2 SANs: Fibre Channel Security

The emergence of SAN technology combined with data protection, privacy, and regulatory concerns has made storage security an important topic. SAN security risks are often misunderstood and/or underestimated. Furthermore, the critical issues associated with SANs, combined with the lack of communication concerning defenses, has created a security gap in storage.

The purpose of this chapter is to discuss Fibre Channel SAN security risks (iSCSI security risks will be discussed in Chapter 8). Each risk will be described and then fully discussed to allow organizations to make decisions based on their SAN data, its implementation, and the organization's risk-tolerance level.

Chapter 2 is the first of three chapters (Chapters 2, 3, and 4) where SAN security risks and the correlating attacks will be discussed. After a detailed description of the security risks, we discuss the details of the SAN attacks. Several sections in the next three chapters will be followed by a self-assessment exercise, allowing administrators to test their own exposures, vulnerabilities, and exploits.

The following topics are the primary focus of this chapter:

- SAN risks
- Risks of Fibre Channel
- Fibre Channel frame weaknesses (session hijacking)
- Fibre Channel address weaknesses (Man-in-the-Middle attacks)

SAN RISKS

In order to discuss the risks in SAN architectures, we must evaluate it on the six areas of security discussed in Chapter 1, "Introduction to Storage Security." Table 2.1 lists each of the sections, as well as their security presence in SANs.

Table 2.1 SAN Risk and Security

Security	SAN Risk
Authentication	Authentication aspects in most SAN environments do not exist. Fibre Channel Authentication Protocol (FCAP), DH-CHAP (Diffie-Hielman CHAP), and Fibre Channel Security Protocol (FC-SP) have emerged to fulfill a significant gap for authentication; however, most SANs are designed with the assumption that authentication has taken place elsewhere in the architecture. For example, organizations often assume authentication occurring at file/record layers (databases) should be enough, which ignores network authentication at lower network levels. This would be similar to requiring authentication on a web application but not requiring authentication for a telnet or SSH connection to the web server. In both scenarios, data can be compromised fully. Authentication is indirectly available through some of the applications that have access to the SAN. Management applications, which can be used to administer storage data, usually require some type of username and password. CT Authentication, DH-CHAP, FCAP, and FC-SP, as well as some other authentication modules, have been developed to authenticate node to node, node to switch, and switch to switch (discussed further in Chapter 9, "Securing Fibre Channel SANs").
Authorization	Authorization parameters are usually provided with World Wide Names (WWNs) from the Fibre Channel host bus adapters. WWNs can be port WWNs, which identify the port, or node WWNs, which identify the node on the fabric.
Encryption	Encryption aspects in most SAN environments do not exist unless some third-party at-rest encryption device is used. Natively, Fibre Channel does not use any encryption in any of its layers (layer 0 thru layer 4).
Auditing	Auditing aspects in most SANs are enabled only at the device or application level, such as a Fibre Channel switch or a management application. There is error management via the fabric; however, nothing for typical security auditing.
Integrity	There are currently no native methods for integrity checking in Fibre Channel frames.
Availability	Availability or Quality of Service (QoS) is indirectly available in layer 2 Fibre Channel frames in the Error Control fields of the frame. This aspect provides more QoS aspects than data availability. Availability is arguably the most important aspect of SAN security. If the storage data becomes unavailable, networks as well as applications melt down quickly.

RISKS OF FIBRE CHANNEL

Risks in Fibre Channel? There are no risks in Fibre Channel, right? Wrong. The Fibre Channel communications medium is absent of several entities that are required for secure transmission. Several of the weaknesses are similar to the weaknesses in IP version 4 (IPv4) and have been repeated in Fibre Channel. This section discusses the following topics:

- Description of Fibre Channel
- Clear-text communication

DESCRIPTION OF FIBRE CHANNEL

In order to understand the security issues with Fibre Channel SANs, we should discuss the architecture of Fibre Channel communications. Fibre Channel uses frames between one node to the other (similar to how IP networks use packets). Each frame contains five layers. The layers within each frame work with the layer below and the layer above to provide different functions within a Fibre Channel topology. Most SANs use either a switched Fibre Channel topology, similar to what we use in an IP-enabled switch network, or a Fibre Channel arbitrated loop (FC-AL). In either topology, each layer performs a specific function depending on the architecture that has been deployed. The five different layers of Fibre Channel frames are as follows:

- Upper Layer Protocol Mapping—FC Layer 4
- Common Services Layer—FC Layer 3
- Signaling/Framing Layer—FC Layer 2
- Transmission Layer—FC Layer 1
- Physical Layer—FC Layer 0

Similar to an IP network, Fibre Channel frames work from the physical layer, layer 0, to the upper layers. The similarities of the two communication methods primarily end at the physical layer; however, they do share similar security weaknesses and both have absent security controls. Several IP weaknesses have translated to vulnerabilities and exploits. Unfortunately, several of these attack types are also available in Fibre Channel frames. The weaknesses in Fibre Channel frames specifically target Fibre Channel layer 2, known as the framing/flow control layer (layer 2 in Fibre Channel and the Data/Networking (layer 2/layer 3) layer in an IP packet). The similarities are close in terms of security weaknesses and the lack of authentication, authorization, integrity, and encryption. Figure 2.1 shows the five different layers of a Fibre Channel frame.

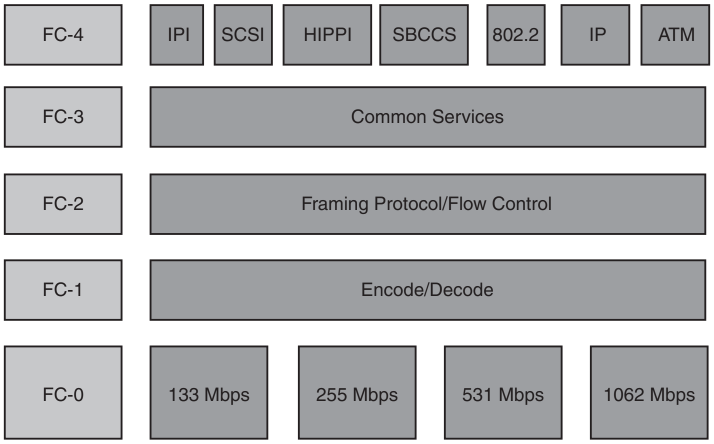

Figure 2.1 Five layers of a Fibre Channel frame.

Fibre Channel layer 2, the Framing Protocol/Flow Control layer, is the primary target when addressing frame security weaknesses. Fibre Channel layer 2 contains the header information for each frame. The header information is the location of several security weaknesses. The contents of the header include a 24-bit address (also known as the port ID) of the source node, the 24-bit address of the destination node, the sequence control number, the sequence identification number, and the exchange information. The following entities are located within the frame header:

- **Source Address (S_ID)**—A 24-bit fabric address used to identify the source address when routing frames.
- **Destination Address (D_ID)**—A 24-bit fabric address used to identify the destination address when routing frames.
- **Sequence ID (SEQ_ID)**—A static number transmitted with each frame in a sequence that identifies the frame as part of a session. Each frame in the same session has the same sequence ID.
- **Sequence Count (SEQ_CNT)**—A number that identifies individual frames within a sequence. For each frame transmitted in a sequence, SEQ_CNT is incremented by 1, allowing the frames to be arranged in the correct order.

- **Exchange ID—**Information that specifies how many frames a node can accept at one time. This information is passed from one node to another.
- **Originator Exchange ID (OX_ID)—**The exchange information of the sender.
- **Recipient Exchange ID (RX_ID)—**The exchange information of the receiver.
- **Type—**The Upper Layer Protocol byte section.
- **Routing Control (R_CTL)—**Contains information such as the routing bits, which contain data values, and the information category, which tells the receiver what type of data is contained in the frame.

Each node on a SAN fabric has a 24-bit fabric address that is used for a variety of things, including routing and name server information. (Note: Do not confuse the 24-bit fabric address with the 64-bit WWN address from the HBA.) Similar to how an IP packet is used to route packets, the 24-bit address is used to route frames from one node to the other. Figure 2.2 shows an example of the header information in Fibre Channel layer 2.

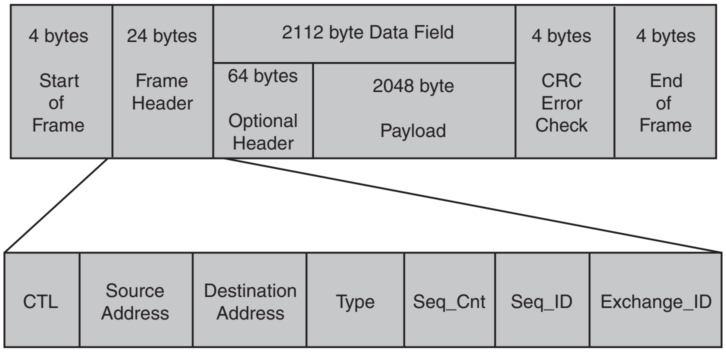

Figure 2.2 Fibre Channel layer 2.

CLEAR-TEXT COMMUNICATION

Fibre Channel communication is clear-text. The lack of security built into the different layers of Fibre Channel frames combined with the fact that it is clear-text allows for certain security threats to be very successful.

