



Brocade SilkWorm Design, Deployment, and Management Guide

SAN DDM

Version 2.0

Copyright © 2003, Brocade Communications Systems, Incorporated.

ALL RIGHTS RESERVED.

Publication Number: 53-0000366-01

Brocade, the Brocade B weave logo, Secure Fabric OS, and SilkWorm are registered trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. FICON is a registered trademark of IBM Corporation in the U.S. and other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners. Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability.

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

Notice: The product described by this document may contain “open source” software covered by the GNU General Public License or other open source license agreements. To find-out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Export of technical data contained in this document may require an export license from the United States Government.

Brocade Communications Systems, Incorporated

Corporate Headquarters

1745 Technology Drive
San Jose, CA 95110
T: (408) 487-8000
F: (408) 487-8101
Email: info@brocade.com

Asia-Pacific Headquarters

9/F, One International Finance
Center, 1 Harbour View Street,
Central, Hong Kong
Tel: +852 2539 0600
Fax: +852 2539 0613
Email: apac-info@brocade.com

European Headquarters

29, route de l' Aeroport
Case Postale 105
CH-1211 Geneva 15,
Switzerland
T: +41 22 799 56 40
F: +41 22 799 56 41
Email: europa-info@brocade.com

Latin America Headquarters

5201 Blue Lagoon Drive
Miami, FL 33126
T: (305) 716-4165
Email: latinam-sales@brocade.com

Document History

The table below lists all versions of the *Brocade SilkWorm Design, Deployment, and Management Guide*.

Document version	Publication Number	Publication Date
First Publication	53-0000366-01	4/30/2003

Table of Contents

Chapter 1 Introduction

Audience	1-2
Guideline Conventions	1-2
Formatting	1-2
Notes and Guidelines	1-3

Chapter 2 SAN Design

SAN Design Background	2-1
Core/Edge Topology	2-1
Redundant Fabric SANs	2-2
Trunking Considerations	2-4
ISL Oversubscription Ratios and Locality	2-5
ISL Oversubscription Ratio Calculations	2-5
Locality	2-7
Recommended ISL Oversubscription Ratios	2-9
Device Attachment Strategies	2-12
Trunk and ISL Connections	2-12
Edge Switch ISL/trunk Connections	2-14
Core ISL/trunk Connections	2-19
Attaching SAN Devices For Availability	2-21
Connecting For Scalability	2-22
Attaching Devices For Performance On SilkWorm 3900 and 12000 Switches	2-23
Platform Specific Design Considerations	2-27
SilkWorm 2000 Series Switches	2-27
SilkWorm 12000 and 3900	2-28
SilkWorm Switch Usage in a Core/Edge Topology	2-30

Zoning Design Considerations & Guidelines	2-35
Zoning and Scalability	2-35
Zoning Database Size	2-35
Designing SANs With Secure Fabric OS	2-36
Switch Location In The Fabric	2-38
Management and Control	2-38
2 Gbit/sec Switch Placement	2-40
Locating A Switch For Fabric and I/O Availability	2-41
Scalability Support and Testing	2-43
Recommended Fabric Topologies and SAN Designs	2-45

Chapter 3 **SAN Deployment**

Planning	3-2
Site Environment Assessment	3-2
SAN Project Checklist	3-3
Planning for Power	3-5
Cable Planning	3-6
Cable Management	3-7
The Rack Layout Plan	3-12
Documentation Guidelines	3-14
Zoning Plan	3-16
Planning the Upgrade of Fabric OS 2.x/3.x/4.x to 2.6.1/3.1/4.1	3-19
Planning Principal Switch Placement	3-20
Planning for Secure Fabric OS Security Measures	3-20
Secure Fabric OS Planning	3-21
Secure Fabric OS Pre-Installation Planning	3-24
Planning the Secure Fabric OS Implementation	3-27
SAN Secure Fabric OS Software Utility Considerations	3-29
LAN Planning Considerations	3-30
Brocade Extended Fabrics Planning	3-30

Staging	3-33
Fabric OS 2.6.1/3.1/4.1 Overview	3-33
Case Study	3-34
Powering the SAN Equipment	3-36
Preparing the Switches for the SAN Fabric	3-36
SAN Fabric Configuration	3-47
Validation	3-54
Sample Script	3-54
Sample Validation Recommendations	3-54
Maintenance and Operations	3-56
Executing the Upgrade of Fabric OS 2.x/3.x/4.x to 2.6.1/3.1/4.1	3-56
Fabric OS Version 4.1 Ultra High Availability (HA)	3-58
Guidelines for using other Fabric OS 2.6.1/3.1/4.1 Commands	3-60
Configuring Extended Fabrics	3-69
Using Fabric OS Troubleshooting Tools	3-70

Chapter 4 **SAN Management**

Event Management	4-1
Configuration	4-1
SNMP Features	4-9
Fabric Watch	4-10
Commands	4-10
Reading Fabric Watch Messages	4-14
Configuring	4-16
Third-Party Software	4-31
Fabric Maintenance With Web Tools, Fabric Manager, and via SNMP	4-33
Firmware Download	4-33
Zoning	4-37
Fabric Manager Sequenced Reboots	4-38
Periodical	4-41

Switch Performance	4-44
Advanced Performance Monitoring	4-44
APM Example.	4-44
APM and Fabric Watch.	4-47
Scripting possibilities with API	4-47

Appendix A Glossary

Terms and Definitions	A-1
------------------------------------	-----

Appendix B Reference Documentation

Brocade Documentation	B-1
Additional Resource Information	B-1

Introduction

The SilkWorm Design, Deployment, and Management (DDM) Guide is focused on covering detailed “how to” information for Brocade products from a design, deployment, and management perspective. The DDM is intended to be used in conjunction with existing Brocade manuals and release notes. The DDM is effective at focusing experienced Brocade SAN professionals on a specific area or subject of the SAN life cycle. The emphasis of the DDM is on breadth, with depth where new features or complexity currently exists. When multiple choices are available for adopting a particular approach or strategy and there are clear benefits to implementing a particular approach, guidelines are provided.

The SilkWorm 2000 series, 3200, 3800, 3900, and 12000 platforms and the available features of Fabric OS 4.1, 3.1, and 2.6.1 are covered in this document. The following Brocade software features are addressed:

- Zoning
- API
- Security
- Performance Monitoring
- Extended Fabrics
- Trunking
- Fabric Manager
- Fabric Watch
- Web Tools

The flow and organization of the document follows the process of first designing a Brocade SAN, followed by the deployment and operation of that SAN. It is important to understand how new features such as Secure Fabric OS, non-disruptive code activation, and scalability impact a SAN and how these topics relate across the SilkWorm and Fabric OS family. Many Brocade features span the disciplines of SAN design, deployment, and management. For example, ISL Trunking influences a SAN design, has specific deployment tips, and can be managed via various interfaces such as the CLI, Web Tools, and Fabric Manager.

Discussed in Chapter 2, *SAN Design*, are topics such as device attachment strategies, switch placement in a fabric, design related Security topics, and zoning guidelines. Once the SAN design and other planning has taken place, the switches and devices require deployment. Chapter 3, *SAN Deployment*, covers subjects such as installation preparation and planning, usage of new features, Security planning, migrating to a secure fabric, staging equipment, validating a fabric, and fabric troubleshooting. A rich set of management interfaces exists for the SilkWorm family of switches. Effectively integrating a particular management interface, such as Fabric Manager or Fabric Watch, into the enterprise management system, capacity planning, and SAN management with SNMP are just a few examples of topics addressed in Chapter 4, *SAN Management*.

FICON™ support is latent in Brocade Fabric OS 4.1. FICON will be supported for deployment after Brocade's OEM partners complete their qualification of Fabric OS 4.1 in their FICON environments. Note that FICON design requirements and practices differ significantly from those used for open systems Fibre Channel (FCP and SCSI) SANs. FICON considerations will be addressed in a future version of this document that will coincide with support for FICON deployments.

1.1. Audience

The DDM is targeted for use by storage administrators, SAN administrators, system administrators, SAN architects, systems engineers, and SAN operators that are involved with the design, deployment, and management of SANs. The DDM is an advanced document and is very concise. Background information and supporting information for a particular topic are kept to a minimum and as appropriate, the reader is referred to supporting documentation. The reader is expected to have working experience with Brocade products. General computer system level troubleshooting skills are always important when configuring complex enterprise solutions. System administration or storage administration experience is also helpful in comprehending this document.

Guidelines are provided throughout the document. Guidelines are recommendations for consideration. The adoption of these guidelines is a function of the user's ability to interpret and correlate relevant SAN information and make decisions based upon their organization and SAN requirements.

1.2. Guideline Conventions

The formatting and conventions used in this document are designed to help the reader locate and comprehend information quickly. In addition to the information provided in standard text, there are Guidelines, Notes, and Cautions to help focus the reader on important information.

1.2.1. Formatting

The following table describes the formatting conventions that are used in this book:

Convention	Purpose
bold text	<ul style="list-style-type: none"> identifies GUI elements identifies keywords/operands identifies menu selections at the GUI or CLI
<i>italic text</i>	<ul style="list-style-type: none"> provides emphasis identifies variables identifies paths and internet addresses identifies book titles and cross references
code text	<ul style="list-style-type: none"> identifies commands in line with text identifies CLI output identifies syntax examples

1.2.2. Notes and Guidelines

Note: Notes emphasize important information.

Guideline: Guidelines are recommendations for consideration. The adoption of these guidelines is a function of the user's ability to interpret and correlate relevant SAN information and make decisions based upon their organization and SAN requirements.

Warning: Warnings alert you to potential damage to hardware, firmware, software, or data.

The red circle with a slash through it (shown in Figure 1-1) indicates that a particular action or type of connection is not recommended. While the action or connection will function, there are better ways to perform the action or make the connection.




Figure 1-1

SAN Design

This chapter discusses several elements of a SAN that relate to SAN Design. The Core/Edge topology is identified as a reference topology. Then topics such as Trunking, ISL oversubscription ratios, locality, device attachment strategies, platform specific topics, zoning, Secure Fabric OS, switch location in the fabric, and finally recommended topologies are presented.

The guidelines established can also apply to other topologies, such as a full mesh or ring. It is up to the reader to interpret the guidelines for topologies other than Core/Edge. Many SAN related elements such as zoning, scalability, extended fabrics, and supportability have some bearing on the design of a SAN and are also discussed in this section.

The Core/Edge topology is preferred for scalable, available, and high performance fabrics for a number of reasons. Other topologies are feasible and supported using SilkWorm switches. Please reference the *Brocade SAN Design Guide* (publication number: 53-000231-05) for further detail supporting the selection of a fabric topology, understanding the language of SAN design (i.e. ISL oversubscription ratio) or creation of a SAN architecture. Regardless of topology chosen, a redundant fabric (i.e. dual fabric) SAN is always recommend.

2.1. SAN Design Background

The Core/Edge topology and the concept of redundant fabrics are fundamental to the area of SAN Design in general and specifically to this chapter. A brief description of the Core/Edge topology and redundant fabrics is provided to support the development and understanding of the guidelines presented. Again, other topologies such as mesh and ring are certainly acceptable and supported.

2.1.1. Core/Edge Topology

A Core/Edge topology is shown in Figure 2-1 and can be built with a variety of switch platforms, such as the SilkWorm 2000 series, 3200, 3800, 3900, and 12000. The type of switch in a fabric does have some bearing on the practical as well as supported size of a fabric. A SilkWorm 12000 is used as the core in Figure 2-1, as the 64-ports per logical switch that this chassis supports enables the size of the fabric to grow to several thousand ports by connecting edge switches. The edge can be built with a variety of switch platforms. Note that the SilkWorm 12000 can contain up to 128 ports in a 14U chassis, configured as two 64-port switches. Each switch is known as a logical switch and may also be referred to as a domain.

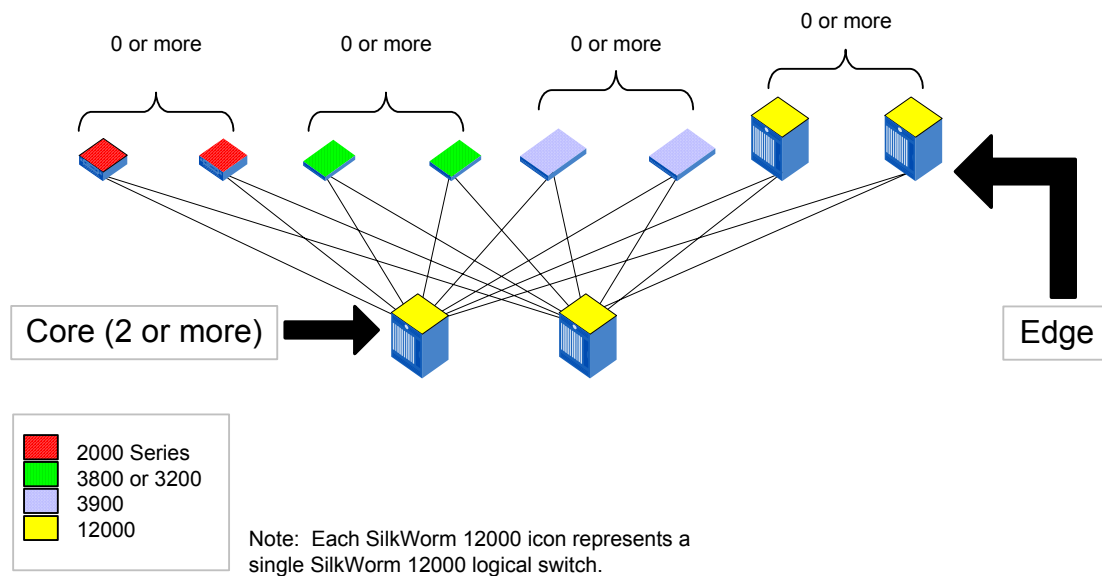


Figure 2-1 The Core/Edge topology

Some benefits of the Core/Edge topology:

- Well-tested and reliable
- Widely deployed in production environments
- Simple and easy to understand
- Able to solve most design problems, fits wells with many SAN solutions, and is an effective choice when design requirements are not well known
- Easy to grow without downtime or disconnection of links and devices
- Pay as you grow
- Flexible
- Capable of exhibiting stellar performance, with full utilization of FSPF load sharing and resiliency features
- Conducive to performance analysis. Because the Core/Edge topology is symmetrical, it is a straightforward process to identify performance issues. Every device has an equivalent path to any other device and the same available bandwidth between any two devices. To identify a performance issue it is only necessary to monitor the core switches. With other topologies, this is not the case.
- The potential to scale to thousands of ports (using high port count switches)

2.1.2. Redundant Fabric SANs

Resilient fabrics and the fault tolerant components that comprise them are very reliable. However, no single fabric can ever truly be a High Availability (HA) solution. The fabric itself is still potentially subject to failures caused by things like disaster, operator error, or software malfunctions. To account for those categories of error, another level of availability must be used: The redundant fabric SAN. This is sometimes known as a multi-fabric or dual-fabric SAN.

Redundancy in SAN design is the duplication of components up to and including the entire fabric to prevent the failure of the SAN solution. Even though an airplane navigation system (e.g. a GPS) is resilient to failures, most jumbo jets also have a redundant navigation system (e.g. a magnetic compass and a map) so that the jet will not get lost even if the resiliency fails to keep the primary navigation system up.

Using a fully redundant fabric makes it possible to have an entire fabric fail as a unit or be taken offline for maintenance without causing downtime for the attached nodes. When describing availability characteristics, what we are concerned with is *path* availability. If a particular link fails, but the path to the data is still there, no downtime is experienced by the users of the system. It is possible that a performance impact may occur, but this is a very small event compared to one or many crashed servers. Two or more fabrics must be used in conjunction with multiple HBAs, multiple RAID controllers, and path switch-over software to be effective for those SAN devices that require the highest availability. Figure 2-2 illustrates the ability of redundant fabrics to withstand large-scale failures. Note that tape drives can be effectively utilized in a redundant fabric environment as well – even if they are single attached. Please reference the SOLUTIONware *Enterprise LAN-Free Backup of Consolidated Storage on a SilkWorm 3800 Based Redundant 96-Port Integrated Fabric* (publication number 53-0000229-01) for a discussion regarding the effective use of tapes in a dual fabric environment.

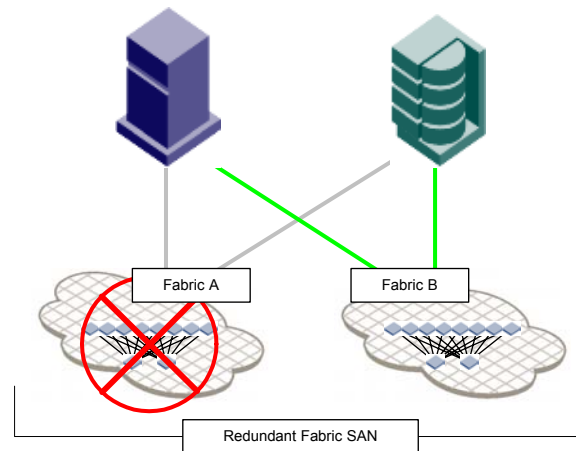


Figure 2-2 Failure of an Entire Fabric

In a redundant SAN architecture, there must be at least two *completely separate* fabrics – just as a high-availability server solution requires at least two completely separate servers. Duplicating components and providing switch-over software is well established as the most effective way to build HA systems. Similarly, multi-fabric SAN architectures are the best way to achieve HA in a SAN.

In addition to enhancing availability, using redundant fabrics also enhances scalability. Using dual fabrics essentially doubles the maximum size of a SAN. If a fabric is limited by vendor support levels to 34 switches / 1200 user ports and a single fabric solution with dual attach devices is utilized, then the SAN is limited to 1200 ports. Twelve hundred dual attach ports is equivalent to 600 devices. However, if a dual fabric with dual attach device solution is utilized, the SAN is capable of supporting 2400 ports or 1200 devices.

Any devices that are dual attached and are capable of supporting an active-active or active-passive dual-path essentially double the potential bandwidth. An active-active dual path means that I/O is capable of using both paths in normal operation. Some devices only support active-passive dual-pathing. With active-passive dual-pathing, the passive path is utilized only when the primary path fails.

Some devices, such as tape drives, are not currently capable of supporting multiple paths. It is possible to address this issue by equally distributing tape devices between the redundant fabrics and configuring the backup applications to use an alternate tape drive should an outage on one of the fabrics occur. However, some elements of a tape backup solution, such as Robot control, do not currently map well into a redundant fabric environment.

Any single attached devices, such as a tape drive, non-critical storage and hosts can be single-attached, by alternately assigning them between the fabrics. When implementing a logical group of single-attached devices, ensure that the devices that access the single-attached devices reside on the same fabric.

2.2. Trunking Considerations

Trunking is a feature that enables traffic to be optimally shared across available inter-switch links (ISLs) while preserving in-order delivery. A trunk group logically joins two, three, or four ISLs into one logical ISL. Use of trunking can minimize or eliminate congestion in the SAN because trunking optimizes ISL utilization. The use of trunking minimizes the effort of managing a SAN since ISLs are now managed as a group instead of individually. The use of trunking optimizes FSPF performance as FSPF does not have to compute as many routes.

Trunking can also increase availability. As long as at least one ISL link remains, I/O continues if an ISL failure occurs -- albeit at a lower bandwidth. It is also possible to dynamically increase bandwidth by adding ISLs to a trunk -- without impacting I/O -- to enable up to 8 Gbit/sec of bandwidth over a single logical link. Trunking is available on the SilkWorm 3200, 3800, 3900, and 12000 platforms. The ports that form a trunk must reside in the same contiguous four-port groups, which are known as quads, and are as shown in Figure 2-3 and Figure 2-4. For additional discussion about Trunking, reference *Exploring Brocade ISL Trunking* (publication number: 53-0000263-01). An octet is a group of two adjacent quads. The SilkWorm 3900 is the only SilkWorm switch that implements octets. Note that the SilkWorm 3900 octets are highlighted in red dotted boxes in Figure 2-3. The way that devices are connected to quads and octets can impact the performance of switches. The use of quads and octets to optimize performance is discussed in *Attaching SAN Devices For Availability on page 2-21*. Octets consist of two quads, each quad capable of supporting an 8 Gbit/sec trunk. An octet has no bearing on trunking and is primarily used as a mechanism to define boundaries to optimize performance

Note: With Fabric OS versions 3.1 and 4.1, Trunking is not supported with LE, L1, and L2 `portcfglongdistance` modes. Please reference the *Brocade Distributed Fabrics User's Guide* for further discussion on this topic.

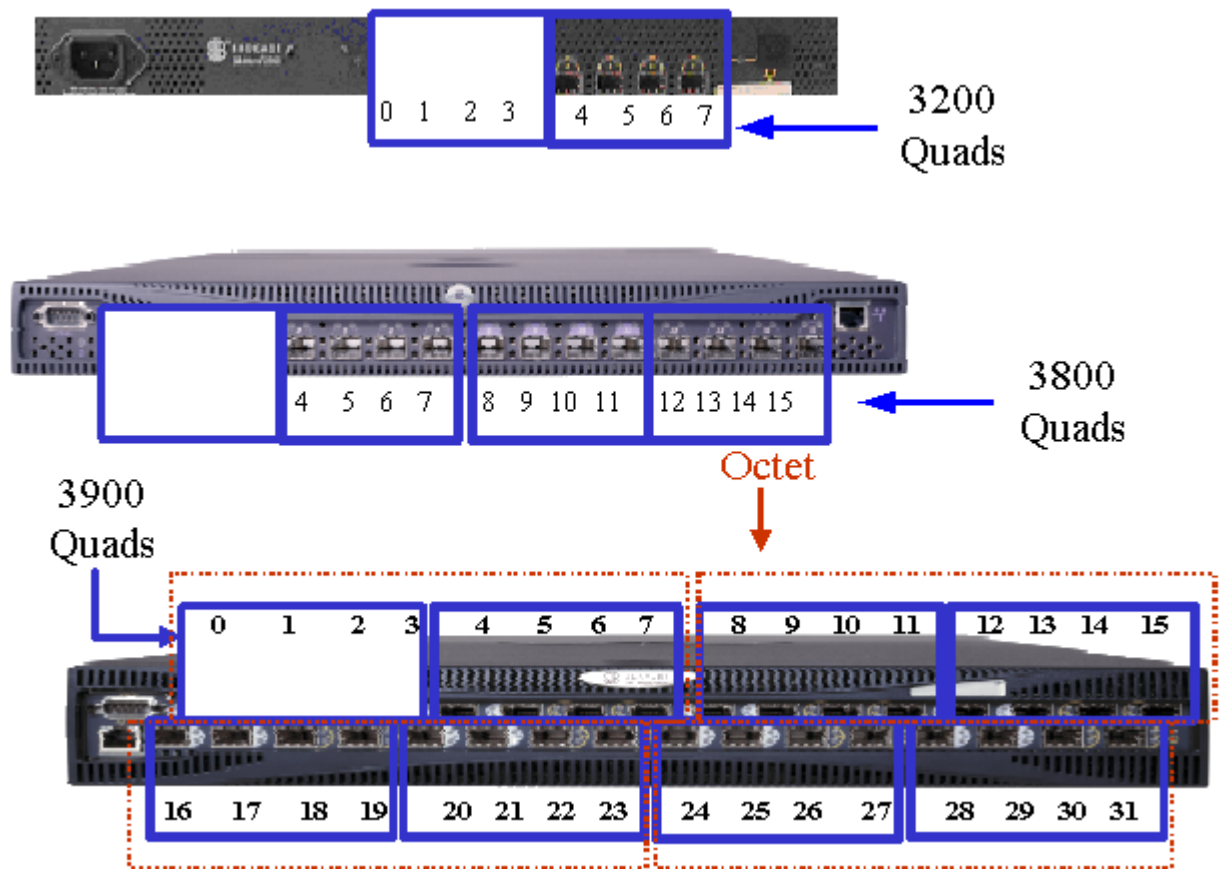


Figure 2-3 SilkWorm 3200, 3800, and 3900 Quads; SilkWorm 3900 Octets




Figure 2-4 SilkWorm 12000 Quads Used for Trunking

2.3. ISL Oversubscription Ratios and Locality

When designing a SAN, it is important to understand the performance boundaries such as storage fan-out ratios and storage performance. While any SAN device that connects to a SAN at 2 Gbit/sec is theoretically capable of 2 Gbit/sec, in reality, that device is most likely capable of a much lower performance. If this device truly is capable of generating 2 Gbit/sec of I/O, then the principles of locality should be applied or sufficient bandwidth should be provisioned for the ISLs. A very popular SAN application is storage consolidation, where many hosts share a storage device or port. Several popular storage vendors target an average of a 6:1 fan-out. This means that on average six hosts are sharing a single storage port. If there were 32 storage ports in a fabric, then one would expect to find an average of 192 hosts. Even if every host requires 1 Gbit/sec or 2 Gbit/sec of bandwidth, the storage devices in the fabric are only capable of delivering 32 Gbit/sec (1 Gbit/sec ports) or 64 Gbit/sec (2 Gbit/sec ports). This equates to 3-6 MB/sec per host. While some ports in the fabric may require maximal bandwidth, not all ports require sustained maximal bandwidth and rarely, if ever, do these ports require maximal bandwidth simultaneously.

2.3.1. ISL Oversubscription Ratio Calculations

When all ports operate at the same speed, ISL over-subscription is the ratio of device, or data input ports that might drive I/O between switches to the number of ISLs over which the traffic could cross. In Figure 2-5, the over-subscription ratio on the switch to the far left is three device ports to one ISL. This is usually abbreviated as 3:1. There are twelve hosts connected to the upper left edge switch and only four ISLs to the core. Thus, there are three hosts for each ISL. If all of these hosts tried to simultaneously use the ISLs at full speed in a sustained manner — even if the hosts were accessing different storage devices — each would receive only about one-third of the potential bandwidth available.

The basic over-subscription formula is “ISL Over-Subscription = Number of Nodes: Number of ISLs”, or $I_o = N_n : N_i$. This is reduced as a fraction so that $N_i = 1$.

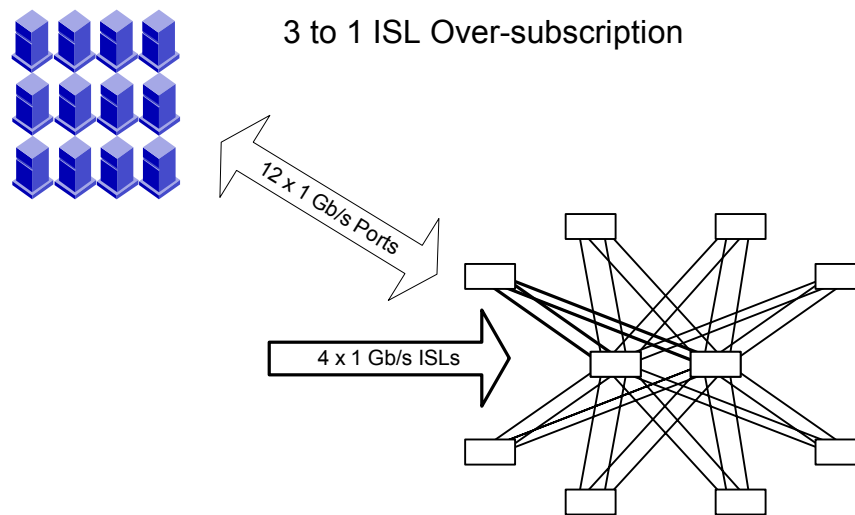


Figure 2-5 ISL over-subscription with 1 Gbit/sec devices

With the advent of 2 Gbit/sec devices today and higher speeds to follow, it is necessary to put some additional thought into calculating ISL over-subscription with variable speed hosts, storage, and ISLs. In [Figure 2-6 on page 2-7](#), six 1 Gbit/sec hosts and six 2 Gbit/sec hosts are depicted. These share access to four 2 Gbit/sec ISLs. To calculate the ISL over-subscription ratio, average the speed of the input ports and divide this result by the speed of the output ports. Multiply the node portion of the ratio by that number.

The mixed-speed over-subscription formula is “**ISL Over-Subscription: Number of ISLs**”

ISL Over-Subscription =

((Average of Node Speeds / aggregate ISL Speed) x Number of Nodes)

Number of ISLs or I_o =

$((Ans/Is)Nn):Ni$.

For Figure 2-6, the ISL over-subscription ratio is 2.25:1.

$Ans = ((6 * 1) + (6 * 2)) / 12 = 1.5$

$Is = 2; Nn = 12$

so $I_o = ((1.5 / 2) 12) : 4$, which reduces to 2.25:1

There are other ways to organize the formula, as shown in Figure 2-6. In that calculation, Nn is reduced out of the formula.

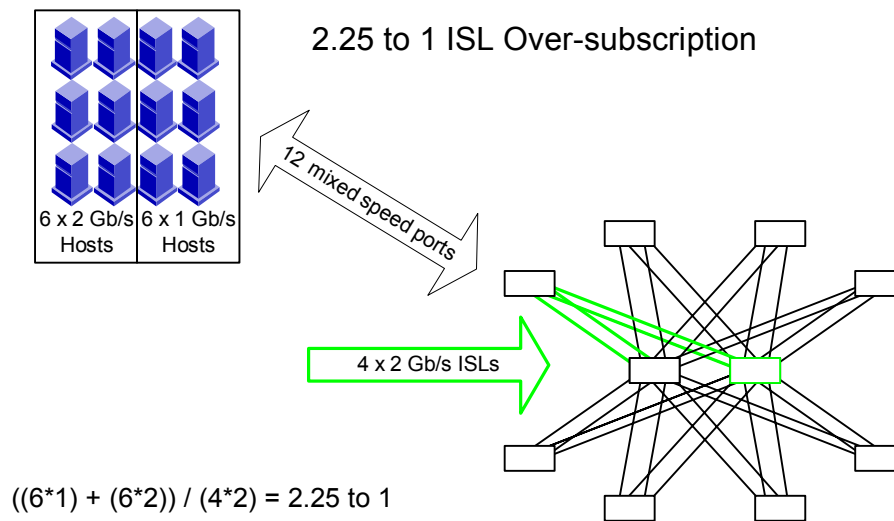


Figure 2-6 ISL Over-Subscription with Mixed-Speed Devices

2.3.2. Locality

If devices that communicate with each other are connected to the same switch or groups of switches then these devices have high locality. If two devices must cross an ISL/Trunk to communicate, then these devices have low locality. The higher the locality, the less traffic crosses ISLs/trunks and therefore, fewer ISLs/trunks are needed. The lower the locality, the more traffic crosses ISLs/trunks and therefore, more ISLs/trunks are needed.

Figure 2-7 depicts the scenario of low locality. When host and storage devices need to communicate in low locality scenarios, all traffic must traverse through ISLs/Trunks. If four 2 Gbit/sec hosts in the Figure 2-7 need to concurrently communicate with four 2 Gbit/sec storage devices/connection at full bandwidth, congestion occurs in the ISLs. This is because eight devices (four hosts, four storage devices) that could potentially generate 1600 MB/sec of I/O, must share only 800 MB/sec of bandwidth. Of course, in reality, most devices cannot sustain full throughput and they would not all peak at the same time. Note that this example involves simplex I/O, meaning only one path is utilized for I/O. One example of a simplex operation is a scenario where all hosts are reading from the storage 100%. Under full duplex operations, the maximum potential bandwidth is 3200 MB/sec. This is why many hosts can share a single storage port, and why many devices can share a single ISL. If all eight devices were connected to the same switch, they could communicate with each other at a potential aggregate bandwidth of 1600 MB/sec without congestion. When a single switch is not large enough to support the number of required devices, a network of switches is needed.

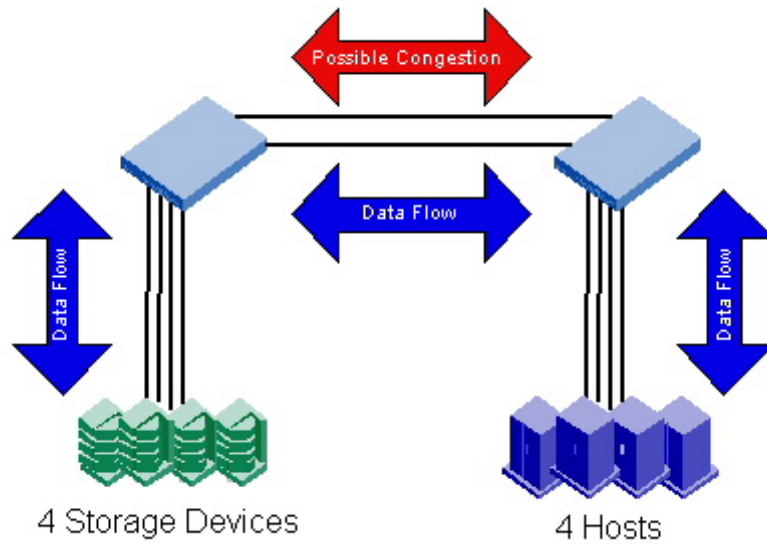


Figure 2-7 Low Locality SAN

With a little planning, it is usually possible to design a SAN with a significant degree of locality, as shown in Figure 2-8. While higher levels of locality are desirable, it is still possible to build very effective SANs with minimal to no locality. In fact, some SANs are deliberately designed with low locality to maximize the administrative simplicity that a low locality design provides. It is a straightforward process to design a tiered SAN that delivers sufficient bandwidth in a low locality environment. The value in doing so is that tiered SANs require minimal planning and management to add hosts or storage - just attach hosts to host-designated switches and storage to storage-designated switches.

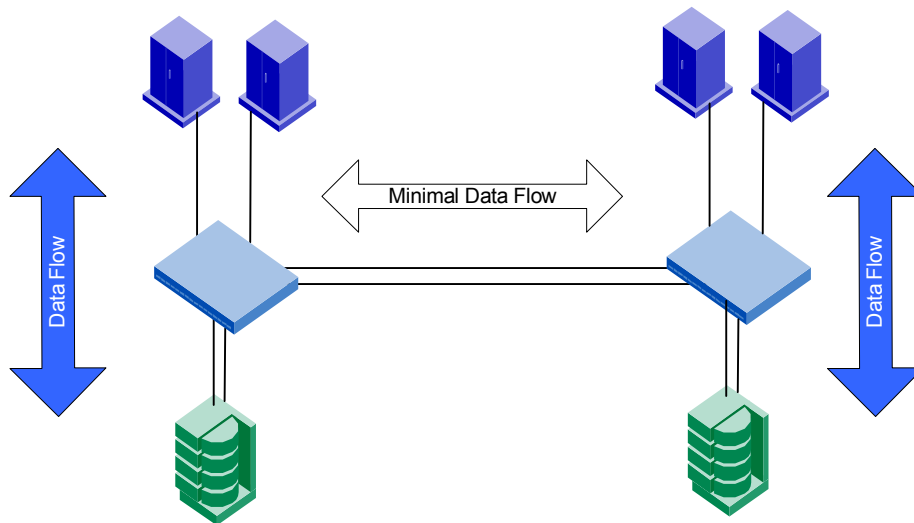


Figure 2-8 High Locality SAN

2.3.2.1. Tiering

Tiering is the process of grouping particular devices by function and then attaching these devices to particular switches or groups of switches based on that function. Tiering is the opposite of high locality: in a highly localized SAN, hosts are attached to the same switches as their storage devices; in a tiered SAN, hosts are rarely attached to the same switches as storage arrays.

It requires some level of effort to plan and manage the layout of a fabric for optimal locality. Sometimes this effort is not necessary if there is a sufficient level of available ISL/trunk bandwidth. For example, if it is known that the peak bandwidth that a host generates is 10 MB/sec and there are fourteen hosts on a switch, it is sufficient to only have one ISL (2 Gbit/sec) connecting that switch to the remainder of the fabric and tiering is a viable design option. However, if those hosts generate 50 MB/sec concurrently, it is probably more appropriate to adopt a device attachment strategy that involves a high degree of locality, or to use more ISLs.

From a cabling and maintenance perspective, tiering is quite effective. In Figure 2-9, a group of switches is designated as the storage switch group, another group designated as the tape group, and a final group is designated as the host group. When it becomes necessary to expand backup, storage, or hosts, it becomes a straightforward effort to attach the new devices to an open port on the appropriate tier and to then enable access (i.e. zoning, configure hosts). If a particular tier requires expansion, add a new switch to that group.




Figure 2-9 A Tiered SAN

2.3.3. Recommended ISL Oversubscription Ratios

ISL oversubscription ratios apply in practice to Core/Edge fabrics. The calculations for ISL oversubscription ratios for a Core/Edge fabric are simple and straightforward, while these same calculations become more complex and less pertinent for other topologies, such as a ring topology. ISL oversubscription ratios principally apply to edge switches, as the role of a core switch is to connect other switches. Connecting devices to a core switch is supported and makes sense for particular scenarios, such as when there are excess ports available on the core switch. When devices are connected to the core switch, the number of ISLs/trunks is usually equal to or greater than the number of devices and the devices actually are undersubscribed or at a 1:1 ISL oversubscription ratio – minimizing the value of this metric for core switches. The ISL oversubscription ratio does become more meaningful for a core switch when devices are connected to the core and there are more devices than ISLs/trunks.

A 7:1 ISL over subscription ratio is aligned with the de facto industry average of 6:1 fan-out. The trend in the storage industry is that the hosts to storage ratios are increasing, as is the performance of storage devices. A 7:1 ISL over subscription ratio should be targeted in SAN designs, with the ISL over subscription ratio being adjusted higher or lower to meet particular performance requirements. While this ISL oversubscription ratio is conservative, it is felt that the pain of not having enough performance and having to reshuffle devices and ISLs is much greater than the pain of having a few extra spare ports that can be used to connect SAN devices in the future if the ISL /trunk experienced usage does not justify such a 7:1 ISL oversubscription ratio. Note that if the SAN devices connected are 1 Gbit/sec devices, the ISL over subscription ratio decreases since the lower bandwidth 1 Gbit/sec SAN devices are now aggregated across 2 Gbit/sec ISLs and 4-8 Gbit/sec trunks. This means that a 7:1 ISL over subscription ratio drops to 3.5:1 and a 3:1 ISL over subscription ratio drops to 1.5:1. The higher the ISL oversubscription ratio, the lower the performance and conversely, the lower the ISL oversubscription ratio, the higher the performance. An ISL oversubscription ratio of 3:1 results in high performance and fewer available ports while an ISL oversubscription ratio of 15:1 results in lower potential performance and more available user ports. The practical boundaries for ISL oversubscription ratios is 3:1 for high performance SANs and 15:1 where lower performance is sufficient. The factors that influence this position include resiliency, number of ports available on a SilkWorm platform, and industry host to storage fan-out ratio. Other ISL oversubscription ratios outside the practical boundaries, such as 1:1 or 31:1 are supported and make sense for specific requirements such as any-to-any high performance connectivity (1:1) or high locality SANs (30:1). Intermediate ISL oversubscription ratios that fall within the practical boundaries are also valid, such as 4.3:1, which can be realized on a 16-port switch with three ISLs, a 32-port switch with six ISLs, or 64- port switch with twelve ISLs.

Guideline: As a starting point or if specific performance requirements are not available, it is suggested to connect at a 7:1 ISL oversubscription ratio and reserve enough ISL ports to achieve a 3:1 ISL oversubscription ratio.

Table 2-1 identifies the number of ports required per platform to achieve a targeted ISL oversubscription ratio and the recommended number of ports to reserve to scale performance. Note that for 8-port switches, the lowest practical ISL oversubscription that maintains resiliency (i.e. more than one connection into the fabric) is 3:1 and two ports are necessary. For the larger switches like the SilkWorm 3900 or 12000, higher ISL oversubscription ratios are possible while maintaining resiliency. With higher ISL oversubscription ratios, it is important to utilize low bandwidth devices or to employ locality. When a particular ISL oversubscription ratio is not recommended for resiliency purposes, the table cell is labeled with N/A (not applicable).

Guideline: When designing a Core/Edge fabric, attach edge switch ISLs for the target ISL oversubscription ratio, but reserve ports to enable the scaling of *performance*. For example, connect ISLs at a 7:1 ISL oversubscription ratio, but reserve ports in the same quad for a 3:1 ISL oversubscription ratio.

Table 2-1 ISL Oversubscription Ratio Port Recommendations for SilkWorm Platforms Used on the Edge in a Core/Edge Fabric

SilkWorm Plat- form	Number Ports Per Switch	3:1 Ports		7:1 Ports		15:1 Ports	
		Required	Reserved	Required	Reserved	Required	Reserved
8-port 2000 Series	8	2	0	N/A	N/A	N/A	N/A
3200	8	2	0	N/A	N/A	N/A	N/A
16-port 2000 Series	16	4	0	2	2	N/A	N/A
3800	16	4	0	2	2	N/A	N/A
3900	32	8	0	4	4	2	2
12000	64	16	0	8	8	4	4

Table 2-2 identifies the number of trunks required per platform to achieve a targeted ISL oversubscription ratio and the recommended number of trunks to reserve to scale performance. The trunk count is expressed as the number of trunks by the number of ISLs the make up that trunk:

#trunks X # ISLs in the trunk

If the table cell is labeled with N/A (not applicable), it is because it is not possible to build a resilient topology with trunks or the resulting configuration would yield an impractical ISL oversubscription ratio. For example, a SilkWorm 3900 configured in a fabric for a 15:1 ISL oversubscription ratio requires two ISLs; however, for this switch to be configured for resiliency, each ISL would need to attach to a different switch. While this configuration is resilient, trunking is not possible since there is only one ISL per switch-to-switch connection.

Table 2-2 Recommended Trunk Configurations to Attain Target ISL Oversubscription Ratios for Switches Used on the Edge in a Core/Edge Fabric

SilkWorm Switch	Number of Ports Per Switch	3:1 #trunks x # ISLs in the trunks		7:1 #trunks x # ISLs in the trunks		15:1 #trunks x # ISLs in the trunks	
		Required	Reserved	Required	Reserved	Required	Reserved
SilkWorm 3800	16	2x2	0	N/A	N/A	N/A	N/A
SilkWorm 3900	32	2x4	0	2x2	2x2	N/A	N/A
SilkWorm 12000	64	4x4	0	4x2	4x2	2x2	2x2

The information in Table 2-1 and Table 2-2 presents a detailed view of the exact number of ports and trunks to configure and reserve to attain a particular ISL oversubscription ratio for switches used on the edge in a Core/Edge topology. In Table 2-1, using the SilkWorm 12000 and a 7:1 ISL target connection rate as an example, it is necessary to connect two 2-ISL trunks and reserve four ports (two 2-ISL trunks worth) of ports to enable the scaling of performance.

Table 2-3 identifies recommended number of ports to reserve for ISLs/trunks on an edge switch when designing a Core/Edge fabric. The resulting configuration yields a 7:1 ISL oversubscription ratio that provisions for scaling to a 3:1 ISL oversubscription ratio should performance requirements dictate.

Table 2-3 Recommended ISL/Trunk Port Allocations for SilkWorm Switches Used as Edge Switches in a Core/Edge Fabric

SilkWorm Platform	Number Ports Per Edge Switch	Number Of Ports Recommended Per Edge Switch For ISLs/trunks
8-port 2000 Series	8	2
3200	8	2
16-port 2000 Series	16	4
3800	16	4
3900	32	8
12000	64	16

2.4. Device Attachment Strategies

How switches connect to other switches and how devices connect to those switches significantly influences the performance and availability of a SAN. Easy to understand and consistent device attachment strategies also simplifies the operation and maintenance of a SAN. This section details effective techniques for connecting devices, and ISLs/trunks to the SilkWorm 2000 series, 3200, 3800, 3900, and 12000 for availability, scalability, performance, and operational efficiency.

2.4.1. Trunk and ISL Connections

Guidelines for connecting ISLs/trunks, which take into consideration a switches architecture, availability, performance, and operational efficiency, are presented in this section.

Note: Previous versions of the DDM provide guidance to connect devices and ISLs/trunks horizontally across the SilkWorm 12000 switch port cards. While horizontal connections are still an acceptable method that offers simplicity, an additional ISL connection strategy, which involves connecting ISLs/trunks diagonally across SilkWorm 12000 switch port cards, is now recommended. Like the horizontal connection recommendation, diagonal ISL/trunk connections enable higher availability. Under certain conditions, the SilkWorm 12000 diagonal connection strategy for ISLs/trunks provides better performance. If a SilkWorm 12000 is already deployed with horizontally connected ISLs/trunk, it is unlikely that changes are required to optimize current implementations. Furthermore horizontal implementations are still valid and supported. If you are considering changing from a horizontal to diagonal strategy, please consult with your switch provider or Brocade field representative.

Connecting ISLs/trunks across switch port cards on the SilkWorm 12000 prevents a switch port card failure from segmenting that switch from the fabric when there is more than one ISL or trunk connecting that switch to the fabric and enables optimal performance. Connecting ISLs/trunks to diagonally opposed corners on the SilkWorm 3900 results in a configuration optimized for performance.

Note: The diagonal recommendations for both the SilkWorm 3900 and 12000 starts in the lower left hand corner of the switch (ports 0-4) and progress to the upper right hand corner (ports 12-15 for the SilkWorm 12000 and ports 28-31 for the SilkWorm 3900). This strategy is chosen for consistency sake between the switches. Connecting in the opposite order and starting in the upper left hand corner of the switch (ports 16-19 for the SilkWorm 3900 and ports 12-15 for the SilkWorm 12000) and progressing to the lower right hand corner (ports 0-3 for the SilkWorm 1200 and ports 12-15 for the SilkWorm 3900) is also an acceptable connection strategy.

Connecting ISLs/trunks to the left or right hand side of a SilkWorm series 2000, 3200, or 3800 switches does not yield any availability or performance benefits; however, doing so does result in a switch that is easier to operate since the cable locations are standardized. For all switches, the adoption of standardized ISL/trunk connections makes it easier to operate and maintain those switches. If it is possible to connect more than one ISL from an edge switch to a core switch and both switches are trunk capable, a choice exists to connect the ISLs to different quads for availability purposes or to use trunking. Under these circumstances, the performance benefits of trunking combined with the resiliency of the Core/Edge topology are felt to result in a simpler and superior solution. For example when connecting a SilkWorm 3800 edge switch to a SilkWorm 12000 core switch with a 3:1 ISL oversubscription ratio, connect using a 2-ISL trunk from the edge to each core instead of connecting two ISLs spread across separate blades on each core.




Figure 2-10 SilkWorm 12000 Horizontal ISL and Device Attachment

Guideline: Place trunking-capable switches adjacent to each other. This maximizes the number of trunking groups that can be created. If using a Core/Edge topology, place trunking-capable switches at the core of the fabric and switches that are not capable of trunking at the edge of the fabric. This allows for the maximum amount of trunking between core switches and edge switches that are capable of trunking.

Guideline: When more than one ISL connects two trunk capable switches, the option exists to spread these connections across different quads for availability purposes and not use trunking or to locate these ISLs on the same quad and to use trunking. Under these circumstances it is recommended to use trunking.

Guideline: If trunking is not possible, or more than one trunk/ISL is required between switches, connect trunks and ISLs from the same switch as follows:

- a) For the SilkWorm 12000, connect the ISLs/trunks diagonally across the quads. Do not connect more than one ISL to the same switch port card unless it is part of a trunk (see Figure 2-11, Figure 2-12, and Figure 2-13).
- b) For SilkWorm 3900 switches connect the ISLs/trunks on corners of the switch that are diagonally opposed (see Figure 2-14)
- c) For 8 and 16-port SilkWorm switches (SilkWorm 2000 series, 3200, 3800), connect the trunks/ISLs to the either the left or right hand side of the switch (for consistency) see Figure 2-15 and Figure 2-16.

Guideline: For the SilkWorm 12000 and 3900, create redundant trunking groups when possible. This protects against multiple ISL failures, optimizes performance, and for the SilkWorm 12000 protects against the rare occurrence of a blade failure. This means that instead of using a single 4-ISL trunk to connect two SilkWorm 12000s or 3900s, utilize two 2-ISL trunks, as shown in Figure 2-11.




Figure 2-11 Use redundant trunk groups when possible for the SilkWorm 3900 and 12000.

2.4.2. Edge Switch ISL/trunk Connections

The recommended ISL/trunk connection guidelines for edge switches that are part of a Core/Edge topology are shown in Figure 2-12, Figure 2-13, Figure 2-14, Figure 2-15, and Figure 2-16. Keeping in line with the recommended ISL oversubscription ratios discussed in *ISL Oversubscription Ratios and Locality* on page 2-5, the connection areas for ISLs/trunks are grouped into blocks of sixteen ports for the SilkWorm 12000, eight ports for the SilkWorm 3900, four ports for

16-port switches, and two ports for 8-port switches. The connection schemes for the SilkWorm 3900, 16-port, and 8-port switches locates the trunks on the left and right hand sides of each switch. While, more ISLs/trunks may be connected for certain configurations, the connection schemes recommended account for SilkWorm switches used as edge switches and provisioning for a of a 3:1 ISL oversubscription ratio.




Figure 2-12 Recommended ISL/Trunk Connection Scheme for SilkWorm 12000 Switches Used as Edges in a Core/Edge Fabric



Figure 2-13 Do not connect more than one ISL to the same switch port card unless it is part of a trunk.

Allocate 8-ports for ISLs/trunks in two groups for Edge switch

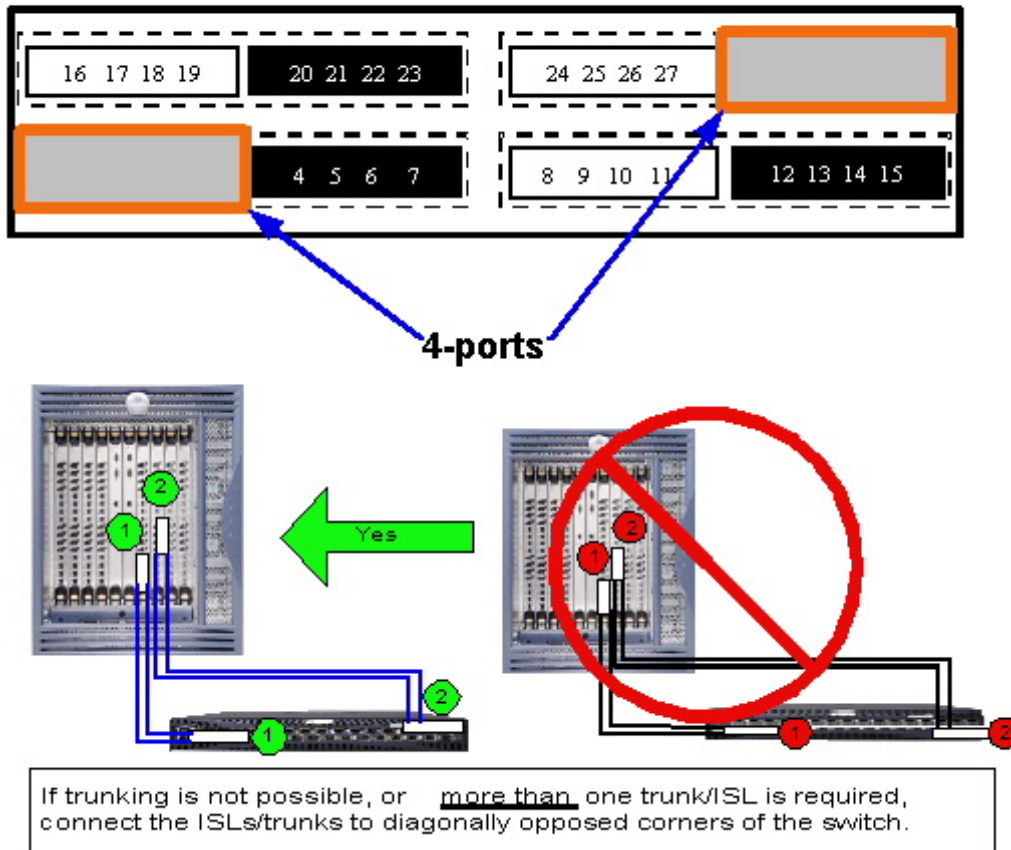


Figure 2-14 Recommended ISL/trunk connection scheme for SilkWorm 3900 switches used as edges in a Core/Edge fabric

Allocate 4-ports for ISLs/trunks in two groups for Edge switch




Figure 2-15 Recommended ISL/trunk connection scheme for SilkWorm 16-port switches used as edges in a Core/Edge fabric

Allocate 2-ports for ISLs/trunks in two groups for Edge switch




Figure 2-16 Recommended ISL/trunk connection scheme for SilkWorm 8-port switches used as edges in a Core/Edge fabric

2.4.3. Core ISL/trunk Connections

The purpose of a core switch is to connect other switches and occasionally SAN devices. When a SilkWorm switch is utilized as a core switch, similar guidelines for ISL/trunk connections as outlined in the previous sections apply. When connecting edge switches to a core switch, there are some attachment strategies for the SilkWorm 3900 and 12000 switches that enable better performance under some circumstances. These connection strategies are similar to the edge switch connection strategies and provision for up to a 3:1 host to storage ratio and assume a tiering / low locality implementation. The premise that hosts do not communicate with hosts and storage does not communicate with storage is the basis for these recommendations. For edge switches that mix host and storage on the same switch, the edge to core connections are not as important; however, if that storage is shared with devices located on other edge switches, it is suggested to connect these edge switches as if they were storage edge switches. Connect the storage edge switches diagonally to the SilkWorm 12000 core and to diagonally opposed corners for the SilkWorm 3900 core, as shown in Figure 2-17. There are no performance or availability benefits for adopting a connection plan for 8 and 16-port switches; however, standardization on a connection plan does enable simpler operation and management of these switches (see Figure 2-18).

Guideline: Connect storage edge switches diagonally to the SilkWorm 12000 core and to diagonally opposed corners for the SilkWorm 3900 core. Fill the remaining quads with host edge switches.

Guideline: Connect 8 and 16-port storage and host edge switches in a groups to realize operation and maintenance simplicity.

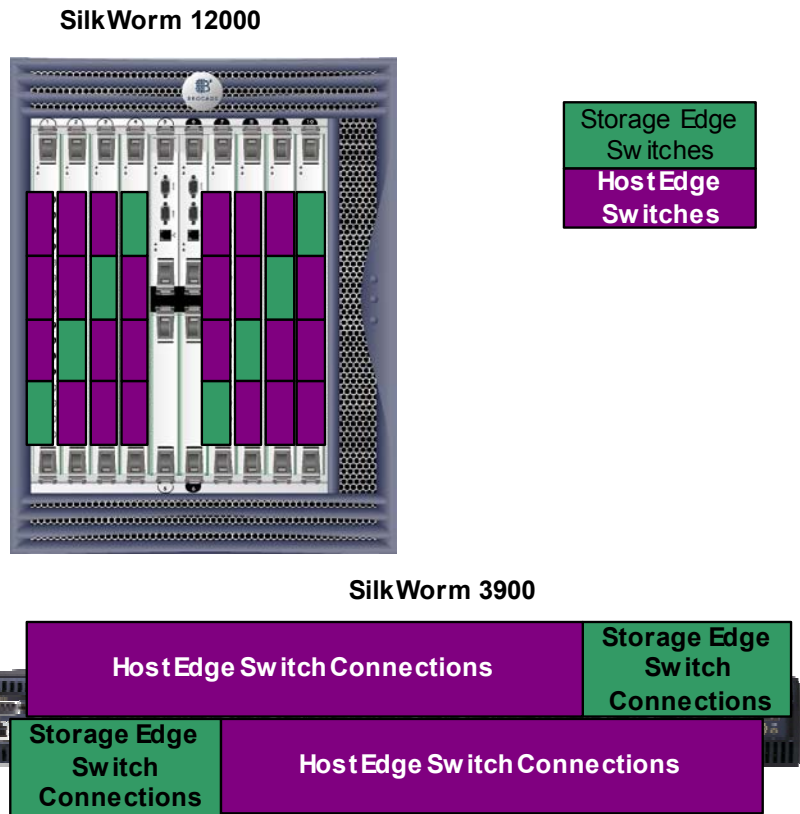


Figure 2-17 ISL/trunk connection method for SilkWorm 3900 and 12000 switches used as core switches in a Core/Edge fabric




Figure 2-18 ISL/trunk connection method for SilkWorm 8 and 16-port switches used as core switches in a Core/Edge fabric

2.4.4. Attaching SAN Devices For Availability

When attaching devices to a SilkWorm 12000, it is recommended to distribute the connections across blades to minimize the impact of a switch port card failure. To effectively distribute the connections, it is important to understand the connection types and relationships. For example a large storage array may support sixteen or more connections. If all sixteen connections were made to a single switch port card, then the failure of that switch port card would sever all sixteen connections to that array. However, if these sixteen connections were evenly distributed across a SilkWorm 12000 switch, the failure of a switch port card would only affect four of the sixteen array ports. Of course, if a redundant (i.e. dual) fabric SAN architecture is implemented, a switch port card failure would have minimal or no impact on operations. Figure 2-19 depicts the attaching devices across switch port cards for availability. Note the high port count devices are distributed across all switch port cards and that devices are partitioned and distributed by type across switch port cards. While it is not necessary to attach devices in groups, as shown in Figure 2-19 it does make it easier to manage the device connections. This same guidance holds true for hosts that connect into the same fabric with multiple connections. For SilkWorm 8-port, 16-port, and 3900 switches, connect devices for availability as shown in Figure 2-20.

