The Complete Reference

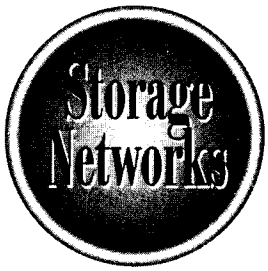

Chapter 25

Security Considerations

As storage networks continue to grow and become more complicated, management processes and security procedures will play an integral role. While management is a key requirement in any type of storage infrastructure, security is the one component that could undermine the availability of both management and storage. This chapter provides an introduction to many of the existing and future security problems in storage networks and the correlating best practices that can be implemented today. Although this chapter will not cover every single storage security problem or provide in-depth coverage on each security exposure identified, it will discuss the architectural and tactical problems that most often plague storage networks and how to approach a solution for each of them.

While various trends have occurred in the digital marketplace over the past five decades, both storage and security have received more attention in the last decade (early 1990s to present) than the first four combined (excluding major governmental agencies' focus on security). Furthermore, with the growth of the Internet, security has been the focus of many conversations concerning network architecture, application design, and now storage technology.

The terms *storage technology* and *security architecture* would not have been mentioned in the same sentence a decade ago, or even five years ago. Today, however, with significant changes in the storage industry (the expansion of the storage network into the WAN and beyond the LAN) and due to education gained from the security mistakes in Internet Protocol (the numerous security problems with IP version 4), the relationship between storage technology and security architecture has become an important one. Today's storage industry has an opportunity to pursue and support a healthy security posture before storage security problems are widely identified and, more importantly, widely exploited.

It's important that you understand why storage and security must coexist so that you can anticipate an attack and understand an attacker's (hacker's) mindset. For example, consider a home with a front door, back door, and garage door as three legitimate entryways. The front door has three security devices: a door-handle lock, a dead-bolt lock, and a chain lock to prevent an intruder from entering. In addition, the homeowner has secured the back door with two security devices: a Master lock combination on the fence leading to the back door and a dead-bolt on the back door. The garage door is opened and closed only with a garage door opener, which is in the vehicle at all times. Additionally, the homeowner has purchased a security system to alarm the authorities in the event of a front or back door break-in (similar to an intrusion detection system, or IDS, in a network). This adds an extra obstacle for access to the front and back doors.

Note *An IDS is a device that passively monitors networks to determine and report any type of malicious activity that is being conducted.*

Now consider the security of the house (or the computer network): The intruder must consider which is the best route, in terms of time and success, to enter the home (or network) and gather priceless goods. The front door contains three security devices. The door lock and dead-bolt can be picked, but this takes a significant amount of time and skill. Furthermore, even if the two locks were picked successfully, the chain lock would have to be cut with an industrial sized chain-cutter, requiring more skill and

time. And then there's the alarm to consider. As a result, the front door sounds like a tough option.

For the back door, the fence Master lock combination will also take a significant amount of time because an infinite amount of combinations would need to be attempted. Even after the combination is finally "brute-forced," the back door dead-bolt would have to be picked, also taking time and skill. Even if this door were opened, the alarm would sound. This option also seems pretty unpromising.

The garage door is secured with the garage door opener. However, most garage door openers do not use an overly complicated infrared architecture to open the garage door. In fact, a local hardware store usually carries a device that spans several channels that might connect on the correct channel and open any garage door. In addition, most hand-held devices, such as Palm and PocketPC, can have applications that capture the infrared signals of a garage door opener to open it exclusively with the hand-held device. The purchase of this device, and the fact that the homeowner will probably be opening/closing her garage at the same time every morning, can result in an intruder using a hand-held device to capture the infrared signal and eventually open the garage. Once inside the garage, the only thing stopping the intruder is the house garage door—but many homeowners do not lock that door because of the assumption that an intruder would not be able to enter the locked garage.

This home example relates to the digital storage network and the mindset of a digital attacker. Although the storage network is not the traditional access point for most attacks, a savvy hacker can avoid the hassle of subverting multiple firewalls, switches, router Access Control Lists (ACLs), Virtual Private Network (VPN) devices, and IDS sensors to gain access to data via the less-protected storage network that has direct access to all of the important data.

Note *Router ACLs are used to allow or deny access to networks based on IP addresses. VPN devices allow for remote networks or individual users to connect to internal networks in a safe and secure manner. A VPN allows multiple networks in different geographic locations to exist as a large virtual network.*

Attackers are not interested in gaining administrator rights or even root access to a given host; rather, they are interested in access to data. The fact that the storage network contains sensitive data and is not adequately protected (similar to the house garage doors) leaves the perfect opportunity for an attacker. Furthermore, many of the storage protocols in place today, such as Fibre Channel and iSCSI, are bandwidth and throughput-focused protocols with security usually absent (similar to the poor infrared channels on many garage door openers). This scenario leaves the "door" wide open for attackers to subvert storage protocols and get direct access to data, without compromising one firewall or encryption device.

Even though the storage network is often not protected thoroughly, it has access to the company's critical data and intellectual property. The myth that the storage network cannot be reached by attackers is easily subverted with a variety of techniques that are usually easier than going through traditional networks with multiple firewalls, switches, router ACLs, VPN devices, and IDS sensors. This fact makes the decision easy for the attackers on which route they should take to access data quickly and easily.

This chapter introduces the concept of security in storage networks and discusses basic principles and best practices to consider. The following topics are discussed:

- Overview of computer security
- Security methods
- Storage security technology
- Storage security challenges
- Fibre Channel SAN security
- NAS security
- Best practices

Overview of Information Security

While security threats used to mean hackers defacing web sites and short-term embarrassment for an organization, new security threats are more serious matters—corrupting or deleting intellectual property, which results in the loss of revenue. The old model of viewing information security as an end result, such as a product, was not successful because it did not support the core principles of a strong security posture. In today's model, security is viewed as a process of good policies and procedures and is more successful and realistic. Security is not a product for an auditor to check off as he or she reviews a list of items. A strong security process in an organization provides a strong foundation to build upon and supports a stable security platform.

To support a good security process, basic security elements should be addressed. These elements can be applied to different aspects of storage networks, such as devices, applications, protocols, appliances, and so on. The following are typical security elements that must be addressed by any secure solution:

- Authentication
- Authorization
- Auditing
- Integrity
- Encryption
- Availability

Authentication

Authentication is the process by which an entity is verified. The entity can be a packet, a frame, a login request, or another entity. In all cases, the entity is identified and then authorized or unauthorized. Authentication and authorization are heavily dependent on each other, because one can be subverted easily in the absence of the other.

Authorization

Authorization is the process of determining which privileges are granted to which authenticated entity. Note that authorization is not the same as authentication. Authorization simply allows or denies actions based on a set of assumed authenticated credentials. Whether the authenticated credentials are valid is not possible to verify with authorization. Authorization views a set of authenticated entities and allocates *rights* to those entities.

Auditing

Auditing is the ability to capture and retain events that occur within the network or specific devices or applications. While auditing is considered a *passive* security element, it can make a network aware of a security incidence, which is often half the battle. In the absence of a pure security technology in the storage network, such as a firewall, it is imperative that auditing on storage devices and applications be increased and enabled to the fullest extent. An unsuccessful attack left unnoticed can leave an organization crippled in security.

Integrity

Integrity is the assurance that unauthorized parties have not modified an entity. Furthermore, integrity confirms that the data has not been altered in transit from the source to the destination. It allows a network to depend on other security elements, such as authentication and authorization.

Encryption

Encryption is the process of protecting information from unauthorized access or modification by converting it into *cipher-text* that can be accessed only through appropriate credentials or keys. Encryption also allows an untrusted entity, such as a network, to be used without additional security elements for support. For example, using encryption with a VPN device allows remote users to use the untrusted Internet as a medium for business operations. Similarly, the use of encrypted protocols, such as Secure Shell (SSH), allows users to use in-band data networks for management functions.

Availability

Availability means ensuring that resources are on hand for legitimate users, applications, or network devices when requested. Security enables availability by ensuring that unauthorized user access or denial-of-service attacks will be unsuccessful in a given storage object. If an attacker's goal is simply to affect a loss of revenue for a given organization, stealing the organization's data is not the only method that can be used to accomplish this; simply making part of the storage network unavailable would result in loss of business operations, which equates to the loss of revenue.

Overall, an organization's storage environment must address these security elements in some form to support a strong storage security posture. The solution must also enable a process to grow and maintain the storage security posture over an undefined period of time. A strong presence in security at one point in time does not necessarily equate to a strong security presence in the future. A successful security plan will foster growth and provide stability to the storage network over a period of time.

