Global Information Assurance Certification Paper

Fibre Channel Storage Area Networks: an analysis from a security perspective

GSEC Gold Certification

Author: José Picó, pico_jose@telefonica.net

Adviser: Joey Niem

Accepted: March 14th 2006

© SANS Institute 2008, Author retains full rights.

Outline

1 Introduction

When I had to decide what the subject of my GIAC GSEC GOLD report was going to be, it took me about 1 or 2 µsec to decide. Coming from the IT storage world, it seems obvious that the amount of resources spent by IT administrators in securing their IP network, and basically their IP perimeter, has nothing to do with the same budget for storage networks. That also means that the technical work that has been made around this topic is extremely low compared to IP security topics. That's why I decided to invest my time in studying what is the real threat about Fibre Channel storage area networks.

Another reason that pushed me in this direction was that being SANS a community of IP security experts, I saw quite probable that such people were interested in something "new", something not heard before and that can be interesting to them. Even if these people never get involved in any storage department, I hope that the reading or just the reviewing of this paper is of interest to them.

The potential scope of the work is very wide, because it should cover a lot of security areas. The paper started as a general awareness report about what SAN security threats and what mitigation techniques exist. This work seemed to be the most appropriated for a GSEC GOLD paper. However, as I was investigating, my curiosity brought me to the place where I

found less documentation: what can be done using only the Fibre Channel vulnerabilities, and what can be done in order to protect the SAN against it. That's why the scope of the paper is finally limited to the Fibre Channel protocol threats only, but analyzed to a greater level of technical detail.

Thus, the following topics are out of the scope of this paper. However, they should also be considered seriously when protecting a SAN:

- IP management interfaces attacks (switch interfaces, management software interfaces, etc.)
- Attacks to IP SANs (iSCSI, FCiP, iFCP, NAS)
- Those attacks where a hardware traffic analyzer physically connected to the SAN is needed are also physical security related and not considered in this paper, including FC man-in-the-middle attack [27] and FC session hijacking [27].
- Attacks against data protection management infrastructure, like backup systems, third copy systems, business continuity infrastructure, etc. These attacks are extremely dangerous because eliminate the option to be recovered from a data corruption attack; an attacker will for sure address them. They are

out of the scope of this report, but extremely important.

- Attacks based on upper-level-protocol weaknesses haven't been studied in this paper.

If the reader is interested in general approaches to SAN security, he can find a lot of work related to these attacks; as an example, the following references can be consulted among others [27] [44] [9] [18] [2] [45].

2 Audience

There are two profiles that can have some kind of interest in reading this report:

- IT Storage Administrators or managers: I hope that this report can help them to evaluate the real risks that they are facing, and to implement the measures suggested in this paper or other measures that could have been missed.

- IP Security professionals: if they're going to get involved in securing a storage network, this report can be a good starting point for them; if they're not, I think that they will enjoy the reading of the problems that "their neighbors" are going to face.

Concepts that are well known for security professionals –what spoofing is, what denial of service attacks are, etc. - won't be explained. Basic knowledge about these topics is assumed to the reader.

I'm also assuming some basic knowledge about SANs; although an introduction about SANs can be found in chapters 4 and 5, some concepts will be explained very concisely, due to length restrictions.

3 Objectives

The objective of this work is simple: analyze what can be done to attack a storage area network using the Fibre Channel protocol weaknesses, evaluate what are the real possibilities to perform such attacks and study what should have been done to avoid the associated threats and to mitigate the risks.

The main problem that I faced when trying to achieve this objective is that this is only a theoretical study: Fibre Channel hardware is quite expensive. Also, trying to put in practice the theoretical threats would have been too time consuming for me. Even if someone would do that, perhaps the results would not be fully valid, because the differences on the protocol implementation between HBA, switches and storage device vendors are significant. That's why if the reader finds any technical error regarding Fibre Channel protocol or especially regarding the implementation possibilities of the attacks described, any suggestion or correction will be very much appreciated.

Although the reader, through the walk across these pages, will follow the steps of an attacker trying to get unauthorized access to the information, it is not my intention to provide a "hacking the SAN" guide, but to prevent that this type of hacking would occur by pointing out what mitigating techniques should be used. That's why the reader will also follow the missed

actions that a storage administrator could have performed in order to protect against each step of the attack. The execution of some of the attacks described requires deep expertise in low-level network programming as well as in the Fibre Channel protocol. As of today, I haven't heard that such a tool has been developed, but when the first tool is distributed over the internet, the risk of the attack to be performed will immediately raise [25]. The consequences for the target company of such attack could be tremendous and I hope that reading this paper will make IT professionals think about it and change the way they configure their storage networks.

Note that all device addresses and names used in this document are fictitious or have been modified to preserve confidentially of real devices.

4 Description of a Fibre Channel SAN

4.1 What is a SAN

A SAN (Storage Area Network) is a term that defines all the Hardware and Software infrastructure that allows a computer to access storage that is not directly attached to it; any storage device connected to the network can be seen –if so configured- from any computer on it.

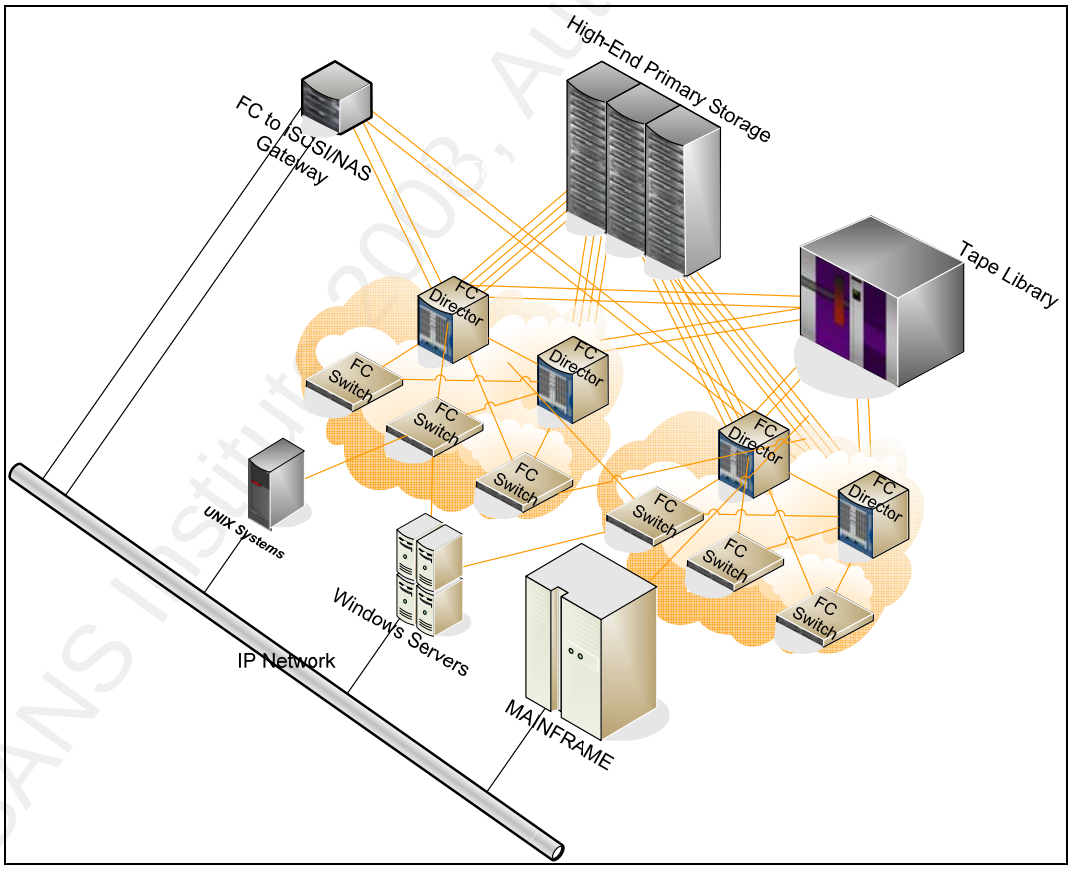

Picture 1: Logical view of a typical SAN

Normally –and traditionally- the back-end mass storage is accessed at block level. The protocol that is working for a long time to access storage at block level is SCSI. This protocol is widely implemented in almost every operating system and it was designed for devices physically connected to the host through an electrical bus. In order to transport such a protocol, Fibre Channel protocol is used. Fibre Channel can be configured in a switched topology, thus becoming a network. Although the most used Fibre Channel encapsulated protocol is SCSI, you can also find IP over FC, FICON or others...

It is still possible to access native SCSI devices through Fibre Channel using FC-SCSI bridges. You may also find elements attached or embedded into a SAN that can build bridges between Fibre Channel and IP networks, basically NAS servers and iSCSI bridges.

Another important aspect of SANs is that they facilitate the implementation of remote replication functionality. This may be done through an extended Fibre Channel network using WDM (Wavelength Division Multiplexing), encapsulating Fibre Channel traffic over existing IP WAN networks (FCIP) or mapping IP address to individual FC devices through iFCP protocol. This point obviously has security implications that will be considered as out of the scope of this report.

Lately, there are some other engines that can be integrated into the SAN and provide storage virtualization services. These engines are starting to be used by customers that have

a large amount of storage and need to perform some special operations between their sites or that have heterogeneous storage that has to be used as single instance to perform operations such us mirroring, 3rd copies or SnapShots.

4.2 Practical concepts and typical tasks when working with SANs

SAN infrastructure, compared to IP networks, is extremely simple to manage. That is a good thing and a bad thing. It is a good thing for administrators because the setup of a SAN is reduced to a set of *relatively* simple tasks that can be performed by themselves almost entirely. However, it is also a bad thing because this simplicity hides the need for the administrator to know Fibre Channel protocol at the level that normally IP network administrators do. This implies that security configuration tasks are normally not considered when setting up a SAN. In my experience, I've seen very few storage administrators that have implemented SAN security measures.

In order to make the reader able to understand the attacking scenarios, the following sections explain what are the typical tasks that a storage administrator usually performs to configure the SAN and to provision storage to hosts.

4.2.1 Aliasing and Zoning

When a Fabric is first started, every device can establish a connection with any other

device connected to the fabric. This is not good for many reasons: security, host storage management, availability, configuration errors risk, etc. The solution for that is a mechanism provided by the fabric, where you can configure what FC devices can see what other devices. This mechanism is called zoning and it is a similar to IP VLANs. All the devices configured in such “zone” can see each other. The rest of the devices connected to the zone that are not included in the zone are not visible by the ones inside and vice versa.

As managing world wide names is quite complicated (a world wide name looks like this: 22:22:00:00:C9:12:34:56), Fibre Channel switches allow to assign symbolic names to World Wide Names in order to make it easier to manage, assign and recognize (a symbolic name might look like this: HOST1_HBA1). This mechanism is called aliasing and is widely used for configuring zones.

4.2.2 LUN Masking

Storage devices have a lot of physical devices in its back-end. In order to provide redundancy, protection and additional functionality, the physical devices are not directly presented to the fabric. Instead, a logical unit is presented which layout is a transformation of a set of physical devices blocks. For example, two mirrored physical disks can be presented as a logical unit.

One storage device has several ports, and it allows a LUN to be presented to the fabric simultaneously by several of these ports, in order to provide multipathing. A FC port can also be zoned to several HBAs (Host Bus Adapters, the Fibre Channel interface cards of the host). That is because the number of host ports (initiators) in a SAN is usually bigger than the number of storage ports (targets). Thus, it is necessary a mechanism that filter the LUNs that a host port can see from a particular storage port. Such mechanism is called LUN Masking and is normally performed at the storage subsystem. By implementing LUN Masking the storage processor filters the access to one specific LUN based on the source address/name of the FC device that is trying to access it. This source address is usually the World Wide Name although it can be the Port ID depending on the vendor (these two concepts will be explained in chapter 5.4).

4.2.3 HBA Masking

Like storage processors, HBAs can implement a similar mechanism to LUN masking. This is mainly used by administrators to filter the access only to the desired LUNs when no LUN Masking has been performed in the Storage Device.

From a security perspective, this mechanism loses its effect in the moment that someone gains root access to the server, since he can modify the configuration file of the HBA driver, allowing access to unauthorized LUNs.

5 Key concepts of the Fibre Channel protocol

5.1 Fibre Channel levels

The Fibre Channel protocol is defined by a set of stackered layers, as IP or OSI protocols are. Here it is a brief description of each level:

- **FC-0: Physical**. This level defines all the physical specifications of the protocol: connectors, cables transmitters, receivers, signal specification, etc. Let's only remark that Fibre Channel is, from this level's perspective, a serial protocol.