Guideline: For the SilkWorm 12000, distribute devices across switch port cards from left to right for optimal availability

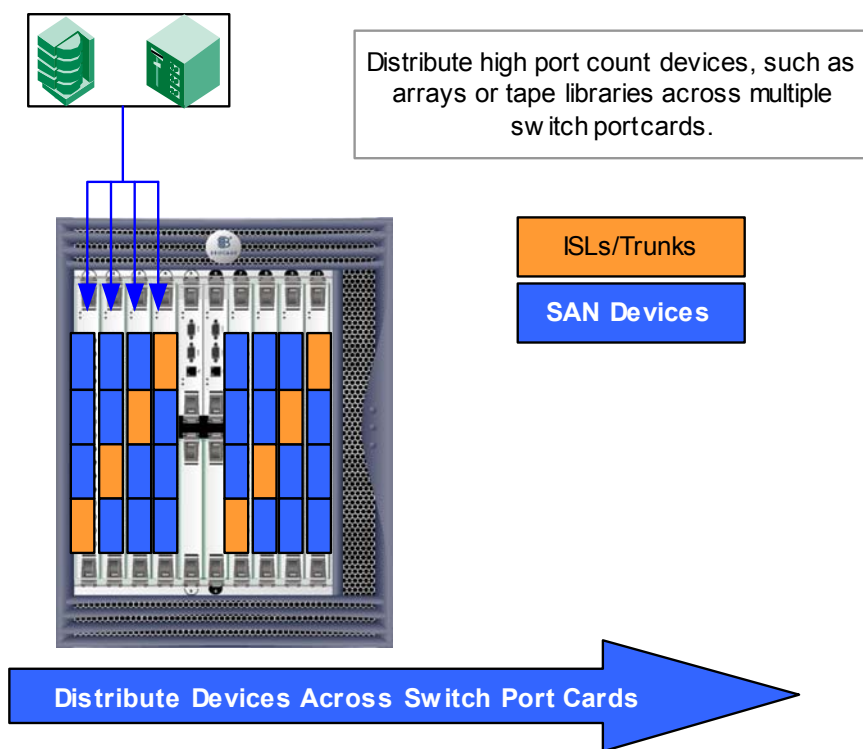


Figure 2-19 Attaching devices for availability on a SilkWorm 12000 used as an edge switch

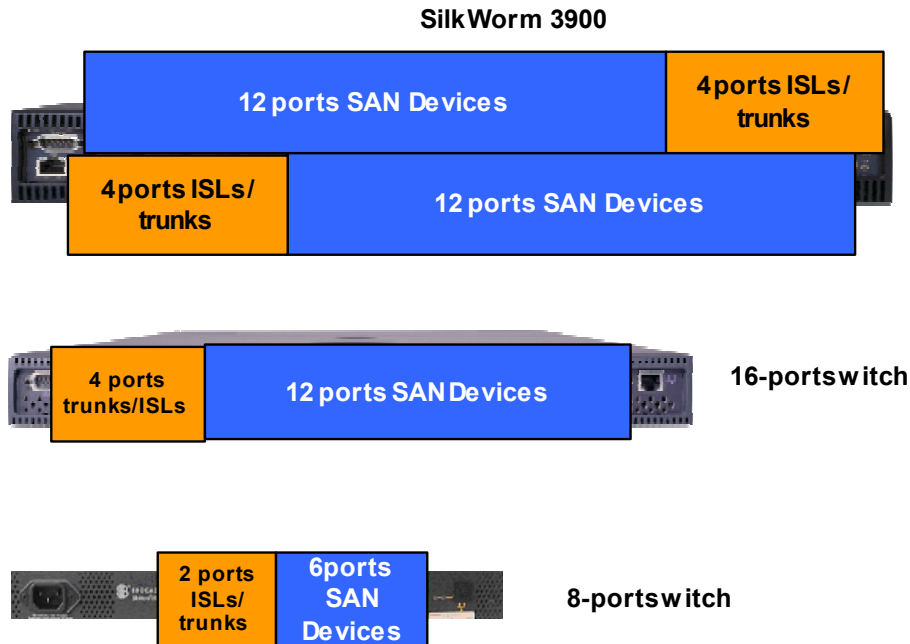


Figure 2-20 SilkWorm 8-port, 16-port, and 3900 device connection plans for edge switches

2.4.5. Connecting For Scalability

When designing SilkWorm based Core/Edge SANs it is key to identify the scaling requirements for the SAN. Once these scaling requirements are known, it is possible to allocate switch ports for ISLs needed for current switch connections as well as future switch connections. If the scaling requirements are not too large and there are unallocated ports, then the remaining ports can be utilized for attaching SAN devices to the core. In many situations, the core switch may be the most highly available switch and it is desirable to attach mission critical devices to the most available switches. Connecting a device that is shared by many other devices (i.e. many to one) may result in improved performance. Scaling for performance is also a consideration. Following the guidelines in Figure 2-20 results in a SAN with enough ports provisioned to enable the scaling of performance.

It is important to understand the implications of attaching SAN devices to the Core. While device placement does not constitute fabric topology, it may very well affect and be effected by topology. The example in Figure 2-21 illustrates how the placement of a device in a fabric can impact scalability. Scenario “A” (Local Attach) in Figure 2-21 depicts a disk system attached to the same switch as the host that needs to access it. This is an effective configuration, because it eliminates the need to manage ISL over-subscription. This configuration is useful when most traffic can be localized and congestion is a greater concern.

Scenario “B” (Core Attach) depicts the case where not all ports on the core are being used by ISLs, and the storage device is directly attached to the core. This configuration has two impacts. First, the number of available ports in the SAN is significantly reduced because core ports are no longer available for connecting additional switches. This means that the connection of a single device to the core could reduce the potential size of the SAN by as many as 64 ports or more. A high performance device that is accessed by many other devices (i.e. a storage device), may realize some performance benefit in connecting to the core. This is because the many devices can spread their load across many ISLs/trunks as opposed to only a few ISLs/trunks if that device were placed on the edge. In Figure 2-21, scenario “B”, if there were many hosts connected across the edge switches, these hosts would spread their load across eight ISLs, instead of two ISL as shown in scenario “C”.

Scenario “C” (Edge Attach) is the typical case. The number of available paths between the host and storage is two. In addition, the core switch ports are available for increasing the size of the SAN by adding new edge switches. The hop count from the server to the storage is 2-hops. Note that hop latency is two microseconds per hop – an inconsequential latency when compared to disk I/O, which is measured in milliseconds.

Guideline: Only connect devices to a core switch after validating your scaling requirements. Attaching devices to a core switch limits the size of your SAN, but can result in higher performance.




Figure 2-21 How Device Placement Can Impact Performance and Scalability in a Core/Edge Fabric

A Core/Edge fabric built with two 16-port core switches, while maintaining a 7:1 ISL oversubscription ratio, can scale to 224 user ports. Using two 64-port core switches and maintaining a 7:1 ISL oversubscription ratio, a Core/Edge fabric can scale to 896 user ports and if four 64-port switches are used in the core, a fabric can grow to 1792 user ports. If a higher ISL oversubscription ratio of 15:1 is used, a fabric built with two 64-port core switches can grow to 1920 user ports and 3840 ports using four 64-port core switches. The connection of devices to the core lowers the potential size of these SANs. If the fabric size is expected to exceed 896 user ports, to maintain the flexibility to scale performance and the size of a fabric or if it is expected to connect devices into the core, it is recommended that four SilkWorm 12000 switches be utilized in the core.

Guideline: Use four SilkWorm 12000 switches in the core of a Core/Edge topology if the size of the fabric is expected to exceed 896 user ports or a large number of devices are going to be connected to the core.

2.4.6. Attaching Devices For Performance On SilkWorm 3900 and 12000 Switches

The term “blocking” is defined as the inability of one device to connect to another device. Brocade Virtual Channel implementation of Fibre Channel does not block and SilkWorm switches are not subject to blocking. The term blocking is often confused with the term congestion. If two or more sources contend for the same destination, performance for each source may decrease; however, available bandwidth is shared fairly by all sources contending for the same destination. Congestion is the realization of the potential of over-subscription. Congestion may be due to contention for a shared storage port or host port, or an ISL.

Brocade 8-port and 16-port switches are non-blocking and congestion free. The SilkWorm 3900 and 12000 are non-blocking switches that can encounter congestion under certain conditions. Recall that a quad consists of four contiguous ports. Quads are components of the SilkWorm 3200, 3800, 3900, and 12000 switches. The SilkWorm 3900 also utilizes a component known as an octet, which consists of two quads. See Figure 2-3, and Figure 2-4 for further detail and background on quads and octets. The whitepaper *The Brocade Channeled Central Memory Architecture: Providing the Building Blocks for Enterprise SANs* details the internal architecture of the SilkWorm 12000 switch and is also helpful in understanding the SilkWorm 3900 internal architecture.

If it is suspected or requirements dictate that SilkWorm 12000 quad-to-quad traffic exceeds 4 Gbit/sec or SilkWorm 3900 octet-to-octet traffic exceeds 8 Gbit/sec, then several mitigation options should be explored. Use tools such `portPerfShow` (part of Fabric OS), Fabric Watch, or Advanced Performance Monitoring to identify high performance 2 Gbit/sec devices. The following strategies only need to be employed if several 2 Gbit/sec devices located on the SilkWorm 12000 need to or are expected to simultaneously communicate with each other in excess of 1 Gbit/sec. Note that if a device is connected at 2 Gbit/sec, it does not necessarily mean that device is utilizing 2 Gbit/sec of bandwidth.

One option involves collocating the high performance 2 Gbit/sec devices together on the same quad or octet – essentially localizing the I/O (see Figure 2-22). This is effective if multiple 2 Gbit/sec devices are communicating with each other simultaneously and at full bandwidth. 2 Gbit/sec traffic within a quad or octet is congestion free. If the high performance 2 Gbit/sec devices are not communicating with each other or are shared by many other devices, then these devices should be collocated on quads with low performing 1 Gbit/sec or 2 Gbit/sec devices (see Figure 2-23). Another option is to move the high performance 2 Gbit/sec devices to a 2 Gbit/sec switch, such as the SilkWorm 3800, which is not subject to 2 Gbit/sec congestion. This is an effective strategy if more than eight high performance 2 Gbit/sec devices need to communicate with each other. Then connect this switch into the fabric to share these high performance devices.

Guideline: Collocate high performance devices that communicate with each other on the same quad for the SilkWorm 12000 and octet for the SilkWorm 3900.

Guideline: Place high performance devices that are shared by many lower performance devices onto quads or octets with low performing devices.

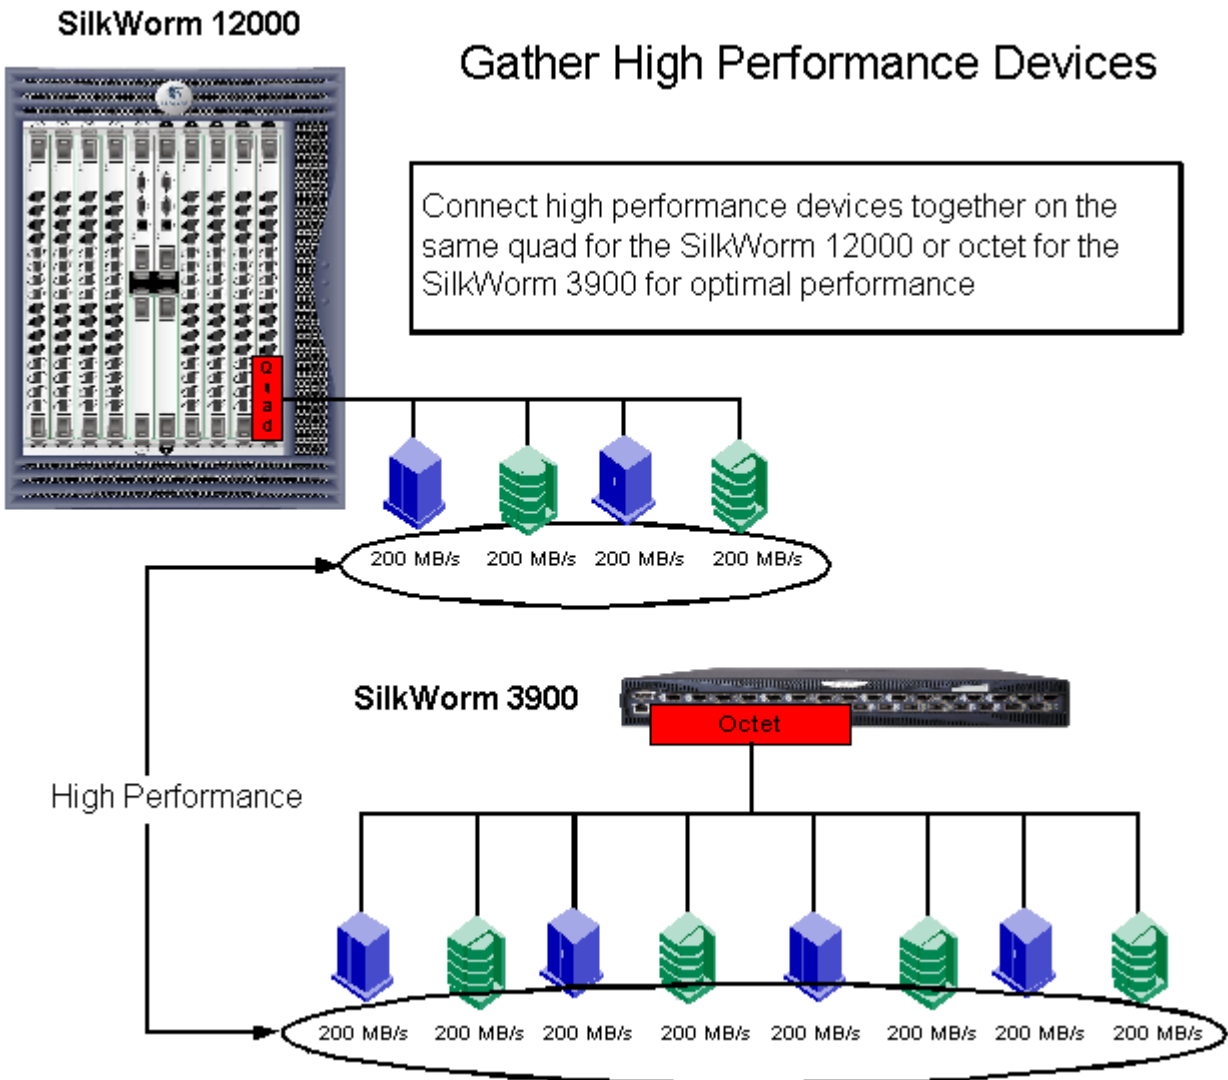


Figure 2-22 Connect high performance devices that communicate with each other onto the same quad or octet

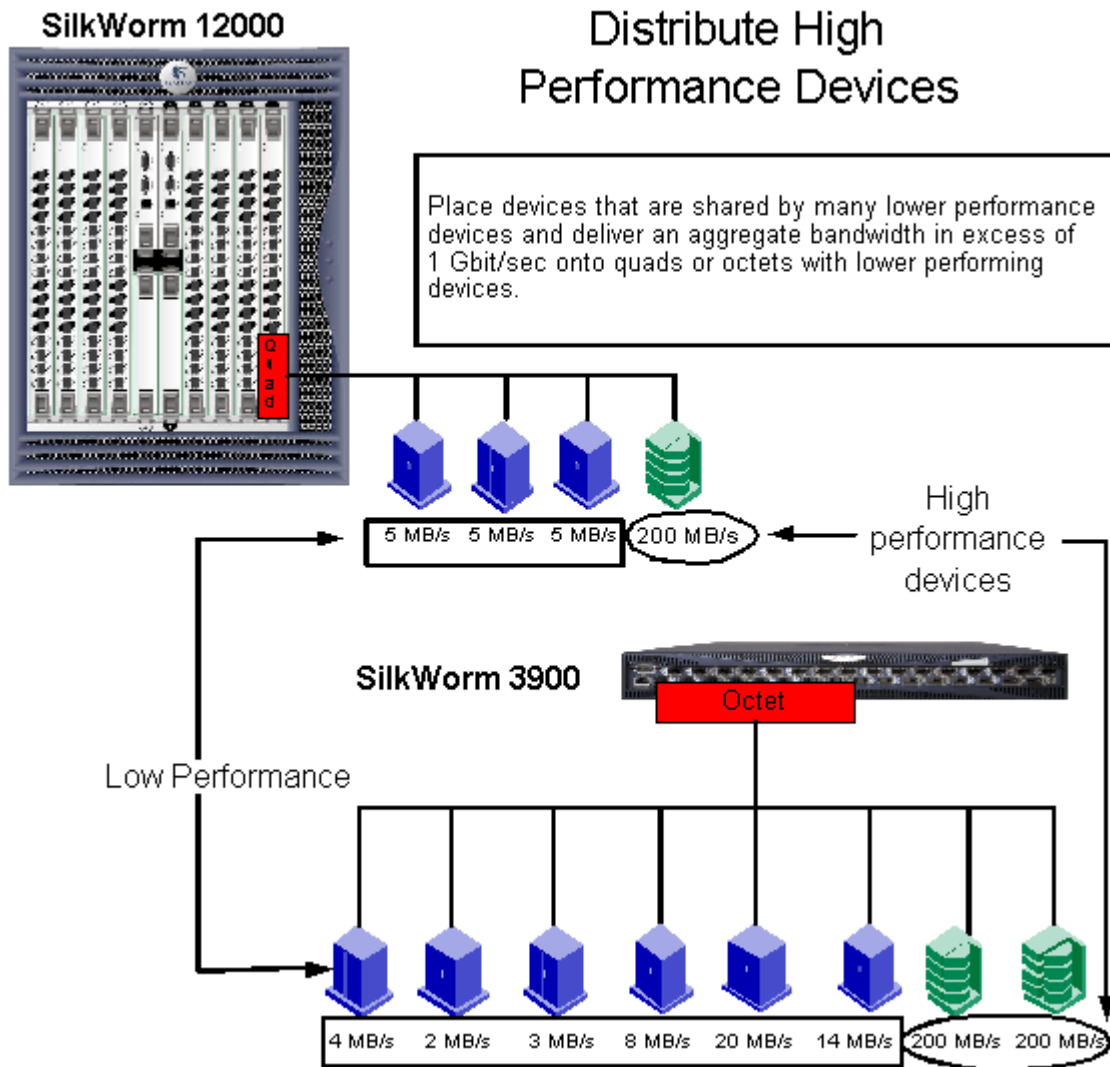


Figure 2-23 Distribute high performance devices across quads or octets with low performing devices

2.5. Platform Specific Design Considerations

Several design considerations such as switch role, usage, availability, and scalability capabilities are covered in this section. All SilkWorm switches can function as standalone switches or as part of a fabric. When participating in a fabric, some switches are better suited for the core than others. All SilkWorm switches can perform in the role of an edge switch; however, depending upon requirements, some SilkWorm switches make better edge switches than others. When building big fabrics, use high port count switches using the latest available Fabric OS. For example, non-disruptive code activation is new in Fabric OS 4.1. What is the optimal placement of Fabric OS 4.1 based switches (SilkWorm 3900, 12000) to maintain availability of the fabric? Tables are provided for each SilkWorm switch to answer the questions posed as well as other questions relating to the relationship between a particular switch platform and SAN design. Additionally, several guidelines are provided to guide the SAN designer in selecting the optimal switch for their needs.

Guideline: Significant scalability enhancements, such as improved performance of zone checking, RSCN distribution, and name server, along with improved quality and new functionality are incorporated into Fabric OS versions 2.6.1, 3.1, and 4.1. It is recommended to utilize these versions of Fabric OS. It is necessary to run Fabric OS versions 2.6.1, 3.1, and 4.1 for large SANs.

2.5.1. SilkWorm 2000 Series Switches

Since its inception, Brocade has been the leader in developing technology and testing practices to expand the limits of Fibre Channel fabrics. As fabrics increase in size, the numbers of switches, inter-switch links (ISLs), and edge devices increase rapidly. This in turn increases the demand on the fundamental computing resources of each switch's Control Processor, as it must rapidly complete tasks such as processing Zoning configuration updates distributed by other switches, analyzing and distributing RSCNs, responding to Name Server queries from hosts logging into the switch, etc. As Brocade assembles and tests fabrics upwards of 1,300 total ports with Fabric OS v3.1.0 and v4.1.0, the SilkWorm 2000 family of switches has approached the limits of its supported capacity. For this reason, the maximum supported size for fabrics containing SilkWorm 2000 family switches remains 500 user ports (see Appendix A, *Glossary* for the definitions of user ports and total ports). This limitation is not expected to change in future releases. Customers planning to build larger fabrics should plan on implementing them solely with the Brocade 2 Gbit/sec switch family. The latest Brocade Fabric OS v3.1.0 and v4.1.0 releases allow the SilkWorm 2 Gbit/sec switches to scale well past this 500 port limit. For fabrics needing to scale beyond 750 user ports, it is recommended to build those fabrics with SilkWorm 12000 and 3900 switches, as the total cost of ownership efficiencies of building large fabrics with large switches are significant. Additionally, the processing power and memory resources of the SilkWorm 12000 and 3900 are significantly better than the SilkWorm 3800 or 3200 switches.

Note: The term control processor is associated with a SilkWorm 12000 component/FRU (field replacable unit). The SilkWorm 2000 series, 3200, 3800, and 3900 series switches do not have a FRU specifically associated with it and when CP is used in the context of other SilkWorm switches, the reference is to the switch CPU and not a FRU.

Guideline: Do not exceed 500 user ports in a fabric that uses SilkWorm 2000 series switches.

2.5.2. SilkWorm 12000 and 3900

Using one fabric per SilkWorm 12000 chassis prevents the same Fabric OS from populating two fabrics. It is not desirable to populate both fabrics of a dual fabric SAN with the same Fabric OS at the same time during an upgrade, since a human error or critical bug could cause both fabrics to be impacted. Once a new version of Fabric OS is loaded and validated, then the new version of Fabric OS can be loaded on the other fabric. If there are two fabrics on the same chassis, it is not possible to sequentially load a new version of Fabric OS. One fabric per SilkWorm 12000 chassis limits operator error to a single fabric. One fabric per SilkWorm 12000 chassis is simpler to manage, as it is not necessary to account for both fabrics when doing administration. The simpler solution will likely be more available. Separation and compartmentalization are fundamental to mitigating single points of failure. A single chassis effectively joins the two fabrics from a risk perspective. When a chassis participates in more than one fabric, there exists a potential to introduce faults into both fabrics. With one fabric per chassis, this risk does not exist.

All SilkWorm switches, except the SilkWorm 12000, implement a fixed chassis design. The SilkWorm 12000 offers a bladed design, which enables the non-disruptive expansion of ports on the chassis as well as hot swap of the CP and switch part cards.

Guideline: Use only one fabric per SilkWorm 12000 chassis.

Note: With Secure Fabric OS, only one fabric per chassis is supported.

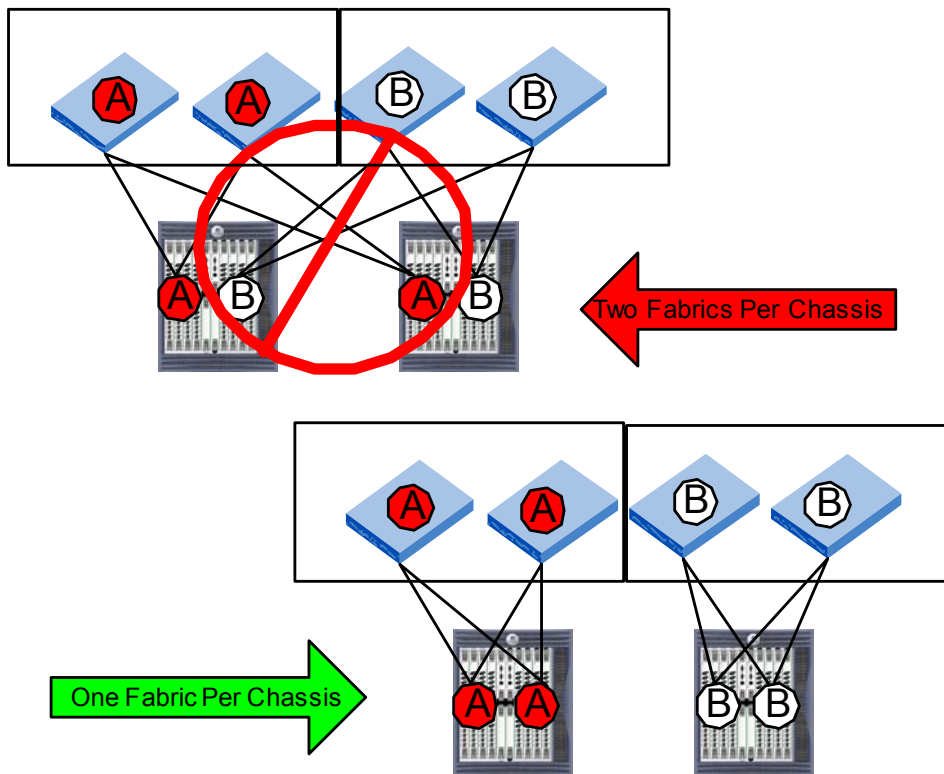


Figure 2-24 One Fabric per SilkWorm 12000 Chassis is Recommended

Hot code activation (HCA) is one of the features in Fabric OS 4.1. The defining characteristic of hot code activation is that while a new firmware image is being activated on a switch there is no disruption of end-to-end data flow between the hosts and storage devices. No disruption means no dropped frames, no retries, and no time-outs. The ASICs on the switch continue to process frames while the new firmware is being activated. Hosts that are logged into their targets will never be aware that anything has happened. This lack of end-to-end data disruption can be achieved with Fabric OS 4.1 on *both* the SilkWorm 12000 and the SilkWorm 3900.

A redundantly configured SilkWorm 12000 can pass control to its standby Control Processor as part of the firmware activation. Critical fabric services are available in about 2 seconds, with all fabric and management services fully available in less than ten seconds. The SilkWorm 3900, with only one Control Processor, must shutdown and reboot Linux as part of the HCA process. This process completes in less than 60 seconds, with the critical fabric services available in about 50 seconds. Again, even though there is a window when fabric services are unavailable, there is no disruption to end-to-end data flow between SAN devices.

Because it does take longer to do hot code activation on the SilkWorm 3900, the switches directly linked to the SilkWorm 3900 need to be tolerant of the 60 seconds when no fabric services from the SilkWorm 3900 will be seen. We have made the necessary modifications to Fabric OS v3.1.0 and v2.6.1 to extend the timeout values so that link and fabric re-configuration is avoided. For this reason, it is strongly recommended that customers deploy neighboring switches (i.e. immediately connected via an E-port) to the SilkWorm 3900 with Fabric OS v2.6.1 on the SilkWorm 2000s series, Fabric v3.1.0 on the SilkWorm 3200/3800, or Fabric v4.1.0 on the SilkWorm 3900/12000.

During the SilkWorm 3900 hot code activation, if earlier releases (Fabric OS v2.6.0x, Fabric OS v3.0.2x, or v4.0.x) are deployed on the neighboring switches to a SilkWorm 3900, a timeout will occur on these neighboring switches resulting in a fabric re-configuration.

The SilkWorm 3900 HCA procedure should always be performed when the fabric is stable. However recovery actions will be taken upon completion of the HCA to ensure that no RSCNs, or fabric configuration changes are missed. If any new hosts or targets were added to the SilkWorm 3900 switch during the HCA reboot time, then the initial FLOGI will time out. After the reboot, the switch will reset that port to cause the FLOGI to happen again. If any device is added elsewhere in the fabric, the RSCN will be delivered after the reboot completes. Finally, if any E-port cables are pulled, or the fabric rebuilds for any reason, then after the HCA reboot completes, the SilkWorm 3900 will cause the fabric to rebuild again so that it may participate in the fabric.

Note: To prevent unnecessary disruptions in the fabric when firmware is activated on a SilkWorm 3900, it is required that any switches directly connected to the SilkWorm 3900 use Fabric OS v2.6.1, v3.1.0, or v4.1.0 (or subsequent versions).

Caution: SilkWorm 2000 series, 3200, and 3800 switches require the setting of the Core Switch PID format to coexist in fabrics with SilkWorm 3900 and 12000 switches.

Guideline: When setting up new SANs, it is recommended to set the Core PID format on SilkWorm 2000 series, 3200, and 3800 switches.

2.5.3. SilkWorm Switch Usage in a Core/Edge Topology

In this section, the various aspects of manageability, availability, performance, and scalability are discussed for each SilkWorm switch platform (see Table 2-6, Table 2-7, and Table 2-8). Table 2-4 and Table 2-5 takes all of these elements into consideration and provides recommendations for switch usage in a core /edge fabric.

Table 2-4 Edge switch recommended usage based on the size of a fabric

Edge Switch Recommended Usage ¹					
	1-96 User Ports	97-224 User Ports	225-500 User Ports	501-750 User Ports	751+ User Ports
SilkWorm 2000 Series	Excellent	Very Good	Good	Not Supported	Not Supported
SilkWorm 3200, 3800	Excellent	Very Good	Good	Good	Not Recommended
SilkWorm 3900 ¹	Excellent	Excellent	Excellent	Excellent	Excellent
SilkWorm 12000	Excellent	Excellent	Excellent	Excellent	Excellent

¹ Brocade scalability testing is an ongoing effort with the maximum size of successfully tested fabrics increasing. Support of a particular SAN configuration is a based upon determinations made by your support provider. Current Brocade testing has validated fabrics in excess of 1200 user ports.

Table 2-5 Core switch recommended usage based on the size of a fabric

Core Switch Recommended Usage ¹					
	1-96 User Ports	97-224 User Ports	225-500 User Ports	501-750 User Ports	751+ User Ports
SilkWorm 2000 Series	Good	Good	Not Recommended	Not Supported	Not Supported
SilkWorm 3200, 3800	Very Good	Very Good	Not Recommended	Not Recommended	Not Recommended
SilkWorm 3900 ¹	Very Good	Very Good	Not Recommended	Not Recommended	Not Recommended
SilkWorm 12000	Excellent	Excellent	Excellent	Excellent	Excellent

¹ Brocade scalability testing is an ongoing effort with the maximum size of successfully tested fabrics increasing. Support of a particular SAN configuration is a based upon determinations made by your support provider. Current Brocade testing has validated fabrics in excess of 1200 user ports.

Table 2-6

Fabric OS 2.x Platforms					
		SilkWorm 20x0	SilkWorm 2400	SilkWorm 22x0	SilkWorm 2800
Number Ports		8	8	16	16
Speed		1 Gbit/sec			
Standalone		This switch is an excellent fit for small, pre-packaged solutions, such as clusters or tape backup.		This switch is an excellent fit for small to medium, pre-packaged solutions, such as clusters or tape backup.	
Edge		This switch can perform as an edge switch; however, to maintain resiliency, two ports need to be dedicated for ISLs.		This switch can perform as an edge switch.	
Core		Can be used as a core for fabrics up to approximately 112 user ports. Because this is a 1 Gbit/sec switch, it is not recommended as a core switch when 2 Gbit/sec switches also exist in the fabric.		Can be used as a core for fabrics up to approximately 224 user ports. Because this is a 1 Gbit/sec switch, it is not recommended as a core switch when 2 Gbit/sec switches also exist in the fabric.	
Scalability		Brocade will not increase the maximum supported fabric size for fabrics containing SilkWorm 2000 family switches beyond the currently supported maximum of 500 user ports.			
Availability	Redundant Power Supplies	No	Yes	No	Yes
	Redundant Cooling	No	Yes	No	Yes
	Hot Swap Power	No	Yes	No	Yes
	Hot Swap Cooling	No	Yes	No	Yes
	Redundant Control Processor	No	No	No	No
	Non-disruptive code activation	No	No	No	No
	Non-disruptive port expansion	No	No	No	No
Platform Specific Comments	Fabric OS	A minimum of Fabric OS 2.6.1 is recommended for Scalability and quality purposes			
	Requires Core PID Format	No, but setting Core PID format is suggested for compatibility purposes.		No, but setting Core PID format is suggested for compatibility purposes.	
	Serial / Front Panel	Serial port	Serial port	Serial Port	Front Panel; No Serial port.

Table 2-7

Fabric OS 3.x Platforms			
		SilkWorm 3200	SilkWorm 3800
Number Ports		8	16
Speed		1-2 Gbit/sec auto-sensing	
Standalone		This switch is an excellent fit for small, pre-packaged solutions, such as clusters or tape backup.	This switch is an excellent fit for small to medium, pre-packaged solutions, such as clusters or tape backup.
Edge		This switch can perform as an edge switch; however, to maintain resiliency, two ports need to be dedicated for ISLs.	This switch can perform as an edge switch.
Core		Can be used as a core for fabrics up to approximately 112 user ports.	Can be used as a core for fabrics up to approximately 224 user ports.
Scalability		The efficiencies of building large fabrics with large switches are significant. As fabrics exceed 500 ports, so does the ability of this switch to operate acceptably in large fabrics. Please consult your support provider to determine the supportability of your fabric.	
Availability	Redundant Power Supplies	No	Yes
	Redundant Cooling	Yes	Yes
	How Swap Power	No	Yes
	Hot Swap Cooling	No	Yes
	Redundant Control Processor	No	No
	Non-disruptive code activation	No	No
	Non-disruptive port expansion	No	No
Platform Specific Comments	Fabric OS	A minimum of Fabric OS 3.1 is recommended for Scalability and quality purposes	
	Serial / Front Panel	Serial port	Serial Port
	Requires Core PID Format	No, but setting Core PID format is suggested for compatibility purposes.	No, but setting Core PID format is suggested for compatibility purposes.

Table 2-8

Fabric OS 4.x Platforms			
		SilkWorm 3900	SilkWorm 12000
Number Ports		32	128 (up to two 64-port domains per chassis)
Speed		1-2 Gbit/sec auto-sensing	
Standalone		This switch is an excellent fit as a standalone SAN of 32 ports. For high bandwidth applications, it is recommended to follow the tuning guidelines discussed in Recommended ISL Oversubscription Ratios on page 2-9 .	This switch is an excellent fit as a standalone SAN of up to 64-ports.
Edge		This switch performs well as an edge switch and is suggested as an edge for larger fabrics.	This switch performs well as an edge switch and is suggested as an edge for larger fabrics.
Core		Can be used as a core for fabrics up to approximately 500 user ports.	Using a Core/Edge architecture, it is theoretically possible to build a fabric that consists of approximately 3840 user ports.
Scalability		The efficiencies of building large fabrics with large switches are significant. If your fabric is expected to exceed 750 ports in size, it is recommended to use the SilkWorm 3900 and 12000 switches.	
Availability	Redundant Power Supplies	Yes	Yes
	Redundant Cooling	Yes	Yes
	Hot Swap Power	Yes	Yes
	Hot Swap Cooling	Yes	Yes
	Redundant Control Processor	No	Yes
	Non-disruptive code activation	Yes	Yes
	Non-disruptive port expansion	No	Yes

Platform Specific Comments	Fabric OS	A minimum of Fabric OS 4.1 is recommended for Scalability and quality purposes	
	Serial / Front Panel	Serial port	2 serial ports (1 per CP) and 2 modem ports (1 per CP)
	Requires Core PID Format	Yes	Yes
	Switch Adjacency	To prevent unnecessary disruptions in the fabric when firmware is activated, it is recommended that any switches connected to the SilkWorm 3900 run V2.6.1, v3.1, or V4.1 of Fabric OS.	Not applicable.
	One Fabric Per Chassis	Not applicable.	<p>The implementation of the SilkWorm 12000 results in two 64-port switches per chassis. It is possible to connect each switch to a different fabric. It is recommended to connect both switches in a chassis to the same fabric. Consider the following when deploying a SilkWorm 12000:</p> <p>The implication of running two fabrics per chassis is that an entire SAN could be lost should a catastrophe be encountered. Splitting the fabrics across two chassis can mitigate catastrophe (such as fire, water damage, human error, etc.)</p> <p>Running one fabric per chassis prevents same Fabric OS from populating two fabrics. It is not desirable to populate both fabrics of a dual fabric SAN with the same Fabric OS at the same time since a critical bug could cause both fabrics to malfunction. If you have two fabrics on the same chassis, you cannot sequentially load a new version of Fabric OS.</p> <p>One fabric per chassis limits operator error to a single fabric.</p>

2.6. Zoning Design Considerations & Guidelines

Zoning is an important element of a secure and healthy SAN. Zoning does have an impact on a SAN designs. The *Brocade Zoning User's Guide Version 3.1/4.1* (publication number: 53-0000523-01) provides a solid overview of how zoning works and guidelines for implementing zoning. This section highlights key elements of zoning that relate to a SAN design.

2.6.1. Zoning and Scalability

Zoning optimizes fabric services, such as RSCN distribution and name server response, and limits unnecessary device discovery. With the new zoning and related name server changes in Fabric OS 3.1 and 4.1, zoning becomes necessary for the proper functioning of large fabrics. For instance, the distribution of RSCNs (registered state change notifications) is reduced to only devices affected by a zone change. In prior releases of Fabric OS 3.x, 4.x, and all versions of Fabric OS 2.x, a zone activation (for example, executing the command `cfgEnable`) resulted in an RSCN being distributed to all devices – regardless of whether these devices were affected by a zone change. Additionally, not using zoning results in unnecessary delays during device discovery for some hosts, especially when a host pointlessly authenticates with hundreds of devices. These delays can last minutes, pause ongoing I/O, and cause unpredictable behavior on a host. Use of zoning on the switches limits the number of devices visible to a host and eliminates this host-based scalability problem.

Guideline: The implementation of zoning is recommended for any SAN and especially critical for any large fabric since zoning is fundamental to the functioning of multi-hundred port fabrics.

2.6.2. Zoning Database Size

Zoning consumes a finite amount of processing and memory resources. As the number of devices in a SAN grows, so do the demands on these same resources. The zoning implementation is optimized to minimize processing resources and leverage ASIC capabilities as much as possible. The zoning database size for SilkWorm 2000 series, 3200, and 3800 switches is 96 KB and 128 KB for SilkWorm 3900 and 12000 switches. To check the size of a zone database, use the command `cfgSize`. A switch with a zoning database size limit of 96 KB limits the size of the zoning database for the whole fabric – even if a SilkWorm 3900 or 12000 switch is present in the fabric. As the size of a SAN grows, it is important to monitor the zoning database. Typically, the zone database size needs to be of concern as the size of a SAN exceeds several hundred ports. The size of an alias name, zone name, or configuration name is limited to 64 characters for Fabric OS versions 2.6.1, 3.1, and Fabric OS 4.1. While it is possible to create 64-character zone, alias, or configuration names, doing so consumes more memory than a shorter name. Additionally, shorter names are easier to remember and less prone to typing errors. Be wary of sacrificing meaning for shortness. See the whitepaper *Zoning Implementation Strategies For Brocade San Fabrics* for effective guidance for naming aliases, zones, and configurations. The variable size of zone objects makes it very difficult to state guidelines as a number of zone entries or alias. A zone database size is similar to disk storage. The usage is not measured so much by how

many files are located on the storage, but by the amount of space taken up by the files. It is recommended to review the zoning configuration of a fabric periodically for unused alias, zoning, and configuration entries and to then delete these unnecessary entries. Unnecessary alias, zone, and configuration entries frequently result from the merging of fabrics or the addition of a switch with predefined zones into an existing fabric.

Note: The maximum zoning database size for SilkWorm 2000 series, 3200, and 3800 switches is 96 KB and 128 KB for SilkWorm 3900 and 12000 switches. A switch with a zoning database size limit of 96 KB limits the size of the zoning database for the whole fabric – even if a SilkWorm 3900 or 12000 switch is present in the fabric.

Guideline: Limit the name of an alias, zone or configuration to as few characters as possible while maintaining meaning of that name. Target 16-characters or less for an alias, zone, or configuration name.

Guideline: For SANs that exceed several hundred ports, monitor the size of the fabric-zoning database with the command `cfgSize`.

Guideline: Routinely review a zoning configuration to identify unused aliases, zones, and configurations and then remove these unused entries.

2.7. Designing SANs With Secure Fabric OS

A secured fabric must be *entirely* secured and all switches in a secured fabric must run a version of Fabric OS that supports security and these switches must be licensed to run security. Scalability testing, which is the testing performed by Brocade of fabrics consisting of hundred and thousands of device ports, with fabrics that have secure mode enabled and with fabrics that have secure mode disabled. Currently, the size of fabrics tested with security enabled are smaller than fabrics tested without security enabled.

Note: Currently the number of switches and devices tested in fabrics with secure mode enabled are lower than those fabrics tested with secure mode disabled.

Fabrics consisting solely of SilkWorm 2000 series switches need to run a minimum version of Fabric OS v2.6.0. If it is desired to run with secure mode enabled in fabrics containing SilkWorm 2000 series switches and any of the following SilkWorm switches: 3200, 3800, 3900, 12000, the minimum version of Fabric OS is V2.6.1. Secure Fabric OS was introduced for the SilkWorm 3200 and 3800 switches in Fabric OS v3.1 and v4.1 for SilkWorm 3900 and 12000 switches.

Note: The following minimum version requirements apply for fabrics that need secure mode enabled:

- SilkWorm 2000 series
 - Fabric OS version 2.6.0 is required for secure fabrics containing only SilkWorm 2000 series switches.
 - Fabric OS version 2.6.1 is required for secure fabrics containing a SilkWorm 2000 series switch and any of the following SilkWorm switches: 3200, 3800, 3900, 12000.
- Secure Fabric OS version 3.1 is required for SilkWorm 3200 and 3800 switches
- Secure Fabric Os version 4.1 is required for SilkWorm 3900 and 12000 switches

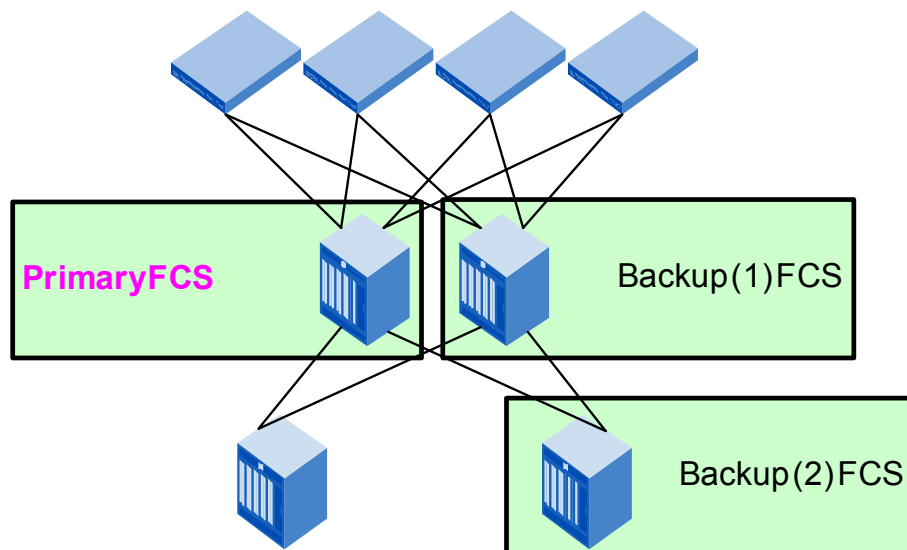
To run with secure mode enabled, requires the use of one or more switches as the Fabric Configuration Server (FCS). FCS switches are “trusted” switches and are used for managing fabrics where secure mode is enabled. These switches should be both electronically and physically secure. You can specify a Primary FCS switch and one or more backup FCS switches, to provide failover ability should the Primary FCS switch fails. All management access to the fabric must flow through the primary FCS switch. Should the primary FCS switch be unavailable, it then becomes necessary to use the first available backup FCS switch for managing the fabric. Please reference *Brocade Secure Fabric OS® User’s Guide Version 3.1 / 4.1* manual for further detail about Secure Fabric OS. The ability exists to use the backup FCS via automatic failover or manually. In case where a backup FCS is utilized manually, the user can pick any switch designated as a backup FCS. In cases of automatic failover, the first backup FCS switch designated will be used.

The primary FCS switch is a central point for distributing fabric configuration information and management changes. Establishing the core switch in a Core/Edge fabric as the primary FCS is recommended since the core switch is optimally located to communicate with all other switches in the fabric. There are several other reasons, which are discussed in [Switch Location In The Fabric on page 2-38](#), for configuring one of the core switches as the primary FCS. For the same reasons it is also recommended to establish the other core as the backup FCS. Note that some core edge topologies utilize three and even four cores. A second backup FCS switch is also recommended and this switch should be an edge switch. One backup FCS is switch is necessary in a Core/Edge topology since a single switch failure does not cause total failure of the fabric, so means to manage the fabric should a single switch fail should exist. Conceivably all the remaining switches could also be established as backup FCS switches. There is a fine balance between being overly cautious and creating unnecessary complexity. For this reason a total of two backup FCS switches are recommended. Figure 2-25 is an example implementation of the recommendations for quantity, type, and location of FCS switches in a secure fabric. If a fabric spans multiple sites, locate at least one backup switch FCS at each site.

Guideline: Use one of the core switches in a Core/Edge topology as the primary FCS switch for secure fabrics.

Guideline: In a Core/Edge topology, designate at least two additional switches as backup FCS switches. It is recommended to use the other core as the first backup FCS and an edge switch as the second backup FCS.

Guideline: When deploying a secure fabric that spans multiple sites, ensure that at least one backup FCS exists at each site.



Note: Each SilkWorm 12000 icon represents a single SilkWorm 12000 logical switch.

Figure 2-25 Recommended Location of Primary and Backup FCS Switches in a Core/Edge Topology

2.8. Switch Location In The Fabric

The Brocade family of switches offers a variety of port counts, availability levels, and performance. The SAN designer has multiple options for placing a new switch into an existing fabric or creating fabric with new switches. Several variables influence switch location in the fabric, including management, control, switch speed, availability, interoperability, and availability. These variables are discussed in this section with recommendation and guidelines relating to:

- Which switch to use for fabric management
- Where to locate 2 Gbit/sec switches as they are integrated into an existing fabric
- Where and what switches to place to match fabric availability with overall availability requirements

2.8.1. Management and Control

It is important that the SAN administration team select one switch in the fabric as the administration switch. Using one switch for access lessens the possibility of multiple administrators making changes to different switches in the fabric at the same time. Keeping in line with the terminology used by Secure Fabric OS, this switch is termed the primary management switch. There are several reasons to select the core switch as the primary management switch for zoning, time services, Fabric Manager, Web Tools, and general administrative access.

The core switch is optimally located to communicate with all other switches in the fabric. The core switch is commonly the larger switch with a more powerful control processor. When zoning information is propagated, the switch where the zone information changed is responsible for distributing the zone information to all other switches in the fabric. The core switch is directly connected to all other switches in the fabric. Typically, the core switch is a later model switch providing the administrator access via the latest Web Tools interface. The primary management switch can also be used by Fabric manager as the main access point for that fabric. Once selected, all the other portions of Fabric manager will primarily access this switch for fabric information. The ability to synchronize time within a fabric is enabled with Fabric OS versions 2.6.1, 3.1, and 4.1.

All switches, in a fabric where secure mode is not enabled, synchronize their time with the principal switch in the fabric. The principal switch in the fabric can synchronize its clock with an NTP timeserver, by identifying the timeserver to the principal switch with the `tsclockserver` command. In a fabric where secure mode is enabled, switches synchronize time with the primary FCS, which may or may not be the principal switch.

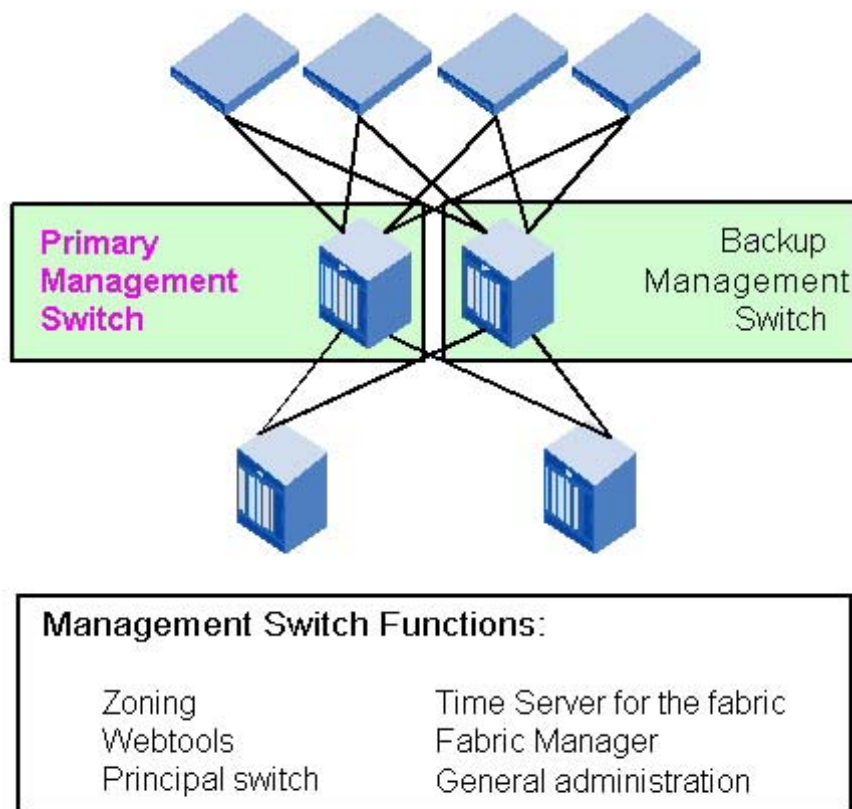
Guideline: Select a core switch in a core /edge topology as the primary management switch for zoning, time services, Fabric Manager, Web Tools, and general administrative access

The establishment of a principal switch in a fabric can vary based on upon the state of the fabric, switch WWN, and whether other switches/fabrics are merging into that fabric. A switch that is the principal switch in a fabric today may not be a principal switch after a new switch or switches are added to that fabric or if that fabric reconfigures. The implementation of the `fabricprincipal` command is based solely on mechanisms specified in the Fibre Channel standards. These mechanisms provide a *preference* for a switch requesting to be the principal switch in a fabric, but they do not provide an absolute guarantee that a switch requesting to be the principal switch will actually achieve this status. Note that the `fabricprincipal` is only available in Fabric OS version 4.1.0 or later.

Guideline: Identify the primary management switch in the fabric as the preferred principal switch by using the command `fabricprincipal`.

With the use of secure Fabric OS, the management of a fabric must be originate from the primary FCS switch and should the primary FCS switch be unavailable, a designated backup FCS switch can be used. The use of a primary switch as a focal point for SAN management is a guideline, and any switch can be used for these purposes. For the reasons stated in this section, it makes sense for the core switch to be the primary switch for management purposes. Should the primary switch fail in a non-secured fabric (i.e. a fabric not running secure Fabric OS), then it is suggested to utilize the remaining core as primary backup.

The primary management switch can also be used as an access point for management server access, access by SNMP software that polls for fabric status, a focal point for fabric related SNMP traps, and as an access point for SAN management software. Figure 2-26 summarizes the recommended location of a primary switch in a non-secured fabric, the recommended functions that the primary should perform, and the location of a backup management switch.



Note: Each SilkWorm 12000 icon represents a single SilkWorm 12000 logical switch.

Figure 2-26 Primary and Backup Management Switch Location and Functions in a Core/Edge Topology

2.8.2. 2 Gbit/sec Switch Placement

When designing a SAN with 2 Gbit/sec switches, the similar guidelines that apply to trunking also apply to 2 Gbit/sec capabilities. Place these switches adjacent to each other to take advantage of 2 Gbit/sec ISLs. Of course, it is also possible to connect a SilkWorm 2000 series switch to a trunking capable switch, as Brocade Trunking capable switches are backwards compatible and will negotiate a 1 Gbit/sec ISL.

For Core/Edge topologies, place 2 Gbit/sec switches in the core. If 2 Gbit/sec connectivity is required, it is acceptable to attach these devices to the 2 Gbit/sec cores if 2 Gbit/sec edge switches are not yet implemented. By placing 2 Gbit/sec switches in the core, it ensures that a 2 Gbit/sec path exists end to end. If a significant number of 2 Gbit/sec devices are required and the performance requirements are high, an effective strategy is to localize the 2 Gbit/sec devices on the same switch or group of switches.

Figure 2-27 depicts 2 Gbit/sec devices with and without 2 Gbit/sec end-to-end paths as well as the localization of some 2 Gbit/sec devices.

Guideline: Place 2 Gbit/sec switches adjacent to each other and for Core/Edge topologies, place the first 2 Gbit/sec switches in the core and subsequent 2 Gbit/sec switches at the edge.

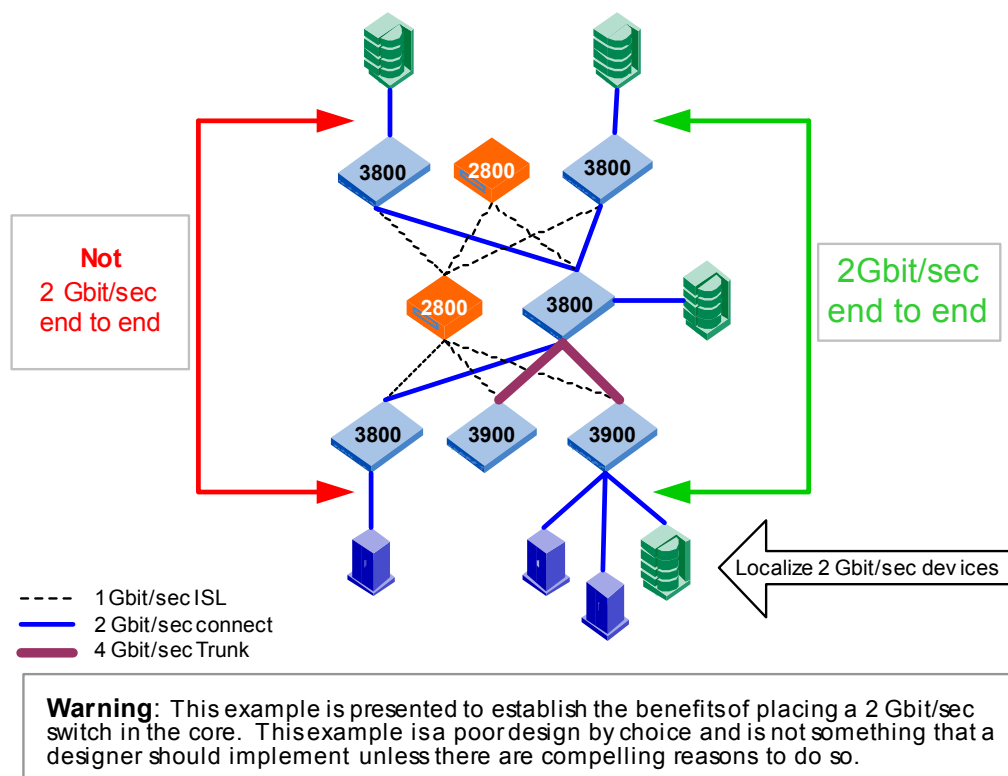


Figure 2-27 End to End 2 Gbit/sec Paths and 2 Gbit/sec Device Locality

2.8.3. Locating A Switch For Fabric and I/O Availability

The type of switch and placement of that switch in a fabric impacts availability of the fabric and I/O operations. While redundant fabrics enable redundancy, failing a path or fabric is the last thing many SilkWorm users want to do, as maintaining continuous I/O is paramount. An effectively designed SAN can match variable availability requirements for storage and hosts. Enabling the SAN administrator to offer different levels of availability provides significant flexibility.