Security Methods

Critical for security is a layered security model, also known as *defense in depth*. In the defense-in-depth model, layers of security are built in conjunction with one another in a complementary fashion. Many networks are built with the "M&M model,"—that is, hard on the outside and soft on the inside. This model crumbles after a single penetration of the outer perimeter. A defense-in-depth model would not crumble if any devices were subverted, such as the outer perimeter, because it would contain security layers behind each device.

An example of the defense-in-depth model in the storage network is an ACL on a storage node, such as an EMC or Network Appliance NAS head, that restricts or permits access according to IP address or subnet. These ACLs complement the required authentication and authorization procedures by the storage appliances and/or operating systems. With this model, if one or both of the security elements were subverted, the attackers would still be denied access if they were not sending the request from the correct IP address. Figure 25-1 illustrates the defense-in-depth model.

The defense-in-depth model allows organizations to protect their critical and sensitive storage data by eliminating any single point of failure. As shown in Figure 25-1, the model can be as simple as enabling security features on operating systems, storage switches, NAS heads, and even back-end disk arrays. Additionally, security controls can be placed on host-bus adapters (HBAs), network interface cards (NICs), client workstations, storage nodes, and storage applications. Because the storage industry has not yet come up with any pure security device, such as a firewall, security features need to be enabled and explored at other devices to support a defense-in-depth architecture to the fullest extent.

A single layer of security does not adequately protect an organization's storage network, proprietary data, or intellectual property. In addition, identified security weaknesses in one area, such as a storage application, can actually nullify strong security measures in other areas, such as a storage switch. A layered security model, in which security is emphasized at key segments throughout the storage network rather than one or two segments, supports a strong storage security posture.

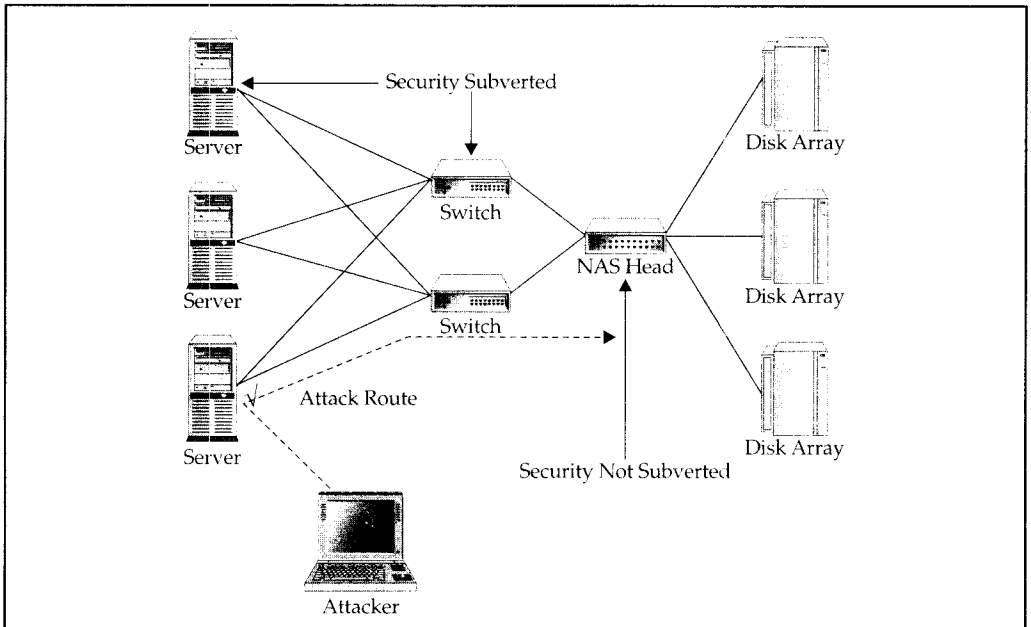

Figure 25-1. *The defense-in-depth model*

Storage Security Technology

The storage security technology industry is currently in the development phase, and in this process more questions than answers have been uncovered. However, emerging storage security organizations are designing solutions to help secure storage networks from both unauthorized users and accidental configuration mistakes. Existing storage vendors are also extending their product features to include security solutions. The following sections discuss the following storage security technologies that are under development as well as those that have already been deployed:

- Inline encryption
- Encryption at rest
- Key-based authentication
- Two-factor application authentication

Inline Encryption

Currently, the T11 committee is working on an initiative to incorporate IP Security (IPSec) technology established in the IP industry into the storage industry with iSCSI and Fibre Channel Security (FCSec). T11 has established a framework that will use ESP (Encapsulating Security Payload) in Fibre Channel layer 2 to protect frame data. ESP within the Fibre Channel frame will be able to provide confidentiality between two end nodes, such as a switch and an HBA.

Additionally, ESP will be able to enhance frame security by providing authentication, integrity, and anti-replay protection for each frame. This will prevent common attacks at the frame/packet layer, such as session hijacking, spoofing, and Man-In-The-Middle (MITM) attacks (described later in the chapter). The adoption of FCSec technology into the storage community will allow for greater flexibility for a storage network. Specifically, the key goals for the FCSec architecture are as follows:

- Node-to-node authentication
- Node-to-switch authentication
- Switch-to-switch authentication
- Frame-level encryption

Node-to-node authentication will allow each entity to authenticate frames transmitted between two nodes. This allows for the integrity of each frame between the two nodes and ensures against spoofing or replay attacks in the fabric.

Node-to-switch authentication will be similar to node-to-node authentication in that it ensures integrity in each frame. Additionally, node-to-switch authentication will allow switch management tools to contain added strength in their allocation mechanisms (for example, zone allocation).

Switch-to-switch authentication will allow each switch to ensure the integrity of data and management, such as Simple Name Server (SNS) information, to be protected against transport layer attacks. In addition, switch-to-switch authentication will reduce the possibility of a rogue switch gaining instant control of the entire fabric.

In addition to the authentication capabilities of FCSec at the frame level, confidentiality can be provided by using encryption of each frame as it exchanges data between two entities, including both node and switch entities.

Encryption at Rest

The encryption of data at rest uses a different approach than inline encryption. Standards committees are developing technologies such as FCSec to protect data and authentication for inline encryption. Encryption of data at rest is being addressed by new storage vendors and new storage products.

Data at rest is one of the most commonly overlooked risk areas in storage security. While most security efforts focus on protocols and architecture, the security of data at

rest tends to be incorrectly assumed as secure. Data on disk is often considered secure because it is kept deep inside the storage network, with multiple storage devices restricting access. However, these types of security assumptions expose the data both to accidental actions and unauthorized activity.

Several new organizations are designing products to encrypt the data before it reaches the disk. At that level of encryption, if the storage node is compromised, the attacker will not have readable access to any of the data residing on the disk. In the event of a configuration error, the data on the disk may not be fully exposed due to the encrypted format of the data.

Various storage devices protect the data at rest in different ways, but all of them try to accomplish the same goal of protection. The encryption solutions that encrypt data on disk protect against possible configuration errors or mistakes commonly found in storage networks. Configuration errors are the most common problems identified in storage networks today. If controls are not in place to mitigate possible configuration issues, the ability to compromise data in the storage network might be larger than expected. Additionally, since storage networks are complicated by nature, the likelihood of a configuration error incorrectly exposing the SAN (Storage Area Network) is higher than for many other types of networks.

Figures 25-2 and 25-3 demonstrate two scenarios of configuration errors that might lead to the compromise of data. Figure 25-2 shows how configuration errors can occur at the storage switch. Figure 25-3 shows how an encryption device can protect against storage configuration errors.

