

# SAN SECURITY: A BEST PRACTICES GUIDE

Incorporating SAN security  
into the enterprise with the  
Brocade Secure Fabric OS



# Table of Contents

<b>Table of Contents</b> .....	<b>3</b>
<b>Preface</b> .....	<b>5</b>
<b>1 Introduction</b> .....	<b>7</b>
<b>1.1 THE CONCEPT OF SECURITY</b> .....	<b>7</b>
1.1.1 <a href="#">General Security Threats to the Enterprise</a> .....	8
1.1.2 <a href="#">High-Level Design Considerations</a> .....	8
1.1.3 <a href="#">High-Level Threat Mitigation Procedures and Tools</a> .....	9
<a href="#">Policy Documentation</a> .....	10
<a href="#">Essential Security Functions</a> .....	10
<a href="#">SAN Security Tiers</a> .....	10
<b>1.2 WHY SAN SECURITY?</b> .....	<b>11</b>
1.2.1 <a href="#">Vulnerabilities and Typical Threats to a SAN Environment</a> .....	11
<a href="#">Fabric or Remote Device-to-Switch Traffic</a> .....	12
<a href="#">Inter-Switch Traffic</a> .....	13
<a href="#">Traffic to or from Fabric Devices</a> .....	13
<a href="#">Conventional SAN Security Vulnerabilities</a> .....	14
<b>2 Brocade Secure Fabric OS (SFOS)</b> .....	<b>15</b>
<b>2.1 NON-SFOS SAN SECURITY MEASURES</b> .....	<b>15</b>
2.1.1 <a href="#">Zoning</a> .....	15
<a href="#">Zoning Types</a> .....	16
<a href="#">Zone Enforcement</a> .....	17
<a href="#">Zoning-Specific Security Issues</a> .....	17
2.1.2 <a href="#">Locking Down E Ports</a> .....	17
2.1.3 <a href="#">Physical Access</a> .....	18
2.1.4 <a href="#">Remote Access</a> .....	18
<a href="#">Example: Minimizing Remote Access Password Exposure</a> .....	20
2.1.5 <a href="#">Security Summary</a> .....	22
<a href="#">Risks to a Non-SFOS-Configured SAN</a> .....	22
<a href="#">Non-SFOS SAN Security Functionality</a> .....	22
<b>2.2 ADVANCED SECURITY THROUGH THE SECURE FABRIC OS</b> .....	<b>23</b>
2.2.1 <a href="#">Secure Fabric OS Overview</a> .....	24
2.2.2 <a href="#">Fabric Configuration Server (FCS)</a> .....	25
<a href="#">Switch Type Overview</a> .....	25
<a href="#">FCS Switch Functionality</a> .....	26
<a href="#">Supplemental Information on the FCS Group and Passwords</a> .....	26
<a href="#">FCS Policy</a> .....	26
<a href="#">Passwords</a> .....	27
<a href="#">Time Server</a> .....	30
<a href="#">Database Management and Propagation</a> .....	30
2.2.3 <a href="#">Management Access Controls (MAC)</a> .....	31
<a href="#">Example 1: Telnet and HTTP Policies</a> .....	32
<a href="#">Example 2: API Policy</a> .....	35
<a href="#">Example 3: Physical Policies</a> .....	35
2.2.4 <a href="#">Secured Management Channel</a> .....	36
<a href="#">Protecting the Fabric with Certificates</a> .....	40
<a href="#">Obtaining Digital Certificates for Switches</a> .....	40
2.2.5 <a href="#">Switch Connection Controls (SCC)</a> .....	41
<a href="#">SCC Example</a> .....	41
<a href="#">Switch Link Authentication Protocol (SLAP)</a> .....	42
2.2.6 <a href="#">Device Connection Controls (DCC)</a> .....	42
<a href="#">DCC Example</a> .....	42
2.2.7 <a href="#">SFOS-Protected SAN Security Summary</a> .....	44
<a href="#">SFOS SAN Security Functionality</a> .....	45

<b>3</b>	<b>Implementation and Design Examples</b>	<b>46</b>
<b>3.1</b>	<b>CASE STUDY 1: A MID-SIZED ENTERPRISE ENVIRONMENT</b>	<b>46</b>
3.1.1	Current Environment	46
3.1.2	Description	47
3.1.3	Detailed Steps	48
	Pre-SFOS Checklists	48
	Collecting Supplementary Information	49
	Enabling Secure Mode in the Fabric	49
	Checking Default SFOS Functionality	51
	Proposed Usage of SFOS Features	51
	Implementing CorpA's SFOS Security Plan	52
	Allow Remote Access	53
	Lock Down Unused Switch Access	53
	DCC Policy Creation	54
	SCC Policy Creation	55
	Policy Review	56
	Verify Other Configuration Information	56
<b>3.2</b>	<b>CASE STUDY 2: GROWING THE CORPA ENVIRONMENT</b>	<b>57</b>
3.2.1	Current Environment	57
3.2.2	Steps	58
3.2.3	Additional Information	59
<b>4</b>	<b>Conclusions</b>	<b>60</b>
	<b>Appendix A: Securing the Enterprise</b>	<b>61</b>
	<b>MANAGEMENT CONTROLS</b>	<b>61</b>
	Risk Determination	61
	Understanding the Components of Risk	61
	System Characterization	61
	Threat Analysis	62
	Using the Values to Determine Risk Level	63
	Review of Security Controls	64
	Life Cycle Management	64
	<b>OPERATIONAL CONTROLS</b>	<b>64</b>
	Physical Security	64
	Contingency Planning and Disaster Recovery	64
	Documentation	65
	Security Policy	65
	Incident Response Plan	65
	Other Policies	65
	<b>TECHNICAL CONTROLS</b>	<b>66</b>
	Border Router and Access Control Lists (ACLs)	67
	Firewalls	67
	Intrusion Detection Systems (IDS)	68
	Virtual Private Network (VPN)	68
	Out-of-Band (OOB) Networks	68
	Port Scanners and Vulnerability Scanners	69
	<b>Appendix B: References</b>	<b>70</b>
	<b>BOOKS AND WEB RESOURCES</b>	<b>70</b>
	Broad Security Information	70
	Technical Controls	70
	Router Lockdown	71
	Enterprise Architecture Guides	71
	Security Policies	71
	Other Useful Sites	71
	<b>Copyright</b>	<b>72</b>

## Preface

Storage Area Networks (SANs) have experienced explosive growth in the past few years. Increased fiber channel speeds, disk storage size, and the quantity of data that needs to be readily accessible are all key driving factors. Mission critical business requirements have spurred growth in various facets of SAN technology, however the marketplace is just now beginning to identify the need for enhanced security in these solutions. The Brocade Secure Fabric Operating System (SFOS) seeks to address evolving customer security requirements effectively – offering a key differentiator for Brocade switch products.

Organizations understand that the data managed in their SAN environment is often highly sensitive and must be controlled properly to ensure confidentiality, integrity, and availability. A compromise in any of these attributes could have unintended consequences resulting in the loss of proprietary information, capital, or other core business resources. As an industry-leading SAN solution provider, Brocade is incorporating security as an essential component of enterprise architecture.

The intent of this guide is to educate the user on many of the security risks facing SAN implementations and risk mitigation techniques using The Brocade Secure Fabric Operating System. The document begins with a broad introduction to enterprise-wide security “best practice” concepts before narrowing its focus to target the SAN environment. Because each element of an integrated infrastructure can potentially have a serious impact on the other components, a holistic approach to SAN security should first be understood and then implemented as the organization’s needs dictate. This guide illustrates the new security functionality provided by enabling the Secure Fabric OS®, demonstrates how to implement the features correctly, and shows how a protected SAN fits into a comprehensive enterprise security strategy.

The security topics are designed to appeal to a broad base of users ranging from project managers to security architects and SAN administrators. The document presents broad security concepts and goes into some detail on the rationale behind why various preventative measures should be implemented. The document first presents information at a conceptual level to allow the reader to understand many of the dynamics of the security environment, then presents specific examples. The document readdresses important topics iteratively to provide progressively more detail and to allow the reader to *digest* the information.

Since there is no perfect security model that applies to all environments, each organization’s design will be predicated on first understanding various pieces of information and then striking an appropriate balance between them. Not every system designer or administrator will need to implement all of the security functionality, however this document provides information on how to evaluate risks that might exist in their SAN environments and illustrate security features they can employ to mitigate these risks. In this way, the reader can use an established security methodology, then customize it to the particular business needs of the organization.

The subsequent chapters will address the following concepts:

- Industry best practices regarding enterprise-wide computer security
- Information on common system and network security threats
- Threat mitigation tools and procedures
- How SAN security fits into the enterprise security model
- Common risks facing the SAN environment

- **How The Brocade Secure Fabric OS (SFOS) mitigates many security risks facing today's SANs**
- **Case study examples on how to integrate Secure Fabric OS features into representative SAN implementations**
- **Appendices providing supplemental information and “jumping off” points for additional research and reading**

Important Note: Some might wonder why a document that targets SAN security should address generic security topics. The answer is quite simple: an enterprise is only as secure as its weakest element. When a security methodology is used, it should be appropriately applied to all elements. In much the same way that a network PC and an installed modem card can greatly minimize the effectiveness of a properly configured perimeter firewall or border router, so too could a weak security infrastructure render a carefully secured SAN vulnerable. Each security element in the enterprise should complement the other components and add strength to the overall design. Security designs require a careful balancing act between usability, performance, cost, and a variety of other factors. This document presents a broad range of security concepts, along with SAN specific security to illustrate clearly the functionality of The Brocade Secure Fabric OS.

# 1 Introduction

## 1.1 THE CONCEPT OF SECURITY

Security is an important component of today's enterprise computing infrastructures. The process of defining, implementing, and managing security policies is pervasive and often an essential business requirement. Numerous management, operational, and technical controls are available to provide layers of protection for the enterprise computing infrastructure. While these measures can be quite effective when correctly implemented and maintained, they do require understanding of some basic security concepts and an awareness of some common tools that protect against various types of accidents and attacks. The common adage "knowledge is power" is particularly applicable to security environments. Organizations that understand their core business infrastructure, vulnerabilities, and risk exposure levels are better prepared to defend themselves from accidental or intentional misuse or abuse.

The topics in this section provide an introduction to some fundamental security concepts; additional details appear in the ensuing chapters. The document addresses design considerations, security threats, controls, and general SAN security topics. This information forms the foundation for illustrating SAN security best practices as well as initiating the process of developing a comprehensive security methodology.

The guidelines and recommendations in this document are aimed at securing the SAN environment. SAN security divides into three tiers:

- **Low Security:** A conventional SAN environment with minimal security customization
- **Medium Security:** A fabric secured with traditional (that is, non-SFOS) configuration options and third-party tools
- **High Security:** An SFOS-secured fabric

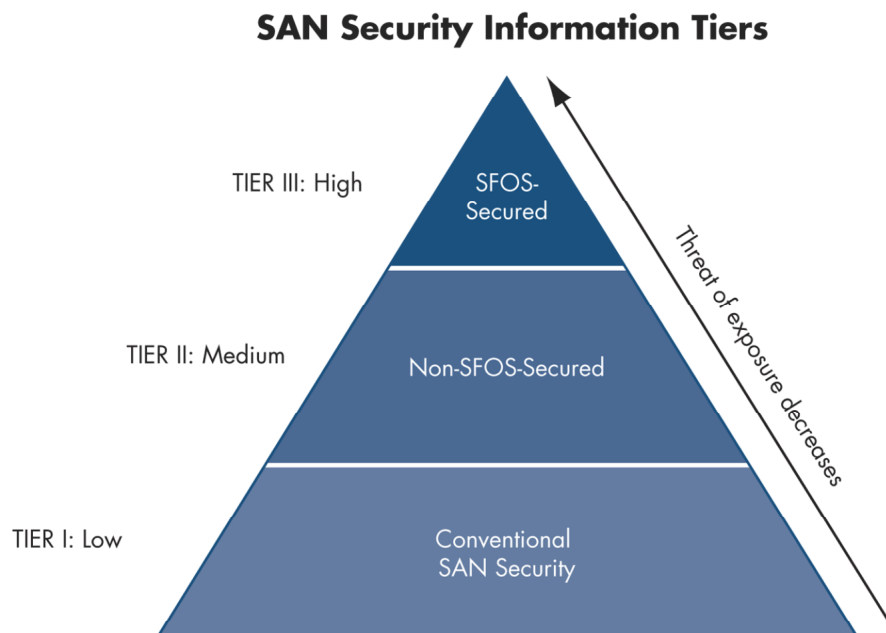


Figure 1: Three SAN Security Tiers

As the fabric progresses from the conventional to the SFOS configuration, the exposure to a variety of threats and attack types decreases greatly. The information below begins by explaining high-level threats to the enterprise, then discusses risk mitigation techniques, and finally shows the relatively significant exposure of a conventional SAN environment.

### 1.1.1 General Security Threats to the Enterprise

A basic assumption in security design is that attackers will target the weakest component in the system. That is to say, why bother trying to thwart complex security measures or attack a hardened system when compromising an easier device would yield the same result? A SAN represents a wealth of valuable data and sensitive information. If the fabric has been configured by conventional methods and lacks enterprise security controls, it might be a fairly easy target. To understand threats against the fabric itself, it is important to understand potential attacks to the environment as a whole.

Would-be attackers have myriad attack vectors, methods, and tools at their disposal. Below are some general categories and common examples of each:

- Inappropriate privilege to change data; for example:
  - Defacing an organization's Web site
  - Editing, deleting, corrupting, or otherwise modifying sensitive data on a JBOD storage system, disk array, or other storage devices
  - Changing switch configurations
- Inappropriate use of an organization's resources (such as, bandwidth, CPU, storage, and so on); for example:
  - Removing or impairing access to a resources with a Denial of Service (DOS) attack
  - Using a compromised dual-homed host with a Host Bus Adapter (HBA) to access SAN resources and read, store, or distribute illegal files; (also called *warez*)
  - Using systems to launch an attack on other organizations
- Inappropriate access to view data; for example:
  - Viewing confidential email or sensitive data files
  - Obtaining access to sensitive business secrets, perhaps stored on SAN-accessible storage devices
  - Accessing salary or other HR information

A huge variety of attack types are available for use against an organization, and newer, faster, better methods are surfacing each day. Different attackers have different goals. Some might be seeking sensitive information, while others are just looking to test their skills or cause a few disruptions. In any case, it is important to have an understanding of what elements or information would be desirable and more likely targeted. It is then imperative to understand how to protect those vulnerable elements. A security framework strives to provide a balance between security and functionality so as to protect high-risk systems adequately without impacting core business processes. The following section contains important considerations when designing a security framework.

### 1.1.2 High-Level Design Considerations

Although security has garnered a high degree of publicity, it can be difficult to know how to protect different systems or data sets appropriately. Each organization has its own unique set of requirements and circumstances that dictate what preventative steps are in line with the level of acceptable risk. A very conservative security posture that might be quite costly and provide diminishing returns for one company might be an absolutely essential part of



another organization's security requirements. Below are some of the most common design considerations for organizations designing a security framework that addresses their unique business needs:

- **Risk Determination:** A value assigned to a system or set of systems (such as, critical, high, moderate, and low) based upon the likelihood of the system being attacked, and the impact to the organization if that system were compromised or rendered inoperable. These values help determine the priority of the various enterprise elements.
- **Coordination of Security Controls:** The ability to manage and maintain multiple security measures logistically and ensure that preventative measures do not conflict with one another, or with core business operations.
- **Cost:** The amount of time and money the organization plans to spend on security controls. The cost should not outweigh the value of the benefits provided by the security measures.
- **Technical Expertise:** The level of sophistication in a security design depends on the technical expertise of the resources available. Expertise is a critical element in ensuring that systems retain the appropriate confidentiality, integrity, and availability intended by the system owner.
- **Corporate Policy:** The purpose of security is often to enforce the organizational policies of an organization with regard to data access, acceptable Web traffic, and the like. The central policy document, typically called a Security Policy Document, provides written guidance on a variety of security related topics.

Refer to Appendix A for additional information on each of these design guidelines.

### 1.1.3 High-Level Threat Mitigation Procedures and Tools

After reviewing high-level design considerations and potential attack methods, system designers can consider defensive measures. Many of the security layers can be configured to ensure that only the appropriate users or systems have access to data. Determining the specific values for these different variables and then correctly implementing steps to enforce the policies can be quite challenging. The following three control types are important to creating a complimentary matrix of checks and balances that appropriately covers the security gamut:

- **Management Controls:** Procedures that ensure that there is proper management oversight with regard to security. These controls seek to answer the question, "What will ensure proper management and control of elements throughout their life cycle?" Addressing these issues requires specific processes and controls that ensure correct overall security management.
- **Operational Controls:** Processes that are typically implemented by security personnel (human controls). These controls work well in conjunction with management controls (oversight) and technical controls (typically non-human hardware or software controls). Security documents, such as a Corporate Security Policy document and Acceptable Use Policy document, are also considered part of this control type.
- **Technical Controls:** Security measures implemented by hardware and software. In conjunction with strong operational controls and management controls, technical controls can detect unauthorized access, track changes, implement security policies, and perform numerous other functions. The Brocade SFOS offers a wide range of technical controls for securing the SAN environment.

This document presents the important aspects of operational, management, and technical controls. However, the majority of the information addresses the various technical controls provided by the SFOS. For additional research or reading in any of these areas, refer to the suggested Web links, references, and supplemental information provided in the appendices at the end of this document.

## Policy Documentation

Proper documentation is an essential part of any complete security program. Security Policy and an Acceptable Use Policy documents are key references in defining security implementation and enforcement. These documents vary widely between organizations. Some might be quite comprehensive and detailed, while others might outline basic policies. What is important is that the documents be usable and that they enhance overall business effectiveness. Security policies must provide more value than the cost and difficulty of implementing them.

Other useful documents that play important roles in maintaining consistent processes and supporting existing controls include, but are not limited to these topics:

- Change control policies
- Testing procedures
- Backup procedures
- Emergency procedures
- Disaster recovery plans
- Security incident policy

Other documentation targeting other specific aspects of the enterprise might be necessary. While not all of these processes fit directly under the umbrella of security, they serve to bolster the overall security infrastructure.

## Essential Security Functions

Examples of security components include: firewalls, Access Control Lists (ACLs), Intrusion Detection Systems (IDS), Virtual Private Networks (VPNs), port scanning software, virus scanners, and The Brocade Secure Fabric OS. Each of these is an example of a technical control that can address specific security concerns in the enterprise. Each element plays an important and distinct role in providing a comprehensive security solution. At a high level, all security components share some common functionality in that they perform one or more of the following functions:

- **Integrity:** Protecting data from unauthorized, unanticipated, or unintentional modification
- **Availability:** Ensuring that resources are accessible on a timely basis when needed
- **Confidentiality:** Protecting information from unauthorized disclosure
- **Authentication:** Proving that an entity is in fact who they claim to be
- **Authorization:** Ensuring that only appropriate entities have access to specified resources
- **Accounting:** Providing appropriate records of specified activities (such as in the form of logs)

## SAN Security Tiers

As mentioned, SAN security divides into three tiers that range from least secure to most secure. The three levels are:

- **Conventional SAN:** This SAN consists of a working fabric in which minimal time and effort has gone into securing or “locking down” the environment.
- **Non-SFOS-Secured SAN:** This SAN uses the available functionality of the Brocade Fabric OS, in conjunction with other security controls, to increase protection of the fabric and its elements.
- **SFOS-Secured SAN:** This is the most secure SAN implementation of the three. It builds on the work done in the previous tier, but adds the significant enhancements available in the SFOS to lock down and secure the environment further.

The following table illustrates the utilization of each security function in each of the three tiers:

## SECURITY FUNCTIONALITY FOR EACH SAN TYPE

Functionality	Conventional	Non-SFOS-Secured	SFOS-Secured
Integrity	No	No	Yes
Availability	No	Limited	Yes
Confidentiality	No	Limited	Yes
Authentication	Limited	Limited	Yes
Authorization	Limited	Limited	Yes
Accounting	Limited	Limited	Yes

Table 1: Essential Security Functions

The conventional and non-SFOS fabrics only utilize a very few of the six important security functions. The SFOS, on the other hand, is specifically designed to incorporate each of these items into various parts of the secured fabric to minimize vulnerabilities and protect against a compromise of the fabric. Later sections in this guide will discuss in greater detail how each of the essential security functions are incorporated into the three SAN tiers and how provide SFOS benefits compared to the other design types.

## 1.2 WHY SAN SECURITY?

In order to appreciate the business value of the SFOS, it is important to answer to the central question: “Why SAN security?” The answer derives from two central design philosophies:

- Defense in depth
- The weakest link

The concept of defense in depth simply states that having multiple, integrated layers of security is better than relying on one monolithic security point. This ensures that the failure of one security control will not compromise the assets under protection. Defense in depth is conceptually analogous to dressing for cold weather, in which case it is better to wear multiple warm layers of clothing rather than just putting on one very thick sweater. SANs have certain inherent risks and vulnerabilities that are unique to their environment as well as sharing more common security problems such as physical security and remote administrative access. Thus, in order to secure the SAN, an organization must implement security measures specific to the SAN as well as leverage the other security implementations in the enterprise.

