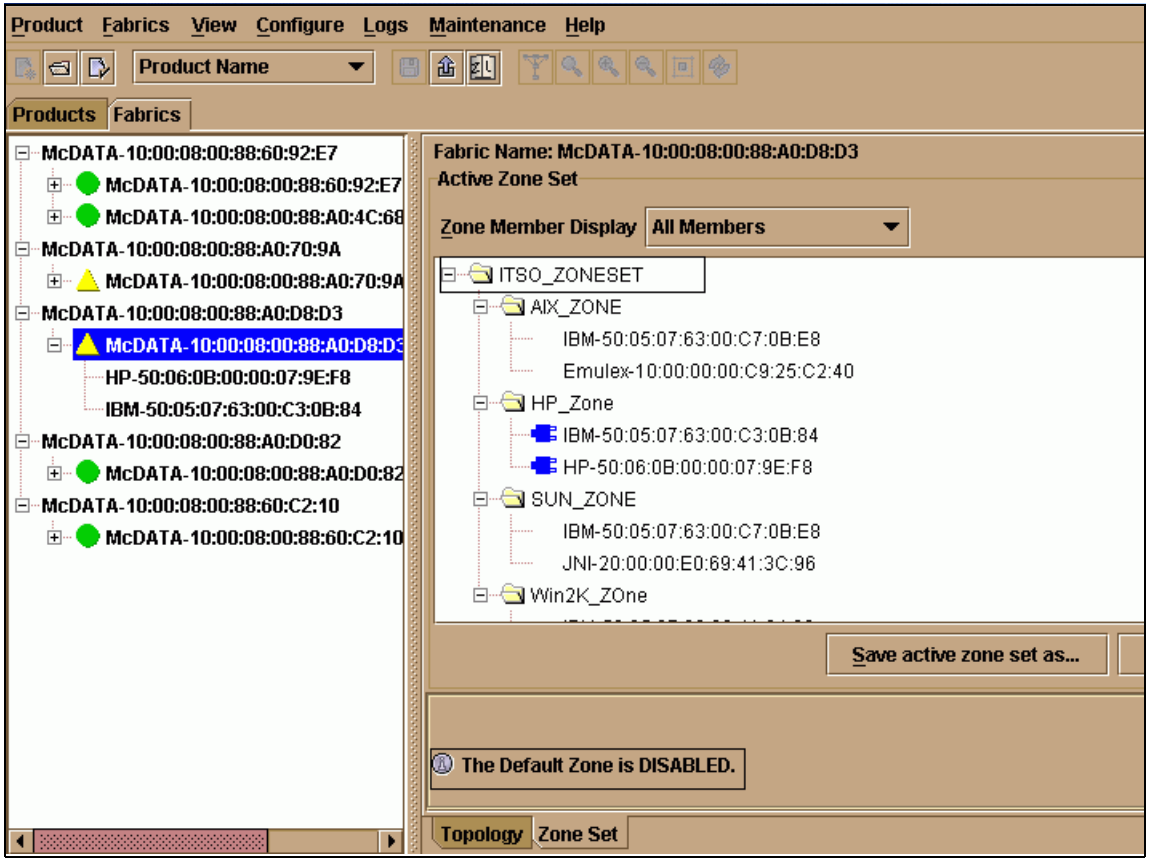


*Figure 4-60 Fabrics active zone set menu*

Notice that there are multiple switches/fabrics managed by this EFC Manager. To be able to view the zone set configuration of a particular fabric or a switch, it has to be selected. The ED-6064(with an IP address of 9.42.164.31) has been selected as it is now highlighted in blue.

Verify the zone set configuration on the ED-6064 (9.42.164.31) by selecting the active zone set menu from the EFCM fabric manager view as shown in Figure 4-60.

The ED-6064 (9.42.164.31) has an active zone set “**ITSO_ZONESET**” with four zones; the HP_Zone has two active devices and the default zone is set to **DISABLED** state.

Now we will verify the zone set configuration on the ED-6064 (with an IP address of 9.42.164.30) by selecting the zone set configuration tab from the EFCM fabric manager view as shown in Figure 4-61.

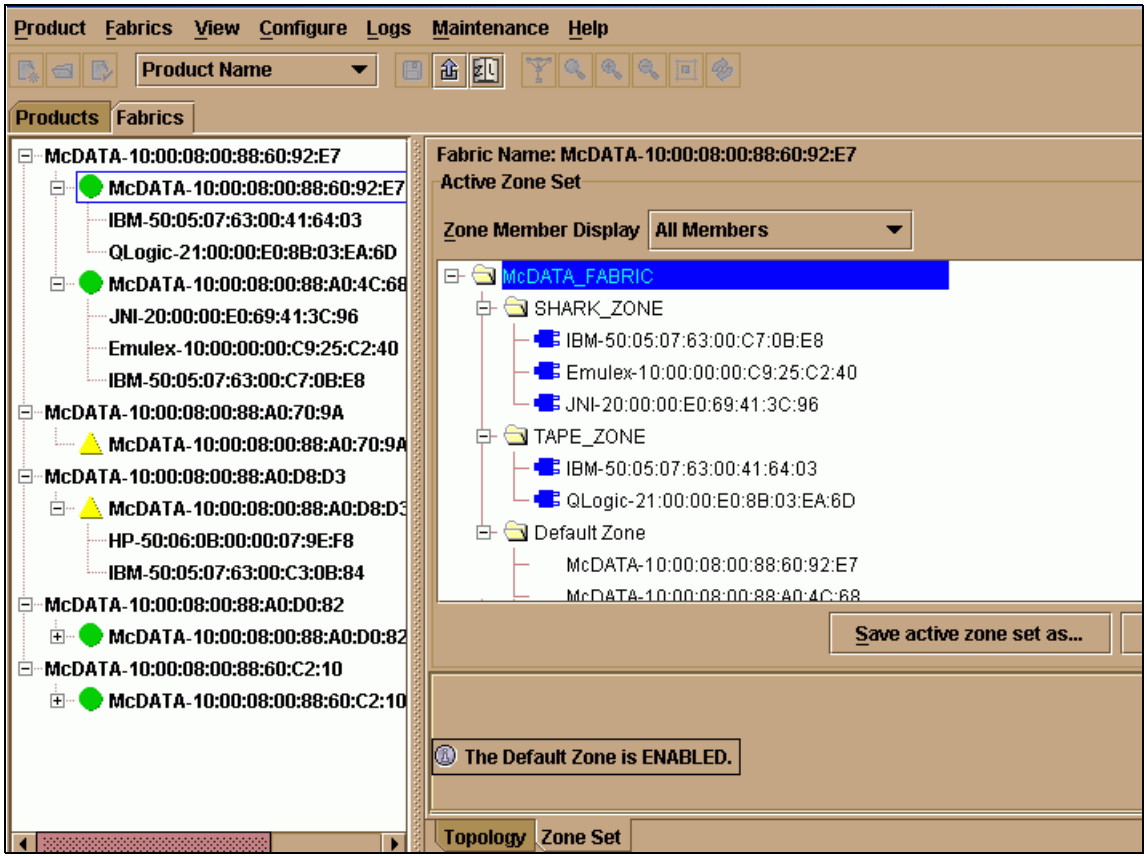


*Figure 4-61 Zone set information and default zone is disabled*

The view shows the active zone set "**McDATA_FABRIC**" has two zones and five devices active, and the default zone is in an **ENABLED** state.

From Figure 4-60 and Figure 4-61, we notice there are two anomalies that are preventing the two ED-6064s from merging:

- Both ED-6064s have conflicting zone sets that are active.
- The ED-6064 (9.42.164.31) has the default zone disabled, while the ED-6064 (9.42.164.30) has the default zone enabled.

To resolve the merge failure issues, two configuration changes are required:

1. Configure the default zone to enabled state on ED-6064 (9.42.164.31).
2. Deactivate the Active Zone Set on any one of the two ED-6064s.

**Attention:** Before making any Zone Set configuration changes from the fabric manager menu, carefully identify and select the switch that requires the default zone **Enabled** and also **Disable** the Active Zone Set. Making configuration changes on a different switch may cause unexpected results.

Identify the switch that requires the configuration changes from the EFC fabric manager, fabrics list on the left-hand side. After the switch is selected, the switch WWN is highlighted in blue as shown in Figure 4-62.

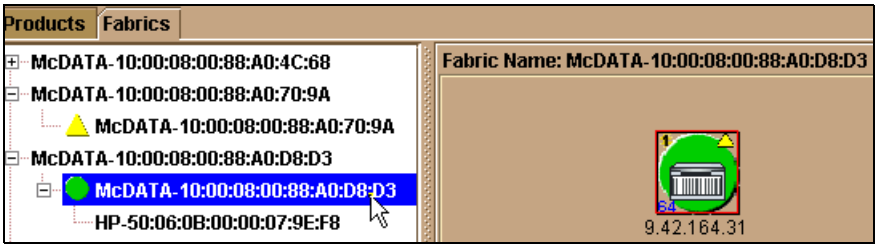


*Figure 4-62 EFCM fabric manager: Fabrics*

After the switch has been identified, then perform the configuration changes by selecting **Configure —> Advanced Zoning —> Configure Default Zone** from the EFCM fabric manager menu shown in Figure 4-63.

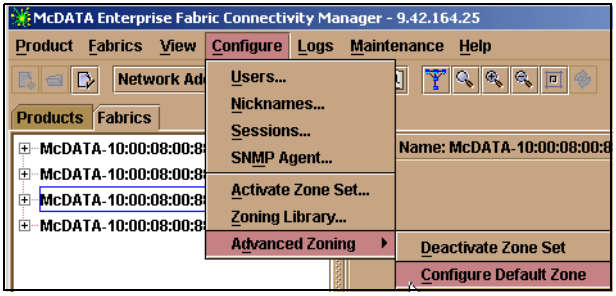


*Figure 4-63 Advanced Zoning configuration menu*

Click **Start** to enable the default zone set as shown in Figure 4-64.

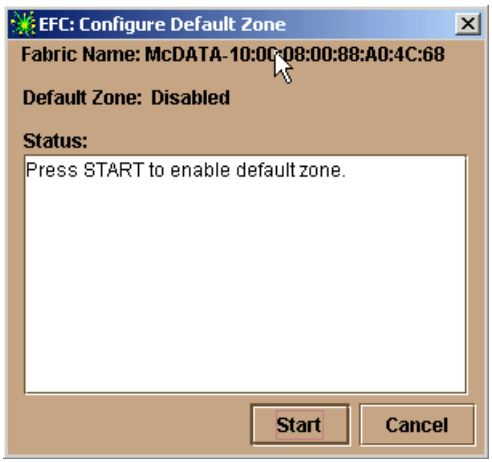


*Figure 4-64 Configure Default Zone menu*

To deactivate the Zone Set, select **Configure —> Advanced Zoning —> Deactivate Zone Set** — then the two fabrics can merge to form one large fabric with 128 ports.

After performing the necessary configuration changes, verify that the two fabrics have merged by analyzing the event log in the product manager menu of any ED-6064, and also by looking at the topology view in the fabric manager, as shown in Figure 4-65.

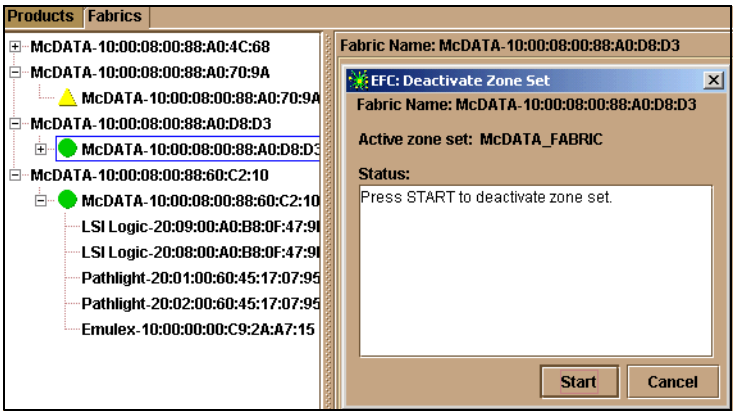


*Figure 4-65 Deactivate Active Zone Set*

## 4.6.3 Segmentation due to domain ID conflict

Another common reason for an ISL to segment is due to a domain ID conflict. Two switches with same domain ID in online state (operational) will not merge even if the **Insistent option** is **unchecked** on both switches. A Segmented E_Port error will appear as shown in Figure 4-66.

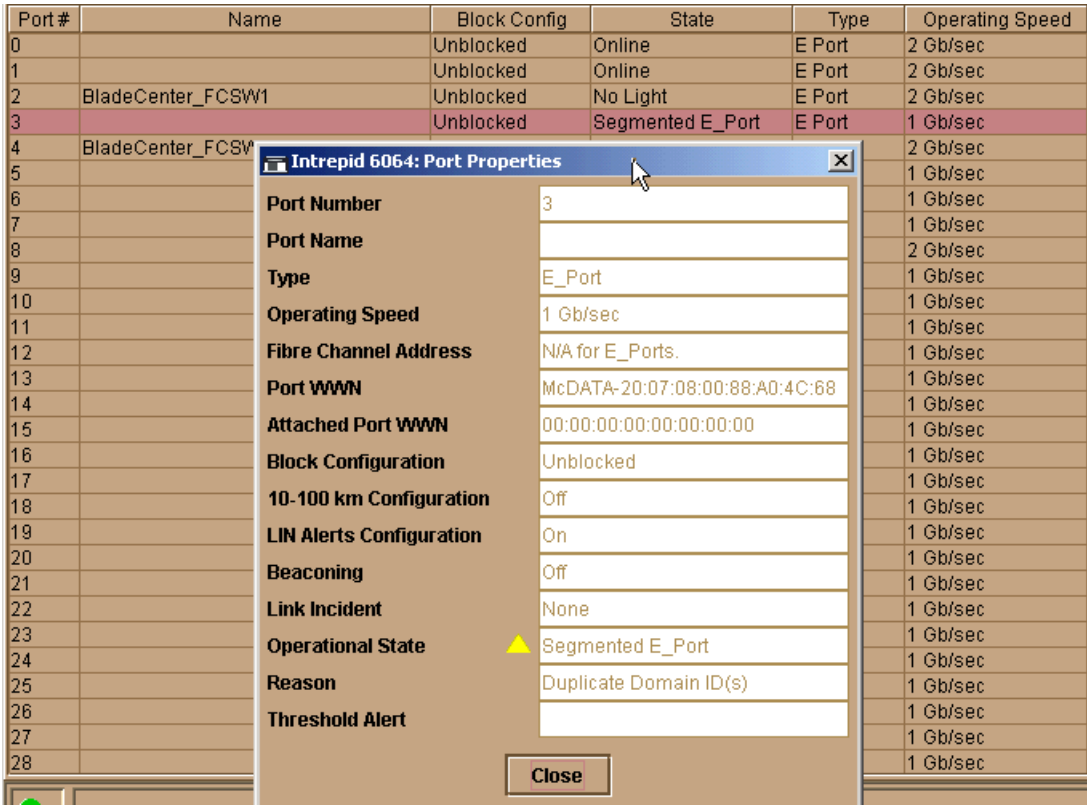


*Figure 4-66 Port Properties menu*

From the **Port Properties** menu **Port Number 3** the **Operational State** is shown as segmented with a **Reason** of Duplicate Domain ID(s).

The Domain ID configuration can be verified from the Switch Operating Parameters by selecting **Configure—Operating Parameters —> Switch Parameters** from the EFC product manager menu as shown in Figure 4-67.

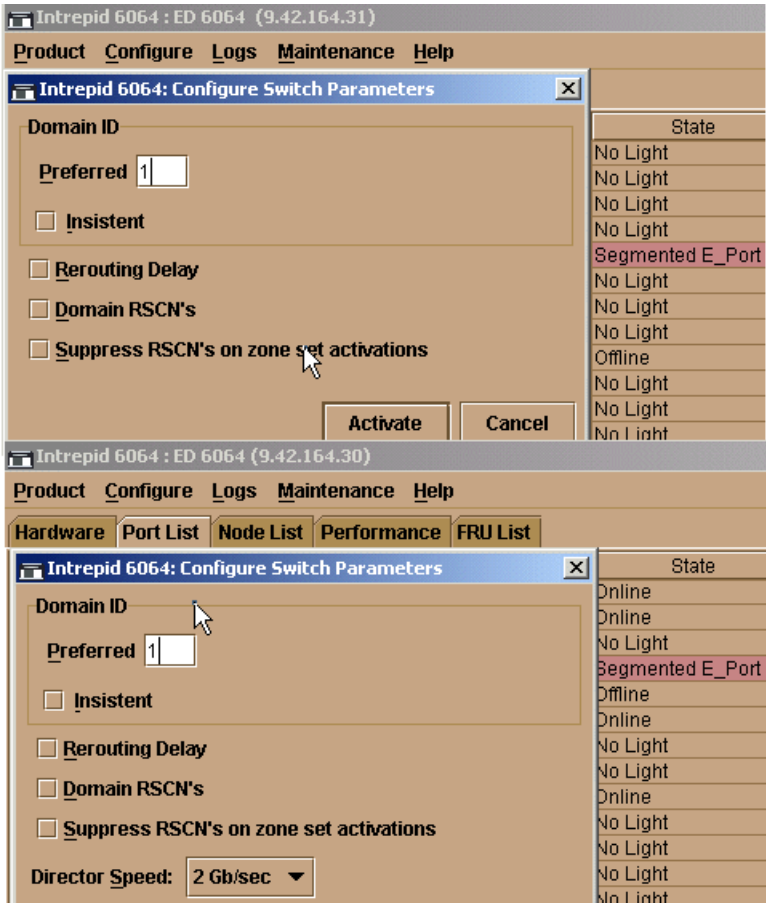


*Figure 4-67 Configure switch parameters*

The **Switch Parameters** under the **Operating Parameters** option are shown, and the window in the background lists the different IP addresses of the two ED-6064s. Both have been configured with a preferred **Domain ID** of **1**. The rest of the options are unchecked, except that the second ED-6064 supports 2 Gb/s speed.

In order to resolve the fabric segmenting problem, set one of the two ED-6064s to **Offline** mode and assign it a unique domain ID. This is done by selecting the **Configure —> Operating Parameters —> Switch Parameters** menu and changing the domain ID on ED-6064 (with an IP address of 9.42.164.31) to 5 and setting the switch to online mode.

The procedure to select **Set Online State** is shown in Figure 4-68.

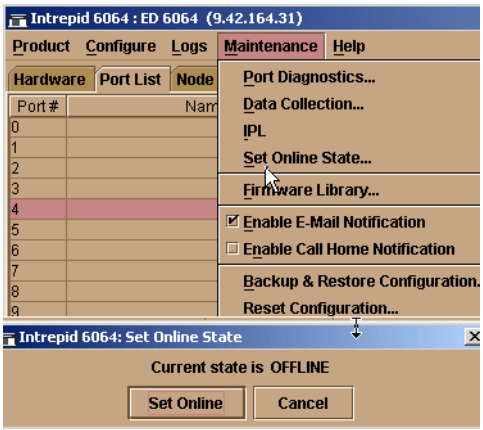


*Figure 4-68 Set Online State*

Set the switch to offline and then configure the **Preferred Domain ID** and click the **Activate** button as shown in Figure 4-69.

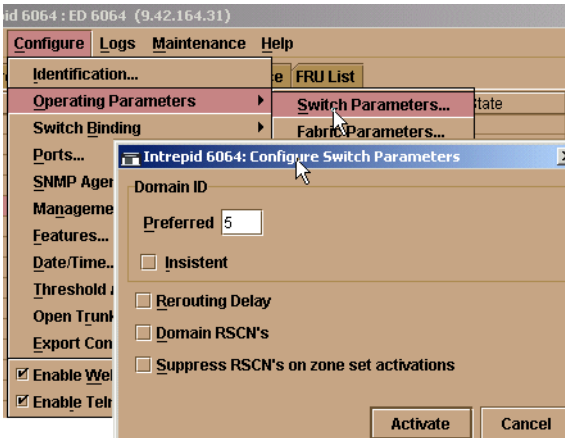


*Figure 4-69 Configure Switch Parameters menu*

The Preferred Domain ID assigned for ED-6064 (9.42.164.31) is now 5. Configure the switch online and verify if the fabric merge was successful from the **Fabric Topology** menu as shown in Figure 4-70.

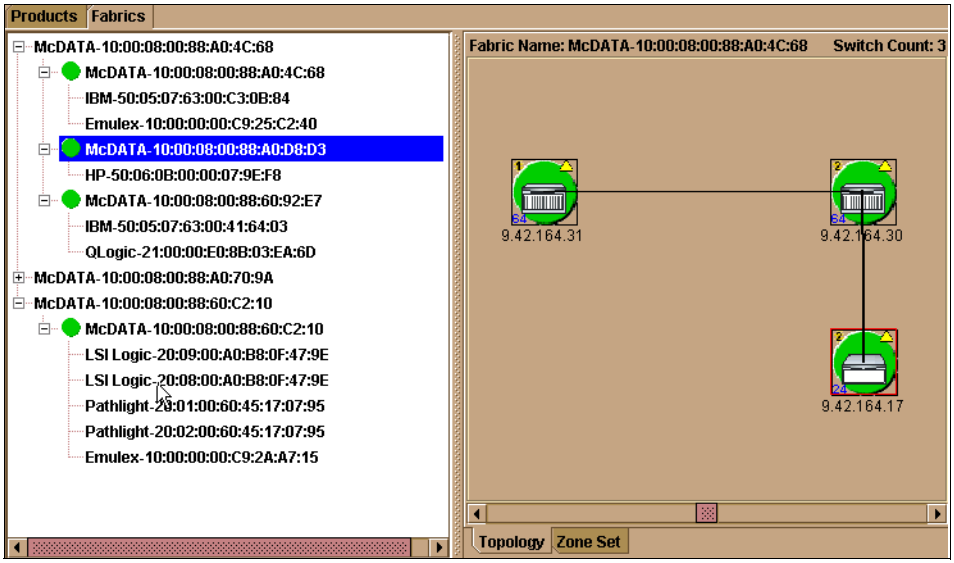


*Figure 4-70 ISL operational and fabrics merged*

The fabrics have successfully merged and are fully operational.

# 4.7 Understanding the McDATA zoning concepts

Fabric zoning is the most common mechanism being implemented in today's SANs to segregate the devices connected to the fabric. Zoning restricts the visibility and connectivity between devices connected to a single or multi-switch fabric.

## 4.7.1 Why we need zoning

In today's heterogeneous SANs, where AIX, Linux, Solaris, HP_UX and Windows hosts can be connected to the same fabric, and have LUNs configured on the same storage device, without zoning it is difficult to guarantee data integrity, security, high availability and fabric stability. For a comprehensive discussion on zoning and the concepts associated with it, refer to the IBM Redbook:

- *IBM SAN Survival Guide*, SG24-6143

## 4.7.2 Zoning implementation

There are different ways to implement zoning for a fabric. One such difference is the implementation of the various zoning definitions. For instance, zoning enforced through the name server table and the access to information about connected node ports, or through additional frame flow control is enforced by the route table in the switch.

### Soft zoning

McDATA uses name server zoning, which is implemented by authorizing or restricting access to name server information. The name server database on McDATA switches stores information on node WWNs and port numbers to identify the devices during the link initialization. The main purpose of the name server is to provide this information to attached node ports about the other node ports in the fabric. The attached node port does not need to probe every destination for information. Instead, it logs in with the name server and requests information on attached node ports.

With name server zoning enforced, the port that asks for information will only receive information about ports from within the same zone. This is also called soft zoning, because the name server enforces the zones, but there is no control of the real data flow. The name server cannot prevent a host communicating with the target bypassing the name server if it discovers the WWPN of the target device from a previous configuration, or if it has been hard configured by the end user (persistent binding is a good example of hard configuring the device address).

### Hard zoning

In contrast to soft zoning, hardware enforced zoning restricts the frame flow to zone members in a route table based in the switch hardware (ASIC). If a source port is not a member of the same zone as the destination port then the routing table for that port is disabled and communication between the two is denied at the entry port.

Hard zoning controls access at the ingress port. When a device attempts to communicate with a destination device outside of its zone by sending a **PLOGI**, the frame is blocked. A Class 2 frame will get fabric rejected, and a Class 3 frame will be dropped.

With 5.01.00 or later release of firmware, hard zoning is enabled by default. Hard zoning is enforced in the software and as such the user can configure a zone based on port WWN, port number, or both. The firmware upgrade from 4.x to 5.01.00 will automatically enforce hard zoning without any manual intervention.

**Restriction:** Hard zoning is not enforced on fabric loop (FL_Port) ports on the Sphereon ES-4500 Switch.

## 4.7.3 Zone member definitions

A zone member is specified either by the switch port number (and with it, the node ports connected to it), or by the WWPN of a node port, or by a mixed approach. Note that the WWNNs are not used for zoning definition.

### Zone member definition by WWPN

The major advantage with WWPN based zoning is that it provides the flexibility to move any device from one port to another port and it still remains the member of the same zone. The WWPN based zoning provides some diagnostic capabilities. For instance, to isolate a bad GBIC or HBA issue on the switch, the device can be connected to one of the spare ports on the switch just by moving the cable from the failing port to another good port, without making any changes to the active zone set.

Each WWPN can span multiple zones.

**Note:** The ESS can now be configured to administer the WWPNs of the ESS FC ports locally, which means they get their WWPN based on the locations in the ESS interface bays. So with zoning based on WWPNs, even in the case of the replacement of an ESS FC adapter, the WWPN does not change and therefore the zoning definitions do not have to be changed.

### Zone member definition by switch port number

Port based zoning is also known as static zoning. It consists of specifying the domain and the port number of the switch. Port based zoning allows greater control to the system administrator.

Another advantage of port based zoning is that a defective host bus adapter can be replaced and reconnected to the same port, even though it has a new WWPN, but the device can resume communication with other members in its zone without any modifications to the active zone set. A single port can also be a member of multiple zones.

### Mixing the two approaches

The two approaches to define FC node ports as zone members can be mixed. Node ports specified by their WWPN or switch ports specified by their number, can be members of more than one zone.

## 4.7.4 Zone management with zone sets

From within the McDATA Fabric Manager, we can specify up to 64 zone sets. A zone set consists of one or more zones that can be activated and deactivated at the same time. Each zone set can contain a maximum of 1024 zones and each zone can contain a maximum of 4096 members. Only one zone set can be active at one time. Activating an inactive zone will deactivate the currently active zone set.

**Restriction:** If all zone names are configured with 64 character names, the number of allowed zones in the zone set is limited to 758.

Node ports that are not configured in a zone within the active zone set are considered as members of the default zone (this takes up one of the 1024 maximum active zones). The default zone can be disabled independently from the active zone. Also, if no zone sets are activated all node ports are in the default zone. If the default zone is disabled while no zone set is active, no node ports can communicate. With the default zone enabled it is possible for all node ports in the default zone to communicate with each other in parallel to the currently active zone set.

There can be multiple zone sets configured for different tasks, for example, if we want to have certain node ports in the same zone for backup, but not during normal operation.

### Our zoning example

An example of how zones and zone sets are related is shown in Figure 4-71.

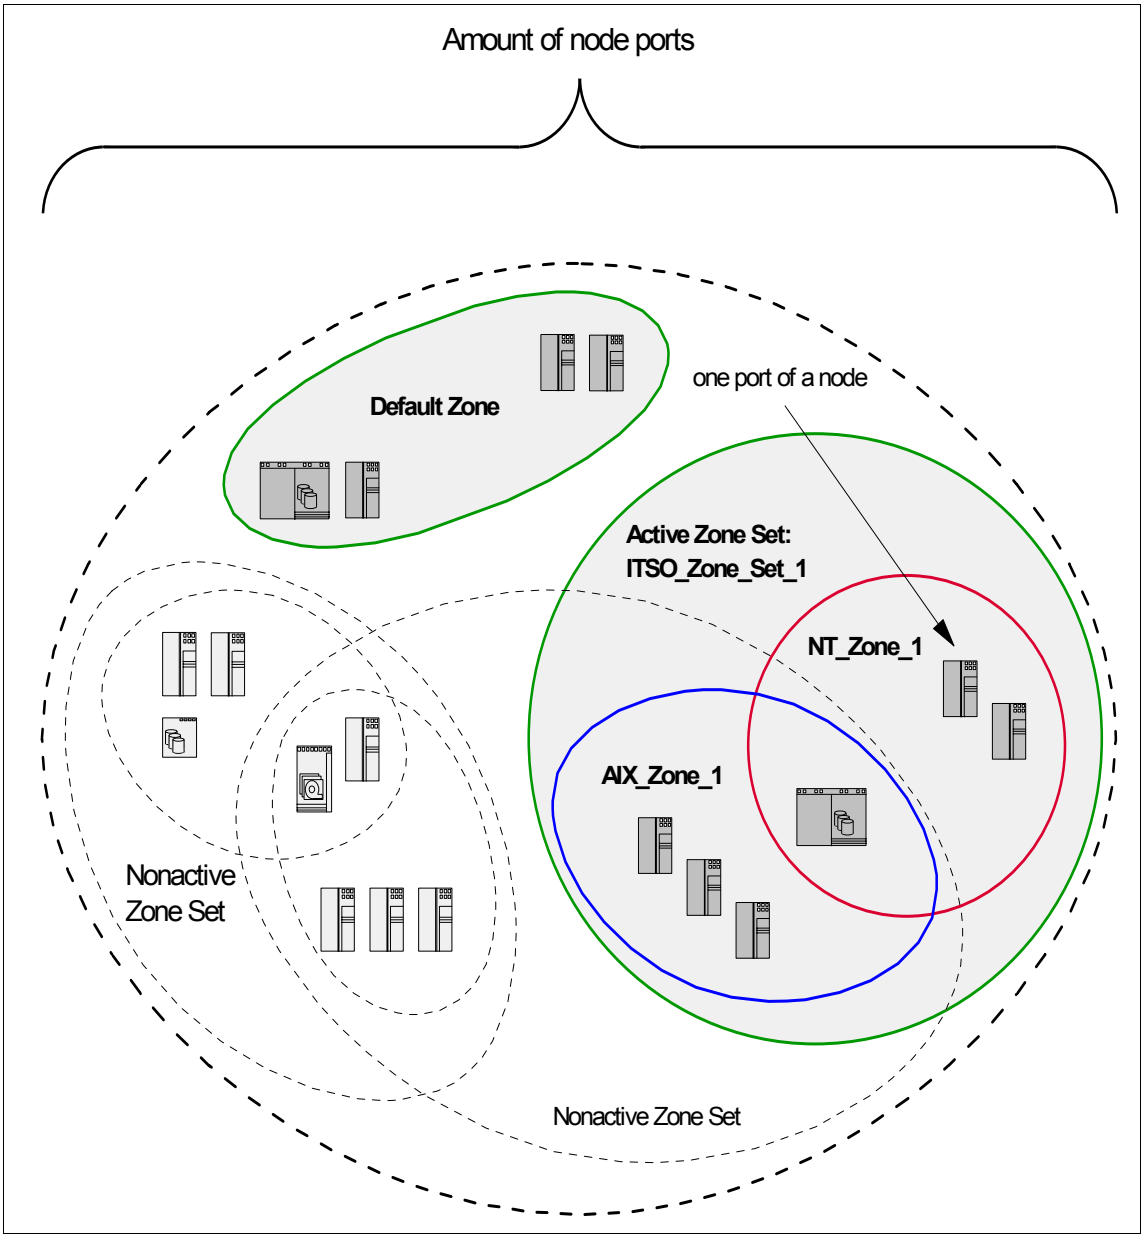


*Figure 4-71 Relationship of zone sets, zones, the default zone and node ports*

The node symbols here (from servers and from the ESS), represent one or more node ports and not necessarily the whole FC node with all ports. This is because zones with McDATA are built up with node ports. For example, all three ESS symbols could be ports of the same ESS.

The solid (blue, red, and purple) areas represent areas where traffic is permitted. The blue and the red zone represent the AIX and the NT zone to be defined in this topic. The green dotted line around the zones represents the active zone set.

The purple area is the default zone. In this example the default zone is enabled, which makes it possible for all node ports, which are not configured in a zone of the currently active zone set, to communicate with each other.

There might be cases where it is appropriate to disable the default zone — for example, if for management and security reasons, the only communicating node ports are those that are explicitly allowed. In this case, connecting node ports without defining them to a zone would prevent them from accessing other ports.

## 4.8 Managing the fabric

To view and manage the zones in the McDATA fabric, we must open the EFC Fabric Manager. The Fabric Manager is accessed by opening the Fabric View of the EFC Manager by selecting the Fabrics tab, as shown in Figure 4-72, and by clicking a fabric icon in the left side of the screen.

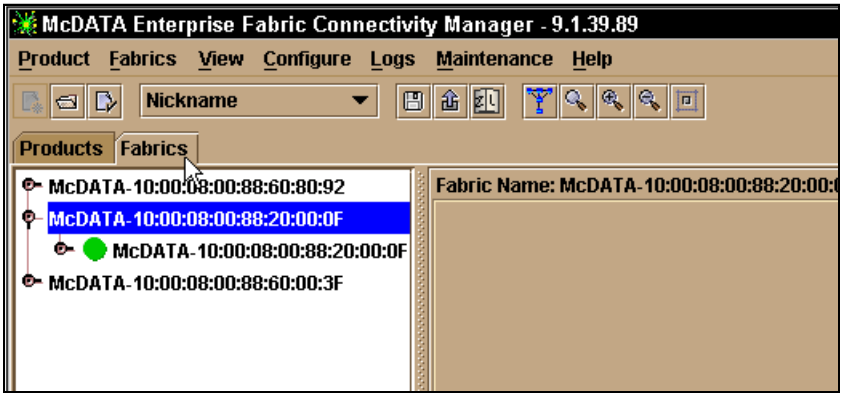


*Figure 4-72 EFC Manager Fabrics View*

The fabrics are listed on the left side of the view, and linked to the Fabric name are the products making up the fabric. Note that in our view we have three Fabrics, and we have selected the first fabric comprising of two products.

## 4.8.1 Using the Fabric Manager views

The EFC Fabric Manager opens as a different view by selecting the **Topology** Tab near the bottom of the window.

### The Topology View

In our case we have one switch installed in this fabric, as shown in Figure 4-73.

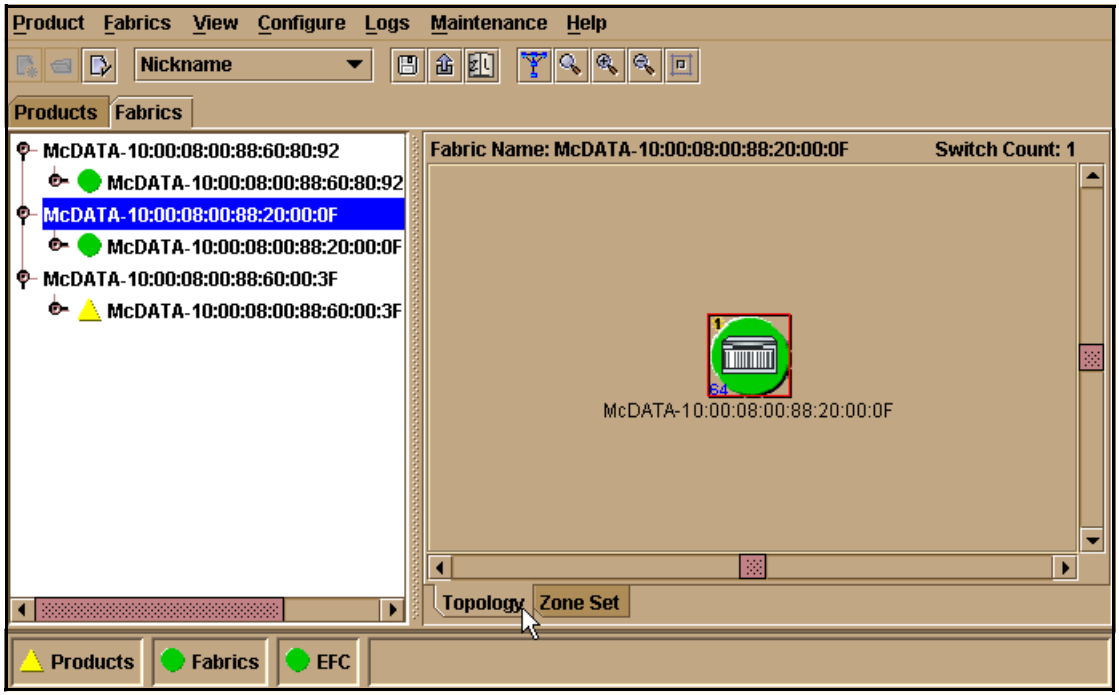


*Figure 4-73 Fabric Manager, Topology View with one device*

Our switches are not connected to each other, so the topology view shows three fabrics in the left window, each with one switch. We will see what this looks like in 4.9.3, “Setting up our zoned multi switch fabric” on page 553

### The Zoning View

To change to the Zoning View, we select the **Zone set** Tab as shown in Figure 4-74, and we then see the Zoning View of the Fabric Manager.

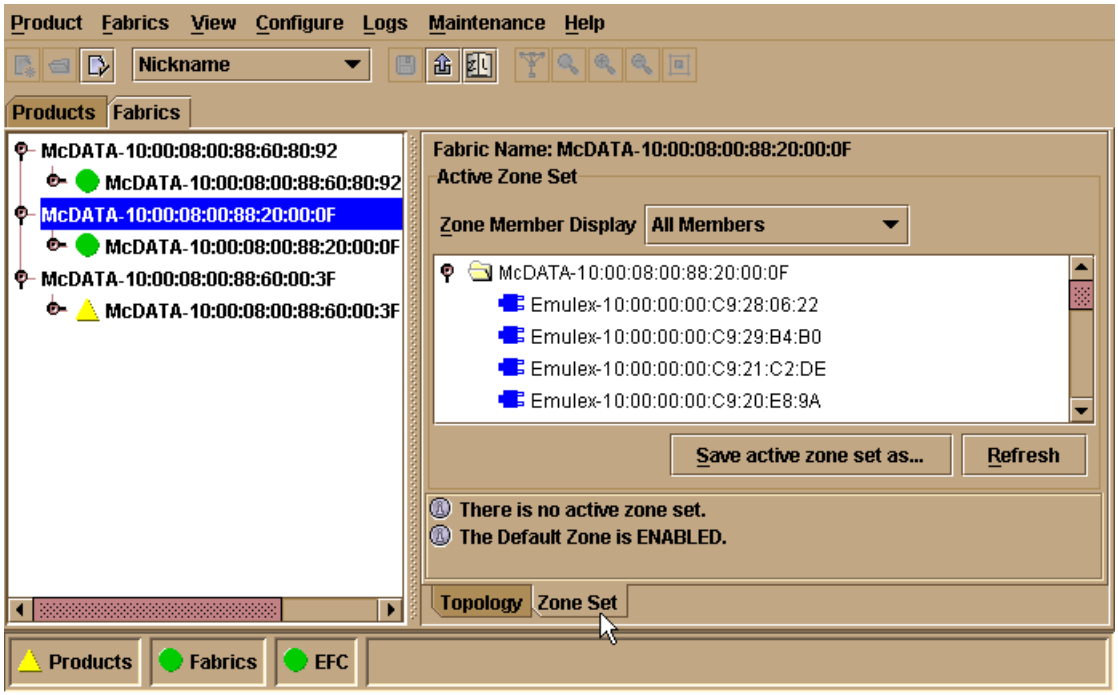


*Figure 4-74 Fabric Manager Zoning View no zone set active*

The main window will display the active zone set, but at this time, there are no active zone sets. From here the active zone set can be saved with another name. Also, we can see if the default zone is enabled or not.

In the topics that follow, we show how we set up zoning in our environment. For these tasks it is not important if we are in Topology view or Zone Set view.

## 4.8.2 Zones, zone sets, and zoning

As an example, we want to have one ESS and two Netfinity® servers defined in a zone, which is shown in Figure 4-75.

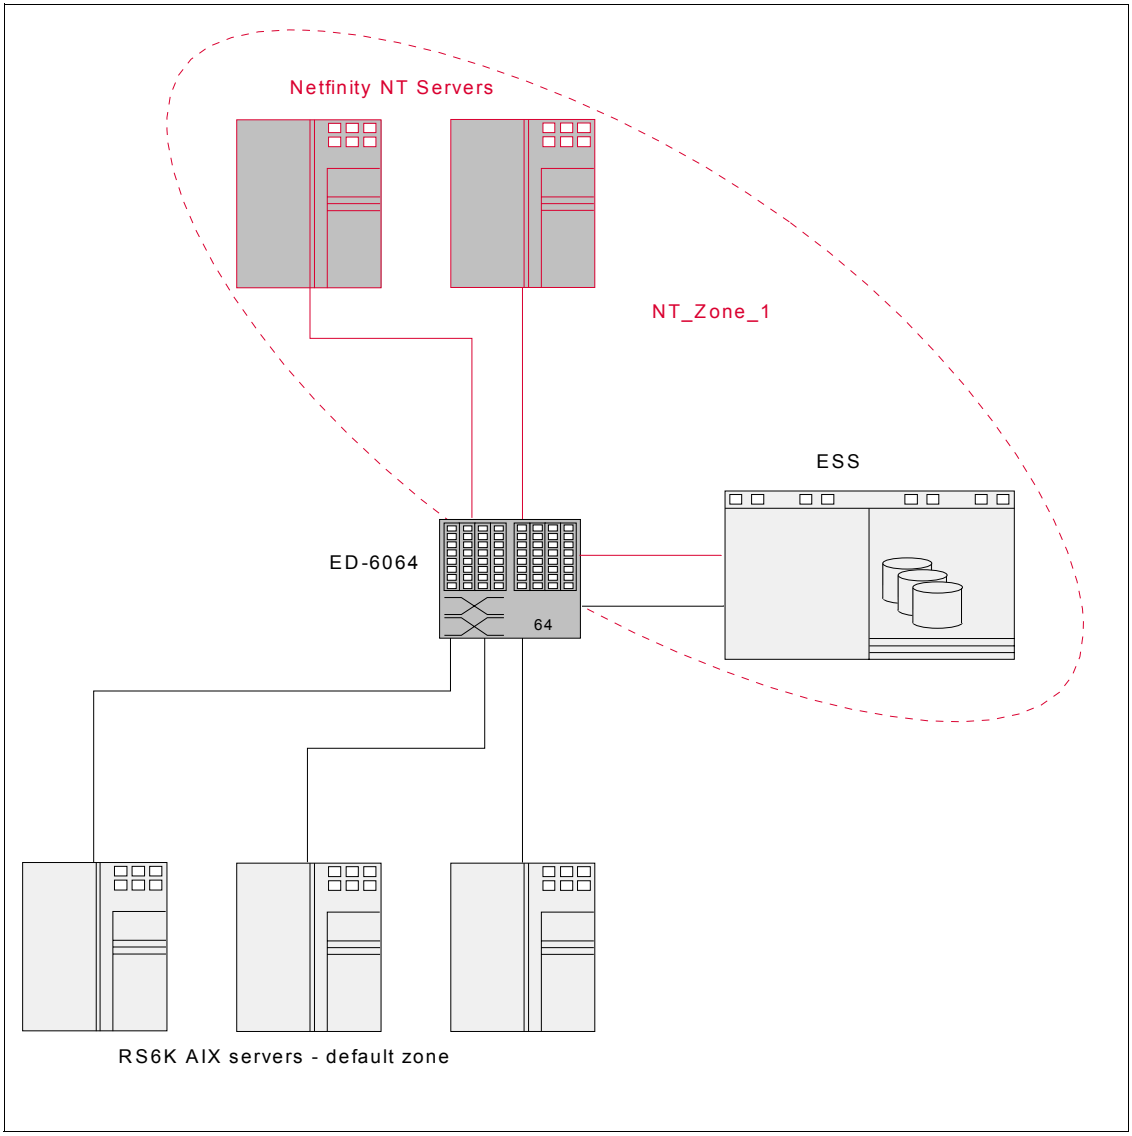


*Figure 4-75 NT zone with two Netfinity node ports and one ESS node port*

There are ESS ports and RS/6000 ports already connected to the switch, but no Netfinity port. In our case we have only one FC adapter on each host, so we only have to add one port of each host to the zones.

## Creating a new zone set

To create a new zone set we select **Configure —> Zoning Library...** This displays the zoning library window and provides us with options for changing our zone set definitions, for example, creating new zone sets, deleting zone sets, and modifying or renaming existing zone sets, as shown in Figure 4-76.

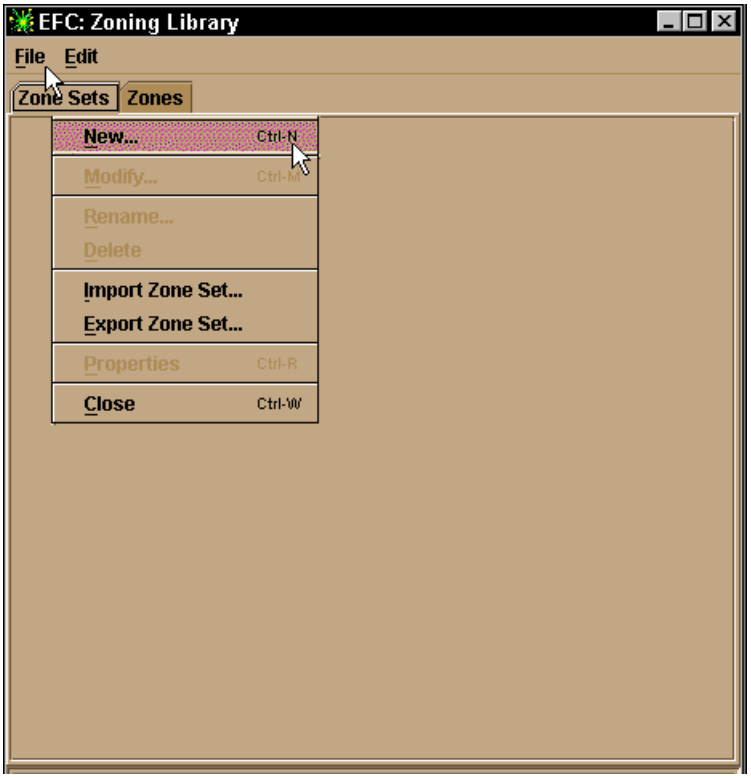


*Figure 4-76 Fabric Manager: Zone Sets menu*

Because there are no zone sets in the library, we will need to create one. We try to start our selection from **File —> New...** However, as we have not yet defined any zones to put into a zone set, we are presented with the error message shown in Figure 4-77.

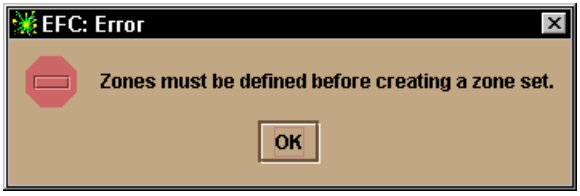


*Figure 4-77 Error adding zone set with no zones*

## Creating a new zone

Therefore, we have to go on and create at least one zone, and in our case, we will create our NT zone by adding some members.

## Adding members to the zone

By clicking the **Zones** tab, and then **File —> New...** we are presented with the New Zone window. At the top of the window we can enter a name for the zone. On the left we have an Attached ports/nodes column, where we can view all of the WWPNs or nicknames of the connected FC ports. Here we need to be very careful, in a multi-fabric environment, that we choose the correct fabric we wish to work on.

This can be selected by clicking the drop-down menu, as shown in Figure 4-97, and we can choose between a Nickname view or WWPN view by clicking the **Display Options...** button. Then we need to expand the folder by double-clicking the folder to view the complete list of WWPNs for that domain.

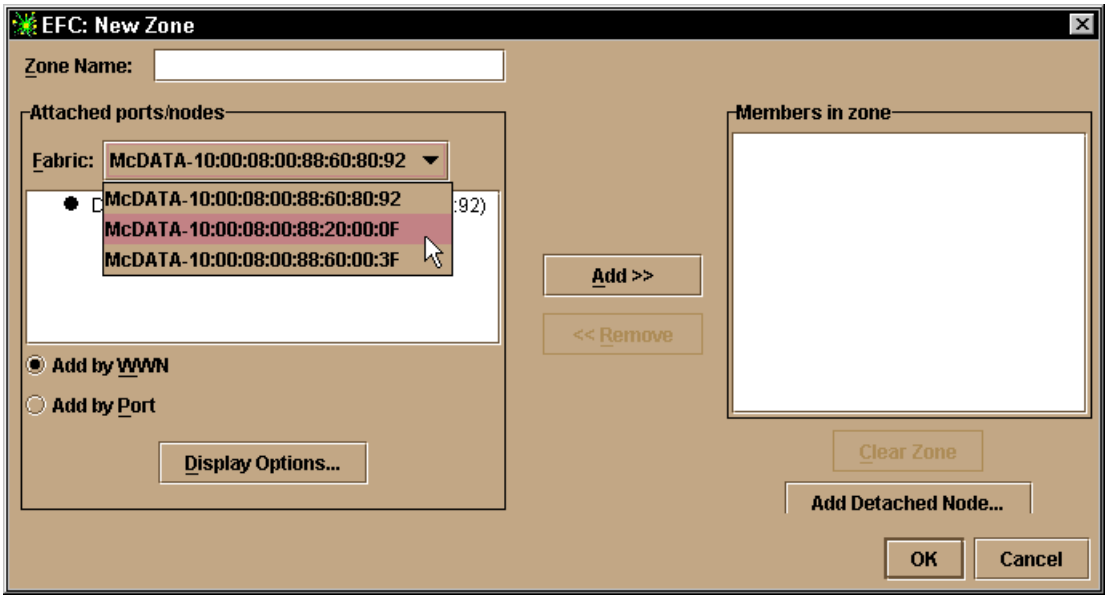


*Figure 4-78 Fabric Manager: New Zone*

Below the left column there are two check boxes. Here we can choose if we want to assign a switch port and all connected node ports to a zone (hard zoning), or if we want to specify the node ports based on their WWPN (soft zoning).

To the right is where we will add the members for the zone. Here, the nicknames become important.

In our example, we chose to add the two Netfinity ports and one ESS port to a zone. To assign one ESS port to the zone, we drag-and-drop the WWPN or nickname associated with the ESS port to the left part of the window or highlight the port on the left and click the **Add>>** button, as shown in Figure 4-79.

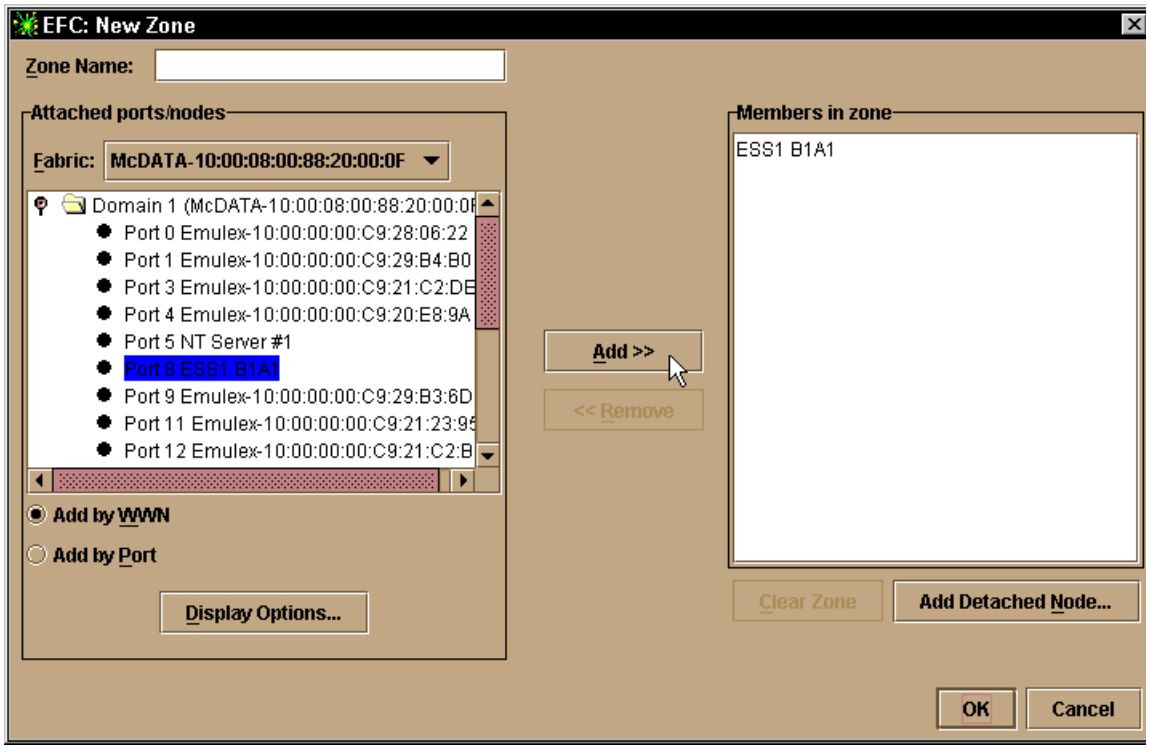


*Figure 4-79 Fabric Manager: Zone definition*

## Adding members by domain and switch port number

We could also add attached or detached node ports to the zone by selecting the **Add by Port** radio button and choosing the domain and switch port they are connected to, or are going to be connected to. We now illustrate that we could specify the members in one zone based on the WWPN, and other members based on the switch port they are connected to, as shown in Figure 4-80.

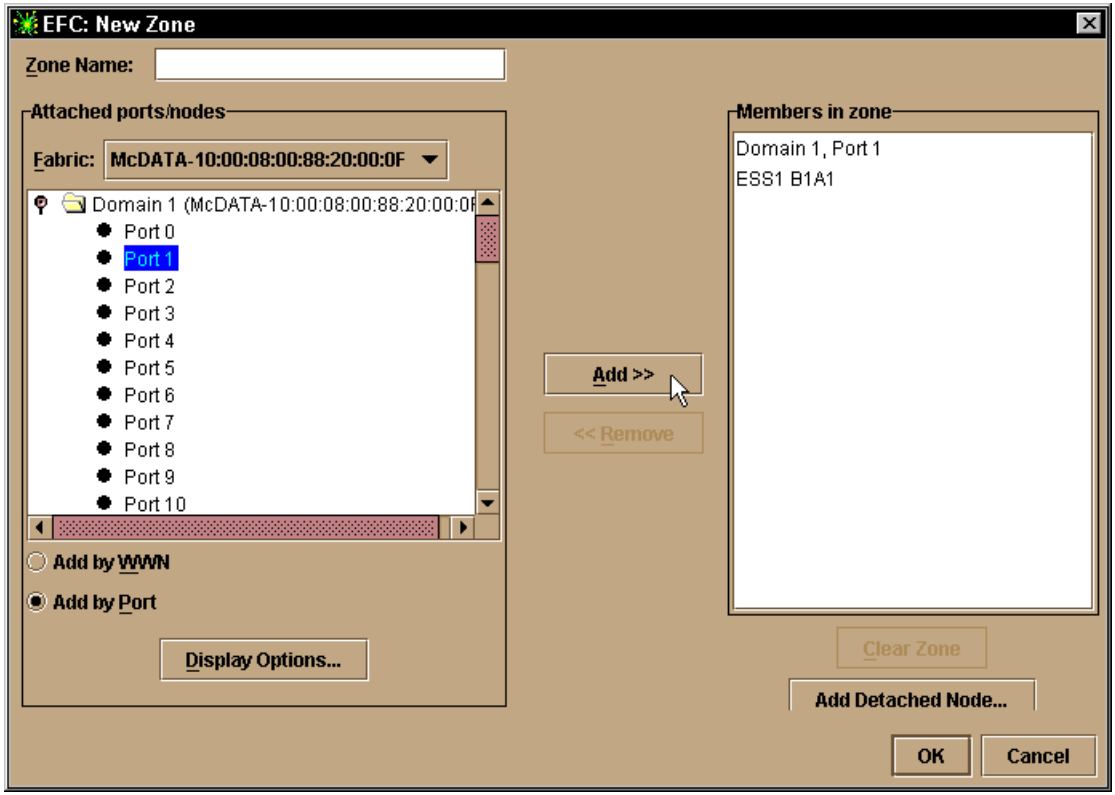


*Figure 4-80 Fabric Manager: Add by port number*

## Adding detached node ports to the zone

We are also able to define a detached Netfinity host port to the zone by using this window. This can be done by specifying the WWPN or nickname of the ports and adding them to the zone.

To perform this task, we click the **Add Detached Node...** button. We get a data entry field where we can insert the WWPN of the port, as shown in Figure 4-81.

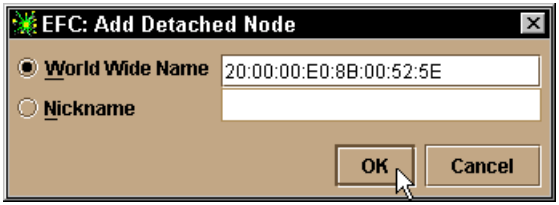


*Figure 4-81 Fabric Manager: Add Detached Node*

The nodes become restricted to the zone in which we have defined them when they are connected. We choose to specify the WWPN to define them to the zone because we want them to be members of the zone regardless of which switch port they are going to be connected to. This will also make it easier should we need to rearrange the cabling at any stage or in the event of a failure.

## Saving the zone

While trying to save the zone using the **OK** button, we got an error message. This is because of we have spaces in the zone name, as shown in Figure 4-82.

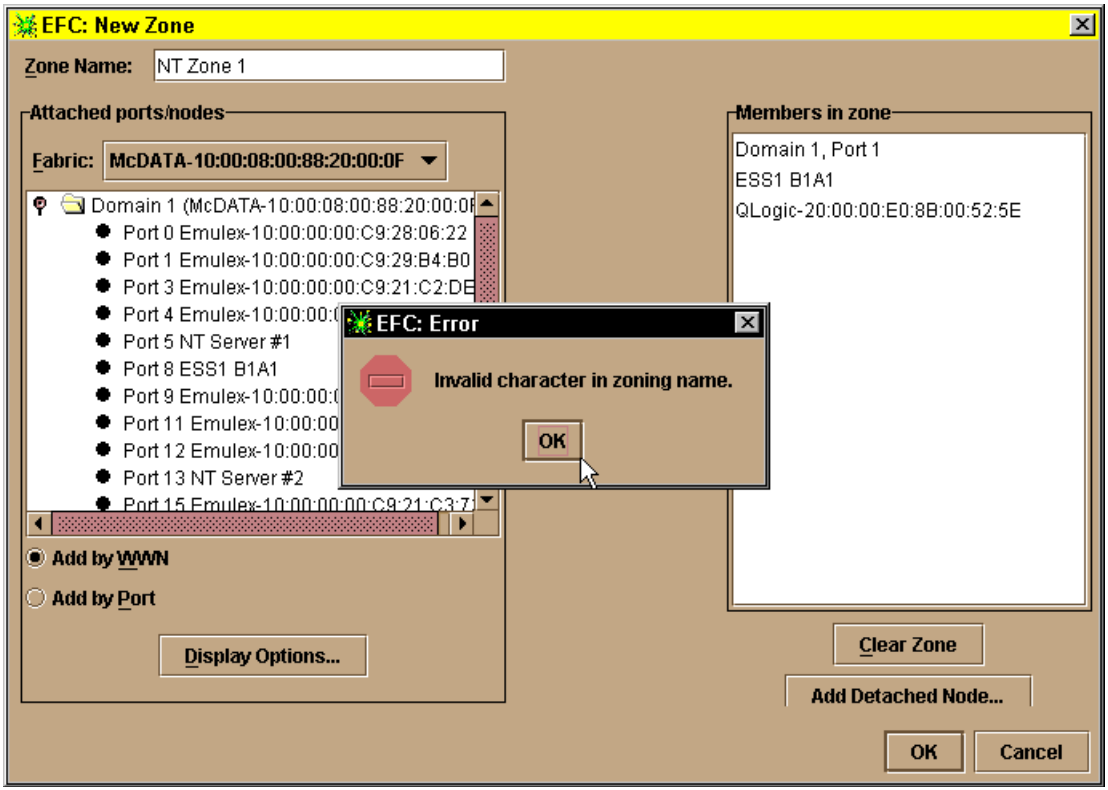


*Figure 4-82 Fabric Manager: Incorrect zone name*

This is not allowed, so to circumvent this problem we changed it to NT_Zone_1. After saving the zone, we return to the Zone Set window where we can view the zones in the zone library.

To view the zone members, highlight the zone in the Zoning library and then go to **File —> Properties** as shown in Figure 4-83.

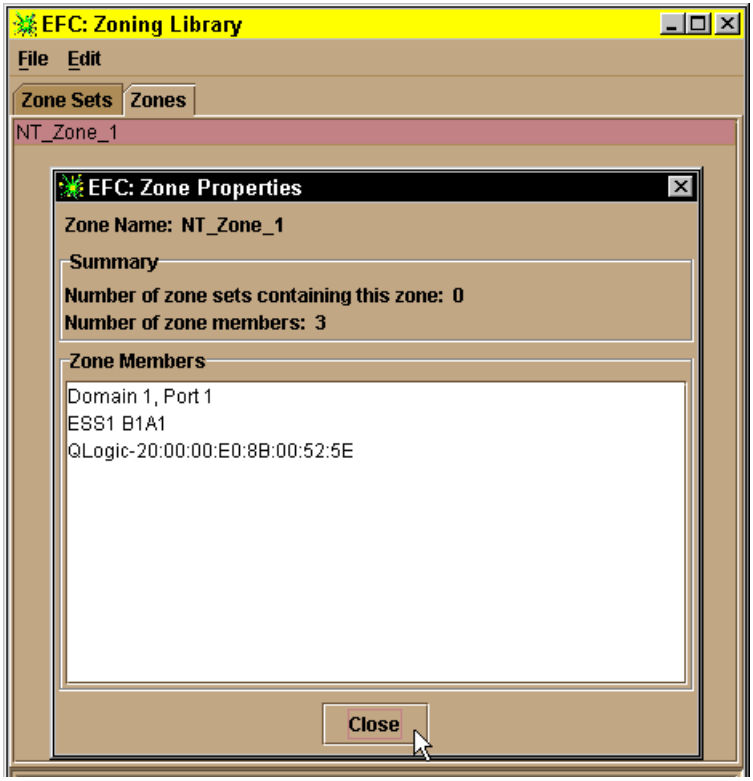


*Figure 4-83 Fabric Manager: View Zone Properties*

## Assigning the zone to a zone set and saving the zone set

To assign the newly created zone, NT_Zone_1, in to a zone set, we select the **Zone Sets** tab and then **File —> New...** opening a New Zone Set window, shown in Figure 4-84. From here we will drag-and-drop the zone from the *zones in library* to the *Zones in Zone Set* window on the right-hand side of the window.

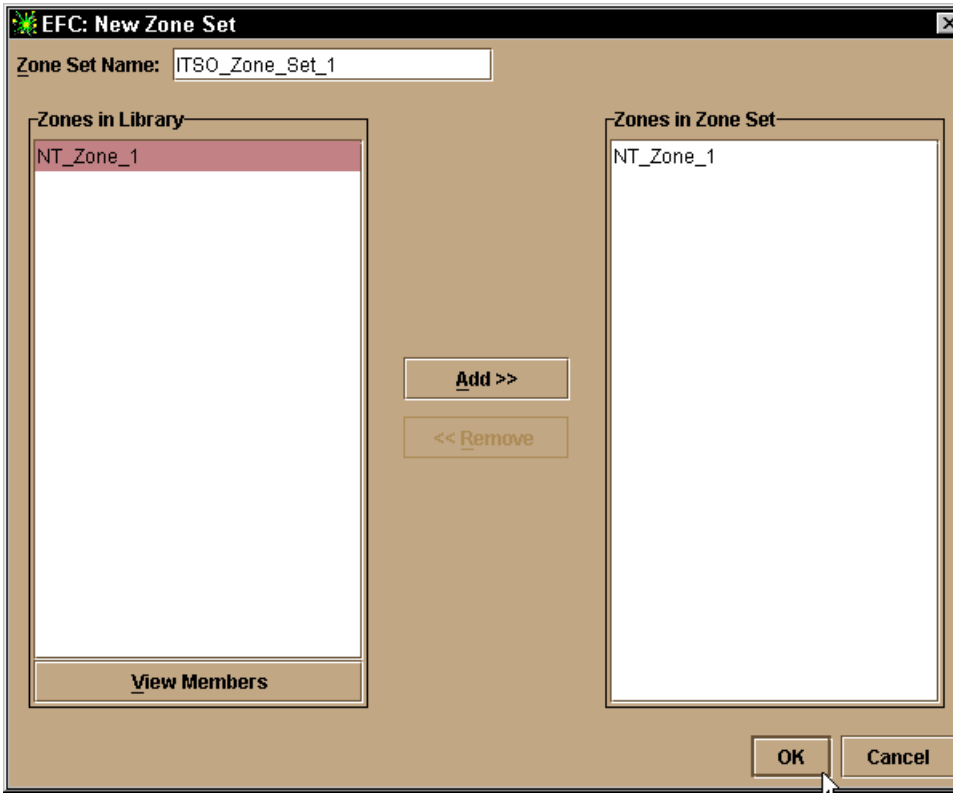


*Figure 4-84 Fabric Manager: Assigning zone to zone set*

Now we save the zone set as a member in the Zone Set Library by giving the Zone Set a name and clicking the **OK** button. The Zone Set Library window now looks like that shown in Figure 4-85.

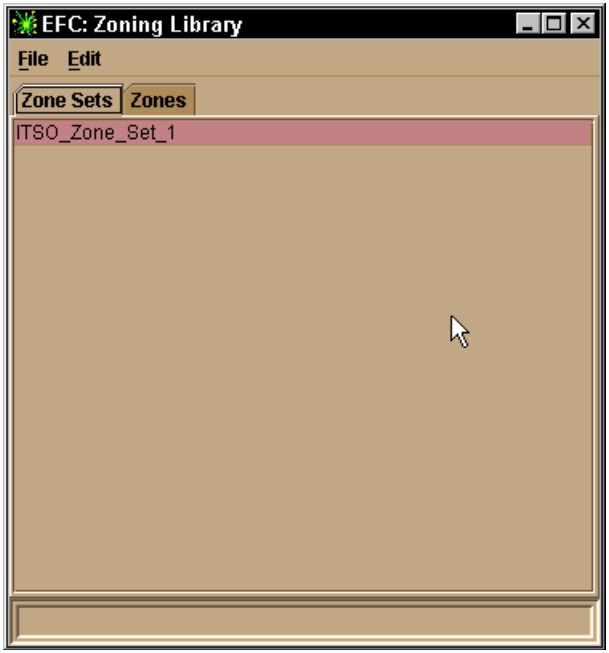


*Figure 4-85 One zone set defined*

We have one zone set (ITSO_Zone_Set_1) with one zone (NT_Zone_1). We have one ESS port and two Netfinity host ports within the zone, one Netfinity port with an Emulex adapter, and one with a QLogic adapter.

The ITSO_Zone_Set_1 zone set, with NT_Zone_1 zone, is now saved. However, it is not active yet, but can be activated to make the fabric zoned. We can now also define more zones in the zone set or create other zone sets.

## Activating the zone set and making the fabric zoned

To finish our zoning example we will activate the zone set now. This is done using the **Configure —> Activate Zone Set...**. With this action we are given a window where we can choose which Zone Set we would like to activate. We highlight the ITSO_Zone_Set_1 and click **Next**. If we have modified an existing zone set and are activating the same zone set, we are given a window displaying what changes are about to be made by the activation. We confirm our changes and click **Next**. Now we are given a confirmation box as shown in Figure 4-86, and after confirming the details, we click **Next**.

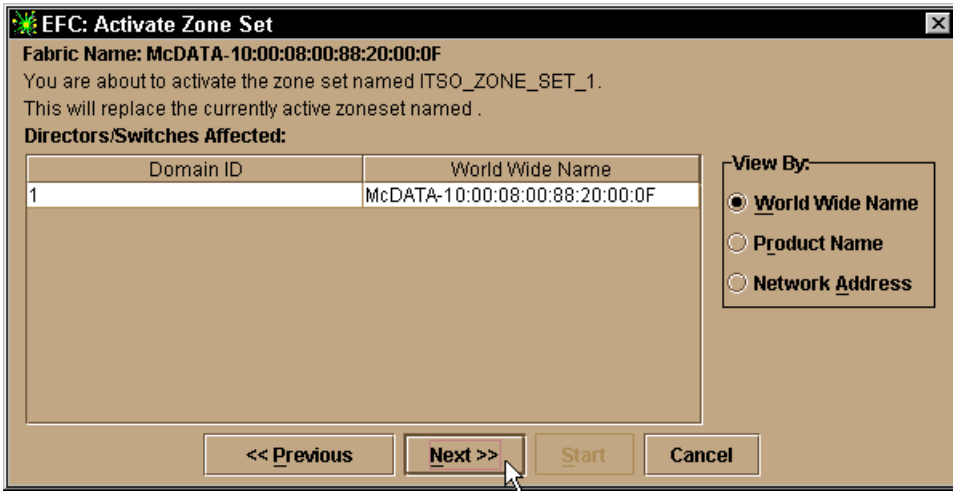


*Figure 4-86 Activate Zone Set confirmation*

Lastly, we get a message window prompting for us to click the **Start** button. After a few progress messages, we receive a message that the activation is complete, as shown in Figure 4-87.

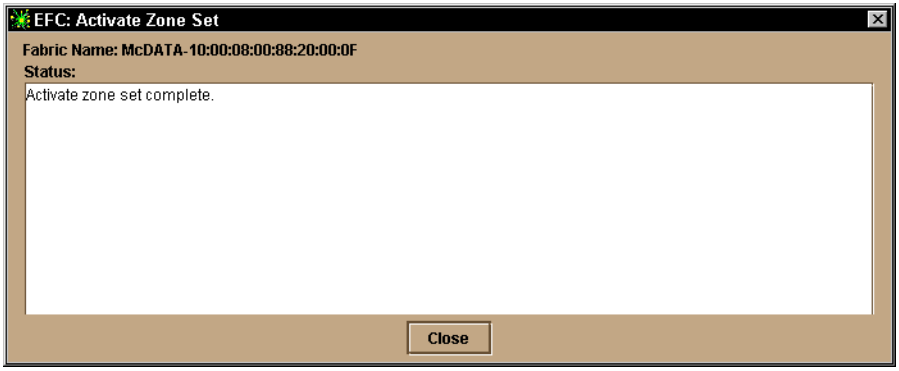


*Figure 4-87 Fabric Manager: Zone set activated*

## Viewing the active zoning configuration

After returning to the Zoning View window and expanding the zone set by clicking the small symbol to the left of ITSO_Zone_Set_1, and then NT_Zone_1, the Zoning View of our Fabric Manager looks like that shown in Figure 4-88.

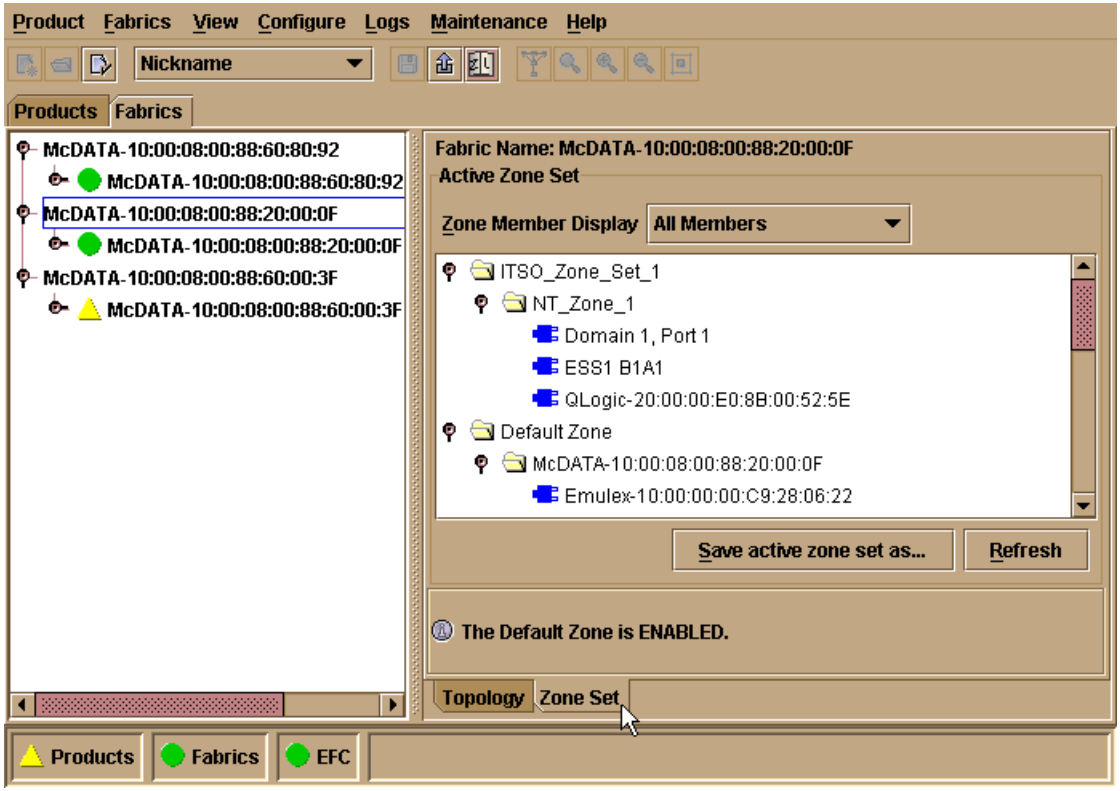


*Figure 4-88 Fabric Manager: Active Zone Set with one zone shown*

We can see the zone set and the associated zone by its name and the configured host ports by their manufacturer's name and the host adapter ports associated WWPN. Also, we see the nickname for the ESS port which we configured earlier. If we had not configured a nickname for the ESS port, we would only see it as another Emulex adapter port, as these are used as the Fibre Channel adapters in the ESS.

