Storage Area Networks Security Protocols and Mechanisms

Whitepaper for information security practitioners

Franjo Majstor
franjo@cisco.com

April 2004

Ver v1.0

Storage Area Networks Security Protocols and Mechanisms

Index of Content

1 Introduction and scope 4

2 SAN technology and protocols overview 4

2.1 DAS vs. NAS vs. SAN 4
2.2 Small Computer Systems Interface known as SCSI 6
2.3 Internet SCSI 6
2.4 Fibre Channel 7
2.5 Fibre Channel over TCP/IP 8
2.6 Other SAN Protocols 9

3 SAN Security Threats Analysis 10

3.1 Availability 10
3.2 Confidentiality and Integrity 10
3.3 Access Control and Authentication 11

4 SAN Security Mechanisms 11

4.1 Securing FC fabric 11
4.1.1 Zoning 11
4.1.2 LUN Masking 12
4.1.3 Fibre Channel Security Protocols 13
4.1.3.1 FC-SP Authentication and Key Management Protocols 13
4.1.3.1.1 Diffie-Hellman Challenge Handshake Authentication Protocol 13
4.1.3.1.2 Fibre Channel Authentication Protocol 14
4.1.3.1.3 Fibre Channel Password Authentication Protocol 14
4.1.3.1.4 FC-SP Authentication protocols comparison 14
4.1.3.2 FC-SP per frame confidentiality and integrity 15

4.2 Securing Storage over IP Protocols 16
4.2.1 IP Security Protocol overview 16
4.2.2 iSCSI Security Mechanisms 18
4.2.3 iFCP, FCIP and iSNS Security Mechanisms 19

5 Storage Security Standard Organisations and Forums 19

6 Future directions 20

7 Summary 20

8 References 21

Index of Exhibits

Exhibit 2: NAS Architecture 5
Exhibit 3: SAN Architecture 5
Exhibit 4: iSCSI Encapsulation. 6
Exhibit 5: iSCSI Solution Architecture 7
Exhibit 6: Fibre Channel Protocol Stack 8
Exhibit 7: FCIP Encapsulation. 8
Exhibit 8: FCIP and iSCSI Solution Architecture 9
Exhibit 9: FC Zoning Example 11
Exhibit 10: FC-SP Authentication and Key Management Protocols 14
Exhibit 11: Fibre Channel Security Protocol Frame 16
Exhibit 12: IPsec Transport and Tunnel Mode 18
Exhibit 13: FC SP Policy Distribution and Key Management options 20

Storage Area Networks Security Protocols and Mechanisms

1 Introduction and scope

Storage devices were up to fairly recently locked into a glass room and hence was the data stored on them enjoying privileges of the physical data center security and protection mechanisms. With a development of a Storage Area Network (SAN) technology, hard drives and tape drives are not necessarily directly attached to a host any more but could be rather physically distant up to several hundred kilometers or even around a globe. Such a flexibility of logically instead of physically attached storage devices to a host made them remotely accessible and highly available, however it brought into a consideration all security elements of the modern network environment like privacy, integrity of the data in transit and authentication of the remotely connected devices. From the data perspective, we could distinguish the storage network security, which refers to protection of the data while it is in transit versus storage data security to which we refer when the data is stored on the tapes or the hard drives. Focus of this article is to make the information security professionals aware of the new communication protocols and mechanisms for storage network security, explain threats and their security exposures as well as describe guidelines for their solutions.

2 SAN technology and protocols overview

2.1 DAS vs. NAS vs. SAN

Historically, storage devices, such as disk drives and backup tapes, were directly attached to a host, hence the name Direct Attached Storage or DAS. This was typically performed via SCSI (Small Computer Systems Interface) parallel bus interface with a speed of up to 320 MBps. This approach of attaching storage devices is coming from internal computer architecture which has obviously got to its limits in several ways. Number of devices which could be attached to one bus is limited even in latest version of SCSI protocol to only 16 devices while the distances are not bigger than 15 meters. Sharing disk or tapes drives amongst multiple hosts were due to architecture of DAS impossible or required specialized and typically expensive software or controllers for device sharing. On the other side, utilisation of the storage spread across the multiple servers was typically lower than on one single pool. Often necessary expansions of storage volumes and replacement of the failed hard drives have in DAS architecture frequently generated system downtimes. DAS Architecture is illustrated in Exhibit 1.

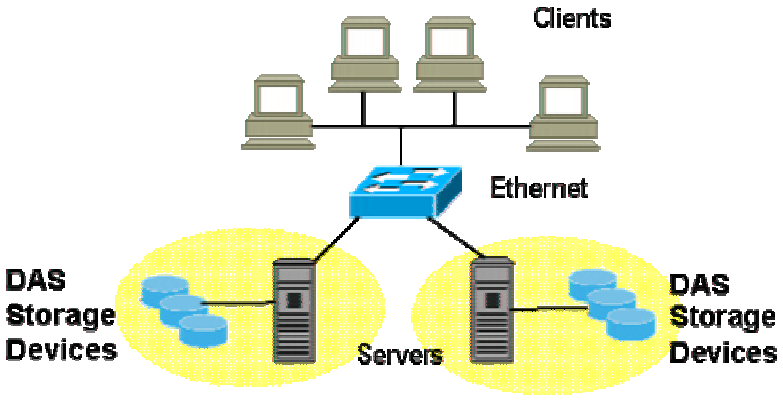

Exhibit 1: DAS Architecture.

The effort to get a better usage of storage devices by the multiple hosts has generated specialized devices for shared storage access on the file level. This architecture is commonly referred as Network Attached

Storage or shortly NAS. NAS architecture consist of a dedicated device named Filer which is actually a stripped down and optimized host for very fast network file sharing. Two most typically supported file systems on Filers are NFS (Network File Systems) for a Unix world and CIFS (Common Internet File System) for the Microsoft world. . While NAS solution has its main advantage in simplicity in maintenance and installation, its main drawback is limited file and operating system support or support of future new file systems. Architecture of a NAS is illustrated in Exhibit 2.

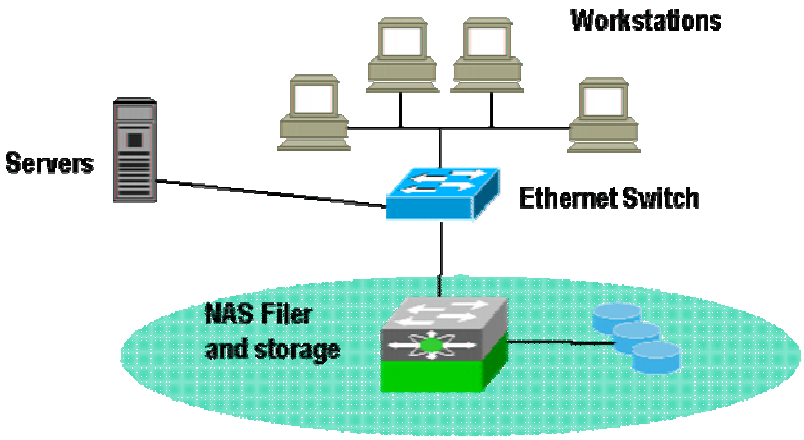

Exhibit 2: NAS Architecture

The latest mechanism of attaching storage remotely with a block level access is commonly referred as Storage Area Network or SAN. SAN consist of hosts, switches and storage devices. Hosts equiped with Host Bus Adapters (HBA) are attached via optical cable to a storage switches which act as a fabric between the hosts and the storage devices. SAN architecture is illustrated in Exhibit 3.