Security designers should integrate the SAN into the overall enterprise security model and incorporate SAN-specific security measures in any secure infrastructure. Because SAN-specific security is a relatively new facet of the overall security model, it is important to understand the potential risks to that environment. The information in the next section discusses some of the common vulnerabilities and threats to the SAN environment.

### 1.2.1 Vulnerabilities and Typical Threats to a SAN Environment

There are myriad ways, both intentional and accidental, in which SAN integrity can be compromised. Incidents can range from the unintentional, such as a device accidentally plugged into the wrong switch port, to the deliberate attack of SAN components. If a mission critical infrastructure element goes down, either by accident or through malicious activity, the consequences can be quite severe. The Brocade Secure Fabric OS is designed both to minimize mistakes that impact production and mitigate intentional hostile threats.

A SAN is subject to numerous potential vulnerabilities. The diagram below illustrates common SAN threat points. SAN threat types fall loosely into three categories:

- Device-to-switch traffic (1)
- Inter-switch communication (2)
- Traffic to or from fabric devices (3)

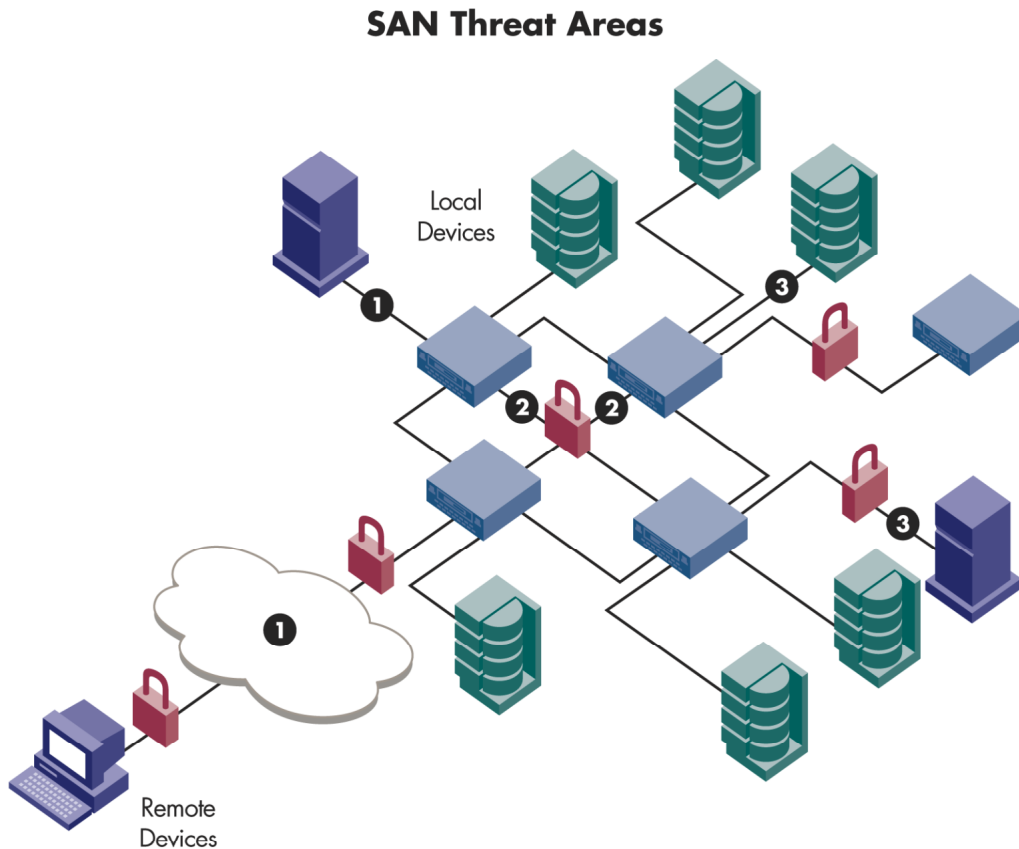


Figure 2: Vulnerabilities and Typical Threats to a SAN Environment

The following sections address each of these categories and detail how each of these potential risk areas can cause problems in a conventional SAN environment. Later sections of this document readdress the same SAN security risks, but will compare them to the more carefully designed non-SFOS-secured and SFOS environments.

### Fabric or Remote Device-to-Switch Traffic

This category includes any traffic that goes between a non-switch device and a fabric switch using Fiber Channel, IP, or other means. These devices can be systems that plug directly into the SAN fabric via a Host Bus Adapter (HBA) as well as other systems, such as PCs used by remote administrators to manage the SAN. This type of risk can include traffic from invalid sources, attempted Denial of Service (DOS) attacks, and the like.

## DEVICE-TO-SWITCH TRAFFIC

Risk	Description
<b>An unauthorized, local (non-switch) device is connected to the fabric or a device is connected to or from the wrong port</b>	A problem arises if an unauthorized device is plugged into a switch and tries to log into the fabric. Because switch ports are not locked down to a specific system, the device might be able to access sensitive data, corrupt file systems, invalidate hard zoning configurations, or cause other problems. Managing physical access to the switch and providing the ability to lock a port or group of ports down to a set of valid World Wide Names (WWNs) eliminates the possibility of a device accidentally connecting to the fabric from an unspecified set of ports.
<b>Invalid management device connection attempt</b>	In order to manage a switch, some sort of connectivity is required. Regardless of whether the access method uses telnet, console, Web, or front panel access, limiting which devices have access to configure the fabric or switches is essential. Having the ability to disable or enable access, or lock the systems down to specific IP addresses for remote access methods would add an additional layer of security and help prevent hostile entities (such as "crackers") from breaking into a switch.
<b>Unprotected management application authentication information</b>	One common problem with telnet (and numerous other network tools) is that they send the login name and password information over the network "in the clear," that is, non-encrypted. Packet sniffers are commonplace and because this user/password combination provides the majority of remote access security in today's SANs, it requires additional protection.
<b>Device Denial of Service (DOS) attack</b>	Attackers might connect unauthorized devices to the SAN in order to disrupt the fabric. That is to say, generating traffic that expends significant switch resources (such as CPU, memory, bandwidth, and so on) can impact the effectiveness of the SAN. This risk illustrates another reason for carefully configuring switch ports to limit fabric connectivity to authorized devices.

**Table 2:** Device-to-Switch Traffic Risks

## Inter-Switch Traffic

Inter-switch traffic risks are limited to issues posed by the interconnection of SAN switches in a fabric. Inter-switch traffic includes the traffic generated when switches attempt to create E\_ports (also called ISL links) as well as any traffic originating and terminating on a fabric switch. This category does not include traffic that only traverses the switches such as data from a fabric host to a JBOD storage system.

## INTER-SWITCH TRAFFIC

Risk	Description
<b>Invalid switch connected to the fabric</b>	In current SAN environments, if an invalid switch is plugged into an existing fabric it can cause significant disruptions, including modifications to zoning, access to fabric devices or resources, and a variety of other problems. Switches should not allow new ISL (E_port) connections unless the port is configured as an E_port, it allows the connection of a specific World Wide Name (WWN), and the partner switch can properly authenticate itself in a secure fashion to minimize "name spoofing," that is, pretending to be someone else.

**Table 3:** Inter-Switch Traffic Risks

## Traffic to or from Fabric Devices

This category includes risks that come from devices that do not talk directly to the switches, but instead use the fabric as the medium over which to talk to other connected devices, such as storage elements or hosts. The primary concerns here are unauthorized access and the possibility of a denial of service (DOS) attack.

#### TRAFFIC TO OR FROM FABRIC DEVICES

Risk	Description
Unauthorized access to data	If a device is able to avoid zoning limitations by physically plugging into a port or by spoofing a valid WWN, sensitive data might be accessible resulting in unauthorized changes, deletions, or data corruption.
Device Denial of Service (DOS) attack	This risk is analogous to a DOS on the fabric, however it targets a specific device or set of devices. For example, an attacker might perform an endless series of resource intensive operations on a target device.

Table 4: Traffic to or from Fabric Devices Risks

#### Conventional SAN Security Vulnerabilities

Although other types of attacks exist, many SAN vulnerabilities fit into one of the risk categories just described. The table below illustrates the risk exposure that these scenarios can introduce into enterprise security management.

#### SECURITY FUNCTIONALITY IN A CONVENTIONAL SAN DESIGN

Functionality	Description
Integrity	The fabric OS performs no significant integrity checking on configuration files.
Availability	Limited lockdown of access to the switches poses a greater risk of DOS attacks, physical access problems (such as powering the switches off), and other fabric-wide disruptions.
Confidentiality	The fabric OS uses no encryption on traffic to or from the SAN.
Authentication	The fabric OS uses weak, "clear text" username/password combination to authenticate remote users connecting to a switch and does not authenticate or validate switch-to-switch links.
Authorization	In spite of some ability to differentiate between administrator and user levels, changes made to one switch can still effect the configuration of the whole fabric.
Accounting	The fabric OS has limited ability via SNMP or Fabric Watch to manage switch parameters. Because the OS does not specifically address security, switch parameter information is limited.

Table 5: Security Functionality in a Conventional SAN Design

With an understanding of potential risks in the context of the conventional SAN environment, the next step is to understand how to minimize the risks presented by the attack types. The next section provides detailed information on the remaining two tiers: the non-SFOS-secured and SFOS-secured environments. This information is necessary to achieve a well-secured SAN that is integrated with the overall security framework.

## 2 Brocade Secure Fabric OS (SFOS)

Having covered a broad overview of enterprise security measures and conventional SAN security vulnerabilities, a more detailed description of securing the SAN environment follows. In current SAN environments, certain best practices and third-party tools are available to help minimize mistakes and prevent disruption or compromise of the fabric. Unfortunately, even a well-designed SAN using all available pre-SFOS functionality is still vulnerable to abuse. To address the shortcomings of the current environment, Brocade has introduced the Secure Fabric OS, which offers an extensive suite of security tools and enhancements to the switch operating system.

The Secure Fabric OS provides the following enhancements:

- Centralized management functions that effect the entire fabric
- Better granularity in fabric management and administration
- Allowance of only predefined, authenticated switches to log into the fabric
- Control of which devices can connect to specific ports on specific switches
- Improved remote management access security
- Support for an enhanced authentication method
- Minimized opportunity for device spoofing, incorrectly attached (rogue) devices, and Denial of Service (DOS) attacks

The previous section, “Vulnerabilities and Typical threats to a SAN Environment,” highlights the problems of an unsecured fabric in a conventional, non-secured SAN environment. The following sections, “Non-SFOS SAN Security Measures” and “Advanced Security Through the Secure Fabric OS,” present information on the next two tiers. The first of these sections provides guidance on how to lock down the current SAN environment using existing tools (such as zoning, password authentication, and so on) and a non-SFOS-enabled fabric. The second section then extends these ideas and design philosophies to building a robust SAN security environment with the SFOS.

Security architects and administrators can create security frameworks iteratively, adding new layers of protection with each pass. For example, some users might wish initially to verify that they are getting the most security out of their existing SAN configuration. Others might want to implement additional open source or proprietary tools to enhance SAN access or management while migrating toward the SFOS. Still others might want to maximize their SAN security by implementing The Brocade Secure Fabric OS. These steps complement one another and can progress toward a tightly secured, sophisticated SFOS SAN environment.

### 2.1 NON-SFOS SAN SECURITY MEASURES

A number of measures are available to minimize problems and enhance the security of the SAN environment before progressing to an SFOS implementation. Some measures leverage existing functionality within the Fabric OS, while other steps involve third-party products or specific network designs. The topics below describe these security options.

#### 2.1.1 Zoning

In existing SAN infrastructures, administrators can decide what resources should be reachable from which devices and then use zoning to create these relationships. Zoning is a licensable feature that allows for the logical grouping of devices connected to a fabric. It works in much the same way as a limited Access Control List (ACL): zoning

regulates traffic between devices based on whether the source and destination systems are in the same zone. If a device, such as a JBOD or a tape backup system, is not in the same zone as a host, that host cannot access the device or the resources on it. Here are some of the reasons that a SAN administrator might use zoning:

- **Mixed Operating System Environment:** If storage devices are not properly segmented from devices running different operating systems, a Windows NT host could accidentally corrupt a disk that is in use by a UNIX host.
- **Restricting Resource Access:** Because typically not all data should be readable by all users or systems, an administrator might want to partition off sensitive data to specific hosts in order to increase security.
- **Better Resource Granularity:** A single multidisk device can have its disks configured in separate zones. Thus, a single JBOD could be selectively storing and providing data to multiple hosts.

## Zoning Types

Brocade offers three zoning types, each with its respective pros and cons. These include soft zoning, hard zoning, and broadcast zoning.

- **Soft Zoning** uses the name server to enforce zoning. The World Wide Name (WWN) of the elements enforces the configuration policy.

### *Pros*

- Administrators can move devices to different switch ports without manually reconfiguring zoning.

### *Cons*

- Devices might be able to spoof the WWN and access otherwise restricted resources.
  - Device WWN changes, such as the installation of a new Host Bus Adapter (HBA) card, require policy modifications.
  - Because the switch does not control data transfers, it cannot prevent incompatible HBA devices from bypassing the Name Server and talking directly to hosts.
- **Hard Zoning** uses the physical fabric port number of a switch to create zones and enforce the policy.

### *Pros*

- This system is easier to create and manage than a long list of element WWNs.
- Switch hardware enforces data transfers and ensures that no traffic goes between unauthorized zone members.
- Hard zoning provides stronger enforcement of the policy (assuming physical security on the switch is well established).

### *Cons*

- Moving devices to different switch ports requires policy modifications.
- **Broadcast Zoning** has many unique characteristics:
    - This traffic allows only one broadcast zone per fabric.
    - It isolates broadcast traffic.
    - It is hardware-enforced.



## Zone Enforcement

Zones are either hardware-enforced or software-enforced depending on the zoning scheme. When *either* port number or WWN—but not both—specify a zone, the zone is a hardware-enforced zone. When *both* port number and WWN specify a zone, it is a software-enforced zone.

Hardware-enforced zoning is enforced at the Name Server level and in the ASIC. Each ASIC maintains a list of source port IDs that have permission to access any of the ports on that ASIC. Software-enforced zoning is exclusively enforced through selective information presented to end nodes through the fabric Simple Name Sever (SNS).

- ❖ Note: Brocade Silkworm 12000, 3900, 3800, and 3200 fabric switches enable port-level and WWN hardware enforcement.

## Zoning-Specific Security Issues

Even when it is well implemented, zoning can still expose the SAN to exploitable risks. These include physical access, zone merging, and WWN spoofing issues, which are addressed below:

- **Physical Access:** Physical control is absolutely essential. Attackers can easily bypass hard zoning and access isolated resources if they have the ability to change the switch cable physically from one device to another that they control. Further, if it is possible to plug one switch into another and successfully join it to the fabric, a variety of other problems can occur.
- **Zone Merging:** Because it is possible to integrate new switches and their respective zoning configurations into an existing fabric, a potential security issue exists. If any switch in the fabric clears the zoning configuration, it eliminates zoning for the entire fabric once the information propagates to the other switches. This can occur by accident or by design if an administrator account password has been compromised, or if a rogue element has inserted a switch into the fabric.
- **WWN Spoofing:** Address and name spoofing are common attack methods in the IP world. In the SAN environment, WWN spoofing enables a device to masquerade as a different entity and thus bypass soft zoning. There are a variety of ways to capitalize on this situation, but many of them require physical access to the fabric or remote access to a fabric device.

### 2.1.2 Locking Down E\_Ports

The high risk of the unauthorized or accidental addition of a switch to a fabric calls for software controls that augment existing physical security measures. One recommended way to limit the addition of a new switch to the fabric is to limit the creation of switch-to-switch ports. Connecting together any two switches and successfully negotiating parameters between them forms an Inter-Switch Link (ISL), which is represented as an E\_port. Using the `portCfgEport` command precludes this kind of occurrence. Be aware that such locking down of E\_ports is persistent across reboots. The example below shows how to prevent port 3 from becoming an E\_port on the switch using `portCfgEport`.

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

Figure 3: portCfgEport Example

One way to enforce proper change control steps and minimize accidental or intentional misuse from “rogue” switches is to disable (mode 0) all the ports on the switch that are not intended to be ISL ports. As new switches are authorized for the production environment, one of the steps in integrating them into the fabric is to re-enable the proper port(s) in the fabric to support ISL.

Using similar logic, the *portDisable* command disables unused ports on the switch to prevent any device from being plugged into them and connecting to the fabric.

- ❖ **Important Note:** With Fabric OS version 2.6.0 and earlier, the effects of *portDisable* and *portEnable* are not persistent across reboots. Upon switch reboot, the ports become available once again.
- ❖ **Important Note:** Beginning with Fabric OS versions 2.6.1, 3.1, and 4.1, the effects of *portDisable* and *portEnable* are persistent across reboots. Administrators can also use the *portcfgpersistentdisable* command to disable a port so that the effect is persistent across reboots.

### 2.1.3 Physical Access

Physical security is an absolutely essential component of any comprehensive enterprise security plan. Even with excellent software controls in place, physical access to enterprise elements opens the door to a whole range of security issues. To ensure physical security, fabric devices should reside in environments where physical access controls provide adequate protection.

### 2.1.4 Remote Access

There are a variety of ways to obtain information from fabric switches. Common management access methods involve the use of telnet for command line functionality, HTTP for Web-based access, in-band Fiber Channel for management server access, and console access for direct switch connectivity. Each of these access methods carries associated security issues.

- **Telnet:** The essential problem with telnet access is that it transmits “in the clear” the username, password, and all data going between the management system and the switch. Any user with a “promiscuous” Network Interface Card (NIC) and “data sniffing” programs can capture the account and password data. If the strength of an organization’s SAN security lies solely with the administrator username and password, then the fabric is vulnerable to complete compromise.
  - ❖ **Note:** With Fabric OS 4.1 or greater, administrators can disable telnet using the *configure* command.

- **HTTP:** Similar to the telnet issues mentioned above, when a system uses a Web-based application such as Brocade WEB TOOLS to authenticate to the switch in order to run privileged commands, it passes the login information “in the clear.”
- **Management Server:** This remote management method uses an in-band Fibre Channel connection to administer or obtain information from the fabric switches. By default, this method grants access to any device. However, it is possible to create an access control list to limit the WWNs of devices that can connect to the switch using this method. Administrators can use the *msConfigure* command to display information and configure the management server.
  - ❖ **Note:** Administrators can address some of the issues described above using policies (such as API, SNMP, and HTTP policies). Additional information on setting up policies is available in the Brocade Fabric OS Reference Manual.
- **Console Access:** Although not usually thought of for remote access, it is possible to adapt console connections to remote use through the use of terminal server devices. Thus, an organization can use telnet, secure shell (SSH), or a similar application to connect to the terminal server, which then in turn connects to the selected device through the console interface. This solution has the potential to provide additional security through the use of third-party products.
- **SSH:** With Fabric OS 4.1 or greater, clients running version 2 of the Secure Shell protocol can also access Brocade fabric switches.

Different designs and tools are available to help secure remote access. Though none of these solutions is as convenient as being able to provide an end-to-end encrypted password directly to the switch, they do help mitigate some of the common risks of clear text password transmissions. The example below illustrates some of the methods for providing more secure remote SAN management.

## Example: Minimizing Remote Access Password Exposure

### **Situation:**

An organization currently uses telnet to access the SAN environment and implement configuration changes. SAN administrators reside on the same Local Area Networks (LANs) as other non-privileged users. Administrators are concerned that the clear text password used by the management applications could be picked up by a packet analyzer, and thereby compromising the SAN environment. What would help minimize this problem?

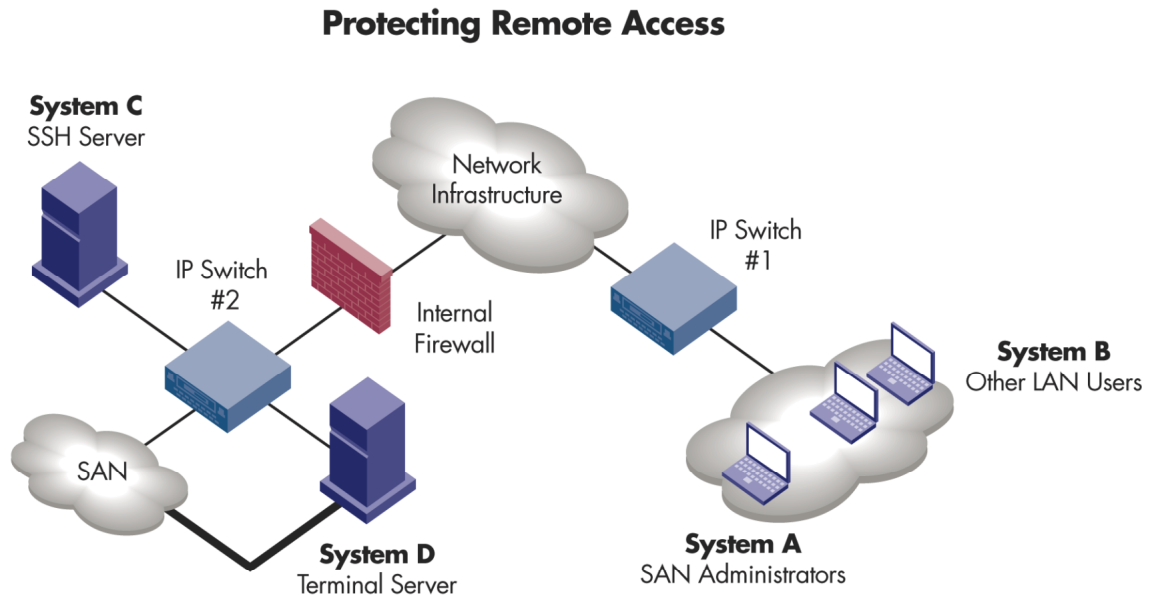


Figure 4: Remote Access – A Sample Environment

### **Discussion:**

There are a variety of ways to address this concern and protect administrative access to the SAN. A combination of tools and re-architecting can provide better data security. Here are a few recommendations:

- In order to address the telnet concern, the administrators (shown as System A in the figure above) could connect to an intermediate system (System C) using Secure Shell (SSH). SSH fully encrypts the data stream, including passwords, between the source and destination. Once connected to this system, the administrator can access the SAN, which is connected to the same switch via telnet. The internal firewall, which could be a router with ACLs, can deny normal telnet sessions and limit the source IP addresses of the systems that are allowed to initiate SSH sessions to the SSH server. This combination provides the following:

#### *Pros*

- SSH encrypts all connection data traveling through the local LAN (Switch 1) and the network infrastructure from System A to System C. This prevents data sniffing from other local LAN users (System B) and from any systems on the network infrastructure.
- The internal firewall can lock down access to devices on Switch #2. This blocks normal telnet attempts and limits SSH access to known SAN administrator source IP addresses.