### Modifying zone sets

From within the Zoning View window, we can also manipulate the zone sets, for example, deactivating a zone set or saving the zone set. As an example, by using the **Save active zone set as...** button, we can copy the same zone set, but assign it a different name.

## 4.8.3 Adding an AIX zone to the existing zone set

We also have AIX hosts, and we want to add a zone with the AIX systems and another ESS, as shown in Figure 4-89.

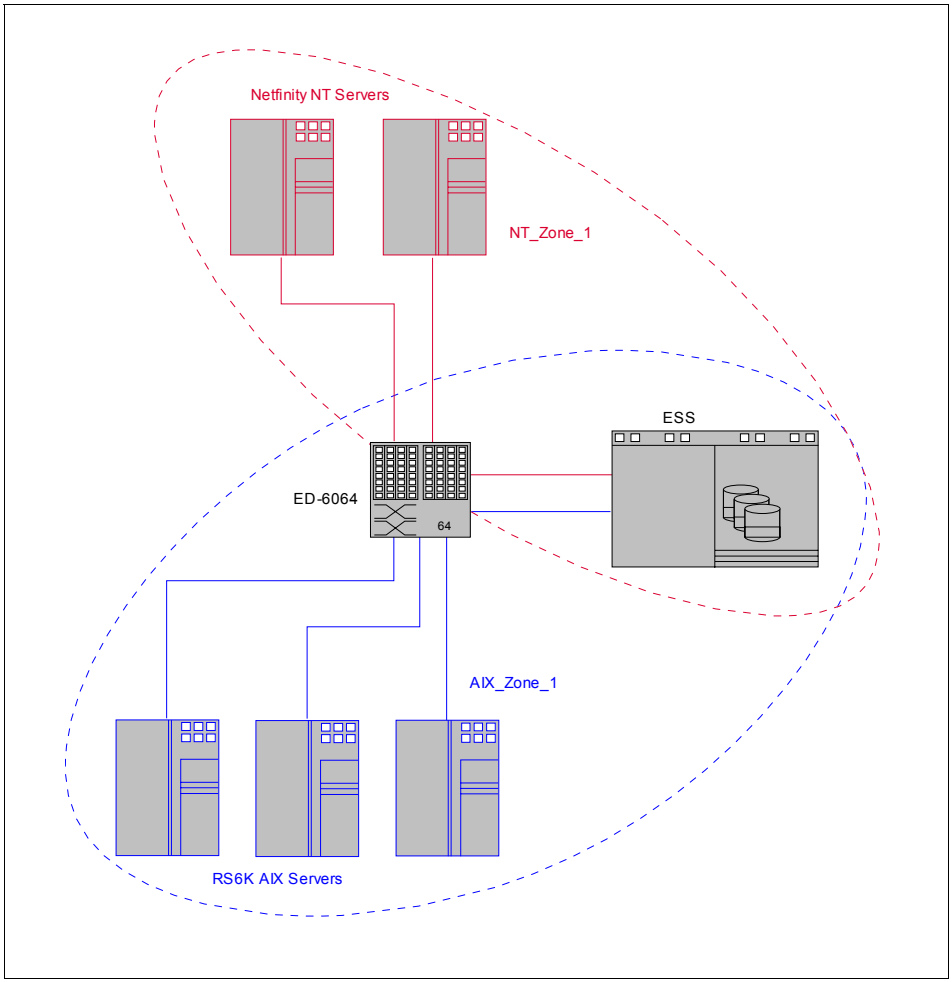


*Figure 4-89 NT zone with three node ports and AIX zone with four node ports*

Our AIX hosts are already connected to our ED-6064, so all we need to do is to define the zone in the zone set which already includes our NT zone.

First we need to create the new zone, and as we did so previously, navigate to the zoning library, select the **Zones** tab and click **File —> New...**, which opens the New Zone window.

We drag-and-drop the adapters of our AIX hosts to the Members in zone window, as shown in Figure 4-90. We have not assigned nicknames to our AIX hosts and so we must be familiar with the WWN of each adapter.

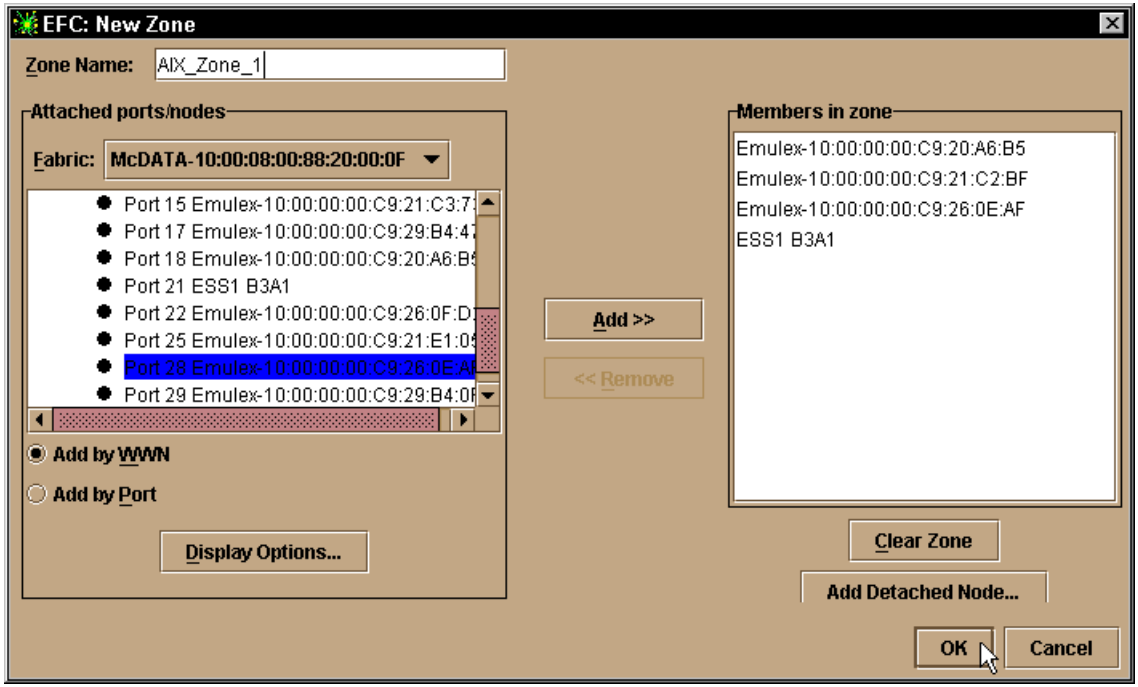


*Figure 4-90 Fabric Manager: Defining an AIX zone*

Then we need to add the new zone to our zone set, by selecting the **Zone Sets** tab, then right-clicking our target zone set, 'ITSO_Zone_Set_1', we select **Modify...** from the context menu, as shown in Figure 4-91, which allows us to add or remove zones in the zone set.

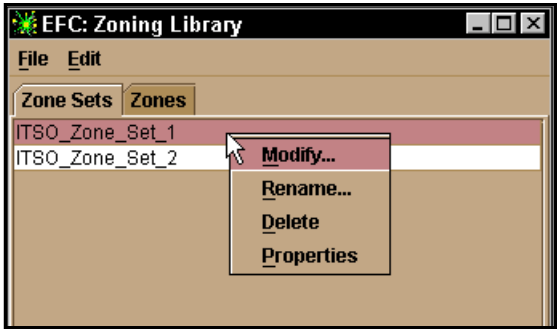


*Figure 4-91 Fabric Manager: Modify selected zone set*

From the Modify Zone Set window, just as we did previously with the NT zone, we select our newly created AIX_Zone_1, from the left column and drag-and-drop to the Zones in Zone Set window on the right, as shown in Figure 4-92.

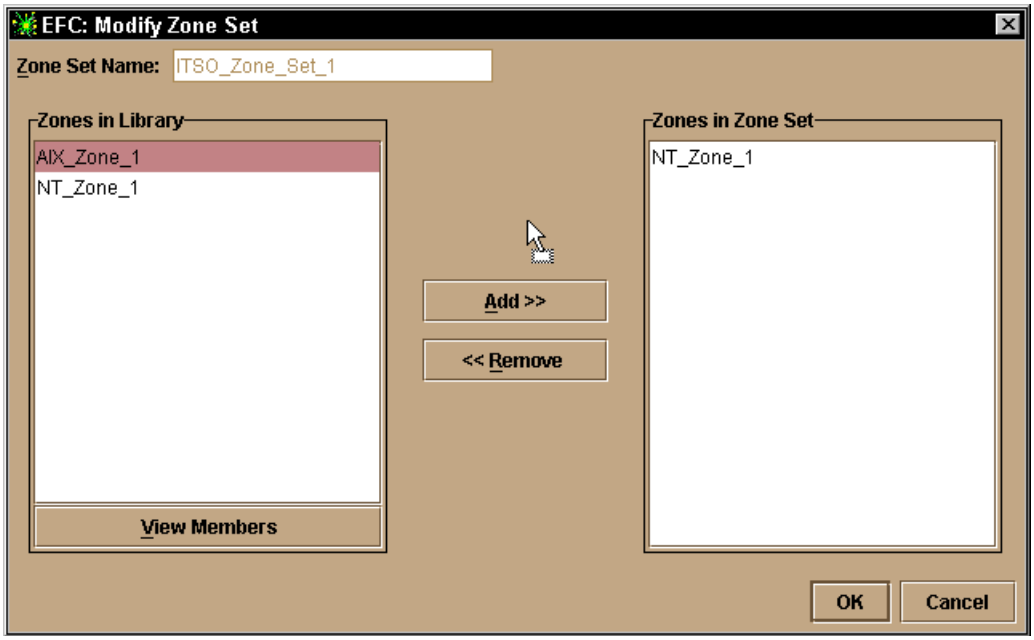


*Figure 4-92 Fabric Manager: Dragging the AIX zone to the zone set*

After we have added AIX_Zone_1 to the Zones in Zone Set column, we save the updated zone set by selecting the **OK** button.

After saving the zone set, we have two zones, NT_Zone_1 and AIX_Zone_1. Both of these are in our zone set ITSO_Zone_Set_1 as shown in Figure 4-93.

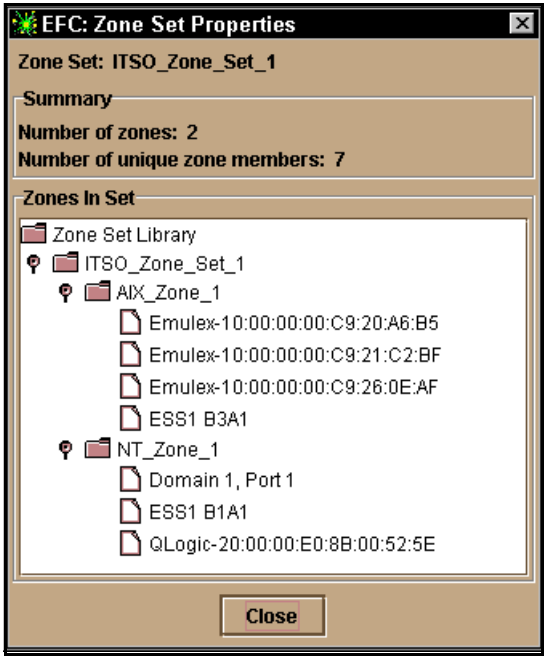


*Figure 4-93 Two zones in one zone set*

To apply the changes to the fabric, we must activate the zone set again. This is performed in the same way as when we activated the zone set with a single zone earlier. By exiting from the Zoning Library, we first ensure that we are looking at the fabric view and that we have highlighted the correct fabric in the left column. Then we select **Configure —> Activate Zone Set** as shown in Figure 4-94.

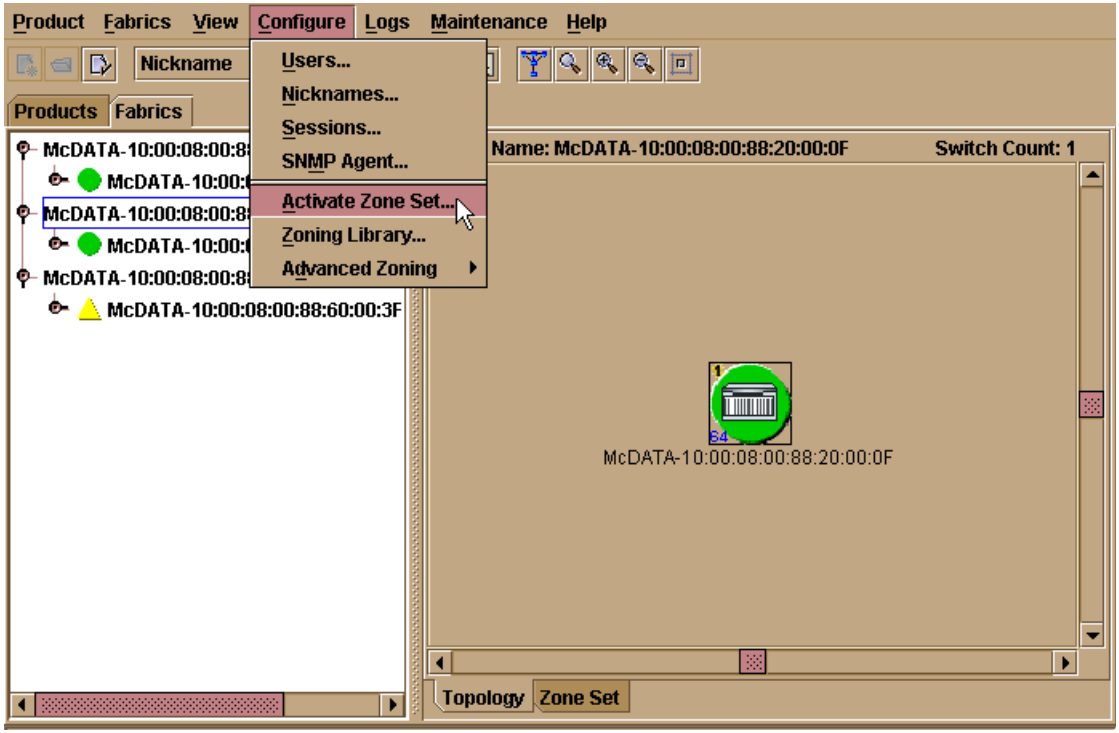


*Figure 4-94 Selecting Activate Zone Set for a Fabric*

In the selection window that appears, we highlight ITSO_Zone_Set_1 and click **Next**, as shown in Figure 4-95.

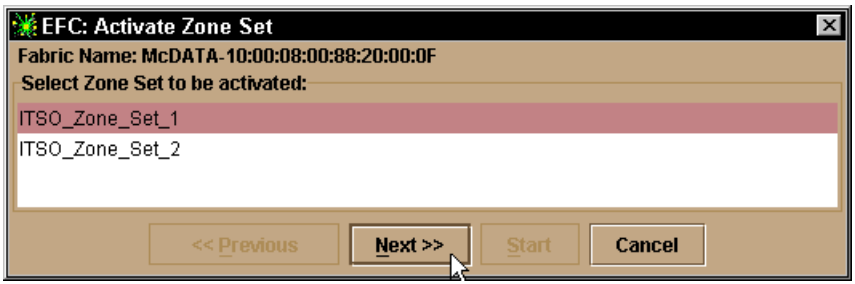


*Figure 4-95 Activate Zone Set selection window*

The activation procedure then displays the differences from the currently active zone set we are making, as shown in Figure 4-96.

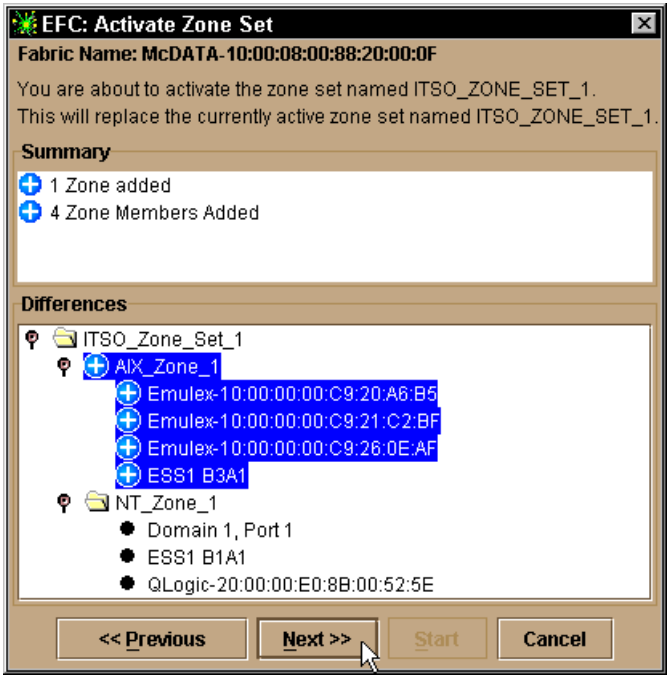


*Figure 4-96 Differences from currently active zone set display*

After we have confirmed that the changes are what we expect, the next step is to confirm that we are activating this zone set on the correct Fabric.

The window shown in Figure 4-97 displays the domain ID and the WWN of the switch that we are targeting for activation. It is wise to double check that these details are correct if you are running a multi-fabric environment.

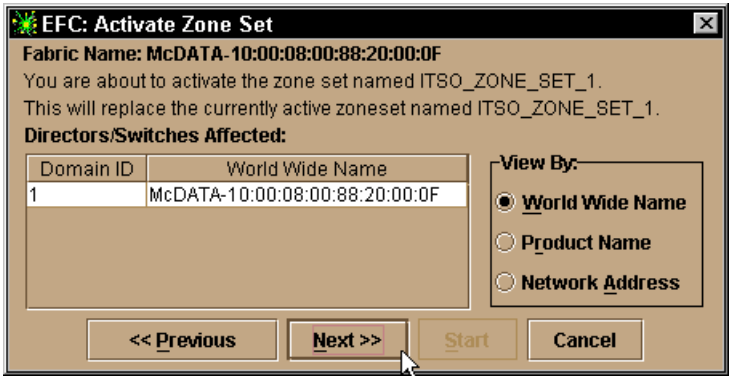


*Figure 4-97 Zone set activation, confirming the correct fabric*

After ensuring the Fabric is correct we click **Next**>> and then **Start** to begin the Zone set activation. On successful completion our changes are reflected in the Fabrics, Zone Set view, which looks like that shown in Figure 4-98.

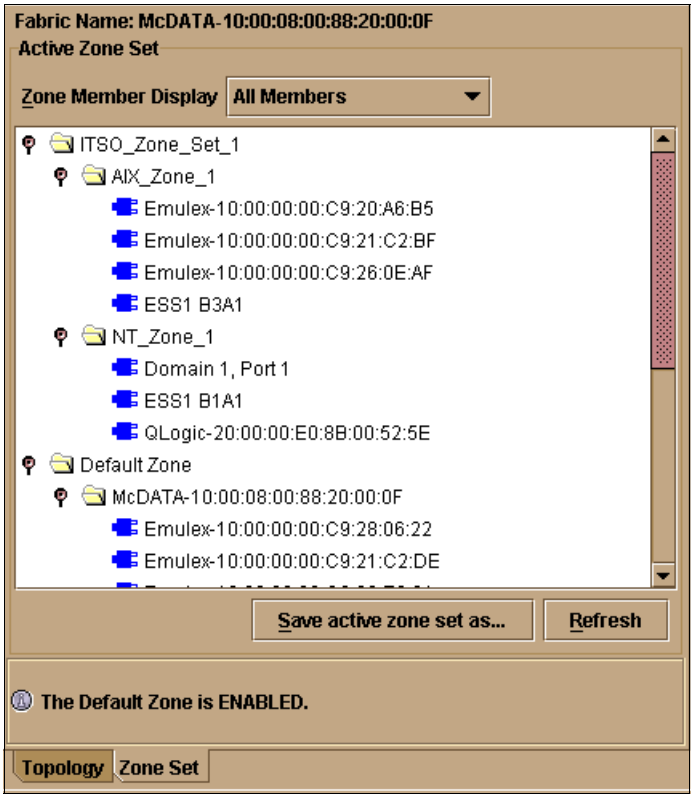


*Figure 4-98 Fabric Manager with two zones shown*

We have now successfully created and activated a zone set with two zones.

## 4.9 Building a multi-switch fabric

The focus of McDATA is highly available connectivity in a data-centric approach where the director is the core for connectivity of nodes. However, there is also a need for core-to-edge connectivity, and this is provided with the E_Port capability of the McDATA G_Ports.

### 4.9.1 Multi switch fabric considerations

The planning of multi switch fabrics depends on many things. Are we going to have a local SAN in one site with up to 64 node ports connected? Then we might not consider cascading our switches. If we want to build a SAN to connect two sites together, or if we want to have more ports in a single fabric, cascading becomes a valued commodity. Also, if we want to extend the SAN to provide departmental user groups with access to centralized storage devices or to establish a centralized backup which does not affect the LAN, cascading becomes a necessity.

Nevertheless, we still might think about whether or not, or to what extent, we want to cascade switches. The reason for this is that by using E_Ports we will sacrifice F_Ports. Also, with an extended fabric, the ISLs can possibly become a bottleneck. This will lead to the use of more ISLs which means even fewer F_Ports available for the attachment of devices. That which seems easy, in the first instance, can get more complicated once we add the zoning concept, load balancing, and any bandwidth issues that may appear.

#### Examples for multi switch fabric solutions

There are many solutions which are only possible by using a multi switch fabric. For example, disaster tolerant solutions that are using a SAN can be built upon a McDATA SAN but only when connecting two sites. We need switches at both sites to back up one site completely.

Disaster tolerance *and* high availability of the host systems and the storage can be established together using a multi switch fabric, and open system hosts using Logical Volume Manager (LVM) mirroring together with clustering software, such as HACMP for AIX or Veritas Cluster Server. To further extend the availability, two footprints (parallel independent fabrics) could be used.

Building upon the disaster tolerant and highly available approach, the SAN can be extended to become a core-to-edge approach, especially if more hosts in the company need access to the SAN — for example, for storage consolidation where there is a need to provide access to distributed hosts to resources in the data center. For hosts that do not need the RAS or bandwidth provided by directors, the McDATA switches connected to the directors serve as connectivity to the SAN backbone.

This is useful if a company wants to get rid of those hundreds of smaller departmental servers (for example, file servers). Disk consolidation which is possible with a corporate-wide SAN can be seen as the first step for server consolidation. Of course, just the connectivity of user groups to centralized disk storage does not replace a file serving solution.

McDATA directors and the 16-port and 32-port switches do not support loop devices directly, but by extending the fabric with the loop switch, this makes the attachment of legacy loop only devices and loop only tapes possible.

## 4.9.2 Solutions for high availability and disaster tolerance

An example of a solution that provides high availability with disaster tolerance is shown in Figure 4-99.

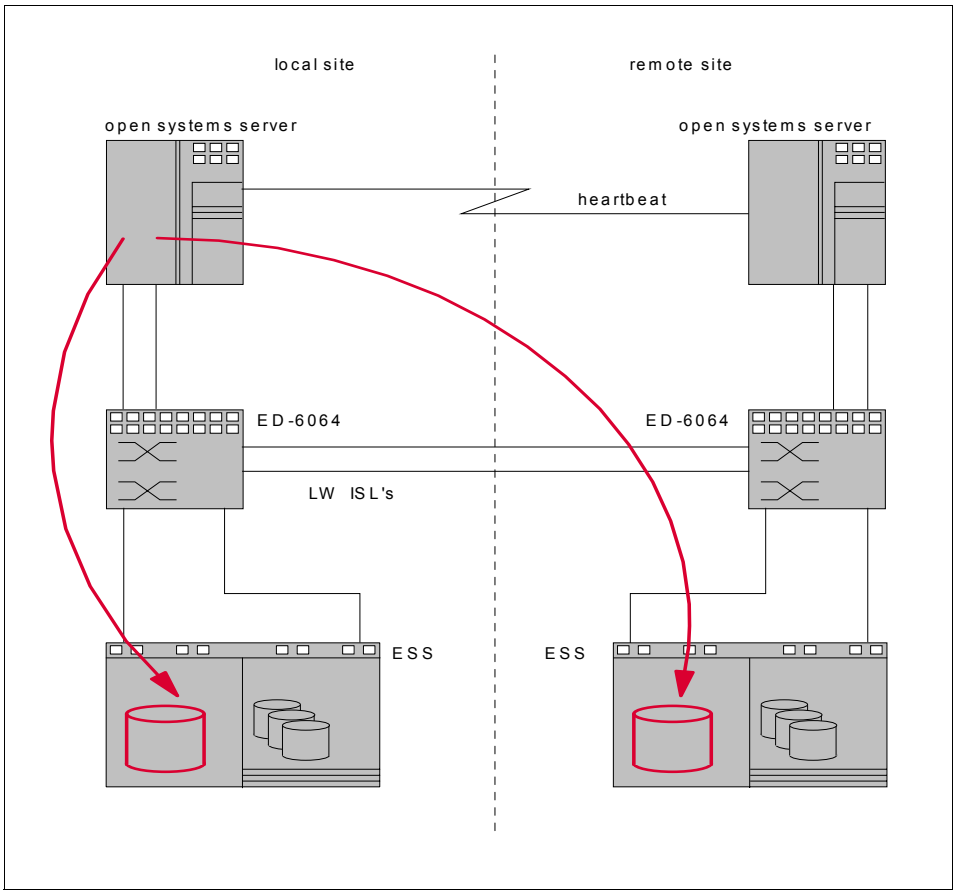


*Figure 4-99 LVM mirroring using the SAN*

This is a setup which consists of the same configuration at both the local and the remote site. Both sites can be up to 10 km apart when using 9 micron fiber optic cable. The open systems server cluster, for example, can consist of two or more RS/6000 with HACMP. The mirroring can be done with the native LVM on AIX.

Another solution could be SUN servers running the Veritas Cluster Server and the Veritas Volume Manager, for example. Due to the high availability of the McDATA ED-6064, one may be sufficient, but only if that leaves enough ports to accommodate the rest of the environment and its expansion.

When more ports and even higher availability are desired, this solution can be extended with another director at each site. Even though a director is highly available, using two independent fabrics (red and blue) removes the director itself as an single point of failure and may not always be regarded as a paranoia.

This is shown in Figure 4-100.

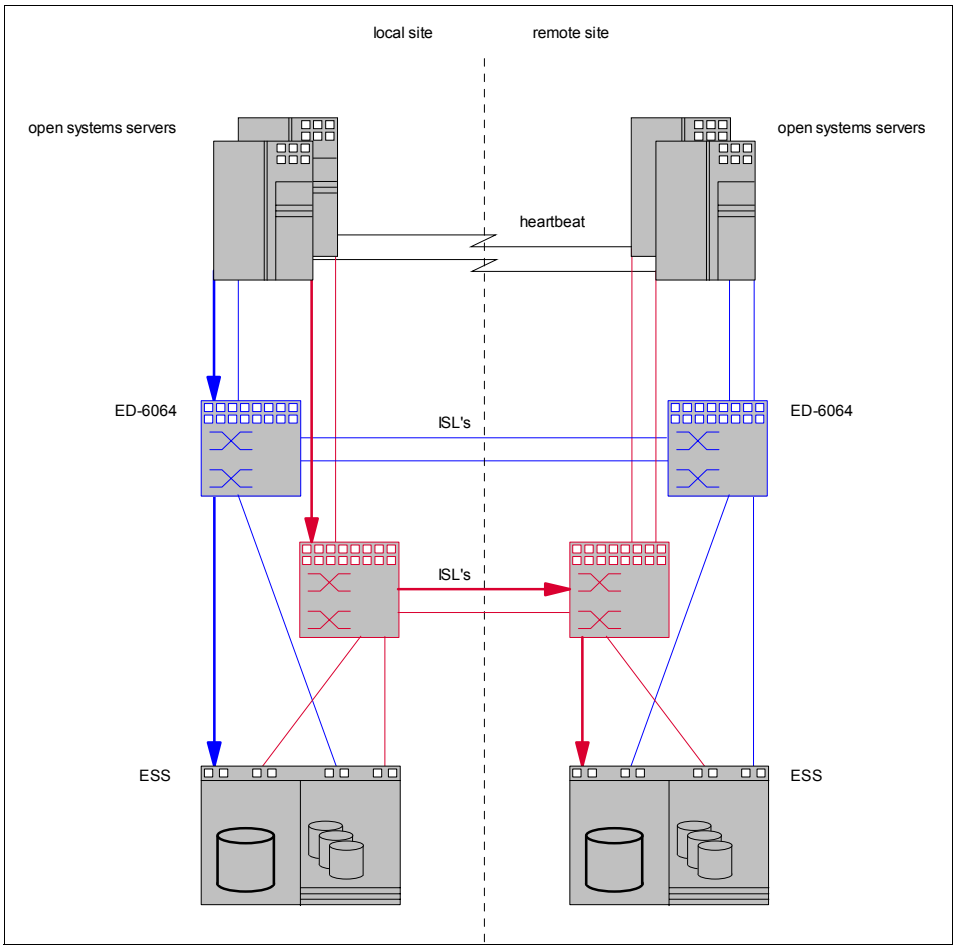


*Figure 4-100 Using two independent fabrics for high availability*

The arrows indicate a possible route for the data to get to both parts of the mirrored sets. In this setup there is no single point of failure at a device level, and even if one site completely fails, the other site will take over operation immediately.

In our example for a multi switch fabric, shown in Figure 4-101, we are not focusing on clustering. What we want to show is how to apply zoning in a multi switch fabric.

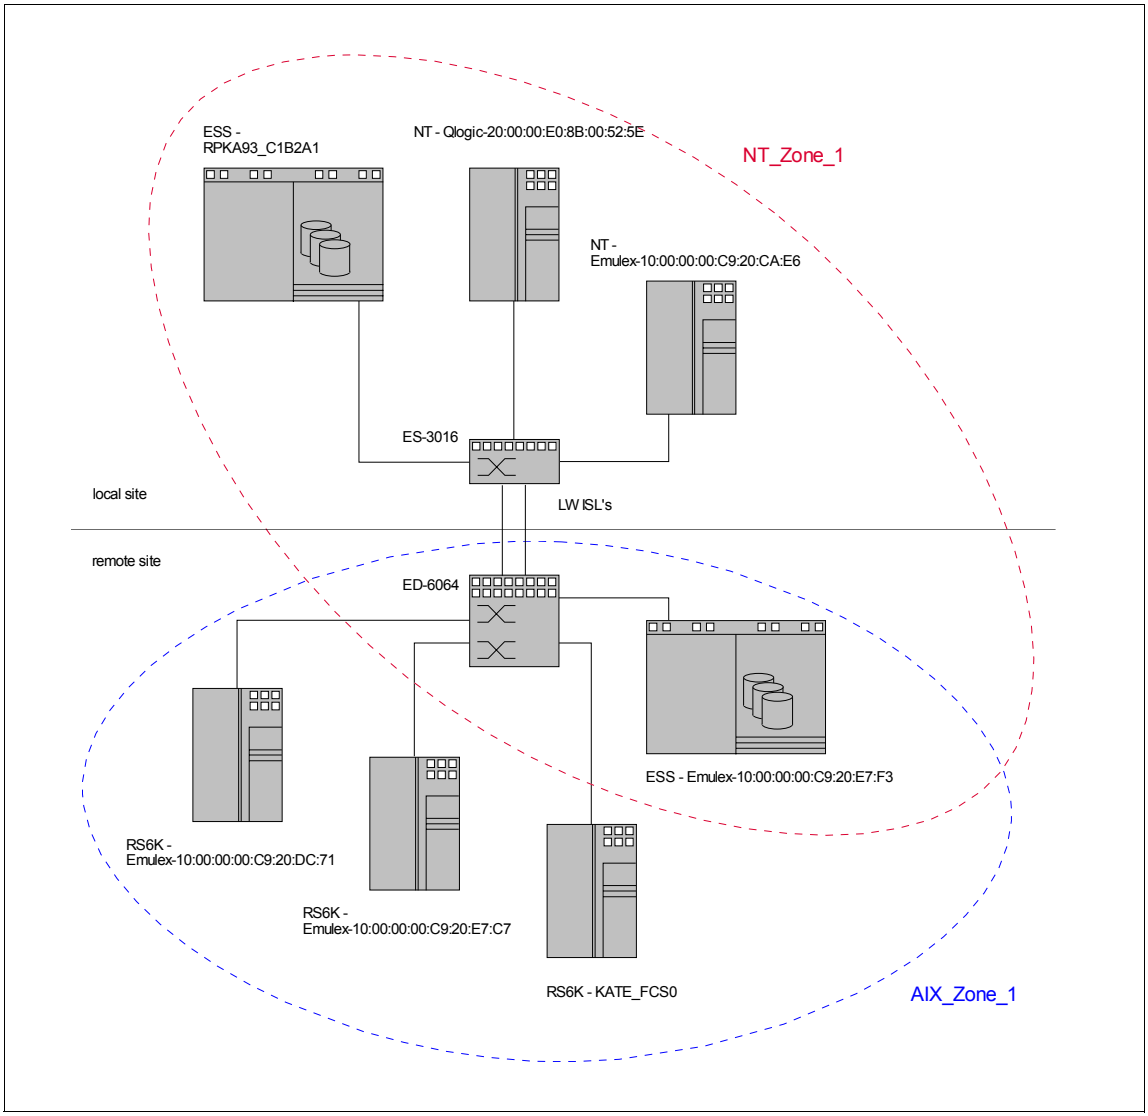


*Figure 4-101 Our zoned multi switch fabric*

Our NT zone spans over both sites, with two ESSs and two Netfinity's. One ESS is at the local site and the other is at the remote site. Both sites are connected with three longwave ISLs between one ED-6064 and one ES-3016 (it could also be an ES-3232 or ES-3216). At the remote site are the AIXD zone with three AIX servers and one ESS that is also a member of the NT zone. This example can be used to establish a mirrored set from within the Windows NT Disk Administrator, with one local copy of the data and one remote. Conversely, the AIX zone is limited to the devices at their site.

## Limits for McDATA multi switch fabrics

The McDATA fabric supports up to 31 interconnected switches managed from one EFC Server (the domain IDs range is from 1 to 31). Although we can connect many switches, the hop count supported by McDATA is limited to three, due to the delay that is applied traversing every switch. The hop count with McDATA is equal to the number of ISL connections traversed between the source and the destination.

**Note:** In IP networking a hop count means the number of connectivity devices (for instance routers) between the source and destination. This makes up the difference of one more hop in IP networking than in FC networks with the same amount of interconnected devices.

These are some additional requirements:

- Every McDATA product should be configured with a unique domain ID and IP address.
- The two fabric devices (director or switch) will not merge if they have the same configured domain ID.
- Only one principal switch is elected per single fabric.
- The flow control parameters (BB_Credit, RA_TOV, ED_TOV) must be the same on every switch that joins the fabric.

IBM supports, with its McDATA products, a homogenous SAN environment.

## 4.9.3 Setting up our zoned multi switch fabric

We will use one ED-6064 and one ES-3016 for our zoned cascading example. We configure both switches as we did before. First, we define the director with its EFC Server and then we define the switch to the same EFC Server and configure it with the Product Manager. After defining the switch to the EFC Server, the Product View now looks like that shown in Figure 4-102.

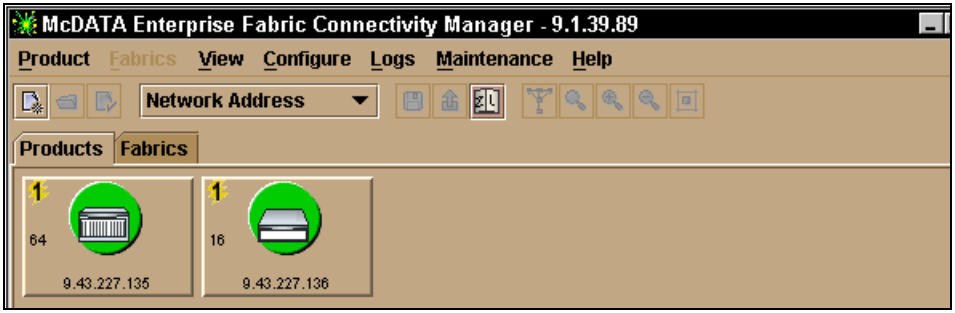


*Figure 4-102 EFC Manager: with two managed switches*

To include the second switch in the fabric of the first, we basically need to connect the switches with longwave or shortwave Fibre Channel cables. The fabric building process itself is transparent to us, as the switches will recognize the connection and automatically configure the G_Ports to be used as E_Ports. However, there are some configuration options that need to be set up or reviewed before connecting the switches.

### Setting the switch priority

In every multi switch fabric, one switch has responsibility for the domain address manager functionality. This switch is known as the principal switch. It controls the allocation and distribution of the domain IDs for all connected switches in the fabric, there must always be a principal switch in a fabric.

A switch can be manually set to be the principal switch, or it may be set to never be principal. This may be done in a core-to-edge environment, for example, where it makes sense for a core switch to normally be principal. If switches are set to the “default” priority, the one with the lowest numerical WWN value becomes the principal switch. To change the switch priority, we also use the Configure Operating Parameters window from the switch view, as shown in Figure 4-103.

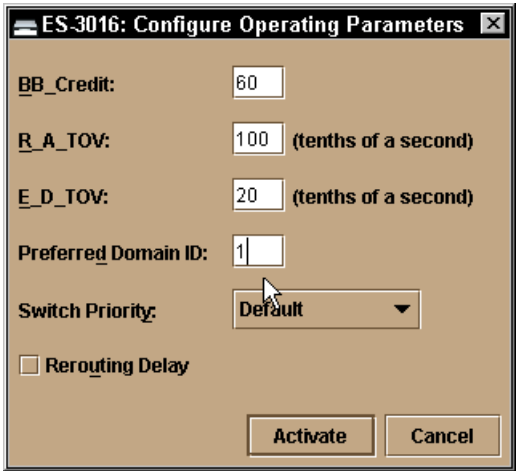


*Figure 4-103 EFC Manager: Configure, Configure Operating Parameters*

## Setting the domain ID

Each switch is recognized in the fabric as a domain and is identified with a domain ID. Domains are used for the 24-bit FC addresses that identify the switch ports in a fabric. Every domain ID in the fabric must be unique ranging from 1 to 31.

To view or to change the domain ID, we go to the Product Manager of the specific switch. Then we select **Configure —> Operating Parameters...**. In the next window, as shown in Figure 4-103, we can change the preferred domain ID and other Fibre Channel parameters for the director, for instance, the Switch Priority.

**Tip:** To change any operating parameters, the switch must be offline.

The domain ID is requested from the principal switch once the switch comes online to the fabric. The preferred domain ID is only used if it does not exist in the fabric. If it is in use already, an unused ID is assigned. This can be seen in the Switch properties display, found by **Product —> Properties...** in the product manager view, as shown in Figure 4-104.

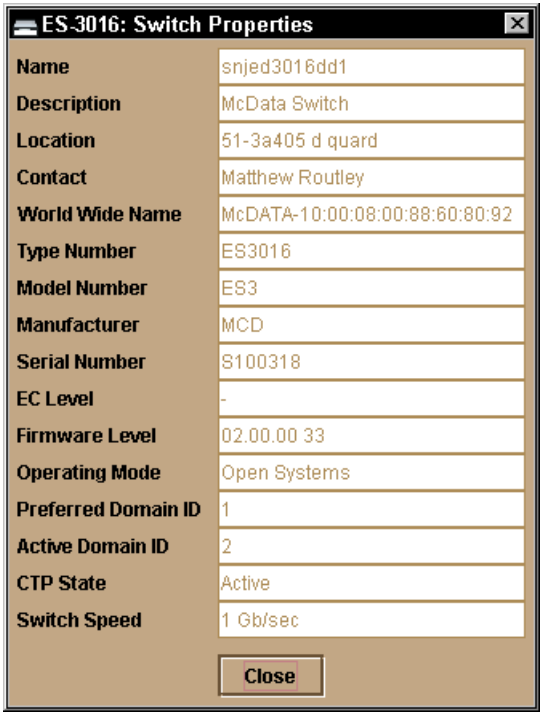


*Figure 4-104 Switch properties, Active Domain ID*

We recommend manually setting the domain IDs prior to building the multi switch fabric and prior to zoning. One reason is that when two switches are joined while active, they will determine if the domain ID is already in use, but if there is a conflict it cannot be changed in an active switch. This conflict will cause the fabric merging process to fail.

The second reason is that the domain ID is used to identify switch ports when zoning is implemented using the domain and switch port number. If domain ID's are negotiated at every fabric start up, there is no guarantee that the same switch will have the same ID next time, and therefore any zoning definitions may become invalid.

## Configuring the ports for the ISLs

The ports for the ISLs can be configured just like the other ports as we described in "Configuring the FC ports" on page 496. From here we can assign a name reflecting the usage of the ports, we can check the check box for extended distance buffering, and we can verify and change the port definitions.

In our example, illustrated in Figure 4-105, ports 0 and 1 are both able to build ISLs. Port 1 is already connected and defined as an E_Port and port 2 is a G_Port that will recognize that it has to act as an E_Port when connected to another switch.

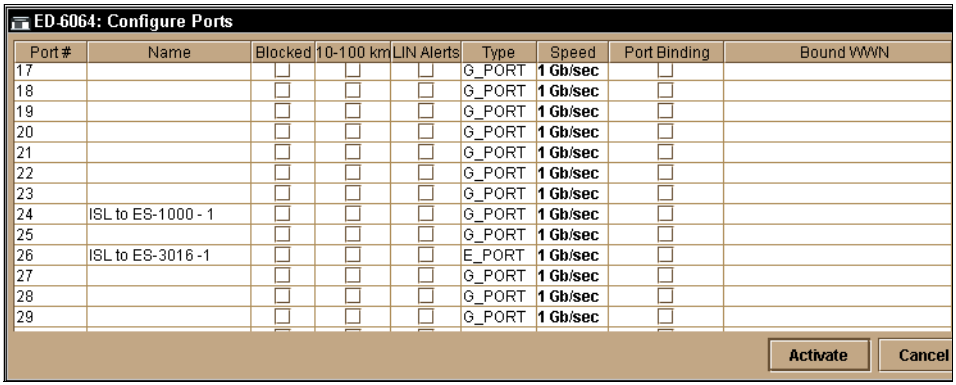


*Figure 4-105 EFC Manager: Configure Ports*

## Other prerequisites for a multi switch fabric

To be able to successfully establish a multi switch fabric, other prerequisites also apply. The operating parameters, resource allocation time-out value (R_A_TOV) and error detection time-out value (E_D_TOV) must be the same and the zoning configuration must be compatible.

## Verifying the compatibility of the zoning configuration

Once the switches are connected with ISLs the adjacent switches exchange their zoning information and merge it to a single active zone set. This resulting zone set now applies to every switch of the merged fabric.

Fabrics can be joined when none of them is zoned, when one of them is zoned or when both of them are zoned. Not zoned means no zone set is active and the default zone is enabled:

1. If none of the fabrics are zoned, no zoning information will be exchanged and the result will be a multi switch fabric with no zoning.
2. If one of the fabrics is zoned, the active zone set will propagate across the fabric and the result will be a multi switch fabric with the zoning information of the former standalone fabric which was zoned before.

3. If both of the fabrics are zoned, the zoning will only work if the configurations are compatible. If the zone configurations are not compatible, the E_Ports of the switches become segmented, which means they cannot carry traffic from node ports, but they can still carry management traffic. Zoning configurations are compatible if one of the two requirements are met:
   - The active zone names of each fabric to be merged are unique, if the zone members are not identical.
   - The active zone names of each fabric to be merged can be identical, if the zone members are identical as well.

In our case, the director fabric is zoned, the switch fabric is not. This means that the active zoning information will propagate across the fabric and the two independent fabrics will join to form a multi switch fabric. If the switch was also zoned, we may end up with a segmented fabric when attempting to merge due to “incompatible zoning”.

Prior to connecting the switches, the Fabric View of the EFC Manager looks like that shown in Figure 4-106.

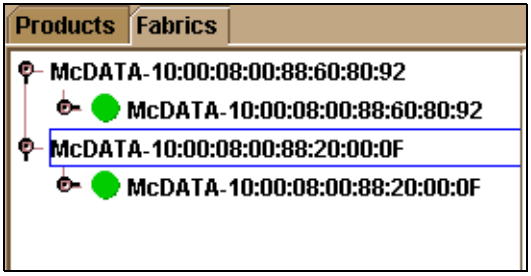


*Figure 4-106 EFC Manager: Two independent fabrics*

This shows that we have two independent fabrics which are not connected with ISLs. It could also look the same, for instance, if the ISL ports have been blocked or have been configured as an F_Port, or if the zoning configuration was incompatible.

## Connecting the switches

Now we can connect the two switches with ISLs. We are using one longwave ISL between the two switches.

To see what the topology looks like now, we select the **Topology** tab of the Fabrics view, as shown in Figure 4-107.

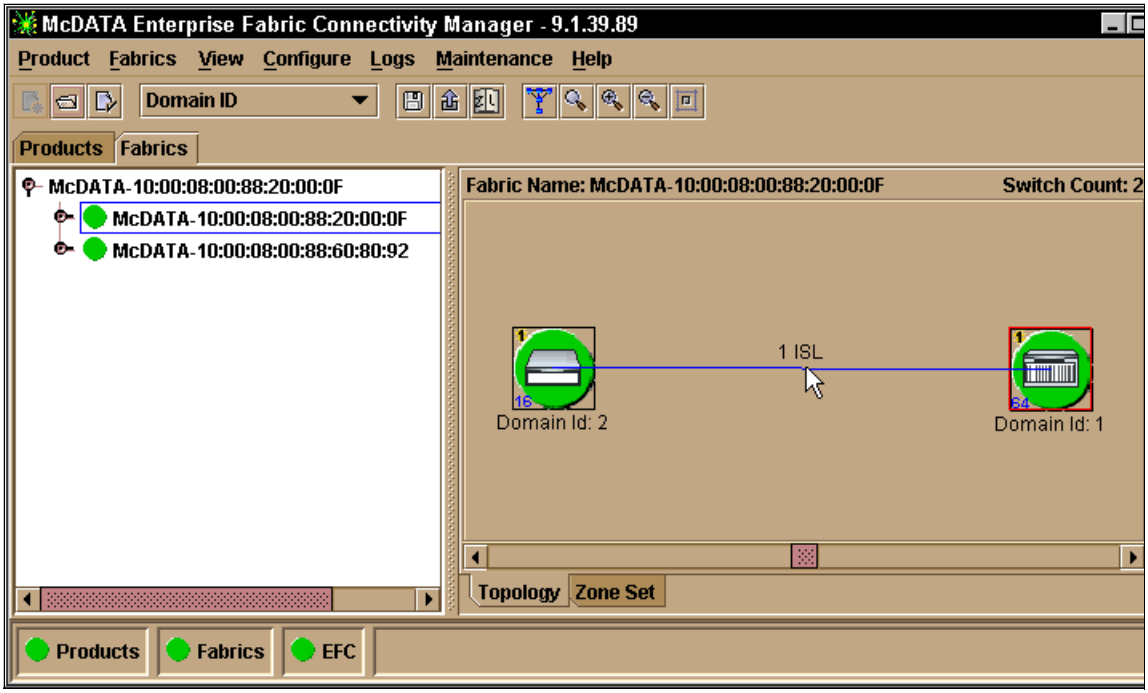


*Figure 4-107 Fabric Manager: Topology View, one merged fabric*

Notice now that the left column shows two fabric switches under the same principal WWN.

Moving the mouse pointer over the link indicates that the connection of the two switches consists of one ISL. This would change as we add more ISLs between these two switches. Clicking either product icon will open the Product Manager for the associated switch. Also note that we have chosen to view the products by their domain ID.

A new feature of EFCM is “Persist Fabric”, which allows us to be notified of changes to the fabric, for example, in the event of a switch failure, or an ISL failure. To turn on “Persist Fabric”, we right-click in the background area of the fabric display; this opens a context menu as shown in Figure 4-108.

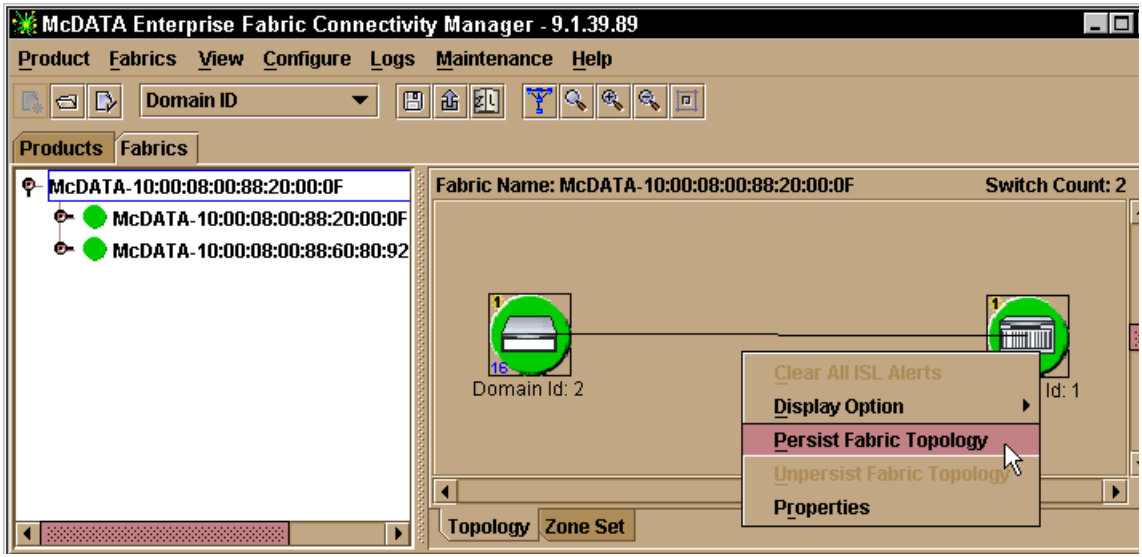


*Figure 4-108 Fabric Manager: Persist Fabric Topology*

Selecting **Persist Fabric Topology** prompts us to give the fabric a name, and we will call it “ITSO Lab51”. Similarly, we can give each product a nickname by right-clicking each product, and this is very helpful to simplify the identification of each switch, as shown in Figure 4-109.

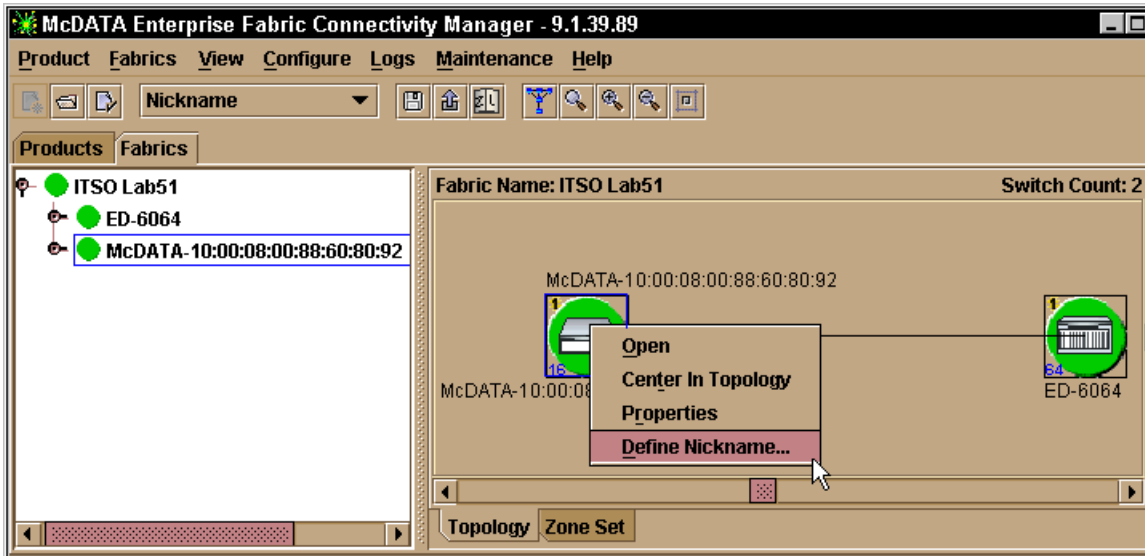


*Figure 4-109 Fabric Manager: Product Nicknames*

Note that we have selected in the line below the menu bar, to view our fabric by Nickname.

Now, with Persist Fabric turned on, a failure of the ISL between our switches would be shown with the yellow triangle attention icon and the ISL changing to a broken yellow line.

We show this in Figure 4-110, and note also the Fabrics Icon in the bottom line is alerting us to a problem in our fabric. Further detail of why the fabric failure occurred can be seen by selecting **Logs —> Fabric Log...** from the pull-down menu.

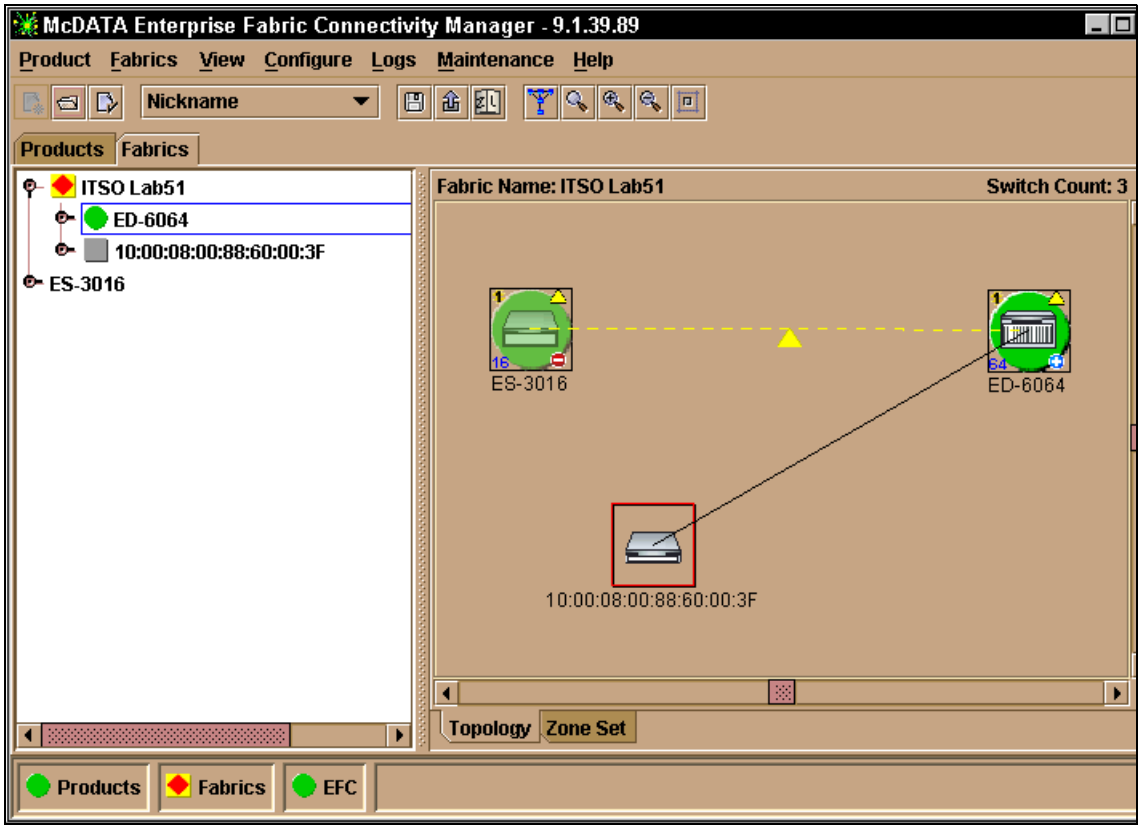


*Figure 4-110 Three switches cascaded, one not defined, and a broken ISL*

If we connect another switch to the fabric without defining the switch to EFCM, or if another EFC server is managing that switch, we see an ISL and a switch in the fabric view, but there is no green circle around the third switch. This is because the EFC Manager is not able to retrieve any information about the device and is unable to manage this device.

We have now successfully completed all the steps necessary to cascade McDATA switches with zoning.

## 4.10 Open Trunking

Open Trunking addresses ISL over-subscription and under-subscription problems experienced in the fabrics due to the load distributed across multiple ISLs in a round-robin fashion. The current load sharing mechanism does not have the capability to detect the link utilization in terms of bandwidth usage. There is no traffic monitoring and sampling done on ISL(s), it is purely done by dedicating a switch input/output ports to route traffic for single or multiple devices.

With this mechanism of static allocation of fabric ports to end devices, this can result in over-utilization when a single or multiple high end servers and storage device are directed to use the same switch port for data flow. Another scenario is link under-utilization, which could be performed by dedicating an ISL(s) for low end devices that may only use 15% of the link's capacity. This static distribution of load remains constant as long as the fabric is stable. If the end device reboots or if the fabric reconfigures due to a new link being added or removed, this will result in re-discovery of the routes and assigning new paths to end devices. The chances of eliminating link congestion in a logical fashion are minimal with the static load sharing mechanism, even by adding new ISL(s) between the two switches.

The Open Trunking feature monitors the average data rates of all traffic flow on ISLs (from a receive port to a target domain), and periodically configures routing tables to reroute data flow from congested links to under-utilized links and efficiently uses bandwidth. The objective of Open Trunking is to make the most efficient possible use of redundant ISLs between neighboring switches, even if these ISLs have different bandwidths.

Open Trunking is performed using the FSPF shortest-path routing database. This solution uses McDATA patented technology to provide real-time traffic monitoring. The feature controls Fibre Channel traffic at a flow level, rather than at a per frame level in order to achieve optimal throughput. This feature may be used on McDATA switches in homogeneous as well as heterogeneous fabrics. This feature complies with current Fibre Channel ANSI standards.

Open Trunking is an optional, user-purchasable software feature that provides automatic, dynamic, statistical traffic load balancing across ISLs in a fabric environment. This feature is available with EO/S 5.1 and EFCM 7.1 and can be enabled on a per-switch basis. It operates transparently to the existing FSPF algorithms for path selection within a fabric.

Open Trunking is discussed more in the IBM Redbook:

- *IBM SAN Survival Guide*, SG24-6143

## 4.10.1 Configuring Open Trunking

The load-balancing aspect is not a user configurable feature. The user can enable/disable OpenTrunking on the switch and configure the settings for congestion thresholds (per port) and the low BB_Credit threshold for fine tuning purposes if required. The ISLs between two switches cannot be manually configured as “trunk groups” and there is no concept of master trunk. The least cost paths are already stored in the path selection table and will be used to redistribute traffic automatically when congestion is experienced on the ISL(s). This means that flow may be rerouted onto a link that goes to a different adjacent switch, as long as that link is on the least cost/shortest path to the destination domain ID.

### Installing the Open Trunking feature key

The Open Trunking feature key can be installed using the EFC product manager, CLI and also from the SANPilot interface.

A feature key is a string of alphanumeric characters consisting of both uppercase and lowercase. The following is an example of a feature key format: **AUY2-9t7A-D7qs-D4**.

**Note:** The total number of characters may vary. The key is case sensitive and it must be entered exactly, including the dashes.

To purchase the Open Trunking feature key, you must supply the device type and serial number of the device that you want to install the feature on. The easiest way to retrieve the serial number is using the SANpilot interface listed under the **Unit Properties** menu as shown in Figure 4-111.

The Unit Properties menu lists the switch type, serial number, WWN, firmware and so on. If you are trying to install the feature on an unsupported firmware level you will be notified that a firmware upgrade is required in order to use the feature.

In Figure 4-111 we show the Unit Properties view.

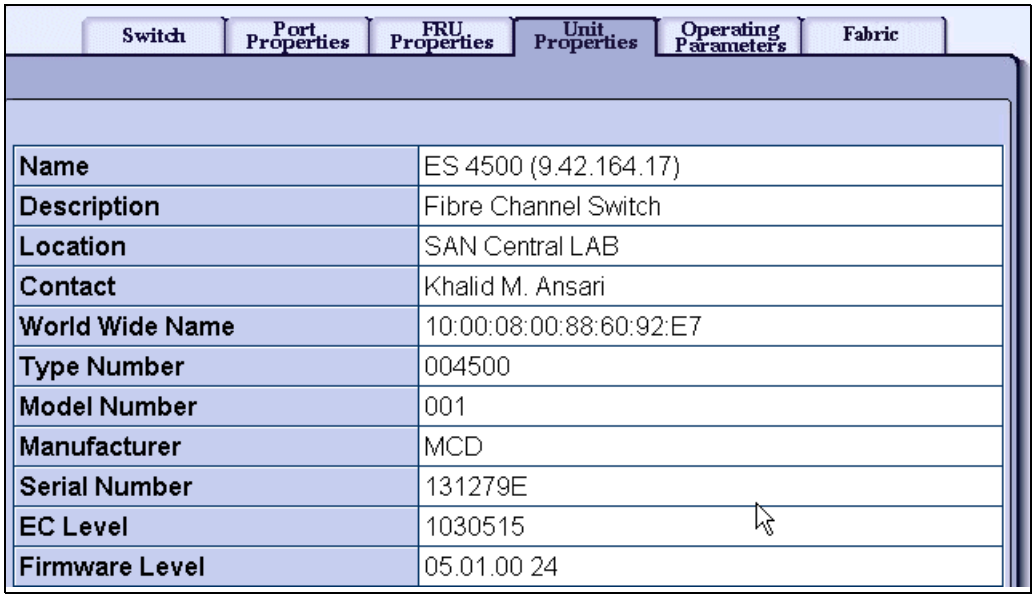


*Figure 4-111 Unit Properties menu from SANpilot interface*

Once you have received the feature key you can proceed to install the feature key as shown in the following topics.

**Attention:** The feature key, which is encoded with a device's serial number, can only be installed on the device to which it is assigned. You can enable the feature key with the director online. However, if a current feature is disabled by activating a new feature key, you should take the director offline before enabling the new feature key.

## Installing the feature Key using SANpilot

Access the switch or director on which you want to install the feature using the Web browser, then select **Feature Installation** tab under the **Operations menu** as shown in Figure 4-112.

In this example the ES-4500 Sphereon switch is used to demonstrate the procedure.

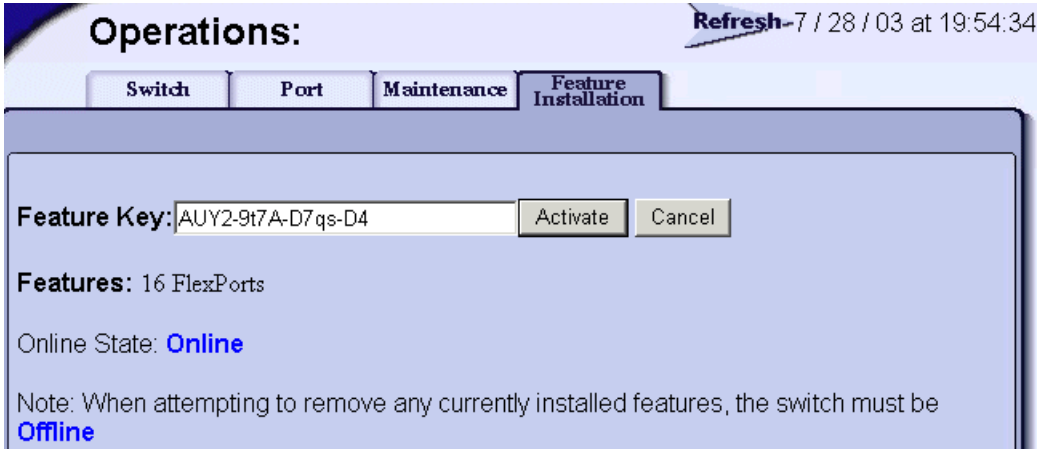


*Figure 4-112 Feature key installation tab under Operations menu*

Insert or paste the alpha-numeric feature key in the text box and click the **Activate** button. You will be prompted with a message to verify and confirm the New features that will be installed and any old features that may be deleted after the activation is successful.

Notice that the new features, as shown in Figure 4-113, that will be installed are Open Trunking and McDATA SANtegrity. In this case the two optional features were purchased for the ES-4500.

Click the **Activate** button after verifying the current and new feature information as shown in Figure 4-113.

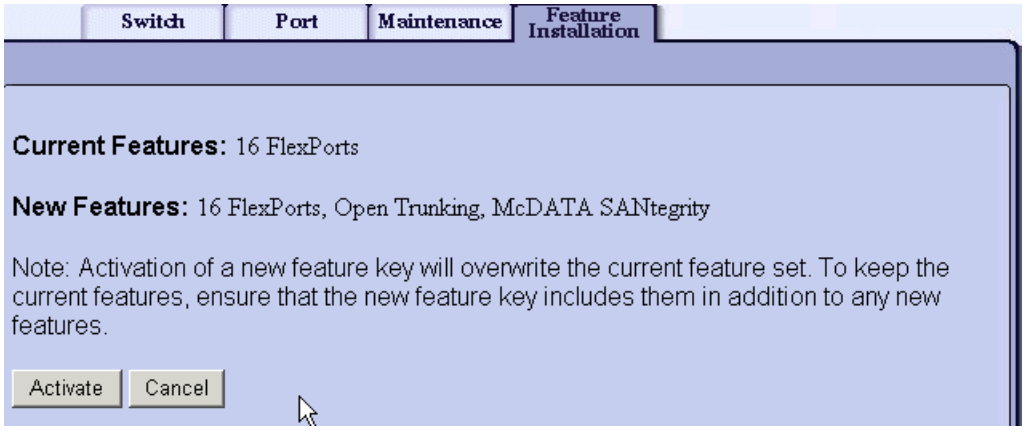


*Figure 4-113 Activating the new features*

If the correct key was supplied, then the feature installation is successful, as shown in Figure 4-114.


*Figure 4-114 The successful feature installation and activation menu.*

## 4.10.2 Enabling Open Trunking

The OpenTrunking feature can be enabled for a specific switch using the EFC Product Manager, CLI and the SANpilot Interface.

The ED-6140 director is used to demonstrate the procedure to enable Open Trunking using the SANpilot interface. Access the switch using the browser and select the **Performance** tab under the Configure menu and switch the **Open Trunking State** to **Enable** from the drop-down list as shown in Figure 4-115.

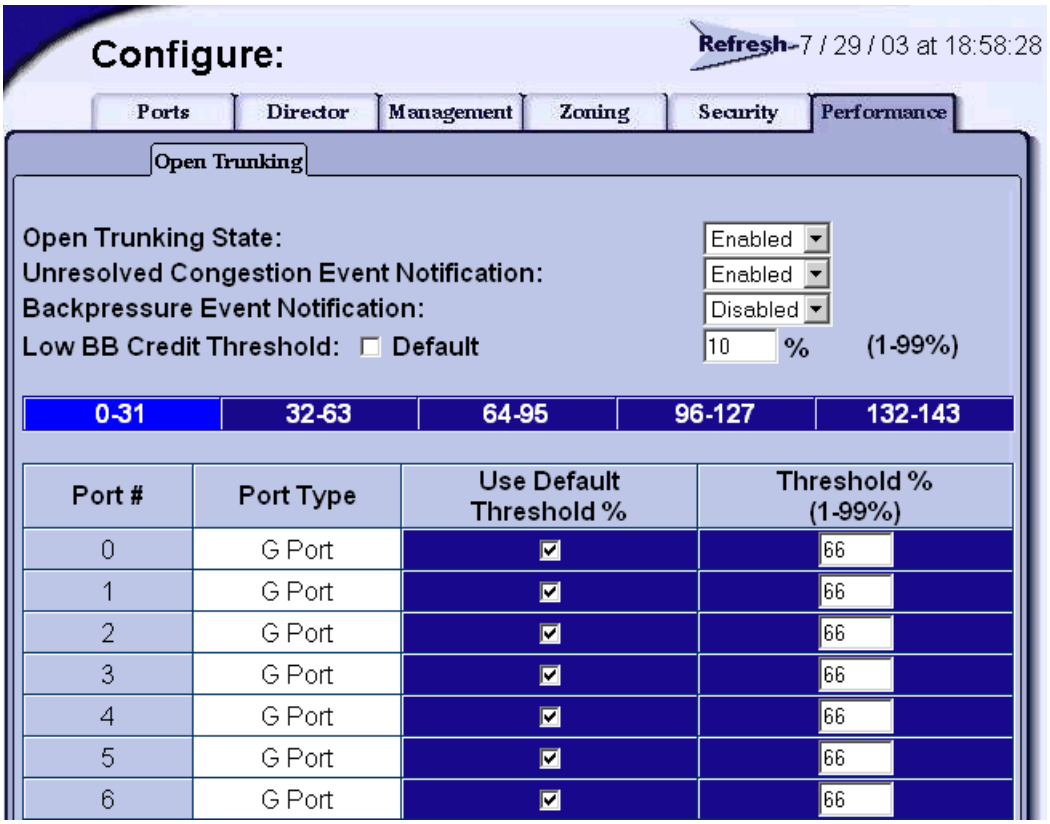


*Figure 4-115 Open Trunking State option*

More detail regarding fine-tuning the other options can be found at the Web site:

`http://www.mcdata.com/knowcenter/techpubs/index.html`

## Open Trunking log

The Open Trunking log is available from the EFC product manager and log flow redistribution data. From the EFC Product Manager, select the **Logs—> Open Trunking Log** option and the window that opens will list data for any rerouting that is experienced on the director or switch, as shown in Figure 4-116.

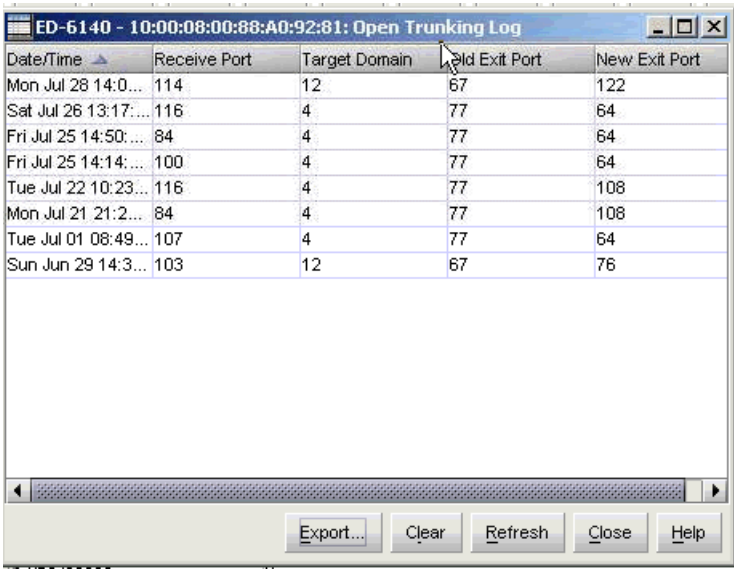


*Figure 4-116 Open Trunking Log view*

# 4.11 SANtegrity

SANtegrity binding enhances data security in large and complex SANs and consists of Fabric and Switch Binding features. These features provide permit and deny operations for connecting a switch to the fabric, and end device attachment to the switch or fabric. SANtegrity, and therefore the binding features, can be enabled by purchasing a feature key and then installing and activating that feature key.