The lack of encryption at the frame level is not a significant negative issue, considering the amount of performance impact the storage network would have if all frames were encrypted. Furthermore, sniffing is a difficult task in a Fibre Channel SAN since it can only take place if a hardware device is connected to a node in the SAN or if a Cisco MDS switch is comprised and configured to send traffic remotely to the software only sniffer called Ethereal. Nevertheless, the lack of data obfuscation that contains sensitive information can allow unauthorized users to view information that is required to complete an attack. In fact, a key starting point for successful attackers is the ability to sniff clear-text communication, which can be conducted with any traffic analyzer.

Clear-text communication can be viewed as the Achilles' heel of data networks. It satisfies the enormous performance and capacity issues, but it also exposes untrusted entities to sensitive information, including SAN information. For example, clear-text protocols in IP networks, such as Rsh, Rsysnc, Rlogin, FTP, Telnet, SNMP, POP3, SMTP, ARP, and even iSCSI, allow many IP risks and attacks to either be possible or escalated. The fact that sensitive information, such as usernames/password, community strings, message challenges/hashes, and/or route information, traverse clear-text communication mediums allow untrusted users to gain sensitive information without doing anything but tapping the connection.

Many IPv4 administrators overlook clear-text communication due to the false sense of security of switched networks. In IP networks, switch technology makes it more difficult to sniff network communication; however, many attacks, such as the Man-in-the-Middle (MITM) attack, can subvert switched networking, including Fibre Channel switched networking.

Fibre Channel networks can use Fibre Channel Arbitrated Loops (FC-AL) or Fibre Channel switched networks. Sniffing Fibre Channel Arbitrated Loops does not require any MITM tricks because the fabric is a loop (ring) topology, where every connected node on the same loop can view the communication of every other node on the loop. Furthermore, using similar techniques used in IPv4 network, sniffing on a Fibre Channel switch fabric is not an impossible task, but significantly more difficult than an IPv4 network. More discussion of the MITM attacks are discussed later in this chapter, but it is important to note that sniffing on a Fibre Channel fabric is a security risk that may expose the sensitive information that traverses the network in clear-text.

The risk and weaknesses of Fibre Channel start with the clear-text transmission of sensitive information, which directly results in enumeration (the first basic step for an attacker). Enumeration is a phase where an unauthorized user would gather information about the network, architecture, device, or application they want to compromise. The result from this phase is the actual fuel that is used to perform an attack. You'll notice that the enumeration phase is not something shown in Hollywood security films, but the truth is that the enumeration phase of an attack is usually 60 to 80 percent of the process itself. The actual act of performing an attack is less than a quarter of the work. As stated earlier, sniffing the network is the first step in the enumeration phase of attacks, which is used to reveal weaknesses in the network itself.

The results of the enumeration phase determine how triumphant the actual attack will be. For example, if the enumeration phase was able to gain significant information about the network, devices, applications, operating systems, routers, WWNs, and IQNs, then the penetration phase will not only be successful, but might also be far more damaging. Conversely, if the enumeration phase does not yield favorable results for an attacker, the actual penetration phase would be short and probably unsuccessful. Figure 2.3 is a graph that shows the relationship of the enumeration and penetration phase of an attack.

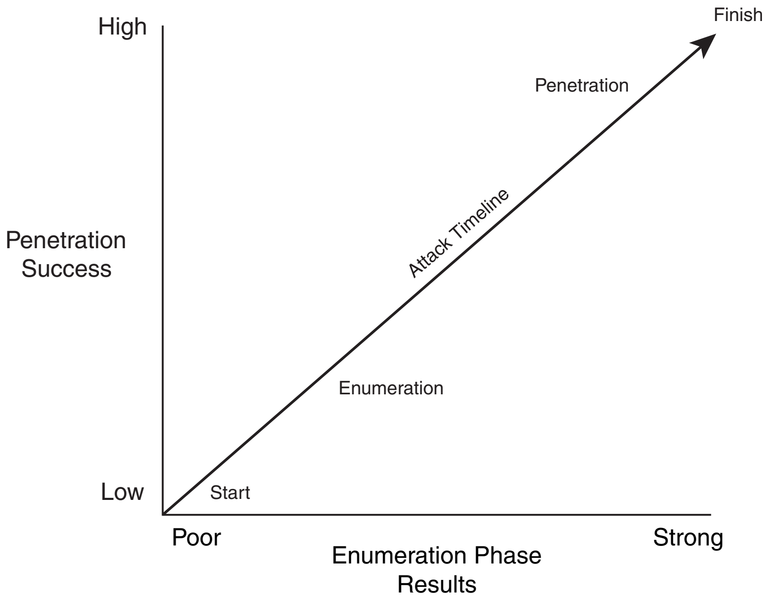

Figure 2.3 Example of a sample attack timeline.

In Figure 2.3, notice the direct relationship between the enumeration phase results and the attack success. As more success occurs in the enumeration phase, the likelihood of success in the attack process increases.

Now that we have established that enumeration is a very critical step in an attack, the problems with clear-text communication leaking an abundance of sensitive information should be understood. The next question to address is exactly what sensitive information in the Fibre Channel frame can actually be used in a possible attack? The following list describes several of the items that an unauthorized user can enumerate from a node connected to the SAN. Each of these entities gives ammunition to attackers to complete a successful attack:

- Fabric name
- Domain identification
- Switch name server information
- Session sequence control number
- Session sequence IDs
- World Wide Names used in the fabric
- Layer-2 frame information
- 24-bit addresses
- Routing information (destination and source IDs)
- Management information (such as SES and FC-SNMP)

The enumeration of a Fibre Channel SAN does not equate into data compromise, but it does significantly help the process. As an attacker tries to gain enough information to perform an attack, he or she will need to enumerate the target before any attack can be executed. Conversely, not all enumeration is negative. An organization may send clear-text information over the network that is not considered to be sensitive; such as Exchange IDs from Fibre Channel frames. The proper exercise of data classification should be conducted, as discussed in Chapter 1, "Introduction to Storage Security," to determine what type of data that traverses the network is consider public or private.

HACKING THE SAN

Hacking the SAN translates to unauthorized access to an entity or data in a storage area network. In the next three chapters, we discuss the following items.

- Session hijacking
- Man-in-the-Middle attacks
- Name server pollution
- WWN spoofing
- LUN masking attacks
- Zone hopping
- Switch attacks

Table 2.2 is summary of the weaknesses that are discussed in the next three chapters and their correlating attacks.

Table 2.2 SAN Security Weaknesses and Correlation SAN Attacks

SAN Weaknesses	**SAN Attacks**
Sequence weaknesses	Session hijacking
Fabric address weaknesses	Man-in-the-Middle attacks
FLOGI/PLOGI weaknesses	Name server pollution
HBA weaknesses	LUN masking attacks/WWN spoofing
FC switch weaknesses	Zone hopping

A key idea to introduce at this time before we begin our discussion on SAN attacks is the difference between a valid attack and a valid risk. In a given network, there are several hundred attacks that are fully possible to execute, but only a handful of them may actually pose a valid risk due to the nature of the network or the business. Hence, for each attack described in this section, a chart is used to describe how easy or difficult the execution of the attack will be, and its risk level also will be discussed. See Figure 2.4 for the example chart.

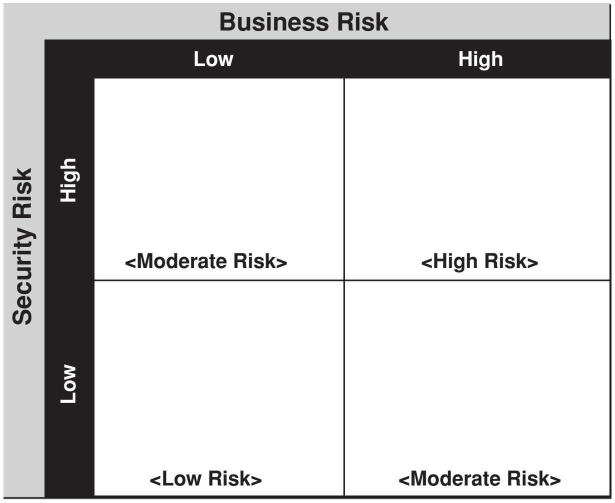

		Business Risk	
		Low	High
Security Risk	High	<Moderate Risk>	<High Risk>
	Low	<Low Risk>	<Moderate Risk>

Figure 2.4 Security and Business Risk chart.

The primary purpose of the SBR chart is to place each threat described in some type of security risk context. This chapter covers many risks and threats in Fibre Channel SANs; many of the threats are easy to perform, but many are very difficult to execute due to the need for physical access to the network or a hardware analyzer for sniffing. It would not be in the best interest of the book to simply skip the threats that are hard to actually perform, but use the SBR chart to appropriately show the risk level of each attack after it has been described.

In Figure 2.4, notice that each area of the chart represents a different security and business risk value. Items in the upper-left corner are high security risk, but low business risk. Risks in this area should be technically mitigated from a security perspective only since the business risk is low. Items in the upper-right corner are high security risk and high business risk. Risks in this area should be resolved immediately since they present a high business and security risk. Conversely, items in the lower-left corner are low security risk and low business risk. Risks in this area can often be accepted (bearable) since the impact is relatively low. Finally, items in the lower-right corner are low security risk and high business risk. Risks in this often need a process solution rather than a technical solution. The type of summary in the Security and Business Risk (SBR) chart will help readers understand what valid attacks are and the risks associated with them.