- From System C, it is possible to establish telnet access to the SAN environment. The proximity of these systems and the relative isolation helps limit the exposure of the data.

*Cons*

- This solution does not support end-to-end encryption of passwords. The SAN login and password still travel in the clear on Switch #2 from the SSH server to the SAN switches. This exposure is significantly lower, however, than having the passwords traverse the entire enterprise.
- This solution might require additional planning and hardware to segment the systems properly from the rest of the network.
- A second way to address the telnet concern is to have the administrators (System A in the figure above) connect by way of SSH to a terminal server (System D) and connect the system to the console ports on the SAN switches.

*Pros*

- SSH encrypts all connection data traveling through the local LAN, Switch 1, and the network infrastructure from System A to System C. This greatly reduces the effectiveness of any data sniffing attempts.
- The internal firewall can lock down access to devices on Switch #2, blocking normal telnet attempts and limiting SSH access to known SAN administrator source IP addresses.
- Console cable access keeps the login information from having to travel in the clear over Switch #2.

*Cons*

- The SAN switches must be close enough to the terminal server to connect by cable.
- Any SANs using Brocade 2800 systems would not fit this model because they use front panel access instead of console access.

## 2.1.5 Security Summary

### Risks to a Non-SFOS-Configured SAN

The information presented in the previous section discusses important security methods that are available to protect the SAN environment. It is useful to readdress the list of potential SAN threats discussed earlier in this document in order to review how each of these methods helps to deter different types of attacks. The table below shows to what extent the current tool set of zoning, securing E\_ports, physical access controls, and remote access can secure the SAN environment.

#### NON-SFOS SAN SECURITY VERSUS COMMON SAN RISKS

Device-to-Switch Traffic	
Risk	Description
An unauthorized, local (non-switch) device is connected to the fabric or a device is connected to or from the wrong port	A combination of physical access controls and disabling of unused ports can greatly reduce the likelihood of many of these attacks. Physical access controls limit unauthorized personnel from moving or switching cables. WWN spoofing is still an issue as it involves reconfiguring the name of an existing device and does not involve physically moving cables. Soft zoning would remain vulnerable in this case.
Invalid management device connection attempt	Current SAN implementations do not allow for device authentication based on source IP address. Thus, there is no way to create a group or to specify authorized or unauthorized management systems. It's possible to implement this control through other components.
Unprotected management application authentication information	The fabric switches rely on protecting the integrity of the administrator username and password combination, which can be difficult to implement with a clear-text-based remote access method. Network design and third-party software can help to mitigate this risk.
Device Denial of Service (DOS) attack	The best way to stop intentional DOS attacks on the SAN is to limit the connection of new devices to the fabric by physical security, port lockdown measures, and carefully created zones. Additionally, devices currently connected to the fabric should be configured and protected to minimize compromise or unauthorized use. Although it is very difficult to protect against malicious authorized users, policies and procedures should be in place to identify abuses and provide appropriate disciplinary action.
Inter-Switch Traffic	
Invalid switch connected to the fabric	The best way to prevent this situation with the current tool set is by careful planning, physical security, and the disabling of E_ports on all unused ports. Because these switches do not have the functionality to authenticate other switches or their WWNs, there is no way to prevent the replacement of one switch with another unauthorized switch.
Traffic to and from Fabric Devices	
Unauthorized access to data	Because this tool set does not allow for the lockdown of a port to a specific WWN, name spoofing is still a problem.
Device Denial of Service (DOS) attack	Protecting against this risk is analogous to a DOS attempt on the fabric itself.

Table 6: Non-SFOS SAN Security Versus Common SAN Risks

### Non-SFOS SAN Security Functionality

The Non-SFOS secured environment provides an improved security environment over conventional SAN solutions. Even the flexibility provided by the Brocade Fabric OS still leaves room for improvement and better leverage of security functionality. The table below summarizes the security enhancements provided by some non-SFOS techniques:

## SECURITY FUNCTIONALITY IN A NON-SFOS SAN DESIGN

Functionality	Description
Integrity	The fabric OS performs no significant integrity checking on configuration files.
Availability	Preventing physical access and locking down ports will help minimize accidental cable pulls, unintentional powering off of fabric devices, and DOS attacks that result from the connection of unauthorized devices connecting to the fabric.
Confidentiality	The fabric OS still provides no encryption, however SSH can help minimize data exposure by providing encryption for a portion of the data path to the switches.
Authentication	Much like in the conventional SAN, the fabric OS still uses a weak, clear text username/password combination to authenticate remote users connecting to a switch and does not authenticate or validate switch-to-switch links.
Authorization	In spite of some ability to differentiate between administrator and user levels, changes made to one switch can still effect the configuration of the whole fabric.
Accounting	Much like in the conventional SAN, the fabric OS has limited ability via SNMP or Fabric Watch to manage switch parameters. Because the OS does not specifically address or incorporate security, switch parameter information is limited.

**Table 7:** Security Functionality in a Non-SFOS SAN Design

Recall that defense in depth is an essential element of a comprehensive security strategy. If one of the defenses fails, additional layers of protection secure important resources. While the non-SFOS SAN environment still presents some serious security flaws due to its limited tool set, the use of certain tools and configurations does help to protect the environment. The information below lists non-SFOS problems that are addressed by SFOS capabilities:

- No centralized management functions: any switch has the potential to create a fabric-wide configuration change
- Requires device ACLs or a firewall to limit management traffic to the switches
- No end-to-end switch password or session encryption
- No way to authenticate switches before allowing them to log into the fabric
- No way to limit which devices can connect to specific ports on specific switches

Proper lockdown the fabric requires true enterprise-strength SAN security. The functionality provided by The Brocade Secure Fabric OS not only fills the existing gaps in SAN security measures, it also provides additional security features, configuration granularity, and a consistent framework to extend enterprise SAN security. The next section details SFOS functionality and illustrates how it fills in the gaps left by the other SAN security implementations.

## 2.2 ADVANCED SECURITY THROUGH THE SECURE FABRIC OS

The SAN environment has enjoyed a relatively secure existence due to the comparative inaccessibility and isolation of older SANs. Recently, however, the SAN environment has evolved. Its role in the enterprise has changed and it has taken on ever-increasing importance in delivering mission critical business functionality. These changes call for appropriate security controls that are tailored to meet the unique challenges of a secure SAN environment. To this end, Brocade has enhanced their Fabric OS with significant security improvements and more consistent, centralized fabric management. This enhanced version product is called the Brocade Secure Fabric OS.

This guide has emphasized how integral each aspect of the security infrastructure is to the other component parts. Proper SAN security relies on proper operational, technical, and management controls to limit access and provide additional layers of protection. Even with these defenses, the SAN should be able to support security controls

directly. Clearly, traditional SAN measures still present significant issues. The inability to lock switch ports down to specific WWNs or to identify authorized management stations by their IP addresses, the lack of a strong switch-to-switch authentication method, and the risks of sending access passwords “in the clear” are all critical issues. Fortunately, the Secure Fabric OS not only mitigates these vulnerabilities, but also provides the SAN administrator with a whole host of additional functionality and configuration granularity that currently does not exist in the environment.

### 2.2.1 Secure Fabric OS Overview

What is the SFOS? It is The Brocade solution for comprehensive, enterprise-strength SAN security. This section illustrates the specific security enhancements and prerequisites of the SFOS in detail.

In order to use the SFOS functionality, administrators must enable secure mode. Before doing so, the following steps are required:

- **Ensure that all switches in the fabric are Brocade switches.**
  - ❖ **Note:** Brocade 1000 Series SilkWorm switches are not security capable and cannot participate in a fabric that will be secure-mode-enabled
- **Install the Secure Fabric OS on all switches in the fabric.**
  - ❖ **Note:** Fabric OS versions 2.6.1, 3.1, and 4.1 or greater support Secure Fabric OS functionality.
- **Purchase the security feature and install the licenses for security and zoning on all switches in the fabric.**
- **Obtain an appropriate digital certificate for all switches. (Refer to the following sections for more information.)**

Once all these conditions have been met and secure mode is enabled, the system is now a legitimate SFOS fabric.

- ❖ **Note:** For a more detailed list of steps, consult the Brocade Secure Fabric OS User Guide.

Secure Fabric OS functionality falls into five basic categories:

- **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 what 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\_ports..
- **Device Connection Control (DCC)** allows only specific devices into the fabric (per their WWNs) from a specific port or group of ports.

Much of the detailed information provided in this section illustrates the benefits of the SFOS explains how the security features work. Although many of the security mechanisms are transparent to the administrator, it is helpful to understand how the product functions. In addition, the best practice recommendations in the following sections emphasize guidelines or design suggestions on how to implement SFOS features.



## 2.2.2 Fabric Configuration Server (FCS)

Most infrastructure designers and administrators consider a large, decentralized management environment to be undesirable. Complex management structures often lead to inconsistent policy implementation, password management issues, and the possibility of applying simultaneous and conflicting configurations that adversely affect the fabric. The SFOS tackles the issue of centralized management by creating a multi-tiered switch configuration infrastructure.

### Switch Type Overview

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. 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.

### Secure Fabric OS Switch Hierarchy

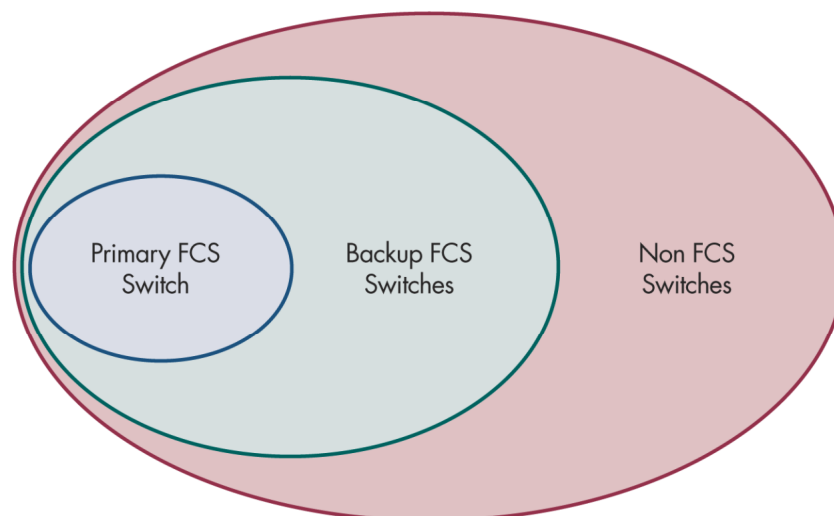


Figure 5: Switch Type Overview – Three Switch Types

## FCS Switch Functionality

Understanding the power of the Primary FCS Switch requires a clear definition of fabric-wide configuration parameters. To maintain fabric effectiveness, the Primary FCS Switch is solely responsible for applying configurations consistently to all switches. The SFOS architecture centrally manages four databases, each of which is discussed at greater length in the following sections of this chapter. Here is a preview:

- **Fabric Management Policy Set (FMPS):** The FMPS database includes five sub-groups.
  - **FCS Policy:** This sub-group contains the WWNs of all the FCS switches and requires a non-empty group (that is, the policy configuration cannot be empty) for the operation of the fabric in secure mode. This group tells all the switches in the fabric which switch is the Primary FCS and the priority order of the Backup FCS Switches.
  - **Management Access Control (MAC):** This information defines which devices can access the fabric in what manner to perform management functions. (Detailed information on the MAC is in a later section.)
  - **Device Connection Control (DCC):** DCC policies determine what device WWNs can insert into the fabric and to which switch ports. (Detailed information on the DCC is in a later section.)
  - **Switch Connection Control (SCC):** Only the WWNs of the switches in this sub-group can participate in the fabric. (Detailed information on the SCC is in a later section.)
  - **Options:** This sub-group allows for future functionality. Today, it prevents WWN-based zones through the NoNodeWWNZoning function. By default, this option is turned off. This enables the creation of WWN-based zones (soft zoning) on devices that support it.
- **Zoning Configuration:** The creation, modification, and application of changes to the zoning configuration policy must take place on the Primary FCS Switch. This is similar to traditional zoning except that changes can only come from the Primary FCS Switch. This eliminates the possibility of plugging in a new switch and wiping out or changing zoning in a fabric.
- **Passwords:** The new secure fabrics have two tiers of passwords: FCS and non-FCS. These password groups are consistent for all switch members across the fabric.
- **Community Strings:** The Primary FCS Switch is solely responsible for managing these community string values consistently across the fabric.

## Supplemental Information on the FCS Group and Passwords

With the exception of the FCS group and password scheme, each of the items just described are either very straightforward or they are further described in later sections. While FCS policy and passwords are not unduly complicated, they do deserve some special consideration.

### FCS Policy

In order to be an FCS device, the WWN of each FCS candidate must be on the FCS\_POLICY list. This list reflects the FCS priority order from top to bottom. The first device on the list represents the Primary FCS Switch and any subsequent devices are Backup FCS Switches. If the Primary FCS Switch is not available, the next switch in the list takes over. If none of the switches in the FCS\_POLICY list are available to the fabric, then no fabric-wide changes can occur until an FCS device becomes active and resumes control.

### FCS Policy Recommendations

The actual number of FCS switches can vary from fabric to fabric. Having more than one FCS switch in the fabric offers significant advantages. Because a valid Primary FCS Switch is required to implement any fabric-wide changes,

having one or more Backup FCS Switches increases fault tolerance provides options when the Primary FCS Switch is undergoing maintenance or troubleshooting. No absolute guidelines govern the number of FCS switches in a fabric. However, because these devices have elevated privilege levels on the fabric, the following guidelines apply:

- All FCS switches should be physically well secured and controlled. A switch at a small remote office is typically not an ideal FCS member.
- Choosing switches from different data centers or geographic locations provides better control options in case of a power outage at a site or in the event of a catastrophic event.
- Due to the centralized policy management of the SFOS, some large fabrics might benefit from having a Primary FCS Switch that is located near the other switches.

**Passwords**

In secure mode, the SFOS has two tiers of passwords and five actual password values. The SFOS makes password distinctions between FCS and Non-FCS switch tiers. The individual accounts break out as follows:

- FCS switch accounts include:
  - **Admin (FCS level):** Provides administrative access to the switch and enables full fabric-wide changes when connected to the Primary FCS Switch
  - **User:** Provides limited access to the switch
  - **Root:** An account reserved for special troubleshooting operations on the switch
  - **Factory:** An account reserved for special servicing of the switch
- Non-FCS switches have only two options:
  - **Admin (non-FCS):** Allows full switch-specific administration. Fabric-wide operations are available only from the Primary FCS Switch.
  - **User:** Provides limited access to the switch. This is the same account password as the FCS switch user account.
    - ❖ Note: On non-FCS switches, the Factory and Root accounts are disabled and not available. However, they can be temporarily activated on non-FCS switches for maintenance purposes.

The following diagram illustrates the password structure:

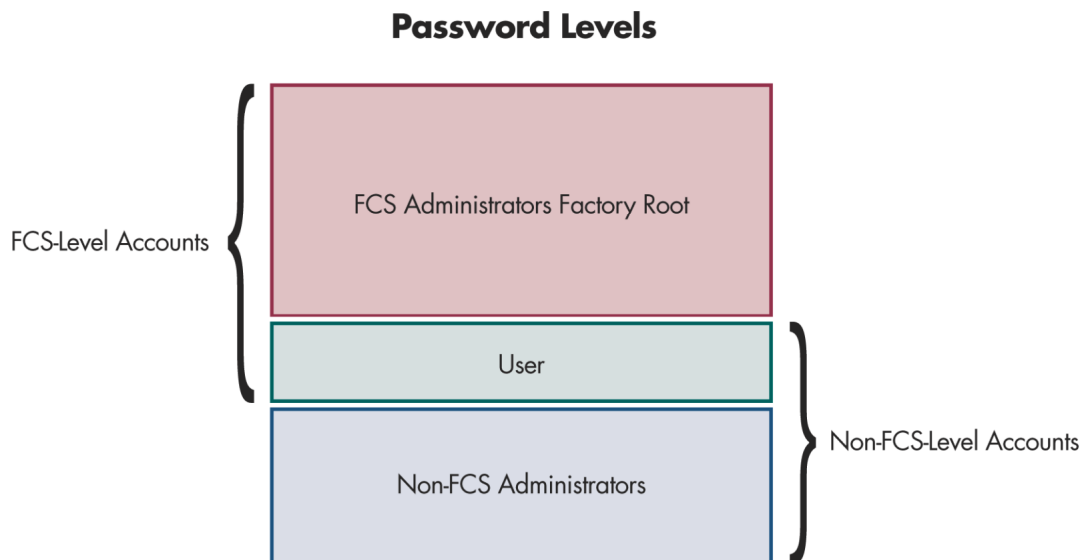


Figure 6: Passwords – A Visual Representation

The following are notable facts regarding passwords:

- The Primary FCS Switch manages the password database, which remains consistent across the fabric.
- It is possible to assign a temporary password to a switch. This is useful when a switch requires maintenance and the administrator does not want to divulge the fabric-wide password. This temporary password is not persistent across reboots.
- Passwords must be at least eight characters long to ensure adequate password complexity.

**Recommendations:**

The best time to reset all fabric-wide passwords is when secure mode is enabled in the fabric. Because all SFOS security functionality is dependent upon having secure mode enabled, SAN administrators run the *secModeEnable* command the first time they want to access additional account granularity. The following steps and the figure below show how to initiate this process:

1. Connect to the Primary FCS Switch and using Secure Telnet and log in as “admin”.
2. Run the *secModeEnable* command to initiate the process.
3. Enter the WWN of the FCS switches (Primary FCS Switch first). Because this example is just a two-switch SAN, it designates only one FCS switch.
4. When prompted, securely change the passwords for the Root, Factory, FCS Switch Admin, Fabric-wide User, and Non-FCS Switch Admin passwords.
5. Note: As a safeguard, the new passwords must differ from the old pre-secure-mode passwords by at least one position.
6. Upon completing these steps, save the information. Upon reboot, the switch comes back up in secure mode.

```

192.168.149.29 - sTelnet
Sessions Edit Terminal Help
Fabricshow
Switch ID   Worldwide Name           Enet IP Addr   FC IP Addr     Name
-----
1: fffc01 10:00:00:60:69:10:69:78 192.168.149.28 0.0.0.0        >"SecureSAN28"
2: fffc02 10:00:00:60:69:11:F9:94 192.168.149.29 0.0.0.0        "SecureSAN29"

The Fabric has 2 switches

SecureSAN29:admin> secModeEnable 1

This is an interactive session to create a FCS list.

The new FCS list is empty.

Enter WWN, Domain, or switch name(Leave blank when done): 10:00:00:60:69:11:F9:94 2
New Switch WWN is 10:00:00:60:69:11:F9:94.

The new FCS list:
 10:00:00:60:69:11:F9:94

Enter WWN, Domain, or switch name(Leave blank when done):
Are you done? (yes, y, no, n): [no] yes
Is the new FCS list correct? (yes, y, no, n): [no] yes
Each encryption/decryption of password takes a while
New FCS switch root password: 3
Re-enter new password:
New FCS switch factory password:
Re-enter new password:
New FCS switch admin password:
Password must differ by at least 1 position 4
New FCS switch admin password:
Re-enter new password:
New fabric wide user password:
Re-enter new password:
New Non FCS switch admin password:
Re-enter new password:
Saving passwd...done.
Saving Defined FMPS ...
done
Saving Active FMPS ...

```

Figure 7: Resetting Passwords During secModeEnable

Proper password management is a very important part of the overall security framework. Only those personnel that require certain login information should receive password information. The following steps are guidelines for securing and managing passwords:

- Early SAN security systems transmitted fabric passwords “in the clear.” As a result, administrators should assume that they have been captured and compromised. The *secModeEnable* process sends passwords only in encrypted form. With a newly-secured fabric, changing the passwords to new unique values at the earliest opportunity limits exposure.
- It is important to change the passwords for all five accounts to strong, unique values.
- The recovery of lost passwords requires a very detailed recovery procedure. Administrators should carefully record all passwords – out of view of other personnel – and store them in a secure location, such as a locked safe. The idea of writing critical passwords down is anathema to most security-minded individuals, so the process requires strict controls, such as the following:
  - The safe containing the passwords requires two different keys, much like a bank safety deposit box, with each key assigned to a different authorized individual.
  - The organization’s security policy lists conditions and procedures for allowing access to the passwords.
  - Administrators monitor and log all access to the passwords in a book that resides in the same safe as the passwords.