## 4.11.1 Fabric Binding

SANtegrity Fabric Binding gives access control tools across the fabric through which the system administrator can permit or deny switches from connecting to the fabric in a SAN. Without the Fabric Binding feature enabled, the fabric/zone configuration can be easily modified or deleted by connecting a new switch to the fabric, and there are no built-in mechanisms to permit or deny any switch from merging into the fabric. It gives greater control to the system administrator and gives protection from hacking into the fabric.

Once Fabric Binding is activated, the Fabric Membership List (FML) automatically includes all the switches that are members of the fabric at the time of Fabric Binding activation. Switches and directors not in the Fabric Membership List at the time of activation are prohibited from joining, and raise alerts and attention indicators as invalid attachments.

In order to add a new switch to an existing fabric that has Fabric Binding activated, the existing Fabric Membership List must be updated with the WWN and domain ID of the switch or director that will be added to the fabric. The new switch or director must also have Fabric Binding activated (prior to joining the existing fabric) and a Fabric Membership List containing the WWN and domain ID of every switch in the existing fabric.

The list identifies switches by WWN and domain ID, so domain ID's must be statically allocated while Fabric Binding is active. Because of this, the Insistent Domain ID feature is automatically enabled on each switch in the fabric when Fabric Binding is activated, and it cannot be disabled while Fabric Binding is active.

EFCM will provide Fabric Binding configuration options in the Fabric Manager (that is to say, for a specific fabric), and not in the Product Manager. Fabric Binding can also be configured using the embedded CLI interface.

## General rules for Fabric Binding

These are some general rules that apply to Fabric Binding:

- Not surprisingly, Fabric Binding activation is disallowed if SANtegrity is not installed.
- Fabric Binding activation is disallowed if the switch is offline. Switches can only be removed from the Fabric Membership List if they are not currently in the fabric.
- If the Fabric Binding configuration in the two fabrics is incompatible (that is to say, the Fabric Membership list is not identical), then the fabrics will not join. This is resolved by adding the attached switch to the Fabric Membership list or changing the Fabric Binding state to Inactive. The Fabric Membership list should be identical on all the switches in the fabric.
- Fabric Binding deactivation is prohibited if the Enterprise Fabric Mode is set to Active.

## Configuring Fabric Binding

We will use EFCM 7.1 fabric manager to demonstrate the procedure to configure Fabric Binding. From the EFC fabric manager menu, select the fabric on which the Fabric Binding feature needs to be activated from the Fabric Tree menu in the left hand column, as shown in Figure 4-117.

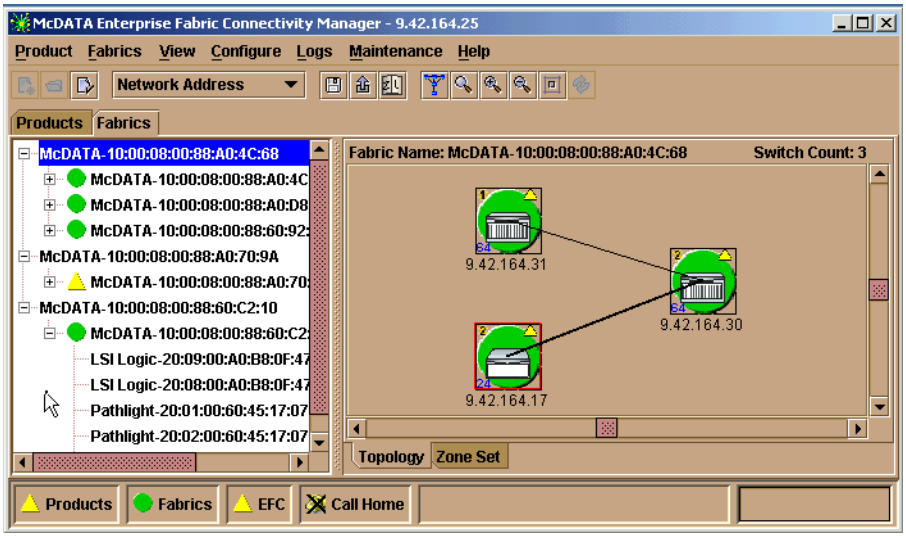


*Figure 4-117 Fabric Tree list*

The fabric WWN has been highlighted (in blue), and once selected, the topology view shows the number of switches in the fabric. Figure 4-117 shows that there are three switches in the fabric, so the Fabric Binding feature will be activated on those three switches and they will automatically be included the Fabric Membership List.

From the EFC Fabric Manager menu, select **Fabric—> Fabric Binding**, and the menu to enable Fabric Binding appears as shown in Figure 4-118.

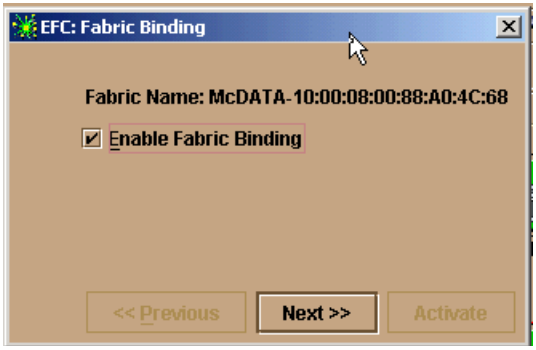


*Figure 4-118 Configure Fabric Binding menu*

Click **Next** to proceed to activate the Fabric Binding feature. The new menu that appears lists the Fabric Membership list to be activated once the Fabric Binding feature is enabled, as shown in Figure 4-119.

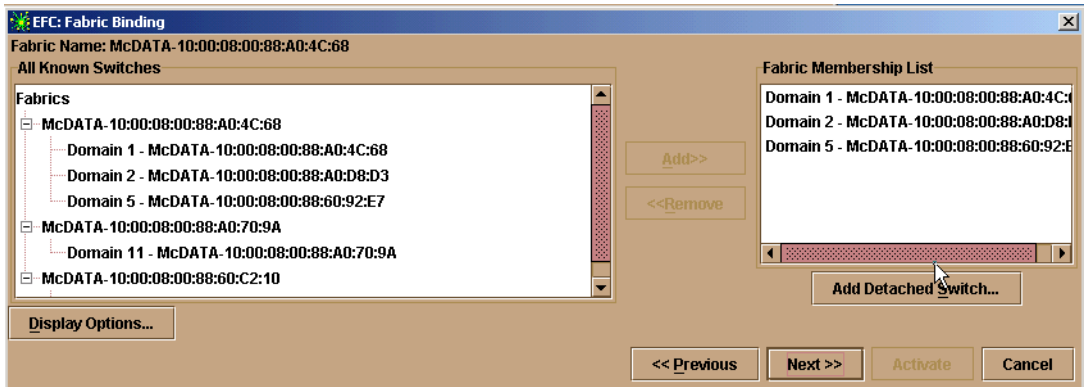


*Figure 4-119 Fabric Binding menu displaying the Fabric Membership List*

Members (switches) can be added or removed from the list before Fabric Binding activation. It also allows you to add detached nodes to the list for future use. Select **Next** for the Fabric Binding menu as shown in Figure 4-120.

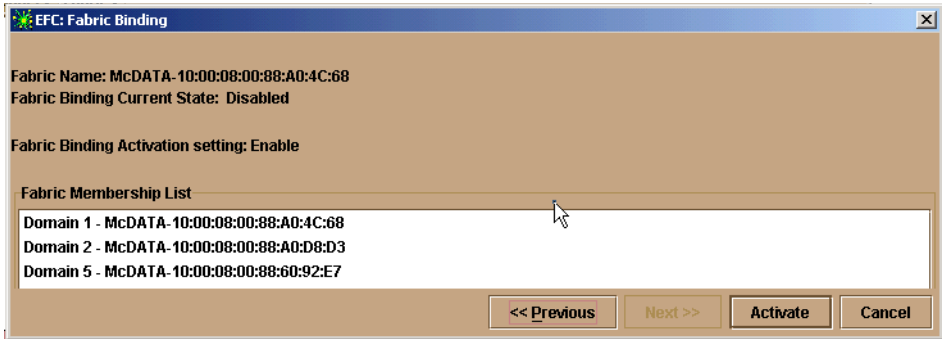


*Figure 4-120 Fabric Binding Activation List*

Verify and confirm the Fabric Membership List and then click the **Activate** option.

Upon successful activation of the Fabric Binding feature a “Fabric Binding Update” complete message, as shown in Figure 4-121, appears.

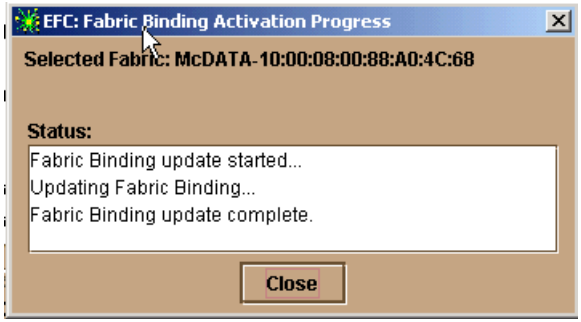


*Figure 4-121 Fabric Binding Activation complete*

At this point the Fabric Binding feature has been activated and the fabric is now locked. Any new switch will be denied access to join the fabric without manual intervention. The System Administrator must edit the Fabric Membership List and add the domain ID and WWN of the new switch to enable it to join the fabric. Furthermore, the new switch should have SANtegrity installed, the Fabric Binding feature enabled, and also have the same Fabric Membership List currently active in the fabric.

More details about SANtegrity can be found at Web site:

`http://www.mcdata.com/knowcenter/techpubs/index.html`

### Edit Fabric Membership List procedure

To Add a new member (switch) to the list, from the EFC Fabric Manager menu Select **Fabric —> Fabric Binding**, then select **Next** in the **Fabric Binding Enable** menu.

From the Fabric Membership List menu, all the switches managed by the EFC Fabric Manager are listed in the left column under the **All Known Switches** list. Select the switch to be added to the Fabric Membership List and select **Add**.

## 4.11.2 Switch Binding

SANtegrity Switch Binding allows an administrator the option to permit/deny which end devices can be connected to which director or switch ports by specifying the WWN of the devices in the Switch Membership List. Without the Switch Binding feature active on the switch, any device can connect to any switch port and there is no built-in mechanism to prohibit end device connectivity. This feature provides an additional layer of security and greater access control tools for the system administrator managing complex environments that include a large number of devices.

When Switch Binding is enabled, only devices that are connected and online are identified and added to the Switch Membership List automatically. Thus the devices in the Switch Membership List are allowed to connect. Servers, storage, and other switches *not* in the Switch Membership List while Switch Binding is enabled are prohibited from connecting, and will raise alerts and attention indicators as invalid attachments. Switch Binding can be implemented for all connections (switch or director binding) or for individual connections (port binding) to give greater granularity.

## Switch Binding enforcement modes

Switch Binding has different enforcement modes:

### *Restrict E_Ports*

E_Ports are blocked from forming ISL connections with any switch WWN not explicitly identified in the Switch Membership List. There is no restriction for F_Ports from connecting to the switch.

### *Restrict F_Ports*

F_Ports prohibit connections from any end device not explicitly identified in the Switch Membership List. There is no restriction for E_Ports to form ISL connections with other switches.

### *Restrict All*

Both E_Ports and F_Ports are prevented from connecting if the switch and end device WWN is not explicitly in the Switch Membership List.

## Switch Binding rules

The following rules apply to the Switch Binding feature:

- The Switch Binding feature cannot be enabled if SANtegrity is NOT Installed.
- If the switch is online and Switch Binding is disabled, the switch will automatically add the WWN of currently connected/online devices to the Switch Membership List (SML) if they are not already in the list.
- If the switch is online and Switch Binding is already enabled, then the user is only allowed to change the enforcement mode (**Restrict E_Ports, Restrict F_Ports, Restrict All*)***. In this case, the switch must automatically add currently attached devices to the SML if any are not already in the list.
- If the switch is offline when Switch Binding is enabled, then the switch does not automatically add attached devices to the Switch Membership List.
- WWNs can only be removed from the list only if the switch is either offline, or Switch Binding is disabled, or if the WWN is not currently connected to the switch. A WWN can also be removed if Switch Binding is not enabled for the

same port type as the WWN, meaning a WWN for an E_Port can be removed if Switch Binding is enabled and in Restrict F_Ports mode. Error message "WWN is already connected on port number [N] and cannot be removed from the list. You must first block the port or disconnect the device."

- If Switch Binding is enabled and restricting either E_Ports or All ports, then the switch searches for the WWN in the Switch Membership List. If the WWN is not in the list, an Invalid Attachment Reason Code is returned indicating a Switch Binding violation.
- If the WWN is not authorized, the port is placed in the **Invalid Attachment** state, and an Event Log entry (WWN Not Authorized) is generated. This is resolved in several different ways, such as adding the attached switch to the Switch Membership List, changing the Switch Binding state from Restricting E_Ports to Restricting F_Ports, or changing the Switch Binding state to Disabled.
- When a new device attempts to login to the fabric, the switch determines if the Port WWN of the attached device is authorized to connect in the following order:

  a. The WWN is verified against the current Port Binding configuration.
  b. The WWN is verified against the current Switch Binding configuration.

- If Switch Binding is enabled and restricting either F_Ports or All ports, then the switch searches for the WWN in the Switch Membership List. If the WWN is not in the list, the switch returns an Invalid Attachment Reason Code indicating a Switch Binding violation. If the WWN is not authorized, the port is placed in the Invalid Attachment state, and an Event Log entry (WWN Not Authorized) is generated.
- Switch Binding Disablement is prohibited if Enterprise Fabric Mode is Active and the switch is online. User interfaces will display an error message.

### 4.11.3 Configuring Switch Binding

The Switch Binding configuration can be performed from the EFCM Product Manager (Switch Binding is configured independently on each switch) and also from the embedded CLI.

Before the Switch Binding feature is enabled, it is best to verify the Switch Membership List to ensure that all the devices are attached to the Switch and you can permit or deny any device from the Edit Membership List menu.

From the EFCM Product Manager menu, select **Configure —>Switch Binding —> Edit Membership List** as shown in Figure 4-122.

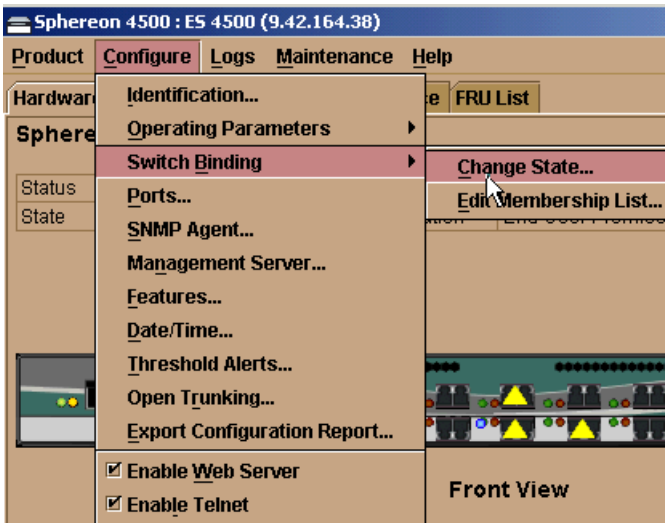


*Figure 4-122 Configure Switch Binding Change State*

The Edit Membership List menu is displayed. It lists all the end devices that are currently connected/online to the switch as shown in Figure 4-123.

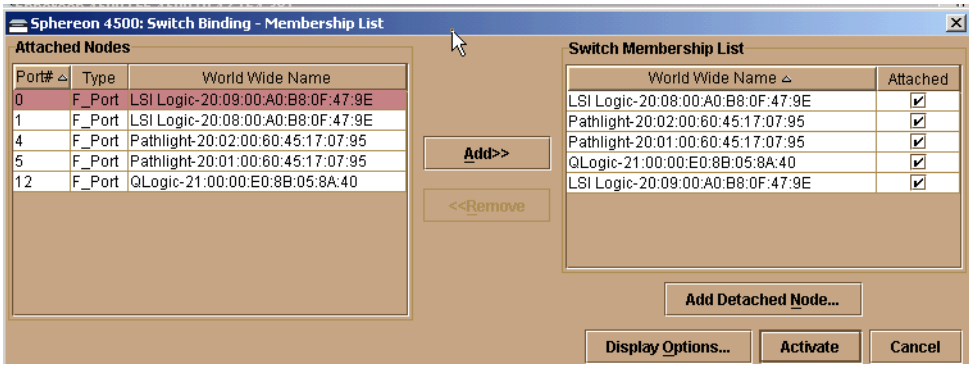


*Figure 4-123 The Switch Binding Edit Membership List menu*

**Attention:** The Switch Membership List can be edited only if the Switch Binding feature is disabled.

From the Edit Membership List menu, you can **Add** and **Remove** members from the Switch Membership List. To **Add** a device that is currently attached but not in the Switch Membership List, select the WWN of the device under the Attached Nodes list and it will enable the **Add** button, which can then be clicked as shown in Figure 4-123.

Similarly, the end devices can be removed from the Switch Membership List by selecting the device under the Switch Membership List, as it will enable the **Remove** option button, as shown in Figure 4-124.

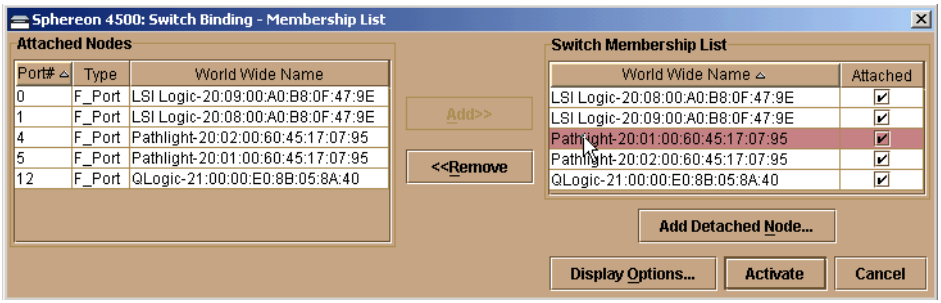


*Figure 4-124 Switch Binding Edit Membership List*

The Switch Binding Change State and the enforcement mode configuration options are available from the EFCM Product Manager view by selecting **Configure —> Switch Binding —> Change State** as shown in Figure 4-122 on page 574.

Once in the Switch Binding Change State menu, check the **Switch Binding Enable** option, and by default the **Restrict All Ports** option is selected as shown in Figure 4-125.

Once Switch Binding is enabled, the option to edit Switch Membership List is not available, but it will allow you to change the enforcement mode.

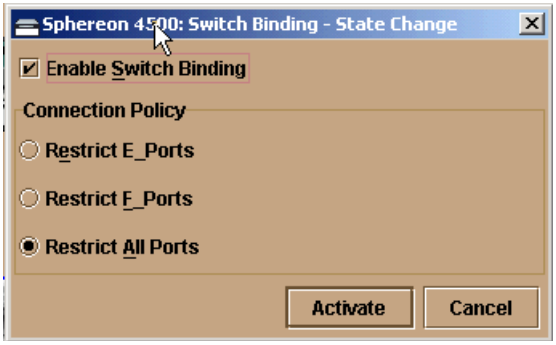


*Figure 4-125 Switch Binding Change State and Enforcement mode*

## 4.12 Firmware download procedure

Before proceeding to download and activate any new firmware, it is a best practice to carefully read the firmware release notes to understand the implications and also verify the fix list for any known problems. The release notes (and other documentation) are available (once registered) at the McDATA File Center site found at Web site:

http://www.mcdata.com/filecenter/template?page=docs.search

The EFCM product manager is used to demonstrate the procedure to download the firmware on the ED-6064.

We recommend that a maintenance window is scheduled in order to activate the new firmware and/or to negate any loss of connectivity issues that may occur, or be required, during the install.

These are the steps that we took to update the firmware:

1. Upgrade EFCM software on the EFC Server to version **07.01.00**. The procedure to upgrade the EFCM software is given below:
   a. Insert CD with **mcdataServerInstall.exe** and **EOSv05[1].01.00.bin** into EFCM Laptop CD Drive.
   b. Exit EFC Manager completely before installing **EFCM 07.01.00**.
   c. Execute the **mcdataServerInstall.exe**. Click **Start --> Run**, then Browse to the CDROM drive and select the **mcdataServerInstall.exe**, click **OPEN** and then **OK**.
   d. Follow the on-screen prompts at this point to install EFCM v07.01.00, select all default options... When complete, proceed to step # 2.
2. This step enables you to revert to the old configuration, in case of configuration lost or corruption issues due to CTP hang or incomplete firmware download.

   The EFC Server uses the product manager application to back up and restore the configuration data stored in the nonvolatile random-access memory (NV-RAM) on a director or switch CTP card on the EFC Manager data directory. The location and file name of the saved configuration cannot be modified. It only allows you to restore the configuration on the director by specifying the correct IP address by setting the director in OFFLINE state.
   a. Back up the device configuration by following the procedure given below.
   b. From the ED-6064 product manager menu, select **Maintenance —> Backup & Restore Configuration**, then select the **Backup** option as shown in Figure 4-126.

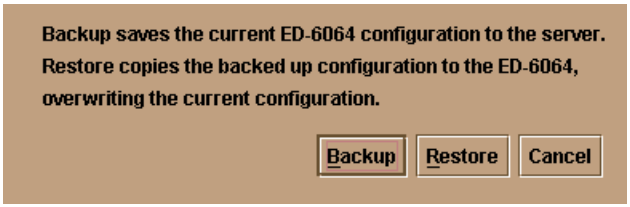


*Figure 4-126 Backup and Restore Configuration menu*

The following configurations are backed up to the EFC Server:

- Identification data (director name, description, and location).
- Port configuration data (port names, blocked states, and extended distance settings).
- Operating parameters (BB_Credit, E_D_TOV, R_A_TOV, director priority, preferred domain ID, rerouting delay, and director speed).
- SNMP configuration (trap recipients, community names, and write authorizations).
- Zoning configuration (active zone set and default zone state).

A backup is immediately attempted when you click the **Backup** button on this dialog box.

- A dialog box displays to confirm that the backup to the server is complete.
- If the backup fails, a dialog box displays to inform you that the backup to the server failed.

3. Download the firmware image file and transfer it to the firmware library:

   a. From the ED-6064 product manager menu, select **Maintenance —> Firmware Library** and then select the new option shown in Figure 4-127.

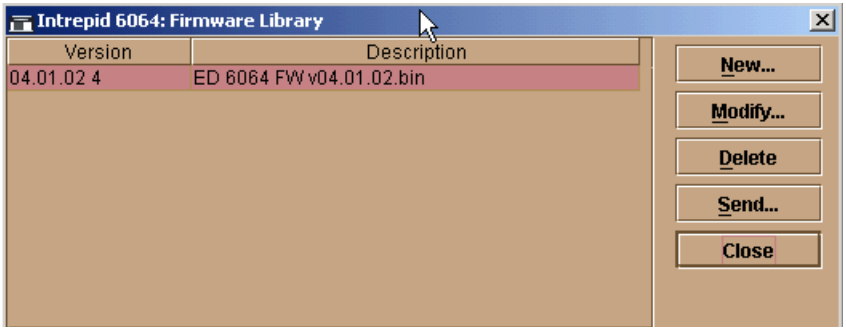


*Figure 4-127 EFCM Firmware Library*

b. Now browse and select the firmware image file and select **Save** as shown in Figure 4-128.

The **Save** option will transfer the image file into the firmware library database.

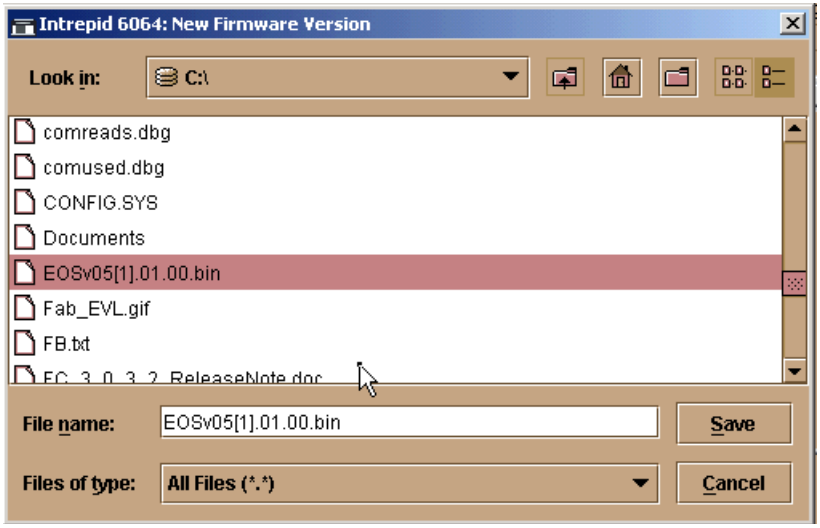


*Figure 4-128 New firmware version transferred to firmware library*

The message box is displayed with a *transfer complete message*. This means the firmware library has stored the new firmware.

4. The active CTP Card must be swapped to ensure that once the firmware is activated the CTP cards can successfully synchronize and a possibility of the hang symptom is ruled out.

   Using the Product Manager, execute a CTP swap:

   a. Verify that an amber LED indicator does not display for either CTP card.

   b. Verify the active and backup CTP cards from the hardware menu of the ED-6064 product manager view by double-clicking the CTP cards shown in Figure 4-131.

   c. Right-click the active CTP, and select **Switchover** from the pop-up menu

   The CTP card in slot 0 is active, as shown in Figure 4-129.

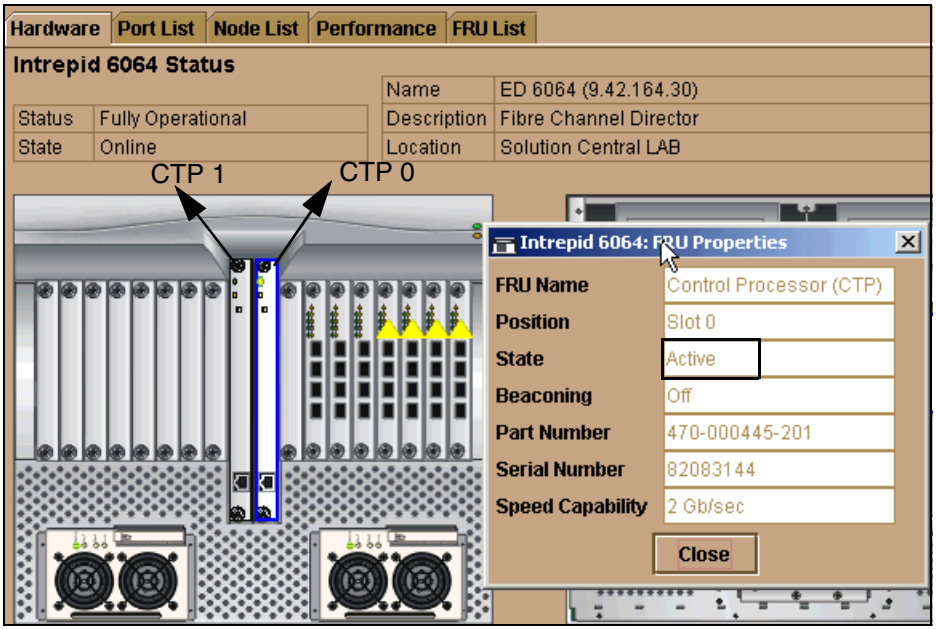


*Figure 4-129 CTP card status*

d. Right-click the CTP 0 in order to show the **Switchover...** button as shown in Figure 4-130.

e. On the Switchover CTP dialog box, click the **Switchover...** button to switch operation to the backup CTP card. When switchover occurs, the green LED illuminates on the backup CTP card to indicate that it is now the active card. Note that the director will lose its Ethernet connection for a short period during the switchover process.

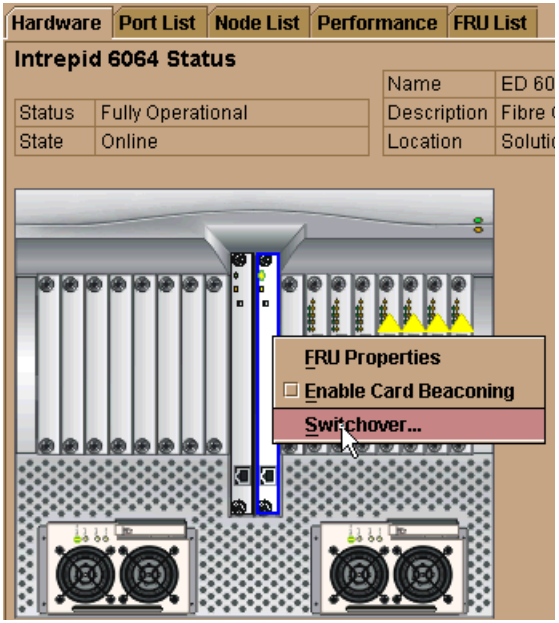


*Figure 4-130 CTP Switchover button*

5. Download and Activate Firmware EOS 5.01.00:

   a. From the firmware library menu, select the firmware that was stored previously and click **Send...** and it will prompt for confirmation to send the firmware, as shown in Figure 4-131.

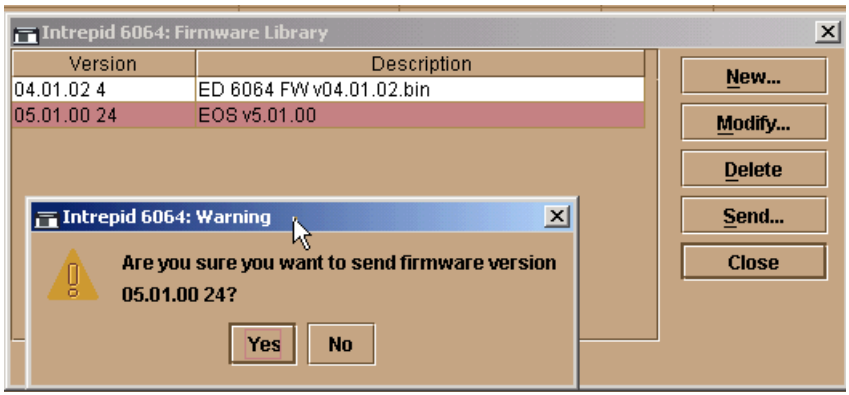


*Figure 4-131 Send firmware download confirmation prompt*

The send function verifies the existence of certain director conditions before the download process is initiated. If an error occurs, a message is displayed indicating the problem must be fixed before the firmware is downloaded. Conditions that terminate the download process include these:

- There is a redundant CTP2 card failure.
- The firmware version is being installed to the director by another user.
- The director-to-EFC Server link is down.

Select **Yes** if all is satisfactory.

As the download begins, a *Writing data to FLASH* message is displayed at the top of the window, followed by a *Sending Files* message. This message remains as the progress bar travels across the window indicating the percent completion of the download. The bar progresses to 50% when the last file is transmitted to the first CTP2 card. The bar remains at the 50% point until the director performs an IPL (indicated by an IPLing message).

During the IPL, the director-to-EFC Server link drops momentarily and the following occurs at the Product Manager:

- As the network connection drops, the ED-6064 Status table turns yellow, the Status field displays No Link, and the State field displays a message stating the reason for this.
- In the Product View, the director icon displays a grey square, indicating that the director status is unknown.
- Illustrated FRUs in the Hardware View disappear, and appear again as the connection is re-established.
- After the IPL, a *Synchronizing CTPs* message displays. This message remains as files are transmitted to the second CTP card and the progress bar travels across the window to 100%. When the download reaches 100%, a *Send firmware complete* message is displayed as shown in Figure 4-132.

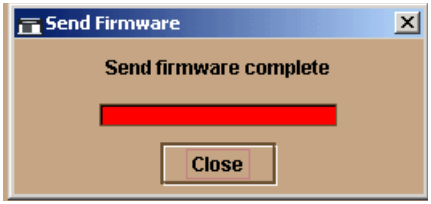


*Figure 4-132 The firmware download progress menu*

The firmware update is now complete and normal service will be resumed.

## 4.13 BladeCenter interoperability

Before discussing the configuration details to successfully merge the IBM BladeCenter into a McDATA fabric, we will first identify the specific parameters critical to successfully merging the two fabrics and detail some limitations in an interop mode fabric.

Table 4-1 lists the configuration options that must be true on each switch for the fabric to be able to merge successfully.

*Table 4-1 BladeCenter and McDATA ED-6064 parameters*

| Configuration options | BladeCenter switch | McDATA fabric |
|---|---|---|
| Recommended Firmware levels | 1.4.0.49 or later | 5.01 or later |
| Domain IDs allowed range | **97 - 12**7 | **1 - 31** |
| Assign Static Domain ID | Domain ID Lock = True | Insistent |
| IO Stream Guard | Disabled On E Port | Not Applicable |
| R_A_TOV | **10000** milliseconds | 100 (tenths of a second) |
| E_D_TOV | **2000** milliseconds | 20 (tenths of a second) |
| Principal Priority Value | 254 | Principal |
| Interop Mode | Default [SW2 Standard] | Open Fabric 1.0 |
| Zone Configuration | WWPN based only | WWPN based only |
| Default Zone Set | disabled / inactive | disabled / inactive |

### Configuration limitations

These are some of the limitations that currently exist:

- When merging McDATA and BladeCenter fabrics, a maximum of 31 interconnected switches per fabric are supported.
- Only WWPN based zoning is supported. Port based or WWNN based zoning is not supported in an interop fabric.
- On McDATA switches with operating mode set to open fabric 1.0, the default zone set needs to be disabled, and it is not supported.

### Important guidelines

These are some of the important guidelines to adhere to:

- While performing the zoning configuration, do not use IBM BladeCenter SAN Utility and EFCM 7.1 *simultaneously*. Using multiple management applications at the same time may cause corruption.
- It is strongly recommended to download EOS 5.01 on ED-6064 and V1.4.0.49-0 on the BladeCenter FC switch because of critical fixes available in the current firmware release on both products.
- The I/O StreamGuard feature should be enabled on the E_Port for the IBM BladeCenter switch.

## 4.13.1 Configuration process

The integration of the BladeCenter into a McDATA fabric is a disruptive process and requires configuration changes such as unique domain IDs, time out values, Domain ID (insistent/lock), interop mode to be done individually on each switch.

**Note:** The two fabrics will segment If the time out values (R_A_TOV, E_D_TOV) are not identical, the Domain ID is not unique, and both switches have conflicting zone set configuration active.

### ED-6064 configuration

In the following examples EFCM 7.1 is used to configure the ED-6064 operating parameters and configure zoning.

To change the ED-6064 switch and fabric operating parameters, the ED-6064 should be set to offline state from the EFC product manager by selecting **Maintenance —> Set Online state —> Set Offline** as shown in Figure 4-133.

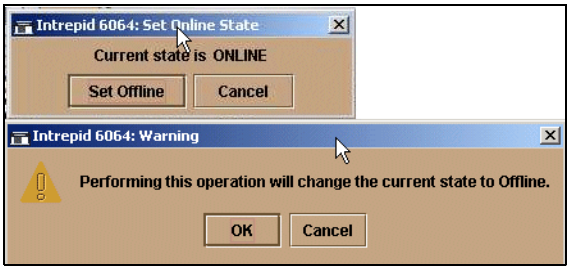


*Figure 4-133 Set the ED-6064 to Offline state*

Once the switch is in offline mode, configure the switch and fabric operating parameters, and take the following steps:

1. **Disable Default Zone:** Select **Configure —> Advanced Zoning —> Configure Default Zone** and select **Start** to disable the default zone as shown in Figure 4-134.

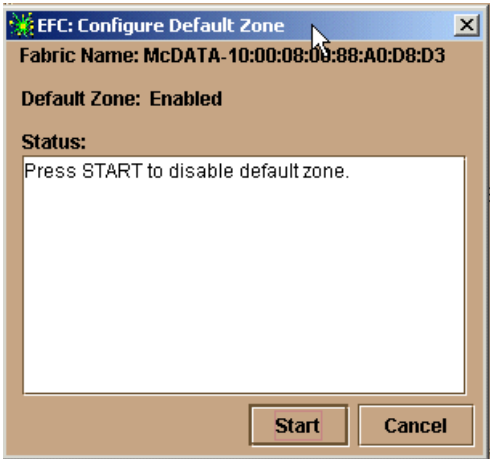


*Figure 4-134 Disable default zone on ED-6064*

2. **Switch Operating Parameters** Configure **Domain ID** =1 and check the **Insistent** option so that it retains the Domain ID = 1 following the fabric reconfiguration. By manually configuring the domain IDs on every switch in the fabric, the system administrator eliminates any possibility of a fabric segmentation due to a domain ID conflict. Also ensure that the **Rerouting delay** option is disabled, as it is not supported in an interop mode fabric.

   This is shown in Figure 4-135.

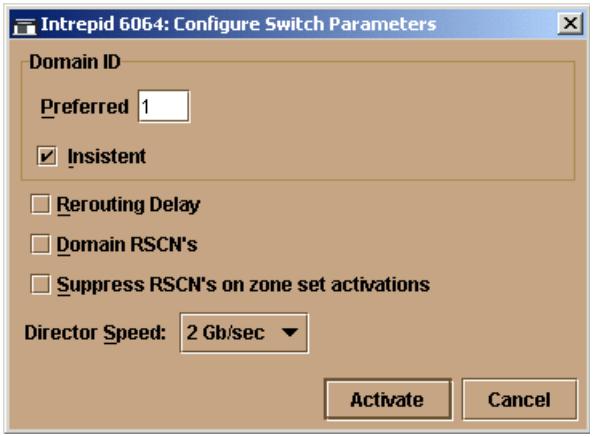


*Figure 4-135 Configure Switch Parameters*

3. **Fabric Operating Parameters**: Configure the ED-6064 fabric operating parameters by selecting **Configure -> Operating Parameters -> Fabric Parameters** as shown in Figure 4-136.

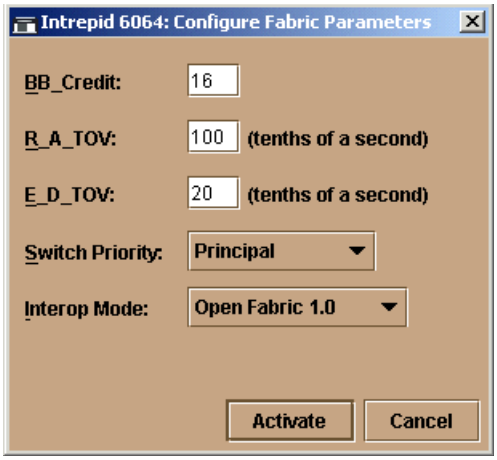


*Figure 4-136 The ED-6064 fabric operating parameters*

4. If an active zone configuration exists on the two switches, then ensure that the zoning configuration on the two connecting fabrics is compatible; otherwise the E_Port will segment due to a zone merge failure.
5. Set switch to **Online** state.
6. At this point the BladeCenter can be connected on any of the available ports of the ED-6064.

7. Verify that the BladeCenter and the ED-6064 merged successfully using the topology menu as shown in Figure 4-137. This can also be verified from the SANUtility and CLI interfaces.

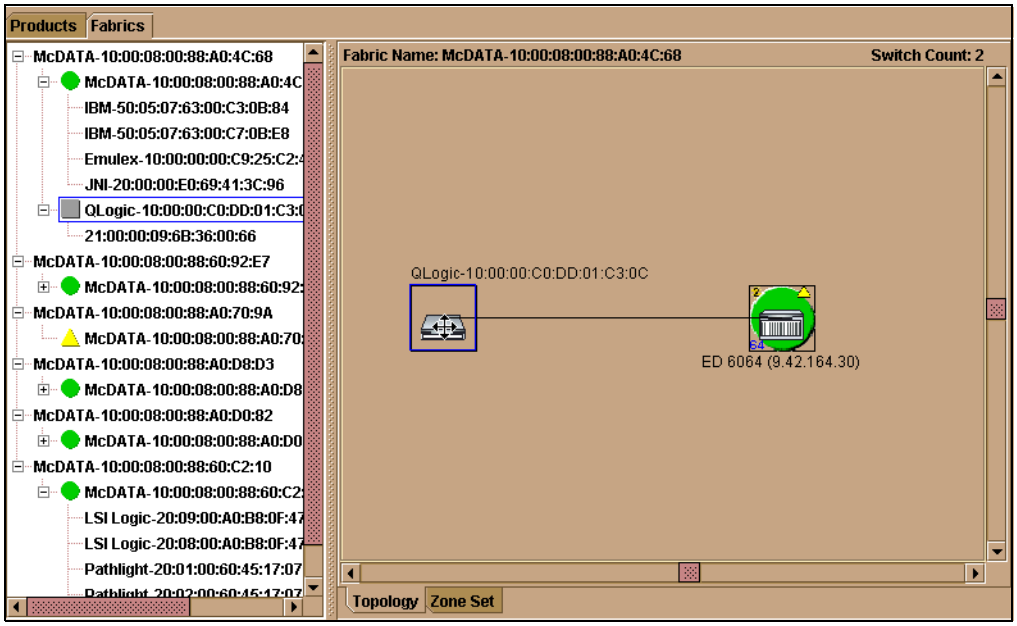


*Figure 4-137 BladeCenter and ED-6064 fabric*

## Configuring zoning from the EFCM

After verifying that the two fabrics merged successfully and that all the devices and ports have become active in the interop mode fabric, proceed to configure and activate the zoning configuration using EFCM 7.01. This topic and more details about zoning are covered in 4.7, “Understanding the McDATA zoning concepts” on page 522.

5

# Implementing BladeCenter

In this chapter we provide the basic details to successfully implement the IBM BladeCenter in a Fibre Channel based Open IBM SAN environment. This includes information on BladeCenter interoperabilty with multi-vendor Fibre Channel switches. We also explain the various procedures and tools available to perform device configuration and management in a multi-vendor fabric.

**Note:** It is not our intent in this redbook to provide an in-depth discussion of the BladeCenter. We will show the steps that we took to implement the BladeCenter in our environment. In the individual SAN fabric product chapters, we show the steps that we took to connect those products to the BladeCenter.

## 5.1 BladeCenter overview

BladeCenter is based on a blade architecture, which is a highly modular design that converges computing resources into cost-effective, high-density enclosures.

The major benefits available from BladeCenter are:

- Simplified management
- Fast installation and redeployment
- Modular scalability
- High availability

The IBM BladeCenter uses the Intel® Xeon™ line of processors providing the latest processor technologies. BladeCenter has the potential to dramatically reduce your IT infrastructure costs.

In Figure 5-1 we show how BladeCenter fits into the fabric.

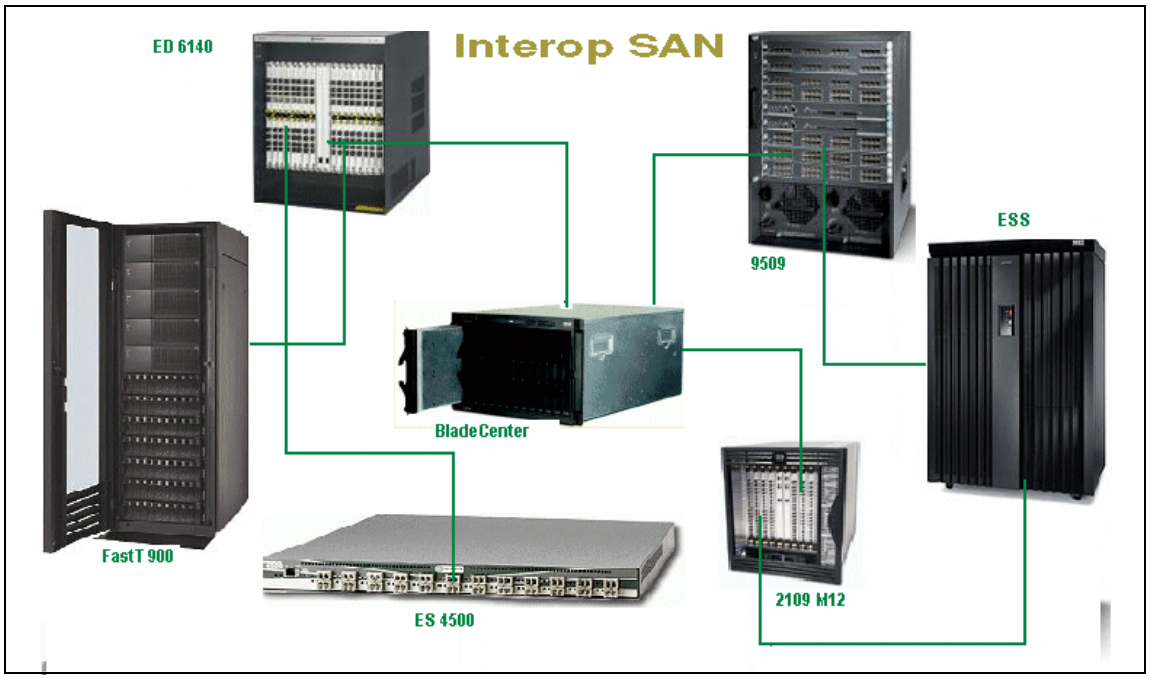


*Figure 5-1 The IBM BladeCenter integrated in a multi-vendor fabric*

## 5.2 IBM BladeCenter in a SAN

The IBM @server™ BladeCenter Fibre Channel switch can be deployed in any enterprise SAN. The BladeCenter Fibre Channel switch has dual ports for redundancy and can connect internally to the FC host bus adapter (daughter cards) installed on the blade servers. The blade server then uses the IBM @server™ BladeCenter HS20 dual port Fibre Channel host bus adapter (F_Port) to connect to the Fibre Channel switch module at a fixed speed of 2 Gb/s and fixed F_Port mode. The port speed of 2 Gb/s and F_Port mode cannot be changed on the host bus adapter or on the switch internal ports.

The Fibre Channel link on the E_Port enables high speed connectivity at 2 Gb/s speed and it can also auto negotiate to 1 or 2 Gb/s by auto-sensing the speed of the connecting device on the other end of the E_Port. It supports multiple topologies such as Point-to-Point, Fabric, and Arbitrated loop. Through Fibre Channel SANs, the blade server can connect to enterprise level storage devices over multi-vendor fabrics allowing the ability to share storage among many servers with high availability.

The host bus adapter with dual F_Ports that sits on the blade server is shown in Figure 5-2.

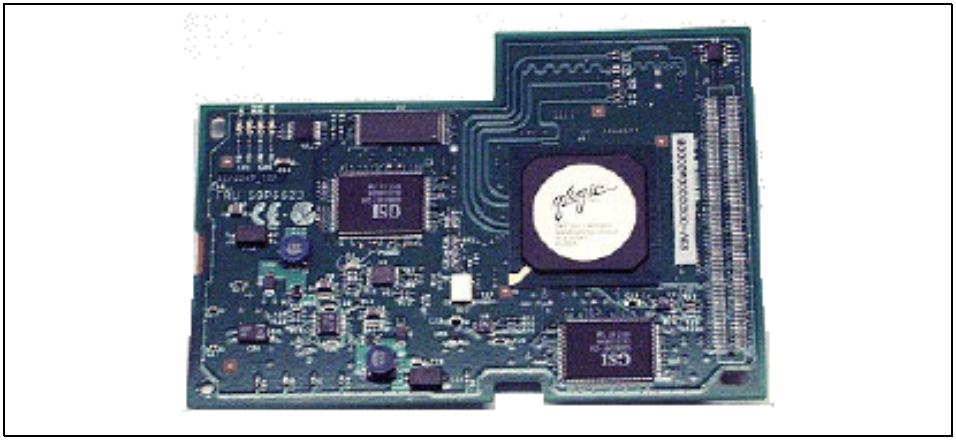

*Figure 5-2 Fibre Channel host bus adapter*

Since there are no PCI slots available in the blade servers, the Fibre Channel host bus adapter uses the same connector to which the local IDE drive usually connects. The blade server can either use the local IDE drive or it can boot from the remote disks on ESS or FAStT storage subsystems over the SAN.

## 5.2.1 BladeCenter switch ports and port types

The port icons on the faceplate display reflect the current information for each port on a switch module according to the type of port view selected in the View menu. To view different types of port information on the port icons, open the View menu and select View Port Modes, View Port States, View Port Speed, or View Port Media. Port view information is displayed within the port icon itself, such as 1G, 2G, IA, TL, and F. To view port information without selecting a port, rest the cursor over a port and a small popup text box displays port information.

On the faceplate display, there are 14 internal ports (numbered 1-14) and 2 external ports (numbered 1-2). Internal ports are fixed at 2 Gb/s speed and F_Port mode *only*, but can be configured online or offline using the internal Port Properties window. The external ports (which are not to be confused with *expansion* ports) have another Port Properties window in which to configure port states, speeds, modes, and media.

**Note:** The two external port icons are numbered 1 and 2 in the faceplate display, but are indexed as port #0 and port #15, respectively.

### *F_Port*

A Fabric Port that allows an end device (N_Port) connectivity.

### *FL_Port*

An Fabric Loop port (NL_Port). Supports a loop of up to 126 public devices.

### *TL_Port*

A Translated Loop Port that supports a loop of up to 126 private devices (no FLOGI capability). Supports a loop of up to 126 private target or initiator devices capable of communicating with up to 63 public or private devices on other ports.

### *G_Port*

A Generic Port that can sense the connecting device type and auto configures as an F_Port or an E_Port.

### *GL_Port*

A Generic Loop Port that auto detects and configures by itself as an F_Port, FL_Port, or an E_Port.

### *Don*

A Donor Port, it allows its buffer credits to be allocated to another port. It is designated as unused so that its buffer credits can be used by another port to increase the distance data can be transferred.

***E_Port***

An expansion port (ISL) between two switches.

### 5.2.2 I/O StreamGuard

When enabled, the I/O StreamGuard port suppresses the generation of RSCN messages for normal fabric events such as plugging and unplugging ports, and name server registrations. I/O StreamGuard also ensures that the port is not returned in any N_Port name server queries.

The I/O StreamGuard feature can be enabled on the expansion (E_Port) and internal port (F_Port) to suppress RSCN events. The I/O StreamGuard feature should always be disabled on the expansion port in order to receive RSCN messages and also to advertise the local name server table in the fabric.

## 5.3 BladeCenter management

BladeCenter can be managed via the telnet session using a hyperteminal application and also from the BladeCenter SAN Utility graphical interface. We will demonstrate BladeCenter configuration and management using both methods.

### 5.3.1 Telnet access

The BladeCenter CLI (Command Line Interface) is accessed using the telnet session. The procedure to access the BladeCenter using the hyper terminal application is shown in Figure 5-3.

Type the **IP address** of the BladeCenter in the **Host address** box and select **TCP/IP (Winsock)** from the **Connect using:** drop down list.

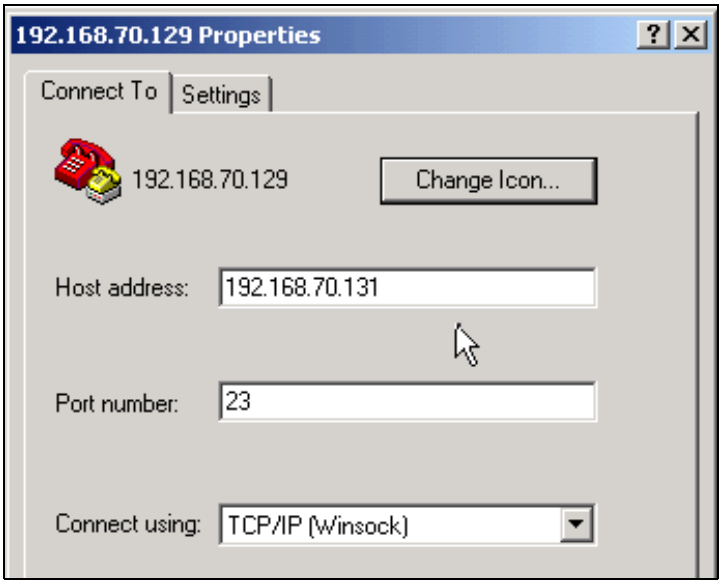


*Figure 5-3 The hyper terminal setting*

At the login prompt, type the default user-id:*USERID,* and password: *passw0rd* for administrator level access. The user-id and password are case sensitive.

## 5.3.2 Command line interface

The first screen that appears after logging into the BladeCenter lists the IP address, MAC address, WWPN, switch firmware version and diagnostic status.

The IBM BladeCenter Fibre Channel switch initial login shell interface is shown in Figure 5-4.

```
      ******************************************************
      *                                                    *
      *       Command Line Interface SHell  (CLISH)        *
      *                                                    *
      ******************************************************

      SystemDescription  IBM BladeCenter(TM) 2-port Fibre Channel Switch Mod

      Eth0NetworkAddress 192.168.70.131 (use 'set setup system' to update)
      MACAddress         00:c0:dd:01:c3:0b
      WorldWideName      10:00:00:c0:dd:01:c3:0c
      SymbolicName       FCSM
      SWImageVersion     V1.4.0.49-0
      SWImageBuiltDate   Tue Jun  3 16:52:53 2003
      DiagnosticsStatus  Passed
      SecurityEnabled    False

  Alarms history ...
  ------------------

  [1][Wed Jan 06 01:53:12.349 1988][A4101][0xdd01c30b.329][cmon: unable to se
w addr]

FCSM: admin>
```

*Figure 5-4 Command Line Interface shell*

The switch configuration can be retrieved by using the various **show** command options available as shown in Figure 5-5.

```
FCSM: admin> show

  Usage: show about       config     log        panel   slot      version
              alarm       domains    lsdb       perf    steering
              blade       donor      mem        port    switch
              broadcast   fabric     ns         post    topology
              chassis     interface  pagebreak  setup   users
```

*Figure 5-5 The show command options*

The **set** command options are used to change the configuration and are listed in Figure 5-6.

```
FCSM: admin> set

  Usage: set alarm     config       port
             beacon    log          setup
             blade     pagebreak    switch
```

*Figure 5-6 The set command options*

The **show switch** output lists vital configuration and diagnostic data such as switch WWN, domain ID, active image, last restart reason and the operational state, as shown in Example 5-1.

*Example 5-1 The show switch output*

```
show switch
  Switch Information
  ------------------
  SymbolicName                      FCSM
  SwitchWWN                         10:00:00:c0:dd:01:c3:0c
  SwitchType                        SwitchBlade
  PROMVersion                       V1.4.0.2-0 (Tue Jun  3 16:52:53 2003)
  CreditPool                        0
  DomainID                          125 (0x7d)
  FirstPortAddress                  7d0000
  FlashSize - MBytes                128
  LogLevel                          Critical
  MaxPorts                          16
  NumberOfResets                    39
  ReasonForLastReset                NormalReset
  SWImageVersion (1) - build date   V1.4.0.42-0 (Mon Feb  3 22:33:00 2003)
  SWImageVersion (2) - build date   V1.4.0.49-0 (Tue Jun  3 16:52:53 2003)
  ActiveConfiguration               default
  ActiveSWImage                     2
  AdminState                        Online
  AdminModeActive                   False
  BeaconOnStatus                    False
  OperationalState                  Online
  PrincipalSwitchRole               False
  BoardTemp (1) - Degrees Celsius   45
  BoardTemp (2) - Degrees Celsius   45
  SwitchDiagnosticsStatus           Passed
  SwitchTemperatureStatus           Normall
```

Before you change the default configuration in the BladeCenter, verify the existing configuration from the command line by using the **show config switch** command, as shown in Example 5-2.

*Example 5-2 The show config switch output*

```
FCSM: USERID > show config switch
  Configuration Name: default
  Switch Configuration Information
  --------------------------------
  AdminState            Online
  BroadcastEnabled      True
  InbandEnabled         False
  DomainID              1 (0x1)
```

```
  DomainIDLock          False
  SymbolicName          FCSM
  R_T_TOV               100
  R_A_TOV               10000
  E_D_TOV               2000
  FS_TOV                5000
  DS_TOV                5000
  PrincipalPriority     254
  ConfigDescription     IBM BladeCenter(TM) 2-port Fibre Channel Switch Module
  ConfigLastSavedBy     Initial
  ConfigLastSavedOn     Initial
```

To be able to merge the BladeCenter switch into a multi-vendor fabric the R_A_TOV, E_D_TOV values have to be same on all the switches in a fabric and the domain ID must be unique and should be locked, and the principal switch priority should be configured to the lowest priority value of 254 so that the IBM BladeCenter switch is *always* the subordinate switch in a fabric where it is interoperating with multi-vendor SAN switches. In the following topics we cover the procedures to change the BladeCenter switch configuration so that it can merge successfully.

### 5.3.3 Initial configuration on BladeCenter

The initial configuration to ready the BladeCenter to interconnect to our multi-vendor fabric is done by entering into Admin mode and from the Edit Configuration prompt as shown in Example 5-3:

**Tip:** The values in square brackets (such as [Online]) are default values. The values that are highlighted in bold (**Domain ID** and **Domain ID lock**) will be changed from the default configuration.

*Example 5-3 The set config switch command*

```
Login: USERID
Password: xxxxxxxx
FCSM: USERID> admin start
FCSM (admin): USERID> config edit
FCSM (admin-config): USERID> set config switch

A list of attributes with formatting and current values will follow. Enter a
new value or simply press the ENTER key to accept the current value. If you
wish to terminate this process before reaching the end of the list press 'q' or
'Q' and the ENTER key to do so.
AdminState         (1=Online, 2=Offline, 3=Diagnostics) [Online ]
  BroadcastEnabled   (True / False)  [True ]
  InbandEnabled      (True / False)     [True ]
```

```
  DefaultDomainID    (decimal value, 1-239) [1 ] 126
  DomainIDLock       (True / False)   [False ]  True
  SymbolicName       (string, max=32 chars)  [FCSM ]  FCSM
  R_T_TOV            (decimal value, 1-1000 msec)  [100 ]
  R_A_TOV            (decimal value, 100-100000 msec) [10000 ]
  E_D_TOV            (decimal value, 10-20000 msec) [2000 ]
  FS_TOV             (decimal value, 100-100000 msec)  [5000 ]
  DS_TOV             (decimal value, 100-100000 msec)  [5000 ]
  PrincipalPriority  (decimal value, 1-255)  [254 ]
  ConfigDescription  (string, max=64 chars)  [IBM BladeCenter(TM) 2-port Fibre
Channel Switch  Module]
Finished configuring attributes.
  This configuration must be saved (see config save command) and activated (see
config activate command) before it can take effect.
  To discard this configuration use the config cancel command.
FCSM2 (admin-config): USERID> config save
FCSM (admin): USERID> config activate
The configuration will be activated.  Please confirm (y/n): [n] y
```

### 5.3.4 BladeCenter SAN Utility and BladeCenter FabricView

The BladeCenter SAN Utility and BladeCenter Fabric View are a graphical user interface based tool that provides menus and radio buttons to manage BladeCenter switche(s) from a remote management station on Linux or Windows platforms. The BladeCenter SAN Utility allows you to configure switch IP address, load firmware, zone configuration and management, port configuration and so on. Multiple fabrics can be managed from a single management station by adding the devices IP address. BladeCenter SAN Utility displays the real time status of the fabric and any changes in the operational fabric are detected. The version of BladeCenter SAN Utility that is used in our environment is release 1.02.20.

The BladeCenter Fabric View application displays port performance using graphs. Fabric View provides a method to visually monitor the real-time traffic for each port on a switch. Traffic for a port is displayed in its own graph that is continually updated to reflect changes as they occur, and is based on the number of kilobytes (Kbytes), or on the number of frames that pass through that port per second. One kilobyte is equal to 1024 bytes per second. We do not cover Fabric View in this redbook, however the install of BladeCenter SAN Utility will result in its installation.

**Important:** Before Installing the Fibre Channel Switch Module, be sure that the IBM @server™ BladeCenter Management Module firmware is build ID BRET30A or later. The Fibre Channel Switch Module might not work properly with an earlier version of the management module firmware. Use the management module Web-based user interface to verify the firmware build ID.

More information and the update is available at this Web site:

http://www-3.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-45487

The BladeCenter SAN Utility is available on the Web and can be downloaded from the Web site:

http://www-3.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-46373

We saved the `windows_1.02.20.exe` to disk, and we show how to install it in the sections that follow.

### 5.3.5 Installing BladeCenter SAN Utility and BladeCenter Fabric View

The installation of BladeCenter SAN Utility and BladeCenter Fabric View is a straightforward process and is shown below:

We selected the executable file from **Start-—>Run**, then browsed our way to the file and selected **OK** as shown in Figure 5-7.

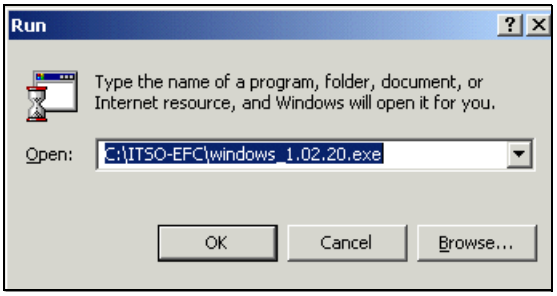


*Figure 5-7 The BladeCenter SAN Utility installer file*

A window is displayed asking you to select the language. We chose English (the default language) and clicked **OK**.

The **Introduction window** appears. Click **Next** as shown in Figure 5-8.

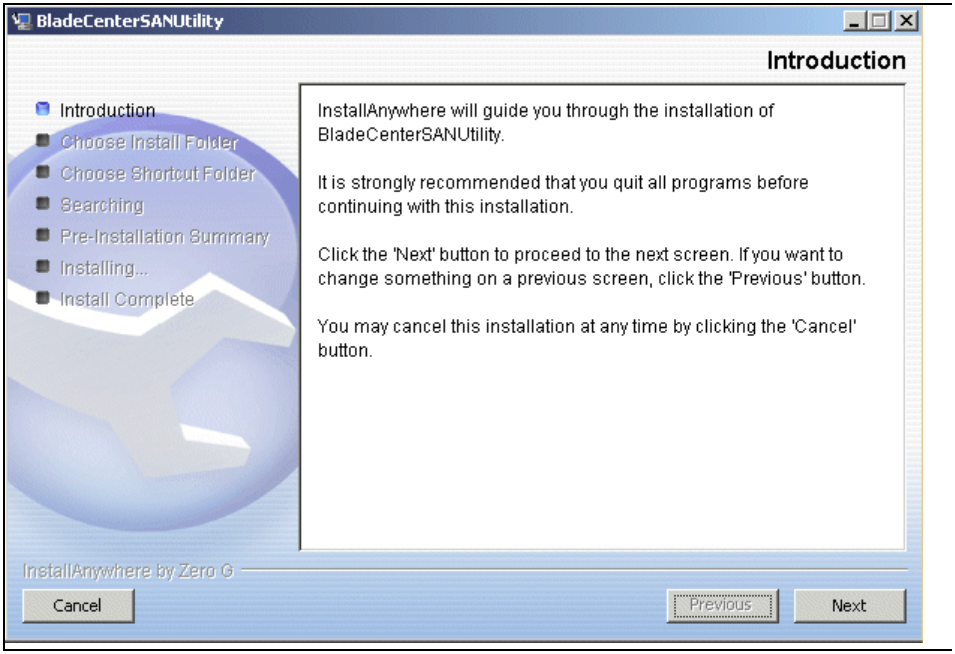


*Figure 5-8 Introduction window*

The **Choose Install Folder** window is displayed, where we can confirm the default installation path or change to a different path if needed. Click **Next** when the path has been chosen as shown in Figure 5-9.

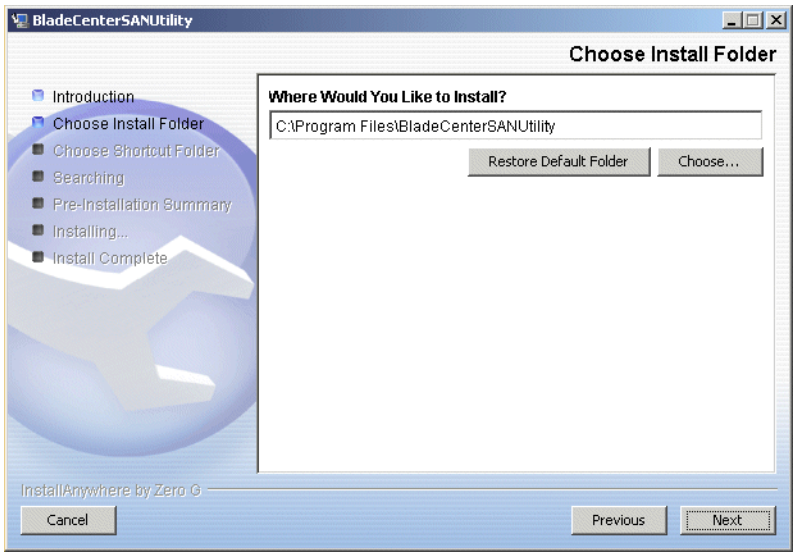


*Figure 5-9 Default installation path*

The **Pre-Installation Summary** window is displayed so we can review the install summary and confirm it. Click the **Install** button to initiate the installation process as shown in Figure 5-10 if all is acceptable.

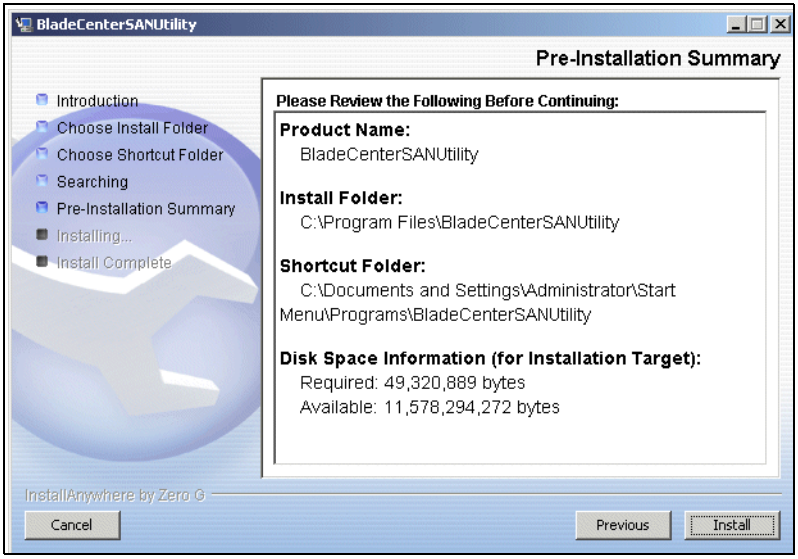


*Figure 5-10 Pre-installation summary menu*

The **Install Complete** window is displayed after the installation process has completed as shown in Figure 5-11.

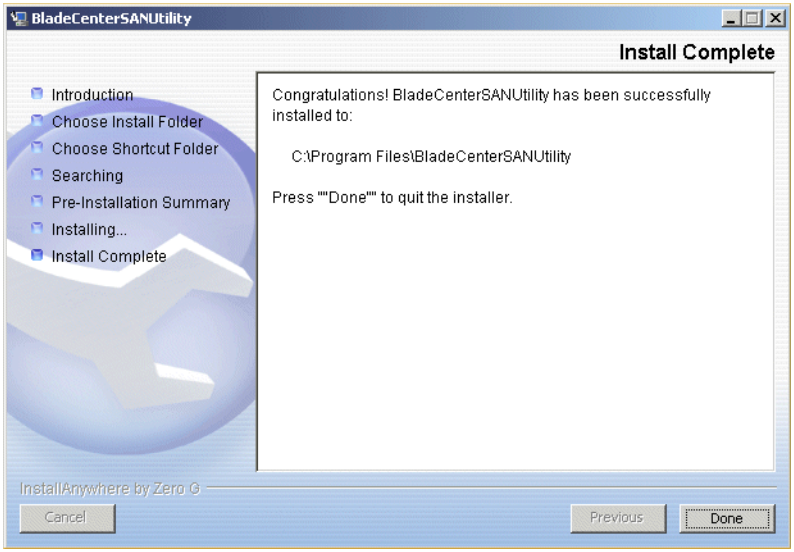


*Figure 5-11 Successful installation*

Check for successful installation of the BladeCenter SAN Utility and BladeCenter Fabric View utilities as shown in Figure 5-12.

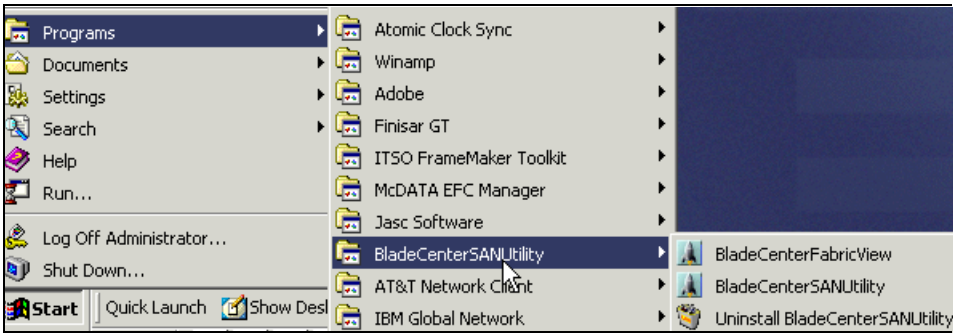


*Figure 5-12 The BladeCenter SAN Utility and Fabric View installed path*

Throughout this chapter we will use the BladeCenter SAN Utility to demonstrate the configuration and management of the BladeCenter.

## 5.3.6 Adding the new fabric

Once the IBM BladeCenter SAN Utility is installed, it is now available to add switches and manage the fabric. From the management station, select **Start—> Programs—> BladeCenter —> BladeCenterSANUtility** to open the application.

The **Topology** view is displayed, and selecting the **Add** option will allow you to add the new switch or multiple switches to be managed. This is done by assigning a name to the view, the IP address of the switch and supplying userid / password, from the **Add a New Fabric** menu as shown in Figure 5-13.

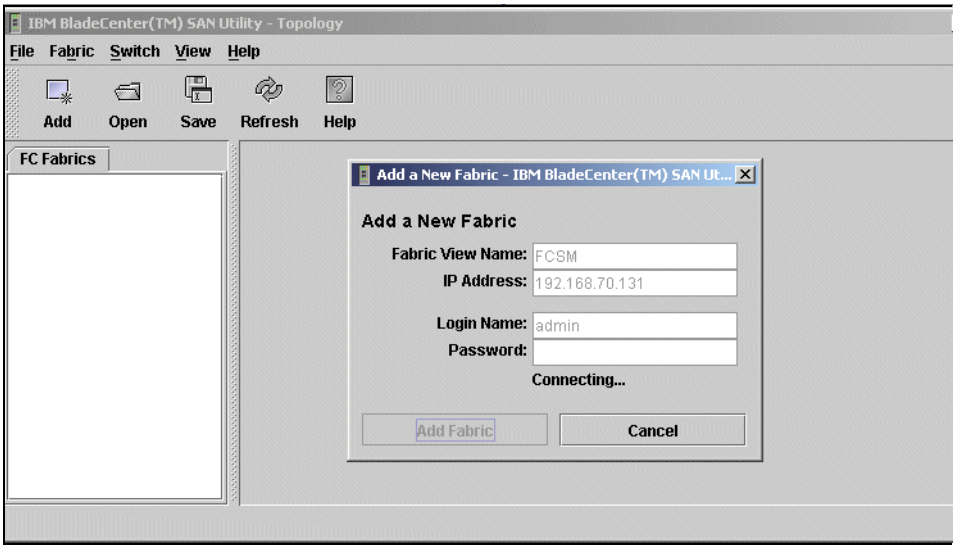


*Figure 5-13 Topology view*

Once the switch is discovered by the utility, it creates a green icon resembling a cloud under the **FC Fabrics** tab, and the **Faceplate** window is displayed. The cloud indicates that the fabric has been discovered and shown on the upper left corner of the topology view. This is shown in Figure 5-14 to the left of **FCSM3.** If the fabric could not be connected to the cloud, it would have remained blue in color.

If we select the green cloud by clicking on it, the switch icon is displayed below the fabric icon. In our environment we will use one BladeCenter switch to merge with the fabric, and one blade server is installed in Bay 1. There is one E_Port and one F_Port active as shown in Figure 5-14.

