BROCADE

SECURING FIBRE CHANNEL FABRICS

SECOND EDITION

SAN Protection for Storage and Security Professionals

ROGER BOUCHARD

BROCADE

SECURING FIBRE CHANNEL FABRICS

SECOND EDITION

SAN Protection for Storage and Security Professionals

ROGER BOUCHARD

This book is dedicated to Nicole, my wife, whose support and understanding throughout the years would not have made this book possible. I would also like to offer a special dedication to Peter Carucci, a wonderful person, a father, and a husband, who left us all much too soon. Peter was an avid supporter of the Brocade encryption solution and SAN security assessment engagement and was instrumental to their success. He is dearly missed by all.

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 Bookshelf Series designed by Josh Judd

Securing Fibre Channel SANs
Written by Roger Bouchard
Edited by Victoria Thomas
Design and Production by Victoria Thomas
Illustrated by Jim Heuser, David Lehmann, and Victoria Thomas

Contributors: Josh Judd (SAN basics), Marcus Thordal (Brocade Encryption Switch), Scott Kipp (key management), Jim Davis (zoning), and Thomas Scheld and Martin Sjoelin (lab experiments for myths)

Reviewers: Greg Farris, Tom Clark, Josh Judd, Marcus Thordal, Scott Kipp, Jim Davis, and Mark Dietrick

Printing History

First Edition, April 2009
First Edition, Rev. A, September 2009
Second Edition, May 2012

Important Notice

Brocade Corporate Headquarters
San Jose, CA USA
T: (408) 333 8000
info@brocade.com

Brocade European Headquarters
Geneva, Switzerland
T: +41 22 799 56 40
emea-info@brocade.com

Brocade Asia Pacific Headquarters
Singapore
T: +65 6538 4700
apac-info@brocade.com

DISCLAIMER

The author is not an attorney and this book in no way represents any legal advice or legal opinion. For legal advice or opinion on data protection measures, consult an attorney.

Acknowledgements

Very special thanks go to Martin Skagen, my friend and Brocade mentor, for his generosity in sharing his extensive technical knowledge with me and his support of my advancement in SAN security. Special thanks to Victoria Thomas, the copyeditor for the first edition, and Patty Barkley, whose support and creativity made the second edition of this book possible. Special thanks to Greg Farris who was the primary proofreader for the second edition and committed many hours to ensure the quality of this edition.

This book would not have been possible without the help of several other contributors and reviewers that shared their knowledge and expertise. For this, I would like to thank the following contributors for the first edition of this book: Tom Clark, Josh Judd, Marcus Thordal, Scott Kipp, Jitendra Singh, Jim Davis, Thomas Scheld, Martin Sjoelin, and Mark Dietrick. Finally, thanks to Ron Totah who provided me with the opportunity to dedicate the time and created the environment essential to complete this project.

About the Author

Roger Bouchard has been in the computer industry since 1978 with a wide range of experience in programming, analysis, consulting, education and management. He has taught IT security courses since 1994 and has been focused exclusively on the storage industry since 1996.

Since Mr. Bouchard joined Brocade in 2000, he has obtained his BCFP, BCSD, and BCSM certifications as well as the CISSP certification in 2005 and an M. Sc. in Information Assurance (MSIA) from Norwich University. His role evolved within the company from a Sales Engineer (SE) Subject Matter Expert (SME) on Security to founding and leading the Security Practice in the Services organization. There he developed processes for SAN Security Assessments and SAN Hardening engagements delivered across North America.

He is currently a Global Solutions Architect, and in this role has written several white papers on SAN security and is a frequent speaker at storage/SAN conferences.

Contents

1 Introduction

As today's IT organizations face more and greater security threats with a growing number of industry and government regulations, securing SAN environments has become an increasingly important aspect of overall data security. This is especially the case as storage area networks continue to grow in size and extend across multiple sites. A key factor in security is that many SANs use protocols other than the Fibre Channel (FC), with many different protocols now carrying storage traffic. Some are upper-level protocols (such as FICON in the mainframe world) while others run over IP (such as Fibre Channel over IP (FCIP) for tunneling FC between sites and iSCSI for fanning out to low-cost servers). The introduction of FC over Ethernet (FCoE) protocol based on the Data Centre Bridging standard, also introduces new security concerns in the SAN.

At a very basic level, security measures need to balance the probability of a threat occurring, the impact of a security breach, the cost of implementing countermeasures, and the value of the assets. The tolerated risk level varies significantly from one organization to another and depends on several factors. It is often dictated by government legislation and industry standards targeted at specific verticals, such as:

- Gramm-Leach Bliley Act (GLBA) for the financial and insurance industries
- Health Insurance Portability and Accountability Act (HIPAA) guidelines for the healthcare industry and the HITECH Act of 2009
- Payment Card Industry Data Security Standard (PCI-DSS) for companies dealing with large volumes of credit card transactions

Other countries also regulate the privacy of information, such as:

- Canada's Personal Information Protection and Electronic Documents Act (PIPEDA)

- The European Union (EU) Data Protection Directive (EU Directive 95/46/EC)
- The Monetary Authority of Singapore

Some legislation is regional such as the precedent-setting California Senate Bill (SB) 1386 and similar laws currently in effect in 46 states[1] at the time of writing. This legislation requires organizations to disclose security breaches of unencrypted personal information belonging to their state residents. This means that a security breach might be made public and have serious business consequences, including customer attrition and loss of brand equity.

Regardless of the specific legislation, the more valuable the data is to an organization, the lower the tolerated risk level will be when it comes to protecting it. This trend will most likely continue, especially as data security becomes an increasingly global issue. SAN security can no longer be overlooked by security and storage professionals, since every day the volume and value of mission-critical data in their storage environments increases. In fact, judging by a significant increase in SAN security-related questions since 2010, SAN security has taken on a more visible role with security professionals.

[1] National Conference of State Legislatures website (March 25 2012): http://www.ncsl.org/programs/lis/cip/priv/breachlaws.htm

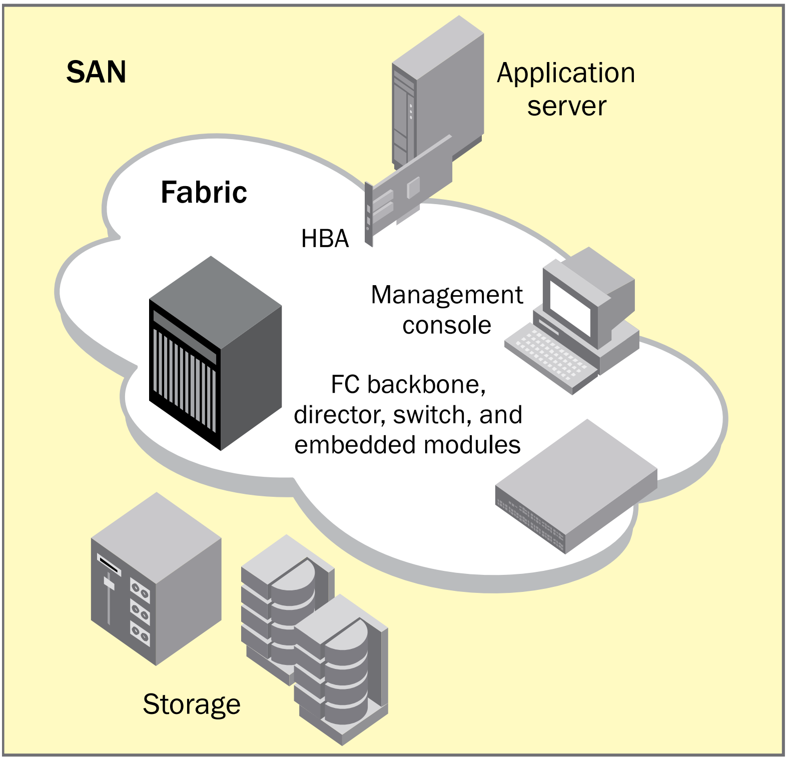

Figure 1. Fabric and SAN

The storage area network (SAN) has been defined in many ways and the limits of where it begins and ends can vary depending on an individual's or organization's perspective. For the purpose of this book, a fabric, often depicted as a cloud in illustrations, refers to the Fibre Channel infrastructure that makes up a storage network, namely, the FC switches, directors, routers, and backbone devices. The Host Bus Adapter (HBA) on the host and the storage controllers are also included in this definition. The *SAN* includes the fabric (network infrastructure) and the storage devices on which the data resides, including disk arrays, tape libraries, and both disk and tape media. Figure 1 illustrates a simple fabric and SAN.

This book discusses the actual data residing on the SAN (classic data protection concepts) at a high level only—it mainly addresses the issue of data confidentiality. For those interested in information about protecting data in greater depth, consult an excellent book entitled, *Strategies for Data Protection*, First Edition, 2008, by my late esteemed colleague, Tom Clark.

The SAN Security Dilemma

The individuals responsible for managing the storage environment typically have a limited knowledge of IT security. In many cases, security is actually viewed as an impediment to performing the daily activities of the storage and SAN administrators.

Conversely, the individuals responsible for ensuring security of information are generally less knowledgeable in the storage and SAN environment than they are of a conventional TCP/IP-based local area or wide area network (LAN or WAN). There is often an assumption on both sides of the fence that the SAN does not need to be secured, since it is a closed and physically protected environment that is not accessible to outsiders. Although this line of thinking is not entirely false, a closed environment does not offer any protection against attack from insiders, which poses the primary threat against a SAN and storage environment.

IT managers and decision makers with limited IT budgets need to make important choices regarding which projects receive funding and how much of the budget goes to each project. Network security certainly receives a great deal of attention and funding, but there is still a misconception that SANs and storage require only minimal security measures, since they are isolated from the outside world and protected from outside threats. As suggested earlier, outsiders are not the primary threat to a SAN but insiders, whether malicious or otherwise, pose the greatest threat. More and more cases are reported of insiders stealing backup tapes or disk drives containing sensitive company information such as medical, research, financial, and customer information. Many cases have been reported of employees actually copying information and taking it with them before they leave their employer and then selling it to criminal elements or using the information in their next position with a competitor. A black market has developed for certain types of data, particularly credit card and other financial information. Organized crime has become increasingly involved in cybercrime, as well as foreign government-sponsored hacker groups seeking to steal highly confidential information, such as intellectual property, in order to quickly gain a competitive advantage.

The second edition of this book is primarily intended to continue to raise awareness among the storage, security, and IT management professionals of the need to secure their SANs. If successful, understanding more about security issues raised in this book will help bridge the knowledge and cultural gap between the storage and security groups within an organization, which in turn will help IT managers better understand the risks and potential liability issues associated with their SAN.

To accomplish this, basic security concepts are introduced for those overseeing the storage environment and then basic storage concepts are presented to those involved in securing IT assets and electronic information. Of value to IT managers may be a review of some of the regulations and legislation in effect throughout the United States and other countries and how they apply to the SAN environment. With the advent of clouds, the storage and networking teams are more likely to work more closely together to protect assets and information residing within the cloud.

Although this book is focused primarily on Brocade® B-Series (classic Brocade) and M-Series (formerly McDATA) technology, the basic SAN security principles introduced here can be applied to any fabric or storage environment regardless of the vendor implementation. While there may be differences in feature availability and implementation among vendors, the general concepts and requirements are comparable. The information in this book is based on current research being performed by many organizations (full list in Appendix B) and real-world experience gained from performing actual security assessments, audits, and hardening engagements with Brocade customers throughout North America.

Why SAN Security?

Although SAN security is a specialized field dealing with issues specific to the storage industry, it follows the same established principles found in all modern IT security. It involves the continuous process of evaluating an environment's current state of security against the constant evolution of technology and an increase in awareness concerning security issues. As a result, a SAN security strategy is integral to an overall IT security strategy and should address all possible threats facing data within a SAN environment.

Since 2002, Brocade has been a leader in Fibre Channel SAN security. Based on years of real-world experience deploying SANs of varying sizes and architectures, Brocade developed a special licensed version of Fabric OS® (FOS), called Secure Fabric OS, designed to meet the specific requirements of the most security-sensitive environments. For instance, Brocade introduced the first access control lists (ACLs) in the Fibre Channel industry and provided the first Fibre Channel authentication mechanism using Public Key Infrastructure (PKI), which has since been replaced with the standards-based DH-CHAP (Diffie Hellman - Challenge Handshake Authentication Protocol), a forthcoming

Internet standard for the authentication of devices connecting to a Fibre Channel switch, as defined in the FC-SP/FC-sec standard specification defined by the ANSI T11 committee.

Most of the security features originally available in Secure Fabric OS have since been replaced with either equivalent or more powerful and flexible functionality in the base Fabric OS (version 5.3.0 or later), so they no longer require a special license. Appendix A provides a comprehensive list of technical security features that can be implemented in a Brocade-based SAN environment. As new security vulnerabilities are discovered or required, Brocade is continually enhancing existing features and creating new security features to help ensure that FC fabric infrastructures and data moving through them remain secure and highly available.

Security represents a delicate balance among factors such as the type of threats and risks, the likelihood that a vulnerability can and will be exploited, the effort and cost associated with implementing countermeasures, the impact on fabric management, and the value of the asset being protected. With more than 100 FC fabric security features available, not all features available should be implemented in all environments. Different organizations have different security requirements and levels of tolerance to risk. A detailed analysis and assessment of the state of security for a given environment should be performed to fully understand the risks and how to best mitigate them. There should be enough detailed information in this book to gain the knowledge necessary to conduct this assessment. Nevertheless, there may be advantages in hiring the services of a third-party organization with expert knowledge in the subject as is frequently done with conventional TCP/IP-based networks. Brocade offers such a service to help customers evaluate and assess the current state of security of their SAN.

Who Needs to Know About SAN and Storage Security?

SAN security is a relatively new field and many organizations have just begun to consider and integrate this area of security into their organization. Many stakeholders within organizations now have an interest in SAN and storage security. Each individual may be interested in different aspects of SAN security to different degrees and to different depths but SAN security can no longer be ignored. The roles of these stakeholders is varied as you can see in the following list,

Chief Executive Officer (CEO). The company CEO obviously has a high-level concern for SAN and storage security, but her focus is in two specific areas: the potential liability resulting from security breaches and non-compliance with industry and government regulations. For some executives, liability may in fact equate to jail time.

Chief Information Officer (CIO). The CIO is usually responsible for the IT department, which owns network, SAN, storage, and other technologies. Protecting these assets and minimizing risk and liability due to a security breach is paramount in this role. This role may extend beyond the technology and, in some cases, may include the actual information that is stored, processed, and managed.

Chief Financial Officer (CFO). The CFO is typically concerned from a compliance and regulatory perspective. The auditing department often falls under the CFO, making sure that appropriate controls are in place to guide the construction of policies and enforce them.

Chief Compliance Officer (CCO). The CCO's role is to ensure that the company is complying with local, state, federal, and industry regulations. He reviews the various regulations and creates the necessary programs to comply with these regulations. He often works in collaboration with the audit team to ensure that all regulations are being followed. The CCO frequently reports to the CFO.

Chief Security Officer (CSO)/Chief Information Security Officer (CISO). The CSO, or the CISO, is directly responsible for the protection of the IT assets and sometimes this extends to protecting all company assets including facilities and personnel. The SAN is of particular concern to the CSO/CISO since the data residing on it is one of the company's most valuable assets.

IT Security Director/Manager. The IT security director or manager's primary concern is with the IT assets, applications, and personnel that she is responsible for. Her concern with the SAN and storage environment is more detailed and she is responsible for implementing many of the controls and policies established by the C-level executives.

Security Professional. The security professional can be responsible for creating security policies, implementing security measures, managing the security aspects of the IT environment, monitoring the state of security of the IT environment, and responding to security incidents. He should have a direct involvement in the SAN and storage security just as he would with the corporate LAN and server environment. Quite often, the security professional will conduct audits or penetration tests/scans on the FC SAN to detect possible vulnerabilities, an increasingly common practice in the past two years.

Storage Professional. The storage professional, which includes SAN administrators, storage administrators, backup administrators, operators and managers, is more concerned with following the security policies during the course of their daily activities managing and running the storage environment. The storage professional will often be called upon by the security team to provide advice on how to implement specific security measures in a SAN and storage environment. Questions the storage professional may be required to answer include: “What is the best way to encrypt backup data on tapes?” and “Which secure protocols can be used to securely manage the SAN switches and storage devices?”

Within a given organization, many individuals will be involved in SAN and storage security at different levels. Each has a vested interest in the due diligence and care required to protect the data residing on the SAN environment. The storage professional is often asked to respond to questions from the security team regarding the security of their SAN. They frequently need to respond to the results of penetration tests or audits, as well as general security concerns around the SAN.

Chapter Summary

In conclusion, although a significant gap has existed between the storage and security worlds, both sides are learning from each other as organizations are faced with more compliance, regulations, and attacks on their electronic data. Organizations such as SNIA (Storage Networking Industry Association), SSIF (Storage Security Industry Forum), IEEE, and OASIS (Organization for the Advancement of Structured Information Standards) are developing best practices and standards to help address these issues.

2

SAN Security Myths

Over the past ten years, this writer has had the opportunity to discuss SAN and storage security issues with thousands of security and storage professionals as well as IT managers and decision makers. These people represent businesses and industries spanning the entire spectrum from financial to health to telecommunications and also government, military, and intelligence communities. Although each organization has its own unique perspective on the subject of SAN security, some issues are common to all groups. Some people immediately understand the need for SAN security and recognize the hole in their IT security strategy. At the other extreme, others simply believe that there is no need to address SAN security at all.

Several misconceptions have developed from the early days of the SAN, which unfortunately have become integrated and accepted into IT folklore and culture and are now perceived as fact. As with all folklore, myths can persist over time and take on a life of their own. Although they can be entertaining to some, it is important to understand the true facts so as not to fall into the “security through obscurity” way of thinking. This line of thinking can lead to a false sense of confidence in the security of a SAN environment. There is nothing more humbling to an organization than an actual security breach that becomes public and highly visible, creating a huge impact on customer and market perception. Hopefully, most organizations will take the bull by the horns and address SAN security issues before tragedy strikes or they become the next target of Wikileaks.

As a result of the writer's contact with real-world people and environments, some of these myths were identified to raise awareness and set the record straight. Although these myths may be quite entertaining to some readers, others may find them quite disconcerting.

SAN Security Myth Number 1

Myth. SANs are inherently secure since they are in a closed, physically protected environment.

Reality. It is generally true that a SAN is installed within a secure, access-controlled data center. Appropriate physical security measures help prevent unauthorized outsiders from gaining access to the computer equipment.

However, most security incidents affecting storage and SAN environments are attributed to insiders or outsiders with the assistance of insiders. Adequate physical security does not prevent insiders from causing security breaches, intentional or otherwise. Protection against insider threats is most likely the greatest challenge facing security professionals. However, specific measures can be implemented to help prevent or mitigate the risks associated with insider threats. It is important to note that insider threats are not always malicious; in fact most often they are not. On the issue of employee trust, employees can unexpectedly “go rogue” and it is difficult to predict which ones will do so.

Finally, most insider incidents are the result of errors during the course of daily operations. Measures can be implemented to reduce the number of errors and to mitigate the risks associated with them, such as using well-documented procedures and monitoring tools.

SAN Security Myth Number 2

Myth. The Fibre Channel protocol is not well known by hackers and there are almost no avenues or tools available to attack fabrics.

Reality. There is unquestionably some merit to this statement and FC-based networks are undoubtedly more secure than conventional TCP/IP networks. For many reasons, some organizations prefer to separate storage traffic from production traffic on isolated networks specifically for this purpose. This exemplifies the concept of separation of duties and isolating different functions from one another within a common environment. There is also some value in utilizing different technologies in the same environment. An attacker with a malicious intent may be quite knowledgeable about TCP/IP networks and would be able to get past the first hurdle but would be stumped when reaching the FC network, hampered by lack of skill with this technology.

Nevertheless, every FC device uses the TCP/IP protocol for management interfaces. Given that TCP/IP is well known by the “black hat” (hacker) community and many exploits are readily available on the

Internet for free, significantly less effort is required to conduct an attack on a SAN device management interface to compromise the SAN from this entry point. For this reason, it is important to apply similar security practices normally used in conventional TCP/IP networks to secure these interfaces, such as the use of secure communications channels (SSH, SSL, etc.), VPNs, and RBACs.

Additionally, some SANs use other protocols based on TCP/IP, such as iSCSI and FCIP. iSCSI is commonly used as a low-cost, lower performance SAN solution and allows organizations to leverage existing TCP/IP infrastructure. FCIP is usually used to connect two or more data centers over distance to enable replication of data or to perform remote backups. Converged networks using IP and FC would also present a new set of problems that extend to both protocols. A vulnerability in one protocol could potentially open the door to the other protocol.

Both iSCSI and FCIP use the TCP/IP protocol and the usual security measures deployed to protect the LAN or WAN should also be implemented with these protocols.

SAN Security Myth Number 3

Myth. You can't “sniff” optical fiber without cutting it first and causing disruption.

Reality. Devices that “sniff” are called “sniffers” (network data monitoring tools). There are several devices that can sniff optical fiber right through the jacket without requiring a splice in the cable. A simple microbend in the cable allows enough light to leak out of the jacket to be captured by a highly sensitive photosensor. Many of these devices can be purchased on the Internet for under $1,000 US. For a data center administrator, the only noticeable change would be a slight decibel (dB) loss on the signal, but dB loss is not usually monitored regularly by SAN administrators. Even if it were detected, it could be attributed to normal dB loss or fluctuations from the distance solution provider.

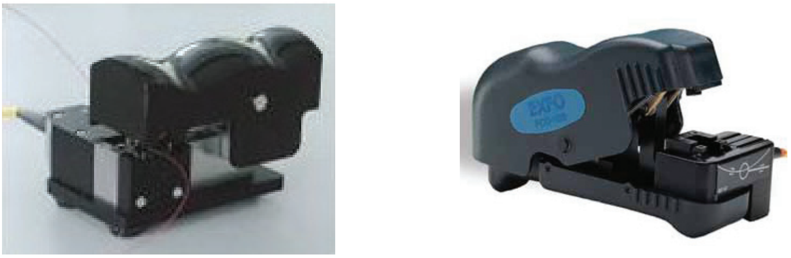

Figure 2. Examples of optical fiber sniffers

SAN Security Myth Number 4

Myth. The SAN is not connected to the Internet so there is no risk from outside attackers.

Reality. This may be the case in many organizations, but consider email and Web application servers that are placed in a demilitarized zone (DMZ). The DMZ is a “safe” zone protected by a series of firewalls in which one end of the application host is connected to the internal production network (sometimes referred to as the clean network) and the other end is connected to the Internet (sometimes referred to as the dirty network). It is entirely possible for a server within a DMZ with SAN-attached storage to be used as an entry point into the SAN from the outside unless the proper precautions have been taken.

This can be done safely but special precautions must be taken to do so. An entire section in this book is dedicated specifically to this complex issue. The reader may refer to Chapter 7 of this book for further details on securing Fibre Channel-attached devices in a DMZ.

SAN Security Myth Number 5

Myth. Even if fiber cables could be sniffed, there are so many protocol layers, file systems, and database formats that the data would not be legible in any case.

Reality. This is simply not true. Although some data may be difficult to read, a considerable amount of data is in pure ASCII format. One argument is that if the data is compressed, it is therefore unreadable. Compression algorithms are well known and the data can be uncompressed easily using one or other of these algorithms. Another argument is that data may be formatted using non-ASCII coding as in databases or specific applications. Some data may certainly be stored in non-ASCII format in various databases and applications, but a significant amount of data remains in ASCII format. Think of a credit card number or a social insurance number, for instance. These types of information are only a few bytes in size and would easily fit into a standard FC frame.

A simple experiment was performed to demonstrate this using a basic SAN and an inexpensive software-based FC trace analyzer in lieu of using an inexpensive sensitive photosensor, as described in Myth Number 3. The diagram in Figure 3 (SAN Test A) illustrates the setup of the test equipment used for this experiment.

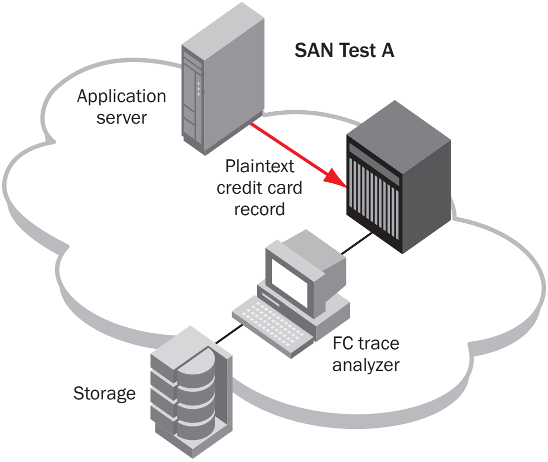

Figure 3. Sniffing FC frames in a SAN

In this setup, a fictitious credit card record was created and sent in unencrypted cleartext to the FC fabric, where the record was then written to the disk array. The FC trace analyzer captured and recorded the FC frames involved in this transaction. The screen capture from the trace analyzer, shown in Figure 4, displays the different frames captured. The frame containing the credit card record is highlighted with the box.

Look at the bottom-right corner of the screen capture to see the ASCII version of the frame contents, within the circled portion of the screenshot. The credit card record can clearly be read from this screen.

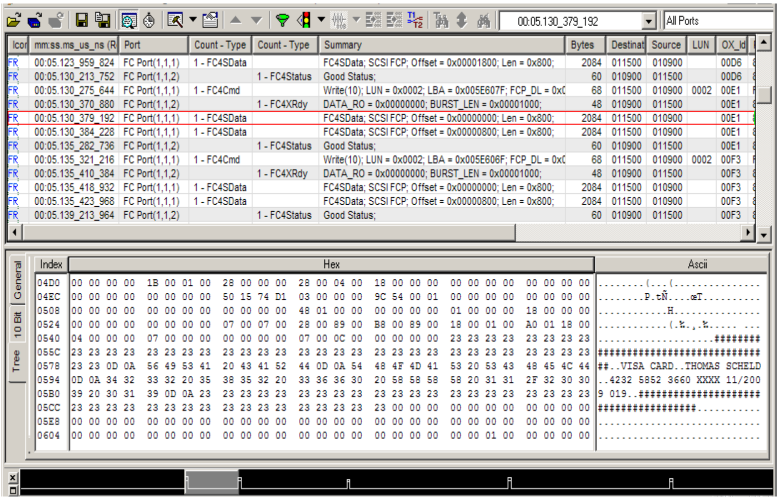

Icon	mm:ss.ms_us_ns (R	Port	Count - Type	Count - Type	Summary	Bytes	Destinat	Source	LUN	OX_Id
FR	00:05.123_959_824	FC Port(1,1,1)	1 - FC4SData		FC4SData; SCSI FCP; Offset = 0x00001800; Len = 0x800;	2084	011500	010900		00D6
FR	00:05.130_213_752	FC Port(1,1,2)		1 - FC4Status	Good Status;	60	010900	011500		00D6
FR	00:05.130_275_644	FC Port(1,1,1)	1 - FC4Cmd		Write(10); LUN = 0x0002; LBA = 0x005E607F; FCP_DL = 0x	68	011500	010900	0002	00E1
FR	00:05.130_370_880	FC Port(1,1,2)		1 - FC4XRdy	DATA_RO = 0x00000000; BURST_LEN = 0x00001000;	48	010900	011500		00E1
FR	00:05.130_379_192	FC Port(1,1,1)	1 - FC4SData		FC4SData; SCSI FCP; Offset = 0x00000000; Len = 0x800;	2084	011500	010900		00E1
FR	00:05.130_384_228	FC Port(1,1,1)	1 - FC4SData		FC4SData; SCSI FCP; Offset = 0x00000800; Len = 0x800;	2084	011500	010900		00E1
FR	00:05.135_282_736	FC Port(1,1,2)		1 - FC4Status	Good Status;	60	010900	011500		00E1
FR	00:05.135_321_216	FC Port(1,1,1)	1 - FC4Cmd		Write(10); LUN = 0x0002; LBA = 0x005E606F; FCP_DL = 0x	68	011500	010900	0002	00F3
FR	00:05.135_410_384	FC Port(1,1,2)		1 - FC4XRdy	DATA_RO = 0x00000000; BURST_LEN = 0x00001000;	48	010900	011500		00F3
FR	00:05.135_418_932	FC Port(1,1,1)	1 - FC4SData		FC4SData; SCSI FCP; Offset = 0x00000000; Len = 0x800;	2084	011500	010900		00F3
FR	00:05.135_423_968	FC Port(1,1,1)	1 - FC4SData		FC4SData; SCSI FCP; Offset = 0x00000800; Len = 0x800;	2084	011500	010900		00F3
FR	00:05.139_213_964	FC Port(1,1,2)		1 - FC4Status	Good Status;	60	010900	011500		00F3

```
Index  Hex                                                                                                    Ascii
04D0   00 00 00 00  1B 00 01 00  28 00 00 00  28 00 04 00  18 00 00 00  00 00 00 00  00 00 00 00   ........(...(...............
04EC   00 00 00 00  00 00 00 00  50 15 74 D1  03 00 00 00  9C 54 00 01  00 00 00 00  00 00 00 00   ........P.tÑ....œT..........
0508   00 00 00 00  00 00 00 00  00 00 00 00  48 01 00 00  00 00 00 00  01 00 00 00  18 00 00 00   ............H...............
0524   00 00 00 00  00 00 00 00  07 00 07 00  28 00 89 00  B8 00 89 00  18 00 01 00  A0 01 18 00   ............(.‰.¸.‰..... ...
0540   04 00 00 00  07 00 00 00  00 00 00 00  07 00 0C 00  00 00 00 00  23 23 23 23  23 23 23 23   ....................########
055C   23 23 23 23  23 23 23 23  23 23 23 23  23 23 23 23  23 23 23 23  23 23 23 23  23 23 23 23   ############################
0578   23 23 0D 0A  56 49 53 41  20 43 41 52  44 0D 0A 54  48 4F 4D 41  53 20 53 43  48 45 4C 44   ##..VISA CARD..THOMAS SCHELD
0594   0D 0A 34 32  33 32 20 35  38 35 32 20  33 36 36 30  20 58 58 58  58 20 31 31  2F 32 30 30   ..4232 5852 3660 XXXX 11/200
05B0   39 20 30 31  39 0D 0A 23  23 23 23 23  23 23 23 23  23 23 23 23  23 23 23 23  23 23 23 23   9 019..#####################
05CC   23 23 23 23  23 23 23 23  23 23 23 23  23 23 23 23  23 00 00 00  00 00 00 00  00 00 00 00   #################...........
05E8   00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00   ............................
0604   00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00  00 00 01 00  00 00 00 00   ............................
```

Figure 4. FC trace analyzer screen

SAN Security Myth Number 6

Myth. Even if fiber cables could be sniffed, the amount of data to capture is simply too large to capture realistically and would require expensive equipment to do so.

Reality. The amount of data to capture can be reduced using simple intelligent filtering technology. Furthermore, given the relatively small size of a credit card or social insurance number, a considerable amount of information can exist inside a single frame's payload. From a security perspective, the greatest challenge with sniffing networks is that this security breach may go unnoticed. An attacker can literally steal information off the cable and no one would be aware of the breach. Suppose that somehow the breach is detected; it would be difficult to know with certainty which specific records were stolen and the entire database would have to be deemed compromised. In the case of credit card information, this could result in the cancellation and a reissue of all credit cards, which could number in the hundreds of thousands or millions in some cases.

There is also the possibility of an insider walking away with a tape drive or disk media containing sensitive information. There have been several documented cases of this type of breach, highlighting the importance of encrypting the data-at-rest on the storage media.

To demonstrate how easy it is to capture sniffed data and rebuild an entire file, a second experiment, similar to the one in Myth Number 6, was conducted. This time, another feature was exploited - the ability to mirror a port. A simple laptop was used to store the frames captured by the FC trace analyzer.

The storage array port (port 0) was mirrored to port 15 using the port mirroring feature built into the switch, as shown on the screen capture in Figure 5 and the diagram in Figure 6.

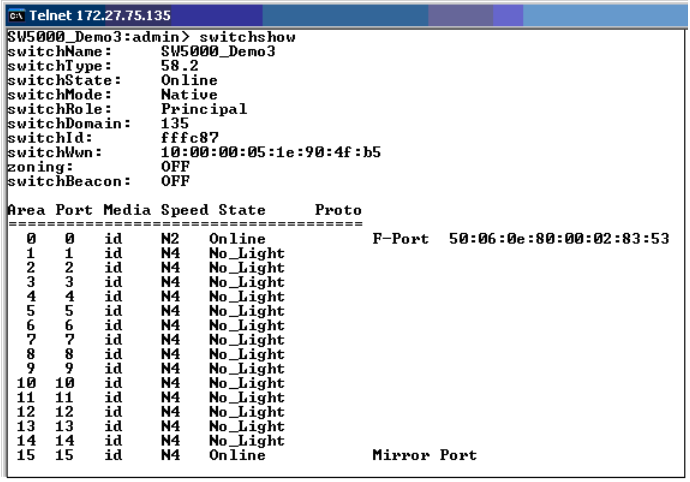

```
SW5000_Demo3:admin> switchshow
switchName:     SW5000_Demo3
switchType:     58.2
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   135
switchId:       fffc87
switchWwn:      10:00:00:05:1e:90:4f:b5
zoning:         OFF
switchBeacon:   OFF

Area Port Media Speed State     Proto
=====================================
  0   0   id    N2   Online           F-Port  50:06:0e:80:00:02:83:53
  1   1   id    N4   No_Light
  2   2   id    N4   No_Light
  3   3   id    N4   No_Light
  4   4   id    N4   No_Light
  5   5   id    N4   No_Light
  6   6   id    N4   No_Light
  7   7   id    N4   No_Light
  8   8   id    N4   No_Light
  9   9   id    N4   No_Light
 10  10   id    N4   No_Light
 11  11   id    N4   No_Light
 12  12   id    N4   No_Light
 13  13   id    N4   No_Light
 14  14   id    N4   No_Light
 15  15   id    N4   Online           Mirror Port
```

Figure 5. Port mirroring screen capture

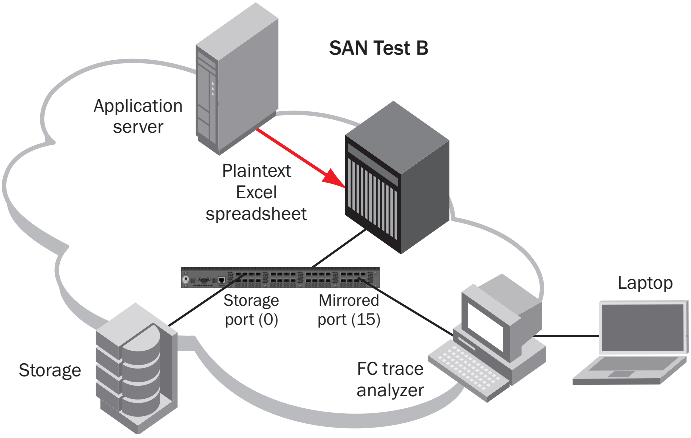

Figure 6. Sniffing FC frames using a mirrored port and storing them on a laptop

An FC trace analyzer was attached to the mirrored port (port 15) and captured the data going to the storage array. Again, as in the previous experiment, an inexpensive photosensor could be used for this demonstration.

A spreadsheet containing fictitious payroll information was sent across the SAN from the host to its LUN (logical unit number) on the disk storage array. The data captured by the trace analyzer was dumped to a binary file on a laptop, which was subsequently used to reconstruct a second disk. The file system on the new disk was mounted and the spreadsheet, shown in Figure 7, could be read as though it were the original copy.

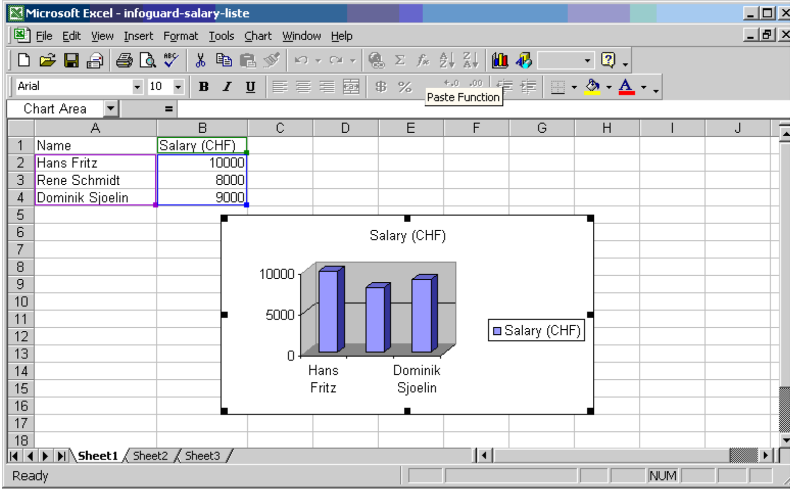

	A	B
1	Name	Salary (CHF)
2	Hans Fritz	10000
3	Rene Schmidt	8000
4	Dominik Sjoelin	9000

Figure 7. Excel spreadsheet reconstructed from FC sniffing

Clearly, FC frames can be selectively captured from a Fibre Channel network and easily stored on a storage device as simple and compact as a laptop. The information contained in the captured frames can be used to reconstruct entire files. Of course, even partial files or partial information could contain enough sensitive data to result in a significant security breach.

SAN Security Myth Number 7

Myth. If the switches already come with built-in security features, why should I be concerned with implementing security features in the SAN?

Reality. Similar to any other IT product, many of the built-in security features are not enabled by default. For example, Brocade switches have about 100 security-specific features available (see "Appendix A: Fabric OS Security Features Matrix" starting on page 199), but very few of them are enabled when the switch is installed out of the box. No two organizations have the same business or security requirements and each has a different risk tolerance level. Even when over 100 security features are available that doesn't mean that all 100+ features must or should be implemented in a given environment.

A careful risk analysis and a comprehensive assessment of the state of security of a SAN environment should be performed first. Subsequently, a SAN security policy should be developed, which will become the blueprint for implementing appropriate countermeasures for that environment. The cost and impact of implementing certain countermeasures on a production environment should be factored in. If the cost of implementing the countermeasures and negative impact on performance or operational efficiency exceeds the benefit gained from the higher security, then consider not implementing that countermeasure.

Chapter Summary

Common SAN security myths include the notion that since a storage network is physically isolated, it is secure; and that the Fibre Channel Protocol is impervious to attack both because it is a complicated protocol with no avenues in and cannot be sniffed. There is also a belief that even if data were to be sniffed, it would be incomprehensible and unusable; however simple tests using an inexpensive optical fiber sniffer show that to be entirely false. Because every SAN environment has its own operational and business requirements, default built-in security features on FC switches are not going to ensure SAN security.

Certainly more security and storage professionals are asking about SAN and storage security than ever before worldwide. The subject comes up in conversations every day and both storage and security professionals alike are craving more information so that they can come up to speed quickly and take the appropriate measures to secure their SAN.

3

SAN Basics for Security Professionals

Although a SAN is a network, it differs significantly from the conventional local area network or TCP/IP-based network. Since security professionals tend to be unfamiliar with SANs, they often overlook or ignore security issues for these networks. This chapter is for IT security professionals with little or no knowledge of storage and the SAN. Storage professionals may elect to skip this chapter and continue with the next chapter, which discusses security basics for storage professionals.

One of the first questions you might ask is this: Why do you need a SAN to begin with?

The original model for open data storage was direct-attached storage (DAS), in which each server had its own storage directly attached using the Small Computer Systems Interface (SCSI, pronounced “skuzzy”) or other protocol. While the DAS model worked well in the early days of the data center, it became clear that DAS had reached its limits when the importance and scale of IT infrastructure outgrew it in the 1990s.

DAS was inefficient, since some disks had large amounts of empty space while others were completely full, and additional disks had to be purchased. This led to a need to pool disk storage into one central location and share the resources with all hosts to optimize utilization of the storage devices. Broadly, this category of solution is known as white space optimization.

SCSI also had distance limitations and could be used only to connect devices in the same rack or at best to another device in an adjacent rack. This precluded its use for most high-availability applications and all disaster recovery solutions. SCSI performance was also an issue for storage applications, particularly backup applications, which demanded increased bandwidth to meet growing business requirements and shrinking backup windows.

The purpose of the SAN is to provide a means of transporting data through a network between a server and any storage devices it requires. The SCSI protocol is still the foundation for the SAN, in which SCSI commands can be transported via a network protocol instead of via directly attached SCSI cables. Although a SAN can be implemented using different network protocols, most SANs have been implemented using the Fibre Channel protocol. (There are other protocols used such as iSCSI and now FCoE, currently under development.)

Note that the spelling of the word “Fibre” is in fact correct here and was written this way intentionally to distinguish it from the word “fiber,” which usually implies fiber-optic cable. This difference appears subtle, but is important. Although the FC protocol is often implemented using fiber-optic cable, it can also be implemented using copper cabling. Since Fibre Channel is the predominant protocol in the SAN market, for the remainder of the book, this will be assumed.

Evolution of the FC Protocol

The FC protocol has evolved through three generations as described in Table 1 and illustrated in Figure 8. The first generation, point-to-point, was a simple non-networked, direct-attached protocol that connected a server directly to a disk or tape device.

The second generation, arbitrated loop, allowed servers and storage devices to communicate in a network with a ring topology (similar to token ring), but any change to the network generated a loop initialization process (LIP), which suspended all communications in the network. This was particularly annoying for backup applications, which, following a LIP interruption, would force the restart of a backup from the beginning. Although the second generation FC allowed for multiple devices to share resources, theoretically it supported a maximum of 127 devices in the entire network and, practically, it did not scale to half that number.

The third generation, switched fabric (FC-SW), allowed for even greater distances, speed, and scalability with a possibility of 2^{24} addressable devices. Although there is a theoretical possibility of addressing over 16 million devices, the modern Brocade fabric, for example, can support approximately 6,000 devices. As a practical matter, if more than a few thousand devices are required, network address translation (NAT) routers are used to interconnect multiple fabrics. A switched fabric creates dedicated connections between two devices as opposed to sharing the connections between all devices as with FC-AL.