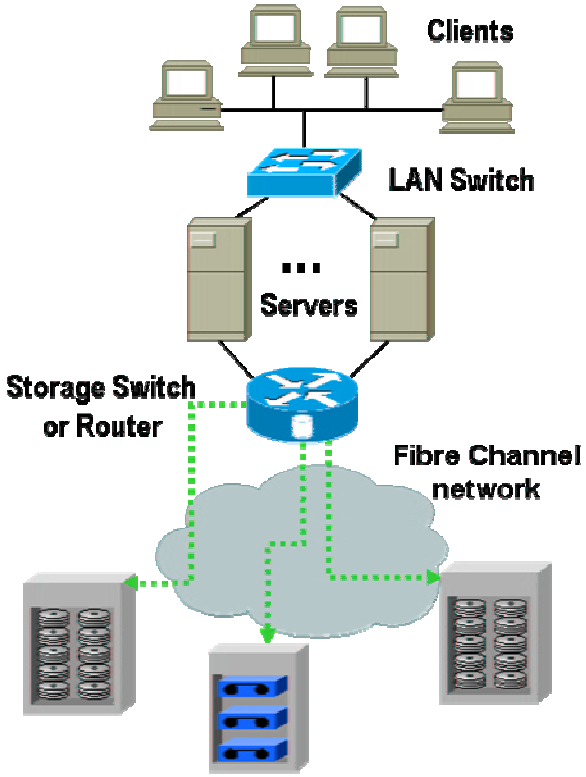

Exhibit 3: SAN Architecture

The invention of a Fibre Channel (FC) has opened a complete new era in the way the storage devices got connected to each other and to hosts. First advantage was the greater distance with up to 10 km, while the different topologies also opened a much bigger number of storage devices that could get connected and shared amongst the multiple hosts.

2.2 Small Computer Systems Interface known as SCSI

In the long history of adaptations and improvements, the line sometimes blurs between where one Small Computer System Interface (SCSI) ends and another begins. The original SCSI standard approved in 1986 by the American National Standards Institute (ANSI), supported transfer rates of up to 5 MBps (megabytes per second) which is, measured by today's standards, slow. Worse yet, it supported a very short bus length. When original SCSI was introduced, however, it represented a significant improvement over what was available at that time, but the problem was the compatibility - since many vendors offered their own unique SCSI options. The next generation of SCSI standard SCSI-2, incorporated SCSI-1 as its subset. In development since 1986, SCSI-2 gained its final approval in 1994 and resolved many of the compatibility issues original SCSI-1 faced. With SCSI-2, it was possible to construct more complex configurations using a mix of peripherals. The most noticeable benefit of SCSI-2 over SCSI-1 was its speed. Also called Fast SCSI, SCSI-2 typically supported bus speeds up to 10 MBps but could go up to 20 MBps when combined with fast and wide SCSI connectors. Fast SCSI enabled faster timing on the bus (from 5 to 10 MHz), thereby providing for higher speed. Wide SCSI used an extra cable to send data that's 16 or 32 bits wide, which allowed for double or quadruple the speed over the bus versus standard, narrow SCSI interfaces that were only 8 bits wide. The latest specification of SCSI protocol, SCSI-3 was among other improvements the first one that did a separation of the higher level SCSI protocol from the physical layer. This was the prerequisite of giving alternatives to run SCSI commands on top of different physical layers than the parallel bus. Hence the SCSI-3 specification was the basis of porting the SCSI protocol to different media carriers such as Fibre Channel or even other transport protocols as TCP/IP.

2.3 Internet SCSI

The SCSI-3 protocol has been mapped over various transports such as parallel SCSI, IEEE-1394 (firewire) and Fibre Channel. All these transports have their specifics but also all have limited distance capabilities. The Internet SCSI or shortly iSCSI protocol is the IETF draft standard protocol that describes means of transporting SCSI packets over TCP/IP. The iSCSI interoperable solution can take advantage of existing IP network infrastructure which have virtually no distance limitations. Encapsulation of the SCSI frames in the TCP/IP protocol is illustrated in Exhibit 4.

Exhibit 4: iSCSI Encapsulation.

The primary market driver for the development of the iSCSI protocol was to enable broader access of the large installed base of DAS over IP network infrastructures. By allowing greater access to DAS devices over IP networks, storage resources can be maximized by any number of users or utilized by a variety of applications such as remote backup, disaster recovery, and storage virtualization. A secondary driver of iSCSI is to allow other SAN architectures such as Fibre Channel to be accessed from a wide variety of hosts across IP networks. iSCSI enables block-level storage to be accessed from Fibre Channel SANs using IP storage routers or switches, furthering its applicability as an IP-based storage transport protocol. iSCSI defines the rules and processes to transmit and receive block storage applications over TCP/IP networks. Although iSCSI can be supported over any physical media that supports TCP/IP as a transport, most iSCSI implementations runs on Gigabit Ethernet. iSCSI protocol can run in software over a standard Gigabit Ethernet network interface card (NIC) or can be optimized in hardware for better performance on an iSCSI host bus adapter (HBA).

iSCSI enables SCSI-3 commands to be encapsulated in TCP/IP packets and delivered reliably over IP networks. As it sits above the physical and data-link layers, iSCSI interfaces to the operating system's standard SCSI access method command set to enable the access of block-level storage that resides on Fibre Channel SANs over an IP network via iSCSI-to-Fibre Channel gateways such as storage routers and switches. iSCSI protocol stack building blocks are illustrated in Exhibit 5.

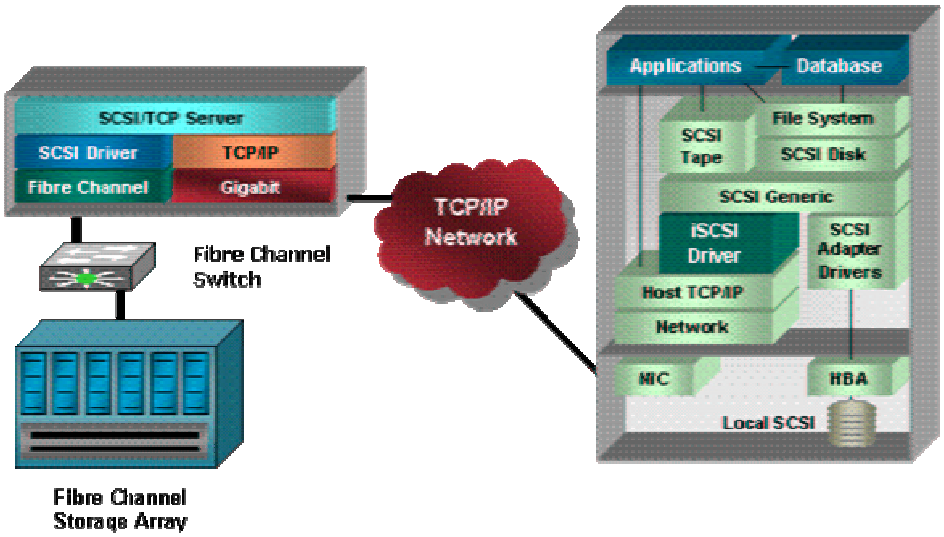

Exhibit 5: iSCSI Solution Architecture

Initial iSCSI deployments were targeted at small to medium-sized businesses and departments or branch offices of larger enterprises that have not deployed Fibre Channel SANs yet, however iSCSI is also an affordable way to create IP SANs from a number of local or remote DAS devices. If there is Fibre Channel present, as it is typically in a data center, it could be also accessed by the iSCSI SANs via an iSCSI-to-Fibre Channel storage routers and switches.