The SilkWorm 3900 and the SilkWorm 12000 are both capable of non-disruptive operation during hot code activation when running Fabric OS 4.1 or subsequent versions. The SilkWorm 12000 is capable of non-disruptive operation during a failover. A failover can be initiated manually or happen automatically due to the detection of a fault. While a new firmware image is being activated on a switch or during a failover there is no disruption of end-to-end data flow between hosts and storage devices. No disruption means just that – no dropped frames, no retries, no time-outs. The SilkWorm 2400, 2800, 3800, 3900 and 12000 all implement hot swap and redundant power and cooling. The SilkWorm 12000 is the only switch that implements redundant control processors. This means that the SilkWorm 12000 is the only SilkWorm switch capable of operating in the event of a control processor failure. The SilkWorm 12000 provides the highest availability in the SilkWorm family of switches. The next highest level of availability offered by a SilkWorm switch is provided by the SilkWorm 3900. Table 2-9 ranks the SilkWorm family of switches by availability levels.

Table 2-9 SilkWorm switch availability levels

SilkWorm Switch	Level Of Availability	Availability Features
SilkWorm 12000	Ultra High	Redundant control processors, Non-disruptive failover, Non-disruptive port expansion, hot swap WWN card, hot code activation , redundant and hot swappable power and cooling
SilkWorm 3900	Very High	Hot code activation , redundant and hot swappable power and cooling
SilkWorm 3800, 2800, 2400	High	Redundant and hot swappable power and cooling
SilkWorm 3200, 2x00, 20x0	Medium	No redundant or hot swappable power and cooling

A dual Core/Edge fabric SAN is redundant and resilient. A Core/Edge SAN is resilient and can sustain operations in the event that a single core fails. I/O operation can continue when an edge switch fails if a redundant fabric with multipathing software is implemented. From an availability perspective, any SilkWorm can perform as a core switch in a Core/Edge fabric. To eliminate disruptions in the fabric and to I/O due to firmware activation or during a failure, a SilkWorm 12000 is recommended. The SilkWorm 12000 offers the highest availability and up to 128-ports per chassis. If scaling requirements do not dictate a several hundred port SAN, consider still using the SilkWorm 12000 in the core, populating the core with devices that require the highest availability of I/O, and connecting devices that require lower availability to lower availability edge switches. Select and connect edge switches based upon the SAN device availability requirements, as shown in Figure 2-28.

Guideline: For the highest fabric and I/O availability, deploy SilkWorm 12000 switches in the core and either SilkWorm 3900 or 12000 switches on the edge of a Core/Edge fabric. Devices requiring lower availability can connect to SilkWorm 2000 series and 3800 edge switches.

Guideline: Using two separate SilkWorm 12000 chassis in the core increases availability even further.

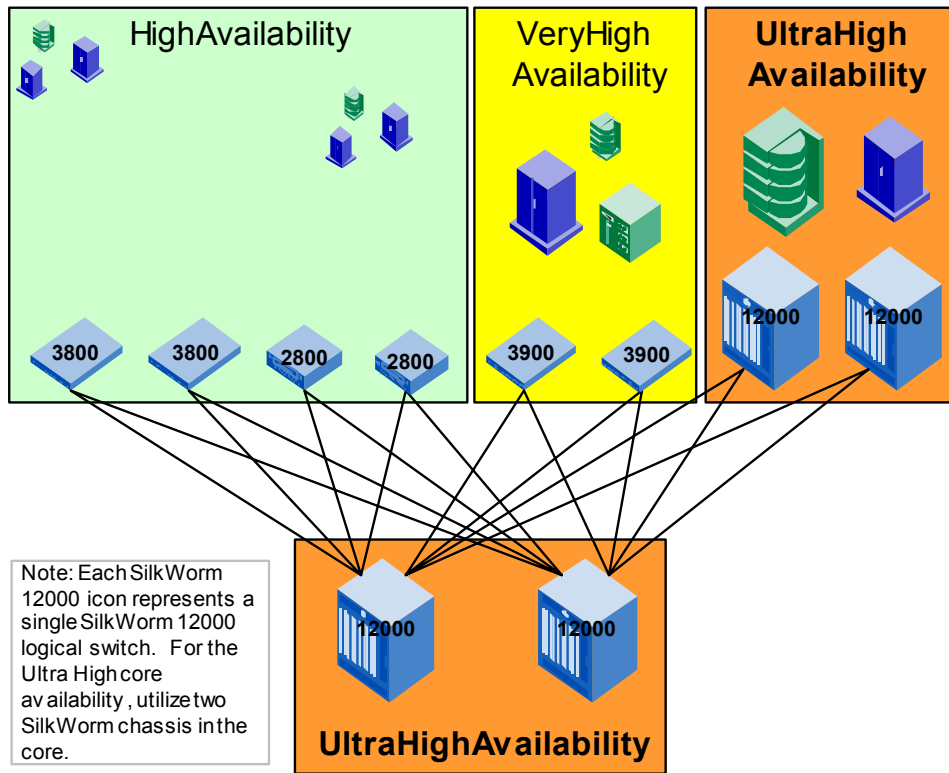


Figure 2-28 Connect Devices According to Availability Requirements

2.9. Scalability Support and Testing

Brocade is currently testing and validating fabrics consisting of hundreds of ports and has plans for testing multi-thousand port fabrics. The ultimate limitation in a fabric design today and as defined in the Fibre Channel standards is a maximum of 239 physical switches, whether they are 8, 16, or 64-port versions. As a practical matter, no vendor has yet tested networks of this size due to the expense and complexity of implementing such a network. The current practical switch-count limit is lower than 239 switches, based upon empirical testing. Other limits on a SAN design are the count and number of ISLs/trunks and the size of the zoning database.

Brocade partners with many OEMs and resellers who supply switches to end-users. Many of these partners provide direct support for the switches they sell. These partners extensively test and support specific configurations, including switches, storage, HBAs, and other SAN components. In some cases, the large fabric configurations supported by Brocade partners will differ from the guidelines Brocade presents in this document. In fact, several Brocade switch partners have developed their own SAN design guidelines, including in-depth support matrixes that specify support configurations and firmware versions.

To determine whether or not a SAN is supported, it is necessary to work with your support provider to determine if your SAN design is valid. Important variables that determine the supportability of a particular SAN are the number of switches, version of Fabric OS, number of ISLs/trunks, number of connected devices, and the hop count.

The Brocade approach to testing fabrics consisting of hundreds and even thousands of ports establishes an effective balance between quality and building a solid foundation from which a wide range of fabric configurations can be derived. Brocade scalability testing validates a range of ISL oversubscription ratios (from 3:1 to 15:1), mixture of switches, hop counts (up to 6), and ISL densities. Many customers implement varying degrees of ISL oversubscription in the same fabric based upon their performance needs. Some customers out of necessity deploy SAN archipelagoes that are composed of loosely connected

2.10. Recommended Fabric Topologies and SAN Designs

When building big fabrics use big switches. Table 2-4 and Table 2-5 outline the recommended usage of SilkWorm switches based on the size of a fabric. A list of guidelines and recommended topologies follows. Port counts in this section are specified as user ports, meaning that ISLs ports are subtracted from the total number of ports in the fabric.

Guideline: When designing a fabric that exceeds 750 user ports or is expected to grow past 750 user ports, it is suggested to exclusively utilize SilkWorm 3900 and 12000 switches in the core and the edge.

Guideline: The practical limit for a simple Core/Edge fabric using 16-port core switches is 224 user ports. Use a SilkWorm 12000 or 3900 in the core if the Core/Edge fabric is expected to exceed 224 user ports.

Guideline: For fabrics that are expected to exceed 896 user ports while maintaining a 7:1 ISL oversubscription ratio, it is recommended to utilize four SilkWorm 12000 switches in the core. A Core/Edge fabric built with SilkWorm 12000 and 3900 switches is capable of scaling to 1792 user ports and yielding a 7:1 ISL oversubscription ratio.

The fabric topologies recommended are all Core/Edge topologies. The formats for these recommendations are templates. The SAN designer can use these templates to adapt to meet their requirements. The key is to map these requirements to a supported SAN design. The definition of whether or not a SAN design supportable is up to the support partner and there are several variables involved with determining this support (see [Scalability Support and Testing on page 2-43](#)). Six recommended fabric topologies are provided:

Highest Availability Fabric Topology: A fabric built with the most highly available switches, implementing a 7:1 ISL oversubscription ratio, and offering the highest fabric and I/O availability (see Figure 2-31)

Low Cost Per Port Fabric Topology: A fabric topology built with high ISL oversubscription ratios (low performance, fewer ports used for ISLs), resulting in a higher number of user ports (see Figure 2-32)

High Performance Fabric Topology: A fabric built with low ISL oversubscription ratios (high performance, more ports used for ISLs), resulting in a lower number of user ports (see Figure 2-33)

Very Large Fabric Topology: A fabric topology using four SilkWorm 12000 cores that is capable of scaling to 4352/3840 total/user ports (see Figure 2-34)

Small Fabric Topology: A fabric that uses smaller switches and is capable of scaling to 288/224 ports (see Figure 2-35)

Extended Distance Fabric Topology: A fabric topology recommended for connecting SANs that span multiple sites and where the availability of connections between sites is limited (see Figure 2-36)

A topology can be selected based on cost, performance requirements, availability requirements, or scaling requirements. Once a topology is selected, it is up to the SAN designer to determine the quantity and type of edge switches, locality model, and to customize the ISL oversubscription ratios.

Note: All Core/Edge topologies can start small (see Figure 2-30) and grow as large as the maximum port counts specified. In some topologies, the maximum port count is specified as a range, since this count is a function of the ISL oversubscription ratio implemented and the type of switch utilized.

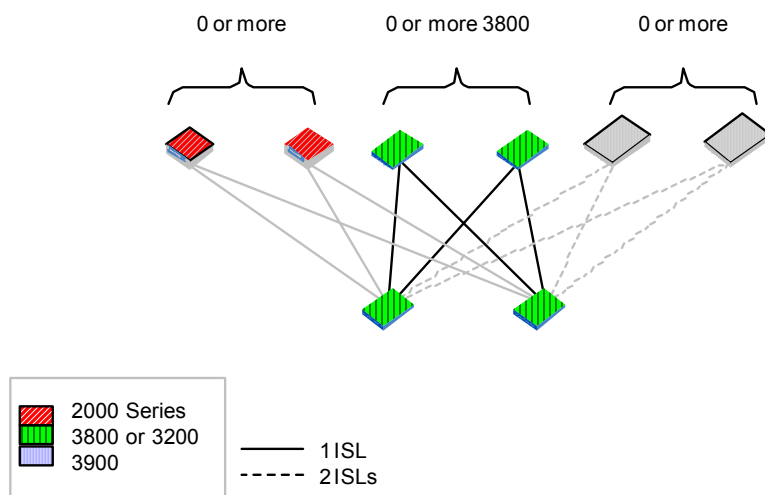


Figure 2-30 A Four Switch Mesh is the Start of Every Core/Edge Fabric Topology

Guideline: Regardless of selected Core/Edge fabric topology, it is recommended to implement the selected topology as a dual fabric SAN.

Note: Use these topology templates in conjunction with your support provider to develop a supported SAN design. The port counts indicate what the topology is capable of scaling to, not what a support provider will support.

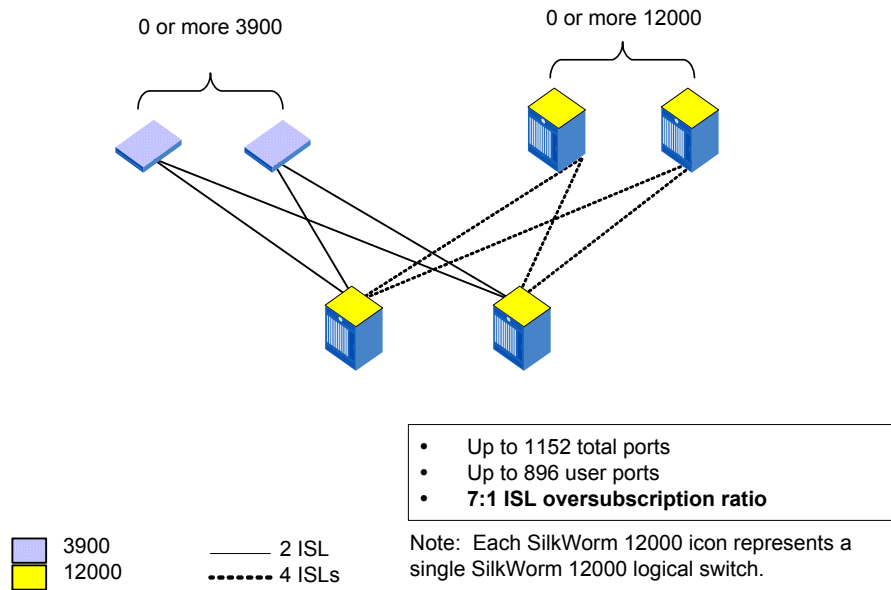


Figure 2-31 Highest Availability Fabric Topology

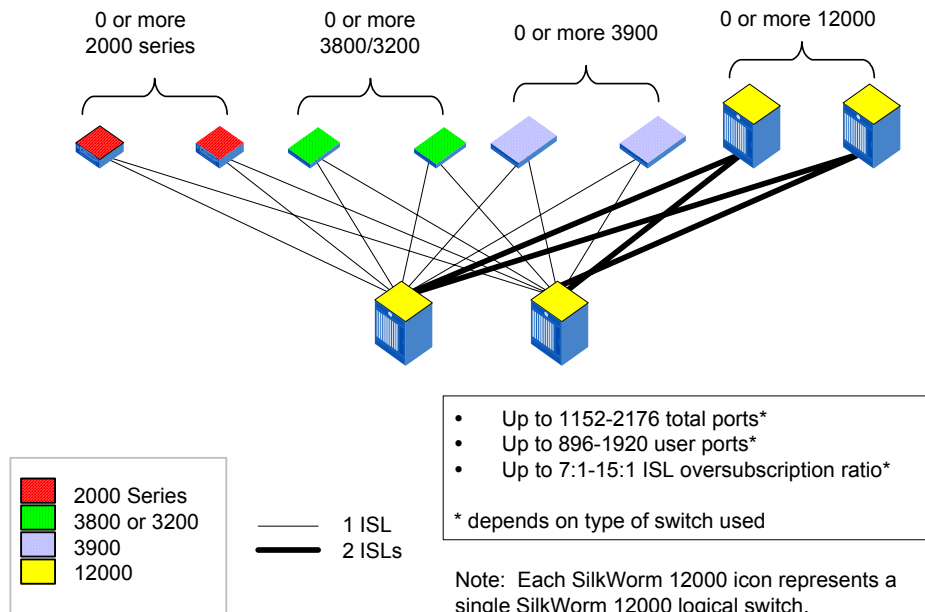


Figure 2-32 Low Cost Per Port Fabric Topology

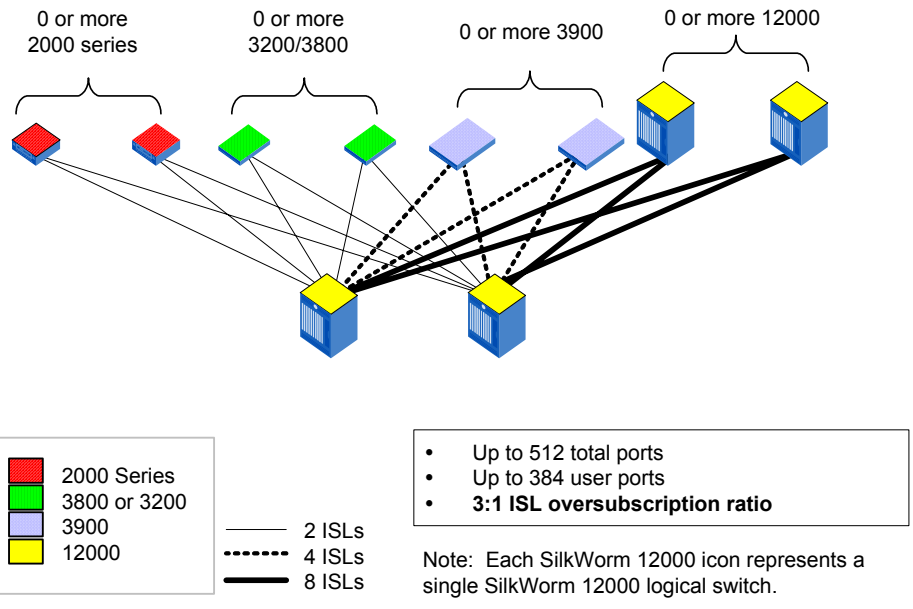


Figure 2-33 High Performance Fabric Topology

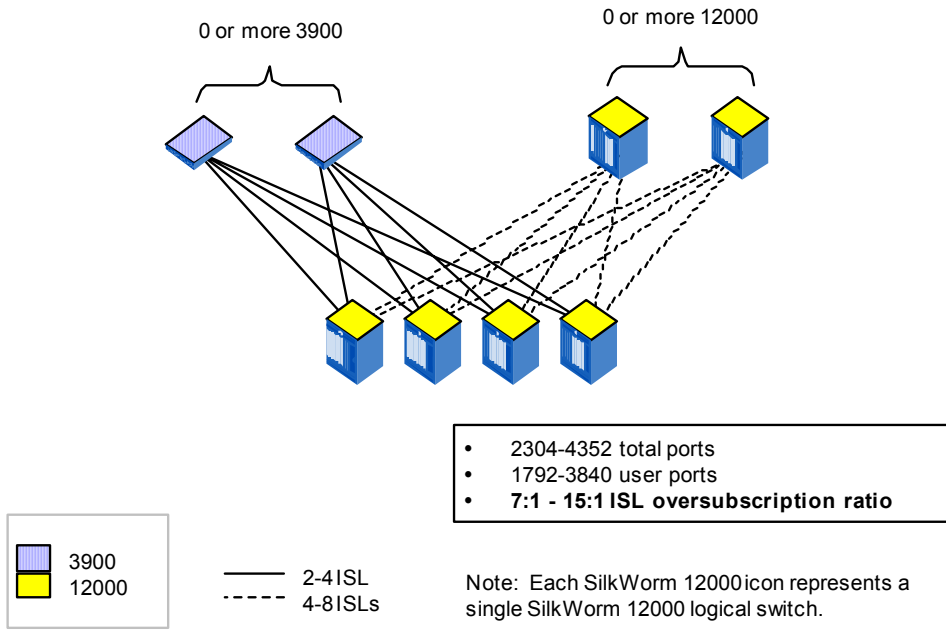


Figure 2-34 Very Large Fabric Topology

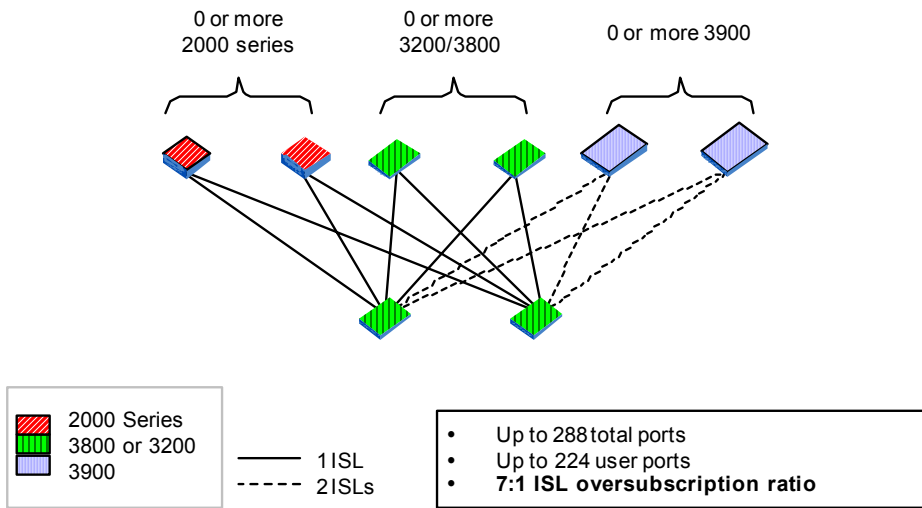


Figure 2-35 Small Fabric Topology




Figure 2-36 Extended Distance Fabric Topology

SAN Deployment

Once the SAN is designed using sound principles as detailed in [SAN Design on page 2-1](#), it needs to be deployed. This is more than plugging in cables, turning on the power and setting IP addresses. In reality, the SAN deployment process can be broken down into four discrete phases:

1. **Planning** - Proper planning allows for estimation of time and effort and provides justification for resources. In addition, a plan provides a means of measuring progress.
2. **Staging** - Staging is about doing the manual work of putting it all together. Staging covers everything from uncrating and racking the switches to configuring Brocade Fabric OS and the applications that will run on the hosts attached to the SAN.
3. **Validation** - The entire SAN configuration needs to be validated to verify it is ready for operation, once staged. The tests should verify device connections, check for the SAN robustness, and most importantly, test the application availability.
4. **Maintenance and Operations** - Once the SAN transitions to an operational state, changes are likely to occur, such as the addition of hosts or storage. This may require more switches if all user ports are allocated. The Fabric OS or other firmware and software may need to be upgraded.

Like any complex project, there are many different ways of doing a SAN deployment. Although complex, there are still general guidelines that should be followed. This document provides many of them in the form of checklists. Using these checklists allows sound decisions to be made and thus optimizes the IT investment. The checklists can be used as is or modified for specific environments. In addition to the checklists, many tips are pointed out. These tips help simplify the deployment process. Other ways of performing the activities discussed in this Deployment chapter are possible.

A case study has been created to illustrate the deployment process. This case study is an actual SAN designed and tested at Brocade. In addition to illustration purposes, the case study SAN provides a validation point for the concepts and procedures discussed in this chapter. Throughout the document, the case study SAN demonstrates key operational tips and points out important caveats.

It is crucial that the production SAN be supportable, manageable, maintainable and easily scaled. Proper planning and documentation is critical to make this a reality. Good up front planning makes the actual staging that much easier and the final production environment simpler to maintain. Effective documentation reduces the potential downtime, whether it be scheduled or not. Since these benefits, and others, are so important, there is heavy emphasis on them.

Besides planning and documentation, this section focuses on the new features and functions of Fabric OS 2.6.1/3.1/4.1, to assist the user who is familiar with Brocade products to quickly understand how to use these new features. New features and enhancements covered include Advanced Zoning, Trunking, Security, and Hot Code Activation (HCA). Other functions, such as firmware updates, are more user friendly. There are also quite a few new commands that allow for better diagnostic information. Guidelines as to which options to choose and under what circumstances to use them are covered as appropriate.

3.1. Planning

Having an effective plan prior to the staging of the equipment is critical for overall SAN deployment success. This success is measured in many forms. The greatest benefit of a good plan is that it gets the SAN deployed on time and within budget. Doing this ensures the ROI to be realized in the shortest time possible.

Planning is all about understanding the requirements and allocating necessary resources. This requires some extra up-front effort. In fact, for higher port-count SANs, it is critical that there be at least one person doing this. With a good plan, progress towards completion can be measured, the right persons identified, the roles and responsibilities defined, the site and SAN documented, and SAN resources can be efficiently utilized. With a good plan, less effort is needed to maintain the infrastructure. When the staging requirements are known, obstacles that may impede progress can be avoided.

This section provides some guidelines on the what essential information is needed for putting an effective plan together. Checklists are used throughout this section as a framework for gathering requirements. To clarify key points, the case study is used.

3.1.1. Site Environment Assessment

The first step is to prepare the site for the SAN. One effective way of doing this is to survey the site. The survey information provides a list of items that require completion before the hardware arrives. As a side benefit, once filled out, the site will then be documented. This is highly desirable in that if any issues arise, the pertinent information can be located quickly and the issue can be proactively solved. For a smaller port-count fabrics, the effort may be minimal. Even for smaller SANs, adherence to this process still provides good benefits, especially through the creation of the documentation.

Here are some questions to think about when defining a survey:

- Is sufficient power available?
- Are the appropriate cables and receptacles in place?
- Are patch panels required?
- Is there sufficient cooling?
- Is there sufficient rack space and footprint available?
- How should the SAN be attached to the LAN infrastructure?
- What management station will be used to administer the fabric? Is there is serial port available?
- What are the special site-specific policies?

It is critical to get the answers to these and other questions, so that when the Brocade fabric switches arrive, everything that is needed will be in place and the installation process can begin immediately and go smoothly. Knowing the requirements ahead of time really helps in planning the activities. As an example, storage switches would require extra security. This will prevent the highly sensitive data from being compromised. The switches, hosts and storage need to be locked down, to prevent unauthorized physical access. Now that this is known, the SAN Project Manager can plan for the space and budget.

3.1.2. SAN Project Checklist

Table 3-1 is an itemized SAN project checklist that provides guidelines on gathering the essential information for driving the project to completion. Following the checklist, the items that require further explanation are expanded upon. For an example of a complete survey form, refer to the Appendix for *Deployment Planning Templates*.

Table 3-1

SAN Project Checklist	
1.	Identify SAN Project Manager
2.	Site Contacts, Name, Location
3.	Site Preparation/Space Planning
4.	Project Plan <ul style="list-style-type: none"> • Site Specific Policies • Required Product Training
5.	SAN Equipment Assessment. Consider choosing a management switch.
6.	Appropriate licenses ordered
7.	Application Software Considerations
8.	Approval Sheet

1. SAN Project Manager

This is the most important item on the checklist. The SAN Project Manager co-ordinates the entire effort. The paper deliverable is a project plan. The co-ordination effort includes holding regular meetings, creating action items and driving the decision making process to resolve them. Many times, it also involves getting the right folks to talk with each other, and informing everyone of the plan and progress. Without this role clearly defined, the SAN project will quickly turn into chaos.

2. Site Contacts, Name, Location

Identify each person on the deployment team, their roles and responsibilities, and the site where the equipment will be staged. The SAN Project Manager should put together this list.

3. Site Preparation / Space Planning

The list of items should include (but is not limited to): site air-conditioning, power, fiber optic cable requirements (including location, lengths, etc.), any fiber patch panels, proposed rack layout, etc. As an example, the project manager may have to work with the site facility staff to make sure that the proper cooling vents and cut tiles are available for those locations that have cables run under raised floor tiles.

Once the equipment list is captured, having adequate space planned out is extremely helpful. By knowing where things go ahead of time, it becomes easier when talking about cabling, patch panels, DWDM attachments, etc. At a minimum, a high-level site schematic should illustrate the location of the Brocade SilkWorm fabric switches and other SAN-related equipment. Give consideration to including host, storage devices and other auxiliary equipment such as the power infrastructure and air conditioning as well.

Label each piece of equipment for quick and easy identification. A scale drawing is really helpful if a large number of racks and other equipment are required. Another good guideline is to create a rack layout diagram. This will illustrate all the equipment locations within the 42 U of space. For large deployments where the same equipment will be replicated at multiple locations, this becomes invaluable. The detailed information, like the port to which each device is attached, is not important at this time. This will be covered in a later section. It is important that all of the equipment is accounted for so that the other preparation steps can be completed.

4. Project Plan

A project plan is the tool used by the SAN project manager to get the SAN in place on time and within budget. The survey is used as a starting point to determine what is needed and when. Keep in mind any site-specific policies that may affect the installation. Include diagrams of proposed rack layouts and an installation schedule, which shows all of the dependencies. Training on the Brocade product family may be required for the staff who will own and operate the SAN. Put this in the plan as well.

When the project plan is complete, the SAN project manager will understand the staging and other requirements. Armed with this information, assigning and scheduling resources to meet the project milestones should be easily justified.

5. SAN Equipment Assessment

Find out whether there is an existing SAN in place or if this is a new installation. In the case of an existing SAN, a migration plan is required, especially if hosts are using the 24-bit PID address for persistent binding. For details about planning and implementing a migration, reference the *SAN Migration Guide* (publication number: 53-0000360-01).

As part of the assessment, generate a high level list of all fabric switches, hubs, hosts and storage. For the Brocade SAN Fabric switches, define the roles. As an example, consider choosing a management switch. This switch will be the management point into the fabric. Core switches are ideal for this choice. As for the other equipment, make sure the SFPs, GBICs, fiber optic cables, and any media converters are available. Include any DWDMs, Gateways, and other devices for connecting SANs over distance. Include any LAN Ethernet hubs/switches that require Brocade switch attachment. Refer to *LAN Guidelines For Brocade SilkWorm Switches* (publication number: 53-0000350-01) for effective LAN Integration guidance.

6. Appropriate licenses ordered

Make sure that the appropriate licenses are ordered given the customer requirements.

7. Application Software

Put together a high level list of the application software to be used such as: Database, CRM, E-mail, Web Services, and the associated host. This may drive switch and device placement within a data center. The objective is to make sure that the application storage requirements are taken into consideration, so that the storage placement is optimal.

8. Approval sheet

Once the site survey and SAN Project plan are complete, a substantial amount of the pre-work will be complete. For larger port-count SANs used in the enterprise, a survey and project plan becomes critical for successful deployment. SANs with smaller port-counts may not require an extensive survey, or even a robust project plan. Perhaps all that is required is to plan power and rack space. In all cases, it is important to focus on getting everything documented so that information is readily available.

Table 3-2

Site Prep Tips
<ul style="list-style-type: none">• Some equipment may have long order lead times, so it is important to identify what is needed as soon as possible.• For cables that will be run underneath a raised floor, not having enough cut tiles is a potential stumbling block.• Heavy equipment, like some RAID devices and the SilkWorm 12000, should be near bottom of the rack for maximum stability.

3.1.3. Planning for Power

This section provides some power guidelines to assist with the deployment planning of Brocade switch platforms. Please refer to the appropriate Brocade SilkWorm hardware reference manual for all the geo-regional power requirements. A list of all manuals is provided in Appendix B, *Reference Documentation*. Cable types also differ by region and are required to be a separate line item in the Brocade product order.

3.1.3.1. Notes for the SilkWorm 3200, 3800, 3900, and 2000 Series

The hardware reference manuals provide the information to select the appropriate cord for the amperage level that complies with local electrical code requirements. When ordering Brocade products, be sure to work with the local Brocade switch supplier to get the appropriate power cord set.

The SilkWorm 3800 and 3900 each have two power supplies. For maximum availability be sure both supplies are used. Each power supply is a FRU and can be hot swapped. The SilkWorm 3200, being a low cost 8-port product, has a single power supply. If it fails, the switch must be replaced.

Note: The power supplies are not interchangeable between the SilkWorm 2000 family of switches and the SilkWorm 3000 family, and the power supplies between the SilkWorm 3800 and SilkWorm 3900 cannot be interchanged.

Guideline: If only one power supply is used, the default settings of switch environmental policies will cause the switch to be in a marginal status. Use `switchstatuspolicyshow` at the command line to see the current environmental settings and `switchstatuspolicyset` to change them. Web Tools will show the switch image highlighted in an amber color as a warning. This setting can be configured administratively as discussed in [SAN Management on page 4-1](#).

3.1.3.2. SilkWorm 12000 Power Requirements

For the SilkWorm 12000, the voltage required is 200 to 240 VAC, 50 Hz or 60 Hz. Two dedicated branch circuits are required for redundancy. Confirm the power cords ordered with the system to ensure the correct power receptacles are installed.

- In the US and most of North America, the available voltage is usually 208 or 240 VAC, the receptacles are NEMA L6-20R, on individual branch circuits, (each rated 20 amps.)
- In UK, Ireland, Hong Kong, two UK-standard 13 amp receptacles, on individual branch circuits, are required.
- In most of continental Europe, two CEE7/7 “Schuko”-compatible receptacles are required, each rated 16 amps, OR two IEC60309, 230V~16A-6h receptacles. Verify with the system order.
- In Australia and New Zealand, 2 Australian-standard, 15 amp receptacles, on individual branch circuits rated 15 amps each, are required.
- For any location in which the US, UK, “Schuko” or AUS/NZ receptacles types are not accepted, the International standard, IEC-60309, 230V~, 16A-6h, receptacles should be planned, and the system order confirmed.

For proper power on the SilkWorm 12000, check for a steady green LED on all power supplies on the cable side of the unit. For all products other than the SilkWorm 12000, the single power LED should be a steady green. If it blinks or is a different color, this indicates there is a problem. The WWN card on the non-cable side of the chassis also displays power supply status as good with four steady green LEDs.

3.1.4. Cable Planning

One of the most important - but somewhat overlooked - deployment aspects is cabling. More than 50% of problems in IT involve a cabling issue. While there is no fail-safe way to prevent problems, this section contains some tips and guidelines to help reduce the risk.

All Brocade products use LC for 2 Gbit/sec and SC for 1 Gbit/sec type connections. For older hosts and storage that support only 1 Gbit/sec, an SC to LC cable is required. The switch will automatically detect the speed and provide the appropriate connection speed during link initialization.

3.1.4.1. Patch Panel Guidelines

Patch panels are great for providing a centralized method for managing host and storage connections to the SAN and other IT equipment. Here are a few guidelines for making effective use of them. Also pointed out are some common problems and solutions.

Guideline: Every fiber optic connection point generates a few dBs of signal loss. Keep the number of connections to a minimum, generally eight or less. If a connection seems intermittent, too many patched connections may be the cause.

Guideline: Be wary of distance limitations. The total distance can add up quickly as cables are run through conduits and patched. For 2 Gbit/sec connections, using shortwave SFPs and 50 micron cables, the maximum distance is 300 meters. Like too many patch panel connections, exceeding this distance limit may also cause an intermittent loss of signal.

Guideline: Avoid using older fiber optic cables with a 62.5 micron core. For 2 Gbit/sec connections, a 62.5 micron core *only* yields 90 meters as the maximum total distance. The newer standard fiber optic cable, 50 micron core, yields a maximum of 300 meters.

Guideline: At times, when plugging in a fiber optic cable from a patch panel, there may be a no light indication on the switch. This may indicate that a Tx and Rx swap has occurred. For Tx and Rx connections to be lined up properly between the switch and devices, a net $\frac{1}{2}$ twist is required along the connection path. Often times, the fiber optic cable connecting between patch panels adds an extra whole number plus $\frac{1}{2}$ number of twists, yielding a net of 0. One way to fix this is to swap the Tx and Rx on the LC or SC connector at the switch port. If using patch panels, be weary of this and use straight through fiber optic cables if at all possible.

Guideline: Do not mix single (longwave) and multimode (shortwave) cables in Patch Panels. The light wavelengths are different in each type. This make the fiber cables incompatible for connecting together.

Guideline: Do not patch together fiber optic cables with a 62.5 micron core to fiber cables with a 50 micron core.

Keeping these tips in mind should allow for more effective and reliable use of a patch panel infrastructure for connecting the hosts and storage.

3.1.5. Cable Management

Proper cable management reduces the cost of SAN maintenance and helps connectors last longer, as there is weight relief at the point where the cable sheath joins at the connector (see Figure 3-1). For fiber optic cables, which are made of high quality glass, proper management reduces the risk of link failure by preventing the cable from exceeding the maximum bend radius.




Figure 3-1 Cable management provides weight relief for the cable sheath at the connector.

Proper cable management facilitates easy cable identification, which is important for accommodating the growth of the SAN. Good cable management is an art. Everyone has a different style that is applied to the site's unique structural characteristics. Thus, there is no single solution for proper cable management. There are endless variations on doing it right. The next few sections will highlight some general guidelines on managing cable connections that can be applied to almost every situation.

3.1.5.1. Cable Management on SilkWorm 2000 and 3000 Series Switches

When racking SilkWorm 2000 and 3000 series switches, it is highly recommended to have space to alleviate the bending of the fiber optic cable connector sheaths while they are plugged into the switch ports. Effective cable management prevents this problem, which occurs over time, as the weight of the cable not only bends the connector but can loosen the connector-cable attachment. If this is not managed, eventually there will be signal loss that could cause loss of frames and, potentially, data corruption. Cables that obstruct the face of the switch make it more difficult and time consuming to replace them. Good management prevents this. As a side benefit, good cable management also makes it easier to identify the cables and to read each label. Figure 3-2 illustrates poor cable management and shows the difficulty of replacing the SilkWorm 3900.




Figure 3-2 Improper Cable Management

Use cable guides to solve these problems. There are horizontal and vertical cable guides for standard EIA racks. Horizontal cable guides come in a variety of shapes and generally take up at least one rack unit. Vertical cable guides are generally much wider and up to 42 U in height, which is the standard size of an EIA rack. These generally are mounted along the side, so extra floor space is required for them. Telco racks have their own methods of management. Be sure to ask the Telco rack vendor for cable management options. Both types have “fingers” that allow the cables to be held while being run across the rack or plugged into a switch port. This may be expensive when considering the value of rack space so the customer may choose not to do it. However, the organizational benefits are enough to justify the expense alone. The next picture shows how effective cable guides are at achieving proper management.

Guideline: Use horizontal and vertical cable guides for cable management. Include them in your rack planning.

Guideline: Plan out the floor space and install the guides before staging the racks with the equipment. Often times, the cable guides are secured to the racks themselves. Once the equipment is plugged in and the cables run, installing cable guides in the racks become much more costly as power cables, racks, etc. all have to be moved or removed.

Guideline: If using seismic platforms, allow enough cable slack for effective movement.




Figure 3-3 Proper Cable Management

Note that the cables in Figure 3-3 are lined up nicely and the strain on the connectors is minimized. The cables are not blocking the switch face and ports. If a switch needs replacement, there is no need to unplug unnecessary cables from surrounding switches or the switches themselves. Downtime is reduced, and money is saved.

Guideline: For Trunk groups, use Velcro tape and spacers (sometimes called pillars) to bundle the cables into groups. Not only does this alleviate the weight, it also allows for effective cable organization. As an added plus, trunk group identification is easily attained when this is accomplished.

Guideline: Rack the cable side of the switches on the same side where the HBAs and storage ports are located. Typically these ports are located on the “back” of those devices. This will prevent cables having to be run along the inside and side of an EIA rack. Be weary of potential cooling issues when doing this.

Guideline: Label the fiber optic cables or use fiber optic cables that are already numbered, such as with serial numbers. This information can be used to build a spreadsheet of devices, switch ports and the cables that connect them together.

Guideline: Allow manageable cable slack for sliders on rack mount kits.

Guideline: Cable guides should always be used in patch panels. This is because patch panels are a natural cable collection point and the scene of many potential problems.

Figure 3-4 illustrates an effective use of cable trays.




Figure 3-4 Patch Panel With Cable Guides

Note: Some of the cable trays are open to show how the cables are run.


3.1.5.2. *SilkWorm 12000 Cable Management Planning*

A prime consideration of cable management for the SilkWorm 12000 is the ability to remove and replace the field-replaceable unit (FRU) components without having to unplug cables. Since the port cards, CP cards, power supplies, and blower assemblies are all hot-swappable, ensuring easy access to each component is vital toward contributing to the uptime of the SAN. In addition, it is important that the LEDs on the individual components and the WWN card remain visible.

If Trunking is being used, the ports and cables used in trunking groups must meet specific requirements. For a list of these requirements, refer to the *Brocade ISL Trunking User's Guide*. Figure 3-5 shows an example of effective cable management, with no cables crossing in front of the FRU and an example of how not to cable a SilkWorm 12000. Note that any single FRU is impossible to replace without have to unplug devices and/or other switches.



Proper Cable Management



Improper Cable Management

Figure 3-5 Proper and Improper SilkWorm 12000 Cable Management

3.1.5.3. *SilkWorm 12000 Cabling Guidelines*

Guideline: Leave at least one meter of slack for each fiber optic cable. This provides room to remove and replace the port card, allows for inadvertent movement of the rack, and helps prevent the cables from being bent to less than the minimum recommended bend radius.

Guideline: Route fiber optic and other cables down along the front of the card to which they are connected, to prevent having to disconnect them when neighboring cards are replaced. Do not route across adjacent cards or in front of the power supplies.

Guideline: Use the Cable Pillars provided with the rack kits to bundle the fiber optic cables from each quad of ports. These guides help to keep individual ports accessible by keeping the cables evenly spaced, and also help to provide clearance for the replacement of a port card or CP card.

Guideline: Use Velcro wraps (not included with the switch) to further bundle the cables on a per-card basis. Tie wraps are not recommended for optical cables because they can easily be overtightened, damaging the optical fibers.

Guideline: Use the cable management tray to route the bundled fiber optic cables, Ethernet cables, and any serial cables down the front of the chassis.

Guideline: The power cord requires a minimum service loop of six inches at the switch to ensure enough freedom of movement to plug and unplug it.

Guideline: Route the power cables to each side of the switch instead of through the cable management tray. The power cable connectors are designed with right and left bends to facilitate cable management.

Guideline: Label the fiber optic cables or use fiber optic cables that are already numbered, such as with serial numbers.

Guideline: Use a spreadsheet to track which devices are connected to which switch, port card, and port in the SilkWorm 12000. The serial numbers or other identifiers of the fiber optic cables can also be tracked.

Guideline: Keep the chassis door on the SilkWorm 12000 closed to protect cables from inadvertent movement.

3.1.6. The Rack Layout Plan

When it comes to racking, usually some creativity is required. There is no single correct technique. The following section presents some guidelines to keep in mind when planning and racking the switches and other equipment.

From the site survey, get the equipment list and put together a spreadsheet showing the number of required rack units (U) that are required for each device. Include switches as well as any hosts and storage. Once again, there is no one right way of doing this. Here is one template that has proven effective with the case study equipment shown. Table 3-3 shows the essential information for a standard 42 U rack. If desired, a scale diagram can be drawn to show the actual rack space constraints.

Table 3-3 Sample Rack Layout Spread Sheet Template.

Device Name	Number of Devices	Rack Unit per Device	Total Used Rack Units
Fiber Patch Panel	1	3	3
KVM Unit	1	3	3
Host 1	1	1	1
Host 2	1	3	3
Cable Guide	1	1	1
SilkWorm 2800	2	2	4
Cable Guide	1	1	1
SilkWorm 3800	2	1	2
Cable Guide	1	1	1
SilkWorm 3900	1	1.5	1.5
Cable Guide	1	1	1
SilkWorm 12000	1	14	14
SW 12000 Cable Guide	1	2	2
Total Units Available			42
Total Units Used			35.5
Total Units Remaining			6.5

3.1.6.1. Racking Guidelines

There are two styles of racks, EIA and Telco. EIA racks typically have four posts, while Telco racks have only two. All Brocade switches are available with EIA rack kits as a separate option as part of the Bill of Materials (BOM). Telco racks typically provide better cooling potential as there is typically more open space. If Telco racks are used, contact the Telco rack vendor for appropriate shelving. Assuming sufficient cooling and airflow exists, which should have been determined by the site survey, then there is no physical need for having extra space between each rack mounted switch.

Brocade SilkWorm switch airflow is from the non-cable side to cable side. When racking switches, totaling ten or more rack units, make sure sufficient air flow is available to cool the switches. This can be done by spacing the switches appropriately or using a rack fan. Extra rack space is always desirable from a cooling standpoint, as more surface area will be exposed. However, a better reason for having extra space between the switches is to allow for proper cable management.

Here are some guidelines to consider when racking Brocade SilkWorm Switches. For additional SilkWorm 12000 racking tips, please see the *SilkWorm 12000 Design, Deploy and Management Guide* (publication number 53-0000251-03).

Guideline: Avoid crossing over rack unit boundaries. Count them before installing the equipment. Some EIA racks have rack units pre-numbered, order these if available.

Guideline: For larger chassis, such as the SilkWorm 12000, create a hole schedule template to assist with planning the rail locations. This may be good idea if a variety of different sized equipment is installed.

Guideline: When racking, use cage nuts with built in threads for the screws. To facilitate installation, mark the unit boundaries on the rack before installing the nuts.

Guideline: It may be simpler to rack the cable side of the switches on the same side where the HBAs and storage ports are located.

Guideline: Consider ordering pre-numbered fiber optic cables.

Guideline: Cable guides should always be used in patch panels.

Guideline: When racking switches, totaling ten or more rack units, make sure sufficient air flow is available to cool the switches. This can be done by spacing the switches appropriately or using a rack fan.

Guideline: Forced airflow is required when using solid doors in the rack, such as those made of Plexiglas.

3.1.7. Documentation Guidelines

This section will provide some guidelines regarding documenting the SAN. Knowing what documentation is needed allows for planning IP addresses, switch domain numbers, what ports should be used for ISLs, etc. Once created, having the documentation readily available allows all components to have good change management. In addition, being organized with updated documentation saves time and effort when referencing equipment for service calls. Even while not being serviced, when information is needed, it can be referenced quickly and easily. Here is a checklist that provides a recommendation as to what documentation should be created before and during the staging of the equipment. Refer to the Appendix for *Deployment Planning Templates* for templates that can be used to assist in documentation.

Table 3-4 Documentation Checklist

Documentation Checklist	
1.	Get an Equipment Binder
2.	Logical Design Diagram
3.	Switch Spreadsheet
4.	ISL Port Map
5.	Device Spreadsheet
6.	Label All Cables
7.	SAN Verification Test Plan Requirements
8.	SAN Verification Test Plan

1. Equipment Binder

One good practice is to put together a binder for each rack. Keep all of the configuration information in the binder, including the rack layout, make and model of each piece of racked equipment, serial numbers, physical cable connections, the names on the labels, etc. As part of the binder, keep all of the service logs to track that history for each piece of equipment. If at all possible, keep this attached to the rack physically.

2. Switch Spreadsheet

Keep track of all Brocade switches. Include switch name, IP address and domain name as well as the role. An example of a spreadsheet is shown in Figure 3-5. This may be reflected in the switch name, for example A-C1, designates this switch as the first core switch in fabric A.

Guideline: Set unique domain numbers for each switch in the SAN. This allows for simpler fabric merges.

Guideline: As a convention, consider setting the domain ID of each switch to the last octet of its IP address.

Table 3-5 Switch Device SpreadSheet

Fabric A			Fabric B		
Switch Name	Domain	IP Address	Switch Name	Domain	IP Address
A-C1	100	192.168.122.100	B-C1	110	192.168.122.110
A-C2	101	192.168.122.101	B-C2	111	192.168.122.111
A-E1	102	192.168.122.102	B-E1	112	192.168.122.112
A-E2	103	192.168.122.103	B-E2	113	192.168.122.113
A-E3	104	192.168.122.104	B-E3	114	192.168.122.114
A-E4	105	192.168.122.105	B-E4	115	192.168.122.115
A-E5	106	192.168.122.106	B-E5	116	192.168.122.116
A-E6	107	192.168.122.107	B-E6	117	192.168.122.117

3. ISL Port Map

Create a port map that shows how each switch is interconnected. This is handy for tracing cables and can be used to reproduce the configuration easily. Figure 3-6 provides an example.

Table 3-6 ISL Port Map

Fabric A			Fabric B		
From	Cable	To	From	Cable	To
FabA-E1-15	I1	FabA-C1-15	FabB-E1-15	I25	FabB-C1-15
FabA-E1-14	I2	FabA-C1-14	FabB-E1-14	I26	FabB-C1-14
FabA-E1-13	I3	FabA-C2-15	FabB-E1-13	I27	FabB-C2-15
FabA-E1-12	I4	FabA-C2-14	FabB-E1-12	I28	FabB-C2-14
FabA-E2-15	I5	FabA-C1-13	FabB-E2-15	I29	FabB-C1-13
FabA-E2-14	I6	FabA-C1-12	FabB-E2-14	I30	FabB-C1-12
FabA-E2-13	I7	FabA-C2-13	FabB-E2-13	I31	FabB-C2-13
FabA-E2-12	I8	FabA-C2-12	FabB-E2-12	I32	FabB-C2-12

4. Device Spreadsheet

The other piece of information that should be included is a device spreadsheet from the host perspective. Include hostname, IP address, HBA make model and WWN, host and storage physical connections, storage LUN assignment map etc. For an example of a device spreadsheet, refer to the *Planning Templates Appendix*.

5. Logical Design Diagram

Create a picture of the logical design. Keep it simple and be careful of adding too much information. The idea is to represent the SAN topology as well as host and storage connections. As an example, see Figure 3-7, which depicts the case study.

6. Label All Cables

Label all of the switches as well with IP addresses and/or switch host names. The SilkWorm 12000 has four IP addresses and requires up four hostnames, one per logical switch and one per CP. Place a label on the side of the chassis for quick and easy access. Pick a naming convention, keep it simple and be consistent. As for information, ISLs should have the destination switch name and port number, for example for a SilkWorm 3800 connected to a SilkWorm 12000 on sw1 physical slot 7 and port 15, use “sw12000 7/15”. Devices should have its name (IP hostname works well), the fabric plugged into and port number, for example “int124 A14”. For the SilkWorm 12000 use the slot/port designation. These can be handwritten on stickers or generated by a labeling machine. Use the same techniques for other types of cables. Once complete, create a spreadsheet with all of this info and put into a binder with all of the other documentation.

7. SAN Verification Test Plan Requirements

At this stage it is important to plan for validating the SAN. What are the requirements for this? Well, it really depends upon the specific SAN application and its purpose. Start thinking about this ahead of time so that a test plan can be put in place.

8. SAN Verification Test Plan

Once the requirements are understood, put together a verification test plan. If at all possible (especially for enterprise class SANs) it is highly recommended to do this with a proof of concept test SAN that is not in production. Focus on testing the software application availability as this what the SAN is being used for.

3.1.8. Zoning Plan

Zoning allows the hosts to access specific storage devices on the SAN. For those SANs with multiple OS platforms, zoning allows for OS separation and co-existence. With no zoning defined on the SAN, any device can see any other device. This is the default setting. Once zoning is in place, all devices must be members of a defined zone. Those devices that are not will be blind to all others. This section will provide some guidelines as to zoning plan definition. For additional information reference the *Brocade Zoning User's Guide*.

Zoning requires careful thought and planning. Armed with the documentation created in the previous section, and understanding the requirements, allows the creation of a zoning plan. Creativity is important here as there is no one “correct” zoning configuration for a given SAN fabric configuration. In general, follow any specific zoning recommendations provided by the switch vendor.

Table 3-7 Zoning Plan Checklist

Zoning Plan Checklist	
1.	Gather the list of host and storage devices to be zoned from the device spreadsheet.
2.	Define the storage requirements for each host based upon software application requirements.
3.	Adhere to recommended storage device configurations such as LUN masking, LUN security, and other specific features supported by the vendor.
4.	Consider specific host requirements for storage value-added feature sets such as Server Free backup, LUN snapshots, or LUN mirroring over distance.

1. List host and storage devices to be zoned

Gather the list of host and storage devices to be zoned from the device spreadsheet.

2. Consider specific host requirements -

Consider specific host requirements for storage value-added feature sets such as Server Free backup, LUN snapshots, or LUN mirroring over distance.

3. List the storage requirements for each host -

Define the storage requirements for each host based upon software application requirements.

4. Define the recommended storage device configurations -

Adhere to recommended storage device configurations such as LUN masking, LUN security, and other specific features supported by the vendor.

Some key points when planning the zoning:

- Clearly understand the storage requirements for each host. This means understanding specifically what storage is needed for the software application that will be in production. To understand this, the number and size of LUN presentations on each storage array Fibre Channel port must be clearly defined.
- Be sure to adhere to recommended storage configurations by the switch vendor for LUN masking and other storage specific features. Some vendors may recommend using a separate HBA for tape devices. For this case, define the zoning configuration to fence off tape devices from any other HBAs, which see the disk storage, in the same host.
- Keep in mind the different OS platforms, backup application requirements and the number of paths to each LUN which may drive the zoning plan. There may be specific host requirements for storage value-added feature sets such as Server Free backup, LUN snapshots, or LUN mirroring over distance.
- As a general rule, have overlapping zones in all cases. An overlapping zone has the HBAs share one more storage ports, but with the HBAs separate from each other. This is sometimes referred to as a single initiator zone. For specific guidelines please see [Table 3-8 on page 3-19](#). Also, be aware that device placement is critical for WWN based hard zoning. Refer to [Zoning Enforcement Notes on page 3-18](#) for details

3.1.8.1. Zoning Guidelines

Guideline: Use persistent binding on the host. This will provide consistent controller, target and LUN numbers for each storage LUN. Backup applications are especially sensitive, as these numbers map directly to the backup application device identities.

Guideline: If using WWNs, zone by World Wide Port Name (WWPN) rather than World Wide Node Name (WWNN). This is because a WWPN uniquely identifies a port to which a target is attached. Some Multipathing software may get confused and not be able to discover targets properly. This is especially true when using multi-port HBAs.

Guideline: Be wary of mixing different HBA vendors in a single zone. Each vendor HBA responds differently to RSCNs, a method to notify an HBA for device discovery, and may cause one of the HBAs to lose the zoned device.

Guideline: In addition it is recommended to have single initiator zones, that is one HBA per zone.

Guideline: Separate HBAs from each other for clustered hosts. Allow each HBA to see the same storage but not each other. Once again, RSCNs, may cause the clustered host HBA to lose the storage array.

3.1.8.2. Zoning Enforcement Notes

All zoning employs the Name Server to limit the information returned to an initiator in response to a Name Server query. This is referred to as soft zoning. When hardware enforced or hard zoning is active, the Brocade switch will monitor Fibre Channel traffic and block any frames that don't comply with the effective zone configuration on the port to which targets are attached. Since all hard zoning is enforced at this location, only targets attached to the SilkWorm 3X00 and 12000 switches support WWN based hard zoning. This is due to hardware limitations in the older SilkWorm 2x00 switches. For these older switches, hard zoning can only be accomplished with Domain/Port numbers. Thus to do hardware zoning in a mixed fabric, use WWN for the SilkWorm 3X00 and 12000 switches and Domain/Port for the SilkWorm 2X00. For more detail, refer to the whitepaper titled: *Zoning Implementation Strategies for Brocade SAN Fabrics*.

3.1.8.3. Zone Port Map

It is a good idea to create a zoning port map. This shows which hosts and storage belong in a particular zone. Using color is really effective in displaying zone membership. For example, Table 3-8 shows the HPUX zone in blue, the Solaris zone in yellow and un-zoned storage in white. The device aliases are labeled as shown. For larger fabrics this approach has its limitations due to the color spectrum limits. For this case, build out a table that checks off each device in a zone. This may or may not be part of the device spreadsheet. For really large fabrics, a large database that houses each device and its zone membership. With this method, any zone can be selected out of the table space. If possible, include the last 4 or 6 digits of the WWPN, this will help in identifying unique devices that will require zoning.

Table 3-8 Partial Port Map and Zoning Table for a SilkWorm 12000 Switch

Brocade SilkWorm 12000 Zones		
Port	Blade 1	Blade 2
15	STORAGE-1A	STORAGE-1C
14	STORAGE-1B	STORAGE-1D
13	STORAGE-1G	STORAGE-1J
12	STORAGE-1H	STORAGE-1K
11	SOLARIS 1-0	SOLARIS 1-1
10	SOLARIS2-0	SOLARIS2-1
9	HPUX1-0	HPUX1-1
8	HPUX2-0	HPUX2-1
7	AIX1-0	AIX1-1
6	AIX2-0	AIX2-1

3.1.9. Planning the Upgrade of Fabric OS 2.x/3.x/4.x to 2.6.1/3.1/4.1

This section will provide some high level guidelines when defining a strategy for doing a firmware upgrade of Fabric OS on an existing Brocade fabric. This activity typically happens in the maintenance phase of a SAN deployment. For the purposes of this discussion, it is assumed that no new switches will be added or removed from the existing SAN infrastructure. For that case, and more details, please see the *SAN Migration Guide* (publication number: 53-0000360-01). For detailed instructions on all upgrades, refer to the *Fabric OS Procedures Guide* for the specific version of Fabric OS used on the switch. Be sure to upgrade to the firmware qualified or recommended by the switch provider.

Each Brocade-based SAN is unique. This is due to the wide variety of OS platforms, HBAs and storage arrays that may be attached. This uniqueness means that each upgrade needs to be carefully planned to minimize the risk of unscheduled downtime. *All updates* to Fabric OS 2.x and 3.x are and will be disruptive. This includes updates to Fabric OS 2.6.1 and 3.1 and later versions. Scheduled downtime is required for single fabric SANs containing switches that run those firmware versions. For dual fabric SANs, upgrade one fabric at a time. For Fabric OS 2.x or 3.x, either RSH or the FTP protocol can be used to execute the upgrade. Unlike Fabric OS 4.x, there is no `firmwaredownloadstatus` command. This is due to the limitation of a single telnet session for Fabric OS 2.x and 3.x.

For Fabric OS 4.x, only the FTP protocol is supported for firmware upgrades. Upgrades from 4.0.x to 4.1 ARE disruptive so be sure to schedule downtime. Upgrades from Fabric OS 4.1 to 4.1.x and beyond will be non-disruptive by default. The SilkWorm 3900 allows two admin telnet sessions. The SilkWorm 12000 allows two telnet sessions per logical switch. Use one for `firmwaredownload` and the other for `firmwaredownloadstatus`. `Firmwaredownloadstatus` is a handy command that shows a log of each upgrade phase. When complete, use `firmwareshow` to display the firmware version on each compact flash partition.