Table 1. FC protocol generations

FC Generation	Description	Characteristics
Point-to-point (FC-P2P)	Non-networked model	Direct-attach between server and a storage device
Arbitrated loop (FC-AL)	Shared network with arbitration protocol	• Implemented using FC-AL hubs • Uses ALPA for addressing • Requires arbitration to initiate conversations • Uses LIP for change notification • Up to 127 addressable devices
Switched fabric (FC-SW)	Fully switched networked model	• Implemented using FC switches/ directors/backbones • Uses WWN for addressing • Uses RSCN for change notification • Up to 2^{24} addressable devices • Can scale higher using NAT routers

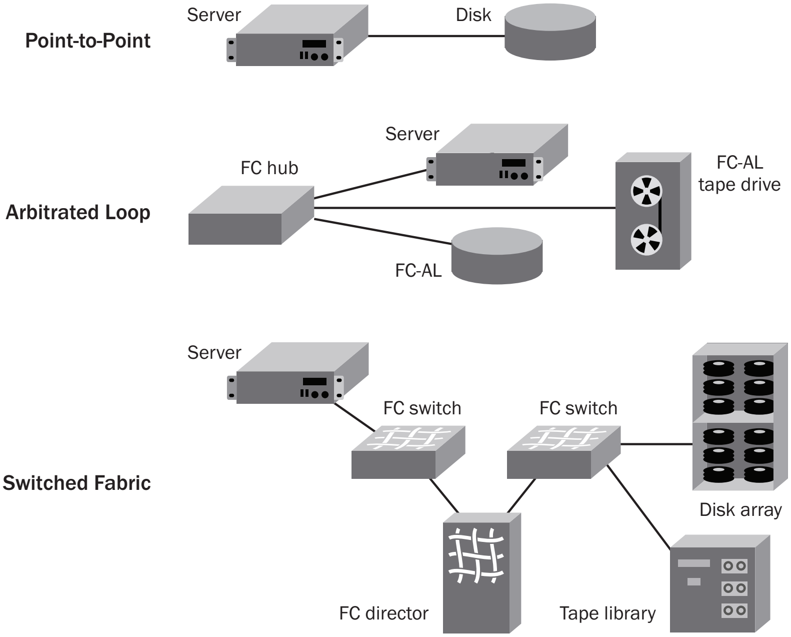

Figure 8. FC protocol generations

The FC switched fabric protocol itself has also evolved through several generations. Its first implementation offered 1 Gigabit per second (Gbps) speeds using a removable transceiver-type device in the switch port, called a gigabit interface converter (GBIC). Subsequent generations, starting with 2 Gbps implementations, used a smaller transceiver called a small form factor pluggable (SFP) device and the latest 16 Gbps technology uses SFP+ transceivers. The successive generations doubled the previous generation's bandwidth to - 2 Gbps, 4 Gbps, 8 Gbps, and now 16 Gbps speeds; 10 Gbps FC is also supported.

Fibre Channel Basics

FC Frames

A switched FC network is called a fabric and the packets transported through the fabric are called frames. An FC frame begins with a start-of-frame (SOF) delimiter and ends with an end-of-frame (EOF) delimiter. Just before the EOF is a Cyclic Redundancy Check (CRC) for integrity checking. The frame header, located immediately after the SOF, provides other information, including the source and destination address. The actual payload can vary in size from 0 to 2,112 bytes for a maximum frame size of 2,148 bytes. Figure 9 illustrates the simplified FC frame format. There are also optional headers used for special cases but these are not discussed in any detail here.

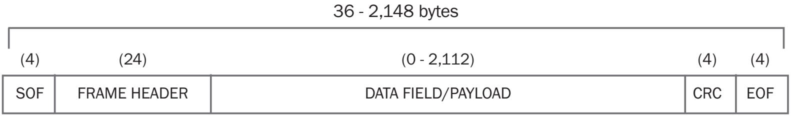

Figure 9. FC frame format

FC Protocol Layers

Similar to TCP/IP, Fibre Channel is a multi-layer protocol with five defined layers as described in Table 2 and illustrated in Figure 10. If you are familiar with the TCP/IP protocol you may notice some similarities, particularly at the lower layers. This is not surprising, since the Gigabit Ethernet standard actually “borrowed” its lower layers from FC. However, many of the functions found in the TCP/IP model are performed in different layers in the FC model and some layers do not exist at all (routing, for example).

Table 2. FC protocol layers

FC Layer	Name	Description
FC-4	Non-networked model	Maps the different upper-level protocols (ULP), such as SCSI and IP, to the lower-level protocols
FC-3	Shared network with arbitration protocol	Includes hunt groups, name server, multicast, alias server, clock synchronization, future services
FC-2	Fully switched networked model	Includes class of service, frame format, sequencing, exchange management, address assignment, multicast management
FC-1	8b/10b encoding 64b/66b encoding	Performs 8b/10b or 64b/66b (for 10 and 16 Gbps FC) serial encoding and decoding
FC-0	Physical	Defines the physical links in the system: connectors, electrical/ optical parameters, data rates

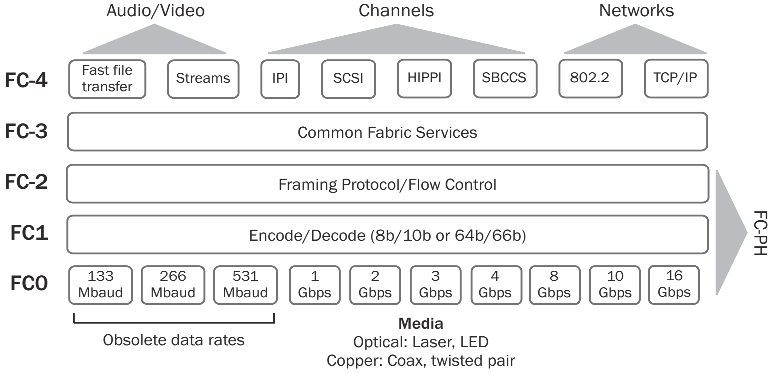

Figure 10. FC protocol layers

Types of Switches

There are different types of switches such as backbones, directors, and routers. Although there are no official definitions for switches, the terms “director” and “backbone” are generally accepted in the industry.

FC-AL hub. A hub is used to connect devices using the arbitrated loop (FC-AL) generation of the FC protocol. Although rarely used today, it can be seen occasionally in older environments or sometimes integrated into a low-end storage arrays (JBOD, or “just a bunch of disks”) or tape library that uses FC-AL disk or tape devices in the background.

FC switch. An FC switch is a networking device that supports the FC protocol and allows hosts and storage devices to communicate with each other. Generally, an FC switch has a 1U or 2U form factor. Some may have redundant power supplies and/or fan modules while others may not. The number of ports in current devices varies from 8 ports all the way to 80 ports.

FC director. This term was borrowed from mainframe ESCON (Enterprise Systems Connection) technology. The ESCON protocol was implemented using highly robust and redundant networking devices called directors, which allowed storage devices to be connected to the mainframe. When the manufacturers of ESCON directors, McDATA and InRange, adapted ESCON directors to support the FC protocol, they simply used the same term.

FC backbone. A backbone has a similar architecture to a director but adds greater performance and advanced functionality to support requirements of next-generation, consolidated data centers. Backbones may offer support for advanced features, such as:

Encryption   Virtualization   Adaptive Networking services

Integrated FC routing   Support for future protocols (FCoE and CEE)

FC router. An FC router is a switch with the ability to connect two or more separate fabrics and allow devices in each fabric to communicate with devices in other fabrics according to user-definable rules. Since there is no routing layer in the FC protocol, FC routers must use a special abstraction layer to present virtual switches to physical switches.

FC gateway. An FC gateway allows devices using different protocols to be connected to an FC fabric. For example, servers connected to a TCP/IP network can be connected using an iSCSI gateway at one end and to an FC fabric at the other end.

FC Fabrics

FC fabrics are implemented using FC switches, directors, backbones, and routers. Each element is addressed by a domain ID (DID) ranging from 1 through 239; no two switches in a fabric can have the same DID.

Enterprise-Class Platforms

Over time, a director has generally become accepted as a switch but with higher reliability, scalability, performance, and flexibility, as described in more detail below. More recently, backbone platforms join directors in a category known as enterprise-class platforms.

Reliability

- Highly redundant hardware architecture with hot swappable components
- Five nines (99.999 percent) availability or better; Brocade directors/backbones tend to be closer to six or seven nines (99.99999 percent) of availability
- Non-disruptive firmware upgrades
- Non-disruptive failover of control processors

Scalability

- Bladed architecture to add blades as needed
- Higher port count, supporting from 32 ports to 384 ports with a possibility of 768 FC ports with a dual chassis configuration using ICLs; Brocade directors/backbones have multi-terabit backplanes

Performance

- High performance backplane architecture

Flexibility

- Support for specialized blades (application, routing)
- Support for other protocols (Ethernet, iSCSI, FCIP, FICON)

FC End Devices

There are two basic types of end device that can be connected to a fabric: initiators and targets. Hosts and servers are known as initiators and are the only devices capable of initiating conversations with other devices in the fabric. Data is written to target storage devices, which cannot generally initiate a conversation on their own with other devices-they require an initiator to do this for them. Hosts are connected to the fabric using a special card, called a Host Bus Adapter (HBA). Storage devices are connected to the fabric through a storage or device controller.

Host and storage devices are connected to FC switch ports, each of which contains an SFP (GBIC in older switches).

Fibre Channel ports are classified into two basic categories, node ports and switch ports:

- The **node ports** identify the ports on the end devices such as the host and storage ports. Switched fabric nodes are called N_Ports. Arbitrated loop nodes are called NL_Ports.
- All **switch ports** begin life as universal ports (U_Ports) and take on specific personalities depending on what they are connected to. When a host or storage device is connected to a switch, the universal port becomes a fabric port, or F_Port. When a switch or router is connected to a switch port, it becomes an extended port, or E_Port. When an FC-AL device is connected to a switch, it becomes a fabric loop port, or FL_Port.

FC switches can be connected to other FC switches via E_Ports on the switch using an inter-switch link (ISL) to merge and form a cohesive fabric of switches. An ISL is simply the connection between two switches or directors. An inter-fabric link (IFL) is used to connect a router to a switch that is in a different fabric. An inter-chassis link (ICL) is used to interconnect two physically separate backbone chassis via a special kind of connector, an ICL cable. In this case the switch port remains an E_Port. Table 3 describes the FC port types and their functions, and Figure 11 illustrates the different FC devices and links in a fabric.

Table 3. FC port types

Type	Category	Name	Description
N_Port	Node	Node port	Port on host or storage device
NL_Port	Node	Node loop port	FC-AL port on host or storage device
F_Port	Switch	Fabric port	Switch port that connects to an N_Port
FL_Port	Switch	Fabric loop port	Switch port that connects to an NL_Port
E_Port	Switch	Extended port	Switch port that connects to other switches forming an ISL

Type	Category	Name	Description
M_Port	Switch	Mirrored port	Switch port that mirrors the data going through another port
U_Port	Switch	Universal port	Switch port with no devices connected to it; will become an E_Port, F_Port, FL_Port, or EX_Port
EX_Port	Router	Extended route port	Switch port on a router connecting to the E_Port on a switch, forming an inter-fabric link (IFL)
D_Port	Switch	Diagnostic port	Used for running link-level diagnostics between two switches

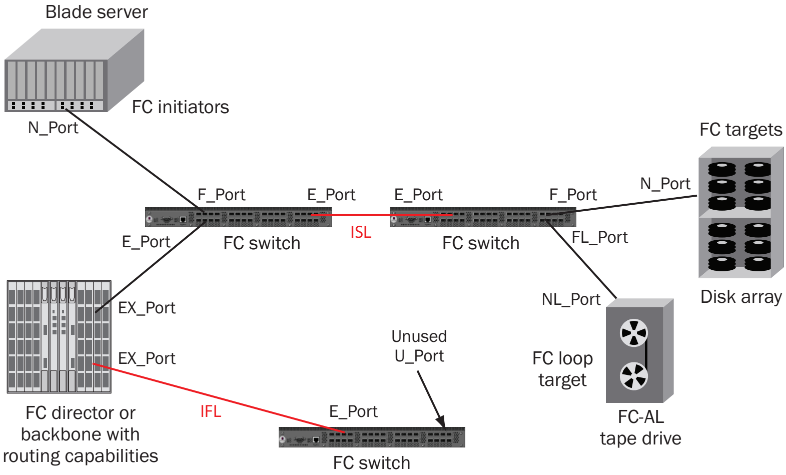

Figure 11. FC devices and port link types

Targets and initiators are identified by a unique 8-byte address called a World Wide Name (WWN), which is equivalent to an Ethernet MAC (Media Access Control) address. There are two types of WWNs:

- A node WWN (nWWN) refers to the actual node or device or host
- A port WWN (pWWN) refers to an actual port on an HBA

Some HBA cards have more than one port, in which case each port has a different pWWN, but there is still only one nWWN for the entire host, as shown in Figure 12.

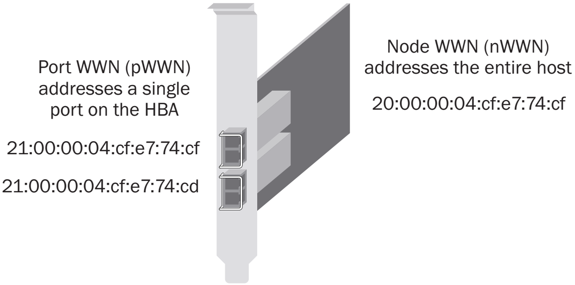

Figure 12. FC device WWNs

RSCN

When new devices are added to a fabric and old ones removed, there must be a means of informing the other devices that a change has occurred. In FC switched fabrics, this is accomplished using a registered state change notification (RSCN). An RSCN is similar in some ways to a LIP in the FC-AL protocol, but it is much less disruptive-particularly with modern HBAs and drivers. When a new device is added to a fabric, an RSCN is broadcast throughout the fabric. With Brocade switches, the RSCN is limited to the affected zone so as not to disrupt the rest of the fabric (see “Zoning” on page 29). This is called RSCN scoping and is similar to the broadcast scoping function provided by virtual LANs (VLANs) in the Ethernet space.

Flow Control

One significant advantage the FC protocol has over other network protocols is the way data flow is controlled. In Ethernet, no meaningful flow control is provided at all. Packets are sent, and if the receiving switch or device is too busy to process them they are discarded. Ethernet LANs rely on higher-level protocols such as TCP to handle flow control. TCP uses rate-based flow control by sending a group of packets and waiting for an acknowledgement back from the other end before sending the next group of packets. If an error occurs and even one packet is not received, then the entire group of packets is resent.

Fibre Channel, on the other hand, uses credit-based flow control and continuously sends frames without waiting for an acknowledgement from the other end. To achieve this, FC uses buffer credits (BB credits) to indicate whether or not there is sufficient memory available to store each transmitted frame.

An FC switch will not transmit data to another switch port until that port has advertised a BB credit. The credit is essentially a promise that the receiving port will be able to deliver the frame to its next destination, either by forwarding it immediately or by storing it and forwarding it later. When a port sends a frame, it cannot use that credit again until the receiving device returns it as an R_RDY call. Once a frame is sent to its next destination, the buffer is freed up. At that point, the associated BB credit is released to notify the other switch that memory is once again available on that port to receive another frame.

FC Fabric Features and Services

The simple name server (SNS), a repository of all known devices attached to the fabric, acts as the fabric's "directory assistance." When a device is connected to a switch, one of the steps it must go through is registering its WWN with the SNS. The SNS may also be responsible for enforcing zoning rules via the WWN on some vendor switches, but Brocade enforces all zoning using the ASIC component of the hardware.

Zoning

Fabrics can become quite large, span great distances, and consist of thousands of nodes. To avoid one flat network allowing every device to be aware of every other device, zoning is used to isolate groups of devices. Zoning is a fabric-based service that groups devices that need to communicate with each other. Once a device is assigned to a zone, it can communicate only with other devices in that zone.

Zoning terminology can be confusing, mainly for historical reasons. The terms "hard and soft zoning," "port and WWN zoning," and "hardware-enforced" and "software-enforced zoning" are often used interchangeably. To be clear, there are two basic methods by which zone members are identified and enforced within a fabric:

- By the hardware via the ASIC: hardware-enforced zoning
- By software as a service: software-enforced zoning

Why Are Zoning Terms Confusing?

In the early days of switched fabrics, hardware-enforced zone members were assigned by the physical port on a switch to which the device was connected. This was defined using the switch domain ID (DID) and the port ID (PID) on the switch. Software-enforced zone members were assigned using the WWN of the device. In this case, the software enforcing the zoning rules is the SNS. For this reason, zoning using the WWN was sometimes called soft zoning and zoning using the DID/PID was sometimes called hard zoning. One of the advantages of zoning by WWN was the flexibility of moving devices to any other port within a fabric without incurring zoning changes. On the other hand, zones defined with DID/PID would not require a zoning change when a faulty device was replaced by a new one with a different WWN. This was particularly useful for tape drives which have a tendency to fail more frequently than other devices such as HBAs.

In 2001, Brocade introduced a technology in the Bloom ASIC, which enabled the ASIC to enforce zoning based on the device WWN, not only the DID/PID. (Although “Bloom” was initially an internal cod name, it is now used externally to identify this generation of ASIC.) Hardware-enforced zoning could now include member definitions using the WWN. This was a significant enhancement from a security perspective. Although Brocade was the first to implement hardware-enforced zoning, most FC switch vendors today enforce zoning through hardware.

For example, with software-enforced zoning, some hosts may cache the WWN of the devices in the zone with which it communicates. If a storage device is removed from the zone and placed in a different zone, the host could still access the storage device even though it is no longer in the same zone. If the WWN in the cache is removed either through a power cycle or cache timeout, the host would not be able to obtain the WWN from the SNS since it is now in a different zone. This is comparable to unlisted telephone numbers. Even though a person delists their phone number, someone who knows their phone number can still call them. If the caller loses the number, however, they would not be able to get if from directory assistance.

With hardware-enforced zoning in Brocade switches, although the host may cache the WWN, the ASIC will block access to the device if it is not in the same zone as the host. This is equivalent to using the call-blocking feature on a telephone. Even though someone has a person's unlisted phone number, if the caller's number is blocked at the central office (CO), then the call would not be allowed.

As a best practice, it is also recommended that you define zones using the pWWN instead of the DID/PID Organizations that implemented WWN zoning definitions were very pleased with their choice when they migrated the SAN from a fabric consisting of several 16-port switches to a single director. The migration is quite simple and involves copying the zoning database to the new director. Those organizations that used DID/PID definitions had to convert all zone definitions manually, since the DIDs and port numbers on 16-port switches did not map to the 256-port director.

Path Selection

Path selection refers to the algorithm for selecting the path frames will follow, given a possible choice. To avoid confusion, this is not the same as FC routing, which is discussed in the routed fabrics section.

There are several types of path selection protocols in FC.

- Fabric Shortest Path First (FSPF)
- Dynamic Path Selection (DPS)
- Trunking

FSPF

Fabric Shortest Path First was invented by Brocade and is now an accepted standard. It is a link-state path selection protocol similar to the TCP/IP Open Shortest Path First (OSPF) routing protocol. FSPF does two things:

- It keeps track of the state of all the links in a fabric
- It calculates a cost to each path

In FSPF, paths are calculated by summing the cost of all the links traversed by the path. Each time a switch is traversed, it is called a hop. Only the lowest-cost path between a source and destination is kept in the routing table. The routing table contains information about which switch port to use when forwarding a frame to its final destination. Table 4 lists the default path costs for the different link speeds.

Table 4. FSPF path costs

Link Speed	Path Cost
< 1 Gbps	2000
1 Gbps	1000
2 Gbps	500
4 Gbps	500
8 Gbps	500
10 Gbps	500
16 Gbps	500

The example in Figure 13 illustrates how FSPF works. A server is connected to a four-switch fabric. The link cost between each switch is set to 500. There are four possible paths in the fabric from the server to the disk array: A-C, A-B-D, A-C-D, or A-B-D-C. The paths through switches A-B-D and A-C-D represent two hops, so the path cost is 500 + 500 = 1000. The cost through path A-B-D-C has three hops, for a cost of 1500. The cost of the path through switches A-C is 500, since there is only one hop. FSPF drops paths A-B-D, A-C-D, and A-B-D-C from its routing table, since they have a higher cost than path A-C, and frames will never follow these paths when A-C is available. In the event that switch C fails, however, FSPF will recalculate the paths and route the frames through A-B-D.

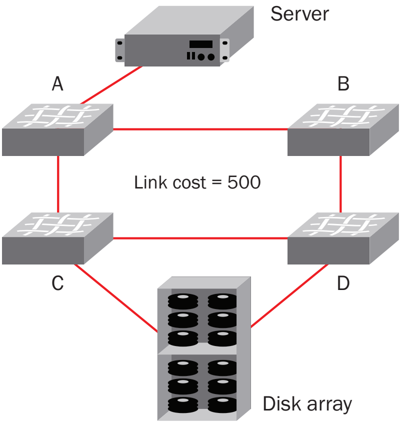

Figure 13. FSPF path selection

When paths with equal costs are available, paths are assigned in a round-robin fashion. Say that two paths are available between three hosts on a switch to the disk storage on another switch. The first host will be assigned one path, the second the other path, and the third will round-robin back to the first path. The Brocade implementation of this feature is called Dynamic Load Sharing (DLS). DLS does not use active load feedback, so the paths remain fixed regardless of the load on the link.

Exchange-Based Routing or DPS

Once a path is determined for a connection between two devices by FSPF, it does not change until the devices log in once again. This may cause load balancing issues in which one path approaches saturation and other available paths carry almost no load. Exchange-based routing improves load balancing over FSPF by sending an entire exchange, sequence of FC frames, through the most efficient path. An entire conversation that consists of several exchanges can be transmitted through different paths. The Brocade implementation of this feature is called Dynamic Path Selection.

Trunking

The term trunking refers to consolidating several links into one link resulting in a higher-bandwidth link. In the LAN world, trunking is often used to consolidate several ports on a network interface card (NIC) card to provide a larger pipe, or trunk, to the switch. In the FC world, trunking currently applies only in the consolidation of ISLs between two switches. Trunking can be implemented in several ways and some FC switch vendors actually use the term “trunking” for exchange-based routing, which can be very confusing.

Brocade implements trunking at the ASIC level in the hardware, which we believe is the only way to truly implement trunking. Hardware-based trunking takes load balancing to the highest level by providing the capability to spread frames across multiple links simultaneously and obtain the best load balancing across available links.

Frame Redirection

In Brocade FOS 5.3 and M-Enterprise OS (M-EOS) 9.8, Brocade introduced a new technology with the capability to redirect frames using a different route than originally intended. This technology was necessary to add a virtualization layer for certain types of applications that would not normally be in the direct data path in the fabric. The Brocade Data Migration Manager (DMM) and EMC RecoverPoint, both of which run on the Brocade 7600 Application Platform, essentially behave as appliances in a fabric, and frames need to be redirected through these devices to perform their intended function. The Brocade encryption solution also requires this technology to allow a Brocade Encryption Switch or Brocade FS8-18 Encryption Blade to be introduced anywhere in a fabric and encrypt from any host to any LUN or tape drive.

Frame redirection, also called nameserver redirection, actually redirects frames to an alternate destination before they reach their final destination. Frame redirection creates an abstraction layer on top of the physical fabric and its configuration. One advantage of this abstraction layer is that it has no impact on pre-existing zoning configurations or the physical hosts and storage devices in the SAN.

An association between a source initiator and a storage port is created in a redirection zone (redirection zones are not the same thing as conventional fabric zones). The redirection zone presents a virtual target to the physical initiator and a virtual initiator to the physical target. The physical initiator believes it is communicating with the physical target but is in fact talking with a virtual target. Once the host or initiator sends a frame to the virtual target, the redirection zone sends the frame to the alternate device, the encryption device in this case. The encryption device encrypts the payload of the frame and sends it back to the fabric where it gets redirected to the destination target device, as shown in Figure 14.

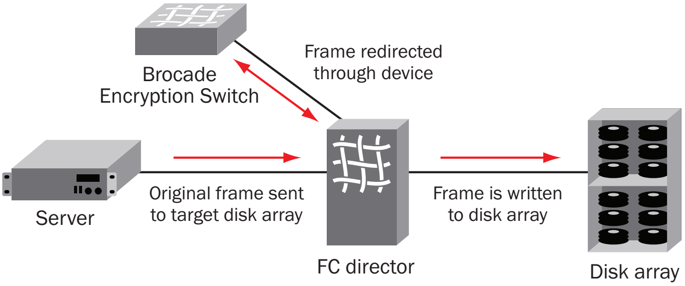

Figure 14. Frame redirection

Fabric Topologies

Switches can be connected together in many ways to form simple fabrics or complex, resilient, multi-tiered fabrics. The topology chosen will depend on the business requirements of each organization. When choosing a topology that is best suited to meet specific business requirements, consider these four factors: performance, scalability, redundancy, and cost.

And although you can architect a SAN for two or three of these factors at the same time, it will usually be at the expense of the other factors. For example, you can have a highly redundant, high-performance fabric but it will most likely not be very scalable. Finding the right balance among these factors is more an art than a science.

Table 5 shows how these four design factors are interrelated for the different topologies described in this section.

Table 5. Fabric topology design factors

Topology	Performance	Scalability	Redundancy	Cost
Cascade	Poor	Poor	Poor	$
Ring	Good	Poor	Good	$
Full mesh	Excellent	Poor	Excellent	$$
Partial mesh	Good	Good	Good	$$
Core-edge	Excellent	Excellent	Good	$$$
Resilient-core-edge	Excellent	Excellent	Excellent	$$$$

For further information on this topic, *Principles of SAN Design, Second Edition*, by Josh Judd, is highly recommended.

Dual Fabrics

As with any network, FC fabrics should be designed without any single points of failure. From an architectural and design point of view, this redundancy is accomplished by using a dual-fabric architecture, as shown in Figure 15. Servers must also have multipathing input/output (MPIO) software running to load balance the traffic between the two paths and to fail over to one path in the event of a path failure. If any hardware component failure occurs in a fabric, starting from the host HBA through to the disk controller, no production downtime will be incurred since there is an alternate path for the traffic.

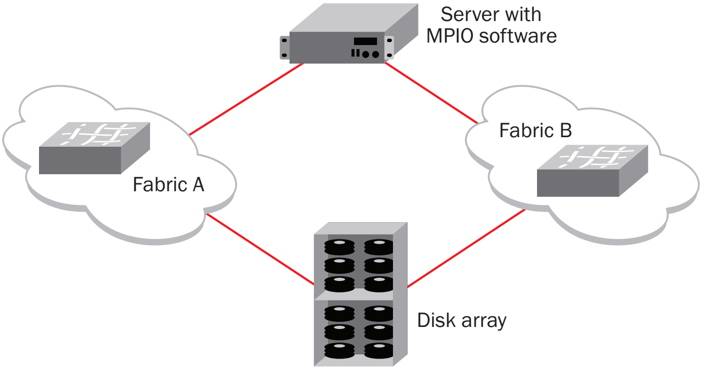

Figure 15. Dual-fabric design

Dual-fabric design is a best practice and should always be used with disk environments. A dual-fabric architecture provides redundant paths to avoid any single points of failure. This is very different from a typical LAN architecture but the impact of a failed access to a disk drive cannot be tolerated as it takes more time to recover from such an event. Converged networks, IP/FC, are particularly problematic as the architecture of each type of network is very different. Tape environments, however, would not benefit from a dual-fabric architecture, since tape drives and backup applications do not have the capability of being dual attached.

Cascade Topology

A cascaded fabric is the simplest architecture; switches are daisy-chained together to form a string of switches. Middle switches are connected to two other switches and the two end switches are connected only to one other switch. Figure 16 illustrates a four-switch (on the left) and a six-switch (on the right) cascade topology. A disadvantage of this topology is that a server attached to Switch A in the six-switch topology would have to traverse four other switches to get to its storage if the storage device were attached to switch F. These multiple hops can degrade performance. For example, if all of the storage devices were attached to switch F and every port on switches A to E were connected to a host, the traffic between switches E and F could become highly congested.

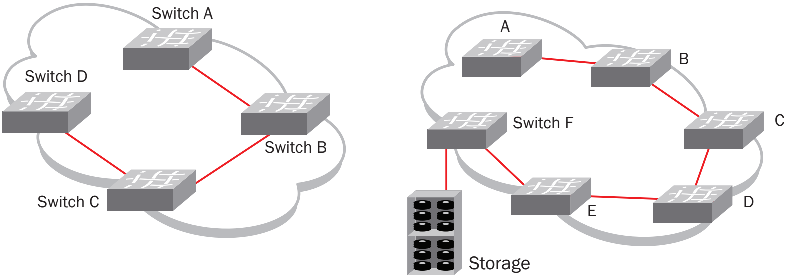

Figure 16. Cascade topology (four and six switches shown)

This is not a scalable design and offers little redundancy. A failure of any switch with this topology would result in isolation of the devices on either side of the failed switch.

Ring Topology

A ring topology is created when every switch in a fabric is connected to two other switches in the same fabric, as shown in Figure 17. A failure of any switch in this topology would still allow all other switches to continue communicating with each other using a path in the opposite direction.

This topology is also not very scalable since the number of hops increases as you add new switches, but it does provide some redundancy with the dual paths.

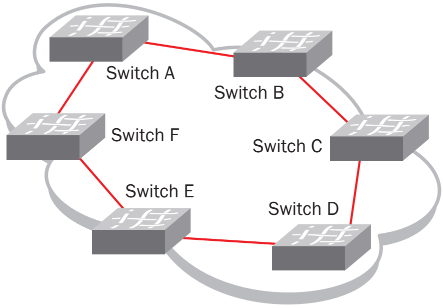

Figure 17. Ring topology

Mesh Topology

The full-mesh topology is created when every switch participating in a fabric is connected to every other switch in the fabric, as shown in Figure 18. This provides the highest level of path redundancy and excellent performance, since there is only one hop between any two switches in the fabric. However, this is the least scalable fabric topology due to the exponential increase in links required as the number of switches increases.

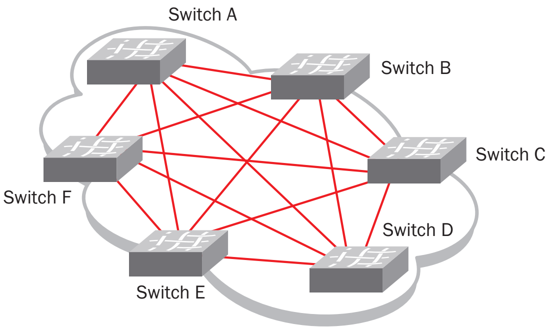

Figure 18. Full mesh topology

The formula for the number of ISLs required to create a full mesh topology is:
1 + 2 + ... + (N-1)

where “N” is the number of switches.

For example: a five-switch full mesh fabric would require:
1 + 2 + 3 + (5-1) = 1 + 2 + 3 + 4 = 10 ISLs

Table 6. Full mesh topology ISL and port requirements

# of Switches	# of ISLs	# of Ports
2	1	2
3	3	6
4	6	12
5	10	20
6	15	30
7	21	42
8	28	56
9	36	72
10	45	90
11	55	110
12	66	132

As can be seen from Table 6, the number of ISLs required increases significantly for each additional switch. It is important to note that each ISL also requires two ports; one at each end switch. Eventually, there will be more ports using ISLs than actual hosts and storage devices in the fabric and switches simply won't have enough ports to connect all the switches.

To improve scalability, a mesh topology can also be constructed as a partial mesh, as shown in Figure 19. In this case, most, but not all, switches are connected to all other switches in the fabric. This is a more scalable alternative to a full mesh design, but at the expense of path redundancy and possibly performance.

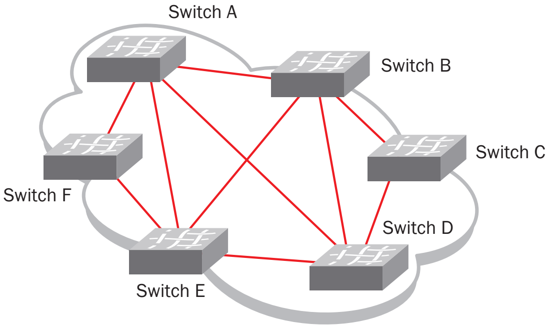

Figure 19. Partial mesh topology

Core-Edge Topology

The core-edge topology is the most commonly implemented and it represents the best compromise among redundancy, scalability, performance, and cost. There are several variations of a core-edge architecture.

Pure core-edge architectures are designed so that all of the traffic must go through the core switch, hence only other switches can be connected to a core switch. The hosts and storage devices are connected to edge switches, as shown in Figure 20. In reality, most core-edge implementations actually connect some storage or host devices to core switches to maximize the cost efficiency of the fabric and make the best utilization of available ports, as shown in Figure 21.

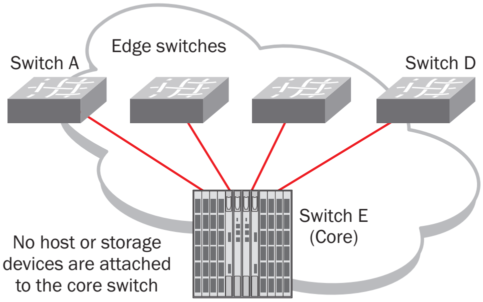

Figure 20. Pure core-edge topology

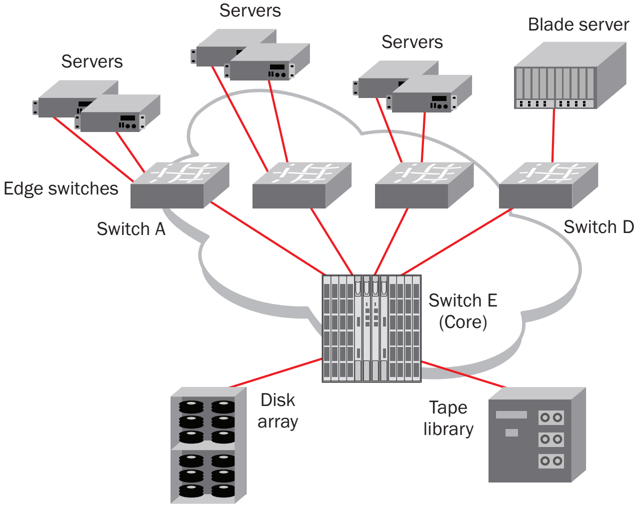

Figure 21. Typical core-edge topology

Resilient Fabrics

A resilient core-edge topology simply means that the core switches are in a redundant configuration, as shown in Figure 22, making the fabric design more resilient to a failure of a core switch. The typical resilient fabric has two core switches with multiple edge switches connected to both core switches. In the event of a core switch failure, there is an alternative path from any edge switch to the other core switch.

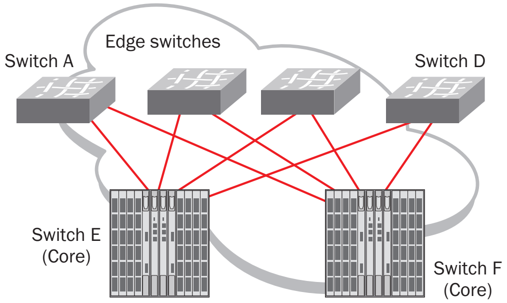

Figure 22. Resilient core-edge topology

Multi-Tiered Fabrics

Multi-tiered fabrics are used for very large fabrics. Typically, one tier is used to connect the storage devices and another tier is used for the hosts. There can be several variations and uses for this topology, including one with a resilient core, as shown in Figure 23.

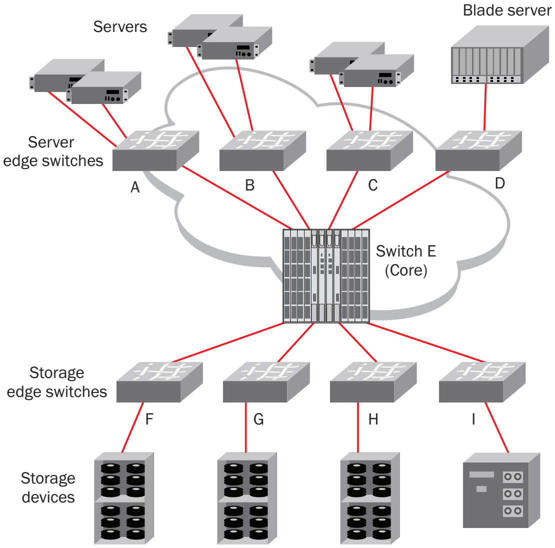

Figure 23. Multi-tiered fabrics

Routed Fabrics

A routed fabric, also called a metaSAN and shown in Figure 24, allows devices in two or more fabrics to communicate with each other without requiring all switches to merge into one flat fabric. The FC protocol, however, was not designed with a routing layer similar to the IP layer in a TCP/IP network. Routing FC fabrics is accomplished by adding an extra abstraction layer and "tricking" switches into believing they are connected directly to a specific physical switch. When a router connects to a switch in another fabric, the connection is referred to as an inter-fabric link (IFL) instead of an ISL. The port at the router end of the IFL is also called an EX_Port and the port at the switch end of the IFL remains an E_Port.

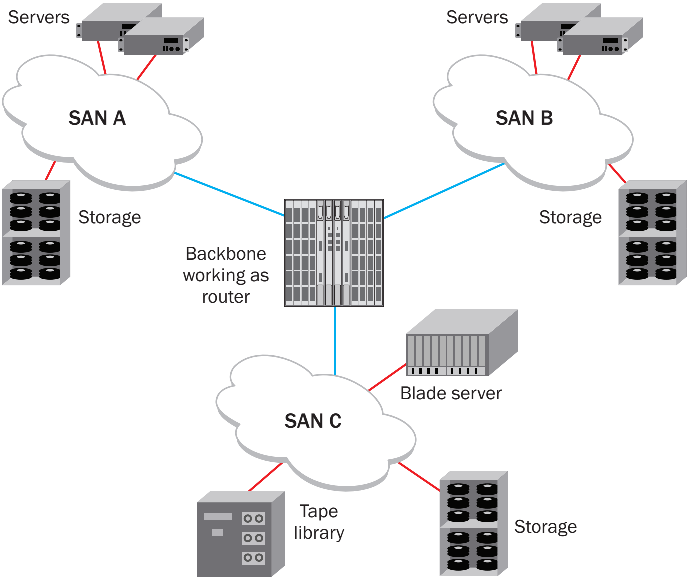

Figure 24. Routed fabrics

Extended Fabrics

In recent years, disaster recovery and business continuity have taken center stage in most IT organizations as a way to protect critical data and prevent potential business outages. Storage networks have played a prominent role in this trend; data replication, remote mirroring, and remote backup are represented in some of the most commonly deployed solutions utilizing long-distance SAN connectivity. Today's organizations typically use two data centers to exchange data between SANs over long distances. Cost, distance, and performance are the primary factors in deciding what technology to use in a long-distance deployment.

As shown in Figure 25, dark fiber is the first method that offers the highest performance for connecting two sites over distance, although this solution comes at a higher price and has distance limitations. The other method uses FCIP, shown in Figure 26, which is a tunneling protocol that can be used to connect to sites over practically any distance using standard WAN connections.

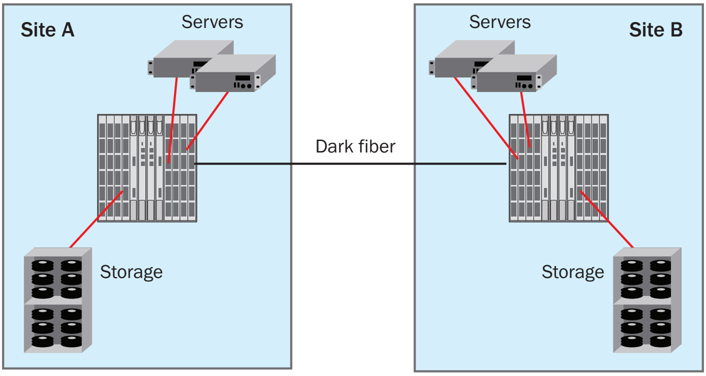

Figure 25. Extended fabric using dark fiber

Implementing dark fiber usually has the greatest initial cost if the organization has to lay the dark fiber or obtain a right of way to do so. Several providers, particularly utility companies, already have a dark fiber infrastructure in place and sell or lease strands of fiber to their customers. Although this option is less expensive than laying your own fiber, it is still quite expensive.

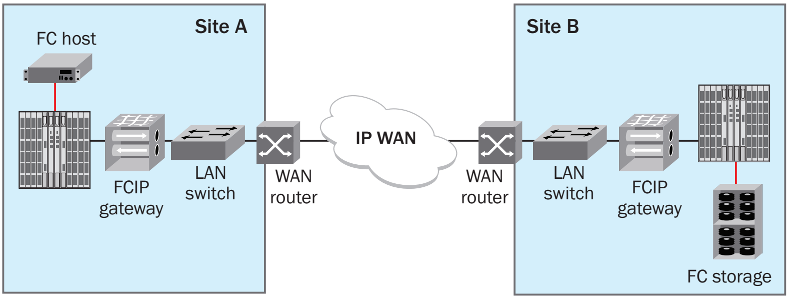

Figure 26. Extended fabric using FCIP

When two separate fabrics at different sites are connected in a standard extended fabric, the link between the two sites becomes a long-distance ISL. Since two switches connected together using an ISL must be part of the same fabric, the fabrics at each site merge to form one fabric. This is important to note given that both sites now share all of the fabric configuration information.

In some cases, it may be preferable to isolate each fabric from the other. A hybrid implementation can be used in this case by using FC routing to maintain isolation between the fabrics at each site. This allows for the sharing of resources between fabrics while maintaining separate configuration and management information.