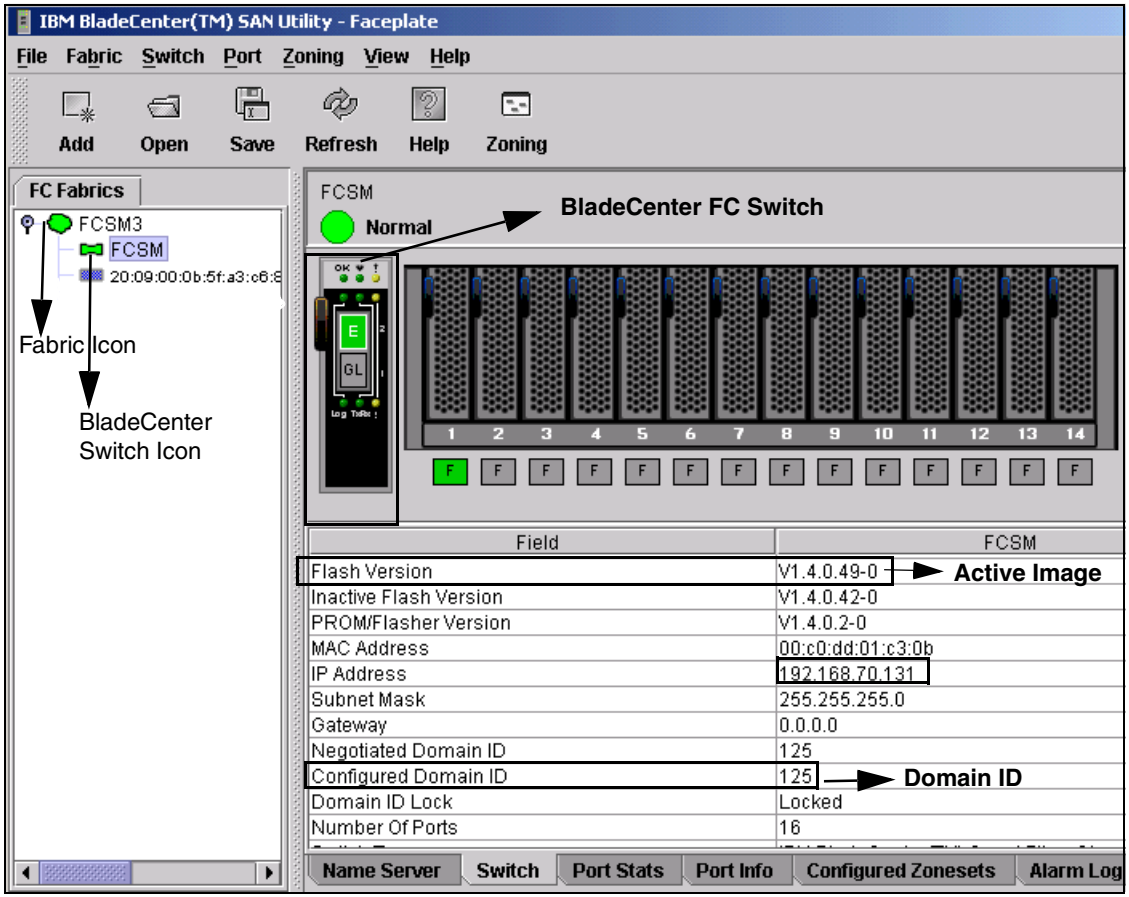


*Figure 5-14 BladeCenter Fabric and Switch view*

In Figure 5-14 we have selected the switch information view. As the arrows indicate, this view lists information such as the active image (firmware), Configured/Negotiated Domain ID and IP address. There are two external ports on the switch and 14 internal ports for the blade server. The gray icon under the green icon indicates that it is a non-native switch.

The expansion port has flexible properties, and it can be configured as F_Port, FL_Port, TL_Port, G_Port, GL_Port, Don types.

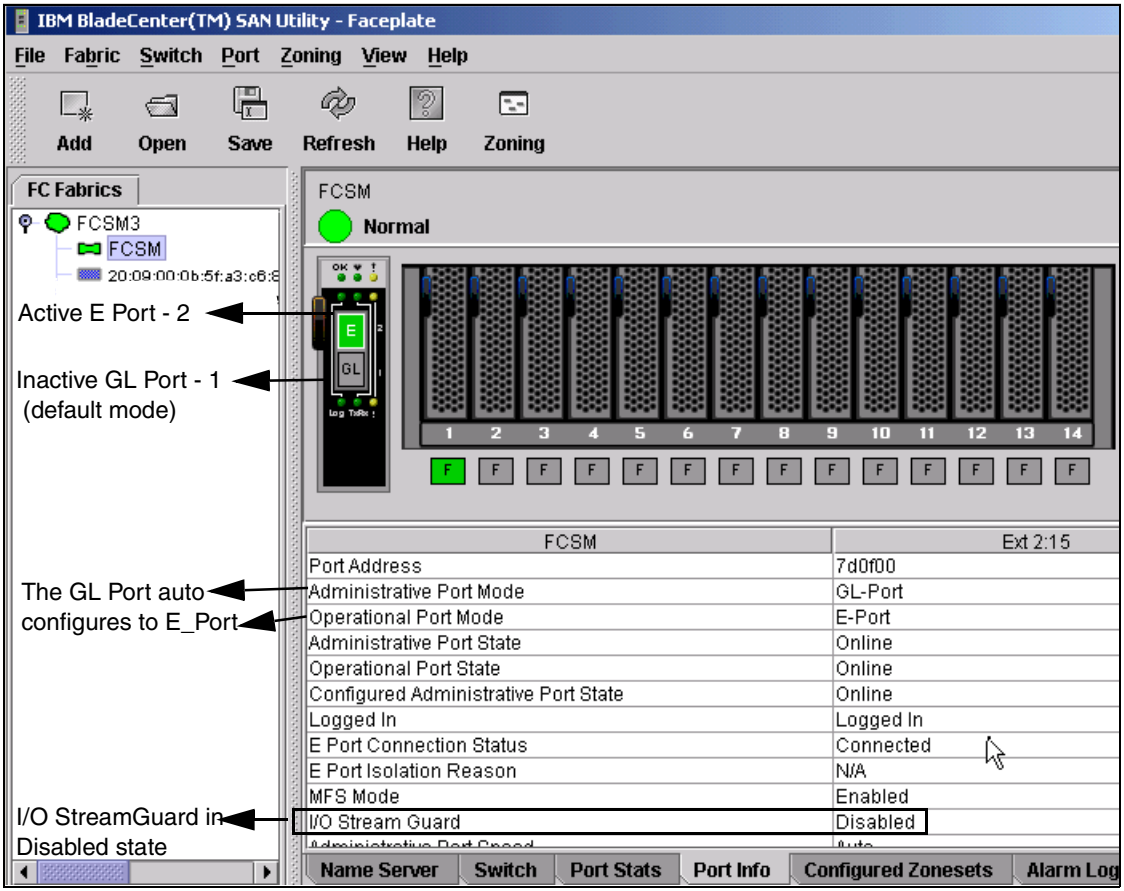


*Figure 5-15 BladeCenter Switch E_Port Info*

By default, the external port mode is set as **GL** mode. If connected to another switch, it automatically configures the GL_Port to E_Port mode as shown in the Operational Port State. The expansion port (the IOStreamGuard) is in disabled state as shown in Figure 5-15.

The blade server is installed in Bay 1 or slot 1 in the BladeCenter chassis. The blade server has fixed F_Port mode and the speed is locked at 2 Gb/s, the F_Port states are online, offline, or test mode.

F_Port

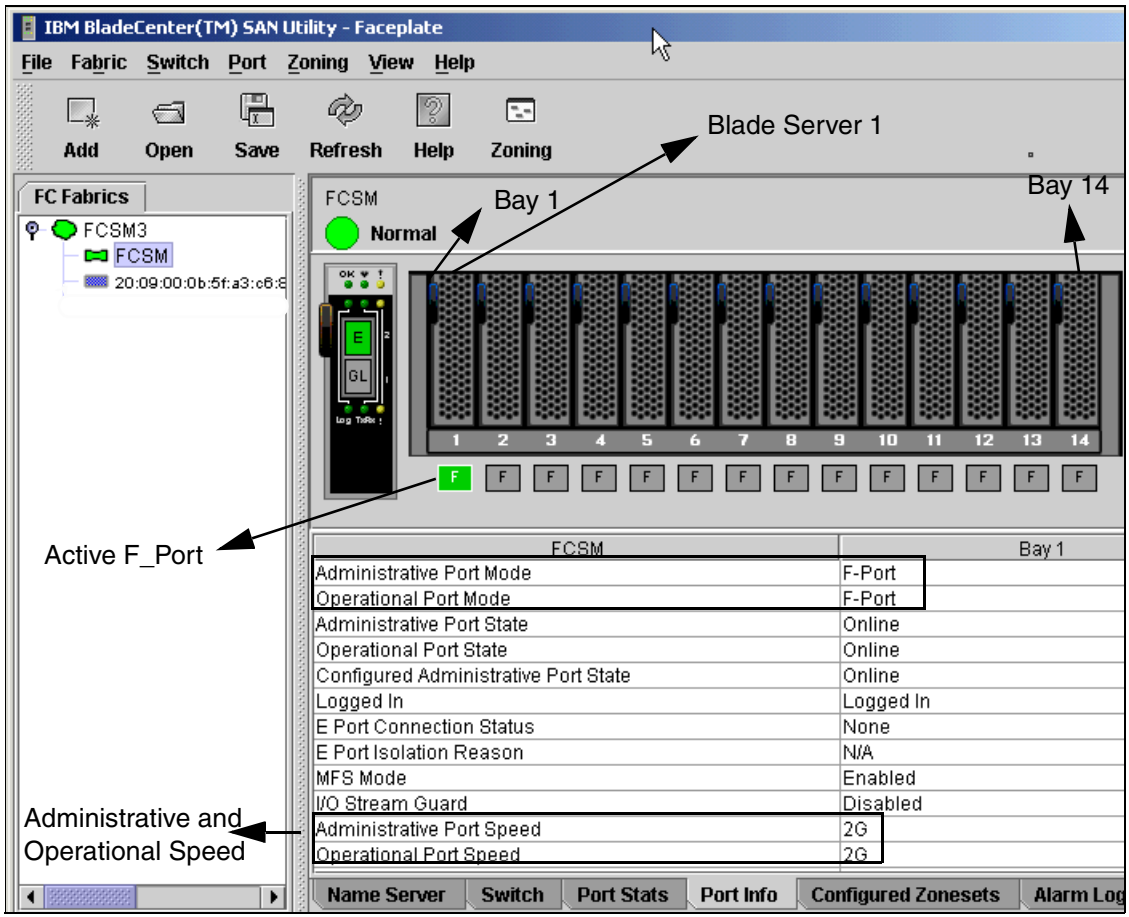


| FCSM | Bay 1 |
|---|---|
| Administrative Port Mode | F-Port |
| Operational Port Mode | F-Port |
| Administrative Port State | Online |
| Operational Port State | Online |
| Configured Administrative Port State | Online |
| Logged In | Logged In |
| E Port Connection Status | None |
| E Port Isolation Reason | N/A |
| MFS Mode | Enabled |
| I/O Stream Guard | Disabled |
| Administrative Port Speed | 2G |
| Operational Port Speed | 2G |

*Figure 5-16 BladeCenter SAN Utility F_Port Info*

With the BladeCenter SAN Utility, various configuration and management options are available, such as switch/port/zone configuration and a firmware download option. The various procedures available are demonstrated in the following sections, where we show how to set options at the switch, fabric, and port levels.

In Figure 5-17, the name server information is listed. To view the name server information, select the **Name Server** tab from the fabric view. The fabric icon should be selected in order to view the fabric wide name server list.

If the switch icon is selected, then the name server information is limited to the blade server ports only.

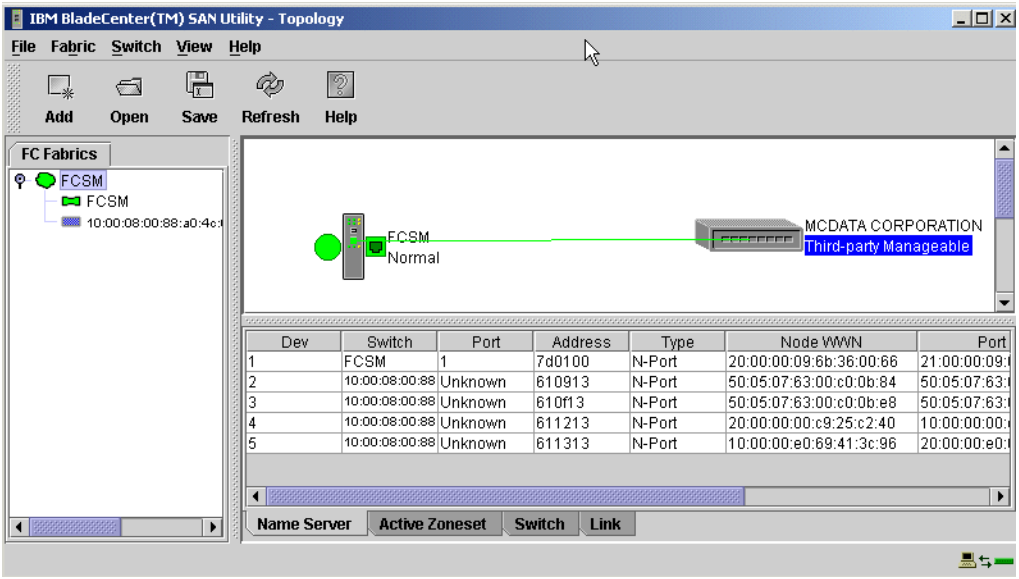


*Figure 5-17 Name Server and the fabric view*

The display shows the name server table of the entire fabric, which includes the devices attached to the BladeCenter and the adjoining vendor switch forming a fabric.

To view the E_Port properties menu, first select the BladeCenter switch icon then select the port by clicking the E_Port as shown in the faceplate menu. From the Face Plate menu bar, select **Port —> Port Properties as** shown in Figure 5-18.

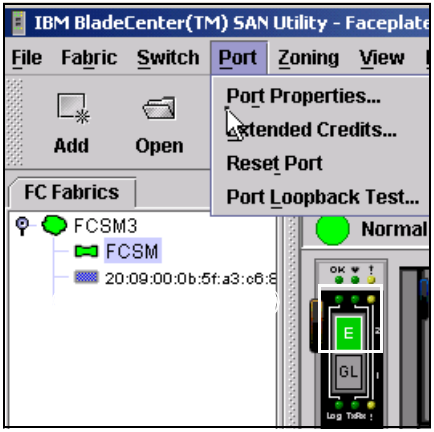


*Figure 5-18 E_Port properties select menu*

In Figure 5-19 the E_Port properties are shown.

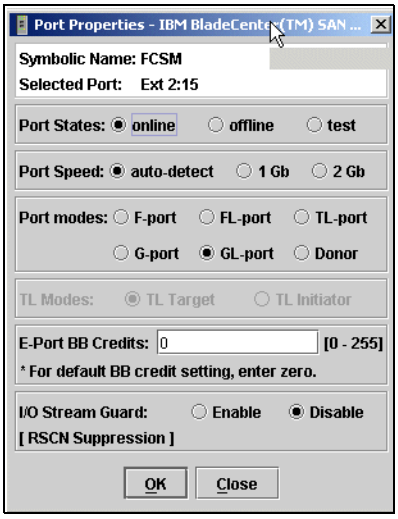


*Figure 5-19 E_Port properties menu*

In Figure 5-20 we show how to select the Switch Properties menu.

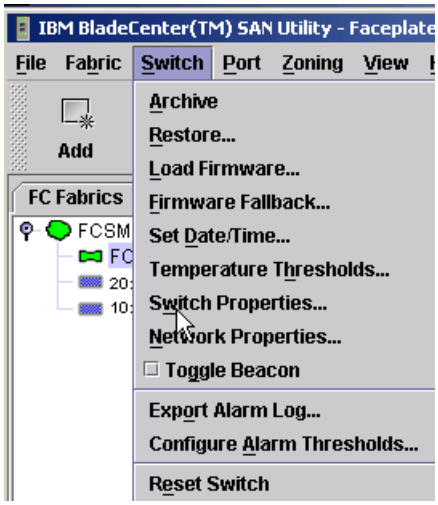


*Figure 5-20 Switch Properties menu*

After selecting the **Switch Properties...** from the drop down list, the properties window is displayed, giving the configuration options as shown in Figure 5-21.

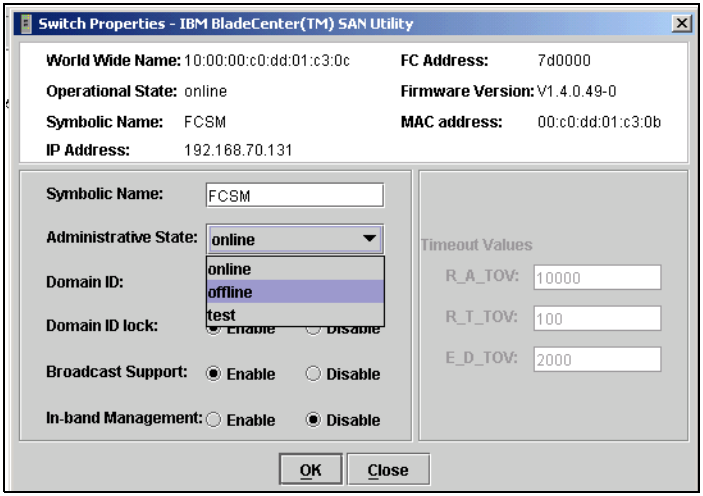


*Figure 5-21 BladeCenter Switch Properties View*

## Switch properties

Most of the fields are self explanatory, but one of special interest is the **Domain ID lock**.

Switches come from the factory with the domain IDs unlocked. This means that if there is a domain ID conflict in the fabric, the switch with the highest principal priority, or the principal switch, will reassign any domain ID conflicts and establish the fabric. If you lock the domain ID on a switch and a domain ID conflict occurs, the switch with the higher WWN will isolate as a separate fabric and the Logged-In LEDs on both switches will flash to show the affected ports.

If you connect a new switch to an existing fabric with its domain ID unlocked, and a domain conflict occurs, the new switch will isolate as a separate fabric. However, you can remedy this by resetting the new switch or taking it offline then back online. The principal switch will reassign the domain ID and the switch will join the fabric.

**Note:** Domain ID reassignment is not reflected in zoning that is defined by domain ID and port number pair. You must reconfigure zones that are affected by domain ID reassignment.

## 5.4 Zoning

A zone is a named group of ports or devices that can communicate with each other. Zoning is available for all FC-SW-2 switches in the fabric. Zoning enables you to divide the ports and devices of the fabric into zones for more efficient and secure communication among functionally grouped nodes.

Managing zones involves creating a zone and adding and removing member port/devices. Zone members can communicate only with members of the same zone. Zones can overlap; that is, a port or device can be a member of more than one zone. Membership in a zone can be defined by port number, device Fibre Channel address, or device world wide name (WWN).

The zoning concepts are such that:

- Orphan zones and aliases are allowed.
- You cannot edit an active zone set.
- The three types of zones are Soft, VPF (hard zone), and ACL (hard zone).
- The entire zoning database can be saved to either temporary or permanent memory.
- The zoning limits for a fabric are:
  - Maximum number of zonesets: 256
  - Maximum number of zones: 256
  - Maximum number of aliases: 256
  - Maximum number of members per zone: 2000
  - Maximum number of members per aliases: 2000
  - Maximum total number of members: 2000

A zone is a named group of ports or devices that can communicate with each other. Membership in a zone can be defined by port number, device Fibre Channel address, or device World Wide Name (WWN). Zone members can communicate only with members of the same zone. Zones can overlap; that is, a port or device can be a member of more than one zone.

Three types of zones are supported. The following zone types define increasing restrictive levels of communication:

- Soft zone
- Access Control List (ACL) - hard zone
- Virtual Private Fabric (VPF) - hard zone

### 5.4.1 Soft zones

Soft zoning divides the fabric for purposes of controlling discovery. Members of the same soft zone automatically discover and communicate freely with all other members of the same zone.

The soft zone boundary is not secure; traffic across soft zones can occur if addressed correctly. Soft zones that include members from multiple switches need not include the ports of the inter-switch links. Soft zone boundaries yield to ACL and VPF zone boundaries. Soft zones can overlap; that is, a port can be a member of more than one soft zone. Membership can be defined by Fibre Channel address, port ID and domain ID, or worldwide name. Soft zoning supports all port modes.

### 5.4.2 Access Control List zones

Access Control List (ACL) zoning divides the fabric for purposes of controlling discovery and inbound traffic. ACL zoning is a type of hard zoning that is hardware enforced. This type of zoning is useful for controlling access to certain devices without totally isolating them from the fabric. Members can communicate with each other and transmit outside the ACL zone, but cannot receive inbound traffic from outside the zone.

The ACL zone boundary is secure against inbound traffic. ACL zones can overlap; that is, a port can be a member of more than one ACL zone. ACL zones that include members from multiple switches need not include the ports of the inter-switch links. ACL zone boundaries supersede soft zone boundaries, but yield to VPF zone boundaries. Membership can be defined only by port ID and domain ID. ACL zoning supports all port modes except TL_Ports.

### 5.4.3 Virtual Private Fabric zones

Virtual Private Fabric (VPF) zoning divides the fabric for purposes of controlling discovery and both inbound and outbound traffic. This type of zoning is useful for providing security and reserving paths between devices to guarantee bandwidth. VPF zoning is a type of hard zoning that is hardware enforced. Members can only transmit to and receive from members of the same VPF zone.

The VPF zone boundary is secure against both inbound and outbound traffic. VPF zones that include members from multiple switches must include the ports of the inter-switch links. VPF zones cannot overlap; that is, a port can be a member of only one VPF zone. VPF zone boundaries supersede both soft and ACL zone boundaries. Membership can be defined only by port ID and domain ID. VPF zoning supports all port modes.

**Note:** Domain ID conflicts can result in automatic reassignment of switch domain IDs. These reassignments are not reflected in zones that use domain ID and port number pairs or Fibre Channel addresses to define their membership. Be sure to reconfigure zones that are affected by a domain ID change. To prevent zoning definitions from becoming invalid when the membership is defined by domain ID/port number or Fibre Channel address, you must lock domain IDs.

### 5.4.4 Aliases

To make it easier to add a group of ports or devices to one or more zones, you can create an alias. An alias is a named set of ports or devices that are grouped together for convenience. Unlike zones, aliases impose no communication restrictions between its members. You can add an alias to one or more zones. However, you cannot add a zone to an alias, nor can an alias be a member of another alias.

### 5.4.5 Zoning database

Each switch has its own zoning database. The zoning database is made up of all aliases, zones, and zone sets that have been created on the switch or received from other switches. The switch maintains two copies of the zoning database: one copy is maintained in temporary memory for editing purposes; the second copy is maintained in permanent memory. Zoning database edits are made on an individual switch basis and are not propagated to other switches in the fabric when saved.

### 5.4.6 Zoning example

In the following example we show the procedure to create a zone and zone set for our fabric. We will create a zone set consisting of a single zone with two members.

Before the new zone set is configured we will verify the zone set configuration properties by selecting **Zoning —> Edit Zoning Config** from the Faceplate menu as shown in Figure 5-22.

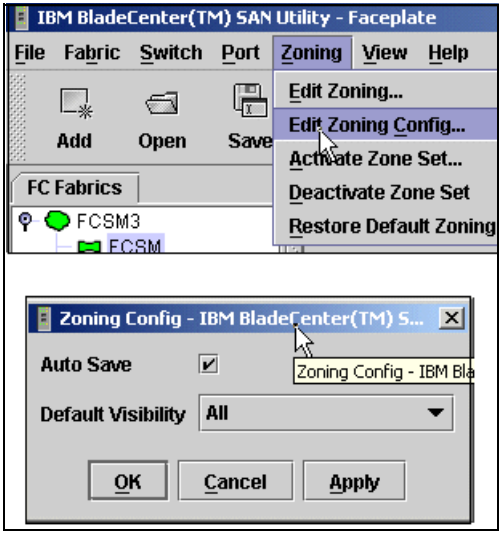


*Figure 5-22 Edit Zoning Config menu*

We recommend that the **Auto Save** option should always be **enabled** and **Default Visibility** should be set to **All**.

The **Auto Save** parameter determines whether changes to the active zone set that a switch receives from other switches in the fabric will be saved to permanent memory on that switch. Changes are saved when an updated zone set is activated. Zoning changes are always saved to temporary memory. However, if Auto Save is enabled, the switch firmware saves changes to the active zone set in both temporary and permanent memory. If Auto Save is disabled, changes to the active zone set are stored only in temporary memory.

The **Default Visibility** parameter determines the level of communication that is permitted between devices when there is no active zone set. The default visibility parameter can be set differently on each switch. When default visibility is enabled (All) on a switch, all ports on the switch can communicate with all ports on switches that also have Default Visibility enabled. When Default Visibility is disabled (None) on a switch, none of the ports on that switch can communicate with any other switch in the fabric.

To perform the actual zone set configuration, the BladeCenter switch icon (we have selected *FCSM*) should be highlighted first, and then you click the **Zoning** option (as indicated by an arrow) as shown in Figure 5-23.

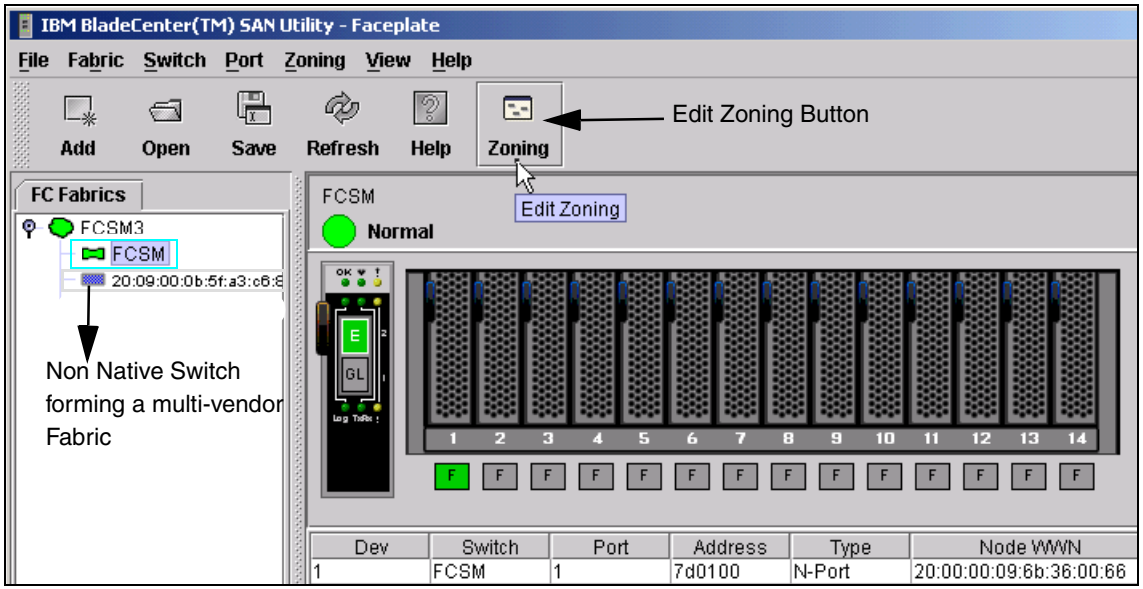


*Figure 5-23 Edit Zoning menu*

The **Edit Zoning** window is displayed. Click the **Zone Set** button to create a new zone set as shown in Figure 5-24.

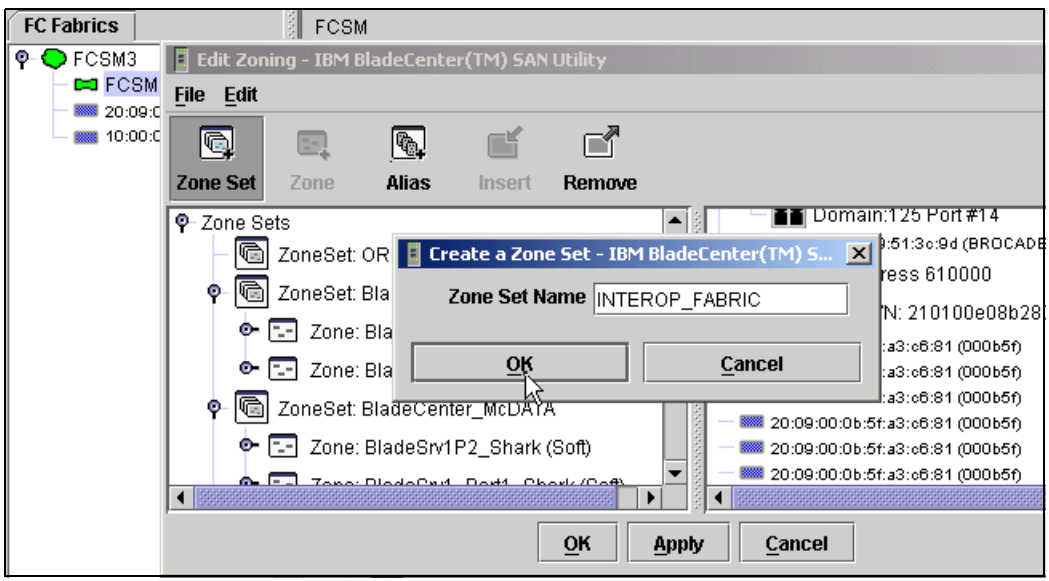


*Figure 5-24 The BladeCenter SAN Utility Create New Zone Set menu*

The name of the new zone set is “INTEROP_FABRIC” to identify that this is an multi-vendor fabric. Click **OK** as shown in Figure 5-24.

After the new zone set has been created successfully, proceed to create a Zone(s) as required by highlighting the “INTEROP_FABRIC” zone set and then clicking on the **Zone** button to create a zone as shown in Figure 5-25.

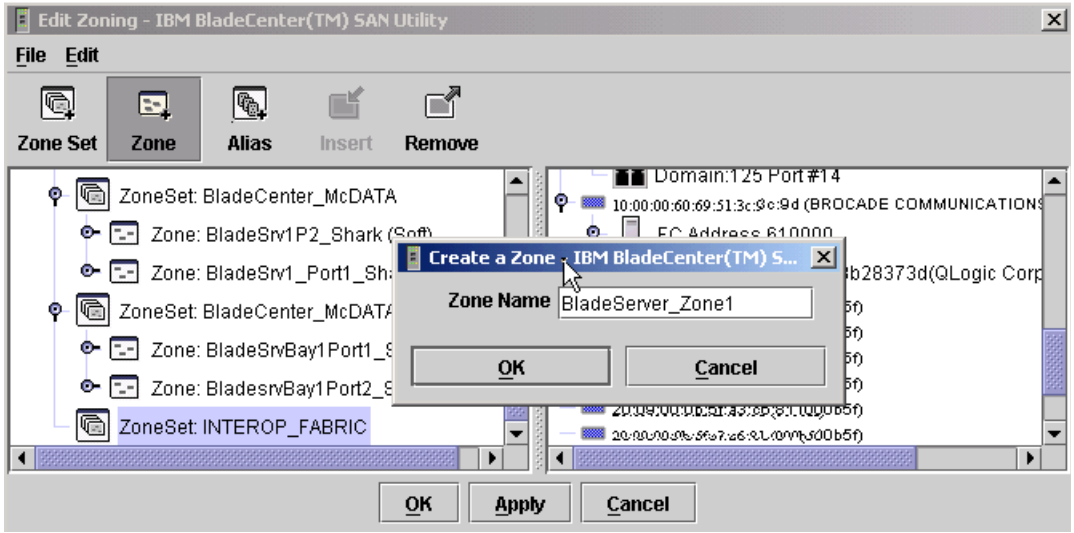


*Figure 5-25 The Create new Zone option from the Edit Zoning menu*

The new zone with the name “BladeServer_Zone1” is created.

After the new zone has been created successfully, highlight the zone “BladeServer_Zone1” found within (under) the “INTEROP_FABRIC” zone set. This will display a list of devices. Select a port-by-port number, Fibre Channel address, or WWN in the port/device tree, and drag it into the zone. To select and drag multiple ports/devices, press and hold the Control key while dragging.

This is shown in Figure 5-26.

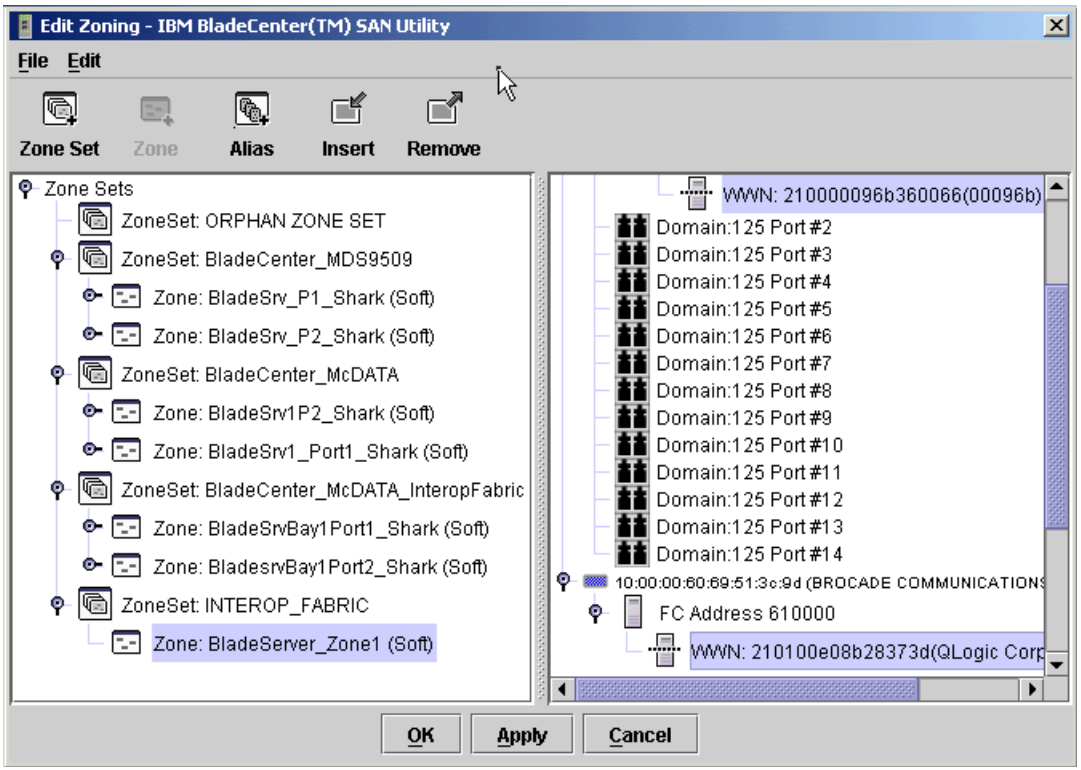


*Figure 5-26 WWN member add*

After the devices have been identified from the list, click the **Insert** button as shown in Figure 5-27.

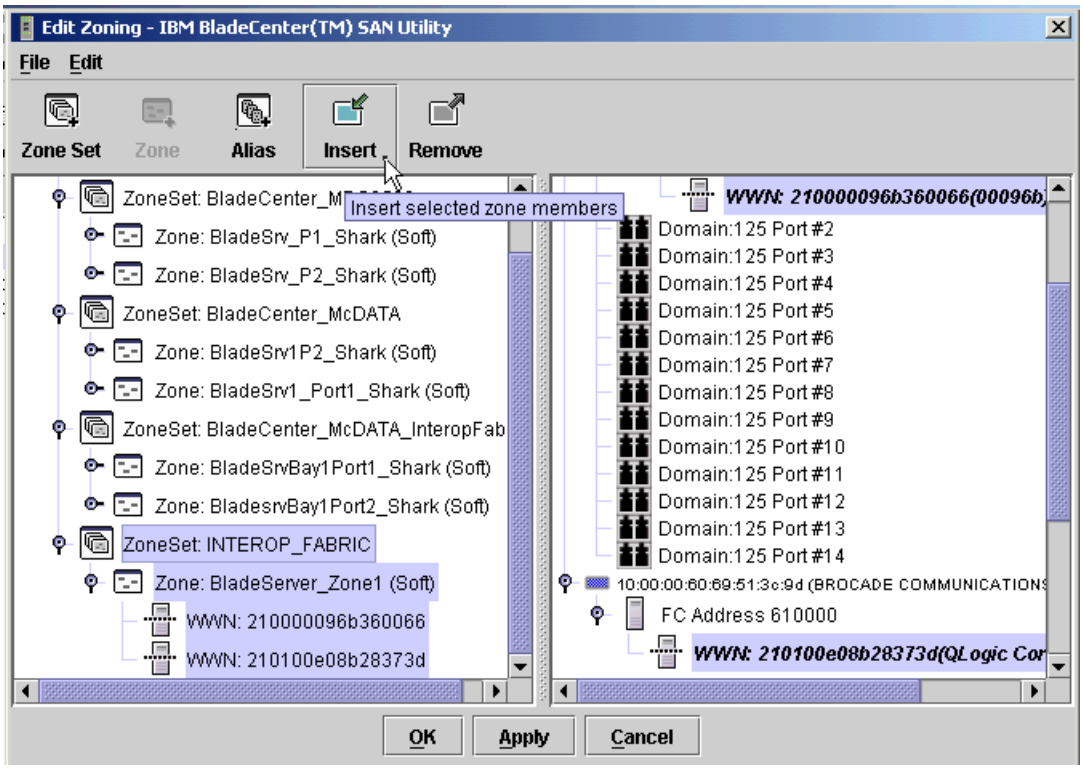


*Figure 5-27 Insert zone members*

The devices are now listed under the “BladeServer_Zone1” zone.

In Figure 5-27 the entire zone set configuration comprising the zone set, zones and zone members is displayed. Now verify and confirm the changes by clicking on **OK** to complete the process prior to activation. The window confirming the completion of the zoning configuration is displayed, reminding us to activate the new zone set for changes to be made effective. Doing so will deactivate the currently active zone set. This is shown in Figure 5-28.


*Figure 5-28 Zone configuration prompt*

We will now activate the “INTEROP_FABRIC” zone set by selecting the **Zoning—> Activate Zone Set** option as shown in Figure 5-29.

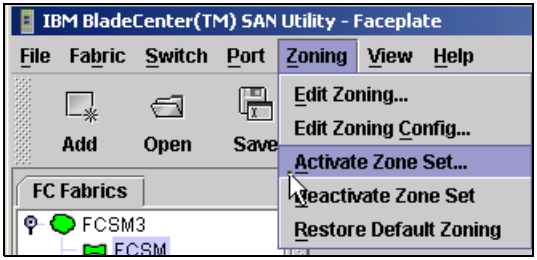


*Figure 5-29 Activate Zone Set menu*

Select the **“INTEROP_FABRIC”** zone set from the drop down list and click **OK** as shown in Figure 5-30.

**Attention:** The currently active zone set will be deactivated prior to activation of the new zone set. Before activating the new zone set, ensure that you consider the implications of deactivating the currently active zone set.

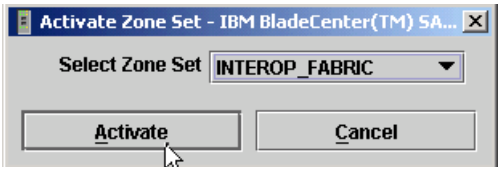


*Figure 5-30 Activate Zone Set selection*

If the zone set activation is complete and successful, a message similar to **Zone Set INTEROP_FABRIC Activated** will be issued. Click **OK**.

To view the active zone set status, select the fabric icon (**FCSM3**) to display the **Active Zoneset** tab and click it, as shown in Figure 5-31.

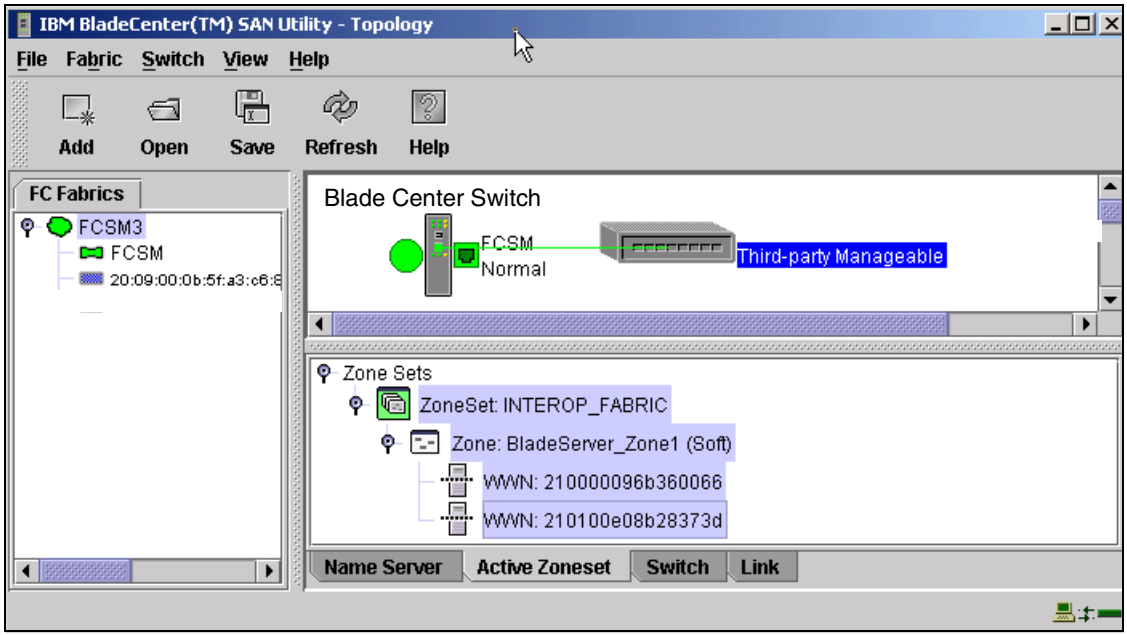


*Figure 5-31 Active Zoneset option*

Our zone set with its zone and members has been created.

## 5.5 Firmware upgrade

In this topic we show how to upgrade the firmware.

**Attention**: The firmware upgrade procedure is disruptive and the switch must be restarted in order to activate the new firmware. A maintenance window should be scheduled. We recommend that you back up the switch configuration before any upgrade takes place.

To download and activate the new firmware on a switch module, the following steps need to be performed:

1. Select one or more switch modules in the topology display.
2. Open the switch menu and select **Switch ->Load Firmware...** as shown in Figure 5-32.

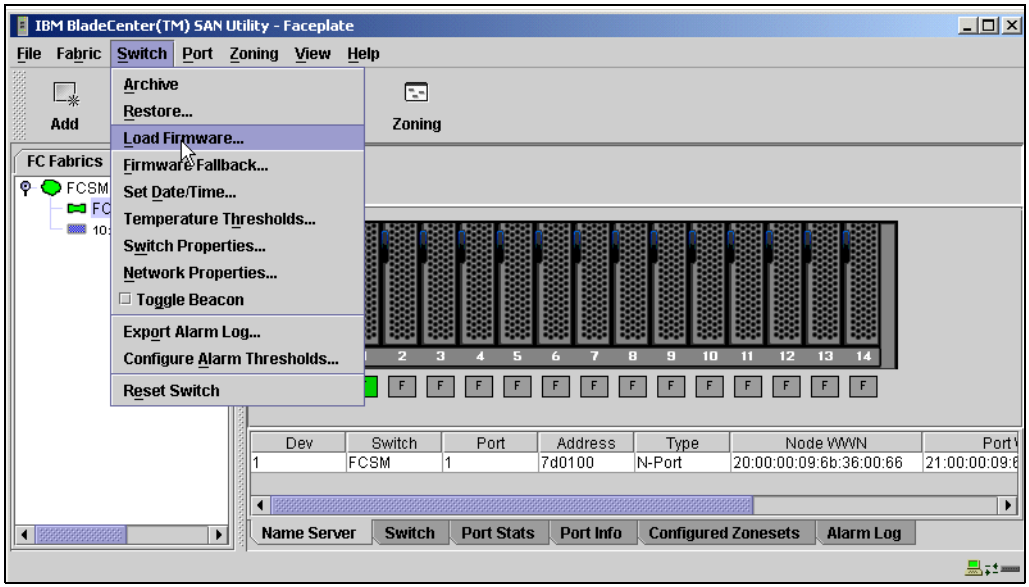


*Figure 5-32 Load Firmware*

3. In the firmware upload window, click the **Select** button to browse and select the firmware file to be uploaded.
4. Highlight a firmware file and click the **Open** button as shown in Figure 5-33.

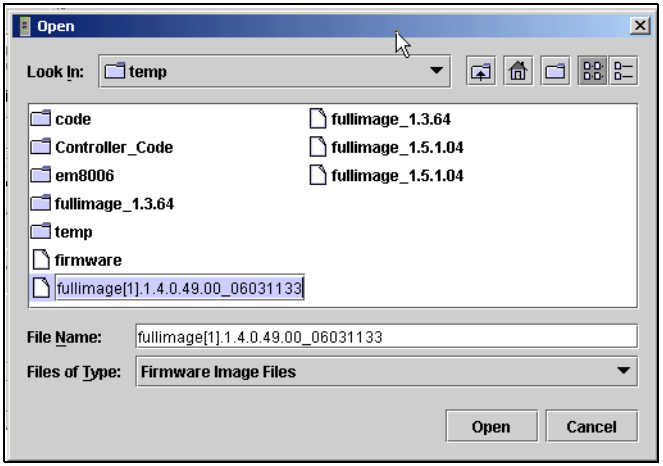


*Figure 5-33 Firmware image selection*

5. Click the **Start** button in the firmware upload window to begin the firmware upload process as shown in Figure 5-34.

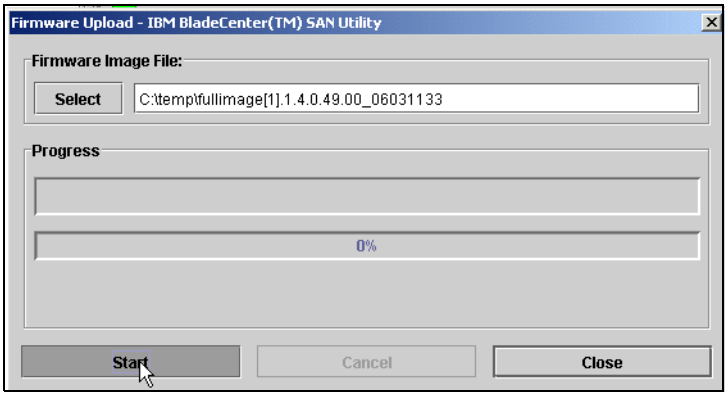


*Figure 5-34 Start firmware download menu*

As shown in Figure 5-35 the progress bar reaches 100% indicating that the firmware download process has completed.

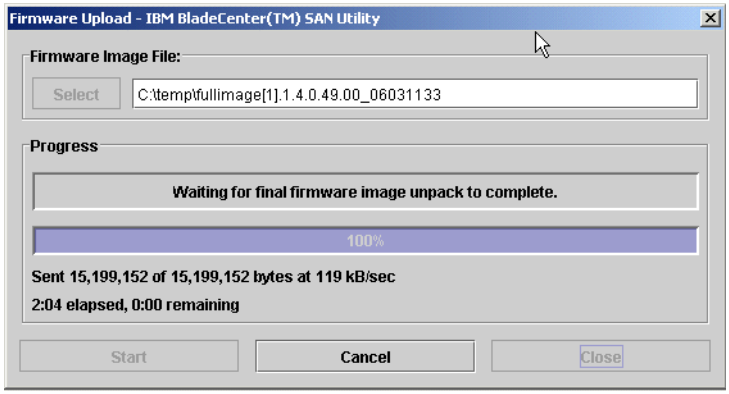


*Figure 5-35 Firmware download complete*

6. When the firmware upload is complete, click the **Close** button to close the firmware upload window.
7. The application will prompt you to reset the switch module in order to activate the new firmware. Click the **OK** button to reset the switch and activate the new firmware.

The active firmware version can be verified from the switch properties menu or from the Switch Info tab and also by using the **`show switch`** command from the CLI.

6

# Implementing the IBM TotalStorage SAN Controller 160

The IBM TotalStorage SAN Controller 160 enables all IBM 7133, 7131, and 3527 SSA Serial Disk Systems to attach to host systems using Fibre Channel host adapters and drivers. This allows you to protect your investment in SSA disk, while still being able to create and build a SAN infrastructure.

The IBM TotalStorage SAN Controller 160 replicates data across or within serial disk systems — simultaneously mirroring two or three copies of data without host involvement. With global hot disk sparing, data is automatically rebuilt if a mirrored disk fails. In this way, the IBM TotalStorage SAN Controller 160 improves performance and data availability while simplifying storage operations.

## 6.1 SAN Controller 160 features

The Instant Copy function can create a separately addressable copy of mirrored data that can be used for tape backup. After the backup has completed, data is resynchronized with the primary copy.

The IBM TotalStorage SAN Controller 160 also can create composite drives by concatenating up to 16 physical disks.

Using these functions, physical drives become members of larger or more complex logical drives.

A diagram to depict a single host to Controller configuration is shown in Figure 6-1.

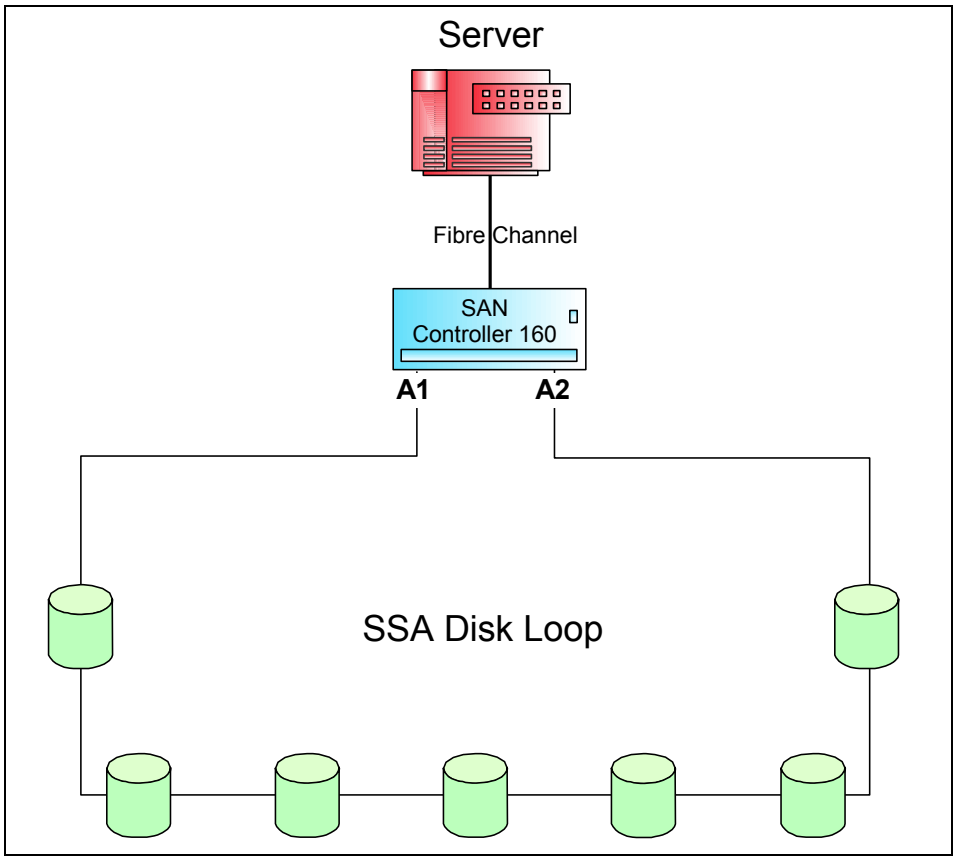


*Figure 6-1 SAN Controller 160 with a single host*

## 6.2 Installing the SAN Controller 160

To install and establish a Controller 160 storage system, the Controller and all the disks to be used must be set up in a proper sequence. During this sequence, only *a single* SAN Controller 160 must be used to do the configuration. Once configured, other SAN Controller 160s can be added to the loop. This sequence is described below.

**Note:** For information and a description of how to understand the LED codes that will be discussed, refer to the *SAN Controller 160 Installation and User's Guide*, 310-605759.

1. Power on SSA disks.

   The SSA disks must be powered on to ensure that all disks spin up and pass the power sequence. Make sure all drive LEDs are on solid to indicate a functioning device. Any faulty or suspect disk drive should be replaced.

2. Clear Controller Node Mapping.

   To begin the Controller installation, the first step is to clear the Node Mapping table. This is done by *shorting* the SSA ports on the rear of the Controller. Plug in an SSA cable from port A1 to A2 on the back of the Controller and power it on. Clearing the table will take only seconds, and when completed, the Status LED on the front of the Controller will flash a code **060**. The Controller is now powered off, the SSA shorting cable is removed and the SSA disks are attached.

3. Connect SSA disks to Controller.

   All the disks to be used are to be connected together to form a complete SSA loop with the Controller included. All dip switches in SW1 should be in the down position. On SW2, dip switches 0 and 1 are set in the down position, all other switches in dip switch 2 should be turned up. This is considered *mode 3*. Power on the Controller, the Status LED will begin to flash rapidly as it searches the SSA loop to recognize all the disk drives. This may take approximately 1 minute to complete. Once the Status LED has stopped flashing and is solid, the process is complete and the Controller is powered down.

4. Perform Controller Subsystem Diagnostic test.

   A Subsystem Diagnostic test is now run on the disk drives called *mode 15*. This tests the disk drives for spindle spin up, read tests and nondestructive write tests. The Controller is set for mode 15 by setting switches 0, 1, 2, and 3, on SW2, to the down position, and the rest turned up. The Controller is now powered on, the Status LED will flash rapidly. The test will be done on each disk drive in the SSA loop separately and will begin with the drive closest to

the A1 port on the back of the Controller. As the test is completed on a drive, the LED on the SSA drive will flash and then it will move to the next drive. This test should continue until all drives have been tested. The test runs in a continuous cycle, so once all drives have been tested at least once, the Controller is powered off.

If a drive fails the test, the testing will stop, and the Controller's Status LED will flash a diagnostic code. A code map with a description of the errors can be found in the *SAN Controller 160 Installation and User's Guide*, 310-605759.

5. Assign Fibre Channel target.

   With the Controller powered off, you can now assign a Fibre Channel target ID number to the Controller. Any number can be selected, however, this number must be a unique ID. No other device can have the same Fibre Channel target ID once it is set on the Controller.

   This is done by setting selected dip switches in SW1 to the down position. The switch is set up in binary notation: a switch that is down represents a 1 and a switch up represents a 0. Figure 6-2 shows the switch numbers and their corresponding value.

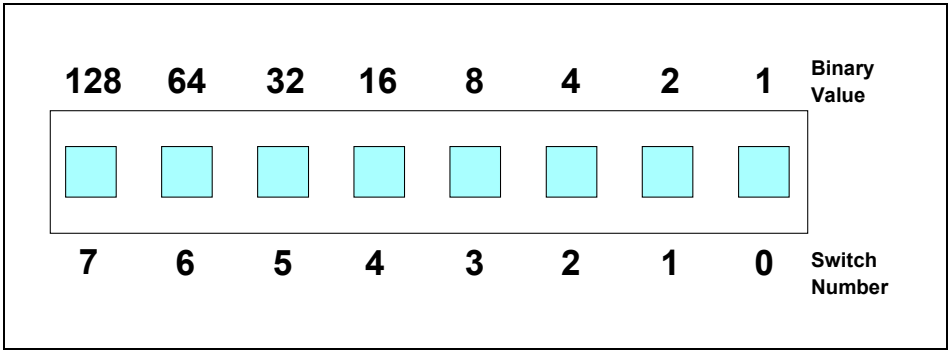


*Figure 6-2 SW1 dip switches*

   For example, by setting switch 2 and 0 down, a value of 5 is represented. All other switches would be in the up position. By setting a switch down, its value is added. To change a number that was previously set on a Controller, power must be cycled to the Controller for the change to take effect.

6. Map the physical drives.

   Before powering on again, SW2 must be set to mode 3 with switch 0 and 1 set down, and all other switches set up. The Controller is powered on, the Status LED will flash rapidly to rediscover the disks and the SSA drive LEDs should be on solid. Once completed, the status LED will be solid, and the drives are now considered to be mapped physical drives. The Controller is powered off.

7. Create general spares.

   The mode on SW2 is changed to mode 12 to set all drives to be general spares. Mode 12 is represented by setting switch 2 and 3 down and the rest turned up. Power on the Controller again, the Status LED will flash. After approximately one minute, the LED will flash code **100** to indicate it has completed. The Controller is again powered off.

8. Format the drives.

   The Controller is set to mode 14, switch 1, 2, and 3 down on SW2, to format all disk drives. Power on the Controller, depending on the number of drives, the format process will take anywhere from 30 to 60 minutes. During this time, the Status LED will flash rapidly and the SSA drive LEDs will flash. When it has completed, the Status LED will flash code **100**. Power off the Controller.

9. Clear the node map.

   Once completed, the node map *must* be cleared. This is done as described earlier by *shorting* ports A1 and A2 with an SSA cable. Power on, wait for code **060** to flash and then power off.

   The drives are now ready to be assigned and used on a host.

   You can also now set up mirror drives or composite drives within the Controller. This is done by setting the switches in SW2 to other modes. For detailed information on setting the switches and selecting the other modes, please refer to the *SAN Controller 160 Installation and Service Guide*, GC26-7433.

10. Perform host attach and power up sequence.

    For a host to now recognize and use the disks, set the dip switches in SW2 back to mode 3, this is normal host operation mode. The Fibre Channel cable from the host can be connected to the Controller. If the SSA drives are not powered on, do this now, and this should be done before the Controller. Next, the Controller is powered on, wait for the Status LED to stop flashing and remain on solid. At this point the host can be powered on.

    A check can be done to see that the SAN Controller 160 is being recognized by the host adapter card. On a Windows 2000 with a QLogic Fibre Channel adapter, during boot up look for a prompt to enter the QLogic BIOS by entering in **ALT Q**. At the BIOS window, select **Scan Fibre Devices**. A list of the Fibre Channel target IDs are presented, scroll down to the ID that you set in SW1. You will see the WWN of the SAN Controller 160. Exit the BIOS and the system will reboot.

Once the system has started, ensure that the host has access to all the drives. This is different depending on the operating system of the computer. For example on Windows 2000, select **Start -> Programs -> Administrative Tools ->Computer Manager**, then select **Disk Management**. This tool will report that new disks have been found and will be added to the system.

## 6.3 Controller 160 Manager software

Rather than using the dip switches to configure the features, another option is to use the Controller 160 Manager software. The Controller 160 Manager also provides configuration, monitoring and management capabilities of the SAN Controller 160 and the SSA drive loop. The Controller 160 Manager can be setup to allow remote access if desired.

The Manager software consists of server and client portions. The server includes a daemon service and a user interface. The client has the user interface only.

The server portion must be loaded on the host that is directly attached to the Controller, as the daemon service is started from here. The daemon must reside on the host that is directly connected to the Controller. This host can also be used to run the Manager software for local access.

The client software can be loaded on to any computer, running a supported operating system, that can communicate to the host with the daemon service running. It must communicate to the server host using TCP/IP. This allows remote access to the Controller and the storage loop. See Figure 6-3.

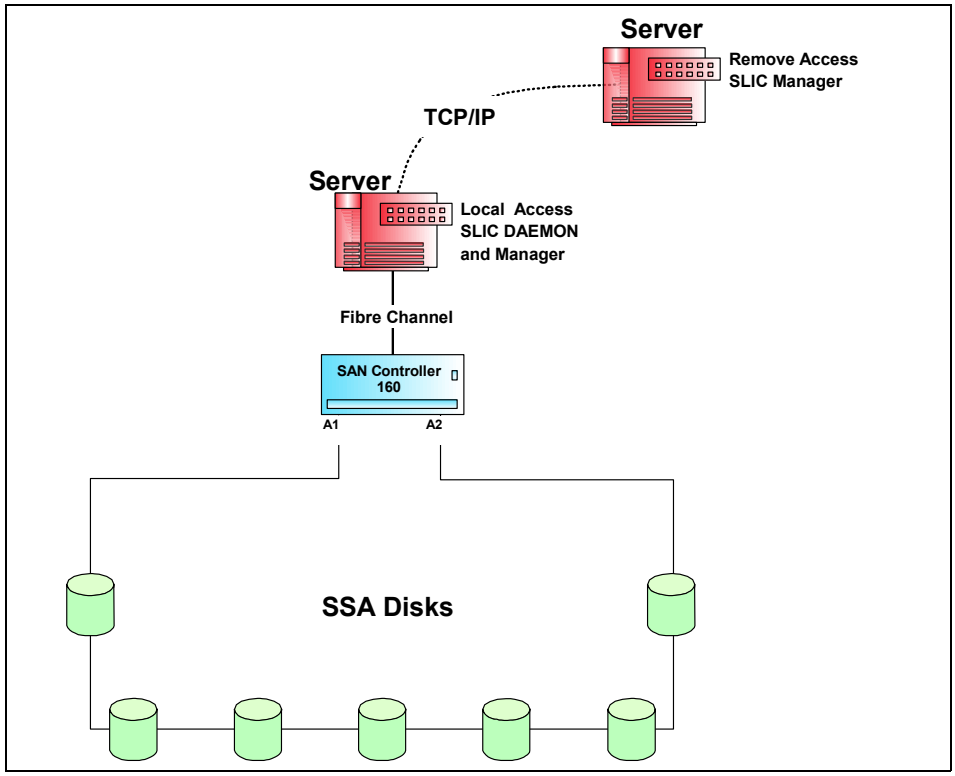


*Figure 6-3 Controller 160 Manager access*

The Controller 160 Manager has a graphical user interface (GUI) and a command line interface (CLI) available for Windows 2000 systems. UNIX systems will only have the command line interface available.

The following installation and configuration examples will detail using the GUI from a Windows 2000 platform. To review the commands that are available for UNIX platforms and installation instructions on other operating systems, refer to the *SAN Controller 160 Manager Installation and User Guide*, GC26-7432.

## 6.3.1 Installing the Controller 160 Manager software

The Controller 160 Manager software can run on many operating systems. The following discussion describes an installation on a Netfinity x330 with Windows 2000.

To install the Controller 160 Manager server software for local access, the SAN Controller 160 Utilities CD-ROM is placed in the CD drive. Select **Start -> Run** and **Browse** the CD drive. Go to `i386\server\setup.exe` and click **OK**. Follow

the prompts displayed on the window to install the Manager software. This will install the daemon service also.

For remote or client access the daemon service is not required. To load the Manager software only, go to `i386\client\setup.exe instead`.

## 6.3.2 Communicating to the Controller

For the Controller 160 Manager server software to communicate to the Controller, it requires space on a disk or several disks that are within the SSA loop. This is referred to as a Controller 160 Zone. To create space on a disk, a file or partition — depending on the operating system used — is created for the Manager software to use. To create this Controller 160 Zone, a configuration file must be created or edited.

### Editing the configuration file

When the software is loaded, a sample configuration file called **`7140.cfg`** is added in the `C:\ibm7140\sdus directory`. This is a text file that can be viewed and edited by simple text editors, such as Windows Wordpad. Open up the *7140.cfg* file and it will contain a sample of how the file should look. Also note that on the left hand side, the `#` sign is entered in every line to mark it out as a comment. This is shown in Figure 6-4.

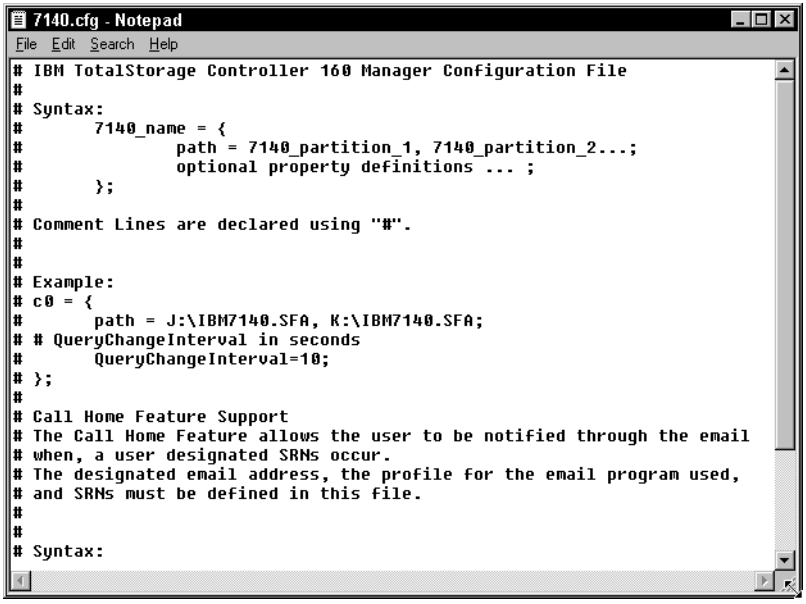


*Figure 6-4 Sample configuration file*

This file can now be edited to be used as the configuration file for your Controller 160 system. Begin by deleting the # sign on the lines that contain the sample configuration. The rest of the information can be entered as shown in Figure 6-5. A description of each entry field is also provided.

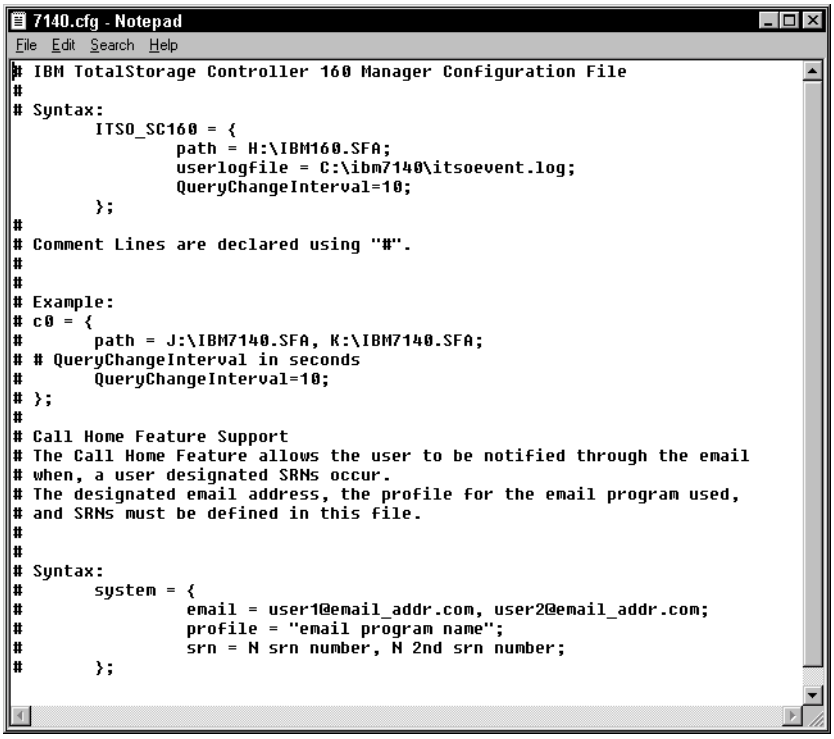


*Figure 6-5 Edited configuration file*

## Configuration file information

The **7140_name** can be any name that you would like to use to identify the Controller.

### *Creating a Controller 160 Zone*

The **path** refers to the Controller 160 Zone, file or partition, used for the Manager to communicate to the Controller. To edit this option, it requires that a drive on the SSA loop has been recognized by the host and that the drive has been formatted. In the example above, a Windows 2000 host was used. The Windows 2000 Disk Administrator was used for the host to recognize the drives, the first drive assigned the next drive letter, H, and it was formatted.

The file naming for a Controller 160 Zone depends on the type of operating system running.

For Windows 2000, the naming is `<drive letter>:\IBM160.SFA`.

You can enter in many Controller 160 Zones, but only one is required to get access at the beginning. After the other drives have been configured as mirrors or composite drives, then Controller 160 Zones can be created for these drives if desired.

Including many Controller 160 zones in the path statement will allow the Manager to access a zone on another drive. This is helpful to protect against when a drive fails, and that drive has a Controller 160 zone defined to it. If the Manager cannot access the first Controller 160 zone, it would try the next zone in the order it was entered in the path statement.

For the naming conventions used on other operating systems to create a Controller 160 Zone, refer to the *SAN Controller 160 Manager Installation and User Guide*, GC26-7432.

The *userlogfile* will define a file with which you can view logged events.

The *QueryChangeInterval* sets the time in seconds that the daemon will poll the Controller. The recommended time set here is 10.

Ensure that at the end of every line a semi-colon ';' is used, and that, if several Controller 160 Zones are created, a comma separates them. Save and exit the file.

## Installing the Controller 160 Manager daemon

With the configuration file edited and a Controller 160 Zone created, the daemon service can be installed and run. To install the service in Windows 2000, open a DOS prompt and go to C:\ibm7140\sdus. Type in **`slicd -install`**, and the daemon will be installed.

## Starting the Controller 160 Manager daemon

To start the daemon service, select **Start —> Programs —> Administrative Tools —> Services** from Windows 2000. Scroll down until you see **IBM TotalStorage Controller 160 Manager Server**; select and highlight it. You will see two columns to the right to indicate its status.

To start the service, click the **Start** button and it will take a few moments to complete. Once it is done, you will see the word **`Started`** in the Status column. If the Startup column contains the word **`Automatic`**, no further action is required. If not, click the **Startup** button, and change the Startup Type to **`Automatic`**. This will have the daemon service start automatically during a reboot. This is shown in Figure 6-6.

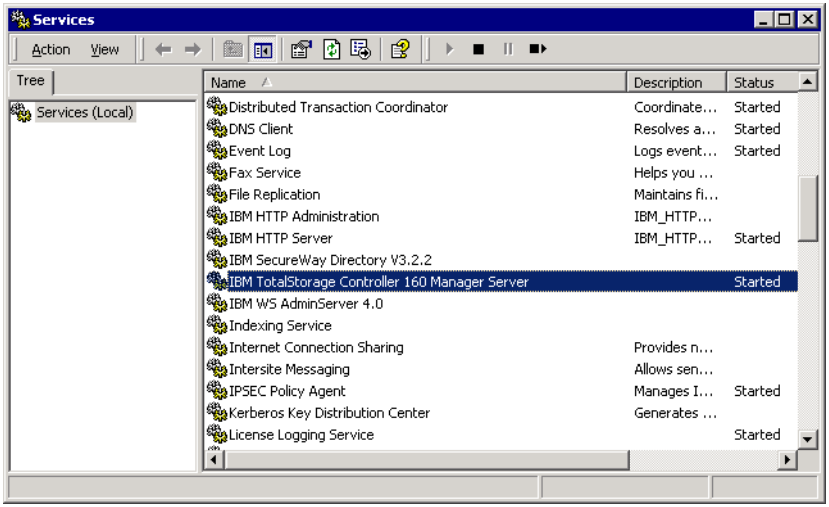


*Figure 6-6 Controller 160 daemon start up in Windows 2000*

## 6.3.3 Starting the Controller 160 Manager

To start the SAN Controller 160 Manager software, select **Start —> Programs —> IBM TotalStorage Controller 160 Manager -> IBM TotalStorage Controller 160 Manager.** The software will load, and a dialog box will appear. In the box with the heading `Hostname`, enter in the name or IP address of the host the daemon service is running. Enter in the Controller 160 name you entered in when editing the **`7140.cfg`** file. An example is shown in Figure 6-7.

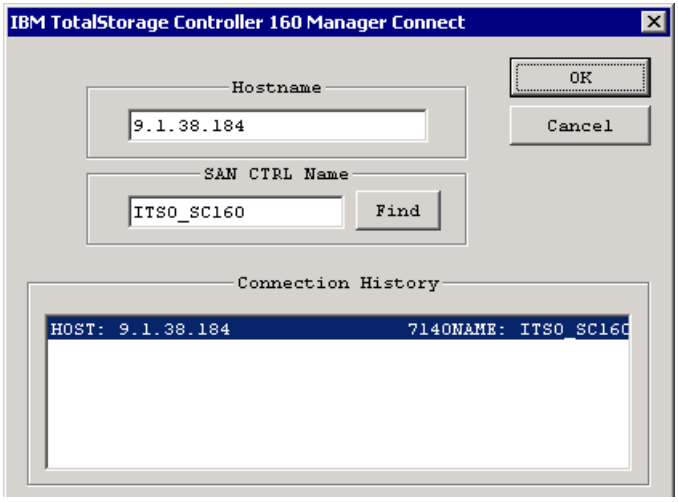


*Figure 6-7 Controller 160 connection window*

Click **OK** and the software will begin to communicate with the Controller. You will notice that the top title bar of your window will now include the host name and Controller 160 name as in Figure 6-8.