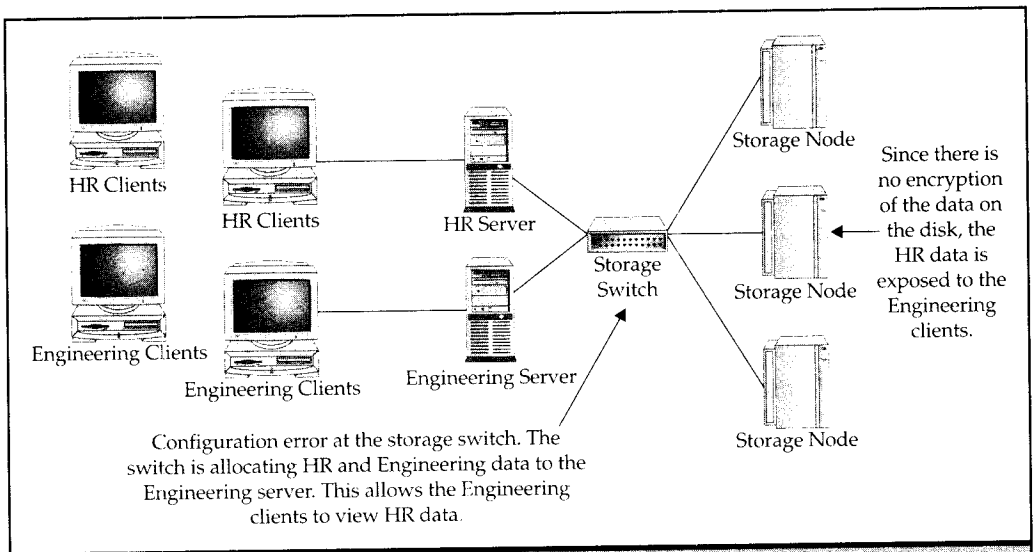

Figure 25-2. *Configuration errors at the storage switch*

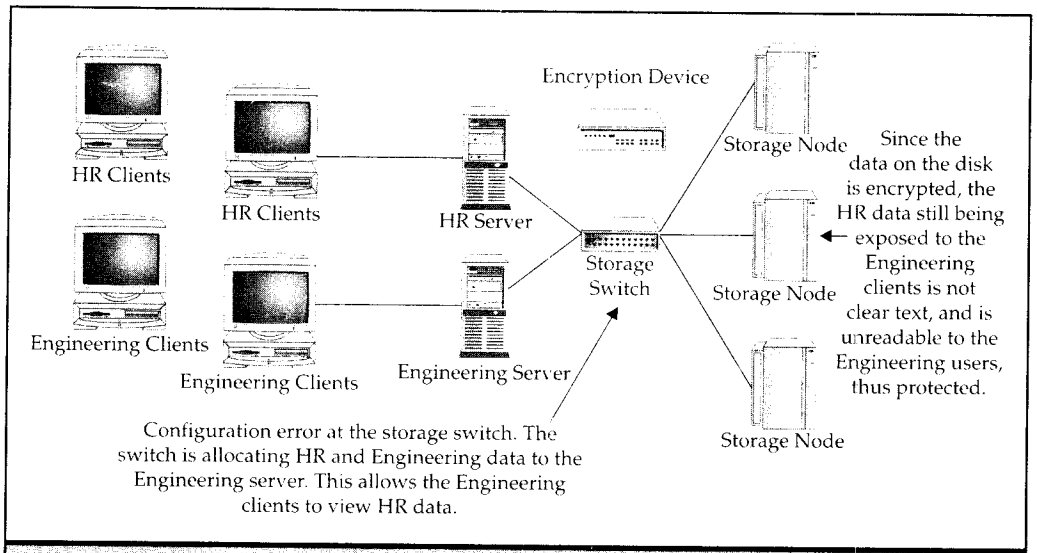

Figure 25-3. *Encryption protects against storage configuration errors.*

Key-Based Authentication

Fibre Channel-Generic Services-3 (FC-GS-3) is another technology that T11 standards organizations have developed. Key server technology is discussed in FC-GS-3 to enable key-based authentication, which supports the best practice of two-factor authentication in storage devices, such as switches, appliances, and clients. Unlike the FCSec project, the FC-GS-3 project is not solely a project about security. FC-GS-3 includes specifications dealing with management, directory service, time service, alias service, and other services. However, we will be referring to FC-GS-3 only in terms of the key server technology that it addresses.

The Key Distribution Service discussed in FC-GS-3 will provide a method for secure distribution of public and private keys in a storage network used for authentication, similar to a KDC (Kerberos Distribution Center) in IP networks. The security "key server" in the storage network would use a level of encryption to protect the keys and underlying protocols for verification purposes. By using a key server, data would be signed, providing a secure means of authenticating the original node as well as verifying the data was not modified during transit. This method of using a key server is often referred to as *CT (Common Transport) authentication.*

CT authentication is the process of authenticating requests between a client and a server. The CT authentication mechanism creates a hash based on an algorithm and secret key to represent the message signature associated with the sender. The hash is transferred as an extended preamble to the receiver. Upon receiving the hash, the receiver computes the hash before authenticating it, using the same algorithm and secret key, to make sure the hashes match.

The key server in the storage network would contain a copy of each node's unique distribution key. The key server would then authorize the security connections between two nodes and would generate a key to use for authentication. This authentication key is then used by one node to authenticate to the other node. The single point of failure in the specification is the key server. If the key server does not contain the proper amount of security, the authentication model that the key server supports cannot be considered secure. Figure 25-4 shows an example of an authentication procedure using FC-GS-3.

A practical example of FC-GS-3 would be the authentication of a World Wide Name (WWN) of an HBA in a storage network. In the "Fibre Channel SAN Security" section of this chapter, the security problems of WWN are discussed. It would not be best practice to use a WWN as the sole entity for authentication. The process of using FS-GS-3 to identify HBAs uniquely would allow the use of WWN without concerns of spoofing or replay attacks being performed successfully. Since WWNs are not reliable as a secure means of host identification, using a key server will minimize the risk of spoofing or session hijacking and allow the use of WWN with an acceptable level of integrity in the authentication process.

Currently, many switch vendors, such as Brocade and McData, are using a hybrid of the FC-GS-3 specification by using switch-to-switch key authentication. In this approach, each switch contains a unique key to authenticate to another switch. Instead of using a third key server device, the two nodes are directly authenticated to each other using their own keys. While this approach may not be as secure as having a key server, it does prevent a variety of attacks, primarily the ability for a rogue switch to take control of a given fabric, its simple name server, and its zoning architecture.

Two-Factor Application Authentication

Two-factor authentication for key applications, such as storage management applications, is becoming a significant requirement for protecting storage networks. The fact that storage management applications can subvert the majority of other security controls within the storage environment makes the management applications a significant target for most attackers.

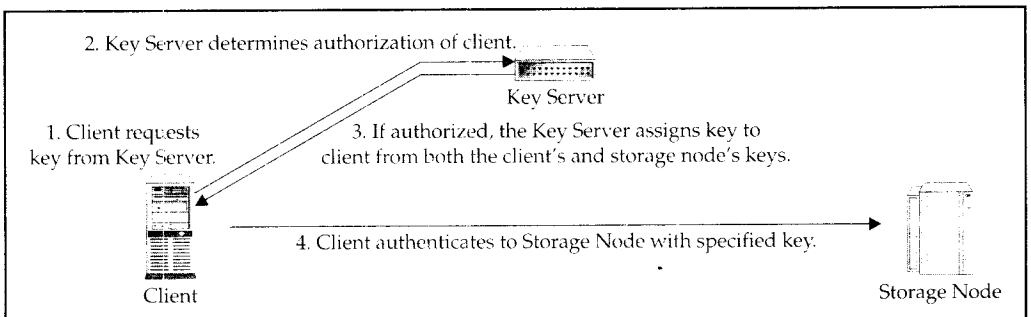

Figure 25-4. *Key distribution service*

Currently, many management applications rely on a username/password combination for authentication. As best practice, a single username and password should not be the only security entity protecting a storage network. For example, an unauthorized user should not be able to plug into a network with his favorite backup program and successfully manage storage agents because he was able to guess the username/password combination or spoof an IP address.

Requiring more than a username and password or an IP address for authentication is possible. To address this issue, certain vendors have included a second level of authentication within the management application itself. These management applications use public and private keys for authentication with authenticated Secure Sockets Layer (SSL) or Kerberos technology to mitigate some of the exposures described so far. Therefore, if an unauthorized user has plugged into the network and successfully guessed a password or spoofed a management IP address, the attacker's rogue backup software would not be able to gain access to the storage agents since a second level of authentication is required. The second level of authentication required from a client application to the server application and/or agents could use tokens, digital certificates, Kerberos tickets, public keys, or other forms of security. This type of application design may utilize usernames/passwords and IP addresses, but it also utilizes credentials implicit in the client application to authenticate to the server and/or storage agents.

Note *Tokens, digital certificates, Kerberos tickets, and public keys can be used for authentication. All these security entities can be delivered digitally to authenticate a sending node to a receiving node, such as storage management software, in place of or in addition to username/password authentication.*

Storage Security Challenges

As the storage industry grows, with both SAN and NAS architectures, the requirements for storage security will also grow. Storage security will need to accomplish a number of milestones and challenges to become mainstream in storage networks.

Unfortunately, a typical "milestone" that uncovers the need for security in a particular industry is often some type of successful and well-publicized attack. For example, wireless security technology, specifically 802.11b, would not have been viewed as such a severe security issue if hackers were not easily able to break Wireless Equivalency Protocol (WEP) encryption. Despite the weakness of 802.11b, many organizations did not invest in wireless security until well after the fact of the discovery of the security problem.