Disk Storage and LUNs

A LUN is the fundamental unit of disk storage to which the I/O operations are addressed to. The term LUN is often used to refer to a logical partition on a disk or group of disks used to build a file system. A LUN can be composed of an entire disk, a group of disks, or a subset of either.

The term LUN is really a misnomer since it actually stands for logical unit number. The LUN in reality is the specific identifier for a Logical Unit (LU). The correct term referring to the disk partition is LU, but LUN is used ubiquitously throughout the storage industry and the term LU is very rarely used.

Chapter Summary

The Fibre Channel protocol is in common use in storage area networks today. FC frames can carry a payload of 0 to 2,112 bytes-with a maximum frame size of 2,148 bytes. FC devices in the fabric include backbones, directors, switches, routers, and embedded switches. Hosts, called initiators, connect to devices in the fabric via N_Ports to F_Ports. FC devices connect to each other via E_Ports and EX_Ports. ISL are created by connecting FC switches together and IFLs connect fabrics.

FC fabric services improve performance and include path selection via FSPF, exchange-based routing, and trunking. Frame redirection is a Brocade proprietary technology that allows data to be redirected for a particular purpose, such as encryption, and then returned.

Although there are a number of different fabric topologies, the simplest are not robust enough for most SANs, and so variations of a core-edge are commonly used. For very large fabrics, multi-tiered fabrics are used for scalability and resilience. Routed fabrics form a metaSAN, which allows devices to communicate without merging to form a single large fabric.

Enterprises with multiple data center sites take advantage of extension using dark fiber or a long-distance fabric extension solution.

SAN storage resides on disk or tape and the terms that describe storage include disk-based storage, disk array, LUN(s), and tape-based storage.

4

Security Basics for Storage Professionals

To the uninitiated, security may seem like a highly complex concept with specialized jargon, but security really boils down to common sense applied through the use of some basic principles. Certainly, implementing security solutions may not be quite that simple, but understanding the general concepts can go a long way toward understanding the issues. SAN security must be approached from a holistic perspective. There is no point in implementing strict access controls and mechanisms in the SAN if the management interface is relatively unprotected. All components of the SAN--from the infrastructure itself to management tools and physical security--must be considered if you want to create a comprehensive SAN security plan.

This chapter is addressed primarily to the storage professional who may have little or no knowledge of security concepts. Security professionals may also find this chapter useful to better understand how basic security concepts apply specifically to the world of Fibre Channel fabrics.

IT security is an extensive field consisting of multiple domains of knowledge. According to the International Information Systems Security Certification Consortium ((ISC)2), which is responsible for the Certified Information Security Professional (CISSP) certification, there are ten fundamental domains composing a body of knowledge for IT security:

- Access Control and Methodology
- Applications and Systems Development
- Business Continuity Planning
- Cryptography
- Law, Investigations, and Ethics
- Operations Security

- Physical Security
- Security Architecture and Models
- Security Management Practices
- Telecommunications, Network Security, and Internet Security

These ten domains apply directly to the SAN and storage environments and must be addressed in a comprehensive SAN security program.

Security Models

SAN security involves more than just guarding against a malicious outsider with sophisticated hacking tools and the intent to destroy or steal data. In fact, most IT security threats are based on internal threats from employees or other people with access to networks and physical equipment inside the firewall. As a result, best practice IT security strives to achieve several basic security objectives, which vary depending on which model is being followed.

At a minimum:

- Data must always be available to authorized users whenever it is needed
- To maintain its integrity, data must not be modified in any way
- Sensitive data such as personal information, intellectual property, and data pertaining to national security, must remain strictly confidential

As you will see, there are several models in current use and they are described in the next few security sections.

The CIA Triad

One of the most commonly used security models is the famous CIA triad.

Confidentiality

Confidentiality as it pertains to electronic data is the protection of information from being disclosed to unauthorized users. There are several reasons why confidentiality must be considered in IT security, ranging from protecting the right to privacy of individuals to sensitive financial information to social security numbers and other pieces of personal information, which can be used to steal someone's identity. Several laws in place today, particularly in the United States, enforce the protection of confidentiality of Personally Identifiable Information (PII) of the citizens of a state by requiring notification of security breaches involving personal information. As of April 2012, 46 states, as well as the District of Columbia, Puerto Rico, and the Virgin Islands have enacted such legislation. “Chapter 9: Compliance and Storage” starting on page 155, discusses compliance and breach disclosure laws in greater detail.

Confidentiality of electronic information is usually accomplished using cryptographic methods such as encryption of data-at-rest or data-in-flight (see “Chapter 5: Elementary Cryptography” starting on page 73). Authentication methods and access controls are other methods used to address the confidentiality issue.

Integrity

Data integrity ensures the accuracy and consistency of electronic information to provide an assurance that the information has not been modified, deleted, destroyed, or tampered with in any way. For example, it is important to ensure data integrity to prevent attackers from modifying data by inserting unwanted code into an application, or to delete pieces of information before they are stored on a disk.

Integrity verification is generally achieved using methods such as hashing algorithms and check sums. These methods are described extensively in Chapter 5.

Availability

Organizations have become highly dependent on their computer systems and any loss of availability of critical applications can have far-reaching and direct repercussions on the company's livelihood. Maintaining availability of applications, and particularly to the data used by these applications, has become essential. High availability (HA), clustering, and fault-tolerant systems are examples of technology used to maintain application availability. Disk mirroring, RAID (redundant array of independent disks), and remote data replication are used to maintain availability of data stored on disks. Software and specialized appliances such as anti-virus, anti-malware, anti-spam, and intrusion detection systems, can prevent attackers from creating a denial-of-service (DoS) attack.

CIANA

This model expands the basic CIA model by adding two more security elements: non-repudiation and authentication. It is most often used in Information Assurance, which is primarily used by the military. This model is taught as part of a course to reach the NSTISS (National Security Telecommunications and Information System Security) 4011 Certification in the US.

Non-Repudiation

Non-repudiation is used to prevent someone who has performed an action from refuting it and claiming they have not performed action in question. For example, someone makes a purchase on the Internet and then claims they never made the purchase once they receive the goods. Non-repudiation is an essential element in conducting business. This also applies in the other direction in a situation in which an e-commerce website provides proof of payment to the customer. Historically, these functions have been performed using physical signatures and receipts, which then become legal and binding contracts for both parties. The same actions are performed electronically using digital signatures and signed certificates, and other methods such as the Confirm button on some Web forms.

Authentication

Authentication is the process of verifying that people really are who they claim to be. There are several ways to authenticate an individual, including user accounts and passwords. Authentication methods can be quite sophisticated with biometric technology such as fingerprint scanners, face/voice recognition, and iris/retinal scanners. Each of these methods is known as a factor of authentication and can be used in combination, known as multi-factor authentication, to provide greater certainty of authenticity. Factors of authentication will be discussed in greater detail in the physical security section (see “Physical Security” on page 113).

The Parkerian Hexad

The Parkerian Hexad is a set of six fundamental concepts of information security, initially proposed by Donn S. Parker. The term was actually coined by M.E. Kabay from Norwich University. The Parkerian Hexad is an extension of the CIA triad discussed previously and introduces three new elements: possession or control, authenticity, and utility.

Possession or Control

If possession is nine-tenths of the law, it has never been more true than in IT security. Loss of control or possession of data must be prevented at all costs, since it must be assumed that once the owner no longer has control, the data is necessarily compromised. Suppose that a backup tape containing customer and credit card information is lost or stolen-a frequent occurrence in recent times. Even if the tape was simply misplaced and no data has actually been read, the assumption must be that the data on the tape is now known and appropriate measures must be taken according to company or industry guidelines, or regulations. Customers must be advised, and credit cards must be reissued, to prevent unauthorized use of the credit card information.

Authenticity

The origin or source of information can be spoofed or forged. Authenticity refers to validating that information does in fact come from the source claimed. Someone can forge an e-mail header to appear like it was sent from someone else. Fields in a database can have incorrect information inserted into them.

Utility

Information has value only if it can be used. If a database file is corrupted, then it is no longer useful and fails the utility test. Data encryption is a very useful method of protecting confidentiality, but if the key is lost the encrypted data is no longer useful since it will no longer be readable. Utility is not the same as availability, but a breach in utility may result in a loss of availability.

Common Security Terms

Asset. Any item having value such as an IT system, data, personnel, or hardware.

Attack. The act of compromising or breaching the security of an asset.

Security threat. A person or event that has the potential to cause harm, including an employee, malware (software used to harm IT systems), or a natural disaster.

Security vulnerability. A flaw or defect in an asset that can allow an attacker to appropriate, gain control, or otherwise prevent the systems' owner from using the system as intended.

Risk. The likelihood or probability that an asset will be compromised, lost, or destroyed. A risk can be accepted, mitigated, transferred, or ignored. If the value of the asset is of little value or the probability of it being attacked is low, then the risk can be accepted or tolerated. Risks can be mitigated by implementing controls and countermeasures to reduce the risk. The risk can be transferred to another entity such as an insurance company or an outsourcing firm. Finally, a risk can simply be ignored which can cause a new risk in itself-remember security through obscurity?

Exploit. Methods and techniques used to take advantage of security vulnerability to perform an attack on an asset. For example, a computer virus exploits flaws in an operating system to attack the computer system.

Countermeasure. Techniques and tools implemented to protect assets or mitigate risks associated with an attack.

Controls. Measures taken to avoid, counteract, and protect against security risks against an asset.

Preventive measure. Similar to a countermeasure but usually non-technical such as policies, procedures, training, and awareness.

Audit trail. Logging mechanism that tracks user and event activity on a system.

Types of Threats

A threat is anything that can cause harm. An IT security threat is anything that can cause harm to IT assets. Threats against IT assets specifically can be classified into three basic categories.

- Disasters
- Technology
- Human

Of course, technology threats and sometimes disasters are created and executed by humans, so arguably there are only two categories.

Threats from Disasters

Disasters are unique threats; they are not generally aimed at a specific target but can indiscriminately wipe out an entire data center and its personnel. Disasters usually cause the most amount of damage, have the greatest impact on a company's operation, and take the longest to recover from. One of the greatest impacts of a disaster is the impact on the personnel. Technology can be replaced relatively easily, but personnel may need to be reassigned to a temporary facility, or in the worst case scenario, be replaced.

Disasters can occur from natural or man-made causes, including:

- Earthquake
- Hurricane
- Thunderstorm
- Tsunami and tidal wave
- Wildfire
- Flood
- Landslide
- Tornado
- Volcano
- Winter and ice storms

Man-made disasters include:

- Terrorism
- Fire
- Dam failure
- War
- Chemical emergency
- Hazardous material spill or leak
- Nuclear plant emergency (could also be technology)

Protecting against disasters that impact IT assets and business requirements can be accomplished in many ways. The key to successfully protecting against disasters is proper planning, implementation of plans, and dry runs.

The first step is to conduct a business impact analysis (BIA) by system to determine the impact of a disaster on each system in the company, and not only computer or IT systems. Once the BIA is completed, a plan must be created, which is usually known as the Business Continuity (BC) plan. Part of the BC plan addresses the recovery of data systems, which is usually referred to as the Disaster Recovery (DR) plan.

Once the plan has been created, it must be executed or implemented. The DR plan is generally implemented using a combination of procedures and technology.

A DR plan can include the following:

- Backups
- Replication
- Mirrored sites (hot/warm/cold)
- Procedures
- Computer Security Incident Response Team (CSIRT)

Finally, once the plan has been deployed, it must be tested on a regular basis. Performing a scheduled or planned failover from the primary site to a secondary site is not for the fainthearted, but it is necessary to demonstrate that procedures and systems will function properly in the event of a real disaster.

Technological Threats

The technological threats to IT assets are created by people and used by people to exploit vulnerabilities in IT systems. The software used to harm IT systems is called malware and includes:

- Viruses
- Worms
- Spyware
- Rootkits
- Trojans/Trojan horses
- Zombies
- Botnets, or bots
- Spam

Besides malware, there are other technological threats used by the “black hat” community to exploit system vulnerabilities and to learn and perfect the skills necessary to attack systems. There are several Web sites and discussion groups for the underground hacking community, from which attack tools can be downloaded. On these sites, information is exchanged among hackers so that they can discover new vulnerabilities and develop the exploits to abuse these vulnerabilities.

One significant threat is the widespread availability of open source software that has hidden malware built into the application. Peer-to-peer (P2P) sites used for sharing software, music, and video files are renowned for installing spyware (malware that captures information and relays it back to another computer) on the unsuspecting downloader's computer. Some spyware may contain key-logging software to capture key strokes from the remote user for the purpose of obtaining passwords, account numbers, and other sensitive private information.

In the case of a SAN, it is possible that a computer used to manage a SAN is infected with spyware. The information collected by the spyware could be used later to compromise the entire SAN and its data.

Threats from the Human Element

By far, the greatest threat to an IT system is the human element. Ultimately, people create the technology used to attack IT systems. Individuals that attack systems fall into two basic categories: insiders (employees or persons authorized to have access to facilities or systems) and outsiders. There are two subcategories for each based on intent: malicious and non-malicious. Table 7 lists the common threats found in each category.

Table 7. Classification of human threats

	Internal	**External**
Malicious	• Disgruntled employees • Contract workers • Third-party providers/ vendors • Opportunity • Coercion • Financial gain	• Hackers • Industrial espionage • Cyber-terrorists • Criminals • Curious individuals • Script kiddies (a pejorative term used by hackers to describe those who use technology developed by others to attack computer systems and networks)

	Internal	External
Non-malicious	• Carelessness • Lack of training • Lack of security awareness • Improper zoning • Misconfigured HBAs • Inadequate backups • Inadequate or non-existent operational procedures • Reduced budgets	N/A

It is interesting to note that this table does not include non-malicious external threats. It is the writer's opinion that all external threats are malicious regardless of the intent since the result is always malicious. For example, even if a curious individual breaches a system and only browses around various directories, the security administrator who detects this breach must now investigate. Who is the person that breached the system? What was his intention? Was she simply collecting information in preparation for a more significant attack in the future?

Addressing these questions during an investigation takes time and costs the company money, resulting in a loss. Hence, all external threats, no matter how benign they may seem at first, have a negative effect and are considered malicious.

Protecting from External Threats

Attackers from the outside come in many forms with different motivational factors. Some hackers attack systems for fame and bragging rights within the "black hat" community. Terrorists attack systems to cause maximum damage and loss to organizations. Others attack systems for profit and personal gain, such as organized criminals. Terrorist organizations have used cyber-crime to finance their terrorist activities. Some attackers are just curious individuals who want to see what they can do. These "script kiddies" may be young hackers without sophisticated computer knowledge, who download hacking scripts and tools from the Internet and try them out on random organizations and systems for amusement.

Isolating the systems and assets from the outside world is the primary method used to protect against external threats. The defense-in-depth strategy works well to provide multiple layers of protection from outside attacks such that each layer adds an additional barrier and challenge to the attacker. (See also “The Brocade SAN Security Model” on page 91.)

There are two access points for an outsider to gain access to an organization's IT assets. Attackers can breach one or both of the following:

- Physical security to gain physical access to the assets
- The network to gain access to the servers and other assets connected to the network

Protecting assets from physical access requires appropriate physical security measures to restrict access to authorized persons only. Protecting assets from being accessed through the network is much more difficult, since there can be more than one entry point into the network. As with any technology, networks have many vulnerabilities with new ones discovered on a regular basis. Although protecting conventional LAN networks is out of scope for this book, if you are interested there are many excellent resources available on this topic.

Protecting from Internal Threats

It is a well established that the majority of attacks are perpetrated by insiders or by an insider who may assist an outsider, deliberately or inadvertently. Protecting against internal threats is arguably the greatest challenge a security professional faces. Insiders are individuals that have been granted physical access to systems and facilities. They are often given passwords to super accounts such as root and admin. Even in the most secure facility, there is really nothing that can be done to prevent an insider from causing physical damage to equipment if they decide to do so. They will most likely get caught doing it, but they cannot be stopped before the damage is done.

Non-malicious insider threats are probably the most common cause of service disruptions in a SAN. Several factors can contribute to this problem, including:

- Lack of knowledge and training
- Undocumented or non-existent operational procedures
- Bypass of operational procedures
- Fatigue caused by long or nighttime working hours
- Misidentification of hardware
- Simple human errors

The key to minimizing the risks associated with this type of threat is to develop solid, well-documented operational procedures and restrict administrator privileges to only the tasks that are required for an administrator's job functions. Organizations should not grant additional privileges to a trusted, long-term, or favored administrator when those privileges are not required for that administrator's job functions.

Malicious insider threats typically involve employees or contractors who have something to gain from exploiting a weakness in the system. These threats are the most difficult to manage and control, since they involve people who have legitimate access to the targeted systems. The key to mitigating risks from this type of threat is to limit the privileges a specific individual has and to distribute workload and responsibilities among multiple administrators. In the event that a security incident occurs, it is important to have a proper incident response procedure in place, with clear methods to track administrator activities and provide evidence for any potential criminal or civil investigation.

The following list, while not comprehensive, provides important points to consider when defending against insiders:

- Proper hiring and screening practices
- Limited access to facilities and assets
- Personal identifiers, physical and digital
- Appropriate controls
- Monitoring
- Procedures and policies
- Incident response
- Training and awareness

The first step, and probably the most important, is to perform appropriate background checks on employees before they are given the “keys to the kingdom.” Background checks can be basic or exceptionally comprehensive, depending on the nature of the systems they will be granted access to and the nature of the organization’s requirements. For military and intelligence positions involving national security, a top secret clearance or higher may be required.

A top secret clearance requires the investigation of a person's history, relationships, lifestyle, financial positions, and includes a polygraph (lie detector) test. For other employees, a simple verification of references from previous employers or a credit check may be sufficient. A credit check may not seem relevant at first, but if a potential employee has considerable financial difficulties, then this could indicate a weakness in that person's life, which could be exploited by a criminal element.

Once hired, employees should be given access only to assets or facilities they need to perform their job function. Providing an access card to allow an employee to enter a building should not necessarily imply that the employee can now access all areas within the building. The same applies to accounts and passwords. A database administrator may be granted root privileges on the database servers for which they are responsible, but they should not have similar powerful privileges on the backup server, Web servers, or any other applications/servers they are not directly responsible for managing. This general concept is also known as “separation of duties”.

Each individual employee should have a unique identifier assigned to them. A building access card, for example, should be unique and have a photo of the employee on it. When employees log into a system, they should use their personal account with the appropriate privileges instead of the generic root or admin accounts, which could be used by anyone. The intention is to be able to associate an action with a person in a manner that cannot be repudiated.

Appropriate controls should be put in place to limit access and detect anomalies or inappropriate behavior. These could be in the form of access control lists (ACL) or role-based access control (RBAC) assigned to individual users restricting what they can do. Programs can log all access to files and file systems, computer systems, facilities, and so on.

Once controls are in place, they must be monitored. There is obviously no sense in capturing valuable access information in log files if no one looks at the log files. A recommendation on the frequency of monitoring varies depending on the type of assets being protected. Some events need to be monitored only occasionally, while others need to be monitored in real-time to provide an immediate response to a breach. Fire and burglary alarm systems are examples of real-time monitoring systems as are credit card fraud detection systems. Many, if not most, security breaches result from operator error. Creating well-documented and detailed operations procedures helps mitigate risks associated with operator error. Security policies also mitigate these risks by establishing guidelines and rules for employees to follow. Policies also serve to protect the company from liability in the event of an employee acting against approved company policy.

Once policies and procedures are in place, they need to be enforced. If a policy is established but infractions are always without consequence, then it will lose its effectiveness over time. Infractions must be flagged in some way, even if it is only a friendly reminder that a certain behavior has been observed with a link back to the policy for the employee to review. Of course, for significant or repeated infractions, sanctions may be more drastic and include employee dismissal or even criminal charges in extreme cases.

Finally, one of the most overlooked aspects with insiders is training and awareness. Training provides improved knowledge resulting in greater efficiency and reduction of operational errors. Awareness training also reduces the frequency of the type of error caused by not realizing the impact of certain actions, as described in the examples below.

A classic technique used by hackers to gain access to systems is called social engineering. This technique involves manipulation of trust when a person impersonates or assumes authentic-seeming characteristics. A common social engineering technique is to impersonate a help desk person and ask an employee to update their company profile. During this process, the unsuspecting employee will be asked to provide their password so that the “help desk person” can log in and make the necessary changes. Another commonly used social engineering technique is phishing. A hacker may send an e-mail to an individual requesting them to update their account profile for their investment bank, for example. They are asked to follow a link which leads them to a phony, but authentic-looking, website. As the user logs in to update their profile, their account and password information is captured and subsequently used to perform unauthorized transactions in their account. Raising the awareness of all employees of the schemes and strategies used by hackers to obtain information is an effective method of combating hacking via social engineering.

Attacks

Attackers have many options and strategies at their disposal to attack IT assets. They can be very simple or highly sophisticated attacks depending on the skill of the attacker and the target that under attack. The first step in any attack usually involves collecting information to determine the best strategy to perform a successful attack on a system.

Preparing for an Attack

A typical technique used by hackers to collect information is a port scan. Port scanning refers to searching for open network ports on a target system. This enables a hacker to know what services are running on the system, information that can subsequently be used in an attack based on known vulnerabilities for these services. Another technique, known as OS fingerprinting, involves analyzing ping responses from systems, which can provide clues to the type of operating system the target uses.

A commonly used technique to obtain information used by more daring and sophisticated hackers is social engineering, discussed in the previous section. Social engineering is highly effective since it does not require sophisticated tools, technology, or access to systems to obtain information-but goes directly to the individuals that have the information already at hand.

Browsing is another common method of collecting information. An attacker can search a person's workspace for passwords written on post-it note or a piece of paper, files on a computer, or activate a GUI in read-only mode. For example, the Brocade Web Tools GUI prior to FOS 5.3.0 displayed all switch information by default, once a switch's IP address was entered into a browser window.

Types of Attacks

Hackers can be very creative individuals and there are many ways in which they can attack and compromise a system. There is an extensive “black hat” community whose members share information across the Internet and make it available to any interested person. The list of attacks is quite long; here are a few attacks that can be used in a SAN environment:

- Back doors
- Sniffing
- Denial-of-service (DoS)
- Man-in-the-middle (MITM)
- Spoofing

Back Door

A back door allows someone to bypass the normal access methods to get into a system. It can have many forms, such as a program with hidden code that allows its creator to enter a system at a later date. Sometimes a host can be bypassed by placing it in single-user mode and bypassing the operating system authentication mechanism. A back door can also be a default account, such as those used by maintenance technicians to gain access to a system when users have forgotten their password to access the system. This is one reason why it is extremely important to change all default account passwords for a new system. A simple Web search reveals default account passwords for most major IT equipment vendors (including Brocade).

Sniffing

Sniffing is the act of capturing traffic on a network. It can be accomplished using highly sophisticated and expensive equipment such as a trace analyzer. Or it can use inexpensive, readily available equipment such as software on a computer that places the network interface card (NIC) in promiscuous mode to capture all traffic that reaches it. As seen in “Chapter 2: SAN Security Myths” starting on page 9, sensitive optical couplers can be purchased for under $1,000 to sniff traffic on an optical fiber cable without having to splice the cable. The data itself can be stored on any computer, including a laptop, and with packet filtering software, unnecessary traffic or noise can be filtered out and only the interesting traffic is kept.

Denial of Service

A denial-of-service (DoS) attack aims at disabling systems or preventing them from performing their intended function. Powering off an FC switch or storage array is a simple form of a DoS attack. A distributed DoS (DDoS) attack is more sophisticated and requires the collaboration of large numbers of computers, usually infected with a sleeping process called a “zombie,” which simultaneously sends a large number of requests to a Web server, resulting in congestion that may bring the system down. The first such attack of significance was performed by an adolescent with the aid of several programs he downloaded from the Internet, and he managed to bring down several Web sites including CNN, Yahoo!, Ebay, Amazon, E*Trade, and Dell.

Man-in-the-Middle

A man-in-the-middle (MITM) is an active form of sniffing in which an unauthorized third party is introduced between two legitimate parties communicating with each other. Often, the MITM pretends to be one of the parties during the authentication process and then relays information between the two parties. The result is that the two parties believe they are communicating directly with each other, but in fact they are communicating through a third party. The third party can then store the traffic exchanged between the two parties and use the information for a subsequent attack. For example, a GUI using HTTP to manage a switch can be compromised by an MITM attack. To prevent this, an end-point authentication mechanism such as SSL can be used to secure the channel between the GUI and the switch.

Spoofing

Spoofing refers to taking on the identity of another device or person. Spoofing can be used in SANs by assigning the WWN of a known device in a fabric to another host's HBA and introducing it into the fabric. The FC protocol does not have any mechanism to prevent duplicate WWNs in a fabric. This may seem odd at first, but it is similar to the Ethernet protocol, in which duplicate MAC addresses are allowed. In fact, some NICs come with several Ethernet ports and by default, each port shares the same MAC address. This is usually done to reduce the number of entries in the arp table where the MAC addresses are cached on the server.

As of FOS 7.0, Brocade has implemented measures (discussed later) to modify the behavior of an FC switch when a duplicate WWN is detected at login. One possibility would be to configure switches to reject any devices attempting to login with a duplicate WWN.

As shown above, there are many techniques a hacker can use to breach a system. All SANs have vulnerabilities that can be exploited, and special measures are required to protect against these attacks. The next section looks at how to protect against these attacks and mitigate the risks associated with them.

Identification and Authentication

One of the great challenges in IT security is providing a method to allow users access to IT resources and to prove that they really are who they claim to be before granting them access. This is usually a two-part process involving identification (stating who you are) and authentication (proving you really are that person). There are several methods available to accomplish both of these functions.

Authentication

When only one method is used to authenticate a person, it is called a single-factor authentication. When more than one method is used to authenticate a person then it is called multi-factor authentication.

The four different factors of authentication are:

- *Something you have* such as a key, an access card, an employee badge, or a user account
- *Something you know* such as a password, a personal identification number (PIN), or an access code
- *Something you are* that is a part of your physical person such as a fingerprint, retina or iris, voice, or facial features (biometrics)
- *How you do something*, such as the way you write your signature or how you type on a keyboard

Using more than one factor of authentication provides stronger authentication. For example, if an employee's access card to the company building is stolen, then the thief would be able to use that card to access the building without any further challenges. On the other hand, if the same employee was also required to enter a 4-digit PIN on a keypad, then that would provide additional protection against someone trying to use a lost or stolen access card. Nevertheless, one could argue that an employee could be coerced to giving someone their PIN. For more sensitive environments, biometrics could help protect against coercion, since it would be very difficult to simulate another person's biometric characteristics, like a fingerprint or retinal pattern. Some devices are quite sophisticated and also measure temperature or other parameters to prevent using body parts that have been removed from their rightful owners.

Biometrics

Biometrics is the science and technology of measuring biological information. In IT, biometric technology is used as an authentication mechanism to identify and verify the identity of individuals via:

- Fingerprints
- Palm prints
- Hand geometry
- Retinal scans
- Iris scans
- Facial patterns
- Voice patterns

The following two biometric characteristics are different from the others, since they do not identify a body part, but rather analyze how an individual performs a specific task:

- Signature dynamics
- Keyboard dynamics

Signature dynamics measure writing speed and pauses at different points in the signature. Keyboard dynamics measure a person's typing patterns, that is, how fast they can type, delays in typing two separate letters, and so on.

One of the challenges of biometrics is balancing the error rate. There are two types of errors in biometrics: false positives and false negatives. A false positive (type I error) occurs when a biometric system falsely confirms a person's identity. A false negative (type II error) occurs when a biometric system fails to identify a person. Of the two types of errors, a false positive is more serious, as this could allow an unauthorized person to gain unauthorized access.

On the one hand, when a biometric system generates too many false negatives, it becomes a source of frustration and nuisance to users of the system, since they are not identified and not authenticated when they should be. It may take several attempts to get a valid authentication and users get annoyed with the entire system, not to mention the time wasted and resulting loss of productivity. A false positive, on the other hand, could be a real problem when an invalid user is identified as a valid user and is authenticated. Biometric systems are tuned in such a way to achieve a good balance between type I and type II errors. In some cases, a biometric system may favor false negatives if false positives are not tolerated.

From a storage perspective, biometrics are often used to access secure computer rooms and are sometimes used for authentication to gain access to a management workstation.

User Accounts and Passwords

The user account is the principal method to identify a user who requests access to an IT system. The password is the primary method of authenticating the identity of a user. At first glance, the user account and password authentication method would appear to be a two-factor authentication method, but in fact both items are something a person knows, so they are two aspects of the same factor.

When the user account and password method is used in combination with another authentication method such as a smart card, a common access card (CAC, used in military and intelligence communities), or a fingerprint reader, then it becomes a two-factor authentication method. Another popular two-factor authentication method uses a piece of hardware called a token, which generates a new authentication code at regular intervals, usually ranging from 30-60 seconds. Since this authentication code continually changes, the user does not need to memorize an access code or change it periodically.

Passwords are a bit like chewing gum in some respects. You don't want to share it with other people, it gets stale after a while, and it makes a big mess if you leave it lying around!

Passwords should be unique to an individual and not shared between groups of individuals. For example, the root or other super user account (like admin) should not be used by multiple system administrators. Pre-defined system accounts with a default password should always be changed when the system is first installed. One of the first things an experienced hacker might do when attempting to break into a system is to use the factory default passwords for that particular system. These passwords are very easy to obtain and a simple Web search for “vendor_name root password” will most likely generate multiple hits with several sources offering comprehensive lists of vendors and passwords.

As time goes by, a password has a higher probability of being discovered and compromised; therefore it is important to change passwords on a regular basis. How often the password should be changed depends on the environment. If the password is changed too often, then it becomes more difficult for the user to remember. As a result, many users simply resort to writing their password down and keeping it somewhere handy such as under their keyboard or on a post-it note on the side of their monitor, not the safest places to keep a password secret.

One challenge with passwords is a situation in which users have to memorize different passwords for each system they are required to manage. The ability for a user to use one password and account for all of the systems they are required to access is called single sign-on. Programs are available that allow a user to create or change a password and automatically update all the systems a user has access to.

Utilities and protocols such as RADIUS (Remote Authentication Dial-In User Service) and LDAP (Lightweight Directory Access Protocol) can perform this function, as well as provide more sophisticated user account management. When a user logs into a system using one of these methods, the authentication request to the system is redirected to the RADIUS or LDAP server, which performs the authentication and sends a confirmation back to the system if the authentication is successful.

Physical Security

The first line of defense to protect IT assets from external threats and many internal threats is physical security. Physical security not only involves preventing and detecting access to assets but also addresses safety concerns affecting the personnel, facilities, and equipment in the data center. This section introduces general concepts of physical security relevant to IT and storage environments.

Physical security controls come in the form of physical and psychological deterrents. Deterrents can be visible or invisible and real or perceived. For example, a guard dog can be used as a real physical deterrent, but a “Beware of Dog” sign (with no dog) can provide the illusion of protection and acts as a psychological deterrent. Lighting can also be used as both a physical and psychological deterrent. When lighting is used in strategic locations with the proper intensity, it provides a disorienting glare effect, which can be a physical deterrent. Lighting is used most often as a deterrent to make intruders feel as if they are being observed and could be discovered-particularly if it is combined with a plainly visible video surveillance system.

To ensure physical protection of assets, the following groups of countermeasures should be considered:

- Policies and procedures
- Personnel
- Barriers
- Equipment
- Records

As with any security program, policies and procedures provide the general guidelines and establish the spirit in which physical security is implemented. The policies and procedures also provide liability protection to an organization when employees do not follow them and incidents occur as a result of not following policy.

Personnel include not only the obvious security guards. System administrators, operators, and employees need to be involved in contributing to effective physical security. System administrators and operators need to follow published procedures and policies to ensure that the systems for which they are responsible are not left unprotected. Employees should be alert for suspicious looking individuals or situations, such as when the exterior door to their office building is left propped open.

Barriers or access control systems can be structural, human, or natural. A door to a computer center with an electronic access control system is a structural barrier. A security guard posted at the entrance of a data center is an example of a human barrier. A natural creek can be a natural barrier if crossing a bridge is required to access a building.

Equipment and technology are heavily used in modern physical security, which include electronic access control systems, locks, fire and intrusion detection systems, and communication systems.

Records and logs are also an important part of physical security to detect patterns, flag anomalies, provide evidence, and record events and activity. Records can be in paper format such as sign-in sheets and incident reports, or they can be in electronic format such as video tapes and electronic access databases.

Physical access controls are put in place to allow authorized individuals to gain access to specified areas. These include barriers of all types such as fences, gates, and doors. These controls can combine multiple mechanisms to provide additional layers of security.

Physical access controls include:

- Electronic access systems
- Intrusion detection systems
- Surveillance systems

Electronic access systems are frequently used to control individual access to buildings and areas within a building. The typical electronic access system uses a card, which can be either swiped in a card reader or placed near a proximity sensor to be read. Information contained in the card identifies the individual user and if the user is authorized, access is granted and recorded in a database with a time stamp. Other electronic access control systems may use biometrics to identify an individual or a combination of methods for multi-factor authentication.

Some electronic access systems use special entry mechanisms to prevent piggybacking, for example. Piggybacking occurs when an individual physically follows an authorized user, knowingly or otherwise, to gain access to a location. Social engineering techniques are often used to bypass this system to convince authorized users to let them piggyback. Piggyback-prevention systems include turn styles, double doors in a system with a second door that won't open until the first is closed, and weight-sensitive floors.

Intrusion detection systems are used to detect unauthorized access to designated areas. These systems include motion sensors, infrared sensors, pressure-sensitive switches, and so on.

Surveillance systems monitor activity in designated areas using security personnel, electronic systems, such as closed-circuit television (CCTV) systems, and computer equipment.

To develop a comprehensive physical security plan, other factors need to be considered:

- Temperature and humidity control
- Power management
- Uninterruptable power supply (UPS)
- Generators

As explained, physical security is the first line of defense in protecting IT assets and is an important component of a comprehensive IT security program.

Information Disposal and Sanitization

Eventually, all storage media will end its useful life and will need to be disposed. A storage media may end its life prematurely when a disk drive or tape cartridge is defective. Typically, a storage vendor will replace a defective disk drive with a new one then refurbish and recycle the old one. In most cases, the defective disk drive is sent to a testing facility, where it is run through diagnostic tests and, in many cases, refurbished and sent back to the field. A proper refurbishing process should also wipe out pre-existing data according to a specified method, which may vary between vendors. Vendors generally take great measures to ensure that no customer data remains on refurbished disks, but there have been reported cases of customers receiving new drives that contained live data from a previous customer.

There is also the technology refresh issue. Once an organization chooses to refresh their storage arrays with newer models, or with a different vendor's products, older arrays are often swapped out as part of a deal. As with failed disk drives, all disks should be properly sanitized, but sometimes old units are put on the second-hand market without prior data sanitization. With numerous reported cases of storage media on the second-hand market containing live data, data disposal and sanitization has gained public attention from the media, risk management teams in the corporate world, and government organizations alike.

Data Sanitization

Data disposal and sanitization deals with maintaining confidentiality of information. Evidently, not all stored data needs to be destroyed or sanitized, and the degree to which it needs to be sanitized depends on the sensitivity and importance of the data as well as the risk of exposure to the company if the data were stolen. Certain industries regulate how certain types of data should be sanitized, while other industries are governed by legislation specifying what and how data should be destroyed.

The first step in developing a data destruction and sanitization strategy is to classify the data to identify which types of data require special sanitization and/or destruction requirements. Once the data has been identified, the level of sanitization to be performed should then be determined.

There are several accepted methods to sanitize and destroy data. The NIST Special Publication 800-88 provides some useful guidelines on sanitizing media. This publication proposes four basic types of data sanitization methods, described in the following sections.

Disposal

Discarding media with no sanitization concerns is appropriate only for non-confidential or non-sensitive information. Simply deleting files and emptying the recycle bin or reformatting a disk drive could meet this requirement.

Clearing

Acceptable for non-sensitive data, clearing protects confidentiality by clearing information using an accepted overwriting method to protect against attacks using data scavenging tools. Simple file deletion is not acceptable at this level of sanitization. However, overwriting does not work on failed or defective media, making it inappropriate for certain environments. Data clearing is also referred to as data shredding, erasure, or wiping.

The clearing method uses one of several techniques to overwrite data on a functional disk drive. Clearing can be accomplished in a variety of ways and several standard algorithms have been developed to accomplish this. Although this method is sufficient for moderately sensitive data, it is usually not appropriate for highly sensitive data. The read/ write mechanisms of disk drives are not precise enough to exactly overlay new data over old data. It is entirely possible to see small bands of residual data underlying the new data using sophisticated forensic equipment such as magnetic force microscopes. Clearly, such forensic equipment is not available to the average hacker, but it certainly could be used by a foreign government, for example, if an enemy's sensitive disk drive should fall into their hands.

There has been controversy around this subject as a result of conflicting research data on the ability to recover overwritten data. Using special microscopes, some researchers were able to demonstrate that overwritten data could be recovered. More recent work has demonstrated that modern drives are more accurate and it is no longer possible to perform such an attack. Nevertheless, it is entirely possible that even modern drives could encounter calibration issues resulting from routine wear and tear, which could allow residual data to be observed.

Purging

Data purging is used to protect against sophisticated laboratory attacks using specialized equipment such as electron microscopes and sophisticated diagnostic and forensic tools. Degaussing, passing a magnetic field through a magnetic media, is an acceptable method of purging data, although certain types of degaussers are more effective than others depending on their energy rating. Clearly degaussing will not work on non-magnetic media such as optical media.

Destruction

Physical destruction of the media is the only accepted method to completely prevent the recovery of data on magnetic media; once the media has been destroyed it can no longer be reused. Physical destruction can be accomplished by disintegrating, incinerating, pulverizing, shredding, and melting. These methods are usually reserved for the most sensitive data and are the most common methods used by military and intelligence agencies to destroy media containing confidential data. They are also often used in combination with each other, for example, a disk may be first crushed then incinerated or melted.

Data sanitization procedures should also include verification processes to ensure proper confidentiality is maintained. Random samples of sanitized media should be tested by persons not involved in the actual sanitization process.

Electronic Data Shredding Methods

Several methods and algorithms have been developed to electronically shred data. Some of these algorithms are standards used by military and other government agencies for clearing certain types of data. Some algorithms may only be used to shred non-classified or non-sensitive data, while others are acceptable for confidential or top secret data. Commonly used data cleaning algorithms are listed in Table 8.

Table 8. Data cleaning algorithms

Algorithm	Passes	Description
US Army	3	Pass 1- random bytes; passes 2 and 3 with certain bytes and with its compliment
US DoD 5220.22-M	3	Pass 1– zeroes; pass 2 – ones; pass 3 – random bytes
US Navy NAVSO P-5239-26	3	Overwriting with pass verification

Algorithm	Passes	Description
US Air Force System Security Instruction 5020	3	Pass 1– zeroes; pass 2 – ones; pass 3 – any character with pass verification
NATO Data Destruction Standard	5	5 passes
US DoD 5220.22-M (ECE)	7	Passes 1 and 2 – certain bytes and its compliment; passes 3 and 4 – random character; passes 5 and 6 – character and its compliment; pass 7 – random character
Canadian RCMP TSSIT OPS-II	7	Alternating passes of ones and zeroes and last pass with random characters
NSA/CSS Policy Manual 9-12	7	Alternating passes of ones and zeroes
Bruce Schneier	7	Pass 1 – zeroes; pass 2 – ones; passes 3 through 7 – random characters
Peter Guttman	35	35 passes of pre-defined patterns (considered excessive given today's drive technology)

Chapter Summary

When securing a SAN environment, it is important to consider a holistic approach. A defense-in-depth strategy presents attackers with multiple layers of challenges and hardens all aspects of the environment. Technological defenses, although important, do not necessarily address issues related to the human element such as human error. Security policies, training, operation procedures, and raising awareness can go a long way to address these issues and are unfortunately often overlooked.

5

Elementary Cryptography

This chapter is an introduction to some of the general concepts of cryptography for an audience that has limited familiarity with cryptography. Many examples are simplified in order to present often highly complex concepts in a more palatable format for those unfamiliar with this field.

Cryptography can be used in a SAN environment to solve several problems. Here are some examples of where cryptography can be used in a SAN:

- Exchanging data between the management interfaces on the switch and the management server
- Exchanging sensitive data between two data centers over public networks
- Protecting data-at-rest on a tape or disk media
- Authenticating devices joining a fabric using DH-CHAP

The word “cryptography” is derived from the Greek words “kryptos,” which means hidden, and “graphia,” which means writing, so it is the art of hidden writing. Stated more completely, it is the art, science, skill, or process of communicating in or deciphering messages written in code. Scholars certainly have speculated about the first use of cryptography, but one fact is indisputable. The need to exchange or store sensitive information in a manner that only the parties involved could understand has been around for a very long time-certainly several centuries. One of the earliest known ciphers was used by Julius Caesar and is appropriately known as the Caesar cipher or the shift cipher. It is based on the concept of shifting the alphabet by a pre-determined number of letters.

For example, if the Latin/Roman alphabet is shifted by five letters, the following cipher results.

Original alphabet: A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

Cipher code: F G H I J K L M N O P Q R S T U V W X Y Z A B C D E

Using this cipher, the word "RETREAT" would be encoded as "WJYWJFY." This type of cipher is also known as a transposition cipher. A substitution cipher is another type of cipher, which mixes up the letters in no particular order.

For example, if the order of the Latin/Roman alphabet is randomized, the following cipher results.

Original alphabet: A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

Cipher code: Q F B O R X K U G W I P A N S Z H T D J C Y M E L V

Using this cipher, the word "RETREAT" would be encoded as "TRJTRQJ." Although these basic ciphers can probably be decoded easily by most weekend puzzle enthusiasts, they were nevertheless useful in their time.

Mechanical devices have been developed to refine the encoding and decoding of messages. One of the best known encoding devices is the German Enigma Machine used in World War II, which used multiple passes of a simple alphabet substitution cipher. The electronic age introduced computers and electronic devices, which further increase the complexity and speed of the encoding process and subsequently the difficulty of decoding messages without the key.