- Change control procedures for changing the passwords are in place for when individuals leave the company or in the event of password rotation.
- This enables the organization to control sharing of the new passwords in a secure fashion among only the individuals that require them.

### Time Server

When numerous security tools and implementations are in play, it is a good idea to have consistent time service on all systems. The Brocade Secure Fabric OS uses a message format similar to Simple Network Time Protocol (SNTP) v.4 to synchronize all switches in the fabric with the Primary FCS Switch. Two important notes on time server functionality and usage are:

- Each timestamp includes a signed hash of the time update, which supports data integrity and sender authentication.
- Database updates from the Primary FCS Switch use time stamps to minimize potential replay attacks on the fabric.

### Database Management and Propagation

Any change to one or more of the databases on the Primary FCS Switch prompts the switch to send newly updated files to all switches in the fabric. These updates ensure the enforcement of a single, consistent policy on all devices. This also ensures that only one device at a time that can edit or update the policies.

The Primary FCS Switch signs every database before distributing to the fabric. Here is how the SFOS handles policy modifications on the Primary FCS Switch:

- When the administrator saves changes to the Defined Policy set, the SFOS retains the information but does not apply it to the fabric.
- When the SAN administrator chooses to apply the changes, thereby incorporating them into the Active Policy set, the SFOS saves the configuration into flash memory; this enables them to persist through a reboot or power loss.

The figure below shows the information available in the secModeShow view. This example shows the Primary FCS Switch (SecureSAN28) in secure mode, and one Backup FCS Switch (SecureSAN\_36), the version stamp (28592), the time stamp, and other information.

```
SecureSAN28:admin> secModeShow
Secure Mode: ENABLED.
Version Stamp: 28592, Tue Nov 13 10:23:08 2001.
Pos   Primary  WWN                               DID  swName.
-----
  1   Yes     10:00:00:60:69:10:69:78          1   SecureSAN28.
  2   No      10:00:00:60:69:10:97:83          5   SecureSAN_36.
SecureSAN28:admin>
```

Figure 8: secModeShow Information

- ❖ **CAUTION:** Brocade strongly recommends against organizations allowing proxy server access to the secure fabric. When a proxy server is included in a security MAC policy for IP-based management, such as HTTP\_POLICY, all IP packets that leave the proxy server will appear to originate from the proxy server. This could allow unwanted hosts that have access to the proxy server to access the secure fabric.

### 2.2.3 Management Access Controls (MAC)

System designers and administrators must strike a balance between flexibility and security when locking down the enterprise environment. Because management access is the primary means of making configuration changes, it represents a significant risk to any mission critical system. On the other hand, locking down change control to local command line interface (CLI) access would only eliminate the very useful ability to effect change remotely or to use and GUI-based management tools. The optimal balance is likely a blend of carefully limited management access that is consistent with an organization's acceptable risk threshold.

The Management Access Control (MAC) facet of the SFOS enables the fabric administrator to choose selectively how to manage the SAN. The MAC targets three broad categories, each with many sub-categories:

- **Remote Access Limitations:** Compares the source IP address of the remote connecting device to the respective policy.
- **Port-Based Access:** Uses the device WWN of the requesting system to compare against the respective policy.
- **Physical Access Connections:** Uses the switch WWN of the requesting switch to compare against the respective policy.

The tables below provide specific information about the policies in these three categories.

#### REMOTE ACCESS LIMITATIONS

Access Method	Policy Name	Description
Telnet	TELNET_POLICY	Limits secure telnet (sectelnet) client access to the fabric (refer to the Secured Management Channel section below)
HTTP	HTTP_POLICY	Limits Web client access to the fabric
RSNMP	RSNMP_POLICY	Limits SNMP read access to the fabric
WSNMP	WSNMP_POLICY	Limits SNMP write access to the fabric
API	API_POLICY	Limits access to the Application Programming Interface (API) from remote sources

Table 8: Remote Access Limitations

#### PORT-BASED ACCESS

Access Method	Policy Name	Description
SCSI Enclosure Services (SES)	SES_POLICY	Lists device WWNs that can access SES services
Management Server (MS)	MS_POLICY	Lists of acceptable device WWNs that are allowed to manage the switches in band

Table 9: Port Based Access

#### PHYSICAL ACCESS CONNECTIONS

Access Method	Policy Name	Description
Serial Connection	SERIAL_POLICY	Enables or disables the serial port on switches in the fabric
Front Panel Access	FRONTPANEL_POLICY	Enables or disables the front panel on switches in the fabric that have front panel access (such as the SilkWorm 2800)

Table 10: Physical Access Connections

**Recommendations:**

Administrators can take a few different approaches when implementing MAC policies. The extent of SAN lockdown can vary according to the risk tolerance and business needs of the organization. Because the policies are very permissive by default, locking down each access control as tightly as possible maximizes the security benefits. In this case, the security principle of “least privilege” applies: an entity should have only the privilege (or functionality) needed to perform its authorized tasks. For example, if an organization does not use HTTP to manage switches, then the administrator should create the HTTP\_POLICY with no entries. Once applied, this configuration will block all incoming Web access.

The following are design guidelines for implementing MAC policies:

- Remote access policies should be as explicit as possible and should allow only the static IP addresses of known administrative consoles. Including whole subnets in the remote access policies makes the group more permissive and potentially less secure.
- If an organization knows that it does not use some access method to the switches, then the administrator should modify the appropriate policy to limit or block that kind of access.
- When creating policies that use WWNs through sectelnet, the following steps help to minimize typos:
  - On the Primary FCS Switch, type *fabricShow* to see information about the switches in the fabric.
  - After identifying the proper name, hold down the left mouse button and drag the pointer to highlight the WWN.
  - Click the right mouse button to paste the highlighted information into a command.
- Test any policy created in order to verify that it works as intended.
- Exercise caution. Configuring policies incorrectly can block all access to the switch. If this happens, refer to the “Recovery Processes” section of the Secure Fabric OS User Guide.
  - ❖ **Caution:** Running *secmodedisable* without backing up policies can cause the loss of policies when Secure Mode is disabled. Be sure to back up policies when making changes.
  - ❖ **Note:** Refer to the Secure Fabric OS User Guide for more specific information on MAC policy states.

MAC implementation is very much like a traditional Access Control List (ACL). When a device attempts to make a management connection, the SFOS refers to fabric policy and determines whether to allow or deny access. The following examples help illustrate the creation and function of these policies:

### Example 1: Telnet and HTTP Policies

**Situation:**

The organization has just added a third SAN administrator (shown as System C in the figure below) to the team and plans to grant remote command line and HTTP access to this person. The organization wants to limit remote access to the SAN in order to minimize potential security holes. What steps enable this functionality?



## Telnet and HTTP Policies

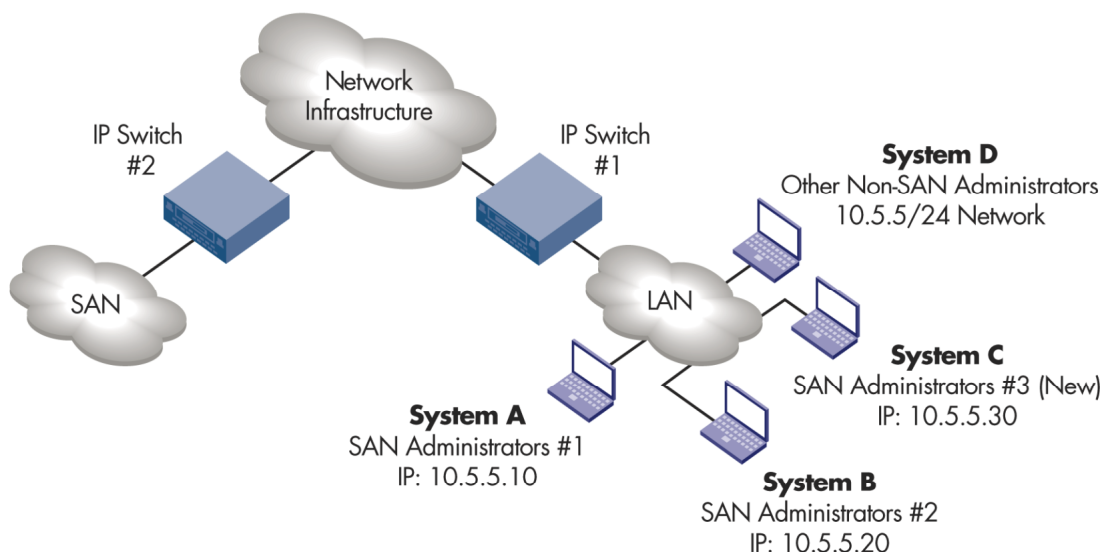


Figure 9: Telnet and HTTP Policy Example

### Discussion:

Because the client has a well-defined understanding of which systems should be accessing the fabric remotely, it is possible to make the access list for the TELNET\_POLICY and HTTP\_POLICY quite restrictive. If the SAN administrator systems use DHCP, they must be tied down to specific IP addresses to ensure constant proper performance. Locking down IP addresses also minimizes the risk posed by having a non-SAN administrator's system obtain an IP address that allows remote access via the TELNET\_POLICY, or by a similar scenario.

The table below shows the IP addresses of authorized workstations.

Policy	Entries
TELNET_POLICY	10.5.5.10 10.5.5.20 10.5.5.30 <- NEW
HTTP_POLICY	10.5.5.10 10.5.5.20 10.5.5.30 <- NEW

Table 11: IP Addresses of Authorized Workstations

In this case, the administrators could have added the whole 10.5.5.0/24 subnet to the policies. The SFOS commands allow certain additions of subnets as well as specific IP addresses. Although this might have been expedient, it is not a good security practice. Using the "least privilege" methodology, a more restrictive and less permissive rule set is the better.

The new policies shown in the table tightly restrict telnet and HTTP access to the fabric to only three source IP addresses. One potential risk that still exists in this environment is the possibility that one of the three valid IP

addresses could be spoofed. Although spoofing would foil the policies that test connections against the source IP of the device, it would be a difficult attack to implement. Remember that the power of Defense in depth relies on multiple layers of security. An attacker would still need the fabric's authentication information in order to effect change, and when the fabric is in secure mode, the passwords only traverse the network in encrypted form.

**SFOS Commands:**

The figure below shows an example of how to modify the HTTP\_POLICY to add the new administrator in this situation.

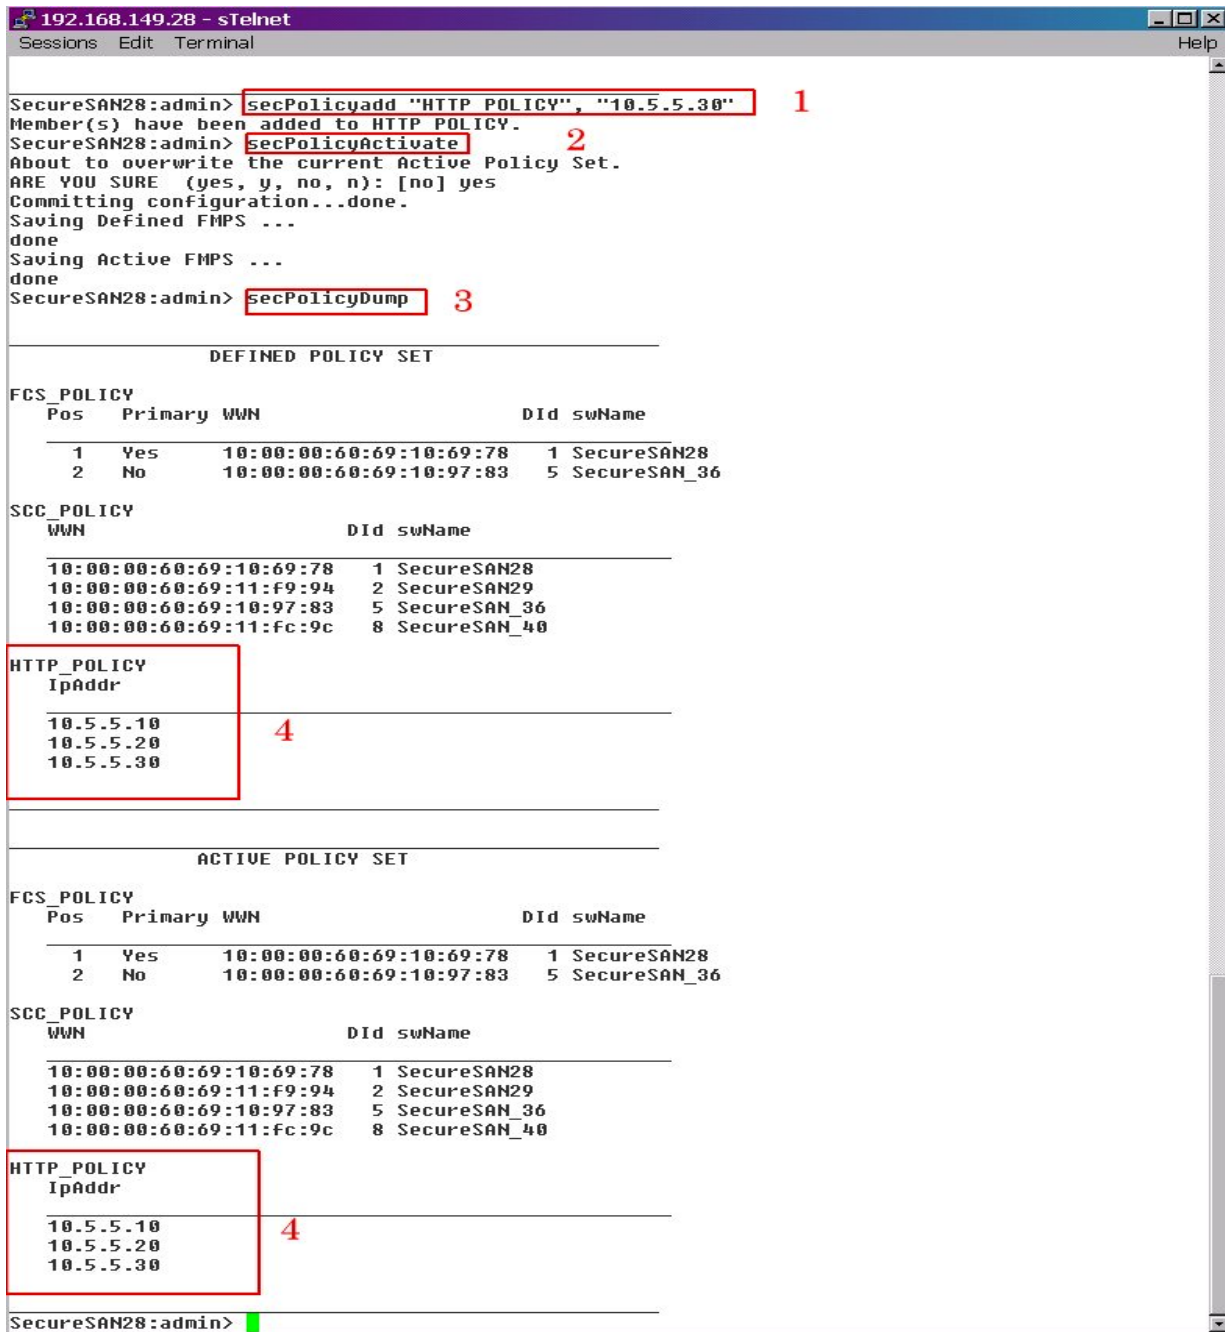


Figure 10: HTTP\_POLICY Creation Example

The process for modifying the HTTP\_POLICY is as follows:

- ❖ Note: The HTTP\_POLICY already exists and contains the 10.5.5.10 and 10.5.5.20 IP addresses.
- The *secPolicyAdd* command modifies the HTTP\_POLICY to add the new 10.5.5.30 administrator.
- The *secPolicyActivate* command saves and propagates the changes throughout the fabric.
- The *secPolicyDump* command produces a view of the current policies.
- The HTTP\_POLICY now contains three admin IP addresses and since it has been activated, the changes appear in both the Defined and Active policy sets.
  - ❖ Note: Refer to the Secure Fabric OS User Guide to obtain additional information on the commands and their functionality.

### Example 2: API Policy

**Situation:**

The same company as shown in Example 1 has just installed Brocade Fabric Manager 3.0 in order to add a GUI management tool to their SAN admin suite. After installation, the fabric is inaccessible to the tool. What changes to the fabric policies are required in order to accommodate the tool?

**Discussion:**

The new Fabric Manager tool uses certificates to protect authentication passwords, much like sectelnet. In order to query the switches for information, the tool establishes a Remote Procedure Call (RPC) connection to the Primary FCS Switch. The tool uses the hooks in the API to query and change data in the fabric. In order to facilitate this data exchange, the API\_POLICY group must include the IP addresses of the SAN administrators that will be running the tool on their desktops.

**SFOS Commands:**

The process for the API\_POLICY follows the same logic as in the HTTP\_POLICY in Example 1 above.

### Example 3: Physical Policies

**Situation:**

Assuming a similar situation to the previous two examples, suppose that two of the switches in the SAN (shown as System E and System F in the figure below) are physically located in a very secure facility, while the other switches in the fabric are located in less secure areas. The SAN administrators become concerned that the current physical controls on the other switch devices are not adequate. They are particularly concerned about console or front panel access because the rooms are not well monitored. The administrators feel fairly comfortable about the other layers of security (port lockdowns, the SCC policy, and other SFOS control mechanisms), but would like to control local switch access. What measures, besides tightening physical controls, would improve security on the other SAN switches?

## Physical Access Policies




Figure 11: Physical Access Control Policy Example

### Discussion:

Fortunately, the SFOS does address the need to secure physical access to the switches by way of two specific policies. The SERIAL\_POLICY and FRONTPANEL\_POLICY use the WWN of the switches to determine which devices will have access by way of the console port of front panel enabled. Creating any of these policies will enable local access for only the switches whose WWNs are in the group. The following table illustrates the creation of two of these policies:

Policy	Entries
SERIAL_POLICY (used on switches with console ports)	10:00:00:60:69:10:66:88
FRONTPANEL_POLICY (used on 2800 model switches)	10:00:00:60:69:10:60:78

Table 12: New Policy Entries

Once applied and distributed, these policies will disable the console and front panels on all other switches in the fabric with the exception of System E and System F, respectively.

- ❖ Note: Prior to customization, the default behavior of these policies is **not** to restrict access. (This is also known as “default allow.”)

### SFOS Commands:

The commands needed for these policies follow the same logic as the HTTP\_POLICY in Example 1 above. When creating these types of policies, be sure to use the correct WWN in the policy.

## 2.2.4 Secured Management Channel

Even with the additional layer of security provided by the various MAC policies, having protected authentication information is paramount. Management access in the traditional SAN environment is not secure without third-party applications. The SFOS uses the public key of the target switch in order to encrypt the password during transport over the network. This minimizes the possibility of divulging sensitive username and password information.

To illustrate, the following screen shots show passwords captured with a commonly used packet capture application called Ethereal. The images show how the data looks using an unencrypted password and then using sectelnet to encrypt the password. The sectelnet is an application developed by Brocade to provide remote, “telnet-like command line interface connectivity to the switches in the fabric. A few of sectelnet’s salient features include:

- Administrators can download the tool from The Brocade Web site.
- The application can be installed on both Microsoft and Sun operating systems.
- Secure telnet uses asymmetric encryption and leverages the digital certificate infrastructure to authenticate the target switch and then transmit encrypted passwords.

Below are examples of encrypted and unencrypted remote access.

**Telnet Example:**

When an administrator types the username and password, telnet transfers the information one character at a time. The following figures show packet capture information for a user authenticating to a fabric switch that is not in secure mode. Here are some pointers to help decipher the capture information:

- The destination is a switch at the IP address 192.168.149.28
- The login password (provided ahead of time for this example) is: abc12345

The figure below shows the switch prompting the user for a password. The data string “password” is a common target search word in data captures and could arouse the interest of a user that is “sniffing” traffic. Figures 13 and 14 then show the initial characters of the password being typed.