*Figure 6-8 Controller 160 Manager title bar*

## 6.4 Using Controller 160 Manager

You can now look to see that all communications are working properly by going to the toolbar and selecting **Tools -> Control Center**. A dialog box will appear, as shown in Figure 6-9.

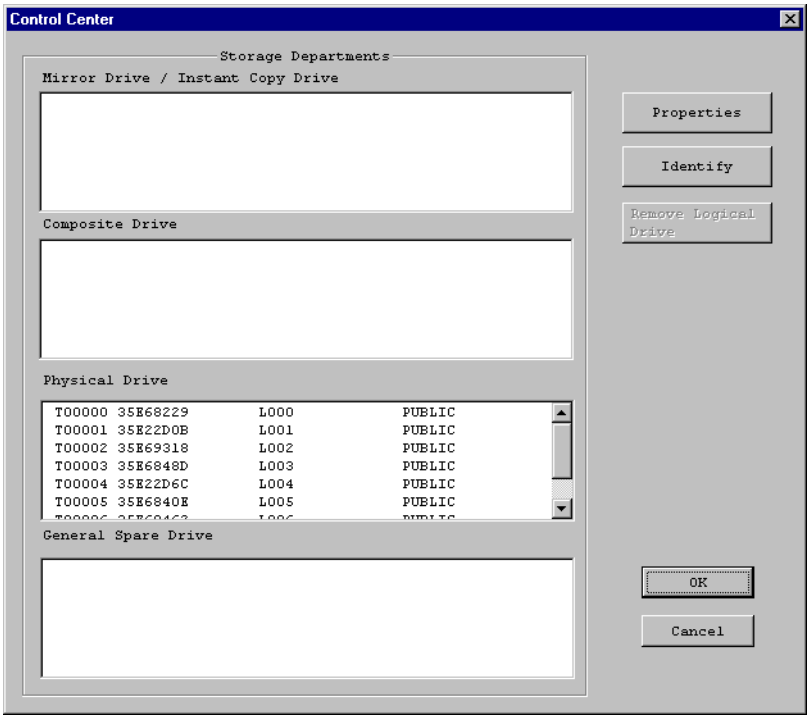


*Figure 6-9 Control Center window*

In the Physical Drive box, the drives that are on the SSA loop can be seen. This window will be useful as you start to create mirrors and composite drives, since it provides a summary of all drives.

## 6.4.1 Drive properties

You can get detailed information on each drive. Select the drive so that it is highlighted and then select **Properties**. A dialog box will appear with the drive's information as shown in Figure 6-10.

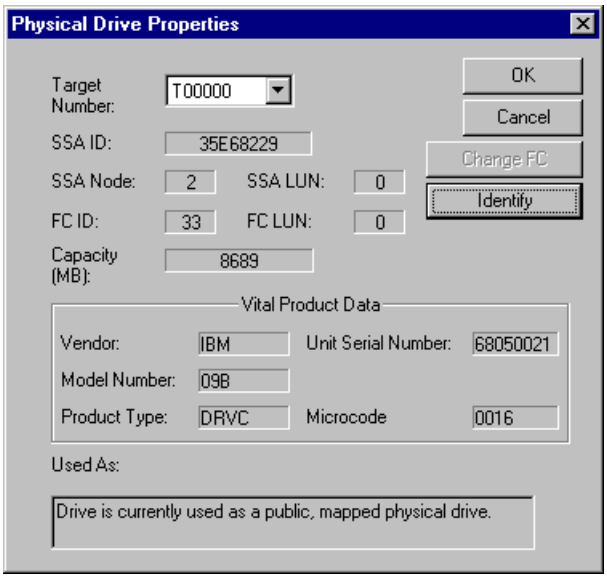


*Figure 6-10 Disk drive properties*

Here you can see its SSA attributes, its Fibre Channel attributes and its model type and serial number. By clicking the **Identify** button, the LED on the selected drive will begin to flash.

## 6.4.2 Controller properties

To view the information on the Controller, go to the toolbar and select **Properties —> TotalStorage Controller 160 Properties**. As shown in Figure 6-11, you will see the serial number of the Controller, its ID that was set in SW1, and its supported features.

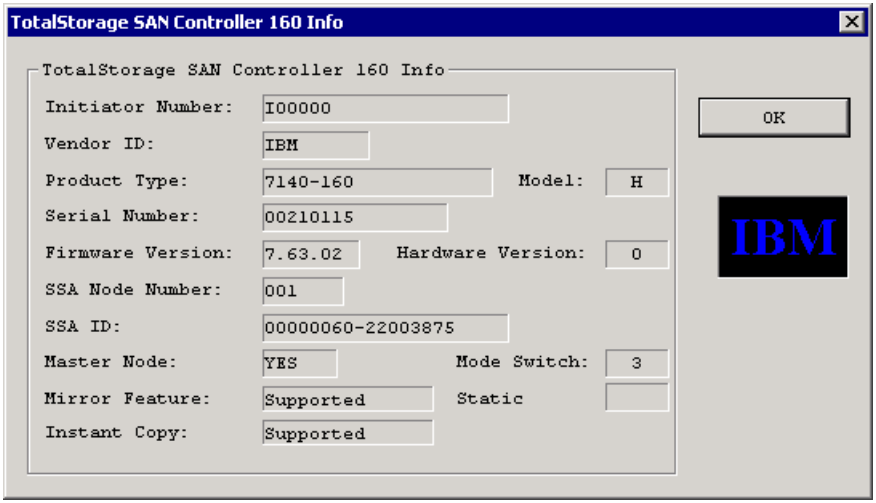


*Figure 6-11 SAN Controller 160 properties*

## 6.4.3 Setting Controller to master

As you move through the toolbar, you may notice that most selections have been grayed out. This is due to the fact that the Controller is currently in a subordinate role and does not have access to create mirrors or composite drives. This function is done by a Master Controller. There can be only one master in a Controller 160 loop. This is used as more Controllers and more disks can be added to the loop. With several Controllers in the same loop, there needs to be a requirement where one system acts as the control, and the others will follow and understand any configuration changes that may occur.

To set the Controller into a master role, select **File -> Program Option** from the top toolbar. You will be presented with a dialog box as shown in Figure 6-12.

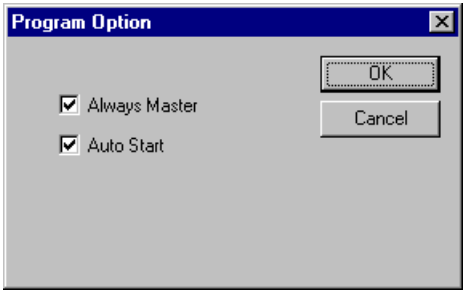


*Figure 6-12 Setting the Controller to master*

Click in the **Always Master** box so that a check mark appears. Once you click **OK**, the Controller will then be set as the master. You will notice that now all options in the toolbar are available and can start to use the features of the SAN Controller 160.

By placing a check mark in the Auto Start box, the Controller 160 Manager will automatically connect to the Controller defined in the Connection window as shown in Figure 6-7 on page 631.

### 6.4.4 The SignOn drive

When the Controller 160 zone was created to be used as the communication path, a disk file or partition was created on a specific disk within the SSA loop. As you begin to access the features of the SAN Controller 160, it should be known which disk was used to create the Controller 160 zone. This disk is considered to be the SignOn drive.

In the topics 6.5, “Composite drive” on page 636 and 6.6, “Mirror drive” on page 641 we describe creating composite and mirror drives, and you will see that the properties of the individual physical drives may change. As they become part of a logical drive, they take on the properties of this logical drive.

If the SignOn drive is used to create a logical drive, its attributes may change and you may lose the communication path that was created in the Controller 160 zone. When you select the **SignOn** drive as a member of a logical drive, a dialog box will be displayed as shown in Figure 6-13 to remind you that the attributes of this drive may be affected.


*Figure 6-13 Selecting SignOn drive dialog box*

As long as the LUN number of the SignOn drive becomes the LUN of the new logical drive, the communications from the Manager to the Controller will not be affected.

Another way to be certain that you do not lose your SignOn drive is not to use the SignOn drive to create logical drives. However, once some logical drives have been created, a Controller 160 zone can be created to one or more of the newly created logical drives. This logical drive can now be used as the SignOn drive and the previous drive is now available to be configured without any problems.

## 6.5 Composite drive

A composite drive is a large drive that consists of two or more smaller drives. The capacity of the composite drive is an aggregate of the capacities of all the smaller drives that are used to comprise this one large drive.

### 6.5.1 Creating a composite drive

To create a composite drive from the Controller 160 Manager, select **Tools -> Composite Drive Setup Wizard**. A dialog box, Composite Drive List, will appear. Currently, the list will be blank, because there are no composite drives created. Once there are composite drives created, you will see a list of the drives. Click the **Next** button and you will see the Members Selection window as shown in Figure 6-14.

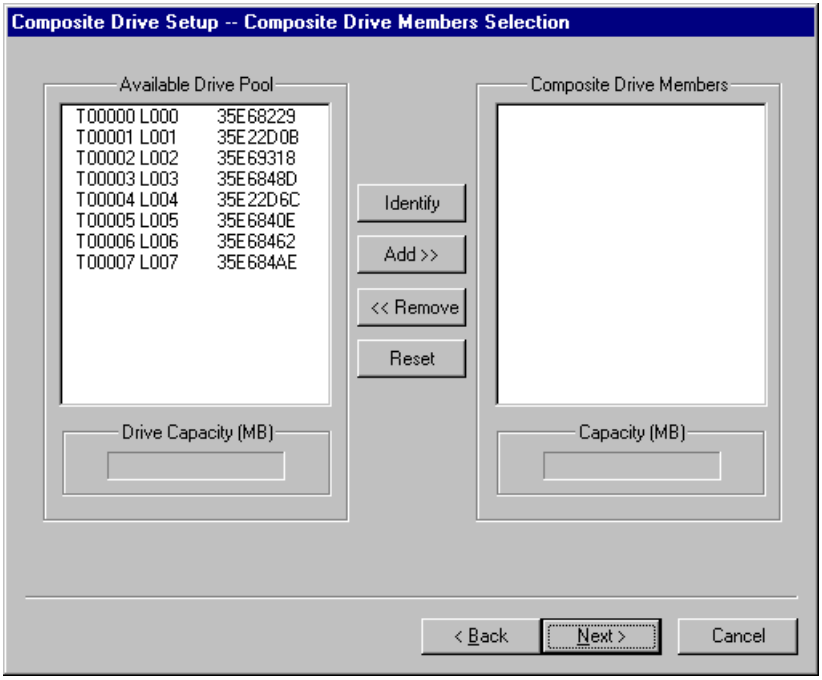


*Figure 6-14 Composite Drive Member Selection window*

From the Available Drive Pool list, click a desired drive and then click the **Add>>** button. The drive name will be added to the Member window. An asterisk will appear on the left hand side of the drive that was selected in the Available Drive window, to denote that the drive has been selected. Each drive is added one at a time. To remove a drive from the Member window, select the desired drive and click the **Remove<<** button.

Below each window there is a *Drive Capacity* box. As a drive is selected, its capacity in megabytes is shown. As you add more member drives to the Member window, the *Drive Capacity* box will add all drive sizes together to provide a total capacity in megabytes. This is shown in Figure 6-15.

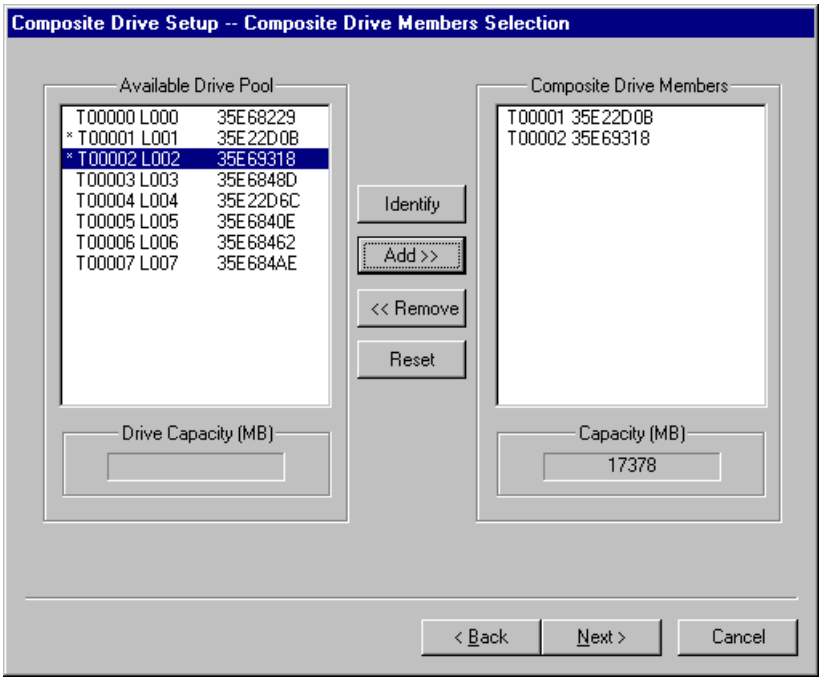


*Figure 6-15 Creating composite drive from available drives*

When all the desired drives are added, click **Next>**. The Assigning Properties window opens, as shown in Figure 6-16.

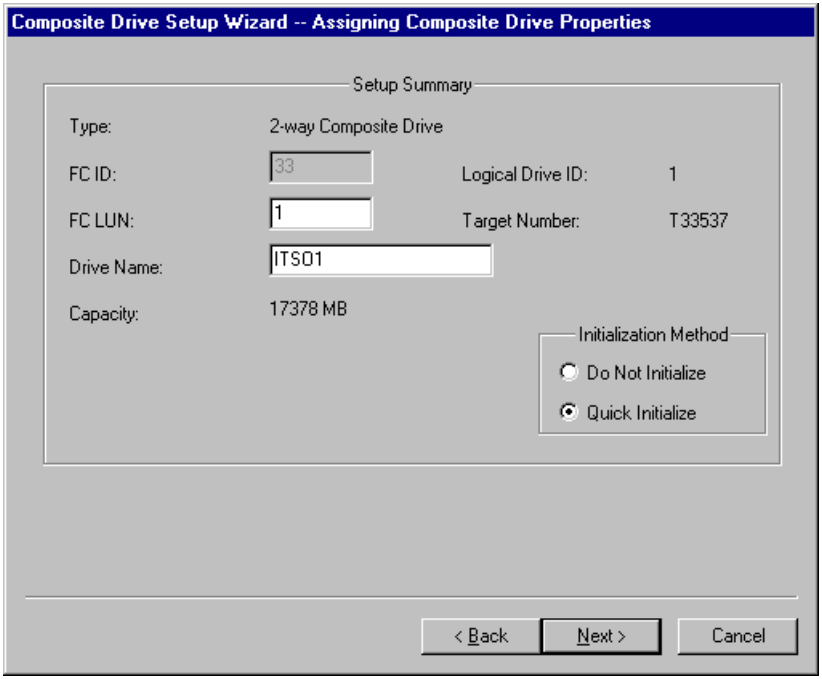


*Figure 6-16 Assigning Composite Drive Properties window*

The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. The name can also be defined to the composite drive for easier identification, with a limit of up to eight characters.

The Initialization Method box refers to whether or not to allow the operating system to write its signature on the composite drive.

Click the **Next>** button and a dialog box will appear, as shown in Figure 6-17, to allow you to create another composite drive. Click **Yes** If you would like to create another composite drive. The Composite Drive List window opens, and the steps described above can be repeated.

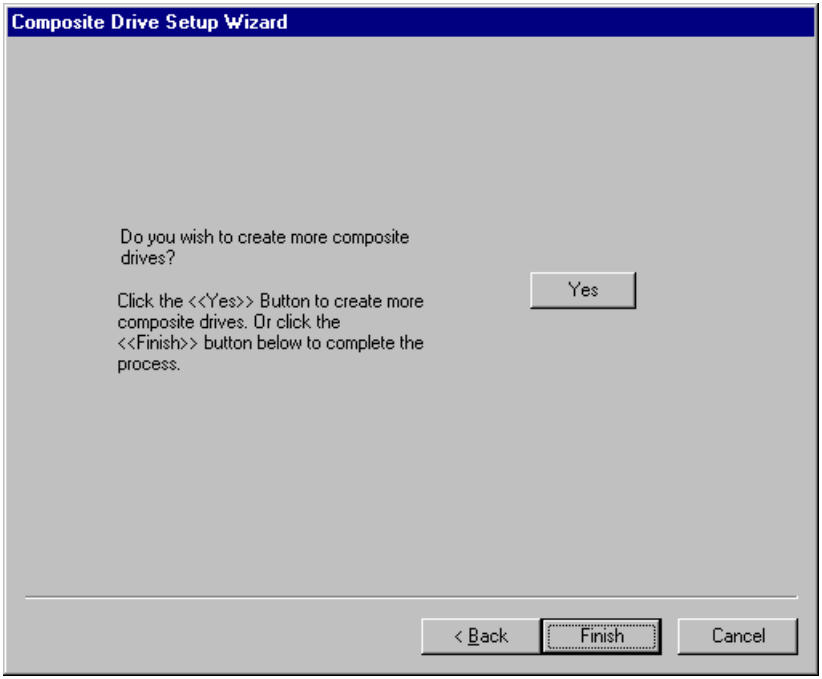


*Figure 6-17 Completing the Composite Drive setup*

Click **Finish** when you have created all the desired composite drives. Up to this point, the configuration has been kept within the Controller 160 Manager software. When the **Finish** button is clicked, the Controller 160 Manager will now communicate to the Controller to complete the process and update the Controller to control the drives.

The Host system must re-scan for devices, or restart, to be able to see the composite drive.

## 6.5.2 Composite drive properties

If you view the Control Center again, by selecting **Tools -> Control Center,** as shown in Figure 6-18, the newly created composite drive is listed in the Composite Drive box.

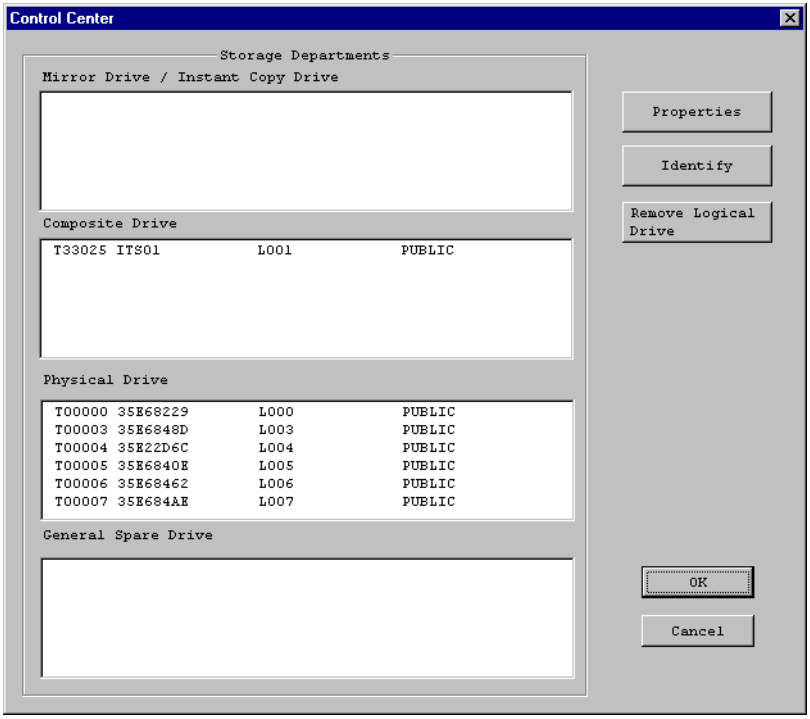


*Figure 6-18 Control Center with composite drive*

Select the composite drive and then click the **Properties** button. The Composite Drive Properties dialog box opens as shown in Figure 6-19.

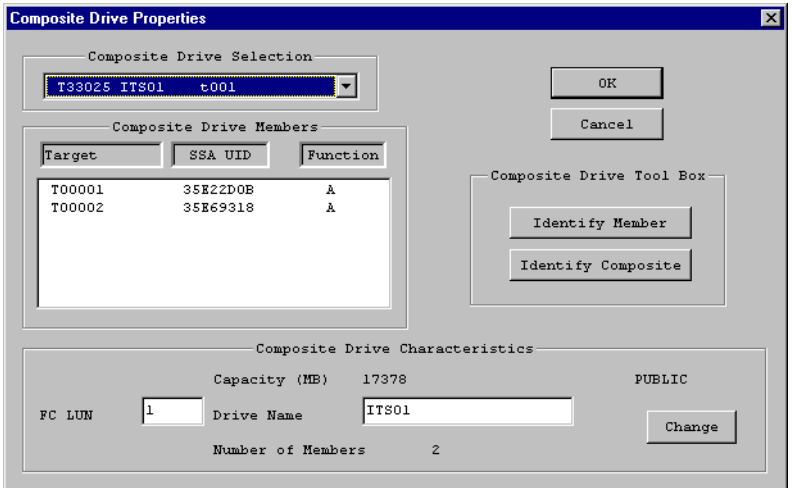


*Figure 6-19 Composite Drive Properties*

Here you can find information on the Composite Drive about the member drives that make up the composite, as well as the Composite Drive characteristics. Clicking on the **Identify Composite** button will cause the LED on the actual SSA drives, that belong to the Composite, to flash. If a member drive is selected and the **Identify Member** button is chosen, then the LED only on that drive will flash. In both cases a dialog box will appear to allow you to stop the flashing.

## 6.6 Mirror drive

A 2-way mirror drive has two drives that contain exactly the same information. The SAN Controller 160 can also support a 3-way mirror or 1-way mirror. A 3-way mirror consists of three drives with the same information. A 1-way mirror is a single drive, or single composite drive, that is used with an Instant Copy Drive that can attach to the single drive mirror to synchronize the data. The Instant Copy Drive can then be split off from the mirror to perform a backup or other action.

The Instant Copy Drive feature can be used with 2-way and 3-way mirrors as well.

### 6.6.1 Creating a mirror drive

To create a mirror using physical drives, from the toolbar, select **Tools -> Mirror Drive Setup Wizard**. You will see a dialog box, Mirror Drive List, that will be blank. If there were mirror drives created, then it would display the names of the drives. Click the **Next>>** button and the **Mirror Drive Members Selection** window opens. The window on the left named *Available Drive Pool* contains a list off all drives that are candidates to participate in a mirror drive.

Select a drive by highlighting it and click the **Add>>** button. The drive name will be added to the Member window. An asterisk will appear on the left hand side of the drive just selected in the Available Drive window, to denote that the drive has been selected. A second or third drive can be added to create a 2-way, or 3-way mirror, respectively. Each drive is added one at a time. To remove a drive from the Member window, select the desired drive and click the **Remove<<** button.

An example of adding two drives to create a 2-way mirror is shown in Figure 6-20.

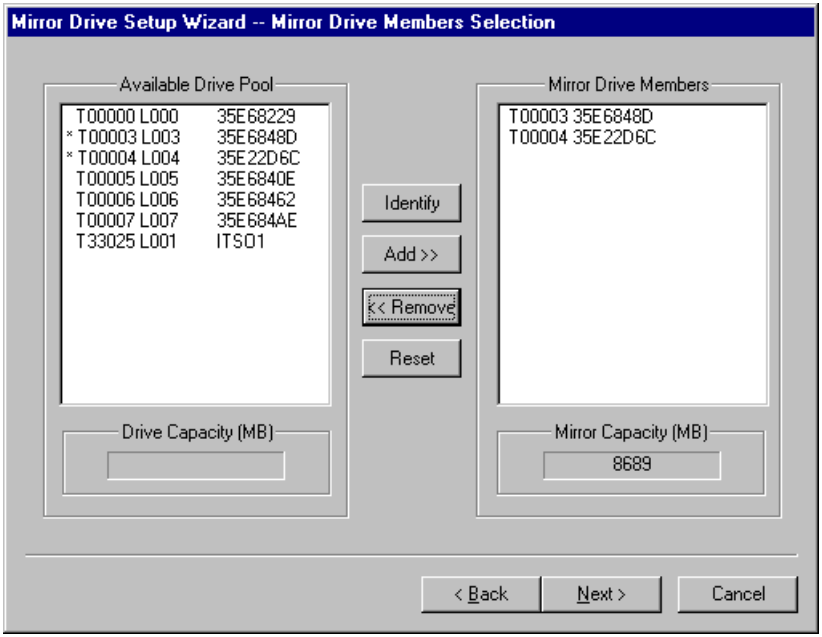


*Figure 6-20 Mirror drive member selection*

Below each window there is a Capacity window that will display the size of the available drive, or of the mirror drive. Each drive that participates in a mirror should be of equal capacity. If you select a drive that has a greater capacity and add it to the mirror, the mirror capacity will still be the smaller of the two, and the rest of the capacity of the larger drive will be unused. For example, if you added a 18 GB drive to the mirror in Figure 6-20, the Mirror Capacity window would still show the capacity of 8,696 MB. Approximately half of the 18 GB drive will be unused.

After all drives have been added, click **Next>** and you will be able to add a dedicated spare drive to the mirror if desired. Highlight one of the remaining available drives, click **Add>>** and its name will appear in the Mirror Drive Dedicated Spare window as shown in Figure 6-21.

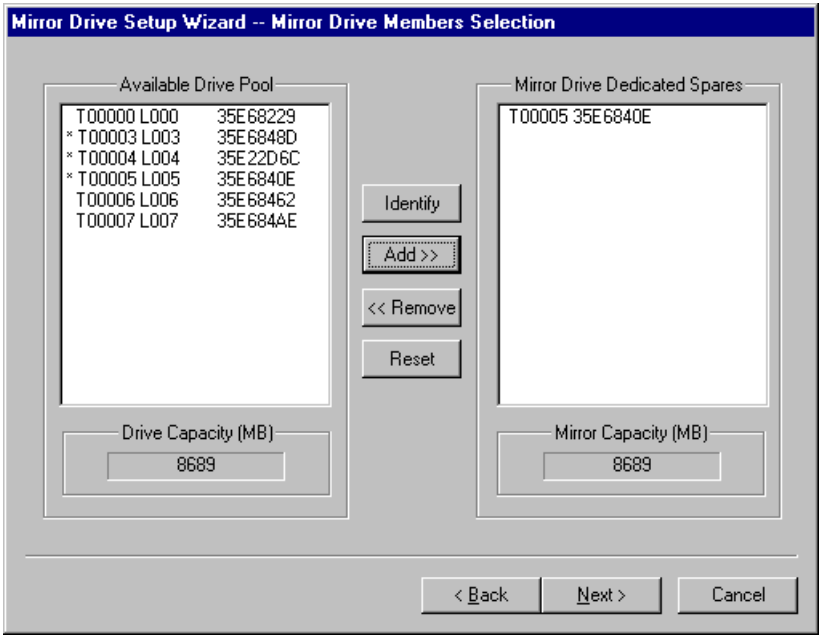


*Figure 6-21 Adding a dedicated spare*

Click the **Next>** button and the properties of the mirror drive can be changed. The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. A name can also be defined to the mirror drive for easier identification, with a limit of up to eight characters.

The Initialization Method box refers to whether or not to allow the operating system to write its signature on the Mirror drive.

The Assigning Mirror Drive Properties window is shown in Figure 6-22.

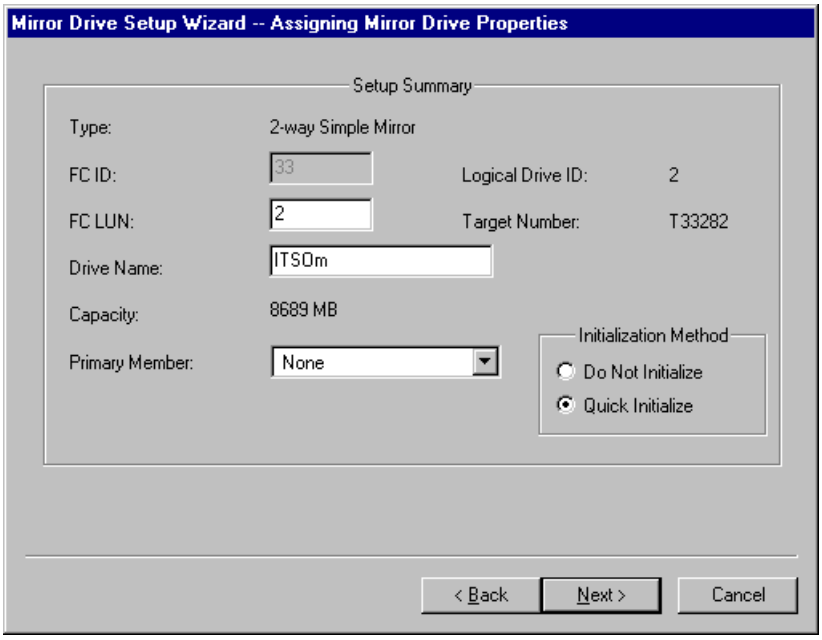


*Figure 6-22 Mirror drive properties*

Click the **Next>** button and a dialog box appears to allow you to create another mirror drive. Click **Yes**, if you would like to create another mirror drive, and the Mirror Drive List window opens, and the steps described above can be repeated.

Click **Finish** when you have created all the desired mirror drives. Up to this point, the configuration has been kept within the Controller 160 Manager software. When the **Finish** button is clicked, the Controller 160 Manager will now communicate to the Controller to complete the process and update the Controller to control the drives.

If **Quick Initialize** in the Initialization Method box was selected, the Controller will take a short period of time to write the host signature and build the mirror. During this time, if you try to communicate to the Controller, you may experience a slower than normal response.

## 6.6.2 Mirror drive properties

If you go to the Control Center window by selecting **Tools -> Control Center**, you will see that the mirror drive is now displayed in the Mirror Drive window. This is shown in Figure 6-23.

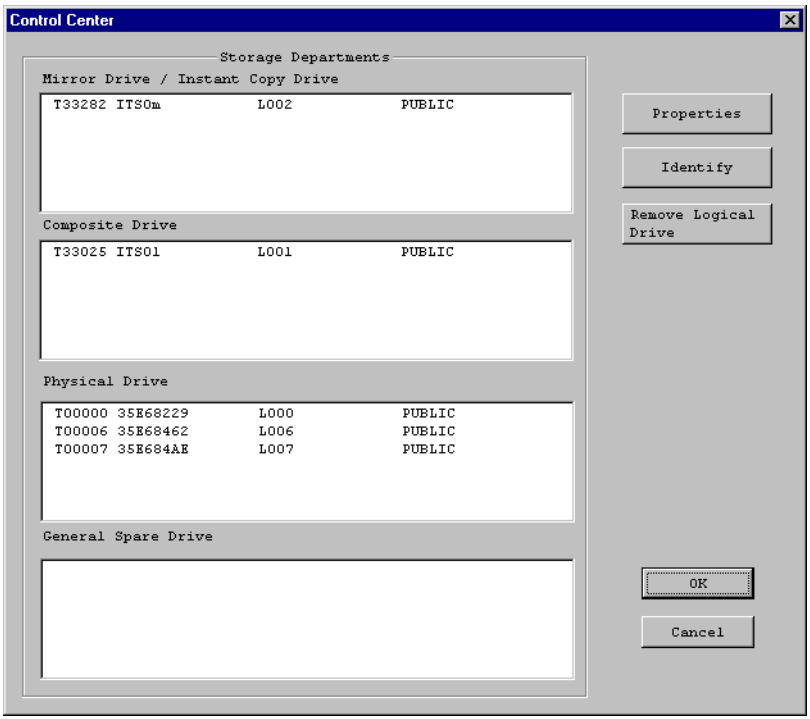


*Figure 6-23 Control Center with Mirror Drive*

If you select and highlight the mirror drive and then click the **Properties** button, the Mirror Properties window opens, and you can see the information in the mirror drive. Figure 6-24 shows an example of the properties of the mirror drive.

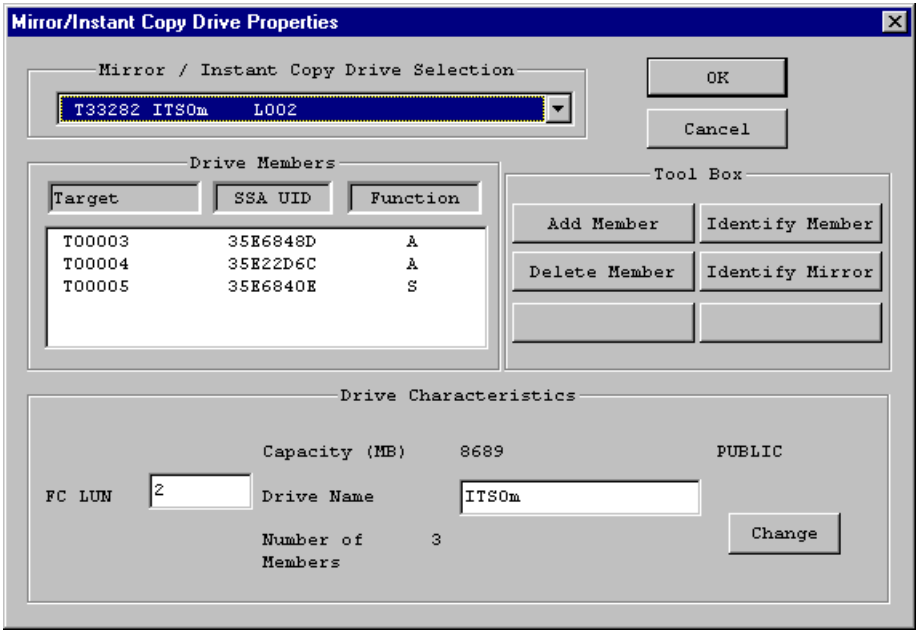


*Figure 6-24 Mirror Drive Properties*

Clicking on the **Identify Mirror** button will cause the LED on the actual SSA drives that belong to the mirror, to flash. If a member drive is selected and the **Identify Member** button is chosen, then the LED on that drive only will flash. In both cases, a dialog box appears to allow you to stop the flashing.

# 6.7 Instant Copy drive

Instant Copy is a feature that allows a drive to become part of a mirror, synchronize to the latest data, and then detach from the mirror. The drive can then be used to back up the data or used elsewhere if desired.

## 6.7.1 Creating an Instant Copy drive

To create an Instant Copy drive, select **Tools -> Instant Copy Drive Setup Wizard**. You will see a dialog box, Instant Copy Drive List, that will be blank. If there were copy drives created, it would display the names of the drives. Click the **Next>>** button and the **Instant Copy Drive Members Selection** window is displayed. The window on the left named *Available Drive Pool* contains a list off all drives that are candidates to become a copy drive.

Select a drive by highlighting it and click the **Add>>** button. The drive name will be added to the Member window. An asterisk will appear on the left hand side of the drive just selected in the Available Drive window to denote that the drive has been selected. An example is shown in Figure 6-25.

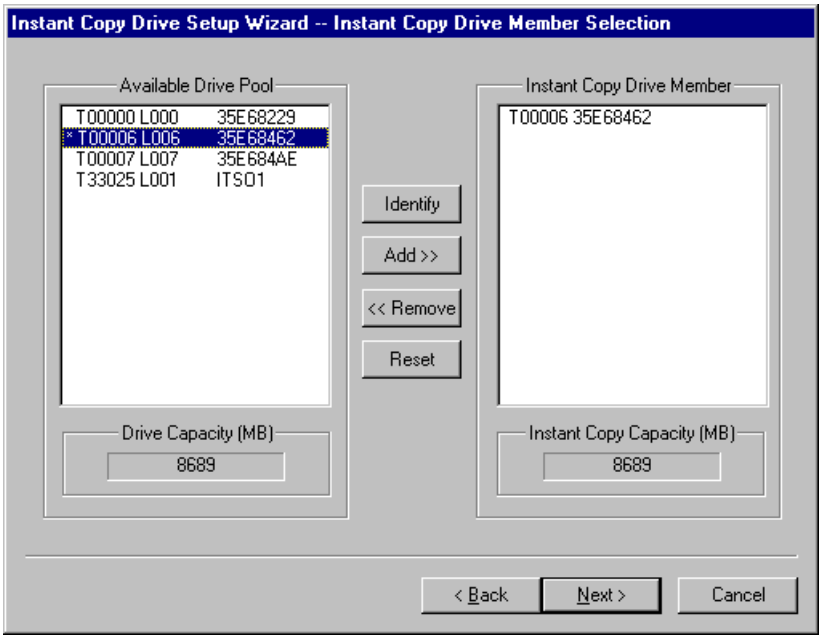


*Figure 6-25 Instant Copy Drive Member Selection*

To remove a drive from the Member window, select the desired drive and click the **Remove<<** button. Below each window there is a Drive Capacity box. As a drive is selected, its capacity in megabytes is shown. Click the **Next>** button to continue to the Assigning Instant Copy Drive Properties window as shown in Figure 6-26.

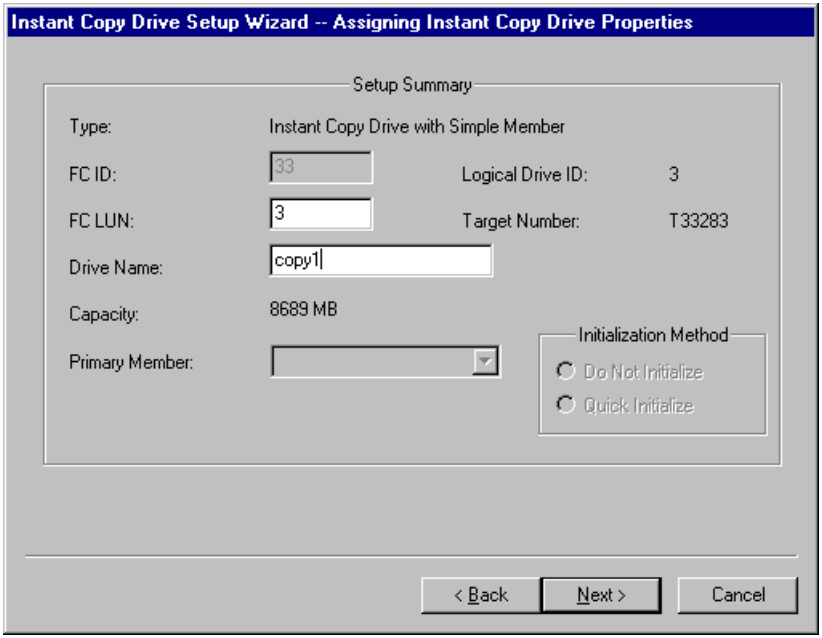


*Figure 6-26 Instant Copy Drive Properties*

The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. A name can also be defined to the copy drive for easier identification, with a limit of up to eight characters.

Click the **Next>** button and a dialog box appears to allow you to create another copy drive. Click **Yes** if you would like to create another copy drive and the Instant Copy Drive List window opens, and the steps described above can be repeated.

Click **Finish** when you have created all the desired copy drives. Up to this point, the configuration has been kept within the Controller 160 Manager software. When the **Finish** button is clicked, the Controller 160 Manager will now communicate to the Controller to complete the process and update the Controller to control the drives.

## 6.7.2 Instant copy drive properties

You can go to the Control Center window by selecting **Tools -> Control Center**. The copy drive that was created above can now be seen in the Mirror Drive/Instant Copy window as shown in Figure 6-27.

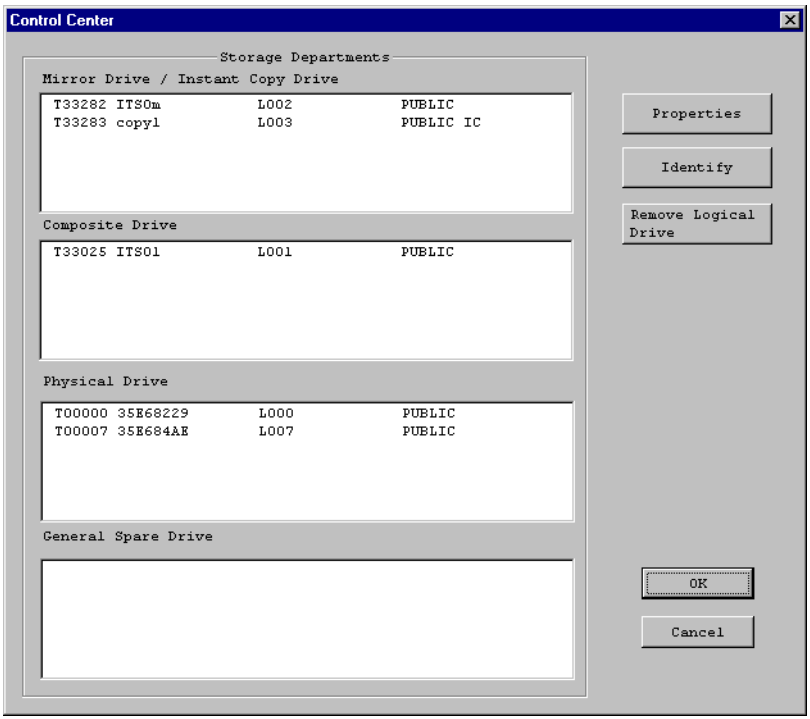


*Figure 6-27 Control Center with Instant Copy Drive*

Notice that in the information provided for the copy drive, there is an **`IC`** included to distinguish between mirror drives and copy drives within this window.

## 6.7.3 Adding an Instant Copy Drive to a mirror

To add or detach the copy drive from a mirror, you select and highlight the mirror drive, and then click the **Properties** button. The Mirror Drive Properties window opens as shown in Figure 6-24 on page 646. Click the **Add Member** button and the Add Mirror Member window opens as shown in Figure 6-28.

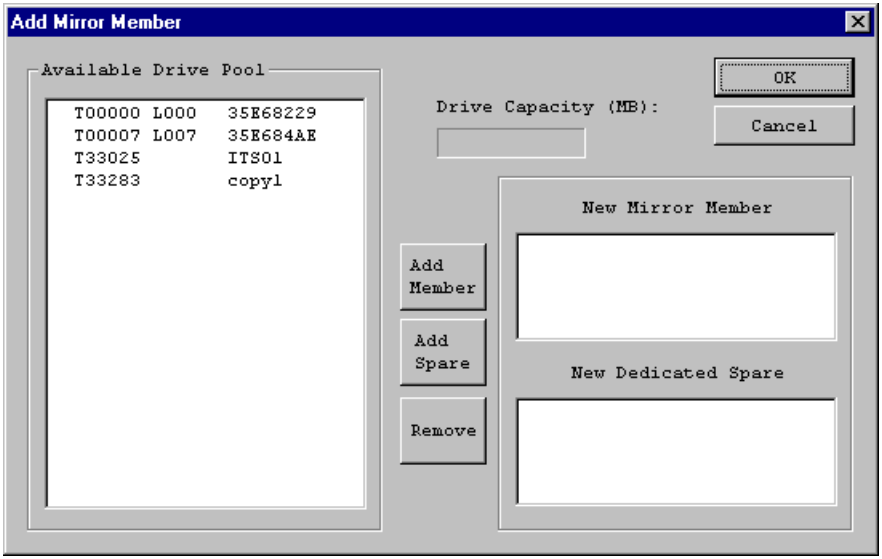


*Figure 6-28 Add Mirror Member display*

Select and highlight the copy drive from the Available Drive Pool window, click the **Add Member** button, and the name of the copy drive will appear in the New Mirror Member window. This is shown in Figure 6-29.

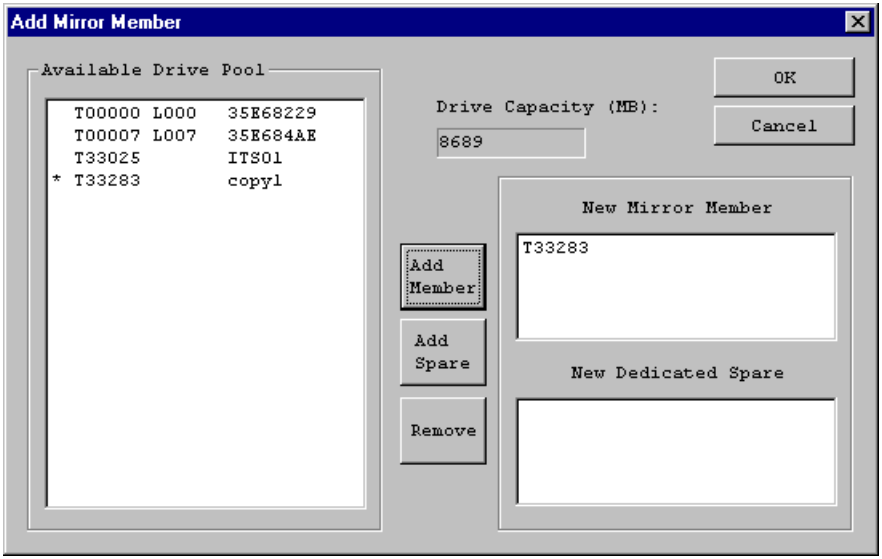


*Figure 6-29 Adding drive members to a mirror*

Click the **OK** button, and the Mirror Drive Properties will now reflect the change as shown in Figure 6-30.

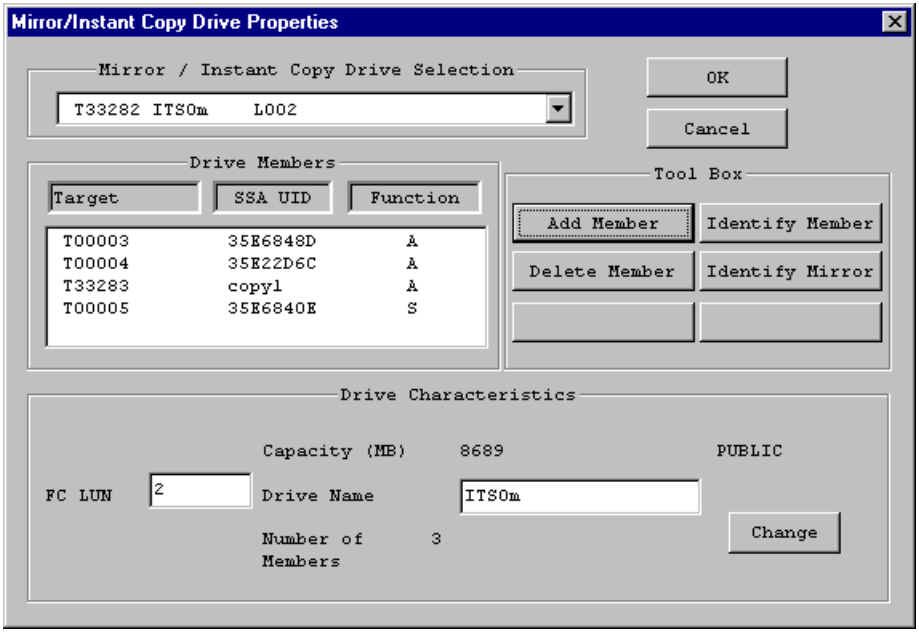


*Figure 6-30 Mirror drive properties with copy drive attached*

Click **OK** to complete the process.

## 6.7.4 Detach Instant Copy Drive from a mirror

To detach, or split off the copy drive from the mirror, the procedure is similar except at the Mirror Drive Properties window, select **Delete Member**. A window will appear that displays all current members of the Mirror. Select the Copy drive, and then delete it from the Mirror. The Copy drive can now be accessed by another host.

# 6.8 Combining composite and mirroring

The Controller 160 Manager can also be used to combine the two features of the Controller. You can create a mirror drive using composite drives. A mirror can have drive members of different sizes, but the actual mirror capacity will be the smaller of the drive sizes.

## 6.8.1 Creating a second composite drive

To provide an example of a mirror using only composite drives, another composite drive is required. The example shown in Figure 6-31 shows that drive 6 and 7 were used to create another composite drive.

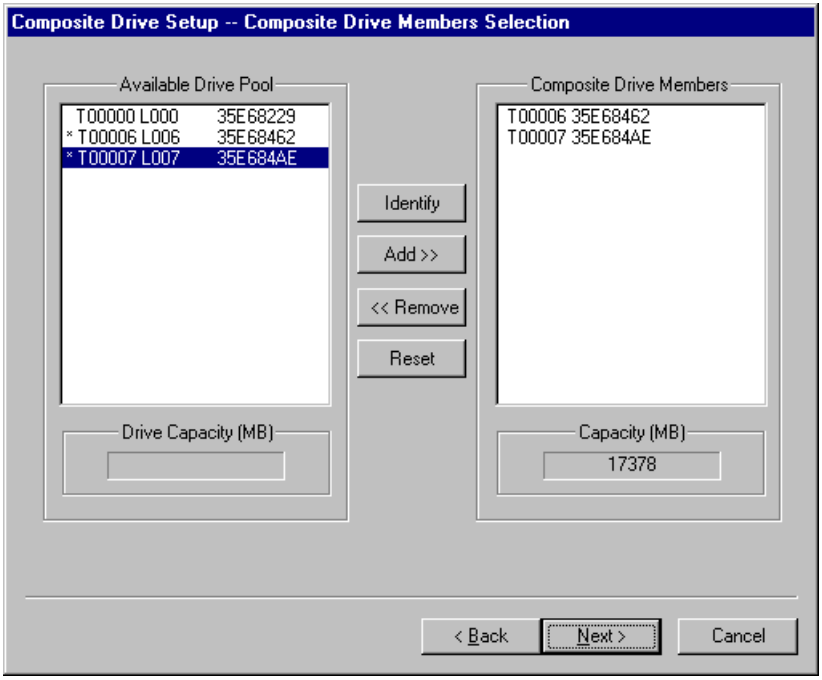


*Figure 6-31 Creating composite drive to be used in a mirror*

Follow the steps to create a composite drive as described in 6.5, “Composite drive” on page 636. Once created, you can view the Control Center window by selecting **Tools -> Control Center** from the toolbar. Figure 6-32 shows that there are now two composite drives.

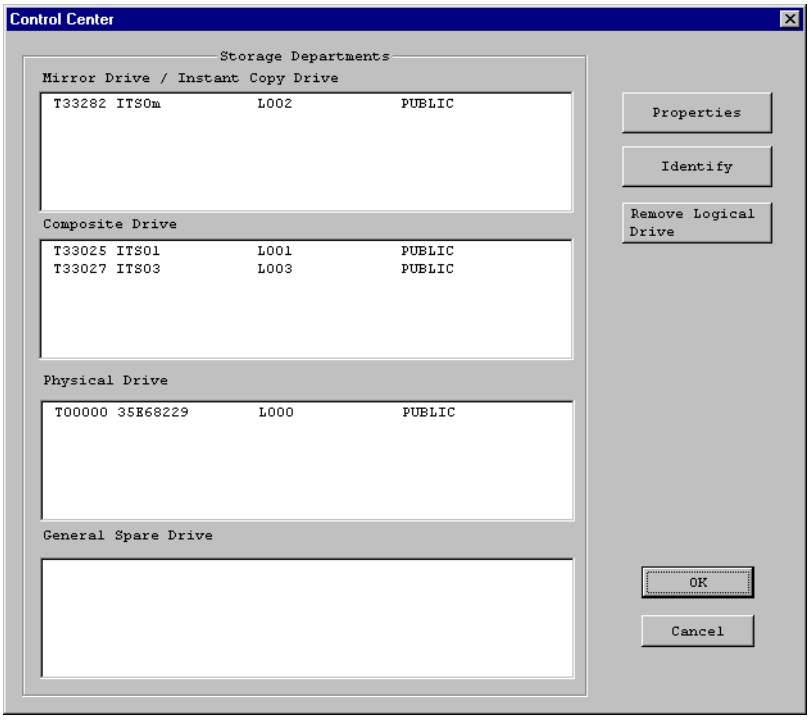


*Figure 6-32 Control Center with two composite drives*

## 6.8.2 Creating the mirror

The mirror can now be created by selecting **Tools -> Mirror Drive Setup Wizard**. When the Member Selection window appears, select the composite drives as members of a mirror. Figure 6-33 shows where composite drives ITSO1 and ITSO3 are selected as members of a mirror.

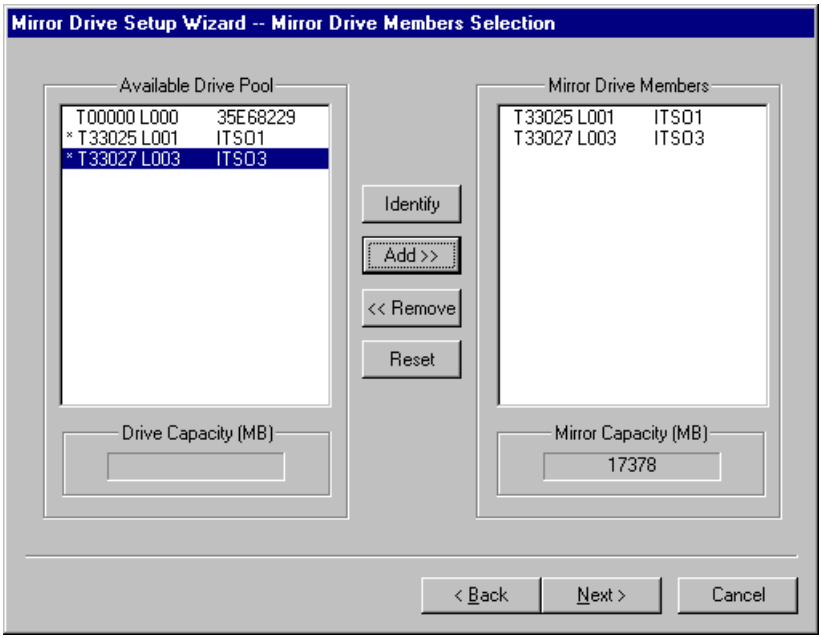


*Figure 6-33 Creating mirror drive from two composite drives*

Continue through the Mirror Drive Setup Wizard to complete the process as described in 6.6, “Mirror drive” on page 641.

## 6.8.3 Viewing mirror drive using composite drives

With the Mirror Drive Setup Wizard completed, you can now view the Control Center window once again, as shown in Figure 6-34.

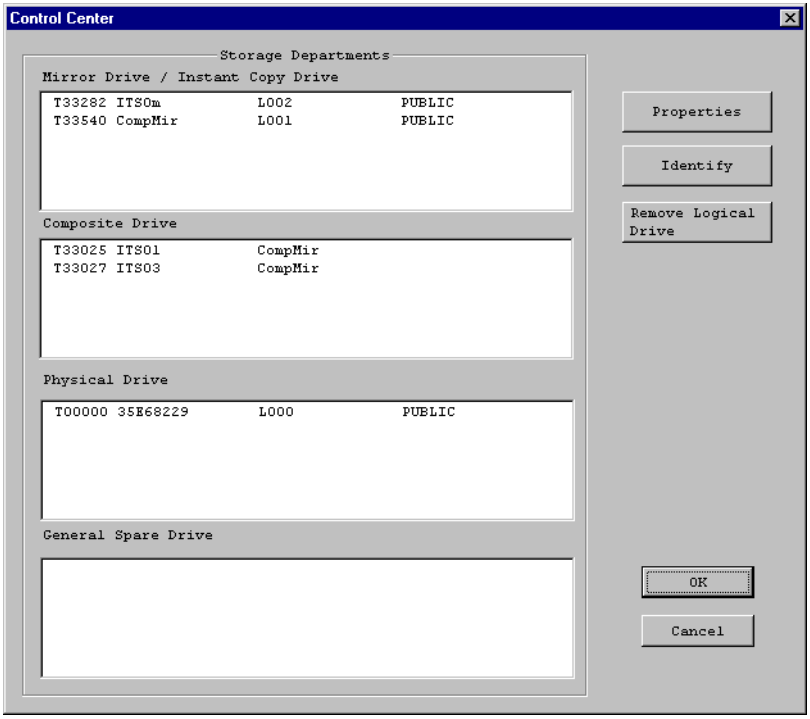


*Figure 6-34 Control Center with mirror drive using two composite drives*

In the Mirror Drive window, you can see the new mirror drive that was created above and named *CompMir*. In the Composite Drive window you can see that there are still the two composite drives, but instead of having Fibre Channel LUN numbers assigned to them, they are shown as belonging to a mirror with the name *CompMir*.

You can highlight the CompMir drive and click the **Properties** button. All the same functions that were described in 6.6.2, “Mirror drive properties” on page 645 are available.

## 6.9 Reusing logical drives

At some point the composite, mirror, and instant copy logical drives that have been created may no longer be required. The logical drive can be removed so that the member drives that made up the logical drive can then be used individually or reconfigured to make new logical drives.

### 6.9.1 Remove a logical drive

To remove a logical drive, you access the Control Center by selecting **Tools -> Control Center** from the top toolbar. At the Control Center window, select the logical drive (composite, mirror, or copy) that you want to remove. Select the **Remove Logical Drive** button on the right hand side and a dialog box appears that will ask you to confirm that you want to remove the logical drive.

Once it is removed, the member drives will become general spares and will show up in the General Spare Drive window of the Control Center. This is shown in Figure 6-35.

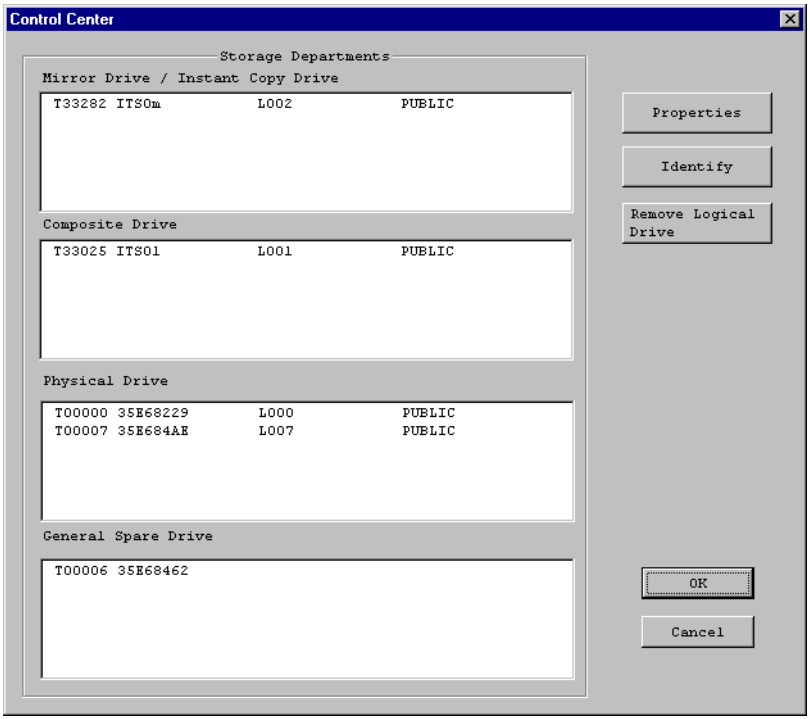


*Figure 6-35 Removing a logical drive*

The logical drive that was created as an Instant Copy Drive in 6.7.1, “Creating an Instant Copy drive” on page 646, has been removed and is now a general spare.

## 6.9.2 Mapping a general spare

You will notice in Figure 6-35 that the general spare does not have a LUN number assigned to it. To get a new LUN number for this drive, you select the drive and click the **Properties** button.

The Drive Properties window appears; select the **Change FC** button. A dialog box opens as shown in Figure 6-36.

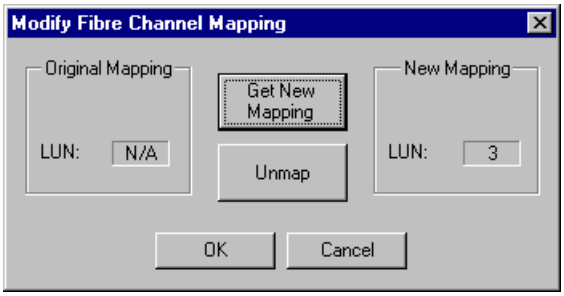


*Figure 6-36 Mapping a general spare*

The Original Mapping box will indicate that no LUN was assigned previously. Click the **Get New Mapping** button, and the next sequential available LUN number will appear in the New Mapping box. Click **OK**. The drive will now appear in the Physical Drive window of the Control Center.

This can also be used to modify the LUN number on an existing mapped drive, as well as remove the LUN number to “unmap” a drive and create a general spare.

It is not necessary to map a general spare. A general spare can be used to create a composite, mirror, or copy drive. Mapping a general spare will create a drive that has a LUN number that can then be used by the host.

## 6.9.3 Removing a mirror containing composite drive

The mirror in this case was made from logical drives on their own. Once the mirror is removed, the composite drives that made up the mirror will return to the Composite Drive window as viewed from the Control Center.

However, since each composite drive had its attributes changed as it became a member of the mirror, it will no longer be mapped. The composite drives will show up as **`UnMapped`** in the Control Center window. This is shown in Figure 6-37. The mirror created in 6.8.2, “Creating the mirror” on page 654 was removed.

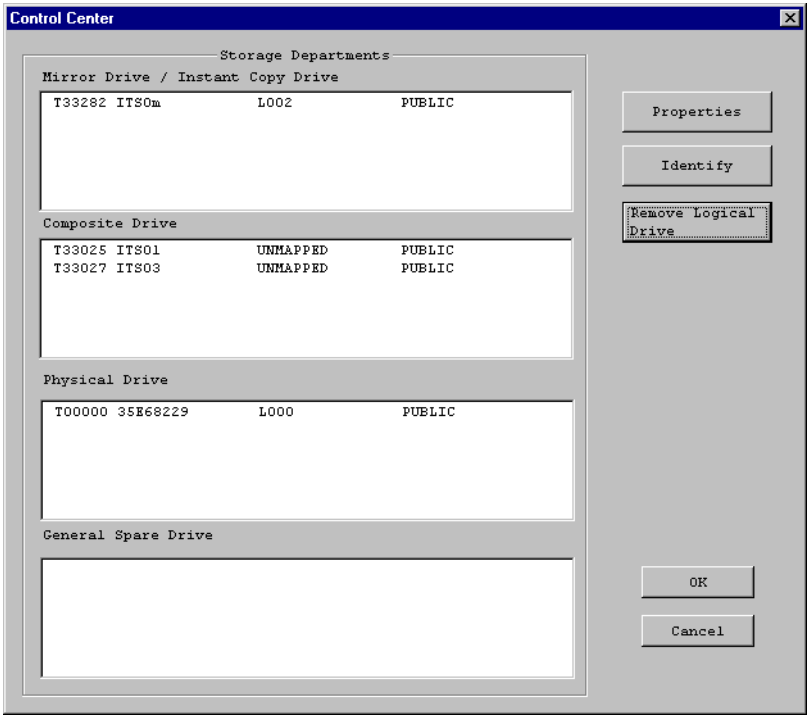


*Figure 6-37 UnMapped composite drives*

The existing composite drives ITSO1 and ITSO3 cannot be mapped or given a new LUN number at this point. Each logical composite drive *must* be removed as well. This will cause the composite drive to be removed and its member drives to become general spares. Once they are general spares, the drives can be used to recreate the composite drive or to create new logical drives.

## 6.10 Expanding the Controller 160 system

The Controller 160 storage system can be expanded to add more SSA disks or more SAN Controller 160s. Each Controller 160 storage system can support up to 64 SSA disks and have 16 Controllers.

## 6.10.1 Adding disks

To add disks to an existing Controller 160 system is very easy because they are SSA disks, and the rules for SSA disks apply here as well.

If there is a dummy drive in an existing SSA drawer, then it can be replaced by a real drive. If a new drawer has to be added, the SSA cabling is changed to include this new drawer.

Because this is SSA, this can be done *on the fly.* As the SSA loop is broken, the Controller will still access all disks due to the structure of the SSA loop. If possible, we recommended that you stop host access and power down the loop. In any case, the rules regarding SSA disks and cabling must be adhered.

As disks are added to an existing loop, the new disks will be recognized. If all disks in the loop are used as single disks (JBOD) and have LUN numbers assigned, the new disks added will have LUN numbers assigned to them automatically. If there are any composite, mirror, instant copy, or spare drives in the loop, then the new disks will not have LUN numbers assigned and become general spares.

## 6.10.2 Adding Controllers

By adding Controllers we can increase the amount of storage a host can access and increase throughput. On the rear panel of the Controller, there are two Fibre Channel GBIC ports that are available and act as a mini-hub.

You can add a Fibre Channel cable from the second port on the existing Controller to one of the ports on the second Controller. You are basically daisy-chaining the Controllers. But since the ports on the Controller act as a hub, an arbitrated loop is created. However, in this scenario there is only one Fibre Channel cable from the host to the Controller and it is a single point of failure.

Another option is to add a second Fibre Channel host adapter that will connect to the other Controller. This provides a high availability feature, because there are now two paths to the storage system. Software must be used for automatic failover and load balancing between the two Fibre Channel host adapters. Failover also can be done manually if so desired.

On the SSA side, there are a few options available. Each Controller can have its own SSA loop so that each one can support 64 SSA disks. In this way, storage capacity is scalable, because it can be increased by adding more Controllers. This is shown in Figure 6-38.

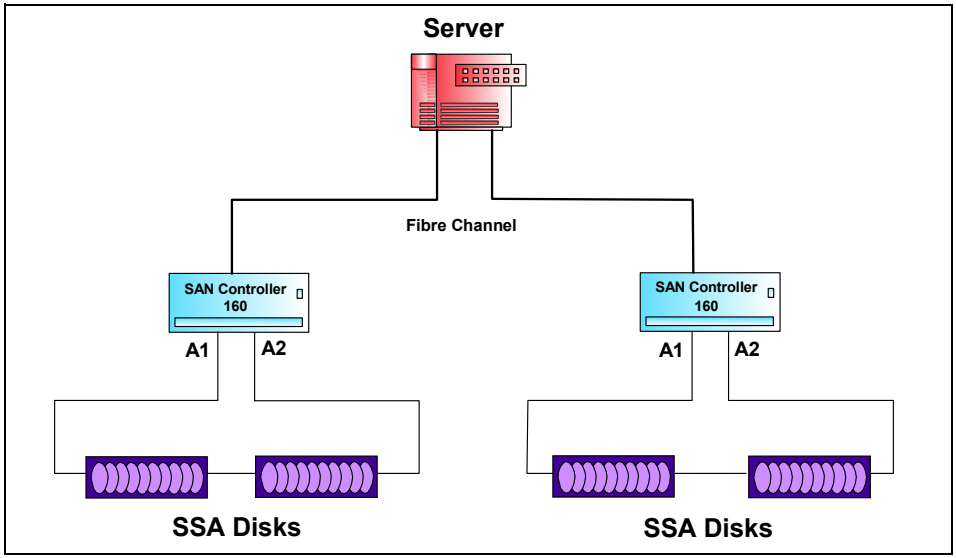


*Figure 6-38 Increasing storage capacity*

The other option is to have each additional Controller added to the same SSA loop. Throughput to the SSA loop will increase, because each Controller can access the disks for multiple simultaneous operations. This configuration is shown in Figure 6-39.

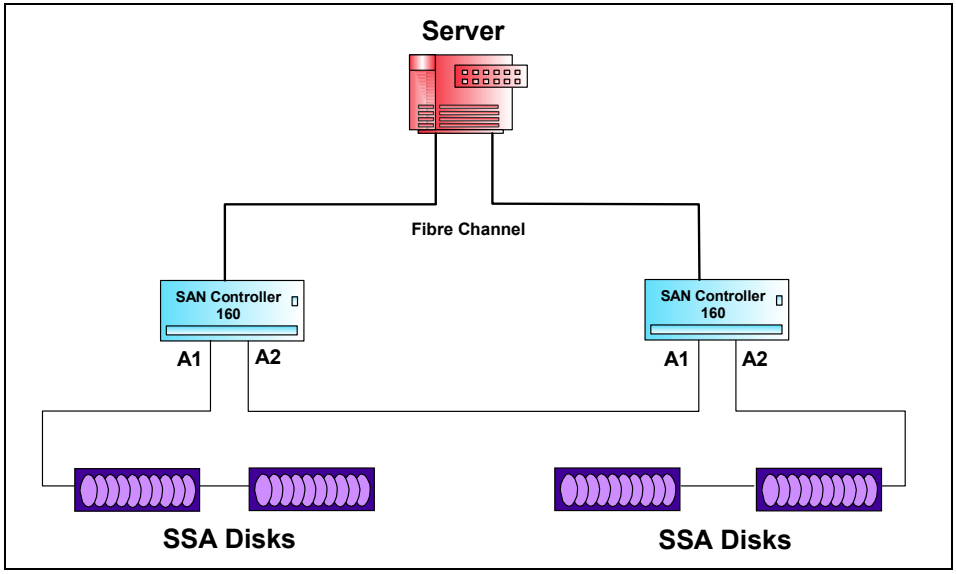


*Figure 6-39 Increasing throughput*

## Installing additional Controllers

With the existing Controller and storage system powered on:

1. The node map on the new Controller must be cleared first.
2. Connect the Controller to the existing system with the power off.
3. Set SW2 to mode 3 and set SW1 to an unique Fibre Channel target ID.
4. Power on the new Controller.
5. When the Status LED on the new Controller is on (solid lit), the install is complete.

When the new Controller is powered on, communication between the two Controllers will occur to query and update the new Controller. The new Controller will be added as a subordinate so that the first Controller will maintain its master status.

Any configuration changes to the storage system is always done on the master. As this is done, the changes are also communicated to the subordinate Controllers so that all systems in the loop are aware of what is happening.

## Using Controller 160 Manager on additional Controllers

Since the master Controller does all the work, it is not necessary to use Controller 160 Manager to view the new Controller. However, depending on how the new Controller is used this may become a requirement. The same daemon used to connect to the first Controller can be used to connect to several Controllers.

The configuration file is edited again and the process to name and create a Controller 160 zone on this new Controller can be added within the same file. The Controller 160 Manager software can now be used to communicate to the new Controller.

The Controller 160 Manager can only communicate to a single Controller at a time. Multiple Controller 160 Manager sessions can be started to communicate to each Controller.

## Master failover

If at some time the Controller that is designated as the master within the Controller 160 storage system fails, the master designation will failover to next nearest Controller. This is accomplished within the communications between the Controllers and it is done automatically.

When the new Controller accepts becoming master, it will maintain the master role if even the failed Controller is replaced and rejoins the storage system. The master role can be changed back to the original Controller, or to another Controller if desired, using the Controller 160 Manager software.

There is no capability to select a specific “failover” Controller.

## 6.10.3 Adding hosts

The Controller 160 storage system can be expanded to include more hosts whether they are homogeneous or heterogeneous. It is recommended that as hosts are added, each host is connected to its own and separate SAN Controller 160.

If more than one host was connected to a single Controller, there will be arbitration and performance issues. Also, it would have a single point of failure with the possibility of losing data access to many systems.

### Homogeneous hosts

If another host of is added and you would like to have both hosts access the same disks, then some sort of access sharing software must be loaded onto both hosts.

If other hosts are added to the storage system and they will not share data, but are connected for storage consolidation, there are a few issues to be considered as the Controller does not provide a LUN masking capability.

In UNIX systems, the hosts will see all disk in the loop. But, if the specific volume is not mounted there will be no data integrity problems.

For Windows 2000, each host will write its own signature on all available disks. Adding another Windows 2000 host to the loop will cause problems. To allow a specific Controller, and host attached to that Controller, access to a specific disk or set of disks, you can set Private Attributes on the disks.

Private Attributes is a setting within the Controller 160 manager that can set a disk to only be accessed by a certain Controller and in turn the host attached to that Controller.

**Note:** For more information and operation on the Private Attributes setting, please refer to the *Controller 160 Manager Installation and User Guide*, 310-605807.

In all cases, if extra control for disk access is required, third party software must be used.

### Heterogeneous hosts

As the Controller does not provide for LUN masking, you must use the Controller 160 Manager Private Attribute setting or third party software to restrict and control host access to the disk.

7

# Implementing the SAN Data Gateway

In this chapter, we describe the steps involved in planning and implementing the IBM Storage Area Network Data Gateway. The SAN Data Gateway is a hardware solution to allow connection of Fibre Channel ready host systems to attach to SCSI storage systems.

**Attention:** While the IBM 3583 SAN Data Gateway Module (SDGM) is very similar to the 2108-R03, we do not cover the SDGM here. For specific SDGM information, refer to the appropriate 3583 documentation.

# 7.1 SAN Data Gateway

The SAN Data Gateway provides several benefits to bridge the legacy gap as storage products migrate from SCSI based attachments to Fibre Channel.

A diagram to show a SAN Data Gateway configuration using a single host is shown in Figure 7-1.