For as long as cryptography has been around, there has also been an equivalent aspiration to decode messages. The process of deciphering messages without access to the key is called cryptanalysis.

Common Cryptography Terms

Here are some basic terms used in cryptography.

Cryptographic algorithm or cipher. The actual procedure used to manipulate a readable message and render it unreadable.

Transposition cipher. A code that shifts or slides the characters in a sequence (such as an alphabet) either to the left or to the right by a specified number of places to encrypt data.

Substitution cipher. A code that mixes up the characters in a sequence (such as an alphabet) in a random order to encrypt data.

Cleartext. Readable data that is transmitted or stored in an unencrypted or readable format.

Plaintext. A cryptographic term that refers to the input to an encryption algorithm. The difference between cleartext and plaintext is subtle and they are often used interchangeably, albeit incorrectly.

Ciphertext. The output of an encryption algorithm, that is, encrypted data that is unreadable (the opposite of plaintext).

Encryption. The process of converting readable data into an unreadable format.

Decryption. The process of converting unreadable data into readable data (the opposite of encryption).

Key. The secret code made up of a sequence of characters, bits, and/or instructions used in conjunction with an encryption algorithm to encrypt and decrypt messages made up of a sequence of characters, bits, and/or instructions.

Key space. The total number of possible keys that exist using a given key size and algorithm.

Cryptographic system or cryptosystem. The hardware or software implementation used to convert plaintext into ciphertext and vice versa.

Symmetric vs. Asymmetric Cryptography

One of the enduring problems in cryptography is the distribution of keys. How do you distribute a secret key and minimize or eliminate the risk of the key being compromised if it is intercepted? This problem is compounded when the key used to encrypt the message is the same as the one used to decrypt it. Before the electronic era, the only way to exchange keys was to meet in person or deliver them the old-fashioned way, and exchange the keys verbally or via printed copy.

Symmetric Keys

Symmetric cryptography uses the same key or a secret key to encrypt and decrypt messages, such as the Cesar cipher. Since the same key is used for both encryption and decryption, anyone in possession of the key can decrypt the message encoded using that key. Distributing the keys to the authorized persons poses a particular challenge and extreme measures sometimes need to be taken for what is termed a secure key exchange. If the key is stolen or intercepted during the transfer process, the code is broken and the encrypted message no longer deemed secure. Examples of well-known symmetric key algorithms are Data Encryption Standard (DES) 3DES (pronounced “triple DEZ”), and Advanced Encryption Standard (AES).

Asymmetric Keys

Asymmetric cryptography has been developed to address the key exchange problem. Exchanging keys in times of war on the battlefield certainly posed a challenge, but the Internet and e-commerce present even greater challenges. How can you conduct millions of transactions per day at wire speeds across the world and make sure you authenticate each transaction?

Asymmetric cryptography is also referred to as public key cryptography, since it makes use of keys that are known publicly. A public key exchange system works on the principle of encrypting a message using a combination of a known public key and a secret private key. Each party has their own public and private key pairs, which are different but mathematically related. Examples of familiar asymmetric key algorithms are used with Public Key Infrastructure (PKI) and RSA (represents the family names of the inventors: Rivest, Shamir, and Adelman).

There are several ways of implementing public key exchanges. Below is a high-level example of how this works, without going into too many details of how it is actually accomplished.

Say that Jim sends Maria a message that only Maria will be able to read. Both Jim and Maria have a private key but Jim does not know Maria's private key and Maria does not know Jim's. They also have a public key that is available on a public server containing the public key repository. Jim queries the repository to obtain Maria's public key and uses it with his own private key to encrypt the message. The message is sent to Maria and she then retrieves Jim's public key. Using the combination of Jim's public key and her private key, she can then decrypt the message and read it.

Bob is a bad guy and he intercepts the message between Jim and Maria. Since Bob does not know either Maria or Jim's private key, he is unable to decrypt the message even if he has Maria and Jim's public keys. Figure 27 illustrates this example.

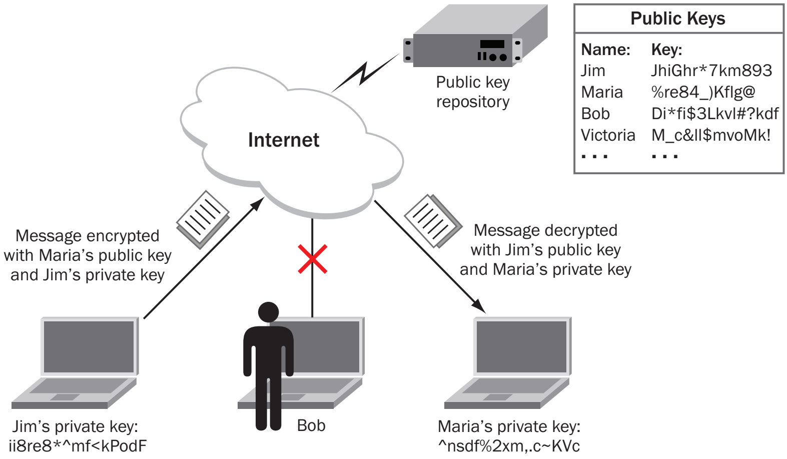

Public Keys

Name:	**Key:**
Jim	JhiGhr*7km893
Maria	%re84_)Kflg@
Bob	Di*fi$3Lkvl#?kdf
Victoria	M_c&ll$mvoMk!
...	...

Figure 27. Simplified public key exchange

Hybrid Systems

Asymmetric keys are computation intensive and are not well suited to processing large volumes of data. Hybrid cryptosystems can be used when an asymmetric algorithm can be used for authentication and key distribution along with a symmetric algorithm for the actual data encryption process.

Cryptographic Algorithms

A cryptographic algorithm or cipher is the actual procedure used to manipulate a readable message and render it unreadable. The readable message that is input to a cipher is called plaintext and its output is called ciphertext.

Early thinking around ciphers encouraged security through obscurity. Proprietary algorithms were kept secret for fear of their being discovered and subsequently broken. With certain exceptions, notably military-grade applications, this thinking has been replaced by the use of open algorithms that withstand public scrutiny. August Kerckhoff proposed six rules for military cryptography in 1883 such that if an encryption algorithm were to fall into enemy hands, it would not result in a compromise of the message as long as the key was not discovered.

Proprietary encryption algorithms are generally not considered as secure, since they do not benefit from being scrutinized by either the cryptographic community at large or the general public. These algorithms are usually analyzed by a group of elite professional cryptographers, who sometimes have tunnel vision and see things from only one perspective, a situation which could result in a gaping flaw that is overlooked.

An open algorithm, on the other hand, has this advantage: at some point thousands of individuals attempted to break it. If thousands of people from different professions and viewpoints are unable to break the code, then the algorithm certainly can be considered more secure than without having been through such a rigorous process. When someone eventually breaks the code, it will become public knowledge and the algorithm will have ended its useful life.

Designing a cryptographic algorithm is very complex and should take the factors listed below into consideration, so it can be used efficiently in practical commercial applications:

- **Speed of encryption.** A highly complex and completely unbreakable algorithm would have no practical commercial use if it also required an inordinate amount of processing power to compute, which would drastically impact performance.

- **Memory usage.** Algorithms that use too much memory to perform their computations and manipulations may require memory components too large to physically fit into certain portable devices, restricting their practical application.

- **Range of applications.** Ability to implement in a wide range of devices, from supercomputers and disk arrays to Smart Cards and radio-frequency identification (RFID) devices, can determine the value of an algorithm.
- **Cost.** If the cost to implement the cryptosystem is too high, it may not find commercial relevance. Military and intelligence applications sometimes warrant the high cost in exchange for stronger cryptographic capabilities.
- **Openness.** In support of Kerckhoffs' principle stated by Auguste Kerckhoffs in the 19th century: a cryptosystem should be secure even if everything about the system, except the key, is public knowledge.

There are three basic categories of ciphers: block ciphers, stream ciphers, and hashing algorithms.

Block Ciphers

Block ciphers are used to encrypt data as an entire block as opposed to one bit at a time. An entire block of data is processed at the same time by the block cipher. A plaintext message is broken down into fixed-length blocks and passed to the block cipher as plaintext. Each plaintext block is encrypted with they key to create a ciphertext block that is the same size as the input plaintext block. The decryption process takes the ciphertext message and breaks it down into fixed-size blocks. Each ciphertext block is decrypted using the key to produce a plaintext block the same size as the input ciphertext block, as shown in Figure 28.

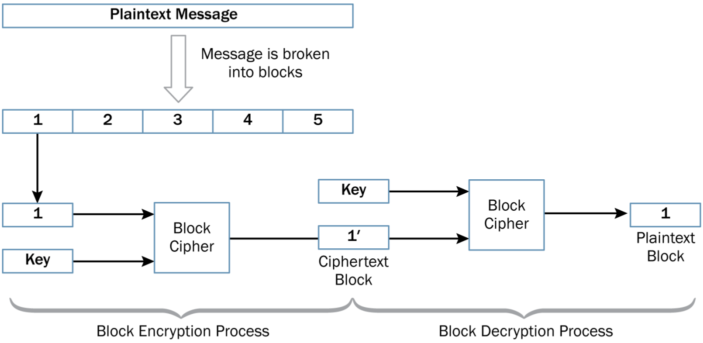

Figure 28. Block cipher encryption/decryption

Stream Ciphers

Stream ciphers process plaintext one bit at a time, as shown in Figure 29. Generally, stream ciphers are considered less secure, since there is a higher risk of having repeating patterns. For this reason, block ciphers are more commonly used. Block ciphers can, however, be used on streaming data when they are operating in a streaming mode of operation, such as the counter (CTR) mode discussed later in this chapter.

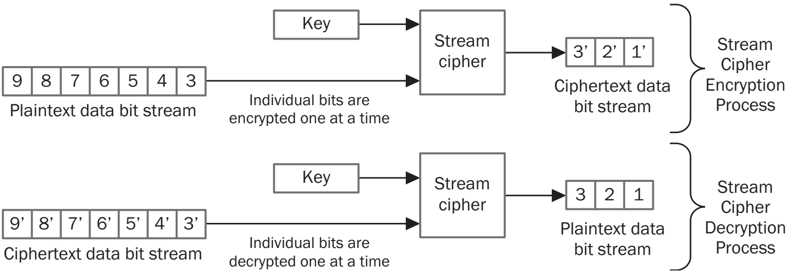
Key
Stream cipher
3' 2' 1'
Ciphertext data bit stream
Stream Cipher Encryption Process
9 8 7 6 5 4 3
Plaintext data bit stream
Individual bits are encrypted one at a time
Key
Stream cipher
3 2 1
Plaintext data bit stream
Stream Cipher Decryption Process
9' 8' 7' 6' 5' 4' 3'
Ciphertext data bit stream
Individual bits are decrypted one at a time

Figure 29. Stream cipher encryption/decryption

Both the block and stream ciphers address the data confidentiality issue by rendering the data unreadable without the key. Hashing algorithms, on the other hand, address the integrity issue by providing a means to verify that data has not been modified.

Hashing Algorithms

Hashing algorithms, shown in Table 30, are used to convert a message of variable length to a shorter, sometimes fixed, length or numerical value. The resulting value is sometimes called a message digest (MD). These algorithms are also called one-way hashing functions since they only work in one direction, and it is not possible to reconstruct the original message from the message digest. As a real-world example, a book could not be reconstructed in its entirety from a simple summary of the book.

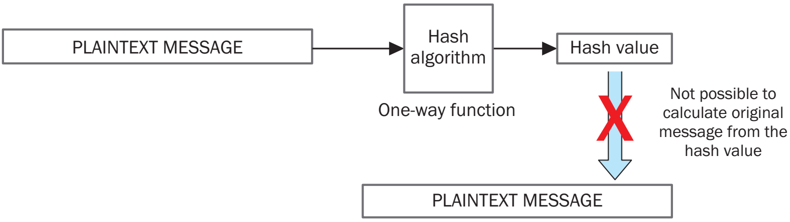

Figure 30. Hashing algorithm

Hashing algorithms are often used for error-checking, but in IT security, they are generally used to verify the integrity of a message. For example, hackers may sometimes modify code, particularly freeware, and add a back door, virus, or some other type of malware into the code. When the original software package is passed through a hashing algorithm, a hash value is generated, which can then be posted in a public location. If someone downloads this software package and puts it through the same hashing algorithm, the resulting hash value should match the one posted. If they do not match, then it can be assumed that the software has been modified and cannot be trusted to be secure.

An MD by itself only provides integrity verification, but an MD can be encrypted with a symmetric key to provide authentication of the provenance of the data. This technique is known as a message authentication code (MAC).

Digital Signatures

A digital signature, shown in Figure 31, is exactly what it says: it is the equivalent of a person's paper signature but for digital transactions. Digital signatures cannot be repudiated later, that is, it would not be possible to deny that someone or something actually sent a message or made a transaction.

A digital signature is created as follows:

1. A message is created.
2. The message is passed through an algorithm to generate a hash value.
3. The hash value is encrypted using a private key from some public/ private key authority.
4. The resulting encrypted hash is the digital signature.

The validation process at the other end goes as follows:

1. The message is passed through the same hashing algorithm.
2. The digital signature is decrypted using the public key of the sender.
3. The resulting decrypted hash is compared with the newly calculated hash.
4. If the hash values match then the message is deemed valid.

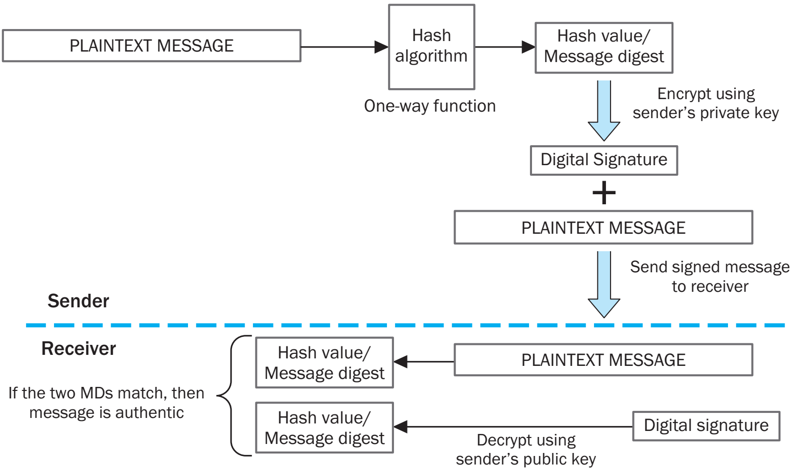
PLAINTEXT MESSAGE
Hash algorithm
One-way function
Hash value/ Message digest
Encrypt using sender's private key
Digital Signature
+
PLAINTEXT MESSAGE
Send signed message to receiver
Sender
Receiver
If the two MDs match, then message is authentic
Hash value/ Message digest
PLAINTEXT MESSAGE
Hash value/ Message digest
Decrypt using sender's public key
Digital signature

Figure 31. Digital signature

Digital signatures provide non-repudiation and integrity to prevent someone from claiming that they did not perform an action or approve a transaction, and to confirm that the message has not been modified.

Modes of Operation

A cryptographic algorithm can be applied in different ways depending on the type of data and specific requirements of its application. For example, some data is fixed length and must remain exactly the same size after it has been encrypted, as is the case with block data written to disks. In other contexts, such as tape backup applications, the data is streaming to the device very rapidly on a flexible media. Encrypting data bit by bit as it is transported serially through a wire requires yet another method of encryption. Instead of creating a different cryptographic algorithm for each application and type of data, the same algorithm is used in different ways to accommodate each specific requirement. These methods are called modes of operation.

The following describes common modes of operation in use today:

- **Electronic Codebook (ECB).** Divides the message into equal-size blocks that are encrypted separately. ECB is not very good for hiding patterns, since identical plaintext blocks encrypt to identical ciphertext blocks.
- **Cipher-Block Chaining (CBC).** A message is divided into equal-size blocks and the entire block is encrypted. The first block is also encrypted using an initialization vector (IV) to randomize the

encryption process. Furthermore, all subsequent blocks are chained in such a way that the encryption process depends on all previously encrypted blocks.

- **Counter (CTR).** Converts a block cipher into a stream cipher by encrypting successive blocks in a data stream using a counter to change the value for each block.
- **Galois Counter Mode (GCM).** A similar mode of operation to the Counter mode with the addition of an authentication component called the Galois mode. Authentication is usually a computing-intensive process, which would not be acceptable for streaming data. Authentication is also necessary to prevent certain types of attack on a data stream. The Galois mode was developed to authenticate a message at very high speeds with minimal performance impact on the data throughput. GCM is used by the Brocade encryption solution to encrypt streaming tape data.
- **XEX-based Tweaked Codebook with Stealing (XTS).** This mode of operation was designed for data formats that are not evenly divisible by a given block size, as is the case for disk drives with sectors not evenly divisible by their block size. XTS is used by the Brocade encryption solution to encrypt block data on disk drives.

DES/3DES

The National Standards Bureau (NSB) recognized the need for a government-wide standard for the encryption of non-classified, sensitive data and developed a cryptographic algorithm to address this requirement. The first draft of the algorithm was written by IBM and was called LUCIFER. The name was eventually changed to the Data Encryption Standard (DES) and it was adopted as an official standard in 1976. The algorithm is a symmetric-key algorithm with 56-bit keys that determine which bits will be transposed and substituted in the original message.

DES was broken by a brute force attack in 1999 by the Electronic Frontier Foundation (EFF), making it imperative to come up with a new cryptographic standard for the Federal Government. Selecting a new cryptographic standard is a complex and lengthy process, since proposed algorithms must be given the test of time and should provide the opportunity to have as many people attempt to break it as possible. In the interim, it was crucial to replace DES with a new algorithm with a larger key space, since DES was no longer secure. The simplest solution was to use a modified application of DES until the new standard could be adopted.

3DES (pronounced “triple DEZ”) became the DES interim replacement. 3DES increased the key space from 256 to 2168 by simply performing three consecutive encryption passes using DES and a different 56-bit key for each pass. Effectively, this created an algorithm that used three 56-bit keys, which is equivalent to a 168-bit key size.

AES

The Advanced Encryption Standard (AES) was developed by the National Institute of Standards and Technology (NIST) to replace DES through a competitive process, in which 15 competitors submitted proposed algorithms. The Rijndael algorithm proposed by Vincent Rijmen and Joan Daemen, two Belgian engineers, was selected as the new encryption standard in 2000. The AES is defined in the Federal Information Processing Standards (FIPS) publication 197. The Rijndael algorithm is a symmetric key block cipher which supports keys with 128 bits, 192 bits, and 256 bits (AES-128, AES-192, and AES-256 respectively). It was rapidly adopted by the industry and most commercial applications for encryption of data-at-rest use AES-256.

The AES standard is the first to use an open cipher that is available to anyone, distinguishing it from its predecessor DES. Although there had been some controversy around DES, which was co-developed by the National Security Agency (NSA), as to whether the NSA had created a back door into the algorithm, the open nature of the AES standard has all but eliminated this possibility.

Diffie-Hellman

Whitfield Diffie and Martin Hellman were the first to publish the concept of public key cryptography in 1976. In actual fact, the public key-private key theorem was first developed independently by James Ellis in 1969 and the algorithm problem was solved by Clifford Cox in 1973. However, their work was not published before the publication of the work of Diffie-Hellman. Without going into too many details of how this algorithm works, it is based on the process of factoring very large prime numbers, which is very difficult to do.

Diffie-Hellman (DH) was the first practical implementation of public key cryptography and is ubiquitous in the IT security industry. It is an integral part of several standards and protocols. In the FC industry, the FC-SP (Fibre Channel-Security Protocol) uses DH-CHAP (DH-Challenge Handshake Authentication Protocol) to authenticate devices or switches joining a fabric.

RSA

At around the same time Diffie and Hellman were completing their work on public key cryptography, three researchers at MIT were also working independently on the same problem. Ronald Rivest, Adi Shamir, and Leonard Adelman found a practical implementation of the public key cryptography algorithm and published their findings in 1977. They obtained a patent for their discovery and subsequently formed a company in 1982 bearing the first initial of their last names: RSA. Their patent expired in September 2000 and is now in the public domain. The RSA algorithm is so widespread that it has become a de facto standard.

Digital Certificates

A digital certificate is sometimes confused with a digital signature but it is very different. A digital certificate is the equivalent of an ID card and is issued to an individual (or device) by a trusted certification authority (CA). It is composed of the owner's name, a serial number, an expiration date, a copy of the owner's public key, and the digital signature of the CA. Some digital certificates use the standardized X.509 format defined in RFC 2459.

Starting in FOS 4.2, Brocade switches came pre-loaded with a digital certificate. Digital certificates are no longer pre-loaded (since the release of FOS 5.1), but one can still be installed, although it needs to be acquired from a source outside Brocade. The digital certificate was originally implemented to authenticate switches joining a secured fabric using the Switch Connection Control (SCC) policy.

PKI

The Public Key Infrastructure (PKI) is a set of programs, hardware, data formats, procedures, and policies required to manage digital certificates. It is a general concept with different implementations offered by multiple vendors. PKI emerged from the necessity to provide a secure means of exchanging information and performing commercial transactions over the Internet. The challenge was to ensure that digital certificates used in commercial transactions were authentic. To accomplish this, it was necessary to build a "web of trust" and provide the necessary authorities to attest to the validity of a digital certificate. Figure 32 illustrates the PKI scheme and its components.

At the heart of the PKI is the certification authority (CA) or trusted third-party (TTP) that generates and distributes the digital certificate. Part of the digital certificate includes a digital signature from the CA attesting to the validity of the digital certificate.

The next step is to establish the identity of the user of the digital certificate, which is accomplished by the registration authority (RA). The RA does not issue certificates but acts as an intermediary between the user and the CA. The role of the CA may be carried out by an actual human or by software running on a CA device.

What happens when a digital certificate expires or is revoked because it has been compromised? In that case, a certificate revocation list (CRL) is maintained at the CA and consulted each time a transaction takes place using a digital certificate.

Where Brocade is the CA, a PKI is used to distribute Brocade digital certificates.

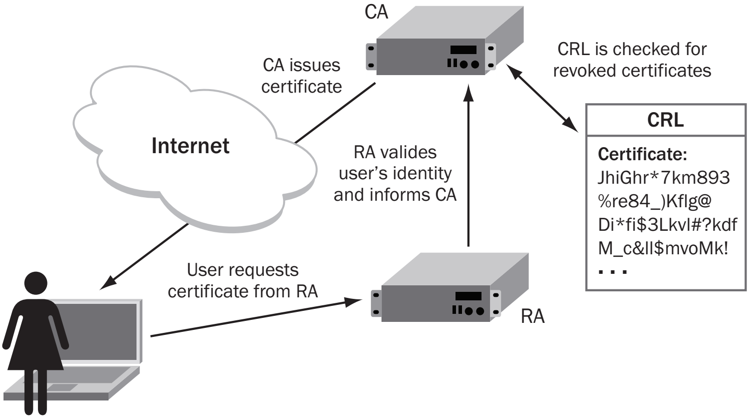
CA
CA issues certificate
CRL is checked for revoked certificates
Internet
RA valides user's identity and informs CA
CRL
Certificate:
JhiGhr*7km893
%re84_)Kflg@
Di*fi$3Lkvl#?kdf
M_c&ll$mvoMk!
. . .
User requests certificate from RA
RA

Figure 32. Public key infrastructure

SSL

Secure Socket Layer (SSL) was developed out of a need to encrypt communications over the Internet and addresses only the confidentiality of data-in-flight. It was originally developed by Netscape in 1994 and SSLv3.0 is the most widely used today.

It is a hybrid encryption system using both symmetric and asymmetric cryptography. Public key cryptography is used to authenticate between clients and servers, whereas symmetric cryptography is used to encrypt the application data. The application data can be encrypted using a 40-bit or 128-bit symmetric key version. The authentication is performed using digital certificates obtained from a CA in a PKI framework.

SSL can be used with several protocols, although it is used primarily with HTTP. For example, SSL is used to secure communications between a Brocade management graphical user interface (GUI) and a Brocade switch. The secured version of HTTP in this case is called HTTPS.

IPSec

IPSec (IP security) is a framework that performs encryption at the routing layer (IP - Layer 3) in the TCP/IP stack. IPSec is commonly used to secure communications in a virtual private network (VPN), but it can be used simply to encrypt communications between two devices on a network. IPSec can either encrypt only the payload or data (transport mode) or it can encrypt the payload and the header information (tunnel mode).

Since it is a framework, it does not actually specify which encryption or hashing algorithms to use but leaves this decision to the user. For example, IPSec can be used to encrypt communications between two Brocade 7800s using FCIP to replicate data between two data centers. Brocade's implementation of IPSec supports the following encryption and hashing algorithms.

Encryption algorithms:

- 3DES
- AES-128 (default)
- AES-256

Hashing algorithms:

- SHA-1 (default)
- MD5
- AES-XCBC

Key Management

The decision to encrypt information residing on disk or tape creates a long-term commitment and a dependence on the encryption keys. After being created, keys need to be backed up and managed. Keys can be lost, stolen, destroyed intentionally, or expired after a pre-determined period of time-all potential security vulnerabilities.

Loss of the encryption keys is comparable to losing the data. Unlike data-in-flight, the keys for data-at-rest must be available for as long as the data needs to be read. In the case of patient health records, information may need to be retained for seven years after the death of a patient, which could be over 100 years. Keys can also be stolen or compromised, in which case the information would have to be re-encrypted or rekeyed using a different key to ensure the confidentiality of the information.

Media such as disk and tape also have a limited shelf life and they go through evolution cycles to an eventually incompatible format. Encrypting data-at-rest requires management of the encryption key for the life of the data. Encryption keys are usually managed by a comprehensive key management system, because keys need to be managed for an extended period of time. A key management system is used to manage the lifecycle of keys. Encryption key information needs to be refreshed as the media expires and the data needs to be re-encrypted using a different key.

Finally, encryption keys need to be backed up in a secure manner to avoid being compromised in the process. Keys can be backed up to a key vault, which can be a feature of a comprehensive key management solution used to establish policies and manage the keys throughout their lifecycle. For redundancy, a typical key vault is implemented with two or more units to prevent single points of failure. If the primary key vault becomes unavailable, the secondary or clustered key vault can accept or provide keys to the encryption device.

Key management solutions are implemented using two basic methodologies to exchange the keys between the encryption device and the key management solution: trusted and opaque.

Trusted Key Exchange

Trusted key managers have the ability to securely obtain cleartext keys. To protect the keys during the transfer, a trusted relationship must be established between the two devices. For example, the device performing the encryption must be able to store the encryption key in the key vault. The encryption device and key vault must authenticate each other to ensure that both are authorized to exchange keys. When each device is authenticated and authorized, then the trusted relationship is established. An example of a trusted key exchange is shown in Figure 33.

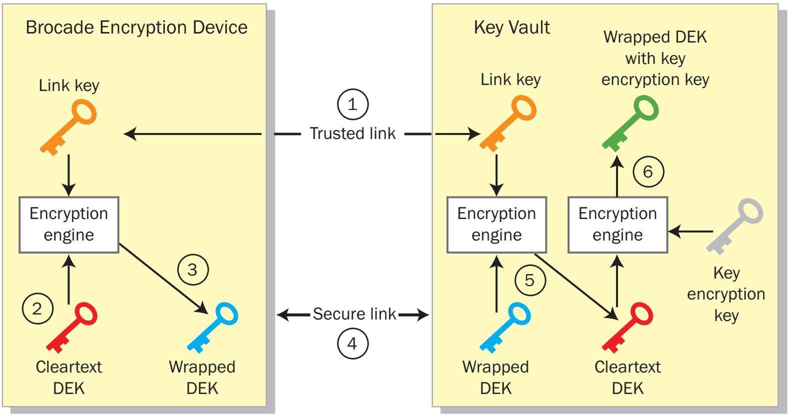

Figure 33. Trusted key exchange

To prevent key exchanges from being sniffed or intercepted during transmission between encryption devices and key vaults, most vendors use secure channels for the key exchange or wrap the key using a symmetric key before sending it over the channel. Many variations for the key exchange process exist. For example, one vendor uses a secure channel (SSL) and wraps (encrypts) the key before sending it across the secure channel.

Opaque Key Exchange

With opaque key managers, the key management application never has access to cleartext keys; the keys are always encrypted. One of the advantages of the opaque key management solution is that it does not require a hardened chassis and can be implemented using a software-only solution in a conventional server. Figure 34 illustrates the simplified opaque key exchange process.

One of the primary distinctions between an opaque key exchange and a trusted one is that the DEK is wrapped prior to sending it to the key vault, where it is stored as is. With a trusted key exchange, the wrapped key is unwrapped at the key vault and then re-wrapped using a different key encryption key. An opaque key vault does not contain information on how the DEK was initially encrypted.

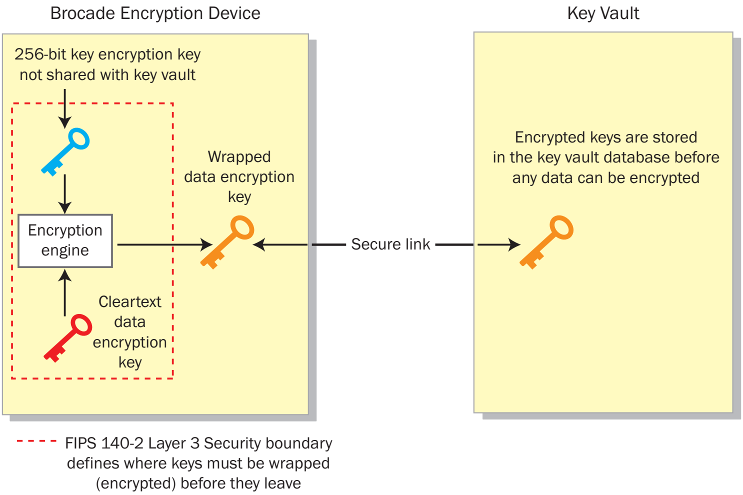
Brocade Encryption Device
Key Vault
256-bit key encryption key not shared with key vault
Encryption engine
Wrapped data encryption key
Cleartext data encryption key
Secure link
Encrypted keys are stored in the key vault database before any data can be encrypted
FIPS 140-2 Layer 3 Security boundary defines where keys must be wrapped (encrypted) before they leave

Figure 34. Opaque key exchange

Chapter Summary

Cryptography has been used for centuries and has evolved considerably with the arrival of computers and high technology. The development of networks and the Internet made it critical to develop new methods to exchange information securely across these new media—and the SAN is no exception. There are several vulnerabilities in a SAN environment that must be addressed using cryptography, such as the secure exchange of data across data centers, between switches and their management servers, or to ensure confidentiality of data-at-rest on disk or tape media.

Many of the technologies commonly used in conventional TCP/IP-based networks can also be used in SAN environments, particularly when protecting the management interfaces. Specific solutions exist to address requirements unique to SAN environments, such as authenticating devices joining a fabric using DH-CHAP.

6

FC Security Best Practices

In previous chapters, basic security and encryption concepts were introduced in addition to FC SAN basics. This chapter puts it all together so that you can apply these security concepts in the form of best practices in your FC SAN environment. It offers guidelines that can be used by SAN administrators and security professionals to help build a SAN security policy and decide which features should be deployed. Implementation of these features will be explained in greater detail in "Chapter 8: Securing FOS-Based Fabrics" starting on page 131.

When you design a SAN security policy, it is not necessary to implement and enable every available security feature. Some security features add performance overhead, others may affect administrator productivity, and yet others may have associated implementation costs. A balance must be struck between the features and the value of the assets being protected, and the probability that the system vulnerability will actually be exploited.

The Brocade SAN Security Model

The concept of defense-in-depth was discussed previously and can be extended to a security model for protecting a SAN. A complete SAN security strategy should provide multiple layers of challenge to an attacker to provide the best protection against all types of threats.

The diagram in Figure 35 illustrates the Brocade SAN security model. This is not an authoritative model but more of a guide to help visualize the components necessary to build a strong SAN security program.

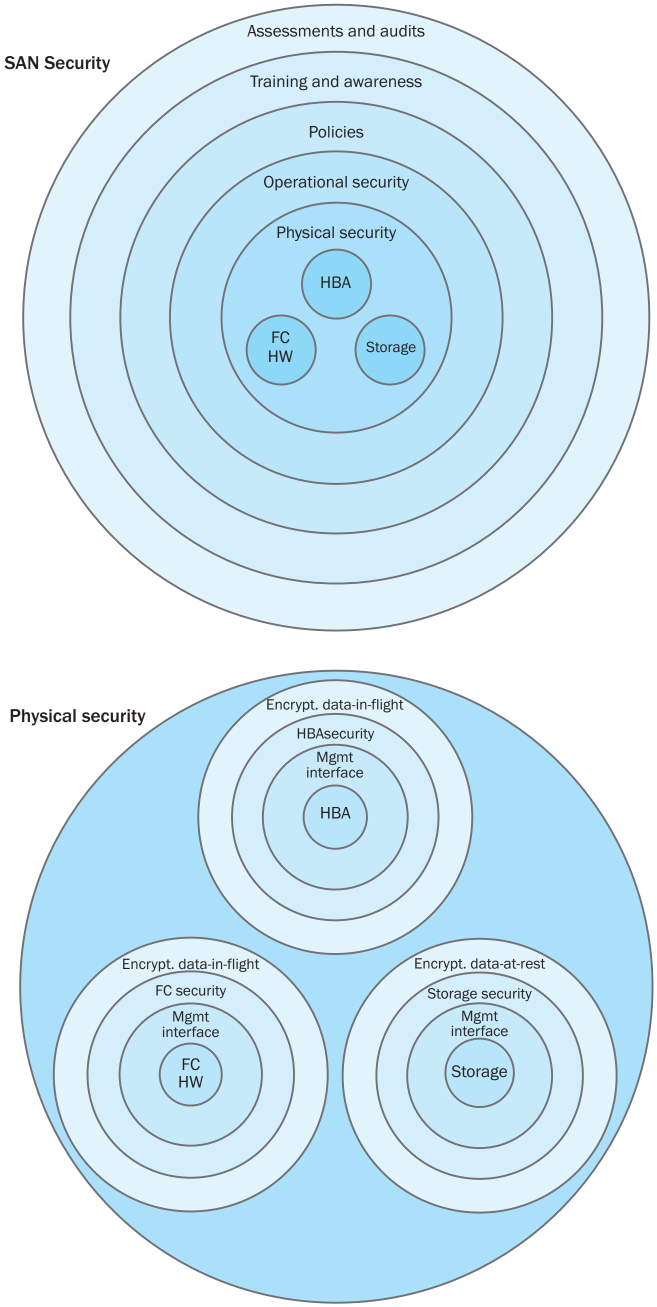

Figure 35. SAN Security Model

The model resembles an onion with layers of skin around its core (OK, perhaps a “mutant” onion with three cores). At the center, the three basic types of SAN hardware are represented by the three smaller circles for the HBA, storage hardware (disk and tape), and fabric infrastructure hardware (switches, directors, and routers). Surrounding the SAN hardware is a series of concentric circles outlining the layers of security to protect that specific type of hardware.

Protecting the HBA

The HBA is the link between the hosts running applications and the SAN, in which the application data is stored. The HBA requires specific protection measures since it has its unique set of vulnerabilities. To help mitigate the risks associated with these vulnerabilities, each HBA vendor offers specific security features and enhancements.

HBA security features can include the following:

- **LUN masking**, generally implemented at the storage device (disk array) level, however, some HBA vendors provide the ability to perform LUN masking at the HBA level.
- **DH-CHAP support**, as defined in the ANSI T11 FC-SP standard, to authenticate an HBA when joining a fabric. This is usually done to protect against WWN spoofing attacks.
- **Secure management interfaces and protocols** such as SSH to secure the CLI interface and SSL to secure the GUI interface.

How real is WWN spoofing? This question certainly comes up frequently and the answer is not always simple. It is possible to change the WWN on any HBA; tools are available from most HBA vendors that allow users to do this and some are also available as freeware.

For example, an attacker from the outside could theoretically compromise a server in a DMZ (although originally a military term, in computer security, Data Management Zone), reconfigure the compromised server's WWN to the WWN of a production server’s WWN, and capture data intended for the other host.

How likely is this to occur? With proper security measures in place, this type of attack is very unlikely. However, if the value of an asset is particularly high or attractive to someone, then the likelihood of an attack on that asset increases, no matter how difficult or sophisticated the attack needs to be. Generally though, only the most sensitive environments, such as military and intelligence organizations, as well as some private organizations, require protection against WWN spoofing attacks.

As a best practice, assign only one host, or initiator, per zone. Single-initiator zoning (SIZ) serves two basic purposes. A SIZ restricts host-to-host communications and limits RSCNs to the zones requiring the information.

HBA Best Practices Summary

- Use single-initiator zones
- Use secure management protocols when accessing the storage device management interface
- For the most sensitive environments, use DH-CHAP to authenticate storage devices joining a fabric

Protecting the Storage Devices

Disk storage devices actually store the data, which is the most valuable asset in a SAN environment, and thus requires the most rigorous security considerations.

Storage device security features can include the following:

- **LUN masking** is usually implemented at the storage device level using vendor-provided tools available through either a CLI or GUI.
- **DH-CHAP support**, as defined in the ANSI T11 FC-SP standard, is used to authenticate a storage device joining a fabric. This is usually done to protect against WWN spoofing attacks.
- **Secure management interfaces and supported protocols** such as SSH to secure the CLI interface and SSL to secure the GUI interface.

As a best practice, LUN masking should always be used to contain the visibility of a LUN to a specific host and to prevent other hosts from seeing LUNs not assigned to them. LUN masking combined with switch-based zoning offers the best protection to a LUN within a fabric.

Furthermore, it is advisable to create separate zones between the HBA and the disk or tape storage. Separating the disk and tape storage prevents RSCNs destined for disk devices from being propagated to tape devices, which are more prone to disruption resulting from an RSCN.

Storage Device Best Practices Summary

- Secure all LUNs with LUN-masking
- Assign disk and tape devices to separate zones to reduce RSCN risks
- Use secure management protocols when accessing the storage device management interface
- For the most sensitive environments, use DH-CHAP to authenticate storage devices joining a fabric

Protecting the FC Infrastructure

The FC infrastructure, composed of the switches, directors, and routers, is the heart of the SAN and all SAN data must pass through the FC infrastructure. The FC infrastructure is one of the most complex components of the SAN and, the more complex it is, the greater the potential vulnerability. Most FC switch vendors offer an extensive set of security features to secure the FC infrastructure.

FC Device Access Controls

The devices connected to the fabric are also vulnerable and require specific protection. For example, one of the advantages of a SAN is the ability to easily add a new switch into the fabric. A SAN administrator only needs to connect a new switch to an available port on an existing switch in a fabric using an ISL and power up the new switch. Automatically, a unique domain ID is assigned and the configuration files are downloaded to the new switch. From a security perspective, however, this time-saving administrative feature could be the security professional's worst nightmare; anyone with a switch could potentially connect to an existing fabric and gain control of the fabric. If an attacker with admin or root access on the rogue switch were to use this technique, they would now have admin and root privileges for the entire fabric.

There are several features available to prevent this scenario. The best defense is a defense-in-depth approach with multiple layers of challenges, described as follows.

The first and simplest line of defense is to persistently disable all unused ports. This will prevent someone without management access to the fabric from connecting a new switch and joining the fabric. It is important to use the persistent disable option to ensure that disabled ports remain disabled following a reboot or power off cycle. Otherwise, an attacker could simply cause a power failure on a switch to enable the unused ports.

The next line of defense could be to prevent ports from becoming E_Ports. In the event that an unused port is enabled, a switch would still be unable to join the fabric since the port will not be allowed to become an E_Port.

The next line of defense could be the creation of an ACL, specifying by WWN and/or switch port which devices are allowed to join the fabric. In Brocade fabrics, ACLs are used to control device access to a fabric:

- **The Switch Connection Control (SCC) policy** is used to specify the WWN of the switches allowed to join the fabric.
- **The Device Connection Control (DCC) policy** is used to specify which hosts and storage devices are allowed to join the fabric. The DCC policy can specify members by WWN but it can also lock down a WWN to a physical port on a switch so that only a specific WWN can connect to a specific port in the fabric, and all other devices connecting to that port will not be authorized. This could be considered a WWN anti-spoofing countermeasure.

The Brocade DCC and SCC policies are also known as fabric-wide consistency policies, since they can be distributed throughout an entire fabric and used in two different modes:

- In **strict mode**, all devices participating in the fabric must be defined in the DCC or SCC policy.
- The **tolerant mode** allows some switches to join the fabric without requiring them to be defined in the ACL. This is useful when the fabric contains switches running older versions of firmware (prior to FOS 5.2.0), which cannot use the FOS DCC and SCC policies.

As a best practice, it is advisable to use the strict policy as much as possible. A fabric is only as secure as its weakest link, and if one switch does not participate in an SCC policy then that would be the weak link and easiest target for an attacker.

The final and most sophisticated line of defense to prevent device access is to use a device authentication mechanism. The ANSI T11 standard FC-SP (FC security protocol) defines the DH-CHAP protocol for this purpose. Devices supporting DH-CHAP can be configured with a shared secret between the device and the switch, and only the device with the corresponding shared secret will be allowed to join the fabric. (see “Diffie-Hellman” on page 84). On Brocade FC switches, Fibre

Channel Authentication Protocol (FCAP) may also be used but it is not part of the defined FC standards. FCAP is considered more secure than DH-CHAP since it uses a stronger digital certificate instead of the shared secret used in DH-CHAP.

Clearly, it is not necessary to implement every one of these lines of defense to prevent FC device access to a fabric. The number of layers an organization decides to implement will depend on their business requirements, the sensitivity of the environment, and amount of risk accepted as tolerable. In reality, very few organizations implement all of these proposed levels of security. It is up to each organization to establish the risk and decide which features should be implemented.