While a break-in might have encouraged organizations to invest in wireless security, to approach security after a breach costs an organization significantly more time, money, and resources. Unlike the wireless industry, the storage industry has more to protect than a possible credit card number going over the wireless network in semi-clear-text. The storage system has an entire organization's intellectual property, sensitive data, and possibly customer information to protect and maintain. Therefore, a successful storage security attack would be more severe than a typical network attack.

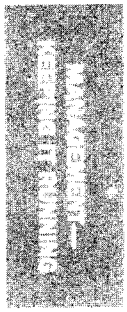

Another major challenge in storage security is industry acceptance. Many storage administrators and experts believe security issues in Fibre Channel, storage devices, and storage architectures are not important issues, for a variety of reasons. However, because storage security has significantly more to protect than other types of networks, the storage industry should be more aggressive in fixing and solving security issues.

For example, consider the security problems encountered in Cisco routers and switches in the early 1990s. These routing and switching devices were developed based on availability and functionality. However, due to the security weakness of IPv4 combined with IP device weaknesses, the devices were easy to subvert and/or compromise. The same argument can be made for storage products that focus on bandwidth or speed and not necessarily security. The weakness of IPv4 and Fibre Channel, combined with potential storage device weaknesses, may well lead to the same results that the IP industry faced.

Industry must face the following key storage security challenges:

- Products and solutions
- Encryption
- Security standards and specifications
- Bandwidth/functionality tradeoffs

Products and Solutions

The first and major storage security challenge is the lack of security product solutions and storage security features in current products. As described earlier, no pure security products are currently available for storage networks to permit or deny access to certain storage nodes. Furthermore, a firewall or similar device would have to support 2 Gbps throughput to hold any type of acceptance and success in the storage industry. Currently, only certain firewalls can support near gigabit throughput, and these are available only for IP networks. There is no such thing as a Fibre Channel firewall.

An additional storage security challenge is the lack of security features in storage products today and the quality of security features that are included in existing products. A good example to demonstrate this problem is the zoning and Logical Unit Number (LUN) masking capabilities in Fibre Channel networks. The zoning capabilities in Fibre Channel switches and LUN masking capabilities in storage nodes were originally developed for segmentation, not security. While the segmentation tools are good for secondary security, they are not adequate to use as primary security tools. In addition, using zoning and LUN masking as the only security tools does not support a best practice storage architecture. In fact, using zoning and LUN masking as security tools is similar to using virtual LANs (VLANs) as the only security tools in an IP network—which would exclude the use of better security devices, such as firewalls, router ACLs, and VPN devices to protect the network. Cisco has stated many times that VLANs should not be viewed as security tools, but rather as segmentation tools. VLAN's *hopping* ability (the ability to jump across one VLAN to another) has been demonstrated by security professionals in a variety of tests.

In addition, most storage products provide only Telnet and web (HTTP) capabilities for a management, both of which are clear-text protocols that can easily be sniffed. The use of encrypted management protocols, such as Secure Shell (SSH) or encrypted web (HTTPS), is still in the process of being adopted.

Encryption

Encryption technology is another major challenge in the storage industry. Inline encryption in both IP and Fibre Channel mediums is difficult to implement without significant bandwidth penalties. IP networks that have bandwidth capabilities of Gbps are reduced to Mbps transfers once encryption technology is in place. Similarly, Fibre Channel networks with 1.0 to 2.0 Gbps capacities would also be reduced to probably less than half that amount. Considering the fact that Fibre Channel is often deployed specifically because of bandwidth capacities, the fact that encryption would directly negate those capabilities is a significant security problem for storage security engineers.

As mentioned, encryption of data at rest is another security challenge. The major challenge of data at rest is interoperability with different types of storage devices. In addition, interoperability of all types of storage appliances should be a requirement supported by standards bodies; unfortunately, competing vendors may not share this perspective.

Security Standards and Specifications

Lack of security standards and specifications is also a significant problem in the storage security industry. A good example is zoning definitions in Fibre Channel switches. The term *hard zoning* is defined one way for certain Fibre Channel switch organizations and another way for other organizations. Some switch vendors refer to *hard zoning* as the ac of locking physical port numbers to a particular zone on a Fibre Channel switch. Other switch vendors refer to *hard zoning* as a routing tool in which routes will not be broadcasted to non-participants of the zone. This disparity between definitions is a standards problem. The lack of storage security standards, from encryption specifications to zone definitions, leads to technologies that do not coexist easily across various storage vendors. In addition, the lack of security standards and specifications will lead to competing vendors developing technology that will most likely not interoperate with one another's products. This will make an end user's storage security architecture a frustrating process to design and implement.

Standards in storage security need to define terms and protocol specifications before storage vendors grow impatient. For storage security features, tools, and products to be successfully adopted by end users, a clear classification of any type of architecture needs to be in place. Standards need to be developed in the following areas:

- Encryption standards
- In-line and at-rest standards
- Standardization of storage security terminology and definitions
- Authentication standards

Bandwidth/Functionality Tradeoffs

Bandwidth and functionality tradeoffs are key challenges in creating storage security. Any type of functionality or bandwidth loss incurred to gain added security will not likely be successful in the storage industry, especially since many storage networks are deployed for the mere purpose of bandwidth and functionality offerings. Any type of solution must have a minimal affect on storage functionally while demonstrating clear interoperability with existing storage features.

Fibre Channel SAN Security

Fibre Channel SAN security deals with a variety of issues ranging from Fibre Channel frame authentication to Fibre Channel switches. Fibre Channel SAN is an emerging technology that still has many items under development, such as management, virtualization and interoperability. Furthermore, although some progress has been made regarding security in Fibre Channel networks, several opportunities exist for growth.

In this section, we discuss some of the security problems in Fibre Channel SANs, at both the frame and device levels. We continue to address some solutions to these security problems and best practices. The following items will be discussed:

- Zoning
- WWN spoofing
- Fibre Channel (FC) frame weaknesses
- Sequence ID and control number
- Disruption of flow control
- MITM attacks (SNS pollution)
- E_port replication
- LUN masking

Zoning

Zones are to Fibre Channel switches as VLANs are to IP switches. Zones are used to segment the fabric into groups, in which certain nodes are a part of a zone and have access to other nodes in the same zone. Storage nodes could be a part of multiple zones or part of a single zone.

Two types of zoning are used, soft zoning and hard zoning. In soft zoning, the SNS in a switch uses an HBA's WWN for zoning. An SNS is a table located in each switch and shared among switches in the same fabric, which separates WWN into the correct zones. Figure 25-5 shows an example of soft zoning. In the figure, WWN 9382108xxxx, 3859658xxxx, and 3582968xxxx have full access to each other. However, if any of these WWNs try to access zone B members, such as 0038283xxxx, access would be denied.

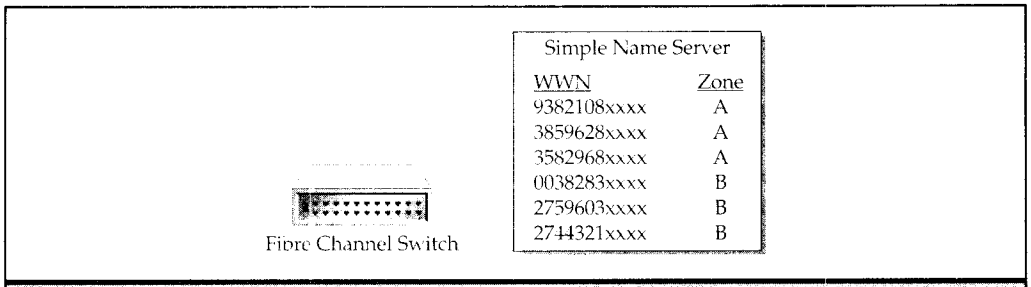

Figure 25-5. *Soft zoning in SANs*

Hard zoning is not only different from soft zoning but is another term that has multiple definitions. One definition for hard zoning is "the act of locking physical ports to certain zones." For example, physical ports 1 through 4 would be zone A, physical ports 5 through 8 would be zone B, and so on. If a WWN was connected to port 2 and tried to access port 6 in zone B, it would be denied access. Figure 25-6 shows an example of this definition of hard zoning.

Another definition of hard zoning is a routing procedure. If route-based zoning is implemented, certain routes will be publicized for certain storage nodes. For example, if node 1 were allowed to access node A, node 1 would be notified about the route to node A. If node 1 were not allowed to access node A, node 1 would not be notified about node A. However, if node 1 knew the existence and route to node A, hard zoning, according to this definition, would not prevent node 1 from this access. Hard zoning, in this reference, is not a restrictive tool; rather, it's an information tool.