Table 3-9

Upgrade Planning Checklist	
1.	Analyze the potential risks and impact to each device on the SAN.
2.	In order to maximize fall back capability, preserve each fabric switch configuration with configupload.
3.	Use Fabric Manager for larger multi-fabric SANs.
4.	Verify the upgrade version is supportable.
5.	Gather the documentation and readme notes for the firmware release.
6.	Schedule downtime for single fabric updates of 2.x and 3.x and 4.0.x to 4.1.
7.	For dual fabrics, update one fabric at a time.

3.1.10. Planning Principal Switch Placement

One of the new features of Fabric OS 4.1 is the ability to hard set a preferred principal switch in the fabric. For core-edge topologies, the principal switch should be a core switch for optimal fabric operation. If using a SilkWorm 12000 in the core, pick it as the principal, as it has the greatest availability. Refer to *Switch Location In The Fabric* on page 3-38.

Warning: Hard setting a preferred principal switch is not completely deterministic, especially in large fabrics. Secure Fabric OS also affects its reliability. Use of sequenced reboot feature of Fabric Manager does mitigate these two limitations. After an un-managed reboot (such as a power failure recovery) do a sequenced reboot under Fabric Manager control to ensure an orderly switch-by-switch reboot.

3.1.11. Planning for Secure Fabric OS Security Measures

There are some SAN Security measures that should be in place before implementing Secure Fabric OS (SFOS). Here are some guidelines in the form of a checklist to assist with the planning process. These steps can be taken to provide some initial restrictions on accessing the SAN and to provide some control over change management. Best of all, there are no additional licenses required. For maximum SAN Security, these measures should be used in conjunction with Secure Fabric OS. Secure Fabric OS provides a single point of management and policies which allow complete control over what switches, devices and management stations are allowed to access the SAN.

Table 3-10

Planning for Secure Fabric OS Security Measures Checklist	
1.	Prevent Physical Access
2.	Prevent Remote Access through IP security measures
3.	Hard Zone the devices
4.	Lock Down E_port creation with portCfgEport.

1. Prevent Physical Access

Use a cage or some other method to only allow authorized personnel to access the switches physically. This is important in that serial port access is a real security risk.

2. Prevent Remote Access through IP security measures

Follow the IP-based security policies for the Brocade Ethernet port attachments and the IP Subnet they are located on. For example, lock down IP access to the switches by putting them on a separate (Virtual LAN) VLAN segment with separate IP router ACLs. This prevents illicit access on the switch Ethernet connections.

3. Hard Zoning

Hard zoning works in much the same way as a limited Access Control List (ACL). It works by restricting the hosts and storage access. For a mixed Operating System environment, this is generally required anyway.

4. Lock Down E_ports

This limits the creation of switch-to-switch ports. Use `portCfgEport` to lock down all ports except those that are to be ISLs. Note that if Trunking is used, and additional bandwidth is required in the trunk group, the port must be re-enabled to allow E-Port configuration. An alternate method is to use the command `portCfgPersistantDisable` to persistently disable a port. See the section titled: *New Persistently Disabling a Switch or Port (Fabric OS 3.1/4.1)* on page 3-73.

```
switch:admin> portCfgEport 3, 0
Committing configuration...done.
switch:admin> portCfgEport
Ports:   0   1   2   3   4   5   6   7
-----
        -   -   -   NO  -   -   -   -
```

Figure 3-6 Using PortCfgEport to disable E_port creation

3.1.12. Secure Fabric OS Planning

Planning for SAN security and change management is important. Organizations understand that the data managed in their SAN environment is often highly sensitive and must have controlled access properly to ensure confidentiality, integrity, and availability. A compromise in any of these areas could have unintended consequences, resulting in the loss of proprietary information, capital, or other core business resources. Proper SAN Security planning with Secure Fabric OS (SFOS) mitigates these risks by ensuring the right SAN security access controls are in place and enforced. Because SAN security is, by its own right a separate subject, a comprehensive treatment will not be discussed in this document. In order to be effective in the implementation of Secure Fabric OS, there are two assumptions that are made for the duration of the discussion. One is that significant non-SFOS security measures are already in place. And two, good security practices exist within the IT infrastructure.

Rather than talk about broad security objectives, the goal of this section will be to provide an overview of SFOS and provide essential planning checklists. These planning checklists and other guidelines can be applied to any organization wishing to implement SFOS and should be used as a foundation for tailored SFOS planning. To clarify important steps, the case study SAN will be referenced as needed.

For all the details on Brocade SAN Security background information, theory, and SFOS implementation details please refer to the *SAN Security: A Best Practices Guide* (publication number GA-RG-250-00). This is an excellent document that covers current industry Security practices as well as illustrates practical examples on SFOS features usage.

Secure Fabric OS is a separately licensed Fabric OS (FOS) product that provides a comprehensive SAN security solution for Brocade fabric member switches and the devices that are attached to them. All SilkWorm switches are supported *except for* the SilkWorm 1000 series of products. To use SFOS in a mixed environment, the minimum firmware version must be Fabric OS versions 2.6.1, 3.1, or 4.1. Once the SAN is properly prepared, it is just a matter of enabling security mode and implementing the previously planned switch roles and policies that have been planned out in advance.

3.1.12.1. Secure Fabric OS Concepts

Secure Fabric OS functionality falls into five basic areas:

- **Fabric Configuration Server (FCS)** provides a centralized way to manage fabric-wide configurations and policies.
- **Management Access Control (MAC)** adds additional layers of granularity when enforcing which devices can access SAN switches by way of which applications.
- **Secure Management Channel** provides a more secure method for running management applications that use encrypted passwords and certificates for authentication.
- **Switch Connection Control (SCC)** improves switch-to-switch authentication by allowing the use of digital certificates as well as locking down which ports can become E_port.
- **Device Connection Control (DCC)** allows only specific devices into the fabric (per their WWNs) from a specific port or group of ports.

Each of these areas are enforced by using policies. If familiar with Windows or UNIX security related administration tasks, the concept is very similar. The best way to do is to use the command line. When complete, the security information is stored in several databases distributed and enforced by the Primary FCS switch.

3.1.12.2. Secure Fabric OS Switch Roles

The implementation of secure mode on a fabric requires grouping the switches logically into three areas:

- **Primary FCS Switch:** This label applies to a single, uniquely powerful switch that is the sole owner of read/write privilege for fabric-wide operations. Design criteria for selecting this switch include:
 - The switch that is in the most secure, best controlled physical location (typically not at a remote office)
 - The most robust switch in the fabric
 - A core switch that is physically near the largest number of switches in the fabric
- **Backup FCS Switches:** One or more switches that can become the Primary FCS Switch if it becomes unavailable. All FCS switches conform to a conventional order, in which the first switch is the Primary FCS Switch, the second switch is the first backup FCS switch to take over in the event of a failure, and so on. Automatic failover occurs to the first backup FCS switch. Manual failover can be initiated by the SAN administrator to any backup FCS switch defined in the FCS list. These switches do not have the ability to make changes to fabric-wide configurations unless they become the Primary FCS Switch.
- **Non-FCS Switches:** This third class of switches encompasses all the remaining switches in the fabric. Any device not designated as an FCS switch type simply functions as a member switch that will never have the ability to modify fabric-wide configuration parameters.

Think of the Primary and Backup FCS switches as members of a trusted switch group. The Primary FCS (Trusted) switch becomes the management point for fabric-wide configuration changes. The Backup FCS switches are trusted members of the fabric and are there to provide redundancy in the event the Primary FCS switch goes down. To maintain Secure Fabric OS policy settings, the Primary FCS switch is solely responsible for applying configurations, which are contained in four separate databases, centrally managed on it. One of these contain fabric-wide passwords for each user account. A temporary password can be assigned to a single switch for maintenance purposes without having to give away the FCS password. There is a separate password list for FCS and Non-FCS switches. Non-FCS switches have only admin and user level account access. The Non-FCS switches are just members of the fabric with policies enforced by the Primary FCS switch.

In order to join a SFOS fabric, new switches must have their digital certificates authenticated. Brocade has developed an additional protocol, called Fibre Channel Authentication Protocol (FCAP) that is used as part of E_port link initialization when Secure Fabric OS is turned on.

When enabling secure mode for the first time on the primary FCS switch, be aware that each switch in the fabric must be rebooted for the Primary FCS switch to propagate all of the information under its control to the backup FCS switches. The Primary FCS switch propagates the following fabric-wide configuration databases as shown in Table 3-11.

Table 3-11 Propagated Primary FCS Configuration Databases

Propagated Primary FCS Switch Configurations
1. Security Polices
2. Zoning Configuration Database
3. Fabric Password Database
4. SNMP Community Strings
5. Date and Time using the primary FCS Time Server

Guideline: For core-edge topologies, such as the case study, pick a core switch as the Primary FCS. Use a SilkWorm 12000 core fabric switch as the Primary FCS or trusted switch if available. It has the maximum uptime, and with Fabric OS 4.1 much greater availability.

Warning: All the items in Table 3-11 **will be wiped out on switches that are allowed to be added** to a SFOS enabled fabric. It is highly recommended to **NOT** add switches that contain zoning information, before making a backup with `configupload`, since all zoning configurations will be written over.

3.1.12.3. Secure Fabric OS Default Policy Settings

With Secure Fabric OS (SFOS) enabled, the default policy settings allow for any switch or device attached to fabric access to any resource. For reference Table 3-12 describes each of the SFOS features and the applied default settings once SFOS is turned on.

Table 3-12 Default Secure Fabric OS Policy settings

Secure Fabric OS Security Feature	Setting
FCS_POLICY	This policy must exist and cannot be empty. It is a list of all FCS switches. Any switch that is not in the FCS_POLICY list is a Non-FCS switch member. The FCS_POLICY lists members by switch name and WWN. This information is contained on the Primary FCS switch and it is the first switch in the list displayed by <code>scemodeshow</code> .
MAC Policies: SNMP, Telnet, HTTP, SES, Management Server (MS) Serial, Front Panel (SilkWorm 2800 only)	These policies are in the “No Policy” state by default; as such, they do not limit or block access.
Switch Connection Control (SCC)	The SCC is set to “No Policy” by default: a switch with any WWN can connect to the fabric.
Device Connection Control (DCC)	The DCC is set to “No Policy” by default: any device can connect to any port.
Options	By default, options are not enabled and therefore allow WWW zoning, for example.

3.1.13. Secure Fabric OS Pre-Installation Planning

As with any new security measures, there needs to be a plan in place before staging Secure Fabric OS (SFOS). This section will provide a SFOS Pre-installation checklist that outlines at a high level what needs to be accomplished before the SAN is ready for SFOS. Once these steps are complete, the SAN is prepared.

An SFOS Implementation Plan checklist follows. This checklist provides high-level guidelines as to what is needed for a successful staging of SFOS. The details of each checklist item follow the high level recommendations and is meant to help with putting together a plan that is tailored for a specific implementation.

3.1.13.1. *Secure Fabric OS Pre-Installation Checklist*

Here are some pre-installation steps that require completion prior to enabling Secure Fabric OS. It is important to do this ahead of time so that the SAN is prepared for SFOS security to be enabled. Table 3-13 contains a checklist that will help in defining a tailored plan. Note that the SilkWorm 1000 series of switches do not support SFOS.

Note: If more information is desired, each checklist item is expanded upon in the subsequent sections.

Table 3-13 Secure Fabric OS Preparation Checklist

Secure Fabric OS Preparation Checklist	
1.	Obtain and read Secure Fabric OS documentation. See this section below for recommended documents.
2.	Be safe. Backup switch configurations with configupload. Prior information such as zoning will be wiped out when a switch or fabric is allowed to join a Secure Fabric OS enabled fabric.
3.	Verify PKI Objects Exist. This is required for Secure Fabric OS implementation. <ul style="list-style-type: none"> • <code>pkishow</code> (Fabric OS 4.1) • <code>configshow "pki"</code> (Fabric OS 3.1/2.6.1)
4.	If the PKI objects do not exist, obtain the PKICert tool to install Digital Certificates. This is required for older switches. This utility runs on Windows and Solaris only.
5.	Download and Install Brocade SecTelnet and Secure Shell (SSH) Security Software Utilities
6.	Verify Fabric OS Version. Update as required. The SilkWorm 1000 series of switches do not support Secure Fabric OS.
7.	Install Security and Zoning Licenses on all switches in the SAN. This is required for Secure Fabric OS.
8.	Schedule downtime when secure mode is enabled. A reboot of each fabric in the SAN is required as the firmware update is disruptive.
9.	Highly Recommended: Set Core PID on all switches not running Fabric OS 4.0 or greater.

1. Obtain and Read Secure Fabric OS Documentation

For comprehensive information on Brocade Secure Fabric OS implementation, it is highly recommended to obtain and read the following documentation. Refer to *Reference Documentation* on page B-1 for additional references.

- *Secure Fabric OS Quickstart Guide* (publication number 53-0000352-01)
- *Secure Fabric OS Users Guide* (publication number 53-0000526-01)
- *SAN Security: A Best Practices Guide* (publication number GA-RG-250-00)

2. Backup the Brocade Switch Configurations

When any big change is going to take place within the SAN, it is a good idea to be prepared with a backup of each switch configuration. Implementing Secure Fabric OS (SFOS) definitely fits into that category. The fact of the matter is that even with planning, something may go astray. Use `configupload` to preserve a backup of the switch before enabling SFOS. For larger fabrics of more than four switches, use Brocade Fabric Manager, as it can backup all switches at once. Give strong consideration to making this a baseline configuration.

3. Verify PKI Objects Exist

For new and existing switches, verify the certificates, public and private keys exist by issuing `pkishow` at the command line for Fabric OS version 4.1. Use `configshow "pki"` for all other versions of Fabric OS. Use an ordinary telnet client since Secure Fabric OS is not enabled yet. Before Secure Fabric OS can be enabled and SecTelnet used, these objects must exist.

4. Obtain PKICert tool to Install Digital Certificates

If the PKI objects do not exist, get PKICert to install them. Usually these tasks only need to be performed on older switches that were factory installed prior to January 2002. Use the PKICert tool to install the digital certificates. There are three overall tasks to perform with PKICert. The overview of them is below. PKICert can be used to install the objects on an entire fabric. When using this tool, it is recommended to do one fabric at a time. Scheduled downtime in single fabrics is not required, but is recommended as part of Secure FOS implementation preparation.

- Request certificates with the Certificate Signing Request (CSR).
- Obtain the Certificate Signing Request from the fabric. This is an XML file and is required to get the digital certificates. PKICert will prompt the administrator for a name.
- Submit the CSR, obtain the digital certificates. These are contained in a second XML file that will be E-mailed after submission to a secure site. The file will have a unique number in the filename.
- Load Digital Certificates onto the fabric. New switches from the factory with Fabric OS 2.6.1, 3.1 and 4.1 have them pre-installed.

5. Download and Install Brocade Security Software Tools

Install the following tools on the hosts that will be used for Secure Fabric OS administration purposes. Refer to the procedure offered by the switch supplier for the location and download instructions.

- Brocade SecTelnet 1.0 or higher
- SSH Client – any client that supports version 2 of the protocol

6. Verify Fabric OS Version

The SilkWorm 1000 series of switches do not support Secure Fabric OS secure mode implementations. To verify the firmware versions, use the `version` command on all switches except the SilkWorm 12000. For it, use the `firmwareshow` command. Before updating, be sure to audit the current end-to-end configuration to understand potential impacts during the upgrade. Update to Fabric OS 2.6.1, 3.1 and 4.1 versions that are *supported* by the switch provider.

7. Verify All Switches Have Zoning and Security License Keys

Both the Zoning and Security licenses are required for enabling Secure Fabric OS. This is because the Primary FCS (Trusted) switch is responsible for doing all zoning management and updates to other switches in the fabric. These other switches must be able to accept these changes. If not installed, follow the switch provider procedure to get them installed.

8. Schedule Downtime When Secure Mode is Enabled

Select the Primary FCS switch in the fabric. After enabling secure mode, all the switches will be rebooted in the fabric in order for the primary FCS switch to propagate and enforce the Secure Fabric OS SAN security features.

9. Set the Core PID Format

While not required for implementing Secure Fabric OS, now is the time to plan for it. Setting the Core PID format on Fabric OS 2.x and 3.x is required for attachment to the higher port count switches that run Fabric OS 4.x. If Fabric OS 4.x based switches are to added as a core, this will save time and effort later on. It is strongly recommended that all initial configurations of Fabric OS 2.x or 3.x based switches have the Core PID set.

Guideline: Setting the Core PID at the initial staging phase, on Fabric OS 2.x or 3.x based switches, will allow for a seamless introduction of a SilkWorm 12000 or 3900 into the SAN fabric.

Now that these steps are complete, the SAN is ready for Secure Fabric OS deployment. In order to be effective at doing that, there needs to be an understanding of what is going to be done. The next section provides some guidelines to consider when putting together this plan.

3.1.14. Planning the Secure Fabric OS Implementation

Now that the SAN is prepared for Secure Fabric OS, here are some guidelines for planning the implementation. Like all planning, it may seem like a headache and not worth the effort at the start. Nothing could be further from the truth, especially when it comes to creating the most secure environment possible. It is highly recommended that the checklist of activities in Table 3-14 be used as a basis for any specific plan, which should be followed rigorously. To provide examples when providing the detail of each checklist item, the case study SAN will be used as an example.

Table 3-14 Secure Fabric OS Implementation Plan Checklist

SFOS Implementation Plan Checklist	
1.	Create a SFOS switch and device list for SFOS policies. This list should contain hostnames, switch names, IP addresses, and WWNs
2.	Plan the FCS placements for each switch. Select a primary FCS switch. For larger port count Brocade SANs, there should be a minimum of three FCS switches chosen.
3.	Covertly mark FCS Switches. Use a small physical mark so that FCS switches are easily located.
4.	Determine policy requirements for each device and host.
5.	Pick the SFOS management hosts. These may require Brocade SecTelnet and a freely available Secure Shell (SSH) client.
6.	Perform a final review of all configuration selections. Make sure last minute changes are included in the plan.
7.	Disable Telnet on Fabric OS 4.1 switches only
8.	Set Recovery Password and Boot Password
9.	Enable SFOS and verify its operation.
10.	Backup Primary FCS switch configurations with <code>configupload</code>. The databases are not saved with a <code>secmodeenable</code> .

Warning: The Primary FCS databases are not automatically backed up after doing a `secmodeenable`. This means that if a `secmodedisable` is done on the Primary FCS switch, all the data is gone. Re-doing a `secmodeenable` *does not restore* the database information. **To backup the Secure Fabric OS data, do a `configupload`** on the Primary FCS. This will preserve the `FCS_LIST`, passwords and other policy configurations that are set on the Primary FCS. If required, use `configdownload` to restore the data to the Primary FCS.

Warning: When enabling Secure Fabric OS, ALL switches in the fabric will be rebooted as part of the initialization process. ALL I/O will be disrupted. This happens when you enable Secure Mode with `secmodeenable`.

1. Create a Secure Fabric OS Switch and Device List

The list should contain switch and host Names, WWNs and IP addresses that will have policies applied to them. This is where having a complete record of SAN components is helpful. If this information is readily available, it is easy to consolidate that information from previously created spreadsheets.

2. Plan FCS Switch Placements

Pick the Primary FCS and backup FCS switch locations within the fabric. Do this based upon the location within the fabric topology. For larger port counts, there should be a minimum of three FCS switches chosen. All remaining switches will be Non-FCS switches. Since the Primary FCS will become the administrative point in the fabric, select a core switch. If there is a second core, make it a backup FCS switch. Should the primary FCS switch become unavailable it then becomes necessary to use the first available backup FCS switch to manage the fabric. The SilkWorm 12000 makes the best choice as the primary FCS switch given it is the highest available switch platform in the Brocade SilkWorm product family.

Note: With secure mode disabled, any switch can be used to manage the fabric.

Guideline: Consider a locked closet to physically secure the primary and secondary FCS switches and/or the management station.

3. Covertly mark FCS Switches

Use a small physical mark so that FCS switches are easily located.

4. Determine the Policy Requirements

Think about what devices should be under access control ahead of time. Knowing this before it is done will provide some flexibility later on. Record each device and host, define the Secure Fabric OS policy to be applied to each. When making these choices, be sure to work within general corporate security guidelines being applied to other IT equipment.

5. Pick Secure Fabric OS Management Hosts

These hosts will have SSH and SecTelnet installed for administrative purposes. Select at least two but not more than five. To maximize overall security effectiveness, these hosts should already be locked down. At a minimum, they should have enterprise LAN based security and be barred from physical access. Two are needed for redundancy purposes. More than five yields to many access points into the fabric. These should be on the device list.

6. Perform a Final Review

To make sure all the policies are in place, do a quick review with the cross-functional team responsible for making the decisions and implementing the configuration. This will prevent last minute changes from being ignored.

7. Disable Telnet on Version 4.1 Only

For the tightest security, it is recommended to disable the telnet daemon. This works with Fabric OS 4.1 or higher only. In Fabric OS v4.1, you have the ability to disable telnet on a per switch basis. Disabling telnet is done with the `configure` command. The Brocade switch does not require the switch to be brought offline with `switchdisable` when doing this task. Refer to *(Optional) Disabling the Telnet Daemon When Secure Mode is Enabled (Fabric OS 4.1 only)* on page 3-45.

8. Set Recovery Password and Boot Password

It is highly recommended to change these from the defaults for maximum security. Write these down and put in a secure place, like a safe. Follow the procedures in the *Secure Fabric OS Users Guide*.

9. Backup the Brocade Switch Configurations

The Primary FCS databases are not automatically backed up after doing a `secmodeenable`. This means that if a `secmodedisable` is done on the Primary FCS switch, all the data is gone. Re-doing a `secmodeenable` **does not restore** the database information. **To backup the Secure Fabric OS data, do a `configupload` on the Primary FCS.** This will preserve the `FCS_LIST`, passwords and other policy configurations that are set on the Primary FCS. If required, use `configdownload` to restore the data to the Primary FCS.

10. Verification Plan

After the SAN security measures and secure mode is enabled with policies set, they need to be tested using a verification plan. Put together a plan that makes sense for the specific SFOS environment that is going to be implemented. Here are some suggestions for the verification plan:

- Add some switches and devices that do not have access rights.
- Try adding a switch to the fabric, which is locked out by an SCC policy.
- Add an HBA that is not in the DCC list.
- Try to telnet to the Primary FCS with a host that should not have access.

3.1.15. SAN Secure Fabric OS Software Utility Considerations

This section will provide an overview of SecTelnet and SSH. SSH can be used as a standalone utility for secure management (with or without Secure Fabric OS). SecTelnet requires Secure Fabric OS (SFOS) for secure management of the SAN. Both are referred to as “out of band” management. To get these applications, contact the switch provider.

3.1.15.1. PKIcert Utility

Only binaries for Windows and Solaris have been developed. Use Version 1.0.5 or higher to do the certificate installation procedure. Generally older fabrics with switches that pre-date Brocade Secure Fabric OS are suspect. If required, this procedure only has to be done once per fabric. To find out if the required digital certificates and other PKI infrastructure items exist, such as public and private keys on Fabric OS 2.6.1 or 3.1, use `configshow "pki"`. For the SilkWorm 3900 and SilkWorm 12000 which run Fabric OS 4.1, use `pkishow`. This procedure requires obtaining and submitting a CSR.

3.1.15.2. Brocade Secure Telnet (SecTelnet)

The Brocade SecTelnet client allows for password encryption upon logging in. The SecTelnet binaries are available for Windows and Solaris only. In addition, use of SecTelnet requires that digital certificates be installed on each switch that is accessed. If not installed, use the Brocade PKIcert application to do the install. Use of SecTelnet outside of Secure Fabric OS does work but is *not* supported by Brocade. SecTelnet works for all versions of Fabric OS, however SFOS must be enabled for support. SFOS is available on versions 2.6.x, 3.1 and 4.1 or higher.

Note: Sectelnet must be used to administer the Primary FCS switch when running Secure Fabric OS. Brocade Sectelnet requires digital certificates are installed on each switch to be administered. This utility only encrypts the passwords sent over the LAN, all other commands etc., are sent as clear text.

Note: The maximum security policy size is:

- 62 KB for active db (for Fabric OS 4.1 and 3.1)
- 62 KB for defined db (for Fabric OS 4.1 and 3.1)
- 16 KB for Fabric OS 2.6.1

Caution: When using a switch with Fabric OS 2.6.1, and secure mode enabled, as an FCS switch the maximum database size is 16KB.

3.1.15.3. Using Secure Shell (SSH)

Secure Shell (SSH) is a standards based secure method for accessing SilkWorm switches running Fabric OS 4.1 or later. Any SSH client that supports version 2 of the protocol can be used. There are literally hundreds of freeware SSH clients available that have this capability. On the switch side, SSH is only supported on Fabric OS 4.1 or higher. The only Brocade products that run Fabric OS 4.1 are the Brocade SilkWorm 3900 and SilkWorm 12000. As far as testing goes, two popular clients, Putty and Fsecure have been used. It is anticipated that Brocade will test with more clients in the future.

Beginning in Fabric OS v4.1, you have the ability to disable telnet on a per switch basis. Disabling telnet is done with the `configure` command. The Brocade switch does not need to be disabled to do this.

3.1.16. LAN Planning Considerations

There are a few guidelines to consider when attaching the fabric to the corporate LAN infrastructure. In general, it is highly recommended to configure a separate VLAN for each fabric. If at all possible, avoid the use of proxy servers from the SAN management stations outside the local subnet. In fact, it is recommended to use a management station on the same VLAN. For detailed guidelines on how to connect Brocade switches to the corporate network, please refer to the *LAN Guidelines for Brocade SilkWorm Switches* (publication number 53-0000350-01).

3.1.17. Brocade Extended Fabrics Planning

This section will provide the essential information for the required for planning connections of Brocade SilkWorm fabrics over longer distances. One common reason is for data replication which gives site redundancy. In this way, if one site goes down due to a disaster, the data can be recovered and brought online in minutes rather than days. A huge cost savings. One other typical use is consolidated remote data archival to tape.

SAN long distance connectivity may involve single mode fiber or through more sophisticated network equipment that allows for greater line availability. For Metro Area Networks (MAN), those distances up to 120 Km, DWDM equipment is generally used. Longer distances that go up to thousands of kilometers generally require FC protocol conversion. This means a Wide Area Network (WAN) transport method is required. Different equipment, such as an ATM long haul switch, maybe required. These types of devices are out of scope for this document. The focus will be on the Brocade SilkWorm long distance implementation and the essential information for planning purposes.

3.1.17.1. Overview

Brocade SilkWorm 2000 Series, 3000 Series, and the 12000 support a separate buffer-to-buffer flow control circuit for each of the eight virtual channels (VCs) used on ISLs. Of the eight VCs used, four are used for management of Fibre Channel unicast frames and each are loaded with five credits on a Normal (L0) ISL. This allows switches to be interconnected at distances of up to 5km for 2 Gbit/sec links and 10km for 1 Gbit/sec links.

The Brocade Extended Fabrics license allows ISLs to be connected at up to 60km for 2 Gbit/sec links and up to 100km for 1 Gbit/sec links, while maintaining maximum bandwidth. This is accomplished by compacting credits on all four virtual channels normally used for unicast frames onto a single virtual channel (VC 2). Fabric OS v2.6 has three EF modes that can be used to configure the amount of credits available to ISLs on long distance links. An additional LE mode that supports 2 Gbit/sec FC performance at 10 km without requiring an EF license was introduced in Fabric OS v3.0/v4.0. Two new modes have been made available in Fabric OS v3.1 and v4.1 that allow for further enhancements to extended distance links.

3.1.17.2. Compatibility and Interoperability

If an Extended Fabrics port is to be installed on a SilkWorm 2000 Series switch, the fabric wide configuration parameter `fabric.ops.mode.longDistance` must be set to **1** on all switches operating within the fabric. Additionally, each long distance port must be set using the `portCfgLongDistance` command. Each of the two ports within a long distance ISL must be configured identically, otherwise fabric segmentation will occur.

SilkWorm 3000 Series and 12000 switches have new features that do not require long distance ports to be configured at the fabric level. Only individual port configuration is necessary.

In a mixed fabric configuration, where long distance ports are installed on the SilkWorm 3000 Series and/or 12000 switches, only port level configuration is required. If a long distance ISL is created between two SilkWorm 2000 Series switches in a mixed fabric, then the fabric wide long distance parameter must be set on all switches within the fabric, as well as port level configuration on the long distance ports.

Note: Long distance links are not supported from a SilkWorm 2000 series switch to a SilkWorm 3000 series or 12000. For example: SilkWorm 2800-to-3800 or 2800-to-12000.

The Brocade Trunking feature, which allows multiple ISLs within a quad to load balance traffic, is not currently supported on ports configured for Extended Fabrics.

Brocade Security features are supported on Extended Fabric links. This includes connections over dark fiber and DWDM networks.

3.1.17.3. Extended Fabrics Modes

Table 3-15 on page 3-31 lists the various Extended Fabrics modes and requirements. Two new modes available in Fabric OS v3.1 and Fabric OS v4.1 are L0.5 and LD. L0.5 can support distances at up to 25 km. LD, or Dynamic long distance mode, can automatically configure the required amount of credits based on the actual link distance. This is determined by calculating latency on the link using the link round trip timer in the ASIC.

Table 3-15 Extended Fabrics Modes

EF Mode	Buffer Allocation		Distance @ 1 Gbit/sec	Distance @ 2 Gbit/sec	Fabric OS Release	EF License Required
	1 Gbit/sec	2 Gbit/sec				
L0	5 (26)	5 (26)	10 km	5 km	All	NO
LE	13	19	~	10 km	3.x, 4.x	NO
L0.5	19	34	25 km	25 km	3.1, 4.1	YES
L1	27	54	50 km	50 km	All	YES
L2	60	64	100 km	60 km	All	YES
LD*	Auto	Auto	Auto	Auto	3.1, 4.1	YES

*The maximum distance for LD is the same as for L2.

For dynamic long distance links, the number of credits can be approximated using the following formula. The data rate is 1.0625 for 1 Gbit/sec and 2.125 for 2 Gbit/sec Fibre Channel. This equation can only be used as an estimation for the number of credits that will be allocated to a given port. The actual amount will most likely be slightly higher.

$$\text{Buffer Credits} = ((\text{Distance in km}) * (\text{Data Rate}) * 1000) / 2112$$

Since only 108 credits are available for use on ports within each quad, configuring long distance ports may cause other ports to become disabled if there are not enough credits available. For example, if two 2 Gbit/sec ports in a quad are configured for L1 mode, each will be allocated 54 buffer-to-buffer credits and cause the other two ports within the quad to become disabled.

Table 3-2 lists some possible port configurations that can be utilized with static long distance modes.

Table 3-16 Port Configurations

Speed	Port a	Port b	Port c	Port d	Total Credits
1 Gbit/sec	L2(60)	L1(27)/E(27)	Fx(16)	X	103
1 Gbit/sec	L2(60)	Fx(16)	Fx(16)	Fx(16)	108
1 Gbit/sec	L1(27)/E(27)/ Fx(16)	L1(27)/E(27)/ Fx(16)	L1(27)/E(27)/ Fx(16)	L1(27)/E(27)/ Fx(16)	64-108
2 Gbit/sec	L2(108)	X	X	X	108
2 Gbit/sec	L1(54)	L1(54)/E(27)	X	X	108(81)
2 Gbit/sec	L1(54)	E(27)	E(27)/Fx(16)	X	108(97)
2 Gbit/sec	L1(54)	Fx(16)	Fx(16)	Fx(16)	102
2 Gbit/sec	E(17)Fx(16)	E(17)/Fx(16)	E(27)/Fx(16)	E(27)/Fx(16)	108(64)

Guideline: For switches running Fabric OS 3.1/4.1 or greater, use LD mode for long distance connectivity.

3.2. Staging

Once the plan is complete and the resources readied, the new SAN can be built and prepared for production. Staging the SAN is more than just uncrating, racking and installing the Brocade switches and other components. Most of the work actually goes into configuring each component's firmware and software. Guidelines on how to accomplish those tasks be presented in this section in the form of checklists. The case study will be used extensively in this section to provide real examples of typical commands used in the staging phase of SAN deployment.

This section pre-supposes staging a new core-edge fabric with clean Brocade switches. A clean switch has no defined or active zoning configuration and has all default settings. It assumed throughout this section that a Core/Edge topology has been chosen as part of the design criteria. The guidelines in this section may need to be modified for other topologies. This section will also not explicitly cover an existing SAN fabric migration to other Brocade switch platforms. Please refer to the *SAN Migration Guide* (publication number 53-0000360-01) for the recommended guidelines and procedures.

Staging a new SAN Fabric requires essentially two tasks. The first is uncrating, racking, cabling and providing power to the Brocade switches. The second is configuring the Brocade Fabric Operating System firmware. This section will not provide recommendations on racking, cabling or power installation for the Brocade SilkWorm Fabric Switch Family. Guidelines for these tasks were provided in [The Rack Layout Plan on page 3-12](#). Nor will this section focus on the devices, such as HBAs and storage targets, attached to the Brocade SAN Fabric. The focus of this section will be on new commands used in staging Brocade SAN fabrics with Fabric OS 2.6.1, 3.1 and 4.1, with older commands referenced as needed. New high-level troubleshooting functions are also introduced within these Fabric OS releases and will be discussed at a high level in this section. Since most of the new functionality is in Fabric OS 3.1 and 4.1, Fabric OS 2.6.1 will get limited coverage.

All of these functions will be shown as a case study example using Telnet based commands. Differences among commands in each of the Fabric OS releases will be pointed out as each task is discussed. Many of these same functions can be accomplished with the Brocade Web Tools GUI interface and Fabric Manager. These will not be discussed in detail in this initial release of the DDM guide. Please refer to the management section. Please see the *Fabric Manager User's Guide* and the *Web Tools User Guide* for detailed instructions on usage. Brocade Fabric Manager and Web Tools will be referenced at a high level when the situation presents itself, as some tasks are easier to do with these software tools, especially with larger port count SANs.

Guideline: For an online list of commands, use the help command. There is a man style online command reference. Simply use `help <command>` for this reference. For example, `help tsclockserver`. Some commands are logically grouped. Those command groups have there own help command. As an example, for the Secure Fabric OS command set, use `sechelp`. Note that the Security License and PKI objects, including the digital certificates must be present for this to work.

3.2.1. Fabric OS 2.6.1/3.1/4.1 Overview

Brocade Fabric OS 2.6.1 runs on the second generation SilkWorm 2000 family of Fibre Channel switches. Brocade Fabric OS 3.1 is for the third generation SilkWorm 3200 8-port entry level and 16-port SilkWorm 3800 switches. Both of these models are based on an embedded VXWORKS UNIX kernel. Brocade Fabric OS 4.1 is for the third generation SilkWorm 3900 and SilkWorm 12000 core fabric switch. These models use a Linux kernel and have 256 MB flash memory cards used to store the Linux kernel and Brocade Fabric OS. Note that the Brocade SilkWorm 12000 switch has two 64-port domains in a single chassis. Before getting into the SAN staging details, below is an overview of the major new features and functions of Brocade Fabric OS 2.6.1, 3.1 and 4.1.

For complete list of all commands, please refer to the *Brocade Fabric OS Reference Guide* for each release. For typical Fabric OS configuration tasks, please refer to the appropriate *Brocade Fabric OS Procedures Guide* for version of Fabric OS. This product guide provides excellent depth on how to do most of the configuration activities discussed in this section.

3.2.1.1. Fabric OS Version 2.6.1

Fabric OS (FOS) version 2.6.1 is a maintenance release that has Secure Fabric OS (SFOS) enhancements for compatibility with Fabric OS versions 3.1 and 4.1. There are a few other new features worth noting. A NTP time server can be set that will propagate periodic time updates to the entire fabric. A new command, `nsaliasshow`, allows zone alias names to be shown along with name server entries. There is a new `trackchanges` function that provides a better means of doing change management. A SilkWorm 2000 series switch or port can now be persistently disabled across reboots with Fabric OS 2.6.1. In Fabric OS 2.6.1 only, Fabric Watch can be configured to notify the administrator via E-mail when a message is generated. These are some of the significant changes in Fabric OS 2.6.1.

Note: Many new features such as hard setting a principal switch, FDMI, and naming ports are NOT available with Fabric OS 2.6.1. Some new features like Ultra HA, are only available on Fabric OS 4.1 or higher. These will be pointed out as required.

3.2.1.2. Fabric OS 3.1

The most important new feature Fabric OS (FOS) Version 3.1 adds is Secure Fabric OS (SFOS). Secure Fabric OS has new methods for controlling access to SAN resources and allows for much more effective change management. In addition to this and the Fabric OS 2.6.1 new features discussed above, there are some other major changes. These include new troubleshooting commands such as setting up persistent error logs, new `supportshow` command groupings, and FDMI for online management of Emulex HBAs. Command line zoning management has some new search tools with `nodefind` and `nszonemember`.

Note: Fabric OS 3.1 DOES not contain the new HA enhancements in Fabric OS 4.1. This means that updating firmware on the SilkWorm 3200 and SilkWorm 3800 fabric switches will continue to be disruptive in single fabric SANs.

3.2.1.3. Fabric OS 4.1

Two of the new features in Fabric OS 4.1 are Secure Fabric OS (SFOS) and greatly enhanced HA. Firmware updates from Fabric OS 4.1 to 4.1.x will now be non-disruptive on the SilkWorm 3900 and SilkWorm 12000. However, note that updates from Fabric OS 4.0.x to 4.1 WILL be disruptive. So be sure to schedule downtime for this. Fabric OS 4.1 contains all of the new functions of Fabric OS 3.1 with a new `switchreboot` command and online WWN Card swapping capability for the SilkWorm 12000.

3.2.2. Case Study

As concepts are introduced and new features explained a case study is used as an example to provide context.

3.2.2.1. Case Study SAN Description

The case study SAN contains two fabrics. Each fabric contains six switches in resilient core-edge topology. Each fabric has two SilkWorm 2800 fabric switches, two SilkWorm 3800 fabric switches, a SilkWorm 3900 switch, and a SilkWorm 12000 core fabric switch. The SilkWorm 2800 switches are connected with two 1 Gbit/sec ISLs. The SilkWorm 12000 switches are deployed at the core to provide the maximum scalability and availability. The SilkWorm 3800 and SilkWorm 3900 switches are connected at the edge using two ISLs per trunk group. This yields 4 Gbit/sec bandwidth for each logical connection. The

design today provides 86 device ports per fabric, with a total of 172 ports in the SAN. Brocade Trunking technology provides up to four ISLs that form a single 8 Gbit/sec logical ISL. If the design recommendations are followed, bandwidth can be added on the fly, just by adding a cable to any trunk group. With the SilkWorm 12000 fabric switch at the core, the SAN used in this solution can scale quickly and easily to hundreds of ports using SilkWorm 16 and 32 port switches on the edge.

The attached hosts include Solaris, Windows, and HP/UX systems. These hosts share the RAID array connections that are connected on the SAN. To maximize the SAN availability, all hosts are running multi-pathing software. Figure 3-7 shows the SAN design.




Figure 3-7 Case Study SAN Design

The assumptions used in the case study are:

- The installation and configuration of the LAN infrastructure, HBA drivers, multi-pathing software, and host applications in the case study SAN is beyond the document scope. These areas are covered only as it pertains to this chapter.
- All host bus adaptors (HBA) are installed, the appropriate drivers loaded and configured for F_port (point-to-point) attachment.
- The host and storage devices are powered on, initialized and are ready to be attached to each fabric.

The SAN has been designed following Brocade recommendations in *SAN Design* on page 2-1.

3.2.3. Powering the SAN Equipment

In the event of that the SAN needs to be brought online, it is recommended to follow the power on sequence in Table 2-15. This will ensure a clean bring up of all devices. If the sequence is not followed, it is possible that some hosts may not see the storage devices.

Table 3-17

High Level Power On Sequence
1. IP Infrastructure
2. Brocade Switches
3. Storage Devices
4. All Hosts

3.2.4. Preparing the Switches for the SAN Fabric

There are many possible ways to configure the switches that make up fabric. For one thing, how a switch is configured depends on the role of the switch. During the planning phase of deployment, the following questions should have been addressed:

- Is it a principal switch?
- Is it a core or edge switch?
- Are long distances are required?
- Are the right license keys in place?

Now that the answers are known, the goal of this section is to provide guidance on preparing Brocade Fabric OS for production in a multi-switch SAN fabric. There are two major steps. First, prepare each switch for attaching to the corporate LAN infrastructure and joining a fabric. The second step is to do the fabric wide configuration, this being primarily zoning. Each will have a separate checklist. As with the rest of the document, this section first provides “bare bones” guidelines. Each checklist item is then expanded upon to provide detail into how each task is completed. Examples will be used throughout. Most of these steps apply to all versions of Fabric OS. Those steps that are unique to a particular version will be called out separately. Some guidelines will be considered optional and will be noted as such in Table 3-18. All commands discussed in this section are available to the admin user.

Table 3-18 Switch Preparation Checklist

Brocade Switch Preparation Checklist
1. Gather Planning Information (Switch Spreadsheet)
2. Check Fabric OS version and status
3. Set IP Address(es)
4. Set Date and Time
5. Set the Switch Name
6. Set the Domain ID Number and Set Core PID on non-Fabric OS 4.x switches
7. Add Fabric OS licenses
8. Optional: Name devices with PortName
9. Optional: Setting up a preferred principal switch in the fabric (Fabric OS 4.1 Only)
10. Set Telnet Session Timeout Value with the timeout command
11. Optional: Disabling the telnet daemon when secure mode is enabled (Fabric OS 4.1 only)
12. Optional: Setup ports for Extended Fabrics (see Extended Fabrics Section)
13. Optional: Set up Fabric Watch, SNMP Traps, <code>switchstatuspolicyset</code> (Pointer to Management Section)
14. Baseline and backup the switch with <code>configupload</code>

3.2.4.1. Gather Planning Documentation (Switch Spreadsheet)

Before getting started, gather the switch spreadsheet put together during the planning phase. This shows the planned IP addresses and domain names as well as the switch roles (Core Switch, Edge switch, Management Switch, etc.) Now that the IP addresses and domain numbers are known, its just a matter of executing the appropriate commands to set these values on the associated switch.

3.2.4.2. Check Fabric OS version and Environmental Status

There are several commands that are recommended to be run to check the Fabric OS (FOS) version and the overall switch environmental status after accessing the switch through a serial cable (see the next section for how to do this). The recommended commands to execute on Fabric OS 2.6.1 and 3.1 are: `version` (`firmwareshow` for Fabric OS 4.x), `uptime`, `switchstatusshow`, `sensorshow` and `portcfgshow`. Outputs from the case study are shown below. `Switchstatusshow` displays any triggered messages set by `switchstatuspolicyset`. `Switchstatuspolicyset`, discussed in the management section, allows the environmental settings that trigger warning messages to be customized.

```
sialab89:admin> version
Kernel:      5.4
Fabric OS:   v3.1.0
Made on:     Thu Feb 20 15:21:32 PST 2003
Flash:      Thu Feb 20 15:22:24 PST 2003
BootProm:   Tue Oct 30 10:24:38 PST 2001
```

Figure 3-8 Version Command

```
sialab89:admin> uptime
8:01am up 1 day, 12:32, 1 user, load average: 1.29, 1.45, 1.27
int219:admin> switchstatusshow
The overall switch status is HEALTHY/OK
poc166:admin> firmwareshow
```

```
Local CP (Slot 5, CP0): Active
    Primary partition:    v4.1.0
    Secondary Partition:  v4.1.0
Remote CP (Slot 6, CP1): Standby
    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.
```

Figure 3-9 Uptime Command

```
int195:admin> switchstatusshow
The overall switch status is Marginal/Warning
Contributing factors:
* 1 bad power supply and 0 missing power supply triggered the Marginal/Warning status
```

Figure 3-10 Switchshow Command

Sensorshow displays the status and values of each FRU. It works for 2.6.1, 3.1 and 4.1.

```
int219:admin> sensorshow
sensor 1: (Temperature) is Ok, value is 46 C
sensor 2: (Temperature) is Ok, value is 43 C
sensor 3: (Temperature) is Ok, value is 32 C
sensor 4: (Temperature) is Ok, value is 45 C
sensor 5: (Temperature) is Ok, value is 43 C
sensor 6: (Fan      ) is Ok, speed is 3308 RPM
sensor 7: (Fan      ) is Ok, speed is 3341 RPM
sensor 8: (Fan      ) is Ok, speed is 3308 RPM
sensor 9: (Fan      ) is Ok, speed is 3409 RPM
sensor 10: (Fan     ) is Ok, speed is 3308 RPM
sensor 11: (Fan     ) is Ok, speed is 3341 RPM
sensor 12: (Power Supply ) is Ok
sensor 13: (Power Supply ) is Ok
```

Figure 3-11 Sensorshow Command Output

```
int195:admin> portcfgshow
Ports          0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
Locked L_Port  .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Locked G_Port  .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Disabled E_Port .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Persistent Disable .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
where ..:OFF, ?:INVALID.
```

Figure 3-12 Portcfgshow Command Output for Fabric OS 2.6.1

```
sialab89:admin> portcfgshow
Ports          0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
Speed          AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Trunk Port     ON ON ON ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance  .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
VC link init   .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Locked L_Port  .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Locked G_Port  .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Disabled E_Port .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Persistent Disable .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
ISL R_RDY Mode .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
where AN:AutoNegotiate, ..:OFF, ?:INVALID.
LM:L0.5
```

Figure 3-13 Portcfgshow Command Output for Fabric OS 3.1

3.2.4.3. Set IP Address

To set IP Addresses on all of Brocade SilkWorm 2 Gbit/sec family of switches, console access to the serial port is required. To do this, attach a straight through serial cable to the console port on a host and the switch serial port. Configure the terminal program with the proper serial port settings. These are shown in Table 3-19.

3.2.4.4. Set Date and Time

To set the date and time use the `date` command. Here is an example that sets the date and time to 4:02pm March 22, 2003. Using `date` with no arguments provides the current day and time.

```
sialab89:admin> date "0322160203"
Sunday March 22 16:02:00 UTC 2003
```

Figure 3-14 Date Command

Table 3-19 Brocade Switch Serial Port Settings

Serial Port Settings
<ul style="list-style-type: none"> • 8-bit • No parity • No Flow Control • One stop bit • 9600 baud

Once attached, and with the terminal program running, hit enter a few times to initiate the connection. The switch will respond with a login prompt.

1. At the prompt, log on as admin and type the command `ipAddrSet` to set the IP address. For example:

```
sw0:admin> ipAddrSet
```

2. Enter the IP Address at the prompt. The example below uses 10.1.23.92.

```
Ethernet IP Address [192.168.155.89]: 10.1.23.92
```

3. At the Ethernet subnet mask prompt enter the subnet mask. The example uses 255.255.255.0.

```
Ethernet Subnet Mask [255.255.255.0]: 255.255.255.0
```

4. For the Fibre Channel IP address and subnet mask, make sure these are at the default value which is [none]. All factory installed units are pre-set with this. To reset to the default values enter 0.0.0.0 as shown.

```
Fibre Channel IP Address [none]: 0.0.0.0
```

```
Fibre Channel Subnet Mask [none]: 0.0.0.0
```

5. Enter the Gateway Address. Since this is on a private VLAN with the Management Station, keep the default of [none] by hitting **enter**. If not at the default value, use 0.0.0.0 to reset it to [none], this is shown below. Enter **Y** to the configure the IP addresses immediately.

```
Gateway Address [none]: 10.1.23.1
```

```
Set IP addresses now?
```

```
[y = set now, n = next reboot]: y
```

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

The IP configuration is complete. Repeat 1-5 for the remaining switches in each fabric.

When the switch has been configured with an IP address, a telnet, Web Tools or Fabric Manager session can be initiated from a management host for administrative purposes. An entire session is shown in Figure 3-15, for reference purposes.

```
sialab89:admin> ipaddrset
Ethernet IP Address [10.77.77.77]: 192.168.15.89
Ethernet Subnet Mask [255.255.255.0]: 255.255.255.0
Fibre Channel IP Address [none]: 0.0.0.0
Fibre Channel Subnet Mask [none]: 0.0.0.0
Gateway Address [10.1.0.1]: 192.168.15.1
Committing configuration...done.
Set IP addresses now?
[y = set now, n = next reboot]: y
sialab89:admin>
```

Figure 3-15 Setting IP Addresses on Brocade Switches

Guideline: Like Fabric OS 2.x and 3.x, Fabric OS 4.1 now contains `ifmodeset` which allows the Ethernet port speed to be set at 10 or 100 Mb/sec. This may be required when attaching to some brands of Ethernet switches.

Guideline: If using a terminal server to manage serial port connections to all Brocade switches in the SAN be sure to *DISABLE* flow control on every virtual serial port on the terminal server side.

3.2.4.5. Set the Switch Name

Once the IP address is set, use `switchname` to set the name of the switch. As an example,

```
sw77:admin> switchname "sialab89"
sialab89:admin>
```

3.2.4.6. Set Domain ID and Core PID format in the same session (on non Fabric OS 4.x switches)

As part of the introduction of Fabric OS 4.0, for greater scalability and compatibility, Brocade changed the 24 bit PID format in Fabric OS versions 2.x and 3.x to support attachment to the higher port count switches. As an example, this means that an old PID 091500 now becomes 090500. Changing the core PID only needs to be done once and thus it is highly recommended to change it during the initial staging of each SAN fabric.

This section will just provide a high level recommended procedure done on the case study SAN. Changing the core PID is disruptive, so SANs with single fabrics will require scheduled downtime. For all the recommendations and details for setting the core PID format, reference the *Fabric OS Procedures Guide*.

To set the switch Domain IDs follow the procedure below. Note that the core PID can be set at the same time. Do one fabric at a time. For each switch do the following:

1. Open a telnet session and do a `switchdisable`. This will cause the switch to go offline.
2. Type `configure` and hit enter. The configure menu appears.
3. Select the **Fabric Configuration** submenu by typing **Y** and hitting **Enter**. Set the domain ID by entering a domain number. Hit enter until the Core PID prompt appears and set the core PID to 1. The default value is 0. DO NOT bring the switch online.
4. Repeat 1-3 for all non 4.x switches in Fabric A only.

5. Once all of the switches have been set, disable with `switchdisable` then enable each SilkWorm 12000 core switch with a `switchenable`. Ignore the any error messages that appear.
6. Now enable each edge switch with `switchenable`
7. Verify the trunk groups form by using `trunkshow` on each SilkWorm 12000 core switch.
8. Repeat procedure 1-7 for Fabric B.

The Core PID format is now set for each switch in the SAN fabrics.

Set Domain ID Number Only

While it is not necessary to set a domain, it is recommended to do so. If no domain is set, the switch will automatically derive a domain as part of the initialization process. To automatically obtain a domain from the fabric, it is necessary that the switch connect to the fabric in a disabled state and then be enabled once the connection is complete.

Setting the domain ID number of the switch applies to all versions of Fabric OS. The domain ID number is the first 8-bits of the 24-bit port ID (PID). The default domain number is 1 for all switches. The following steps provide a high level outline of the procedure, for more in depth information please refer to the *Fabric OS Procedures Guide*.

1. Disable the switch with `switchdisable`. This must be done for all versions of Fabric OS.
2. Use the configure command to set the Domain ID. Running `configure` will invoke several prompts.
3. Enter Y at the Fabric parameters prompt. The first entry will be the Domain ID. Type in the number. The valid range is 1-239.

Guideline: Use the last octet of the switch IP address for the number.

4. Hit Enter to take the other default values.
5. Once the configuration is committed, re-enable the switch with `switchenable`.

3.2.4.7. Add Fabric OS Licenses

Follow the instructions given by the switch provider. This normally entails using the paper instructions and going to the Brocade web site an entering in a registration number. A high level procedure for adding a trunking license is illustrated in the following example.

1. Get the license ID of the switch. All licenses are based on the WWN. To find out the switch license ID for Fabric OS versions 2.x and 3.x, use `wwn`. For Fabric OS 4.x use `licenseidshow`.

Example for Fabric OS 2.x and 3.x:

```
sialab89:admin> wwn
10:00:00:60:69:51:10:42
```

Example for Fabric OS 4.x:

```
poc166:admin> licenseidshow
10:00:00:60:69:80:0f:ac
```

2. Add a license with `licenseadd`. Use the key as generated by the instructions in the paper pack.

```
sialab89:admin> licenseadd "SeQedReQRSbfRfeB"
```

3. Now check the licenses with `licenseshow`. Note that the Trunking License now exists.

```
sialab89:admin> licenseshow
zyzze9b9b0z00fh:
  Web license
  Zoning license
  Fabric license
```

```

SeQedReQRSbfRfeB:
  Trunking license
SebSeyQy9cafcTft:
  Web license
  Zoning license
  SES license
  QuickLoop license
  Fabric license
  Remote Switch license
  Remote Fabric license
  Extended Fabric license
  Entry Fabric license
  Fabric Watch license
  Performance Monitor license

```

Note: After adding a Trunking License, use `switchofgtrunk` to enable all ports for trunking. The switch does not need to be disabled with `switchdisable` to do this. Most other licensed features are not activated automatically and require an extra step. Please refer to the *Brocade Fabric OS Procedures Guide* for specific information about licensed options.

3.2.4.8. (Optional) Name devices with PortName

In Fabric OS 3.1 and 4.1 there is a new command, `portname`, that lets the administrator label a port. For all switches except the SilkWorm 12000, use the port number (or area number). For the SilkWorm 12000 use physical slot/port number as the port number argument.

Use `portname` to label port 26 on a SilkWorm 3900 as shown.

```
int219:admin> portname 26, "int124_HBA_B"
```

View the port name labels by using `portname` with no arguments.

```
int219:admin> portname
port 26: int124_HBA_B
```

To remove the label, use an "" as shown.

```
sialab90:admin> portname 8, ""
Committing configuration...done.
```

Use `portname` again to display the port name. Note that the name no longer exists.

```
sialab90:admin> portname
sialab90:admin>
```


Port name can also be viewed using from the Name Server window within Web Tools, as shown in Figure 3-16.

Domain #	Port #	Port Name	Port ID	Port Type	Fabric Port WWN	Device Port WWN	Device Node WWN	Device ID
2	28	portname	021c00	N	20:1c:00:60:69:80:04:04	10:00:00:00:c9:2a:b5:ec	20:00:00:00:c9:2a:b5:ec	[35] *Em
2	25	portname	0219ef	NL	20:19:00:60:69:80:04:04	21:00:00:00:87:00:03:59	20:00:00:00:87:00:03:59	[28] *HIT
2	25	portname	0219e8	NL	20:19:00:60:69:80:04:04	21:00:00:00:87:04:21:0b	20:00:00:00:87:04:21:0b	[28] *HIT
2	25	portname	0219e4	NL	20:19:00:60:69:80:04:04	21:00:00:00:87:04:21:bb	20:00:00:00:87:04:21:bb	[28] *HIT
2	25	portname	0219e2	NL	20:19:00:60:69:80:04:04	21:00:00:00:87:00:2f:7a	20:00:00:00:87:00:2f:7a	[28] *HIT

Figure 3-16 Names Server Window Showing Port Name Values

Note: Port names do not change with the `configdefault` command. However, they can be cleared on port by port basis with `portcfgdefault`. The `portshow` command displays port name in the first line of the output. `Switchshow` does NOT display the `portName`. The port name label is persistent across switch reboots and power cycles. `Nsaliasshow` displays *the zoning alias*, not the `portname` value.

3.2.4.9. (Optional) Setting up a Preferred Principal Switch in the Fabric (Fabric OS 4.1 Only)

The principal switch is responsible for handing out domain IDs to the rest of the fabric upon a fabric build. In some cases it may be desirable to hard set a switch to always be the principal switch. In Fabric OS 4.1 only, a preferred principal switch can be selected. Brocade has done extensive testing with this feature. The results have shown that with large fabrics and/or Secure Fabric OS enabled the selection is not completely deterministic. For the vast majority of fabrics, there is no issue.