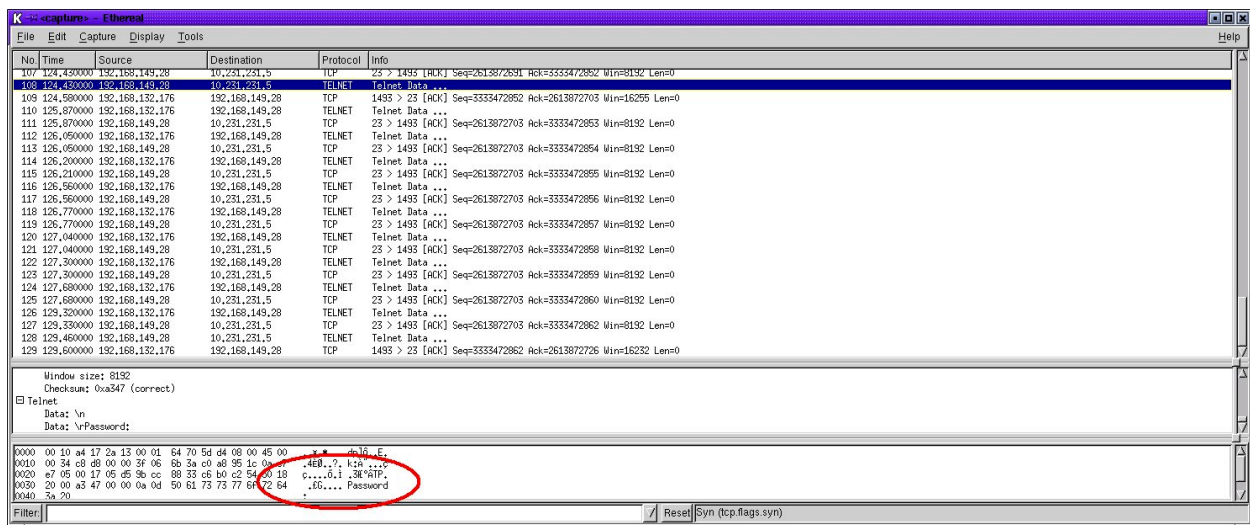


Figure 12: Password Prompt

The next figure shows the beginning of the password as it passes “in the clear” over the IP network. The first character is an “a” and appears as the last item in the circled characters.

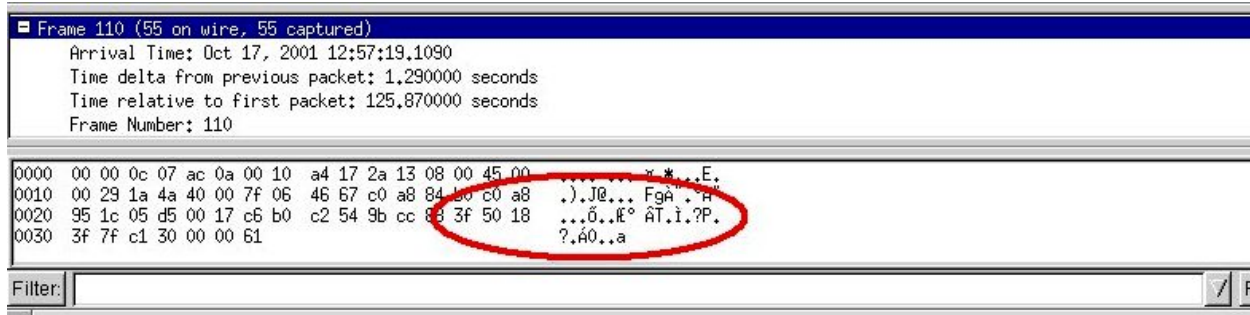


Figure 13: The First Password Value

Figure 14 shows the second character captured in the password string. The letter “b” appears as the last item in the circled characters.

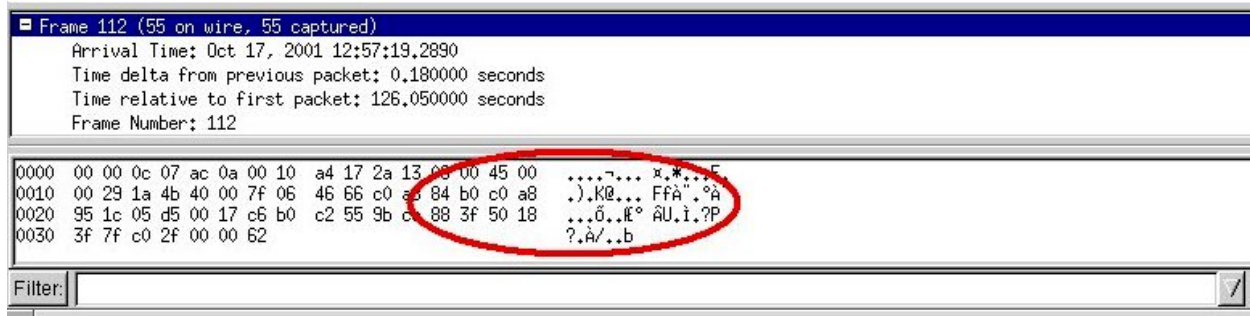


Figure 14: The Second Password Value

The sniffer can detect the remaining password characters in the same manner. In this way, an observer could quickly obtain the admin password “abc12345” for the SAN and have full access to all fabric devices.

**Secure Telnet Example:**

Secletelnet uses information obtained from the switch to encrypt the password properly before transmitting it over the wire.

The figure below shows how, with sectelnet, the SFOS protects the same password prompt as seen in the unsecured telnet session. The SFOS encrypts the actual password value.

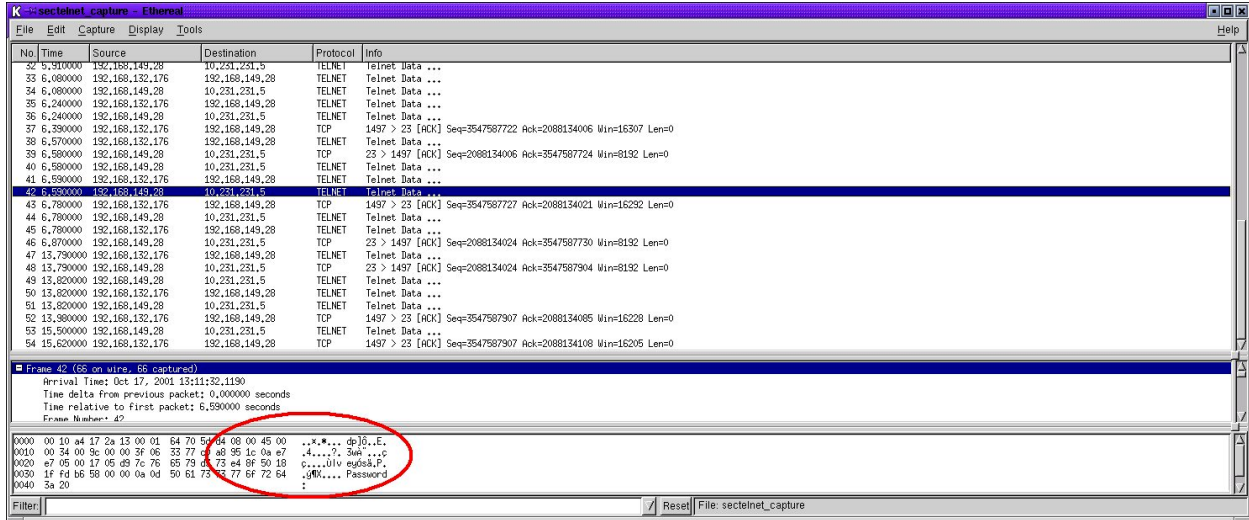


Figure 15: The Password Prompt

This time, however the public key of the switch encrypts the password value. So, the password value that appears on the IP network looks like gibberish.

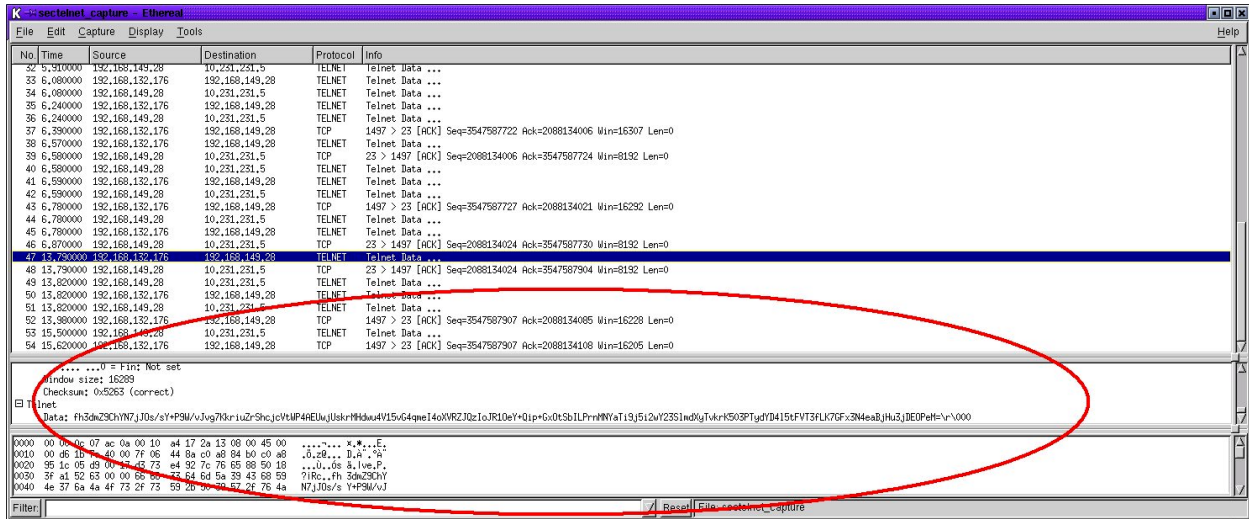


Figure 16: The Encrypted Password

Once the switch gets the encrypted password it replies back to the sender that it is decrypting the password.



```

telnet
Data: \n
Data: \rDecrypting the password.\n
Data: \rPlease wait a few seconds...\n
Data: \r

0000 00 10 a4 17 2a 13 00 01 64 70 5d d4 08 00 45 00  ..x.*... dp]ô..E.
0010 00 62 00 d2 00 00 3f 06 33 13 c0 a8 95 1c 0a e7  .b.ô..?. 3.A"...ç
0020 e7 05 00 17 05 d9 7c 76 65 8b d3 73 e5 43 50 18  ç...ùlv e.ôsâCP.
0030 1f fd ee b4 00 00 0a 0d 44 65 63 72 79 70 74 69  .yî'.... Decrypti
0040 6e 67 20 74 68 65 20 70 61 73 73 77 61 72 64 2e  ng the p assword.
Filter:

```

Figure 17: Decrypting the Password

Once the password is validated, the rest of the session occurs “in the clear.” In this session, the admin connects to the switch SecureSAN28.

```

Frame 53 (77 on wire, 77 captured)
Arrival Time: Oct 17, 2001 13:11:41.0290
Time delta from previous packet: 1.520000 seconds
Time relative to first packet: 15.500000 seconds
Frame Number: 53

0000 00 10 a4 17 2a 13 00 01 64 70 5d d4 08 00 45 00  ..x.*... dp]ô..E.
0010 00 3f 00 d4 00 00 3f 06 33 34 c0 a8 95 1c 0a e7  .?.ô..?. 34A"...ç
0020 e7 05 00 17 05 d9 7c 76 65 c5 d3 73 e5 43 50 18  ç...ùlv eAôsâCP.
0030 20 00 33 0c 00 00 0a 0d 0a 0d 53 65 63 75 72 65  .3..... ..Secure
0040 53 41 4e 32 38 3a 61 64 6d 69 6e 3e 20          SAN28:ad min>
Filter:

```

Figure 18: Other Data After Password Validation

## Protecting the Fabric with Certificates

SFOS uses digital certificates as one of the building blocks in establishing a secure fabric. Digital certificates are used in the Switch Link Authentication Protocol (SLAP) to authenticate switches and by setelnet and other clients to protect sensitive password information.

### Obtaining Digital Certificates for Switches

By following the certificate generation steps in the Fabric OS User Guide, administrators can gather the CSRs from each of the switches. Brocade then processes the CSRs to create proper Digital Certificates, which are then installed back on each switch. At this point, each switch has a certificate with which it can authenticate itself (using switch WWN) and a method to authenticate other switches. As previously discussed, authentication is the process of proving that users are who they say they are. Validation can take a number of forms, but it usually consists of one or more of the following elements:

- **What you know:** For example, the admin name and password.
- **Who you are:** A method that uses biometric data to verify physical characteristics, such as a fingerprint or palm scan.
- **What you have:** For example, a Digital Certificate or an RSA SecurID card.

In this case, the switches use their certificates to prove their identities to other switches and management devices.

- ❖ Note: Any switch that comes from the factory with a Fabric OS that supports Secure Fabric OS (such as version 4.1) already has a Digital Certificate loaded on the system.



## 2.2.5 Switch Connection Controls (SCC)

One way to prevent unauthorized switches from joining the fabric is to use of the SCC\_POLICY. Much like the MAC policies, the SCC is simply a group of switch WWNs that have permission to join the fabric. The SCC\_POLICY contains one WWN for each valid switch in the fabric.

The SCC\_POLICY provides an additional “sanity check” to the SAN creation or modification process. Administrators can take two approaches when configuring the fabric to use the SCC security layer.

- The administrator can add the asterisk (\*) character to the *secPolicyCreate* command to create and configure the SCC\_POLICY. This form of the command adds all the WWNs of the switches currently in the security enabled fabric. In this case, the administrator must verify the WWN names and the number of devices to ensure that there are no incorrect or unexpected devices in the fabric.
- The administrator can create the policy manually by adding the WWNs of the switches that are to participate in the fabric. Note that upon creation, the SCC automatically adds all FCS switches to the group, giving the administrator a useful starting point.
  - ❖ Note: Use the SCC\_POLICY to prevent unauthorized switches from joining the fabric and not as a tool to isolate authorized switches that are already in the secure fabric.

When the administrator validates the WWNs of the switches in the SCC\_POLICY against a list of approved names, this compliments the controls implemented by the use of certificates and SLAP. Technical controls alone can only be used to verify switch information (via certificates) and enforce policies. Without management and operational controls, technical controls alone cannot actually create an organization’s policy or validate the information. The combination of management, organizational, and technical controls provides a useful system of checks and balances in the SAN environment and the enterprise as a whole.

### SCC Example

#### **Situation:**

Administrators have configured the SCC\_POLICY for a fabric consisting of switches A, B, and C. What happens when they connect switch D to switch C?

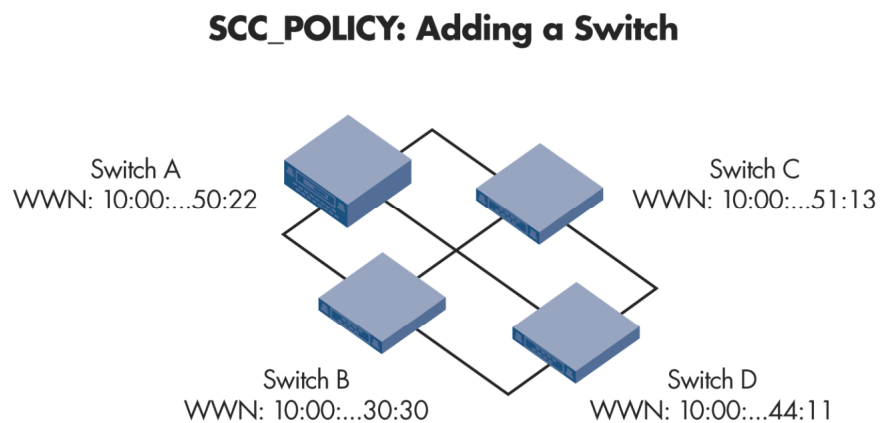


Figure 19: SCC Example – Adding a Switch

**Discussion:**

In the initial SAN configuration, the SAN administrators verified the WWNs of switches A, B, and C and then added them to the SCC\_POLICY group. With the addition of the new switch, the *secPolicyAdd* command makes it easy to modify the policy to accommodate switch D. The changes will take effect with policy activation.

Policy	Entries (WWNs truncated for example only)
SCC_POLICY	10:00:....:50:22
	10:00:....:30:30
	10:00:....:51:13
	10:00:....:44:11 ← New (Switch D)

**Table 13:** SCC\_POLICY with New Switch

- ❖ Note: If the administrator does not modify the SCC\_POLICY to include Switch D, then the SFOS will reject the switch when it attempts to join the fabric.

### Switch Link Authentication Protocol (SLAP)

To prevent a non-valid switch from spoofing a valid WWN in order to join the fabric, switches conduct mutual authentication on their E\_ports. Using digital certificates and the SLAP protocol helps make a switch's WWN much more trustworthy.

- ❖ **Note:** For details regarding the capabilities and functions of certificates, refer to the FCAP and SLAP protocol proposals documented in ANSI T11. Also, the Brocade Web site ([www.brocade.com](http://www.brocade.com)) contains a white paper entitled *Driving Advancements in Storage Ecosystem Management and Security*, which discusses Fabric Device Management Interface (FDMI) and Fibre Channel Authentication Protocol (FCAP).

### 2.2.6 Device Connection Controls (DCC)

Similar to a port lockdown methodology in an IP switch environment, DCCs allow the SAN administrator to select what device WWNs can connect to which switch ports. By creating various unique policies using the DCC\_POLICY\_XXX name format, administrators can lock down a fabric to varying degrees of granularity. To achieve extreme control (and high management), the administrator can connect a fabric so that each switch port can connect to only a single WWN. DCCs are more flexible: a group of switch ports can support numerous WWNs on any given port. This configuration range allows the SAN administrator to strike the balance between security and flexibility by layering the amount of restrictions on the fabric.

#### DCC Example

**Situation:**

The administrators have configured the DCC\_POLICY for a fabric consisting of switches A and B. They have locked down ports 1, 2, and 3 on Switch A to support a group of device WWNs. What happens if a new host tries to plug into Port 1?

## DCC\_POLICY\_hrgroup Example

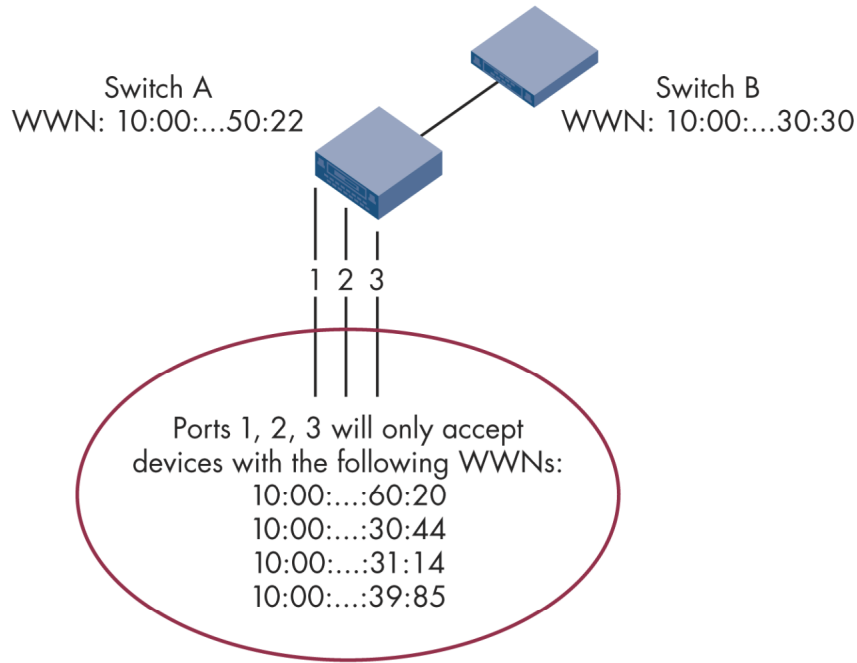


Figure 20: DCC\_POLICY Example

**Discussion:**

Because the Human Resources team primarily uses the devices connected to ports 1, 2, and 3 and switch A, the administrators create a policy called “DCC\_POLICY\_hrgroup” that contains the three ports and allows any of four WWNs to connect to them. This allows flexibility in physical cabling and allows the option of physically swapping one of the three devices with a fourth allowed device, such as if a backup HBA, without disruption.

Policy	Entries (WWNs truncated for example only)
DCC_POLICY_hrgroup	10:00:....:60:20
	10:00:....:30:44
	10:00:....:31:14
	10:00:....:39:85

Table 14: DCC\_POLICY\_hrgroup

So, if someone plugs a new host with a WWN of “10:00:....:95:66” into ports 1, 2, or 3 of switch A, the SFOS will deny the device. Locking ports down to specific WWNs minimizes the possibility of an unauthorized device entering the fabric on the wrong port or of name spoofing. The DCC policy fortifies hard zoning in that not only can administrators group specific ports into a zone, those ports can also represent specific WWNs. Keep in mind that locking down ports to this level does remove some of the flexibility of hard zoning, which normally allows any device plugged into a given port to become a part of a zone. SAN administrators must calculate the pros and cons of different switch configurations before implementing the policies.

## 2.2.7 SFOS-Protected SAN Security Summary

The Secure Fabric OS greatly improves the ability to minimize SAN-specific vulnerabilities. The tool set provides significant flexibility and numerous configuration options. The table below compares the improved functionality provided by the Secure Fabric OS against corresponding SAN risks:

SECURE FABRIC OS SAN SECURITY VERSUS COMMON SAN RISKS	
<b>Device-to-Switch Traffic</b>	
Risk	Description
<b>An unauthorized, local (non-switch) device is connected to the fabric or a device is connected to or from the wrong port</b>	Using a combination of physical access controls and disabling ports greatly reduces the risk of many of these attacks. WWN spoofing can now be virtually eliminated if ports are locked down to specific names.
<b>Invalid management device connection attempt</b>	The SFOS implementation does allow for device restrictions based on the source IP address. This enables the creation of groups of authorized management systems that use encrypted passwords when authenticating to the fabric.
<b>Unprotected management application authentication information</b>	The SFOS augments the integrity of the admin username and password combination by transmitting only encrypted fabric login passwords. Network designs and third-party software can segment the unencrypted non-password data to minimize the risk of session hijacking.
<b>Device Denial of Service (DOS) attack</b>	The SFOS provides numerous options for limiting the connection of invalid devices to the fabric. Once a device has established a valid connection, the SFOS does not identify or block DOS traffic. So, devices currently connected to the fabric should be configured and protected to minimize compromise or unauthorized use. Keep in mind that it is very difficult to protect against malicious authorized users. Policies should include methods for identifying abuses and taking appropriate disciplinary action.
<b>Inter-Switch Traffic</b>	
<b>Invalid switch connected to the fabric</b>	The SFOS provides multiple ways to authenticate and limit a rogue switch from joining the fabric. First, each switch must have a valid certificate and must use that certificate to authenticate to its peer switches by way of SLAP. Further, if the SCC list has been configured, the switch WWN must appear in it.
<b>Traffic to and from Fabric Devices</b>	
<b>Unauthorized access to data</b>	Port lockdowns and physical security should help minimize WWN spoofing. Tough system- and application-level controls must still be in place to protect data access and integrity properly.
<b>Device Denial of Service (DOS) Attack</b>	Protecting against this risk is analogous to a DOS attempt on the fabric.