Now that we understand the architecture of Fibre Channel frames and the problems associated with clear-text communication, we will now discuss the security weaknesses with Fibre Channel frames. The following list describes each weakness that we will discuss:

- Sequence weaknesses
- Address weaknesses
- Fabric, port, and node login weaknesses
- FLOGI, PLOGI, and address spoofing

FIBRE CHANNEL FRAME WEAKNESSES

The following sections discuss the weaknesses with Fibre Channel at the frame level.

SEQUENCE WEAKNESSES

A *sequence* is a set of frames transmitted unidirectionally from one entity to another in order to maintain a session between two nodes. A frame uses a Sequence ID (Seq_ID) and a Sequence Count (Seq_CNT) in each frame to identify, control, and maintain the session. Each frame includes the Sequence ID (Seq_ID) that identifies the unique session that the frame belongs to. For example, if a node were communicating with several different entities, each session would have a unique Seq_ID to identify which frame belongs to which session. In addition to the Sequence ID, the Sequence Count is used also. The Sequence Count is used to ensure frames are placed in the right order by the entities. Each Seq_CNT is incremented by 1 for each subsequent data frame.

The Sequence ID and Sequence Count have similar responsibilities in a Fibre Channel frame as the Initial Sequence Number (ISN) in an IP packet. The ISN in an IP packet is also responsible for maintaining a session between two nodes on an IP network.

In order for a session to be maintained between two nodes, all session information must be maintained. The security weakness with IP networks is the predictable (guessable) ISN. An ISN is the core component to allow packets to become a part of an established session. If the ISN is predictable, it would potentially allow unauthorized packets to join or hijack the session. A third-party entity could inject packets with the next relevant ISN and take control of the established session. For example, for simplicity's sake, let's say two of the first three packets in a session have the ISN of 123 for packet number one and 456 for packet number two. An attacker could probably predict that the third packet should have an ISN of 789 to be part of the existing session. If the attacker sends their packet to the target first and uses the ISN of 789, the session will then be handed over to the attacker and not the legitimate node. This means that if an entity was able to guess the next ISN in the sequence and the entity was able to inject packets in an established session, the entity would own the session. See Figure 2.5 for more details.

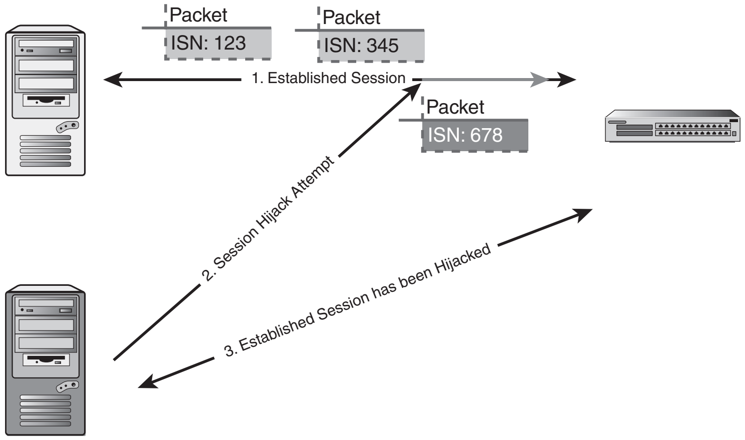
Packet
ISN: 123
Packet
ISN: 345
1. Established Session
Packet
ISN: 678
2. Session Hijack Attempt
3. Established Session has been Hijacked

Figure 2.5 ISN session hijacking attack.

As shown in Figure 2.5, a weak or predictable ISN can leave an established and trusted session vulnerable to a session hijacking attack. A predictable value to allow/deny entities into a trusted session should not be used. ISN values must be unpredictable *and* unique, not unpredictable *or* unique. This premise has also plagued HTTP (web) applications. Session identifiers in cookies used in HTTP applications, known as the Session ID, are being used to maintain sessions in the stateless HTTP protocols. While applications distribute cookies with unique Session IDs, the Session IDs are not necessarily unpredictable. This allowed unauthorized users to guess or predict the Session ID and log into applications using another user's existing sessions. For example, if a legitimate user and a malicious user logged on to their favorite webmail service, they would receive a cookie from the webmail site that contains a Session ID. Both users would present their cookie to the site each time they want to check email. If the site granted a Session ID of 100 to the legitimate user and 101 to the malicious user, then an attacker could guess/predict that the Session IDs is a three-character digit being incremented by one for any new session. The attacker could access the site under the legitimate user's profile by changing their Session ID in their cookie to 100 and send it to the site. Once the site receives the cookie from the malicious user that contains the Session ID of 100, it would recognize the attacker as the legitimate user because the site recognizes the user based on the Session ID in a user's cookie. This would allow the attacker to log to the legitimate user's session and give them access to profile pages, credit card pages, and account information.

As stated in the Preface, attacks don't change, but they do get modified. Similar to the session hijacking weaknesses in IP ISN packets, which was introduced over 15 years ago, and Session IDs in HTTP cookies, the Sequence ID and Sequence Counts in Fibre Channel frames are unique values, but they are predictable. The Sequence Count value is incremented by one, a very predictable pattern. Furthermore, the Sequence ID is a unique number, but is also a static number that does not change within the session. Therefore, of the two entities that maintain the session in a Fibre Channel frame, one is a static value and the other is a value that increments by one, both of which can allow an unauthorized entity to predict or guess values and inject their own frames to hijack a session. For example, let's say two nodes are communicating on a Fibre Channel network—they have a Sequence ID of 12, and the first frame has a Sequence Count of 1117171342 and the second frame has a Sequence Count of 1117171343. An unauthorized node could inject frames to the target node with the Sequence ID of 12, since the frame is in clear-text and the value does not change, and predict the next Sequence Count of 1117171344 in order to hijack the session. Notice that although 1117171342 is a long and unique number, the first four digits is the date (November 17th or 1117) and the last six digits is time (5:13 and 42 seconds or 17:13:42). By doing some simple pattern matching and predictions, it is easy to figure out that the next few frames will have the Sequence Count of 1117171344, 1117171345, and 1117171346.

What does this mean? Well, this does not mean that you can go download Hunt, Ettercap, or WebProxy and begin to perform session hijack attacks on Fibre Channel SANs (Hunt, Ettercap, and WebProxy are IP/Application tools to hijack sessions in IP networks or web applications). Although the weaknesses are there in a Fibre Channel frame, the threat and exposure is relatively low. A Fibre Channel analyzer is required to modify frames and inject them into established session. With speeds up to 2gb/sec, this is not an easy task. Nonetheless, the weakness of session management with the Seq_ID and Seq_CNT fields in a Fibre Channel frame do exist, which tells us that security may have not been a significant factor when developing session management for Fibre Channel or Fibre Channel-enabled products.

SESSION HIJACKING

Session hijacking is the act of an untrusted third party intercepting and controlling (hijacking) a valid session between two trusted entities. Telnet is a good example of a trusted session between two entities that can be hijacked by an anonymous third party on the segment if weak ISNs are being used for the TCP packets. Figure 2.6 shows a high-level example of session hijacking.

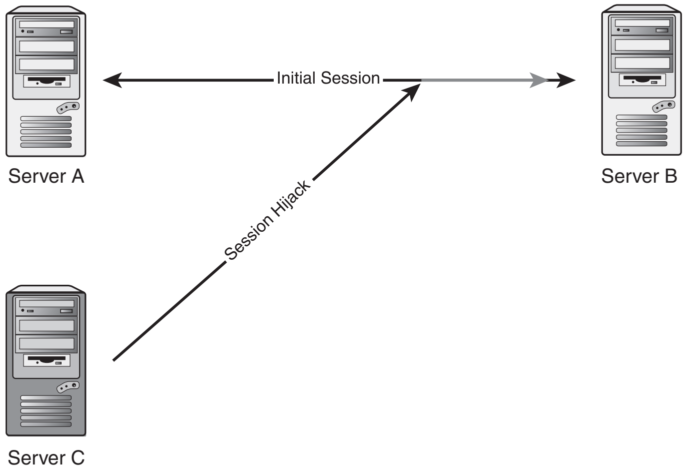

Figure 2.6 Sample session hijacking between two trusted entities.

Session hijacking was first introduced many years ago in the IP networking world for weak and predictable Initial Sequence Numbers in TCP headers for IP packets.

The attack became quite easy to execute with IP tools such as Hunt (http://lin.fsid.cvut.cz/~kra/#HUNT) and Ettercap (http://ettercap.sourceforge.net/). Session hijacking resurfaced in the application world when weak and predictable session IDs became apparent in application cookies to maintain state in web (HTTP) communication. Similar to our discussion in Chapter 1 on how several attacks don't change, but get modified, the idea of session hijacking can be applied to Fibre Channel frames also.

ASSESSMENT EXERCISE

In order to better understand session hijacking, we will demonstrate an example using an IP network and also using the tool called Hunt, which works solely on the Unix/Linux platform. While Hunt is a great tool, you might have to try it several times to get it to work correctly. Based on Figure 2.6, server C will be the malicious user hijacking a session from server A to server B. The following steps outline the method to perform session hijacking:

1. On server C, download Hunt from http://lin.fsid.cvut.cz/~kra/#HUNT.
2. Unzip Hunt:
   a. cd /usr/local/bin
   b. gunzip --c hunt.tar.gz | tar xvf --
3. Compile Hunt:
   c. cd /usr/local/bin/hunt-1.5
   d. make
   e. make install
4. Execute Hunt:
   f. ./hunt
5. On server A, telnet to server B, using the telnet command:
   g. telnet serverB
6. On server C, choose option A (see Figure 2.7).
7. On a hub (or a switch using arpspoof), you should see the Telnet session from server A (10.10.80.122) to server C (128.101.81.1) with Hunt.
8. At the `choose conn>` prompt, enter **0** to hijack the session (see Figure 2.8).