- **FC-1: Encoding**. The Fibre Channel protocol encodes 8-bit character into 10-bit characters for transmission efficiency purposes.

- **FC-2: Framing and Signaling**. This is the level where transport protocol is specified. It contains the transmission structure, frame definition, procedures, etc.

- **FC-3: Extended Link Services**. This level implements some services used by the FC devices to obtain and provide information to the FC network infrastructure itself and to other FC devices.

- **FC-4: Upper Level Protocols (ULP)**. This level defines how upper level protocols have to be mapped with Fibre Channel, defining mapping entities, etc. This is the level where FCP (SCSI over Fibre Channel) is defined.

5.2 Fibre Channel Topologies

Fibre Channel can be implemented on some topologies:

- Point-to-point topology: consists of a FC port directly connected to another FC port. The TX connector of one node is directly connected to the RX connector of the remote node, resulting in one FC cable connecting both nodes.

- Arbitrated-loop-topology: all the nodes are connected to a common media, which is physically shared for all FC Ports.

- Switched Topology: it is similar to an IP switched network, where HBA's connect to Storage Devices or Controllers through devices that allow simultaneous traffic between its ports. Each group of interconnected switches is called a fabric and defines a unique naming, zoning and addressing space within a Fibre Channel network.

- Public loop topology: it is an arbitrated loop connected to a fabric, where the

devices in the loop can be accessed from those in the fabric.

We will consider only Switched Topology, because is the common one in current production SANs, and also because it is the most probable attack scenario that we must consider.

5.3 Fibre Channel entities: nodes and ports.

A Fibre Channel port, as defined in the FC-FS-2 standard [4], is “a Hardware Entity that has a Link Control Facility that is compatible with Fibre Channel Protocol”. In other words, hardware that has a port where you can connect a Fibre Channel cable and observes the Fibre Channel Level 0, Level 1 and Level 2 specifications, providing their defined functions. A node, as defined in that standard is “a collection of one or more Nx_Ports controlled by a level above FC-2”. In other words, it is an entity that has one or more FC-Ports. The following diagram illustrates this concept [4].

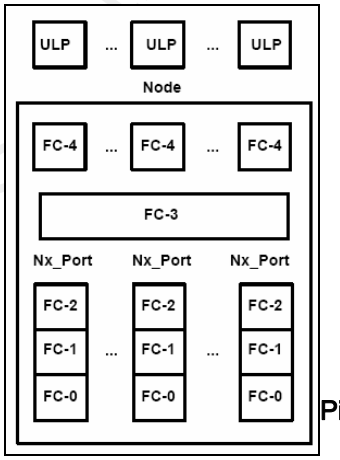

Picture 2: Fibre Channel Node and Port

There are many types of Fibre Channel ports. The following table summarizes the main Fibre Channel port types and its purpose [38]:

Type	Description
E_Port	An E_Port is an expansion port. A port is designated an E_Port when it is used as an interswitch expansion port to connect to the E_Port of another switch, to build a larger switched fabric. These ports are found in Fibre Channel switched fabrics and are used to interconnect the individual switch or routing elements. They are not the source or destination of IUs, but instead function like the F_Ports and FL_Ports to relay the IUs from one switch or routing elements to another. E_Ports can only attach to other E_Ports. An Isolated E_Port is a port that is online but not operational between switches due to overlapping domain ID or nonidentical parameters.
F_Port	An F_Port is a fabric port that is not loop capable. Used to connect an N_Port to a switch. These ports are found in Fibre Channel switched fabrics. They are not the source or destination of IUs, but instead function only as a "middle-man" to relay the IUs from the sender to the receiver. F_Ports can only be attached to N_Ports.
FL_Port	An FL_Port is a fabric port that is loop capable. Used to connect NL_Ports to the switch in a loop configuration. These ports are just like the F_Ports described above, except that they connect to an FC-AL topology. FL_Ports can only attach to NL_Ports.
G_Port	A G_Port is a generic port that can operate as either an E_Port or an F_Port. A port is defined as a G_Port when it is not yet connected or has not yet assumed a specific function in the fabric.
L_Port	An L_Port is a loop capable fabric port or node. This is a basic port in a Fibre Channel Arbitrated Loop (FC-AL) topology. If an N_Port is operating on a loop it is referred to as an NL_Port. If a fabric port is on a loop it is known as an FL_Port. To draw the distinction, throughout this book we will always qualify L_Ports as either NL_Ports or FL_Ports.
N_Port	N_Port is a node port that is not loop capable. Used to connect an equipment port to the fabric. These ports are found in Fibre Channel nodes, which are defined to be the source or destination of information units (IU). I/O devices and host systems interconnected in point-to-point or switched topologies use N_Ports for their connection. N_Ports can only attach to other N_Ports or to F_Ports.
NL_Port	An NL_Port is a node port that is loop capable. Used to connect an equipment port to the fabric in a loop configuration through an FL_Port. These ports are just like the N_Port described above, except that they connect to a Fibre Channel Arbitrated Loop (FC-AL) topology. NL_Ports can only attach to other NL_Ports or to FL_Ports.
U_Port	U_Port is a universal port. A generic switch port that can operate as either an E_Port, F_Port, or FL_Port. A port is defined as a U_Port when it is not connected or has not yet assumed a specific function in the fabric.

5.4 Fibre Channel Addressing

5.4.1 Fibre Channel Address

In the context of a fabric, the Fibre Channel address of a FC port is a 24-bit number that identifies a FC device connected to it. The switches use these addresses as an index into its internal routing tables, in order to route the FC frames from one device to another by passing the frame to the next switch in the shortest path (see chapter 8.2).

The meaning of these three bytes is described in the following table [8]:

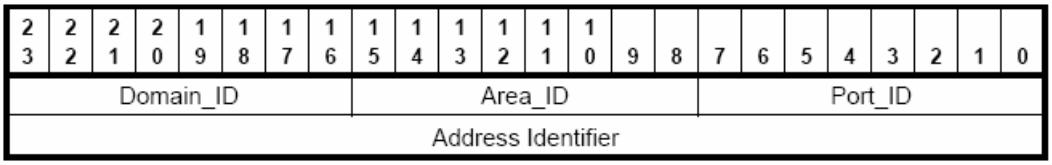

Picture 3: Fibre Channel Address structure

- The "Domain ID" is a number that identifies one or more switches in the same domain of a fabric. Usually, the domain id is the number of a switch inside a fabric, but if there is more than one, they should be directly connected through a direct attached cable called an Inter Switch Link (ISL).

- The "Area ID" identifies an Nx_Port that is connected to a switch. If the device connected to a switch is not an Nx_Port, but an arbitrated loop, then this number identifies the whole public loop.

- The "Port ID" identifies a single Nx_Port (except for those reserved) or the AL-PA (Arbitrated Loop Physical Address) of a device port connected to a public loop in that port of the fabric.

Some of the identifiers can't be used because they are reserved by the fabric to provide access to some of its services.

5.4.2 World Wide Names

In order to uniquely identify a device (Node or Port) connected to a Fibre Channel infrastructure, the device has a static address or "name" that is unique to it and is called the World Wide Name (WWN). There is one name associated to the FC Node (World Wide Node Name –WWNN) and one associated to each FC Port(s) of the node (WWPN). The WWN is static and configured on the device; this concept is quite similar to the MAC address of an Ethernet device.

The World Wide Name is a 64-bit structure, usually represented by 8 pairs of hexadecimal digits in the form HH:HH:HH:HH:HH:HH:HH:HH. It has a particular structure that reveals information about the device [4] [33].

5.5 Fibre Channel services

Fibre Channel fabrics provide a set of services to its clients (i.e. the Fibre Channel nodes), to allow them to interact with the storage network to exchange information like connection state, connection parameters, configuration or topology changes, etc.. These services can be accessed by login into ports that hold a Well Known Address (WKA). WKA are port FC IDs that are reserved for internal use of the fabric, usually fabric services.

The following table enumerates the special ports and the service that the fabric provides through it [4]:

Address Value	Description
FF FC 01h to FF FC FEh	Reserved for Domain Controllers
FF FF F0h	N_Port Controller (see FC-LS)
FF FF F1h to FF FF F3h	Reserved
FF FF F4h	Event Service (see FC-GS-5)
FF FF F5h	Multicast Server (see 22.3.4)
FF FF F6h	Clock Synchronization Service (see clause 25)
FF FF F7h	Security Key Distribution Service (see FC-GS-5)
FF FF F8h	Alias Server (see 23.2)
FF FF F9h	Quality of Service Facilitator - Class 4 (QoSF) - Obsolete
FF FF FAh	Management Service (see FC-GS-5)
FF FF FBh	Time Service (see FC-GS-5)
FF FF FCh	Directory Service (see FC-GS-5)
FF FF FDh	Fabric Controller
FF FF FEh	F_Port Controller
FF FF FFh	Broadcast Alias_ID (see 22.4)

Picture 4: Fabric Well Known Addresses

Each service can have one or more sub-services. Not all services are available in all fabrics; it basically depends on the switch vendor and version that is implementing a particular fabric. Usually switch vendors provide information about what services are available on their products (see [15] as an example).

6 The attacking scenario

As we are considering analyzing only Fibre Channel attacks, the first thing that we must study is under what scenarios such attack would be possible.

When talking about Fibre Channel security the most common mistake that I've found is to assume that Fibre Channel is secure simply because the Fibre Channel ports are always inside the Data Center so you cannot directly plug your computer to a Fibre Channel network. Although essentially true, this argument is seriously incomplete. Usually in enterprise organizations, there are some very well protected servers (back-end production servers) connected to the SAN that very few people has direct access to and that are very difficult to compromise. But there are also a lot of non-production servers (development, test or preproduction environments, DMZ servers [17], etc.) that also have access to the SAN. It is normal to find that hundreds of people have an account on them; many of them even don't belong to the company, but to a contractor. Usually, the security protection of these servers is not as strong as the production ones -many times the root/administrator password of these servers is "root" or similar, the security perimeter around these severs is not as strong as the production ones, etc.-. So our starting point is that an attacker has compromised such a system, probably first with a non-privileged user and then he has switched to a user with administrator rights. We'll not cover how this is possible. Probably the reader has a lot of

expertise on that matter. We're not considering that the attacker has compromised one of the back-end servers, because that would be an even better case for the attacker; he would already have access to the company's core information which is his objective.

The system will probably have two FC HBA's with some kind of path failover and load balancing driver on top of the Fibre Channel driver. Also, the system is probably configured with monitoring capabilities that will detect an HBA failure. This is something that an attacker should take into account in order to not being detected.

I've seen this scenario quite a few times among customers using SANs. Very often, the development/test environments not only have access to the same SAN as production ones, but also they have access to the same disk arrays as the production ones (sometimes this is even a requirement of the production environment due to the need of advanced data replication techniques used for development/preproduction environment refreshment or activities like data warehouse extraction).

Once the attacker has gained root access to one of these systems, the attacker will need some more things to perform its attack. Depending on the attack he wants to perform, he will need some of the following:

- The tools provided by the OS and/or the HBA driver or other commercial tools. This kind of tools can be used in most cases, although they offer limited

functionalities, depending on the vendor. Just to name a few examples: HP-UX HBA's have fcmsutil [35], QLogic HBA's have SANSurfer application [43], and Emulex HBA's have lpfc and HBAnyware applications [29]. There are other management tools that make use of some of the techniques explained in this report for management purposes [47]. During the analysis of the attacks we'll make reference to some of these tools to illustrate the concepts explained.

- A self-developed tool that, as an application, makes use of the installed driver interface to access the fabric in a special way, constructing customized Fibre Channel frames and sequences (we will call "*fctool*" to such kind of tool from now on)
- Manually modify driver configuration data or HBA internal configuration data, mainly to perform WWN Spoofing as explained in chapter 8.
- In some cases it would be needed to install a modified version of the HBA FC Driver. In this case the attacker would prefer a Linux machine, because the driver source code is usually available, thus easiest to modify. I'm not aware of the existence of such a tool yet, and writing it would be a difficult task, but not impossible.

In this document we'll try to consider the difficulty grade of every attack which. among other factor, depends on the difficulty to develop the needed tools.

Finally, it is important to be aware of the attacker's motivations. This kind of attack is not "casual", but very badly intentioned: the attacker's objective is to steal, corrupt or destroy the company's core information. The attacker has planned his actions for a long time and he knows the consequences of what he's doing. So if a company suffers such an attack, the objective pursued by the attacker should always be considered as serious.