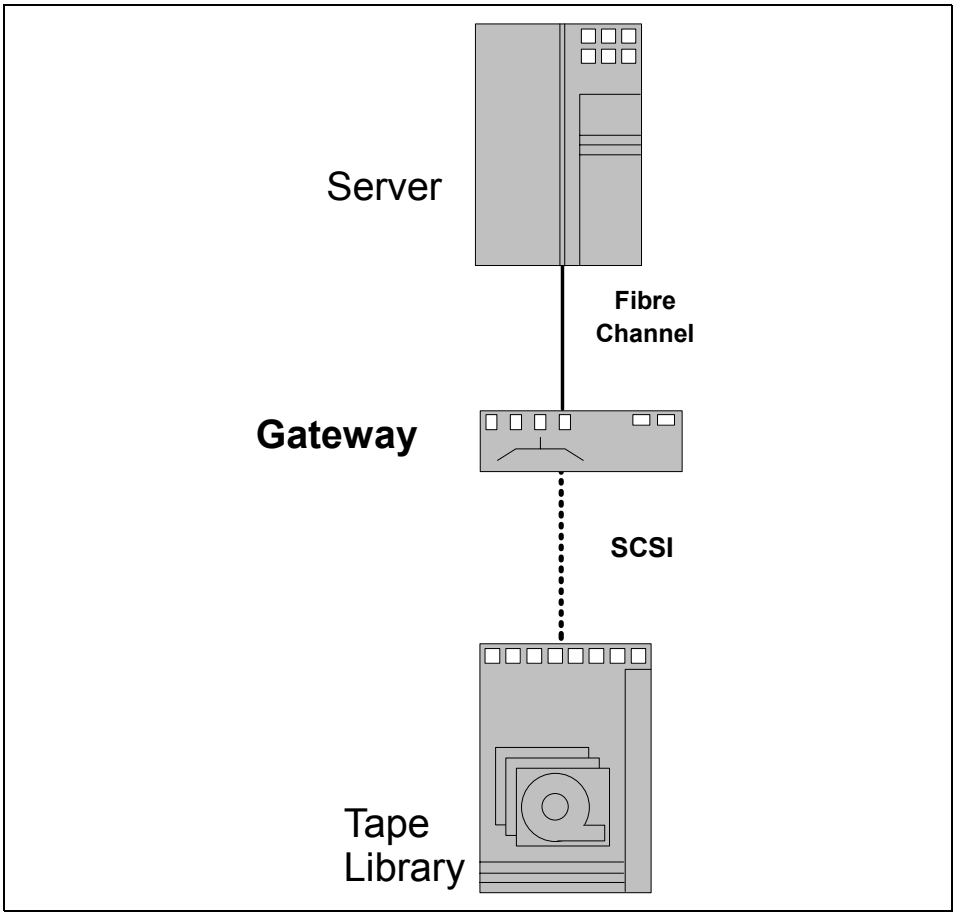


*Figure 7-1 SAN Data Gateway configuration*

The IBM Storage Area Network Data Gateway allows you to:

- Protect legacy storage equipment while utilizing the latest host servers with Fibre Channel support
- Expand connectivity to storage devices with use of IBM SAN hubs, switches, and directors

- Perform channel zoning and LUN masking capability to allow access at a volume level
- Overcome the distance limitations of SCSI based host systems using longwave ports that support distances up to 10 km
- Utilize the StorWatch SAN Data Gateway Specialist, which is an easy-to-use interface for managing and controlling access of host systems to storage devices

The SAN Data Gateway is available as a rack-mount unit or as a stand-alone tabletop unit. The gateway model provides two shortwave Fibre Channel ports and four Ultra SCSI Differential ports to attach disk or tape storage devices. One or two Fibre Channel cards — dual-port, shortwave and/or single port, longwave — may be added for a maximum of six shortwave ports, or two shortwave and two longwave ports. If you are using the dual-port shortwave cards, Figure 7-2 depicts the port assignment numbers for the optical interfaces.

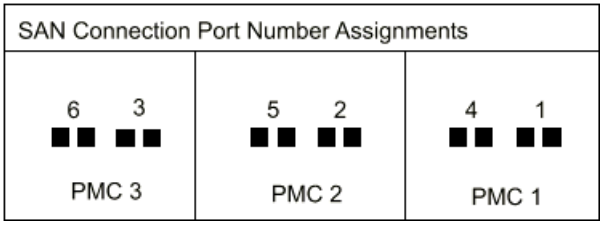


*Figure 7-2 SAN connection port assignment*

In the following section we give you an overview of:

- Basic install setup of the SDG
- Installation and usage of the StorWatch SDG Specialist
- How LUN mapping works
- How to use channel zoning
- Virtual Private SAN

For more information on the SAN Data Gateway, we refer you to these books:

- *2108 Model G07 Installation and User's Guide,* SC26-7304
- *2108 Model R03 Installation and User's Guide,* SC26-7355

## 7.2 Installation

In this section we briefly summarize the main steps for setting up the SAN Data Gateway, detailed information is provided in the Installation and User's Guides listed in the previous section.

1. After installing the hardware, set up a connection between the SDG's serial port and a serial port on a PC, using the supplied null modem cable and a PC terminal program, such as Netterm or Hyperterm. Use the following parameters to configure the terminal session:

   – Bits per second: 19200
   – Data bits: 8
   – Parity: None
   – Stop bits: 1
   – Flow Control: Xon/Xoff
   – VT-100 compatible terminal

   Once connected, power on the SAN Data Gateway and the start up messages will appear and scroll across the window. When the power on sequence has completed, a prompt **`Gateway>`** or **`Router>`** appears on the window as shown in Figure 7-3.

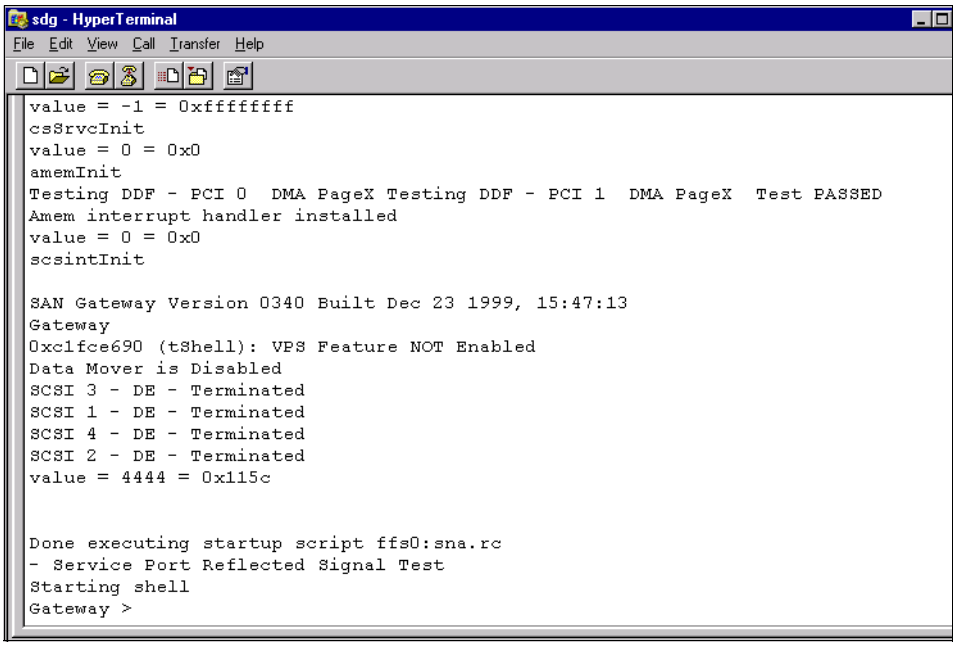


*Figure 7-3 IBM Storage Area Network Data Gateway startup*

If you type in **`help`** and then press Enter, a list of available commands is provided. The commands are case sensitive and *must* be entered as they appear.

2. We suggest the following basic commands sequence to install a new or an already used SDG:

   a. **initializeBox**

      This restores the factory defaults by deleting all configuration files including persistent address map and VPS databases, and then automatically restarts the Gateway.

   b. **ethAddrSet**

      Once restarted, the Ethernet port must be configured and attached using network information provided by the network administrator. To set the IP address, use the **ethAddrSet** command. The address must contain the double quotes (“):

```
Gateway > ethAddrSet “9.111.24.66”
Network not Enabled
Write complete
Host Address set to 9.111.24.66 for Ethernet interface
value = 0 = 0x0
Gateway >
```

      If a subnet mask is required then add it after the IP address and separate the two addresses using a comma. For example:

```
Gateway > ethAddrSet “9.111.24.66”, “255.255.255.0”
```

   c. **gateAddrSet**

      If a gateway or standard router is to be specified, then issue the **gateAddrGet** command to view if there is one set and the **gateAddrSet** command to set or change it. For example:

```
Gateway > gateAddrGet
No current gateway address set
value = 0 = 0x0
Gateway > gateAddrSet “193.1.1.11”
Write complete
```

   d. **ethEnable**

      The Ethernet port on the SAN Data Gateway comes from the factory disabled. To enable it, you must issue the **ethEnable** command:

```
Gateway > ethEnable
Write complete
Ethernet will be enabled on next boot
value = 0 = 0x0
Gateway > _
```

      This will not take effect until the unit is rebooted. The reboot can occur from a power off, or by issuing the **reboot** command. During the reboot, you will see that the IP address is set and now enabled.

e. **userAdd "admin","password"**

If a user would prefer to telnet to the SAN Data Gateway rather than by using the service terminal port after initial setup, this can be done. First you must create a user from the service terminal by using the **userAdd** command. Enter the login name and password using the quotes and comma, the user name variable must be 3 to 80 characters, and the password variable must be 8 to 40 characters, for example:

```
ITSO > userAdd “itso”, “itsouser”
value = 0 = 0x0
ITSO > _
```

**Restriction:** You cannot telnet to the Gateway and use the service port at the same time. When you telnet to the Gateway, the service port on the rear of the unit will stop its communications. After you end the telnet session, then the service port will become available again.

f. **disableCC** or **enableCC:**

The **disableCC** command disables the command and control interface (LUN 0). The SDG is addressable as a SCSI target device for command and control support. On a Fibre Channel interface, this device is seen as logical unit number 0, (LUN 0). If the command and control support is enabled, then you might see problems on AIX machine when running **cfgmgr**:

```
cfgmgr: 0514-621 WARNING: The following device packages are required for
device support but are not currently installed. devices.fcp.array
```

Therefore we recommend to **disableCC** if the SDG is connected to an AIX server. For other operating systems you can **enableCC**.

g. **setHost X,HHH** (where X can be FC Port 0,1,2,3,4,5,or 6 and HHH can be "solaris", "nt", "hpux", "aix", ”switch”, ...)

The **setHost** command sets the operating system type for the specified SAN interface. This customizes the way the SDG is presented to the particular operating system. If the port is 0, the change applies to all SAN connections; otherwise, the host type is applied only to the SAN connection on the specified interface. The default setting is NT. Currently, the OS can be specified as NT, AIX, Netware, HP-UX, or Solaris. You have to put the hosttype in double quotes as shown in Example 7-1.

*Example 7-1 setHost for FC 4 to AIX*

```
Router > setHost 4,"aix"
value = 0 = 0x0
```

**Note:** Ports 1 and 4 belong to the first FC Card; 2 and 5 belong to the second FC Card; 3 and 6 belong to the third FC Card inside the SDG, as these were shown in Figure 7-2 on page 667.

3. **reboot:**

   A reboot makes the above changes effective. You can now connect to the SDG via a telnet session, using the Ethernet address, and the userid you defined.

## 7.2.1 Startup sequence

You must start up the SAN Data Gateway and the attached host and target devices in a specific order. When you add or remove SCSI devices or update firmware, you must restart. The following procedures describe the situations and order of procedure when you restart the SAN Data Gateway.

Before you restart the SAN Data Gateway, you must stop all input and output (I/O) activity between the host and SCSI devices.

1. **SCSI devices:**

   Turn on the SCSI devices. You must turn on all SCSI devices attached to the SAN Data Gateway before you initially turn on or restart the SAN Data Gateway.

2. **SAN Data Gateway:**

   The SAN Data Gateway scans the SCSI buses when it starts. If you add or remove SCSI devices after the Gateway has started, the Gateway will not detect the changes. You can invoke an SCSI rescan or restart operation from either the StorWatch SAN Data Gateway Specialist client or the service terminal.

3. **Fibre Channel host:**

   Before you turn on or restart the hosts that are connected with Fibre Channel to the SAN Data Gateway, you must wait until the SAN Data Gateway has finished starting. You will know the Gateway has finished starting when the ready light on the front panel blinks at frequency intervals of one second.

   – Some operating systems provide you with software methods that allow you to add or remove SCSI devices dynamically after the host has started. To ensure reliable operation, restart the host.

   – If you update SAN Data Gateway firmware, you must restart the Gateway to use the new firmware. To ensure compatibility between the firmware features or functions and the host, restart the host.

– If you update SCSI device firmware, the SAN Data Gateway Explorer application does not display the new firmware version until the SAN Data Gateway has issued an SCSI inquiry. The SCSI inquiry occurs when the Gateway rescans the SCSI buses. The SCSI inquiry also occurs when the StorWatch SAN Data Gateway Specialist client application or the service terminal rescans the SCSI buses.

Currently, up to eight different hosts can be attached to each Fibre Channel port. If all six ports are installed, then 48 different hosts can attach to the SAN Data Gateway.

## 7.2.2 Displaying devices

Using the **fcShowDevs** command, we can display information about the devices that are accessible from each Fibre Channel interface. The display shows the LUN that the SDG has assigned to each device, the SCSI Channel that the device is attached to, the actual SCSI ID and LUN of the device, the vendor, product, revision and serial number of the device.

The **fcShowNames** command displays the node and port names (addresses) of the Fibre Channels. If the output does not meet your physically installed devices, then execute **scsiRescan** or **reboot**. Collect the output from **fcShowDevs** and **fcShowNames** for further use (**fcShowNames** shows you the WWN, the assigned LUN and the serial number of the tape drives as shown in Example 7-2).

*Example 7-2 Output of fcShowDevs and fcShowNames*

```
Router > fcShowDevs
FC 1:
LUN Chan  Id  Lun  Vendor   Product          Rev   SN
-----------------------------------------------------
  0   0    0    0  PATHLGHT SAN Router       32aC  21081341573
  1   2    6    0  IBM      ULT3583-TL       2.50  IBM7801954
  2   2    0    0  IBM      ULT3580-TD1      16E0  6811020764
  4   2    1    0  IBM      ULT3580-TD1      16E0  6811007030
value = 6 = 0x6

Router > fcShowNames
------------------------------------------------------------------
 Ctlr : PCI Addr : ISP   :        Node       :        Port
  Id  : Bs Dv Fn : Type  :        Name       :        Name
------------------------------------------------------------------
   1  : 00 06 00 : 2200  : 10000060.45161ff5 : 20010060.45161ff5
```

## 7.3 StorWatch SAN Data Gateway Specialist

The StorWatch SAN Data Gateway Specialist software provides remote capability for all management, configuration, and event notification. It is comprised of three parts:

- **Agent:** The agent is embedded in the operating system of each SAN Data Gateway to provide a stand-alone manageable host. The StorWatch SAN Data Gateway Specialist software uses SNMP to set and retrieve information that controls the operation of the Agent. The Specialist also uses SCSI over TCP to allow updates to the Gateway and target device.
- **Server:** The server is a Java application that runs on a host and is used to maintain communication with the agents and acts as an intermediary between the agent and the client. The server coordinates the request from multiple clients to manage multiple gateways or agents. Multiple clients can share data the server already knows about, and the server receives all traps from the agent and forwards them to the clients that are registered to receive them.
- **Client:** The client is a Java application that operates from any compatible computer as long as it has a TCP/IP connection to the server. One or more clients can connect to a server. The client provides the user interface to allow the management and configuration of the SAN Data Gateway.

The server and client can be installed on to the same computer.

The StorWatch SAN Data Gateway Specialistis available for the following operating systems:

- Windows 95,98, 2000, and NT 4.0 with SP5 or later
- AIX ver 4.3.3 or later
- Solaris 2.6 or later

### 7.3.1 Installing StorWatch Specialist

The Specialist software is not bundled with the IBM Storage Area Network Data Gateway. The Specialist software is downloaded using a Web browser by going to the IBM Storage Area Network Data Gateway Web site:

http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/sdgcdr.html

The Download Main Page window will load. Then select the specific operating system platform. Review the readme.txt file for the latest information and instructions, before installing.

This Web site also contains all the latest firmware for the SAN Data Gateway and drivers for supported host bus adapters.

**Important:** Your 2108-G07 should not be below v3.43.07 level of firmware, and your 2108-R03 should not be below v3.42.12 to include some important Data Path Protection improvements.

## Install on Windows

The StorWatch SAN Data Gateway Specialist software file is a self-extracting file. Once it has been downloaded, execute or run the file and it will automatically load onto your computer. During the install process you will be asked if you wish to install the Server, Client, or both, as shown in Figure 7-4.

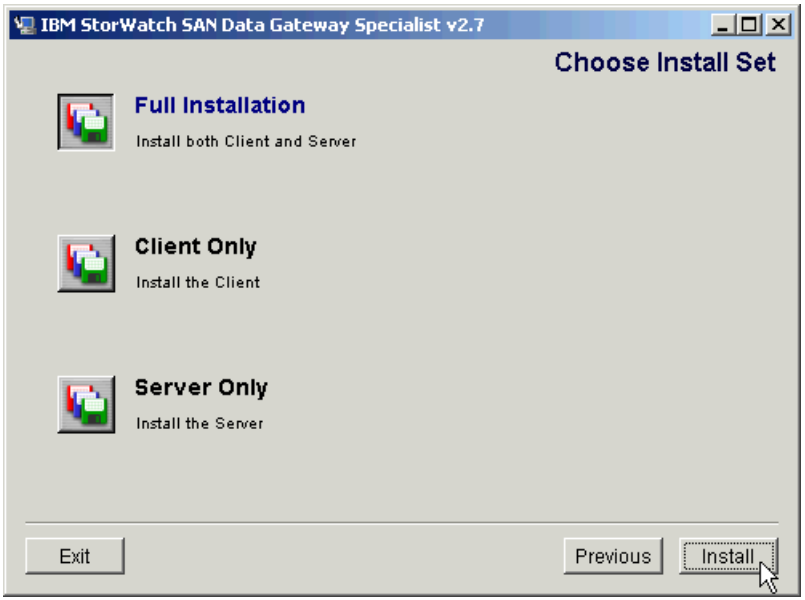


*Figure 7-4 Install of the SDG Storwatch Specialist on Windows*

## Install on AIX

Download the latest version and the readme file for AIX from the Web site above. If needed decompress and/or untar the files. The install images will be called something similar to storwsdg_bff. Install the server code and client code with:

```
installp -a -d storwsdg_bff all
```

If you want to install only the server code, type:

```
installp -a -d storwsdg_bff storwsdg.server
```

Or, to install the client code only, type:

```
installp -a -d storwsdg_bff storwsdg.client
```

### Install on Sun Solaris

Download the latest version and the readme file for Sun Solaris from the Web site above. If needed, decompress and/or untar the files. For installation of the files, enter:

```
pkgadd -d . IBMswsdg
```

During the installation you were asked if you want to install the client and if you want to install the server:

```
Install StorWatch SAN Data Gateway Specialist Client [y/n] ? y
Install StorWatch SAN Data Gateway Specialist Server [y/n] ? y
```

## 7.3.2 Starting the Specialist

To start the Specialist, the server must be started first, and then the client can be launched. Figure 7-5 provides an example of the StorWatch SAN Data Gateway Specialist with server and client loaded onto the same Windows NT computer.


*Figure 7-5 Starting the StorWatch SAN Data Gateway Specialist server*

Once the server has been launched, you should see a window similar to Figure 7-6.

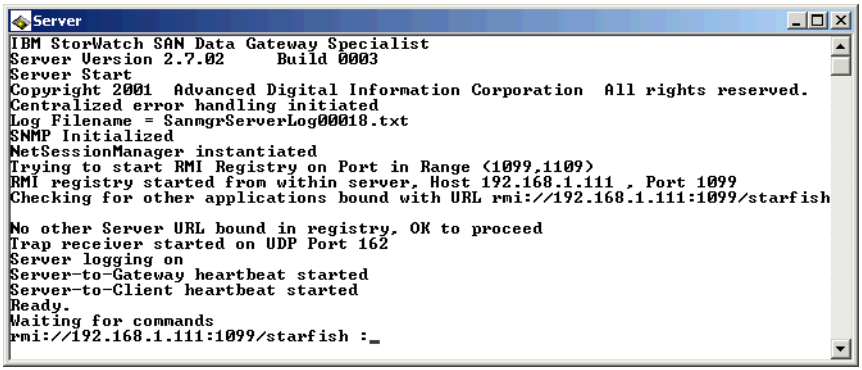


*Figure 7-6 StorWatch SAN Data Gateway Specialist server*

Similarly, we can start the server in the background with AIX or Sun Solaris:

```
/usr/bin/ServerLaunch -background &
```

The client software can now be launched. We start the client in Windows by clicking **Start —> Programs —> IBM StorWatch SAN Data Gateway Specialist —> Client**, as shown in Figure 7-5.

To start the client on AIX or on Sun Solaris, we type:

```
/usr/bin/ClientLaunch
```

If the server and client are not on the same computer, then we need to enter in the IP address of the computer that has the server software loaded in the dialog box that will appear, as shown in Figure 7-7. If the server and client are on the same computer, you will be automatically connected to the server.

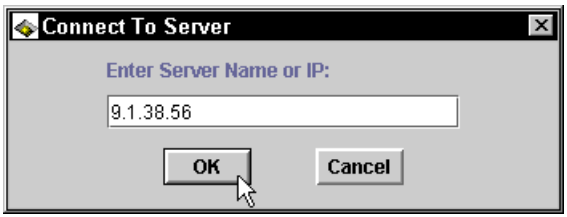


*Figure 7-7 SDG StorWatch connect to server*

After connection to the server is successful, a dialog box will appear, in which you can enter a user name and password as shown in Figure 7-8.

The IBM Storage Area Network Data Gateway provides a default user with administrator privileges:

- User Name: StorWatch
- Password: StorWatch

The fields are case sensitive, so they must be entered in as shown.

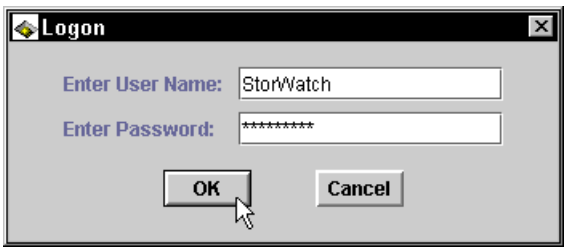


*Figure 7-8 SDG StorWatch logon*

A new administrator account can be set up by selecting **Admin —> Add User** from the pulldown menus. After a new administrator account is created, then the default user *StorWatch* is deactivated.

**Note:** If a new administrator account has been created, and the password is lost, and no other account has administrator access, then you will need to contact a service representative.

### 7.3.3 Using the StorWatch SAN Data Gateway Specialist

Once you are logged in to the Specialist, you must now connect to the SAN Data Gateway. At this point a dialog box, as shown in Figure 7-9, should appear, requesting the IP address of a SAN Data Gateway. As it connects, it will download the information from the SAN Data Gateway and be presented on your window.

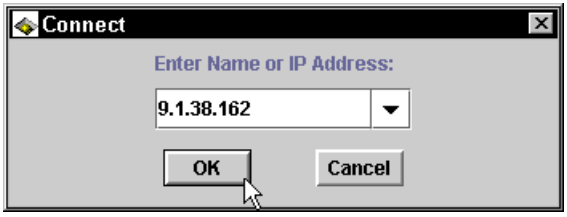


*Figure 7-9 Connect to SDG*

If the dialog box does *not* appear automatically, select **Tools —> Connect** from the pulldown menu. This can also be used to connect to several Gateways or Routers from a single client session.

In Figure 7-10, we show the initial view once a connection to a SAN Data Gateway is established and the data has been downloaded.

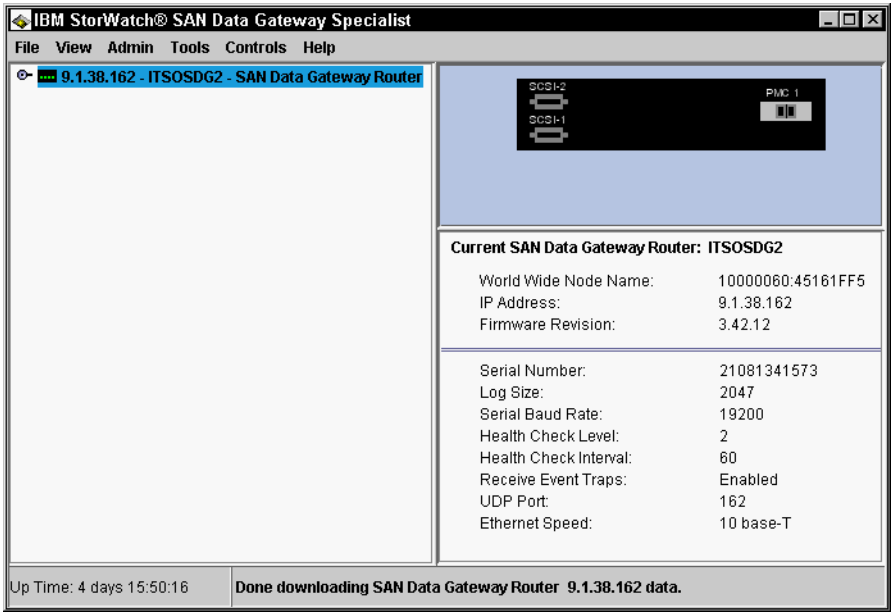


*Figure 7-10 StorWatch SAN Data Gateway Specialist initial view*

The left side of the window shows the IP address and name of the SAN Data Gateway unit we are connected to, and the right side provides a graphical representation of the unit, and product data information. You will also notice that the pull-down menus will have options available that were previously greyed out. You can now connect to another SAN Data Gateway, disconnect from a SAN Data Gateway, enable and access the Zoning and VPS features, restart the Gateway, and also refresh the data to your window by downloading it again.

These options become available when a SAN Data Gateway is highlighted. As you begin to add SAN Data Gateway systems or drill-down into a particular Gateway by selecting and highlighting different channels or ports, different options will become available and other options will become greyed out and unavailable. Be aware of what system, channel, or port is highlighted as you move through the menus.

As we can connect to several SAN Data Gateway systems from one client session, select the particular Gateway you want and it will be highlighted in blue as shown in Figure 7-11.

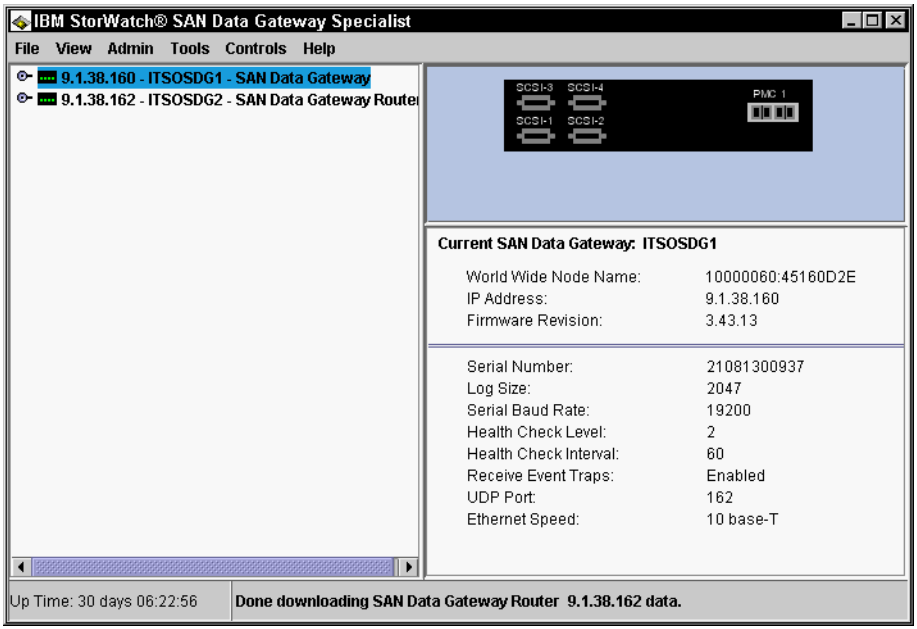


*Figure 7-11 Selecting from multiple SAN Data Gateways*

On the left-hand side of the highlighted Gateway, there is a small toggle key, and by selecting this, it expands the view to show you all SCSI ports and any installed Fibre Channel ports. For example, Figure 7-12 shows a Gateway with two SCSI ports and one Fibre Channel port.

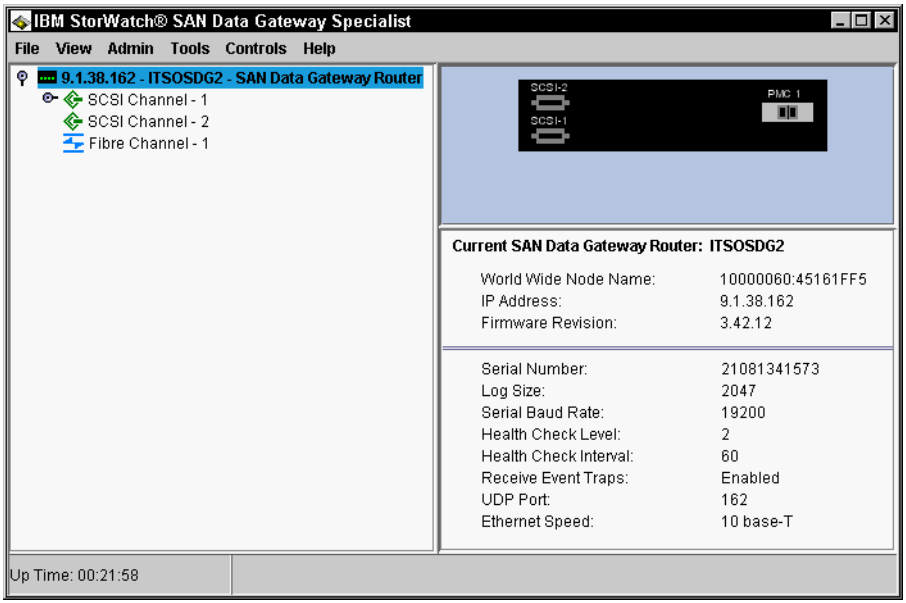


*Figure 7-12 Expanded Gateway view*

## SCSI settings

If you want to view or change the SCSI settings, first select the desired SCSI Channel on the left side of the client display. Then select **Controls —> SCSI Channel** from the menu. Alternatively, you can right-click the selected SCSI Channel and select **SCSI Channel** from the context menu as shown in Figure 7-13.

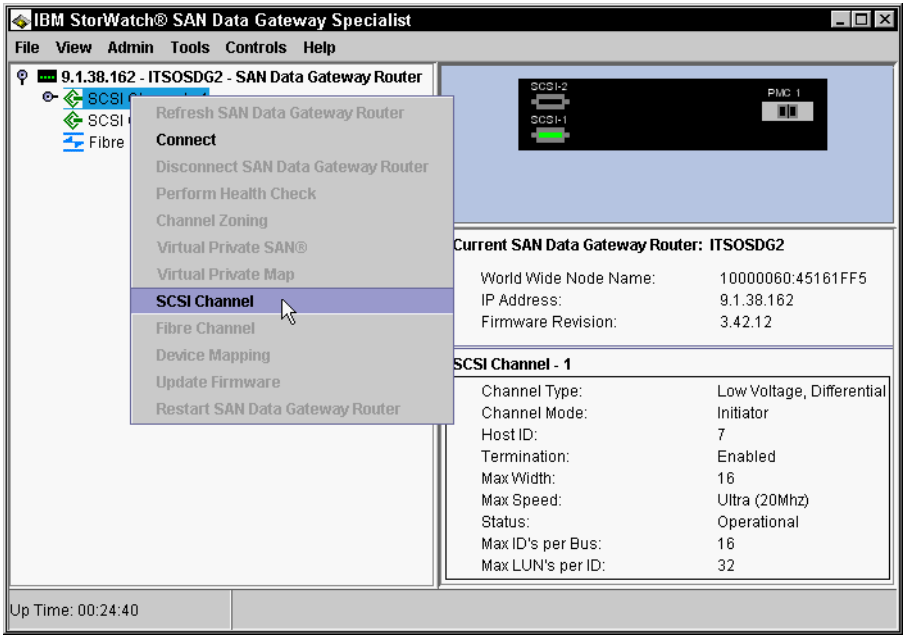


*Figure 7-13 Select the SCSI option*

As we select and highlight each SCSI Channel, you will notice that the information window on the right side will provide data that is unique to that SCSI channel, and the port is highlighted in the graphical representation.

If you change the attached SCSI devices, the SAN Data Gateway is not automatically aware of the new configuration until a SCSI inquiry is performed. The SCSI inquiry occurs when it rescans the SCSI buses. If you want to rescan the SCSI bus select the button **Re-scan SCSI Bus** as shown in Figure 7-14.

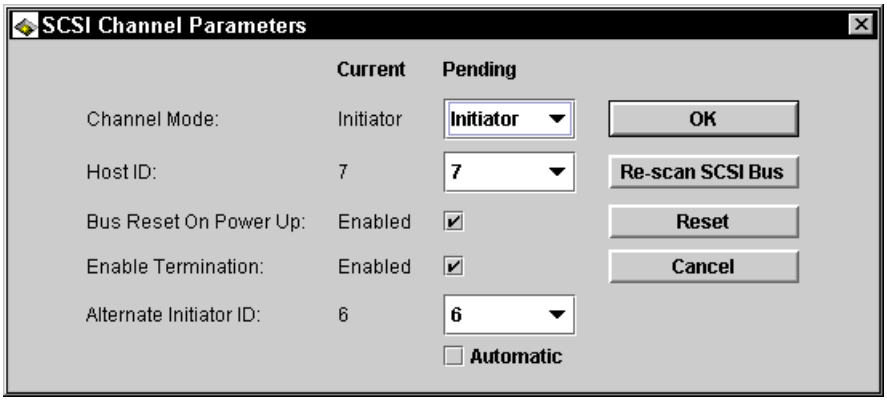


*Figure 7-14 SCSI channel parameters*

SCSI channel - 1 has a toggle key on the left-hand side to depict that there are devices attached. By clicking the key we will expand the tree and view the different SCSI devices attached.

Then, by selecting a specific device from the list on the left, the right panel will display information pertaining to that SCSI device, as shown in Figure 7-15

As you select and highlight the different ports or devices, there are different options available from the pulldown menus. If a SCSI channel is highlighted, select **Controls** from the pulldowns. You will notice that all options are grayed out except for **SCSI Channel**. Once selected, a dialog box will appear as shown in Figure 7-14, and display the settings for the SCSI channel.

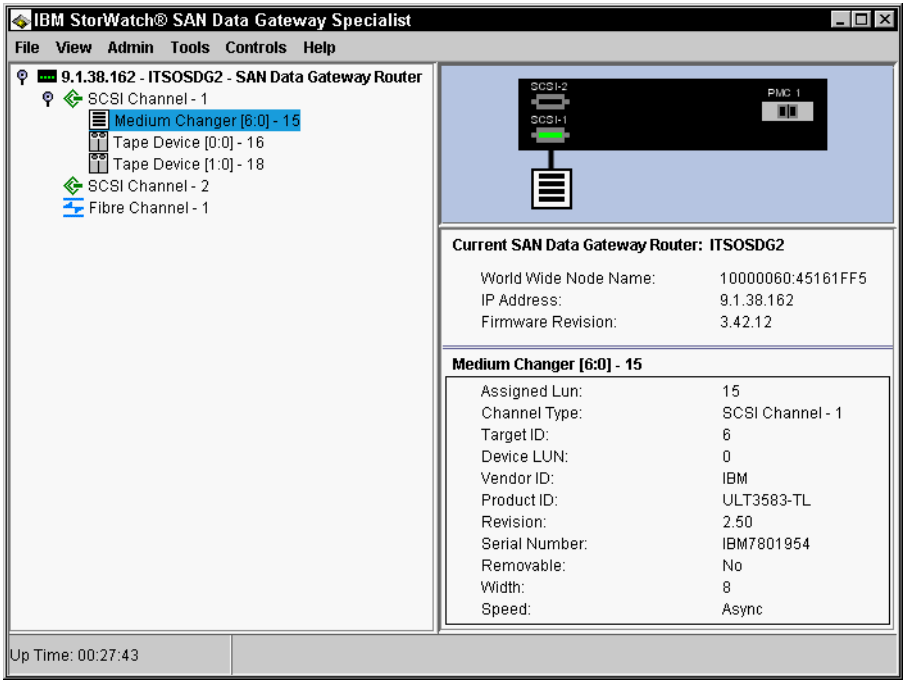


*Figure 7-15 Attached SCSI device data*

## FC settings

Similarly, you can perform the same drill-down function on the Fibre Channel host as we did with the SCSI channel and SCSI devices. Select one of the Fibre Channel port connections, and its data will be shown on the right-hand view panel.

If you want to view or change the FC options, with the desired FC Channel on the left side selected, we then select **Controls —> Fibre Channel**, as shown in Figure 7-16.

A detailed description of the SCSI and Fibre Channel settings can be found in the *IBM Storage Area Network Data Gateway Installation and User's Guide*, SC26-7304.

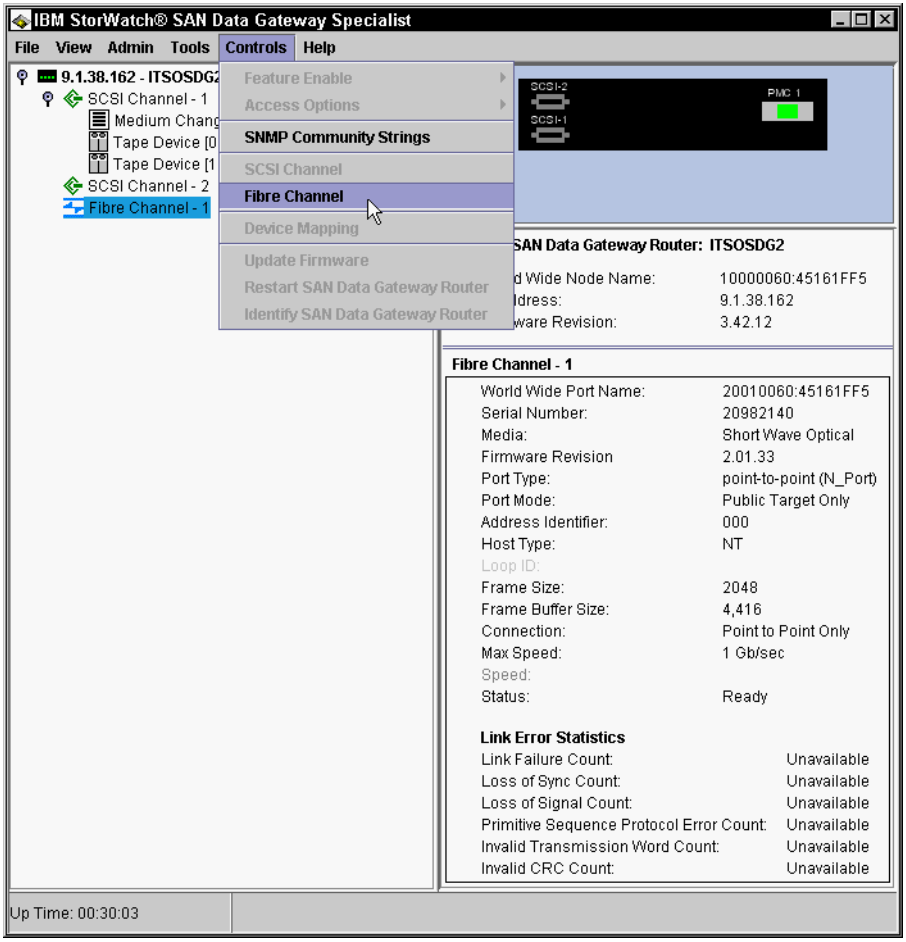


*Figure 7-16 Select the Fibre Channel options*

By selecting this, a dialog box will display the parameters that can be changed for the Fibre Channel port selected as shown in Figure 7-17.

You can change the host type to match your configuration. The connection options should be *point-to-point* or *point-to-point preferred* if the SDG is connected to a switch. If the SDG is directly connected to a server, both server HBA and SDG need to have the same connection options.

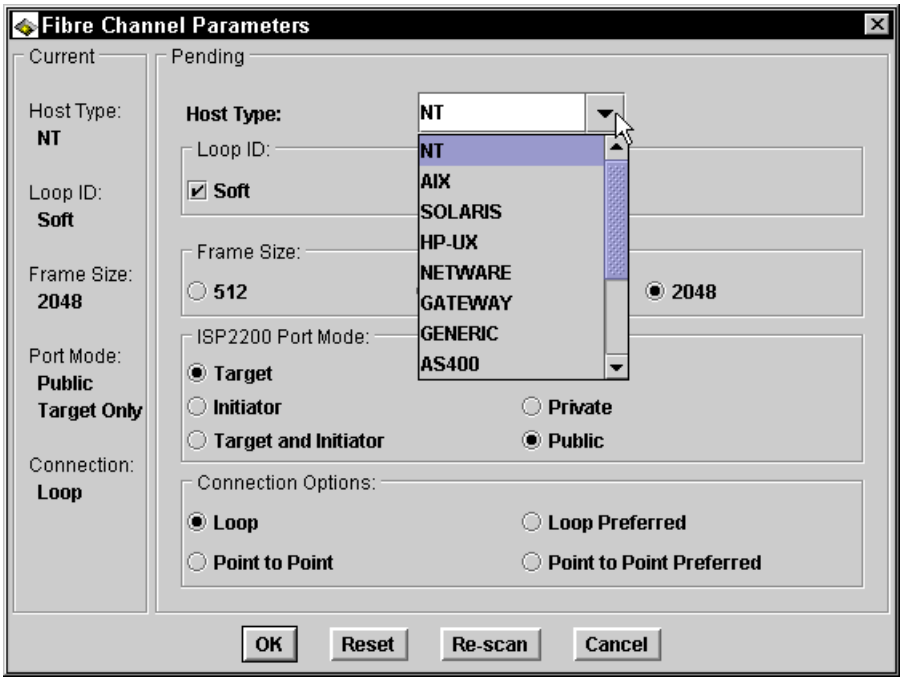


*Figure 7-17 Fibre Channel parameters*

By selecting the toggle key to the left of the Fibre Channel, we can expand the tree for that Fibre Channel, and select the host system attached to that port. Figure 7-18 shows the detail of the specific host.

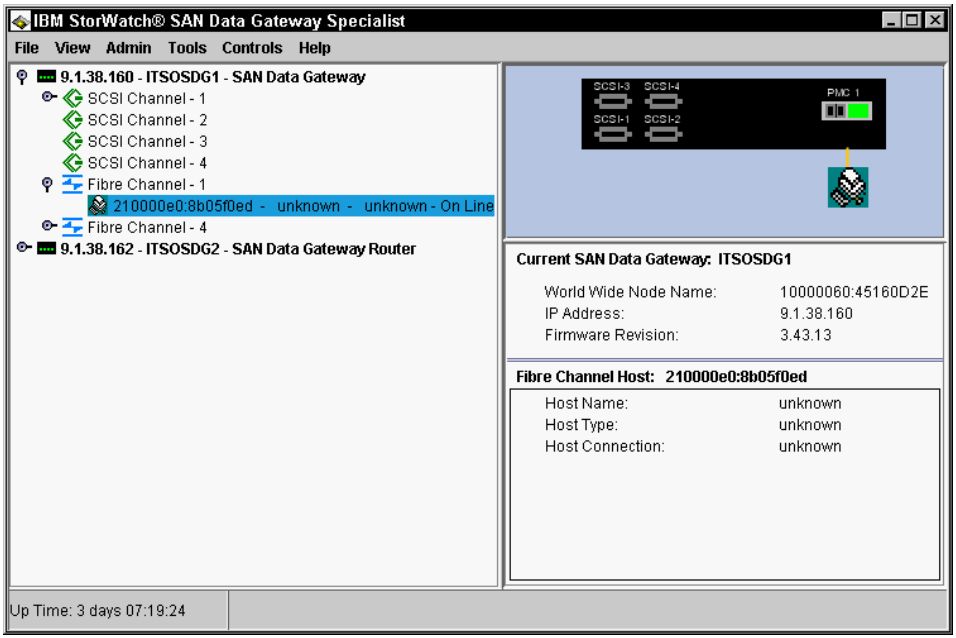


*Figure 7-18 Fibre Channel host data*

## 7.3.4 Upgrading the firmware

New versions of the SAN Data Gateway firmware can be downloaded from

http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/sdgcdr.html

This site has a link to a Downloads section where you will find a list of various downloads, including documentation, firmware, and platform dependent drivers. We selected the *Firmware Download - Main Page* link. You will get a window similar to that shown in Figure 7-19.

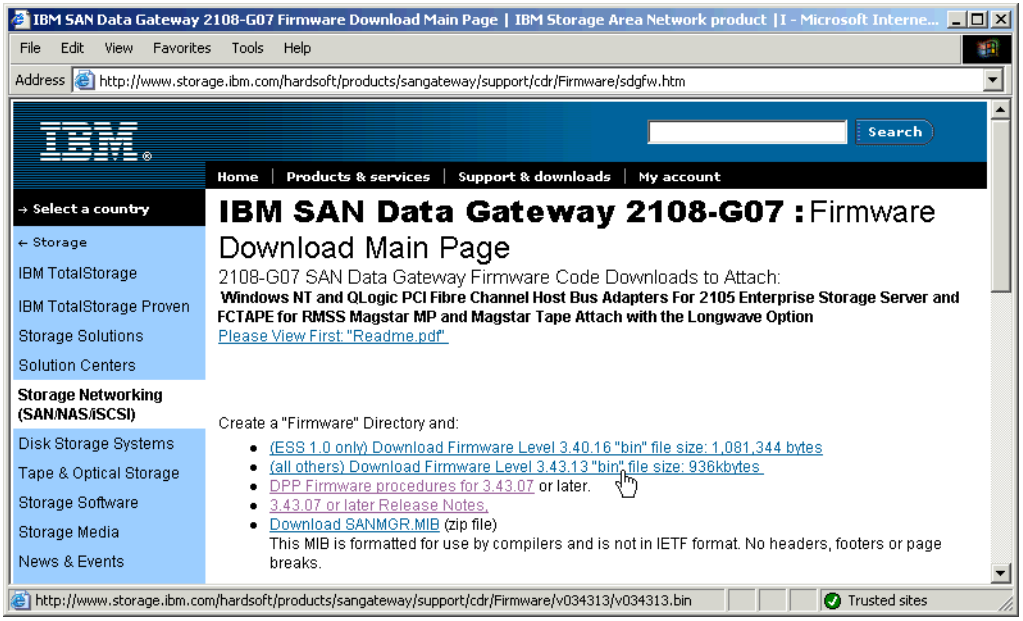


*Figure 7-19 Downloading the SAN Data Gateway firmware*

From this window, we selected the Windows NT version, and we will save this into a directory for downloading into the SAN Data Gateway at a later stage.

From the main window of the SAN Data Gateway Specialist, you can check the current Firmware Revision level as shown in Figure 7-20.

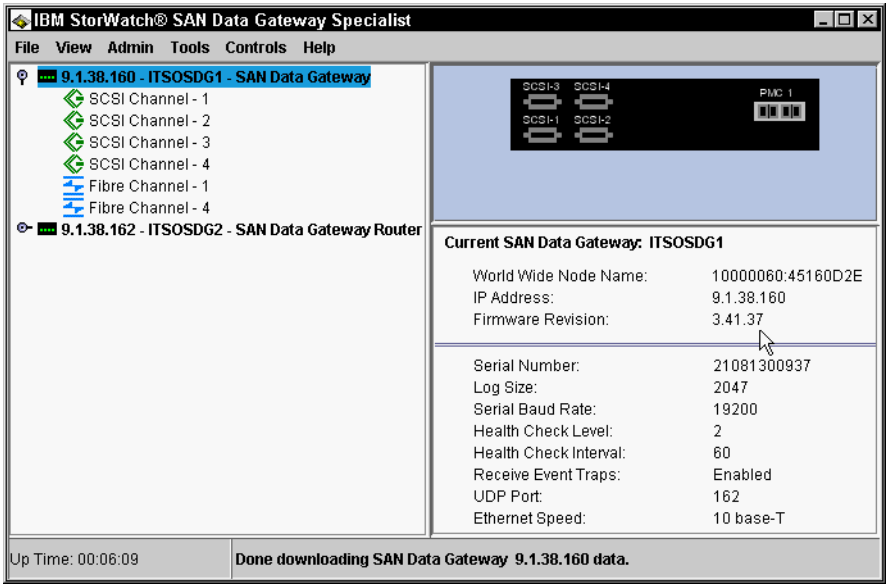


*Figure 7-20 SAN Data Gateway Firmware Revision Level*

New versions of the SAN Data Gateway can be downloaded, as shown in Figure 7-21, by selecting the **Controls** option and then selecting the **Update Firmware** option.

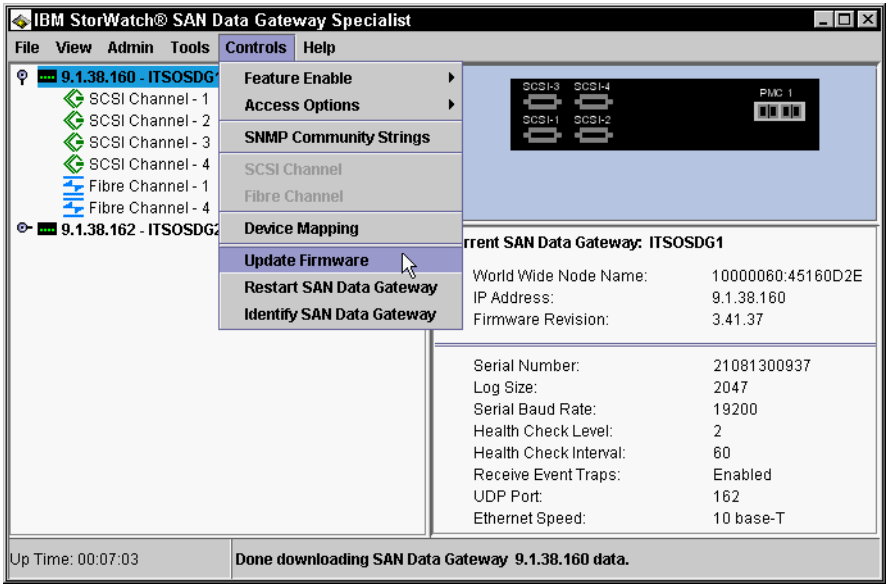


*Figure 7-21 Updating the SAN Data Gateway firmware*

The SAN Data Gateway Specialist will then prompt for the location of the new firmware as shown in Figure 7-22. This is the file that was downloaded from the SAN Data Gateway Web site previously.

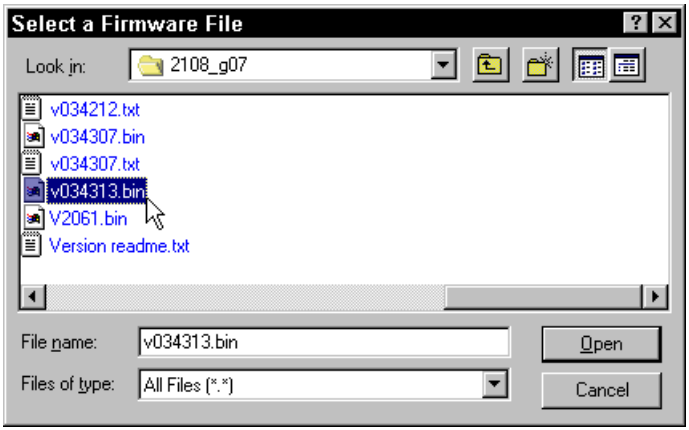


*Figure 7-22 Specifying location of the firmware*

Downloading the firmware into the SAN Data Gateway is a disruptive process, so the Specialist displays a warning message as shown in Figure 7-23.


*Figure 7-23 Warning message prior to downloading the firmware*

Once all I/O activity is stopped, you can click **Yes** to continue the process.

Once the firmware process has been started, the SAN Data Gateway Specialist displays a message in the bottom status line as shown in Figure 7-24.

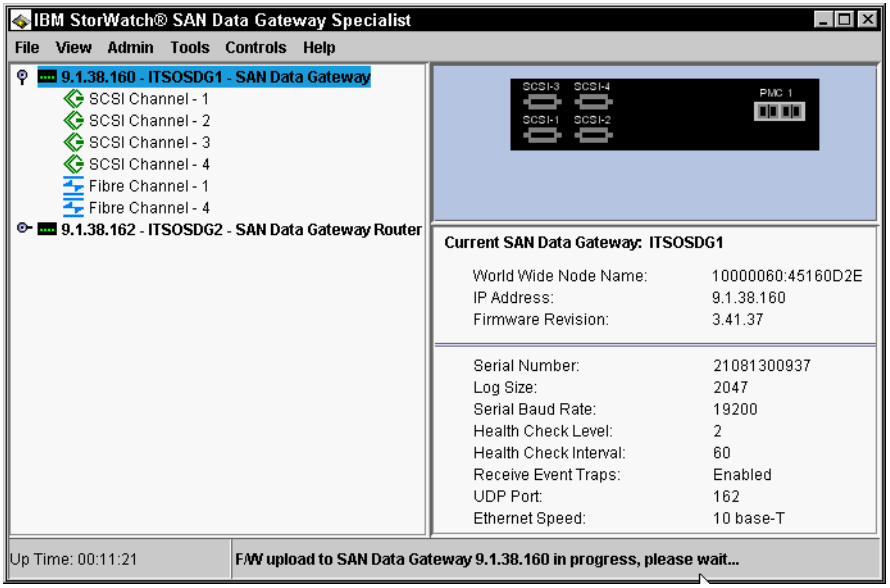


*Figure 7-24 Download in progress*

This indicates that the firmware is being downloaded to the SAN Data Gateway.

The Specialist will then give the option to restart the SAN Data Gateway as shown in Figure 7-25.

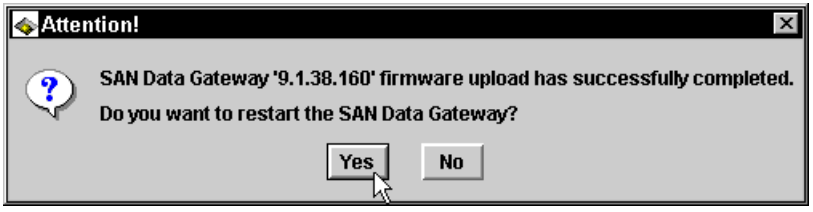


*Figure 7-25 Message prior to restarting the SAN Data Gateway*

As restarting the SAN Data Gateway is a disruptive process, the Specialist issues a warning as shown below in Figure 7-26.

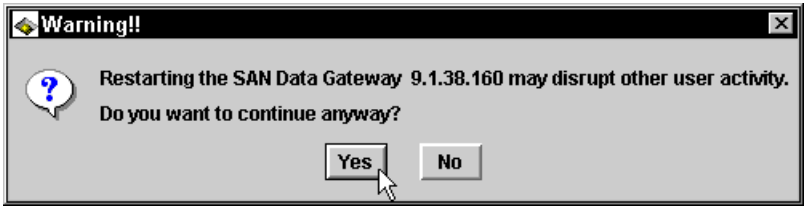


*Figure 7-26 Warning message prior to restarting the SAN Data Gateway*

The Specialist will now return to the main window, and a message is displayed at the bottom of the window as shown in Figure 7-27 indicating that the Gateway is in the process of restarting.

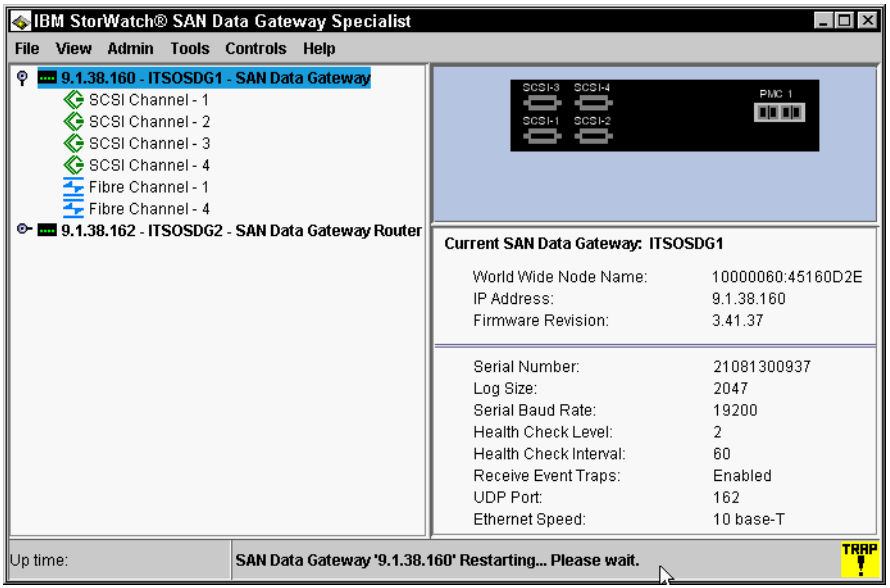


*Figure 7-27 SAN Data Gateway now restarting*

Once the restart process is completed, the Specialist displays the prompt shown in Figure 7-28 before refreshing the display.

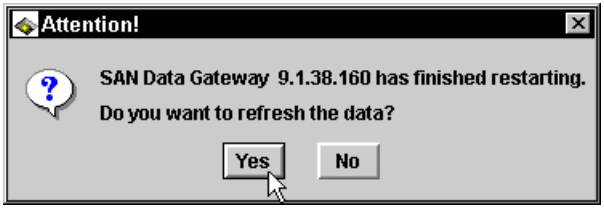


*Figure 7-28 Restart completed message*

Once the Specialist has finished refreshing, we can see that the firmware revision level has been updated successfully, as shown in Figure 7-29.

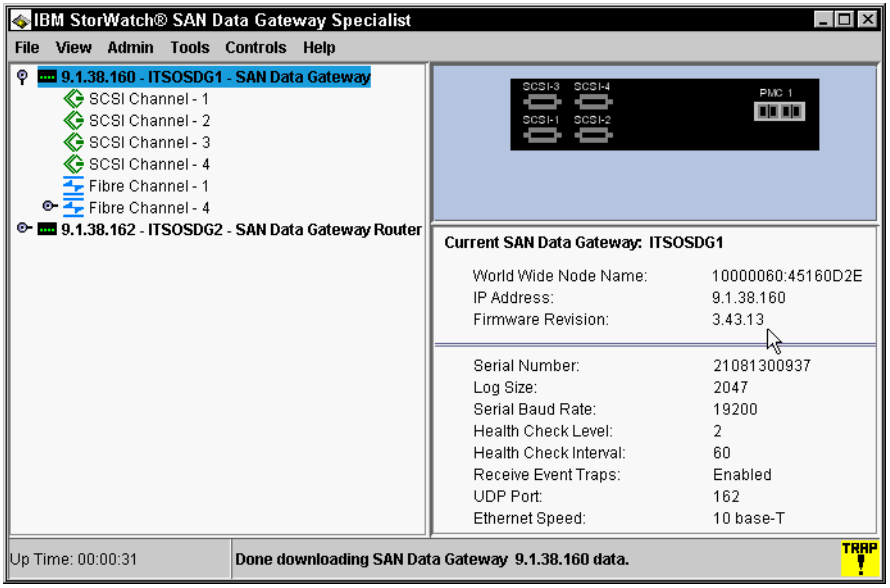


*Figure 7-29 New firmware revision level*

## 7.3.5 LUN mapping

In SCSI terminology, a tape drive or a disk drive is attached to a “bus” and has a unique address on that bus. There are three parts of the address in a conventional SCSI ID:

- Bus (or channel)
- Target ID
- LUN

A simple case of two tape drives attached to a single bus is shown in Figure 7-30.

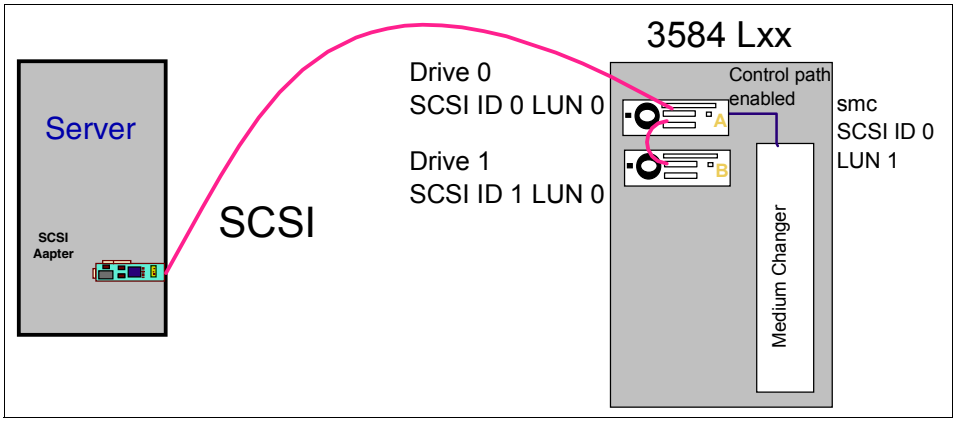


*Figure 7-30 Basic SCSI connection to a system*

This configuration has a device map (from the host point of view) that looks as shown in Table 7-1. Some tape library devices, like the 3584, support LUNs for the devices; the tape drive is addressed as LUN 0, the SCSI Medium Changer (**smc**) is available as LUN 1.

*Table 7-1 Target ID and device mapping — native SCSI*

| Target ID | LUN | Device |
|---|---|---|
| 0 | 0 | Drive 0 |
| 0 | 1 | smc |
| 1 | 0 | Drive 1 |

If a product such as a SAN Data Gateway is placed between the system and target devices, the addressing has another layer. This is because the targets (the tape drives) are not directly attached to the host but are connected to a SCSI adapter installed in the gateway instead. Figure 7-31 shows the device mapping with the additional layer due to the SAN Data Gateway.

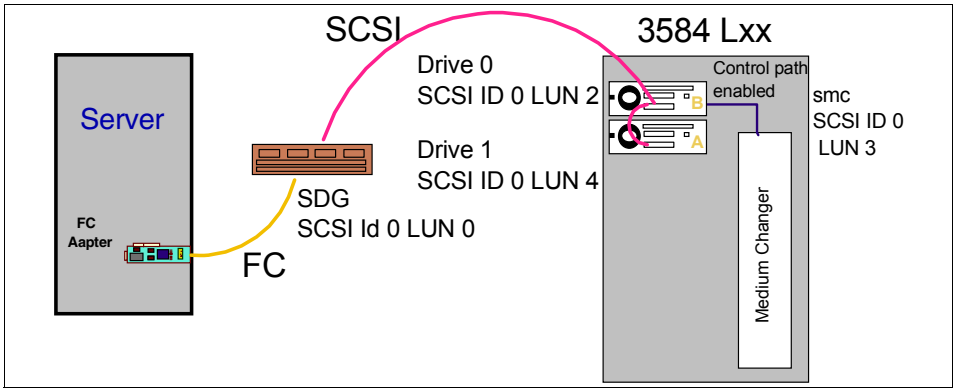


*Figure 7-31 SAN Data Gateway attached through Fibre Channel — host view*

The system now only has a single target ID (target 0 in this case) directly attached to the SAN Data Gateway. The gateway forwards commands to and from its targets, the tape drives. However, there is a need to map the devices (tapes) so the host system can use them. This is achieved by using another layer of mapping, LUNs. The device map might now look like Table 7-2 from the host perspective.

*Table 7-2 Device map from host perspective — with SAN Data Gateway*

| Target ID | LUN | Device |
|---|---|---|
| 0 | 0 | SAN Data Gateway |
| 0 | 2 | Drive 0 |
| 0 | 3 | smc |
| 0 | 4 | Drive 1 |

Note that LUN 0 points to the SAN Data Gateway. This allows you to send commands to control the gateway. This is referred to as the Command and Control Interface.

Tape drives are always assigned an even LUN number. If a control path is enabled (3584) for this drive, its LUN is one higher than the drives. The odd-assigned LUN number that follows the tape-drive even number is reserved for the **smc**. This algorithm has been chosen for the best compatibility with existing applications and operating systems. The **smc** in a 3583 has its own SCSI ID, but will still be assigned an odd LUN number.

If the resulting map is not suitable for your environment, then you can edit the mapping done by the SDG with StorWatch Specialist. First select your desired SDG, then select **Controls —> Device Mapping** as shown in Figure 7-32.

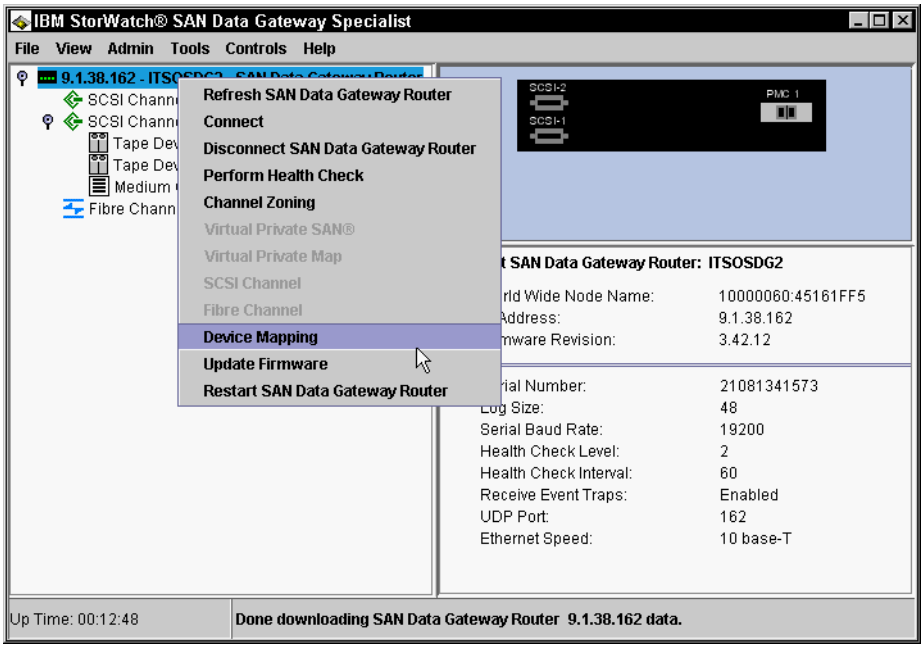


*Figure 7-32 Select Device Mapping*

Figure 7-33 shows the current configuration if you use instead the command line interface (telnet or serial connection) with the command **`fcShowDevs`**. Devices that have already been entered into the persistent device map or devices which were automatically assigned by the SDG are shown in black type on the left. The assigned LUN for each device is shown in the left-most column. On the left side you see devices which are currently not assigned. You can delete a drive from the right side by dragging it down to the Recycle Bin icon.

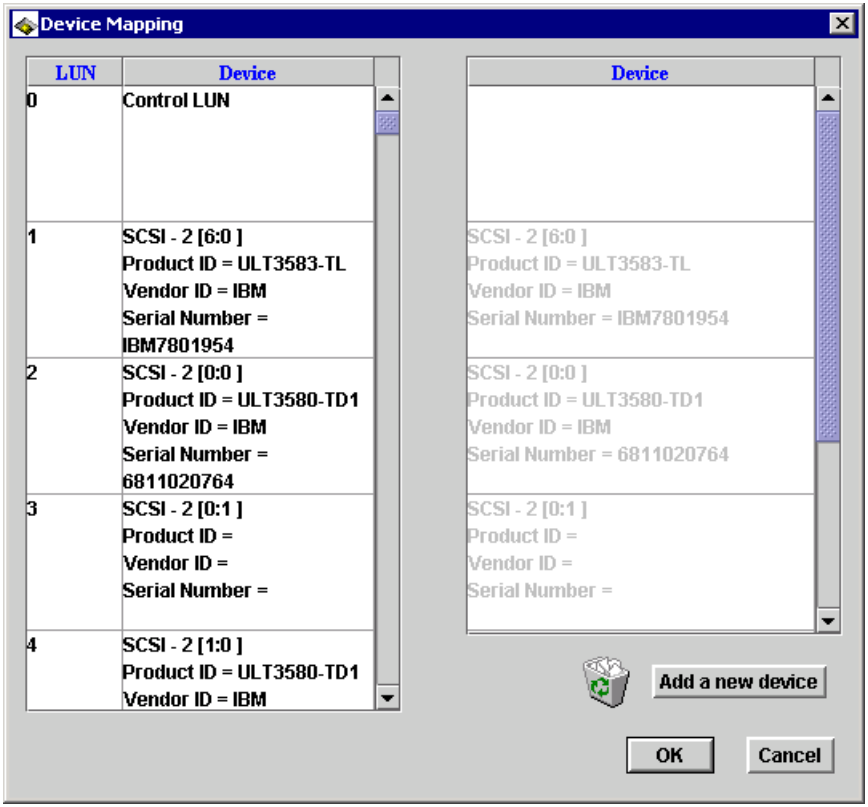


*Figure 7-33 Device Mapping*

You can easily change this configuration by drag-and-drop. Select the device which you want to move to a different LUN. Hold your left mouse button and move the device to an empty LUN. After you have made any changes you need to reboot the SDG as shown in Figure 7-34.

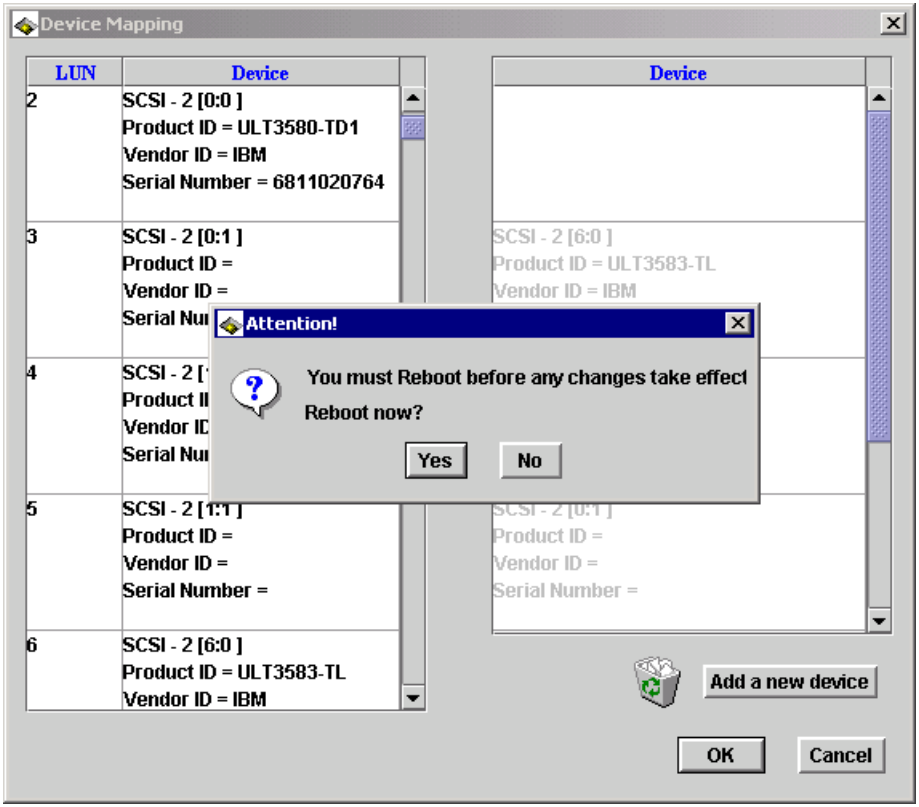


*Figure 7-34 Device Mapping required a reboot*

## 7.3.6 Access control by channel zoning

Click **Controls —> Access Options —> Channel Zoning** (as shown in Figure 7-35) to configure zones to restrict access between SAN connections and SCSI channels. The default settings allow all SAN connections to access all SCSI channels.