The recommended way to set the principal switch is shown in the next example. This sets the current switch as the preferred principal switch and forces a fabric build to enable it. Fabric builds are not disruptive.

```
poc165:admin> fabricprincipal -f 1
fabric: Reconfiguration due to Principal Selection Mode
fabric: Reconfiguring at Mon Mar 17 15:54:59 2003
Principal Selection Mode enabled (Forcing fabric rebuild)
poc165:admin>
5 4 3 2 1

10 9 8 7 6 5 4 3 2 1

fabric: Principal switch
fabric: Domain 165
```

To enable preferred principal switch selection without doing a rebuild do the following:

```
poc165:admin> fabricprincipal 1
Principal Selection Mode enabled (Activate in next fabric rebuild)
```

To turn off the persistent selection mode on the switch, use 0 as the argument as shown:

```
poc165:admin> fabricprincipal 0
Principal Selection Mode disabled
```

Use fabric principal with no argument to check the current mode. This example shows the current state as enabled.

```
poc165:admin> fabricprincipal
Principal Selection Mode: Enable
poc165:admin>
```

Note: When another switch that is configured to be the preferred principal switch the lower WWN switch will win and the Higher WWN switch will log an error saying “PSS principal failed (Lost to WWN: <wwn>)”. The preferred principal switch remains selected across reboots, power cycles, and upon a fabric rebuilds.

3.2.4.10. Set Telnet Session Timeout Value

It is highly recommended to set the admin telnet session timeout values to 10 minutes. This prevents a telnet session from locking up access to a switch.

Note: Table 3-20 shows the number of allowed telnet sessions per switch model.

Table 3-20

SilkWorm Switch Model	Number of Telnet Sessions
SilkWorm 2xxx series	1
SilkWorm 3200 (Fabric OS 3.x)	1
SilkWorm 3800 (Fabric OS 3.x)	1
SilkWorm 3900 (Fabric OS 4.x)	2
SilkWorm 12000 (Fabric OS 4.x)	4 (2 per logical switch)

This example shows how to set the timeout value on Fabric OS 2.6.1 and 3.1.

```
sialab89:admin> timeout
TimeOut is Disabled

sialab89:admin> timeout 10
Committing configuration...done.
TimeOut is now 10 minutes

sialab89:admin> timeout
TimeOut is 10 minutes

sialab89:admin> timeout 0
Committing configuration...done.
TimeOut is now Disabled
```

Figure 3-17 Timeout Command Output in Fabric OS 2.6.1 and 3.1

Note, when issuing the command in Fabric OS 4.1 you will get the following messages:

```
IDLE Timeout Changed to 10 minutes
The modified IDLE Timeout will be in effect after NEXT login
```

3.2.4.11. (Optional) Disabling the Telnet Daemon When Secure Mode is Enabled (Fabric OS 4.1 only)

To maintain a high level of security, it is recommended to disable the telnet daemon using the `configure` command on all switches with Fabric OS 4.1. The example below illustrates the procedure.

Note: Switches running Fabric OS 2.x or 3.x do not have the capability to disable the telnet daemon.

Guideline: In most cases the switch must be disabled with `switchdisable` prior to using the `configure` command, however when disabling the telnet daemon the switch remains online while the `configure` command is invoked.

```
int219 login: admin
Password:
Please change your passwords now.
Use Control-C to exit or press 'Enter' key to proceed.

Password was not changed. Will prompt again at next login
until password is changed.
int219: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/0) Thu Apr 17 17:39:11 2003...

Security policy change: TTY pts on switch instance 0 will be logged out.

Connection closed...
```

Enabling a disabled telnet daemon:

To enable a disabled telnet daemon connect to the switch serial port and follow the procedure below.

```
int219: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): [off] on
int219:admin> 0x2b2 (fabos): Switch: 0, Warning FW-STATUS_SWITCH, 3, Switch st
atus changed from Marginal/Warning to HEALTHY/OK

int219:admin>
```

3.2.4.12. (Optional) Setup Ports for Extended Fabrics

The guidelines on how to set those parameters are in *Configuring Extended Fabrics* on page 3-69.

3.2.4.13. (Optional) Setup Fabric Watch, SNMP Traps, `switchstatuspolicyset`

Although optional, it is highly recommended to define the environmental and SNMP management items at this point in time. For guidelines on setting these parameters please see the Management Section.

3.2.4.14. Baseline and Backup the Switch With `configupload`

This is very important. Once all the settings are complete, use `configupload` to backup the switch configuration. This uploads all of the `configshow` parameters that define the switch configuration. Doing this is highly recommended in the event the original settings need to be restored. If minor changes to the configuration is required, just edit the configuration file and re-download the configuration to the switch. One further benefit is that a “standard” switch configuration can be defined and uploaded to all of the remaining switches in the fabric. This is known as baselining. Refer to *Configupload* on page 3-41.

Guideline: Define a “golden” switch. For larger SAN fabrics, or for the staging of many smaller fabrics, use `configupload` to baseline the “golden” switch. The “golden” switch configuration can be downloaded with Fabric Manager to all others in the fabric. Do baselines by Fabric OS version. In other words, if the fabric contains switches running Fabric OS 3.1 and 4.1, create two “golden” switch configurations. Doing this really helps when it comes to change management.

3.2.5. SAN Fabric Configuration

Now that each switch has been prepped, the SAN fabric can be built and configured. A checklist provides some essential high-level guidelines when performing these tasks. Most of what is done is to setup the zoning. Once these steps are complete, the staging is essentially complete. Later sections will point out other new functions and features of Fabric OS 3.1 and 4.1 not covered here.

Table 3-21 SAN Fabric Configuration Checklist

SAN Fabric Configuration Checklist	
1.	Gather Planning Documentation (ISL Map, Device Spreadsheet, Logical Design Diagram)
2.	Cable ISLs. Check Fabric with <code>fabricshow</code>
3.	Cable Host and Storage Devices. Check with <code>switchshow</code> , <code>nsshow</code>
4.	Use documentation to label all cables
5.	Optional: Dummy Zone to prevent device access
6.	Setup NTP Time Server on the Fabric. Use the Principal or Primary FCS Switch only
7.	Run <code>PortTest</code> to validate fabric plumbing
8.	Pre-Zoning Check. Verify hosts and storage are in the fabric that are to be zoned. Use <code>nscamshow</code> on a core switch to do this.
9.	Implement Zoning
10.	Profile each SAN fabric

3.2.5.1. Gather Planning Documentation

Gather the ISL Port Map, Device Spreadsheet and Logical Design Diagram that was put together during the planning phase. These items will be used as the SAN map that shows how to build the fabric and where to attach all the devices. Determining this was the hard part. Now that the required information is readily available, its just a matter of plugging in the cables and verifying the fabric has formed.

3.2.5.2. Cable ISLs

Cable the ISLs as documented by the ISL port map. *Cable Planning* on page 3-6 provides guidelines on cabling.

3.2.5.3. Cable Hosts and Storage Devices

Cable the Hosts and Storage devices as documented by the Device Spreadsheet. *Cable Planning* on page 3-6 provides guidelines on cabling. Please refer to it for some good suggestions if desired. Once cabling complete use `fabricshow` to verify the fabric is built. Use `switchshow` and `nsshow` to verify devices are online and registered with the Name Server. Some example outputs from the case study SAN is shown in the following illustrations. Note that `fabricshow` shows the six switches as expected.

```
sialab89:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr   FC IP Addr     Name
-----
89: fffc59  10:00:00:60:69:51:10:42  192.168.155.89  0.0.0.0        "sialab89"
90: fffc5a  10:00:00:60:69:51:0f:2b  192.168.155.90  0.0.0.0        >"sialab90"
166: fffc6  10:00:00:60:69:80:0f:ad  192.168.173.166 0.0.0.0        "poc166"
195: fffc3  10:00:00:60:69:10:93:14  192.168.162.195 0.0.0.0        "int195"
196: fffc4  10:00:00:60:69:10:62:2c  192.168.162.196 0.0.0.0        "int196"
219: fffcdb  10:00:00:60:69:90:04:1a  192.168.162.219 0.0.0.0        "int219"
```

Figure 3-18 Fabricshow Output for Fabric OS 3.1

```

sialab89:admin> switchshow
switchName:      sialab89
switchType:      9.2
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:     89
switchId:        fffc59
switchWwn:       10:00:00:60:69:51:10:42
switchBeacon:    OFF
Zoning:          ON (bkup_cfg_B)
port 0: id N2 Online      E-Port 10:00:00:60:69:80:0f:ad "poc166" (upstream)
  (Trunk master)
port 1: id N2 Online      E-Port (Trunk port, master is port #0)
port 2: -- N2 No_Module
port 3: -- N2 No_Module
port 4: id N2 No_Light
port 5: id N2 No_Light
port 6: -- N2 No_Module
port 7: -- N2 No_Module
port 8: id N2 No_Light
port 9: id N2 No_Light
port 10: id N2 No_Light
port 11: id N2 No_Light
port 12: id N2 No_Light
port 13: -- N2 No_Module
port 14: -- N2 No_Module
port 15: id N2 Online      L-Port 15 public
sialab89:admin>

```

Figure 3-19 Switchshow Output for Fabric OS 3.1

There is a new option for `nsshow`. The `-r` option shows only the devices that have done a State Change Registration (SCR). These devices accept Registered State Change Notifications (RSCNs) which allow for device discovery when changes occur in the fabric.

```

int219:admin> nsshow -r
{
Type Pid      COS      PortName                               NodeName                               TTL(sec)
N   db0e00;    3;50:00:60:e8:02:ee:78:18;50:00:60:e8:02:ee:78:18; na
FC4s: FCP [HITACHI OPEN-E      -SUN0118]
Fabric Port Name: 20:0e:00:60:69:90:04:1a

N   db0f00;    3;50:00:60:e8:02:ee:78:19;50:00:60:e8:02:ee:78:19; na
FC4s: FCP [HITACHI OPEN-E      0118]
Fabric Port Name: 20:0f:00:60:69:90:04:1a

N   db1a00;    2,3;10:00:00:00:c9:30:d0:62;20:00:00:00:c9:30:d0:62; na
FC4s: FCP
Fabric Port Name: 20:1a:00:60:69:90:04:1a

N   db1b00;    2,3;10:00:00:00:c9:30:d0:9d;20:00:00:00:c9:30:d0:9d; na
FC4s: FCP
NodeSymb: [35] "Emulex LP9002 FV3.90A7 DV5-4.82A4 "
Fabric Port Name: 20:1b:00:60:69:90:04:1a

```

Figure 3-20 NsShow Output with `-r` Option Displaying Devices Registered for State Change Notification

To show all the devices in the fabric, use `nsallshow`. An example output is shown below.

```
sialab89:admin> nsallshow
{
  590fd1 590fd2 590fd3 590fd4 590fd5 590fd6 590fd9 590fda
  590fdc 590fe0 590fe1 590fe2 590fe4 590fe8 590fef 5a0800
  c30600 c30800 c40d00 db0e00 db0f00 db1a00 db1b00
23 Nx_Ports in the Fabric }
```

Figure 3-21 NsallShow Output Showing All Devices Participating in the Fabric

3.2.5.4. Use Documentation to Label All Cables

This is important enough to be called out as a separate step. Having labels on every ISL and device connection is one of the most invaluable tasks that can be completed. When troubleshooting or servicing a device, this allows for greatly reduced downtime.

3.2.5.5. Optional: Dummy Zone Configuration to Prevent Device Access

The default zoning configuration allows any device to see any other device. If the devices are plugged in and the desire is to have them locked out, define a dummy-zoning configuration. To do this, create a zone configuration with an unused port. The example in Figure 3-22 shows the four steps that creates a dummy zone for switch domain 101 and port 5. With this complete, no devices are allowed to see any other device.

```
int219:admin> zonecreate "dummyzone","101,5"
int219:admin> cfgcreate "dummycfg","dummyzone"
int219:admin> cfgenable "dummycfg"
int219:admin> cfsave
```

Figure 3-22 Creating a Dummyzone

3.2.5.6. Setup NTP Time Server on the Fabric

An extension of the Time Service capability is now available on Fabric OS 2.6.1, 3.1 and 4.1 with `tsclockserver`. The purpose of this command is to provide synchronization of the fabric time with an external NTP server. The time server only needs to be configured on one switch in the fabric. Configure the time server on either the Principal switch or while in Secure Fabric OS mode, the Primary FCS (trusted) switch. Once setup, the time server configuration will be propagated in band to each other switch in the fabric.

The default value for `tsclockserver` is `LOCL`, as shown in Figure 3-23. This means no time server has been setup for the fabric and each switch updates its own local time value. If no argument is used, the current configuration is shown. The CLI Command has the following use: `tsClockServer [<NTP Server address>]`. Any switch in the fabric can be used to set the parameter, as it is a fabric wide setting. An example of how to use `tsclockserver` is shown in Figure 3-24.

Guideline: If desired, use the primary management switch in the fabric, generally this would be the principal or Primary FCS Switch for NTP service setup.

```
int219:admin> tsclockserver
LOCL
```

Figure 3-23 Default Setting

```
sialab89:admin> tsclockserver "192.168.126.60" update this
Updating Clock Server configuration...
```

Figure 3-24 Making the Fabric-wide Change

3.2.5.7. Run Fabric Diagnostics with PortTest

After building the fabric it is a good idea to run some diagnostics. There is a new function in Fabric OS 3.1 and 4.1 that allows diagnostics to be run on a local switch that has both online devices and E_ports. The command that performs this function is called `porttest`. Run this command for a bit on the fabric to see if there are any errors. The test only runs on connected ports, whether it is a device or ISL. This is a good test to run if a faulty SFP is suspected on an online port. There are lots of options to this command. The man page is a good reference. Please note that the parameters are slightly different in Fabric OS 3.1 and 4.1. An example that illustrates the use of this command is shown for both versions of Fabric OS are below.

For Fabric OS 3.1 the following `porttest` runs for 1000 iterations on all ports with `-1` option. The test lasts about seven minutes. `Porttestshow` displays the current status. Check for PASS on the first line after the port number. This indicates a successful test run. To avoid lots of output, it is a good idea to run the test only on ports with the connected devices.

```
sialab90:admin> portTest -1,1000
sialab90:admin> porttestshow
Port 0 : PASS
PortType: E PORT(SLAVE)           PortState: TESTING
PortInternalState: TX             PortTypeToTest: ALL_PORTS
Pattern: 0xb                      Seed: 0xaa                UserDelay: 10
TotalIteration: 1000             CurrentIteration: 575
TotalFail: 0                     ConsecutiveFail: 0
StartTime: Mar 22 03:58:06
StopTime: NONE
Timeout: 0                       ErrorCode: 0
Port 8 : PASS
PortType: F PORT                 PortState: TEST DONE
PortInternalState: INIT          PortTypeToTest: NO_TEST
Pattern: 0xb                      Seed: 0xaa                UserDelay: 10
TotalIteration: 1000             CurrentIteration: 1000
TotalFail: 0                     ConsecutiveFail: 0
StartTime: Mar 22 03:58:06
StopTime: Mar 22 04:03:58
Timeout: 0                       ErrorCode: 0
```

Figure 3-25 PortTestShow

This is the same test in Fabric OS 4.1. The arguments are a bit more intuitive. The following `porttest` runs for 1000 iterations on all ports with `-1` option as before. As in the previous example `porttest` lasts about seven minutes with these arguments. Once again, to avoid lots of output, it is a good idea to run the test only on ports with the connected devices. If there are lots of devices, use the terminal emulator software to save the output to a log file.

```
int218:admin> portTest -iteration 1000
```

If `porttest` is configured to run indefinitely, use the `stopporttest` command to stop the test.

3.2.5.8. Pre-Zoning Check

Before doing the zoning, perform a sanity check on the device connectivity. Use the command `nscamshow`. Capture the output to a file and compare to the Device Spreadsheet. All devices should be present.

`Nscamshow` displays all the devices for each non-local switch domain that have been registered with each of the other fabric member switches name server. For a core-edge topology, such as the case study, use a core switch to execute this command, as there no end devices registered with the local name servers. If devices are attached locally, just use `nsshow` in addition to capture devices registered the local name server database. The partial output from a core switch in the case study is shown. Note that private loop HBAs will not be captured.

```
int63:admin> nscamshow
nscam show for remote switches:
Switch entry for 75
  state rev  owner
  known v260 0xfffc3f
Device list: count 1
```



```

Type Pid    COS      PortName                               NodeName
N    4b0f00;  3;20:05:00:a0:b8:07:5d:c7;20:04:00:a0:b8:07:5d:c6;
    FC4s: FCP
    Fabric Port Name: 20:0f:00:60:69:10:10:9d

Switch entry for 88
state rev    owner
known  v310 0xffffc3f
Device list: count 1
Type Pid    COS      PortName                               NodeName
N    580800; 2,3;10:00:00:00:c9:29:04:8f;20:00:00:00:c9:29:04:8f;
    FC4s: FCP
    Fabric Port Name: 20:08:00:60:69:51:0e:0a

Switch entry for 217
state rev    owner
known  v410 0xffffc3f
Device list: count 4
Type Pid    COS      PortName                               NodeName
N    d90e00;  3;50:00:60:e8:02:ee:78:08;50:00:60:e8:02:ee:78:08;
    FC4s: FCP
    Fabric Port Name: 20:0e:00:60:69:90:03:fa
N    d90f00;  3;50:00:60:e8:02:ee:78:09;50:00:60:e8:02:ee:78:09;
    FC4s: FCP
    Fabric Port Name: 20:0f:00:60:69:90:03:fa
N    d91a00;  2,3;10:00:00:00:c9:30:d0:66;20:00:00:00:c9:30:d0:66;
    FC4s: FCP
    Fabric Port Name: 20:1a:00:60:69:90:03:fa
N    d91b00;  2,3;10:00:00:00:c9:30:d0:c2;20:00:00:00:c9:30:d0:c2;
    FC4s: FCP
    Fabric Port Name: 20:1b:00:60:69:90:03:fa

```

Figure 3-26 Nscamshow Output

3.2.5.9. Implement Zoning

This section provides guidelines and tips for implementing zoning. The SAN fabric case study will be used to illustrate the command line. The zoning configuration can be done from the command line or with a GUI such as Fabric Manager or Web Tools. These methods will not be discussed in this section, for more information refer to *Zoning on page 4-37* and the *Brocade Zoning User's Guide* for your specific version of Fabric OS. Please note that although it is possible to have many zoning configurations defined on a fabric, there is only one active zoning configuration allowed.

Note: The steps shown in the example below will not re-create the configuration.

1. To configure a hardware zone on a SilkWorm 2000 series switch, first use `alcreate` with domain ID, port ID as the second argument. In the example below it is "218,8". As discussed in the Zoning section earlier, the WWN can be used for hardware zoning on all switches that run Fabric OS 3.x or 4.x. The example will use the WWN for the following steps.

```

int195:admin> alcreate "E250_OrA","218,8"
int219:admin> alcreate "int122_HBA_B","10:00:00:00:c9:24:f5:f9"
int219:admin> alcreate "stor_port_B","20:04:00:a0:b8:07:5d:c7"

```

2. Use the command `zonecreate` to define a zone using the aliases. The `zonestow` command is used to check the newly created aliases are correct.

```

int219:admin> zonecreate "int122_B_zone","int122_HBA_B; stor_port_B"
int219:admin> zonestow
Defined configuration:
alias: int122_HBA_B    10:00:00:00:c9:24:f5:f9
alias: stor_port_B    20:04:00:a0:b8:07:5d:c7

```

- Now use `cfgcreate` to define a configuration with one or more zones. This example has only one, defined as "E250_A_ZONE". Enable the new zone configuration with the command `cfgenable`.

```
int219:admin> cfgcreate "bkup_cfg_B","int122_B_zone"
int219:admin> cfgenable "bkup_cfg_B"
zone config "bkup_cfg_B" is in effect
```

- Use `cfgsave` to write the configuration to flash memory.

```
int219:admin> cfgsave
Updating flash ...
```

- Verify the new configuration with `zoneshow`.

```
int219:admin> zoneshow
Defined configuration:
cfg:  bkup_cfg_B      int122_B_zone
zone:  int122_B_zone  int122_B_zone; stor_port_B
alias: int122_HBA_B   10:00:00:00:c9:24:f5:f9
alias: stor_port_B   20:04:00:a0:b8:07:5d:c7
Effective configuration:
cfg:  OracleBackup_cfg
zone:  int122_B_zone
           10:00:00:00:c9:24:f5:f9
           20:04:00:a0:b8:07:5d:c7
```

- There is a new command called `cfgActvshow` that works that shows current active configuration. This command does not exist in Fabric OS 2.6.1.

```
sialab89:admin> cfgActvshow
Effective configuration:
cfg:  bkup_cfg_B
zone:  int121_B
           10:00:00:00:c9:24:f5:f9
           20:04:00:a0:b8:07:5d:c7
```

- To display the amount of memory the zone database has and the amount used, invoke `cfgsize`. In Fabric OS 2.6.1 and 3.1 the maximum allowed size is 96 KB. For Fabric OS 4.1 the maximum size is 128 KB. The output from each is shown below.

```
sialab89:admin> cfgsize
Zone DB max size - 98232 bytes
  committed - 274
  transaction - 0
int219:admin> cfgsize
Zone DB max size - 130956 bytes
  committed - 274
  transaction - 0
value = 0
```

For more details on the recommended zoning practices please reference the whitepaper titled *Zoning Best Practices*.

Warning: For mixed fabrics, the maximum zoning database size will be about 96 KB. The vast majority of fabrics will not even be close to this limit. The largest ones deployed today on the order of a thousand ports, use about 40 KB for the zoning database.

Guideline: If two or more fabrics make up the SAN, when creating aliases in each fabric do the following on a non-production fabric and use clean switches, with no zones enabled or defined.

- Configupload to a host.
 - Edit the zoning configuration appropriately.
 - Download the zoning configuration from the host.
-

3.2.5.10. Profile The SAN Fabrics

Once the SAN fabrics are built and the zoning is in place, it is a good idea to capture a profile of the fabric. The show commands are ideal for this. The example below will provide a sample checklist of commands that was used to capture the profile of the SilkWorm 12000 that is in the case study SAN. This checklist can be modified for the SilkWorm 3000 series by replacing some of the commands.

Table 3-22

SilkWorm 12000 Profiling Commands	
firmwareshow	sensorshow
hashow	switchshow
licensidshow	nsshow
licenseshow	nsallshow
ipaddrshow 4	islshow
portcfgshow	trunkshow
slotshow	fabricshow
chassisshow	zonestshow

These commands can be easily scripted. As changes take place in the fabric, consider periodically updating the profile information. This will proactively simplify change management. Plus this will provide a living documentation set that can be referenced for technical support.

3.3. Validation

Once the SAN is staged, it is highly recommended to verify its functionality and robustness before going into production. While less important for the entry-level environment, validation becomes critical for SANs with higher port counts. All pertinent tests for a particular implementation will not necessarily be discussed. The real idea of this content is to get the thought process going for any specific case.

Some sample validation guidelines and procedures will demonstrate how to check the SAN stability, High Availability (HA), and Security. If at all possible, it is a good idea to do these tests with generated I/O, preferably with the application up and running. An extensive SAN qualification was not the objective of this section. Rather these tests are meant to be used as guidelines and a proof point that the SAN is operating properly before it is put into production. This section will focus on validating a core-edge SAN. If core-edge is not used, all tests in this section can be tailored for other fabric topologies. Separate tests will be required for the particular application in use in the SAN. This is out of scope and will not be covered.

3.3.1. Sample Script

With no application, it is possible to generate I/O. If using UNIX hosts following sample script can be used. For Windows hosts, use an I/O tool such as Iometer. This script creates and writes to a file and then does continuous reads of it. The path and size in blocks need to be specified. As an example, `sbtest /hds01 1000` will create a file of size 1000 blocks in /hds01 and once created, it will do successive reads until terminated. More I/O can be generated on a single host using multiple instances of this script running in the background. `Portperfshow` is a handy command line tool that can be used to quickly check I/O integrity. The sample script is shown below.

```
#!/bin/sh
PATH=$1
SIZE=$2
COUNT=`/usr/bin/expr $SIZE \* 2`
TMPFILE="$PATH/sbtest.$$"
RUN=0
echo "Building test file ($TMPFILE)..."
/usr/bin/dd if=/dev/zero of=$TMPFILE bs=512k count=$COUNT > /dev/null 2>&1
echo "Done."
while [ 0 -eq 0 ];
do
    DATE=`/usr/bin/date`
    echo "Run #: $RUN Timestamp: $DATE"
    /usr/bin/dd if=$TMPFILE of=/dev/null bs=512k count=$COUNT >/dev/null 2>&1
    RUN=`/usr/bin/expr $RUN + 1`
done
```

Figure 3-27 Sample Script

3.3.2. Sample Validation Recommendations

Table 3-23 Switch Preparation Checklist

Validation Checklist
1. Fabric Stability Validation - refer to <i>Fabric Stability Validation</i> on page 3-55
2. High Availability ISL Failure Simulation - refer to <i>High Availability ISL Failure Simulation</i> on page 3-55
3. High Availability Switch Failure Simulation - refer to <i>High Availability Switch Failure Simulation</i> on page 3-55

3.3.2.1. Fabric Stability Validation

Run the application on the hosts for a period of time. In the case study, the script was run for 72 hours to check for I/O stability. No problems should be observed.

3.3.2.2. High Availability ISL Failure Simulation

All hosts are attached to one edge switch as shown in Figure 3-7. From this switch, one trunk with two ISLs are attached to each core switch to form a trunk group with 4 Gbit/sec of available bandwidth. These trunk groups contain an ISL trunk master and a second ISL as a trunk member. A maximum of 400 MB/sec should be generated across them to provide maximum ISL stressing. The switch tests simulated failures on an edge and core switch. Hot code load should also be tested before going into production.

Note: When attaching ISLs to the SilkWorm 2800 there will be no trunking as it is not supported in Fabric OS 2.x and the connections will be auto-negotiated at 1 Gbit/sec.

Case 1 Member ISL

With I/O, the non-trunk master ISL was removed for ten seconds and then replaced. This was repeated three times. I/O should continue without any effect.

Case 2 Trunk Master ISL

The Trunk master ISL was removed for ten seconds and then replaced. Repeat this test three times. I/O traffic may be paused briefly on the second fabric, however the I/O should not timeout on the host.

3.3.2.3. High Availability Switch Failure Simulation

Case 1 Edge Switch Failure

A `fastboot` was initiated on a switch with no I/O to simulate edge switch failure. Repeat this test three times in succession. The switch should fully recover after the fastboot and the fabric should rebuild. Verify this with the `fabricshow` command.

Case 2 Core Switch Failure

A `fastboot` was initiated on a core switch. Repeat this test three times in succession. The switch should fully recover after the fastboot and the fabric should rebuild. Verify this with the `fabricshow` command.

Case 3 CP Failover (SilkWorm 12000 only)

Initiate an `hafailover` on the active CP with I/O running in the fabric. Full recovery without an I/O pause should occur in all cases. Refer to *Fabric OS Version 4.1 Ultra High Availability (HA)* on page 3-58 for expected behavior.

Case 4 Non-Disruptive Code Load (Fabric OS 4.1 or higher only)

Initiate a `firmwaredownload` on a switch with Fabric OS 4.1 or higher. Full recovery without an I/O pause should occur in all cases. Refer to *Fabric OS Version 4.1 Ultra High Availability (HA)* on page 3-58 for expected behavior on the SilkWorm 39000 and 12000.

3.4. Maintenance and Operations

3.4.1. Executing the Upgrade of Fabric OS 2.x/3.x/4.x to 2.6.1/3.1/4.1

This section will provide some high level guidelines when updating firmware on Brocade SilkWorm Switches. The context of this section will be using the command line to update a single switch. For updating multiple switches, Brocade Fabric Manager should be used. This will not be discussed here. Please see the Management section for details on using Fabric Manager. There are some significant differences between updating Fabric OS 2.x/3.x and 4.x. These will be pointed out. For detailed instructions on updating Fabric OS 3.x and 4.x reference the *Fabric OS Procedures Guide*. There is an excellent table that shows supported update methods on the SilkWorm 12000. Pay attention to the information presented in this section, as there have been some significant changes in the behavior of `firmwaredownload` for Fabric OS Version 4.x.

3.4.1.1. Fabric OS Upgrade Overview

All updates from Fabric OS versions 2.x and 3.x to later versions are disruptive. This includes updates to Fabric OS versions 2.6.1 and 3.1 as well. For those versions of firmware, either RSH or FTP can be used to execute the update. Due to the fact that only one telnet session is allowed for those firmware versions, there is no `firmwaredownloadstatus` command.

In Fabric OS 4.x, `firmwaredownload` only allows the FTP protocol to be used for updates. Updating from Fabric OS 4.0.x to 4.1 is disruptive. So schedule downtime for these upgrades. All further updates from Fabric OS 4.1 to 4.1.x will be non-disruptive. For Fabric OS 4.0.0d or higher, `firmwaredownload` performs a check on the HA status. If it fails the check then the firmware cannot be updated automatically.

3.4.1.2. Fabric OS Upgrade Guidelines for 4.0.x to 4.1

For the SilkWorm 3900, only Fabric OS version 4.0.2 or higher is supported. When updating from Fabric OS 4.1 and beyond, the disruption will last longer on the SilkWorm 3900 as there is only one processor and the firmware must failover to itself. Essentially, this is because the Fabric OS Linux kernel and other processes that run in user space must be stopped and restarted through a `fastboot`. By default, a `firmwarecommit` is launched after the switch is rebooted. This process runs in the background and copies the new firmware from the flash memory primary partition (just updated) to the backup partition.

`Firmwaredownloadstatus` can be used on another telnet session to check the update progress. This essentially plays back the `firmwaredownload` log. The only message that will be received during the update is “Firmwaredownload has started.” After the `fastboot`, which happens automatically, run `firmwaredownloadstatus`. The total update time is about 13 minutes, including the `firmwarecommit`. Even though it takes 13 minutes on a SilkWorm 3900 to complete the upgrade, it will be non-disruptive. The fabric state will be saved in non-volatile memory, and after the reboot, all of the fabric services start up first.

```
int219:admin> firmwaredownloadstatus
[0]: Sun Mar 16 20:41:07 2003
Firmwaredownload has started.
[1]: Sun Mar 16 20:47:08 2003
Firmwaredownload has completed successfully.
[2]: Sun Mar 16 20:48:52 2003
Firmwarecommit has started.
[3]: Sun Mar 16 20:53:50 2003
Firmwarecommit has completed successfully.
[4]: Sun Mar 16 20:53:50 2003
Firmwaredownload command has completed successfully.
```

On the SilkWorm 12000, for Fabric OS versions 4.0.0 to 4.0.0c the `firmwaredownload` command must be used individually on each CP. To minimize downtime, start with the Standby CP and then do a failover with `hafailover`. Find out which is active by executing `hashow`. These procedures are outlined in the *Brocade Fabric OS Procedures Guide*.

For Fabric OS versions 4.0.0d or higher, `firmwaredownload`, which is invoked from a logical switch, will automatically update both logical switches in a SilkWorm 12000. This includes all updates for Fabric OS 4.1 or higher as well. When running it, the output of `firmwaredownload` does not show each package being downloaded like on the SilkWorm 3900. This makes monitoring the status difficult. Before doing the update, open a second telnet session on the logical switch and run `firmwaredownloadstatus`. This will show a log of each step in the process.

Guideline: One other change in Fabric OS 4.0.0d or higher is the `-s` option. This allows a single CP to be updated at a time. If more control over the update is desired, use this command first on the Standby and then the Active CP. This option is not supported on the SilkWorm 3900.

When complete, about 20 minutes, use `firmwareshow` to display current firmware versions.

```
poc165:admin> firmwareshow
Local CP (Slot 5, CP0): Active
    Primary partition:      v4.1.0
    Secondary Partition:    v4.1.0
Remote CP (Slot 6, CP1): Standby
    Primary partition:      v4.1.0
    Secondary Partition:    v4.1.0
```

When updating from Fabric OS 4.1 and beyond, the process is non-disruptive. This is much faster on the SilkWorm 12000 as there are two CPUs and two fabric state images, one per CP. The standby simply becomes the active during the failover. Thus, the impact on a production fabric is much less, when using a SilkWorm 12000. When updating from Fabric OS 4.1 and beyond, use `hashow` first before launching `firmwaredownload`. The `hashow` output should display as follows on Fabric OS 4.1.

```
poc165:admin> hashow
Local CP (Slot 6, CP1): Standby
Remote CP (Slot 5, CP0): Active
HA enabled, Heartbeat Up, HA State synchronize
```

This means that a `firmwaredownload` will be allowed to start non-disruptively. If `hashow` displays anything else, the `firmwaredownload` will not start.

Caution: Fabric OS 4.1 is supported only on the SilkWorm 3900 and 12000. Do not try to load 2.x or 3.x firmware on those switches or vice versa. Reference the *Brocade Fabric OS Procedure Guide* for Fabric OS upgrade process for your switch type. Fabric OS version 3.1 has a new feature called *firmware watermarking*. This prevents other Fabric OS firmware versions from writing over it.

3.4.2. Fabric OS Version 4.1 Ultra High Availability (HA)

With the introduction of Fabric OS version 4.1, there have been significant enhancements to the high availability characteristics. Brocade has worked with partners to perform extensive testing of this. There are new warnings that are provided by error logs, and SNMP traps. This section will provide some detail as to the nature of these enhancements and guidelines on how to verify that the new HA is functional.

Non-disruptive means that data traffic continues to flow during a failover or firmware updates. The following bullets provide the non-disruptive highlights that Fabric OS version 4.1 provides. Be aware that no other version of Fabric OS has non-disruptive capability. For this reason, it is highly recommended that those environments that require minimum downtime use the SilkWorm 3900 and SilkWorm 12000.

- Non-disruptive fail-over for the SW12000 Active and Standby CP Cards
- Non-disruptive firmware activation on both the SW12000 and SW3900 platforms.
- Introduction of a standby CP health monitor for the SW12000. This new function watches out for conditions that might prevent the fail-over mechanism from working.

Table 3-24

Ultra HA Feature	Data Flow Delay (Reads, Writes)	Fabric Services Delay (Name Server, FLOGI, etc.)
CP Failover (SW12000)	0 (seconds)	Between 5-10 seconds
Firmware Activation (SilkWorm 12000)	0 (seconds)	About 2 seconds
Firmware Activation (SilkWorm 3900)	0 (seconds)	Between 30-50 seconds

3.4.2.1. Ultra HA Functionality Overview

The mechanism for deciding when to fail-over is identical to Brocade Fabric OS 4.0.2. The major difference is that in Fabric OS 4.1 the fail-over process will continue to allow data traffic to flow. On the SilkWorm 3900 there is only one processor, so failover does not apply. For the SilkWorm 12000, the two CPs now maintain a synchronized state through Ethernet data packets transmitted across the backplane. This means that there is essentially a mirror image of the fabric configuration that is stored on each CP.

In order to facilitate this, the environment must be stable for synchronization to occur. After initial power up, the SilkWorm 12000 waits for a stable environment before attempting to synchronize the state between the two CPs. If the switch is standalone, then the system requires 3 X FS_TOV or 15 seconds. After this time, whenever a change is detected the Active CP updates the standby with the required information.

3.4.2.2. HA Caveats

Here is some key information and caveats about what Non-Disruptive means. Greatly enhanced HA only works for Fabric OS 4.1 and any updates from Fabric OS 4.0.x will be disruptive. Telnet commands run during Non-Disruptive failover should be reissued. Within a maximum 30 seconds, IP will restart. I/O continues through the fabric. Port LED Lights may stop blinking but I/O will still continue. The `firmwaredownload` process will still take the same amount of time as before. On the SilkWorm 3900 Non-Disruptive Code load, takes longer since there is only one processor. Linux must reboot

3.4.2.3. New HA Commands

As a result of the new HA capability there are several new commands that have been developed. The most important command is `hashow`. This tells the state of the HA present on the local switch. The older HA mode is still supported but not recommended. The four new commands are, `haShow`, `haSyncStart`, `haSyncStop`, and `hadump`. The rest of this section will talk to these commands and provide examples of usage.

3.4.2.3.1 Using haShow

This is `haShow` when both of the CPs on the SilkWorm 12000 are properly synchronized.

```
poc165:admin> hashow
Local CP (Slot 5, CP0): Active
Remote CP (Slot 6, CP1): Standby, Healthy
HA Enabled, Heartbeat Up, State Synchronized
```

When the SilkWorm 12000 does not have synchronized CPs, `haShow` will display the following output:

```
poc165:admin> hashow
Local CP (Slot 5, CP0): Active
Remote CP (Slot 6, CP1): Standby, Healthy
HA Enabled, Heartbeat Up, HA Not Synchronized
```

3.4.2.3.2 Guidelines with using haSyncStop, haSyncStart, and hadump

The command `hasyncstop` should be used only for troubleshooting purposes, as it will temporarily turn off non-disruptive failover or non-disruptive code activation. This means that the next fail-over or firmware download and activation *will* be disruptive. When `haSyncStop` is run, all ports WILL be reset. The command will only be effective until next fail-over or if the command `haSyncStart` is issued. Note that the two CPs will re-synchronize as normal only after the next failover. Therefore this command should not be executed command during day to day operations. Here is an example of using the command `haSyncStop`.

```
poc165:admin> hasyncstop
Stop synchronize 0x228 (fabos): Switch: 0, Info FSS_ME-FORCELOG, 4, Software out of sync!
```

`HaSyncStart` is used to re-activate synchronous HA between the CPs. It only needs to be executed after an `haSyncStop`. Here is an example of usage and output.

```
poc165:admin> hasyncstart
Start synchronize ...
done
poc165:admin> 0x223 (fabos): Switch: 0, Info FSS_ME-FORCELOG, 4, Software is in sync!
```

`Hadump` is a supportability command that displays HA related logging information. `Hadump` should not have to be used by the general SAN administrator. It is really meant to be a troubleshooting tool for technical support.

3.4.3. Guidelines for using other Fabric OS 2.6.1/3.1/4.1 Commands

This section contains additional new features and functions within the new Fabric OS family. Many new commands and the usage will be discussed. Most new features and functions apply to Fabric OS 3.1/4.1 only.

3.4.3.1. Nscamshow (Fabric OS 2.6.1/3.1/4.1)

Nscamshow shows the non-local device name server registration on all of the other switches in the fabric. This is useful to identify other devices on other switches. The count variable shows the number of devices that are registered. Note that the core switch, int63 in the example below, does not have any name server entries, as it only has other switches attached to it. Not every switch is shown, to conserve space.

```
int170:admin> nscamshow
Switch entry for 63
state      rev  owner
known      v410 0xffffcaa
Device list: count 0
No entry is found!

Switch entry for 75
state      rev  owner
known      v260 0xffffcaa
Device list: count 1
Type Pid  COS      PortName      NodeName
N 4b0f00;  3;20:05:00:a0:b8:07:5d:c7;20:04:00:a0:b8:07:5d:c6;
FC4s: FCP
Fabric Port Name: 20:0f:00:60:69:10:10:9d

Switch entry for 88
state      rev  owner
known      v310 0xffffcaa
Device list: count 1
Type Pid  COS      PortName      NodeName
N 580800;  2,3;10:00:00:00:c9:29:04:8f;20:00:00:00:c9:29:04:8f;
FC4s: FCP
Fabric Port Name: 20:08:00:60:69:51:0e:0a

Switch entry for 154
state      rev  owner
known      v260 0xffffcaa
Device list: count 2
Type Pid  COS      PortName      NodeName
N 9a0800;  2,3;10:00:00:00:c9:24:f6:5d;20:00:00:00:c9:24:f6:5d;
FC4s: FCP
Fabric Port Name: 20:08:00:60:69:10:68:fa
N 9a0a00;  3;50:06:0b:00:00:0a:48:4a;50:06:0b:00:00:0a:48:4b;
FC4s: FCP
Fabric Port Name: 20:0a:00:60:69:10:68:fa
```

3.4.3.2. Portloginshow (Fabric OS 2.6.1/3.1/4.1)

There is a handy command called `portloginshow` that displays the Fibre Channel services for which a device has registered. In the command output, there will be one login record per service. As shown in the example, which in this case is an HBA, there would be two entries. One is for the FLOGI to the well-known address FFFFEE, which is the Fabric Port Login service for fabric registration, (that provides the 24-bit PID address). The other is for the PLOGI to the well-known address FFFFEC, which is the name server.

```
int195:admin> portloginshow 6
Type  PID      World Wide Name      credit df_sz cos
-----
fe  c30600  10:00:00:00:c9:24:f5:f9  64  2048  c  scr=3
fc  c30600  10:00:00:00:c9:24:f5:f9  12  2048  c  c2sema=0x100ba770
```

The credit on the first entry is in reference to BB_Credit (buffer-to-buffer credit), on the second it references the EE_Credit (end-to-end credit). COS represents the class of service a bit map column that is displayed in hex. The SCR is in reference to a state change registration with a class of 3.

This next example shows an HBA. The `df_size` refers to the FC maximum payload size, for the HBA it is 2048 bytes. For HBAs, this is a typical value.

```
sialab90:admin> portloginshow 8
Type  PID      World Wide Name      credit df_sz cos
-----
fe  5a0800  10:00:00:00:c9:2b:4f:75  16  2048  c  scr=3
```

This command is also useful to display all of the devices that are logged into a port. For JBODs, which are Fibre Channel loop devices, this is invaluable as each PID (with the `AL_PA`) is shown. In the example, a JBOD is attached to port 15 on a SilkWorm 3800 with Fabric OS 3.1. Note that the `AL_PA` and port WWNs are in a tabular list and thus are easily identified. `Df_size` refers to the maximum FC payload size, for the JBOD it is 2112 bytes, the largest possible.

```
int170:admin> portloginshow 15
Type  PID      World Wide Name      credit df_sz cos
-----
ff  aa0fef  22:00:00:20:37:15:1f:f7  0  2112  8
ff  aa0fe8  22:00:00:20:37:e6:9a:a3  0  2112  8
ff  aa0fe4  22:00:00:20:37:e6:9a:84  0  2112  8
ff  aa0fe2  22:00:00:20:37:e6:9a:7d  0  2112  8
ff  aa0fe1  22:00:00:20:37:e6:99:be  0  2112  8
ff  aa0fe0  22:00:00:20:37:15:1f:d6  0  2112  8
ff  aa0fdc  22:00:00:20:37:15:1d:51  0  2112  8
ff  aa0fda  22:00:00:20:37:15:21:21  0  2112  8
ff  aa0fd9  22:00:00:20:37:15:0b:a5  0  2112  8
ff  aa0fd6  22:00:00:20:37:e6:9a:6c  0  2112  8
ff  aa0fd5  22:00:00:20:37:e6:9a:8c  0  2112  8
ff  aa0fd4  22:00:00:20:37:15:0c:0b  0  2112  8
ff  aa0fd3  22:00:00:20:37:15:0a:a5  0  2112  8
ff  aa0fd2  22:00:00:20:37:15:1a:f9  0  2112  8
ff  aa0fd1  22:00:00:20:37:15:1f:d0  0  2112  8
int170:admin>
```

3.4.3.3. PortCamShow (Fabric OS 3.1/4.1 Only)

In some environments, the total cached limit of 64 SIDs and 512 DIDs per quad may be an issue. In order to address this need, Brocade created a new command to monitor these values. This command is called `PortCamShow`. It displays the number of FC SIDs and DIDs used out of the maximum possible. One useful tip is to use `nsaliasshow` along with `PortCamShow` to identify what devices are allowed to see each other. In this example, there is an HBA zoned to a single storage port by port WWN.

First use `nsshow` to find out the PID of the HBA and in which quad it is located. For this case, it is 5a0800.

```
sialab90:admin> nsaliasshow
{
  Type Pid      COS      PortName          NodeName          TTL(sec)
  N     5a0800;   2,3;10:00:00:00:c9:2b:4f:75;20:00:00:00:c9:2b:4f:75; na
    FC4s: FCP
    Fabric Port Name: 20:08:00:60:69:51:0f:2b
    Aliases: int124_w2k_B
The Local Name Server has 1 entry }
```

`PortCamShow` with no argument displays the used SID/DID on every port as shown.

```
sialab90:admin> portcamshow
Port 0 to Port 15:
-----
Port      SID used  DID used
0         0         0
1         0         0
2         0         0
3         0         0
4         0         0
5         0         0
6         0         0
7         0         0
8         2         2
9         0         0
10        0         0
11        0         0
12        0         0
13        0         0
14        0         0
15        0         0
-----
Quad ports (SID Free, DID Free)
0-3 (64, 512)  4-7 (64, 512)  8-11 (64, 512)  12-15 (64, 512)
```

Using the port number as the argument for `PortCamShow` (in this case 8) shows the number of SID/DID used, 24-bit address of each device that is in the SID and DID tables and the total SID/DID entries that are free.

```
sialab90:admin> portcamshow 8
-----
Port      SID used  DID used  SID entries  DID entries
8         2         2         5a0800      5a0800
          c40d00      5a0800
-----
Quad ports (SID Free, DID Free)
8-11 (62, 510)
```

From the PortCamShow table, note the C40d00 device. To find out what it is, telnet to the switch with domain ID C4 hex (196). From the output of nsaliasshow, the device is identified.

```
int196:admin> nsaliasshow
{
  Type Pid      COS      PortName      NodeName      TTL(sec)
  N   c40d00;    3;20:04:00:a0:b8:07:5d:c7;20:04:00:a0:b8:07:5d:c6; na
      FC4s: FCP [LSI      INF-01-00      0401]
      Fabric Port Name: 20:0d:00:60:69:10:62:2c
      Aliases: LSI_B

The Local Name Server has 1 entry }
```

3.4.3.4. Portzonestow (Fabric OS 3.1/4.1 Only)

Portzonestow is a new command that shows how the zoning on the ports is enforced. An example of its usage is shown below. From the output shown in the example, note that port 8 is HARD enforced by WWN. This is shown in bold. Also note that although not able to be zoned E_Ports are displayed as well.

```
sialab88:admin> portzonestow

Local Zoning Port-level information:

PORT: 0  enforcement: E-Port      defaultSoftZone: 0      defaultHardZone:0
PORT: 1  enforcement: E-Port      defaultSoftZone: 0      defaultHardZone:0
PORT: 2  enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 3  enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 4  enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 5  enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 6  enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 7  enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 8  enforcement: HARD WWN    defaultSoftZone: 0      defaultHardZone: 0
PORT: 9  enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 10 enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 11 enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 12 enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 13 enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 14 enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
PORT: 15 enforcement: Not Zoned   defaultSoftZone: 0      defaultHardZone: 0
```

This is the output on a SilkWorm 3900. Note there is some additional information, as the port type is shown. Only partial output is shown.

```
int219:root> portzonestow

PORT: 26  Enforcement: HARD PORT      defaultHard: 1  F-port: 1
PORT: 27  Enforcement: HARD PORT      defaultHard: 1  F-port: 1
PORT: 28  Not Zoned
PORT: 29  Not Zoned
PORT: 30  Not Zoned
PORT: 31  Not Zoned
```

3.4.3.5. NodeFind (Fabric OS 3.1 and 4.1 Only)

Nodefind is a new command in Fabric OS 3.1/4.1. It provides device searching capability within a single fabric. An example is shown below. In order to find the location of a node, the WWN of interest must be known. The command returns the 24-bit port ID of the device. The first 8-bits is the switch domain. The next 8-bits is the port number.

```
sialab89:admin> nodefind "10:00:00:00:c9:30:d0:62"
{
  db1a00
  1 device with wwn 10:00:00:00:c9:30:d0:62 }
sialab89:admin>
```

In this example the Domain ID is db and the port number is 1a. To get effective use of the command, issue `fabricshow`. `Fabricshow` assists with finding the switch the device is on. Looking at the ID column (bold) reveals that the switch name where the device is located is int219. The port number is 1a hex which is 26 in decimal. Thus the device is on switch int219, port 26.

```
sialab89:admin> fabricshow
Switch ID      Worldwide Name          Enet IP Addr    FC IP Addr      Name
-----
64: fffc40 10:00:00:60:69:80:4d:fd 192.168.162.64  0.0.0.0         "int64"
89: fffc59 10:00:00:60:69:51:10:42 192.168.155.89  0.0.0.0         "sialab89"
90: fffc5a 10:00:00:60:69:51:0f:2b 192.168.155.90  0.0.0.0         "sialab90"
195: fffcc3 10:00:00:60:69:10:93:14 192.168.162.195 0.0.0.0         "int195"
196: fffcc4 10:00:00:60:69:10:62:2c 192.168.162.196 0.0.0.0         "int196"
219: fffcdb 10:00:00:60:69:90:04:1a 192.168.162.219 0.0.0.0         >"int219"
```

The Fabric has 6 switches

3.4.3.6. Nsaliasshow (Fabric OS 3.1 and 4.1 Only)

Nsaliasshow is very much like `nsshow`. The only difference is that the zone alias is now displayed in the output. Be aware that this is not the port name assigned by `portname`. The example shows the Alias name assigned through zoning in bold.

```
sialab90:admin> nsaliasshow
{
  Type Pid      COS      PortName          NodeName          TTL(sec)
  N      5a0800;    2,3;10:00:00:00:c9:2b:4f:75;20:00:00:00:c9:2b:4f:75; na
  FC4s: FCP
  Fabric Port Name: 20:08:00:60:69:51:0f:2b
  Aliases: int124_w2k_B
}
```

The Local Name Server has 1 entry }

3.4.3.7. Cfgactvshow (Fabric OS 3.1 and 4.1 only)

One of the complaints about `cfgshow` (or `zonestow`) was that the output was too long. It showed all the alias names and member, zone names and members and every zoning configuration defined. `Cfgactvshow` addresses this concern. It displays only the active configuration and the zones that are defined that are part of it. As shown in the example, this provides a summary of which devices are defined in the active zoning configuration.

```
sialab90:admin> cfgactvshow
Effective configuration:
  cfg:  bkup_cfg_B
  zone:  int121_B
          10:00:00:00:c9:24:f5:f9
          20:04:00:a0:b8:07:5d:c7
  zone:  int124_B
          10:00:00:00:c9:2b:4f:75
          20:04:00:a0:b8:07:5d:c7
  zone:  int202_B
          50:06:0b:00:00:0a:48:84
          20:04:00:a0:b8:07:5d:c7
```

3.4.3.8. NsZoneMember (Fabric OS 3.1 and 4.1 Only)

`NszoneMember` displays the zoned members, both local and remote, for a particular device. For the purposes of this command, the port ID identifies the device for which the zoned members are to be identified.

First determine the port ID (PID). To find this information for a particular device on the fabric that is of interest use `nsaliasshow`. This will not only show the PID but the device zoning alias name as well. In the example, an HBA was picked on fabric A that is in the host int124. The desire was to find out what storage was supposed to be seen.

```
sialab88:admin> nsaliasshow
{
  Type Pid      COS      PortName                               NodeName                               TTL(sec)
  N      580800;   2,3;10:00:00:00:c9:29:04:8f;20:00:00:00:c9:29:04:8f; na
  FC4s: FCP
  Fabric Port Name: 20:08:00:60:69:51:0e:0a
  Aliases: int124_w2k_A
}
```

The Local Name Server has 1 entry }

Once the PID of the device is known, use it as the argument of `NszoneMember`. Note that in this case there are no local or remote storage devices zoned with the HBA.

```
sialab88:admin> nszoneMember "580800"
No local zoned members
No remote zoned members
```

3.4.3.9. Port Swapping (Fabric OS 4.1 Only)

A new function of Fabric OS 4.1 is the ability to do port swapping on a local switch in a fabric. For users of Operating Systems that use the PID for persistent binding, this new function adds a lot of value. Now, if for troubleshooting purposes, the port of the host HBA needs to be moved on the same switch, the cable can be moved without a reboot of the host. An example will be shown to illustrate the steps required to perform a swap. This section will provide an overview of port swapping and provide guidelines on its usage.

Port Swap Functionality

Swapping ports re-assigns the area numbers for the pair ports being swapped. That is, after the swap, port A will get port B's area number. In order to accomplish this, both switch ports must be disabled first. When complete, the result of the `Portswap` operation is persistent. This means that the swapped area numbers stay swapped across reboots and power cycles. `Portswap` information is kept in separate data base and can not be manipulated by editing the switch configuration database. By default, port swapping is turned off. To turn on port swapping, use `portswapenable`. This allows port swapping to be configured. Disabling the ability to swap ports is done with `portswapdisable`. This does not allow further ports to be swapped BUT does keep the current swapped area numbers in place. To reset the ports with the original area numbers, port swap must be done again to undo the configuration. To view the current ports being swapped, use `portswapshow`. The next two examples illustrate these points and show how to use the command.

Port Swap Examples

1. Enable port swapping with the command `portswapenable`.

```
int219:admin> portswapenable
```

2. Disable the ports that are to be swapped. In this case, the ports are 14 and 15. For the SilkWorm 12000 use physical slot number/port number. As an example, `portdisable 1/14` will disable port 14 on physical slot 1.

```
int219:admin> portdisable 14
int219:admin> portdisable 15
```

3. Swap the ports with `portswap`. A few moments pass before the configuration completes.

```
int219:admin> portswap 14 15
portswap done
```

4. Check to see if the port area numbers are swapped with `portswapshow`. Note that the swapping is persistent across power cycles and switch reboots.

```
int219:admin> portswapshow
PortSwap is enabled
Port          Area
=====
14             15
15             14
```

This completes the swapping of the ports. The cables can now be moved.

5. The final step is to enable both of the ports. Note that the devices will redo the FLOGI and PLOGI as well as go through the normal registration process.

```
int219:admin> portenable 14
int219:admin> portenable 15
```

To disable the creation of new port swapping configurations, use `portswapdisable`. It is important to note that the original swap still is in effect as shown with `portswapshow`. Once again, this configuration is persistent across power cycles and reboots.

```
int219:admin> portswapdisable
```

```
int219:admin> portswapshow
PortSwap is disabled.
Existing Portswap condition is still effective.
Only future Portswap operations are not allowed.
```

```
Port          Area
=====
14             15
15             14
```


Undoing the port swap

This next procedure shows how to undo the port swap done earlier.

1. Make sure that port swapping is enabled.

```
int219:admin> portswapenable
```

2. Disable both ports that are to be re-configured to their original Domain IDs.

```
int219:admin> portdisable 14
int219:admin> portdisable 15
```

3. Do the swap to return to the original area numbers.

```
int219:admin> portswap 14 15
portswap done
```

4. Check the swap configuration to verify that the ports now have their original area numbers.

```
int219:admin> portswapshow
PortSwap is enabled
Port          Area
=====
No ports have been swapped
```

5. Optional, disable port swapping using portswapdisable. This is the default setting for Fabric OS 4.1.

```
int219:admin> portswapdisable
int219:admin> portswapshow
PortSwap is disabled.
Existing Portswap condition is still effective.
Only future Portswap operations are not allowed.
```

```
Port          Area
=====
No ports have been swapped
```

6. Move the cables back to the original locations.
7. Enable each of the affected ports. Once again, the devices will do the appropriate FLOGI and PLOGI.

```
int219:admin> portenable 14
int219:admin> portenable 15
```

Note: The portswap information is kept in separate database and thus cannot be manipulated by editing the regular switch configuration database.

3.4.3.10. Switch Rebooting Methods

This section will provide a high level overview of methods of rebooting a switch. Also, there is a table that summarizes the amount of time each method requires. All methods of rebooting are meant to be disruptive.

The command `reboot` works on all platforms. When it is issued, the switch shuts down gracefully and restarts. As it comes online, all of the POST and self-diagnostic tests are run. There are a about 12 of them. This takes quite a bit of time.

```
int219:admin> reboot
```

The command `fastboot` works on all switch platforms as well. This is the same as doing a `reboot` with diagnostics turned off with `diagdisablepost`. Diagnostics can be re-enabled with `diagenablepost`. This reduces the time for it to come back online significantly.

```
int219:admin> fastboot
```

New SwitchReboot for the SilkWorm 12000 Only

For the SilkWorm 12000 only there is a new command, `Switchreboot`. Used to reboot a single logical switch image and will always be disruptive. `Switchreboot` was designed to clean up and restart all Fabric OS applications so the disruption was planned. An example of a `switchreboot` is shown below. The SilkWorm 3900 only supports the `reboot` or `fastboot` command.

```
int64:admin> switchreboot
int64:admin> Selecting i2c bus...Done.
Stopping all switch daemons...Done.
Releasing i2c bus...Done.
Powering off slot 7...Done.
Powering off slot 8...Done.
Powering off slot 9...Done.
Powering off slot 10...Done.
Checking all slots are powered off....Done.
Cleaning up kernel modules...Done.
Done.
Powering on slot 7...Done.
Powering on slot 8...Done.
Powering on slot 9...Done.
Powering on slot 10...Done.
Checking diagnostics.....
.....Done.
10 9 8
fabric: Subordinate switch
fabric: Reconfiguration due to Fabric Merge(port 47)
fabric: Reconfiguring at Tue Mar 25 06:48:15 2003

5 4 3 2 1
fabric: Subordinate switch
fabric: Domain 64
```

Table 3-25 summarizes the reboot methods for the each switch in the SilkWorm family.