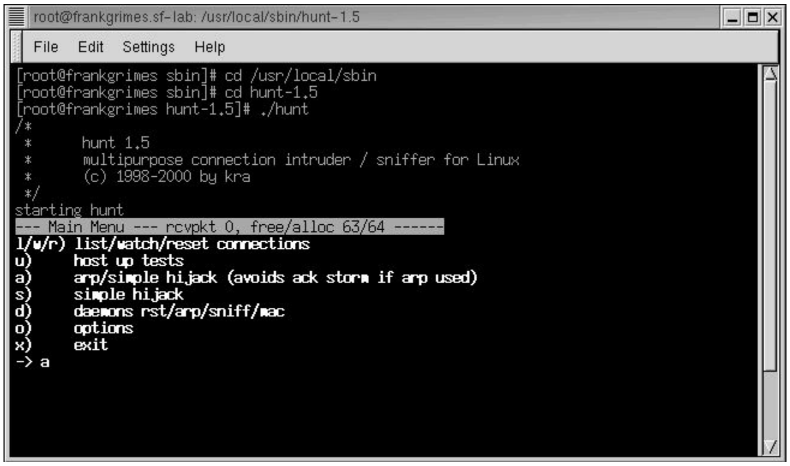

```
[root@frankgrimes sbin]# cd /usr/local/sbin
[root@frankgrimes sbin]# cd hunt-1.5
[root@frankgrimes hunt-1.5]# ./hunt
/*
 *      hunt 1.5
 *      multipurpose connection intruder / sniffer for Linux
 *      (c) 1998-2000 by kra
 */
starting hunt
--- Main Menu --- rcvpkt 0, free/alloc 63/64 ------
l/w/r) list/watch/reset connections
u)     host up tests
a)     arp/simple hijack (avoids ack storm if arp used)
s)     simple hijack
d)     daemons rst/arp/sniff/mac
o)     options
x)     exit
-> a
```

Figure 2.7 Hunt's options.

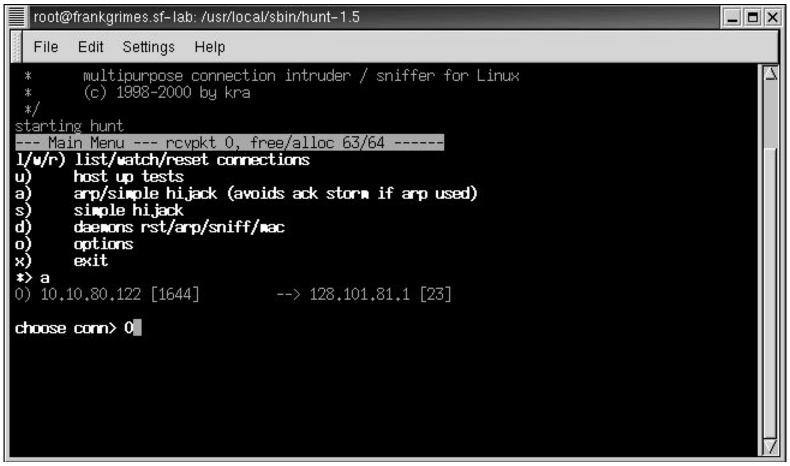

```
 *      multipurpose connection intruder / sniffer for Linux
 *      (c) 1998-2000 by kra
 */
starting hunt
--- Main Menu --- rcvpkt 0, free/alloc 63/64 ------
l/w/r) list/watch/reset connections
u)     host up tests
a)     arp/simple hijack (avoids ack storm if arp used)
s)     simple hijack
d)     daemons rst/arp/sniff/mac
o)     options
x)     exit
*> a
0) 10.10.80.122 [1644]        --> 128.101.81.1 [23]

choose conn> 0
```

Figure 2.8 Telnet session with Hunt.

9. Accept all the defaults (labeled in brackets) until you see the target session (see Figure 2.9).

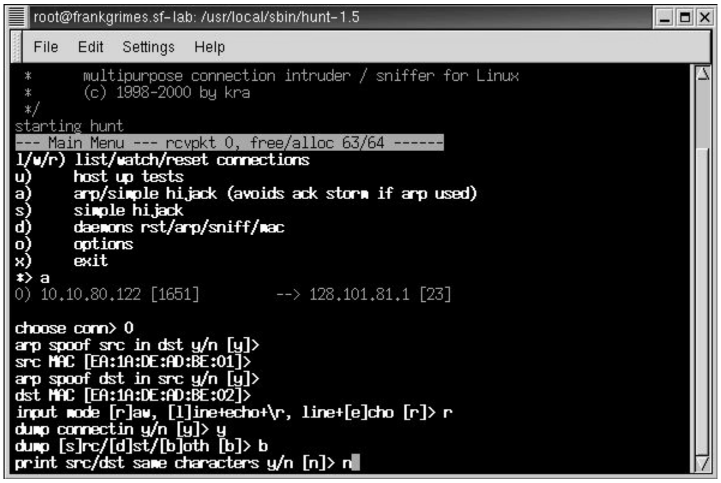

```
 *      multipurpose connection intruder / sniffer for Linux
 *      (c) 1998-2000 by kra
 */
starting hunt
--- Main Menu --- rcvpkt 0, free/alloc 63/64 ------
l/w/r) list/watch/reset connections
u)     host up tests
a)     arp/simple hijack (avoids ack storm if arp used)
s)     simple hijack
d)     daemons rst/arp/sniff/mac
o)     options
x)     exit
*> a
0) 10.10.80.122 [1651]          --> 128.101.81.1 [23]

choose conn> 0
arp spoof src in dst y/n [y]>
src MAC [EA:1A:DE:AD:BE:01]>
arp spoof dst in src y/n [y]>
dst MAC [EA:1A:DE:AD:BE:02]>
input mode [r]aw, [l]ine+echo+\r, line+[e]cho [r]> r
dump connectin y/n [y]> y
dump [s]rc/[d]st/[b]oth [b]> b
print src/dst same characters y/n [n]> n
```

Figure 2.9 Hunt's defaults.

10. Ensure the user on server A continues to type and is not sitting idle with the telnet prompt.
11. Back on server C, if the user on server A is using the Telnet session, you will be able to either watch the session or hijack it at any time by pressing the Enter key a couple of times (preferably after the user has su to root!!!). Once you have hijacked the session, Hunt will inform you that "you took over the connection."
12. Done. Server C has now hijacked the Telnet session from server A to server B due to poor session information in the Initial Sequence Number of the TCP header (see Figure 2.10).

Based on our attack here, the IP session hijacking will be classified as a high security risk, since authentication can be subverted from a malicious attacker. Additionally, the attacker will also get a high business risk, since a successful IP session hijack attack will allow an unauthorized user to gain access to machines containing sensitive data. See Figure 2.11 for an SBR chart.

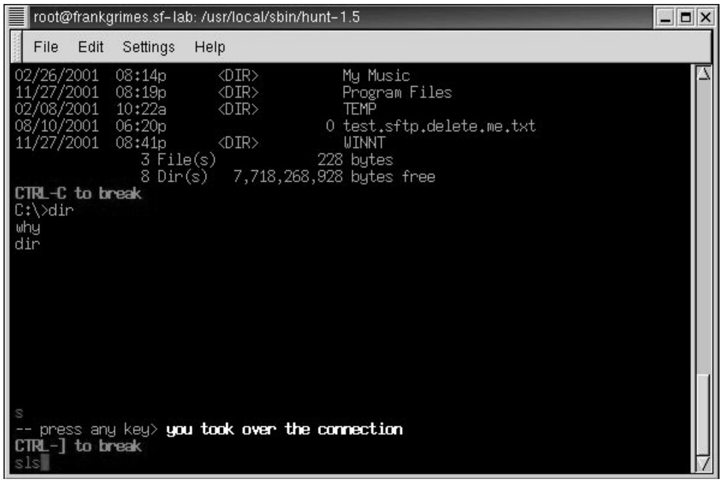

```
02/26/2001  08:14p      <DIR>          My Music
11/27/2001  08:19p      <DIR>          Program Files
02/08/2001  10:22a      <DIR>          TEMP
08/10/2001  06:20p                   0 test.sftp.delete.me.txt
11/27/2001  08:41p      <DIR>          WINNT
               3 File(s)            228 bytes
               8 Dir(s)   7,718,268,928 bytes free
CTRL-C to break
C:\>dir
why
dir








s
-- press any key> you took over the connection
CTRL-] to break
sls
```

Figure 2.10 Successful hijacked session from server A.

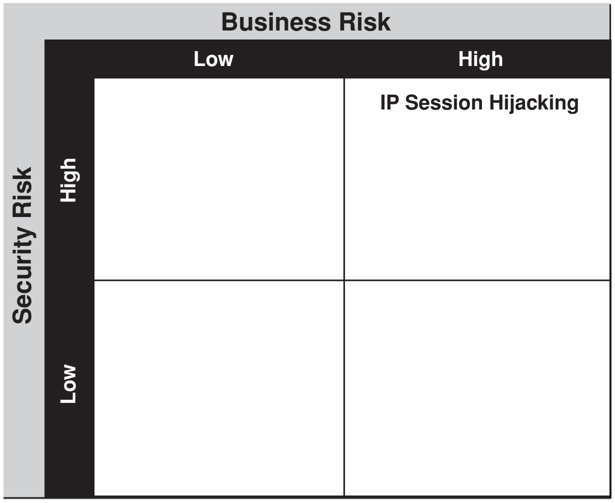

		Business Risk	
		Low	High
Security Risk	High		IP Session Hijacking
	Low		

Figure 2.11 SBR chart for IP session hijacking.