7 Enumeration phase: gathering target information

7.1 Attacker's goal

When launching an attack against a company's information through the Fibre Channel network, the first thing the attacker needs to do is to gather information about the actual SAN topology, configuration and most important, the list and type of devices connected to it and how to route to them, i.e. its Nx_Port_ID associated with its WWN's and the type of device it represents [27]. Not having all the needed information will make much difficult the attacker's success. An example table is presented to illustrate the data needed:

Type	Pid	WWPN	WWNN	COS	PortSymb	Fabric Port Name
N	030500	50:06:0b:01:23:45:67:89	50:06:0b:01:23:45:67:89	3		20:05:00:60:69:01:23:45
N	030B00	50:06:0b:01:23:45:67:89	50:06:0b:01:23:45:67:89	3		20:0B:00:60:69:01:23:45
N	032500	10:00:00:00:c9:01:23:45	20:00:00:00:c9:01:23:45	2	[] "Emulex LP9802 FV1.90A4 DV HOST1 v5-2.23a6 "	20:25:00:60:69:01:23:45
N	033800	50:06:04:81:23:45:67:89	50:06:04:81:23:45:67:89	3	[] "EMC SYMMETRIX SnSnSnSnSnSn SAF- 2b EMUL xxx 5670_083+ "	20:38:00:60:69:01:23:45
N	034500	50:06:01:61:23:45:67:89	50:06:01:61:23:45:67:89	3		20:45:00:60:69:01:23:45
N	034A00	50:06:04:81:23:45:67:89	50:06:04:81:23:45:67:89	3	[] "EMC SYMMETRIX SnSnSnSnSnSn SAF-44c EMUL xxx 5670_083+ "	20:4A:00:60:69:01:23:45
N	034F00	50:06:01:61:23:45:67:89	50:06:01:61:23:45:67:89	3		20:4F:00:60:69:01:23:45
N	035700	20:00:00:e0:8b:01:23:45	20:00:00:e0:8b:01:23:45	3		20:57:00:60:69:01:23:45
N	037100	50:06:01:61:23:45:67:89	50:06:01:61:23:45:67:89	3		20:71:00:60:69:01:23:45
N	037500	10:00:00:e0:02:01:23:45	10:00:00:e0:02:01:23:45	3		20:75:00:60:69:01:23:45
N	037700	50:06:04:81:23:45:67:89	50:06:04:81:23:45:67:89	3	[] "EMC SYMMETRIX SnSnSnSnSnSn SAF- 18b EMUL xxx 5670_083+ "	20:77:00:60:69:01:23:45
N	037E00	20:00:00:e0:8b:01:23:45	20:00:00:e0:8b:01:23:45	3		20:7E:00:60:69:01:23:45
N	039100	50:06:01:61:23:45:67:89	50:06:01:61:23:45:67:89	3		20:91::00:60:6901:23:45
N	039900	50:06:01:61:23:45:67:89	50:06:01:61:23:45:67:89	3		20:99:00:60:69:01:23:45
N	03AB00	10:00:00:00:c9:01:23:45	20:00:00:00:c9:01:23:45	2	[] "Emulex LP9802 FV1.90A4 DV HOST2 v5-2.23a6 "	20:AB:00:60:69:01:23:45
N	03AC00	10:00:00:e0:02:01:23:45	10:00:00:e0:02:01:23:45	3		20:AC:00:60:69:01:23:45

N	03B500	50:06:04:81:23:45:67:89	50:06:04:81:23:45:67:89	3	[] "EMC SYMMETRIX SnSnSnSnSnSn SAF- 18b EMUL xxx 5670_083+ "	20:B5:00:60:69:01:23:45
N	03C700	50:06:0b:01:23:45:67:89	50:06:0b:01:23:45:67:89	3		20:C7:00:60:69:01:23:45
N	03DF00	20:00:00:e0:8b:01:23:45	20:00:00:e0:8b:01:23:45	3		20:DF:00:60:69:01:23:45
N	03E500	10:00:00:00:c9:01:23:45	20:00:00:00:c9:01:23:45	2		20:E5:00:60:69:01:23:45
N	03E700	50:06:04:81:23:45:67:89	50:06:04:81:23:45:67:89	3	[] "EMC SYMMETRIX SnSnSnSnSnSn SAF-44c EMUL xxx 5670_083+ "	20:E7:00:60:69:01:23:45
N	03FA00	50:06:01:61:23:45:67:89	50:06:01:61:23:45:67:89	3		20:FA:00:60:69:01:23:45

Another very interesting thing to know about the Fabric we want to attack is its topology: how many switches are there, how are they interconnected, the zoning scheme and configuration, etc. This information can be obtained through the Fabric Configuration Service (Management Service), defined in FC-GS-5 standard [5].

7.2 Technique description

The first try the attacker would perform is to access the management server, simply because if he succeeds he will get all the information needed!

Fabrics provide a mechanism called management service that is defined in FC-GS-5 standard [5]. Most current fabric implementations are compliant with it [12] [49] [19] [23]. The reason why is that it is extremely useful because it allows in-band management applications to control the fabric and avoid the need to connect switch-by-switch to perform discovery and configuration procedures. This is especially true when talking about big fabrics, due to the operational overhead that these tools can avoid to storage administration teams.

The services offered by the fabric management server as specified in FC-GS-5 [5] are:

- Fabric Configuration Server
- Unzoned name server
- Fabric Zone Server
- Security Policy Server
- Fabric Device Management Interface

Depending on the fabric implementation, some of these management server services are supported or not. Fabric Configuration Server and Unzoned Name Server are two services supported by most fabric implementations [49]. In most cases, they're enabled and accessible by any host by default [12] [23] [49] [47].

Depending on what you want to access, there are commercial applications that can access this service [47] [29] [43]. However, let's assume that the attacker uses its own-developed application "*fctool*", which could be considered as the most dangerous situation.

The attacker is using an HBA that is already logged into the Fabric and has registered its name into the Directory Server. Although these concepts haven't been explained yet (they're explained in chapter 8.2) for now it is enough for the reader to know that an HBA first must login into the Fabric to get its Fibre Channel address and negotiate operating

parameters, and in order to communicate to any other FC node he needs to login into the node also in order to also set up communicating parameters.

The Management Service uses the Well Known Address (WKA) *FFFFFAh* as the FC port address that acts as the fabric interface to communicate with FC nodes. Prior to communicate to Management Server, the attacker must perform the Port Login (PLOGI) operation to the destination address *FFFFFAh*, as show in the picture (see also chapter 8.2 for PLOGI procedure description).

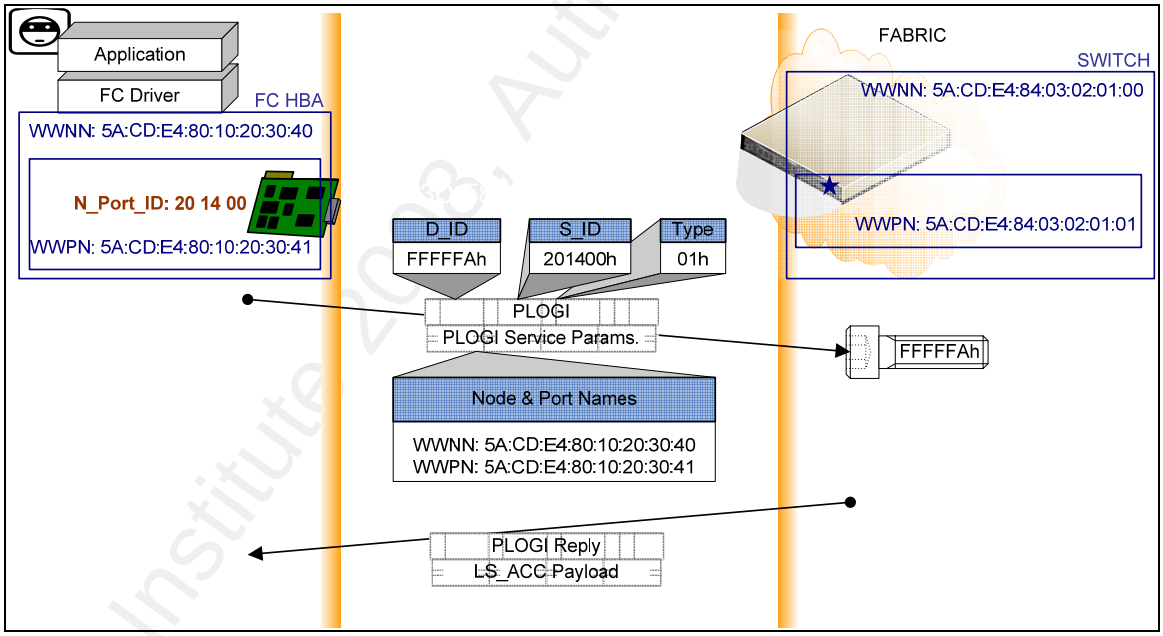

Picture 5: Attacker's application login in Management Services WKA

The management server, as the rest of generic services provided by the fabric, uses a common interface called “Common Transport” that employs an ad-hoc defined Fibre Channel

Sequence to communicate with: the Common Transport Unit (CT_IU). All this definition can be found in FC-GS-5 standard [5]. The attacker, that knows the CT_IU structure, must construct the appropriate CT_IU (Common Transport Information Unit) and sent it to the WKA *FFFFFAh*. The subtype of the CT_IU will determine the service that he's trying to access. The picture shows the access request to the unzoned name server that the attacker's HBA performs after the previous PLOGI to *FFFFFAh*. The attacker obtains the list of all port identifiers of some type of FC-4 defined protocol (typically FCP).

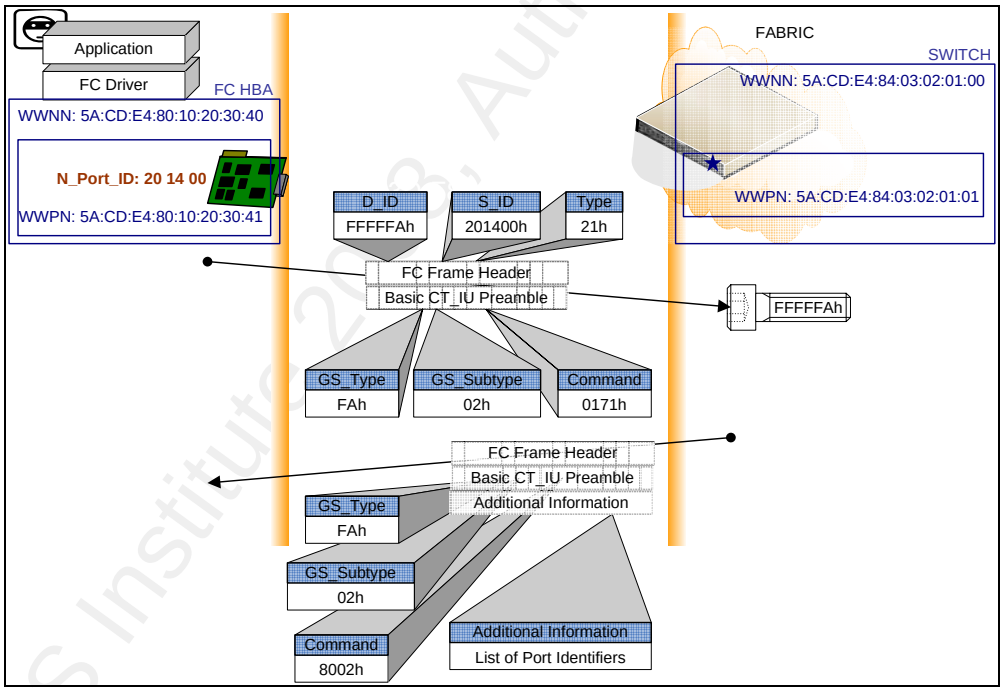

Picture 6: Attacker's application accessing the unzoned name server

Our attacker would continue using the Management Service and query the Fabric Configuration Service in order to obtain information like switch type and configuration, fabric

routing configuration, fabric zoning parameters, etc.

This is the way that some FC management tools can use to control a whole fabric with In-Band philosophy [47], while other tools use the out-band IP method through device's API's. There are two weaknesses regarding this service:

- By default, there is no access control: accessing to Management Service is not restricted to zoning enforcements and no other control mechanism is enforced by default, although switches usually provide that functionality [12] [23] [49] [47] (see next section, chapter 7.3, for related mitigation techniques explanation).
- Lack of authentication if FC-SP is not configured: standard FC-GS-5 includes a recommendation saying "Use of CT Authentication or of other methods to ensure message integrity and authentication (see FC-SP) is recommended for use within the management service". Unfortunately, from my experience, it is very rare to find it implemented nowadays in production fabrics.