Table 3-25 Reboot Methods

Switch Model	Fabric OS Version	Reboot Command
SilkWorm 2800	Fabric OS 2.6.1	reboot fastboot
SilkWorm 3200	Fabric OS 3.1	reboot fastboot
SilkWorm 3800	Fabric OS 3.1	Reboot fastboot
SilkWorm 3900	Fabric OS 4.1	Reboot fastboot
SilkWorm 12000	Fabric OS 4.1	Reboot Fastboot switchreboot

Guideline: On the SilkWorm 12000, it is recommended that if an Active CP requires a reboot, that an hfailover is executed. This makes the process non-disruptive. That way, the other CP becomes the Active and the new standby can be rebooted non-disruptively.

3.4.4. Configuring Extended Fabrics

Fabric OS v3.0.2 and above contains an additional optional parameter, VC translation link initialization, to the `portCfgLongDistance` CLI command. When set to **1**, this parameter indicates that the enhanced link reset protocol should be used on the port. The default value for this parameter is **0** and is compatible with earlier Fabric OS v3.x implementations. For optimal performance, specify **1** when E_port links are between switches with Fabric OS v3.0.2 and above. Specify **0**, or nothing, when connecting to previous releases of Fabric OS.

To configure Extended Fabrics on SilkWorm 2000 Series switches, the following procedures must be done.

- Set the fabric wide configuration parameter `fabric.ops.mode.longDistance` to **1** using the `configure` command. This parameter must be set on all switches within the fabric.
- Set the proper EF mode for each long distance port using the `portCfgLongDistance` command.

When configuring Extended Fabrics on SilkWorm 3000 Series switches and above, only port level configuration is necessary.

- Set the proper EF mode for each long distance port using the `portCfgLongDistance` command. At the same time, set the VC translation link initialization bit to **1** if the long distance switches are at Fabric OS v3.0.2 or above.

For mixed SilkWorm 2000 and 3000/12000 Series fabric configurations, where the long distance ports are located on SilkWorm 2000 Series switches, the fabric wide parameter `fabric.ops.mode.longDistance` must be set to a value of **1**. This does not apply to mixed fabrics where long distance ports are located on SilkWorm 3000 or 12000 series switches.

The following output shows how to configure port 0 on a SilkWorm 3800 switch running Fabric OS v3.1 for L2 Extended Fabrics mode with VC translation link initialization set to 1.

```
mw123:root> portcfglongdistance
Usage: portCfgLongDistance port ,"distance_level",<vc translation link init>
distance_level:      L0      - normal
                    LE      <= 10km
                    L0.5    <= 25km
                    L1      <= 50km
                    L2      <= 100km
                    LD      - auto
vc trans link init:  0 normal
1 vc translation
mw123:root> portcfglongdistance 0,"L2",1
Committing configuration...done.
```

Figure 3-28 Configuring a port for Long Distance

Use the `portCfgShow` command to determine which ports are configured as long distance ports. The following switch output shows that port 0 has been configured for L2 Extended Fabrics and VC Translation Link Initialization has been turned on.

```
mw123:root> portcfgshow
Ports          0  1  2  3   4  5  6  7   8  9 10 11   12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
Speed          AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Trunk Port     ON ON ON ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance  L2 .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
VC link init   ON .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked L_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked G_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Disabled E_Port .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Persistent Disable .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
               where AN:AutoNegotiate, ..:OFF, ?:INVALID.   LM:L0.5
```

Figure 3-29 Portcfgshow Output

3.4.4.1. Fibre Channel Performance over Distance

Performance over long distance links can vary for multiple reasons. The number of buffer-to-buffer credits determines the number of FC frames that a switch can transmit on a link at one time before requiring an acknowledgement back from the receiver, thus performance degradation may occur if not enough credits are available. This may or may not be a concern depending on application performance requirements. For example, at 25 km, L0.5 EF mode utilizes 34 credits for 2Gbit/sec Fibre Channel. If the application using the link is only capable of transmitting 100 MB/sec, only 17 credits would be necessary and LE mode, which does not require a license, could be used.

Another performance factor on long distance links is response times for SCSI transactions. For example, in class 3 traffic, every write transaction requires two round trips; one for the SCSI write request command, and another for the transmission of data. Transmission of data cannot occur until the initiator receives a transfer ready response. Response time can be calculated by multiplying the link propagation time by 4. The effect of response times can be minimized if an application allows the number of outstanding I/Os to be increased.

3.4.5. Using Fabric OS Troubleshooting Tools

This section will just provide an overview of the new and changed Brocade Fabric OS troubleshooting tools. Some examples will be given to demonstrate the practical use. There have been big changes with error logging and the use of `portlogdump`. For details and examples for setting up persistent error logs and the new uses of `portlogdump`, please see the *Fabric OS Procedures Guide*.

3.4.5.1. Support Show Command Groups (Fabric OS 3.1/4.1)

`Supportshow` has been an invaluable troubleshooting tool. In previous versions of Fabric OS, `Supportshow` would provide output that was not configurable. This would generate more output than needed. In operating the new version of `Supportshow`, it can now be run with any combination of 11 command groups. This makes `Supportshow` more flexible and easier to capture the desired information.

If required for support reasons, `supportshow` can always be reconfigured to display more information. As an example of a command group, here are the commands for the FC Fabric command group 4.

Table 3-26 Command Group 4

Command Group 4: FC Fabric	
<code>fabricShow</code>	<code>qlShow</code>
<code>islShow</code>	<code>cfgShow</code>
<code>trunkShow</code>	<code>fabStatsShow</code>
<code>topologyShow</code>	<code>fabLogDump</code>
<code>faShow</code>	

This section will provide some guidelines on setting up `supportshow` command groups and make a recommendation on which ones to use. By default, eight command groups are enabled for `supportshow`. These are shown by `supportshowcfgshow` output below. This command can be run as `admin` or `root`. Changes to the `supportshow` configuration can only be made as the `root` user.

```
int219:root> supportshowcfgshow
os          enabled
exception  enabled
port       enabled
fabric     enabled
services   enabled
security   enabled
network    enabled
portlog    enabled
system     enabled
extend     disabled
filter     disabled
perfmon    disabled
```

Figure 3-30 `supportshowcfgshow` output

If you have the `root` password consider configuring `supportshow` to disable the following groups. For general information about the fabric and the switch `supportshow` is running on, these groups are really all that is needed. For most SAN administrative troubleshooting cases, the data provided by the remaining groups will do.

```
int219:root> supportshowcfgdisable "os"
Config update Succeeded
int219:root> supportshowcfgdisable "port"
Config update Succeeded
int219:root> supportshowcfgdisable "security"
Config update Succeeded
int219:root> supportshowcfgdisable "portlog"
Config update Succeeded
int219:root> supportshowcfgdisable "network"
Config update Succeeded
```

Note: If using Secure Fabric OS, do not disable the security command group as shown. Even the remaining groups will generate a lot of output.

Finally, verify the setting with `supportshowcfgshow`. The output that should be seen is shown in the following example.

```
int219:root> supportshowcfgshow
os          disabled
exception  enabled
port       disabled
fabric     enabled
services   enabled
security   disabled
network    disabled
portlog    disabled
system     enabled
extend     disabled
filter     disabled
perfmon    disabled
int219:root>
```

Figure 3-31 `Supportshowcfgshow` Command Output

To clarify the groups being used, please refer to Table 3-27. The security group is for Secure Fabric OS support information and it is group 6.

Table 3-27 Group Name and Number

Group Name	Group Number
exception	2
fabric	4
services	5
security	6
system	10

Below is a summary of the `supportshow` configuration commands and the function they perform. Once again, only the root user is allowed to make changes.

Table 3-28 Supportshow Configuration Commands

Supportshow Configuration Command	Function
<code>supportShowCfgShow</code>	Displays list of command groups and whether they are enabled.
<code>supportShowCfgEnable</code>	Allows root user to enable a single command group
<code>supportShowCfgDisable</code>	Allows root user to disable a single command group

For all `supportShow` output, no matter how its configured, there are certain commands that will always be executed. These are, in the order of execution date, version, and `supportshowcfgshow`.

Note: Enabling and disabling is persistent except for filter and extended groups. This is not a big deal in that those two command groups are rarely needed anyway. There are two command groups that have only one member. The exception group has `errdump`. The portlog group only has `portlogdump`.

3.4.5.2. New Track Changes (Fabric OS 3.1/4.1)

There is a new change management tool called `trackChangesSet`. This command allows the SAN administrator to track successful and unsuccessful logins, logouts and configfile changes via a log file or SNMP on a local switch. It is disabled by default and must be administratively turned on by `trackChangesSet`. Refer to *trackChangesSet* on page 4-4.

This example shows `trackchangesshow`. It displays the default setting.

```
int170:admin> trackChangesShow
Track changes status: OFF
Track changes generate SNMP-TRAP: NO
```

The example shows how to turn on track changes but not use SNMP to send traps.

```
int170:admin> trackChangesSet 1, 0
0x102d7450 (tShell): Mar 28 04:12:08
    INFO TRACK-TRACK_ON, 4, Track-changes on

Committing configuration...done.
0x102d7450 (tShell): Mar 28 04:12:11
    INFO TRACK-CONFIG_CHANGE, 4, Config file change from task:tShell
```

Now `trackchangesshow` now shows the logging capability enabled. Note that SNMP is still turned off.

```
int170:admin> trackchangesshow
Track changes status: ON
Track changes generate SNMP-TRAP: NO
```

With failed logins, the information is logged to the screen after the unsuccessful login as shown .

```
Login incorrect
0x102d7450 (tShell): Mar 28 04:23:50
    INFO TRACK-FAILED_LOGIN, 4, Unsuccessful login
```

Use `errshow` to display the changes. As stated in the *Fabric OS Procedures Guide*, by default there are only 1024 entries. By default, the `errlog` is cleared after a reboot. With Fabric OS 3.1 and 4.1, this log can be expand to 2048 entries maximum and set to be persistent.

```
int170:admin> errshow
Error 27
-----
0x102d7450 (tShell): Mar 28 04:25:31
    INFO TRACK-LOGIN, 4, Successful login

Type <CR> to continue, Q<CR> to stop:
Error 26
-----
0x102d7450 (tShell): Mar 28 04:23:50
    INFO TRACK-FAILED_LOGIN, 4, Unsuccessful login

Error 25
-----
0x102d7450 (tShell): Mar 28 04:23:50
    INFO SEC-SECVIOL_LOGIN, 4, Security violation: Login failure attempt via TEL
NET. Peer IP: 192.168.162.211
.
Error 15
-----
0x102d7450 (tShell): Mar 28 04:20:02
    INFO TRACK-LOGIN, 4, Successful login

Error 14
-----
0x102d7450 (tShell): Mar 28 04:19:53
    INFO TRACK-LOGOUT, 4, Logout
```

Note: With Fabric OS 3.1/4.1 three unsuccessful logins will disconnect the current telnet session.

3.4.5.3. New Persistently Disabling a Switch or Port (Fabric OS 3.1/4.1)

There are four new commands to allow for persistently disabling ports or switches. When configured, the state of the switch or port will remain disabled through power cycles or reboots. There are two reasons why this may be done. First, there may be a bad SFP or switch that causes fabric instability. These may need to be brought down temporarily until a replacement is found. The second reason is that unused ports maybe persistently disabled for security concerns. When in this state, no device or switch will be allowed to join the fabric on that port. Below are examples of usage.

To disable a switch persistently, use `switchCfgPersistentDisable`. After a few moments, the switch is disabled. To verify, use `switchshow` to display the current state. Note that the `SwitchRole` is now `Disabled (Persistent)`. This indicates that the command has taken effect.

```
int170:admin> switchCfgPersistentDisable
Committing configuration...done.
```

```

int170:admin> switchshow
switchName:      int170
switchType:      9.1
switchState:     Offline
switchMode:      Native
switchRole:      Disabled (Persistent)
switchDomain:    170 (unconfirmed)
switchId:        fffcaa
switchWwn:       10:00:00:60:69:50:10:90
switchBeacon:    OFF
Zoning:         ON (bkup_cfg_A)
port 0: id N2 In_Sync      Disabled
port 1: id N2 In_Sync      Disabled
port 2: -- N2 No_Module    Disabled
port 3: -- N2 No_Module    Disabled
port 4: id N2 No_Light     Disabled
port 5: id N2 No_Light     Disabled
port 6: -- N2 No_Module    Disabled
port 7: -- N2 No_Module    Disabled
port 8: -- N2 No_Module    Disabled
port 9: id N2 No_Light     Disabled
port 10: id N2 No_Light    Disabled
port 11: id N2 No_Light    Disabled
port 12: id N2 No_Light    Disabled
port 13: id N2 No_Light    Disabled
port 14: -- N2 No_Module    Disabled
port 15: id N2 In_Sync     Disabled

```

Figure 3-32 Disabling a switch

Note that portcfgshow still shows all ports as not disabled. This is fine, since the switch as a whole is disabled.

```

nt170:admin> portcfgshow
Ports          0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
Speed          AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Trunk Port     ON ON ON ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
VC link init   .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked L_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked G_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Disabled E_Port .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Persistent Disable .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
ISL R_RDY Mode .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
              where AN:AutoNegotiate, ..:OFF, ?:INVALID.
              LM:L0.5

```

Figure 3-33 Portcfgshow Output

To re-enable the switch persistently use `switchCfgPersistentEnable`. Note that the fabric must do a re-configure after it is brought back online.

```

int170:admin> switchCfgPersistentEnable
Committing configuration...done.
Command in progress

fabric: Subordinate switch
fabric: Domain 170
. . . . . done

```

Figure 3-34 switchCfgPersistentEnable Output

Note: If the switch is re-enabled with `switchenable` it will be enabled temporarily. The next power cycle, reboot or `fastboot` will cause the switch to be disabled, persistently. This is because the state of the switch is now stored in the flash non-volatile memory.

Use `portCfgPersistentDisable` to persistently disable a port. Use `portcfgshow` to check the status. An example of this that shows port 7 being disabled persistently is shown next.

```
int170:admin> portCfgPersistentDisable 7
Committing configuration...done.

int170:admin> portcfgshow
Ports          0  1  2  3    4  5  6  7    8  9 10 11    12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed          AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Trunk Port     ON ON ON ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
VC link init   .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked L_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked G_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Disabled E_Port .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Persistent Disable .. .. .. ..  .. .. .. ON  .. .. .. ..  .. .. .. ..
ISL R_RDY Mode .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
              where AN:AutoNegotiate, ..:OFF, ??:INVALID.
              LM:L0.5
```

Figure 3-35 portCfgPersistentDisable Output

To re-enable a port persistently, use `portcfgpersistentenable`. The port may be re-enabled temporarily by `portenable`.

```
int170:admin> portcfgpersistentenable 7
Committing configuration...done.
int170:admin> portcfgshow
Ports          0  1  2  3    4  5  6  7    8  9 10 11    12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed          AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Trunk Port     ON ON ON ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
VC link init   .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked L_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked G_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Disabled E_Port .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Persistent Disable .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
ISL R_RDY Mode .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
              where AN:AutoNegotiate, ..:OFF, ??:INVALID.
              LM:L0.5
```

Please note that a `portcfgdefault` will turn off persistent disablement of a port. Note that the port will not be enabled automatically. A `portenable` is required. `Portcfgdefault` will set all other port settings to default values. This sequence is shown in Figure 3-36.

```
int170:admin> portCfgPersistentDisable 7
Committing configuration...done.
int170:admin>

int170:admin> portcfgdefault 1
Committing configuration...done.
int170:admin> portcfgshow
Ports          0  1  2  3    4  5  6  7    8  9 10 11    12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed          AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Trunk Port     ON ON ON ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
VC link init   .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked L_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Locked G_Port  .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Disabled E_Port .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
Persistent Disable .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
ISL R_RDY Mode .. .. .. ..  .. .. .. ..  .. .. .. ..  .. .. .. ..
              where AN:AutoNegotiate, ..:OFF, ??:INVALID.
              LM:L0.5
```

```

Speed                AN AN AN AN    AN AN AN AN    AN AN AN AN    AN AN AN AN
Trunk Port           ON ON ON ON    ON ON ON ON    ON ON ON ON    ON ON ON ON
Long Distance        .. .. .. ..    .. .. .. ..    .. .. .. ..    .. .. .. ..
VC link init         .. .. .. ..    .. .. .. ..    .. .. .. ..    .. .. .. ..
Locked L_Port        .. .. .. ..    .. .. .. ..    .. .. .. ..    .. .. .. ..
Locked G_Port        .. .. .. ..    .. .. .. ..    .. .. .. ..    .. .. .. ..
Disabled E_Port      .. .. .. ..    .. .. .. ..    .. .. .. ..    .. .. .. ..
Persistent Disable   .. .. .. ..    .. .. .. ..    .. .. .. ..    .. .. .. ..
ISL R_RDY Mode       .. .. .. ..    .. .. .. ..    .. .. .. ..    .. .. .. ..
                    where AN:AutoNegotiate, ..:OFF, ?:INVALID.
                    LM:L0.5
int170:admin> portenable 1

```

Figure 3-36 portCfgPersistentDisable Output

Warning: When a `portcfgpersistentdisable` is done on an enabled E_port, a fabric reconfiguration will occur if the conditions are right. This is the same behavior as the `portcfgdisable` command. Persistently disabling groups of ports is not supported. Each port must be persistently disabled at a time.

3.4.5.4. New FDMI (Fabric OS 3.1/4.1)

FDMI stands for Fabric Device Management Interface. This new function provides the ability to manage the HBA's in-band using Fibre Channel. This capability is a result of a Brocade partnership with Emulex. FDMI is supported on all versions of Fabric OS and is fully backward compatible with all other versions of Fabric OS. There are two new commands `fdmishow` and `fdmicacheshow` that support this function. `FdmiShow` displays local FDMI information that include, local HBAs, HBA port list, HBA attributes and Port attributes. `FdmiCacheShow` displays remote FDMI information such as Remote HBAs and Remote HBA port lists

To use the new FDMI feature both the firmware on the HBA and switch must support it. For the Brocade SilkWorm Switches, the Fabric OS version must be at 3.1 or higher or 4.1 or higher. If the Emulex HBA does not support FDMI, the following entries will exist from `fdmishow` and `fdmicacheshow`.

```

int170:admin> fdmishow
Local HBA database contains no entry.
Local Port database contains no entry.
Remote HBA database contains no entry.
Remote Port database contains no entry.

int170:admin> fdmicacheshow
Switch entry for domain 63
  state:   known
  version: v410
  wwn:    10:00:00:60:69:80:4d:fc

  No devices.
  Total count of devices on the switch is 0

Switch entry for domain 75
  state:   unsupported
  version: v260
  wwn:    10:00:00:60:69:10:10:9d
No devices.
  Total count of devices on the switch is 0

```

SAN Management

This chapter covers the methods and procedures that pertain to the SAN after it has been setup and in a functional and stable state. Specifically, the following topics are discussed: event management, switch maintenance, and switch performance.

- **Event management** addresses SNMP setup, new MIB features, traps, as well as focuses the administrator toward the most pertinent alerts. Brocade Fabric Watch concepts are covered with setup and new features, such as Switch Availability Monitor (SAM). Security concepts will be covered for both SNMP and Fabric Watch. Third-party SAN Management software is discussed as it pertains to successful integration of MIBs and receiving traps.
- **Switch maintenance** addresses the methods of firmware download, zoning and reboots as it pertains to best practices and tool selection. Several guidelines are mentioned for methods of documenting and testing the fabric.
- **Switch performance** addresses Advanced Performance Monitoring. A Fabric Watch example is also provided.

4.1. Event Management

Events are messages, generally in the form of an SNMP trap, that are sent from a switch. The messages can pertain to a broad range of switch areas from port performance to component health to protocol integrity. Being able to configure, receive, understand, and address the messages is called Event Management.

4.1.1. Configuration

The configuration section starts with a look at setting up SNMP as well as Fabric Watch. Several new features of Fabric OS 3.1/4.1 are also discussed.

4.1.1.1. SNMP Commands

One of the most standard methods for monitoring and managing a network device is through Simple Network Management Protocol (SNMP). It is a universally accepted protocol that is portable, lightweight, and is widely deployed. SNMP allows an administrator to monitor the health and performance of countless devices locally or remotely.

All Brocade Fibre Channel switches have an Ethernet port and take advantage of SNMP by providing a myriad of data for the administrator. The data can be retrieved both actively via SNMP queries or can be received passively by SNMP traps.

With so much data available, it can be difficult to know what information is needed to maintain a healthy switch and fabric. Several key pieces will be pointed out later, but first a look at the commands needed to set up SNMP on a Brocade switch.

The settings altered using the following commands are issued all from the command line. Settings can also be changed using Brocade Web Tools, however understanding the underlying command is always helpful.

There are four commands that are helpful in configuring SNMP on a Brocade switch as listed:

- *snmpMibCapSet* on page 4-2
- *agtcfgset* on page 4-2
- *trackChangesSet* on page 4-4
- *agtcfgDefault* on page 4-4

snmpMibCapSet

This command allows users to turn ON/OFF support for certain MIBS and traps. It displays current settings and then prompts the user to change the values for each parameter. By default, support for SW-MIB, FA-MIB, SW-TRAP and FE-MIB are all enabled in Fabric OS 3.1/4.1. During a firmware upgrade all settings are persistent and will be saved through the Fabric OS upgrade.

All MIB support can be turned off except for SW-MIB and FE-MIB. Ensure the SW-TRAP is enabled as this is the support required to send Brocade traps. Figure 4-1 shows what a user would see after issuing the `snmpMibCapSet` command.

```
Switch1:admin> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB SW-MIB SW-TRAP
FA-MIB (yes, y, no, n): [no]
FICON-MIB (yes, y, no, n): [no]
HA-MIB (yes, y, no, n): [no]
SW-TRAP (yes, y, no, n): [yes]
FA-TRAP (yes, y, no, n): [no]
SW-EXTTRAP (yes, y, no, n): [no]
FICON-TRAP (yes, y, no, n): [no]
HA-TRAP (yes, y, no, n): [no]
no change
switch1:admin>
```

Figure 4-1

Notice several new MIBs in the list and their corresponding traps. MIB and trap support for specific features such as High Availability (HA) can easily be turned on or off depending upon the fabric.

SW-EXTTRAP allows Brocade 6400 group as well as soft serial number (SSN) information to be sent in traps.

Guideline: If using third-party SAN management software, it is advisable to turn on support for the FA-MIB as several companies use values within that MIB to discover the presence of switches. Some later versions of Fabric OS have the FA-MIB selected by default.

agtcfgset

This command allows the administrator to change the configuration of the SNMP agent in the switch. The major settings in this command are the trap destinations, Access Control List (ACL), and the value of `swEventTrapLevel`.

Traps can be sent to six different destinations or IP addresses. The community string should be between 2 and 16 characters. Make sure the trap receiver has the same community string.

The ACL has six IP entries, if all entries are empty any IP address may access the switch via SNMP. Access to a specific IP addresses can be given, as well as access to an entire subnet, by using a zero as a wildcard in the appropriate IP address octet (please note that you must conform to the standard IP subnetting rules for this wildcard functionality to work). For example, the entire subnet 192.168.162.0 has access to the SNMP agent on this switch, as shown in Figure 4-2.

```
switch:admin> agtcfset
sysDescr = Fibre Channel Switch.
sysLocation = End User Premise
Current SNMP Agent Configuration
Customizable MIB-II system variables:
sysDescr = Fibre Channel Switch.
sysLocation = End User Premise
sysContact = Field Support.
swEventTrapLevel = 0
authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
Community 1: Secret Code (rw)
No trap recipient configured yet
Community 2: OrigEquipMfr (rw)
No trap recipient configured yet
Community 3: private (rw)
No trap recipient configured yet
Community 4: public (ro)
No trap recipient configured yet
Community 5: common (ro)
No trap recipient configured yet
Community 6: FibreChannel (ro)
No trap recipient configured yet
SNMP access list configuration:
Entry 0: No access host configured yet
Entry 1: 192.168.162.0
Entry 2: No access host configured yet
Entry 3: No access host configured yet
Entry 4: No access host configured yet
Entry 5: No access host configured yet
switch:admin>
```

Figure 4-2 Using the agtcfset command to configure an ACL to limit access to one subnet.

swEventTrapLevel

When running a secure fabric in Fabric OS 3.1/4.1, the community strings establish functionality, and can only be set on the Principal Fabric Configuration Switch (PFCS). Once the communities are configured on the PFCS, they are propagated to all the other secure switches in the fabric. SNMP security is discussed in *SNMP with Secure Fabric OS* on page 4-8

One setting that can cause confusion is the `swEventTrapLevel`. The value `swEventTrapLevel` is set to define the severity level of traps to be sent through specific trap 4 based on Table 4-1. For example if `swEventTrapLevel` is set to 3 then all warning messages and more severe will be sent as a specific trap 4 `swEventTrap`. There are specific traps (which are discussed later and shown in Table 4-1) that come from the Brocade SW-MIB and `swEventTrapLevel` is specific trap number. By adjusting the value of this it is possible to filter specific severities of traps. A setting of zero only means that no traps will be received from this category.

It is critical to understand that `swEventTrapLevel` does NOT control all traps from the switch only trap category 4. Trap 4 works by taking any `errlog` entry that meets the `swEventTrapLevel` severity and wraps it up into a trap and sends it out.

Another area of confusion is when the switch is configured for Fabric Watch and `swEventTrapLevel` is set to 3 or higher. Fabric Watch can place events in the `errlog` so that, in essence, two traps can be received about the same error. This can cause problems if a specific trap kicks off a batch or perl script.

Table 4-1 `swEventTrapLevel` Settings for message severity.

0 – none
1 – critical
2 – error
3 – warning
4 – informational
5 – debug

trackChangesSet

This command enables the track-changes feature that keeps track of successful and unsuccessful logins, as well as any changes to the configuration file. There are two parameters for this feature:

- Mode – toggles track-changes on/off.
- SNMP-trap – toggles the trap on/off.

For both parameters a value of 0 means off and a value of 1 means on. The syntax is:

```
trackChangesSet mode SNMP-trap
```

The following example would turn on the feature and place the messages in the errlog but not send out a trap:

```
trackChangesSet 1 0
```

Note: The trackchanges message placed in the errlog will be sent as specific trap 4 (Table 4-4) if `swEventTrapLevel` is set to severity level 4. The trap which is specific trap 6 will be sent regardless of the `swEventTrapLevel` setting.

agtcfgDefault

This command allows the administrator to reset the configuration of the SNMP agent to factory defaults. The command will prompt the user for confirmation of the action and will only proceed to reset if accepted.

4.1.1.2. SNMP setup in Web Tools

The ability to change SNMP within Brocade Web Tools is functionally the same except for the commands

- `snmpMibCapSet`
- `trackChangesSet`

These commands and their corresponding parameters must be issued and configured from the command line. Figure 4-3 shows a Web Tools GUI of all the values previously discussed.

Guideline: All SNMP configurations can be done from the Web Tools GUI except for `snmpMibCapSet` and `trackChangesSet` which must be done from the CLI.

SwitchName: ddmA20 DomainID: 20 WWN: 10:00:00:60:69:90:04:cd Tue Apr 15 2003, 2:20 PM

License Admin | Port Setting | Routing | Extended Fabric | Configure | Trunk Information

Switch Information | Network Config | Upload/Download | SNMP

SNMP Information

Contact Name: Description:

Location: Trap Level:

Enable Authentication Trap

Community/Trap Recipient

Community Stri...	Recipient	Access Control
Secret C0de	0.0.0.0	Read Write
OrigEquipMfr	0.0.0.0	Read Write
private	0.0.0.0	Read Write
public	0.0.0.0	Read Only
common	0.0.0.0	Read Only
FibreChannel	0.0.0.0	Read Only

Access Control List

Access Host	Access Control List
0.0.0.0	Read Write
0.0.0.0	Read Write
0.0.0.0	Read Write
0.0.0.0	Read Write
0.0.0.0	Read Write
0.0.0.0	Read Write

Apply Close Reset Refresh

[Switch Administration opened]: Tue Apr 15 2003, 2:18 PM
 [Switch Administration closed]: Tue Apr 15 2003, 2:19 PM
 [Switch Administration opened]: Tue Apr 15 2003, 2:20 PM

Enter SNMP Contact Name between 4 and 255 characters

Figure 4-3 SNMP configuration from within Web Tools.

4.1.1.3. MIBs

The definition of the latest Brocade suite of MIBs is shown in Table 4-2. The latest nomenclature of MIB versions is shown in Table 4-3 along with appropriate Fabric OS support. For Fabric OS MIB support and how it interacts with third-party management software, refer to *Third-Party Software* on page 4-31.

Table 4-2 MIB Definitions

MIB Name	Explanation
SW-MIB	Brocade enterprise MIB.
SW-TRAP	Brocade trap definitions.
FA-MIB	Fibre Alliance MIB of which Brocade supports version 3.0.
FE-MIB	Fabric Element MIB was recently accepted by the Internet Engineering Task Force (IETF). In Fabric OS 3.1/4.1, Brocade supports the experimental version and the IETF accepted version located under the MIB-2 branch.

Table 4-3 Nomenclature of MIB Versions

Fabric OS v4.1.0	v3.1.0	v2.6.1
BRCD_v5_0.mib	BRCD_v5_0.mib	BRCD_v5_0.mib
SW_v5_0.mib	SW_v5_0.mib	SW_v5_0.mib
HA_v5_0.mib	N/A	N/A
FA_v3_0.mib	FA_v3_0.mib	FA_v2_0.mib
FE_RFC2837.mib	FE_RFC2837.mib	FE_EXP.mib
ENTITY_RFC2737.mib	N/A	N/A

Note: The latest version of the FE-MIB references the FRAMEWORK.MIB, some MIB browsers may require loading this MIB before the FE-MIB is loaded.

Without a Fabric Watch license, or if it is turned off, the other main source of traps comes from trap four, or the `swEventTrap`, which was previously mentioned. Without Fabric Watch, a recommended setting for `swEventTrapLevel` is either 3 or 4, as shown in MIBs for Fabric OS 3.1/4.1

4.1.1.3.1 Entity MIB

The Entity MIB module is a way of representing multiple logical entities supported by a single SNMP agent. Essentially the Entity MIB is way of breaking down each logical component, such as a FRU, within a switch and having an associated OID and value assigned to it.

Note: The Entity MIB module implemented only in Fabric OS v4.x and is not supported in Fabric OS v3.x or Fabric OS v2.6.1.

The following is the list of possible entities that can classify a logical unit:

1. Other
2. Unknown
3. Chassis
4. Backplane
5. Container - chassis slot or daughter-card holder
6. Power Supply
7. Fan
8. Sensor
9. Module - plug-in card or daughter-card
10. Port
11. Stack - stack of multiple chassis entities

The Entity-MIB has one trap, `entConfigChange`, and is generated when the value of the `entLastChangeTime` changes. It can be utilized by a network management station to trigger logical/physical entity table maintenance polls.

Guideline: An NMS should periodically check the value of `entLastChangeTime` to detect any missed `entConfigChange` notification-events possibly due to throttling or transmission loss.

4.1.1.3.2 HA MIB

The High Availability MIB tracks not only HA status but the status and history of all Field Replaceable Units (FRUs) within a switch. Since only the SilkWorm 12000 switch has a dual CP HA design, only this switch supports the HA functionality.

Three new traps are defined in the HA-MIB which include a FRU status, FRU change, and CP status traps. The corresponding trap OID is .1.3.6.1.4.1.1588.2.1.2.2.0, which is different than the SW trap OID .1.3.6.1.4.1.1588.2.1.1.1.0 (for Traps 1 through 6).

The HA-MIB has three traps:

- fruStatusChanged
- cpStatusChanged
- fruHistoryTrap

There are similarities between the Entity MIB and the Brocade proprietary HA-MIB. The relationship is between the FRUTable and the entPhysicalTable, however the FRUTable handles only a subset of entries from the entPhysicalTable

4.1.1.4. Traps (SW MIB)

Each fabric and management process is different, but there are a couple of traps that can be used as flags to alert an administrator to facilitate troubleshooting.

From the Brocade enterprise MIB there are six different traps defined and listed in Table 4-4.

Table 4-4

Name	Specific	When	Configuration Command
swFault	1	During boot, if diagnostics fail	Always on
swSensorScn	2	Obsolete	N/A
swPortScn	3	Port changes state	Always on
swEventTrap	4	Switch Event	Command: agtcfgset Variable: swEventTrapLevel
swFabricWatch	5	Threshold breached	Command: fwconfigure
swTrackChanges	6	Login/Logout	Command: swtrackChangesSet

One of the most succinct but informative traps from Table 4-4 is swPortScn which simply states that a port has either gone online or offline. As there are not many changes occurring in the fabric, most ports that are connected to storage or hosts are always online so a trap reporting a port as offline would be an important flag to pursue.

One change to note for the Brocade enterprise MIB is that the SW and SW-TRAP MIBs are now combined into one MIB that is backward compatible with all previous versions.

Guideline: Without Fabric Watch, many severity level 3 and 4 messages might not be seen so it is advisable to turn up the swEventTrapLevel to 3 or 4 which will allow more of those severity messages to be delivered from the swEventTrap.

Because of the nature of trap 4, where the entire errlog entry string becomes one variable of the trap, it can be difficult to set up network management software (NMS) to parse or correlate two related traps. If further parsing of the trap is needed, pull out the desired information.

The other solution is to configure Fabric Watch which has more distinct and precise variables sent in each trap. Parsing and correlating traps tends to be much easier with a swFabricWatch trap and is discussed later.

Guideline: If parsing messages is a part of your NMS strategy or if this is a desirable activity, use the swFabricWatch trap.

4.1.1.5. SNMP with Secure Fabric OS

When running Secure Fabric OS, there are several subtle changes to SNMP that need to be considered. The concept of Fabric Configuration Servers is introduced and discussed in more detail in the *Secure Fabric OS User's Guide*.

Essentially, there is a list of switches, based on their WWN, which are in charge of controlling all fabric management functions or parameters. The first switch from this list is the Primary FCS and controls all propagation of settings.

4.1.1.5.1 Community Strings

The community strings are one of the management functions that are controlled by the Primary FCS (PFCS). The configuration of these strings is performed only on the PFCS. The PFCS then propagates the community strings to all other participating Secure Fabric OS switches.

4.1.1.5.2 SNMP MAC Policy

With a switch in a secure Fabric, the ACLs that could be configured via the agtcfset command are now obsolete, replaced with SNMP Management Access Control (MAC) Policies. They are as follows:

- RSNMP_POLICY (read access)
- WSNMP_POLICY (write access)

These policies contain a list of host IP addresses from which connections or messages are accepted by any switch in the secure fabric. The SNMP MAC Policies can be used to limit switch access to specific, trusted workstations in the customer's environment, and is done so with differing read/write access levels. The SNMP host must send its request to the primary FCS switch to perform write operations.

Each policy has three different settings, Non-existent, Empty and some defined parameter listing any number of IP addresses. There are nine different combinations listed in Table 4-5. Take note of two cases which cause ambiguity and are therefore illegal.

Table 4-5

	RSNMP	WSNMP	[RO] Result	[RW] Result
1.	Non-existent	Non-existent	All hosts allowed to read-only [RO]	All hosts allowed to read-write [RW]
2.	Non-existent	Empty	All hosts allowed to read-only [RO]	No host allowed to read-write [RW]
3.	Non-existent	Having host B in policy	All hosts allowed to read-only [RO]	Only host B allowed to read-write [RW]
4.	Empty	Non-existent	This is ambiguous so it is illegal. Warning is displayed and cannot be saved with secPolicysave.	
5.	Empty	Empty	No host allowed to read-only [RO]	No host allowed to read-write [RW]
6.	Empty	Having host B in policy	No host allowed to read-only [RO]	Host B allowed to read-write [RW]
7.	Having host A in policy	Non-existent	This is ambiguous so it is illegal. Warning is displayed and cannot be saved with secPolicysave.	
8.	Have host A in policy	Empty	Host A allowed to read-only [RO]	No host allowed to read-write [RW]
9.	Having host A in policy	Having host B in policy	Host A allowed to read-only [RO]	Host B allowed to read-write [RW]

Note: If a switch running Fabric OS (FOS) 2.6.0 with illegal policy combinations tries to join into a fabric consisting of switches running Fabric OS 2.6.1/ 3.1/4.1, the Fabric OS 2.6.0 switch will not be allowed to merge and become active on the fabric. The reason for this is the security behavior will not be consistent across the fabric. For consistent behavior, it is recommended to upgrade to Fabric OS v2.6.1 or turn off security on that particular switch.

When a switch goes from secure mode to non-secure mode, the community names stay the same, but can be configured on a per switch basis. The ACLs that had configuration parameters are now reinstated and the security policies RSNMP and WSNMP no longer apply.

4.1.2. SNMP Features

4.1.2.1. Security and Port Naming

New functionality introduced in Fabric OS v3.1 and v4.1 is Security and the Port Name feature. The Port Name feature allows you to give up to a 32 character string name to the port. When using the command `agtCfgDefault` in Secure mode, primary FCS allows the setting and configuration for all agent attributes. The non-Primary FCS allows the setting and configuration for all agent attributes except for community strings and ACLs.

In secure mode, the community strings can ONLY be changed on the *Primary switch* regardless of whether the switch is online/offline. The *Primary switch* then distributes the community strings to other switches in the secure fabric. In secure mode, SNMP ACLs are disabled and the access control lists are maintained by security module policies WSNMP_POLICY and RSNMP_POLICY, as described earlier.

4.1.2.2. MIB Support

- During a failover or firmware upgrade there is no disruption to the SNMP agent which allows for continuous monitoring of all FRUs and overall health of the switch.
- Setting community strings from SNMP and Web Tools – same type of behavior change as described above for telnet commands when Security is enabled
- New Single SW MIB will be introduced that will be backward compatible to all Fabric OS 2.6, v3.x /v4.x.
- HA, FRU, FICON MIBs are not available in Fabric OS v3.1.

4.2. Fabric Watch

Another powerful management tool is a licensed Brocade product called Fabric Watch. This tool has the ability to monitor many different switch parameters and send alerts on user defined threshold breaches.

The alert delivery mechanisms available cover the aforementioned SNMP trap along with an errlog message, port log lock, E-mail, and an alert through the Brocade Application Programming Interface (API).

Like SNMP, the information can be overwhelming at first so this section will not only cover some basic configuration but address how to read and understand what the Fabric Watch messages mean. Several guidelines are provided and some of the Fabric OS v3.1/4.1 features are discussed.

4.2.1. Commands

There are six commands used to configure Fabric Watch, some of the commands will only need to be issued once and then can be ignored.

4.2.1.1. *fwclassinit*

If a switch has a Fabric Watch license but it has not been initialized, it can be quickly initialized by typing the command *fwclassinit*. This will initialize all Fabric Watch classes.

4.2.1.2. *fwalarmsfilterset*

This command globally turns on or off all **non-environmental** Fabric Watch alarms. Environmental alarms are all controlled by Fabric Watch but do not need a Fabric Watch license, they include sensors fan, power supply and temperature. All other facets of Fabric Watch, including Field Replaceable Unit (FRU) alarms, are turned on or off with this command.

When trying to configure Fabric Watch turn off the alarms by typing:

```
Switch1:admin> fwalarmsfilterset 0
```

Guideline: To prevent false alarms, turn off Fabric Watch messages when configuring or doing maintenance to the SAN or Fabric Watch.

With Fabric Watch alarms now disabled, an administrator, from the command line, can configure Fabric Watch or other switch features without an over abundance of messages filling up the screen. If working from Web Tools to configure Fabric Watch, this command is not necessary as the messages are not seen from the Web Tools configuration GUI.

To turn all of the alarms back on, type:

```
Switch1:admin> fwalarmsfilterset 1
```

4.2.1.3. fwconfigure

To configure and change Fabric Watch settings from the command line based menu system, type `fwconfigure`. From there it is fairly straightforward navigation through the different classes listed in Table 4-6.

Table 4-6

```
Switch1:root> fwconfigure

1 : Environment class
2 : SFP class
3 : Port class
4 : Fabric class
5 : E-Port class
6 : F/FL Port (Optical) class
7 : Alpa Performance Monitor class
8 : EE Performance Monitor class
9 : Filter Performance Monitor class
10 : Security class
11 : Switch Availability Monitor class
12 : Quit
Select a class => : (1..12) [12]
```

- The environmental class is partially enabled without a Fabric Watch license. Alarms will be sent out based on default values yet the ability to configure the thresholds are only available for the licensed version.
- The SFP class deals with the Short Form-factor Pluggable characteristics shown in Table 4-7. The RXP and TXP refer to the receive and transmit power, respectively.

Table 4-7

SFP Monitored Class
1. Temperature
2. RXP
3. TXP
4. Current
5. Voltage

- The Port, E-port and F/FL port (Fabric and Fabric Loop) classes monitor the same categories, shown in Table 4-8, but for a specific type of port. The Port class covers all types of ports including the E-port and F/FL port.

Table 4-8

Port, E-port and F/FL Port Class
1. Link loss
2. Sync loss
3. Signal loss
4. Protocol error
5. Invalid words
6. Invalid CRCs
7. RXPerformance
8. TXPerformance
9. State Changes

- The AL_PA performance class monitors invalid CRCs.
- The EE performance class monitors CRCs, RXPerformance, and TXPerformance.
- The security class monitors 21 different areas of security policy when secure mode is enabled and will be discussed in more detail in section 2.1.2.6.
- The Switch Availability Monitor (SAM) class monitors the efficiency of ports. The SAM class provides statistics on switch downtime and uptime and will be discussed in section 2.1.2.5.

For more information about Fabric Watch classes, reference the *Brocade Fabric Watch User's Guide*.

Note the different sub categories presented in Table 4-6. They will come in handy when trying to figure out what any particular Fabric Watch trap means.

4.2.1.4. *fwconfigreload*

This command reloads the Fabric Watch configuration and should only be used after downloading (`configdownload`) a new Fabric Watch configuration file from a host. For Fabric OS v4.1 this is no longer needed as the command is automatically committed after the configuration is downloaded.

4.2.1.5. *fwfrucfg*

This command allows the administrator to configure the states and actions of four FRU classes shown in Table 4-9. There are seven different states that each FRU can be in, shown in Table 4-10. Based on the configuration, Fabric Watch generates an action when the FRU enters a particular state. Unlike the other classes of Fabric Watch, the FRU class only has two alert mechanisms, errlog and E-mail.

Table 4-9 FRU Classes

1	: Slot
2	: Power Supply
3	: Fan
4	: WWN.

FRU Classes**Table 4-10** FRU States

1.	Absent
2.	Inserted
3.	Ready
4.	Up
5.	On
6.	Off
7.	Faulty

FRU States

Some of the states do not pertain to certain FRUs, for example, the Ready, Up, and On states are pertinent to the slot class as when a blade is plugged in, it goes through several Power On Self Tests (POSTs). As each test is passed, the slot is placed into another state.

The fans and power supplies are much less sophisticated, so many of the states are reached less than a second apart from the previous state making the two indiscernible. For these FRUs, using a slightly abbreviated list of alerts will make the messages more meaningful. For example, if a power switch is turned off/on it will turn off/on two power supplies (each power switch controls two of the four power supplies) and the following messages will be received:

Faulty, Inserted, Ready, Up, and On.

If one of the four power supplies is pulled out and put back in, the following messages will appear:

Absent, Inserted, Ready, Up, On.

Combining the two lists gives:

Absent, Faulty, Inserted, **Ready, Up, and On.**

As mentioned before, there are several states that are not pertinent to power supplies and fans which occur almost simultaneously. The last three, Ready, Up, and On are examples of this. So to reduce the number of messages received eliminate Ready and Up from the alert configuration.

Guideline: Eliminating Ready and Up out of the FwFruCfg configuration can help eliminate extraneous messages.

The final set of alerts to use for all power supply issues is:

Absent, Faulty, Inserted, On.

Without a Fabric Watch license, the environmental monitoring agent will send an abbreviated list of messages to the errlog, including, Removed, Faulty, and Inserted.

4.2.1.6. *switchstatuspolicyset*

The `switchstatuspolicyset` command sets the current policy parameters for calculating the overall status of the switch (enclosure). The status of the switch only affects the color of the Web Tools GUI, Fabric Manager GUI and a message displayed in the errlog and telnet console.

Most third-party management software do not use this policy to determine the state of a switch. This can cause confusion if the policy is set differently than the way vendor software states the health of a switch. It could be helpful to synchronize health status policies.

Guideline: To avoid confusion, make sure the policy is set the same as the third-party management software states the health of a switch.

Typing the command will print the current parameters in a 3-column format, as shown in Table 4-11. The first column consists of the contributor; the second column specifies the minimum number that contributes to the Down/Failed status which is red; the third specifies the minimum number that contributes to the Marginal/Warning status which is yellow or orange.

Table 4-11 Example switch status policy parameters

	<i>Down</i>	<i>Marginal</i>
<i>FaultyPorts</i>	2	1
<i>MissingSFPs</i>	64	0
<i>PowerSupplies</i>	3	1 (change to 2 in example below)
<i>Temperatures</i>	2	1
<i>Fans</i>	2	1
<i>PortStatus</i>	64	0
<i>ISLStatus</i>	0	0

For example, it would be necessary to alter the configuration when the maximum number of power supplies in a SilkWorm 12000 chassis are not used. If there are only three power supplies instead of the maximum of four that a chassis can house then, per the default configuration in Table 4-11, the switch would be in a marginal state.

By altering the Marginal level to two, the switch will return to a normal health state. Having two of the three power supplies fail would now give two overall bad power supplies and cause the switch to be in a marginal state.

Also note that by setting the PortStatus variables both to zero will cause port online/offline status messages to not appear in the errlog. If either the *Down* or *Marginal* variable is set to a non-zero value then a message will be placed in the errlog.

4.2.2. Reading Fabric Watch Messages

From a logistical standpoint, knowing how to read Fabric Watch messages should come after having setup and configured Fabric Watch. However, since all switches have some Fabric Watch alerts and messages turned on as default, it will prove effective to learn some error message syntax now.

4.2.2.1. Message Types

There are four types of Fabric Watch messages.

Table 4-12

Type	Meaning	Configuration Command	Possible Alerts
STATUS	Switch Health Status	SwitchStatusPolicySet	Errlog SNMP Trap 4
FRU	Field Replaceable Unit	FwFruCfg	Errlog E-mail
THRESHOLD	Threshold Breach	FwConfigure	Errlog SNMP Trap 4
EM	Environmental	Free – no configuration possible	Errlog SNMP Trap 4

There are many messages that can be sent from a Brocade switch, the best way to understand what they mean is to read a few of them.

4.2.2.2. Message Key Points

Below are four messages that were delivered to the errlog when a power supply was pulled out.

FRU Message

```
0x29a (fabos): Mar 05 15:46:48
Switch: 0, Info FW-FRU_ABSENT, 4, Power Supply #002 state has changed to FRU_ABSENT
```

Status Message

```
0x368 (fabos): Mar 05 15:46:48
Switch: 0, Warning FW-STATUS_SWITCH, 3, Switch status changed from HEALTHY/OK to Marginal/Warning (---
1 missing power supply;)
```

Threshold Breach

```
0x368 (fabos): Mar 05 15:46:48
Switch: 0, Warning FW-BELOW1, 3, envPS002 (Env Power Supply 2) is below low boundary. current value: 0
(1 OK/0 FAULTY). (faulty)
```

Environmental

```
0x25a (fabos): Mar 05 15:46:48
Switch: 0, Info EM-FRU_REM, 4, Power Supply #2 removal detected.
```

The key component to knowing if the message is from the Fabric Watch agent is to recognize the “FW” at the beginning of each message. All messages except the last one are delivered by Fabric Watch. If a Fabric Watch license wasn’t present on the switch then only the last message would be received.

The threshold messages can be a little more difficult to understand so referring back to the classes can be helpful in understanding them. For example:

```
FW-ABOVE, 4, portTXPerf (Port TX Performance) value has changed. current value: 105831 KB/s. (info)
```

From the message above, the first box has portTXPerf and could be confusing but there is more information from the message. The second box expands on the term by saying it’s a performance breach.

For more detailed explanations of the messages refer to the *Brocade Fabric Watch User’s Guide*.

If a message being received is not wanted then there is enough information in the message to be able to figure how to turn it off. Configuring, tuning and honing Fabric Watch is discussed next.

4.2.3. Configuring

4.2.3.1. Tuning Fabric Watch

Fabric Watch is a threshold monitoring and alerting engine, and by default, Fabric OS versions before v3.1/v4.1 deliver an abundance of information to the administrator. This can be overwhelming and confusing so the following section will cover several ways of configuring and honing Fabric Watch to a user's desired level.

As mentioned earlier, one way to turn off Fabric Watch messages completely, while, for example, work is being done at the command line, is to type the command

```
Switch1:admin>fwalarmsfilter 0
```

To turn the messages back on, replace the zero with a one.

4.2.3.1.1 Tuning Fabric Watch With Web Tools

Configuring Fabric Watch from the Web Tools GUI is straightforward. There are two approaches

1. Receiving a message and wanting to turn it off.
2. Wanting to turn on a particular message

To approach the first concept of receiving a message let's take the message from the previous section and pursue turning it off.

```
FW-ABOVE, 4, portTXPerf8 (Port TX Performance008) value has changed. current value: 105831 KB/s. (info)
```

After pulling up the switch in a browser (Figure 4-4) click on the magnifying glass icon labeled **Watch** and login with the required credentials.




Figure 4-4 Web Tools Main Page

4 SAN Management

Fabric Watch is broken down into the categories of Class and Area. Sometimes it is obvious as to which Class. A quick browse through the GUI should reveal the error message source. Looking at the two boxes in the original message one can ascertain that this is a Port Class error (E-Port and F/FL Optical are the other Port classes). To get to this error Class click on **Port** under the PORTS as shown in Figure 4-5.

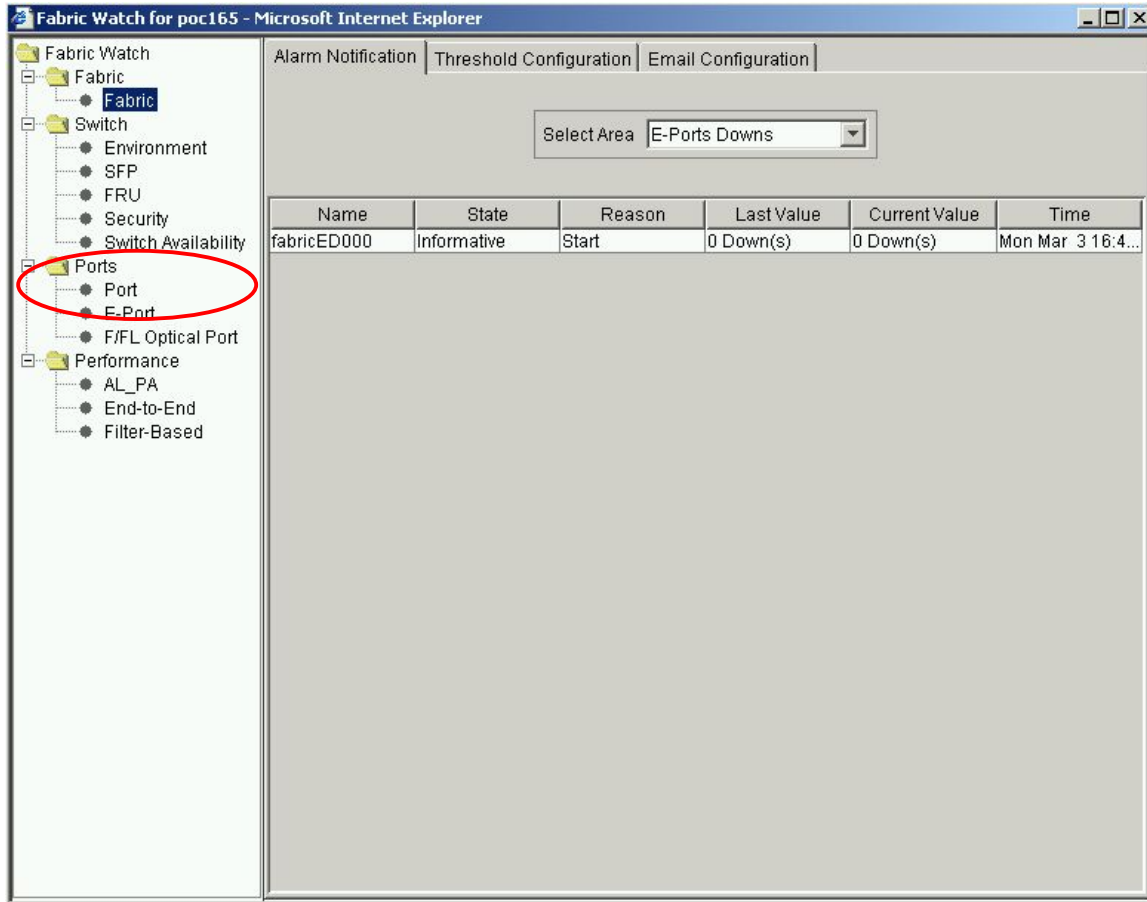


Figure 4-5 Fabric Watch Main Page

Another menu will open, as shown in Figure 4-6. This is the Fabric Watch GUI which, by default, opens with the class FABRIC. Under the class of FABRIC are sub classes called areas.

With the Port class now in the window all ports can be seen with their current state, the reason they are in that state and more

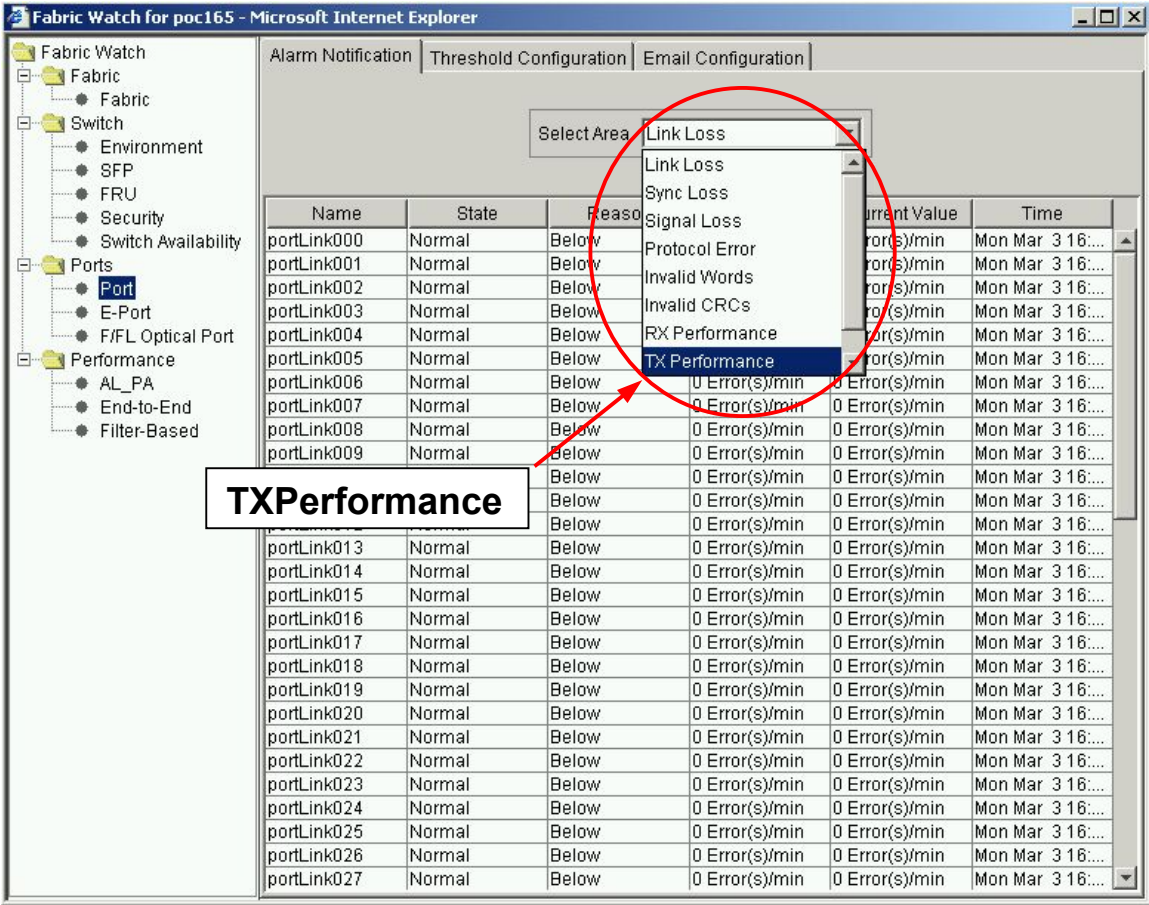


Figure 4-6 Ports Class With Port Area Selected

FW-ABOVE, 4, portTXPerf8 (Port TX Performance008) value has changed. current value: 105831 KB/s. (info)

Glancing once more at the original message (above) more information can be pulled from this as to the source. TXPerf or TXPerformance gives just the information needed to locate the Area type of this message. TXPerformance (Transmit Performance) can be selected by clicking on the menu, as shown in Figure 4-6.