SESSION HIJACKING—FIBRE CHANNEL

In a Fibre Channel architecture, in order for two Fibre Channel nodes to communicate with each other, an established session must be made. The session information is managed by the Sequence Count number (Seq_CNT) and the Sequence Identification number (SEQ_ID). The definitions for Sequence Count numbers and Sequence Identification numbers are as follows:

- **Sequence Count (SEQ_CNT)—**A number identifies individual frames within a sequence. For each frame transmitted in a sequence, SEQ_CNT is incremented by 1, allowing the frames to be arranged in the correct order.
- **Sequence ID (SEQ_ID)—**A static number transmitted with each frame in a sequence that identifies the frame as part of a session. Each frame in the same session has the same Sequence ID.

The session information between two FC nodes is the entity that is in charge of maintaining a session. For example, if 100 frames were delivered from several nodes to another single node, there needs to be a method to understand which frame came from which node and to organize the frames in their correct order. In order to complete this, the Seq_ID and Seq_CNT are used. The Seq_ID and Seq_CNT will tie each frame to a particular session and place it in its correct order.

The issue with session management starts with the lack of Fibre Channel authentication when sending or receiving frames. Similar to the IP and Hunt example earlier, in order to break in and hijack a session, a malicious user could send frames to an authorized node with the correct Seq_ID and Seq_CNT (using the source address (S_ID) as the attacker and not the original session holder), thus, transferring the session's control to the malicious user. Furthermore, since the Seq_ID never changes (which makes it very easy to guess), and the Seq_CNT number increments by the value of one (which makes it very easy to predict), the hijacking process of the session is quite trivial.

Although the attack is very trivial, as demonstrated with IP/Hunt as well as web applications and session identifiers, currently there are no automated tools to perform this type of attack. A Fibre channel analyzer would have to be used to actually perform session hijacking on FC frames, rendering the attack as a high security threat, but a low risk item (see Figure 2.12).

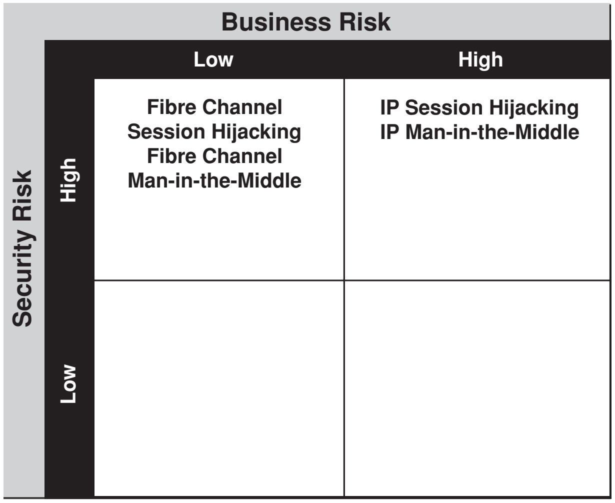

		Business Risk	
		Low	High
Security Risk	High	Fibre Channel Session Hijacking Fibre Channel Man-in-the-Middle	IP Session Hijacking IP Man-in-the-Middle
	Low		

Figure 2.12 SBR chart for Fibre Channel session hijacking.

ASSESSMENT EXERCISE

In order to fully understand the Fibre Channel hijacking attack, the following steps describe the attack process according to Figure 2.13:

1. Node Kusum makes an established connection with node Neeraja.
2. Node Kusum and node Neeraja exchange frames for communication.
3. Using a Fibre Channel traffic analyzer, the malicious user, Lakshman, identifies the static value for the Seq_ID and the Seq_CNT number.
4. Lakshman then injects frames to Neeraja with the Seq_ID, taken from the frames between Kusum and Neeraja, and then increments the Seq_CNT number by one, which will identify the next frame in the session.
5. Neeraja receives the frame(s) from Lakshman, and because the frames have the correct Seq_ID and the correct value for the Seq_CNT, the frames are regarded as the next set of frames in the session.
6. Because the S_ID of the Lakshman's frames are from a different address, the session is then handed to that node wherever the session last left off.

7. Lakshman has hijacked the session from Kusum and now has an established connection to Neeraja without any authentication or authorization.
8. Despite the fact that Lakshman has hijacked the session, Neeraja still thinks the established connection is with Kusum.

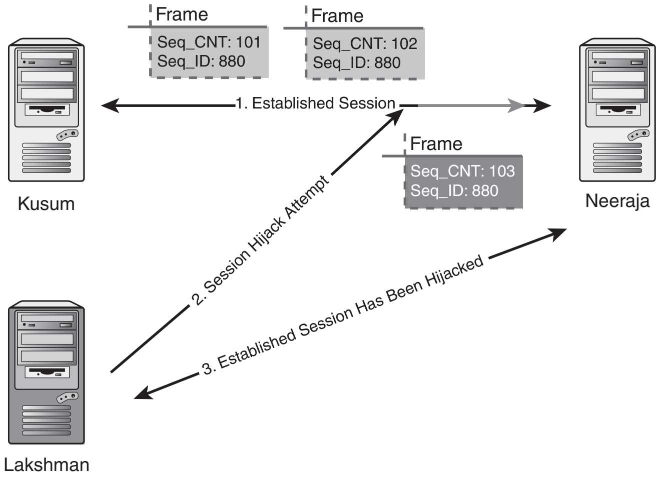

Figure 2.13 Topology for session hijacking.

The technical details of the attack are obviously more complicated, but generally follow the steps outlined beforehand. The first step would be to enumerate the frame information from the two trusted entities. Using any type of Fibre Channel fabric analyzer or IP sniffer with an IP to Fibre Channel connector, there are a variety of methods to do this attack depending on the type of architecture. For example, if a fabric loop topology had been deployed (FC_AL), the analyzer can see all the traffic in the loop of every node connected to the fabric. In a switched architecture, the analyzer would need to be connected to the core FC switch and also within the same routing segment of the target. Once the targets have been identified and enumerated by the traffic analyzer, the following fields in the header part of the frame are important for the attack: S_ID, D_ID, OX_ID, RX_ID, Seq_ID, and Seq_CNT.

Within Fibre Channel layer 2, you would modify the header information with your traffic analyzer of the frame that you will generate in order to complete the attack. When crafting the frame, the S_ID would change from the original source fabric address to the fabric address of the attacker. The D_ID would remain the same, which is the fabric address of the target (while both entities are technically targets, the entity that is on the receiving end of the session is the real target). The OX_ID and RX_ID values would have to be consistent with the original source (Kusum in our previous example) since the malicious node would need to be able to send and receive the same amount of frames as the original source and consequently send the correct amount of frames to the target specified in the original RX_ID field. The Seq_ID field will also need to remain identical to the original in order to ensure the target node considers the frame(s) as part as the legitimate session. Unlike the Seq_ID, the Seq_CNT field will not remain identical but rather will need to be incremented by one in order for the target to consider that frame as the next legitimate frame in the session. This is probably the trickiest part of the attack; even though the act of incrementing the Seq_CNT by one is a trivial procedure, it is not as easy to determine what Seq_CNT is the last one. For example, using your traffic analyzer, you are able to view the Seq_CNT number of all sessions, but by the time you send your frame(s), the legitimate source may already have sent a frame with that Seq_CNT, thus leaving your frame useless. Although this attack takes some trial and error, a good way to ensure that the attack works correctly is to estimate what the Seq_CNT number will be by the time you have the opportunity to send your frame(s) and to set up multiple instances of malicious frames, each using a Seq_CNT that could possibly be the next legitimate one by the time it reaches the target. After you have successfully done this with your traffic analyzer, you will notice that your node will start receiving the frames from the target with the same Seq_ID and Seq_CNT from the original session, despite the fact that you have not officially logged into the node (NLOGIN).

ATTACK SUMMARY: SESSION HIJACKING

Attack description—Hijacking a data or management session by guessing the sequence control number (SEQ_CNT) and the sequence identification number (SEQ_ID) of a Fibre Channel frame.

Risk level—High. An unauthorized entity could gain access to an authorized management session or simply modify the sequence numbers randomly and attempt to perform a denial-of-service attack.

Difficulty—High. This is a sophisticated attack that requires deep knowledge of Fibre Channel frames and the use of a hardware and software traffic analyzer.

Best practice—None to date; however, the use of strong or unpredictable SEQ_CNT or SEQ_ID would mitigate this issue in the future. Ask your storage vendor about frame authentication or integrity options.

FIBRE CHANNEL ADDRESS WEAKNESSES

Now that we have established that attacks don't change, but they do get modified, let's discuss another attack that stems network and application history. Manipulation of the 24-bit fabric address can cause significant damage and denial of service in a SAN.