If the previous technique fails, because the Management Service is not enabled or is well protected, the attacker still has some chances: an HBA always has access to the fabric name server. The fabric name server is provided by the fabric directory service, which can be accessed through the *WKA FFFFFCh* subtype 0x02. This is the easiest way to obtain name server information, because usually HBA driver tools provide functionality to perform such

queries. However, the information that can be obtained with it is incomplete; if the SAN is zoned, which is the most probable situation, trying to get this information directly from the name server service will result in view of the devices zoned to the current WWN or Port_ID of the HBA that the attacker is using. In our example, if the attacker's HBA FC_ID is *03AB00h* and it is zoned only to storage ports with FC_ID *030900h* and *030D00h*, the result will be the following subset of information:

Type	Pid	WWPN	WWNN	COS	PortSymb	Fabric Port Name
N	034F00	50:06:01:61:23:45:67:89	50:06:01:61:23:45:67:89	3		20:4F:00:60:69:01:23:45
N	037100	50:06:01:61:23:45:67:89	50:06:01:61:23:45:67:89	3		20:71:00:60:69:01:23:45
N	03AB00	10:00:00:00:c9:01:23:45	20:00:00:00:c9:01:23:45	2	[] "Emulex LP9802 FV1.90A4 DV HOST2 v5-2.23a6 "	20:AB:00:60:69:01:23:45

The example shows the most probable situation: if zones are designed with care –i.e. following vendor's best practices [16]-, usually one zone contains only one initiator. That means that from the attacker's HBA only one or more target devices can be seen and no other host's HBA data would be seen. The attacker needs to know data about the hosts that he's going to impersonate afterwards.

Having only the zone-restricted name server table, an attacker still can do some things to try to get more information: he can try to guess what WWN other hosts have based on the WWN's he can see. The first difficulty in this task is the guessing itself. This task is not as unaffordable as it would seem. The number of possibilities for brute force attack is in the order of 16^5= 1048576, due to the WWN format definition (see [4] [10]) and due to the fact that companies usually buy more than one HBA at a time we could start guessing from the

"closer" WWN to ours [27].

However, this task is difficult, costly and detectable. Not mainly because the brute force procedure, but because every try will probably imply a reboot of your machine, because to know if the guessed WWN is valid within the fabric you will probably have to change your HBA WWN. Remember that the attacker machine is a development machine that is being monitored by the systems team and unplanned reboots should be detected. Of course things can be done to hide it (which are out of the scope of this report), but anyway the effort seems detectable and costly.

Much easier to get is the WWN of ports of the storage array or tape library that our attacker's system is connected to. Most vendors use a naming scheme or algorithm that allows the attacker to infer other WWPN's of the system based on the WWPN that he can see [10]. Usually, a storage system has a unique WWNN and the WWPN are numbered from it.

Another technique that can be used is Fabric address port scanning [25]. It will only work if the zoning enforcement type is soft zoning (zoning enforcement types are explained in chapter 8; for now, let's only say that soft zoning will not prevent the attacker's HBA to try a PLOGI to any FC Port in the fabric). This technique consists in sending a PLOGI request to every D_ID you expect to find an Nx_Port and wait for the response.

Depending on the response you can conclude that the D_ID is or is not in the fabric. If successful –which should be the case-, the PLOGI negotiation will also provide you with the WWN's of the devices (among other information); if not, at least you know that a device is at that address.

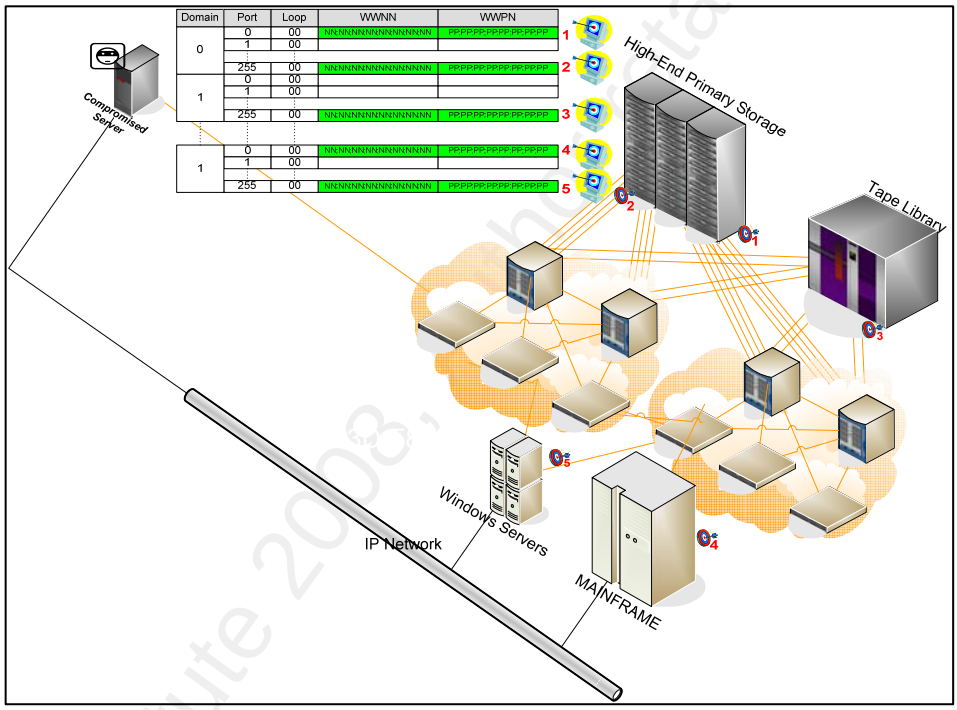

Picture 7: D_ID Port Address Scanning

7.3 Mitigation techniques

The storage administrator should have performed the following actions in order to hide as much as possible its SAN configuration information:

- Study whether the management server is needed or not. First, understand what management server is and what it is used for. Then, identify if the fabric has management applications that use in-band management features. Then, enable management server only when needed. If management server is disabled (through the use of ACL's or by disabling it whenever possible, depending on the switch vendor) the attacker will only have access to the zone-restricted name server information.
- If the management server is needed, configure access control through the use of ACL's in order to configure what WWN's or ports are authorized to access it [28] [50]. Most switch implementations provide a mechanism to filter what WWN's can perform particular management operations through fabric management server. If this is done, the attacker will only have access to the zone-restricted name server information. Even if he would successfully try WWN spoofing techniques (see chapter 8.2), he has less probabilities to succeed, mainly because he doesn't know what WWN to spoof yet.

- Configure CT_IU Authentication [35]: The authentication of the device's WWN trying to get access to the unzoned name service of the management service is defined by FC-GS-5 standard [5] as CT_Authentication and requires the use of Extended CT_IU preamble (see chapter 11.3). Most Fabric implementations include the implementation of FC-SP as a built-in functionality or as an additional feature [12] [49] [23] but, as I've remarked before, unfortunately it is very rare to find it implemented nowadays in production fabrics. If the administrator does this, it will make spoofing techniques much more difficult or even impossible to perform.

- Configure hard zoning (see chapter 8) to avoid Fabric Address port scanning [50] [28] [26] [27] [41] [45] [16] [35] [18].

- Consider the use of Virtual Fabrics [18] [35] [40]: most current switch implementations now support virtual or logical fabrics. That means that you can have more than one fabric inside the same switches. The switch logic maintains the fabric absolutely separated, except for the explicitly configured routes between virtual fabrics. If it is possible to configure Nx_Ports belonging to different environments in different virtual fabrics, then the explained techniques are only dangerous within virtual fabric boundaries, thus reducing risk.

8 Impersonation phase: bypassing zoning

8.1 Attacker's goal

The attacker has succeeded in the previous phase, so he now knows what FC devices are present in the fabric, the fabric topology and zoning configuration. He can now plan what storage subsystems attack, what zone restrictions apply to them and what are the production hosts that are authorized to access it. The goal of the attacker is to clean its way to the front-end Fibre Channel ports of the storage subsystems where the data lives. Accessing these ports at Fibre Channel level 2 is the next goal to achieve.

If the storage subsystem shares its ports between the attacker's machine and the production machines, then this step is not needed. This will probably not be the case; the production machines would likely have dedicated storage ports. In that case zoning will for sure be there preventing the attacker to see the target Fibre Channel ports from its HBA; the attacker's next problem to solve is how to bypass zoning.

8.2 Technique description

First, the attacker needs to know what type of zoning enforcement the fabric is using. Zoning enforcement is usually applied on per-zone basis [12].

If the zoning put in place is soft zoning, then the way is quite easy. By definition, soft zoning enforcement only hides routing information to non-members of the authorized zone [8] [28] [18] [41]. In other words, queries to the fabric name server are zoning restricted as explained in chapter 7.1. Apart from that, nothing will avoid communication between two Nx_Ports, even if they're not in the same zone, provided that you have somehow been able to know the D_ID of the port you want to communicate with [27] [41]. It is that simple. So subverting soft zoning is basically being able to know routing and device information (see chapter 7) thus being able to send frames to that destination D_ID.

Soft zoning is still used in many production fabrics. There are several reasons for that. First, there is a lot of misunderstanding among storage professionals about the differences between hard and soft zoning [27]: there is an extended belief that says that hard zoning is port zoning and soft zoning is WWN zoning, which is totally wrong; there is a lot of literature where you can read that erroneous concept and also switch vendors that only support port-based hard zoning usually tend to associate both concepts in their manuals. Secondly, some management applications don't work appropriately with hard zoning [27]. Also, hard zoning implementation has limitations depending on the vendor and the zone configuration.

If the zoning enforcement is hard zoning, then the attacker should try a couple more things. Hard zoning enforces frame filtering on a frame by frame basis [8] [28] [18] [41], so it

doesn't allow any frame to reach another FC port if it is not permitted in the zoning rules. The way this mechanism is implemented is absolutely vendor-dependent [25].

The first try for the attacker would be to try to modify the zoning configuration. He could do that through IP-based management interfaces of the switch (although out of the scope of this report, storage administrator must secure their fabric's IP management interfaces), or trying to have access to the Fabric Zone Server through the management server [5] as explained in chapter 7.2. This access is going to be possible depending on the switch vendor; some of them support this fabric management server interface while others don't [49].

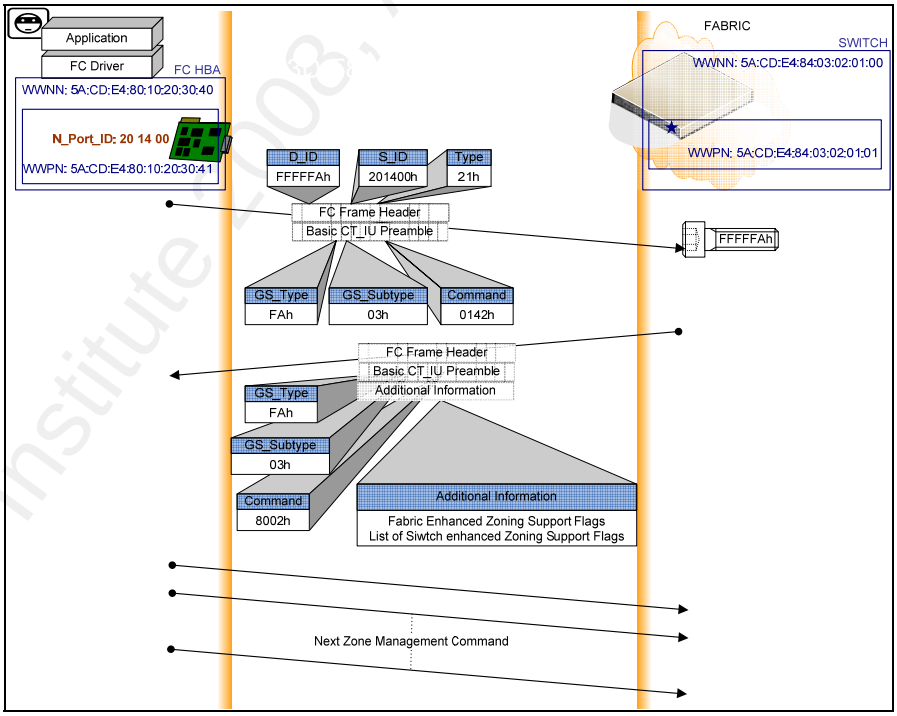

Picture 8: Accessing the Fabric Zone Server

The CT_IU constructed must have the subtype set to 0x03 to specify the Fabric Zone Server. Getting access to this service will provide the attacker full control of the Fabric's Zone Configuration. Our attacker could make all these changes from the application level and grant zone access to its HBA.