FC Devices Best Practices Summary

- Persistently disable unused ports
- Prevent switch ports from becoming E_Ports
- Use the SCC and DCC policies to restrict device access by WWN and/or by a physical port on a switch
- For more sensitive environments, use DH-CHAP to authenticate devices joining a fabric
- Use a strict fabric-wide consistency policy where possible
- Distributed/fabric-wide consistency policies: strict vs. tolerant

Separation of Duties

The principle of separation of duties is used to restrict individuals to performing only the tasks necessary to perform their work-related activities and nothing more. When a backup administrator is given admin privileges on the backup and restore system, this does not imply that he or she should also have the same privileges on other systems. In some cases, a single task requires two or more people to complete it (similar to co-signing a check). An extreme example of this is nuclear missile launch codes, which must be entered by two different individuals to launch a nuclear missile. In a less extreme but still pertinent case in the storage world, requiring more than one person to destroy all of the keys on a key vault would be another example of separation of duties through the use of quorum. Quorum occurs when a pre-established consensus from a pool of individuals is required to accept or perform a task. A typical quorum would be two out of three or three out of five required persons to authorize a task.

Another method, discussed in greater detail in “Protecting Management Interfaces” on page 107, is to restrict the privileges granted to a user via a role-based access control (RBAC) by assigning specific roles to a user account. A role could be read-only and allow a user to only view information but not modify or delete it. At the other end of the spectrum is a role that grants full admin privileges; other roles are somewhere in between. Typically these roles are customized for specific types of functions such as an operator or a security administrator.

Physical isolation and routing. Separation of duties can also be accomplished through isolation of systems, each managed by a different individual. This method is used frequently in the data center where a separate SAN is built for each different group or project within an organization. It is often used by outsourcing firms with shared multi-tenant environments. Many of their customers prefer not to become part of the collective but rather being physically isolated from any other customer. A separate SAN can be constructed for each customer requiring physical isolation and a restricted group of administrators can be assigned to manage each environment.

While it's true that creating physically isolated SANs provides the ultimate form of isolation, it is also true that use of storage resources is not as optimized as in a shared environment. A compromise between a fully shared SAN and a fully isolated SAN is a logical SAN (LSAN), which places an FC router between the two SANs. For example, one environment may require its own hosts, applications, disk storage, and FC infrastructure to be managed independently from the shared environment. However, to avoid the additional cost of a tape backup system, an administrator can create an LSAN to enable sharing of the tape backup resources with the isolated SAN. This implementation provides physical isolation but has the advantage of sharing some resources according to strict pre-defined rules.

Note that FC routing can also be implemented to create LSANs using the Integrated Routing (IR) feature on Brocade 8 Gbps switches, available since FOS 6.1.1.

Zoning. When FC SANs first emerged more than a decade ago, there was no real access control mechanism to protect storage used by one host from being accessed by another host. This was not a significant issue at the time, since the original SANs were relatively small. Over time, however, as SANs became larger, more complex, and mission-critical to most data centers, this became a risk. To help secure particular devices and data, Brocade invented the concept of “zoning,” or restricting device communication only to member devices inside a given zone. Today, zoning is an accepted standard and plays an integral role in SAN security.

For Brocade switches, there are two ways to identify zone members, and zone enforcement is performed either in the switch Name Server or inside the switch ASICs. Identification methods include DID/PID (DP), the switch domain ID and the switch port number; and port World Wide Name (pWWN), the storage or host port WWN. Brocade recommends pWWN identification because of the management flexibility it provides; also, several advanced Brocade features require pWWN zoning.

Brocade switches that operate at 2 Gbps or faster enforce both DP and pWWN zones in hardware. This was not the case with Brocade 1 Gbps switches, and users frequently chose DP identification because it was the only hardware-enforced zoning method at the time. Now, a zone with all DP identification or all pWWN identification uses the more secure hardware enforcement.

However, there are some cases when mixing identification methods results in software enforcement. These cases include mixing DP and pWWN identification within a zone or using a DP identification for one zone and the pWWN attached to that DP in another zone. For this reason, Brocade recommends using the same zoning identification method (preferably pWWN) across the entire SAN to ensure that:

- All zoning is hardware enforced
- Advanced features such as Fibre Channel Routing are usable
- Zoning management methods are consistent

As a security best practice, organizations should use single-initiator, or single-HBA, zones. This means that each zone should have only one host defined, although it can have multiple target storage nodes. Single-HBA zoning improves security, helps contain RSCNs, and makes the SAN much easier to manage and troubleshoot. An extension of this best practice for mixed disk and tape traffic on the same HBA is to utilize two zones for each HBA: one for disk nodes and one for tape nodes. This approach isolates the disk and tape devices, even though they continue to communicate through the same HBA.

Another best practice is to activate Default Zoning. By default, if no zones are defined or the current zoning configuration is disabled, all devices can see each other in the SAN, which can create a variety of problems. First, the SAN is more vulnerable from a security perspective.

Second, HBA drivers can have difficulty discovering an entire SAN. The Default Zoning feature ensures that devices not already assigned to an active zone will be assigned to the Default Zone and will not be seen by other devices when an administrator disables a zoning configuration.

RSCNs are required for a SAN to function properly, but RSCNs can be potentially disruptive if not managed properly by the SAN switch. Brocade switches forward RSCNs only to zones with devices affected by the addition or removal of a device. Also, Brocade switches forward only one RSCN if identical RSCNs occur within a half-second window, an approach that limits the impact of a device sending hundreds or thousands of RSCNs per second. Furthermore, organizations can entirely suppress RSCNs on specific ports. Some applications, particularly in the video imaging and multimedia industries and tape backups, actually require this capability.

Finally, it is possible for a switch to obtain a new domain ID after a reboot, particularly when a switch is added to a new fabric or after a massive power failure. To prevent this from occurring, it is a best practice to assign a domain ID to a switch using an insistent domain ID (IDID). An IDID will survive reboots or power failures and will never change once assigned. Table 9 explains the domain ID behavior in insistent and non-insistent domain IDs.

Table 9. Domain ID behavior

DID Assigned?	Non-Insistent Domain ID	Insistent Domain ID
DID not in use	DID is assigned	DID is assigned
DID already assigned	New DID is assigned	Switch won't join fabric

Virtual Fabrics and Administrative Domains. Virtual Fabrics (VF) allows a physical switch to be partitioned into multiple Logical Switches, each with its own unique fabric ID (FID). Logical Switches can be connected to physical or Logical Switches, similar to a physical switch using ISLs, to form Logical Fabrics. This feature is very useful for multi-tenant environments and for environments that can benefit from a logical separation of data and management on a common physical fabric.

The Administrative Domain (AD) feature was introduced in Brocade FOS 5.2.0 and provides another method of partitioning a fabric into separately managed domains. An AD is a logical grouping of devices that can be managed separately either by the same or different system administrators. An administrator may have different privileges in one AD than in another AD. For example, a SAN administrator may have a read-only user role in AD 12 and an Admin role in AD 14.

You can use either the AD feature or the VF feature, but not both at the same time. VFs deliver more complete partitioning since this feature applies to the data, control, and management paths of the fabric, and AD applies only to the management path.

Traffic Isolation. Separation of duties can also be applied to some extent to the type of traffic so that one type of traffic does not affect another type. This is implemented using the Brocade Traffic Isolation (TI) feature. Figure 36 illustrates how TI works. In this example, data is replicated and backed up between Site A and Site B. Since tape backup is highly I/O intensive and data replication is exceedingly important, it is preferable to use TI to handle the data replication traffic on one ISL and the tape backup traffic on another ISL. Without TI, the data replication process would be competing for bandwidth with highly I/O-intensive tape backup. This could result in severely impacted replication performance and could cause problems with synchronous replication applications.

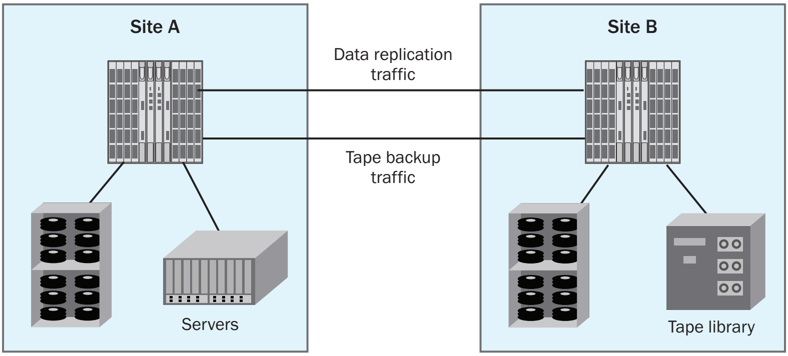

Figure 36. Example of traffic isolation

Separation of Duties Best Practices Summary

- In sensitive environments, separate fabrics fully
- Create LSANs to isolate fabrics but share resources such as tape backup systems
- Use pWWN when defining zoning members
- Use single-initiator zones
- Use default zones
- Use Virtual Fabrics to logically separate sensitive environments from less-sensitive environments
- Use traffic isolation to separate traffic to avoid contention

User and Password Management

Passwords and accounts are the gatekeepers to the management interfaces and must be protected to ensure they are not acquired by unauthorized individuals. The simplest method to prevent unauthorized persons from obtaining a password by sniffing traffic is to use a secure protocol, such as SSH or SSL or both, to access the management interface.

One of the most common ways to break into a SAN is by attempting to log in using the default passwords. One of the first items verified during a SAN security assessment is whether the default passwords have been changed for the standard accounts (user, admin, root, and factory). In recent Brocade SAN security assessments, one out of every four companies assessed had at least two switches still using the default password for at least one account. This is the simplest way to break into a corporate SAN and can be easily prevented by changing the default passwords during initial installation.

As a best practice, all individual users should have their own unique user account. The default accounts, such as root, factory, and admin, should never be used. Instead, individual SAN administrators should be assigned unique user accounts with the appropriate roles to allow them to carry out their daily management responsibilities.

Password policies can be defined to enforce basic rules on how passwords are created and managed. Passwords should be strong in the sense that they would be difficult to guess or break using a dictionary or brute force attack. The use of common words, or common number sequences, do not make good passwords; random combinations of at least eight numbers and alphabetic characters is typically a minimum. Passwords should not be reused on a regular basis and this can be enforced using the password history feature. Accounts should also be locked out after several (usually three) unsuccessful login attempts.

Finally, passwords should be forced to expire after a period of time, even though this is always a sensitive subject. Over time, a password will have a higher probability of being discovered and compromised, therefore it is important to change passwords on a regular basis. How often the password should be changed depends on the organization's specific requirements. If the password is changed too often, it becomes more difficult for users to remember the password and not confuse it with previous passwords. When forced to change their passwords too frequently, some users may simply resort to writing their password down and keeping it somewhere near the computer, and possibly accessible to others.

Since most system administrators are responsible for more than one system, a unique account administrator must be created on each managed device. The same goes for password changes, which must also be changed on each device the administrator is responsible for. To simplify this, a single sign-on method such as RADIUS or LDAP is recommended. These methods allow a SAN administrator to change a user's password for all servers in one centralized location.

User and Password Best Practices Summary

- Use secure channels
- Change default passwords on ALL default accounts
- Use unique user accounts with proper roles and privileges
- Create and enforce password policies (strength, history, expiration, and lockout)
- Use an account and password management method such as RADIUS or LDAP

SAN Availability

SAN availability is an important consideration when designing a SAN security program to protect against a targeted denial-of-service attack, natural disaster, hardware failure, or human error.

The key to maintaining high availability is to eliminate or reduce the number of single points of failure (SPOF). SPOFs can be found throughout the FC fabric, including:

- Hardware devices
- Paths between devices
- Data centers

Switch hardware availability. The hardware itself may have redundant components such as power supplies and fan modules. Some of these components may be “hot swappable” to allow replacement of field replaceable units (FRUs) in the field without bringing down the switch. Another solution for hardware redundancy is to use enterprise-class directors instead of switches. Directors offer greater hardware redundancy and overall robustness for maximum production uptime. The Brocade 48000 Director or Brocade DCX or DCX-4S Backbone, for example, can offer “six nines” (99.9999%) availability or better.

One of the simplest and best ways to eliminate a hardware SPOF is through the use of redundant dual--fabric architectures. In a dual-fabric design any single hardware component could fail without undue impact on the production environment (see “Dual Fabrics” on page 35). All hardware is duplicated in this architecture and there are two or more paths between any host and its associated storage. Of course, a dual-fabric architecture applies only to disk-based SANs, since backup applications cannot handle dual-attached tape devices.

One common availability error observed in many data centers using dual-fabric architectures is to co-locate both fabric A and fabric B in the same physical rack or cabinet in the computer room. Often this is the result of a procurement issue when the switches are initially purchased along with one single rack. Once a fabric has been racked and installed, it most likely wil never move again. Realistically, this should be planned before installing the FC equipment and a technology refresh or move to a new data center provides an excellent opportunity to fully separate the fabrics right from the start.

Data path availability. Redundant data paths between the host and storage devices are part of a dual-fabric architecture. Dual-attached hosts using MultiPathing I/O (MPIO) software can load-balance traffic between the two paths or they can fail over to a single path in the event of the failure of one path.

Data path redundancy can also be built into a fabric, using resilient fabrics or other architectures that provide path redundancy as discussed in “Chapter 3: SAN Basics for Security Professionals” starting on page 19. Some SAN designers simply use dual ISLs for redundancy between the switches instead of using single ISLs.

Data center availability. The data center itself can be an SPOF in the event of a natural or man-made disaster such as an earthquake, fire, or massive local power failure. This problem is addressed with multiple data centers maintaining replicated copies of data between them.

Fabrics in one data center can be mirrored in a second data center to create a “hot” site, which can be used to fail over all activity from the primary data center to the secondary. Exchanging data between the data centers can be done using dark fiber (depending on the distance and cost) or using the FCIP protocol over a public or private WAN.

SAN Availability Best Practices Summary

- Use switches with redundant, hot-swappable components or enterprise-class platforms (Brocade 48000, DCX, DCX-4S) for greater hardware availability
- Deploy disk-based SANs using a dual-fabric architecture
- Install fabric A and fabric B of a dual fabric in separate racks
- Use redundant data paths in the fabric design
- Use secondary data centers for Disaster Recovery (DR) and Business Continuity (BC)

Logging and Monitoring

The ability to track activity in a SAN certainly does not prevent attacks but it may act as a deterrent when it is known this is being done. Logs can be used to detect intrusions and provide evidence in legal prosecution of unauthorized users, as well as simplifying the troubleshooting process.

Log files. Since logs are ubiquitous and all IT systems use them, sophisticated attackers will often try to destroy the log files following an attack to remove any traces or evidence of their activities. To prevent or minimize the risk of this occurrence, the syslog should be redirected to a more secure alternate location, which can be done on most systems, including on Brocade switches.

Different FC equipment vendors offer various levels and types of logging, but they are often not enabled by default. For example, on Brocade switches, the following logging features are not enabled by default:

- Event auditing
- Track changes
- Fabric Watch security class

To obtain more detailed logging, these logging features should be enabled.

An often neglected but important detail with log files is the time stamp. Switches and other FC devices in a SAN run their own internal time clocks. Without any means of synchronization, the clocks on each device will be different and make it virtually impossible to correlate an event in the log file of one device with the log file of another device. This problem can be resolved simply by using the Network Time Protocol (NTP) to synchronize the time on each device. This can be accomplished by specifying an NTP server, either an external or internal one, and each device will then synchronize its internal clock with the NTP server.

Monitoring. When a security breach occurs, it is imperative to detect it as soon as possible to allow for a quick response and prevent or minimize damage caused by the attack.

All FC switch vendors provide a GUI to manage their switches. The GUI is usually the primary management tool to monitor the status of the SAN in real-time. Unfortunately, a critical event may not be observed immediately unless a SAN administrator is posted in front of the GUI at all times and pays constant attention.

To automate monitoring, other tools can provide automated alerts in the form of e-mail notifications or pages. Often, SNMP (Simple Network Management Protocol) is used to send traps to a third-party management framework. Since the SNMPv1 protocol has known vulnerabilities, as a best practice use SNMPv3.

With Brocade switches, the Fabric Watch feature can be used to monitor specific fabric and switch events and generate an SNMP trap or send an e-mail to alert administrators of the event. Specifically, Fabric Watch has a Security class to monitor specific security events such as unsuccessful login attempts and device access control policy violations.

Logging and Monitoring Best Practices Summary

- Redirect the syslog to a secure server
- Enable all event auditing and change tracking features
- Synchronize all switches and directors using NTP
- Use real-time monitoring of important security events
- Monitor all logs on a regular basis
- Use a real-time management tool to monitor security events
- Use an automated alerting method for notification of security breaches

Protecting Management Interfaces

This has been repeated several times so far in this book but it cannot be emphasized enough: management interfaces are one of the most vulnerable points in the SAN. It is important to use appropriate procedures and protocols when using the management interfaces for all components in the SAN including the HBA, storage devices, and FC infrastructure devices (switches, directors, backbones, and routers).

One of the simplest techniques for protecting management interfaces is to use a separate LAN, subnet, or VLAN to isolate the management network from the production network. This limits access to the management network to SAN administrators only and not to the company at large.

Since insiders can be a significant threat, it is also good practice to use secure protocols to encrypt the communications between management workstations and the devices being managed. This can be done using protocols such as SSH and SSL/HTTPS. If secure protocols are used, then it is equally important to disable the equivalent unsecure protocols. For instance, if the security policy now requires administrators to log into switches using SSH, then telnet access (port 23) should be disabled. If HTTPS is the protocol of choice, then HTTP (port 80) should be disabled.

For more secure environments, it is possible to restrict management to one specific management point in a fabric. It is common practice in physical security to have only one entry and exit point into a facility, since it is easier to manage and control a single entry point than it is to control multiple entry points. This is the primary reason why most enterprise computer rooms only have one access door.

Fabric management can be performed using any switch in a fabric, which means that multiple management points are available. Using a fabric configuration server (FCS) policy, administrators can specify a specific switch as the only management point and they can also assign alternative switches as backup management switches in the event the primary management switch fails.

Management Interface Best Practices Summary

- Use a separate LAN or VLAN for the management network—never use the production network for the management interfaces
- Use secure protocols to access management interfaces (SSH, SSL)
- Disable the equivalent unsecure protocols
- Limit the points of entry for management (use FCS policy if necessary)

Maintaining Data Confidentiality

Confidentiality as it pertains to electronic data is the protection of information from being disclosed to unauthorized users. In the context of a SAN, data is either in flight on a cable (data-in-flight) or at rest on a storage media, tape or disk (data-at-rest); both discussed in the following sections.

Encrypting Data-in-Flight

Encrypting data-in-flight uses a different encryption method than encrypting data-at-rest. Data-at-rest on a disk is block based, and when it is written back to disk the encrypted data must be exactly the same size as the cleartext data before encryption. With data-in-flight, the data is streaming over a cable in a serial fashion and needs to be encrypted on the fly as it moves across the cable. The concepts of stream and block ciphers were discussed in "Cryptographic Algorithms" on page 78.

Data-in-flight can be found at three different points within a SAN:

- Between a host and the fabric
- Between two switches
- Between the fabric and a storage device

Host-to-Fabric Encryption

Protecting the confidentiality of the data exchanged between the host server or HBA and the fabric is accomplished by encrypting the data-in-flight and can be implemented in several ways. Software-based encryption applications can be installed on the server, but as with any software-based encryption solution there will be a negative performance impact of 30-50%. This may be acceptable for some applications and environments, while others may not tolerate any performance degradation.

Hardware implementation is the only implementation that does not impact performance, but it is not a feature currently available on HBAs. The cost of implementing host-based encryption is relatively inexpensive for small environments, but the cost increases rapidly as the number of hosts increases in the fabric.

Switch-to-Switch Encryption

The FC infrastructure is a highly intelligent and reliable transport network that moves frames between servers and storage devices. All of the data in a SAN environment moves through this infrastructure and is usually transmitted in cleartext. As was shown in "SAN Security Myth Number 3" on page 11, data moving through a fiber optic cable can be sniffed without splicing the cable or breaching its protective jacket. Data can also be moved across a public network using the FCIP protocol, which uses a TCP/IP tunnel to move an FC frame. FCIP is particularly vulnerable since it uses the TCP/IP protocol, along with all of its associated vulnerabilities.

The primary confidentiality issue with switch-to-switch communications is not over the ISLs used to connect switches in a data center, but between switches that connect two data centers over distance. A dark fiber strand that is owned or leased by the organization is used to connect two data centers.

Securing Dark Fiber

Securing a dark fiber connection is similar to the techniques described earlier. Intercepting communications on a dark fiber, as was discussed earlier, does not require expensive equipment or physically cutting the cable. To prevent an attacker tapping into the fiber, dark fiber cables can be protected in pressurized tubes. Changes in gas pressure in the tube indicate physical tampering. However, this type of attack is extremely rare and this solution is usually implemented only in governmental security organizations and some financial institutions.

With the latest Condor-3 ASIC, the Brocade 16 Gbps FC products are capable of encrypting up to two ISLs per ASIC at full 16 Gbps line rate. Furthermore, the Condor-3 ASIC can also first compress data, then encrypt if both features are required.

Securing WAN Connections

Securing a SAN connection across a TCP/IP transport requires extra consideration, since IP is much easier to access than dark fiber. Protocols such as iSCSI and FCIP are based on TCP/IP, and security in these types of networks should include standard security mechanisms normally used with a conventional LAN. To maintain confidentiality of data-in-flight, encryption protocols such as IPSec are commonly used to encrypt data travelling across TCP/IP networks. IPSec can use different encryption algorithms to perform the actual encryption.

Most FCIP solutions on the market offer high-performance, hardware-based encryption using IPSec, including the Brocade 7500 Extension Switch and FR4-18i Extension Blade.

Fabric-to-Storage Encryption

There are no data-in-flight encryption solutions available today to encrypt the data between a fabric and a storage device other than a specialized encryption appliance, which can be installed in the data path. However, a data-at-rest encryption solution may accomplish this by encrypting the data prior to sending it from the fabric to the storage device, thus ensuring data confidentiality.

Encrypting Data-at-Rest

Data-at-rest includes tape and disk media, which require different encryption methodologies. Disks are block-based devices and tapes are streaming devices, which usually require different modes of operation to perform the encryption. Encryption of data-at-rest can be performed in several places in the SAN, as shown in Figure 37.

- Application
- Appliance
- Fabric/network itself
- Host
- Disk
- Tape

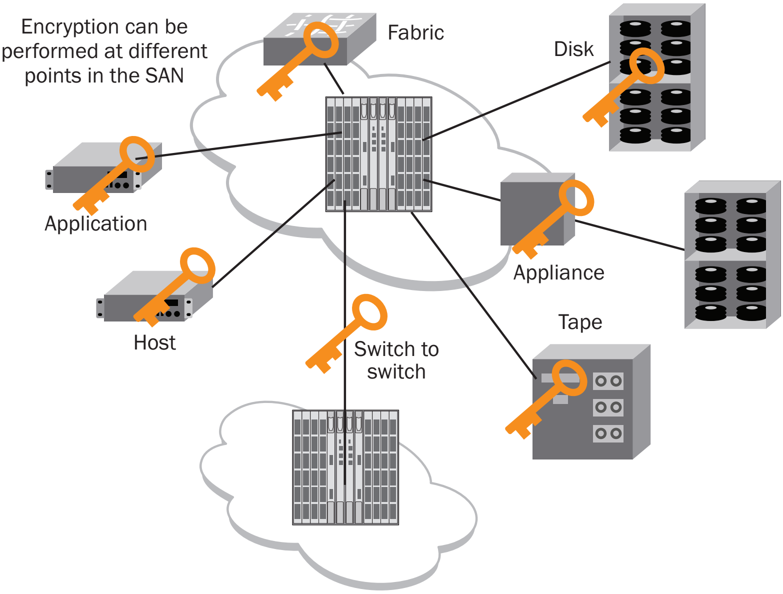

Figure 37. SAN encryption points for data-at-rest

Application-Based Encryption

There are several schools of thought as to where encryption should take place. Some applications actually require the data to be encrypted at the application level to prevent unauthorized users from viewing certain types of data. For instance, it is possible to encrypt an entire column containing sensitive information in a database using encryption software. Of course, any software-based implementation would negatively impact performance. Performance-sensitive and mission-critical applications may not be well-suited to this type of encryption. Furthermore, the cost of implementing application-based encryption can be quite expensive, as it requires modifying production code. Although this may be justified for some applications, it is not easily scalable to other multiple applications across a typical production environment.

Some backup applications offer an encryption module to encrypt the data to the backup media. The encryption module is built into the backup application software, but this method utilizes processing cycles on the backup server resulting in a negative performance impact, which increases the backup window.

There are also specialized applications designed to encrypt data-at-rest to disk or tape media. Several vendors such as RSA and PGP offer such solutions. Again, the main issue with software-based solutions is performance degradation and the impact on production server and application performance.

Appliance-Based Encryption

Appliance-based encryption solutions do not become a part of the fabric and must be inserted in the data path between the host and the storage to encrypt the data. The process of inserting the appliance in the data path may cause a disruption of the production environment.

Fabric-Based Encryption

Fabric-based encryption is accomplished using switches with encryption and compression capabilities. These switches can be added to an existing fabric using standard ISLs and assigned a domain ID, as with any other FC switch. One of the main advantages of the Brocade fabric-based encryption solution over appliance-based solutions is the ability to redirect or reroute frames from anywhere within the fabric through the encryption switch. Brocade FC switches use a technology known as frame or nameserver redirection, which was introduced in Fabric OS 5.3. Frame redirection enables a transparent integration of the encryption solution into an existing fabric. Data can be written from servers to storage devices anywhere in the fabric without requiring direct insertion of the switch into the data path.

Another significant advantage of fabric-based encryption is the ability to encrypt data in a heterogeneous environment. Some solutions, such as the Brocade Encryption Solution, encrypt data directed to both tape and disk devices and also work with a variety of third-party vendor appliances. This provides organizations with greater flexibility and independence from the storage vendors.

Host-Based Encryption

Host-based encryption can be implemented using software installed on the host. The greatest issue with host-based software encryption is the negative performance impact resulting from CPU utilization of the encryption application.

Storage-Based Encryption

Tape-based hardware encryption solutions have the advantage of being implemented in hardware and operating at wire speeds with no observable performance degradation during a backup operation. On the other hand, these solutions require new specialized tape drives with built-in encryption capabilities (such as LTO-4 or -54).

Although this solution addresses the tape encryption problem quite effectively, it does not address disk encryption. Many organizations begin with a data-at-rest encryption project exclusively to address a tape encryption problem. However, even without an internal policy, it is highly likely that regulations or legislation will eventually force the encryption of both disk and tape media. Addressing the disk encryption requirement would require a disk encryption solution that uses different encryption hardware.

Disk-based or array-based hardware encryption solutions are now available from several vendors. Similar to tape encryption, disk-based encryption addresses disk encryption effectively. Disk-based encryption does not, however, address tape encryption. Furthermore, rekeying of data (re-encrypting a LUN with a different key) can only be performed as a data migration process with current disk-based or array-based encryption solutions.

Physical Security

Physical security is a vast subject and this book cannot do justice to the topic. However, some best practices that apply to the SAN environment are highlighted in this section. Most organizations assessed by the author were found to have adequate physical access controls to the computer room and the SAN equipment; hence, this aspect of physical security will not be addressed here. Note that this area of security is generally addressed by a different group than the storage or security administrators.

As mentioned previously, one of the most frequently observed oversights in data centers is to physically install the switches in a dual-fabric configuration in the same rack or cabinet. One particular customer's data center was located in a room on the floor underneath the cafeteria, and as Murphy’s law would have it, a water leak from the cafeteria made its way into the computer room. Fortunately, this leak did not damage the SAN equipment, but if it had, the entire SAN would have failed, along with all of the application servers and storage devices, resulting in a massive outage.

Most of an organization's critical applications reside in the SAN and a loss of the SAN is disastrous. Simply installing fabric A equipment in one rack and fabric B equipment in another rack would address this issue.

In shared environments particularly, it is good practice to lock racks or cabinets containing switches and SAN equipment. In some shared environments with isolated SANs, the entire SAN can be enclosed inside a locked wire cage structure to prevent unauthorized access.

The final aspect of physical security considered during a SAN security assessment are the environmental and utility factors.

Power feeds and circuits should also be redundant to connect one switch power supply into one circuit and the other power supply into a different circuit. The equipment should be protected with an uninterruptable power supply (UPS) system and the UPS system should be tested and batteries replaced regularly.

To protect against a loss of availability resulting from a power failure, data centers should also use power generators, exercised on a regular basis to be certain they will function properly when a power failure occurs. Contracts should be in place with a service level agreement (SLA), guaranteeing a pre-determined response time from identified diesel fuel providers in the event of a power failure. A massive power grid failure similar to the one experienced in northeastern US and Canada on August 13, 2003, could result in hundreds or thousands of data centers within a large area scrambling for available diesel fuel to refuel their power generators.

The equipment in a computer room not only consumes enormous quantities of power but generates so much heat that a complete failure of the cooling system could result in a shutdown of the entire computer room within a few hours. It is important to ensure there is proper cooling in all areas of the computer room and to eliminate any hot spots in an aisle or other area.

Physical Security Best Practices Summary

- Use separate racks for fabric A and B
- Lock cabinets in multi-tenant or sensitive environments
- Use separate power circuits for redundant hardware components
- Use a UPS and test and change batteries regularly
- Use proper cooling and avoid hot spots
- Use power generators and exercise them regularly; establish SLAs with fuel delivery providers

Operational Security and Procedures

According to several studies, insiders are still responsible for the majority of security breaches through inadvertent mistakes. The human element is the least predictable factor in a production environment and mistakes are frequent. Mitigating risks associated with human error usually involves eliminating the human element whenever possible. This can be accomplished through automation of tasks using scripts, third-party management software, and custom applications. These risks can also be reduced by eliminating the guesswork in operations by creating detailed and well-documented operations procedures.

In some organizations, only a single individual understands and knows how to manage the SAN environment. Again, Murphy's Law will invariably ensure that this one SAN administrator will leave unexpectedly or get hit by the proverbial bus. Properly documented procedures will enable another system administrator to at least perform the essential functions to continue operating the production SAN. It is not necessary to document every single procedure on the SAN, but the critical tasks, and those that are used frequently, should be documented.

Switch configuration files should be backed up frequently, depending on how often changes are made to the production environment. The same applies to syslog and other log files. They should all be backed up automatically to a secure server with restricted access.

Operational Security Best Practices Summary

- Document critical and frequently used operations
- Back up configuration files automatically
- Back up log files automatically

Training and Awareness

Properly documented procedures certainly mitigate the risks associated with SAN administrators making mistakes in day-to-day operations. However, personnel must also be trained on how and when to use these procedures. This author has seen many procedures created with the best of intentions and then left in the darkest recesses of the computer room or buried under a pile of manuals on the SAN administrator's desk.

In many organizations, training is left up to the individual and RTFM (Read The "Fine" Manual) is the order of the day. This method may be initially less expensive for the employer, but it can lead to costly errors. Trying out commands that are best left alone or not fully understanding all the consequences of executing certain commands can lead to dire circumstances. One particularly dangerous practice is to use the root account to execute generally undocumented commands. The root account should NEVER be used by system administrators and should be used only by a vendor representative or under the guidance of the vendor.

Most people never read an entire manual from cover to cover but rather focus on the sections that address the immediate requirements to complete a given task. Formal training can provide greater efficiencies in the long run by providing best practices, along with commands to perform complex operations more efficiently.

Training should be aimed not only at the storage administrator but also at the security administrator and IT management.

Security awareness helps prevent security breaches by sensitizing the staff to security issues and attack methods used by hackers. One of the most frequent attacks by hackers to obtain passwords and other sensitive information is social engineering. For example, a hacker calls a corporate user impersonating an official company help desk or support person and requests the user's password or other sensitive information. Users should NEVER divulge their password to anyone, even a real company help desk person.

Training and Awareness Best Practices Summary

- Train staff, security and storage administrators as well as management
- Raise security awareness of the risks/vulnerabilities involved in a SAN environment

Policies and Plans

The SAN security policy outlines the spirit of how the SAN environment should be managed and operated.

All enterprise data centers have an IT security policy in place but very few have a specific SAN security policy. In itself, this is not so terrible, but inexplicably IT security policies are seldom applied to the SAN environment, and the truth is that most IT security policies could be applied directly to the SAN environment without any change. For example, it may be a defined policy requirement to use SSH instead of telnet to access the CLI on a server. This could easily be extended to the SAN environment and SSH could be used to access the CLI on switches and the other SAN devices. In the absence of a specific SAN security policy, SAN administrators should follow the spirit of the IT security policy to manage their SAN.

A disaster recovery or business continuance plan is another useful component of an IT security strategy. As with policies, DR/BC plans in place do not usually include the SAN environment specifically. The SAN environment does have a specific architecture, configuration information, accounts and passwords, and a host of other features that need to be replicated in the event of a disaster. If a specific DR or BC plan for the SAN is not feasible, then the existing DR/BC plan should include the SAN environment.

The last plan of concern is the computer security incident response (CSIR) plan. An incident response plans outlines in detail what needs to be done in the event of a security breach. It usually involves the creation of a CSIR team (CSIRT), which will be mobilized when an incident occurs. The CSIRT is usually composed of employees from various groups across an organization, for whom this is not their primary role; although technical people are required to address the technical aspects of the response. There is a need for Human Resources specialists to deal with HR issues (such as dismissing an employee) or public relations issues resulting from the incident in order to prevent further exposure. Management should be represented to make rapid high-level decisions to minimize the impact of an incident and enable a proper response if unexpected costs are involved.

It is not necessary to have a CSIR plan or a team specifically for the SAN environment, but members of the SAN management team should certainly be involved in building the CSIR plan and participating in the CSIRT.

Policies Best Practices Summary

- Apply IT security policy to the SAN environment
- Develop a specific SAN security policy
- Build a DR/BC plan for the SAN or integrate the SAN into an existing DR/BC plan
- Participate in the company CSIR plan and team

Assessments and Audits

Although the terms “assessment” and “audit” are sometimes used interchangeably, there is a subtle but important difference between the two. An audit is a verification process to establish whether a policy or industry standard is being followed. To perform a SAN security audit, an industry standard or an internal SAN security policy must exist. Since there are no official industry standards in the US currently, this leaves an internal SAN security policy as the basis for an audit. The Storage Network Industry Association (SNIA) sponsors a Storage Security Industry Forum (SSIF) that has been developing a Best Current Practices (BCP) document for storage security. (See “Appendix B: Standards Bodies and Other Organizations” starting on page 205.) This document is not an accepted industry standard, but it is a good reference to help security and storage professionals build an internal SAN security policy.

An assessment, on the other hand, is not formal and its scope is not restricted to boundaries established by a policy or industry standard. An assessment is complementary to an audit. Some organizations with internal security policies perform yearly audits, but they also perform a comprehensive assessment to validate, expand, and update the existing security policy.

Usually, organizations without a SAN security policy in place, or those who want to integrate the SAN environment into the existing IT security policy, have an assessment performed by a third-party vendor specializing in SAN security.

Chapter Summary

Although there are no industry standards currently defining the requirements to secure a SAN environment, organizations such as the SNIA SSIF and Brocade are working on raising awareness around SAN security and have developed storage security best practices to help organizations better understand the security issues surrounding a SAN.

Brocade has been involved heavily in this area and has developed over 100 different security features to help harden an FC switch infrastructure, created several security-related white papers, and provided professional services engagements to assess, audit, and harden a SAN. Hopefully this book will also contribute to SAN security cause by providing a better understanding of the security issues associated with a SAN and raising awareness of these issues.

7

Deploying SAN-Attached Devices in a DMZ

A DMZ (demilitarized zone) is a part of the network that sits between the internal private network and the external network or Internet. The DMZ also acts as a buffer between the inside and outside networks where applications such as e-mail, FTP, and Web servers exchange information between both networks. This buffer is critical for preventing potential attackers from the outside network, or Internet, to communicate with any of the internal systems directly.

A SAN is a separate network from the LAN, which is used to exchange information between servers and storage devices such as disk arrays and tape devices. SANs are currently implemented in the data center using three protocols: Fibre Channel, iSCSI, and the recent FCoE and DCB protocols. This chapter focuses on the FC protocol since it is by far the most widely deployed. From a security perspective, there are clearly concerns with connecting servers located in a DMZ, which are accessible from the Internet and whose storage is connected via a SAN. The greatest fear is that a SAN-attached server in a DMZ will be compromised and somehow used as a stepping stone to gain access to the SAN itself. The next question becomes whether securing SAN-attached devices in the DMZ can be done safely or not.

Certainly, there are risks involved in having a SAN in a DMZ, but with proper design and configuration it can be implemented with a high degree of safety. Note that vulnerable SAN components must be properly secured before attempting this. It has been explained previously that security is not always for preventing criminal activities originating from outside the boundaries of the data center. Security measures must be put into place to prevent unauthorized internal breaches and prevent the propagation of human error beyond a fixed scope.

A server connected to a fabric can potentially see all of the storage devices and servers connected to the fabric unless proper measures are taken. This chapter explores several techniques that accomplish the secure configuration of a SAN containing servers that are within a DMZ. The subject matter is intended to enable experienced SAN or security administrators to address DMZ-related security concerns. If you are an avid reader and intend to read the entire book, you will notice several redundant sections in this chapter. Please feel free to skip over them if desired.

Securing the Management Interfaces

Every FC switch has an Ethernet port used as a primary management interface. Switches are usually managed using a command-line interface (CLI) or a graphical user interface (GUI) via the Ethernet port using an IP address. It is extremely important to ensure that management interfaces are located on a network segment that is isolated from the Internet, and if necessary the production network as well. The switch management interfaces should never be accessible from the Internet, at least not without a secure VPN (Virtual Private Network). This can be implemented in several ways using one or more of the following technologies: a separate physical network, non-routable subnets, VLAN (typically, a private VLAN is used), ACLs, Policy-Based Routing (PBR), and/or firewalls (implementing a VPN).

The management interfaces should also be used in conjunction with secure protocols such as SSH, SSL (HTTPS, SCP, SFTP), and SNMPv3. Conventional protocols such as telnet, HTTP, and SNMPv1/2 exchange data in standard readable or cleartext format and should be disabled once the secure protocols are configured. Information such as passwords and user IDs can easily be captured using network sniffing tools. Secure protocols such as SSH and SSL use encryption algorithms to protect unauthorized viewing of data, including passwords and user IDs.

User accounts and passwords are the first line of defense for a network device management interface and are an important component for preventing unauthorized access. It is important to assign a separate account to each individual administrator who has access to the switches instead of a shared account between all or some administrators. Role-Based Access Controls (RBAC) are used to assign specific rights, which are tied to a person’s user account on the network or on a particular device.

The factory default passwords for all default accounts must be changed before a network device goes into production. This is usually done during the initial switch configuration. Companies should have policies requiring strong passwords and the periodic changing of those passwords. This includes forcing at least eight characters, using a combination of alphabetic, numeric and special characters, and preventing the use of repeating characters and sequences. The password policy should also set a password expiration time and disable accounts after a number of unsuccessful login attempts. To simplify password management, a single place to administer user names and passwords for all users and devices is in large environments. RADIUS (Remote Authentication Dial In User Service) and the LDAP (Lightweight Directory Access Protocol) are tools that provide a simple, centralized method to enable and disable user accounts and change passwords for all switches in a SAN from one centralized location.

Securing the Servers in the DMZ

The servers in the DMZ should be secured using conventional security techniques such as firewalls, anti-virus software, and other methods. If a server in a DMZ is compromised and an attacker manages to gain control of the server, he or she can now access the storage devices attached to that server via the SAN, just as they would if they were direct-attached. Since these servers are connected to the internal network, the internal IP network is now at risk of attack from a compromised server. Firewalls are commonly deployed to provide a barrier to protect the internal network.

Can an attacker use this tactic as a stepping stone into other storage devices and servers on the SAN? Is the DMZ a potential entry point to perpetrate an attack on the SAN itself? These are the most frequent concerns expressed by security professionals when deciding whether a server in a DMZ should be connected to a SAN. The following sections discuss several methods for preventing attackers from using a server in a DMZ to gain access to other servers or devices connected to the SAN.

Securing the Storage Devices

There are several common techniques that prevent a server from being able to see or access storage to which it is not explicitly assigned.

Port Disable, Disable E_ports, Port ACLs

The initial step to control device access to a fabric is to disable any unused FC switch ports and then prevent them from becoming E_Ports. This prevents unauthorized hosts or storage devices from joining a fabric by connecting them to an unused FC port. Port ACLs, such as the Brocade DCC (Device Connection Control) policy, should be used to lock a particular host or storage device WWN to a physical port on the FC switch. This prevents a host with the same WWN as another production host on the fabric from joining the fabric. While these port control methods add some additional management steps to the configuration of an FC switch, they significantly increase the security of the switch and reduce the risk of rogue devices joining the fabric.

Zoning

Zoning is a common technique implemented within an FC fabric. Zoning allows for devices such as servers, disks, and tape drives to be grouped together and isolated from other devices. Devices can only communicate with other devices that are in the same zone. All Brocade FC switches are capable of hardware-enforced zoning, in which an ASIC enforces the decision to allow or prevent devices to communicate with each other, as defined by the zoning configuration. Hardware enforcement is always performed on Brocade FC switches if all zone identification in a zone configuration is DID/PID (port zoning) or pWWN (WWN zoning). Mixing identification methods in a zone configuration will cause the zone enforcement to be performed by the less secure Name Server enforcement method. Brocade recommends using all pWWN definitions when configuring zoning to ensure that all zones are hardware enforced and to enable some advanced Brocade features such as Fibre Channel Routing (see “Zoning” on page 29).

LUN Masking

LUN masking, can be implemented on the HBA (Host Bus Adapter) or on the disk controller. This feature assigns a specific LUN to a specific pWWN in the SAN. No other server will be able to see or access that LUN unless multiple LUN masking mappings are configured. Typically, LUN masking will be configured on the storage subsystem. LUN masking is less effective when it is configured only on the server, since the masking can be disabled if the server is compromised. A server breach is more likely than a storage subsystem breach.