2.4 Fibre Channel

Fibre Channel (FC) is an open industry standard serial interface for high-speed systems. FC is a protocol for transferring the data over fibber cable that consists of multiple layers covering different functions. As a protocol between the host and a storage device, FC was really out of a scope of an average information technology professional for a simple reason that it was point to point connection between the host with a HBA and storage device of typically same vendor which did not require any knowledge or understanding except maybe during the installation process. From the speed perspective, FC is available already in flavors of 1 Gbps and 2 Gbps while specifications for 4Gbps as well as 10Gbps are being worked on and are not that far away.

FC protocol stack is defined in a standard specification of a Technical Committee T11.3 of an INCITS (InterNational Committee for Information Technology Standards) and is illustrated in Exhibit 6.

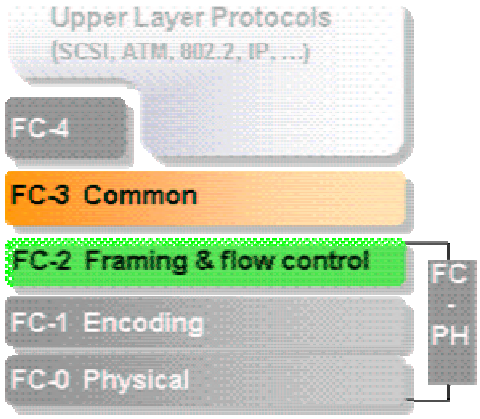

Exhibit 6: Fibre Channel Protocol Stack

The lowest level (FC-0) defines the physical link in the system, including the fibre, connectors, optical and electrical parameters for a variety of data rates. FC-1 defines the transmission protocol including serial encoding and decoding rules, special characters and error control.

The Signaling Protocol (FC-2) level serves as the transport mechanism of Fibre Channel. It defines the framing rules of the data to be transferred between ports, mechanisms for controlling the different service classes and the means of managing the sequence of a data transfer.

The FC-3 level of the FC standard is intended to provide the common services required for advanced features such as:

- Striping -To multiply bandwidth using multiple ports in parallel to transmit a single information unit across multiple links.
- Hunt groups - The ability for more than one port to respond to the same alias address. This improves efficiency by decreasing the chance of reaching a busy port.
- Multicast

FC-3 Layer is the one initially thought to be also used for encryption or compression services, however latest development have put these services to the Layer 2 of a FC architecture as it will be described later.

FC-4, the highest level in the FC structure defines the application interfaces that can execute over Fibre Channel. It specifies the mapping rules of upper layer protocols such as SCSI, ATM, 802.2 or IP using the FC levels below.

2.5 Fibre Channel over TCP/IP

Fibre Channel Over TCP/IP (FCIP) protocol is described in the IETF draft standard as the mechanisms that allow the interconnection of islands of Fibre Channel storage area networks over IP-based networks to form a unified storage area network in a single Fibre Channel fabric. Encapsulation of the FC frames which are carrying SCSI frames on top of the TCP is illustrated in Exhibit 7.

Exhibit 7: FCIP Encapsulation.

FCIP transports Fibre Channel data by creating a tunnel between two endpoints in an IP network. Frames are encapsulated into TCP/IP at the sending end. At the receiving end, the IP wrapper is removed and native Fibre Channel frames are delivered to the destination fabric. This technique is commonly referred to as tunneling, and has historically been used with non-IP protocols such as AppleTalk and SNA. Usage of the FCIP as well as iSCSI protocols is illustrated in Exhibit 8.

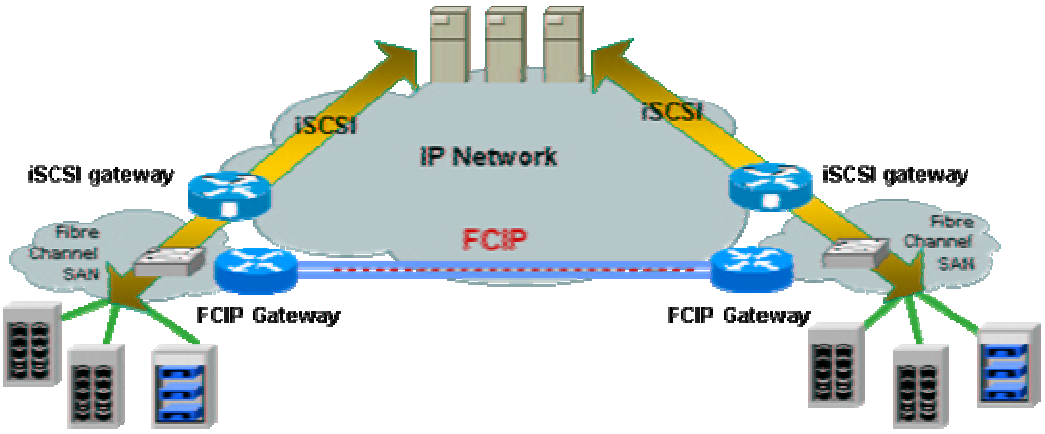

Exhibit 8: FCIP and iSCSI Solution Architecture

The technology is implemented using FCIP gateways, which typically attach to each local SAN through an expansion-port connection to a Fibre Channel switch. All storage traffic destined for the remote site goes through the common tunnel. The Fibre Channel switch at the receiving end is responsible for directing each frame to its appropriate Fibre Channel end device.

Multiple storage conversations can concurrently travel through the FCIP tunnel, although there is no differentiation between conversations in the tunnel. An IP network management tool could view the gateways on either side of the tunnel, but cannot view in on the individual Fibre Channel transactions moving within the tunnel. The tools would thus view two FCIP gateways on either side of the tunnel, but the traffic between them would appear to be between a single source and destination, not between multiple storage hosts and targets.

Connecting Fibre Channel switches creates a single Fibre Channel fabric analogous to bridged LANs or other Layer 2 networks. This means that connecting two remote sites with FCIP gateways creates one Fibre Channel fabric that can extend over miles. This preserves Fibre Channel fabric behavior between remote locations but could leave the bridged fabric vulnerable to fabric reconfigurations or excessive fabric-based broadcasts.