It is also very improbable that storage administrators will detect the attack until it is done, because usually zone changes are not included in operational frameworks as an alarm.

If the attacker success in gaining access to this interface, then no other techniques would be needed in order to subvert zoning, because he will simply change the zoning to authorize his HBA. If he doesn't, then he would need to continue working.

He knows that zoning enforcement is hard, i.e., fabric is filtering communication between nodes on a frame-by-frame basis. Depending on what zone membership policy is being used, the attacker will need to use a different technique.

If the type of zone membership is by WWN (node WWN or port WWN), the switch performs the zoning enforcement by looking up into its directory server table to match every frame S_ID and D_ID with the associated WWN's in its name server table. After that, it will check in its zoning tables if both WWN's are allowed to communicate. This is done on a frame-by-frame basis, so that every frame is filtered according to the defined criteria, and it

should be done at wire speed. Although this is the more flexible way to implement hard zoning, not all switch vendors support hard zoning based on WWN.

In order to subvert this kind of zoning, the attacker should use a technique called WWN Spoofing [27] [41] [39] [50] which consists in changing the WWN of the attacker's HBA in order to impersonate one of the authorized production HBA's.

If you read the classic Fibre Channel literature, it is widely accepted that FC devices' Node WWN's (for HBA's and target devices) can be changed [27] [41] [9]. In some devices, depending on the HBA model that is being used, also World Wide Port Name can be changed [34]. In certain switch implementations, the matching of the WWN for zoning enforcement is made for both WWNN and WWPN. In this case, the attacker can change the WWNN and even if the zoning enforcement is made by WWPN the switch will authorize the changed WWNN [27].

We must evaluate how difficult it is for an attacker to perform such change on a host HBA's. In some driver implementations, HBA's driver and utilities provide this functionality [27], but I haven't found this functionality in more recent HBA driver's documentation. The T11 committee has also a working draft standard called FC-HBA that specifies that the HBA driver should not provide such mechanism [3].

If it is not possible to change the WWN through the HBA's driver tools, I've found evidence that it is still possible to do it, at least in some HBAs. Some cards have this information stored somewhere in the host [32], so changing the HBA WWNN should be done by finding the correct registry key or system file and manipulate it using the correct format. Other HBAs have its WWNN and WWPN in its HBA NVRAM. Corrupting the NVRAM with appropriate values and uploading the HBA with the modified values would change the WWN of those HBAs [37].

All these changes will require a reboot of the machine in order to allow the HBA to make effective its new names [27], thus making this attack a little "noisy". The attacker would try it at the point of the day where he expects less monitoring. He would also try to hide its footprints from the system in order to trick storage administrators about the reason of this reboot.

Once done, no WWN hard zoning enforcements will restrict the attacker's HBA to access the Nx_Ports zoned with the spoofed WWN. On the other hand, the name server table will probably be corrupted with two entries with duplicated WWN's, which has some implications that will be discussed in chapter 9.

If the attacker fails, then probably the zone policy membership is configured by Nx_Port_ID. This kind of zoning enforcement does not require any WWN matching; the

switches simply check into their zoning tables that S_ID and D_ID of the frames are authorized to communicate and, as this is hard zoning enforcement, they make it on a frame-by-frame basis. It is not very common to find it implemented in production environments, just because the Nx_Port ID assigned by the fabric usually –depending on the switch implementation- change every time the switch port where the device is attached changes; this fact will force the storage administrator to change all the device-related zoning every HW topology change, which is not usually considered a very flexible approach.

In order to subvert this kind of zoning, Fibre Channel address spoofing techniques can be used [27] [25]. Similar to the previous technique, all the frames sent to the fabric by the HBA could change its S_ID to one that is convenient to the attacker.

The execution of this technique is possible through a modification of the Fibre Channel driver [17]. If the compromised machine is a Linux box, the attacker has access to the source code of the driver and with root privileges he could install a modified version of the driver (see the qla2x00_fabric_login function on qla2x00.c file on [42]). In this modification the attacker could instruct the driver to handle spoofed S_ID on the PLOGI process and to correctly handle the PLOGI response. If, on the other hand, the compromised machine is not a Linux box, the attacker should modify the driver by using disassembly techniques, which is much more difficult, but still possible.

In order to bypass zoning via S_ID spoofing, the attacker must send Fibre Channel frames with the S_ID of the production HBA, but also modify the fabric routing tables in order to get the responses back. This technique is called name server pollution and takes advantage of the Fibre Channel address weakness [27]. When a Fibre Channel device plugs into a Fibre Channel network –or try to get connected to it- it has to perform an operation that is called Fibre Channel Login (FLOGI) [6]. In the FLOGI process, the HBA performs a login request to the Fabric, providing its Name and its service and class capabilities. It performs such request by accessing to the Fabric's WKA *FFFFFEh* (F_Port controller). The following picture shows the normal FLOGI process with the Fabric acceptation (the Fabric also could reject the FLOGI request).

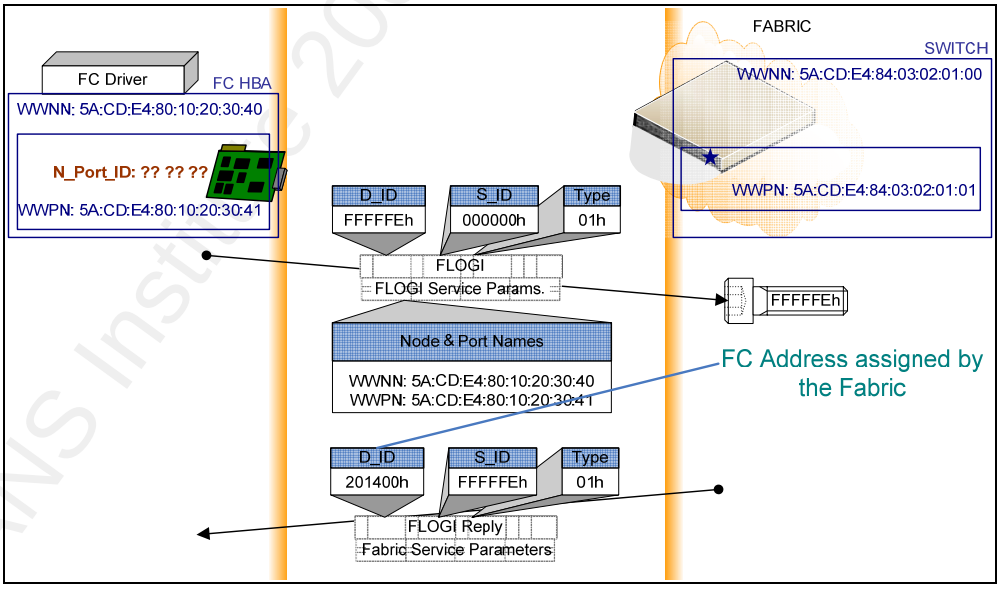

Picture 9: FLOGI procedure example (normal behavior)

Usually, in this request the Fabric is asked to assign a Fabric Channel address to the HBA, although try to get a fixed Fibre Channel address is possible within the protocol, in which case, the Fabric will validate or reject the requested Fibre Channel address.

After the HBA has successfully logged into the Fabric, it must perform a Special Port Login into the Name Server provided by the Fabric's Directory Server [27] [6], in order to register into it its just acquired Fabric address, its name and the information about its service parameters and capabilities, thus making all this information available to the rest of the allowed devices of the Fabric. And here is where the weakness is: the PLOGI information is registered into the name server tables without any validation or authentication process [27] (always considering that FC-SP hasn't been implemented, as explained in chapters 8.3 and 11.1).

The attacker could try to gain benefit from this weakness and, using its modified version of the HBA driver, inject a spoofed S_ID in the PLOGI process in order to make the Directory Server of the fabric have incorrect association between the attacker's HBA device name and S_ID. The S_ID introduced will be the one belonging to the legitimate HBA. The goal of this action is to try to make the switch redirect the legitimate traffic to the attacker's HBA [27].

The following picture illustrates the way to perform such technique.

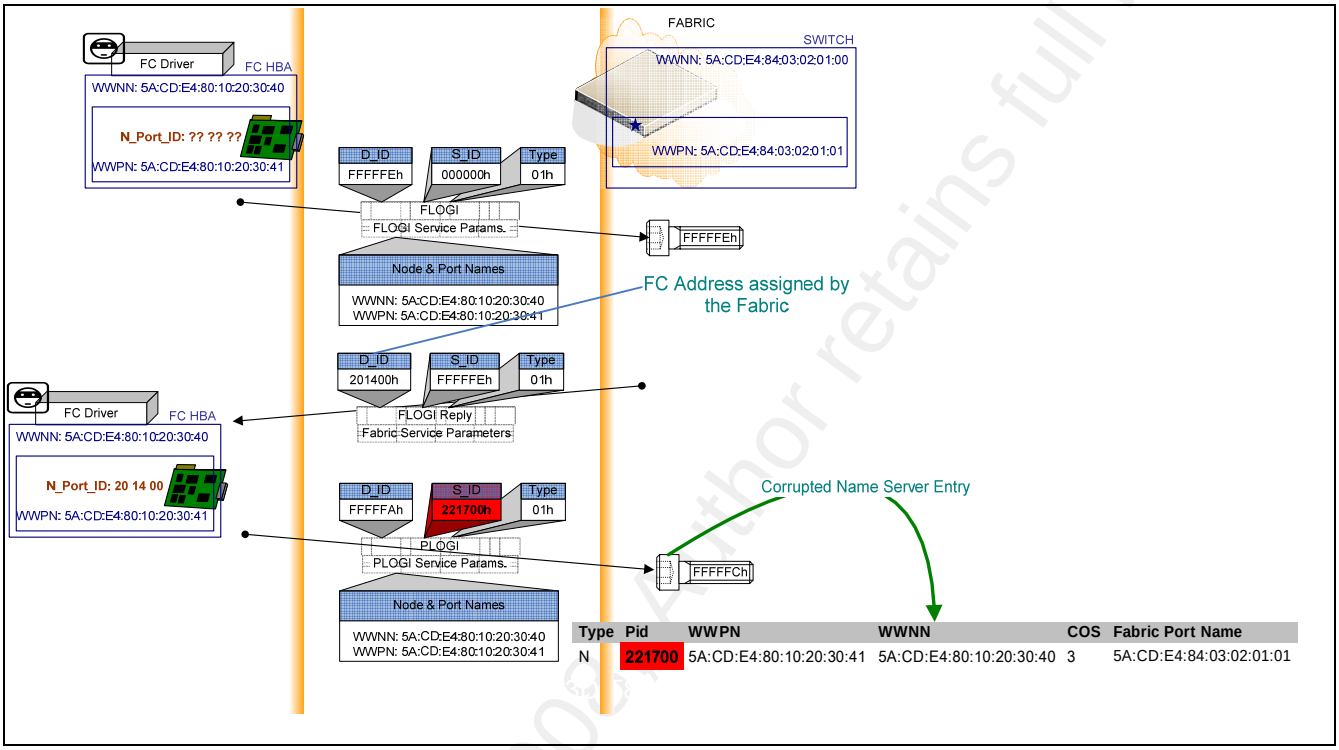

Picture 10: PLOGI procedure against Directory Server (spoofed S_ID)

From this point on, the name server has the polluted information. The consequences of that technique are absolutely up to the switch vendor and configuration [25], because the FC-SW-4 standard [8] doesn't specify the routing tables' contents and/or updates, but only the FSPF exchanges that allow every individual switch to build a FSPF topology map. These exchanges are based on switch Link State Records (LSR) that contains only information about the Domain ID, the neighbor switch and the cost of the link. When a switch routes a

frame, it uses the D_ID as an index into its local routing table: for local entries it should resolve what port index of the local switch is the destination Nx_Port; for remote D_ID the routing table should resolve, based on the Domain_ID contained in the first byte of the D_ID, what is the output port that holds the Inter Switch Link (ISL) for the next hop in the Shortest Path to destination, and pass the frame to the next switch over that ISL. The path selection algorithm is usually Dijkstra's algorithm, but this is not a standard constraint either. What I infer from that is that, unless static routes have been configured in all fabric switches, only HBA's on the same domain as yours are likely to be spoofed, because the routing for domain IDs that are not in the local switch are based on the D_ID's Domain ID.