World Wide Name Spoofing

WWN is used to identify an HBA in a storage area network. WWN is used for authorization of data from one particular zone to another. Additionally, WWNs are used in soft zoning procedures on storage switches to separate certain servers and data. Soft zoning separates WWNs into different zones. Because a switch contains an SNS table that matches up each WWN to a particular zone, a particular WWN would be granted or denied access to another storage node based on the results of the SNS zone table.

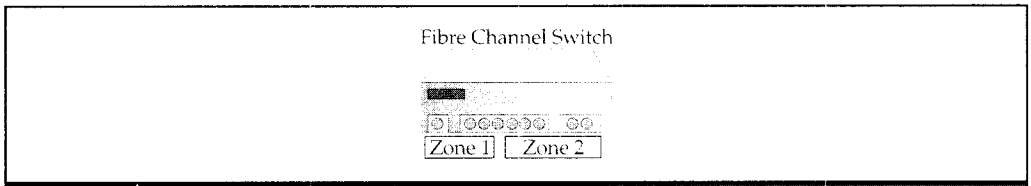

Figure 25-6. *Hard zoning in SANs*

WWN spoofing is a high-risk problem in storage area networks because WWNs are often used as the only tool for granting authorization. WWN spoofing is as simple as loading up the device drivers included with the HBA and changing the WWN. An attacker can then spoof a WWN and gain access to data. The switches that contain the zone information would grant access to this spoofed WWN because it authorizes the WWN and does not authenticate it. Figure 25-7 illustrates this type of attack.

As best practice, hard-zoning procedures can be used to deter spoofing problems. The physical definition of hard zoning would not use WWN for authorization, but rather the physical port numbers. If an attacker spoofs another WWN, the hard-zoning switch would still not grant it access since the physical port on the switch is trying to gain access to another port to which it does not have authorization, regardless of what WWN it claims to be. However, using the route-based definition of hard zoning, WWN spoofing would not be a good alternative. Because the route-based definition is not a restricted tool, but rather a tool used to deliver routing methods, a hard-zoning switch could still be subverted if the attacker knew the address and the route to its target. Therefore, if a WWN were spoofed with route-based hard zoning, access would be permitted.

Frame Weaknesses

Fibre Channel architecture contains five different layers, numbered 0 to 4. Despite the differences between IP and Fibre Channel, Fibre Channel frames contain weaknesses that are similar to current weaknesses in IPv4 packets. These IPv4 weaknesses have been turned into vulnerabilities and exploited at a variety of levels. The weaknesses in Fibre Channel frames are specifically in Fibre Channel layer 2, known as the framing

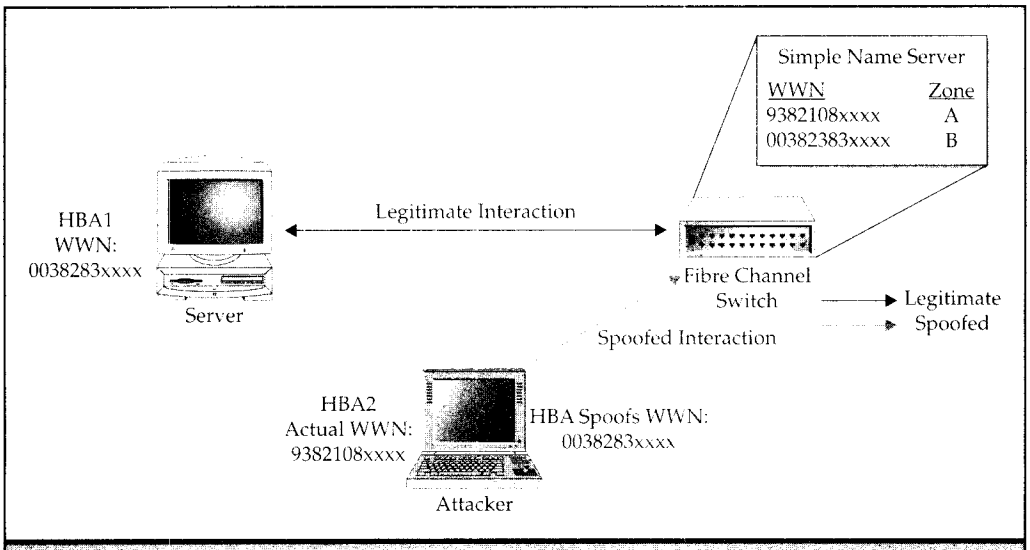

Figure 25-7. *SAN World Wide Name table spoofing*

protocol/flow control layer. Fibre Channel layer 2 is where most security issues have been identified. Fibre Channel layer 2 contains the frame header, which contains the 24-bit source address of the frame, the 24-bit destination address of the frame, the sequence control number (Seq_Cnt), sequence ID (Seq_ID), and the exchange information (Exchange_ID).

As shown as Figure 25-8, layer 2 contains the sequence control number and the sequence ID. A *sequence* is a series of one or more related frames transmitted unidirectionally from one port to another. The sequence series is responsible for keeping the connection alive between two nodes. A given frame must be a part of some sequence between two end nodes. Frames within the same sequence have the same sequence ID (Seq_ID) in the frame header, which is kept to identify frames that belong in the same transmission process between two nodes.

In addition, as each frame is transmitted in a sequence, the sequence count number (SEQ_CNT) is incremented by one. Therefore, that transmission can remain initiated and frames can remain ordered by the incrementing process. However, the fact that the sequence ID is a constant value, thus predictable, and the sequence control number is incremented by a predictable number, which is one, makes the sequence series a predictable value. Because the sequence series is responsible for maintaining the session between two nodes, an attacker would be able to launch a session-hijacking attack and take control of the session.

This attack was made popular with IPv4 packets with similar problems of predictability in Initial Sequences Numbers (ISNs). In ISNs, the hijacking packet would simply need to guess the predictable sequence, similar to the Fibre Channel frame, and take control of a management or data session. Figure 25-9 illustrates this attack at a high level.

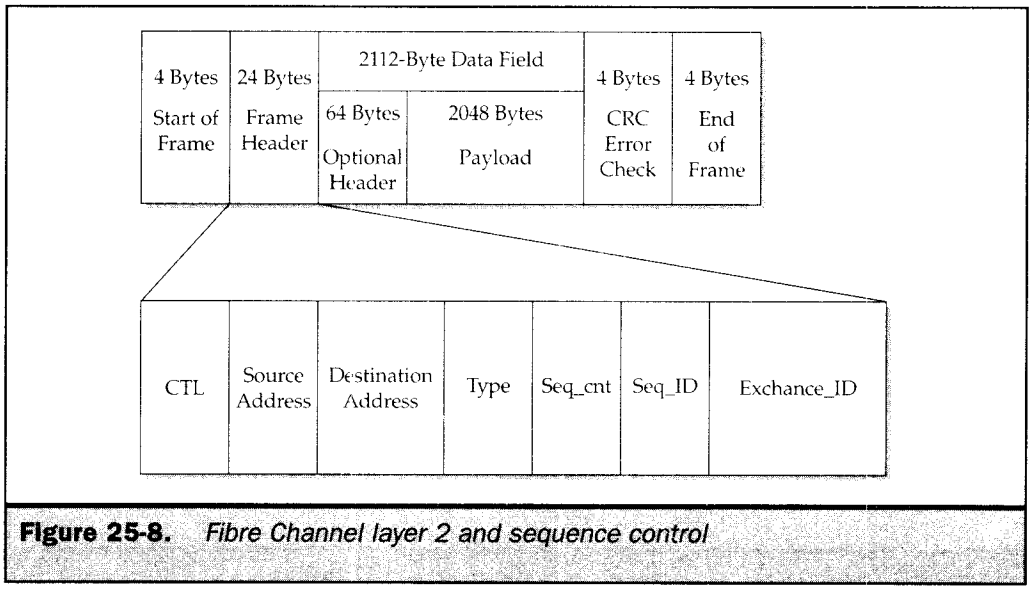

Figure 25-8. *Fibre Channel layer 2 and sequence control*

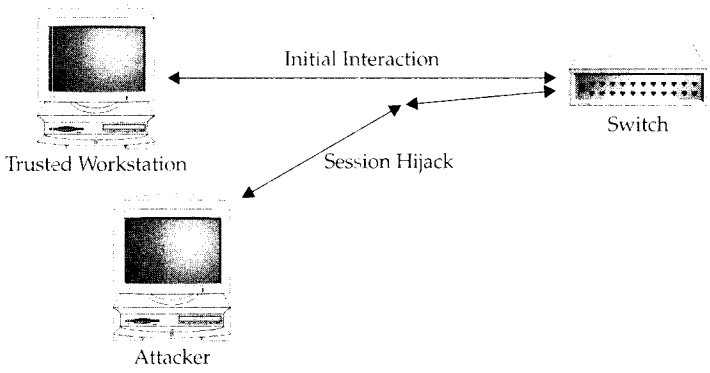

Figure 25-9. *Fibre Channel frame header attack*

The use of predictable session information leaves Fibre Channel frames vulnerable. Additionally, in-band management protocols, such as SCSI Enclosure Services (SES), can allow management sessions to be hijacked and exposed to unauthorized users. As in-band management methods and applications grow, the weaknesses of the session information in Fibre Channel frames could amplify and result in loss of data.