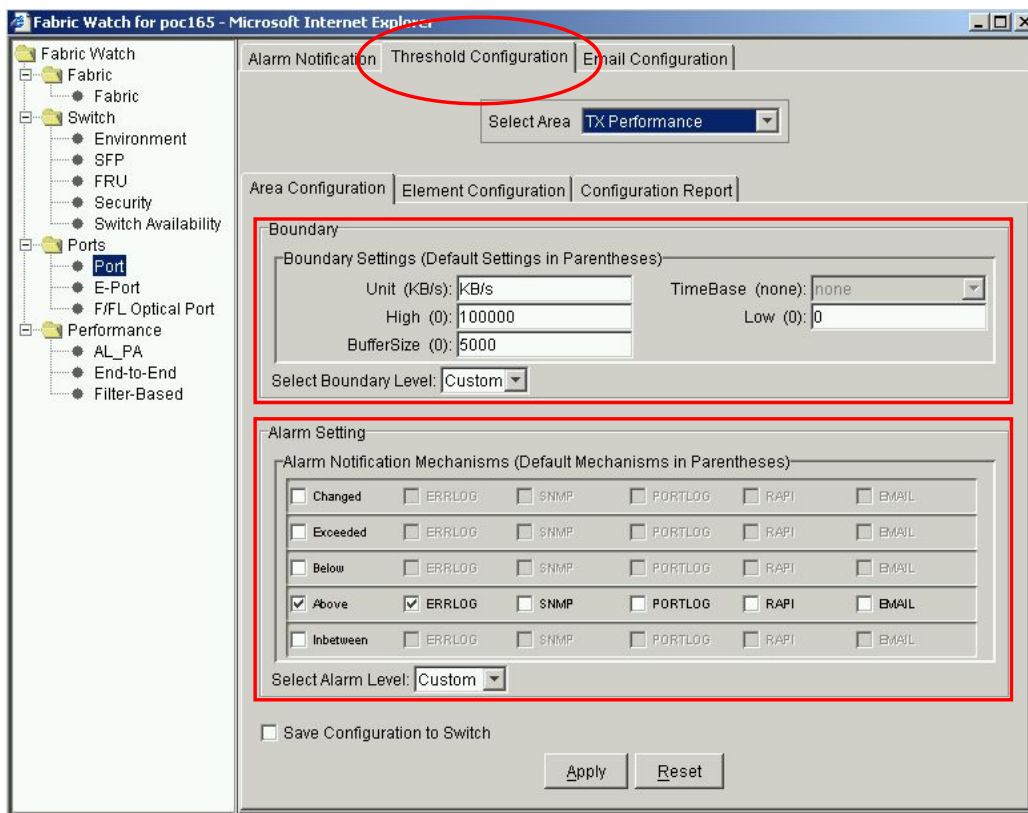


Figure 4-7 TXPerformance Area Threshold Configuration

With the Port area selected and showing in the GUI click on the **Threshold Configuration** tab as shown in Figure 4-7. This will bring up the configuration page also shown in Figure 4-7. There are a couple of things to take note in Figure 4-7, the boundary and the alarm setting boxes. To understand why these areas are important take one more look at the original message.

FW-ABOVE 4, portTXPerf8 (Port TX Performance008) value has changed. current value: 105831 KB/s. (info)

The first boxed word says ABOVE and pertains to the Alarm Setting. This area is the lower area of the Threshold Configuration GUI in Figure 4-8.

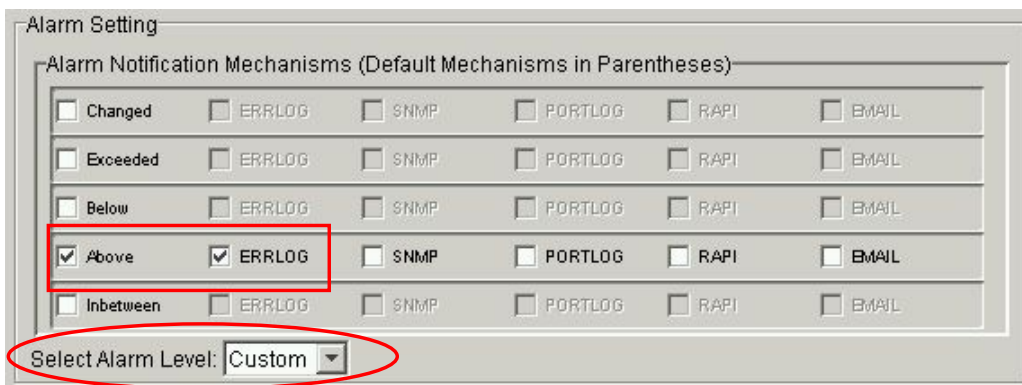


Figure 4-8 Alarm Settings

The alarm setting region contains all different types of threshold breaches along with what alarm to send when a breach occurs. As the original message eludes, an ABOVE threshold breach occurred and according to Figure 4-8 a message was placed only in the ERRLOG. To turn this alert off (or to add other messages) check the ERRLOG box and save the configuration.

To add another alert when the transmit performance goes below a certain value, check the BELOW box and then the appropriate alerts.

The other possible alerts are

1. SNMP – an SNMP trap sent to user defined destinations.
2. PORTLOG lock which locks the port log so port data can be saved for troubleshooting.
3. RAPI – a message sent through the Brocade Application Programming Interface.
4. EMAIL – E-mail sent to an appropriate address.

The Boundary settings, as seen in Figure 4-9, determine the threshold values that cause a breach/trap. The High boundary is set to 10000 KB/sec and the message received had a value of 105831 KB/sec. If more traffic will be flowing through this port then perhaps a higher threshold value is needed.

Make sure that if any changes are made to any of these settings to select **CUSTOM** under the **Boundary Level** and **Alarm Level** menus. This is circled in both Figure 4-8 and Figure 4-9. If the Boundary and Alarm Level settings are set to **default** with custom changes, then unexpected results could occur.

Figure 4-9 Boundary Settings

4.2.3.1.2 fwconfigure

This command was discussed earlier so by applying the techniques used with the Web Tools Fabric Watch GUI the command line interface method of tuning messages will be pursued with equal fervor.

Start by typing the command and bringing up the first menu list and selecting number **3 Port class** (shown below) since that has already been determined from the message.

```
Switch1:admin> fwconfigure

1: Environment class
2: SFP class
3: Port class
4: Fabric class
5: E-Port class
6: F/FL Port (Optical) class
7: Alpa Performance Monitor class
8: EE Performance Monitor class
9: Filter Performance Monitor class
10: Security class
11: Switch Availability Monitor class
12: Quit

Select a class =>: (1..12) [12] 3
```

```
1: Link loss
2: Sync loss
3: Signal loss
4: Protocol error
5: Invalid words
6: Invalid CRCS
7: RXPerformance
8: TXPerformance
9: State Changes
10: return to previous page
```

```
Select an area => : (1..10) [10] 8
```

Figure 4-10 fwconfigure Class and Area Selection

Next select the **8. TXPerformance** area. A list of current port values will scroll by. This is the same information that was seen in Figure 4-6. At the bottom of the list is another menu selection, choose number **4** to go into the advanced configuration.

Figure 4-11 shows an abbreviated view of the advanced configuration menu. (Several ports were left out due to the large number).

The two gray portions represent the same feature information that was configured in the previous section using Web Tools. Looking at the familiar message, the ABOVE alarm settings and the threshold setting of 100000 KB/sec can be seen in Figure 4-11.

```
FW-ABOVE, 4, portTXPerf8 (Port TX Performance008) value has changed. current value: 105831 KB/s. (info)
```



```

Index ThresholdName      BehaviorType      BehaviorInt
  0 portTXPerf000        Triggered         1
  1 portTXPerf001        Triggered         1
  ::::::::::::::::::::: ::::::::::::::::::::: :
  ::::::::::::::::::::: ::::::::::::::::::::: :
  63 portTXPerf063      Triggered         1

Threshold boundary level is set at : Custom

      Unit          Default      Custom
      KB/s         KB/s
Time base
Low          0            0
High        0            100000
BufSize     0            5000

Threshold alarm level is set at : Custom

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16

Valid alarm matrix is 31

      Default      Custom
Changed      0            0
Exceeded     0            0
Below 0      0            0
Above 0      0            1
InBetween    0            0

1 : change behavior type           11 : change threshold alarm level
2 : change behavior interval       12 : change changed alarm
3 : change threshold boundary level 13 : change exceeded alarm
4 : change custom unit             14 : change below alarm
5 : change custom time base        15 : change above alarm
6 : change custom low              16 : change inBetween alarm
7 : change custom high             17 : apply threshold alarm changes
8 : change custom buffer           18 : cancel threshold alarm changes
9 : apply threshold boundary changes 19 : return to previous page
10 : cancel threshold boundary changes

Select choice => : (1..19) [19]

```

Figure 4-11 Advanced Configuration Menu of Fabric Watch using CLI

To turn off the alarm simply select number **15 change above alarm** and set the value to zero. To add additional messages add the Errlog value of 1 to the value of another message. For example if Errlog (value of 1) and an E-mail alert (value of 16) is desired then simply add their values together. For these two messages the proper value would be 17.

4.2.3.1.3 Configuration File Download

One other method for tuning Fabric Watch is to combine one of the previous two methods, either Web Tools or CLI, with the `configupload` and `configdownload` functions.

The idea is to configure one machine, upload the settings, and then download the settings to any appropriate Brocade switch. Having covered configuration techniques for Fabric Watch, pick a method and tune all alarm and threshold settings on one switch.

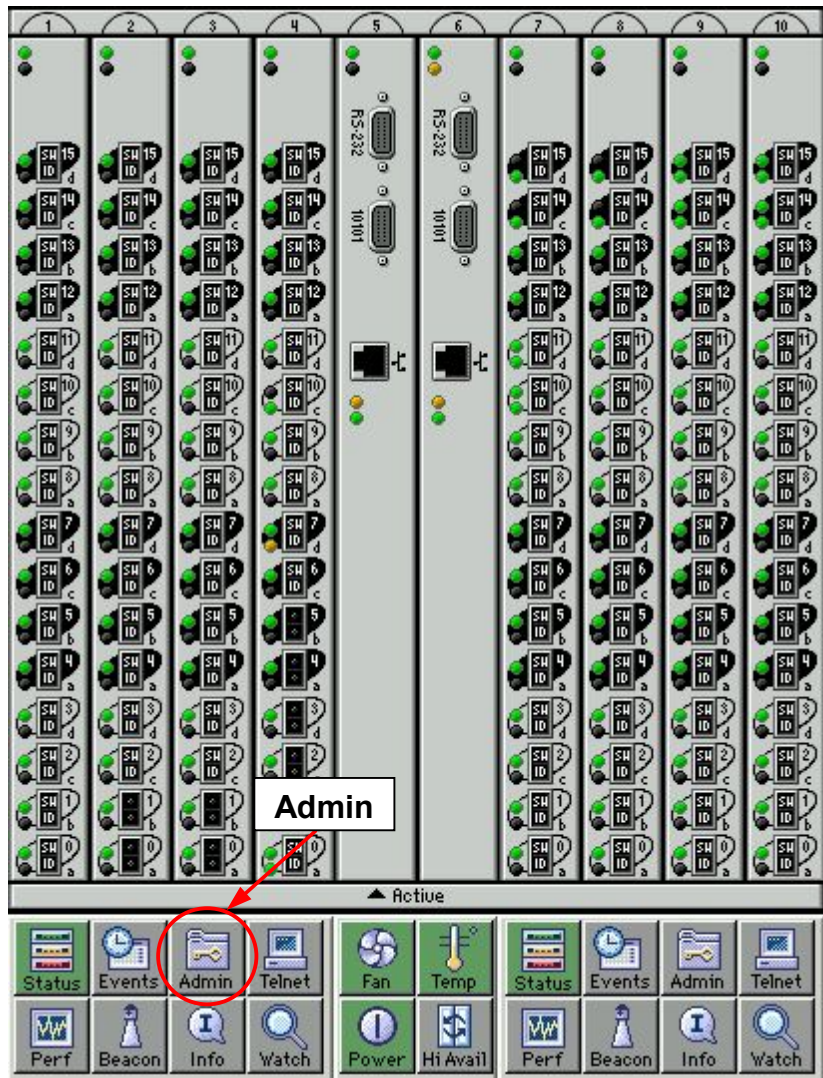


Figure 4-12 Web Tools Administration

To upload the settings to an FTP server, Web Tools GUI will be used to illustrate this. Bring up the main Web Tools page for the switch that contains the configuration and click on the **Admin** button shown in Figure 4-12.

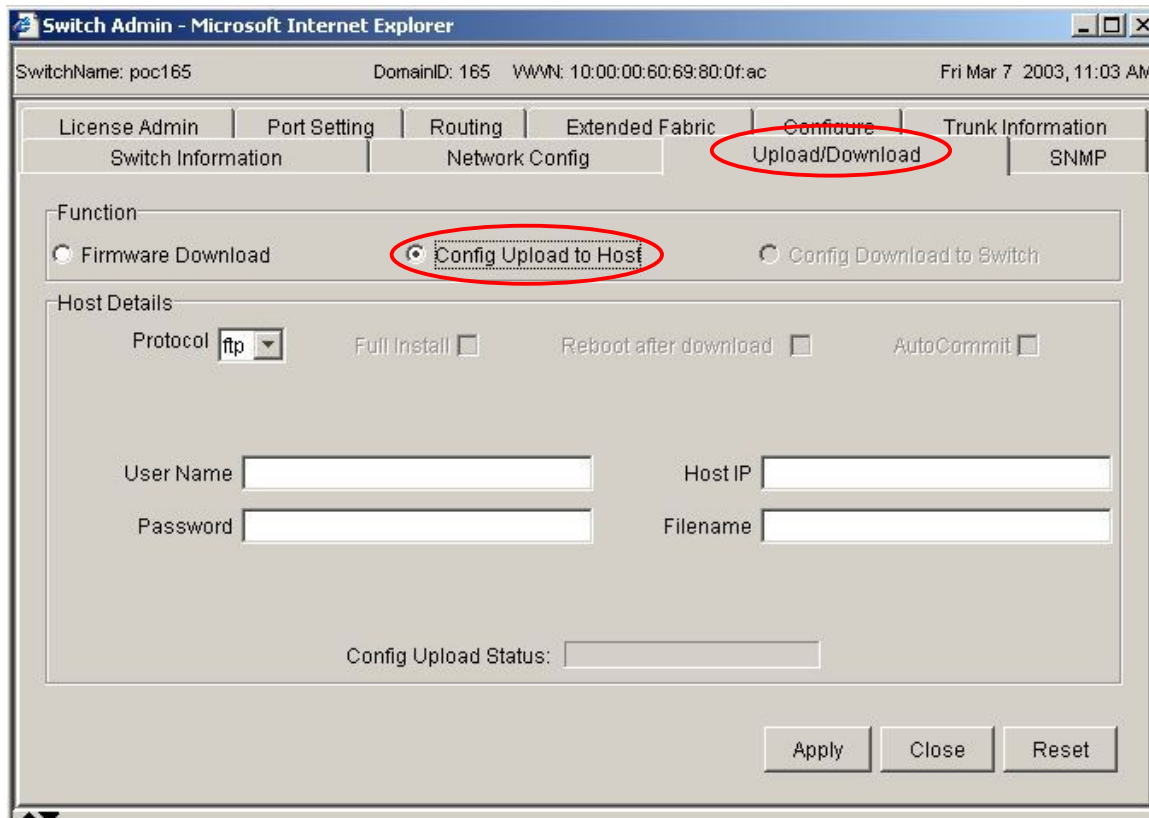


Figure 4-13 Config Upload of Configuration

Login with the proper credentials and select the **Upload/Download** tab. Then click the **Config Upload to Host** button and fill in the appropriate login credentials for the host along with a filename. Click **Apply** when ready to upload the file.

Since the file just uploaded contains not only Fabric Watch Settings but settings for all aspects of the switch it is necessary to eliminate everything but the Fabric Watch parameters. Use a text editor to edit the file, keeping everything with the word **thresh** on the front of it. A sample line to keep within the configuration file looks like this:

```
thresh.cust.port.TXPerf.high:100000
```

This parameter should look familiar as it is the original message that we have been working with as our example. This just happens to be the parameter that causes a threshold breach if transmit traffic goes above 100000 KB/sec. (The **cust** stands for custom parameter).

Note: Make sure to put a carriage return at the end of the config file or else it will not be read correctly and will fail.

To download the now edited Fabric-Watch-only file to another switch, login into that switch's main Web Tools page as instructed before and select the **Switch Information** tab (Figure 4-14).

Guideline: Modify the parameters on one switch, then propagate the profile to other switches in the fabric using `configdownload`.

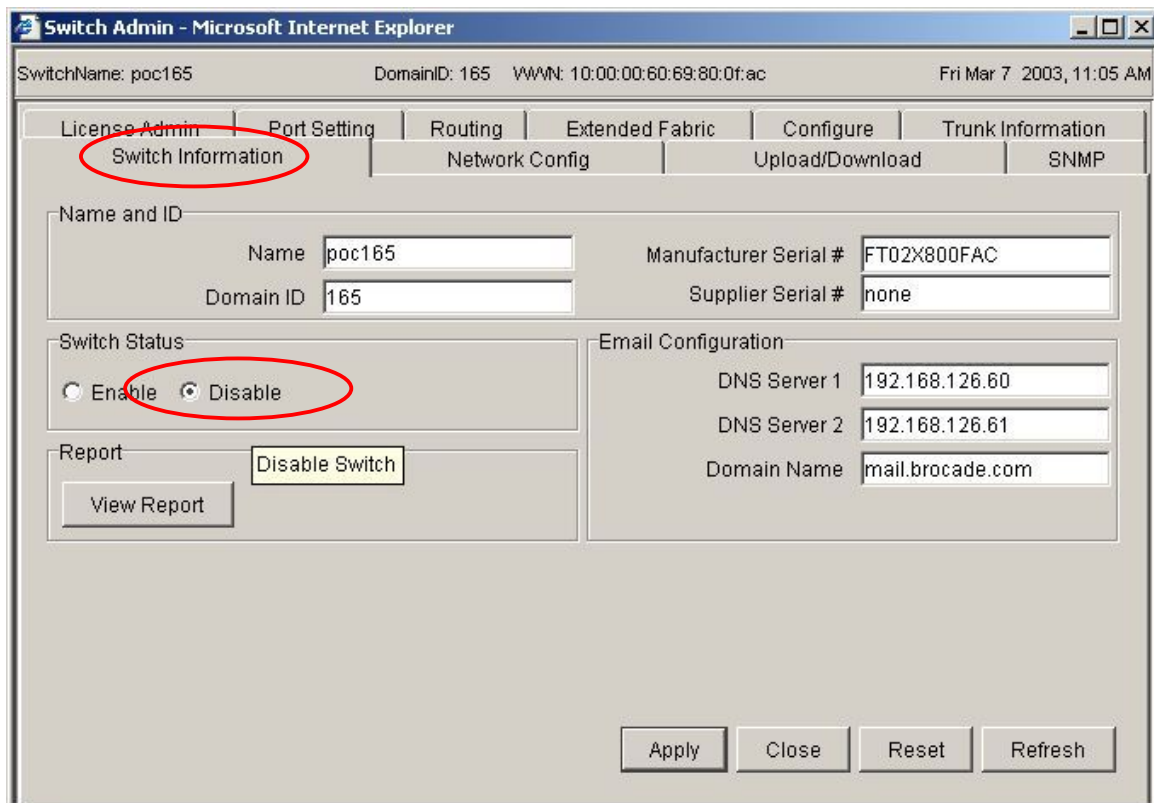


Figure 4-14 Disabling Switch for Configdownload

When doing a configdownload it is necessary to disable the switch so click on the **disable** button and click the **apply** button.

Now go back to the **Upload/Download** tab and select the **Config Download to Switch** button. Fill in the appropriate FTP server information and the file name uploaded and edited from the previous steps and click **apply**.

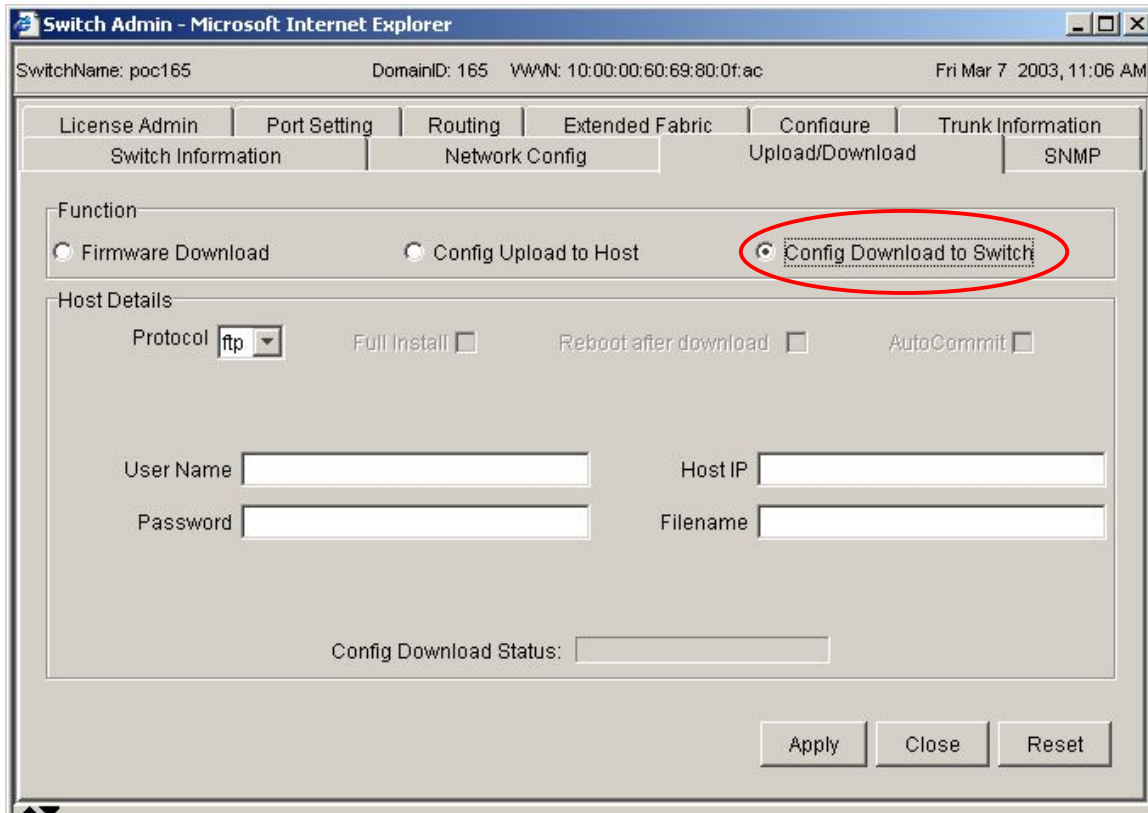


Figure 4-15 Configdownload of Configuration

Now go back to the Switch Information tab and re-enable the switch. The new parameters will be read and the values that were set on the one switch can be propagated to other switches in the same way.

Note: On versions of Fabric OS before 3.1/4.1 it was necessary, after doing a configdownload of Fabric Watch settings, to type the command `fwconfigreload`. This command would reload all the flash settings back into Fabric Watch. This is not needed on Fabric OS 3.1/4.1 because it is done automatically.

4.2.3.1.4 Fabric Watch Best Practices

With several techniques covered on to how read messages and tune Fabric Watch from various interfaces it's time to look at what is *really* important to tune.

The number of messages delivered by default in Fabric OS 3.1/4.1 is scaled down from previous versions of Brocade Fabric OS. Each integration approach with Fabric Watch is different and it tends to be easier to turn on what is wanted rather than try to turn off what is not needed.

With a blank Fabric Watch configuration in mind, several thresholds that can be very helpful in maintaining a Brocade fabric will be investigated.

Fabric Class contains several informative messages, some of the following are useful for overall fabric connectivity like E-Port Downs, Fabric Reconfigure, and Segmentation, others can be more useful for security purposes like Zoning Changes and Fabric Logins. Please refer to Table 4-13 for a list of Areas, their meaning and usefulness

The default alarm in Fabric OS v4.1 for Fabric Areas in Table 3-7 is CHANGED and is also the suggested alarm for other Fabric OS versions.

CHANGED means that any time an E-port goes down, an alarm is sent, or anytime the fabric reconfigures or the zoning changes an alarm is sent. This is not useful for time based monitoring but it is very useful for configuration change alerts.

Note: Since all of these are fabric based messages there is no port based information specified such as the E-port number in an E_port down message. The number shown will be zero because it is a message coming from fabric zero (like port numbers, fabric numbers start from zero).

Table 4-13

Area to Monitor	Meaning	Usefulness
E-Port Downs	E-Port went down	Connectivity
Fabric Reconfigure	Fabric is reconfiguring	Connectivity
Segmentation	Fabric is segmented	Connectivity
Zoning Changes	Zoning has changed	Security
Fabric Logins	Port has logged on to fabric.	Connectivity Security

Another place to enhance the E-Port Downs message is under the Port Class. There are three Port Classes as mentioned previously. Setting the State Change area for any or all of the Port Classes to deliver a message can be used to correlate information.

For example, if an E-Port Down message is received (which has no specific port number associated with it) but an E-Port State Change is also received which contains a port number the two can then be correlated.

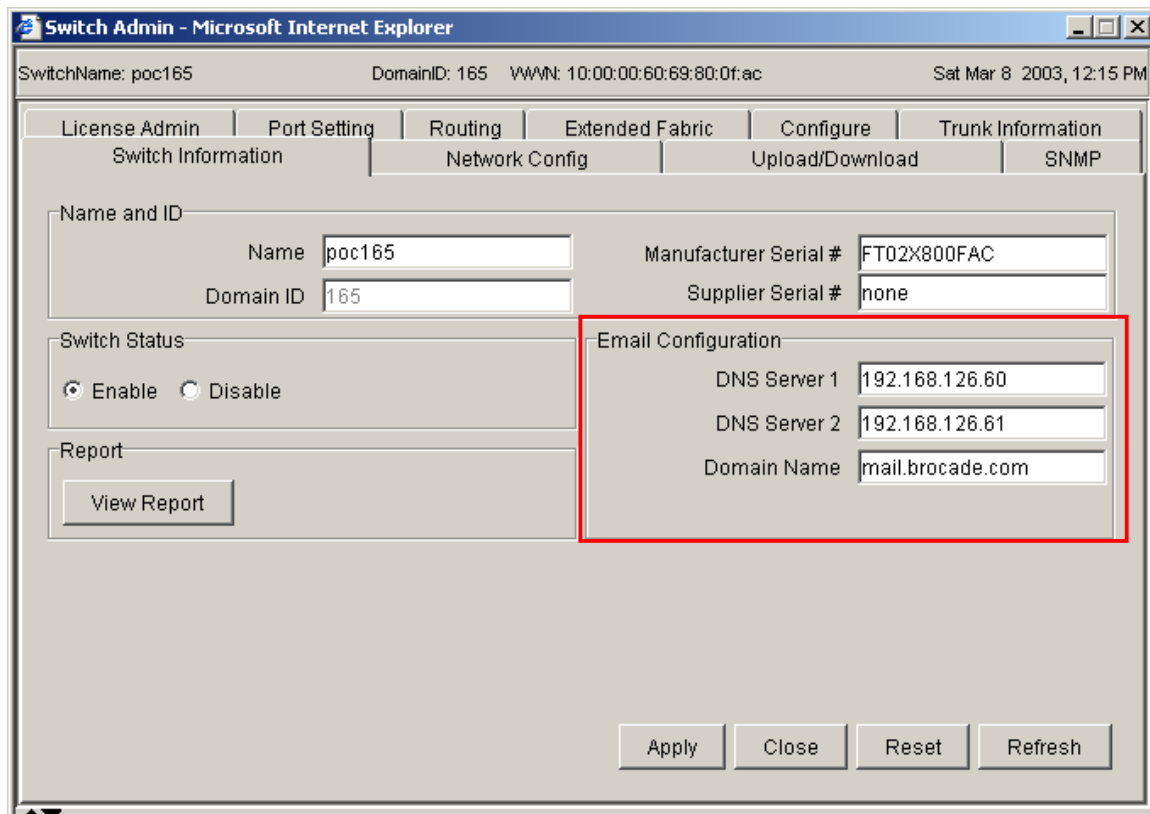
Another beneficial area to monitor in the Port Class is Transmit (TX) and Receive (RX) Performance. These messages can be used to track port performance and allow for improvements in routing and the design of the fabric.

An extended use of the TX/RX is to trigger an ABOVE alert for start of a storage back up exercise and then trigger a BELOW alert when the traffic has dropped off which could signify that the backup is done. This is a way to track the rate of backing up data over a period of time or to confirm if it went as long as it should have.

Additional performance features are available in the Performance Class. One to take note of is the EEPeformance Area which allows for performance monitoring from host to storage as opposed to switch port performance. This allows the administrator to expand storage allocation or allow for applications to be split up from one host onto two separate hosts.

4.2.3.2. Setting Up Fabric Watch E-mail

The easiest way to setup Fabric Watch E-mail is from the Web Tools GUI. Login into **Admin** section and select the **Switch Information** tab. In the lower right portion of the window is the Email Configuration shown in Figure 4-16. Enter in at least one DNS server and the Domain Name of the SMTP server. Click the **apply** button when finished.



The screenshot shows the 'Switch Admin' web interface in Microsoft Internet Explorer. The browser title is 'Switch Admin - Microsoft Internet Explorer'. The page header displays 'SwitchName: poc165', 'DomainID: 165', 'WWN: 10:00:00:60:69:80:0f:ac', and the date 'Sat Mar 8 2003, 12:15 PM'. The navigation menu includes 'License Admin', 'Port Setting', 'Routing', 'Extended Fabric', 'Configure', and 'Trunk Information'. The 'Switch Information' tab is active, showing fields for 'Name' (poc165), 'Domain ID' (165), 'Manufacturer Serial #' (FT02X800FAC), and 'Supplier Serial #' (none). The 'Switch Status' section has 'Enable' selected. The 'Email Configuration' section, highlighted with a red border, contains 'DNS Server 1' (192.168.126.60), 'DNS Server 2' (192.168.126.61), and 'Domain Name' (mail.brocade.com). A 'View Report' button is located below the status section. At the bottom of the page are 'Apply', 'Close', 'Reset', and 'Refresh' buttons.

Figure 4-16

E-mails are enabled/disabled based on class so log into the Fabric Watch GUI and click any class. Now click on the **E-mail** configuration tab (Figure 4-17). Fill out the Mail To address and enable/disable as needed. Check the **Mail Validation** check box to send out a test mail.

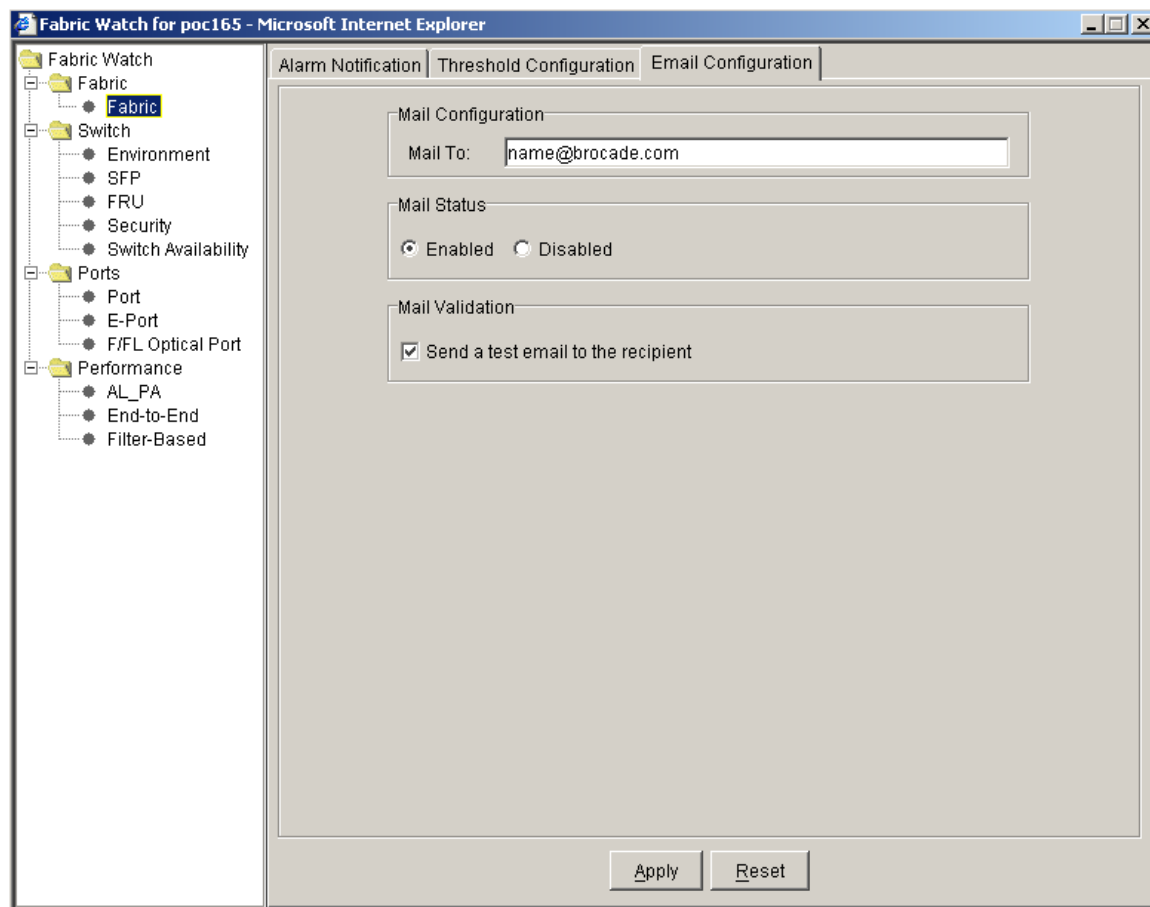


Figure 4-17

4.2.3.2.1 Security Class

This is a new class for 3.1/4.1 and contains 21 different monitoring Areas.

- | | |
|---------------------------|------------------------------|
| 1. Telnet Violations | 11. DCC Violations |
| 2. HTTP Violations | 12. Login Violations |
| 3. API Violations | 13. Invalid Timestamps |
| 4. RSNMP Violations | 14. Invalid Signatures |
| 5. WSNMP Violations | 15. Invalid Certificates |
| 6. SES Violations | 16. SLAP Failures |
| 7. MS Violations | 17. SLAP Bad Packets |
| 8. Serial Violations | 18. TS Out of Sync |
| 9. Front Panel Violations | 19. No-FCS |
| 10. SCC Violations | 20. Incompatible Security DB |
| | 21. Illegal Command |

Note that the Login Violations Area might create an unwanted level of redundancy by sending additional messages about the same error so extra care is needed in setting it in conjunction with the `TrackChanges` trap and the other Fabric Watch Area of Fabric Logins.

The breadth of security areas to monitor is quite impressive, allowing an administrator to keep a tight and secure fabric. As the breadth is large, the detail of each is even finer and is outside the scope of this document. Refer to the *Brocade Secure Fabric OS User's Guide* and the *Brocade Fabric Watch User's Guide* for a more in depth explanation of each Area.

Take note that the configuration, either from the Fabric Watch GUI in Web Tools or from the command line, of these Areas is exactly the same process, as previously described with the other Classes and Areas.

4.2.3.2.2 SAM Class

The SAM (Switch Availability Monitor) class monitors the efficiency of ports and helps identify problems with ports by providing statistics on switch downtime and uptime.

Areas of the SAM class are:

- Total Down time (in percentage)
- Total Up time (in percentage)
- Duration of occurrences
- Frequency of Occurrences

Table 4-14

Area	Description
total downtime	Indicates the total downtime of each F_port and E_port.
total uptime	Indicates the total uptime of each F_port and E_port.
duration of occurrences	Indicates the amount of time a port stays down
frequency of occurrences	Indicates how frequently a port goes down.

The `FwSamShow` command displays information about port availability. This provides total uptime, total down time, number of faulty occurrences and total time offline for each port.

4.2.4. Third-Party Software

There are a myriad of third-party SAN Management, Storage Resource Management, Event Management, and Network Management software available. Choosing one to fit a current need and yet have flexibility for growth is the most difficult part.

Once the software is chosen, there are a few things to remember when integrating it into a Brocade Fabric.

4.2.4.1. Management Server

When installing the software, it is general practice to make one dedicated server the management server to house the software installation. With this method it is easier to incorporate this one server into a dedicated IP subnet or secure fabric for tighter security.

Many software packages have adopted a client/server model which adds to the concept of installing the software or engine onto one server. Not only is the server in a more secure environment but now queries to the fabric can be minimized by having this model. When an administrator wants to modify or monitor the fabric, a lightweight client package is loaded onto a desktop or laptop and the user can log into the one server.

4.2.4.1.1 Loading New MIBs

Whether the MIBs are needed for the software to monitor a Brocade switch or the MIBs are needed in a MIB browser there is a specific order to load MIBs because of certain dependencies, as shown in Table 4-15.

Table 4-15

MIB to Load First	MIB to Load Second
SW	TRP (Brocade Trap MIB before v4.1)
REG-TC	SW (switch enterprise)
FRAMEWORK	FE – Fabric Element
ENTITY	HA – High Availability
FRAMEWORK	Entity
FA – Fibre Alliance	None

The four MIBs listed below must be loaded before the HA MIB:

- REG-TC
- SW
- FRAMEWORK
- ENTITY

4.2.4.1.2 Brocade Features Used as Reference

There are a couple of features that were mentioned in previous sections that third-party software will use to monitor switch health which will be revisited here.

The FA MIB is a standard MIB among Fibre Channel products and is commonly used to determine switch information as well as the operational status of a switch. Making sure the FA MIB support is turned on will ensure that it is discovered by most software packages. Refer to the section 2.1.1.1 on the SNMP command `snmpMibCapSet`.

Another feature is the SwitchStatusPolicy which is a Brocade implementation for determining the overall health of the switch. The health status is configured based on a number of critical components being down, like fans or power supplies. This health status might be used by the third-party software to determine the health of the switch and could cause confusion if not configured accordingly. Refer to the Fabric Watch section 2.1.2.1.6 on the command `SwitchStatusPolicySet`.

4.2.4.1.3 Trap Receivers

All Brocade switches send traps out UDP port 162. To receive such a trap, an application must be listening on the same UDP port. Additionally it is important to be listening to the same SNMP community that the trap originates from.

There are several applications that allow for both UDP port and community configuration on the trap receiver. If traps are being sent out of a community other than the standard Public (Read Only) or Private (Read/Write) it will be necessary to configure the receiver accordingly.

For further information on what value to configure for the public and private community strings, please refer to your site network security personnel. Please remember that these setting will allow anyone on or off of your network to either READ or READ/WRITE configuration setting on your switch. This may be a relatively large security breach. Please do not set the public and private community strings to arbitrary values, nor leave them at their default values.

4.3. Fabric Maintenance With Web Tools, Fabric Manager, and via SNMP

Web Tools and SNMP have been covered extensively as they pertain to event management. Now it's worth looking at them as tools for fabric maintenance.

4.3.1. Firmware Download

From time to time it is necessary and desired to upgrade the Fabric Operating System on a Brocade switch. The following sections provide several approaches and some insight.

4.3.1.1. Firmware Download with Web Tools and CLI

There are two CP cards in the SilkWorm 12000 running in active-standby mode. The proper version of firmware must be loaded on each CP card and it is highly recommended that the same version be running on both.

Note: On the SilkWorm 12000, Brocade Web Tools is only available via either of the two *logical* switch IP addresses, not the two CP card IP addresses.

The `version` command can be used to determine the current version of firmware running on each CP card or Brocade switch.

Firmware download on a per switch basis is just as easy with the CLI as it is with Web Tools. When using Web Tools, login into the switch and click on the **Upload/Download** tab.

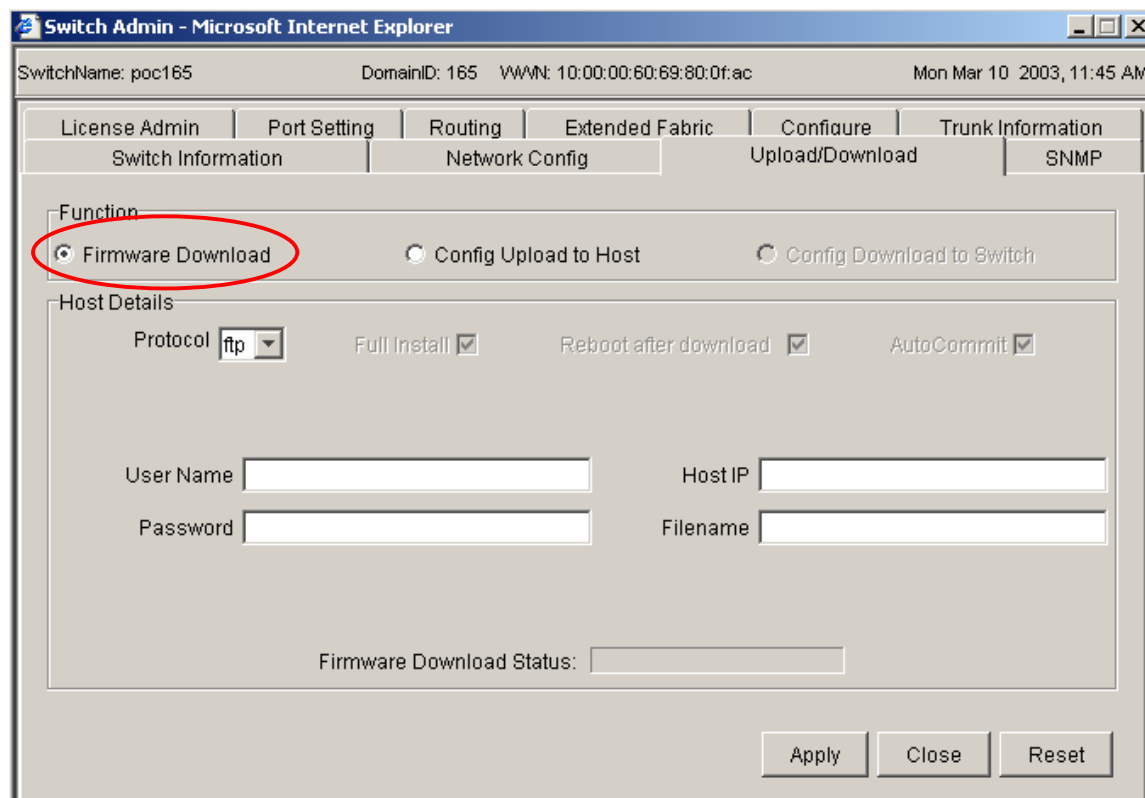


Figure 4-18 Firmware Download With Web Tools

Click the **Firmware Download** button and fill in the appropriate information for logging onto the server and the path/filename of the firmware. For SilkWorm 3900 and 12000 switches, the file name will always be *release.plist*.

The CLI command `firmwaredownload` requires the following responses from the user:

```
Switch1:admin> firmwaredownload
Server Name or IP Address [host]: 192.168.126.111
User Name [user]: johndoe
File Name [/usr/switch/firmware]: /usr/switch/firmware/v3.1
Protocol (RSHD or FTP) [rshd]: ftp
Password:
```

4.3.1.2. Firmware Download with Fabric Manager

A separate Brocade software product called Fabric Manager is a centralized fabric management tool which allows maintenance of multiple switches simultaneously from one location.

Downloading firmware to several switches at the same time is very convenient and time saving. To do this start Fabric Manager and enter one switch in a fabric that needs new firmware.

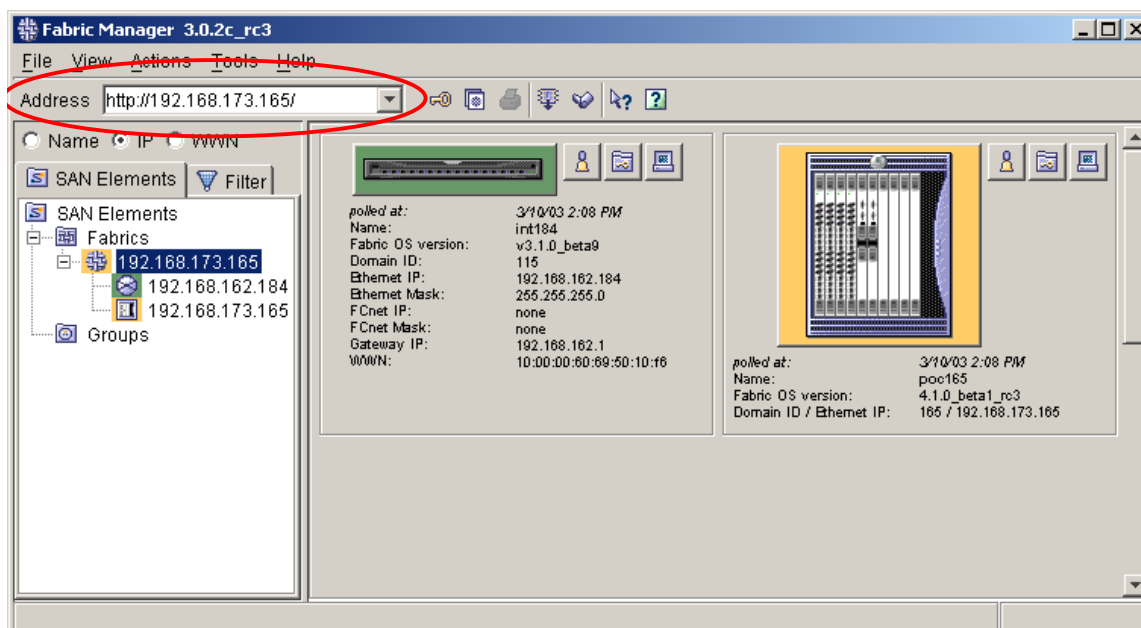


Figure 4-19

To save time with login information, Fabric Manager provides a login tool so that all login credentials can be entered and verified prior to executing any downloads.

To do this, click on **File** and select **Fabric Login** to open up a window as shown in Figure 4-20.

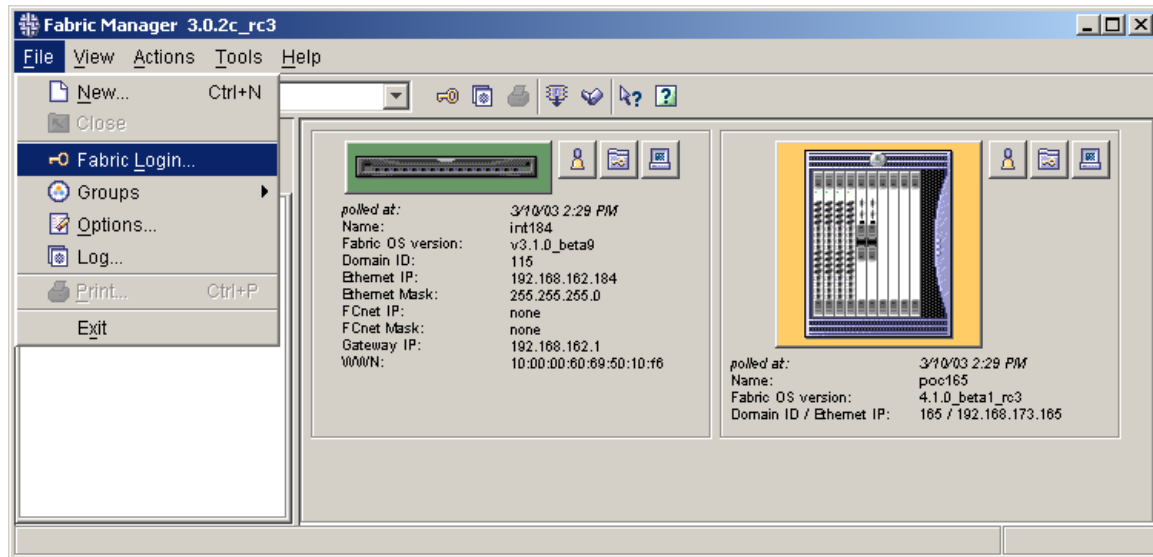


Figure 4-20 Fabric Login Setup Launch

Drag and drop, or click and use the arrows, to place the desired switches into the right hand side of the window.

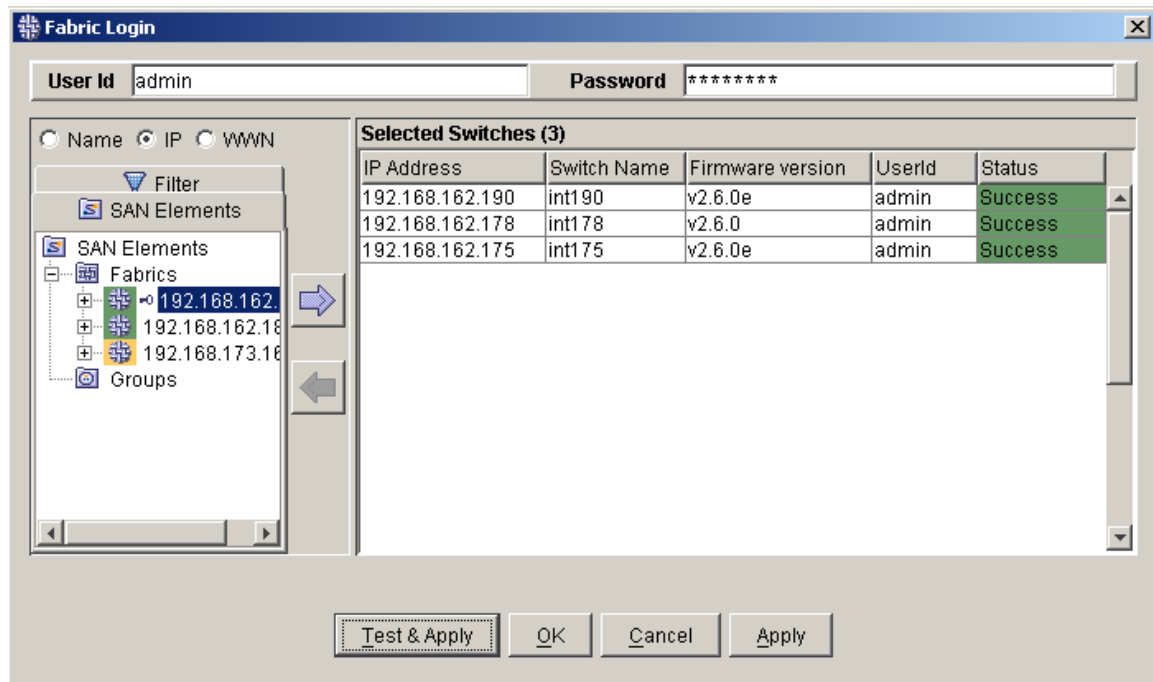


Figure 4-21 Fabric Login Window

Fill in the appropriate login credentials and click **Test & Apply**. If one should fail, confirm the user name and login and try again. If there is more than one password for each switch then the process has to be done for each one individually.

Once the login information is completed, there should be a little key symbol next to each switch. To start a firmware download sequence, click on the **Tools** menu (Figure 4-22) and select **Download Firmware**.

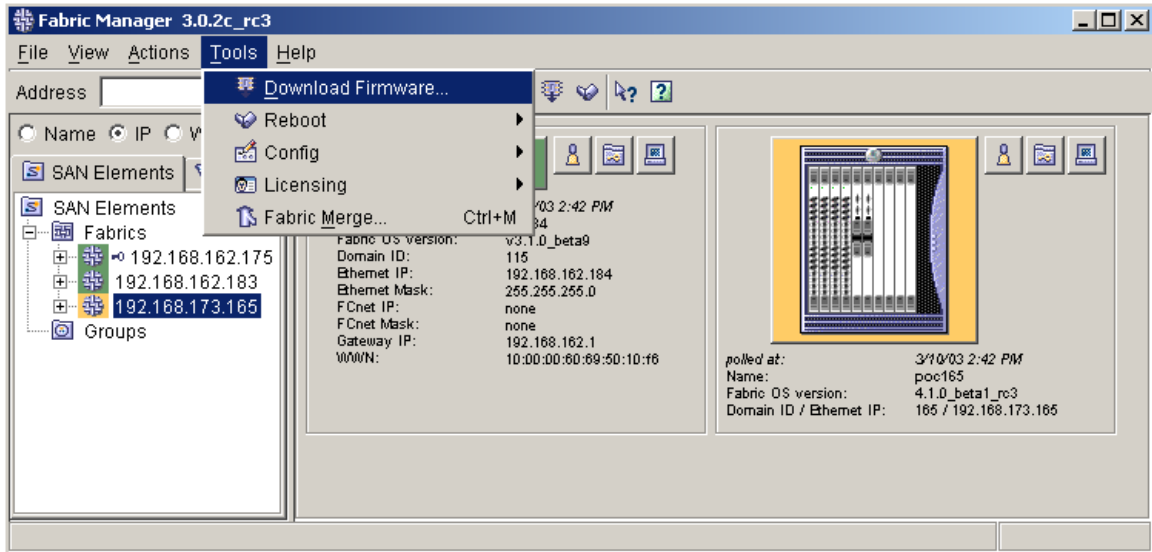


Figure 4-22 Download Firmware Launch

In the download firmware window, drag and drop, or click any like type switches that need firmware changed. Fill in the user name, password and file as in the CLI and Web Tools methods.

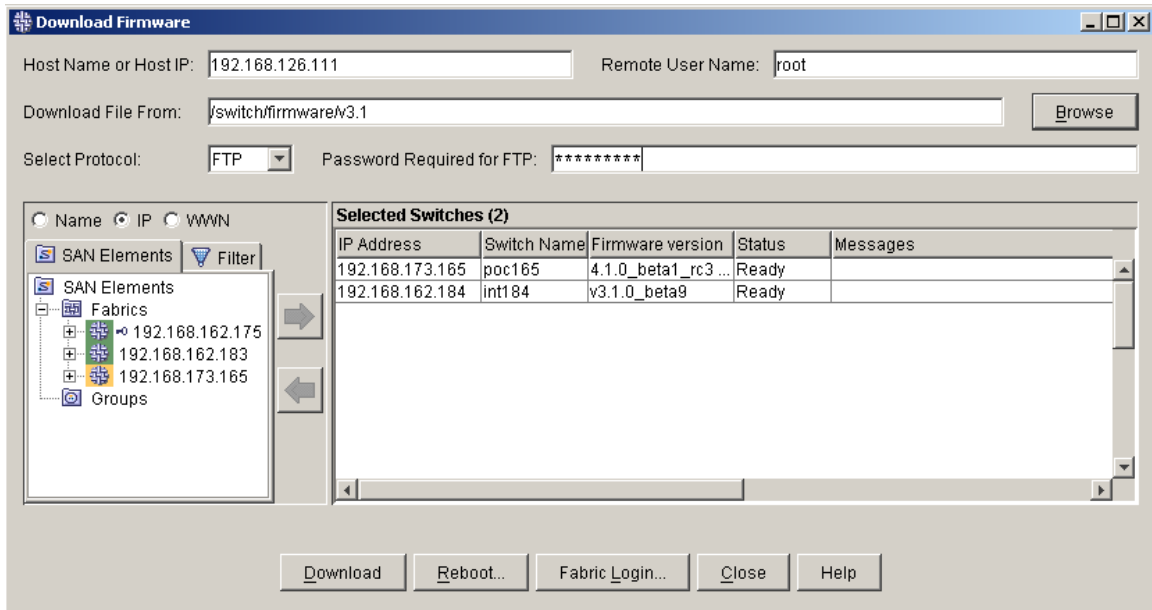


Figure 4-23 Firmware Download Window

Click **download** (Figure 4-23) when all information is complete. The download will begin and the current status of the download will be listed in the Messages line in the window.

A new command on the SilkWorm 12000 switches is `firmwaredownloadstatus` which relays the same information as the Fabric Manager messages, but can be accessed from the command line.