Administrative Domains

The last technique, Administrative Domains (ADs), are used to logically group FC switches, switch ports, and device pWWNs (in a physical fabric) that should be managed separately from other components in the fabric. Zoning logically groups devices that communicate with each other, while ADs create a Logical Fabric, with logically grouped devices, that can be managed independently as though they were separate switches.

For example, a SAN administrator may want to create a Logical Fabric to allow a sensitive project in a shared SAN environment to be managed separately from the rest of the production environment and to further isolate and protect it. Privileges can be assigned to a SAN administrator to manage the special project environment and different privileges could be assigned to the same administrator to manage the shared production environment. One benefit of this approach is that changes in the special project environment will not cause any disruption to the shared production environment and vice versa.

Authentication of Servers

To further enhance security, strong authentication mechanisms should be employed to authenticate servers joining a fabric. The ANSI T11 Technical Committee for FC has a standard that defines the use of an authentication protocol to authenticate end devices to switches. This protocol, DH-CHAP, uses a shared secret to ensure that the pWWN of the HBA joining the fabric has not been spoofed and is in fact genuine. It is possible to change the pWWN on an HBA using tools from HBA manufacturers, and so it would be possible for someone to configure the HBA on a server to have the same pWWN as another server on the SAN. Use DH-CHAP and port ACLs to prevent spoofing of a server HBA pWWN. Since DH-CHAP also requires end-devices that support DH-CHAP and requires additional management overhead at initial configuration, very few organizations truly feel the need to implement this feature.

Physical Separation of the Fabric

Another method for protecting the production SAN is to simply use a separate physical SAN dedicated to the DMZ. A separate group of FC switches could be used to connect all servers inside the DMZ. Storage devices could also have dedicated ports attached to this switch or entire storage devices could be dedicated to the DMZ servers. Although this is probably a more secure solution, it requires dedicated hardware and decreases optimization of the storage devices.

The diagrams in Figures 38, 39, and 40 illustrate both improper and proper methods of connecting servers in a DMZ to a SAN.

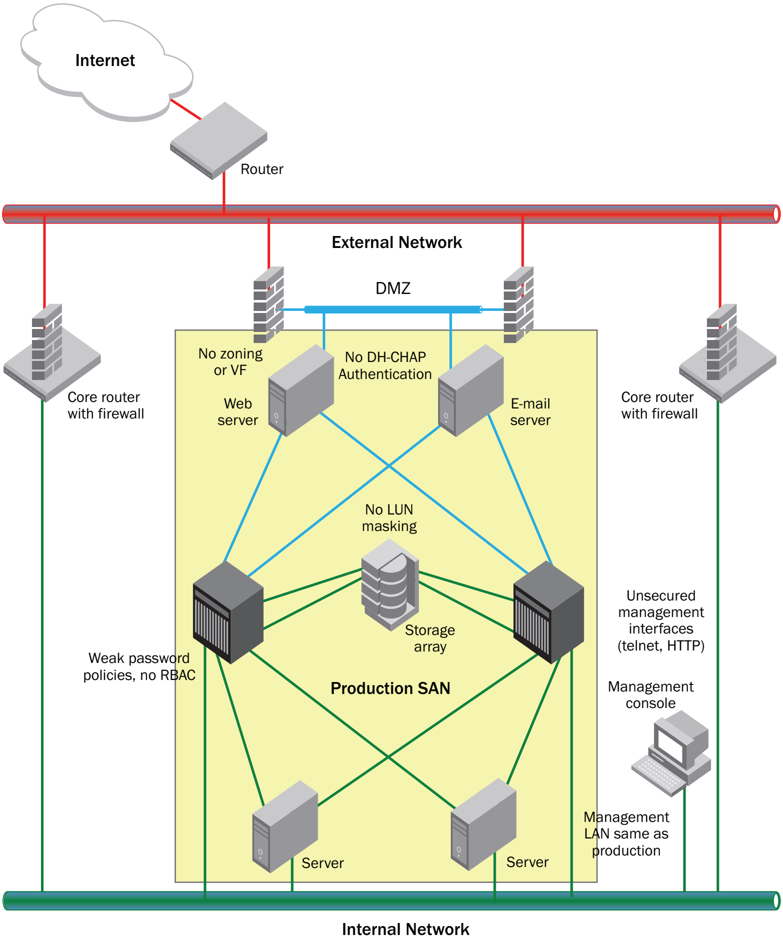
Internet
Router
External Network
DMZ
No zoning or VF
No DH-CHAP Authentication
Web server
E-mail server
Core router with firewall
Core router with firewall
No LUN masking
Storage array
Unsecured management interfaces (telnet, HTTP)
Weak password policies, no RBAC
Production SAN
Management console
Management LAN same as production
Server
Server
Internal Network

Figure 38. Poor example of securing SAN-attached DMZ servers

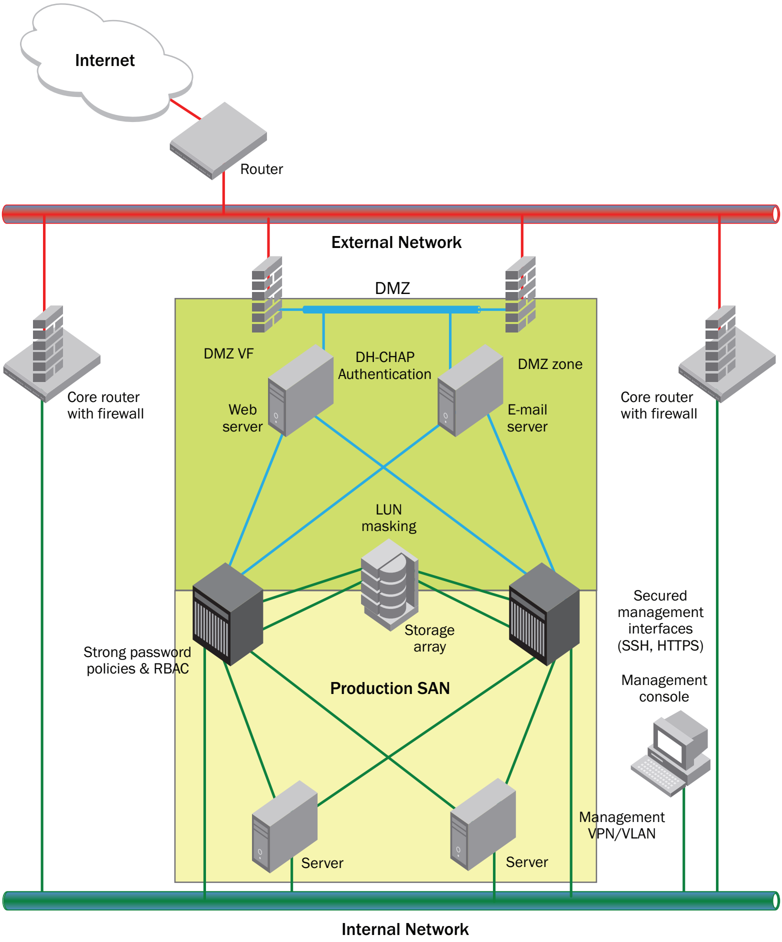

Figure 39. Securing SAN-attached DMZ servers connected to a production SAN

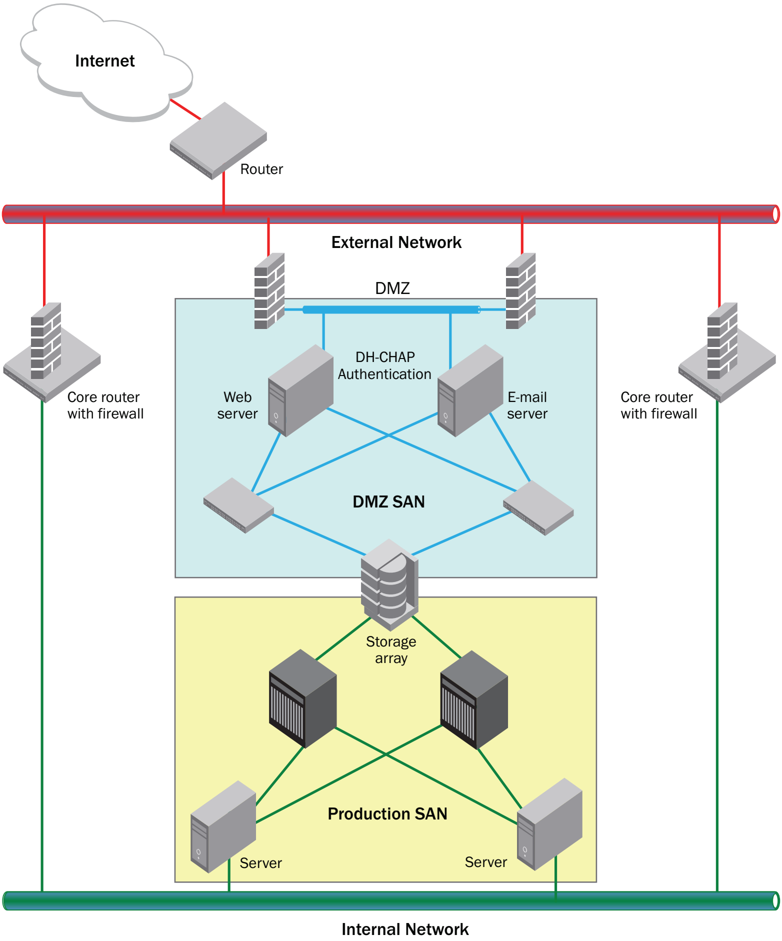

Figure 40. Securing SAN-attached DMZ servers using a physically separate SAN

Auditing and Assessing the SAN

Network security has become more public and is at the top of the priority list for many IT managers. Most companies perform security audits of their networks on a regular basis but SANs are often overlooked. It is important to audit the SAN as well to ensure the protection of centralized data. SAN security audits should go beyond the technological components and include a review of SAN security policies and physical security.

Brocade SAN Health Pro is a free SAN analysis tool that helps document and analyze a SAN to document its topology and uncover misconfigurations. For more information visit: http://www.brocade.com/services-support/drivers-downloads/san-health-diagnostics/index.page.

Best Practices Summary for DMZ

- Use a separate network, subnet, VLAN, or VPN for management interfaces
- Use secure protocols to communicate using management interfaces and disable unused protocols
- Use strong password management policies and implement RADIUS or LDAP
- Persistently disable unused ports, disable E_port connectivity on all unused and node ports, and implement port ACLs
- Use hardware-enforced pWWN zoning
- Use LUN masking
- For more sensitive environments, use DH-CHAP to authenticate servers
- Perform a SAN security assessment to understand the current state of security of the SAN

Chapter Summary

As with any networks, SANs have security vulnerabilities, but with proper design and configuration a SAN can be extremely secure. In order to safely connect servers in a DMZ to a SAN, several security precautions should be taken. The techniques described in this chapter, in combination, can provide a high level of protection to the production SAN in the event that a server in the DMZ ever becomes compromised.

The recommendations provided in this chapter are not meant to be exhaustive. It is possible to further harden a SAN if the reader chooses to do so. In this case, it is recommended that you read the other chapters of this book for further details on hardening a SAN.

8 Securing FOS-Based Fabrics

When it comes to SAN security, each organization has its own requirements and level of tolerance to risk. Although Brocade FC switches can be secured to a very high degree, organizations usually use a more pragmatic approach by finding the right balance between ease of SAN management and minimizing security risks. A FOS-based fabric provides more than 100 security features. Clearly, not all features need to be implemented in a fabric, but you have the flexibility to implement those that are necessary to achieve the precise level of protection that your organization requires.

To protect FOS-based SANs, use the defense-in-depth strategy described earlier in “The Brocade SAN Security Model” on page 91. This chapter covers most of the security features available in Fabric OS (FOS) 7.0 and earlier. It is not intended to be an implementation manual; but it does cover features and their associated commands at a high level. For further details on implementing these features, please refer to the appropriate version of the *Brocade Fabric OS Administrator's Guide* and the *Brocade Fabric OS Command Reference*.

NOTE: The FOS CLI commands in this chapter are written with uppercase letters to make them easier to read. However, when commands are executed, they are case insensitive and are normally typed as all lowercase, as shown in code examples. Additionally, bold text is used for commands in order to highlight them.

Secure Fabric OS: A Historical Overview

In the early 2000s, when the FC protocol was gaining momentum in enterprise environments while more mission-critical and sensitive information was being stored on the SAN, customers began expressing concern about the security of this new environment. Security professionals realized that the SAN was just another type of network and also required specific measures to protect the information residing on it.

To address these new enterprise security requirements, Brocade introduced the first security features for the FC SAN environment with the introduction of Secure Fabric OS (SFOS) in 2002 (released with FOS 2.6.0). SFOS introduced the first ACLs FC switch authentication mechanism (using PKI and digital certificates), and secure IP protocols to access the management interface. Although PKI can still be used, it has now been replaced with the industry standard DH-CHAP authentication method using shared secrets.

Most of the security features previously available in SFOS have since been replaced with equivalent or more powerful and flexible functionality in the base Fabric OS (version 5.3.0 and later). Appendix A provides a comprehensive list of technical security features that can be implemented in an FOS-based SAN environment. Brocade is continually enhancing existing features and creating new ones to ensure that FC fabric infrastructures and the data moving through them remain secure and highly available as new security vulnerabilities are discovered.

It is important to note that even though the ACLs in SFOS and the new base FOS equivalents share the same names, they are not compatible. A secure fabric running SFOS must be converted to the equivalent FOS-based security features according to the procedures detailed in the *Fabric OS Administrator's Guide*.

With SFOS, all switches in a secured fabric were required to be in secure mode in order to join and participate in the fabric. With the standard FOS-based secure mode, fabrics can be in either strict or tolerant mode:

- In **strict mode**, all switches must participate in the fabric, as was the case with SFOS.
- In **tolerant mode**, not all switches need to participate, which is particularly useful when a fabric contains older switches that cannot be upgraded to a firmware release greater than FOS 5.2.0. However, this mode is not recommended from a security perspective, as a fabric will only be as secure as its weakest link. A switch that does not participate in a secure fabric will become the vulnerable point that could be used to gain access to the fabric.

The following is a list of the Secure Fabric OS features:

- FCS (Fabric Configuration Server) policy
- SCC (Switch Connection Control) policy
- DCC (Device Connection Control) policy
- MAC (Management Access Control) policy
- PKI Switch-Switch Authentication

Securing Management Interfaces

As mentioned throughout this book, management interfaces are probably the most vulnerable points in a SAN from a security perspective. The physical interfaces on an FC switch include management Ethernet, USB, and serial ports. With the older Silkworm 2800 switch, there was also a front LCD panel that could be used to manage the switch.

The USB port is used exclusively for downloading firmware to a switch, and as such, is not used for other management purposes. The USB is often considered a security risk and it can be disabled from the active CP using the **usbStorage -d** command. The USB port is disabled by default and must be explicitly enabled each time a USB storage device is inserted in the port using the **usbStorage -e** command. A USB storage device must also be present to enable the port.

The serial interface can be used only to access the CLI and cannot be disabled. (SFOS had a MAC policy to disable access to the serial port but no such policy exists in base FOS today.)

With an Ethernet interface, the following tools can be used to manage a Brocade FC switch:

- **CLI**: telnet or secure shell (SSH)
- **GUI**: Brocade Web Tools or Brocade DCFM
- **FTP** (File Transfer Protocol) or **SCP** (Secure Copy Protocol)
- **SNMP** (Simple Network Management Protocol)

To protect the Ethernet interface, organizations should employ reliable IP network security best practices to isolate management interfaces and ensure that they are accessible only to the appropriate staff. Typically, the Ethernet interface is connected to a dedicated LAN or a VLAN used exclusively for management purposes and is not connected to the production LAN, which provides proper isolation between the two LANs.

Most organizations display a standard welcome message or banner at system login. Although this type of login banner might not be a major deterrent, it can help minimize liability and provide legal support in the event of a security breach. It should be a standard feature of any IT security strategy.

SAN and security administrators have several tools at their disposal to tighten the security around these management interfaces.

Encrypting Management Communications

Communication between an FC switch and its management server is particularly vulnerable to several types of attacks. Secure protocols must be used to prevent someone from sniffing or capturing cleartext information exchanged between the switch and management server.

To encrypt communications when the CLI is used, Brocade FC switches support SSH from FOS 4.1.1. All that is required is an SSH client on a management station to access the CLI; there is no special configuration necessary on the FC switch to implement SSH. Since SSH is the secure equivalent of telnet, the unsecure telnet service should no longer be available to users. This can be accomplished by disabling the telnet service for switches running pre-FOS 5.3.0 or by denying port 23 (telnet) using the IP filter feature on switches running FOS 5.3.0 or later.

One of the problems frequently encountered with SSH occurs when a SAN administrator creates a script and wants to use SSH to secure the communication when executing the script. Prior to FOS 6.1.0, a password had to be hard-coded into the script, which could be considered a security violation, since the password was written in cleartext format in the script. As of FOS 6.1.0, SSH public key authentication was introduced to allow for password-less authentication, specifically for authenticating a user within a script. The procedure for implementing this feature is well documented in the *Fabric OS Administrator's Guide.* It consists of creating a public/private key pair, which is exchanged between the switch and an authorized user.

Encrypting communications between the GUI interfaces such as Brocade Web Tools, DCFM, and Brocade Network Advisor is achieved using the HTTPS service, which is based on SSL/TLS (introduced in FOS 4.4.0). Enabling HTTPS requires a certificate that must be obtained from a third party prior to configuration and installed on each switch using the **secCertUtil** command.

The traffic exchanged on the Ethernet management interfaces between two FC switches or enterprise-class platforms (directors or backbones) can also be encrypted by creating a tunnel using IPSec (configured with the **ipSecConfig** command).

Finally, several FOS commands are used to copy information to and from the switches that use unsecure services such as FTP to exchange data in cleartext format. The FTP service can be replaced in some cases with the SCP service, which is based on SSH. The following commands can be secured using SCP instead of FTP:

- **configUpload** and **configDownload** (since FOS 4.4.0)
- **firmwareDownload** (since FOS 5.3.0)
- **supportSave** (since FOS 5.3.0)

In order to use SCP instead of FTP for the configuration upload and download operations, the **configure** command must be used. To use SCP with the **firmwareDownload** and **supportSave** commands, a parameter must be entered at the command line to indicate that the SCP protocol will be used.

SNMP is a commonly used protocol to monitor and manage Brocade switches. The earlier versions of this protocol, SNMPv1/2, had some security vulnerabilities that could be exploited. It is safer to use the latest version, SNMPv3, since it supports encrypted community strings along with several other capabilities.

If you are using SNMPv1/2, make sure to change the default community strings predefined in FOS, since they are well known and can be used in an attack. Additionally, the security level can be changed and set to:

- No security
- Authentication only
- Authentication and privacy

These settings can be configured using the **snmpConfig** command.

Protecting Login Sessions

One of the first things a hacker may do when planning an attack is to collect information about the targeted device. One of the simplest methods to collect SAN information is to simply type the IP address of a switch in a Web browser using HTTP to immediately display the Web Tools interface. The Web Tools interface would conveniently display a fair amount of useful information that could be used for an eventual attack. A feature called upfront login was introduced in FOS 5.0.1, forcing users to provide a username and password (in addition to switch information) to view the Brocade Web Tools GUI, a feature that has been enabled by default since FOS 5.3.0.

Once a user logs into an FC switch with telnet or SSH to use the CLI, a login banner or message can be displayed. By default, the login banner is not set and should be customized. The login banner can be enabled and created using the banner command. The motd command can be used to set a message-of-the-day. A banner is displayed after the user has logged into the switch and the motd is displayed prior to logging in.

When you create a login banner or motd, use these guidelines:

- Include language indicating that the user is logging in to a private network and unauthorized users will be prosecuted.
- Include language indicating that any user accessing this interface is consenting to be monitored. This is to address privacy issues.
- Do not provide more information than is necessary (do not include organization name, type of OS, and so on).

There is legal precedence of a successful defense from hackers claiming they were not informed they were not authorized to log into a network, which is why this must be stated clearly in the motd before they log in to the targeted system. In other cases, authorized users logged into their employer's network and performed illegal activities, getting caught via network monitoring. They used the defense that their right to privacy had been violated, since they were not informed that they could be monitored. For this reason, it is important to explicitly require consent in the login banner or motd. As mentioned earlier, the first phase in an attack is to collect information. Do not provide unnecessary information, or information that could identify the organization or specifics about the hardware they are logging into, in the login banner or motd. For the most part, laws in the US have adapted to the technology and legal provisions are already in place to address such issues. Nevertheless, it remains good practice to include explicit language, providing additional protection and demonstrating due diligence.

System and SAN administrators sometimes have a tendency to log into a switch and forget to log out when they are no longer using the interface. Enable a telnet and Web Tools session timeout feature using the timeout command (set to 15 minutes by default). The Web Tools session timeout, available since FOS 6.2, can be set from the Web Tools interface.

Filtering IP Traffic

The concept of a "firewall" has existed for quite some time in the conventional LAN world but is a relatively recent feature in FC-based SANs. The IP filter (IPF) feature, introduced in FOS 5.3.0, behaves as a firewall and replaces the MAC policy found in Secure Fabric OS. Using an IPF, a TCP/IP port can be either allowed or denied and a SAN administrator can define a specific IP address or range of IP addresses that are allowed to access a specific TCP/IP port.

There are two IP filter policy types: one for IPv4 and one for IPv6. Table 10 identifies a few well-known ports used with Brocade switches that can be controlled using IPF.

Table 10. Well-known ports and services

Service Name	Well-Known Port Number
FTP	20, 21
SSH	22
SCP (uses SSH)	22
telnet	23
HTTP	80
SNMP	161, 162
HTTPS	443
SYSLOG	514

The IPF is often used to deny the use of an unsecure service, such as telnet, when its equivalent secure version, such as SSH, must be utilized.

Password and User Management

As explained previously in "Chapter 6: FC Security Best Practices" starting on page 91, it is important to be able to associate each user with legitimate access to the SAN by their unique user name. To accomplish this, SAN administrators can create up to 255 customized user accounts. Each account can also have specific roles defined using the RBAC features. Doing so not only improves security but also improves troubleshooting and change tracking, while still clearly defining each administrator's appropriate role and authorization rights.

Brocade switches are factory installed with four user accounts:

- root
- factory
- admin
- user

For the reasons described above, these accounts should never be used except as a last recourse when a SAN administrator's password is lost, for example, and there is no other way to gain access to the switch. Each of these user accounts is also assigned a default password by Brocade. Some OEM partners change the Brocade default password to a different default password. As a best practice, change these default passwords at first login.

The password database is local to each switch. However, as of FOS 5.2.0, it is possible to manually distribute the local password database to other switches in a fabric using the **distribute** command (`distribute -p PWD -d switch_list`). If for some reason you want to exclude one or more switches from this distribution, use the **fddCfg** command entered from the switch to be excluded (`fddCfg --localreject PWD`).

Password Policies

A password policy can be created to ensure that users create strong passwords and follow the organization's password policy. There are four Brocade password policies that can be configured:

- Password strength
- Password history
- Password expiration
- Account lockout

Implement password policies on Brocade switches using the **passwdCfg -set** command.

Password strength refers to how difficult it is for someone else to guess or break a user's password. A hacker often tries to guess a password using real words such as a person's name, spouse's name, pet's name, and so on. Real words should never be used as part of a password since they are too easy to guess. The use of numbers, special characters, and cases all contribute to making a password stronger.

The following lists the different Brocade password strength features:

- **Lowercase.** Minimum number of lowercase characters to use (default = 0).
- **Uppercase.** Minimum number of uppercase characters to use (default = 0).
- **Digits.** Minimum number of numeric characters to be used (default = 0).
- **Punctuation.** Minimum number of punctuation characters to be used (default = 0).
- **MinLength.** Minimum number of characters required (8–40 characters).
- **Repeat.** Maximum length of repeated character sequences that is disallowed (1–40 characters; default = 1).
- **Sequence.** Maximum length of ASCII character sequences that increase by one. For example, “ABCDE” is a 5-character sequence increasing by one (1–40 characters; default = 1) as well as “45678”.

Example:

```
switch:admin> passwdcfg --set -uppercase 3 -lowercase 4
-digits 2 -minlength 9
```

This example sets a password strength policy that required at least 3 uppercase letters, 4 lowercase letters, 2 digits, and an overall minimum length of 9 characters.

Password history prevents users from using passwords that they used previously for a pre-defined number of passwords:

History. Number of previous password values (including the current value) that are disallowed when creating a new password (1–24; default = 1).

Example:

```
switch:admin> passwdcfg --set -history 10
```

This example sets a history policy that prevents the use of any of a user's previous 10 passwords.

Password expiration or aging is used to control how long a password can exist. The following lists the different Brocade password expiration parameters:

MinPasswordAge. The minimum number of days that must elapse before a user can change a password (0–999 days; default = 0). Setting this parameter to a non-zero value discourages users from rapidly changing a password in order to circumvent the password history setting to select a recently-used password.

MaxPasswordAge. The maximum number of days that can elapse before a password must be changed, (0–999 days; default = 0).

Warning. The number of days prior to password expiration that a warning about password expiration is displayed. (0–999 days; default = 0).

Example:

```
switch:admin> passwdcfg --set -minpasswordage 7 -
maxpasswordage 180 -warning 14
```

This example sets a password expiration policy that specifies that users cannot change a password for 7 days after they set a password and must change their password after 180 days (a warning is sent to them 14 days before their password is about to expire).

Password lockout is used to disable an account after a series of unsuccessful login attempts to prevent unauthorized users from entering consecutive password guesses until they guess the right one. The following lists the Brocade password lockout parameters:

- **LockoutThreshold.** The number of times a user can attempt to log in using an incorrect password before locking out the account (0–999; default = 0). Setting the lockout threshold to 0 (“zero”) disables the lockout policy.
- **LockoutDuration.** The time in minutes after which a previously locked account is automatically unlocked (0–99999 minutes; default = 30). Setting the lockout duration to 0 (“zero”) requires administrative action to unlock the account.

Example:

```
switch:admin> passwdcfg --set -lockoutthreshold 5 -
lockoutduration 0
```

This example configures a password lockout policy that gives a user 5 tries to enter the correct password and specified that once an account is locked, it can only be unlocked by an administrator.

The lockout policy can be used as a denial-of-service (DoS) attack when an attacker guesses a user password until the switch locks out the account. Once the account is locked, then the authorized user is no longer able to access his account. The admin account is particularly vulnerable to this type of attack and thus has a special policy. The admin lockout policy can be disabled to prevent a DoS attack on that account; however, it is then vulnerable to a brute-force guessing attack. The admin account lockout policy is enabled or disabled using the **passwdCfg** command (`passwdCfg [- - enableadminlockout] [- - disableadminlockout]`).

When a switch authenticates a user, by default it consults the local password database. However, the Brocade user authentication model allows for two other methods to authenticate users: RADIUS and LDAP.

SAN administrators can manage both passwords and usernames on each switch locally or through a centralized access control administration method, such as the RADIUS authentication protocol or the LDAP. These protocols allow a SAN administrator to change a password or disable a user's account from one central location and that change is applied immediately across all switches to which the user has access.

The authentication method to be used is defined using the **aaaConfig** command (`aaaConfig - - authspec ["radius" | "ldap" | "radius;local" | "ldap;local" - - backup]).`

For redundancy, more than one authentication server can be added using the **aaaConfig - - add** command.

Role-Based Access Control

RBAC can be used to restrict which commands a user can use. For example, a SAN administrator may want to allow summer interns to get their feet wet in SAN management by viewing and monitoring the SAN configuration and status, but does not want to them to be able to change any configuration parameters. A user account can be created with the User role to allow view but not modify permission. Table 11 lists the roles available in Fabric OS and when these roles became available. As of FOS 7.0, users can create their own customized roles with the **roleConfig** command.

Table 11. Brocade RBAC

Role Name	First in FOS	Duties	Description
Admin	All	All administration	All administrative commands excluding chassis-specific commands
BasicSwitch Admin	5.2.0	Restricted switch administration	Mostly monitoring with limited switch (local) commands

Role Name	First in FOS	Duties	Description
Chassis-role permission	6.2.0	Chassis-specific configuration	Role permission only and applied to the user account through the **userConfig** command
FabricAdmin	5.2.0	Fabric and switch administration	All switch and fabric commands, excludes user management and AD commands
Operator	5.2.0	General switch administration	Routine switch maintenance commands
SecurityAdmin	5.3.0	Security administration	All switch security and user management functions
SwitchAdmin	5.0.0	Local switch administration	Most switch (local) commands, excluding security, user management, and zoning commands
User	All	Monitoring only	Non-administrative use such as monitoring system activity
ZoneAdmin	5.2.0	Zone administration	Zone management commands only

Other Password-Related Features

It is possible to bypass the normal login procedure to recover a password by bringing the switch into single-user mode and obtaining special password recovery code from Brocade. This may be viewed as a security hole in some environments. To prevent unauthorized users from entering a switch into single-user mode, a password can be set on the boot PROM. A recovery string can also be defined in case the boot PROM password is lost, to allow Brocade to recover the password.

WARNING: If the boot PROM password is set and forgotten and there is no recovery string defined (or it is also forgotten), then there is no way of regaining management access to the switch if the admin or root passwords are lost.

FC-Specific Security

Brocade has developed several FC-specific security features that would not normally be available in a conventional LAN. For example, devices connecting to a Fibre Channel fabric can be authenticated using a strong protocol with the DCC policy.

FC Port Access Management

The FC ports on a switch are particularly vulnerable for several reasons. They can be used to introduce unauthorized devices into the fabric, such as another FC switch. They can also be used to connect an authorized device prematurely, for example, before an HBA has been configured, which may cause unexpected switch behavior.

The simplest method of protecting unused FC ports is to disable them. Use the **portDisable** command, but note that port status changes do not survive reboots. Changes using the **persistentPortDisable** command, on the other hand, persist and survive reboots.

An additional layer of defense that can be used to prevent unauthorized switches from joining a fabric is to disable the ability of an FC port to become an E_Port using the **portCfgPort** command.

Single Point of Management Access

Managing an FOS-based fabric by default can be performed from any switch. However, it is always simpler to secure one entry point than to secure multiple entry points, and this rule applies to FC fabrics as well. In large fabrics made up of numerous FC switches, there are many possible management points that all need to be secured properly. To create a single point of control for fabric management, Brocade introduced the FCS policy in FOS 5.3.0. The FCS policy identifies one switch as the primary point of control (the fabric configuration server) to manage all switches in the fabric. Administrators must perform changes to zoning, user accounts, passwords, or policies via the primary FCS, thereby reducing the number of possible entry points for a potential attacker.

The FCS policy can be defined using the **secPolicyCreate** command (`secPolicyCreate "FCS_POLICY", "member ;…;"member"`), where the “member” is the switch domain ID.

Example:

```
switch:admin> secpolicycreate "FCS_POLICY", "2;4"
FCS_POLICY has been created
```

Switch and Device Access Controls

Brocade created a set of ACLs to prevent unauthorized access of switches and devices in a fabric in the form of the SCC and DCC policies.

In a FOS 4.4.0 environment or later, use the SCC policy to define which switches are allowed to participate in a fabric. The switches are defined as members of the SCC using their WWN. The SCC policy can be defined using the **secPolicyCreate** command (`secPolicyCreate  "SCC_POLICY", "member ;…;"member")`, where the "member" is the switch domain ID and an asterisk (*) is used to define all switches in a fabric.

Example:

```
switch:admin> secpolicycreate "SCC_POLICY", "2;4"
```

In a FOS 5.3.0 environment or later, use the DCC policy to define which devices are allowed to join a fabric. The DCC policy can identify member devices using their WWN or the physical port in the fabric to which they are connected. To further enhance security, a WWN can be locked down to a specific port (as a WWN spoofing countermeasure) by preventing a device that is configured to mimic an existing device from joining a fabric, unless the device being spoofed is first disconnected and then physically replaced with an unauthorized device.

The SCC policy is defined using the **secPolicyCreate** command (`secPolicyCreate "DCC_POLICY_policyname", "member ;…;"member"`), where the "member" is either a WWN or the switch domain ID (portID). When both the WWN and the switch ID/port ID definitions are used together, this is called "locking down a port" and only the WWNs associated with that port are allowed to join the fabric.

Example:

```
Switch:admin>   secpolicycreate   "DCC_POLICY_server",
"11:22:33:44:55:66:77:aa;1(3)"
```

This example creates a policy called DCC_POLICY_server and locks down the device with WWN 11:22:33:44:55:66:77:aa to port 3 of the switch with domain ID 1.

Switch and Device Authentication

ACLs such as the DCC and SCC policies provide an identification method for devices joining a fabric. Since a WWN can be spoofed, some organizations require more than simple identification and require that devices authenticate to prove they really are what they "say" they are. Authentication in an FC fabric can be accomplished using different protocols such as SLAP, FCAP, and DH-CHAP. Some of these protocols are based on the use of digital certificates and others use shared secrets.

Brocade-supported SLAP (Switch Link Authentication Protocol) is based on digital certificates in SFOS. Today, SLAP is no longer supported on FC switches. FCAP (Fibre Channel Authentication Protocol), based on digital certificates, and DH-CHAP, based on exchange of shared secrets, are the principle authentication protocols used in FC. DH-CHAP is more frequently used, since it is part of the FC-SP standard and does not require obtaining third-party digital certificates.

Brocade introduced the AUTH policy in FOS 5.3.0 to allow SAN administrators to enforce device authentication. The AUTH policy can be set to either of the following:

- OFF: No authentication required (default)
- ON: Strict enforcement of authentication on devices joining F_Ports
- PASSIVE: Authentication is optional and only authenticates devices configured for and capable of authentication

The ON mode of the AUTH policy was introduced recently in FOS 5.2.0. Prior to this, device authentication could not be configured to require authentication.

Isolation and Separation

Some environments or devices require special protection from other environments or devices. SAN administrators may want to prevent a sensitive system from being accessed by the general production environment, for example. Perhaps a test environment needs to be isolated from the production environment, to prevent changes in the test environment from affecting the production systems. Environments and devices can be separated from each other in an FOS environment either physically or logically as follows:

- Physically
  - Physically isolate critical or sensitive systems where appropriate using separate fabrics
  - FC routing can provide isolation and controlled sharing
- Logically
  - Zoning (hardware-enforced pWWN)
  - Virtual Fabrics/Administrative Domains
  - Traffic isolation zones

FC Routing

Fibre Channel Routing (FCR) is a means of isolating two fabrics from each other, while allowing specific devices in separate fabrics to communicate with each other according to a set of pre-defined rules. FCR can be implemented in one of two ways in an FOS-based fabric:

- Brocade 7800 Extension Switch or FX8-24 Extension Blade
- Integrated Routing (IR) feature, available in FOS 6.2.0 and later

The Brocade 7800 and FX8-24 are specialized routing hardware platforms; IR is a licensed feature available on standard Condor 2-based products (“Condor 2" identifies the ASIC type), which include the Brocade DCX/DCX-4S Backbone and Brocade 5100/5300 Switch. With the IR feature, a specific port in a supported switch can be configured to perform FC-FC routing.

Zoning

Zoning provides a logical means to group devices together and to isolate them from other devices. Zoning has been discussed at length in “Chapter 3: SAN Basics for Security Professionals” starting on page 19, as well as “Chapter 6: FC Security Best Practices” starting on page 91. This section discusses zoning in greater detail and how it is implemented and managed in an FOS environment.

As a best practice, it is preferable to implement zones on FOS-based fabrics using the pWWN instead of the domain ID/port ID, since both are hardware-enforced and the pWWN provides more flexibility from a management perspective. However, do not use a combination of the two (mixed zone) within the same zone, as this will result in zone enforcement by the name server, which is less secure.

A set of zones make up a zone configuration, and it is possible to have more than one zone configuration in a fabric. For example, there could be one zone configuration for the day shift, during which most production takes place, and another for the night shift, during which maintenance and backups are usually performed. When a configuration is changed, the effective configuration is disabled and the new configuration is enabled and then becomes the effective configuration. During this transition period, particularly with large fabrics, the name server must indicate to all the servers that there is a change in the devices with which they are allowed to communicate. During this transition, when the effective configuration is temporarily disabled, it is possible for all servers in the fabric to see all devices, since no zone configuration is effectively defined.

To prevent this from happening, default zones were created to ensure that all devices in a fabric cannot see each other during a configuration change. The default zone can be set to NOACCESS mode to prevent devices from seeing each other using the **defzone - - noaccess** command.

Virtual Fabrics and Administrative Domains

Organizations can also employ Brocade Administrative Domains (AD), introduced in FOS 5.2.0, so administrators have access only to the groups of SAN ports, WWNs, and switches required by their job function. Organizations can use ADs and RBACs together to limit an administrator to only the areas of the SAN and the amount of control required to perform their duties. Providing full administrative authority and a complete view of the SAN for administrators who do not need that level of access exposes the organization to accidental or malicious attacks, which can result in downtime or data loss. Brocade switches support up to 256 ADs.

The Virtual Fabrics (VF) feature was introduced in FOS 6.2. VF provides two capabilities: Logical Switches and Logical Fabrics. A physical switch can be partitioned into multiple Logical Switches that are managed and behave like a physical switch. Each Logical Switch is associated with a Logical Fabric. A Logical Fabric is a fabric that contains at least one Logical Switch. Logical Fabrics can include physical switches, support single fabric and shared multiple fabric ISL connections, and IFL connections for FC-FC routing to edge fabrics. VF provides full data, control and management isolation.

Traffic Isolation Zones

Traffic isolation zones were introduced in FOS 6.0.0 to address the problem of shared bandwidth between devices over the same ISL. This problem was particularly apparent when different I/O-intensive applications competed for available bandwidth over a dark fiber connection between sites. For example, data replication between two sites could be competing with a backup application for bandwidth over a pair of dark fibers between the primary site and the DR site. The data replication application can be configured in synchronous mode and is directly related to the performance of the production environment. The backup environment is less critical, since it does not have a direct effect on the production environment.

In this case, it would be preferable to give the data replication traffic priority over the backup traffic, or at least isolate these two applications from each other and assign all of the backup traffic to one ISL and the data replication traffic to a different ISL. This was not possible, since the FSPF routing protocol could not distinguish between the two types of traffic and simply shared the load between the two available ISLs. Traffic Isolation zones were created to address this issue. Traffic isolation can force traffic from one source to be sent over one path and traffic from a different source to another path. In the previous example, the backup traffic could be sent over one path and the data replication traffic over another path.

TI zones can be created using the **zone** command (`zone - - create -t ti zone_name -p "ports"`).

Example:

```
zone --create -t ti red_zone -p "1,1; 2,4; 1,8; 2,6"
```

Logging and Change Management

The primary logging mechanism on Brocade switches is the syslog (system log). The first rule when using logs is to ensure that the clocks on all switches in the fabric are synchronized so that log files have consistent time stamps across the SAN. This is easily accomplished with NTP (Network Time Protocol). An NTP server is used as the primary source of accurate time for the entire SAN. NTP is defined by providing the IP address of an NTP server and for redundancy, more than one NTP server can be specified (up to eight servers). NTP servers are defined using the **tscLockServer** command.

As mentioned in previous chapters, one of the first things a sophisticated hacker may do is try to remove all traces of his or her activity on a system once an attack is completed. The syslog is one file that is typically removed in this process, so it is important to redirect the syslog to a secure server in a different location from the actual switches using the **syslogDIpAdd** command.

As an extra precaution, log files redirected to a secure server should also be backed up regularly. Furthermore, backups for this server should probably be retained for a longer period of time than most other backups. It would be preferable to retain all backups in the event of a security incident that is only detected several months after it occurred. The backup of the log files could be the only way to obtain proof of the incident, if required at trial.

Audit log

Certain classes of events that occur in a SAN may be of great interest to security professionals. These events include login failures, zone configuration changes, firmware downloads, and other configuration changes, all of which may have a serious effect on the operation and security of the switch. These events can be recorded and filtered using the Brocade audit log feature, introduced in FOS 5.2.0. Auditable events using this feature are generated by the switch then sent to an external host through syslogd (the daemon that sends messages to the syslog).

Track Changes Feature

From a security perspective, it may also be important to keep a record of specific changes that cannot be considered switch events but that can provide useful information, such as unsuccessful login attempts. The track changes feature introduced in FOS 4.0.0 tracks these changes and logs them into the syslog. The following list identifies the changes tracked by this feature:

- Successful login
- Unsuccessful login
- Logout
- Configuration file change from task
- Track changes on
- Track changes off

Fabric-Based Encryption

Encryption ensures confidentiality of data, whether it is at rest or in flight. Encryption of data-at-rest in an FOS environment can be performed at the fabric level using the Brocade Encryption Solution. This solution is discussed in greater detail in “Chapter 11: Brocade Data Encryption Products” starting on page 173.

Encrypting data-in-flight can be used to secure communications between two data centers connected through an FCIP tunnel, for example. This solution could be implemented in an FOS environment using the Brocade 7800 or FX8-24, also discussed at length in Chapter 11.

FIPS Mode

As discussed in “Chapter 9: Compliance and Storage” starting on page 155, FIPS 140-2 is a standard that was established to simplify the procurement of security products by providing a simple method to ensure that products meet certain security requirement levels. Brocade switches by default are not compliant with the FIPS standard, but they can be placed into FIPS mode to immediately enhance the security level of the switch. FIPS mode has been available since FOS 6.0.0.

The following is an important distinction: placing a switch in FIPS mode is not the same as making the switch FIPS-compliant. Placing a switch in FIPS mode enhances the security level of the switch according to the compliance requirements specified by FIPS 140-2 Level 2. Enabling FIPS mode is a disruptive action, since it requires a reboot of the switch to take effect.

FIPS mode is enabled and configured using the **fipsCfg** command.