The solutions for predictable sequence numbers in Fibre Channel frames needs to be similar to the solutions for IPv4 vulnerability. Operating system settings and encryption technology enabled IPv4 weaknesses to reduce the likelihood of successful hijacking attacks. Similarly, Fibre Channel frames need to make session information unpredictable. This change can come from the HBA drivers that are installed on the operating systems or within the fabric itself. Either solution should be able to address the static nature of sequence IDs and the predictably of sequence control numbers.

While encryption may be a solution for an IPv4 network, using encryption as a solution for Fibre Channel frames is not as easy, since encryption involves extra overhead and bandwidth costs that may not be acceptable in SANs. As a mitigating solution to encryption, separate out-of-band management networks should be deployed to eliminate the possibility of an attacker hijacking management frames to access storage devices and their data.

Disruption of Flow Control

Flow control is responsible for message flow between two Fibre Channel nodes. In Fibre Channel networks, devices transmit frames only when each device is ready and able to accept them. Before these devices can send and receive frames, they must be logged in to each other and must have established a *credit level* to send and receive the correct amount of frames. The establishment of credit from one storage node to the other is conducted at the Exchange_ID level of the frame header. This *credit* refers to the number of frames a device can receive at a time. This value is exchanged with another device during login, so each node knows how many frames the other node may receive.

The problem with flow control is that Fibre Channel frames are unauthenticated. An attacker can therefore send out incorrect Exchange_ID information to a variety of addresses, thus creating a temporary denial-of-service between two nodes. For example, because frames are unauthenticated, an attacker could replace valid frames with incorrect frames that contain a higher Exchange_ID number between two nodes. Because one node will now be sending more frames than can be received, the receiving node will be unable to accept the frames properly and will not be able to respond appropriately. Similarly, instead of creating a higher Exchange_ID number, an attacker could lower the number for the Exchange_ID information, thus leaving one node waiting for more frames to process. This would result in the frames being passed slower from one node to the other, delaying the process of data communication.

The solution for the disruption of flow control is authenticated frames. Because the process of exchanging ID information requires logging in to the other node, authentication of the node should also be addressed at this point. However, while this attack is successful for the disruption of flow control, the disruption is actually quite minimal and the fabric has the ability to correct itself quickly.

Man-in-the-Middle Attacks

MITM attacks have plagued IPv4 networks for some time. The problem in IPv4 is the use of Address Resolution Protocol (ARP) packets. ARP is a protocol that matches a machine's NIC address—known as the *Media Access Control (MAC) address* of the machine—to an IP address. However, both the IP address on a network and the MAC address can be easily spoofed. Because no authentication of ARP packets is required in IPv4, a spoofed IP address to another MAC address can redirect information to unauthorized users. The attacker could send out ARP packets that match their MAC addresses to a target's IP address. Therefore, when data is sent to the target, the data will be sent to the attacker since the target's IP address matches the attacker's MAC address.

This attack can be partially replicated in Fibre Channel networks. Fibre Channel switches include a table, referred to as the Simple Name Server. The SNS matches up WWN (similar to a MAC address) to the 24-bit address of the storage node. Since Fibre Channel frames are also unauthenticated, an attacker can use any traffic analyzer program to generate frames in the fabric with a target's 24-bit address and the attacker's WWN, thus changing the SNS information. The SNS would have an updated entry for the 24-bit address with a new WWN. When traffic is being sent to the target's 24-bit address, it would go to the attacker's machine since it matches up to the attacker's WWN.

As shown in Figure 25-10, an attacker sends out a modified frame to xFFFFFE to log in to the fabric (FLOGI). The modified frame has a source address of another trusted entity on the fabric, such as another trusted switch, and the WWN of the attacker. The fabric assumes that the attacker is now the legitimate host since the switch's 24-bit address is matched to the attacker's WWN. All frames destined for the real node are passed to the attacker and then to the legitimate node.

An MITM attack is simply *polluting* the SNS table with switches. Polluting the SNS could also be conducted when joining the fabric. An N_Port would send a FLOGI (fabric login) to the well-known address of xFFFFFE (similar to a broadcast in the IPv4). The switch receives the frame at xFFFFFE and returns an accept frame (ACC).

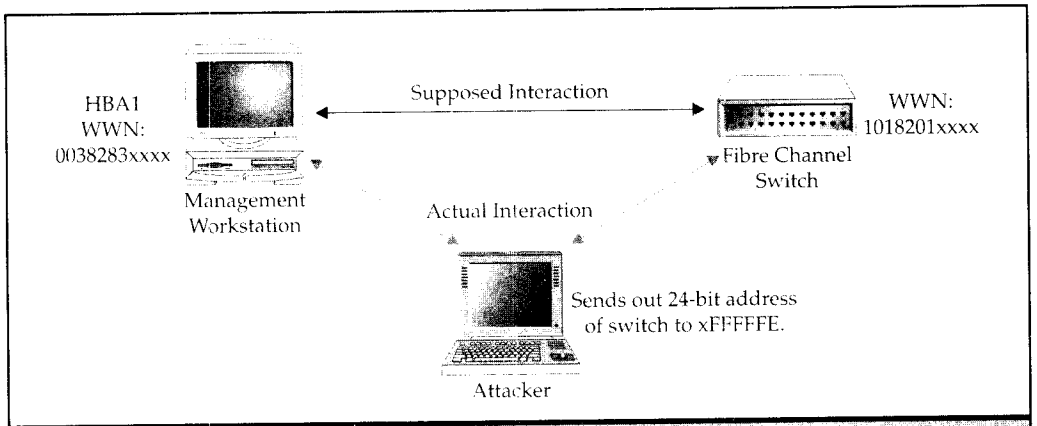

Figure 25-10. *Man-in-the-Middle attack on SANs*

The service information would then be exchanged. Because no validation is required to receive an accept frame, an attacker could send a modified 24-bit address to xFFFFFE to attempt to corrupt the SNS table in the switch. As soon as the attacker receives the accept frame, the attacker knows that the SNS table has been modified.

LUN Masking

A *Logical Unit Number* (LUN) is a unique identifier that is used to differentiate between devices, including physical or virtual disks volumes. LUN masking allows nodes to be aware only of other files or data blocks that they are authorized to access. This is implemented by masking off LUNs that may be available. Similar to zoning, LUN masking is a segmentation tool that has been used for security. LUN masking can be implemented at three levels: at the HBA, at the storage controller itself, and using a third-party device in the storage network.

HBA-based LUN masking uses driver utilities that map WWN to LUNs. The HBA drivers contain masking features that allow the WWN to view only a certain set of authorized hosts. Masking at the HBA level offers easy implementation since a third-party device would not be required in the SAN; however, it does require additional system administration resources since each HBA in a storage network would have to be individually configured. This solution may work well for small storage environments, but it is virtually unmanageable in large SANs. In addition, since LUN masking occurs at the HBA level, if the operating system were compromised, an attacker could change the masking information without any other security controls preventing this. This would allow an attacker to load up device drivers and change the allocation to include as many LUNs as desired, thereby negating any security gained by the masking rules.

Conducting LUN masking at the storage controller is another way of implementing LUN security. The storage controller would map access privileges from each HBA's WWN to the authorized LUN. LUN masking at the storage controller level is probably

the best alternative for most SANs, since it has the ability to scale in large storage networks and does not require an extra storage device to add undue overhead. In this scenario, however, LUN masking is dispersed across many storage devices instead of being centrally managed at one device.

The third method of implementing LUN security is using a third-party storage device for allocation. The third-party device would be placed between the client nodes and the storage nodes. The device would handle requests for LUNs before they proceeded to the storage controller. Once the requests have been made, the device would grant or deny access. If the request is granted, it is forwarded to the storage devices. This solution scales appropriately to enterprise-sized storage networks and provides a central resource for LUN management. However, it places yet another storage device in the storage network. Furthermore, it adds another layer of overhead from the client node to the storage controller.