2.6 Other SAN Protocols

There are several other SAN protocols which are in IETF draft proposal or development like Internet Fibre Channel Protocol (iFCP) or Internet Storage Name Services (iSNS). iFCP is also a gateway-to-gateway approach in which FC frames are encapsulated directly into IP packets and IP addresses are mapped to a FC devices. This is more iP-oriented scheme than the IP tunneled SCSI frames FCIP, but is a more complex protocol that was designed to overcome the potential vulnerabilities of stretched fabrics, enable multi-point deployments and provide native IP addressing to individual Fibre Channel transactions.

iSNS protocol is used for interaction between iSNS servers and iSNS clients in order to facilitate automated discovery, management, and configuration of iSCSI and FC devices on a TCP/IP network. iSNS provides intelligent storage discovery and management services comparable to those found in FC

networks, allowing a commodity IP network to function in a similar capacity as a storage area network. iSNS also facilitates a seamless integration of IP and FC networks, due to its ability to emulate FC fabric services, and manage both iSCSI and Fibre Channel devices. iSNS thereby provides value in any storage network comprised of iSCSI devices, Fibre Channel devices (using iFCP gateways), or any combination thereof. iFCP requires iSNS for discovery and management, while iSCSI may use iSNS for discovery, and FCIP does not use iSNS.

3 SAN Security Threats Analysis

Security is a key source of a wide acceptance when it comes to SAN technologies. According to numerous market surveys, the main reason why most enterprises have not yet deployed SANs is due to security concern. When SAN technology was introduced, security was routinely ignored. This was partly because the largely unknown Fibre Channel protocol used for communication was not a big target for attackers and also mainly because security simply wasn't a priority. Today, when SANs are starting to reach across the country or even around the globe, storing and transferring terabytes of sensitive and confidential data, may quickly draw the attention of potential attackers.When the underlying protocol carrying the data over long distance and out of the glass room does not provide the essential data protecting mechanism, data in transit is exposed to a threat of being stolen, seen by the unintended party, modified or simple being not available when it is needed. Logical instead of physical attachment of the storage devices also opens issues of the access control and an authentication of the remote nodes exchanging the data. Moving SAN communications to IP-based networks makes it even more exposed and vulnerable to many of the attacks made on corporate networks.

3.1 Availability

With a SAN technology, storage device could be reached through a possible several redundant paths as well as be easily shared between multiple hosts and simultaneously accessed by multiple clients. It is not necessary any more to bring critical hosts down to be able to replace broken storage devices or expand their capacity. With such features, we could say that the SAN technology has, by decupling the storage from hosts, achieved the greatest level of the storage availability. However we have to keep in mind that by moving storage communication protocols to run on top of TCP/IP, we have also inherited threats and exposures of the TCP/IP environment. We could look at the threats end exposure from two perspectives: exposures to data running on top of TCP as well as exposure to SAN infrastructure devices. It is important to look at the mechanisms which are available or not available within each of the SAN carrier protocols for protecting the storage devices against the availability attacks. With introduction of the storage switches and routers as new infrastructure devices also managed via TCP/IP protocol, it is vital to have proper availability protection mechanisms in place on their management channels as well as have access control mechanisms and different role levels for their configuration control management.

3.2 Confidentiality and Integrity

IP networks are easier to monitor but are also easier to attack. One of the major issues introduced by running SANs over IP networks is the opportunity to sniff the network traffic. All IP based storage protocols just encapsulate the SCSI frames on top of TCP and do not provide any confidentiality or integrity protection. Same is valid for the Fibbre Channel communication. Although it is much more difficult than sniffing an IP-based network, it is also possible to sniff a Fibre Channel network. Hence both IP as well as FC based SAN’s require additional traffic protection mechanisms regarding the confidentiality as well as integrity of the data.

3.3 Access Control and Authentication

Another critical aspect of SAN security is authorization and authentication, controlling who has access to what within the SAN. Currently, the level of authentication and authorization for SANs is not as detailed and granular as it should be. Most security relies on measures implemented at the application level of the program requesting the data, not at the storage device, which leaves the physical device vulnerable. Moving SAN communications to IP-based networks makes it even more exposed and vulnerable to attacks made on corporate networks, such as device identity spoofing. Each of the technologies, like iSCSI as well as FC or FCIP has its own mechanisms of how to address the remote node authentication requirements or it rely on other protocols such as IP Security protocol (IPsec)

4 SAN Security Mechanisms

The basic rules of security also apply to SANs. Just because the technology is relatively new, the security principles are not. First, SAN devices should be physically secured. This was relatively simple to accomplish when SANs existed mainly in well-protected datacenters. But as SANs grow more distributed and their devices sit in branch office closets, physical security is tougher to guarantee. On top of that, each of the protocols mentioned so far has its own subset of security mechanisms.

4.1 Securing FC fabric

By itself, Fibre Channel is not a secure protocol. Without implementing certain security measures within a Fibre Channel SAN, hosts will be able to see all devices on the SAN and could even write to the same physical disk! The two most common methods of providing logical segmentation on a Fibre Channel SAN are zoning and LUN (Logical Unit) masking.

4.1.1 Zoning

Zoning is a function provided by fabric switches that allows segregation of a node in general by physical port, name or address. Zoning is similar to network VLANs (virtual LANs), segmenting networks and controlling which storage devices can be accessed by which hosts. With zoning, a storage switch can be configured for example to allow host H1 to talk only with storage device D1, while host H2 could talk only to storage device D2 and D3, like it is illustrated in Exhibit 9.

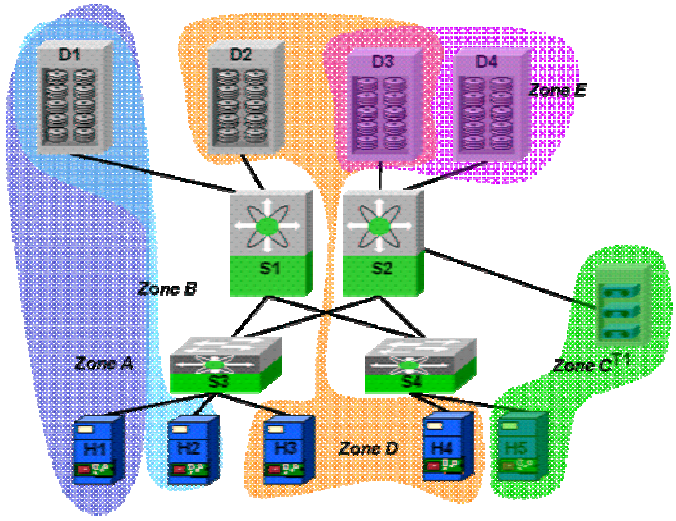

Exhibit 9: FC Zoning Example

One host or storage device could also belong to a multiple zones, like for example on the same exhibit, device D1 belonging to Zone A as well as to Zone B. Zoning can be implemented using either hardware or software, hence we distinguish two main types of Zoning within FC: 'Soft' Zoning and 'Hard' Zoning. Soft Zoning refers to software-based Zoning; that is, zoning is enforced through control-plane software on FC switches themselves - in the FC Name Server service. FC Name Server service on a Fibre Channel switch does mapping between the 64-bit World Wide Name (WWN) addresses to Fibre Channel IDs (FC_ID). When devices connect to a FC fabric, they use the Name Server to find which FC_ID belongs to a requested device WWN. With soft zoning, a FC switch responding to a Name Server query from a device will only respond with a list of those devices registered in the name server that are in the same zone(s) as that of the querying device. Soft Zoning is from the security perspective only limiting visibility of the devices based on the response from the Name Server and does not on any other way restrict access to the storage device from an intentional intruder. This is the job of a Hard Zoning, which refers to hardware-based Zoning.