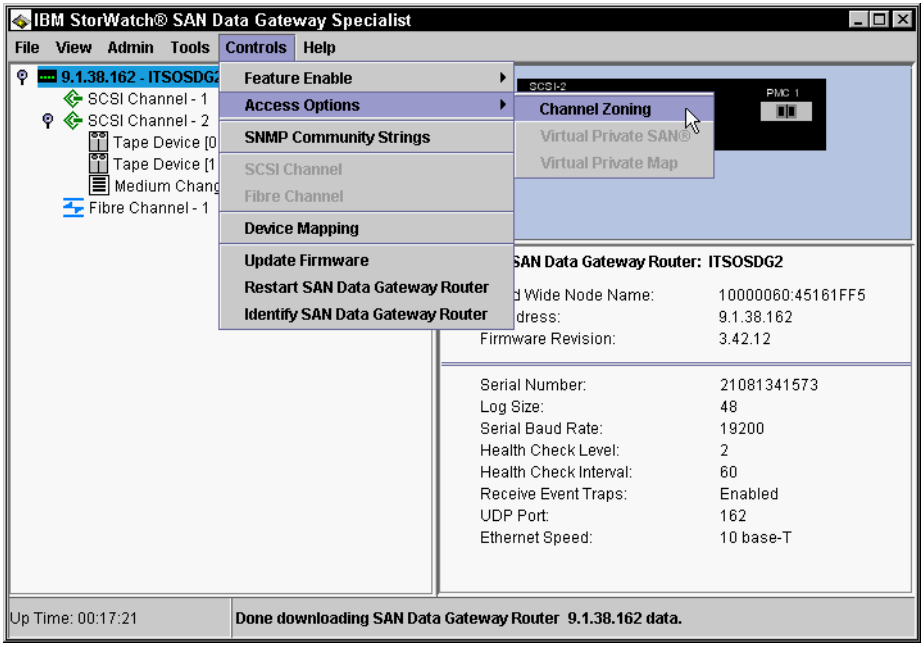


*Figure 7-35 Select Channel Zoning*

When you select this menu option, a pop-up window displays the current channel zoning settings. Figure 7-36 shows the settings for a gateway that has two SAN connections and four SCSI channels. Currently only SCSI Channel 3 and SCSI Channel 4 are assigned to FC 1 and FC 2. Clear the check marks or put additional check marks by clicking in the boxes to create restricted access zones for the desired SAN connections and SCSI channels. All combinations are possible.

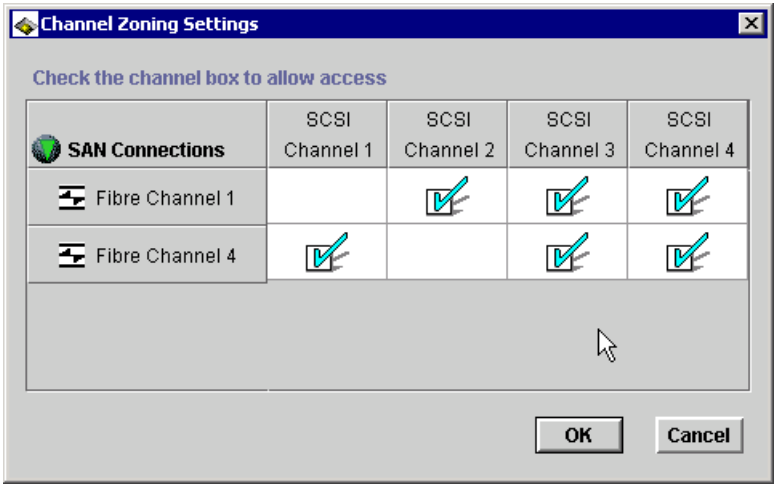


*Figure 7-36 Channel Zoning settings*

## 7.3.7 Access control by Virtual Private SAN (VPS)

Channel zoning provides access control between ports. While channel zoning provides control of paths between host adapters and SCSI storage ports, it does not limit access to specific devices (LUNs) within the storage system. Virtual Private SAN (VP SAN) provides “LUN masking” to limit access between host adapters and LUNs attached to SAN Data Gateway SCSI ports. The IBM StorWatch SAN Gateway Specialist, an easy to use graphical user interface, provides the tools to define SAN Data Gateway channel zoning, the VP SAN LUN-masking, and control which host systems have access to specific storage devices.

VPS is applicable mainly in disk environments, and we will not discuss it further here, as most storage subsystems are native Fibre Channel capable.

It is also not a capability of the Integrated SAN Data Gateway Module for the 3583.

# A

# CNT FC/9000 T_Port mode

In this appendix we discuss the legacy T_Port mode of the CNT FC/9000 family. We consider the different types of zoning available, as well as the translative loop port (TL_Port) mode.

# Zoning in T_Port mode

The following zoning methods are available in T_Port mode:

- Hardware zoning
- Name Server and Broadcast zoning

Hardware zoning actually isolates ports from other ports within the same director. Consequently, hardware zones are never overlapping. Any communication between hardware zones is blocked. This is widely considered as the highest level of security.

Name Server and Broadcast zoning allow you to overlap ports. That means it is possible to place a single port into multiple zones.

Routing information is stored in a central name server table, which will be accessed by initiators before actually starting their I/O.

## Understanding CNT hard zoning

Hard zoning follows physical boundaries within a single-stage switch chassis, and limits the communication of a port to only other ports in the same hard zone.

Hard zoning, in certain circumstances, is the only way to provide the required additional level of security, but careful consideration should be applied prior to activating any hard zones, as it may be possible to isolate devices.

- By no means is it possible to have communication over the boundaries of hard zones. This is also true if malfunctioning fabric initiators try to get around the name server tables.
- Hard zones take precedence over all other kinds of zoning (for example, over broadcast and name server zoning).
- If no hard zone is enabled at all, then all ports are considered as being part of one large default hard zone.
- Name server zones and broadcast zones can be implemented within hard zones. If so, they further limit the connectivity between members of a hard zone.
- Hard zones can be created spanning multiple directors in one fabric.

### Hard zoning rules

There are a number of rules that must be followed to implement hard zoning successfully:

- You can define a maximum of 16 hard zones in an CNT fabric, independently of how many chassis are used in the fabric.

- When a hard zone is created, it must be in a granularity of four ports.
- There is a fixed segmentation of a director into port groups which each have four ports. These port groups will be used when setting up hard zones.
- A single director port can only be part of just one hard zone. It cannot be a part of two hard zones at the same time.
- An *all-or-nothing* rule applies to hard zoning: Either *all* director ports are members of any hard zones, or *none* of them are members of hard zones.
- Any update to hard zone layouts will cause all members of affected zones to perform a fabric login. Hard zone changes should be restricted to initial setup and at maintenance slots.

## Fixed placement of port groups

As mentioned before, the granularity to set up hard zones is in groups of four ports. So one hard zone is built up of one or more multiples of these groups.

The location of these groups is fixed, and you cannot change this. A particular group consists of four ports: two ports on a FIO blade and two ports with the same location on a neighboring FIO blade as shown in Figure A-1.

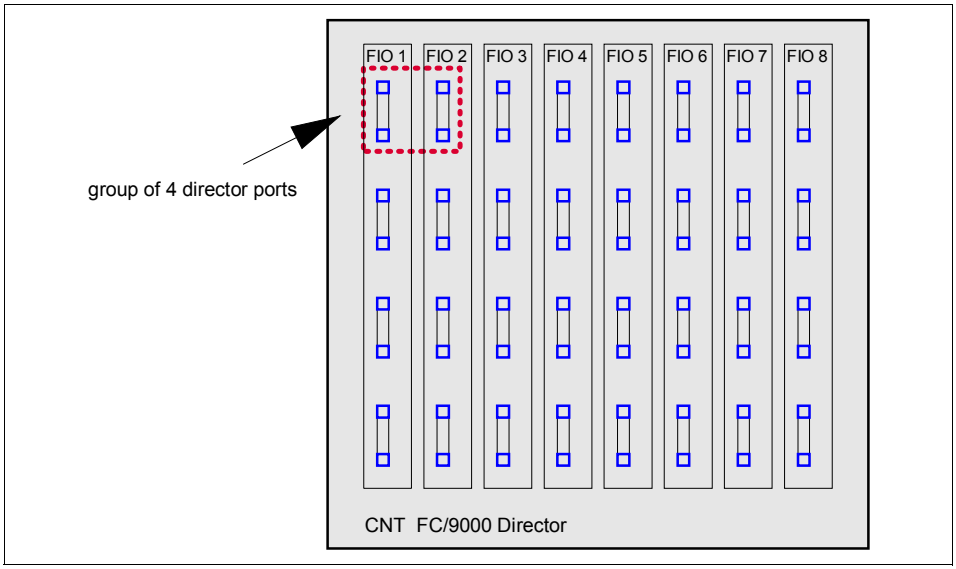


*Figure A-1 CNT hard zoning: fixed location of port groups*

The two ports of one FIO blade being part of such a group are called adjacent ports.

The entire director is always automatically segmented into such groups. Consequently, with a fully equipped 64 port director, 16 port groups are automatically defined.

## Building hard zones using port groups

Now we know that a hard zone must consist of port groups. In our example, we have created two hard zones, as illustrated in Figure A-2.

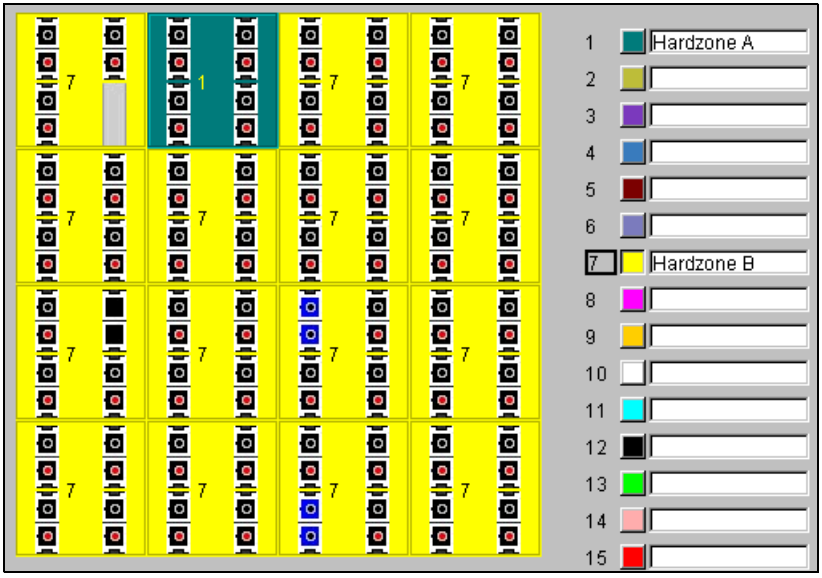


*Figure A-2 CNT hard zoning: basic example with 2 zones*

This picture was taken by using the IN-VSN tool.

Hard Zone A consists of only one port group. Hard Zone B consists of the remaining 15 port groups. Remember the “all-or-nothing” rule: When implementing hard zoning, all ports must be zoned.

In the example above, the four ports of Zone A are allowed to talk to each other. All 60 ports of Zone B are allowed to talk to each other. However, any communication between Zone A and Zone B is blocked.

Remember that hard zoning strictly excludes ports other than those in the same hard zone from communicating together.

When creating hard zones, all port groups in a particular hard zone must be adjoining. In Figure A-3 we show an example of an incorrect implementation.

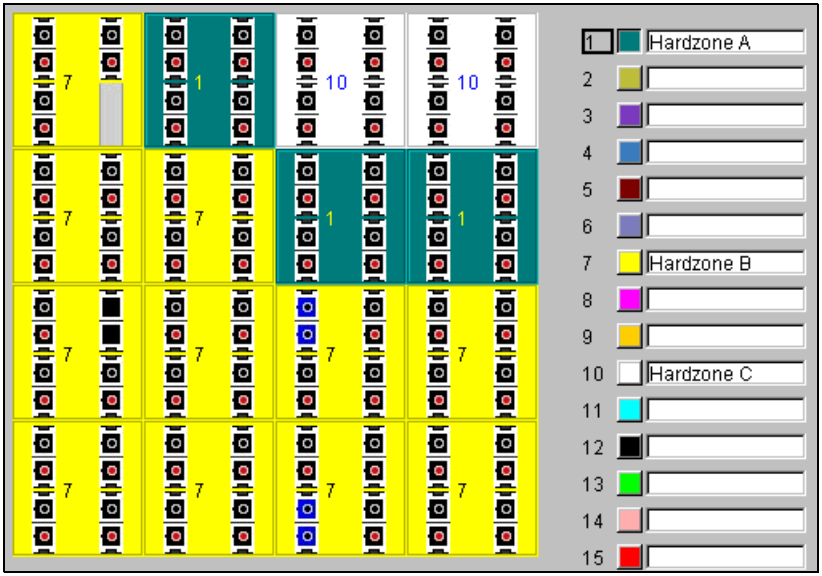


*Figure A-3 Violating the adjoining rule*

The upper port group of Hard Zone A has no adjoining contact to the remaining two port groups of Hard Zone A. This is a violation of the hard zoning rules, and therefore it is not possible to create a hard zone layout such as this.

## Understanding CNT broadcast zoning

Broadcast zones focus the distribution of broadcast messages to only those targets that need to receive them. This results in reduced fabric traffic and eliminates unnecessary message processing. Broadcast zones can overlap and are also assigned by director ports.

The granularity in which broadcast zones are built is one director port.

However, broadcast zoning is used for IP networking. Storage area networking environments do not exploit this kind of zoning.

## Understanding CNT name server zoning

In contrast to hard zoning, there is no actual physical segmentation of fabrics with name server zoning.

Instead, a *name server table* is used to implement that type of zoning.

Name server tables store information about nodes that have logged into the fabric. Name server zones restrict the access of affected ports to this information. All unzoned nodes have full access to the name server table.

This means that zoned ports logging into a fabric will only get name server table information of other ports which are part of the same name server zone. Unzoned ports are considered as being part of an *orphan zone* which is not visible to users, but allows communication for all unzoned nodes.

Because this kind of zoning relies on the proper usage of FCP protocol commands, it is also widely called *soft zoning*.

Name server zoning gives better flexibility then hard zoning:

- A particular director port can be part of multiple name server zones.
- The granularity of ports to build a name server zone is only one port.
- Name server zone members need not be physically adjoining each other.
- The change of name server zones does not enforce a port re-login of affected ports.

However, since name server zoning relies on the correct usage of FCP commands, there is still a risk that malfunctioning nodes would affect other ports even if these other ports are not members of the same zone.

Due to its strength and flexibility, name server zoning is widely used in open systems environment.

### Name server zoning rules

There are several rules that must be adhered to:

- CNT name server zoning refers to physical director ports. A type of name server zoning pointing to WWNs of attached nodes will be available in the future.
- Each name server zone needs to get a unique number and a name.
- As many as 256 zones are possible in a fabric.
- Name server zones cannot cross the boundaries of defined hard zones.
- Director ports which have private nodes attached to it need to be set to TL mode. name server zoning is not effective for TL_Ports.

## Hard zones and name server zones together

Hard zoning can strictly separate port groups, and this can be seen as an effective security feature.

Name server zoning allows us to further define the communication control on a per port granularity.

When name server zones and hard zones are used in conjunction, we need to consider the following principles:

- If you have no hard zone(s) created at all, then all director ports are considered as being part of one big default hard zone.
- However, if at least one hard zone is implemented then this default hard zone is not effective anymore. So, be aware of the fact, that if you add just one hard zone all remaining ports must be added to a hard zone too. For instance you could create a second hard zone containing all the remaining ports.
- Multiple name server zones can exist within one hard zone.
- Name server zones cannot cross the boundaries of hard zones
- Using CNT we do not have to worry about zone sets or active and inactive zones. Creation of zone sets or creation of passive zones is not possible in CNT environment. Either a zone exists and is thereby active, or it does not exist.
- All nodes that are not part of any name server zone have unlimited access to the name server table. This is also true if no name server zone is implemented at all. Consequently, by default, all attached nodes have access to the name server table. Only those nodes that are part of name server zones will have limited access to this information
- The only exception is a TL_Port Config list which explicitly allows access to the specified TL_Port even if name server zones exist that normally would imply otherwise.

So, we can see that we have several different ways to control actual node access with CNT:

- Hard zoning
- Name server zoning
- Translation Entries lists for TL_Ports

The following rules are enforced with CNT directors:

- Both TL-CFG lists as well as name server zones cannot span hard zone boundaries.
- You can use name server zoning to further limit access between nodes

- TL-CFG lists are for TL_Ports only. They overrule the name server principles. So, TL_Port attached nodes can only communicate with ports specified in the TL-CFG lists.

This hierarchy is shown in Figure A-4.

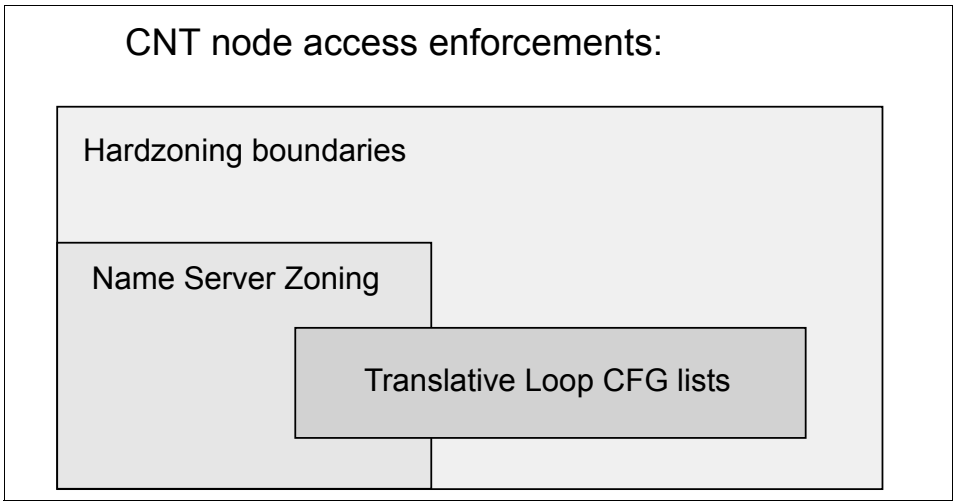


*Figure A-4 CNT access enforcements*

The scenario illustrated in Figure A-5 describes a possible layout.

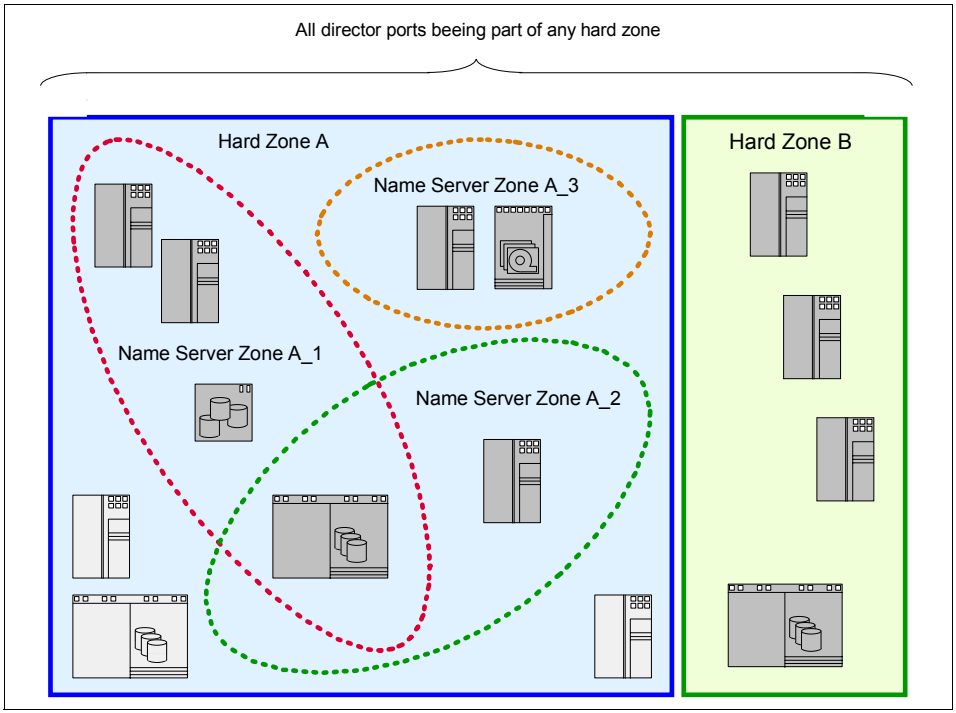


*Figure A-5 CNT fabric scenario: Hard zoning and name server zoning*

In this layout:

- All three defined name server zones are within Hard Zone A. They cannot span over to Hard Zone B.
- Name server zone A_1 and A_2 are overlapping. Both have an ESS node as a member.
- All members of name server zones can only access other members within the same Name server zone.
- All remaining nodes in Hard Zone A which are not part of any name server zone are part of the *orphan zone*. This allows communication between these three nodes.
- There is no name server zone at all in Hard Zone B. All members can have access to any other member in Hard Zone B since they are put into the orphan zone. However, this orphan zone is limited to Hard Zone B.

## Defining hard zoning

In our example shown in Figure A-6, we want to separate eight director ports for usage by the finance department only. In no way should it be possible for anyone else to gain access to their ports. That is why these eight director ports will be put into one dedicated hard zone (the Finance hard zone).

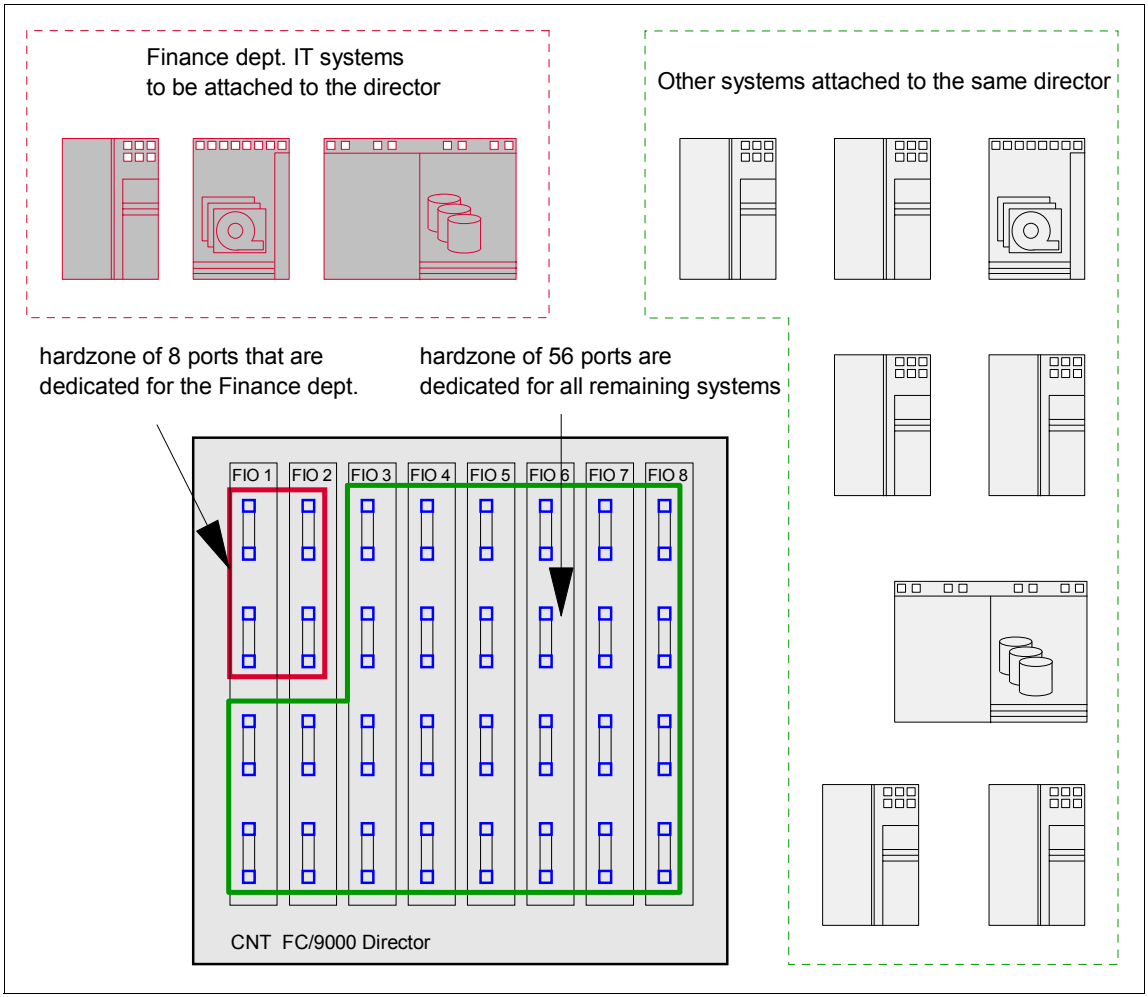


*Figure A-6 CNT hard zoning: Layout scenario*

Considering the hard zoning rules that we have already discussed, we will place all remaining ports into a second zone. This second zone can be used by all other systems. Name server zoning is possible in both hard zones.

To start our hard zoning setup, we first select the particular director in the navigation tree and then click the **Hard Zoning** tab as shown in Figure A-7.

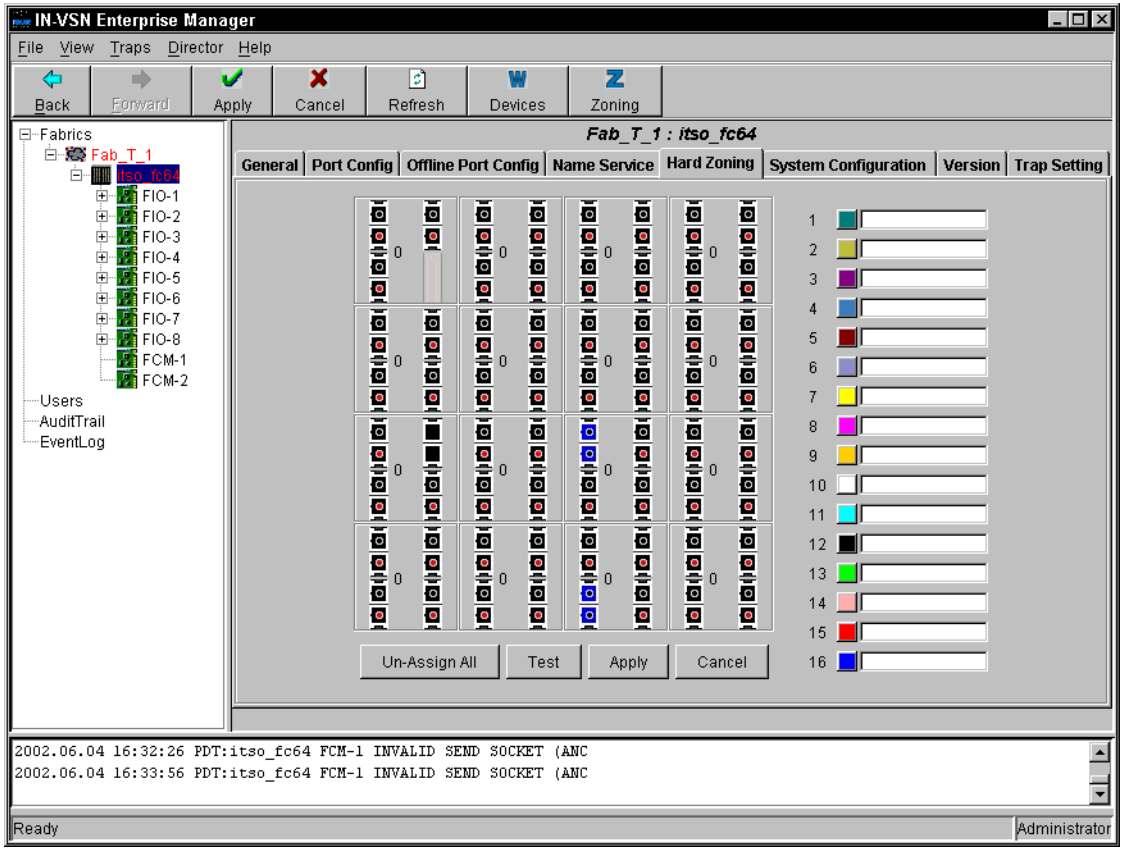


*Figure A-7 IN-VSN: Selecting hard zoning in the director view*

This window is used for all hard zoning configurations related to the selected director.

This include actions such as:

- Renaming a zone
- Adding a zone
- Deleting a zone
- Change the zone layout

To create a specific hard zone, select one of the colored rows on the rightmost side of this window and type in a name for this zone. This is illustrated in Figure A-8.

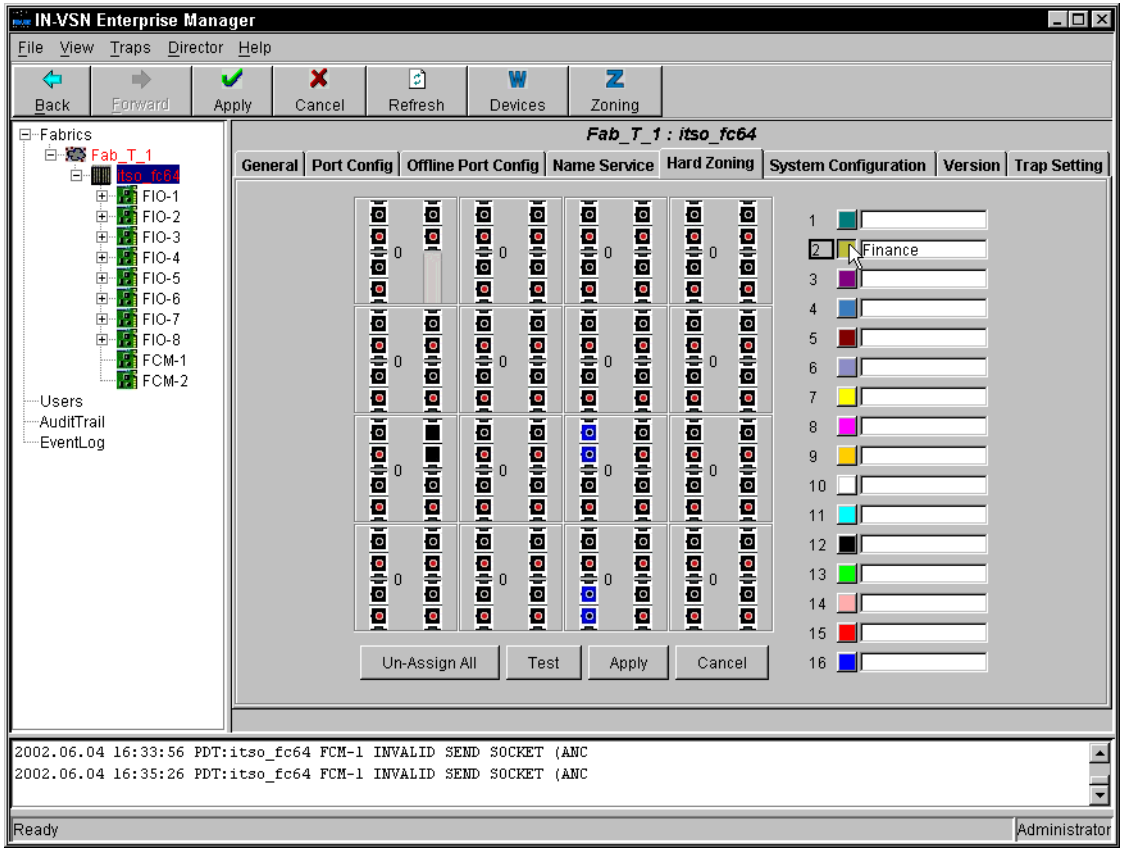


*Figure A-8 IN-VSN: Specifying a name for a hard zone*

We have chosen to name the first zone *Finance*. However, later on, we could change the name again.

The remaining ports will be placed in a second zone, which we will call *Other*.

**Note:** The used color or number for a dedicated zone does not have any effect on the actual behavior or performance. These numbers and colors are solely used to simplify the organization and usage of hard zoning.

Hard zone colors, names, or numbers are not visible in any way by attached nodes.

To include eight ports as members in our first zone, *Finance*, we click the port groups we want to use while the *Finance* zone is still marked in the rightmost table.

We have selected the two upper leftmost port groups to be part of hard zone **Finance** as illustrated in Figure A-9.

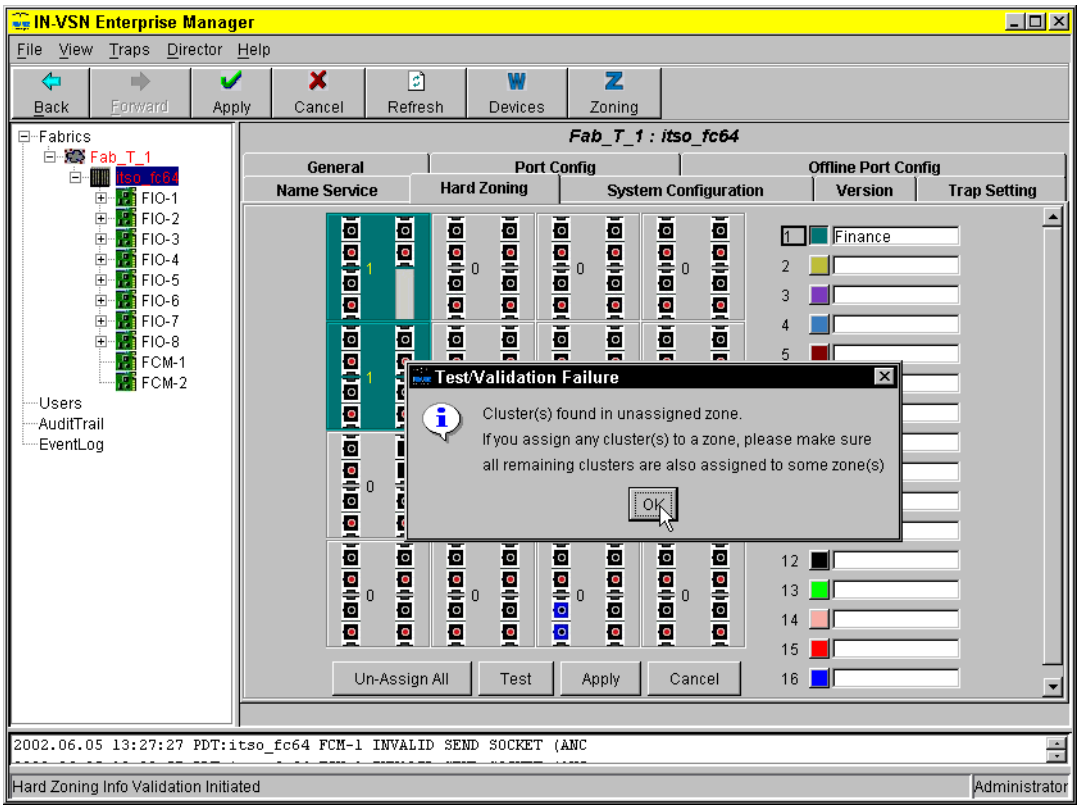


*Figure A-9 IN-VSN: two ports assigned to a hard zone*

To test the integrity of our zone layout, we click the **Test** button. As you can see, we got a **Test Failed** feedback.

This is because we violated a hard zoning rule: Either all ports of a fabric belong to hard zones or none of them are hard zoned.

Consequently, we will place all remaining director ports in a second zone called *Other,* as we have described before.

Again we test this layout by clicking the **Test** button. This time everything looks fine, as illustrated in Figure A-10.

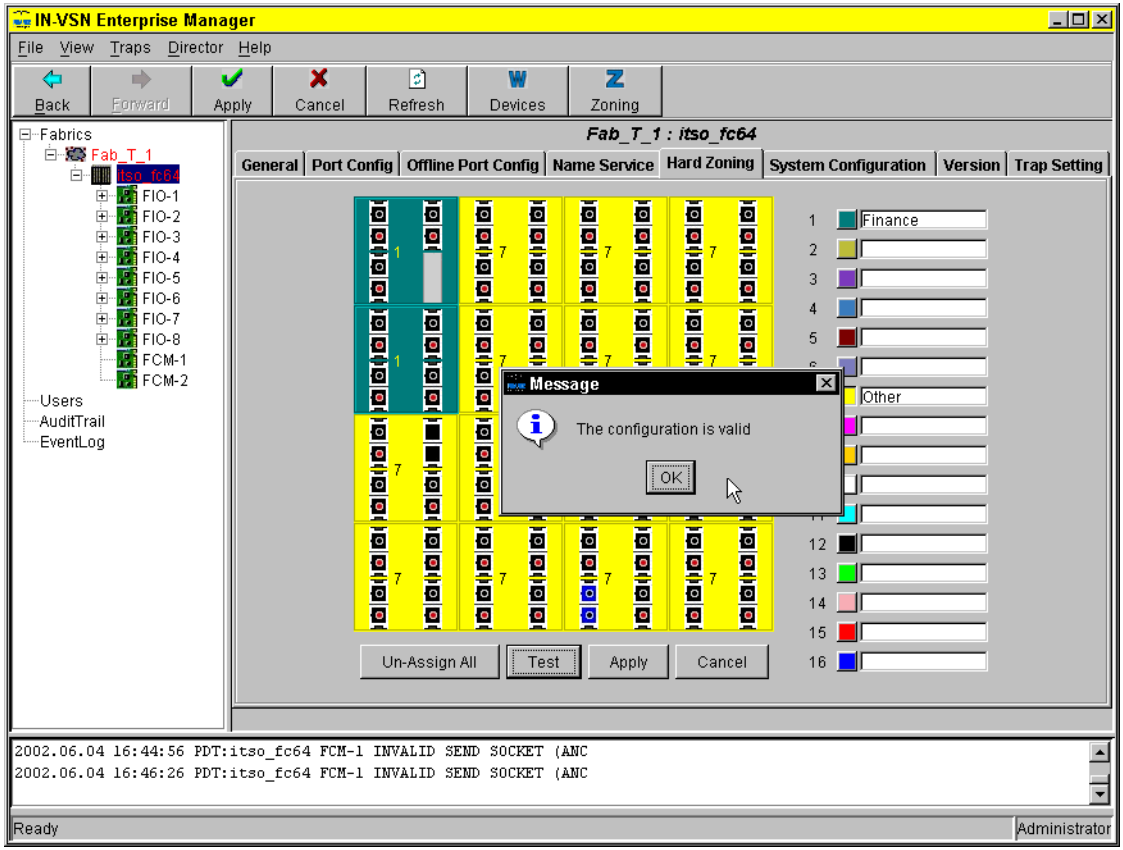


*Figure A-10 IN-VSN: Having two hard zones defined*

We now have two zones defined: *Finance* and *Other*.

The *Finance* zone consists of two port groups giving a total port count of eight. The *Other* zone has the remaining 14 port groups assigned to it having a total port count of 56.

However, this hard zone setup is not yet active since we have not clicked the **Apply** button at this stage.

To demonstrate the adjoining rule, we have added the bottom left port group to the *Finance* hard zone.

Remember that the adjoining rule for hard zoning demands that all port groups of one zone must be adjoining.

As illustrated in Figure A-11, when clicking the **Test** button now we get a violation message, as expected.

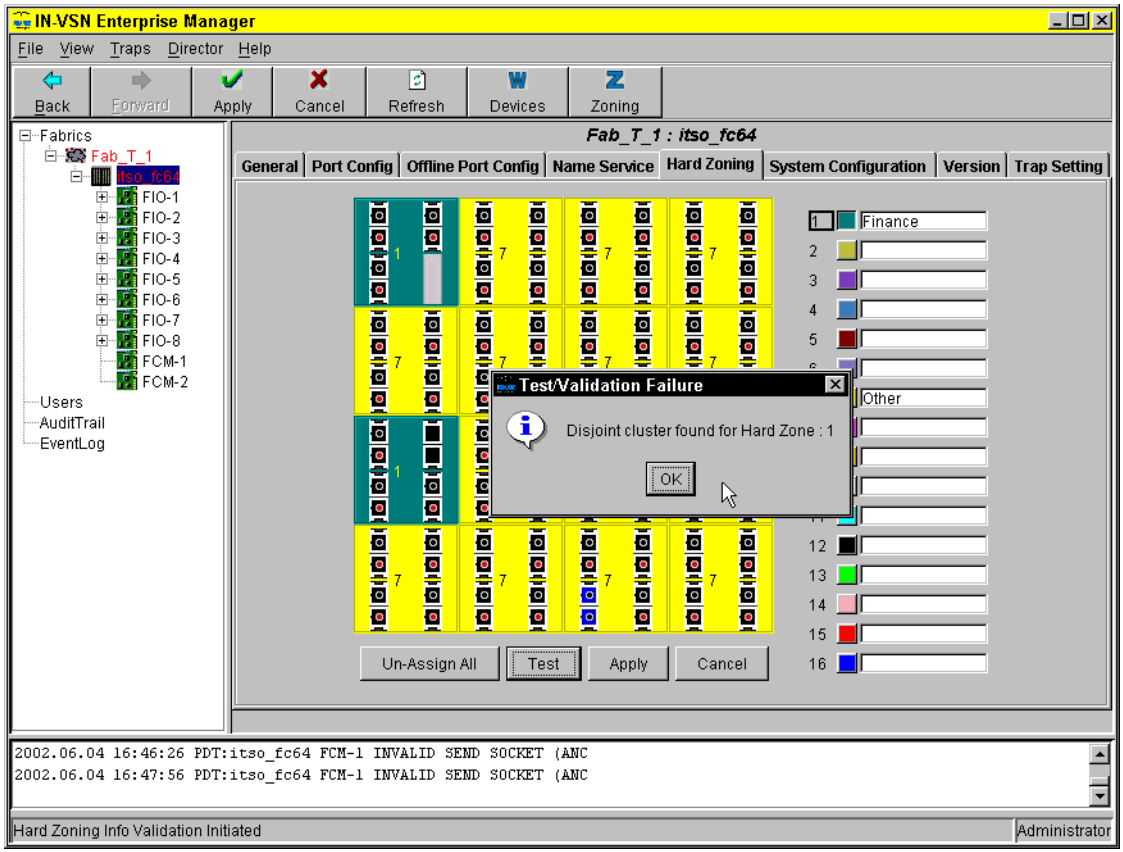


*Figure A-11 Violation of ports*

To successfully finish our hard zone setup as planned, we have put the bottom left port group back into the *Other* zone.

To activate our hard zone layout, we click **Apply** as illustrated in Figure A-12.

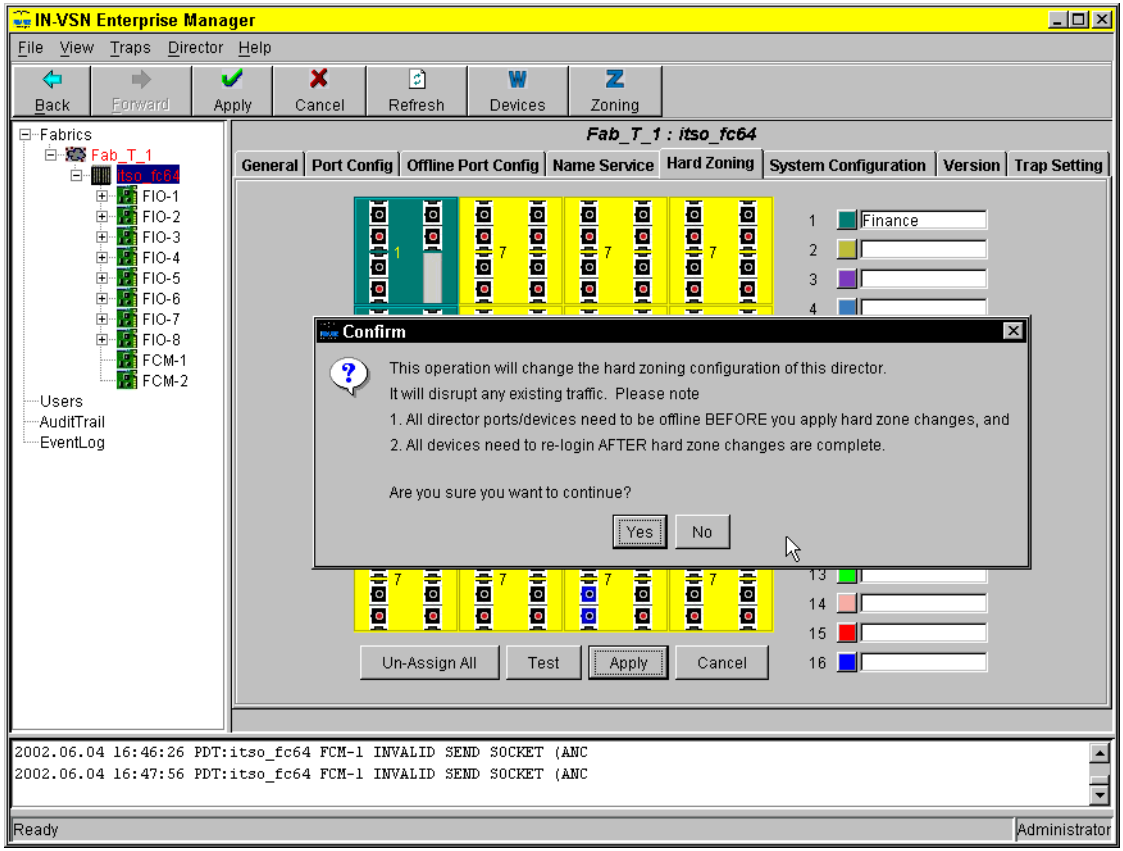


*Figure A-12 IN-VSN: Applying a hard zone setup*

As explained earlier, a director reconfiguration task such as this will affect in-progress I/O. A warning message is issued to indicate this.

To activate our setup, we click **Yes**.

After this hard zone setup is activated, there is no communication possible between the *Finance* zone and the *Other* zone.

This is the most secure way to separate nodes in a fabric environment.

## Defining name server zones

Our goal is to enable all servers to access the ESS.

To achieve a high level of availability, we have chosen to implement a multipathing architecture except for one server (GEODE).

All *primary* FC ports of these servers will be zoned to ESS Bay 1, Port 4, and all *secondary* FC ports of these servers will be zoned to ESS Bay 4, Port 4.

We use **`pvlinks`** as multipathing software. Another possibility is to use the IBM Subsystem Device Driver (SDD).

In Figure A-13 we show an overview of what we will achieve.

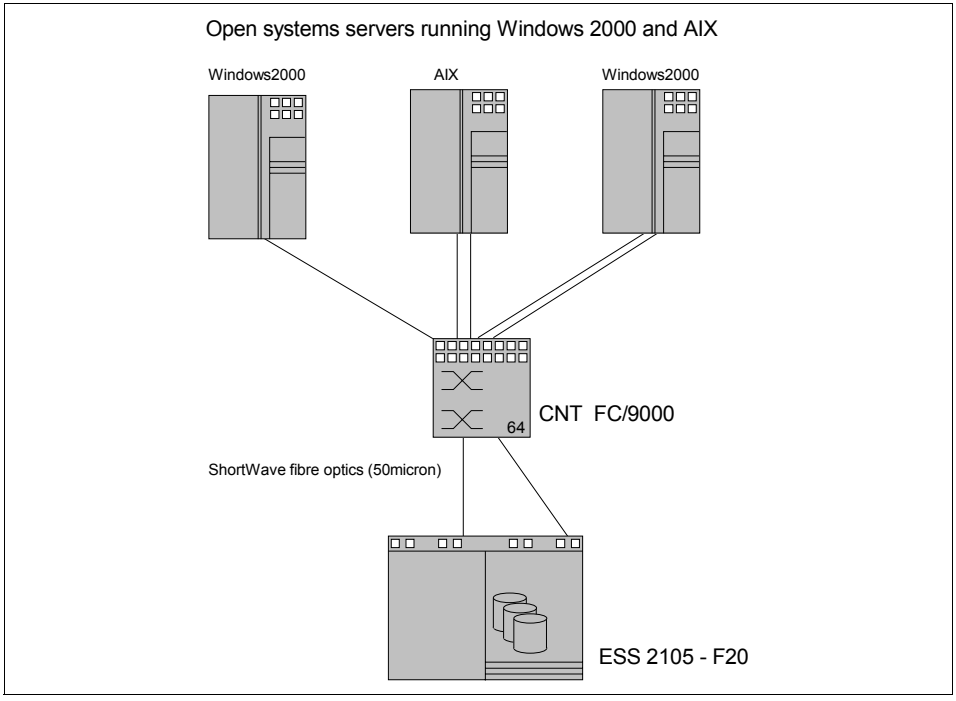


*Figure A-13 Server setup for soft zoning*

Be aware that such a zone-configuration is only advisable in specific environments.

Be sure to have only host adapters of the same type in a zone. This is true in our case. Therefore, we just implement two zones instead of dedicated zones for each server.

By implementing two zones, we limit the access of a particular host FC port to only one ESS FC port. We do this because we like to have only two paths to a ESS logical volume.

These two paths are created by giving the two host FC ports of one server access to the particular ESS logical volume using ESS LUN masking.

In Figure A-14 we show how these ports are physically attached to the CNT director.

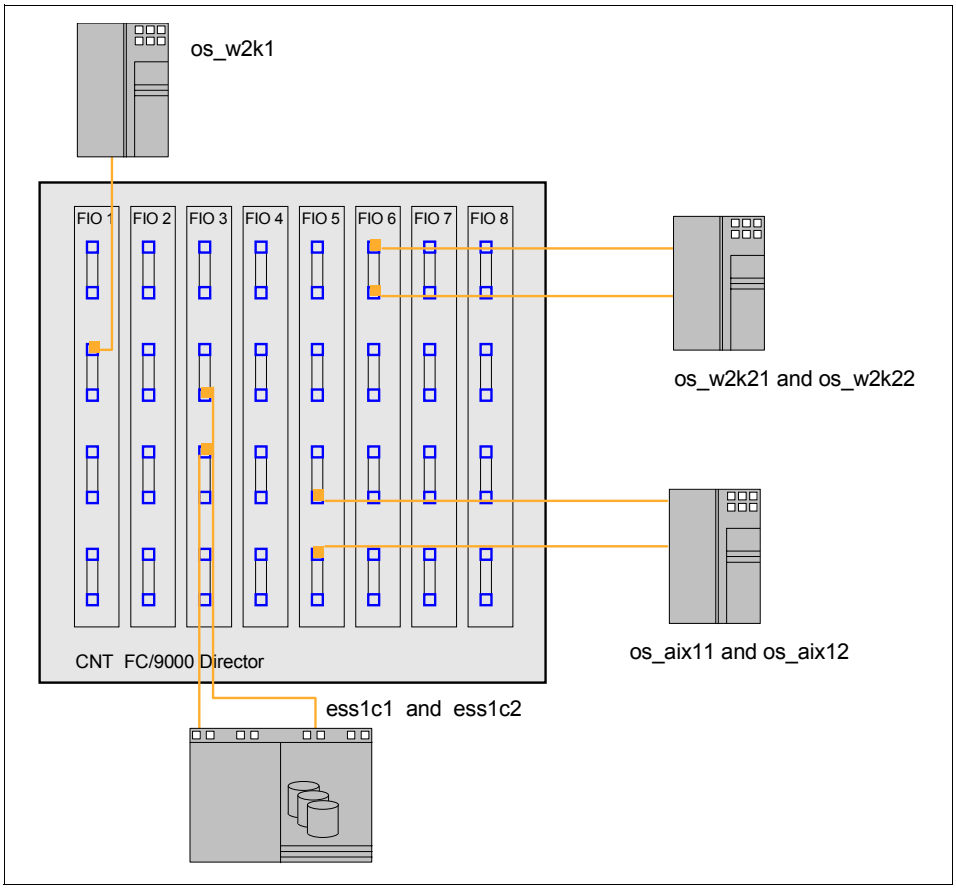


*Figure A-14 Physical cable connection for soft zoning setup*

We shall now explain how our zoning will look logically. As illustrated in Figure A-15, we plan to have two zones.

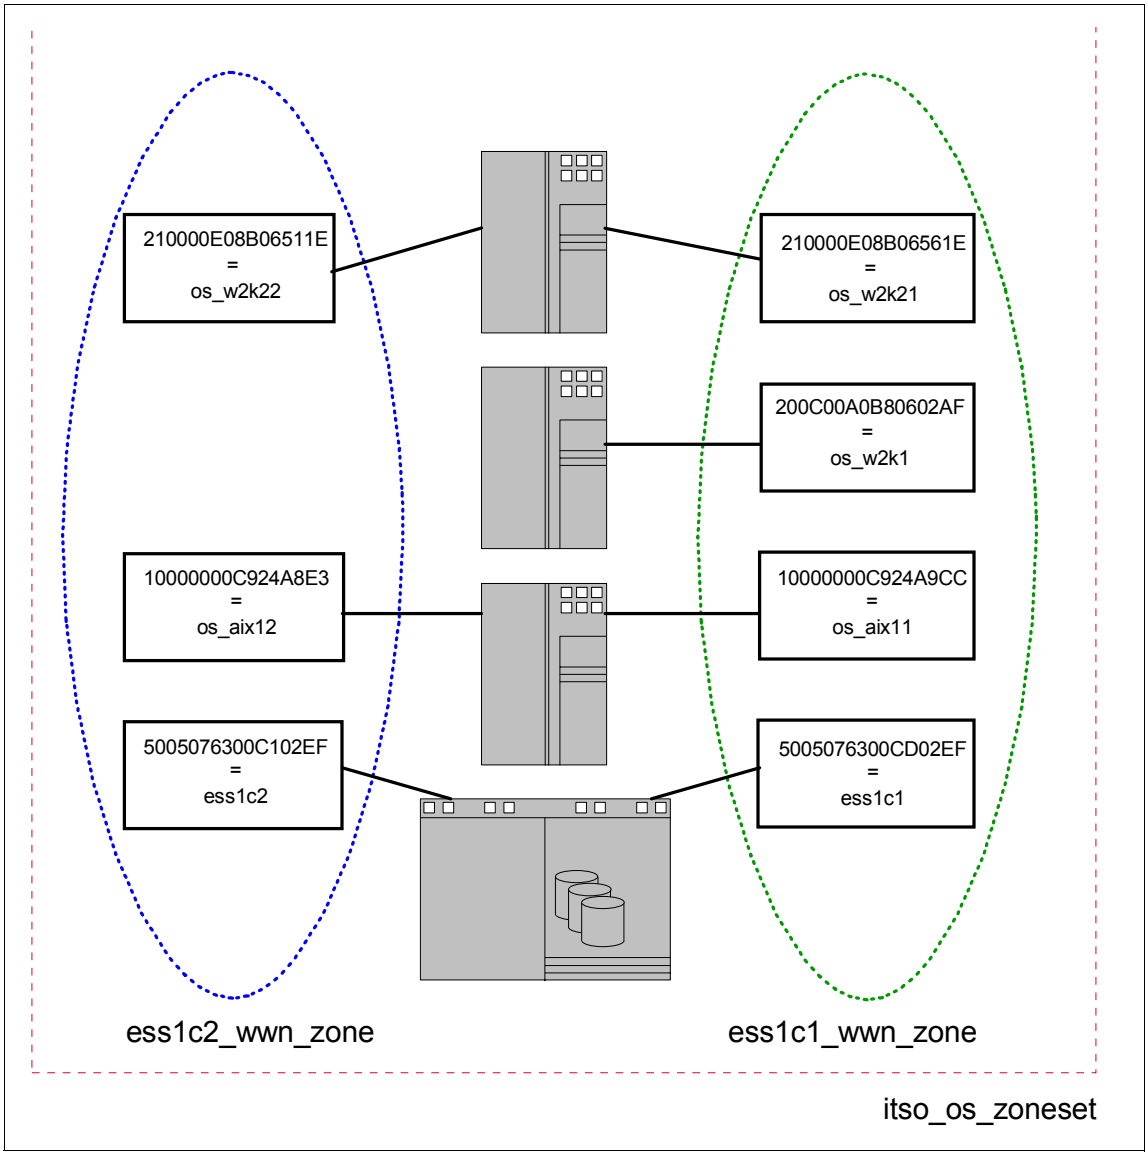


*Figure A-15 Logical view of our two name server zones*

Zone *ess1c1_zone* has four director ports as members, including the one ESS ess1c1 is attached to.

Zone *ess2c2_zone* has three director ports as members, including the port that is used to attach ESS ess2c2.

By implementing such a layout, each server is attached to two zones. Since each zone has a different ESS port as a member, the server will get exactly two paths to the ESS. Use of **`pvlinks`** and ESS based LUN masking will enable sufficient multipathing.

To actually implement name server zoning, click the specific fabric icon, in our case *Fab_T_1*, in the navigation tree, and then select the **Zoning** tab as shown in Figure A-16**.**

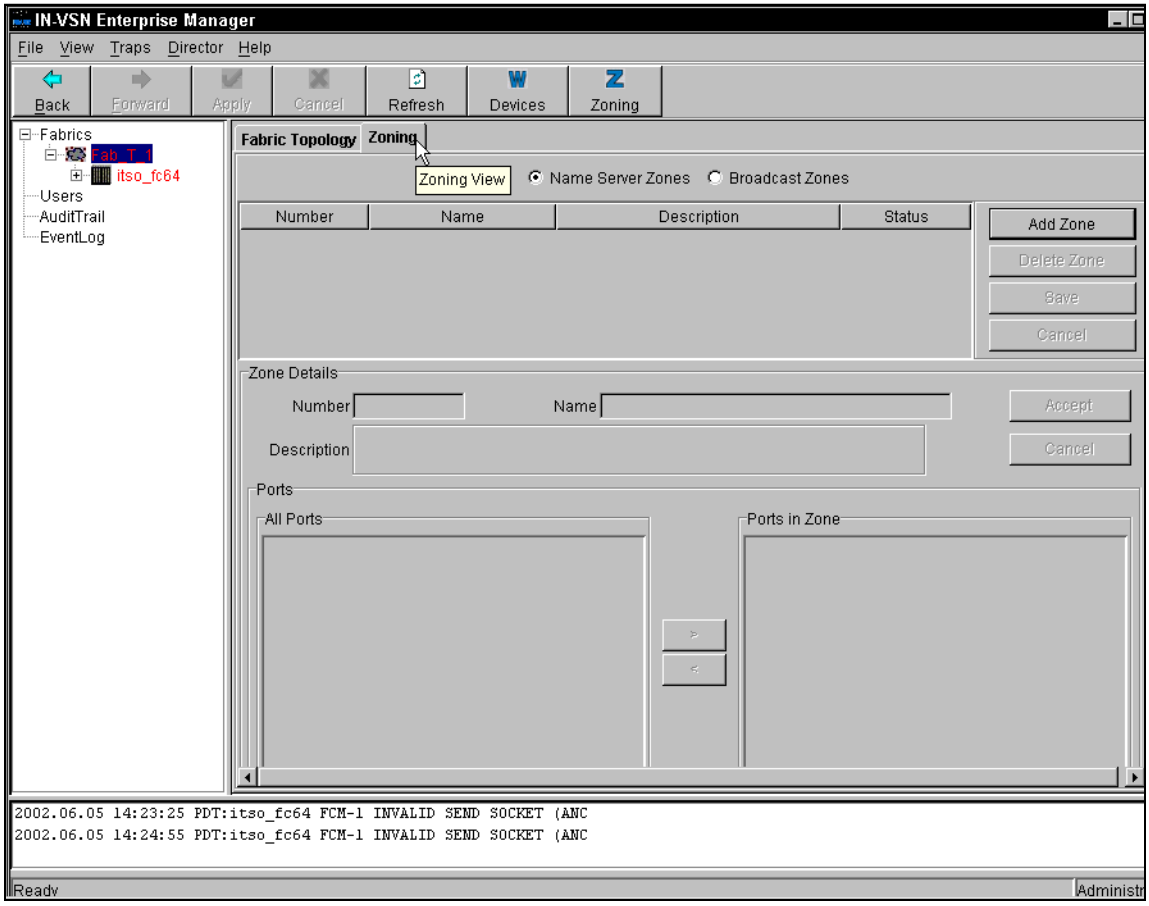


*Figure A-16 IN-VSN: Entering the zone screen in fabric view mode*

Initially, there are no zones defined at all. Remember that in such a case, the only effective zone is the default zone. This enables all ports to communicate with all other ports (any-to-any). Only zoned ports will have limited but controlled access. All unzoned ports are considered as being part of the default orphan zone, which enables them to communicate without any restrictions.

Leave the **Name Server Zone** field selected. Click **Add** to create a new zone. Enter zone number, zone name, and a description as illustrated in Figure A-17.

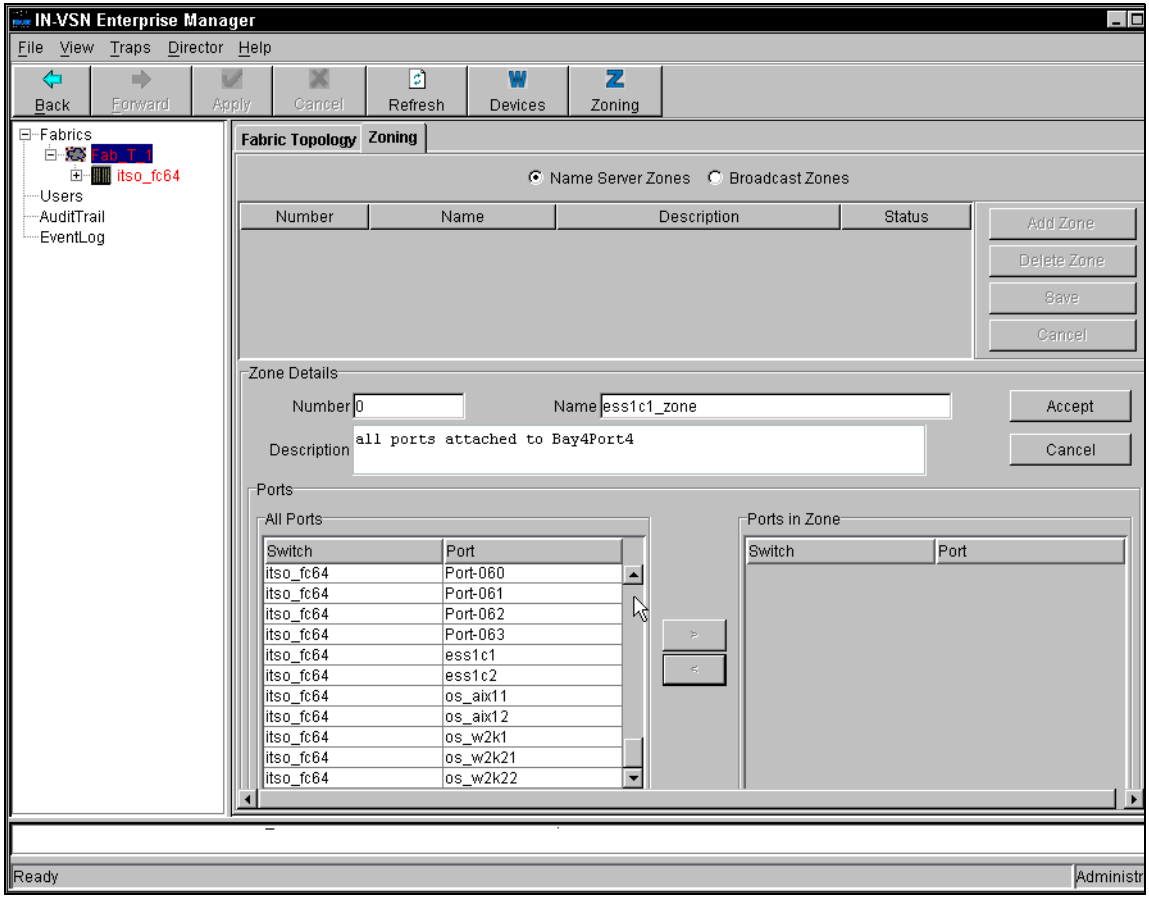


*Figure A-17 IN-VSN: Entering number and name for a new zone*

From the Ports scroll list, select the ones you want to be a member of that zone as shown in Figure A-18.

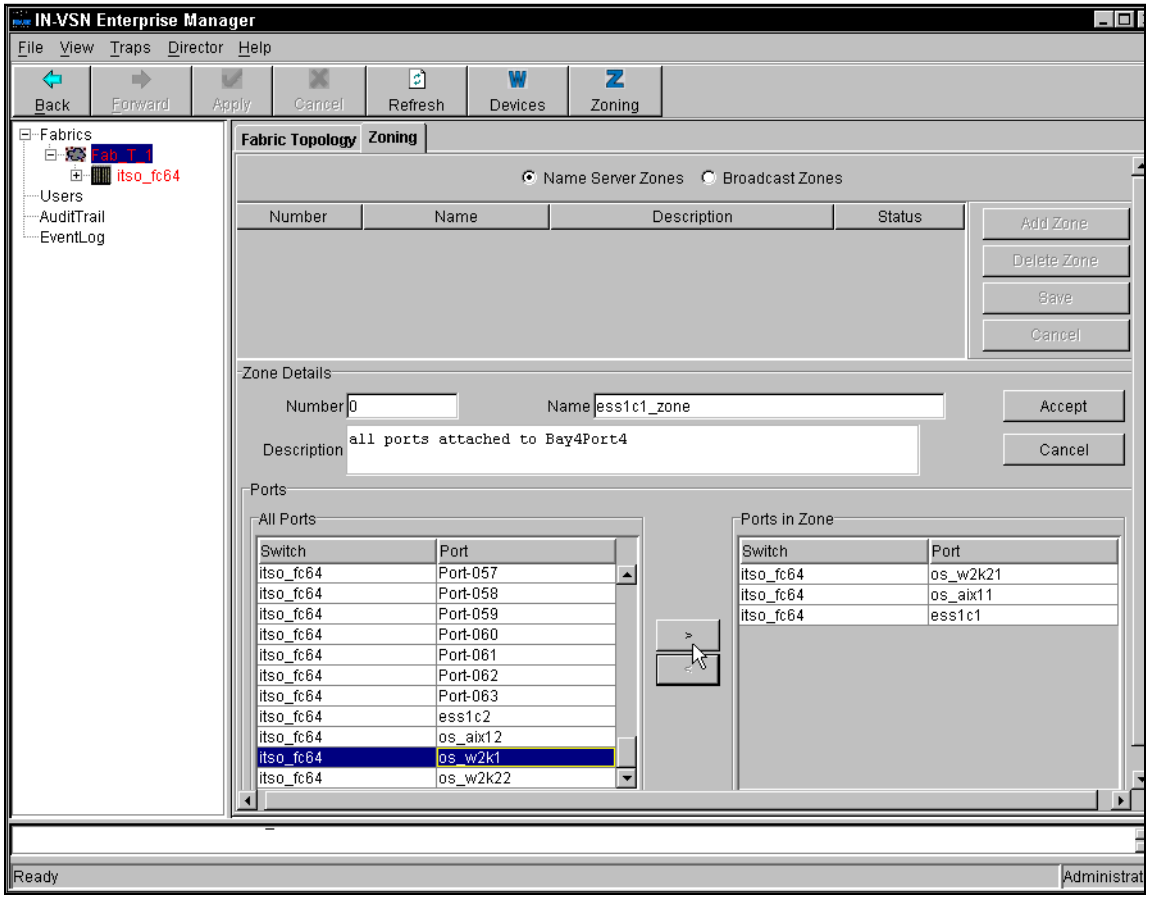


*Figure A-18 IN-VSN: Selecting the members of a new zone*

We added these ports to the zone *ess1c1_zone*:

- ess1c1
- os_aix11
- os_w2k1
- os_w2k21

The effect is that all these server ports will have access only to the ESS port connected to the director on port “ess1c1”.

After putting all needed ports in the **Ports in Zone list**, we continue with clicking **Accept**. This is shown in Figure A-19.

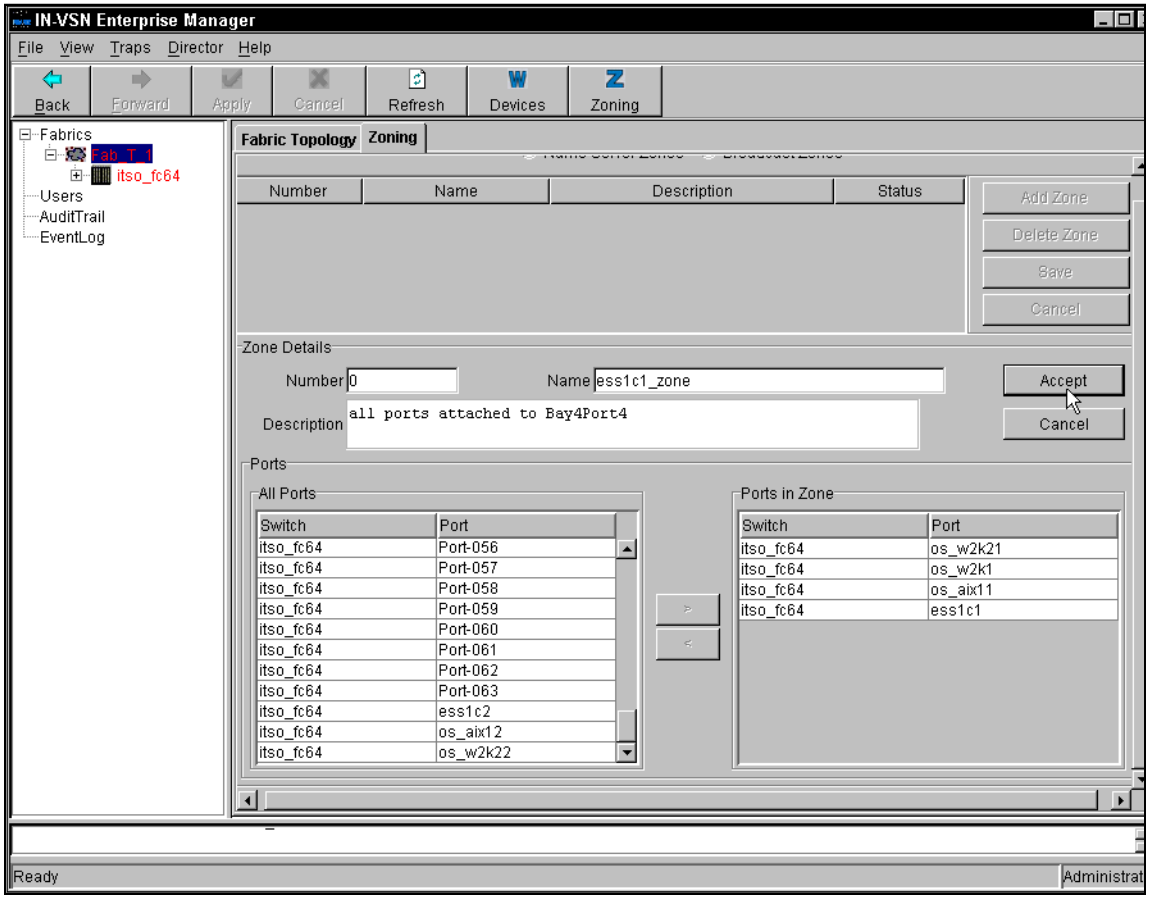


*Figure A-19 IN-VSN: Accepting settings for first new zone*

However, now we are getting an error message, since our previously chosen zone number of **0** is not valid. This is shown in Figure A-20.