Table 25-1 compares and contrasts each type of LUN masking technique and the pros and cons of each.

LUN Masking Type	Pros	Cons
Storage controller	Scaleable to large SANs. No undue overhead with another storage device. Security cannot be subverted by an operating system compromise.	Management of LUN allocation is dispersed throughout the SAN instead of being centrally managed.
Third-party device	Scaleable to large SANs. Centralized management of LUN allocation. Security cannot be subverted by an operating system compromise.	Added overhead by placing another storage device in the SAN.
HBA level	No undue overhead with another storage device.	Only possible in small SANs (not scaleable). Management of LUN allocation is dispersed throughout the SAN instead of being centrally managed. Security can be subverted by an operating system compromise.

Table 25-1. *LUN Masking Techniques Compared*

E_Port Replication

E_port (FC extenstion ports) replication is a feature in many Fibre Channel switches. Each port on a switch has a particular function (refer to Chapters 14 and 15 for additional information on FC ports). Each type of port holds a different function, whether it connects two client nodes together or connects two switches together. An E_port, for example, connects two switches together, enabling the connected switches to exchange routing and zoning information.

Because no authentication is required, E_port replication security is not guaranteed. If a storage server provider (SSP) were extending its fabric to a client using unauthenticated E_port replication, the simple act of plugging in the storage provider's E_port to the client E_port would exchange routing and zoning information for the entire fabric, not just for a specific zone. This would allow an unauthorized switch to access information.

To prevent accidental E_port replication, port-type locking and key-based authentication between switches can be used. Products from major switch vendors, including Brocade, McData, and Qlogic, can lock port types rather easily. For example, if physical ports 1–11 connect only client nodes to storage nodes, the port type can be locked to n-type. Similarly, if only port 12 is used as an E_port, it can be locked using port-type locking. This would prevent an attacker from replacing a client node with a switch and have an E_port exchange information from the two switches automatically.

Key-based authentication between the switches is also a best practice to prevent unauthorized E_port replication. Key-based authentication would require each switch to authenticate before being added to the fabric. This would disable any type of communication to and from the switch until validation keys or certificates have been exchanged from the joining switch to the existing switches.

NAS Security

Network attached storage (NAS) security addresses security at a device level. Traditionally, NAS devices, such as filters, are marketed as pure storage devices. Even though many are also deployed as storage devices, most, if not all, have built-in operating systems, usually based on some sort of UNIX flavor, that include services and protocols that need to be secured in a way similar to any other type of device on a network. Additionally, NAS devices are usually based on IP technology, using Ethernet or gigabit Ethernet, which is threatened by the IPv4 exposures. The two most common implementations of the NAS are Common Internet File System (CIFS) and Network File System (NFS). The following section focuses on security concerns of CIFS and NFS NAS devices.

Common Internet File System

CIFS is a standard protocol that is used for data sharing on remote storage devices on IP networks. CIFS security focuses on two security elements: authentication and authorization. Two types of access rights can be used for CIFS, each with its strengths and weaknesses: these are share-level authentication and user-level authentication.

Share-level authentication is based on share points and passwords. Because share points can be made available to multiple users with required authentication of each user, no accountability or authentication is required on a user-by-user basis.

Following are its specific vulnerabilities:

- User accountability or authentication is not required by individual users (no record of the access).
- A single password, which is shared by design, is responsible for the entire share point.
- Passwords must be changed every time an employee is terminated or resigns.
- Passwords are transmitted in clear-text (however, all Windows 2000/XP machines and NT 4.0 service pack 3 machines support non-plaintext password authentication).

Per-user password authentication for a NAS appliance is best practice in terms of security. Some key options are required to attain a desired level of security. Both LAN Manager and NT LAN Manager (NTLM) have associated severe security weakness, making them capable of being reversed engineered and reduced to the equivalent of a clear-text protocol. A good solution is to use Windows support of NTLM v2.

In addition to NTLM v2 for username/password protection, Server Message Block (SMB) signing should be used for CIFS SMB communication. SMB signing places a digital security signature into each SMB packet, which is then verified by both the client and the server. Without SMB signing, the server authenticates the client, but the client never truly authenticates the server, so mutual authentication is not completed. As best practice, mutual authentication should be enabled, especially in an IP network, to eliminate the possibility of IPv4 attacks.

Network File System

NFS is also a standard protocol used for data sharing on remote storage devices on an IP network. NFS is basically the equivalent of CIFS, but in the UNIX environment. NFS is used to allow a computer to mount (share) files over a network. NFS began as a User Datagram Protocol (UDP), but many Transmission Control Protocol (TCP) implementations exist today.

Like CIFS, most NFS implementations have security weaknesses:

- NFS communication is not encrypted (all traffic is in clear-text).
- NFS clients (hosts) have a limited level of authentication.
- NFS users are not easily authenticated.

As best practice, NFS storage devices should not solely grant or deny access to NFS mounts based on the host name of the client. IPv4 attacks allow host name and IP spoofing to be accomplished quite easily, possibly exposing the entire storage appliance to unauthorized users. Furthermore, this type of NFS mounting does not require passwords for authentication, exposing the data significantly.

Furthermore, per-user password authentication for an NFS appliance is best practice in terms of security. However, some key options need to be enabled to secure this method properly. Obviously, clear-text communication should not be used from an NFS client to NFS appliance. Certain NFS appliance vendors do support RSA/DES encryption with NFS communication. This will eliminate the transfer of sensitive information, such as usernames, passwords, NFS mounts, file handles, and contents of data.

Note *RSA (Rivest, Shamir, Adleman) and DES (Data Encryption Standard) are algorithms used for encryption and authentication systems.*

In addition, Kerberos (v5) is supported by many appliance vendors, which significantly reduces the risk of username and password credentials being compromised or replayed. Here are some specific solutions:

- Encrypt NFS communication with either RSA/DES algorithm (supported by most vendors).
- Do not authenticate based solely on host name.
- Enforce Kerberos (v5) for username/password authentication.

Best Practices

As mentioned, a defense-in-depth architecture protects a storage network's critical data. However, a layered security model can be subverted if simple best practices are not followed when designing storage security. For example, consider the home security example. The homeowner had spent several thousand dollars buying locks and security systems for her home. However, a simple best practice, such as not locking the house garage door, could negate all the other strong security measures that she had placed in the home.

Simple best practices in storage security can help prevent many of the common security attacks that unauthorized users will attempt. Adhering to the majority of these best practices can leave an attacker frustrated or bored and motivate him to move to another target or stop the attack altogether.

Following are some basic best practices to follow:

- Strong passwords
- Configuration
- Strong management
- Clear-text management
- Remote connectivity
- Operating system security
- Storage appliance services

Passwords

Using strong passwords and changing default passwords on storage devices is good security practice. Whether it is an EMC Cellera, a Network Appliance filter, or a Brocade switch, basic and simple passwords can be easily obtained and guessed. In addition, because most storage appliances have some sort of web management capability that is enabled by default, accessing the storage appliance and guessing passwords is not difficult for an accomplished attacker. Several tools exist to brute-force web-based authentication forms that can attempt more than 100 passwords in only a few minutes.

Weak passwords lead to a bigger problem, subscribing to the myth that storage networks are isolated and cannot be accessed from regular corporate networks. Many storage administrators believe that the storage network is not easily accessible. However, one of the primary purposes of a storage network is to connect devices on the network to back up data. Therefore, since these connections are permitted, storage networks *are* accessible, either indirectly or directly.

Using weak passwords can also significantly cripple storage security. While many organizations change weak passwords from *prom* or *password* to *admin* or *manage*, both *admin* and *manage* are common passwords that are contained in most tools that attackers use to crack accounts.

Following is a list of some common passwords. If these passwords are used in your storage network, consider changing your password policy for storage devices:

password	monitor	*<switch vendor>*	Config
admin	temp	*<company name>*	Test
manage	root	Letmein	secret
prom	backup	Secureme	keepout
filer	KuSuM	abcd1234	Test123
netcache	momanddad	Money	green

Configuration

Good configuration practices can build a solid security foundation in a storage environment. Furthermore, poor configuration practices, an abundance of configuration errors, or the lack of configuration options can significantly decrease the security posture in a storage network.