Hard Zoning is enforced through switch hardware access ports or Access Control Lists (ACLs) which are applied to every FC frame that is switched through the port on the storage switch. Hardware zoning hence has a mechanism not just to limit visibility of FC devices but also to control the access and restrict the FC fabric connectivity to an intentional intruder.

FC Zoning should always be deployed in a FC fabric if not from a node isolation perspective, then for the purpose of minimizing the loss of data. In general, it is also recommended that as many Zones are used as there are hosts communicating with storage devices. For example, if there are 2 host each communicating with 3 storage devices; it would be recommend using 2 zones.

4.1.2 LUN Masking

To further protect the SAN, LUN (Logical Unit Number) Masking could be used to limit access to storage devices. LUN Masking is an authorization process that makes a LUN available to some hosts and unavailable to other hosts. LUN Masking is important because Microsoft Windows based hosts attempt to write volume labels to all available LUN's. This can render the LUN's unusable by other operating systems and can result in data loss. LUN Masking goes one step beyond zoning by filtering access to certain storage resources on the SAN and could be as well provided through hardware (i.e. intelligent bridges, routers, or storage controllers) or through software, utilizing a piece of code residing on each computer connected to the SAN. For each host connected to the SAN, LUN Masking effectively masks off the LUNs that are not assigned to the host, allowing only the assigned LUNs to appear to the host's operating system. The hardware connections to other LUNs still exist, but the LUN Masking makes those LUNs invisible. Managing paths by LUN Masking is a reasonable solution for small SANs, however, due to the extensive amount of configuration and maintenance involved, it is cumbersome for larger SANs.

Although Zoning and LUN Masking provide one layer of SAN devices separation, they are not exclusive security mechanisms but rather isolation mechanisms, and as such they do not give any granular control over data access. Overall SAN security depends on the security of the hosts accessing the storage devices, especially if specific controls are not in place to protect the data. Consider the zoning example: If host H1 can access storage device D1, an unauthorized user or an attacker who compromises host H1 will be able to access any data on Storage Device D1. For SANs to be secure, there must be control that requires proper authorization and authentication to access any data on the storage device, regardless of where the request is originating. It is also needed to limit access to a SAN so that only authenticated and authorized nodes could join the FC fabric as well as protect the confidentiality and integrity of the data in transport through the fabric. These security mechanisms are addressed in work in progress under the Fibre Channel Security Protocol (FC-SP) specification.

4.1.3 Fibre Channel Security Protocols

To address additional security concerns of FC fabric, top SAN industry players have developed the Fibre Channel Security Protocols (FC-SP) specification which is the effort of a working group of the International Committee for Information Technology Standards (INCITS) T11.3 committee. The result is the draft of the future FC-SP standard that extends the Fibre Channel architecture with:

- switch-to-switch, switch-to-device, and device-to-device authentication
- frame-by-frame FC-2 level encryption that provides origin authentication, integrity, anti-replay and privacy protection to each frame sent over the wire
- consistent and secure policy distribution across the fabric

With implementing FC-SP, switches, storage devices and hosts shall be able to prove their identity through a reliable and manageable authentication mechanism. FC-SP can protect against impersonation attacks from rogue hosts, disks, or fabric switches, as well as providing protection from common misconfigurations when cabling devices in a fabric. With FC-SP, Fibre Channel traffic can be secured on a frame-by-frame basis to prevent snooping and hijacking, even over non trusted links. A consistent set of policies and management actions are propagated through the fabric to provide a uniform level of security across the entire fabric. FC-SP includes support for data integrity, authentication for both switch-to-switch and host-to-switch communication as well as optional confidentiality.

4.1.3.1 FC-SP Authentication and Key Management Protocols

Authentication is the process by which an entity is able to verify the identity of another entity. As such, authentication is the foundation of security. A Fibre Channel device may authenticate the entity trying to access resources by verifying its identity. Different authentication protocols may be used to validate an entity on the basis of different parameters. Each Fibre Channel entity is identified by a name. The purpose of an authentication protocol for Fibre Channel is to verify, by using some form of digital credentials, that a claimed name is associated with the claiming entity. FC-SP specify three optional authentication mechanisms, whose first role is to address the threat of identity spoofing within or when accessing the FC fabric.

4.1.3.1.1 Diffie-Hellman Challenge Handshake Authentication Protocol

Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) is a password based authentication and key management protocol that uses the CHAP algorithm (RFC 1994) augmented with an optional Diffie-Hellmann algorithm. DH-CHAP provides bidirectional and optionally unidirectional authentication between an authentication initiator and an authentication responder. In order to authenticate with the DH-CHAP protocol, each entity, identified by a unique name, shall be provided with a secret. Each other entity that wants to verify that entity shall know the secret associated with that name or defer the verification to a third party, such as a RADIUS or TACACS+ server that knows that secret. When the Diffie-Hellmann part of the protocol is not performed, DH-CHAP reduces its operations to those of the CHAP protocol, and it is referred to as DH-CHAP with a null DH algorithm. DH-CHAP with a null DH algorithm is the authentication protocol mandatory to implement in each FC-SP compliant implementation, for interoperability reasons. DH-CHAP protocol has the other parameters that are possible to negotiate such are the list of hash functions (e.g. SHA1, MD5) and the list of the usable Diffie-Hellman Group Identifiers. Diffie-Hellman Group Identifiers that are possible are 1, 2, 3 or 4, with group bit sizes of 1024, 1280, 1536 and 2048 respectively.

4.1.3.1.2 Fibre Channel Authentication Protocol

Fibre Channel Authentication Protocol (FCAP) is an optional authentication and key management protocol based on digital certificates that occurs between two Fibre Channel end points. When the FCAP protocol successfully completes, the two Fibre Channel end points are mutually authenticated and may share a secret key. In order to authenticate with the FCAP protocol, each entity, identified by a unique name, shall be provided with a digital Certificate associated with its name, and with the certificate of the signing Certification Authority. Each other entity that wants to participate in FCAP shall be also provided with its own certificate as well as the certificate of the involved Certification Authority for the purpose of the other entity certificate verification. At this time of FC-SP specification only supported format of the digital certificate is X.509v3. FCAP is for the purpose of the shared secret derivation also using the Diffie-Helman algorithm. For the hashing purpose FCAP is using RSA-SHA1 algorithm.

4.1.3.1.3 Fibre Channel Password Authentication Protocol

Fibre Channel Password Authentication Protocol (FCPAP) is an optional password based authentication and key management protocol that uses the Secure Remote Password (SRP) algorithm as defined in the RFC 2945. FCPAP provides bidirectional authentication between an authentication initiator and an authentication responder. For the hashing purpose, FCPAP protocol is relying on SHA-1 algorithm. When the FCPAP protocol successfully completes, authentication initiator and responder are authenticated and by using Diffie-Helman protocol have obtained a shared secret key. Parameters for authentication in the SRP algorithm are a password, a salt, and a verifier. In order to authenticate with the FCPAP protocol, each entity, identified by a unique name, shall be provided with a password. Each other entity that wants to verify that entity shall be provided with a random salt, and a verifier derived from the salt and the password.