CAUTION: FIPS mode is disruptive and may have unexpected implications if you are not familiar with this mode of operation. For example, if you lose the admin password on a switch running in FIPS mode, there will be no way to regain management control of that switch. FIPS mode should be used only if there is a mandatory operational requirement to do so. Again, operating a switch in FIP mode does not imply that the switch is FIPS 140-2 compliant.

When a Brocade switch is in FIPS mode, the following occur:

- Root account disabled
- Telnet disabled, only SSH can be used
- HTTP disabled, only HTTPS can be used
- RPC disabled, only secure RPC can be used
- Only TLS-AES128 cipher suite used with secure RPC
- SNMP read-only operations exclusively, SNMP write operations disabled
- DH-CHAP/FCAP hashing performed only using SHA-256
- Mandatory firmware signature validation
- SCP used exclusively (no FTP) for **configUpload**, **configDownload**, **supportSave**, and **firmwareDownload** commands
- IPSec usage of AES-XCBC, MD5, and DH group 1 blocked
- RADIUS uses only PEAP or MSCHAPv2, CHAP and PAP disallowed
- Only the following encryption algorithms functional: HMAC-SHA1, 3DES-CBC, AES128-CBC, AES192-CBC, and AES256-CBC

Starting in FOS 6.2.0, the following steps are required to prepare a switch to run in FIPS mode:

1. (Optional) Configure RADIUS or LDAP server.
2. (Optional) Configure authentication protocols.

3. (LDAP only) Install SSL certificate on a Microsoft Active Directory (AD) server and CA certificate on the switch for using LDAP authentication.
4. Block telnet, HTTP, and RPC (using IP filters).
5. Disable boot PROM access.
6. Configure the switch for signed firmware.
7. Disable root access.
8. Enable FIPS mode (using **fipsCfg** command).

Refer to the *Fabric OS Administrator's Guide* for the version of firmware you are using before performing the procedure to make sure that you have the most complete and current information. Once FIPS mode is enabled, then several other steps are required to reset and zeroize certain switch parameters.

Other FC Security Features

A few other security features are available in FOS that have not been covered in previous sections and that are worthy of mention.

RSCN Suppression

It was explained earlier that RSCNs are contained to the devices within a FOS-based zone. It is also possible to explicitly suppress RSCNs at the port level. Some specialized applications are very sensitive and can be affected by an RSCN. If the environment is static and never changes once it is installed, RSCNs can be disabled to prevent interruptions. RSCN suppression can be configured using the **portCfg rscnsupr** command.

Signed Firmware

Firmware can be tampered with and a modified version of the firmware installed on a switch. This type of attack, although unlikely on a Brocade switch, is usually performed by modifying the code to adding a “back door,” or malicious code known only by the author of the modified code. To ensure that the code being installed on a switch is in fact the authorized version and has not been modified by a third party, a hash value of the firmware is calculated. This hash value is then digitally signed with a private key at the source using the RSA algorithm and 1024-bit keys. The public key of the source is included in the firmware package to allow the switch to authenticate the firmware. This feature, called signed firmware, was introduced in FOS 6.1.0.

When installing new firmware on a switch that has been configured for firmware signature validation, the public key is retrieved from the local public key file included with the firmware package and the firmware is validated.

A switch must be configured to enforce firmware signature validation and this is done using the **configure** command.

Example:

```
switch:admin> configure
Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
Configure...
System services (yes, y, no, n): [no]
…
cfgload attributes (yes, y, no, n): [no] yes
Enforce secure config Upload/Download (yes, y, no, n):
[no]
Enforce firmware signature validation (yes, y, no, n):
[no] yes
```

Fabric Watch Security Class

Fabric Watch is a Brocade licensed feature in FOS that is used to monitor switch events and send alerts in the form of SNMP traps or e-mails. Fabric Watch events are grouped into classes, one of which is of particular interest to security personnel: the security class.

The Fabric Watch security class includes the following events:

- API violations
- Front panel violations
- Illegal commands
- Security database
- Invalid signatures
- Login violations
- No FCS violations
- SCC Violations
- SES Violations
- SLAP Failures
- TS out of sync
- DCC violations
- HTTP violations
- Incompatible
- Invalid certificates
- Invalid time stamps
- MS violations
- RSNMP violations
- Serial Violations
- SLAP Bad Packets
- Telnet Violations
- WSNMP Violations

Insistent Domain ID

It is possible for a switch to obtain a new domain ID after a reboot, particularly when a switch is added to a new fabric or after a massive power failure. To prevent this from occurring, it is a best practice to assign a domain ID to a switch using an insistent domain ID (IDID). Once it is set, a DID survives reboots or power failures and will never change.

The insistent domain ID is set using the **configure** command:

- Select `y` after the `Fabric Parameters` prompt
- Select `y` again after the `Insistent Domain ID Mode` prompt

Chapter Summary

With over 100 security features and more added in every Fabric OS release, there are many tools at the disposal of SAN and security professionals to increase the security level of their SAN environment. Most of these features are relatively simple to implement and do not add any overhead to the daily management tasks of the SAN administrator. Some features actually simplify management (RADIUS and LDAP), for example, by allowing a SAN administrator to change the password for a user in one convenient location as opposed to every switch in the SAN.

Deciding which FOS security features to implement depends on each individual organization's requirements, which includes factors such as:

- Specific vulnerabilities
- Probability of a vulnerability being exploited
- Value of the asset being protected
- Cost of implementing the countermeasures
- Impact on day-to-day management activities

Once these factors are weighed carefully, a SAN security policy can be created and implemented using appropriate countermeasures.

9

Compliance and Storage

Certainly, most organizations will demonstrate due diligence and implement security measures on their own to protect their sensitive and critical data from loss or theft. Nonetheless, one of the primary driving factors for organizations to implement specific security measures is compliance, particularly mandatory and regulatory compliance. Compliance is the state of being in accordance with established guidelines, specifications, or legislation. These guidelines, specifications, and legislation can be industry-specific, an accepted standard, or government legislation. Guidelines and specifications are not necessarily mandatory; some provide guidelines on which organizations can base their security policies to better protect their IT environments. Legislative specifications, however, are mandatory for certain organizations. Non-compliance is not an option and if prosecuted, organizations face severe penalties, including fines and jail time for executives in some cases.

Guidelines, specifications, and legislation are not generally aimed at one specific area of technology, such as SANs and storage, but usually apply to all technologies and systems. A holistic approach is the best strategy to meet most regulatory compliance requirements.

Payment Card Industry Data Security Standard (PCI-DSS)

The PCI-DSS standard has been one of the most significant drivers for adoption of encryption solutions for data-at-rest and data-in-flight to address the protection of information related to credit card transactions.

The PCI Security Standards Council was formed in December 2004 by its founding members:

- Visa Inc.
- Master Card Worldwide

- American Express
- Discover Financial Services
- JCB international

The Data Security Standard (DSS), first established in September 2006, defines requirements to help prevent credit card fraud and hacking into credit card management systems. Merchants are required to meet minimum security standards. The following describes the general requirement categories but there are many specific requirements within each category.

Build and maintain a secure network:

- Install and maintain a firewall configuration to protect cardholder data
- Do not use vendor-supplied defaults for system passwords and other security parameters

Protect cardholder data:

- Protect stored cardholder data
- Encrypt transmission of cardholder data across open, public networks

Maintain a vulnerability management program:

- Use and regularly update anti-virus software on all systems commonly affected by malware
- Develop and maintain secure systems and applications

Implement strong access control measures:

- Restrict access to cardholder data by business need-to-know
- Assign a unique ID to each person with computer access
- Restrict physical access to cardholder data

Regularly monitor and test networks:

- Track and monitor all access to network resources and cardholder data
- Regularly test security systems and processes

Maintain an information security policy:

- Maintain a policy that addresses information security

Sensitive cardholder data under the PCI-DSS is defined as:

- Primary Account Number (PAN)
- Cardholder name
- Service code
- Expiration date

PCI-DSS uses a multi-tiered approach to managing merchant risks that depends on several factors. Merchants fall into a specified merchant level based on the criteria identified in Table 12.

Table 12. PCI-DSS merchant levels and criteria

Merchant Level	Criteria
Level 1	• All merchants processing over 6 million transactions per year • Merchants whose data has been previously compromised • Any merchant deemed to meet Level 1 compliance
Level 2	• All merchants processing from 1 to 6 million transactions per year • All merchants required by another payment network to report compliance as a Level 2 merchant
Level 3	• All merchants processing from 20,000 to 1 million transactions per year • All merchants required by another payment network to report compliance as a Level 3 merchant
Level 4	• All other merchants

Level 1 merchants, due to the significant number of transactions they process, are required to have an annual onsite audit. All other merchants must complete an annual self-assessment questionnaire and all merchants, including Level 1, must undergo a quarterly network security scan performed by an approved scanning vendor (ASV).

PCI-DSS and Storage

Several requirements defined in the PCI-DSS affect the SAN and storage environments, specifically:

- Requirement 3.4.1 refers to the possible use of disk encryption to protect cardholder data.
- Requirement 3.5 and 3.6 refer to protecting the keys used to encrypt cardholder data.
- Requirement 4.1 addresses encryption of data-in-flight when transmitting sensitive information over open, public networks. Protocols such as SSL/TLS and IPSec are recommended.
- Several other requirements mandate the use of secure management interfaces, such as SSH and SSL.
- Other requirements define system security parameters, such as synchronizing system clocks (10.4).

Breach Disclosure Laws

The recent increase in news articles and public display of security breaches in the US and worldwide is largely attributable to recent laws forcing organizations to disclose security breaches or risk penalties. Several websites are dedicated to publishing these security breaches:

- Privacy Rights Clearinghouse:
  http://www.privacyrights.org/
- Open Security Foundation Data Loss DB:
  http://datalossdb.org/
- Office of Inadequate Security:
  http://www.databreaches.net/

Breach disclosure laws require organizations to disclose specific types of security breaches, particularly those involving personally identifiable information (PII) of individuals of a given state. There is no current federal legislation to address breach disclosure.

The precedent-setting law was the California Senate Bill (SB) 1386 which came into effect on July 1, 2003, as a result of a security breach of California's state website in 2002.

California SB 1386 states that:

1. Any agency that owns or licenses computerized data that includes personal information;
2. shall disclose any breach of the security of the system following discovery or notification of the breach in the security of the data;
3. to any resident of California whose unencrypted personal information was, or is reasonably believed to have been, acquired by an unauthorized person.

California SB 1386 was not perfect, so it was necessary to expand its scope to impose a general security standard on businesses that maintain certain types of personal information about California residents. California Assembly Bill (AB) 1950 came into effect in January 2005 and also requires businesses, and their subcontractors, to maintain "reasonable security procedures and practices."

At the time of writing, there are 46 US states that have implemented breach disclosure laws. The National Conference of State Legislatures (NCSL) website (http://www.ncsl.org/default.aspx?tabid=13489) contains a list of all US states with breach disclosure laws, along with references to them.

Other similar breach disclosure laws have been enacted in other countries including:

- **Canada.** Personal Information Protection and Electronic Documents Act (PIPEDA)
- **UK.** Data Protection Act (DPA)
- **EU.** EU Data Protection Directive (Directive 95/46/EC) and Basel II
- **Japan.** Personal Information Protection Law (PIPL)
- **Australia.** Commonwealth Privacy Act (CPA)

The following website provides a list of international privacy laws: http://www.informationshield.com/intprivacylaws.html.

Breach Disclosure Laws and Storage

One of the most common disclosures affecting the storage industry is the loss or theft of a backup tape. In many cases, a lost or stolen tape media that is encrypted would not require disclosure; and in others, a disclosure would still be required but it would be qualified with the fact that the data was encrypted and does not pose any risks of exposing PII. This is quite significant from a public relations perspective for an organization that has suffered such a breach.

There have also been reported cases of disk subsystems being sold on the open market with actual data still residing on the disk drives. Similarly, there have been cases of disks installed in a customer's environment that still contained data, although they were allegedly refurbished by vendors.

Health Insurance Portability and Accountability Act (HIPAA)

HIPAA was enacted by the US Congress in 1996 to help maintain confidentiality of healthcare transactions or electronic protected health information (EPHI). Title II of HIPAA, the Administrative Simplification (AS) provisions, requires the establishment of national standards for electronic healthcare transactions and national identifiers for providers, health insurance plans, and employers. The AS provisions also address the security and privacy of health data.

Offenses under HIPAA can have the following consequences:

- A fine of not more than $50,000, imprisonment of not more than 1 year, or both
- If the offense is committed under false pretenses, a fine of not more than $100,000, imprisonment of not more than 5 years, or both
- If the offense is committed with intent to sell, transfer, or use individually identifiable health information for commercial advantage, personal gain, or malicious harm, a fine of not more than $250,000, imprisonment of not more than 10 years, or both

A major criticism of HIPAA has been that, in spite of providing well-defined penalties, it has not really been heavily enforced; although there have been recent cases of healthcare institutions being audited by the US Health and Human Services. To address this issue, the US Government enacted the Health Information Technology for Economic and Clinical Health (HITECH) Act in 2009. The HITECH Act now provides specific penalties, both civil and criminal, to enforce HIPAA rules. This has resulted in a spike in health-related organizations adopting encryption solutions to comply with HIPAA.

HIPAA and Storage

When healthcare transactions flow over open networks, as would be the case if replicating data to a secondary data center using FCIP, they must be protected by some technical safeguard such as encryption.

The guidelines are not always clear but there is a reference to: “Implement a mechanism to encrypt EPHI whenever deemed appropriate.”

Gramm-Leach-Bliley Act (GLBA)

The Gramm-Leach-Bliley Act (GLBA) was enacted by the US Congress on November 12, 1999, to open up competition among banks, securities companies, and insurance companies. The Financial Privacy Rule governs the collection and disclosure of customers' personal financial information by financial institutions. The Safeguards Rule requires all financial institutions to design, implement, and maintain safeguards to protect customer information.

US financial institutions (including credit reporting agencies) are required to comply with GLBA. It is enforced by the Federal Trade Commission (FTC) and other government agencies.

Some of the penalties under GLBA include:

- “the financial institution shall be subject to a civil penalty of not more than $100,000 for each such violation”
- “the officers and directors of the financial institution shall be subject to, and shall be personally liable for, a civil penalty of not more than $10,000 for each such violation”

The key term is here is “personally liable”, which certainly gets the attention of the officers and directors of a financial institution. This law is enforced and has lots of bite to it, with several cases tried and currently on trial under this act.

GLBA and Storage

There is a provision in the GLBA to have “a policy in place to protect the information from foreseeable threats in security and data integrity”. An integral part of this policy is to encrypt sensitive financial information and transactions. There is also a requirement to put in place the major components of that which is to govern the collection, disclosure, and protection of consumers' nonpublic personal information or personally identifiable information.

Encryption of data-in-flight, data-at-rest, as well as other SAN and storage security countermeasures can provide the necessary components to protect consumers' nonpublic personal information or PII.

Sarbanes-Oxley Act (SOX)

The Sarbanes-Oxley Act was enacted on July 30, 2002, as a response to several corporate and accounting scandals that shook the business world at the turn of the century. SOX does not apply to privately-owned companies but to public company boards, management, and public accounting firms.

Section 404 treats IT controls that specifically address financial risks. Many companies use the COSO (Committee of Sponsoring Organizations of the Treadway Commission) framework and COBIT (Control Objectives for Information and Related Technologies) to support SOX section 404 compliance.

SOX and Storage

SOX has no direct implications in the storage environment other than general system security implications that apply to the storage equipment. In certain cases, there may be some requirements for a minimum retention period for backup data.

Export Laws for Cryptographic Products

Until recently, cryptographic algorithms and materials were considered to be munitions, and as such fell under specific export regulations as dictated by each country. Although cryptographic material is no longer considered munitions, it is still subject to export regulations in the US.

In the US, export of cryptographic material is controlled by the Department of Commerce Bureau of Industry and Security (BIS). Some countries, known as “rogue states”, are strictly forbidden to export cryptographic material. For the most part, laws around exporting cryptographic material outside of these countries have been relaxed, but there still are some restrictions. It is best to verify with the BIS before exporting any cryptographic material.

Other countries also have restrictions on exporting or importing cryptographic materials. For example, France, at the time of writing, has an import restriction on 128-bit keys, which are subject to special permission.

Federal Information Processing Standards (FIPS)

IT security product consumers may not necessarily have the expertise, knowledge, or resources necessary to fully evaluate products, that is, whether the security of a product is appropriate and meets their requirements. Assertions from the vendors and developers of these products may not provide the highest level of confidence to the consumer. To increase this level of confidence, a consumer can hire an independent organization to evaluate products for them or simply use a pre-established standard that vendors must comply with.

When US Federal and private sector organizations make purchasing decisions for security products that perform a cryptographic function, they must evaluate the proposed products from each vendor. This is sometimes accomplished by creating an evaluation matrix comparing the different product features. A compliant/non-compliant system may be used, while others may prefer a weighted point system to give more importance on some functionality over others. Since this matrix can become quite large and complex when multiple vendors respond to a tender, a standard was created to establish base security criteria levels for all vendors.

The National Institute of Science and Technology (NIST), reporting into the US Department of Commerce, created publication 140-2 on May 25, 2001 (also known as the Security Requirements for Cryptographic Modules) to simplify the acquisition process. FIPS 140-2 was developed primarily for US Federal organizations and provides standard evaluation criteria for cryptographic modules used in certain security products. It is sometimes used by private sector organizations in North America but seldom in other parts of the world. The FIPS 140-2 standard applies specifically to the cryptographic modules used in security products. A cryptographic module consists of the hardware, software, and/or firmware used to implement security functions (including encryption algorithms and key generation) and is contained within a cryptographic boundary that establishes its physical boundaries (see Figure 44 on page 177).

Each organization has different security requirements and requires different degrees of security, hence FIPS 140-2 defines four security levels (see below). The lowest security level begins at 1 and each subsequent level builds upon the previous ones.

The actual certification of the cryptographic module is performed by an independent lab, which validates the product to ensure it meets the criteria required for the Security Level being sought by the vendor. Once the tests are completed, the results are submitted to NIST and upon their approval the product is officially posted on the NIST website at http://csrc.nist.gov/groups/STM/cmvp/validation.html.

Security Level 1

Security Level 1 provides the lowest level of security and it basically defines production-grade equipment with no physical security. Pretty much any product using a cryptographic module would qualify for this level of certification. An example of a Security Level 1 certified product is an ordinary laptop with a software-based encryption module.

Security Level 2

Security Level 2 enhances Security Level 1 with the tamper evidence requirement. Tamper evidence is implemented using special coatings or seals or pick-resistant locks for removable covers and doors. If a protective cover or door is tampered with to allow physical access to critical security parameters or keys stored in the cryptographic module, the coatings or seals will be broken and permanently modified.

Additionally, role-based authentication must be used to authenticate an operator with a specific role that allows them to perform certain tasks, such as deleting keys.

Security Level 3

Security Level 3 builds upon Security Level 2 with the addition of tamper-resistant mechanisms to prevent someone from gaining access to the critical security parameters (CSP) stored in the cryptographic module. This may include tamper detection and response systems, which could, for example, zeroize the keys stored in the local cache when the cover or door is opened.

Security Level 3 must also include identity-based authentication mechanisms to authenticate a specific individual and verify that they are authorized to perform the specified task.

Security Level 3 also requires that plaintext CSPs be exchanged using different ports than those used for other purposes (such as management interfaces). This enforces the principle of separation of duties to allow different individuals to have authority over the different types of functions and prevents one individual from having total control over the entire process.

Security Level 4

Security Level 4 provides the highest level of security and builds upon Security Level 3. The physical security mechanisms at this level must provide a complete envelope of protection around the cryptographic module. All unauthorized attempts to physically access the cryptographic module must be detected and responded to by zeroizing all plaintext CSPs. The cryptographic module must also be protected against extremely vigorous environmental conditions that exceed the normal operating ranges for voltage and temperature.

Only the most demanding environments require products certified to Security Level 4, such as combat zones and highly secure facilities that use equipment containing highly sensitive information. Under these exacting conditions, the equipment must still be able to zeroize the CSPs. For this reason, some people refer to Security Level 4 as a “science experiment,” since the testing process is extremely rigorous, lengthy, and expensive and few products are certified to this level.

FIPS Process

Once a vendor applies to qualify under FIPS 140-2, there is a series of stages to go through. The vendor and product under evaluation are published on the NIST/NIAP website at: http://www.niap-ccevs.org/cc-scheme/vpl/.

There are five basic stages to get to final acceptance and qualification:

1. Implementation Under Test (IUT)
2. Review Pending
3. In Review
4. Coordination
5. Finalization

Common Criteria (CC)

Common Criteria (CC), like the FIPS 140-2 standard, were also developed to simplify the acquisition process of IT security products. It is a standard of evaluation of security properties of IT products and systems. As such, it addresses the three basic tenets of security: protecting the Confidentiality, Integrity, and Availability (CIA) of information. While FIPS 140-2 focuses on the actual cryptographic module, CC deal more with the engineering processes employed in the development of a product including hardware, software, and/or firmware.

Unlike the FIPS 140-2 standard, CC is an international standard developed by the International Organization for Standards (ISO) and the International Electrotechnical Commission (IEC) and is specifically referred to as ISO/IEC 15408:2005. Several countries contributed to developing this standard, including: Australia, New Zealand, Canada, France, Germany, Japan, the Netherlands, Spain, the UK, and the US. It is, however, recognized internationally by 28 countries.

CC also employ various accreditation levels, ranging from the lowest evaluation assurance level (EAL) 1 to the most secure level EAL 7. The Brocade encryption solution is validated to EAL4+, which is the highest level relevant to networking products. The CC validated products list can be downloaded at: http://csrc.nist.gov/groups/STM/cmvp/validation.html.

Vendors seeking CC, or ISO/IEC 15408:2005, accreditation must have their product undergo independent testing by an approved laboratory to obtain the desired EAL accreditation level. A security product under CC evaluation is referred to as a target of evaluation (TOE), which can include hardware, operating systems, computer networks, and applications.

To evaluate a TOE, the security requirements the product or system is designed to address and its security functions must be defined. This requirements and functions definition is referred to as the security target (ST).

Since there are many different security requirements addressing specific security problems, categories are created to simplify classification of individual products. Each category is represented by an implementation-independent structure known as a protection profile (PP). When evaluators evaluate a TOE, they compare the ST for that product or system against pre-defined PPs and make a statement of compliance or non-compliance to the PP.

Evaluation Assurance Levels (EAL)

Consumers may have different security requirements for individual product types and require assurances that a product meets specified criteria to address their requirements. CC uses an increasing hierarchical scale to define these assurance levels: the evaluation assurance level (EAL). Table 13 describes the seven EALs defined by ISO/IEC 15408:2005.

Table 13. Common Criteria evaluation levels

Evaluation Assurance Level	Definition
EAL1	Functionally tested
EAL2	Structurally tested
EAL3	Methodically tested and checked
EAL4	Methodically designed, tested, and reviewed
EAL5	Semiformally designed and tested
EAL6	Semiformally verified design and tested
EAL7	Formally verified design and tested

In some cases, a vendor chooses to evaluate a product to a specific EAL but may not have all of the functionality to achieve the next highest level. In this case, a vendor can "augment" the EAL achieved with some additional assurance components from the next highest EAL level.

Defense Information Systems Agency (DISA)

The US DISA provides real-time IT and communications support to the President, Vice President, Secretary of Defense, military services, and combatant commands. They create a series of security checklists or Security Technical Implementation Guides (STIG). The STIGs provide basic guidelines to implement specific types of technology that certain departments and groups within the US military must comply with. Hence they are also recognized as security policies. The checklist is used to verify that systems are being implemented in compliance with policy and are used as a baseline for audits.

One STIG applies specifically to the SAN environment: Sharing Peripherals Across the Network (SPAN). It addresses the implementation of a SAN infrastructure and devices connected to it. This STIG enforces items such as zoning, documentation, physical security, and management.

A complete list of available STIG checklists can be found at:
http://iase.disa.mil/stigs/checklist/.

Federal Information Security Management Act (FISMA)

As the number of security breaches within the US Federal government grew and raised public awareness of protection of information assets, the government was under pressure to implement standards and provide guidelines around IT security. To address these issues, Congress established the FISMA Implementation Project in January 2003, to bolster computer and network security at specific Federal government agencies and affiliated parties by mandating yearly audits.

This act was somewhat controversial and some critics felt it has become more an exercise in documentation rather than an improvement in the state of IT security within the Federal government. The concern was that government agencies would seek compliance and not security.

FISMA and Storage

Although there are no specific SAN-related standards or guidelines in FISMA, it does apply to the information that is stored in a SAN environment.

Chapter Summary

The storage and SAN component of an IT environment are often subject to compliance requirements. Compliance guidelines and legislation described in this chapter that apply to the storage and SAN environments include PCI-DSS, Breach Disclosure Laws, HIPPA, GLBA, FIPS, Common Criteria, and FISMA. Often third parties are required to ensure the credibility of compliance reports. Cryptographic material, formerly categorized as munitions, is subject to export regulations in the US.

10

Other SAN Security Topics

SAN security is still a relatively new field and has not yet achieved mainstream status. Efforts have been made by various organizations, however, including Brocade, to assemble and disseminate more information on this important subject and develop a more structured approach to security in the SAN and storage space. Other organizations and consortiums are developing new standards, particularly in the key management space, to simplify and enable interoperability among different vendor solutions.

New technologies are emerging in the storage industry that may have a significant impact on how storage will be managed in the future. As these new technologies mature, new vulnerabilities and risks will come along with them.

iSCSI

The iSCSI protocol was designed as an alternative to Fibre Channel, but in reality it is a complementary technology. The attraction of iSCSI was the concept of leveraging an existing LAN infrastructure to also carry block-based storage data and thus reduce the cost of a SAN. However, one of the challenges faced by iSCSI is TCP/IP, which can be a very lossy protocol with associated performance degradation. To compensate for this, TCP/IP offload engines, or TOE cards, were created to offload the CPU processing requirements for the TCP/IP stack on the server. This subsequently results in an increase in costs, which offset a good part of the benefit of using iSCSI. For this reason, iSCSI has been primarily used in environments requiring less performance. In enterprise environments, FC is usually deployed for the high-performance enterprise applications and iSCSI for the less critical, low to mid-performance servers. With the proliferation and reduced cost of 10 Gbps Ethernet, iSCSI has seen a greater adoption rate, as this technology addresses some of the performance issues previously experienced.

Security concerns with iSCSI are similar to those with TCP/IP in general, since it is based on that protocol suite. There are a few storage-specific security features available with iSCSI to authenticate devices when joining a network, for example.

There are other storage-specific security features, such as ACLs, which can be used with iSCSI. Additionally, device authentication can be accomplished using the Kerberos, SRP (Secure Remote Password), CHAP (Challenge-Handshake Authentication Protocol), and SPKM-1/2 (Simple Public Key Mechanism) protocols (which are less secure than DH-CHAP with FC). IPSec is also used with iSCSI, particularly with extended fabrics over public WANs, to maintain data confidentiality by encrypting the data stream.

FCoE/DCB

Fibre Channel over Ethernet (FCoE) has gained a considerable amount of attention in the past few years. The promise of converged Ethernet/FC networks has proved interesting to many organizations and has been the subject of great debate in the industry. The key to exchanging storage data over an Ethernet protocol requires a more robust, lossless version of Ethernet, as storage devices do not tolerate dropped frames.

Since 2009, the currently accepted lossless Ethernet protocol is Data Center Bridging (DCB), replacing CEE. Although the concept of converged networks seems appealing and appears to promise great cost savings resulting from having only one storage and LAN, deployment of a converged network has failed to reach critical mass. At this point in time, there is no cost benefit; however, that may change in the future. One area that has seen some FCoE adoption is for server connectivity (top of rack) to consolidate server I/O and reduce cabling.

The Future of Key Management

Key management is one the greatest challenges for encrypting data that will be retained for an extended period of time. This could be over 100 years in the case of certain medical records, for example, which must be kept for the life of a patient.

Key management solutions today are all proprietary, with each vendor offering different features and functionality.

OASIS and KMIP

OASIS (Organization for the Advancement of Structured Information Standards) is a consortium that drives the development, convergence, and adoption of open standards. OASIS has developed a number of security-related standards for identity management, key management, and Web service security. The Key Management Interoperability Protocol (KMIP) defines an interface between encryption devices that consume keys and the key management system that manages the keys. KMIP is an agreement between members of a consortium to use a common key management interface. KMIP is now the industry-accepted key management interface standard.

11

Brocade Data Encryption Products

Brocade has been offering data encryption functionality since the introduction of the Brocade 7500 Extension Switch, which supported IPSec for encryption of data transported over an FCIP tunnel. In September 2008, Brocade introduced a hardware platform for encryption of data-at-rest for both disk and tape media, which offers unprecedented encryption processing from a single device. This encryption solution is actually based on a switch platform and not a single-purpose appliance. It functions in the same way as a conventional Layer 2 FC switch but with the additional hardware required to support line-speed encryption and compression functionality.

Brocade Encryption for Data-At-Rest

Data-at-rest refers to all data that is no longer in motion and has been recorded on storage media such as a disk drive or a tape cartridge.

The Brocade encryption solution is available in two form factors that share the same internal hardware. The solution is available as a standalone FC switch, the Brocade Encryption Switch, and as a blade for the Brocade DCX 8510 and DCX family of backbone products, the FS8-18 Encryption Blade. The term “encryption device” is used throughout this chapter and refers to either the encryption switch or the encryption blade. The Brocade encryption solution includes the Brocade encryption device, along with all other components required for a production environment, such as the key vault.

The Brocade encryption device features the following:

- Up to 96 Gbps processing bandwidth for disk encryption
- Up to 48 Gbps processing bandwidth for tape encryption with compression
- Encryption using the industry standard AES-256 algorithm
- Compression using a variant of gzip
- 8 Gbps FC port speeds
- Disk encryption latency of 15–20 microseconds
- Tape encryption and compression latency of 30–40 sec
- Brocade-developed encryption ASIC technology
- FC switching connectivity based on the Brocade Condor 2 ASIC
- Dual Ethernet ports for HA synchronization and heartbeats
- Smart Card reader used as a System Card (ignition key optional)

  The ignition key feature is built into the encryption solution at no extra cost and enabled as an option to enhance the level of security on the switch. The ignition key is a Smart Card, which can be inserted into the Smart Card reader to initialize the cryptographic functionality of the switch. The Brocade Encryption Switch behaves as a regular 8 Gbps Layer 2 FC switch only until the ignition key is inserted and encryption enabled.

  If the ignition key feature is used, it is imperative to store the Smart Card in a safe location after the cryptographic functions of the switch have been enabled. The Smart Card must be reinserted in the reader (see Figure 41 and Figure 43) each time the switch is rebooted or power cycled to enable the cryptographic capabilities of the switch.

Brocade Encryption Switch

The Brocade Encryption Switch is the standalone version of the hardware encryption device for data-at-rest. It offers the following features:

- 32 x 8 Gbps FC ports
- Three redundant fan modules
- Two redundant power supplies
- USB port
- One RJ-45 GbE management port

- Two redundant RJ-45 GbE ports for intercluster communication
- FIPS 140-2 Level 3 compliant cryptographic boundary cover
- Smart Card reader used as a System Card (ignition key - optional)

Figure 41 and Figure 42 illustrate the Brocade Encryption Switch and its components.

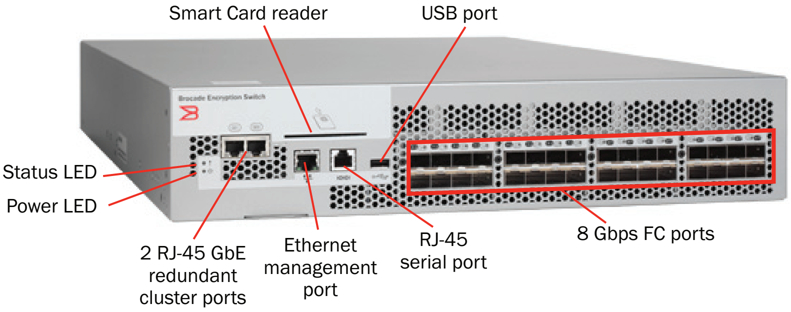

Figure 41. Front view of the Brocade Encryption Switch

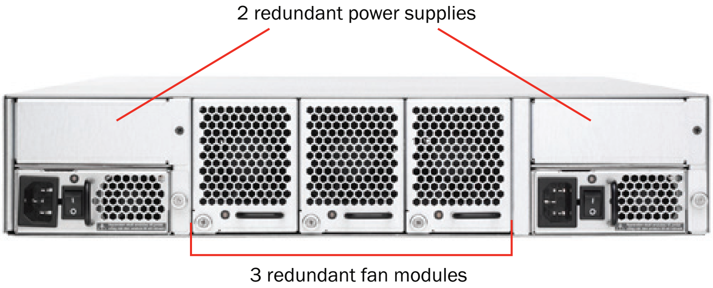

Figure 42. Rear view of the Brocade Encryption Switch

The Brocade Encryption Switch is also available in an entry-level version for disk encryption. Some companies may not require the full 96 Gbps of bandwidth for disk encryption. The entry-level version of the encryption switch was created offering up to 48 Gbps of encryption processing bandwidth at a lower price point. The entry-level version is identical to the advanced throughput version, but with half the encryption processing available for use. All 32 FC ports remain enabled and can be used to connect hosts and storage devices. Later, if the 48 Gbps encryption bandwidth is exceeded, a simple license upgrade to the full 96 Gbps bandwidth version can be purchased.

Brocade FS8-18 Encryption Blade

The Brocade FS8-18 Encryption Blade is the embedded version of the Brocade Encryption Switch for the Brocade DCX 8510/DCX/DCX-4S Backbone. It has the same functionality and performance characteristics as the Brocade Encryption Switch:

- 16 x 8 Gbps FC ports
- USB port
- One RJ-45 GbE management port
- Two redundant RJ-45 GbE ports for intercluster communication
- FIPS 140-2 Level 3 compliant cryptographic boundary cover
- Smart Card reader for the System Card (ignition key optional)
- Up to four FS8-18 blades supported in one Brocade DCX 8510/DCX/DCX-4S chassis

Figure 43 and Figure 44 illustrate the FS8-18 blade and its components.

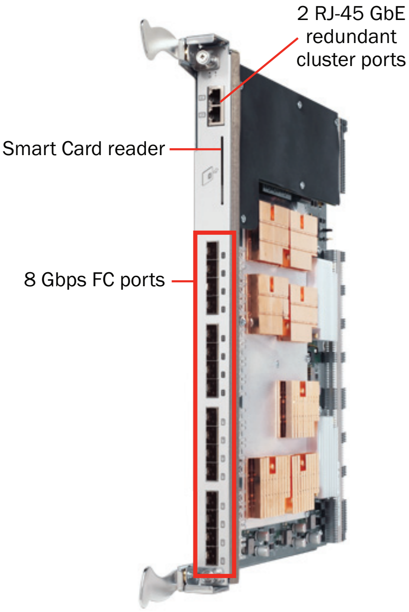

Figure 43. Profile view of the Brocade FS8-18

The FIPS 140-2 Level compliance posed several challenges for the FS8-18. The typical Brocade enterprise-class platform blade has all of its ASICs exposed on the card. To prevent tampering with the components of the blade involved in the cryptographic (crypto) process it was necessary to build a physical crypto security boundary protecting all the memory, true random number generator, encryption, and Condor-2 ASICs. This physical boundary was secured by placing a cover over these components, which in turn posed a new challenge: cooling. The cover cannot have vents for air circulation, since this could allow intruders to access the internal components with specialized tools. Instead, copper heat sinks were placed on the cover to dissipate the heat, as shown in Figure 44.

As with the Brocade Encryption Switch, the FS8-18 Encryption Blade is also available in an entry-level version for disk encryption. The entry-level version of the blade, though, applies to the entire DCX 8510/DCX/DCX-4S chassis. The Brocade DCX family chassis can support from one to four FS8-18 blades per chassis. With the entry-level version, each blade is limited to 48 Gbps of encryption processing bandwidth per blade for disk, regardless of the number of blades installed. The entry-level version affects only the disk encryption processing bandwidth; all 16 FC ports remain enabled and can be used to connect hosts and storage devices. Later, if the 48 Gbps encryption bandwidth is exceeded, either new FS8-18 blades can be added or all the encryption blades in the chassis can be upgraded with a simple chassis-level license upgrade to the full 96 Gbps bandwidth.

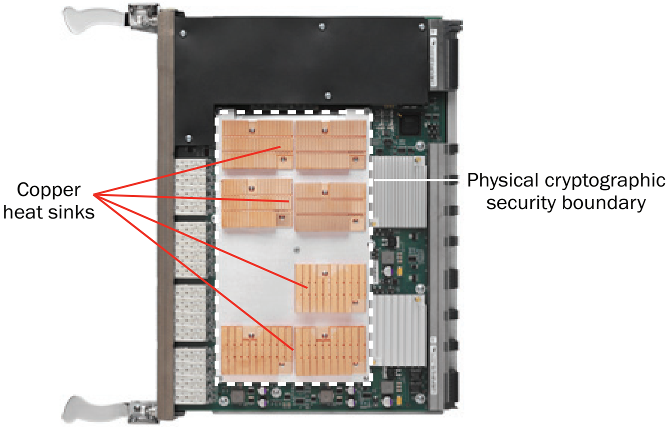

Figure 44. Side view of the Brocade FS8-18

One advantage of using encryptino blades is that you do not need to be concerned with ISLs, since all of the encryption is performed via the backplane. It is not necessary to connect the hosts or storage devices involved in the encryption process into one of the 16 FC ports on the blade. In fact, encryption will take place even though there are no devices directly connected into the blade. This is accomplished using the frame redirection technology described in “Frame Redirection” on page 33.

Brocade Encryption Features

The Brocade encryption solution offers several features that were introduced in releases of Fabric OS subsequent to the initial release of the hardware.

The following are the encryption features of the Brocade encryption solution:

- NetApp KM500 Applicance
- EMC Data Protection Manager (DPM, formerly RKM)
- HP Enterprise Secure Key Manager (ESKM)
- IBM Tivoli Key Lifecycle Manager (TKLM)
- Thales e-Security keyAuthority (formerly TEMS)
- SafeNet KeySecure
- Symantec NetBackup
- EMC NetWorker
- CommVault Data Protection
- HP Data Protector
- BakBone NetVault
- CA ARCserv
- Symantec Backup Exec
- Microsoft System Center Data Protection Manager
- Disk encryption
- Tape encryption
- Offline, in-place, first-time encryption and rekeying
- Online, in-place, first-time encryption and rekeying
- High availability (HA) cluster

- Data encryption key (DEK) cluster
- DataFort compatibility mode
- FIPS 140-2 Level 3
- Common Criteria (EAL-4+)
- Multi-path rekeying to a LUN through an EE
- System card to enable crypto capability
- Quorum authorization of sensitive operations
- Access Gateway for third-party support (switch only)
- LUN Decommissioning

Brocade Encryption Process

The Brocade encryption solution uses the industry standard AES-256 encryption algorithm implemented in hardware:

- **Disk encryption** is performed using the XTS mode of encryption, which is better suited for fixed-block data
- **Tape encryption** is performed using the GCM mode of encryption, which is better suited for variable-length and streaming data

Compression is an important component of a data-at-rest encryption solution for tape. Once data is encrypted, it is no longer compressible. Compression works on the principle of searching for patterns and optimizing them. Encryption takes data and removes all patterns by randomizing the data. Once the data is randomized and all patterns are removed, then the compression algorithm has no patterns to optimize. If encrypted data is sent directly to a tape drive, the native compression capabilities of that tape drive will no longer be effective. Hence, it is important to compress the data first and then send it to the tape drive to prevent an unnecessary increase in the number of tape media used for backups.

The compression algorithm used in the Brocade encryption solution is based on a variant of the standard gzip algorithm. The compression ratio obtained using this compression algorithm may vary, like any other compression algorithm, depending on the type of data and how compressible it is. Data with a a great deal of white space compresses quite well, while some data may not compress at all.

CryptoTarget Containers

A Crypto Target Container (CTC) is created for each storage target port hosted on a Brocade encryption device and is used to set up the encryption to a media. A CTC can be composed of only one storage port target but it can have multiple initiators or hosts associated with it. A CTC can also have several LUNs behind the storage port in the CTC. Furthermore, once a storage port has been assigned to a CTC, it cannot exist or be defined in another CTC. Essentially, this forces all traffic that goes through a specific storage port to be encrypted and to go through the same encryption device.

NOTE: The storage port can still be made accessible (with appropriate zoning) for other hosts in case encryption is not required for their LUNs. In this case, these LUNs are not added to the CTC.

First-Time Encryption and Rekeying

The first-time encryption (FTE) process can generally be performed using two methods. The first is to copy the original cleartext data on the production LUN to a second LUN with an equivalent amount of disk space while encrypting the data at the same time. This method obviously requires an equivalent amount of disk space, which may or may not be available. Furthermore, once the data is copied to the new LUN, the servers must now point to it, which requires rebuilding the device tree on the server and may result in disruption of the production environment.

The other method, which is implemented by the Brocade encryption solution, is to perform the FTE in-place on the same LUN. The process involves reading the first logical block on the LUN (which is in cleartext), encrypting it, and then writing it back to its original location as ciphertext. Subsequent blocks are encrypted in the same manner sequentially until all blocks in the LUN have been encrypted. This process can be performed offline or online, depending on an organization's business requirements.

Figure 45 illustrates the FTE process.