Another possibility that the attacker could find would be that the membership policy is based on physical ports of the switch. The pair Domain_ID + Physical port number determines the zone membership. In this case, there are only two ways to subvert the zoning [27]:

- Get access to the SAN management interface and change zone settings (this has been already discussed)
- Get physical access to the switch and physically connect our HBA into the right port (out of the scope of this paper)

As you can imagine, this can be considered the most secure way to implement zoning.

But it is also the less used in production environments. The reason why is because this kind of configuration are much less flexible from a production perspective. Change management will become much more complicated for storage administrators, as well as binding and other management tasks will. So if you have a good physical security on your datacenter, and your management IP network is well protected, you usually don't implement this kind of membership policy. Storage administrators always have to perform risk management evaluation in order to make their decision between security and flexibility [27] [41].

When finishing trying to apply these techniques, the attacker can try –using his *fctool*- a PLOGI against the storage subsystem front-end Fibre Channel ports, in order to verify that he has been successful. If he receives the PLOGI confirmation (which is LS_ACC reply Sequence with OX_ID equal to the OX_ID of the PLOGI, and the Common Service N_Port/F_Port bit = 0, among some other parameters [6]) he's sure he's been successful.

8.3 Mitigation techniques

The administrator can minimize the risk of getting a successful zoning bypass attack on its fabric by using the following techniques:

- Fabric zone server configuration: the storage administrator should understand what the fabric zone server is and if its switches support it. Then, decide whether the service is needed or not (normally it is only needed if the company

uses a management software that uses in-band procedures for Fibre Channel zoning provisioning services). If it is not needed, simply disable the access to it by disabling it on the switches interface or by configuring the appropriate ACL's.

- If the fabric zone server is needed, configure the ACL's only to allow the management host's HBA WWN's [28] [50]. If the administrator performs these actions, he avoids the risk of an attacker adapting the fabric zone configuration to grant access to its HBA.

- Consider using Virtual Fabrics [18] [35] [40]: if the storage administrator has configured virtual fabrics then zoning subversion techniques are only dangerous within virtual fabric boundaries. This reduces a lot the associated risk and limits it to the devices in the virtual fabric.

- Hard zoning: whenever possible, use hard zoning because, as explained, soft zoning is like no zoning and thus, much easier to bypass [50] [28] [26] [27] [41] [45] [16] [35] [18].

- Implement FC-SP device authentication [35]: FC-SP offers strong device authentication as explained in chapter 11. By correctly configuring it, WWN Spoofing and S_ID Spoofing techniques won't succeed.

- Port locking [28] [41]: consider the possibility to lock a port to a WWN. This is a feature that most switches offer and consists in associate a WWN to a physical port so that the switch only allows an Nx_Port to communicate if it is connected to the allowed physical port. The storage administrator should evaluate the risk of not implementing that against the loss of configuration flexibility. By doing this, WWN Spoofing is not possible in our scenario.

- Configure the switch, whenever possible, to avoid duplicate WWNs. This is a feature that some switches offer. By doing this, WWN Spoofing is much more difficult, because the attacker should need to bring totally down the spoofed device prior to register the spoofed WWN of your HBA.

- Zone membership configuration: The storage administrator should evaluate the security that physical port membership offers against the loss of configuration flexibility [27] [41]. By configuring zones with physical port membership, WWN Spoofing and S_ID Spoofing won't succeed.

- Monitoring techniques and procedures [28]: it is extremely important to quickly detect the facts that can be the symptoms of this attack, like WWN duplicates in a name server, unexpected LOGIN or LOGOUT of HBA from the fabric, unexpected reboots, etc. Having a well-configured monitoring system and/or

IDS or log correlation system is extremely important in order to detect this phase of an FC attack.

- Change management procedures [28] [35]: implement change control procedures that immediately detect an unauthorized change in the SAN.

9 Denial of service phase: turning off legitimate hosts

9.1 Attacker's goal

In order to reach this phase of the attack, the attacker has already used some of the impersonation technique explained in chapter 8.

If the attacker used HBA WWN spoofing, the storage device providing storage to the host that the attacker has impersonated will see two identical WWNs with different Nx_Port_ID trying to access the same set of LUN's. Most storage devices will not correctly handle this situation, causing both attacker and legitimate devices having intermittent access to the volumes presented by the storage device [27]. The legitimate drive will loose real-time access and the host will probably start reporting I/O errors, causing the upper-level application (like databases, file systems, etc.) to crash. The same will occur to the attacker's HBA, causing both systems to have intermittent access to the storage.

The use of these techniques is by itself a successful denial of service attack against the device that really owns the spoofed WWN or FC address [27]. But this denial of service attack is not the main goal pursuit by the attacker. At this point, the attacker doesn't have full control of the storage and he cannot use the operating system SCSI layer to access it. Although other techniques could be used to communicate to devices without using the host

SCSI stack [27], it is more likely that the attacker wants to access the information by mounting the unauthorized disk into its system using the SCSI layer of the host.

That's why the attacker's goal is now to eliminate as maximum the legitimate HBA access, trying to gain enough time to perform its malicious actions on the storage subsystem information. This phase may not be "mandatory" for the attacker, depending on what he wants to achieve, but I would like to make the reader aware of the associated risks.

9.2 Technique description

One of the link services provided by the fabric is RSCN (Registered State Change Notification). When an event occurs in the fabric, the fabric issues this sequence to all affected ports. The RSCN will be limited to the devices in the same zone of the originating device [48]. The RSCN payload may contain the list of Nx_Port affected and the command codes. Every receiving port -depending on the vendor- will react by performing actions that ensure that the fabric information that they have of the fabric is accurate. It usually can include rediscovery methods such ADISC, PDISC or name server queries [6]. During the rediscovery process no data flow is possible in that device. This is not normally a problem because the frequency of this kind of events is not enough to impact in the real data flow of FC devices. If the frequency of these events increases, then the data flow performance could be impacted or even disrupted, causing I/O disruption [46] and thus the crash of upper level

applications. If the upper level application is a path-failover and load-balancing software, then it is possible that this software disable the I/O traffic over that HBA. Although it won't cause the HBA go offline, the HBA will stop sending traffic until an administrator reviews the problem. If the path failover software doesn't automatically stop traffic on the HBA, it is also possible that, if the attack is continuously repeated, the storage administrator wrongly thinks that the HBA is broken and puts it offline until it is changed. Whatever the case, and depending on the storage subsystem, it could provide the attacker more time of full control of the disks, thus gaining information-access time.

There are many ways to cause an RSCN event. The first one is, using the driver command line utilities, putting offline and back online the attacker's HBA continuously; the HBA should be in the same zone as the target one (see chapter 8). This technique will cause disruption for the attacker's HBA also, so it won't probably be used unless the attacker can use a second HBA's of his system or of other compromised system.

Another technique that can be used here is to, from the application level, issue a lot of explicit RSCN request to the fabric ([6]). In order to make more difficult for the fabric to coalesce events, all the RSCN sent should have different payload contents. This behavior is not recommended in FC-LS although no mechanism has been defined in the standards to prohibit it.

Also, the attacker could send RSCN requests directly to the target Nx_Port indicating a list of affected Nx_Ports in the fabric, causing the targeted HBA to respond to the RSCN requests and disrupting real time IO traffic. The standard specifies that RSCN should only be sent to Nx_Nodes that have registered to receive RSCN requests to the originating node. The most probable situation is that the target HBA has registered to the Fabric Controller WKA *FFFFFDh*. The attacker could try anyway to send RSCN directly to the HBA and see whether it rejects the frames or not, because the reaction of the HBA will be different depending on HBA vendor, so for the attacker, it is worth a try.

Next technique the attacker can try if he doesn't succeed in the previous one is the LOGO attack. I haven't read anything out there about it, but studying the standards, it is theoretically possible.

The explicit Logout ELS allows an Nx_Port to transmit a LOGO ELS to one Nx_Port and specifying that another Nx_Port must be logged out from him FC-LS [6]. In this ELS the affected Nx_Port ID is specified in the payload of the ELS. An attacker could use this feature of a storage subsystem Front End Nx_Port in order to force him to initiate a logout of the legitimate HBA. In this case, all the sequences are terminated abruptly and all the resources associated to this connection are liberated. IO stream and storage LUN access is interrupted and upper level application crashes, with risk of data corruption due to in-flight transactions.

The attacker can make the same thing to the HBA's Nx_Port thus attempting both systems to logout from each other. If the attacker continuously repeats this protocol, it probably will have time to gain full access to the target LUN on the storage subsystem, by making the authorized HBA not bother with its traffic thus having time to accessing the unauthorized information with full control.

A variation of the previous attack could be used by the attacker in the case that S_ID spoofing hasn't been used in previous phases. He could use a modified HBA driver in order to send a FLOGO (a fabric logo) to the F_Port controller. In this case, the source ID should be spoofed to the one belonging to the attacked HBA. When the switch's F_Port receives such a frame, depending on the switch implementation, it won't map the S_ID with the port ID where it comes from, and it will release all the resources associated to the S_ID port and send a LS_ACC to the attacked Nx_Port ID. Any further communication from the S_ID will be rejected because the device is not logged in the fabric.

The first LOGO attack is more efficient than the previous one –because if succeeds it forces the HBA to be logged out from the fabric- but it cannot be repeated continuously, because the HBA is not in the fabric anymore. That means that when the HBA re-login into the fabric, it will probably try to PLOGI into the storage subsystem before the attacker repeat the technique, thus interrupting or bothering the attacker's malicious access to the

information.

9.3 Mitigation techniques

The storage administrator should perform the following actions in order to minimize the risk of the attacks described in this section:

- Understand & Configure RSCN behavior in switches. Nowadays, most switch implementations have contention mechanisms for this issue. Zoning contention of RSCN ELS will apply zoning restriction to the devices zoned with the affected Nx_Port; collating RSCN events is another implementation enhancement that some switch provide in order to group many RSCN events sent to the same Nx_Port into one RSCN ELS. RSCN suppression on some ports is another available mechanism. The reader can find documentation on RSCN behavior of its particular switch regarding the vendor manuals. Here are presented some examples [12] [23].

- Have a FC traffic analyzer and FC troubleshooting tool and team ready: it is very rare to find a pure Fibre Channel protocol-related problem. That's the reason why production storage teams don't usually have a FC analyzing team ready to act. Most switches provide the ability to troubleshoot Fibre Channel

[19] and/or the ability to use a Span port in combination with an external traffic analyzer. Having such infrastructure ready (in place or using the service of a provider) will improve a lot the root cause analysis of a Fibre Channel problem and also of a Fibre Channel denial of service attack. If the attacker is detected in 4 days, he will for sure have achieved his objective. If he's detected in 4 hours, it should have been much more difficult for him. I'm conscious that this is something difficult to justify, because it is a quite expensive measure and it will be used very few times; but the storage network administrator should at least evaluate the need for this and present it to the management.

10 Data access phase: reaching the data

10.1 Attacker's goal

If the attacker has reached this point, it is very close to achieve his goal: steal or destroy important company's data. It can be a bank's user data, a government infrastructure company, or a defense organization. The lack of use of the mitigation techniques explained before as facilitated the attacker's way to the information. The data is now exposed and at risk.

10.2 Technique description

As explained in chapter 4.2.2, "LUN Masking", the storage subsystems usually enforce the access to its LUN's by the WWN of the source device. If the attacker hasn't used it yet, it should use now WWN spoofing in order to allow the storage system to grant production LUN access to the attacker's HBA. Although this is an impersonation technique, I've put it here instead of chapter 8.2 because I consider that, in the attacker's mind, bypassing zoning goes before than this action.

Once LUN masking has been bypassed, the first thing that the attacker can do is to access the LUN he wants to corrupt. Using the native OS tools, the attacker will try to acquire the FC targets, acquire the SCSI targets and, if needed, mount the file systems. During the

time that the legitimate HBA is not operative, the attacker tries to get full access to the LUN's that he sees as disks.

If the attacker wants to steal information, normally he would try to find and filter the information he wants to steal, mainly because the size of the information stored in production systems is enormous and he needs to store the stolen information somewhere away. The attacker could search for the information that he wants to steal by mounting the file system, starting an instance of the production database, etc., and then copying it to some media that he owns.

On the other hand, if the attacker's goal is to destroy information, then a random write at block level into all the production LUN's (using, for example, the Unix dd command) will for sure cause a data corruption that will make the whole data repository not functional. In this case the only way to recover for the administrators is to use on the backup mechanisms put in place. At this point, the RPO (Recovery Point Objective) achieved will basically depend on the investment made by the company in such infrastructure. In addition, it must be taken in account that the attacker has already thought about it and will have tried to make these mechanisms not functional before performing the attack. Although this topic is out of the scope of this paper it is very important to think about it seriously.