Each node in a SAN has a 24-bit fabric address that is used for routing, among other things. Along with routing frames correctly to/from their source and destinations, the 24-bit address is also used for name server information. The name server is a logical database in each Fibre Channel switch that correlates a node's 24-bit fabric address to their 64-bit WWN. Additionally, the name server is also responsible for other items, such as mapping the 24-bit fabric address and 64-bit WWN to the authorized LUNs in the SAN. Furthermore, address information is also used for soft and hard zoning procedures (discussed in the Chapter 4, "SANs: Zone and Switch Security"). The 24-bit fabric address of a node determines route functions with soft and hard zoning procedures, specifically if a frame is allowed to pass from one zone to the other. While there are several other uses of the 24-bit address, the use of the address in name servers and zoning procedures are by far the most important in terms of security.

The major issues with the 24-bit address is that it is used for identification purposes for both name server information and soft/hard zone routing, almost like an authorization process, but it is an entity that can be easily spoofed. Using any traffic analyzer, the 24-bit source address of a Fibre Channel frame could be spoofed as it performs both PLOGI (Port Login) and FLOGI (Fabric Login) procedures.

In Fibre Channel, there are three different types of login—Port Login, Fabric Login, and Node Login. Two can be corrupted with a spoofed 24-bit fabric address. Before we discuss how spoofing disrupts these processes, let's discuss the login types first.

FABRIC LOGIN (FLOGI), PORT LOGIN (PLOGI), AND NODE LOGIN (NLOGI)

The Fabric Login (FLOGI) process allows a node to log in to the fabric and receive an assigned address from a switch. The FLOGI occurs with any node (N_Port or NL_Port) that is attached to the fabric. The N_Port or NL_Port will carry out the FLOGI with a nearby switch. The node (N_Port or NL_Port) will send a FLOGI frame that contains its node name, its N_Port name, and any service parameters. When the node sends its information to the address of 0xFFFFFE, it uses the 24-bit source address of 0x000000 because it hasn't received a legitimate 24-bit address from the fabric yet. The FLOGI will be sent to the well-known fabric address of 0xFFFFFE, which is similar to the broadcast address in an IP network (though not the same). The FC switches and fabric will receive the FLOGI at the address of 0xFFFFFE. After a switch receives the FLOGI, it will give the

N_Port or NL_Port a 24-bit address that pertains to the fabric itself. This 24-bit address with be in the form of Domain-Area-Port address from, where the Domain is the unique domain name (ID) of the fabric, Area is the unique area name (ID) of the switch within the domain, and Port is the unique name (ID) of each port within the switch in the fabric. Table 2.3 shows how the 24-bit address is made.

Table 2.3 24-Bit Addresses

24-Bit Address Type	Description
8-bit domain name	Unique domain ID in a fabric. Valid domain IDs are between 1 and 239.
8-bit area name	Unique area ID on a switch within a fabric. Valid area IDs are between 0 and 255.
8-bit port name	Unique port ID within a switch in a fabric. Valid port IDs are between 0 and 255.

A 24-bit address (port ID) uses the following formula to determine a node's address:

```
Domain_ID x 65536 + Area_ID x 256 + Port_ID = 24 bit Address
```

An example address for and node on the first domain (domain ID of 1), on the first switch (area ID of 0), and the first port (port ID of 1), would be the following:

```
1 x 65536 + 0 x 256 + 1 = 65537 (Hex: 0x10001)
```

After the node has completed the FLOGI and has a valid 24-bit fabric address, it will perform a Port Login (PLOGI) to the well-known address of 0xFFFFFC to register its new 24-bit address with the switch's name server, as well as submit information on its 64-bit port WWN, 64-bit node WWN, port type, and class of service. The switch then registers that 24-bit fabric address, along with all the other information submitted, to the name server and replicates that information to other name servers on the switch fabric. Figures 2.14 and 2.15 show the FLOGI and PLOGI processes.

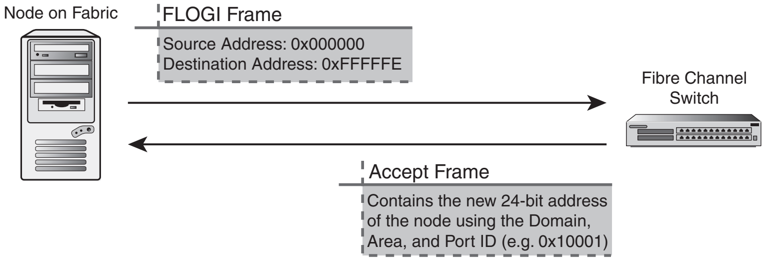
Node on Fabric
FLOGI Frame
Source Address: 0x000000
Destination Address: 0xFFFFFE
Fibre Channel Switch
Accept Frame
Contains the new 24-bit address of the node using the Domain, Area, and Port ID (e.g. 0x10001)

Figure 2.14 FLOGI process.

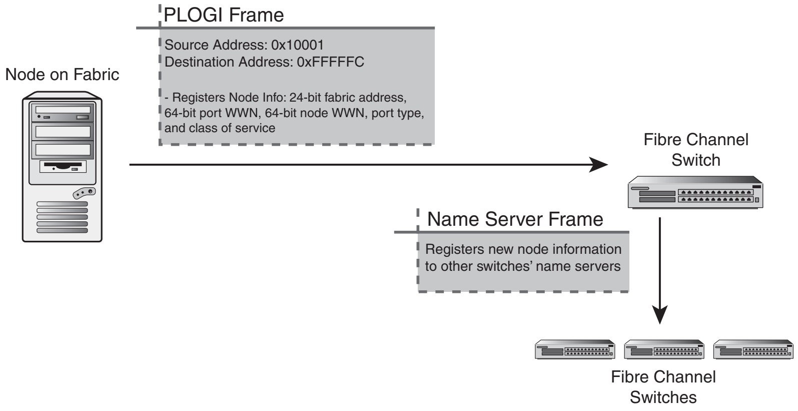
PLOGI Frame
Source Address: 0x10001
Destination Address: 0xFFFFFC
- Registers Node Info: 24-bit fabric address, 64-bit port WWN, 64-bit node WWN, port type, and class of service
Node on Fabric
Fibre Channel Switch
Name Server Frame
Registers new node information to other switches' name servers
Fibre Channel Switches

Figure 2.15 PLOGI process.

A Node Login is somewhat similar to a Fabric Login, but instead of logging in to the fabric, the node would log in to another node directly (node to node communication). The node will not receive any information from the fabric, but will receive information from the other node as it relates to Exchange IDs (OX_ID and RX_ID) and session information (Seq_ID and Seq_CNT). After this information has been exchanged, the two nodes will begin to communicate with each other directly.

FLOGI, PLOGI, AND ADDRESS SPOOFING

Now that we have established facts concerning FLOGI, PLOGI, and address spoofing, let's understand how the weaknesses interrelate them. After performing the FLOGI process, an FC node needs to perform a PLOGI to the well-known address of 0xFFFFFC. The PLOGI then registers the 24-bit address of the node to the Name Server (also referred to as a Simple Name Server) of the switch. If an entity were to spoof their 24-bit fabric address and send it to the address of 0xFFFFFC, the switches would see a node performing a PLOGI. Once the switch receives the information from the PLOGI frame, it will register the spoofed 24-bit address of the node to the name server—thus, polluting the name server with incorrect information. You might wonder what the big deal is since the node has corrupted its own information; however, consider the fact that the 24-bit address is used for hard and soft zoning. For example, let's say the 24-bit address of 65537 (Hex: 0x10001) was allowed to route to nodes in zone A and no other addresses can access that zone. A malicious attacker has the address of 65541 (Hex: 0x10005) and cannot access that zone. The malicious attacker can spoof (change) their 24-bit address to match 65537 (0x10001) and then route frames to the restricted zone A, despite being unauthorized to do so. Spoofing the 24-bit address during PLOGI negates any route-based zoning rules that may have been applied. The simple process of spoofing now creates the ability to route (hop) across hard and soft zoning rules. Figure 2.16 shows the FLOGI/PLOGI spoofing process.

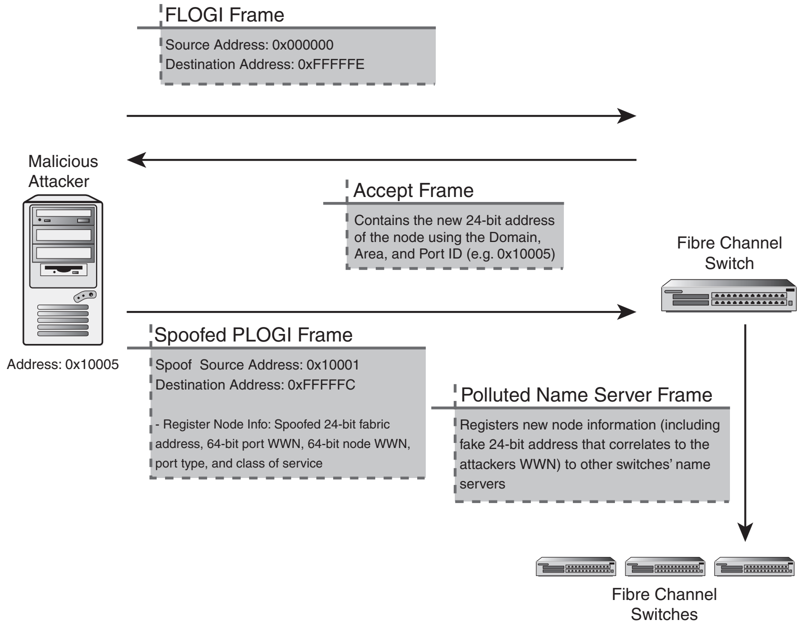
FLOGI Frame
Source Address: 0x000000
Destination Address: 0xFFFFFE
Malicious Attacker
Accept Frame
Contains the new 24-bit address of the node using the Domain, Area, and Port ID (e.g. 0x10005)
Fibre Channel Switch
Spoofed PLOGI Frame
Address: 0x10005
Spoof Source Address: 0x10001
Destination Address: 0xFFFFFC
- Register Node Info: Spoofed 24-bit fabric address, 64-bit port WWN, 64-bit node WWN, port type, and class of service
Polluted Name Server Frame
Registers new node information (including fake 24-bit address that correlates to the attackers WWN) to other switches' name servers
Fibre Channel Switches

Figure 2.16 FLOGI/PLOGI spoofing process.