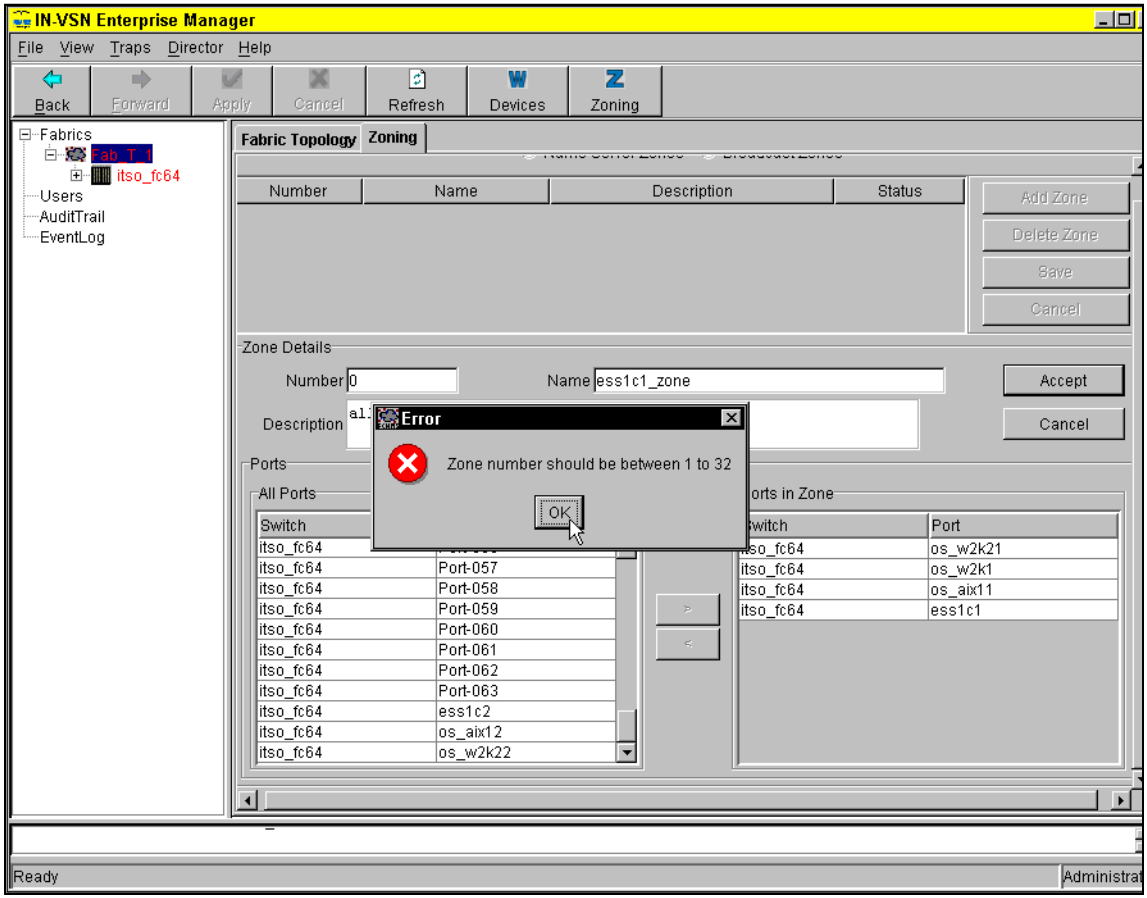


*Figure A-20 IN-VSN: Zone number 0 is not allowed to use*

So we just change this zone number to **1**, and click **Accept** again. Now this action is accepted.

This zone is now listed in the zone list with a status of *Added*.

To make this zone effective, click **Save** as shown in Figure A-21.

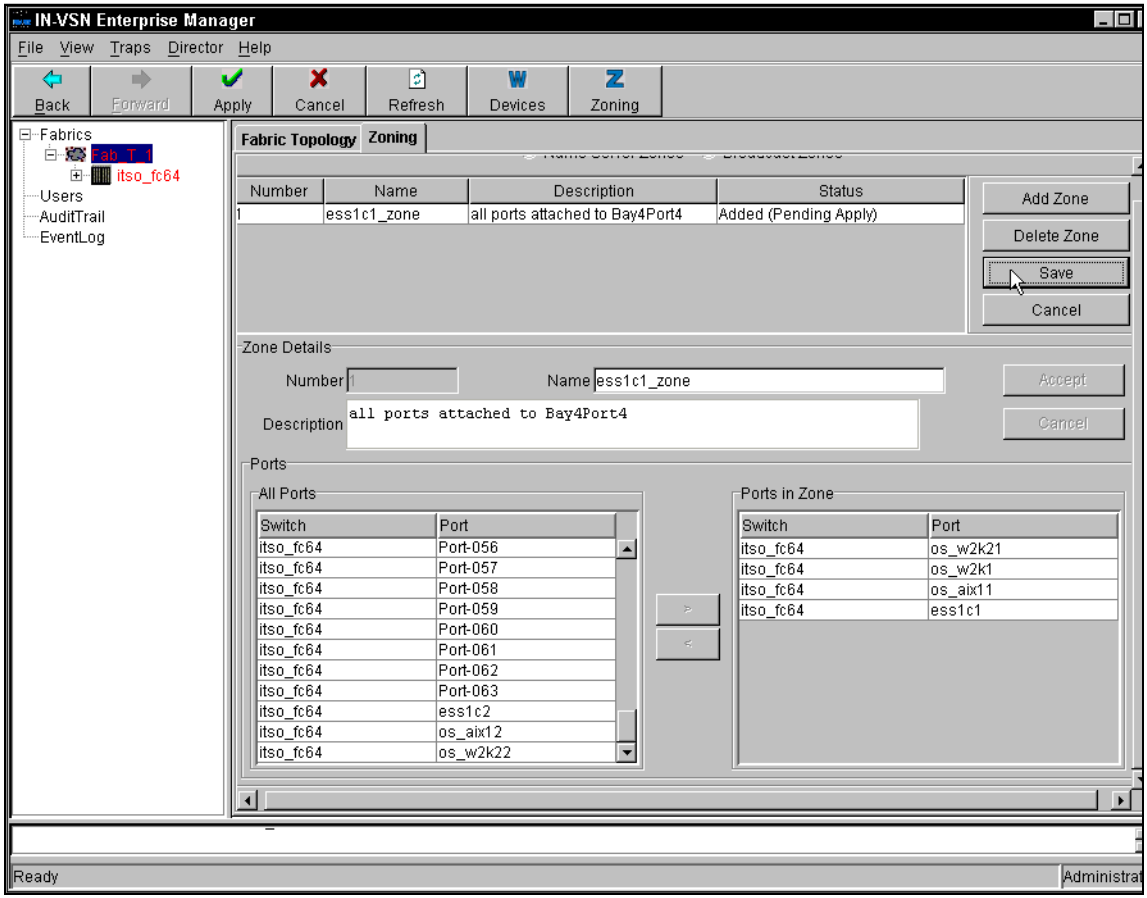


*Figure A-21 IN-VSN: Saving a newly added zone*

Using the same procedure, we now add a second zone called *ess1c2_zone* which has these ports as members:

- ess1c2
- os_aix12
- os_w2k22

After saving this second zone, the zone list contains both zones with a status of Saved as illustrated in Figure A-22.

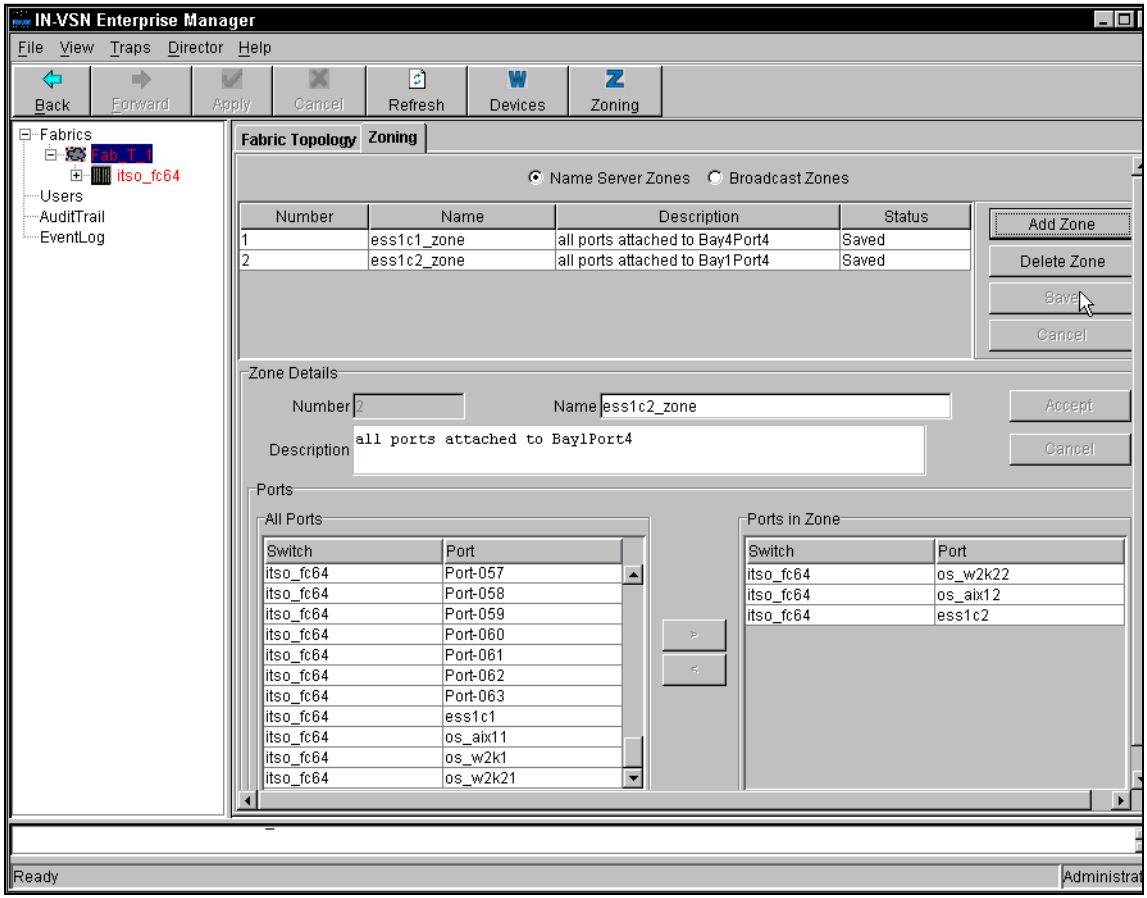


*Figure A-22 IN-VSN: Zone list with both zones added and saved*

We now have two zones defined in a fabric consisting of one director.

Since both zones are saved and all specified external ports are actually attached and online, we can start I/O between these servers and our ESS.

# Translative loop port (TL_Port) mode

Normally, private loop ports are not able to talk with other ports outside of their own loop. Therefore, their addressing is different from normal ports. However, with the *Translative Loop* mode of specific director ports, it is possible to let them talk with other ports in the fabric.

To enable this you have to set the Admin Type of this particular director port to **TL** as shown in Figure A-23.

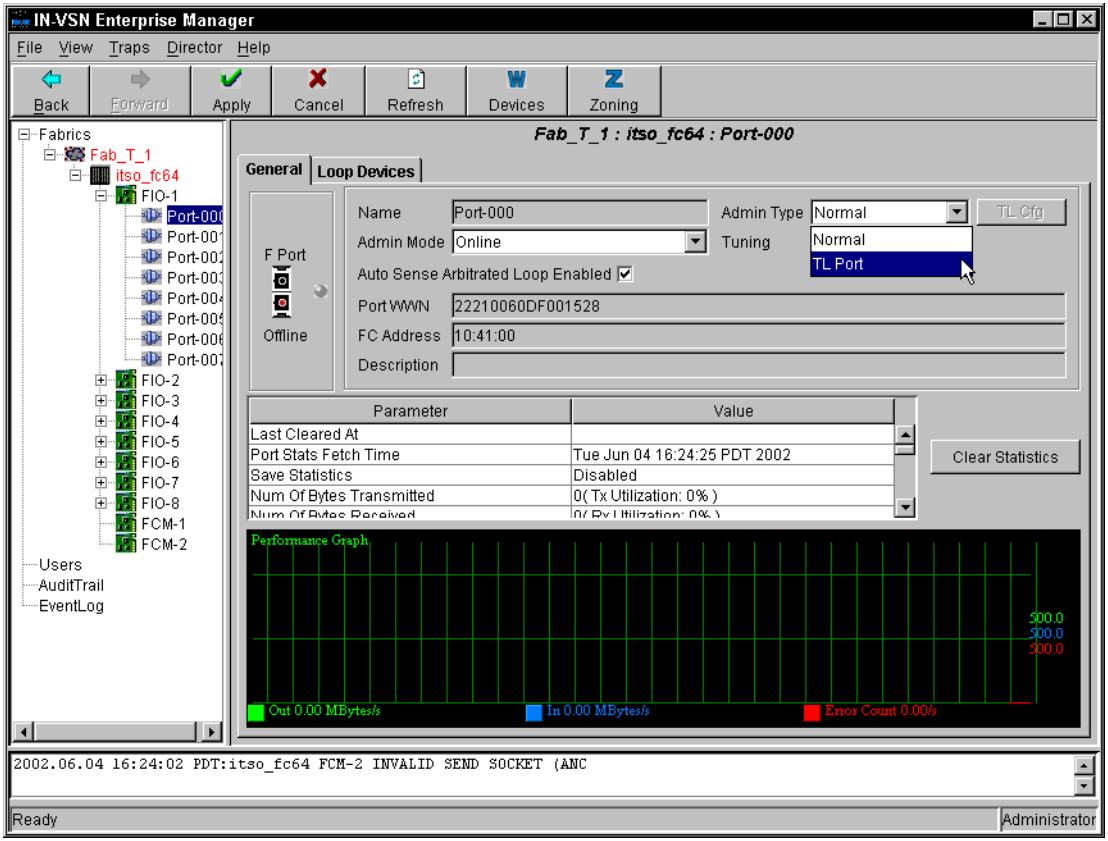


*Figure A-23 IN-VSN: Setting a port to Translative Loop mode*

Each TL port has its own Translation Entries list. This list contains the set of off-loop devices which can communicate with this TL port. Off-loop devices can be:

- Any other N_Port in the fabric (initiators or targets)
- Any other public loop port in the fabric (NL_Port)
- Any other private loop port in the fabric (TL_Port)

The creation and maintenance of this list is different between private targets and private initiators. Therefore, you have to specify whether a particular TL_Port has private targets or a private initiator attached to it. Click the **TL-CFG** button to open the TL Config window as shown in Figure A-24.

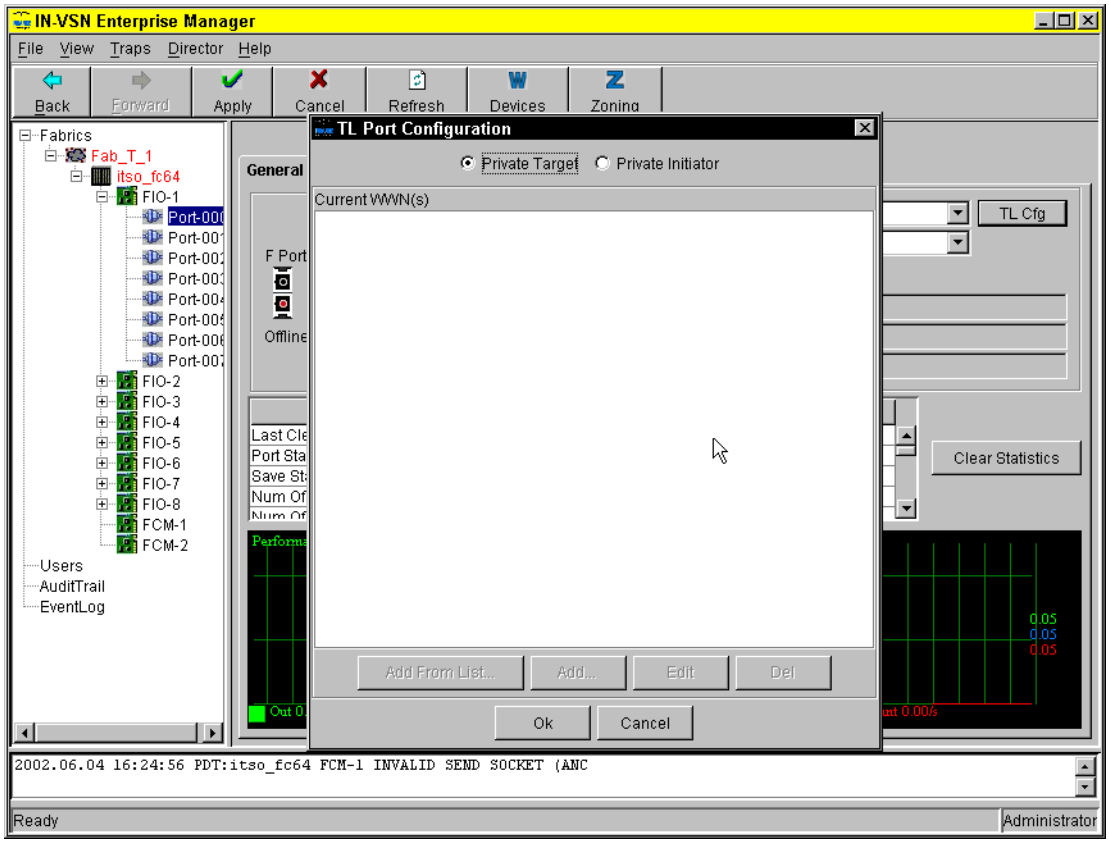


*Figure A-24 IN-VSN: Selecting target or initiator mode for TL ports*

Select **Private Target** if you have targets attached to this TL_Port which should be accessed by other private initiators — for instance, private loop JBODs accessed by another private initiator.

Select **Private Initiator** if you have private loop initiators attached, such as elderly FC host adapters. Also, select this option if you want a private target device on this TL_Port to get accessed by a public initiator.

**Note:** The selection of Private Target or Private Initiator for a TL_Port does not necessarily mean that you have such types attached to it.

The private target option is only useful if a private initiator should access this particular target port.

If a public initiator should get access to this private target you should actually set this private target to Private Initiator. Then you should proceed to specify the WWN of the Public initiator that should have explicit access to this private target.

In such a scenario, the Private Initiator label for the radio button may be misleading.

The translation entries list for private target TL_Ports contains all initiators that try to get access to these targets. An Auto Learning feature enables this translation entries list to be updated automatically. That is why, for private targets, you cannot specify any WWNs in the TL-CFG window. However, you should limit the actual number of initiators trying to communicate with this TL_Port to a maximum of 31.

An Auto Learning feature for TL_Ports attached to private initiators is not available. All devices that this private initiator has to communicate with must be added to the list manually. Once you have selected **Private Initiator** in the TL-CFG window, some additional buttons become active. Click **Add from List** to add devices that are currently known to the fabric, as illustrated in Figure A-25.

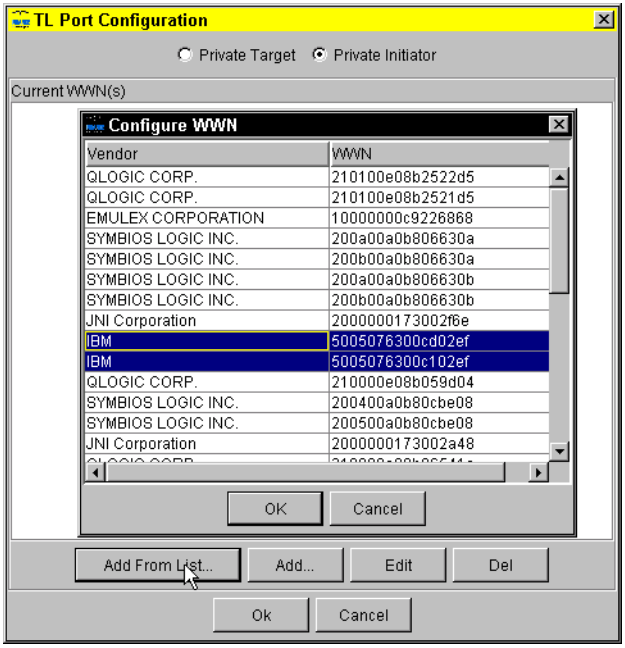


*Figure A-25 IN-VSN: adding possible targets*

If you want to add device addresses that are currently not attached to the fabric, then you can add them by typing in their WWN. Click the **Add** button and enter the WWN as shown in Figure A-26.

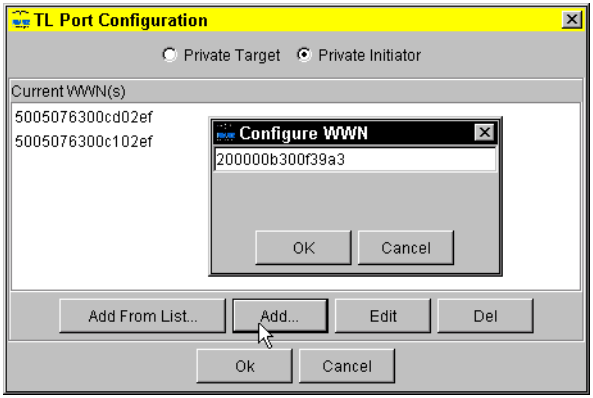


*Figure A-26 IN-VSN: Adding WWN targets*

Each TL_Port has its own translation entries list, as shown in Figure A-27.

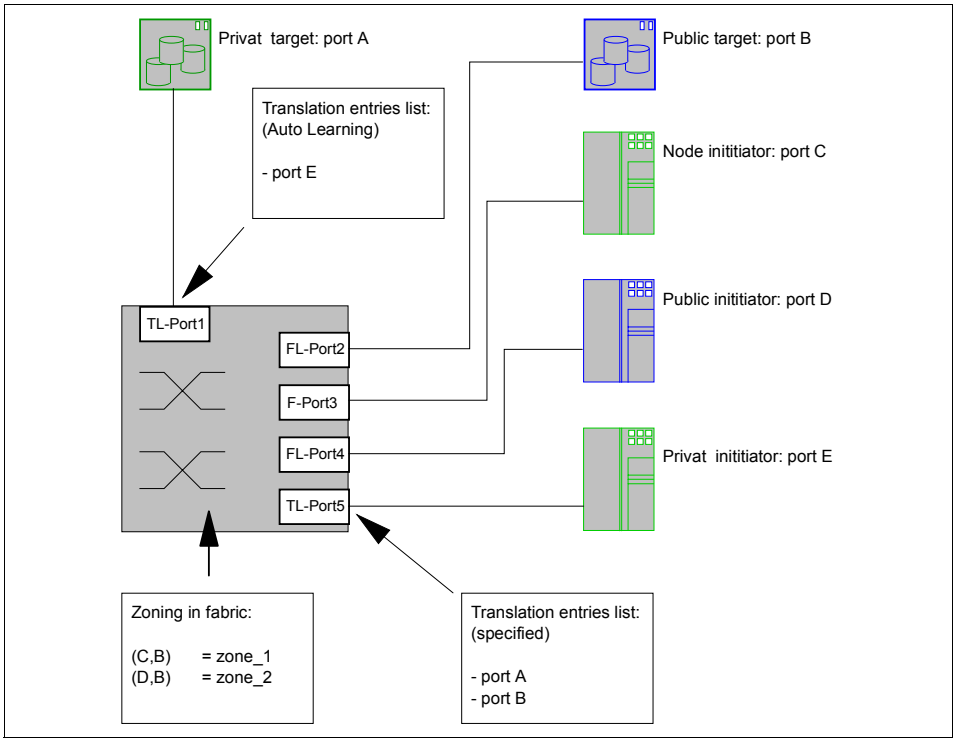


*Figure A-27 Translation Entries List with zoning*

Only TL_Ports have translation entries lists, in our case port1 and port5.

The list of port5 (initiator) was manually created by adding the WWN of portA and portB. It means that port5 with its attached portE could have access to the ports A and B. Since TL-Cfg settings overrule name server zone settings, we do not need to zone these ports.

Furthermore, portD has access to portB. This is not implemented through a translation entries list, but by using zoning. Since portD is not a private loop port, there is no list for this port anyway.

The list for portA is automatically created using the auto learning feature. Once the initiator portE has actually logged in to portA, then portE will be added to the list.

**Note:** Even if the translation entries table for private target ports is updated automatically, you will not see its content. The auto learning feature is always used in the background. You cannot deactivate it or activate it.

## TL_Ports and zoning

Once the translation entries list is created, these TL_Ports can actually talk to other ports in the fabric.

But the question remains: How does this TL_Port configuration interact with name server zoning, or with hard zoning?

First, TL_Port configurations cannot cross the boundaries of hard zoning.

Second, name server zoning does not have any impact on TL_Port communication:

- Private Target TL_Ports propagate all their translative device addresses to the fabric. However, they can just be used by other Private Initiators. To enable communication between them the TL-CFG list of this private initiator must include the address of the private target.
- Private initiator devices have only access to devices specified in their TL-CFG list.
- Even if a TL port is part of a name server zone, it will not see the members of this name server zone unless they are put into the TL-CFG list.
- On the other hand, the TL-CFG list of a TL_Port is effective even if the name server layout would imply it otherwise.

Note that this feature of name server zones is limited to TL_Ports only. Normal public loop ports will be handled as normal director ports with all the normal consequences of name server zoning.

## Impact of LIP in the fabric

Ports in loop networks use a process called *Loop Initialization Primitive sequence* (LIP) to establish their port addresses. All members of that loop are involved in a LIP.

Loop initialization occurs whenever there is a change in the layout of a loop, such as adding a new node, leaving of a node, or breaks in service in the loop.

The start of a LIP causes data transfers in progress to stop, thereby severely affecting the performance and availability of Arbitrated Loops.

However, since CNT is using a *Translation Entries List,* these LIPs will not be propagated to other fabric members. This is true even if multiple private loop ports are zoned together. Therefore, all LIP impact is limited only to the external physical loop (for example, an FC_AL hub).

This feature is a major benefit of using TL_Ports.

## Cascading in T_Port mode

Cascading is a very effective and easy to use method to extend distances and increase maximum port count. However, we need to consider some of the following implications:

- By connecting multiple directors, they will become one fabric inheriting all the rules for one fabric, including:
  - Name server zones are valid fabric-wide.
  - Numbers and names of name server zones must be unique fabric-wide.
  - Either all or none of the director ports in a fabric are hard zoned.
  - Each director or switch in a fabric must have a unique switch ID.
  - When having no zoning implemented at all then all nodes have access to any other node fabric-wide.
- T_Ports cannot be part of name server zones.
- When having hard zones installed that cross chassis boundaries, then both parts of the hard zone must be directly connected by at least one ISL. You must have dedicated ISLs for each hard zone spanning multiple chassis.
- Up to eight FC/9000 directors can be cascaded to create one fabric.

## Migrating from T_Port to E_Port

This section covers the migration process from T_Port to E_Port assuming that your are running firmware version 3.0.x or above.

**Attention:** Migrating from T_Port to E_Port removes all previous zoning information. Hard zoning configuration, Name Server zoning, as well as the Broadcast zoning configuration, are removed. This process is also disruptive for all traffic carried by the director.

To migrate from T_Port mode to E_Port mode, follow these steps:

1. In fabrics with more than one Director, put all Inter-Switch Link (ISL) ports offline. Note that this MUST be done for all ISLs in the FABRIC before starting the address mode change. If preferred, the ISL may be physically disconnected.
2. Put all boards in offline state. Since the T to E conversion is an address mode change, all ports should be put offline until the fabric is updated to avoid transient address mismatch.
3. Click the director icon on the navigation tree and we select the “System Configuration” tab. Then select “E Port” in the “Inter Switch Link Type” as shown in Figure A-28.

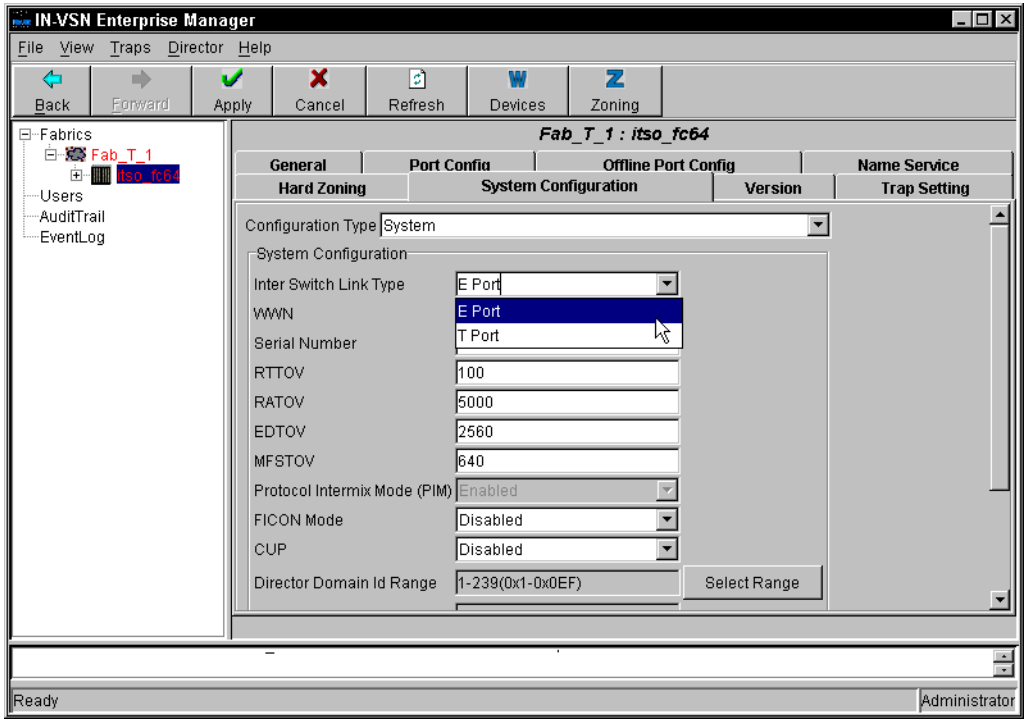


*Figure A-28 IN-VSN: Configuring the director in E Port mode*

4. Click **Apply** to validate the changes.IN-VSN prompts for a confirmation as shown in Figure A-29, as the process will be disruptive for the traffic.

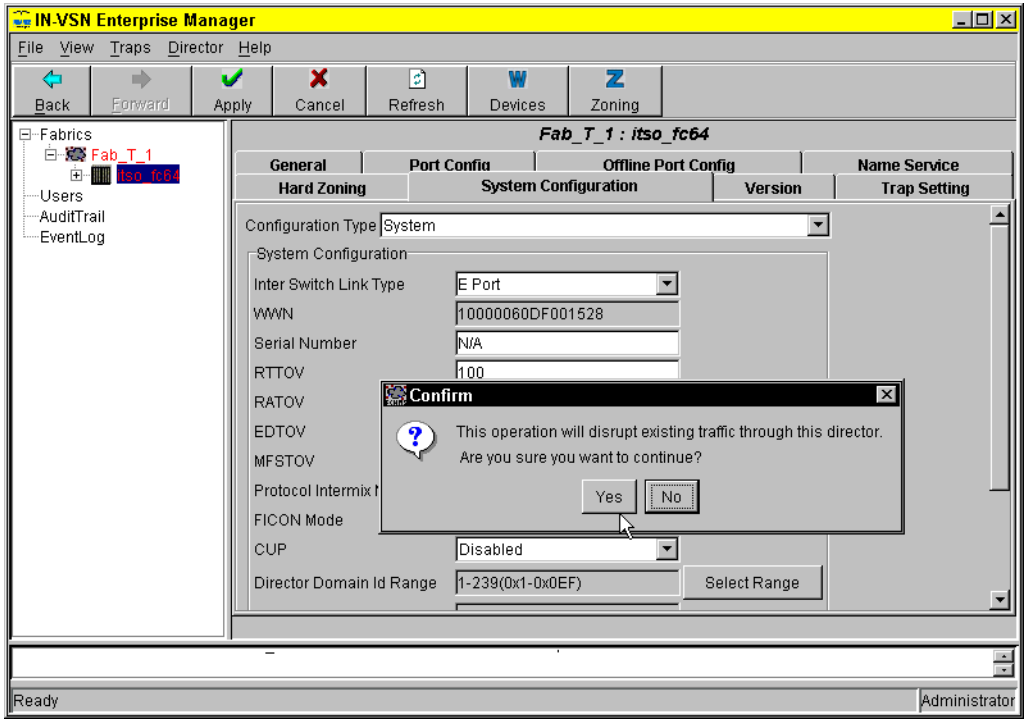


*Figure A-29 IN-VSN: ISL mode configuration confirmation window*

5. Once the configuration process is finished, you can return to the director view. You should wait a minute or so after the end of the process before performing any operation on the director. The IN-VSN tabs are now different as shown in Figure A-30.

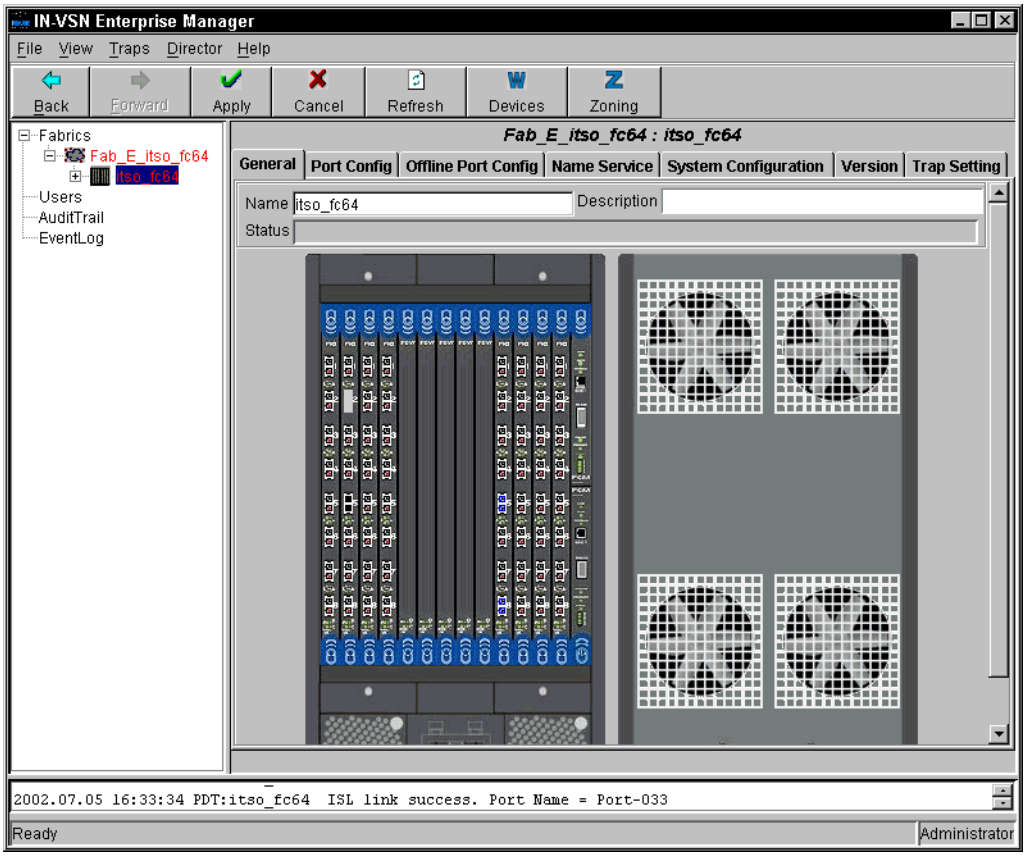


*Figure A-30 IN-VSN: Director view in E Port mode*

The hardware zoning tab at the director level and the Name Server zoning tab at the fabric level are not available in E_Port mode. Note that the fabric name has changed to Fab_E_itso_fc64.

6. Repeat these steps for each Director in the fabric.

# Glossary

**8b/10b** A data encoding scheme developed by IBM, translating byte-wide data to an encoded 10-bit format. Fibre Channel's FC-1 level defines this as the method to be used to encode and decode data transmissions over the Fibre Channel.

**Adapter** A hardware unit that aggregates other I/O units, devices or communications links to a system bus.

**ADSM** ADSTAR Distributed Storage Manager.

**Agent** (1) In the client-server model, the part of the system that performs information preparation and exchange on behalf of a client or server application. (2) In SNMP, the word agent refers to the managed system. See also: Management Agent

**Aggregation** In the Storage Networking Industry Association Storage Model (SNIA), “virtualization” is known as “aggregation”. This aggregation can take place at the file level or at the level of individual blocks that are transferred to disk.

**AIT** Advanced Intelligent Tape - A magnetic tape format by Sony that uses 8mm cassettes, but is only used in specific drives.

**AL** See Arbitrated Loop

**AL_PA** Arbitrated Loop Physical Address

**ANSI** American National Standards Institute - The primary organization for fostering the development of technology standards in the United States. The ANSI family of Fibre Channel documents provide the standards basis for the Fibre Channel architecture and technology. See FC-PH

**Arbitration** The process of selecting one respondent from a collection of several candidates that request service concurrently.

**Arbitrated Loop** A Fibre Channel interconnection technology that allows up to 126 participating node ports and one participating fabric port to communicate.

**ATL** Automated Tape Library - Large scale tape storage system, which uses multiple tape drives and mechanisms to address 50 or more cassettes.

**ATM** Asynchronous Transfer Mode - A type of packet switching that transmits fixed-length units of data.

**Backup** A copy of computer data that is used to recreate data that has been lost, mislaid, corrupted, or erased. The act of creating a copy of computer data that can be used to recreate data that has been lost, mislaid, corrupted or erased.

**Bandwidth** Measure of the information capacity of a transmission channel.

**Bridge** (1) A component used to attach more than one I/O unit to a port. (2) A data communications device that connects two or more networks and forwards packets between them. The bridge may use similar or dissimilar media and signaling systems. It operates at the data link level of the OSI model. Bridges read and filter data packets and frames.

**Bridge/Router** A device that can provide the functions of a bridge, router or both concurrently. A bridge/router can route one or more protocols, such as TCP/IP, and bridge all other traffic. See also: Bridge, Router

**Broadcast** Sending a transmission to all N_Ports on a fabric.

**Channel** A point-to-point link, the main task of which is to transport data from one point to another.

**Channel I/O** A form of I/O where request and response correlation is maintained through some form of source, destination and request identification.

**CIFS** Common Internet File System

**Class of Service** A Fibre Channel frame delivery scheme exhibiting a specified set of delivery characteristics and attributes.

**Class-1** A class of service providing dedicated connection between two ports with confirmed delivery or notification of non-deliverability.

**Class-2** A class of service providing a frame switching service between two ports with confirmed delivery or notification of non-deliverability.

**Class-3** A class of service providing frame switching datagram service between two ports or a multicast service between a multicast originator and one or more multicast recipients.

**Class-4** A class of service providing a fractional bandwidth virtual circuit between two ports with confirmed delivery or notification of non-deliverability.

**Class-6** A class of service providing a multicast connection between a multicast originator and one or more multicast recipients with confirmed delivery or notification of non-deliverability.

**Client** A software program used to contact and obtain data from a *server* software program on another computer -- often across a great distance. Each *client* program is designed to work specifically with one or more kinds of server programs and each server requires a specific kind of client program.

**Client/Server** The relationship between machines in a communications network. The client is the requesting machine, the server the supplying machine. Also used to describe the information management relationship between software components in a processing system.

**Cluster** A type of parallel or distributed system that consists of a collection of interconnected whole computers and is used as a single, unified **computing resource.**

**Coaxial Cable** A transmission media (cable) used for high speed transmission. It is called *coaxial* because it includes one physical channel that carries the signal surrounded (after a layer of insulation) by another concentric physical channel, both of which run along the same axis. The inner channel carries the signal and the outer channel serves as a ground.

**Controller** A component that attaches to the system topology through a channel semantic protocol that includes some form of request/response identification.

**CRC** Cyclic Redundancy Check - An error-correcting code used in Fibre Channel.

**DASD** Direct Access Storage Device - any on-line storage device: a disc, drive or CD-ROM.

**DAT** Digital Audio Tape - A tape media technology designed for very high quality audio recording and data backup. DAT cartridges look like audio cassettes and are often used in mechanical auto-loaders. typically, a DAT cartridge provides 2GB of storage. But new DAT systems have much larger capacities.

**Data Sharing** A SAN solution in which files on a storage device are shared between multiple hosts.

**Datagram** Refers to the Class 3 Fibre Channel Service that allows data to be sent rapidly to multiple devices attached to the fabric, with no confirmation of delivery.

**dB** Decibel - a ratio measurement distinguishing the percentage of signal attenuation between the input and output power. Attenuation (loss) is expressed as dB/km.

**Disk Mirroring** A fault-tolerant technique that writes data simultaneously to two hard disks using the same hard disk controller.

**Disk Pooling** A SAN solution in which disk storage resources are pooled across multiple hosts rather than be dedicated to a specific host.

**DLT** Digital Linear Tape - A magnetic tape technology originally developed by Digital Equipment Corporation (DEC) and now sold by Quantum. DLT cartridges provide storage capacities from 10 to 35GB.

**E_Port** Expansion Port - a port on a switch used to link multiple switches together into a Fibre Channel switch fabric.

**ECL** Emitter Coupled Logic - The type of transmitter used to drive copper media such as Twinax, Shielded Twisted Pair, or Coax.

**Enterprise Network** A geographically dispersed network under the auspices of one organization.

**Entity** In general, a real or existing thing from the Latin ens, or being, which makes the distinction between a thing's existence and it qualities. In programming, engineering and probably many other contexts, the word is used to identify units, whether concrete things or abstract ideas, that have no ready name or label.

**ESCON®** Enterprise System Connection

**Exchange** A group of sequences which share a unique identifier. All sequences within a given exchange use the same protocol. Frames from multiple sequences can be multiplexed to prevent a single exchange from consuming all the bandwidth. See also: Sequence

**F_Node** Fabric Node - a fabric attached node.

**F_Port** Fabric Port - a port used to attach a Node Port (N_Port) to a switch fabric.

**Fabric** Fibre Channel employs a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices. The term is most often used to describe a more complex network utilizing hubs, switches and gateways.

**Fabric Login** Fabric Login (FLOGI) is used by an N_Port to determine if a fabric is present and, if so, to initiate a session with the fabric by exchanging service parameters with the fabric. Fabric Login is performed by an N_Port following link initialization and before communication with other N_Ports is attempted.

**FC** Fibre Channel

**FC-0** Lowest level of the Fibre Channel Physical standard, covering the physical characteristics of the interface and media

**FC-1** Middle level of the Fibre Channel Physical standard, defining the 8b/10b encoding/decoding and transmission protocol.

**FC-2** Highest level of the Fibre Channel Physical standard, defining the rules for signaling protocol and describing transfer of frame, sequence and exchanges.

**FC-3** The hierarchical level in the Fibre Channel standard that provides common services such as striping definition.

**FC-4** The hierarchical level in the Fibre Channel standard that specifies the mapping of upper-layer protocols to levels below.

**FCA Fibre Channel Association.**

**FC-AL** Fibre Channel Arbitrated Loop - A reference to the Fibre Channel Arbitrated Loop standard, a shared gigabit media for up to 127 nodes, one of which may be attached to a switch fabric. See also: Arbitrated Loop.

**FC-CT** Fibre Channel common transport protocol

**FC-FG** Fibre Channel Fabric Generic - A reference to the document (ANSI X3.289-1996) which defines the concepts, behavior and characteristics of the Fibre Channel Fabric along with suggested partitioning of the 24-bit address space to facilitate the routing of frames.

**FC-FP** Fibre Channel HIPPI Framing Protocol - A reference to the document (ANSI X3.254-1994) defining how the HIPPI framing protocol is transported via the Fibre Channel

**FC-GS** Fibre Channel Generic Services -A reference to the document (ANSI X3.289-1996) describing a common transport protocol used to communicate with the server functions, a full X500 based directory service, mapping of the Simple Network Management Protocol (SNMP) directly to the Fibre Channel, a time server and an alias server.

**FC-LE** Fibre Channel Link Encapsulation - A reference to the document (ANSI X3.287-1996) which defines how IEEE 802.2 Logical Link Control (LLC) information is transported via the Fibre Channel.

**FC-PH** A reference to the Fibre Channel Physical and Signaling standard ANSI X3.230, containing the definition of the three lower levels (FC-0, FC-1, and FC-2) of the Fibre Channel.

**FC-PLDA** Fibre Channel Private Loop Direct Attach - See PLDA.

**FC-SB** Fibre Channel Single Byte Command Code Set - A reference to the document (ANSI X.271-1996) which defines how the ESCON command set protocol is transported using the Fibre Channel.

**FC-SW** Fibre Channel Switch Fabric - A reference to the ANSI standard under development that further defines the fabric behavior described in FC-FG and defines the communications between different fabric elements required for those elements to coordinate their operations and management address assignment.

**FC Storage Director** See SAN Storage Director

**FCA** Fibre Channel Association - a Fibre Channel industry association that works to promote awareness and understanding of the Fibre Channel technology and its application

and provides a means for implementers to support the standards committee activities.

**FCLC** Fibre Channel Loop Association - an independent working group of the Fibre Channel Association focused on the marketing aspects of the Fibre Channel Loop technology.

**FCP** Fibre Channel Protocol - the mapping of SCSI-3 operations to Fibre Channel.

**Fiber Optic** Refers to the medium and the technology associated with the transmission of information along a glass or plastic wire or fiber.

**Fibre Channel** A technology for transmitting data between computer devices at a data rate of up to 4 Gb/s. It is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

**FICON** Fibre Connection - A next-generation I/O solution for IBM S/390 parallel enterprise server.

**FL_Port** Fabric Loop Port - the access point of the fabric for physically connecting the user's Node Loop Port (NL_Port).

**FLOGI** See Fabric Log In

**Frame** A linear set of transmitted bits that define the basic transport unit. The frame is the most basic element of a message in Fibre Channel communications, consisting of a 24-byte header and zero to 2112 bytes of data. See also: Sequence

**FSP** Fibre Channel Service Protocol - The common FC-4 level protocol for all services, transparent to the fabric type or topology.

**FSPF** Fabric Shortest Path First - is an intelligent path selection and routing standard and is part of the Fibre Channel Protocol.

**Full-Duplex** A mode of communications allowing simultaneous transmission and reception of frames.

**G_Port** Generic Port - a generic switch port that is either a Fabric Port (F_Port) or an Expansion Port (E_Port). The function is automatically determined during login.

**Gateway** A node on a network that interconnects two otherwise incompatible networks.

**Gb/s** Gigabits per second. Also sometimes referred to as Gbps. In computing terms it is approximately 1,000,000,000 bits per second. Most precisely it is 1,073,741,824 (1024 x 1024 x 1024) bits per second.

**GB/s** Gigabytes per second. Also sometimes referred to as GBps. In computing terms it is approximately 1,000,000,000 bytes per second. Most precisely it is 1,073,741,824 (1024 x 1024 x 1024) bytes per second.

**GBIC** GigaBit Interface Converter - Industry standard transceivers for connection of Fibre Channel nodes to arbitrated loop hubs and fabric switches.

**Gigabit** One billion bits, or one thousand megabits.

**GLM** Gigabit Link Module - a generic Fibre Channel transceiver unit that integrates the key functions necessary for installation of a Fibre Channel media interface on most systems.

**Half-Duplex** A mode of communications allowing either transmission or reception of frames at any point in time, but not both (other

than link control frames which are always permitted).

**Hardware** The mechanical, magnetic and electronic components of a system, e.g., computers, telephone switches, terminals and the like.

**HBA** Host Bus Adapter

**HIPPI** High Performance Parallel Interface - An ANSI standard defining a channel that transfers data between CPUs and from a CPU to disk arrays and other peripherals.

**HMMP** HyperMedia Management Protocol

**HMMS** HyperMedia Management Schema - the definition of an implementation-independent, extensible, common data description/schema allowing data from a variety of sources to be described and accessed in real time regardless of the source of the data. See also: WEBM, HMMP

**hop** A FC frame may travel from a switch to a director, a switch to a switch, or director to a director which, in this case, is one hop.

**HSM** Hierarchical Storage Management - A software and hardware system that moves files from disk to slower, less expensive storage media based on rules and observation of file activity. Modern HSM systems move files from magnetic disk to optical disk to magnetic tape.

**HUB** A Fibre Channel device that connects nodes into a logical loop by using a physical star topology. Hubs will automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

**HUB Topology** see Loop Topology

**Hunt Group** A set of associated Node Ports (N_Ports) attached to a single node, assigned a special identifier that allows any frames containing this identifier to be routed to any available Node Port (N_Port) in the set.

**In-band Signaling** This is signaling that is carried in the same channel as the information. Also referred to as in-band.

**In-band virtualization** An implementation in which the virtualization process takes place in the data path between servers and disk systems. The virtualization can be implemented as software running on servers or in dedicated engines.

**Information Unit** A unit of information defined by an FC-4 mapping. Information Units are transferred as a Fibre Channel Sequence.

**Intermix** A mode of service defined by Fibre Channel that reserves the full Fibre Channel bandwidth for a dedicated Class 1 connection, but also allows connection-less Class 2 traffic to share the link if the bandwidth is available.

**Inter switch link** A FC connection between switches and/or directors. Also known as ISL.

**I/O** Input/output

**IP** Internet Protocol

**IPI** Intelligent Peripheral Interface

**ISL** See Inter switch link.

**Isochronous Transmission** Data transmission which supports network-wide timing requirements. A typical application for isochronous transmission is a broadcast environment which needs information to be delivered at a predictable time.

**JBOD** Just a bunch of disks.

**Jukebox** A device that holds multiple optical disks and one or more disk drives, and can swap disks in and out of the drive as needed.

**L_Port** Loop Port - A node or fabric port capable of performing Arbitrated Loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.

**LAN** See Local Area Network - A network covering a relatively small geographic area (usually not larger than a floor or small building). Transmissions within a Local Area Network are mostly digital, carrying data among stations at rates usually above one megabit/s.

**Latency** A measurement of the time it takes to send a frame between two locations.

**LC** Lucent Connector. A registered trademark of Lucent Technologies.

**Link** A connection between two Fibre Channel ports consisting of a transmit fibre and a receive fibre.

**Link_Control_Facility** A termination card that handles the logical and physical control of the Fibre Channel link for each mode of use.

**LIP** A Loop Initialization Primitive sequence is a special Fibre Channel sequence that is used to start loop initialization. Allows ports to establish their port addresses.

**Local Area Network** (LAN) A network covering a relatively small geographic area (usually not larger than a floor or small building). Transmissions within a Local Area Network are mostly digital, carrying data among stations at rates usually above one megabit/s.

**Login Server** Entity within the Fibre Channel fabric that receives and responds to login requests.

**Loop Circuit** A temporary point-to-point like path that allows bi-directional communications between loop-capable ports.

**Loop Topology** An interconnection structure in which each point has physical links to two neighbors resulting in a closed circuit. In a loop topology, the available bandwidth is shared.

**LVD** Low Voltage Differential

**Management Agent** A process that exchanges a managed node's information with a management station.

**Managed Node** A managed node is a computer, a storage system, a gateway, a media device such as a switch or hub, a control instrument, a software product such as an operating system or an accounting package, or a machine on a factory floor, such as a robot.

**Managed Object** A variable of a managed node. This variable contains one piece of information about the node. Each node can have several objects.

**Management Station** A host system that runs the management software.

**MAR** Media Access Rules. Enable systems to self-configure themselves is a SAN environment

**Mb/s** Megabits per second. Also sometimes referred to as Mbps. In computing terms it is approximately 1,000,000 bits per second. Most precisely it is 1,048,576 (1024 x 1024) bits per second.

**MB/s** Megabytes per second. Also sometimes referred to as MBps. In computing terms it is approximately 1,000,000 bytes per second. Most precisely it is 1,048,576 (1024 x 1024) bytes per second.

**Metadata server** In Storage Tank™, these are servers that maintain information (“metadata”) about the data files and grant permission for application servers to communicate directly with disk systems.

**Meter** 39.37 inches, or just slightly larger than a yard (36 inches)

**Media** Plural of medium. The physical environment through which transmission signals pass. Common media include copper and fiber optic cable.

**Media Access Rules (MAR)**.

**MIA** Media Interface Adapter - MIAs enable optic-based adapters to interface to copper-based devices, including adapters, hubs, and switches.

**MIB** Management Information Block - A formal description of a set of network objects that can be managed using the Simple Network Management Protocol (SNMP). The format of the MIB is defined as part of SNMP and is a hierarchical structure of information relevant to a specific device, defined in object oriented terminology as a collection of objects, relations, and operations among objects.

**Mirroring** The process of writing data to two separate physical devices simultaneously.

**MM** Multi-Mode - See Multi-Mode Fiber

**MMF** See Multi-Mode Fiber - - In optical fiber technology, an optical fiber that is designed to carry multiple light rays or modes concurrently, each at a slightly different reflection angle within the optical core. Multi-Mode fiber transmission is used for relatively short distances because the modes tend to disperse over longer distances. See also: Single-Mode Fiber, SMF

**Multicast** Sending a copy of the same transmission from a single source device to multiple destination devices on a fabric. This includes sending to all N_Ports on a fabric (broadcast) or to only a subset of the N_Ports on a fabric (multicast).

**Multi-Mode Fiber** (MMF) In optical fiber technology, an optical fiber that is designed to carry multiple light rays or modes concurrently, each at a slightly different reflection angle within the optical core. Multi-Mode fiber transmission is used for relatively short distances because the modes tend to disperse over longer distances. See also: Single-Mode Fiber

**Multiplex** The ability to intersperse data from multiple sources and destinations onto a single transmission medium. Refers to delivering a single transmission to multiple destination Node Ports (N_Ports).

**N_Port** Node Port - A Fibre Channel-defined hardware entity at the end of a link which provides the mechanisms necessary to transport information units to or from another node.

**N_Port Login** N_Port Login (PLOGI) allows two N_Ports to establish a session and exchange identities and service parameters. It is performed following completion of the fabric login process and prior to the FC-4 level operations with the destination port. N_Port Login may be either explicit or implicit.

**Name Server** Provides translation from a given node name to one or more associated N_Port identifiers.

**NAS** Network Attached Storage - a term used to describe a technology where an integrated storage system is attached to a messaging network that uses common communications protocols, such as TCP/IP.

**NDMP** Network Data Management Protocol

**Network** An aggregation of interconnected nodes, workstations, file servers, and/or peripherals, with its own protocol that supports interaction.

**Network Topology** Physical arrangement of nodes and interconnecting communications links in networks based on application requirements and geographical distribution of users.

**NFS** Network File System - A distributed file system in UNIX developed by Sun Microsystems which allows a set of computers to cooperatively access each other's files in a transparent manner.

**NL_Port** Node Loop Port - a node port that supports Arbitrated Loop devices.

**NMS** Network Management System - A system responsible for managing at least part of a network. NMSs communicate with agents to help keep track of network statistics and resources.

**Node** An entity with one or more N_Ports or NL_Ports.

**Non-Blocking** A term used to indicate that the capabilities of a switch are such that the total number of available transmission paths is equal to the number of ports. Therefore, all ports can have simultaneous access through the switch.

**Non-L_Port** A Node or Fabric port that is not capable of performing the Arbitrated Loop functions and protocols. N_Ports and F_Ports are not loop-capable ports.

**Operation** A term defined in FC-2 that refers to one of the Fibre Channel *building blocks* composed of one or more, possibly concurrent, exchanges.

**Optical Disk** A storage device that is written and **read by laser light.**

**Optical Fiber** A medium and the technology associated with the transmission of information as light pulses along a glass or plastic wire or fiber.

**Ordered Set** A Fibre Channel term referring to four 10 -bit characters (a combination of data and special characters) providing low-level link functions, such as frame demarcation and signaling between two ends of a link.

**Originator** A Fibre Channel term referring to the initiating device.

**Out of Band Signaling** This is signaling that is separated from the channel carrying the information. Also referred to as out-of-band.

**Out-of-band virtualization** An alternative type of virtualization in which servers communicate directly with disk systems under control of a virtualization function that is not involved in the data transfer.

**Peripheral** Any computer device that is not part of the essential computer (the processor, memory and data paths) but is situated relatively close by. A near synonym is input/output (I/O) device.

**Petard** A device that is small and sometimes explosive.

**PLDA** Private Loop Direct Attach - A technical report which defines a subset of the relevant

standards suitable for the operation of peripheral devices such as disks and tapes on a private loop.

**PLOGI** See N_Port Login

**Point-to-Point Topology** An interconnection structure in which each point has physical links to only one neighbor resulting in a closed circuit. In point-to-point topology, the available bandwidth is dedicated.

**Policy-based management** Management of data on the basis of business policies (for example, “all production database data must be backed up every day”), rather than technological considerations (for example, “all data stored on this disk system is protected by remote copy”).

**Port** The hardware entity within a node that performs data communications over the Fibre Channel.

**Port Bypass Circuit** A circuit used in hubs and disk enclosures to automatically open or close the loop to add or remove nodes on the loop.

**Private NL_Port** An NL_Port which does not attempt login with the fabric and only communicates with other NL Ports on the same loop.

**Protocol** A data transmission convention encompassing timing, control, formatting and data representation.

**Public NL_Port** An NL_Port that attempts login with the fabric and can observe the rules of either public or private loop behavior. A public NL_Port may communicate with both private and public NL_Ports.

**Quality of Service** (QoS) A set of communications characteristics required by an application. Each QoS defines a specific transmission priority, level of route reliability, and security level.

**Quick Loop** is a unique Fibre Channel topology that combines arbitrated loop and fabric topologies. It is an optional licensed product that allows arbitrated loops with private devices to be attached to a fabric.

**RAID** Redundant Array of Inexpensive or Independent Disks. A method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

**Raid 0** Level 0 RAID support - Striping, no redundancy

**Raid 1** Level 1 RAID support - mirroring, complete redundancy

**Raid 5** Level 5 RAID support, Striping with parity

**Repeater** A device that receives a signal on an electromagnetic or optical transmission medium, amplifies the signal, and then retransmits it along the next leg of the medium.

**Responder** A Fibre Channel term referring to the answering device.

**Router** (1) A device that can decide which of several paths network traffic will follow based on some optimal metric. Routers forward packets from one network to another based on network-layer information. (2) A dedicated computer hardware and/or software package which manages the connection between two or more networks. See also: Bridge, Bridge/Router

**SAF-TE** SCSI Accessed Fault-Tolerant Enclosures

**SAN** A Storage Area Network (SAN) is a dedicated, centrally managed, secure information infrastructure, which enables any-to-any interconnection of servers and storage systems.

**SAN** System Area Network - term originally used to describe a particular symmetric multiprocessing (SMP) architecture in which a switched interconnect is used in place of a shared bus. Server Area Network - refers to a switched interconnect between multiple SMPs.

**SANSymphony** In-band block-level virtualization software made by DataCore Software Corporation and resold by IBM.

**SC Connector** A fiber optic connector standardized by ANSI TIA/EIA-568A for use in structured wiring installations.

**Scalability** The ability of a computer application or product (hardware or software) to continue to function well as it (or its context) is changed in size or volume. For example, the ability to retain performance levels when adding additional processors, memory and/or storage.

**SCSI** Small Computer System Interface - A set of evolving ANSI standard electronic interfaces that allow personal computers to communicate with peripheral hardware such as disk drives, tape drives, CD_ROM drives, printers and scanners faster and more flexibly than previous interfaces. The table below identifies the major characteristics of the different SCSI version.

| SCSI Version | Signal Rate MHz | BusWidth (bits) | Max. DTR (MB/s) | Max. Num. Devices | Max. Cable Length (m) |
|---|---|---|---|---|---|
| SCSI-1 | 5 | 8 | 5 | 7 | 6 |
| SCSI-2 | 5 | 8 | 5 | 7 | 6 |
| Wide SCSI-2 | 5 | 16 | 10 | 15 | 6 |
| Fast SCSI-2 | 10 | 8 | 10 | 7 | 6 |
| Fast Wide SCSI-2 | 10 | 16 | 20 | 15 | 6 |
| Ultra SCSI | 20 | 8 | 20 | 7 | 1.5 |
| Ultra SCSI-2 | 20 | 16 | 40 | 7 | 12 |
| Ultra2 LVD SCSI | 40 | 16 | 80 | 15 | 12 |

**SCSI-3** SCSI-3 consists of a set of primary commands and additional specialized command sets to meet the needs of specific device types. The SCSI-3 command sets are used not only for the SCSI-3 parallel interface but for additional parallel and serial protocols, including Fibre Channel, Serial Bus Protocol (used with IEEE 1394 Firewire physical protocol) and the Serial Storage Protocol (SSP).

**SCSI-FCP** The term used to refer to the ANSI Fibre Channel Protocol for SCSI document (X3.269-199x) that describes the FC-4 protocol mappings and the definition of how the SCSI protocol and command set are transported using a Fibre Channel interface.

**Sequence** A series of frames strung together in numbered order which can be transmitted over a Fibre Channel connection as a single operation. See also: Exchange

**SERDES** Serializer Deserializer

**Server** A computer which is dedicated to one task.

**SES** SCSI Enclosure Services - ANSI SCSI-3 proposal that defines a command set for soliciting basic device status (temperature, fan speed, power supply status, etc.) from a storage enclosures.

**Single-Mode Fiber** In optical fiber technology, an optical fiber that is designed for the transmission of a single ray or mode of light as a carrier. It is a single light path used for long-distance signal transmission. See also: Multi-Mode Fiber

**SMART** Self Monitoring and Reporting Technology

**SM** Single Mode - See Single-Mode Fiber

**SMF** Single-Mode Fiber - In optical fiber technology, an optical fiber that is designed for the transmission of a single ray or mode of light as a carrier. It is a single light path used for long-distance signal transmission. See also: MMF

**SNIA** Storage Networking Industry Association. A non-profit organization comprised of more than 77 companies and individuals in the storage industry.

**SN** Storage Network. See also: SAN

**SNMP** Simple Network Management Protocol - The Internet network management protocol which provides a means to monitor and set network configuration and run-time parameters.

**SNMWG** Storage Network Management Working Group is chartered to identify, define and support open standards needed to address the increased management requirements imposed by storage area network environments.

**SSA** Serial Storage Architecture - A high speed serial loop-based interface developed as a high speed point-to-point connection for peripherals, particularly high speed storage arrays, RAID and CD-ROM storage by IBM.

**Star** The physical configuration used with hubs in which each user is connected by communications links radiating out of a central hub that handles all communications.

**Storage Tank** An IBM file aggregation project that enables a pool of storage, and even individual files, to be shared by servers of different types. In this way, Storage Tank can greatly improve storage utilization and enables data sharing.

**StorWatch Expert** These are StorWatch applications that employ a 3 tiered architecture that includes a management interface, a StorWatch manager and agents that run on the storage resource(s) being managed. Expert products employ a StorWatch data base that can be used for saving key management data (e.g. capacity or performance metrics). Expert products use the agents as well as analysis of storage data saved in the data base to perform higher value functions including -- reporting of capacity, performance, etc. over time (trends), configuration of multiple devices based on policies, monitoring of capacity and performance, automated responses to events or conditions, and storage related data mining.

**StorWatch Specialist** A StorWatch interface for managing an individual fibre Channel device or a limited number of like devices (that can be viewed as a single group). StorWatch specialists typically provide simple, point-in-time management functions such as configuration, reporting on asset and status information, simple device and event monitoring, and perhaps some service utilities.

**Striping** A method for achieving higher bandwidth using multiple N_Ports in parallel to transmit a single information unit across multiple levels.

**STP** Shielded Twisted Pair

**Storage Media** The physical device itself, onto which data is recorded. Magnetic tape, optical disks, floppy disks are all storage media.

**Switch** A component with multiple entry/exit points (ports) that provides dynamic connection between any two of these points.

**Switch Topology** An interconnection structure in which any entry point can be dynamically connected to any exit point. In a switch topology, the available bandwidth is scalable.

**T11** A technical committee of the National Committee for Information Technology Standards, titled T11 I/O Interfaces. It is tasked with developing standards for moving data in and out of computers.

**Tape Backup** Making magnetic tape copies of hard disk and optical disc files for disaster recovery.

**Tape Pooling** A SAN solution in which tape resources are pooled and shared across multiple hosts rather than being dedicated to a specific host.

**TCP** Transmission Control Protocol - a reliable, full duplex, connection-oriented end-to-end transport protocol running on top of IP.

**TCP/IP** Transmission Control Protocol/ Internet Protocol - a set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks.

**Time Server** A Fibre Channel-defined service function that allows for the management of all timers used within a Fibre Channel system.

**Topology** An interconnection scheme that allows multiple Fibre Channel ports to communicate. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.

**T_Port** An ISL port more commonly known as an E_Port, referred to as a Trunk port and used by CNT.

**TL_Port** A private to public bridging of switches or directors, referred to as Translative Loop.

**Twinax** A transmission media (cable) consisting of two insulated central conducting leads of coaxial cable.

**Twisted Pair** A transmission media (cable) consisting of two insulated copper wires twisted around each other to reduce the induction (thus interference) from one wire to another. The twists, or lays, are varied in length to reduce the potential for signal interference between pairs. Several sets of twisted pair wires may be enclosed in a single cable. This is the most common type of transmission media.

**ULP** Upper Level Protocols

**UTC** Under-The-Covers, a term used to characterize a subsystem in which a small number of hard drives are mounted inside a higher function unit. The power and cooling are obtained from the system unit. Connection is by parallel copper ribbon cable or pluggable backplane, using IDE or SCSI protocols.

**UTP** Unshielded Twisted Pair

**Virtual Circuit** A unidirectional path between two communicating N_Ports that permits fractional bandwidth.

**Virtualization** An abstraction of storage where the representation of a storage unit to the operating system and applications on a server is divorced from the actual physical storage where the information is contained.

**Virtualization engine** Dedicated hardware and software that is used to implement virtualization.

**WAN** Wide Area Network - A network which encompasses inter-connectivity between devices over a wide geographic area. A wide area network may be privately owned or rented, but the term usually connotes the inclusion of public (shared) networks.

**WDM** Wave® Division Multiplexing - A technology that puts data from different sources together on an optical fiber, with each signal carried on its own separate light wavelength. Using WDM, up to 80 (and theoretically more) separate wavelengths or channels of data can be multiplexed into a stream of light transmitted on a single optical fiber.

**WEBM** Web-Based Enterprise Management - A consortium working on the development of a series of standards to enable active management and monitoring of network-based elements.

**Zoning** In Fibre Channel environments, the grouping together of multiple ports to form a virtual private storage network. Ports that are members of a group or zone can communicate with each other but are isolated from ports in other zones.

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

## IBM Redbooks

- *IBM SAN Survival Guide*, SG24-6143
- *IBM SAN Survival Guide Featuring the IBM 2109*, SG24-6127
- *IBM SAN Survival Guide Featuring the McDATA Portfolio*, SG24-6149
- *IBM SAN Survival Guide Featuring the INRANGE Portfolio*, SG24-6150
- *Designing and Optimizing an IBM Storage Area Network*, SG24-6419
- *Designing and Optimizing an IBM Storage Area Network Featuring the IBM 2109 and 3534*, SG24-6426
- *Designing and Optimizing an IBM Storage Area Network Featuring the INRANGE Portfolio*, SG24-6427
- *Designing and Optimizing an IBM Storage Area Network Featuring the McDATA Portfolio*, SG24-6428
- *Designing an IBM Storage Area Network*, SG24-5758
- *Introduction to SAN Distance Solutions*, SG24-6408
- *Introducing Hosts to the SAN fabric*, SG24-6411
- *Implementing an Open IBM SAN,* SG24-6116
- *Implementing an Open IBM SAN Featuring the IBM 2109, 3534-1RU, 2103-H07*, SG24-6412
- *Implementing an Open IBM SAN Featuring the INRANGE Portfolio*, SG24-6413
- *Implementing an Open IBM SAN Featuring the McDATA Portfolio*, SG24-6414
- *Introduction to Storage Area Network, SAN*, SG24-5470
- *IP Storage Networking: IBM NAS and iSCSI Solutions*, SG24-6240
- *The IBM TotalStorage NAS 200 and 300 Integration Guide*, SG24-6505
- *Implementing the IBM TotalStorage NAS 300G: High Speed Cross Platform Storage and Tivoli SANergy!*, SG24-6278

- *iSCSI Performance Testing & Tuning*, SG24-6531
- *Using iSCSI Solutions' Planning and Implementation*, SG24-6291
- *Storage Networking Virtualization: What's it all about?*, SG24-6210
- *IBM Storage Solutions for Server Consolidation*, SG24-5355
- *Implementing the Enterprise Storage Server in Your Environment*, SG24-5420
- *Implementing Linux with IBM Disk Storage*, SG24-6261
- *Storage Area Networks: Tape Future In Fabrics*, SG24-5474
- *IBM Enterprise Storage Server*, SG24-5465

## Other resources

These publications are also relevant as further information sources:

- *Building Storage Networks*, ISBN 0072120509

These IBM publications are also relevant as further information sources:

- *ESS Web Interface User's Guide for ESS Specialist and ESS Copy Services*, SC26-7346
- *IBM Storage Area Network Data Gateway Installation and User's Guide*, SC26-7304
- *IBM Enterprise Storage Server Configuration Planner*, SC26-7353
- *IBM Enterprise Storage Server Quick Configuration Guide*, SC26-7354
- *IBM SAN Fibre Channel Managed Hub 3534 Service Guide,* SY27-7616
- *IBM Enterprise Storage Server Introduction and Planning Guide, 2105 Models E10, E20, F10 and F20*, GC26-7294
- *IBM Enterprise Storage Server User's Guide, 2105 Models E10, E20, F10 and F20*, SC26-7295
- *IBM Enterprise Storage Server Host Systems Attachment Guide, 2105 Models E10, E20, F10 and F20*, SC26-7296
- *IBM Enterprise Storage Server SCSI Command Reference, 2105 Models E10, E20, F10 and F20*, SC26-7297
- *IBM Enterprise Storage Server System/390 Command Reference, 2105 Models E10, E20, F10 and F20*, SC26-7298
- *IBM Storage Solutions Safety Notices*, GC26-7229
- *Brocade Secure Fabric User's Guide*, 53-0000526
- *PCI Adapter Placement Reference*, SA38-0583

- *Translated External Devices/Safety Information*, SA26-7003
- *Electrical Safety for IBM Customer Engineers*, S229-8124
- *SLIC Router Installation and Users Guide*, 310-605759
- *SLIC Manager Installation and User Guide*, 310-605807
- *Cisco MDS 9000 Family Configuration Guide*, DOC-7814893
- *Cisco MDS 9000 Family Command Reference*, DOC-7814894
- *Cisco MDS 9000 Family Fabric Manager User Guide*, DOC-7814895
- *Cisco MDS 9500 Series Hardware Installation Guide*, DOC-7814900
- *Cisco MDS 9216 Hardware Installation Guide*, DOC-7814901
- *Cisco MDS 9000 Family Troubleshooting Guide*, OL-3450

## Referenced Web sites

These Web sites are also relevant as further information sources:

- IBM TotalStorage hardware, software and solutions:

  `http://www.storage.ibm.com`
- IBM TotalStorage Storage Networking:

  `http://www.storage.ibm.com/snetwork/index.html`
- Brocade:

  `http://www.brocade.com`
- CNT:

  `http://www.inrange.com`
- McDATA:

  `http://www.mcdata.com`
- QLogic:

  `http://www.qlogic.com`
- Emulex:

  `http://www.emulex.com`
- Finisar:

  `http://www.finisar.co`
- Veritas:

  `http://www.veritas.co`

- Vixel:

  `http://www.vixel.com`
- Tivoli:

  `http://www.tivoli.co`
- JNI:

  `http://www.Jni.com`
- IEEE:

  `http://www.ieee.org`
- Storage Networking Industry Association:

  `http://www.snia.org`
- Fibre Channel Industry Association:

  `http://www.fibrechannel.com`
- SCSI Trade Association:

  `http://www.scsita.org`
- Internet Engineering Task Force:

  `http://www.ietf.org`
- American National Standards Institute:

  `http://www.ansi.org`
- Technical Committee T10:

  `http://www.t10.org`
- Technical Committee T11:

  `http://www.t11.org`
- IBM @server™ xSeries® 430 and NUMA-Q Information Center:

  `http://webdocs.numaq.ibm.com`

# How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

**ibm.com**/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

## IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

# Index

## B

## C

## D

## E

## F

## G



## H



## I

## J



## K



## L

## M

## N

## O



## P

## Q



## R

## S

## T

## U



## V

## W



## X



## Y



## Z

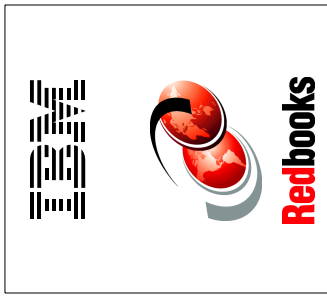

# Implementing an Open IBM SAN

(1.5” spine)
1.5”<-> 1.998”
789 <->1051 pages

# Implementing an Open IBM SAN


**Discover the latest additions to the IBM SAN family**

**Enhance your skills while using an easy-to-follow format**

**Grow with the new technology**

"Do everything that is necessary and absolutely nothing that is not."

In this IBM Redbook, which is an update and major revision of the previous version, we have tried to consolidate as much of the critical information as possible while covering procedures and tasks that are likely to be encountered on a daily basis.

Each of the products described has much, much more functionality than we could ever hope to cover in just one redbook. The IBM SAN portfolio is rich in quality products that bring a vast amount of technicality and vitality to the SAN world. Their inclusion and selection is based on a thorough understanding of the storage networking environment that positions IBM, and therefore its customers and partners, in an ideal position to take advantage by their deployment.

We cover the latest additions to the IBM SAN family, which includes products from companies such as Brocade, Cisco, CNT, and McDATA. We show how they can be implemented in an open systems environment, and we focus on the Fibre Channel protocol (FCP) environment in particular. We address some of the key concepts that they bring to the market, and in each case, we give an overview of those functions that are essential to building a robust SAN environment.

**INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

**BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**
**ibm.com**/redbooks

SG24-6116-03 ISBN 0738499951