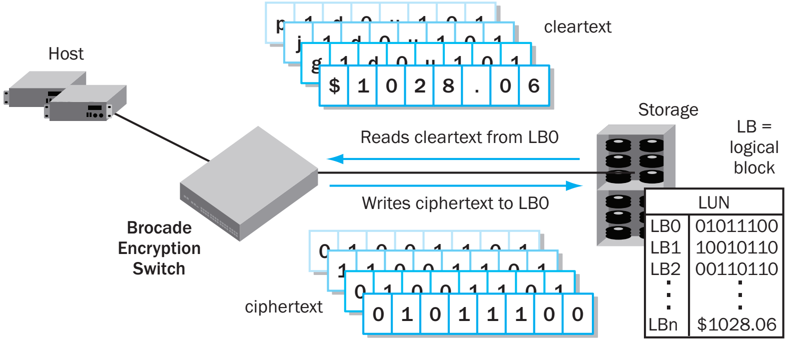

LUN	
LB0	01011100
LB1	10010110
LB2	00110110
⋮	⋮
LBn	$1028.06

Figure 45. First-time encryption operation

The next encryption process to consider is the rekey operation, in which a LUN is re-encrypted using a different key. There are two basic reasons why a rekey operation would be performed: a compromised key or a security policy requirement. If a key is lost or stolen, it is compromised and the data encrypted with this key can no longer be considered secure. The security or risk management department of an organization may implement a policy requiring that all keys must be refreshed on a specified schedule, such as every 36 months. This is often done out of fear that keys may have been compromised without their knowledge and the organization may prefer to err on the side of caution by forcing a rekey of all encrypted data after a defined period of time. However, most of the time, the primary reason organizations perform a rekey operation is that they are mandated to do so as a result of a compliance requirement, such as with the PCI-DSS. Rekeying can be performed automatically by setting an expiration date on a key using the Brocade encryption device, but this is not generally recommended. It is preferable to expire keys manually to control exactly when this is performed and schedule off-peak hours.

In-place rekeying is not possible for tape, since a tape drive is a steaming device and the media itself is flexible. Rekeying data on a tape involves copying it to a new tape and encrypting it with a different key as the data is copied. In the case of disk media, the process is much simpler, since the LUN with the compromised key can be rekeyed in-place and online if necessary.

During the rekey operation, the LUN actually has two keys assigned to it, one used for new writes and one for reading data that has not yet been rekeyed. Once the rekey process is completed, the original key is no longer used. As with a first-time encryption, the rekey operation can be performed online or offline.

Clustering and Availability

One of the principle tenets of security is maintaining availability. Needless to say, downtime can be expensive and precautions must be taken to prevent a loss of availability of the information. This is particularly true for encryption solutions, since there is a complete dependence on the encryption keys to recover encrypted information. Compounding this problem is the importance of the applications that require encryption. Any loss of availability of information that is important enough to require encryption is mostly likely to be disastrous for its owners. Extensive precautions must be taken to protect the keys and to maintain the availability of the encryption solution.

As with any IT solution, there are several ways to ensure availability. Choosing the best method to maintain availability depends on the value of the information (and impact of a loss of availability), the risk and probability of disruption, and the cost of implementing high availability. As with all aspects of IT, it's about getting the best value for your investment.

Clustering is commonly used to ensure protection against hardware failure. There are two types of clusters for Brocade encryption solutions, which can be used independently or simultaneously. The high availability (HA) cluster provides hardware redundancy for the encryption devices. The DEK cluster allows two or more encryption devices to share the same keys.

HA Cluster

The HA cluster is an active-passive clustering configuration in which one encryption device is a warm standby for the other encryption device it is paired with. Only two encryption devices can form an HA cluster and they must exist within the same fabric. Heartbeats are exchanged between the encryption devices using redundant Gigabit Ethernet ports through an out-of-band dedicated network to let the other device know it is still “alive.” This same dedicated network is used to synchronize key state information between the units to allow one device to take over for the other when the HA pair has failed and no longer appears in the nameserver. Unlike the DEK cluster described below, the HA cluster will not result in a path failover following a failed encryption device.

Since the HA cluster uses an active-passive configuration per CTC, it is more efficient to balance the load across both encryption devices instead of having the entire load on one unit with the other being entirely inactive unless the active unit were to fail. It is possible for each encryption device to be active simultaneously and carry its own encryption load. In this case, each unit is active with its own load and, at the same time, can be passive while waiting for the other unit to fail over. In the event that one encryption device fails, it is important to consider the available bandwidth on the other cluster member and its impact on application performance.

For example, let’s say that Encryption Device A in the cluster is currently pushing 52 Gbps of traffic and Encryption Device B is pushing 61 Gbps. If Encryption Device B fails, Encryption Device A will take over the CTCs. Since Encryption Device A is already pushing 52 Gbps and now has an additional 61 Gbps, for an aggregate of 113 Gbps of traffic, this exceeds the 96 Gbps capability of the encryption device. At this point, there will be more I/O going through Encryption Device A than it can handle and a performance bottleneck will occur, resulting in a downgraded performance of the production environment.

DEK cluster

The DEK cluster by definition shares the same data encryption keys as all other encryption devices within a cluster management group. The DEK cluster is composed of encryption devices that are members of the same encryption group. An encryption group contains several encryption devices that share the same DEKs. For each encryption group, there must be one encryption device designated as the group leader. The group leader is responsible for functions such as the distribution of the configuration to the other members of the group, authenticating with the key vaults, and configuring CTCs.

It is important to note that the DEK cluster offers good redundancy. The loss of one encryption device would not necessarily result in a loss of production, given that disk solutions are implemented using dual paths. With a dual path, there is always an alternative path for the data to get to the LUN. For this reason, the HA cluster is very seldom used, other than for the most stringent application requirements and environment, where downtime cannot be tolerated and intra-fabric redundancy is required.

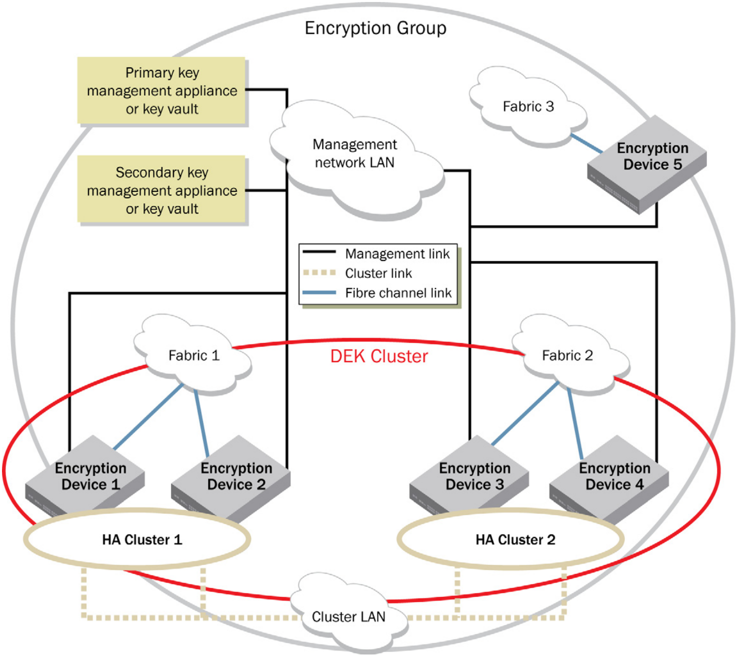

Figure 46. HA and DEK cluster

Key Management

Once data is encrypted onto a storage media, the keys become highly critical and extensive measures must be taken to protect them. Appropriate measures should be taken to manage these keys throughout their lifecycle. Keys need to be backed up as they can be lost, stolen, destroyed intentionally, or expired after a pre-determined period of time.

Loss of the encryption keys is equivalent to losing the data. Unlike data-in-flight, the keys for data-at-rest must be available for relatively long periods of time, depending on the type of information being encrypted. With patient health records, for example, it is possible that information is kept for the lifetime of a patient, which can be over 100 years. Keys can also be stolen or compromised, in which case the information would have to be re-encrypted using a different key to ensure the confidentiality of the information. Media such as disk and tape also have a limited shelf life and may undergo evolution cycles to an eventually incompatible format (remember 8-track tapes and floppy disks?). The information needs to be refreshed as the media expires and must be re-encrypted using the same key (exact replica of tape) or a different key.

For redundancy, a typical key vault will be implemented with two or more units to prevent single points of failure. If the primary key vault becomes unavailable, the secondary or other key vault can accept or provide keys to the encryption device.

The following key management solutions are currently supported:

- NetApp Lifetime Key Management (LKM)
- EMC Data Protection Manager (DPM, formerly RKM)
- HP Enterprise Secure Key Manager (ESKM)
- Thales Encryption Manager for Storage (formerly TEMS)
- IBM TKLM v2
- SafeNet e-Security KeySecure

Brocade supports the OASIS KMIP (Key Management Interoperability Protocol), which has become the industry-accepted key management interface standard.

Brocade encryption devices generate the actual data encryption key and store it locally in its cache. The DEK is used to encrypt data using the AES-256 encryption algorithm. Before any data encryption begins, the key must be backed up to a key vault, or key manager, and then placed in the local cache before it can be used. Subsequently, once the DEK has been committed to the key vault and an acknowledgement has been received from the key vault, the DEK is exchanged with the other members in the encryption group.

When a new LUN, tape media, or LUN with existing cleartext data is encrypted, the Brocade encryption device generates a new DEK. This key is then backed up to the primary key vault, and secondary key vault if it exists. Once the primary key vault has successfully stored the DEK, it confirms this to the encryption device. The DEK is then synchronized with all of the other members in its encryption group, as shown in Figure 47. Only once this has occurred will the new key be used to encrypt actual production data.

Redundant Key Vaults

Key vaults may also be configured in a clustered configuration to provide redundancy. Each key management solution vendor offers different clustering features and functionality , but all of them provide some form of clustering capability. Although clustering the key vault is an optional feature, it is certainly recommended as a best practice. Ideally, a key vault should be located in at least two physically separate locations to provide the maximum redundancy in the event of a catastrophe that destroys an entire site.

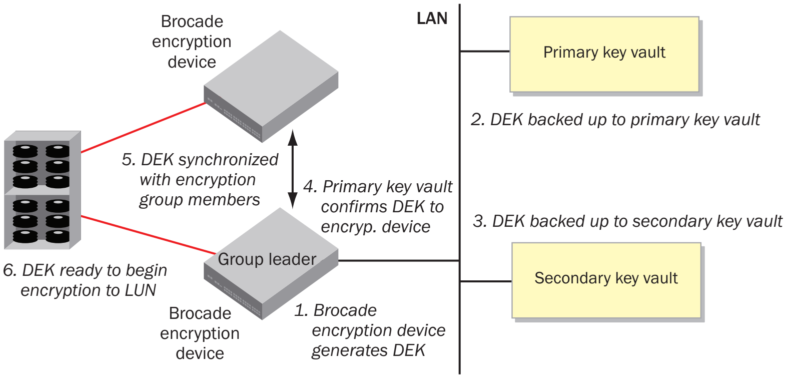
Brocade encryption device
LAN
Primary key vault
2. DEK backed up to primary key vault
5. DEK synchronized with encryption group members
4. Primary key vault confirms DEK to encryp. device
3. DEK backed up to secondary key vault
Group leader
Secondary key vault
6. DEK ready to begin encryption to LUN
Brocade encryption device
1. Brocade encryption device generates DEK

Figure 47. DEK synchronization

DataFort Compatibility Mode. The NetApp DataFort encryption appliance was at one point the market leader in the storage encryption space. NetApp and Brocade established a strategic relationship to use the Brocade encryption solution as the next-generation DataFort. One of the challenges to making this happen was determining what to do with existing DataFort customers who have thousands of tapes previously encrypted using the DataFort product. The solution was to create a DataFort compatibility mode in the Brocade encryption solution to read media previously encrypted with the DataFort appliance. The DataFort compatibility mode can read either disk or tape media and can also write to new tapes or existing LUNs encrypted with the DataFort format.

The DataFort compatibility mode does several things. The Brocade encryption device uses the ECB mode of operation for the AES-256 encryption algorithm, which is used by the DataFort product. The metadata format used by the DataFort product replaces the native format used by the Brocade encryption device. The compression algorithm is the same on both platforms so there is nothing special which must be done for compression.

The DataFort compatibility mode enables an easy migration from the DataFort product to the new Brocade encryption solution, which will also integrate with the NetApp LKM key management solution, or the LKM-compatible SafeNet KeySecure (in SSKM mode), already deployed with the DataFort encryption appliance. However, customers using earlier versions of the LKM, which was software-based, need to upgrade to the SafeNet KeySecure appliance now that the LKM appliance has reached end-of-availability.

Encrypting with Backup Applications. Although only the payload portion of the frame is encrypted, special considerations must be taken to adapt to each backup software vendor. There are two basic elements in a backup solution that an encryption solution must consider.

The first is how the backup application writes its metadata to the tape media. This is necessary to determine where to write the key information on the media for later data recovery. Obviously, the actual cleartext key is not stored on the tape media itself, which would be equivalent to sliding a spare house key under the front porch doormat. In fact, only an index (key ID) referring to the key is written to the tape media as part of the tape header written by the backup application.

The second consideration is how each backup application handles tape pools. Keys can be assigned either on a per-tape media basis or on a per-pool basis. As a best practice, it is preferable to assign one key per physical tape media to reduce the rekey overhead in the event that a key were to be compromised. Nevertheless, for some special corner cases, it may be useful to use one key per pool. For instance, if a set of tapes is planned to be sent to a third party, perhaps for auditing purposes, a single key could be used for the entire tape set to simplify the reading of the tapes at the other end.

The following backup software solutions are supported in FOS 7.0 and later:

- Symantec (formerly Veritas) NetBackup
- IBM Tivoli Storage Manager (TSM)
- EMC (formerly Legato) NetWorker
- CommVault Galaxy Data Protection
- HP Data Protector
- BakBone NetVault
- CA ARCserve
- Symantec BackupExec
- Microsoft System Center Data Protection Manager

Brocade Encryption Internals

The Brocade encryption device is a state-of-the-art hardware product built to integrate seamlessly into an existing SAN infrastructure and integrate with the market leaders of encryption key management. Both the encryption switch and the encryption blade essentially share the same hardware components and offer the same functionality, but in a different form factor. The encryption blade does not have a USB port, serial port, or management Ethernet port and the switch does not have a backplane. Figure 48 and Figure 49 illustrate the simplified internal architecture of the Brocade Encryption Switch and FS8-18 Encryption Blade respectively.

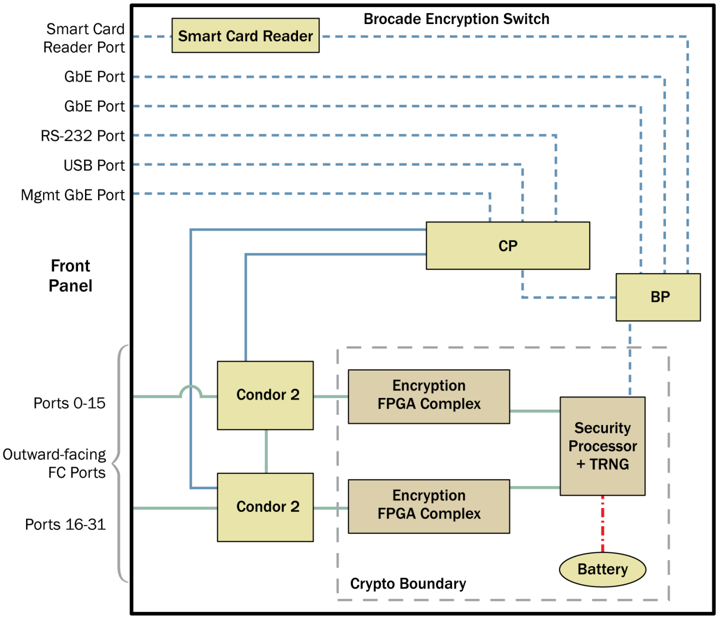

Figure 48. Brocade Encryption Switch internal architecture

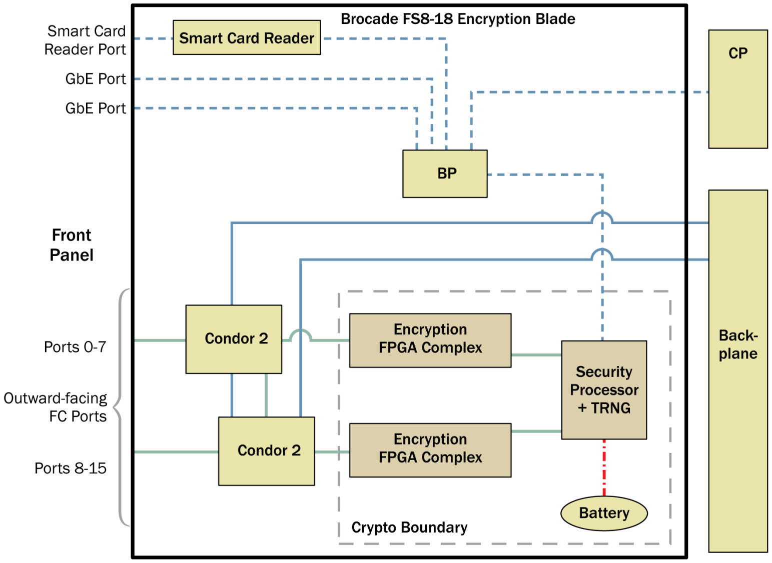

Figure 49. Brocade FS8-18 Encryption Blade internal architecture

The components described in the following three sections are enclosed within a physical crypto boundary. The security boundary is designed to comply with the FIPS 140-2 standard at Level 3 to isolate all hardware components involved in the processing of cleartext keys. The encryption switch cover is the physical crypto boundary for the Brocade Encryption Switch and the encryption blade has a special cover that encloses the necessary hardware on the card.

Encryption FPGA Complex

The FPGA complex is composed of several FPGAs and other hardware components. The principle encryption component is the FPGA (Field Programmable Gate Array). An FPGA is a programmable hardware device that contains instructions to perform specific functions. The advantage of the FPGA is that it is programmable and the instructions can be changed at any time. A new feature or enhancement can be made without requiring a hardware upgrade.

The FPGA complex is where the actual encryption and compression is performed in the encryption device, in addition to a few other functions.

Security Processor + TRNG

The Security Processor provides data security functions such as generating and processing symmetric keys (the DEK) based on the TRNG.

The TRNG (True Random Number Generator) is the hardware component used to generate the random number from which the DEK is generated. A TRNG uses physical phenomena such as transient noise to truly randomize the random number generation process. The TRNG used in this solution meets the FIPS validation requirements.

Battery

A Lithium-ion battery is used when there is no power to the encryption device. This battery has a life span of approximately seven years after power has been removed from the encryption device. It is used primarily to sustain the FIPS 140-2 Level 3 tamper response mechanism, which zeroizes the keys stored in the local cache once tampering has been detected.

The remaining components are found outside the security boundary.

Control Processor (CP)

The Control Processor performs various control and coordination functions such as authentication processes.

Blade Processor (BP)

The Blade Processor acts as a bridge between the Security Processor and the Control Processor, as well as with the Smart Card reader and GbE ports.

Condor 2 ASIC

The Condor 2 ASIC features forty 8 Gbps ports and is the heart of the FC Layer 2 switching. Each encryption device has two Condor 2 ASICs.

Design and Implementation Best Practices

The Brocade Encryption Switch, like any other security product, does not come fully configured out of the box. It must be configured properly and be part of a well-designed architecture with the appropriate operational procedures to ensure continuous and secure operation. This section outlines some best practices for the design and implementation of the Brocade encryption solution.

Encryption is only one component of a comprehensive SAN security program. An organization may have the best encryption solution possible, but if it is installed on a SAN with security holes, then the entire solution may be vulnerable. In security, a system is only as strong as its weakest link, which is usually the place attackers will target first.

The design and best practice recommendations in this chapter are not meant to be comprehensive. For more information on design and implementation best practice for the Brocade encryption solution, please refer to the Brocade Encryption Best Practices Guide, available through the local Brocade contact person.

Management Interfaces

Managing and configuring the Brocade encryption solution can be performed either with the FOS CLI or Brocade DCFM/BNA Enterprise version, as well as DCFM/BNA Pro/Pro+. As a best practice, it is highly recommended to use DCFM/BNA. The CLI requires several commands to perform certain operations, which can be performed with one mouse-click in DCFM/BNA. Furthermore, typing multiple CLI commands increases the risk of typing errors, resulting in potential configuration errors. The DCFM/BNA interface also provides wizards that guide users through the configuration process to further reduce the risk of errors introduced as a result of improper sequencing of commands.

The management interfaces should never be accessed using unsecure protocols such as telnet for the CLI or HTTP for DCFM/BNI. Use secure protocols, such as SSH instead of telnet and HTTPS instead of HTTP, and block or disable their equivalent unsecure services.

For additional protection, the System Card or ignition key feature should be implemented and a Smart Card required, enabling the encryption capability of the switch. This will prevent someone who steals both the switch and the disk media from being able to decrypt the data on the storage media. Of course, it is equally important to store the System Card in a secure location away from the encryption switch and storage media.

Availability

As with any IT solution, there are many ways to ensure availability. Selecting the best method to maintain availability depends on the value of the information (and impact of a loss of availability), the risk and probability of disruption, and the cost of implementing high availability.

Clustering

Clustering is a commonly used method to ensure protection against hardware failure. There are two types of cluster for Brocade encryption solutions, which can be used independently or simultaneously. The high availability (HA) cluster provides hardware redundancy for the encryption devices. The data encryption key (DEK) cluster allows two or more encryption devices to share the same keys.

For tape encryption using a single fabric, a single encryption device could be sufficient, since tape drives are single attached devices (actively attached devices). However, some organizations consider the backup application as mission-critical or high priority due to a service-level agreement that must be respected. If this is the scenario, a business case can be made to justify the use of a second encryption device to form a HA cluster.

For disk encryption using a dual-fabric configuration, the minimum requirement is for one encryption device per fabric. In the event of the failure of one encryption device, the MPIO software on the host automatically fails over the traffic to the remaining path. This may result in degraded performance in some heavily used systems, which may or may not be acceptable. If it is not acceptable, then add a second encryption device in each fabric to form two HA clusters.

For redundancy, it is good practice to implement more than one path from the disk storage device to the fabric. If more than one path exists in the same fabric from a host to a LUN, then it is important to use FOS 6.3 or later when performing a first-time-encryption or a rekey operation. Multipath rekeying operations through a single encryption engine are not supported prior to FOS 6.3.

Redundant Key Vaults

Key vaults can also be configured in a clustered configuration to provide redundancy. Each key management solution vendor offers different clustering features and functionality, but all of them provide some form of clustering capability. Although clustering the key vault is an optional feature, it is certainly recommended as a best practice. Ideally, a key vault should be located in at least two separate locations to provide the maximum redundancy.

Encrypting Disk Storage

Data can be encrypted on disk storage at the LUN level. One single key is used to encrypt the data on a LUN except during a rekey operation, which requires two keys. LUNs on a disk array are discovered through the standard SCSI LUN Discovery process.

Performance

As explained earlier, the latency of the Brocade encryption devices is practically negligible compared to the time it takes to complete an I/O operation. However, a complex fabric may have multiple ISLs and offer many paths between the various devices within the fabric. As discussed earlier, the frame redirection feature can automatically redirect frames to the encryption device regardless of where it is located in the fabric. However, certain locations for the encryption devices offer the best performance.

The basic concept of locality applies to the encryption solution as well as standard FC fabric designs. Locality simply states that a host and its storage devices should be located as closely as possible to one another, given a specific architecture. For example, the highest locality occurs when a host and its associated storage device are connected to the same switch in a fabric or the same blade in a director or backbone. Essentially, SAN placement of the encryption devices should be done as close as possible between the host and its storage devices.

To avoid forcing traffic to pass through ISLs, a backbone can be used to consolidate multiple switches. The Brocade FS8-18 Encryption Blade in a Brocade DCX 8510, DCX or DCX-4S does not require ISLs to perform the encryption and all traffic destined for encryption passes through the backplane.

First-Time Encryption and Rekeying Operations

Many organizations have a policy regarding a sensitive operation such as a data migration or encryption of data on a LUN to quiesce the environment first and then perform this operation offline. Other organizations cannot tolerate downtime and must perform an FTE or rekey operation online.

The Brocade encryption solution allows for online or offline FTE or rekey operations. An online FTE or rekey operation may result in performance degradation of the applications accessing the LUN as a result of I/O contention between the application requirements and the FTE or rekey operation.

A rekey operation could be required after the LUN's DEK has been compromised or after it has expired. It is possible to configure the Brocade encryption device to automatically begin a rekey operation once the DEK expires. However, as a best practice, it is preferable to configure the encryption device manually to perform the rekey operation. Since a rekey operation is very I/O intensive and may negatively impact application performance, a manual rekey would allow the scheduling of the rekey operation when it is more convenient, such as during off-peak hours.

Other Best Practices

Firmware Upgrades

Firmware upgrades on the Brocade encryption device are disruptive to encryption traffic I/O. However, layer-2 FC traffic that is not being redirected will not be affected, but redirected traffic will be affected since the encryption engines and Blade Processor must reset.

To avoid production downtime for disk environments using a dual-fabric configuration, upgrade the switches on Fabric A first and then fail over the traffic back to Fabric A. When both paths are online again, Fabric B is failed over to Fabric A and Fabric B is upgraded.

To avoid impacting production for tape environments attached to a single fabric, it is simply recommended to perform the upgrade during off-hours or in the next available maintenance window.

Key Management

Key Expiration. Part of managing the keys is determining how long a key should exist. Many organizations never expire a key, while others require expiration every six months (or more). There is no general rule as to the frequency of key expiration and it depends entirely on the business requirements and tolerance to the risk that a 256-bit key will go stale or be compromised. Since an online rekey operation can affect application performance and an offline rekey requires downtime, most organizations would rather not perform a rekey too frequently. Generally, it is considered safe to expire 256-bit keys somewhere between every two to four years.

Key-Per-Media vs. Key-Per-Pool. For tape encryption, a single DEK can be assigned to one tape media or to an entire pool of tapes. The best practice is to have one DEK per tape media. In the event the DEK is compromised, it is much simpler to create a new backup for one tape as opposed to an entire pool of tapes.

Brocade Encryption for Data-In-Flight

Data-in-flight refers to data that is in transit. Data-in-flight could be moving across a copper cable, dark fiber, or even through the air using wireless devices. Data-in-flight poses a different problem from a data confidentiality perspective, particularly when it is transported over public networks. Data transported between two remote sites using an FCIP tunnel over a public network can be vulnerable if it is sent in cleartext format. The Brocade 7800 Extension Switch and Brocade FX8-24 Extension Blade support FCIP tunneling and address the data confidentiality issue by encrypting data using the well-known IPSec protocol.

Brocade 7800 and FX8-24

The Brocade 7800 Extension Switch and equivalent FX8-24 Extension Blade are capable of connecting two fabrics over great distances using the FCIP protocol. An FCIP tunnel is created between two sites, which are connected together over a public IP-based WAN. Since the WAN is a public network, there is always a risk of data transferred over such a network being sniffed by an unauthorized user. To protect the FCIP tunnel from a sniffing attack, the data-in-flight over the WAN should be encrypted. This can be done using the IPSec protocol.

Table 14 shows the different encryption and authentication algorithms supported with the Brocade implementation of IPSec for FCIP.

Table 14. IPSec encryption and authentication algorithms for FCIP

Encryption Algorithm	Authentication Algorithm
3DES	SHA-1
AES-128 (default)	MD5
AES-256	AES-XCBC

Data-at-Rest Solution for Data-In-Flight Problem

It is possible to use the Brocade data-at-rest encryption solution to encrypt data-in-flight over distance with proper design.

Data being replicated or sent over a dark fiber from the primary data center to the DR data center can be encrypted using a data-at-rest encryption solution. If the encryption device and host is located in the primary data center and the storage is at the secondary site then the encryption device would encrypt the frame payload before sending it over the dark fiber connection. At this point, the payload is encrypted and cannot be read if captured along the way. This technique is often used for cross-site backups, where data stored at one site is backed up to a tape library located at another site. Figure 50 demonstrates how the data-in-flight for a cross-site backup can be encrypted using a data-at-rest encryption solution.

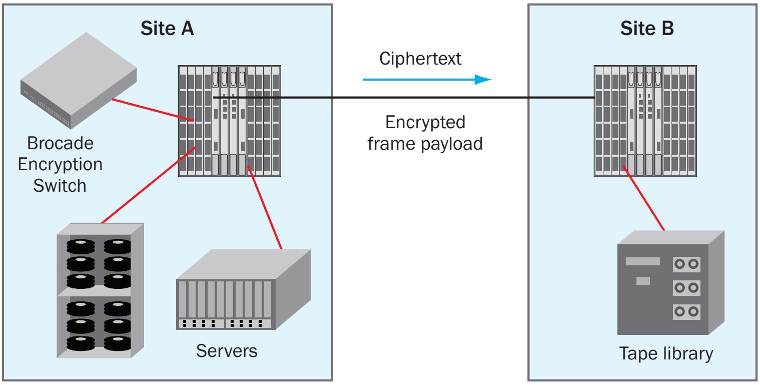

Figure 50. Encrypted cross-site backup

Similarly, this same strategy could be used for data replication between two sites.

Figure 51 illustrates how a data-at-rest encryption solution can be used to encrypt data on the dark fiber during data replication. In this case, the data stored on the primary data center is encrypted using the encryption device. The disk-to-disk replication application (such as EMC SRDF or IBM PPRC) will simply copy the data which is already in ciphertext format to the alternate site where it will be stored as is in ciphertext.

The latest Brocade FC products are based on the 16 Gbps Condor-3 ASIC. This new ASIC has built-in encryption and compression capabilities that allow SAN administrators to configure up to two ISL ports (E_Ports) per ASIC for data-in-flight encryption. This feature may also be used to encrypt replicated disk data between two sites or for cross-site backups when both sites are connected via ISLs using dark fiber. A new 16 Gbps switch at one data center will encrypt outbound frames on the ISL and get decrypted at the other end by another 16 Gbps FC switch. The Condor-3 ASIC is also capable of compressing data. As seen previously, it is not possible to compress encrypted data, so the compression is the first operation to take place when used in conjunction with encryption.

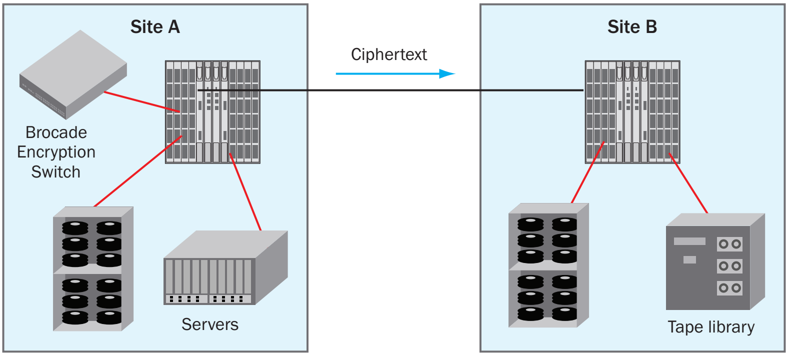

Figure 51. Encrypting data over dark fiber with data-at-rest encryption

Chapter Summary

Brocade provides encryption solutions for both data-at-rest and data-in-flight. The Brocade Encryption Switch and the Brocade FS8-18 Encryption Blade for the Brocade DCX backbone family can be used for both disk and tape media to encrypt data-at-rest. The Brocade encryption switch is a standard 8 Gbps Layer 2 FC platform and, when used in encryption mode, provides robust encryption (and compression) in combination with third-party key management. The addition of a Smart Card reader for an ignition key provides additional security.

Brocade offers data encryption for data-in-flight in the Brocade 7800 Extension Switch and Brocade FX8-24 Extension Blade, both of which support IPSec for encryption of data transported over an FCIP tunnel. The Brocade data-at-rest encryption solution, described in detail in this chapter, can be used to encrypt data-in-flight. The encryption device in the primary data center encrypts the frame payload before sending it over the dark fiber connection.

The latest Brocade 16 Gbps FC technology, based on the Condor-3 ASIC, also offers the capability to encrypt data-in-flight for up to two ISLs.

A

Fabric OS Security Features Matrix

Legend for security level:

B = Basic, I = Intermediate, A = Advanced, O = Optional

Security Feature	FOS 2.x	FOS 3.x	FOS 4.x+	Security Level
SSH (AES, 3DES, RSA)	-	-	4.1.1	I
OpenSSH Public Key Authentication	-	-	6.1	I
TLS/SSL (AES, 3DES, RC4/RSA)	-	N/A	4.4	I
HTTPS (AES, 3DES, RC4/RSA)	-	N/A	4.4	I
PEAP/TLS	-	-	5.3	A
SNMPv3 (AES, 3DES)	-	-	4.4 (DES)	I
SHA-256	-	-	7.0	A
NTP (to synchronize timestamps)	2.6.1	3.2	4.2	B
NTP (up to 8 NTP servers)	-	-	5.3	B
PKI digital certificates (SLAP/RSA) Not factory shipped since May 15, 2005	2.6	3.1.0	4.1	A
DH-CHAP (E-Ports, switch binding)	-	3.1.0	4.4	A
DH-CHAP (F-Ports, port binding)	-	-	5.3	A
DH-CHAP enforcement for HBAs	-	-	6.2	A
MS-CHAPv2	-	-	5.3	A

Security Feature	FOS 2.x	FOS 3.x	FOS 4.x+	Security Level
Secure RPC (for Brocade API using SSL)	-	-	4.4	A
Secure File Copy (SCP) for configUp/Download	-	-	4.4	I
Secure File Copy (SCP) for firmwareDownload	-	-	5.3	I
Secure File Copy (SCP) for supportSave	-	-	5.3	I
SecTelnet	2.6	3.1	4.1	I
SFTP	-	-	7.0	I
Telnet disable (IPfilters from FOS 5.3)	-	-	4.4	I
Telnet timeout	2.6	3.1	4.1	B
Web Tools timeout	-	-	6.2	B
Secure passwords (centralized control via RADIUS/CHAP)	-	3.2	4.4	A
RSA RADIUS Server	-	-	6.1	A
RADIUS password expiration	-	-	6.2	A
RADIUS source IP address information	-	-	6.2	A
LDAP	-	-	6.0	A
LDAP in FIPS mode	-	-	6.1	A
SLDAP	-	-	6.1	A
Multiple User Accounts (MUA – up to 15)	-	3.2	4.4	I
Multiple User Accounts (MUA – up to 255)	-	-	5.2	I
Role Based Access Controls (RBAC) Admin, User, Switch Admin Roles	-	-	5.0.1	I
Operator, Zone Manager, Fabric Admin, Basic Admin Roles (RBAC) added	-	-	5.2	I
Security Admin Role (RBAC) added	-	-	5.3	I

Security Feature	FOS 2.x	FOS 3.x	FOS 4.x+	Security Level
User-defined roles (RBAC) added	-	-	7.0	I
RBC permission violation (message ID : SEC-3047)	-	-	6.0	A
Admin lockout policy	-	-	5.3	I
Boot PROM password reset	-	-	4.1	A
Password hardening policies	-	-	5.1	B
Upfront login in Web Tools	-	-	5.0.1 Default in 5.2	B
Login banner	2.6	3.1	4.1	B
Motd	-	-	7.0	B
Syslog redirection	-	-	6.3	I
Monitor attempted security breaches (via Audit Logging)	-	-	5.2	A
Monitor attempted security breaches (via Fabric Watch – Security Class)	-	-	4.4	A
FC Security Policies - Device Connection Control/Switch Connection Control (DCC/SCC) policies	SFOS ONLY	SFOS ONLY	5.2	A
Management access controls	SFOS ONLY	SFOS ONLY	SFOS ONLY	A
IP Filters (IPF)	-	-	5.3	A
Trusted Switch (FCS) central security management	SFOS ONLY	SFOS ONLY	SFOS ONLY	A
FCS policy (without SFOS)	-	-	5.3	A
AUTH policy	-	-	5.3	
Management Access Controls (SNMP, Telnet, FTP, Serial Port, Front Panel)	SFOS 2.6	SFOS 3.1	SFOS 4.1	A
Zoning	All	All	All	B

Security Feature	FOS 2.x	FOS 3.x	FOS 4.x+	Security Level
Hardware-enforced zoning by WWN and Domain/Port ID	(Port based only)	3.0	4.0	B
Default zoning	-	-	5.1	I
Insistent domain IDs	-	-	4.2	I
RSCN suppression/aggregation	-	3.1	4.1	B
Configurable RSCN suppression by port	-	-	5.0.1	O
Event auditing	-	-	5.2	I
Change tracking	2.4	3.0	4.0	I
Firmware change alerts in Fabric Manager	-	-	4.4	A
E-Port disable (portCfgEPort)	2.6	3.2	4.2	I
Persistent port disable (E/F/FL/Ex/M-Ports)	2.6.1	3.2	4.2	I
Administrative domains	-	-	5.2	A
Virtual fabrics	-	-	6.2	A
Logical Switch/Logical Fabric/Base Fabric/default Fabric (replaces AD)	-	-	6.2	A
IPsec (7500 only)	-	-	5.2	O
IPsec to secure management interfaces	-	-	6.2	O
IPv6	-	-	5.3	O
IPv6 auto-configuration	-	-	6.2	O
IPv6 for IPsec	-	-	6.2	O
Security DB size increased to 1 MB (from 256K)	-	-	6.0	-
FIPS mode (140-2 level 2)	-	-	6.0	A
USB port disable/enable	-	-	6.0	B
Fabric-Based encryption for data-at-rest	-	-	6.1.1_enc	O

Security Feature	FOS 2.x	FOS 3.x	FOS 4.x+	Security Level
Hash authentication of firmware (signed firmware)	-	-	6.1.0	A
Integrated routing	-	-	6.1.1	O
Traffic isolation zones (TI)	-	-	6.0	O
Duplicate WWN Management	-	-	7.0	O

B

Standards Bodies and Other Organizations

FCIA

The Fibre Channel Industry Association (FCIA) is a mutual-benefit, nonprofit international organization of manufacturers, system integrators, developers, vendors, industry professionals, and end users. The FCIA is committed to delivering a broad base of Fibre Channel infrastructure technology to support a wide array of applications in the mass storage and IT-based arenas. FCIA working groups and committees focus on specific aspects of the technology, targeting both vertical and horizontal markets and including data storage, video, networking, and SAN management.

The FCIA is also responsible for managing events such as interoperability testing, such as "plug-tests" held at the University of New Hampshire and Fibre Channel Technology demonstrations at industry events such as SNW (Storage Networking World).

For more information, visit the FCIA Web site at:
http://www.fibrechannel.org

IEEE

The Institute of Electrical and Electronic Engineers (IEEE) has a wide variety of standards developed in relation to security. The IEEE 1619 Security in Storage Working Group (SISWG) develops standards for encrypting storage media for data-at-rest. SISWG has developed standards for disk-drive-based encryption (IEEE 1619) and tape-based encryption (IEEE 1619.1). SISWG operates as a project under the IEEE Computer Society Information Assurance Standards Committee.

For more information on SISWG, visit the SISWG website at:
https://siswg.net/

ANSI T11

The American National Standards Institute (ANSI) is the voice of US standards and its conformity assessment system and was formally recognized as such in 1970.

T11 is the ANSI technical committee defining the Fibre Channel protocols and physical layer. Fibre Channel Security Protocol (FC-SP) defined methods of authorizing, authentication, and encrypting Fibre Channel interfaces for a fabric. To claim compliance with FC-SP, devices need to support authentication via Diffie Hellman-Challenge Handshake Authentication Protocol (DH-CHAP). DH-CHAP is a mutual authentication between end devices and switches. Fibre Channel Framing and Signaling 2 (FC-FS-2) defined the structure of the Fibre Channel frame that conveys the Encapsulating Security Payload (ESP) header as defined in Request for Comments (RFC) 4303. For more information on T11.

For more information, visit the ANSI T11 website at:
http://www.t11.org/index.html

SNIA

The Storage Network Industry Association (SNIA) is a not-for-profit organization which was incorporated in 1997, and although it is not directly involved in the development of standards, it acts as a catalyst for the development of storage solution specifications, the development of storage solution specifications and technologies, global standards, and storage education. It is composed of individuals representing member companies that work together to further advance the storage industry.

For more information, visit the SNIA website at:
http://www.snia.org/home

SNIA also has various technical work groups and forums addressing specific areas of storage. The Storage Security Industry Forum specifically focuses on issues concerned with storage security. This forum has created several valuable documents with the help of various industry contributors.

For more information, visit the SSIF website at:
http://www.snia.org/forums/ssif

The technical work groups support the SNIA mission by delivering information and standards that accelerate the adoption of storage networking.

Specifically, the SNIA Security Technical Work Group (TWG) helps drive some of the standards addressing storage security issues. Its focus is not only with Fibre Channel security but with any security inherent in underlying transports or technologies.

For more information, visit the SNIA TWG website at:
http://www.snia.org/tech_activities/workgroups/

IETF

The Internet Engineering Task Force (IETF) has the large job of securing the Internet. The Security Area of the IETF defines security protocols for a variety of techniques to authorize, authenticate, encrypt, and manage various aspects of data exchanges. From Public Key Infrastructure (X.509) to Mail Security (S/MIME), the IETF addresses many aspects of security.

For more information on security in the IETF, visit:
http://trac.tools.ietf.org/area/sec/trac/wiki

OASIS

The Organization for the Advancement of Structured Information Standards (OASIS) is a consortium that drives the development, convergence and adoption of open standards. OASIS has developed a number of security-related standards for identity management, key management and web service security. The Key Management Interoperability Protocol (KMIP) defines an interface between encryption devices that consume keys and the key management system that manages the keys.

For more information, visit the OASIS website at:
http://www.oasis-open.org/home/index.php

Index

SECURING FIBRE CHANNEL FABRICS

SECOND EDITION

Although Storage Area Network (SAN) security is a specialized field dealing with issues specific to the storage industry, it follows established principles found in other IT areas. This book is primarily intended to raise awareness of the need for SAN security and attempts to bridge the knowledge and cultural gap between the storage and security groups within an organization. The basic SAN security principles introduced can be applied to any corporate storage environment—which typically includes technology from multiple vendors.

ROGER BOUCHARD

$49.95

Brocade Bookshelf
www.brocade.com/bookshelf

ISBN 978-1-61539-277-3
90000>
9 781615 392773

BROCADE