10.3 Mitigation techniques

At this point we're assuming that the information can be accessed. At this point we need to protect the information against corruption/loss and/or unauthorized access. There are a couple of things that can be made:

- Data Encryption [26] [2] [41] [45] [40]: there are solutions in the market that provide data encryption for data-at-rest. The main issue with these solutions is performance, but nowadays there are implementations that provide it at almost-wire-speed. Also, there are different choices for the administrator about where to put encryption in the data path (Host, SAN, storage subsystem). Depending on where the encryption is implemented and what authentication techniques are used, the information protection level will be different. In general, data encryption can minimize the risk of information theft when the information can be accessed by an unauthorized user. The storage administrator must evaluate seriously the data encryption solution to choose and discuss with the company's management about the data encryption need, the associated costs and the risks that the lack of this technique has (physical security, unauthorized access, etc.)

- Data protection [45]: every company has data protection infrastructure. We must understand data protection infrastructure as a wide concept that combines

backup techniques, the use of advanced features like third copies or periodic snapshots, data replication across distance, the use of media externalization, etc. In the data protection plan the storage administrator should have agreed the RPO (Recovery Point Objective) and RTO (Recovery Time Objective) with the company's management. Also, the data protection infrastructure must be protected against insider attacks in order to minimize the consequences of an attack as the one described in the previous chapter. Backup server and all the systems that can access the backup index database should be very well protected; all the systems with access to subsystem advance software functionalities must be well protected. There are high-end storage subsystem in the market that even allow the use of ACL's in order to only allow certain hosts to perform certain advanced operations [29]. If the storage administrator has implemented all these recommendations, and the attacker gets to corrupt the data, the data loss will be minimized because the data protection infrastructure will allow the administrator to recover the data within the agreed parameters. If these recommendations haven't been implemented and/or the attacker has also compromised the data protection systems, then the data loss would be likely a disaster.

- Monitoring techniques [28]: in the attacks explained, the attacker's most hard issue is to have time to perform its operations. Early detection of attacker activities will dramatically reduce the attacker chances to succeed. The storage administrator should set up the monitoring infrastructure and correlation events infrastructure in order to detect this kind of behavior: events such zoning changing, unexpected reboots, name server changes, HBA failures, FC traffic analysis are things that I've never found in the mind of the monitoring systems administrators, which are much more worried about availability and/or perimeter security issues. Usually these departments are different from storage administration departments; both teams should work together to detail and define the correct monitoring events, the correct escalation path and the correct isolating actions. Usually, it would be enough to minimally extend the current monitoring systems to include this FC security monitoring. The attacks described in the previous chapters would have never succeeded if the attacks would have been early detected.

- Response teams: as stated, an FC response team should be defined to react to the attacks described. If the cost of this response team is not affordable, at least it should be evaluated to have this service from an external provider that offers a

reasonable SLA to the company regarding response time and ability to perform particular actions like FC traffic analysis.

11 FC-SP: Fibre Channel Security Protocols

As this is the main recommendation that we've made in order to mitigate the risks exposed in this paper, we will dedicate a full chapter about this topic: implement FC-SP [7] in our SAN. From my experience, very few production SAN have been protected using this technique. There are several reasons for that: not all HBAs and switches support this standard, there is no awareness of Fibre Channel risks, no attacks have been reported in the industry, etc. On the other hand, "Execution of this attack requires expertise in low-level network programming and Fibre Channel functionality. However, as soon as someone develops ... software and distributes it on internet, the level of skill necessary will be greatly reduced." [25]

FC-SP provides the following features: device authentication, per message authentication, integrity, protection, anti-reply, secrecy, and fabric management policy set access control [19] [27].

This chapter tries to be an overview of FC-SP functionality and functioning basis.

11.1 Device Authentication

Many of the attacks against FC networks take advantage on the lack of authentication

of the FC protocol (WWN Spoofing, S_ID spoofing, M-I-T-M attack [27], etc.) and others take advantage of the lack of integrity checks (Session Hijacking [27]), etc.

Device authentication is required in switch-to-switch, host-to-switch and host-to-host communications [19]. This method is defined in FC-SP [7] and it is implemented as an extension of the LOGI procedure. This is accomplished by triggering Authentication protocol between node ports that are trying to perform the LOGI sequence. The security bit parameter of the LOGI sequence is used to specify whether authentication is required or not, as shown in the picture [7].

Service Parameter	Word	Bits	PLOGI and PLOGI LS_ACC Parameter applicability			FLOGI Parameter applicability			FLOGI LS_ACC Parameter applicability		
			Class			Class			Class		
			1*	2	3	1*	2	3	1*	2	3
Common Features	1	31-16									
Security Bit	1	21	y	y	y	y	y	y	y	y	y

Key:
y - indicates yes, applicable (i.e. has meaning);
n - indicates no, not applicable (i.e. has no meaning)
* The Class 1 Service Parameters shall be used for Class 6.

Picture 11: Security bit parameter of the LOGI sequence [7]

The transport of the authentication protocol is made through the use of Authentication Transactions specified in [7] within the authentication protocol (AUTH).

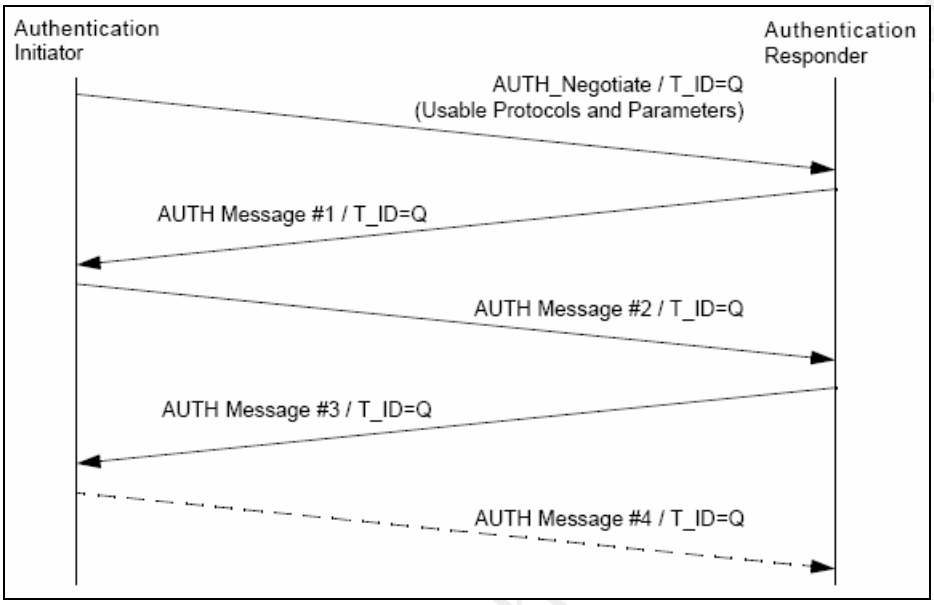

Picture 12: An example of AUTH transaction [7]

The following authentication mechanisms are support by the standard [7]:

- DH-CHAP: Is the only mandatory authentication mechanism within the standard. DH-CHAP can be integrated with a RADIUS server or TCACS+ server in order to centralize the authentication device accounts and passwords. DH-CHAP has some vulnerability: challenge message interval, reflection of CHAP messages and offline dictionary attack [27].

- FCAP (Fibre Channel Authentication Protocol, based on certificates)

- FCPAP (Fibre Channel Password Authentication Protocol, based on SRP –

Secure Remote Password), requiring a password, a salt and a verifier.

The use of authentication mitigates the threats exposed by the use of the techniques described in chapter 8.2 and makes much more difficult –if not impossible- using them to perform attacks that need to bypass zones.

11.2 Secure the FC-2 layer

In order to provide data origin authentication, integrity protection, replay protection and confidentially on a per-frame basis, an extension of the FC Frame header, very similar to IPSec [19], called ESP_Header (Encapsulating Security Payload) has been defined in FC-FS-2 [4].

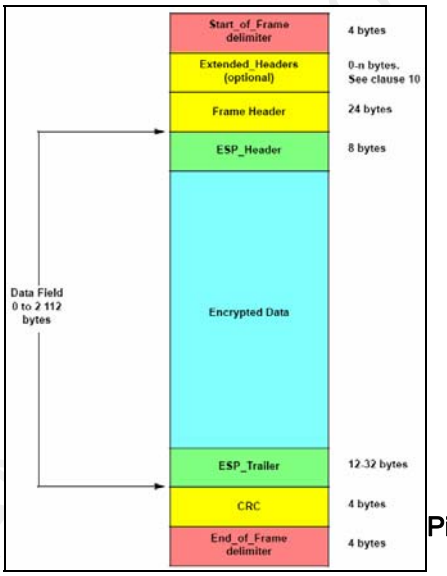

Picture 13: Extended header and CRC in the FC Frame

The ESP header always provides frame integrity check through an ESP_Trailer ICV (Integrity Check Validation) and can also provide confidentiality of the payload by encrypting it. A Security Association agreement must be conducted between both nodes in order to set up encryption parameters. This SA management transaction protocol has been defined as a subset of IKEv2 protocol (RFC 4306). This is an example of one transaction:

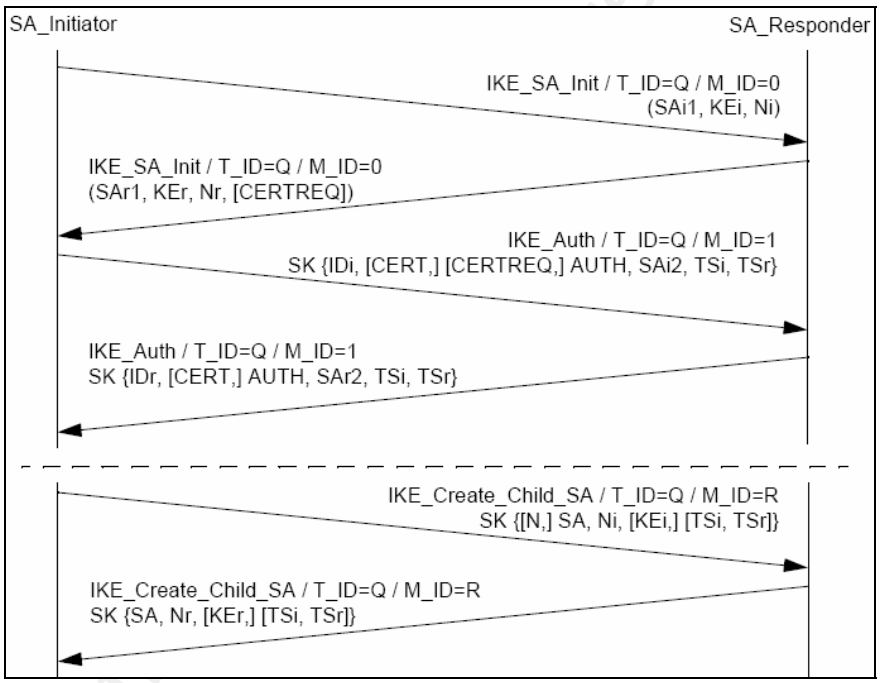

Picture 14: SA Management Transaction Example [7]

The key used for securing the SA exchange can be the one previously originated by the Authentication mechanisms or one pre-shared key [19].

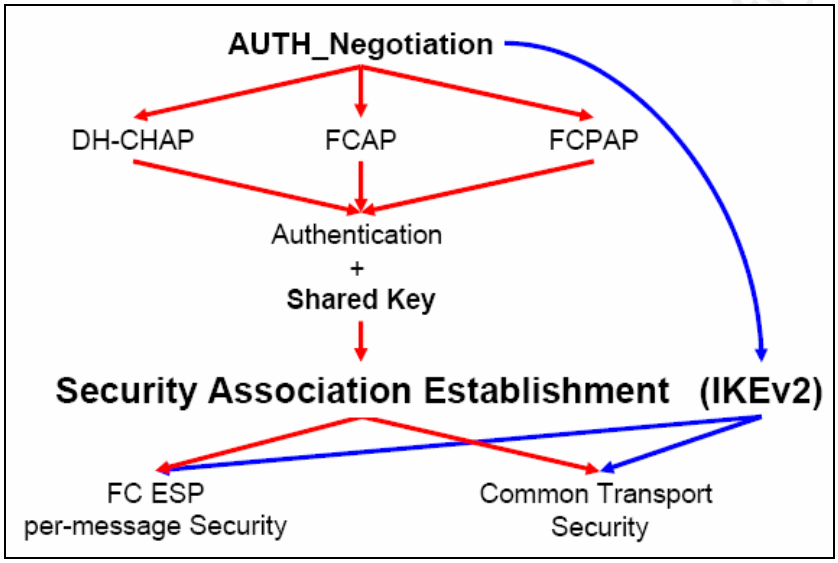

Picture 15: Security Association and Auth Negotiation [19]

Encryption algorithms supported are represented in the following table:

Algorithm (see table 70)	IKEv2 Support	ESP_Header Support	CT_Authentication Support
ENCR_NULL	P	R	R
ENCR_AES_CBC	R	A	R
ENCR_AES_CTR	A	A	A
ENCR_AES_GCM[a]	A	R	A
ENCR_3DES	A	A	A

[a] Support for the 128 bit key length is Required, support for the 192 bit and 256 bit key lengths is Allowed.