We will take this idea a bit further in the next section, "Man-in-the-Middle Attacks," when I discuss the issues of spoofing the 24-bit fabric address and spoofing a node WWN. The fact is that this attack is very severe by breaking the integrity of any hard or soft zoning rules. However, a traffic analyzer is required to perform this attack, thus creating barriers to perform the attack itself.

MAN-IN-THE-MIDDLE ATTACKS

A *Man-in-the-Middle (MITM) attack* is the act of an untrusted third party intercepting communication between two trusted entities. For example, when you call a friend on the telephone, you dial his or her phone number and wait for an answer. When your friend picks up the phone, you then begin communicating with him or her. In a MITM attack, a malicious user would intercept the connection between you and your friend. Instead of talking to your friend directly, you would actually be communicating through a malicious third party. The malicious third party would then connect you to your friend. Both you and your friend would begin communicating, not knowing that an unauthorized entity has connected you both and is listening to every word of the conversation. It is like a three-way call, but two of the three callers don't know that there is a third person listening.

In the digital world, the untrusted third party plays the role of a router, but unlike an authorized router, the untrusted third party should not have permission to view, modify, or intercept any of the communication between the two trusted entities. Figure 2.17 shows a high-level example of a Man-in-the-Middle attack.

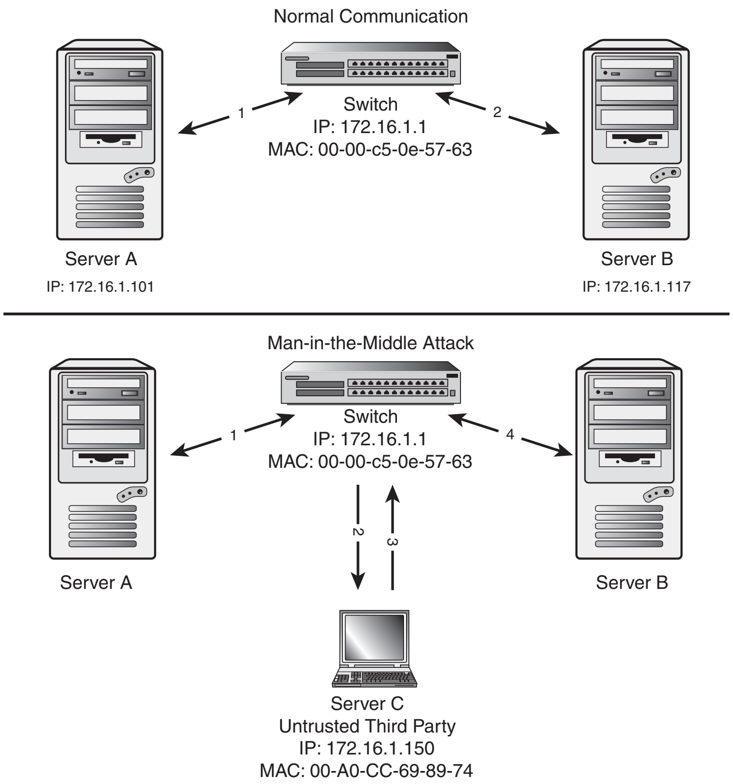

Figure 2.17 Sample Man-in-the-Middle attack.

Man-in-the-Middle attacks were first introduced many years ago in the IP networking world. Unauthenticated OSI layer 2 Address Resolution Protocol (ARP) packets could update ARP tables (tables that match a node's IP address to their machine (MAC) address) in switches and/or operating systems. The purpose of the MITM is to sniff on a switch. A switch will only transmit information to the correct port, not allowing any other ports to see any communication that is not theirs. On the other hand, a hub is a dumber device that allows all ports to see all communication, making it quite easy to sniff a neighbor's traffic. Many switches are layer 2 devices, meaning that they can transmit packets from one port on a switch to another without the need for an IP address, but with a node's machine address (MAC). For example, on a Windows operating system, type **ipconfig /all** from the command line and then press Enter. Notice the physical address is the machine address of your node, which is actually the MAC address of your Network Interface Card (NIC). The MAC address comes from the manufacturer of the NIC to identify it. If your node wanted to speak to another node via a layer 2 switch, it would use the MAC address and not your IP address. Layer 2 routing is common for performance reasons, allowing switches to transfer packets quickly across the network. Once the packets get to a layer 3 device, such as a router, then a node's IP address can be used.

Since ARP is a layer 2 protocol, it uses a node's MAC address to identify nodes and transfer packets. ARP is similar to Name Servers in the Fibre Channel world, where Name Servers match the WWN of HBAs to the 24-bit fabric address (as well as a few other items, such as zones and LUN access).

MAN-IN-THE-MIDDLE ATTACKS—IP

Before we can begin to understand the idea about a Fibre Channel Man-in-the-Middle attack, let's first understand the concept using the IP protocol. An entity using IP, such as a switch or an operating system, will send out ARP requests when it is trying to communicate with other entities. For example, if server A wanted to communicate with server B, which has the IP address of 172.16.1.1 and the MAC address of 00-0A-CC-69-89-74, server A would send out an ARP request asking, "Who is 172.16.1.1?" Then the switch or the operating system would respond, replying with its MAC address, which is 00-0A-CC-69-89-74. The issue with ARP, which we will also address with Fibre Channel name servers, is that any malicious entity could send out an ARP reply instead of the actual server. For example, if you stepped outside your home and yelled out, "What is the address of the post-office," a malicious neighbor could say, "I am the post-office; please send your mail to me." If you believed this malicious neighbor without asking for proof, then your mail would be compromised. This is how ARP works, without any authentication. A malicious user could send out ARP replies with the incorrect information.

Since there is no authentication with ARP, similar to how there is no authentication with PLOGI in Fibre Channel fabrics, an entity receiving an ARP reply from an attacker would update their routing table with the incorrect information. Furthermore, even if a node did not send out an ARP request, which would request the MAC address of a specific IP address, it doesn't mean it won't receive an ARP reply and update its own routing table. For example, a malicious user could send out ARP replies to the entire network segment, telling each entity that the MAC address of the router, which is 172.16.1.1, is actually the MAC address of the malicious entity. When one node tries to communicate to any other node by going through the default router, it will actually be going to the malicious entity first, since it is using the MAC address of the malicious entity for layer 2 routing.

ASSESSMENT EXERCISE

Let's attempt an IP Man-in-the-Middle attack using the following steps to better understand the issue:

1. Ensure you are on a test network and you have full permission from your network administrator because a Man-in-the-Middle attack will cause network disruption.
2. In our example using Figure 2.17, servers A and B will try to communicate with one another and server C will intercept that traffic.
3. Download Cain and Abel to server C from http://www.oxid.it/cain.html.
4. Type **ipconfig** on server A. Notice that is the default route listed as "Gateway Address." In our example, the default gateway is 172.16.1.1. Therefore, when server A tries to communicate to server B, it would go through its router first, which is 172.16.1.1 (see Figure 2.18).
5. Still on server A, ping its gateway address by typing **ping <default gateway>**, such as ping 172.16.1.1. See Figure 2.19. Now that server A has pinged the router, it will have its MAC address in its ARP table. Type **arp –a** to see the dynamic ARP table in the system (see Figure 2.19).
6. Switch to server C (the malicious user). Use Cain and Abel to enable server C to send out ARP replies to the network segment, telling everyone that the address of 172.16.1.1 is actually associated with the MAC address of 00-00-86-59-C8-94, which is the MAC address of server C, not the MAC address of the router. The MAC address of the router is 00-00-C5-0E-57-63, which we know from the `arp –a` command that told us the MAC address of the default router (172.16.1.1). Perform the steps in the next exercise on server C.

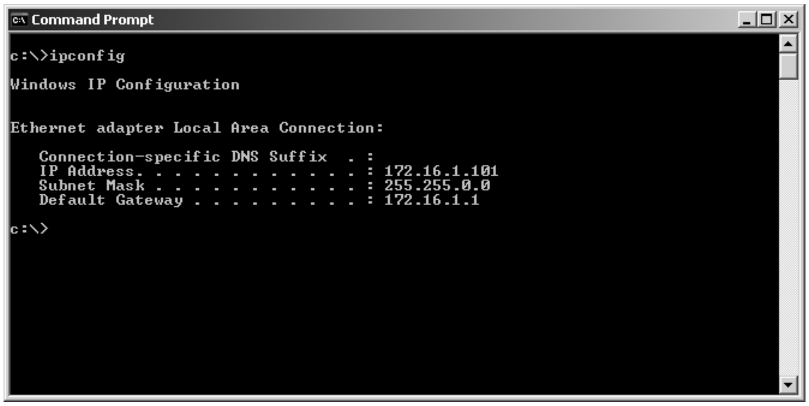

```
c:\>ipconfig

Windows IP Configuration


Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix  . :
   IP Address. . . . . . . . . . . . : 172.16.1.101
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Default Gateway . . . . . . . . . : 172.16.1.1

c:\>
```

Figure 2.18 Server A’s IP address and gateway.