4.3.2. Zoning

The Zoning feature allows grouping of devices and/or ports for creating configuration files. When using this feature on the SilkWorm 12000, the one major change is how a port is selected. To select the appropriate port, it is now necessary to first select the slot on which the port resides as shown in Figure 4-24.

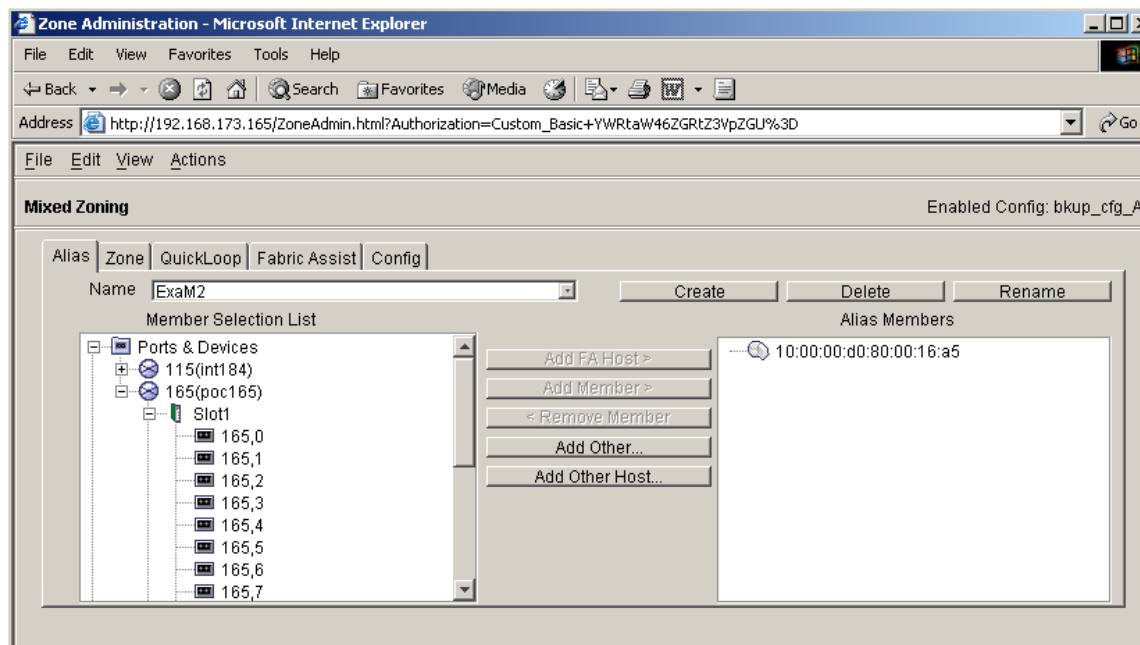


Figure 4-24 Web Tools Zoning Tool

The zoning commands available on the command line are slightly more difficult to master and are left for the power user. Please reference *Brocade Zoning User's Guide* for more information on this interface and other issues.

Fabric manager does not have it's own zoning tool and just references the Web Tools zoning window. To access this within Fabric Manager click on the **Actions** tab and select **Zone Admin**. This will launch the Web Tools zoning GUI.

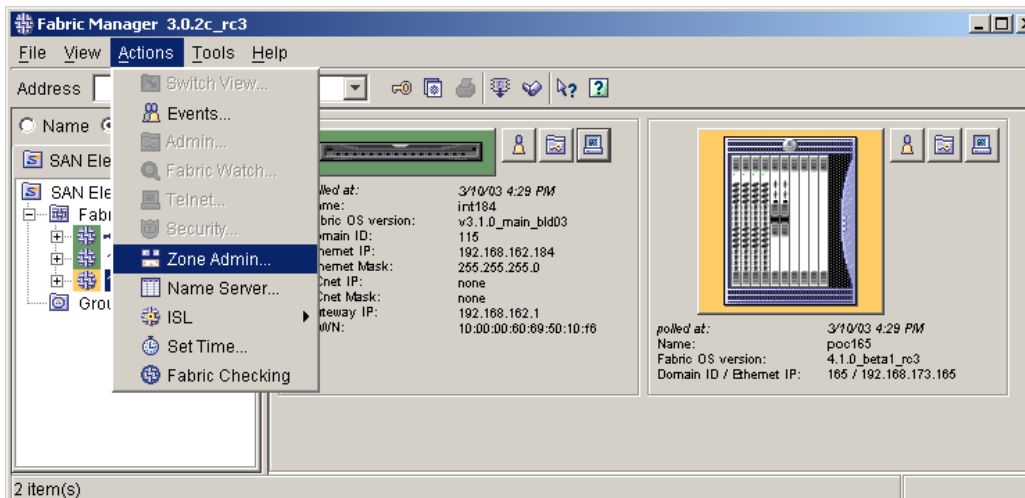


Figure 4-25 Launch of Zoning from Fabric Manager

4.3.3. Fabric Manager Sequenced Reboots

Fabric Manager has another convenient feature that allows for timed and sequenced reboots. The tool allows an administrator to manage switch reboots and minimize interruption of switch access. It is recommended to use the reboot sequence after doing a fabric wide Fabric OS upgrade to maintain the principal switch and to limit fabric disruption starting with the core switches.

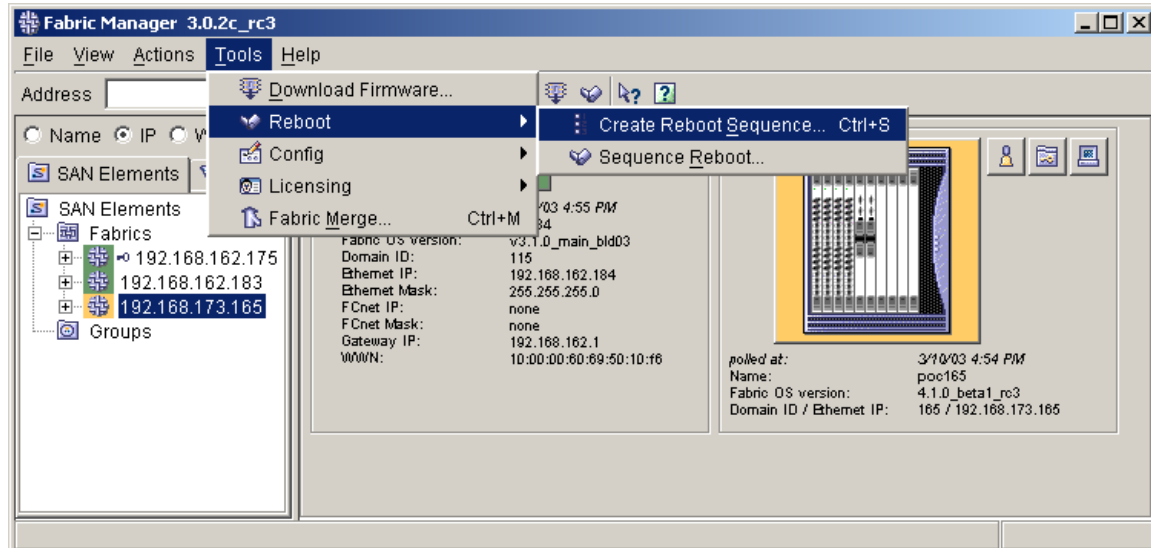


Figure 4-26 Selecting Create Reboot Sequence

To create a reboot sequence, click on the **Tools** menu within Fabric Manager and select **Create Reboot Sequence** under the **Reboot** category (Figure 4-26).

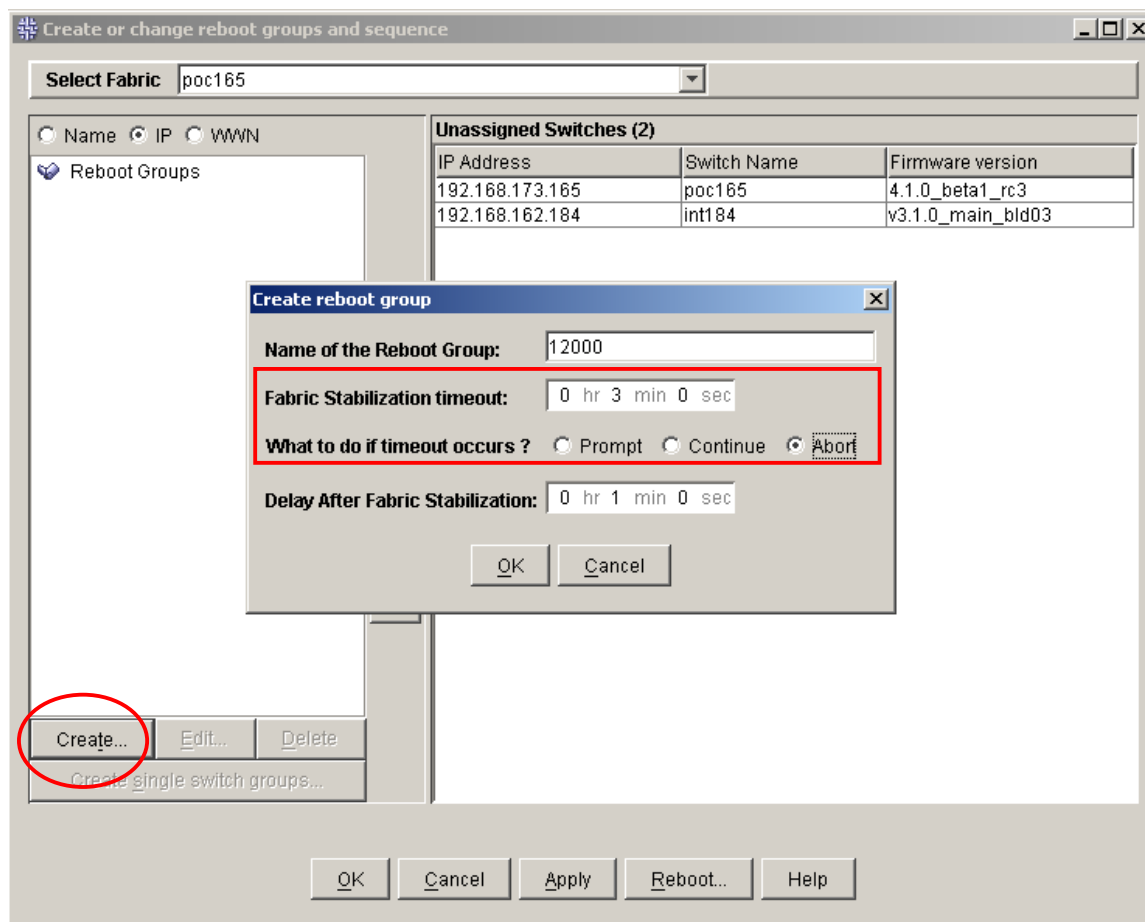


Figure 4-27 Creating Reboot Sequence

After the window appears, click on the **Create** button on the left side (Figure 4-27). There are four options in the small dialog box that also appears in Figure 4-27.

- Name of Reboot Group
- Fabric Stabilization timeout
- Action if timeout occurs
- Delay after Fabric stabilization

The second and third options are the most important to consider. Fabric stabilization can vary from fabric to fabric so it's best to always cushion the time to avoid overlapping reboots. One way to select a reasonable timeout value is by rebooting one switch and timing the length of a reboot all the way through stabilization.

The other choice to consider is what to do if a timeout does occur. Erring on the side of safety by aborting the reboot sequence is always a safe approach so that the timeout issue can be pursued.

Once the sequence is created, highlight it and then drag and drop or use the arrows to move desired switches into the sequence (Figure 4-28). The switch order in the sequence can also be configured by using the up and down arrows.

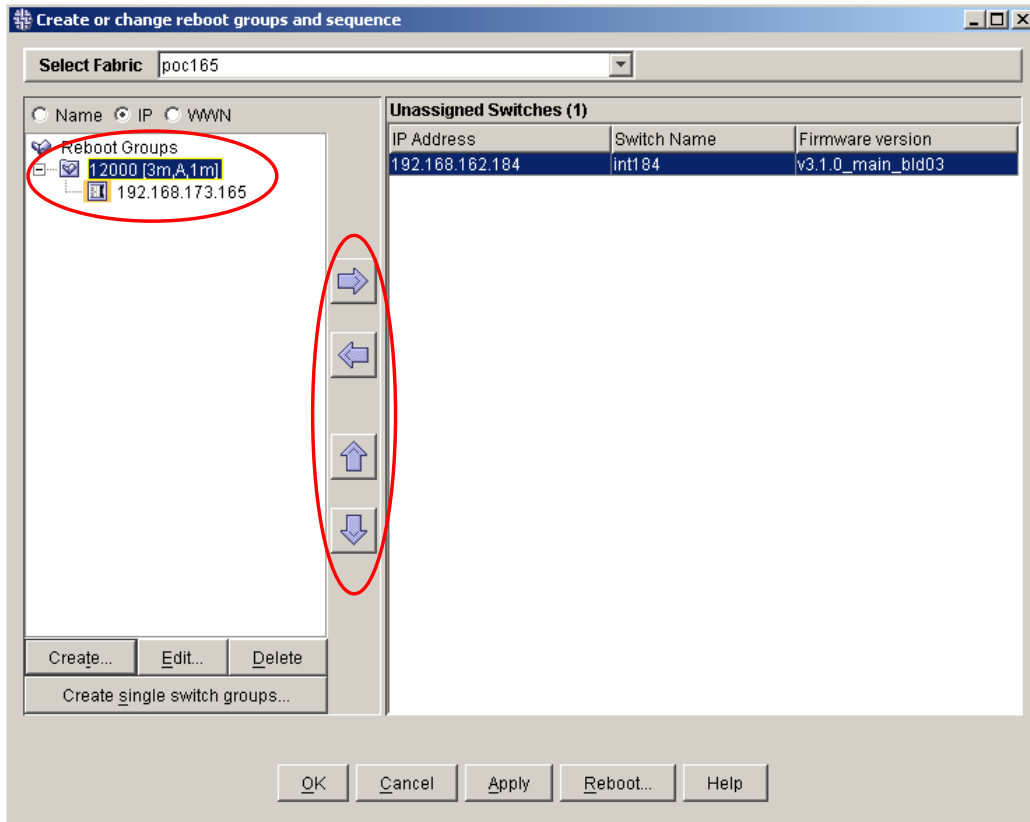


Figure 4-28 Selecting Switches For Reboot Sequence

To execute a Reboot Sequence click the **Tools** menu and select **Sequenced Reboot** under the **Reboot** subcategory. This brings up a Sequenced Reboot window shown in Figure 4-28.

Highlight the reboot sequence to be executed and click the right arrow, or drag and drop, to the right side. Now choose either **Fastboot** or **Reboot**.

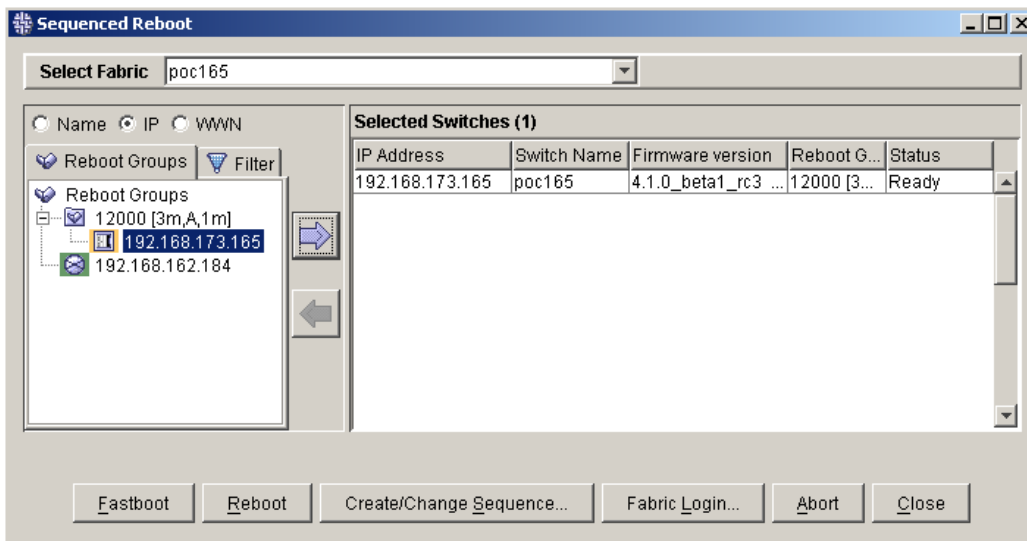


Figure 4-29 Executing a Sequenced Reboot

4.3.4. Periodical

Maintaining a fabric is not just firmware downloads, zoning, and rebooting it's also maintaining a record of current architecture, components, versions and much more. This section will touch on what aspects of a Brocade fabric should be backed up, saved, or catalogued in case it needs to be referenced or reproduced.

4.3.4.1. Configupload

The configupload command was discussed earlier as it pertained to propagating Fabric Watch configurations. Fabric Watch, as was probably noted, is not the only configuration parameter to be uploaded when executing the command.

Everything that is needed to reproduce the configuration on another switch is uploaded with the configupload command. With this, it is prudent to issue this command on all switches and maintain a catalogue of these records on a periodic basis.

In between the periodic uploads it is behooving to issue this command before any configuration changes or firmware downloads. This will ensure that, for example, if the zoning configuration gets mis-configured and/or forgotten, a back up is close at hand.

Fabric Manager also has a feature to upload configurations but has the added ability to filter specific portions of a switch configuration. In the Fabric Watch example that was given earlier, it was necessary to edit out all unwanted data except for the Fabric Watch parameters. This step would not be necessary by using Fabric Manager and specifically selecting only Fabric Watch parameters.

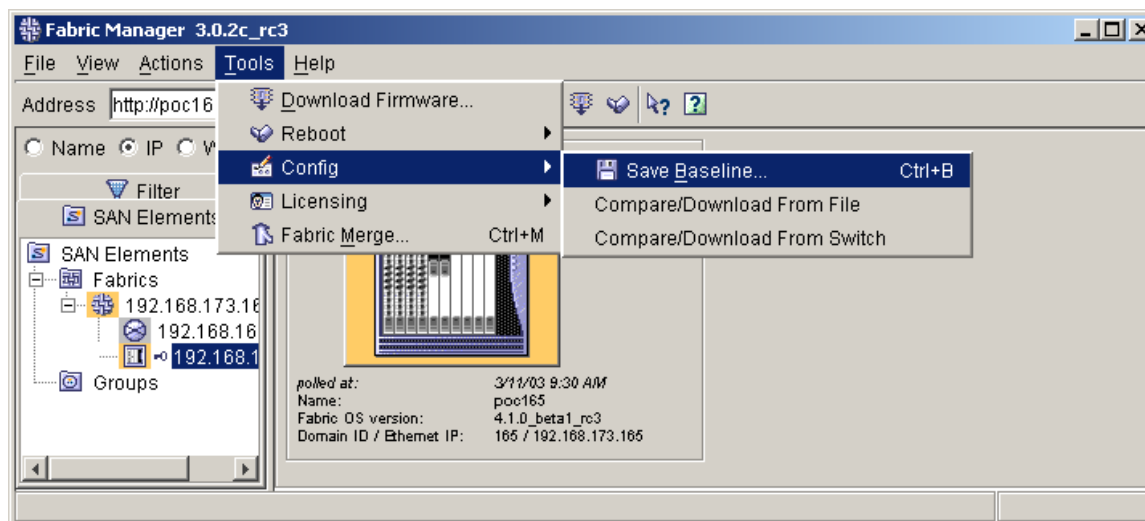


Figure 4-30 Launch Configuration Upload Tool in Fabric Manager

To access the configuration upload tool within Fabric Manager click on the **Tools** menu and select **Save Baseline** under the **Config** subcategory (Figure 4-30).

This opens up a template selection which allows for a quick upload of Fabric Watch and SNMP settings. If a more detailed manual selection of specific parameters is desired then select the **Full Configuration** (Figure 4-31).

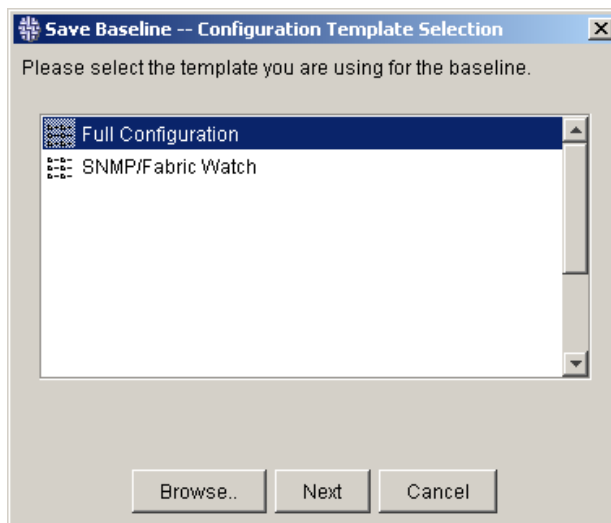


Figure 4-31 Configuration Template Selection

After choosing the **Full Configuration**, the switch selection tool appears. Select the desired switch and enter the proper login credentials if needed. After clicking **OK** a parameter selection tool appears. Open the tree by clicking on the plus symbol (+) to reveal a myriad of parameters as shown in Figure 4-32.

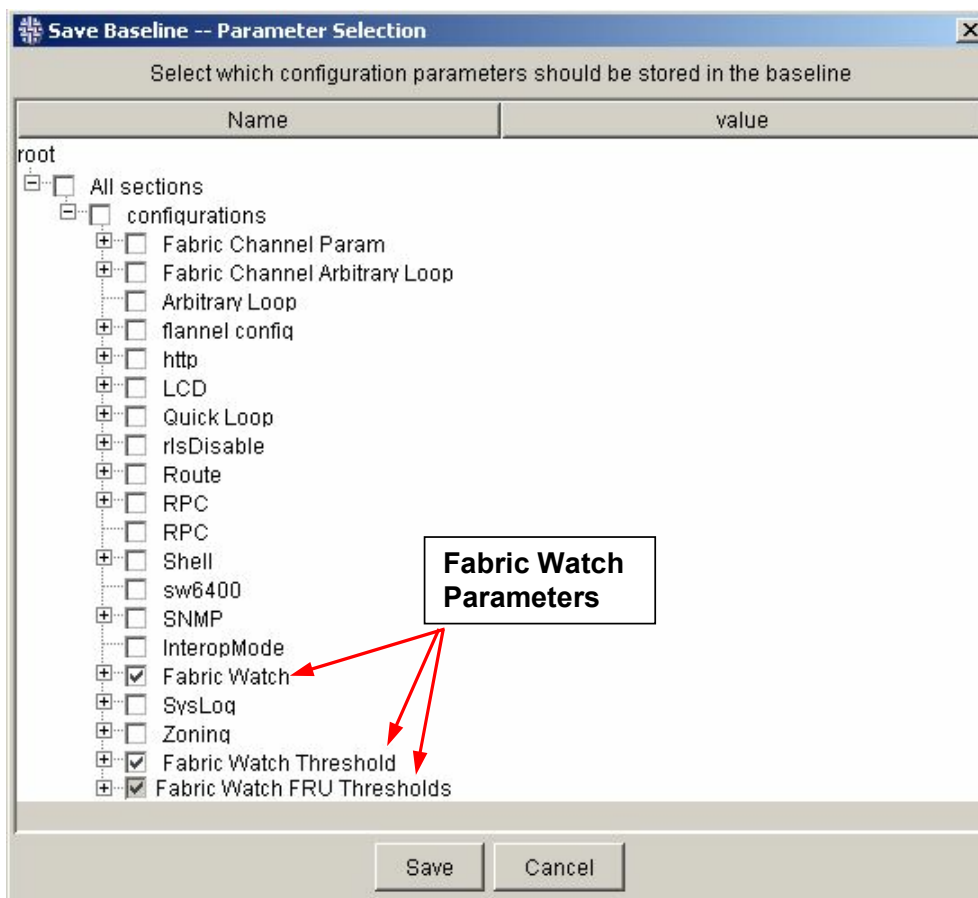


Figure 4-32 Configuration Parameter Selection

There are three separate Fabric Watch parameter selections, opening up each one reveals the specifics of each. The categories contain information accessed from the following CLI commands.

Parameter Selection Category	Equivalent CLI command
Fabric Watch	<code>Switchstatuspolicyset</code>
Fabric Watch Threshold	<code>Fwconfigure</code>
Fabric Watch FRU Thresholds	<code>fwfrucfg</code>

After selecting all desired parameters click **save**. The file will be saved to a location on the local host running Fabric Manager.

4.3.4.2. Other Show Commands

There are several other commands that will help build a complete picture of each switch and the fabric.

Supportshow – executes a number of CLI commands in a predefined order. Some of the commands include

1. *ISLshow* - This command will display the current connections and status of the ISL of each port on this switch. The WWN where the ISL is connected to, the speed of the connection, and whether this ISL is trunked, are displayed.
2. *topologyShow* - Use this command to display the fabric topology, as it appears to the local switch.
3. *nsShow* - This command displays local Name Server information, which includes information about devices connected to this switch, and cached information about devices connected to other switches in the Fabric.
4. *portShow* - This command displays port status information. Some information varies with the switch model and port type.

There are several other commands that are issued with the `supportshow` command and the entire list might contain too much data. If a more abbreviated catalogue of data is desired then individual commands should be executed separately. Refer to the Brocade Fabric OS Reference, v4.1 for more information on these commands.

4.3.4.3. Graphical Topology View

Another convenient way to represent and visualize the fabric topology is with a graphical view provided by third-party software.

The view can be used for easily identifying connections and for locating devices. Most software uses the graphical topology as the main centerpiece for real time monitoring which makes highlighting devices issues easier. Using a tool with such a feature can be invaluable in a SAN environment.

Note: Brocade Web Tools is only available through either of the two logical switch IP addresses. It is not accessible through either of the two CP card IP addresses.

4.4. Switch Performance

4.4.1. Advanced Performance Monitoring

Performance Monitoring is a comprehensive tool used for monitoring the performance of networked storage resources. This tool can help reduce over-provisioning while enabling SAN performance tuning and increasing administrator productivity.

4.4.1.1. New Features

There are some new design features that Fabric OS v3.1/4.1 brings to performance monitoring. The focus of the improvements is on enhancing the integration with the Brocade API.

The Advanced Performance Monitoring (APM) support for Fabric Access API will allow for

1. Ability to setup End to End (EE) monitors based on WWN of the Nx Ports or PID of the Nx ports and make them visible to Fabric Access Layer from any switch in the fabric
2. Ability to setup Filter-based monitors
3. Ability to get values of the monitors once they are set up.

There is also a new distributed framework where APM will have the capability to communicate with the APM on another switch thereby allowing an administrator to create or delete monitors on a remote switch in fabric.

This distributed framework is only available on Fabric OS 3.1/4.1 switches. A switch with older firmware (3.0.x/4.0.x) can be used for APM using 3.0.x/4.0.x level capabilities but it will not participate in the new distributed framework for Advanced Performance Monitoring. The Fabric Access API will also not be able to setup any performance monitors on a switch running Fabric OS v3.0.x/v4.0.x.

The same feature set is available in all Fabric OS versions 3.x and 4.x.

4.4.2. APM Example

To illustrate the functionality of Advanced Performance Monitoring, an example of will be given using a host, a JBOD and a SilkWorm 12000.

Starting from the Web Tools main page, click on the **Perf** button and a performance monitor window appears. Click on the **Performance Graphs** tab and select **SID/DID Performance** (Source ID/Destination ID) from the **Advanced Monitoring** subcategory (Figure 4-33).

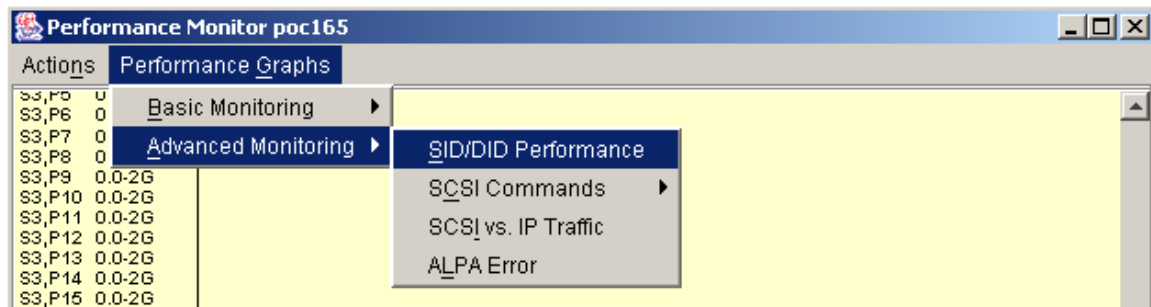


Figure 4-33 Setting up End To End Monitoring

With the SID/DID window open (Figure 4-34) select a source ID and a destination ID. The top box needs to have a slot and port number and this is the SID slot/port number. Click **OK** when finished.

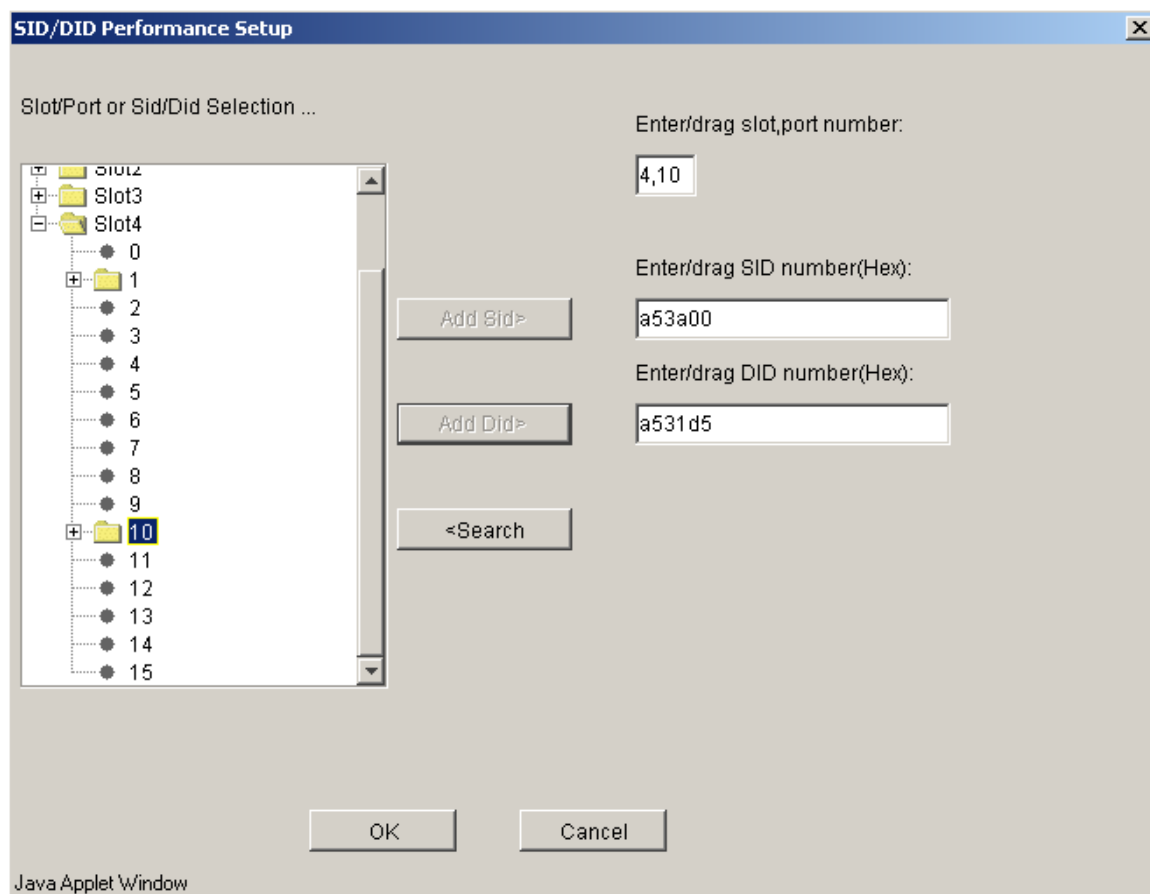


Figure 4-34 SID and DID Selection For EE Monitoring




Figure 4-35 EE Performance Monitoring

Figure 4-35 shows two graphs. The top graph is the default graph that always appears when starting out fresh. This graph measures the throughput for each port and will be used in comparison with the newly created EE monitor that is in the lower graph.

The lower graph is measuring TX/RX in Bytes/sec of slot 4 port 10 with a source ID of the host and a destination ID of disk d5 in the JBOD. The TX or RX is about 5 MB/sec for both and when compared to the upper graph and the overall slot 4 port 10 traffic is about 3.4% (5 MB/sec / 147 MB/sec).

From the CLI there is a similar process for setting up an EE monitor. Please refer to *Brocade Performance Monitoring User's Guide* for more information.

To view the current EE monitor that was setup in the Web Tools GUI type the command

```
Perfshoweemonitor 4/10, 5
```


Where the first argument 4/10, is the slot/port of the source ID and the second argument 5 is the time interval to measure. Figure 4-36 shows the data as it polls every five seconds.

```

Command Prompt - telnet poc165
poc165:root> perfshoweemonitor 4/10, 5
perfshoweemonitor 58, 5: Tx/Rx are # of bytes and crc is # of crc errors

  1
-----
crc   Tx   Rx
-----
0     0     0
0    19m  26m
0    19m  27m
0    26m  24m
0    31m  21m
0    22m  24m
0    22m  25m
0    22m  27m
0    23m  25m
0    26m  25m
0    24m  23m
0    19m  27m
0    20m  26m
0    27m  22m
0    25m  23m
  1
-----

```

Figure 4-36

4.4.3. APM and Fabric Watch

One way to use APM as a complementary tool with Fabric Watch is to have thresholds setup with Fabric Watch. When a threshold is breached the area can be monitored graphically real time by monitoring the area in question.

A straight forward example is when TX/RX traffic breaches some threshold. Having a graphical representation of that particular port by setting up APM can help to monitor traffic patterns and behavior over a period of time.

Being able to pull down the values that are seen in the APM graphs is not possible through APM. One way to get values is poll the number through SNMP or through API scripting which is introduced next.

4.4.4. Scripting possibilities with API

The Fabric Access API Scripting Toolkit provides a management interface to Brocade fabrics through the Brocade Fabric Access API. The API Scripting Toolkit can be employed to automate routine management tasks, collect fabric performance information, control devices and zones in the fabric and update switch configuration and firmware versions.

Management scripts employing the toolkit can access all the Brocade switches in a fabric from a single access point to:

- Discover Fabrics, Switches, Devices, Zones
- Control Switches, Ports, Zones, Licensing information
- Perform Switch configuration uploads and downloads, Firmware downloads
- Obtain Port traffic, Error information.

The API Scripting Toolkit is a Perl adaptation to the Brocade Fabric Access API. Perl is by far the most popular environment for enhancing operational management environments today. Most large organizations today have administrators who are very familiar with Perl and how it may be used to customize management environments.

One advantage to using API scripting instead of a telnet interface, Expect script or SNMP is that they all require a connection to every switch in a fabric. The API Scripting Toolkit on the other hand requires only one connection to the fabric and communicates with the other switches in the fabric via in-band facilities. This eliminates a potentially difficult problem if there are a lot of switches in the fabric or some switches are long distances apart.

The API Scripting toolkit uses a different TCP/IP interface so it does not conflict with the telnet interface. This means that scripts and the telnet interface can be used at the same time.

Another advantage is that Expect enforces some very severe restrictions on the telnet interface. Expect scripts are very sensitive to syntax and format changes in the commands and output they process. Even minor changes in command syntax and output formats will cause an Expect script to fail. The API Scripting toolkit uses an object structure which is not sensitive to format changes making the API approach much more robust. For more information on the API Scripting Toolkit education class go to http://www.brocade.com/education_services/index.jsp.

Glossary

A.1. Terms and Definitions

These terms and definitions are provided to ensure that a consistent language for describing SANs is used throughout the document and so that the reader understands what these terms mean. This section is not intended to be all-inclusive.

Table A-1 Key Terms and Definitions

Terms	Definitions and Impact
“golden” switch configurations	For larger SAN fabrics, or for the staging of many smaller fabrics, use <code>configupload</code> to baseline the “golden” switch. The “golden” switch configuration can be downloaded with Fabric Manager to all others in the fabric. Do baselines by Fabric OS version. In other words, if the fabric contains Fabric OS 3.1 and 4.1 switches create two “golden” switch configurations.
Blocking	The inability of one device to connect to another device. The Brocade Virtual Channel implementation of Fibre Channel does not block. The term blocking is often confused with the term congestion.
Congestion	If two or more sources contend for the same destination, performance for each source may decrease; however, available bandwidth is shared fairly by all sources contending for the same destination. Congestion is the realization of the potential of over-subscription. Congestion may be due to contention for a shared storage port or host port, or an ISL.
Control Processor	The term control processor is associated with a SilkWorm 12000 component/FRU (field replacable unit). The SilkWorm 2000, 3200, 3800, and 3900 switches do not have a FRU specifically associated with it and when CP is used in the context of other SilkWorm switches, the reference is to the switch CPU and not a FRU.
Core PID Format	The 24-bit Switch Fabric Port Identification (PID) also known as SID consists of Domain_ID, Area and AL_PA fields.
Core Switch	Also known as a “core fabric switch.” This is one of the switches at the logical center of a Core/Edge fabric. There are generally at least two core switches per Core/Edge fabric to enable resiliency within the fabric. Ports on a core switch are normally used for ISLs.
Edge Switch	This is one of the switches on the logical outside edge of a Core/Edge fabric. There are generally many more edge switches than core switches. Ports on edge switches are used for SAN device connections.
Fabric	One or more interconnected Fibre Channel switches. The term “Fabric” only refers to the interconnected switches, not to nodes or devices connected to the fabric.

Fabric build (BF)	<p>The build fabric Switch Fabric Internal Link Service requests a non-disruptive configuration to the entire fabric. A BF process shall not cause the Domain_ID list to be cleared. This preserves existing node port addresses and allows open exchanges to be completed.</p> <p>Impact: Fabric build is a non-disruptive process to I/O.</p>
Fabric Port Count	<p>The number of ports available to connect SAN devices in a fabric. ISLs ports (E-ports) are not included in this count. (Also known as user port count.)</p>
Fabric Re-Configuration (RCF)	<p>Fabric reconfiguration is a disruptive fabric operation during which domain IDs may change. If the Domain_ID changes, all attached node ports must re-login with the fabric and be assigned new N-Ports identifiers reflecting the change in Domain-IDs.</p> <p>Impact: Reconfigure causes Class-n frames (1,2,3,4 or 6) to be discarded and class 1 connection to be abnormally removed.</p>
Fabric Segmentation	<p>A fabric is unable to resolve the switch configuration parameters during the rebuild process with one or more switches, and may isolate them from the fabric, causing fabric segmentation.</p> <p>Impact: I/O operations are ceased only on those devices losing their access due to segmentation.</p>
Fabric Topology	<p>A topology is “the logical layout of the components of a computer system or network and their interconnections.” A fabric topology is the layout of the switches that form a fabric.</p>
Fan-in	<p>The ratio of storage ports to a single host port.</p>
Fan-out	<p>The ratio of host ports to a single storage port.</p>
FRU	<p>Field Replaceable Unit</p>
FSPF	<p>Fabric Shortest Path First protocol. The FSPF protocol was developed by Brocade and subsequently adopted by the Fibre Channel standards community for allowing switches to discover the fabric topology and route frames correctly. It is now the industry standard routing protocol for Fibre Channel networks.</p>
HA	<p>High Availability</p>
High Locality	<p>If devices that communicate with each other are connected to the same switch or groups of switches then these devices have high locality. The higher the locality, the less traffic crosses ISLs/trunks and therefore, fewer ISLs/trunks are needed.</p>
Hop Count	<p>For evaluating SAN designs, the hop count is identical to the number of ISLs that a frame must traverse to reach its destination.</p>
Host Edge Switch	<p>Edge switch with host device connections only.</p>
Incremental Upgrade	<p>Replacing one switch at a time in an online fabric.</p>
ISL	<p>Inter-Switch Link. ISLs connect two switches via E-ports.</p>

ISL Over-Subscription Ratio	In networks where all ports operate at the same speed, the over-subscription ratio for an ISL is the number of different ports that could contend for the use of its bandwidth. If there are 14 node ports on a switch and two ISLs, the ratio is 14:2, or 7:1. When there is a mixture of port speeds, the exact calculation is not as simple. The rule of thumb is that the lower the ratio is, the better performance is likely to be.
Latency	The time it takes for a frame to traverse from its source to its destination is referred to as the latency of the link. Sometimes a frame is switched from source to destination on a single switch and other times a frames must traverse several hops between switches before it reaches its destination.
Locality	The degree that I/O is confined to a particular switch or segment of a fabric. If two devices that need to communicate with each other are located on the same switch or fabric segment, then these two devices are said to have high locality. If these same devices are located on different switches or segments of a fabric and these two devices need to communicate with each other, then these devices are said to have low locality.
Logical Switch	The SilkWorm 12000 can contain up to 128 ports in a 14U chassis, configured as two 64-port switches. Each switch is known as a logical switch and may also be referred to as a domain.
Low Locality	If two devices must cross an ISL/Trunk to communicate, then these devices have low locality. The lower the locality, the more traffic crosses ISLs/trunks and therefore, more ISLs/trunks are needed.
NMS	Network Management Software
Node	Any SAN device – usually either a host or storage device – that attaches to a fabric.
Node Count	The number of nodes attached to a fabric.
Octet	An octet is a group of two adjacent quads. The SilkWorm 3900 is the only SilkWorm switch that implements octets. Octets are used primarily to define boundaries for performance tuning purposes.
Offline Fabric	A non-functional state of fabric unsuitable for I/O operation.
Online Fabric	A functional stable state of a fabric performing reliable I/O fabric operations.
Over-Subscription	A condition where more nodes <u>could potentially</u> contend for the use of a resource – such as an ISL – than that resource could simultaneously support, that resource is said to be over-subscribed.
PID bindings	Static mapping between physical and logical devices on a host accomplished via Port_ID (PID).
Radius	The greatest “distance” in hops between any edge switch and the center of a fabric can be thought of as that fabric’s radius. Low radius networks have lower hop counts and latency than high radius fabrics. The unit of measurement for a fabric radius is hops.
Redundant Fabric	A SAN composed of two or more independent fabrics The multiple fabric architecture makes dual fabric SANs redundant. Impact: SAN topology configured to provide two or more alternate paths for high availability.

Resilience	The ability of a fabric to adapt to or tolerate a failure of a component.
SAN	A Storage Area Network (SAN) can consist of one or more related fabrics and the connected SAN devices.
SAN Architecture	The overall design or structure of a storage area network solution. This includes one or more related fabrics, each of which has a topology. Other components may also be included, such as host, storage, and other SAN devices.
SAN Port Count	The number of ports available for connection by nodes in the entire SAN. The SAN Port Count equals the fabric port count in a single fabric SAN and is equal to the sum of each fabric's port count in a multi-fabric SAN.
Scalability	The ease with which a particular design can grow and adapt without requiring a significant change in SAN architecture or requiring a substantial re-layout of existing SAN devices.
Secure Mode Disabled	An operating mode where all switches that participate in the fabric are unable to successfully execute the command <code>secModeEnable</code> or if the command <code>secModeDisable</code> is successfully executed in the fabric.
Secure Mode Enabled	An operating mode where all switches that participate in the fabric are running a version of Fabric OS that supports the security feature, have licenses to run security, and the command <code>secModeEnable</code> has been successfully executed.
Single Fabric	A SAN composed of a single fabric may be configured to provide one or more paths via different switches of the fabric. Impact: Offers no Protection at fabric level. All paths are closed when fabric is offline, completely stopping I/Os.
SPOF	A single point of failure. A SPOF in a SAN is any component – either hardware or software – that could cause a fabric or a SAN to fail.
Storage Edge Switch	Edge switch with storage device connections only.
Tiering	The process of grouping particular SAN devices by function and then attaching these devices to particular switches or groups of switches based on that function
Total Ports	The total number of ports of all the switches in the SAN
User Ports	Total number of switch ports less ports used for ISLs/trunks

Reference Documentation

B.1. Brocade Documentation

The following related publications are provided on the Brocade Documentation CD-ROM and on the Brocade web site. To access Brocade Partner web site go to www.brocade.com and click on the **Partner Login** link.

- **Brocade Fabric OS documentation**
 - *Brocade Fabric OS Procedures Guide*
 - *Brocade Fabric OS Reference*
 - *Brocade Diagnostic and System Error Message Manual*
- **Brocade Fabric OS optional features documentation**
 - *Brocade Performance Monitoring User's Guide*
 - *Brocade Zoning User's Guide*
 - *Brocade Web Tools User's Guide*
 - *Brocade Distributed Fabrics User's Guide*
 - *Brocade Fabric Watch User's Guide*
 - *Brocade ISL Trunking User's Guide*
 - *Brocade Secure Fabric OS® User's Guide*
 - *Secure Fabric OS Quickstart Guide*
 - *Brocade QuickLoop User's Guide (v 3.1 only)*
- **Brocade Hardware documentation**
 - *Brocade SilkWorm 12000 Hardware Reference*
 - *Brocade SilkWorm 3900 Hardware Reference*
 - *Brocade SilkWorm 3800 Hardware Reference*

B.2. Additional Resource Information

The following related publications are provided on the Brocade Partner web site and are an excellent resource for additional information.

- *Brocade SAN Design Guide* (publication number: 53-0000231-05)
- *Core Switch PID Format Update Best Practices* (publication Number 53-0001626-01)
- *Designing Next-Generation Fabrics With Brocade Switches* (whitepaper <http://www.brocade.com>)
- *Exploring Brocade ISL Trunking* (publication number: 53-0000263-01)
- *LAN Guidelines For Brocade SilkWorm Switches* (publication number: 53-0000350-01)
- *SAN Migration Guide* (publication number: 53-0000360-01)
- *SAN Security: A Best Practices Guide* (publication number: GA-RG-250-00)
- *SilkWorm 12000 Core Migration User's Guide* (publication number 53-0000477-02)

B

Reference Documentation

- SOLUTIONware: *Enterprise LAN-Free Backup of Consolidated Storage on a SilkWorm 3800 Based Redundant 96-Port Integrated Fabric* (publication number 53-0000229-01)
- *Zoning Implementation Strategies For Brocade San Fabrics* (whitepaper <http://www.brocade.com>)



Brocade SAN Site Survey

Completed By:

First Name:	Last Name:	Phone:	Mobile
Email:		Pager:	Fax
Address:		Country/City/State/Zip	

Site Name and Location

Name:	Phone:	Fax
Address:	Country/City/State/Zip	

Site Contact #1

First Name:	Last Name:	Phone:	Mobile
Email:		Pager:	Fax
Address:		Country/City/State/Zip	

Site Contact #2

First Name:	Last Name:	Phone:	Mobile
Email:		Pager:	Fax
Address:		Country/City/State/Zip	

Site Contact #3

First Name:	Last Name:	Phone:	Mobile
Email:		Pager:	Fax
Address:		Country/City/State/Zip	

Site Profile Summary & Implementation Schedule:

Include the purpose or purposes of the SAN. For example, LAN-Free Backup, Storage Consolidation, Remote Replication, etc. In addition, please add any other critical information regarding this site. The list of items should include but is not limited to site power, fiber cable requirements (including location, lengths, etc), any fiber patch panels, proposed rack layout, etc.

Attach a project plan with a Visio diagram, proposed rack layouts and an installation schedule. When complete, the SAN project manager/SAN Architect will be able to assess the staging requirements and be able to assign/schedule resources as needed.

Survey Authorization

Site Representative:	Date:
-----------------------------	--------------



Brocade SAN Site Survey

SAN Topology Assessment

Item 1	Please attach a high level list of all hardware devices to be attached.	Completed <input type="checkbox"/>
--------	---	------------------------------------

Notes: Each device should have the make/model identified. If 2 topologies are identical in every way then a single list will be sufficient. This list must include all switches, HUBS, DWDM, Gateways, and other devices for connecting SANs over distance.

Item 2	Please attach a high level topology diagram that illustrates the location of each existing or new fabric member switches in the SAN at this site. If possible, include the edge devices as well.	Completed <input type="checkbox"/>
--------	--	------------------------------------

Notes: Each SAN should have every device represented by an ICON. Each fabric should be represented as a cloud that can be referenced to the correct topology. A single SAN frequently incorporates more than one fabric (redundant fabric architectures for example). When illustrating such a SAN it is important to show how all devices connect to each fabric. Since clouds will be used it is not necessary to identify which switch/port each device has been connected to.

Switch Output

Item 3	If migrating into an existing SAN please attach the supportShow output from each of the existing switches	<input type="checkbox"/>
--------	--	--------------------------

Notes: The output of supportShow should be saved to a file and then attached to this document. If there are any non-brocade switches in the fabric then similar information should be captured.

Device Configuration

Item 4	For each device in the SAN please provide a detailed profile. See below for a list of information that should be included	Completed <input type="checkbox"/>
--------	---	------------------------------------

Complete hardware profile

Server		Storage		Miscellaneous Device	
OS	Patch level	Disk type	Disk Capacity	Make	Model
Bus Speed	Bus Type	Disk Firmware	# of Disks	Firmware	Driver
CPU	Other Data	RAID Controller	Firmware	Other Information	

Port Identification/Information (For Affected Existing SANs)

Identification of each Fibre Channel card in that device. Include the Firmware and Driver
 For each card, identify each port and the switch/port it is connected to.
 For each port, identify what device(s) it has access to across the SAN.
 For each port, identify what fabric zones it is a member of

Notes: Each device should be reference as follows XXXXXXXXXX-yy-zz, where the first segment is a 10 character ID for a device, the second segment is a 2 digit number identifying the slot, and the last segment is a 2 digit number identifying the port.

Software Application Information

Identify each software application required for that device.
 Backup Applications
 Multi-path Applications
 Management Application
 Other applications
 LUN Security Configuration

Survey Authorization

Site Representative:	Date:
----------------------	-------

Brocade SAN Site Survey

SAN Requirements Checklist

Target number of user device ports in SAN environment (with any phasing of implementation)	
Number of Fabrics in SAN	
Number of sites in environment	
Distance requirements	
Number /types of hosts	
Number /types of storage devices	
Number /types of tapes	
Other devices	
Customer requirement for fail over/redundancy, reliability of SAN	
Customer scaling plans	

SilkWorm 12000 Pre-Install Checklist

Task	Complete
The voltage required is 200 to 240 VAC, 50 Hz or 60 Hz. Two dedicated branch circuits required for redundancy. Confirm the power cords ordered with the system to ensure the correct power receptacles are installed. <ul style="list-style-type: none"> o In the US and most of North America, the available voltage is usually 208 or 240 VAC, the receptacles are NEMA L6-20R, on individual branch circuits, each rated 20 amps.) o In UK, Ireland, Hong Kong, two UK-standard 13 amp receptacles, on individual branch circuits, are required. o In most of continental Europe, two CEE7/7 "Schuko"-compatible receptacles are required, each rated 16 amps, OR two IEC60309, 230V~16A-6h receptacles. Verify with the system order. o In Australia and New Zealand, 2 Australian-standard, 15 amp receptacles, on individual branch circuits rated 15 amps each, are required. o For any location in which the US, UK, "Schuko" or AUS/NZ receptacles types are not accepted, the International standard, IEC-60309, 230V~, 16A-6h, receptacles should be planned, and the system order confirmed. 	RECEPTACLE <input type="checkbox"/> NEMA <input type="checkbox"/> UK <input type="checkbox"/> CEE7/7 <input type="checkbox"/> AUS/NZ <input type="checkbox"/> IEC-60309
Appropriate licenses ordered	<input type="checkbox"/>
Soft copy of manuals provided to customer	<input type="checkbox"/>
Customer training scheduled	<input type="checkbox"/>
Brocade personnel have access to customer site	<input type="checkbox"/>
Verify that customer loading docks are adequate for delivery	<input type="checkbox"/>
Verify appropriate shelf kits are available: standard (4 posts) or telco (2 posts)	<input type="checkbox"/>
Assure adequate aisle space from staging to install area for 12000(s)	<input type="checkbox"/>
Verify cooling vents and cut tiles available	<input type="checkbox"/>
Cables on site (LC-LC, SC-LC)	<input type="checkbox"/>
If existing switches are going to be connected to the 12k, are these switches running at least v2.6.0c (2000 series) or v3.0.2c (3000 series) ?	<input type="checkbox"/>
Is core pid format enabled on existing SilkWorm 2000 & 3000 series switches?	<input type="checkbox"/>

Survey Authorization

Site Representative:

Date:



SAN Components Inventory

The worksheets in this Excel workbook are intended to be used as a template and starting point for you to record:

- an inventory of your current storage infrastructure and analysis of whether those existing components should be used in your Storage Area Network (SAN)
- new components needed for your SAN
- component compatibility
- needed port count for your SAN

Enhance the worksheets as you see fit to meet your needs. The completed spreadsheet can be used to help you purchase your SAN components.

For more detailed information on how to use these worksheets, please visit the following steps in the Planning and Design section of the Brocade SAN Info Center (www.brocade.com/san):

[Step 2: Inventory and Analyze Your Environment](#)

[Step 3: Determine Your SAN Components](#)

We appreciate your feedback regarding the usefulness of these worksheets, and how we might improve them for use by others. Please send your comments to:

saninfocenter@brocade.com



SAN Components Inventory: HOSTS

Host Identifier							
New (N) or existing (E)							
Manufacturer							
Make							
Model							
Operating System							
Total Slots							
Type A							
Type (PCI, SBUS, etc)							
Number							
Size							
Speed (Mhz)							
Type B							
Type (PCI, SBUS, etc)							
Number							
Size							
Speed (Mhz)							
Single (S) or Dual (D) Attached Ports							
# Ports required for single attached							
# Ports required for dual attached							
HBA							
General information							
- make							
- model							
- version							
- number of ports							
- number of HBAs							
- slot numbers							
- DMP/Failover Application							
Driver type: (check applicable type)							
-fabric							
-PTP							
-private loop							
-public loop							
Connections supported per HBA							
Fibre Channel Ready? (Y or N)							
Server							
- Dimensions							
- Power Requirements							
- Console location							
- Ethernet interface list							
- Phone line required for mgmt? (Y/N)							
- Physical location							
Application list (with version)							
Application Name 1							
Application Name 2							
Application Name 3							
Application Name 4							



SAN Components Inventory: HOSTS

Host Identifier							
Storage requirements							
Requirement 1							
Application Name							
Initial Requirements							
Projected Requirements							
Requirement 2							
Application Name							
Initial Requirements							
Projected Requirements							
Requirement 3							
Application Name							
Initial Requirements							
Projected Requirements							
Requirement 4							
Application Name							
Initial Requirements							
Projected Requirements							



SAN Components Inventory: SWITCHES

Switch Identifier						
Zoning info						
Firmware Version						
IP Address(es)						
Gateway						
Port configurations						
0						
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						
11						
12						
13						
14						
15						
16						
17						
18						
19						
20						
21						
22						
23						
24						
25						
26						
27						
28						
29						
30						
31						
32						
33						
34						
35						
36						
37						
38						
39						
40						
41						
42						
43						
44						
45						
46						
47						
48						
49						
50						
51						
52						
53						
54						



SAN Components Inventory: SWITCHES

Switch Identifier						
55						
56						
57						
58						
59						
60						
61						
62						
63						
Devices attached to each port (type, WWN, etc.)						
0						
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						
11						
12						
13						
14						
15						
16						
17						
18						
19						
20						
21						
22						
23						
24						
25						
26						
27						
28						
29						
30						
31						
32						
33						
34						
35						
36						
37						
38						
39						
40						
41						
42						
43						
44						
45						
46						
47						
48						



SAN Components Inventory: SWITCHES

Switch Identifier						
49						
50						
51						
52						
53						
54						
55						
56						
57						
58						
59						
60						
61						
62						
63						
Licensing						
Security information						
Passwords						
Control Port info						
Ethernet, telnet						
Special settings						
Operating temperature						
Power requirements						
Phone line required for mgmt?						
Yes or No (Y/N)						
Physical location						
Firmware level						

Device Name	Number of Devices	Rack Unit per Device	Total Used Rack Units
Fiber Patch Panel			0
KVM Unit			0
Host 1			0
Host 2			0
Cable Guide			0
Silkworm 2800		2	0
Cable Guide			0
Silkworm 3800		1	0
Cable Guide			0
Silkworm 3900		1.5	0
Cable Guide			0
Silkworm 12000		14	0
SW 12000 Cable Guide			0

Total Units Available			42
Total Units Used			0
Total Units Remaining			42