Figure 21: SFOS Versus Common SAN Risks

## SFOS SAN Security Functionality

The SFOS-secured environment provides significantly improved security functionality over non-SFOS environments. The table below lists SFOS security enhancements:

SECURITY FUNCTIONALITY IN AN SFOS SAN DESIGN	
Functionality	Description
Integrity	Every time an SFOS process compares a data hash to a signed data hash integrity checking is taking place. If the SFOS detects any traffic modifications, the hash value will appear different and the integrity check will fail. The SFOS does integrity checks on data such as database and time stamp updates.
Availability	In addition to all the functionality provided in the non-SFOS environment, the SFOS can minimize port spoofing through DCC policies and manage valid switches through the SCC policy. With restrictive MAC policies the SFOS can further reduce the likelihood of DOS attack vectors.
Confidentiality	In secure mode, the SFOS provides end-to-end password encryption.
Authentication	The SFOS still uses username/password combinations to authenticate remote users, however the passwords are protected with encryption. Further, switches use certificates and the SLAP protocol to authenticate one another.
Authorization	The SFOS provides additional granularity in the admin accounts. Only changes made on the Primary FCS Switch with the fcs-admin account can effect fabric-wide configuration changes.
Accounting	In addition to the management available in the non-SFOS environment, the SFOS provides improved visibility for managing security events and configurations.

**Table 12:** Security Functionality in an SFOS SAN Design

Adding the Secure Fabric OS to the SAN environment significantly reduces the security holes left by traditional implementations. Each successive layer of security builds on the preceding one to form a unified, comprehensive structure that is stronger than the sum of its parts. The SFOS offers a variety of tools to the SAN or security architect and provides options on how they can lock down different aspects of the environment..

Security knowledge is of little value in a vacuum. Correct understanding and application proves its worth. The next chapter takes all of the security information presented to this point and applies it in a series of examples and case studies.

## 3 Implementation and Design Examples

This chapter pulls together the information presented in the previous chapter into a usable framework. Below are two case studies that show how to implement a secured SAN properly with the Secure Fabric OS. Each case states assumptions and design goals along with a recommended implementation.

### 3.1 CASE STUDY 1: A MID-SIZED ENTERPRISE ENVIRONMENT

#### 3.1.1 Current Environment

CorpA is a medium-sized company with both e-business and retail store sales. The company relies on “24x7” access to mission critical systems and data. The company has already implemented numerous enterprise security measures, but now wishes to upgrade the SAN environment to the enhanced security of the SFOS. The company makes the following assumptions when defining its SAN use:

***Assumptions:***

- The current SAN fabric consists of three Brocade SilkWorm 3900 switches.
- SAN capacity is currently adequate for the fabric requirements.
- All SAN elements are fairly well secured and protected physically.
- Systems accessible from the Internet all reside on a screened subnet.
- Only two SAN administrators require remote configuration access to the environment
- Change control procedures are in place for scheduling SAN downtime.

The following diagram illustrates the logical layout of the environment:

## Current CorpA Environment

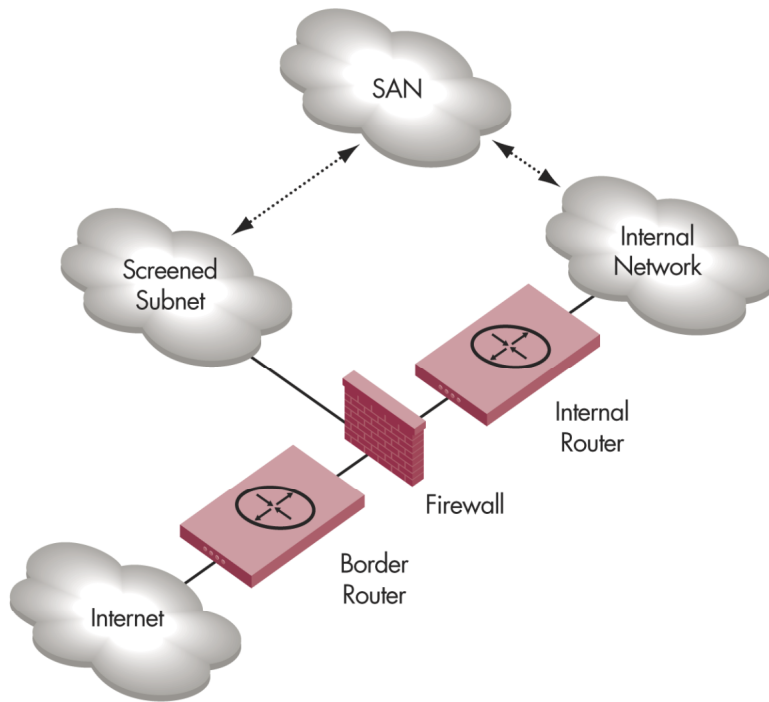


Figure 22: Current CorpA Environment

### 3.1.2 Description

CorpA has already implemented the following significant non-SFOS security measures to help protect their infrastructure:

- Proper security policy, change control, and acceptable use documentation
- Carefully configured router ACL and firewall policies that are consistent with the security policy
- Various IDS sensors strategically placed throughout the enterprise
- Out-of-Band (OOB) networks to segment sensitive data streams
- Internal Access Control Lists (ACLs) or firewalls to limit traffic to sensitive network segments such as the SAN IP management subnet
- Proper physical controls on all wiring closets and the data center
- Zoning and port lockdown implementations to secure the current SAN
- Strong two-factor authentication system where appropriate

Because CorpA has already followed many of the guidelines outlined in this document for their infrastructure design and implementation, most of the remaining work has to do with migrating the SAN to the SFOS. The information below provides a closer view of CorpA's SAN and the infrastructure components that interact with it.

## Closer View of the CorpA SAN

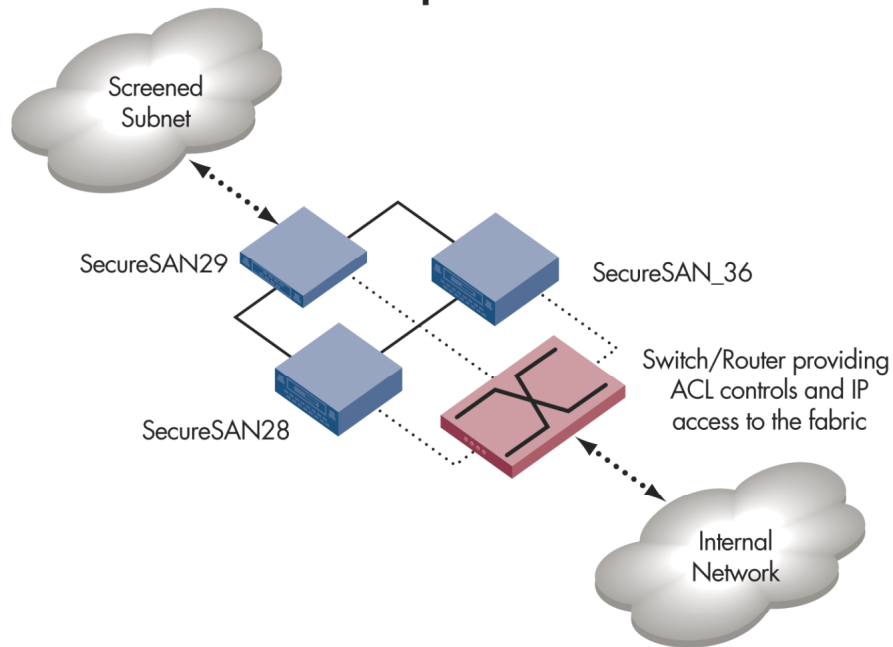


Figure 23: A Closer View of the CorpA SAN

### 3.1.3 Detailed Steps

#### Pre-SFOS Checklists

Before beginning the migration to the SFOS environment, CorpA administrators follow a checklist of tasks required to complete the transition.

- Obtain a copy of The Brocade Secure Fabric OS User Guide, which provides detailed information on enabling secure mode and implementing other SFOS features.
- Obtain the Fabric OS (FOS) code and verify that it will work on every switch model in the fabric.
- Schedule any required downtime for FOS upgrades.
- Back up the configuration information on each switch.
- Ensure that rollback procedures are approved and in place.
- Upgrade the switches to the new FOS, which will generate the following:
  - A unique public/private key pair for each switch
  - A unique Certificate Signing Request (CSR)
- Test the fabric functionality as needed to ensure correct operation.
- Purchase and install zoning and security licenses on all the switches in the fabric.
- Download and install the PKICert utility on an appropriate system.
- Use PKICert and the associated procedures to obtain and install certificates on all the switches.

Once the above steps are complete and the switches are running the correct version of the FOS, the switches are ready for secure mode. The following must be in place before enabling secure mode in the fabric:



- Install sectelnet on each SAN administrator system that requires it. The sectelnet application is required to enable secure mode remotely.
- Upgrade any Web-based administration tools to appropriate revision levels.
- Review and modify any router ACLs as needed to allow appropriate administrative access to the SAN.

### Collecting Supplementary Information

Before enabling secure mode, CorpA’s administrators also consolidate important information about the SAN and associated devices. This includes collecting the WWNs of all the switches as well as the IP addresses of any remote management systems involved in configuring or viewing SAN information. The administrators will use the information to configure various SFOS options after enabling secure mode.

Here is a summary of the information collected:

CORPA SAN COMPONENT INFORMATION	
Device	Device Information
SecureSAN28	10:00:00:60:69:10:69:78
SecureSAN29	10:00:00:60:69:11:F9:94
SecureSAN_36	10:00:00:60:69:10:97:83
SAN Admin 1	192.168.132.84 (static IP address)
SAN Admin 2	192.168.132.100 (static IP address)

**Table 16:** CorpA SAN Component Information

### Enabling Secure Mode in the Fabric

At this point, CorpA is poised to implement the fabric in secure mode and begin using the functionality of the SFOS. The SAN administrators proceed as follows:

- The administrators have read and understood the new security features in the SFOS. They have determined that their three switches would play the following roles:
  - SecureSAN28: Primary FCS Switch
  - SecureSAN\_36: Backup FCS Switch
  - SecureSAN29: Non-FCS Switch
- One administrator uses the sectelnet application to connect to the IP address of the Primary FCS Switch (SecureSAN28) and authenticates using the admin password.
- Note: The fabric must be in secure mode to ensure password security. At this point, the password will pass over the network in clear text. To encrypt the password at this stage requires third-party software.
- The administrator initiates secure mode using the *secModeEnable* command. This specifies the FCS switches and creates the FCS\_POLICY group, which add the switch WWNs in priority order with the first device being the Primary FCS Switch.
- At this point, the system prompts the administrator for passwords to all accounts and passes all password information in encrypted form. The administrator chooses strong passwords, records them carefully and accurately, and immediately locks them in a secure location. (Policies are already in place in CorpA to manage these secured passwords.)

```

192.168.149.28 - sTelnet
Sessions Edit Terminal Help
SecureSAN28:admin> fabricShow
Switch ID      Worldwide Name      Enet IP Addr      FC IP Addr      Name
-----
  1: Fffc01 10:00:00:60:69:10:69:78 192.168.149.28  0.0.0.0      >"SecureSAN28"
  2: Fffc02 10:00:00:60:69:11:f9:94 192.168.149.29  0.0.0.0      "SecureSAN29"
  5: Fffc05 10:00:00:60:69:10:97:83 192.168.149.36  0.0.0.0      "SecureSAN_36"

The Fabric has 3 switches

SecureSAN28:admin> secModeEnable

This is an interactive session to create a FCS list.

The new FCS list is empty.

Enter WWN, Domain, or switch name(Leave blank when done): 10:00:00:60:69:10:69:78
New Switch WWN is 10:00:00:60:69:10:69:78.

The new FCS list:
 10:00:00:60:69:10:69:78

Enter WWN, Domain, or switch name(Leave blank when done): 10:00:00:60:69:10:97:83
New Switch WWN is 10:00:00:60:69:10:97:83.

The new FCS list:
 10:00:00:60:69:10:69:78
 10:00:00:60:69:10:97:83

Enter WWN, Domain, or switch name(Leave blank when done):
Are you done? (yes, y, no, n): [no] yes
Is the new FCS list correct? (yes, y, no, n): [no] yes
Each encryption/decryption of password takes a while
New FCS switch root password:
Re-enter new password:
New FCS switch factory password:
Re-enter new password:
New FCS switch admin password:
Re-enter new password:
New fabric wide user password:
Re-enter new password:
New Non FCS switch admin password:
Re-enter new password:
Saving passwd...done.
Saving Defined FMPS ...
done
Saving Active FMPS ...
done
Committing configuration...

```

Figure 24: Enabling Secure Mode in the CorpA Fabric

- Once the configuration is committed to flash memory, the switches reboot and come back up in secure mode.
- The administrator then reconnects to the Primary FCS Switch and verifies that all the switches are responding correctly. The commands used include:
  - *FabricShow* to verify that all three switches appear in view
  - *SwitchShow* to check that ports are not segmented
  - *NsAllShow* to validate that all devices are registered in the name server

- *secModeShow* to verify that secure mode is enabled and to verify the FCS list and order
- The administrator then uses *sectelnet* to verify that the other two switches are in working order.

### Checking Default SFOS Functionality

CorpA now has a SAN of three Brocade switches all running the correct version of the Fabric OS, with secure mode enabled. The next step in securing the SAN environment is to review the default security policies that are in place. The following table contains the necessary information:

#### DEFAULT SFOS SECURITY FEATURES

Feature	Description
FCS_POLICY	This policy must exist and cannot be empty. The SFOS generates this policy when enabling secure mode; it includes:  SecureSAN28: 10:00:00:60:69:10:69:78 SecureSAN_36: 10:00:00:60:69:10:97:83
MAC Policies: SNMP, TELNET, HTTP, SES, MS, Serial, and FrontPanel	These policies are in the "No Policy" state by default; as such, they do not limit or block access.
Options	By default, options are not enabled and therefore allow WWW zoning.
DCC	The DCC is set to "No Policy" by default: any device can connect to any port.
SCC	The SCC is set to "No Policy" by default: a switch with any WWN can connect to the fabric.

**Table 17:** Default SFOS Features

### Proposed Usage of SFOS Features

CorpA's SAN administrators have a good understanding of what the SFOS can provide to the fabric as well as what problems they must address. The following list contains the issues that CorpA wants to address with the new configuration:

- Eliminate any non-approved switches from joining the fabric.
- Lock down all required remote administration to the IP addresses of the two SAN administrators.
- Eliminate any unneeded access, including physical access to the switches.
- Lock down all switch ports in some fashion.

The following table lists configuration changes required to address these security concerns are provided in the table below:

## PROPOSED SFOS CONFIGURATION

Feature	Description
FCS_POLICY	SecureSAN28 - 10:00:00:60:69:10:69:78 SecureSAN_36 - 10:00:00:60:69:10:97:83
MAC policy subset: TELNET_POLICY, HTTP_POLICY, API_POLICY	Lock down to the SAN administrator IP addresses: 192.168.132.84 (static IP address) 192.168.132.100 (static IP address)
MAC policy subset: SNMP, SES, MS	Create empty in order to disable SNMP, SES, and MS access.
MAC policy subset: SERIAL_POLICY, and FRONTPANEL_POLICY	Create empty in order to disable console and front panel access to the devices; the administrators feel that they can easily re-enable access if and when necessary.
Options	Not critical for CorpA; leave in default mode.
DCC_POLICY_disk1A	Create a policy group to lock down two of CorpA's SAN device WWNs to a specific port group and switch.
DCC_POLICY_sshosts	Create a policy group to lock down the screened subnet SAN host WWN to a specific port and switch.
SCC_POLICY	Create the policy and lock it down to the three switches: SecureSAN28 - 10:00:00:60:69:10:69:78 SecureSAN_36 - 10:00:00:60:69:10:97:83 SecureSAN29 - 10:00:00:60:69:11:F9:94
Other Security Measures: Zoning, Port Lockdown	Configure zoning the same way as pre-SFOS implementation and disable all unused ports.

**Table 18:** Proposed SFOS Security

Now that the SAN administrators have a plan and know what SFOS features they wish to use, implementation can begin. The next section illustrates the steps that CorpA administrators take to lock down the SAN environment.

### Implementing CorpA's SFOS Security Plan

The CorpA SAN administrators follow this progression to implement SFOS changes:

- Create telnet, HTTP, and API policies with the administrator IP addresses.
- Lock down all other unused policies.
- Create DCC policies to restrict identified ports.
- Create the SCC\_POLICY to prevent other switches from joining the fabric.
- Verify existing zoning and port lockdown configurations.
  - ❖ Note: The following commands are helpful in understanding SFOS syntax and functionality:  
**SecHelp** Displays a list of security-related commands.  
**Help XX** provides access to help information for any give command; (replace "XX" with the command in question).

## Allow Remote Access

The administrators follow these four steps to restrict remote access to the switches:

- Create and modify the TELNET\_POLICY to include all SAN administrator IP addresses.
- Create and modify the HTTP\_POLICY to include all SAN administrator IP addresses.
- Create and modify the API\_POLICY to include all SAN administrator IP addresses.
- Activate these policies when complete.
  - ❖ Note: Administrators should check the policies by running *secPolicyDump* and review the entries in the Defined Policy Set section before activating the changes; (This procedure is not shown in Figure 29).

```
SecureSAN28:admin>
SecureSAN28:admin>
SecureSAN28:admin> secPolicyCreate "TELNET_POLICY","192.168.132.84"
TELNET_POLICY has been created.
SecureSAN28:admin> secPolicyAdd "TELNET_POLICY","192.168.132.100"
Member(s) have been added to TELNET_POLICY.
SecureSAN28:admin>
SecureSAN28:admin> secPolicyCreate "HTTP_POLICY","192.168.132.84"
HTTP_POLICY has been created.
SecureSAN28:admin> secPolicyAdd "HTTP_POLICY","192.168.132.100"
Member(s) have been added to HTTP_POLICY.
SecureSAN28:admin>
SecureSAN28:admin> secPolicyCreate "API_POLICY","192.168.132.84"
API_POLICY has been created.
SecureSAN28:admin> secPolicyAdd "API_POLICY","192.168.132.100"
Member(s) have been added to API_POLICY.
SecureSAN28:admin>
SecureSAN28:admin> secPolicyActivate
About to overwrite the current Active Policy Set.
ARE YOU SURE (yes, y, no, n): [no] yes
Committing configuration...done.
Saving Defined FMPS ...
done
Saving Active FMPS ...
done
SecureSAN28:admin>
```

Figure 25: Allowed Remote Access Policies

- ❖ Note: The *secPolicyCreate* command also allows the addition of multiple addresses.

## Lock Down Unused Switch Access

The administrator then goes through and creates empty policies for controls currently not needed. When left empty, these controls disable much of the unused access functionality. The figure below shows the necessary steps.

```

SecureSAN28:admin>
SecureSAN28:admin>
SecureSAN28:admin> secPolicyCreate "WSNMP_POLICY"
WSNMP_POLICY has been created.
SecureSAN28:admin> secPolicyCreate "RSNMP_POLICY"
RSNMP_POLICY has been created.
SecureSAN28:admin> secPolicyCreate "SES_POLICY"
SES_POLICY has been created.
SecureSAN28:admin> secPolicyCreate "MS_POLICY"
MS_POLICY has been created.
SecureSAN28:admin> secPolicyCreate "SERIAL_POLICY"
SERIAL_POLICY has been created.
SecureSAN28:admin> secPolicyCreate "FRONTPANEL_POLICY"
FRONTPANEL_POLICY has been created.
SecureSAN28:admin> secPolicyActivate
About to overwrite the current Active Policy Set.
ARE YOU SURE (yes, y, no, n): [no] yes
Committing configuration...done.
Saving Defined FMPS ...
done
Saving Active FMPS ...
done
SecureSAN28:admin> █

```

Figure 26: Deny Un-needed Policies

### DCC Policy Creation

The administrator decides to lock down some devices to designated ports and switches by using the DCC functionality. This requires the information contained in the table below.