4.1.3.1.4 FC-SP Authentication protocols comparison

As listed, each of the authentication protocols have their similarity and differences depending on what mechanism they use for the authentication as well as hashing which are illustrated in the table in Exhibit 10.

FC-SP Authentication Protocol	**Authentication Mechanism**	**Hashing Mechanism**	**Key Exchange Mechanism**
DH-CHAP	RFC 1994, CHAP	MD5, SHA-1	DH
FCAP	x509v3 certificates	RSA-SHA1	DH
FCPAP	RFC 2945, SRP	SHA-1	DH

Exhibit 10: FC-SP Authentication and Key Management Protocols

As we have also seen, by using a Diffie-Helman algorithm all three authentication protocols are capable of not doing only initial mutual entity authentication but are also capable of doing a key management and deriving the shared secret which could be used for the different purpose such as per frame integrity and confidentiality.

4.1.3.2 FC-SP per frame confidentiality and integrity

Recognizing the need for a per-message protection that would secure each FC frame individually, top storage vendors like Cisco Systems, EMC, QLogic, and Veritas proposed an extension to the FC-2 frame format that allow for frame-by-frame encryption. The frame format has been called the ESP Header, since it is very similar to the Encapsulating Security Payload (ESP) used to secure IP packets in IPsec. Given the overall security architecture is similar to IPsec, this aspect of the security architecture for FC is often referred to as FCsec.

The goals of the FCsec architecture are to provide a framework to protect against both active and passive attacks using the following security services:

- Data Origin Authentication to ensure that the originator of each frame is authentic.
- Data Integrity and Anti-Replay protection that provides integrity and protects against each frame transmitted over a SAN.
- Optional encryption for data and/or control traffic that protects each frame from eavesdropping.

The goal of FCsec is also to converge the storage industry on a single set of security mechanisms, regardless of whether the storage transport was based on iSCSI, FCIP, or FC so that FCsec could be layered onto existing applications with minimal or no changes to the underlying applications.

One of the main benefits behind the use of ESP to secure an FC network is its great flexibility; it can be used to authenticate a single control messages exchanged between two devices, to authenticate all control traffic between two nodes, or to authenticate the entire data traffic exchanged between two nodes. Optional encryption can be added to any of the steps above to provide confidentiality.

A per-entity authentication and key exchange protocol provides also a set of other services including the negotiation of the use of ESP for encapsulation of FC-2 frames, the exchange of security parameters to be used with the ESP encapsulation protocol, and the capability to update keys used by the two entity without any disruption to the underlying traffic flow.

ESP is used as a generic security protocol. Independently from the upper layers, ESP can provide the following:

- Per message integrity, authentication and anti-replay.
  When used with a null encryption algorithm and an HMAC as authentication algorithm it guarantees that the frames have not been altered in transit, authenticated for the originating entity and belong to the same sequence exchange.
- Traffic encryption.
  When used with a non-null encryption algorithm such as AES, triple DES, or RC5, it allows the encryption of the frame content.

The specific fields covered by authentication as well as fields that can optionally be encrypted within the FC-SP frame are illustrated in Exhibit 11.

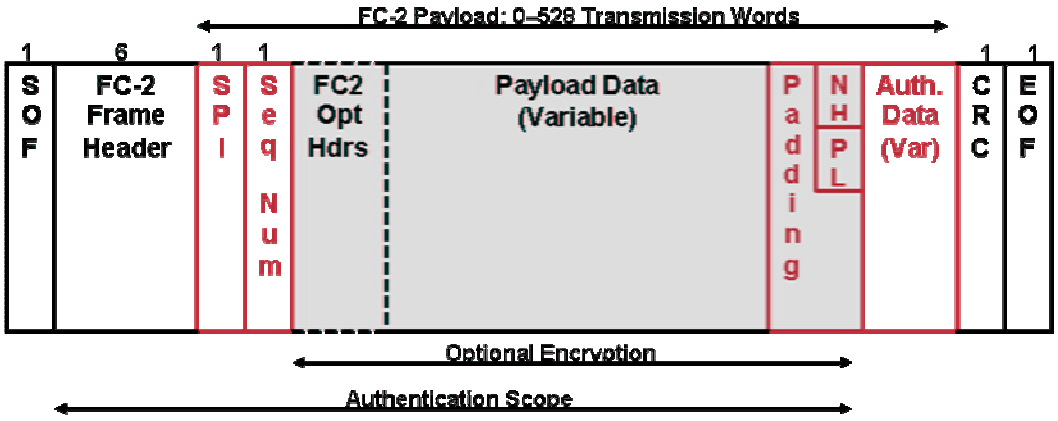

Exhibit 11: Fibre Channel Security Protocol Frame

While IPsec protocol will be briefly discussed later, it is important to notice here the major differences between the IPsec ESP and FCsec in the role of authentication and confidentiality. FCsec frame format gives authentication the complete frame including the header of the frame and has mandatory authentication, while encryption is optional. On the other side, IPsec ESP header does not offer the authentication of the packet header. For that purpose IPsec uses Authentication Header (AH) and while ESP mandates encryption, it has an optional authentication for the rest of the packet payload.

4.2 Securing Storage over IP Protocols

With an exception of initial session login authentication, none of the other IP based SAN protocols: iSCSI, iFCP, FCIP or iSNS does not define its own per-packet authentication, integrity, confidentiality or anti-replay protection mechanisms. They all rely upon the IPsec protocol suite to provide per-packet data confidentiality, integrity, authentication and anti-replay services together with Internet Key Exchange (IKE) as the key management protocol.

The IP Storage working group within the Internet Engineering Task Force (IETF) has developed a framework for securing IP based storage communications in a draft proposal ‘Securing Block Storage Protocols over IP’. The proposal covers use of the IPsec protocol suite for protecting block storage protocols over IP networks (including iSCSI, iFCP and FCIP), as well as storage discovery protocols, iSNS.

4.2.1 IP Security Protocol overview

This chapter is by no means an extensive IP Security (IPsec) protocol description but rather an overview, of the elements that are necessary in order to understand its usage for storage over IP protocols protection. IPsec is applied at the network layer, protecting the IP packets between participating IPsec peers by providing the following:

- Data Confidentiality
  The IPsec sender can encrypt packets before transmitting them across a network.
- Data Integrity
  The IPsec receiver can authenticate packets sent by the IPsec sender to ensure that the data has not been altered during transmission.

- Data Origin Authentication
  The IPsec receiver can authenticate the source of the IPsec packets sent.
- Anti-Replay
  The IPsec receiver can detect and reject replayed packets.

To achieve listed functions, IPsec protocol uses:

- Diffie-Hellman key exchange for deriving key material between two peers on a public network.
- Public key cryptography or pre-shared secret for signing the Diffie-Hellman exchanges to guarantee the identities of the two parties and avoid man-in-the-middle attacks.
- Bulk encryption algorithms, such as DES (Data Encryption Standard), 3DES (Triple DES) or AES (Advance Encryption Standard) for encrypting the data.
- Keyed hash algorithms, such as HMAC (Hashed Message Authentication Code), combined with traditional hash algorithms such as MD5 (Message Digest 5) or SHA1 (Secure Hashing Algorithm 1) for providing packet integrity and authentication.