Most often, an attacker will complete a successful attack not because of some elite security problem that she has discovered on the spot, but rather because of poor configuration decisions and/or configuration errors that lead to wide open and unprotected devices. For example, a minor configuration error, such as a storage

administrator allocating the incorrect LUN to a given WWN, can expose data that is not authorized. A major configuration error, such as a storage appliance backing up an entire file system with world-readable access permissions, can also expose all data to unauthorized users. Both have a significant impact on the security posture of the storage network, despite one being major and one being minor.

The lack of understanding of security options and/or features in existing storage products may also lead to poor configuration decisions. For example, many end users are not aware of several security options that are readily available on their storage appliances and/or storage switches today. Not only are these security features not well advertised by vendors, but end users do understand how to use them. This combination often leads to end users making configuration decisions that don't work for security and not considering configuration decisions that are secure.

Management

Management is a critical component to storage security architecture. Since functional storage networks rely heavily on storage management practices, it is imperative that a strong management environment exist. Because storage network management contains a significant amount of control, a compromise of any management entity would give an attacker a considerable amount of privileges.

Protecting management interfaces for a storage network is a significant best practice. It is possible to attack many storage solutions through the Web, Telnet, or SNMP management interfaces. In many cases, gaining access to management interfaces is as simple a wire sniffing, session hijacking, and replay attacks. In other cases, it is as simple as loading an appropriate management program and logging in to the management application.

Many storage devices and applications rely on unsafe and clear-text protocols (such as Telnet, SNMP, FTP, CIFS, or NFS) to communicate both data and management commands to and from storage devices. Support for encrypted data channels, such as SSH, have not been adopted universally. Issues also exist with the networking devices that are used to support the storage environment. These devices are subject to attack and may be managed in unsafe ways.

In addition to using unsecure protocols, many organizations make a common mistake by plugging their management interfaces on storage devices, such as switches and storage appliances, into the internal corporate network. Connecting the management interface of a storage appliance into the internal LAN potentially gives any internal employee, external VPN user, third-party business partner, or external onsite contractor/consultant the ability to connect to the device and attempt to log in. In addition, if the management methods use clear-text technology, such as Telnet or web browsers (HTTP), the exposure is amplified. Table 25-2 lists a set of best practices for storage security.

Risk	Solution
Insecure channels for management	Use encrypted management protocols such as SSH and SSL. SSH with port forwarding can be used with many storage applications. HTTPS (SSL) is available on some storage devices that offer web management. Also, HTTPS provides the ability to wrap clear-text web management, such as HTTP, around an SSL tunnel.
Hard coded (a user name/password that does not change) or weak username and passwords	Enforce two-factor authentication for all management to reduce the likelihood of compromise due to a username and password being lost.
Shared channels for management	Do not plug management connections to normal, internal networks. Segment the management network by isolating it from any other network in the organization, especially the internal LAN.
Shared accounts	When possible, limit authenticated users to perform functions within the job responsibility (e.g., backup administrators versus storage administrators). Avoid complete authorization of all management functions to every authenticated user.
Share applications	When possible, use filtering to restrict management of storage devices to a limited number of management clients. Filtering can occur at the operating system or application level, limiting the accessibility of any users loading a management application and managing the storage network.

Table 25-2. *Best practices for storage security*

Remote Connectivity

Remote connectivity, such as dial-in lines, modems, and call-home capabilities of storage devices, are often overlooked when considering security issues in storage networks. Many storage network engineers have not explored the remote side of dial-in lines, modems, or call-home devices to ensure that the security of the remote network is not undermining the security of the storage network. Remote connectivity can leave a well-secured storage network vulnerable to an attack.

It is important to know and understand which storage devices can perform what network actions without any user interaction. The following questions should be asked when using any kind of device that uses remote connectivity:

- Are the storage devices using secure protocols?
- Do the remote connections require two-factor authentication?
- What kinds of controls are placed on the remote connections (other than username and password)?
- Are any IP address limitations or requirements necessary to attempt a valid connection?

Operating System Security

Operating system security is a best practice that is often ignored in storage networks. An operating system can act as a gateway into the storage network since it is connected to the internal LAN and to the back-end SAN. However, little concentration of security is placed on these systems—it's similar to having a firewall that connects the outside, untrusted Internet with the internal corporate network with a rule that allows any IP address to access any IP address.

In a storage network, operating systems are often the only "firewalls" that protect access to data. A compromised operating system can enable an unauthorized user to attack storage devices, such as switches and storage appliances, directly, and to attempt sniffing techniques. This makes the operating system an easy attack target. For example, consider a Fibre Channel SAN that has an HBA in an operating system for backup purposes. In addition to the HBA, the operating system also has a NIC that is plugged into the internal LAN. If proper operating system security has not been placed on the server that has both a connection to the storage network and a connection to any other network, the operating system may be responsible for partial or even complete unauthorized access to the SAN.

Many storage networks consist of unsecured and default installations of several types of operating systems, from all flavors of UNIX to all versions of Windows. Many environments do not contain a host hardening or a secure build process for operating systems that exist in storage networks. This gives an attacker an opportunity to compromise a system in such a way that he can be placed directly inside the storage network, making further attacks easier. Figure 25-11 is a graphical representation of the importance of operating system security.

Storage Appliance Services

Storage devices such as Network Appliance (NetApp) and EMC devices contain a significant amount of system services that are not all storage related. For example, most storage vendors support storage protocols such as CIFS and NFS. However, in addition to running CIFS or NFS services, it would not be unusual to see FTP (File Transfer

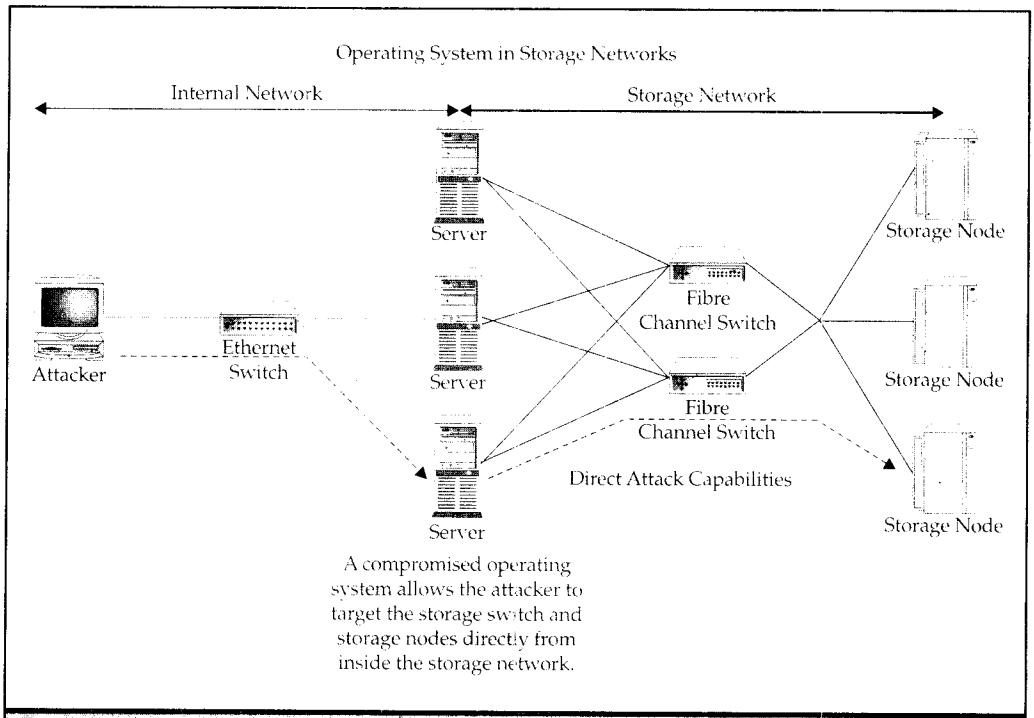

Figure 25-11. *Operating system security in storage networks*

Protocol), Telnet, Simple Network Management Protocol (SNMP), mount, portmapper, Domain Name System (DNS), HTTP, NetBIOS, RSH, syslog, and others running on these storage devices.

Similar to an operating system with default installations, storage devices can be unsecure in the default state. As a best practice, storage devices, such as filers, data movers, and NAS heads, should be deployed with a minimal amount of services and with unnecessary default modules disabled. This process not only secures the storage device, but it also allows the core service of the system, such as the storage protocol, to run as the sole module. A minimalist approach to storage devices leaves less room for error and less possibility of security exposure. An accidentally enabled SNMP daemon could leave the storage appliance vulnerable to information disclosure. Even a general SNMP vulnerability, which would affect all SNMP devices, would leave the storage node vulnerable to direct compromise from this single service enabled. In general, it is best practice to disable all storage device services except the required storage protocols and management protocols.