#### CORPA'S DCC CONFIGURATIONS

Feature	Description
DCC_POLICY_disk1A	Group the following to SecureSAN_36, ports 5/6: Device1 : 10:00:00:60:69:10:99:88 Device2: 10:00:00:60:69:10:90:66
DCC_POLICY_sshosts	Host: 10:00:00:60:69:10:50:55, SecureSAN29, port 8

Table 19: DCC Configuration Information

The figure below shows the steps for creating these policies.

```

SecureSAN28:admin>
SecureSAN28:admin>
SecureSAN28:admin>
SecureSAN28:admin> secPolicyCreate "DCC_POLICY_disk1A","10:00:00:60:69:10:99:88;SecureSAN_36(5)"
DCC_POLICY_disk1A has been created.
SecureSAN28:admin>
SecureSAN28:admin> secPolicyAdd "DCC_POLICY_disk1A","10:00:00:60:69:10:90:66;SecureSAN_36(6)"
Member(s) have been added to DCC_POLICY_disk1A.
SecureSAN28:admin>
SecureSAN28:admin> secPolicyCreate "DCC_POLICY_sshosts","10:00:00:60:69:10:50:55;SecureSAN29(8)"
DCC_POLICY_sshosts has been created.
SecureSAN28:admin>

```

Figure 27: DCC Policy Creation Steps

After verifying the DCC policies, (refer to the figure below) the administrator can activate the policy set.

```

DCC_POLICY_disk1A
Type      WWN                               DId swName
-----
Type <CR> to continue, Q<CR> to stop:
Switch 10:00:00:60:69:10:97:83  5 SecureSAN_36.
=Port=> 5,6.
Device 10:00:00:60:69:10:99:88
Device 10:00:00:60:69:10:90:66

DCC_POLICY_sshosts
Type      WWN                               DId swName
-----
Switch 10:00:00:60:69:11:f9:94  2 SecureSAN29.
=Port=> 8.
Device 10:00:00:60:69:10:50:55

```

Figure 28: Verifying the DCC Policies

### SCC Policy Creation

The administrator knows the SAN environment well and recognizes a special "\*" option that is available when initially creating the SCC\_POLICY group. This automatically places the WWNs of all switches connected to the secure fabric into the SCC list automatically. The administrator performs the following steps:

- Check the fabric to ensure that only the three known switches are currently connected and that there are no "test" or "rogue" switches on the fabric.
- Create the SCC\_POLICY group.
- Activate the new policy configuration.

```

SecureSAN28:admin>
SecureSAN28:admin> fabricShow
Switch ID  Worldwide Name          Enet IP Addr  FC IP Addr    Name          1
-----
1: fffc01 10:00:00:60:69:10:69:78  192.168.149.28  0.0.0.0      >"SecureSAN28"
2: fffc02 10:00:00:60:69:11:f9:94  192.168.149.29  0.0.0.0      "SecureSAN29"
5: fffc05 10:00:00:60:69:10:97:83  192.168.149.36  0.0.0.0      "SecureSAN_36"

The Fabric has 3 switches

SecureSAN28:admin> secPolicyCreate "SCC_POLICY","*" 2
SCC_POLICY has been created.
SecureSAN28:admin>
SecureSAN28:admin> secPolicyActivate
About to overwrite the current Active Policy Set. 3
ARE YOU SURE (yes, y, no, n): [no] yes
Committing configuration...done.
Saving Defined FMPS ...
done
Saving Active FMPS ...
done
SecureSAN28:admin>

```

Figure 29: SCC Policy Creation

## Policy Review

The administrator can now verify the Active Policy Set as a whole for accuracy and completeness. The *secPolicyShow* and *secPolicyDump* commands can display the policy either all at once or one page at a time. The figure below shows the completed security policy for CorpA.

```
Sessions Edit Terminal
FCS_POLICY
Pos Primary WWN Did swName
-----
1 Yes 10:00:00:60:69:10:69:78 1 SecureSAN28
2 No 10:00:00:60:69:10:97:83 5 SecureSAN_36
SCC_POLICY
WWN Did swName
-----
10:00:00:60:69:10:69:78 1 SecureSAN28
10:00:00:60:69:11:F9:94 2 SecureSAN29
10:00:00:60:69:10:97:83 5 SecureSAN_36
SERIAL_POLICY
WWN Did swName
-----
EMPTY
FRONTPANEL_POLICY
WWN Did swName
-----
EMPTY
TELNET_POLICY
IpAddr
-----
192.168.132.84
192.168.132.100
HTTP_POLICY
IpAddr
-----
192.168.132.84
192.168.132.100
API_POLICY
IpAddr
-----
192.168.132.84
192.168.132.100
RSNMP_POLICY
IpAddr
-----
EMPTY
WSNMP_POLICY
IpAddr
-----
EMPTY
SES_POLICY
Device WWN
-----
EMPTY
MS_POLICY
Device WWN
-----
EMPTY
DCC_POLICY_disk1A
Type WWN Did swName
-----
Switch 10:00:00:60:69:10:97:83 5 SecureSAN_36.
=Port=> 5,6.
Device 10:00:00:60:69:10:99:88
Device 10:00:00:60:69:10:90:66
DCC_POLICY_sshosts
Type WWN Did swName
-----
Switch 10:00:00:60:69:11:F9:94 2 SecureSAN29.
=Port=> 8.
Device 10:00:00:60:69:10:50:55
```

Figure 30: CorpA SFOS Security Policies

## Verify Other Configuration Information

As a next step, the administrator verifies the other zoning and port lockdown steps and checks the functionality. (Note: Because these steps do not involve new configurations to the environment, this section does not discuss them in detail.)



3.2 CASE STUDY 2: GROWING THE CORPA ENVIRONMENT

3.2.1 Current Environment

CorpA has grown significantly since the last example. The number of devices requiring access to the SAN has increased to the point that CorpA has purchased a new Brocade switch that will function as a non-FCS switch. The diagram below illustrates the new environment:

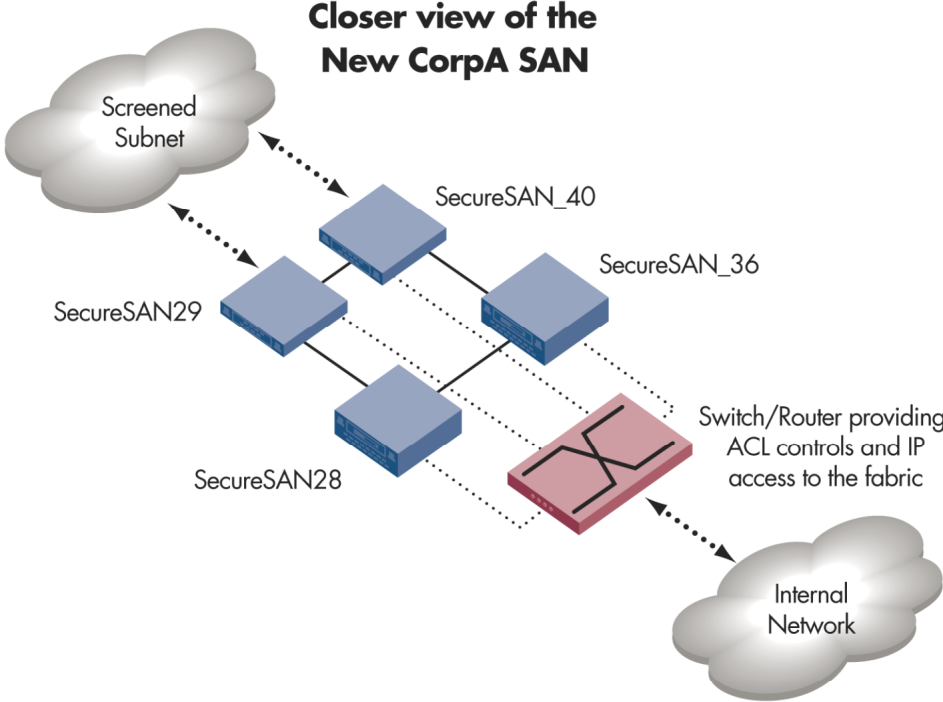


Figure 31: The New CorpA SAN Environment

The new switch comes with a version of the Fabric OS that supports SFOS functionality. The switch also ships with a certificate pre-loaded, which eliminates the need to generate the information. The new switch has the following WWN:

Device	WWN
SecureSAN_40	10:00:00:60:69:11:FC:9C

Table 20: New Switch Information

### 3.2.2 Steps

First, the administrator reviews the “Joining the Secure Fabric” section of the Secure Fabric OS User Guide. CorpA’s SAN administrators then complete the following steps in order to integrate the new switch into the fabric:

- Re-enable one port on SecureSAN29 and one on SecureSAN\_36.
  - ❖ Note: If the two ports have been locked down to prevent them from coming up as E\_ports, re-enabling them will undo this configuration.
- Add the WWN of SecureSAN\_40 to the SCC list and activate the new policy, allowing it join the fabric.
  - ❖ Note: Because SecureSAN\_40 is a non-FCS switch, no modifications to the FCS\_POLICY are required.

```
SecureSAN28:admin> secPolicyAdd "SCC_POLICY","10:00:00:60:69:11:fc:9c"  
Member(s) have been added to SCC_POLICY.  
SecureSAN28:admin> secPolicyActivate  
About to overwrite the current Active Policy Set.  
ARE YOU SURE (yes, y, no, n): [no] yes  
Committing configuration...done.  
Saving Defined FMPS ...  
done  
Saving Active FMPS ...  
done  
SecureSAN28:admin>  
SecureSAN28:admin> █
```

Figure 32: Adding the New WWN to the SCC\_POLICY

- Turn SecureSAN\_40 on. (At this point, the switch is IP-accessible, but not yet connected to the fabric.)
- Enable secure mode on SecureSAN\_40 (*secModeEnable*) and identify the WWNs of the Primary FCS Switch and Backup FCS Switches.
- Reboot the SecureSAN\_40 switch, causing it to come up in secure mode.
  - ❖ Note: The new switch is now connected to the other switches and updates can occur between the devices and SecureSAN\_40.

The new fabric for CorpA now has four switches as illustrated in the figure below:

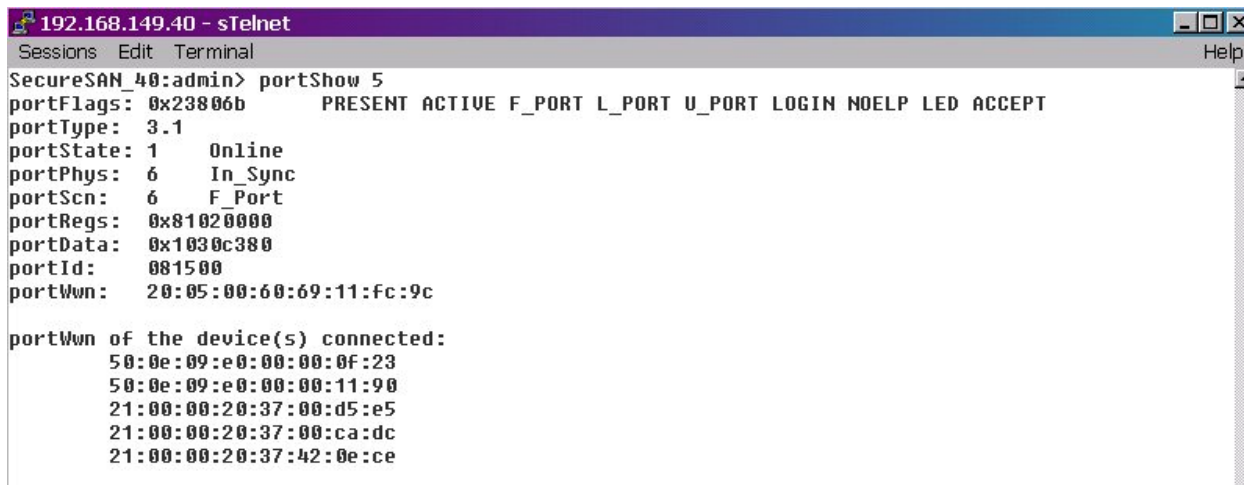
```
SecureSAN28:admin> fabricShow  
Switch ID  Worldwide Name          Enet IP Addr  FC IP Addr    Name  
-----  
1: fffc01  10:00:00:60:69:10:69:78  192.168.149.28  0.0.0.0      >"SecureSAN28"  
2: fffc02  10:00:00:60:69:11:f9:94  192.168.149.29  0.0.0.0      "SecureSAN29"  
5: fffc05  10:00:00:60:69:10:97:83  192.168.149.36  0.0.0.0      "SecureSAN_36"  
8: fffc08  10:00:00:60:69:11:fc:9c  192.168.149.40  0.0.0.0      "SecureSAN_40"  
  
The Fabric has 4 switches  
SecureSAN28:admin> █
```

Figure 33: The New CorpA Switch Fabric

### 3.2.3 Additional Information

Now that SecureSAN\_40 is part of the secure fabric, the administrators plan to attach a JBOD to it. They want to lock down ports with a DCC policy. The thing to note here is that SecureSAN\_40 has multiple WWNs that all connect through one switch port. The steps below show how to handle this situation:

- Plug the device into port 5 of the target switch and start it up. Then use *sectelnet* to connect to SecureSAN\_40 and run the command *portShow 5* to discover the device's WWNs. The figure below shows the results.

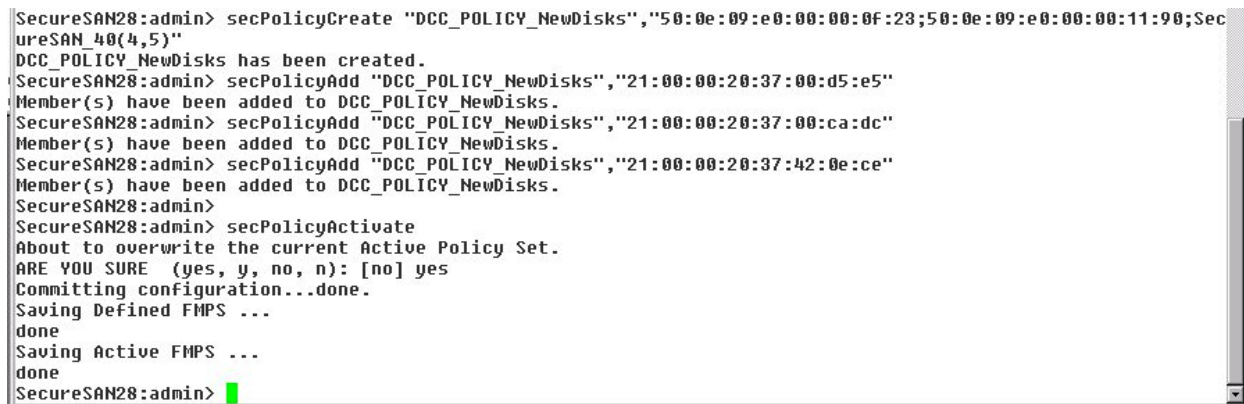


```
192.168.149.40 - sTelnet
Sessions Edit Terminal
SecureSAN_40:admin> portShow 5
portFlags: 0x23806b      PRESENT ACTIVE F_PORT L_PORT U_PORT LOGIN NOELP LED ACCEPT
portType: 3.1
portState: 1      Online
portPhys: 6      In_Sync
portScn: 6      F_Port
portRegs: 0x81020000
portData: 0x1030c380
portId: 001500
portWwn: 20:05:00:60:69:11:fc:9c

portWwn of the device(s) connected:
50:0e:09:e0:00:00:0f:23
50:0e:09:e0:00:00:11:90
21:00:00:20:37:00:d5:e5
21:00:00:20:37:00:ca:dc
21:00:00:20:37:42:0e:ce
```

Figure 34: The Multiple WWNs of SecureSAN\_40

- Next, use *sectelnet* to access the Primary FCS Switch and create the new DCC policy. The administrators want to be able to insert a device into port 4 or port 5 on SecureSAN\_40 if necessary. The figure below shows the steps for completing the configuration.



```
SecureSAN28:admin> secPolicyCreate "DCC_POLICY_NewDisks","50:0e:09:e0:00:00:0f:23;50:0e:09:e0:00:00:11:90;SecureSAN_40(4,5)"
DCC_POLICY_NewDisks has been created.
SecureSAN28:admin> secPolicyAdd "DCC_POLICY_NewDisks","21:00:00:20:37:00:d5:e5"
Member(s) have been added to DCC_POLICY_NewDisks.
SecureSAN28:admin> secPolicyAdd "DCC_POLICY_NewDisks","21:00:00:20:37:00:ca:dc"
Member(s) have been added to DCC_POLICY_NewDisks.
SecureSAN28:admin> secPolicyAdd "DCC_POLICY_NewDisks","21:00:00:20:37:42:0e:ce"
Member(s) have been added to DCC_POLICY_NewDisks.
SecureSAN28:admin>
SecureSAN28:admin> secPolicyActivate
About to overwrite the current Active Policy Set.
ARE YOU SURE (yes, y, no, n): [no] yes
Committing configuration...done.
Saving Defined FMPS ...
done
Saving Active FMPS ...
done
SecureSAN28:admin>
```

Figure 35: Creating the New DCC for SecureSAN\_40

- ❖ Note: The *secPolicyAdd* command can add more than one WWN at a time. However, due to the single-line character limit, typically only three WWNs can be added at a time.

The administrators can now check to ensure that the other ports are properly secured. At this point, the administrators can run checks and make modifications as necessary.

- ❖ Keep in mind that with Fabric OS versions 2.6.0 and earlier, ports disabled using the *portDisable* command are not persistent through reboots.

## 4 Conclusions

The principle of Defense in depth aligns with the adage that “you are only as secure as your weakest link.” Each element in a comprehensive enterprise security solution relies on the strength of the other pieces. Traditionally, SAN security has only seen limited implementation and has relied on utilizing third-party products to bolster security. Today however, organizations demand a comprehensive SAN security solution capable of protecting mission critical fabrics. The Brocade Secure Fabric OS is designed to meet the comprehensive security needs of the SAN community.

The SFOS provides the SAN administrator with a whole suite of tools for implementing SAN security. In addition to core features, such as the centralized database policy management, that are active in secure mode by default, numerous other optional features are available to suit a wide variety of SAN security requirements. Remote access polices, switch connection controls, port lockdown features, and other security measures can mitigate a broad range of security concerns. The SFOS empowers SAN administrators with granular controls to design and implement policies customized to their organization’s needs.

While this document addresses a broad range of non-SAN related security topics and methodologies, keep in mind that there is no one correct process or “silver bullet” for designing and implementing security in the enterprise. It falls to each organization to apply the information to its own unique needs. To review, here are some of the topics covered in this document:

- General security threats
- High-level threat mitigation procedures and tools
- Management controls, operational controls, and technical controls
- Risk level assessment for enterprise elements
- Common risks to the SAN environment
- Minimizing vulnerabilities by implementing security in pre-SFOS fabrics
- Understanding the SFOS features
- Enabling secure mode and implementing the SFOS

Correct assessment and implementation of SFOS features can greatly enhance SAN security. An enterprise-wide security architecture that includes firewalls, IDSs, physical protection, and other controls can ensure a tightly secured SAN environment. Although security procedures require resources to implement and maintain, the potential business cost associated with high security exposure is often much greater. The Brocade Secure Fabric OS offers a comprehensive SAN security solution to organizations that rely on their storage area networks to provide secure, mission critical data and services.

## Appendix A: Securing the Enterprise

This appendix presents information contained in the National Institute of Standards and Technology (NIST) Special Publication 800-18 document. . The appendix addresses three the major concerns in the context of enterprise security – management controls, operational controls, and technical controls – and clearly outlines a methodology and next-steps.

### MANAGEMENT CONTROLS

Management controls ensure proper implementation and maintenance of protective security layers. These controls provide sufficient management oversight power level to filter business information into the controls so that each element is “secure enough” according to the organization’s perspective. These controls also govern security management and control quality levels.

### Risk Determination

When an organization faces the daunting task of securing its enterprise, it might be hard to know where to start. The answer is: start with what is known. In most business environments, time and money are in limited supply. Security resources should be commensurate with the organization’s security posture. That is to say, risk-averse organizations should dedicate more resources to security than a group with a higher risk threshold. The challenge is to create a plan that manages risk exposure and makes efficient use of resources. This challenge is not new to the IT world. One logical way to attack the problem is to perform a risk analysis on all elements requiring security attention and create a prioritized list.

A risk analysis is the process of evaluating a device, set of devices, or other element set in the enterprise against a consistent set of criteria to determine the extent of security required. In many cases, resource limitations suggest a multi-phase approach, in which the first phases address high-risk systems and subsequent phases provide additional protection until all security requirements are addressed.

No two organizations will apply the same process or criteria in assessing risk. Different architectures and business models call for different priorities. For example, an organization’s De-Militarized Zone (DMZ), file servers, and SAN might each garner different levels of attention. One methodology for determining risk is available in the document, *Special Publication 800-30: Risk Management Guide*, published by the National Institute of Standards and Technology (NIST). The guide provides a approach for assigning risk. Below is a summary of key points in the 800-30 methodology.

### Understanding the Components of Risk

The following information is helpful in understanding the components of risk:

#### System Characterization

Security architects and administrators must determine which systems or system groups are going to be considered and evaluated. Excluding any elements from analysis could create a sizeable hole in the final evaluation. Smaller organizations might choose to enumerate each of their chosen elements individually, while larger organizations might wish to group related elements together logically. In either case, it is important to choose and group the elements carefully. Below is a typical grouping for an enterprise.