The IPsec framework consists of two major parts:

- Internet Key Exchange (IKE), which negotiates the security policies between two entities and manages the key material.
- IP Security Protocol suite, which defines the information to add to an IP packet to enable confidentiality, integrity, anti-replay and authenticity controls of the packet data.

IKE is a two phase negotiation protocol based on the modular exchange of messages defined in RFC 2409. It has two phases and accomplishes the following three functions in its Phase 1 and fourth one in Phase 2:

- Protected cipher suite and options negotiation - using keyed MACs, encryption and anti-replay mechanisms
- Master key generation - via Diffie-Hellman calculations
- Authentication of end-points using pre-shared secret or public key cryptography
- IPsec Security Association (SA) management (traffic selector negotiation, options negotiation plus key creation and deletion)

IPsec is adding two new headers to the IP packet:

- AH (Authentication header)
- ESP (Encapsulation Security Payload) header.

AH header provides authentication, integrity and replay protection for IP header as well as for all the upper-layer protocols of an IP packet. However, it does not provide any confidentiality to them. Confidentiality is the task of the **ESP header**, besides providing authentication, integrity and replay protection for the packet payload. Both of the headers could be used in two modes: transport and tunnel modes. The **transport mode** is used when both the communicating peers are hosts. It may also be applied when one peer is a host and the other is a gateway, if that gateway is acting as a host or ending point of the communication traffic. The transport mode has the advantage of adding only a few bytes to the header of each packet. With this choice however, the original IP packet header could only be authenticated but not encrypted. The **tunnel mode** is used between two gateway devices, or between a host and a gateway if that gateway is the conduit to the actual source or destination. In the tunnel mode, the entire original IP packet is encrypted and becomes the payload of a new IP packet. The new IP header has the destination address of its IPsec peer. All the information from the original packet, including the headers, is protected. The tunnel mode protects against attacks on the endpoints due to the fact that, although the IPsec tunnel

endpoints can be determined, the true source and destination endpoints cannot be determined because the information in the original IP header has been encrypted. This is illustrated in Exhibit 12.

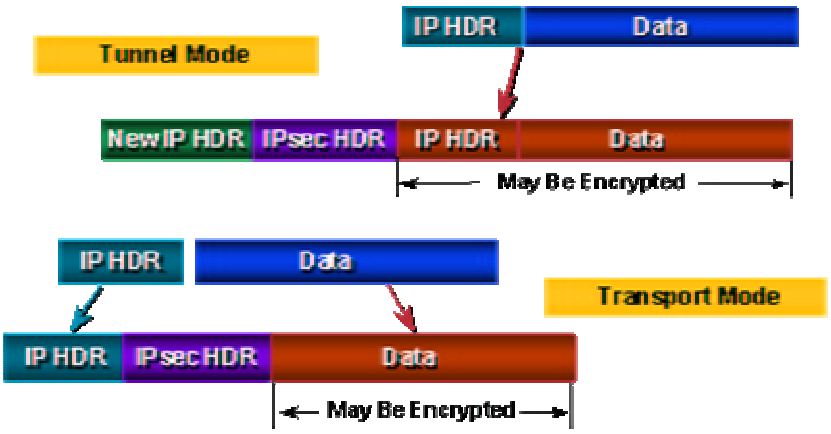

Exhibit 12: IPsec Transport and Tunnel Mode

With IPsec, data can be transmitted across a public network without fear of observation, modification, or spoofing. This enables applications such as Virtual Private Networks (VPNs), including intranets, extranets, remote user access, and remote transport of storage over IP.

The IETF's draft RFC is dictating IPsec and IKE to be used with the IP based storage protocols to provide secure private exchanges at the IP layer. In order to be compliant, an IP storage network element must follow up the specifications and implement IPsec tunnel mode with the ESP where confidentiality is obtained by encrypting the IPsec tunnel using 3DES or optionally AES in cipher block chaining (CBC) mode, integrity checking is done via using SHA-1 and node authentication is done via IKE using pre-shared key or digital certificates.

4.2.2 iSCSI Security Mechanisms

iSCSI draft RFC specifies that although technically possible, iSCSI should not be used without security mechanisms except only in closed environments without any security risk. Security mechanisms defined in the draft standard are the following:

- in-band authentication between the initiator and the target at the iSCSI connection level,
- per packet protection (integrity, authentication, and confidentiality) by IPsec at the IP level.

iSCSI protocol specification defines that during login, the target must authenticate the initiator and the initiator may authenticate the target, which means that mutual authentication is optional but not mandatory. The authentication is performed on every new iSCSI connection during the login process with a chosen authentication method. The authentication method cannot assume an underlying IPsec protection, because IPsec is optional to use and an attacker should gain as little advantage as possible by inspecting the authentication process. Due to listed requirements, chosen authentication method for iSCSI protocol is Challenge Handshake Authentication Protocol (CHAP). The authentication mechanism protects against an unauthorized login to storage resources by using a false identity (spoofing). Once the authentication phase is completed, if the underlying IPsec is not used, all following messages are sent and received in clear. The authentication mechanism alone, without underlying IPsec, should only be used when there is no risk of eavesdropping, message insertion, deletion, modification, and replaying.

An iSCSI node must also support Internet Key Exchange (IKE) protocol to provide per packet authentication, security association negotiation, and key management where a separate IKE phase 2 security association protects each TCP connection within an iSCSI session.

4.2.3 iFCP, FCIP and iSNS Security Mechanisms

iFCP and FCIP are peer-to-peer transport protocols that encapsulate SCSI and Fibre Channel frames over IP. Therefore, Fibre Channel, operating system, and user identities are transparent to the iFCP and FCIP protocols. iFCP and FCIP sessions may be initiated by either or both peer gateways. Consequently, bi-directional authentication of peer gateways must be provided. There is no requirement that the identities used in authentication be kept confidential. Both, iFCP and FCIP as well as iSNS protocol heavily rely on IPsec and IKE for providing security mechanisms for them. In order to be compliant with security specifications in their draft RFCs, storage nodes using any of the three IP storage protocols must implement IPsec ESP in Tunnel Mode for providing data integrity and confidentiality. They may implement IPsec ESP in Transport Mode, if deployment considerations require use of Transport Mode. When ESP is utilized, per-packet data origin authentication, integrity and replay protection also must be used. For message authentication they must implement HMAC with SHA-1 and should implement AES in CBC MAC mode. For ESP confidentiality, they must implement 3DES in CBC mode and should implement AES in CTR mode. For the key management entities must support IKE with peer authentication using pre-shared key and may support peer authentication using digital certificates.

5 Storage Security Standard Organisations and Forums

All IP related protocols are under development within the Internet Engineering Task Force (IETF) working groups. This includes iSCSI, FCIP and iFCP protocols as well as IPsec and interaction of IP storage protocols with IPsec and IKE. On the other side FC, FC-SP and SCSI specifications are developed within American International Committee for Information Technology Standards (ICNCITS) technical committees. The INCITS is the forum of choice for information technology developers, producers and users for the creation and maintenance of formal de jure IT standards. INCITS is accredited by, and operates under rules approved by, the American National Standards Institute (ANSI) and is ensuring that voluntary standards are developed by the consensus of directly and materially affected interests.