Picture 16: Encryption algorithms supported by FC-SP

11.3 Common transport unit protection

As explained before, CT Units are the method to access various critical FC management and control interfaces. The lack of authentication, confidentiality and integrity in CT Units presents a risk for these interfaces.

In order to protect CT Units, CT_Authentication has been defined in FC-GS-5 [5].

The first protection mechanism is the authentication itself, always provided through the use of extended CT_IU preamble, which computes a hash using a shared secret key and a hashing algorithm that is managed through SA transactions in an identical manner as the one explained in the previous chapter. This method also provides anti-reply and integrity protection to the CT_IU communication, avoiding a non-authorized node to send unauthorized CT_IUs to the fabric FC-GS-5 [5].

Additionally, confidentiality con be added to the protection scheme by using a similar mechanism to the ESP_Header, thus encrypting the CT_IU payload, as shown in the picture:

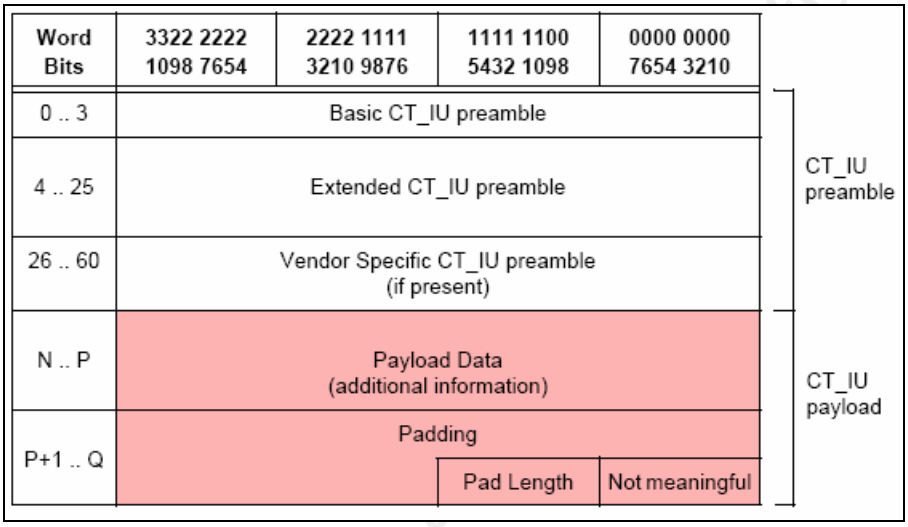

Picture 17: CT_IU with Extended CT_IU preamble & encrypted payload [7]

The use of these techniques mitigates the threats derived from the access to the management server described in chapter 7 [19] [27].

11.4 Policy management

A framework for policy management has been implemented in order to provide a set of policies to manage the fabric and switches. It is an extension of the zone enforcement model in order to expand it to allow the access control of other elements in the fabric. The objects included in the specification are represented in the following table [7]:

Object Identifier	Description	Cardinality	Reference
0000 0001h	Policy Summary Object	One per Fabric	7.1.3
0000 0002h	Switch Membership List Object	One per Fabric	7.1.4
0000 0003h	Node Membership List Object	One per Fabric	7.1.5
0000 0004h	Switch Connectivity Object	A set per Fabric	7.1.6
0000 0005h	IP Management List Object	One per Fabric	7.1.7
0000 0006h	Attribute Object	A set per Fabric	7.1.8

Picture 18: Policy management objects [7]

12 References

1. Abraham, Vineet M. (1999). Design, implementation and evaluation of a Fibre Channel driver for ip on linux. In University of New Hampshire (Ed.), Thesis
2. Ahuja, Vijay (2003). Minding storage. TechTarget, Retrieved from http://http://searchsecurity.techtarget.com/
3. ANSI INCITS. (2004). FC-HBA (386). FC-HBA DRAFT. Latest draft web site: http://www.t11.org
4. ANSI INCITS. (2006). FC-FS-2 (424). FC-FS-2. Latest draft web site: http://www.t11.org
5. ANSI INCITS. (2006). FC-GS-5 (427). FC-GS-5 (final stages of approval, includes ANSI INCITS. (2004). FC-GS-4 (387)). Latest draft web site: http://www.t11.org
6. ANSI INCITS. (2006). FC-LS (433). FC-LS DRAFT. Latest draft web site: http://www.t11.org
7. ANSI INCITS. (2006). FC-SP. FC-SP DRAFT. Latest draft web site: http://www.t11.org
8. ANSI INCITS. (2006). FC-SW-4 (418). FC-SW-4. Latest draft web site: http://www.t11.org
9. Arnold, Gordon (2007). Introduction to Storage Security. SNIA. retrieved from http://www.snia.org/education/tutorials/2007/spring/security/Introduction_to_Storage_Security-v4.pdf
10. BROCADE, (2003).Zoning implementation strategies for Brocade SAN Fabrics. BROCADE

11. BROCADE, (2006). Secure fabric OS administrator's guide. BROCADE.

12. BROCADE, (2007). Brocade Fabric OS v. 5.2.1. Administrator's Guide.

13. BROCADE, (2007). Brocade Fabric OS v. 5.2.1. Reference Guide.

14. BROCADE, (2007). Brocade Fabric OS v. 5.3. Frequently Asked Questions.

15. BROCADE, (2007). Brocade Standards Compliance. Retrieved from Brocade Web site: http://www.brocade.com/products/interop/standards_compliance.jsp

16. BROCADE, (2007).The growing need for security in Storage Area Networks. BROCADE

17. c2olen (2006). SAN Storage in a DMZ, Part 2. Retrieved from http://blogs.rupturedmonkey.com/?p=37

18. CISCO, (2004). SAN Security. In CISCO (Ed.), CISCO Networkers. Web Page: http://www.cisco.com/networkers/nw04/presos/docs/OPT-2054.pdf

19. CISCO, (2005). Maino, Fabio (2005). FC-SP: An Overview of the Standard for Fibre Channel Security. In University of Minnesota DTC – Minneapolis, MN (Ed.), DISC: Third Intelligent Storage Workshop Minessota: CISCO SYSTEMS. Retrieved from University of Minnesota's Digital Technology Center Web site: http://www.dtc.umn.edu/resources/maino.pdf

20. CISCO, (2006). Cisco MDS 9000 Family Fabric Management. Retrieved from Cisco White Papers Web site:

http://www.cisco.com/en/US/netsol/ns514/networking_solutions_white_paper09186a00800c4661.shtml

21. CISCO, (2006). Cisco MDS 9513 Multilayer Director DataSheet. CISCO

22. CISCO, (2007). Cisco MDS 9000 Family CLI Configuration Guide, Release 3.x. CISCO.

23. CISCO, (2007). Cisco MDS 9000 Family Fabric Manager Configuration Guide, Release 3.x. CISCO. Retrieved from Cisco Products Web site: http://www.cisco.com/en/US/products/ps5989/products_configuration_guide_book09186a0080667aa0.html

24. CISCO, (2007). Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide. CISCO

25. DECRU, (2003). SAN Security Threats. DECRU

26. Dwivedi, H. & Hubbard, A. (2003). Securing Storage Networks. @Stake, Retrieved from http://cnscenter.future.co.kr/resource/rsc-center/vendor-wp/atstake/atstake_storage_networks.pdf

27. Dwivedi, Himanshu (2006). Securing Storage: a practical guide to SAN and NAS security. Addison-Wesley.

28. EMC, (2006).Security for Fibre Channel storage area networks. Best practices planning. EMC

29. EMC, (2004). Solutions Enabler Symmetrix Access Control CLI Version 5.5 Product Guide.

30. EMULEX, (2007). Emulex Driver and Utilities for Linux Version 8.0.16.31 or later. HBAnyware Version 3.2 User Manual. EMULEX

31. EMULEX, (2007).Troubleshooting and Maintenance Manual for Emulex HBAs. EMULEX.

32. EMULEX, Solution ID: 12136. Retrieved from Emulex Knowledgebase Web site: http://oak.emulex.com/fcsskb/findsol2.asp?solution=12136

33. Evans, Chris (2003,October,08). How to… interpret worldwide names. Techworld, retrieved July 8th, 2007, from http://www.techworld.com/storage/features/index.cfm?featureid=156

34. GSTInc, (2007). User Guide - D2D & D2D2T. GSTInc. Retrieved from http://www.gstinc.com/support/manuals/d2d2t-software.pdf

35. Hofer, Larry (2007). Best current practices and implementing the FC security protocol (FC-SP). In SNIA (Ed.), SNIA. Retrieved from SNIA Education Web Page: http://www.snia.org/education/tutorials/2007/spring/security#3

36. HP, (2005). fcmsutil(1M) man page. Retrieved from HP Documentation Web site: http://docs.hp.com/en/B2355-60127/fcmsutil.1M.html

37. HP, (2006). Enabling HP MSA1000 Boot from SAN for Red Hat Enterprise Linux on HP

ProLiant BL20p G2 Blade Servers. Retrieved from HP active answers web site: http://h71019.www7.hp.com/ActiveAnswers/downloads/WP_HP_MSA1000_Boot_SAN.pdf

38. IBM, SAN - Explaining different ports. Retrieved from IBM RedBooks Technotes Web site: http://www.redbooks.ibm.com/redbooks.nsf/redbookabstracts/tips0037.html?Open

39. Kipp, Scott (2001). Fibre Channel threat model. In SNIA (Ed.), T11 comitee. Retrieved from http://www.t11.org/ftp/t11/pub/fc/security/01-459v0.pdf

40. Petrocelli, Tom (2003). Guidelines for ensuring storage security. InfoStor, Retrieved August 19, 2007, from http://www.infostor.com/Articles/Article_Display.cfm?Section=Articles&ARTICLE_ID=186304

41. Preston, W. Curtis (2007, May). Protect your SAN from an attack. Storage Magazine, Retrieved August 18th, 2007, from http://searchstorage.techtarget.com/magItem/0,291266,sid35_gci917665_idx2,00.html

42. QLOGIC, (2005). QLOGIC QLA2340 Fibre Channel driver for linux. Retrieved from QLOGIC Support Page Web site: http://support.qlogic.com/support/drivers_software.asp?id=m10 [QLOGIC, 2005]

43. QLOGIC, (2007). SANsurfer FC HBA CLI Application and User's Guide. QLOGIC.

44. Requejo, A. & Zamorano, J. (2004). Buenas prácticas de seguridad en redes de almacenamiento. Germinus SIC, 60, Retrieved from

http://www.germinus.com/sala_prensa/articulos/Buenas%20practicas%20de%20seguridad%20en%20redes%20de%20almacenamiento.pdf.

45. SNIA. (2005). Introduction to Storage Security. SNIA. Retrieved from http://www.snia.org/ssif/documents/Storage-Security-Intro.051014.pdf.

46. SNSEurope (2002). SAN: Providing Speed and Continuous. SNSEurope, Retrieved from http://www.snseurope.com/snslink/news/printer-friendly.php?newsid=1484

47. SYMANTEC, (2007). Veritas CommandCentral Storage Administrator's Guide 5.0. Symantec Corporation.

48. Tate, John, Cartwright, Brian, Croning, John, & Dapprich, Christian (2003).IBM SAN Survival Guide. IBM. Retrieved from IBM Redbooks Web site: http://www.redbooks.ibm.com/redbooks/pdfs/sg246143.pdf

49. TechTarget (2007,March,29). Enterprise switch upgrades specifications. ComputerWeekly.com, retrieved from http://www.computerweekly.com/Articles/2007/03/29/222748/enterprise-switch-upgrade-specifications

50. U.S. Defence Information Systems Agency. (2003). Storage Area Networks (SAN) Security Analysis. Retrieved from http://iase.disa.mil/stigs/whitepaper/san-white-paper.doc