- **Perimeter:** includes the border router, firewall, and any external network IDS sensors.
- **DMZ:** includes Internet-accessible devices segmented from the internal networks.
- **Desktops:** includes all user desktops.
- **SAN devices:** includes SAN storage, hosts, and switches.
- **IP infrastructure:** includes routers, switches, and other appliances that form the IP network.

Risk analysis can proceed when all elements or groups of elements have been determined. The diagram below shows how various factors contribute to a final risk evaluation. The information that follows explains each item in more detail.

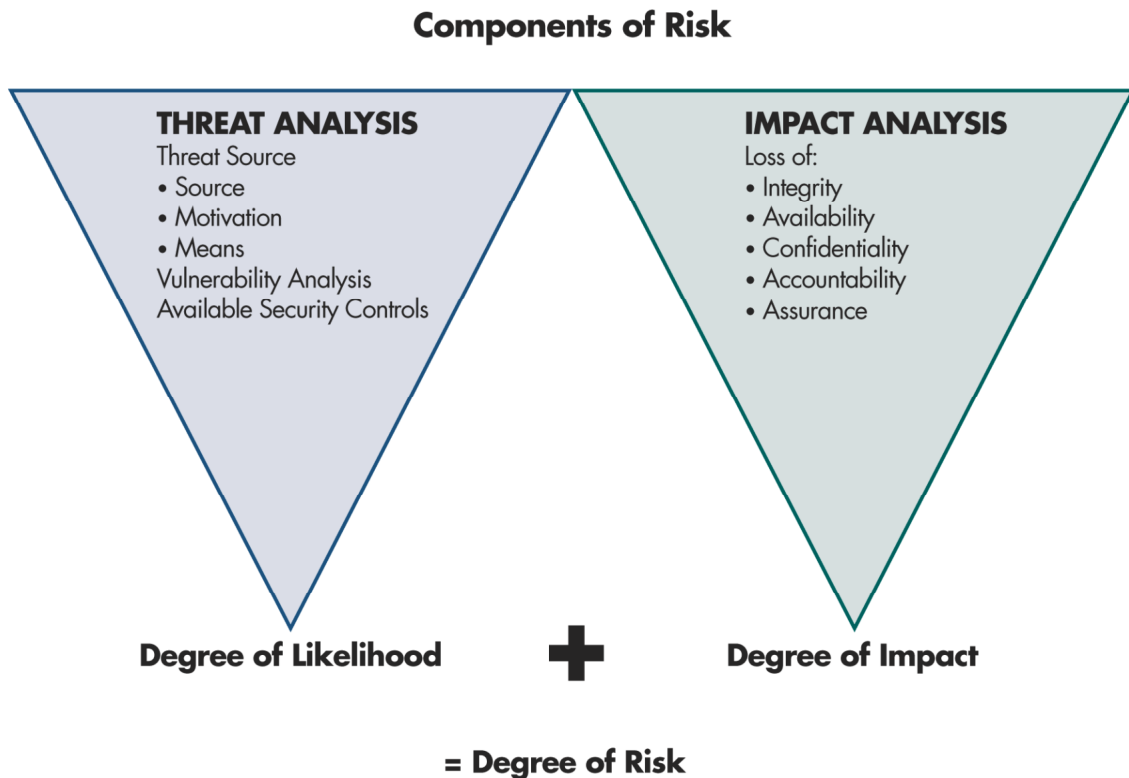


Figure 36: Risk Determination – Components of Risk

### Threat Analysis

Threat analysis refers to the process of determining the degree to which an entity or event could exploit a specific element’s vulnerabilities. The following steps help determine a final likelihood and threat analysis rating for each element.

- **Threat Analysis:** This step identifies who or what might pose a threat. Given the assets that comprise the system, is the problem more likely to happen due to an accident or intentional attack? The motivation and means of the threat source are also important contributing factors.
- **Vulnerability Analysis:** This step identifies the resiliency of an element in the event of an attack. Some systems, such as Web servers, present a significant number of severe vulnerabilities.
- **Control Analysis:** This step determines what means are either being used or available to protect the element from attack. Some systems have a variety of options available to them to secure the element, while others have very few.

- **Impact Analysis:** This step determines the impact of a security incident on core business functions and mission critical operations. The resulting value measures how critical a given element is to the successful continuation of business operations. The following items are helpful in grading these effects.
  - **Loss of Integrity:** Data might be changed and no longer be trustworthy
  - **Loss of Availability:** System or resource is no longer accessible and cannot perform its function, or is operating in a diminished capacity.
  - **Loss of Confidentiality:** Sensitive information might be or might have been viewed by unauthorized entities.
  - **Loss of Accountability:** Individual user actions can no longer be tracked reliably.
  - **Loss of Assurance:** The cumulative confidence associated with an entity based on its ability to control integrity, availability, confidentiality, and accountability.

### Using the Values to Determine Risk Level

The previous section presents a variety of parameters for determining risk values for the likelihood and impact of attacks. This process can be quite difficult and often involves both subjective and objective criteria. Whatever the approach, it is important to make the process as well defined and repeatable as possible. The combined values determine an overall risk value.

Designers and administrators can use a matrix or spreadsheet when combining likelihood and impact values. These values can vary in granularity. One organization might use a scale ranging from 1-10 in which 10 is of the utmost criticality; another organization might use a broader scale of 1-100 to allow for greater refinement and distinction. The chart below presents a sample risk determination matrix in which 1 represents lowest priority and 25 represents the highest priority.

**SAMPLE RISK LEVEL MATRIX**  
(1=LOWEST PRIORITY; 25=HIGHEST PRIORITY)

Impact Level	Likelihood Level (5=High; 3=Moderate; 1=Low)				
	5	4	3	2	1
5 (Critical)	25	20	15	10	5
4	20	16	12	8	4
3 (Moderate)	15	12	9	6	3
2	10	8	6	4	2
1 (Low)	5	4	3	2	1

**Table 21:** Risk Levels

The ultimate goal of the risk assessment process is to create a prioritized list that will help the organization determine how to distribute resources. In a spreadsheet used for the process, each of the elements would have values assigned for likelihood, impact, and risk level. The spreadsheet would then be easy to sort by final risk level values to create an ordered list. This priority list would be a useful starting point for making resource distribution decisions. Keep in mind, however, that a tool such as this should serve as a guide and not an absolute arbiter. It is important to apply a “sanity check” to the final results to ensure that they make intuitive sense.

## Review of Security Controls

Administrators must manage and maintain all security controls properly in order to ensure integrity and effectiveness. Periodic audits of security controls are important. Tools such as port scanners or vulnerability scanners are useful for testing defenses and determining necessary changes. These steps qualify as management controls if they are properly scheduled, conducted, and the findings acted upon. Administrators must coordinate tests among teams so that, for example, port scanning traffic doesn't alarm the IDS team. Scanning control initiatives require proper oversight to ensure success.

## Life Cycle Management

Life cycle management refers to the ongoing implementation of additional security measures. Developing a system from a list of requirements through purchase, testing, production operations, and eventually to disposal requires proper management and supervision. It is up to the organization's management to ensure that each step has been correctly completed and to assure consistent quality. Life cycle management ties in directly to a variety of other steps (presumably documented) such as requirements gathering, change control processes, QA testing, and the like.

## OPERATIONAL CONTROLS

Operation controls deals with the aspects of security that are enforced by people. Often the definition of how to enforce a policy or how to perform a procedure is documented in the organization's Security Policy. The following topics cover important aspects of operational controls.

### Physical Security

Physical protection of important infrastructure elements is absolutely critical to overall security. If attackers have physical access to a system, they can bring that system down (by pulling cables, powering it down, hitting it with a hammer, and the like), or they can bypass many security safeguards by connecting directly to a device. Systems should be installed in a lockable space that supports the environmental requirements of the various components. A small, unventilated closet might be easy to secure, but it could cause an appliance to overheat and fail.

The current generation of SANs are particularly vulnerable to physical attack or accidental configuration update. By attaching new switches or devices to an existing fabric, a potential attacker could:

- Remove or modify existing zoning information
- Access information on the fabric for reading, changing, or destroying data
- Cause a variety of DOS attacks

Many organizations use house their equipment in specially designed data centers. These facilities have backup power, lockable racks, cable management systems, and door lock mechanisms that require an individual ID card and log access information. Physical security also includes front desk guards, cameras, and other features typically associated with a locked down facility. It is important to keep in mind that inadequate physical security could render even a complex technical security infrastructure almost totally useless.

### Contingency Planning and Disaster Recovery

Although this topic could comprise many possible scenarios, suffice it to say that planning for a possible disaster is essential for most organizations. For example, if the emergency sprinkler system were to go off in a primary data center and information is not backed up at an off-site location, the impact would likely be severe and potentially fatal



for the organization. Administrators should refer to the organization's policies and documentation to develop contingency plans that suit the business needs of the company.

## Documentation

The need to document is the bane of many technical professionals, but a requirement for most businesses. Administrators and managers must balance the need for written policies and procedures with the need for efficiency. This section outlines the importance of some key elements of security documentation.

### Security Policy

The Security Policy document is of singular importance to an organization. It defines in writing a company's security posture and can provide a very comprehensive, detailed approach of security structure within the organization. In general, the Security Policy document should:

- Outline how an organization will to protect itself
- List what is and is not allowed
- Discuss appropriate and inappropriate data, traffic, access methods, and so on
- Outline security implementations (with appropriate levels of detail)
- Provide information on incident handling
- Provide a unified code that is consistently enforceable

A security policy can include additional elements. Refer to Appendix B for additional information and resources.

### Incident Response Plan

If and when a security incident occurs, the organization should have a pre-determined process in place for dealing with the problem. Incident handling can be intensive if the incursion goes deep into the infrastructure, and especially complex if legal action is required. Appropriate individuals in the organization must receive training and education to carry out the process. If an organization ignores this issue until after a security incident, greater damage could occur as the scrambles to deal with the problem. A variety of books, papers, and methodologies have been written on incident response. Refer to Appendix B for additional information and resources.

### Other Policies

Below is a summary of other important documents that relate to overall enterprise stability.

- **Acceptable Use Policy:** A document that outlines appropriate, allowed business uses for a resource. This document typically defines what kind of Internet Web use is acceptable and the consequences of misuse or abuse.
- **Change Control Policies:** These policies outline the steps required to make an alteration to the production environment. Such changes would include, adding or removing a device, upgrading the OS, adding patches, or any change that could impact the usability of a device or other production systems. These policies often provide a useful "sanity check" and can help minimize the potential impact of poorly planned implementations.
- **Testing Procedures:** These provide a standard method for determining if a system is functioning properly. Testing procedures are often tied into change control procedures for new production devices.
- **Backup Procedures:** Keeping proper backups of important business data is a fundamental part of good business practice. A well-defined backup procedure helps to ensure that data backups are timely, efficient, and accurate.

- **Emergency Procedures/Disaster Recovery Plans:** As mentioned, even in a well-designed security environment, accidents or natural disasters can occur. With procedures in place, administrators can handle unplanned emergencies more efficiently, thereby minimizing business interruption.

## TECHNICAL CONTROLS


Technical controls refer to security procedures that are enforced by hardware or software tools. Many such tools are commonplace in the enterprise and each represents a layer of security designed to mitigate specific risks. The following list and diagram outlines a representative environment that uses a variety of technical security controls to implement policies and to protect integrity, availability, confidentiality, authentication, authorization, and accounting.

Applying the defense in depth principle, some useful technical controls for architecting a comprehensive security suite include:

- **Border Router:** Often used as a first line of defense, these devices usually represent the outermost part of the perimeter and work well in concert with a firewall.
- **Firewall:** A device used to regulate traffic between an untrusted (or less trusted) area and a protected area. A carefully configured firewall will allow only the kind of traffic required by the business as defined in the organization's security policy.
- **Intrusion Detection System (IDS):** These are sensors that come in two basic flavors, host-based (HIDS) and network-based (NIDS). These sensors monitor data traffic and flag or react to potentially dangerous network patterns.
- **Virtual Private Networks (VPNs):** These are often used to protect (through encryption) data traveling over an untrusted or unsecured pathway, such as the Internet.
- **Out-of-Band (OOB) Networks:** These partition certain traffic types. An organization might use an OOB network to help prevent management traffic or backup traffic from being compromised or to offload some of the bandwidth impact on the rest of the production network. OOB networks must be carefully managed and locked down to avoid compromise.
- **Scanners:** Periodic scanning by authorized internal or consulting resources provides a useful way to benchmark security controls and test functionality. Administrators must coordinate these tests carefully and act upon the results in a responsible manner.
- **SAN Security:** Different methods are available for locking down the SAN environment. The Brocade Secure Fabric OS significantly enhances the options for security implementation.

The following diagram illustrates the technical controls described above.

## Defense In Depth: Technical Controls



**Figure 37: Technical Controls**

- ❖ Note: For additional information on the elements described in this section, refer to the additional links and references in Appendix B.

### Border Router and Access Control Lists (ACLs)

Border routers and ACLs can support a defense in depth strategy. Many organizations maintain control over their perimeter routing device, called a border router. Typically, this device provides a termination point for WAN connections and serves as the interface to the Internet and Telco networks. These border routers can be configured to resist attack. Commonly, an organization will configure the border router with ACLs to minimize the traffic that reaches the border firewall.

ACLs typically reside on routers, or they protect host devices (for example as “TCP wrappers”). While ACLs can offer excellent supplementary protection to an environment, they tend to be difficult to manage in large environments and lack the more sophisticated intelligence of a full-fledged firewall. ACLs can be very useful in minimizing IP address spoofing, dropping source-routed packets, and regulating other high-risk traffic types.

### Firewalls

Usually placed at network choke points, firewalls and ACLs provide administrators a way to regulate traffic. Typically, firewalls reside around the perimeter of the network and limit the types of traffic that can pass through

them. Firewall policy is based on the organization's security policy. There are different types of firewalls, each with its own set of pros and cons. Below is a brief overview of three common types of firewalls.

- **Static Packet Filters:** These systems can be characterized as “limited but fast.” Static filters compare each data packet against the firewall configuration to determine whether to pass or drop the data. These systems are not connection-oriented and as such do not use data they have already seen to determine which action to take.
- **Dynamic (Stateful) Packet Filters:** These filters enable firewalls to use previous data history, in addition to the firewall policy, to determine how to handle information. Dynamic filters tend to be a bit slower than static filters, but they provide significantly more security and flexibility.
- **Proxy Servers:** These are application-specific systems that provide a very high degree of control over traffic, but tend to be difficult to configure. Proxy servers literally break the client-server connection between systems and then carefully analyze the data passing through it. This process causes delays in the network traffic stream.

### **Intrusion Detection Systems (IDS)**

Intrusion Detection Systems (IDS) are common in many security architectures. Network-based Intrusion Detection Systems (NIDS) monitor network traffic and compare the data against known hostile patterns called “attack signatures.” Host-based Intrusion Detection Systems (HIDS) protect a specific monitored system. HIDS monitor network traffic destined for the host system and actions initiated on the system, then checks the local file systems for changes that match a pre-defined rule set. Both NIDS and HIDS can perform a variety of tasks when a pattern or action is identified as being hostile. Some of the most common of these actions are logging and alerting. Many IDS systems provide the additional flexibility of disrupting hostile sessions and altering filter rules on firewalls to block specific traffic.

IDS systems typically consist of two components:

- **Sensors:** HIDS and NIDS are the two available sensor options. These devices are remote listeners that record and analyze network traffic. Sensor placement is of paramount importance.
- **Console Station:** This is a system that controls and manages the remote sensors. The console might also store and aggregate (that is, correlate) data, targeting “low” and “slow” attacks. In some cases, the HIDS reside on the system being protected; however, installing HIDS on a central console station is more appropriate, particularly in large deployments.

Before deploying an IDS system, it is important to understand and evaluate the network design and implementation. Optimal placement depends on the network design and what the organization wants to protect. Very often, IDS systems are an integral part of the overall security architecture.

### **Virtual Private Network (VPN)**

Virtual Private Networks create protective “tunnels” for data across unsecured links. Common uses for VPNs are to connect remote offices with other sites, to provide an encrypted extranet between organizations, or to enable remote users to access sensitive corporate information more securely. VPNs typically use the Encapsulating Security Payload option within the IPSec protocol to provide data confidentiality, integrity, authentication, and anti-replay protection.

### **Out-of-Band (OOB) Networks**

An OOB network segments sensitive traffic from the rest of the network infrastructure. Multi-homed devices have one interface connected to the OOB network and are configured to listen for specific traffic types on that interface. Common uses for OOB networks include:

- **Management Traffic:** Enterprise management systems can use these networks to access information through SNMP or other means. To minimize exposure on more public interfaces, the polled devices can be configured to only listen to this traffic on the OOB network interface.
- **High Volume Traffic:** Certain high volume traffic can be directed to an OOB network to minimize the load on other production or highly utilized networks. The OOB segment can be used to copy large files, provide backups, and the like.
- **Remote Administrative Access:** Admin traffic (via telnet or other methods) can be physically segmented from other network traffic to provide an additional layer of protection. Although switches can limit some data snooping, it is possible to compromise switches and capture data on other switch ports. Fully separating management traffic makes it even more difficult to observe or modify this important data.

## Port Scanners and Vulnerability Scanners

An important aspect of management and operational controls is periodically validating the security of a system or group of systems. Two common tools used in systems testing are port scanners and vulnerability scanners.

- **Port Scanners:** These consist of a software package (such as, nmap) that typically has a variety of network parameters and that determines which ports of a target device are listening for traffic. Port scanners help determine which services or daemons are listening and help identify the OS of the device, (also known as OS/IP stack fingerprinting).
- **Vulnerability Scanners:** Vulnerability scanners are software tools (such as, CyberCop Scanner, NNESSUS, and so on) that determine if a target system is vulnerable to particular types of attacks. Often systems in an enterprise are patched or configured inconsistently, thereby allowing known attacks to compromise some systems and not others. Vulnerability scanners provide a useful way to audit systems and obtain a list of needed fixes.

Refer to Appendix B for links to additional resources.

## Appendix B: References

### BOOKS AND WEB RESOURCES

This appendix contains a list of useful references and resources for researching various topics. The list is far from exhaustive but serves as a starting point for finding additional information about some of the topics covered in this document.

#### Broad Security Information

- NIST: special publications (800 series and 500 series)  
<http://csrc.nist.gov/publications/nistpubs/index.html>
- SANs Reading Room  
<http://www.sans.org/infosecFAQ/index.htm>
- NSTISSC Library  
<http://www.nstissc.gov/html/library.html>
- SecurityFocus  
<http://www.securityfocus.com/>
- Lance's Security Papers  
<http://www.enteract.com/~lspitz/>

#### Technical Controls

- Brocade: SAN-related information  
<http://www.brocade.com>
- OpenSSH  
<http://www.openssh.org>
- TCP wrappers  
<http://www.porcupine.org/>
- Tripwire (Open Source)  
<http://www.tripwire.org/>
- Snort.org (Open Source IDS)  
<http://www.snort.org/>
- Nessus vulnerability scanner  
<http://www.nessus.org/>

- Nmap port scanner  
<http://www.insecure.org/nmap/>
- SAINT vulnerability checks  
<http://www.wwdsi.com/saint/>

### Router Lockdown

- Improving Security on Cisco Routers  
<http://www.cisco.com/warp/public/707/21.html>
- Secure IOS Template  
<http://www.cymru.com/~robt/Docs/Articles/secure-ios-template.html>

### Enterprise Architecture Guides

- Cisco Systems: “Cisco Safe: A Security Blueprint for Enterprise Architectures”  
[http://www.cisco.com/warp/public/cc/so/cuso/epsso/sqfr/safe\\_wp.htm](http://www.cisco.com/warp/public/cc/so/cuso/epsso/sqfr/safe_wp.htm)
- Sun Microsystems: “Building Secure N-Tier Environments”  
<http://sun.com.ar/blueprints/1000/ntier-security.pdf>

### Security Policies

- NIST: Special Publication on Internet Security Policy  
<http://csrc.nist.gov/isptg/>
- SANs Policy Issues  
[http://www.sans.org/infosecFAQ/policy/policy\\_list.htm](http://www.sans.org/infosecFAQ/policy/policy_list.htm)
- PentaSafe  
<http://www.pentasafer.com>

### Other Useful Sites

- RFC document search  
<http://www.rfc-editor.org/rfcsearch.html>
- TCP/UDP Ports Database List  
<http://www.portsdb.org/dump.html>
- Gibson Research Corporation  
<http://grc.com/default.htm>
- Google Groups: Usenet Search  
<http://groups.google.com/>

# Copyright

## IMPORTANT NOTICE

This document is the property of Brocade Communications Systems, Inc. (“Brocade”). It is intended solely as an aid for installing and configuring Storage Area Networks constructed with Brocade switches. This document does not provide a warranty to any Brocade software, equipment, or service, nor does it imply product availability. Brocade is not responsible for the use of this document and does not guarantee the results of its use. Brocade does not warrant or guarantee that anyone will be able to recreate or achieve the results described in this document. The installations and configurations described in this document involved third-party software and hardware. Brocade does not make any warranties or guarantees concerning such third-party software and hardware.

2003 Brocade Communications Systems, Inc.

ALL RIGHTS RESERVED.

(c) 2003 Brocade Communications Systems, Inc. All Rights Reserved.

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. 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.

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

Brocade Communications Systems, Incorporated  
1745 Technology Drive  
San Jose, California 95110  
U.S.A.

Part number: GA-RG-250-01