Multiple specifications in different standard bodies as well as numerous vendor implementations obviously require standards to drive the interoperability of the products. The lack of interoperability among storage devices also creates security problems. Each vendor designs its own technology and architecture, which makes communication between devices difficult, if not impossible.
Forums and vendor associations are luckily smoothening things up. Storage Networking Industry Association (SNIA) is a non-profit trade association established in 1997 which is working on ensuring that storage networks become complete and trusted solutions across the IT community by delivering materials, educational and information services to its members. The SNIA Storage Security Industry Forum (SSIF) is a vendor consortium dedicated to increase the availability of robust storage security solutions. The forum tries to fulfill its mission by identifying best practices on how to build secure storage networks and promoting standards-based solutions to improve the interoperability and security of storage networks.

6 Future directions

Storage security is still evolving topic and security mechanisms defined in the draft standards yet need to be implemented as well as their interoperabity tested and approved from storage security forums. We have also seen that most of the IP based storage network protocols rely for their protection on IPsec. While IPsec is today already well defined and accepted set of standards, it is also developing further with a new key management specification IKEv2 and FC-SP is following its example by allowing in its latest specification use IKEv2 as its security policy distribution and key management protocol. All options of the FC-SP are illustrated in Exhibit 13.

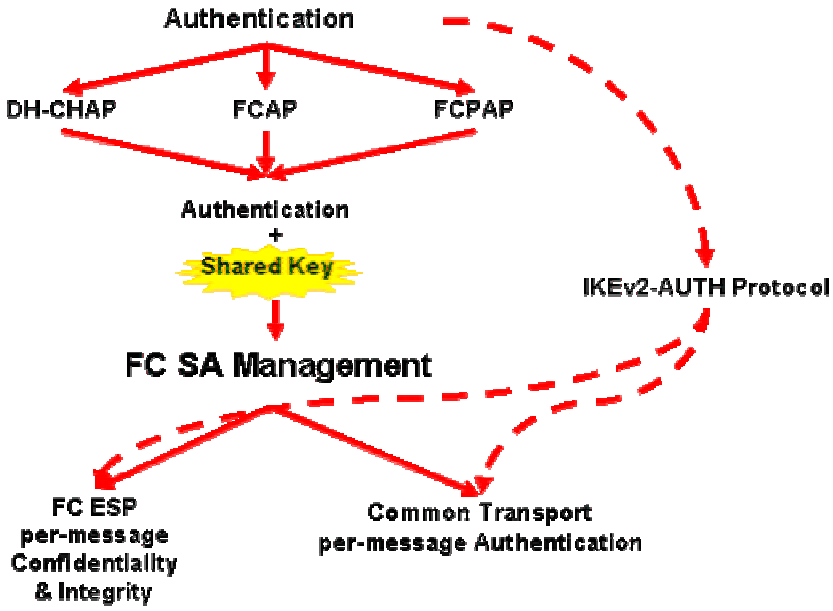

Exhibit 13: FC SP Policy Distribution and Key Management options

An FC Security Association (SA) management protocol is actually simplified version of the Internet Key Exchange protocol version 2 (IKEv2) that builds on the results of the FC authentication and key management protocol. The SA management protocol uses obtained shared secret key as the authentication principal to setup the Security Associations. There are situations where it is acceptable to use IKEv2 to perform both functions of authentication and SA management. This is referred to as a protocol with the name IKEv2-AUTH. On a side of the SAN security protocols development it is also necessary that the hardware implementations follow up the software ones, cause only when the security mechanisms are built-in in silicon will the SAN technology leverage the full benefit of them. The most of the future development in SAN security area lay on the side of protecting the data while it is stored at disk which requires further research of the group key management protocols and their implementation on SAN technology.

7 Summary

Although SAN technologies and protocols are relatively new, security threats they are exposed to are not. This is in particular true once when the storage data leaves the protection space of the data center glass room and traverse the external, most of the time security wise uncontrolled and unprotected network segments. Good news is that SAN technologies and protocols are already fairly equipped with proper security mechanisms in most aspects. Even though that all of the security mechanisms like node authentication, data integrity and confidently do not exist built-in in all storage protocols themselves, specially when they are carried on top of IP, there are pretty matured specifications coming from

international standardization organizations such as IETF and ICNCITS that well define how they should be extended or be used in conjunction with IPsec and IKE protocols as their protection mechanisms. Native SAN fabric protocol FC is on the other side either already leveraging the development of IPsec in a from of FCsec protocol or closely following the development in the key management and policy distribution area with next generation Internet Key Management protocol IKEv2. This all promises unified level of storage data protection traveling over different media carriers and encapsulation protocols. It is now up to industry forums such as SNIA and SSIF to evangelize the security best practices and guidelines to be used when designing, deploying or maintaining the SAN networks. Information security professionals have to be aware that the data stored or traversing the SAN technologies is exposed to security threats and understand and use all possible tools, protocols and mechanisms for their protection.

8 References

[1] Abboba, B., at al., Securing Block Storage Protocols over IP, IETF Internet Draft, <draft-ietf-ips-security-19.txt>, January 2003.

[2] Cyrtis Preston W., Using SANs and NAS, First Edition, O'Reilly & Associates, Inc, February 2002

[3] Dale. L., Whitepaper: Security Features of the Cisco MDS 9000 Family of Multilayer Storage Switches <ftp://ftp-eng.cisco.com/ltd/mds_security_whitepaper16.pdf>, November 2003.

[4] Dwivedi, H. and Hubbard, A., Whitepaper: Securing Storage Networks <http://www.@stake.com/research/reports/acrobat/atstake_storage_networks.pdf>, April 2003.

[5] Doraswamy N., Harkins D., IPSec The New Security Standard for the Internet, Intranets and Virtual Private Networks, Prentice Hall PTR 1999.

[6] Harkins D., Carrel D., The Internet Key Exchange (IKE), RFC 2409, November 1998.

[7] Kaufman, C., Internet Key Exchange (IKEv2) IETF Internet Draft, <draft-ietf-ipsec-ikev2-12.txt>, January 2004.

[8] Monia, C. at al., iFCP - A Protocol for Internet Fibre Channel Storage Networking, IETF Internet Draft, <draft-ietf-ips-ifcp-14.txt>, May 2003.

[9] Satran, J., at al., iSCSI, IETF Internet Draft, <draft-ietf-ips-iscsi-20.txt>, January 19th 2003.

[10] Rajagopal, M., Rodriguea, E., Fibre Channel Over TCP/IP (FCIP), IETF Internet Draft, <draft-ietf-ips-fcovertcpip-12.txt>, February 2003.

[11] Simpson, W., PPP Challenge Handshake Authentication Protocol (CHAP), RFC 1994, August 1996.

[12] Snively R., at al., Fibre Channel Security Protocols (FC-SP) Rev 1.3, INCITS working draft proposed by ANSI, January 31st, 2004

[13] Wu, T., The SRP Authentication and Key Exchange System, RFC 2945, September 2000.

[14] Yongdae K,. at al, Secure Group Key Management for Storage Area Networks, IEEE Communications Magazine, Vol. 41, No.8; p92-99 August 2003.