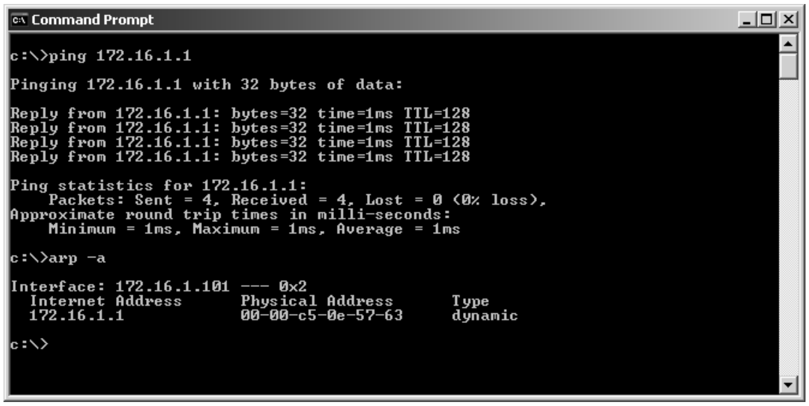

```
c:\>ping 172.16.1.1

Pinging 172.16.1.1 with 32 bytes of data:

Reply from 172.16.1.1: bytes=32 time=1ms TTL=128
Reply from 172.16.1.1: bytes=32 time=1ms TTL=128
Reply from 172.16.1.1: bytes=32 time=1ms TTL=128
Reply from 172.16.1.1: bytes=32 time=1ms TTL=128

Ping statistics for 172.16.1.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 1ms, Average = 1ms

c:\>arp -a

Interface: 172.16.1.101 --- 0x2
  Internet Address      Physical Address      Type
  172.16.1.1            00-00-c5-0e-57-63     dynamic

c:\>
```

Figure 2.19 Server A’s ARP table.

ASSESSMENT EXERCISE

Complete the following steps to perform a MITM attack according to Figure 2.17 with Cain and Abel:

1. Install the Cain and Abel program using its defaults.
2. Install the WinPCap packet driver, if you don’t already have one installed.

3. Reboot.
4. Launch Cain and Abel (Start -> Programs -> Cain).
5. Select the icon in the upper-left corner that looks like a green Network Interface Card.
6. Ensure that your NIC card has been identified and enabled correctly by Cain.
7. Select the Sniffer tab.
8. Select the + symbol in the toolbar.
9. The MAC Address Scanner window appears. This enumerates all the MAC addresses on the local subnet. Hit OK. See Figure 2.20 for the results.

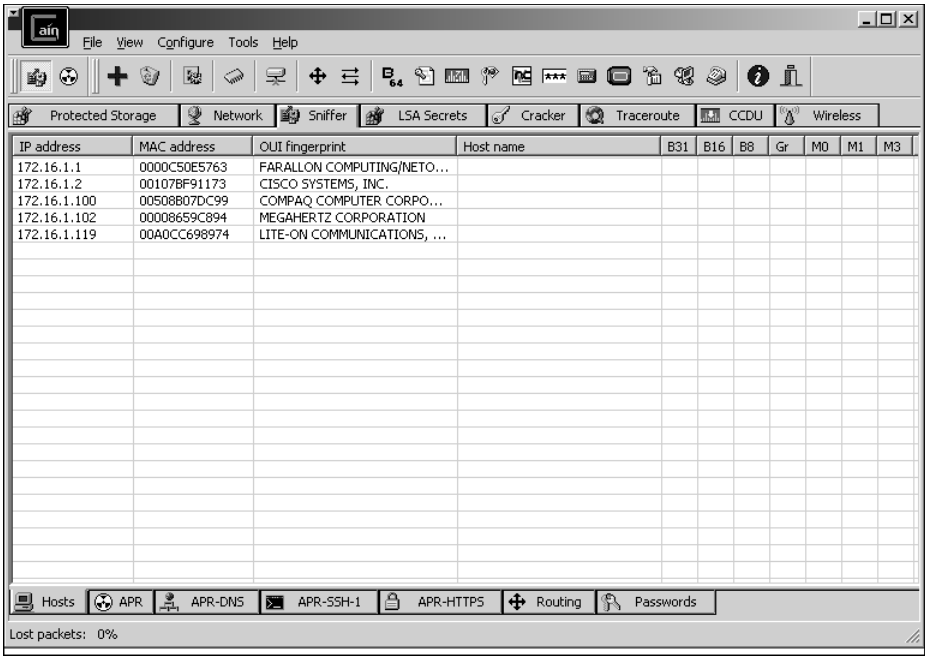

IP address	MAC address	OUI fingerprint	Host name	B31	B16	B8	Gr	M0	M1	M3
172.16.1.1	0000C50E5763	FARALLON COMPUTING/NETO...								
172.16.1.2	00107BF91173	CISCO SYSTEMS, INC.								
172.16.1.100	00508B07DC99	COMPAQ COMPUTER CORPO...								
172.16.1.102	00008659C894	MEGAHERTZ CORPORATION								
172.16.1.119	00A0CC698974	LITE-ON COMMUNICATIONS, ...								

Lost packets: 0%

Figure 2.20 MAC Address Scanner results.

10. Select the APR tab on the bottom of the tool to switch to the ARP Pollution Routing tab.
11. Select the + symbol on the toolbar to show all the IP addresses and their MACs (see Figure 2.21).

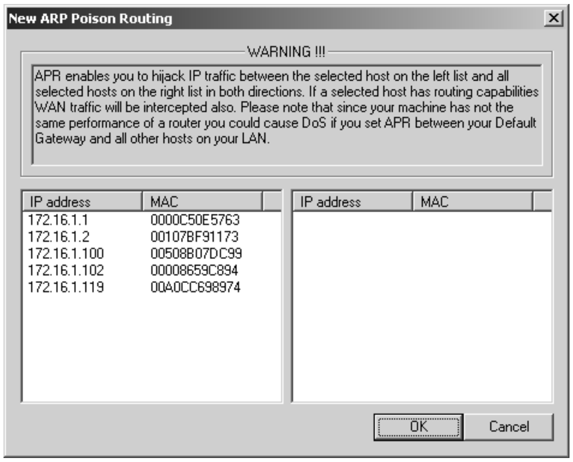

APR enables you to hijack IP traffic between the selected host on the left list and all selected hosts on the right list in both directions. If a selected host has routing capabilities WAN traffic will be intercepted also. Please note that since your machine has not the same performance of a router you could cause DoS if you set APR between your Default Gateway and all other hosts on your LAN.

IP address	MAC
172.16.1.1	0000C50E5763
172.16.1.2	00107BF91173
172.16.1.100	00508B07DC99
172.16.1.102	00008659C894
172.16.1.119	00A0CC698974

IP address	MAC

Figure 2.21 IP addresses and their MACs.

12. On the left hand side of Figure 2.22, choose the target for your MITM attack. Most likely this will be the default gateway in your subnet, so all packets will go through you first before the real gateway of the subnet.
13. Once you select your target, which is 172.16.1.1 in our example, you then select the hosts on the right side that you want to intercept traffic. This value can be all the hosts in the subnet or one particular host. We will choose one host, which will be 172.16.1.119. Select OK (see Figure 2.22).

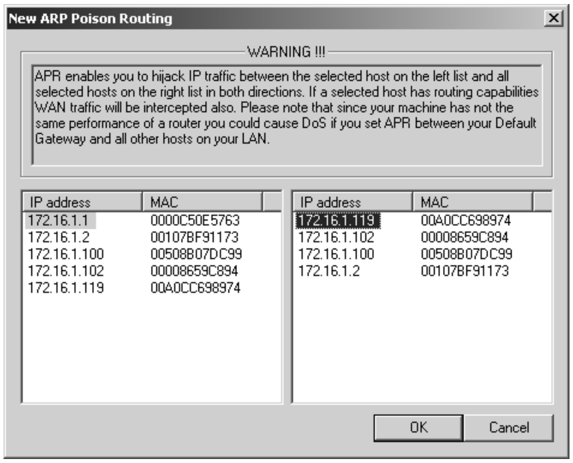

APR enables you to hijack IP traffic between the selected host on the left list and all selected hosts on the right list in both directions. If a selected host has routing capabilities WAN traffic will be intercepted also. Please note that since your machine has not the same performance of a router you could cause DoS if you set APR between your Default Gateway and all other hosts on your LAN.

IP address	MAC
172.16.1.1	0000C50E5763
172.16.1.2	00107BF91173
172.16.1.100	00508B07DC99
172.16.1.102	00008659C894
172.16.1.119	00A0CC698974

IP address	MAC
172.16.1.119	00A0CC698974
172.16.1.102	00008659C894
172.16.1.100	00508B07DC99
172.16.1.2	00107BF91173

Figure 2.22 Man-in-the-Middle targets.

14. Now select the yellow and black icon (second one from the left) to officially start the MITM attack. This will allow server C to start sending out ARP responses on the network subnet, telling 172.16.1.119 that the MAC address of 172.16.1.1 has been updated to 00-00-86-59-C8-94 (see Figure 2.23).

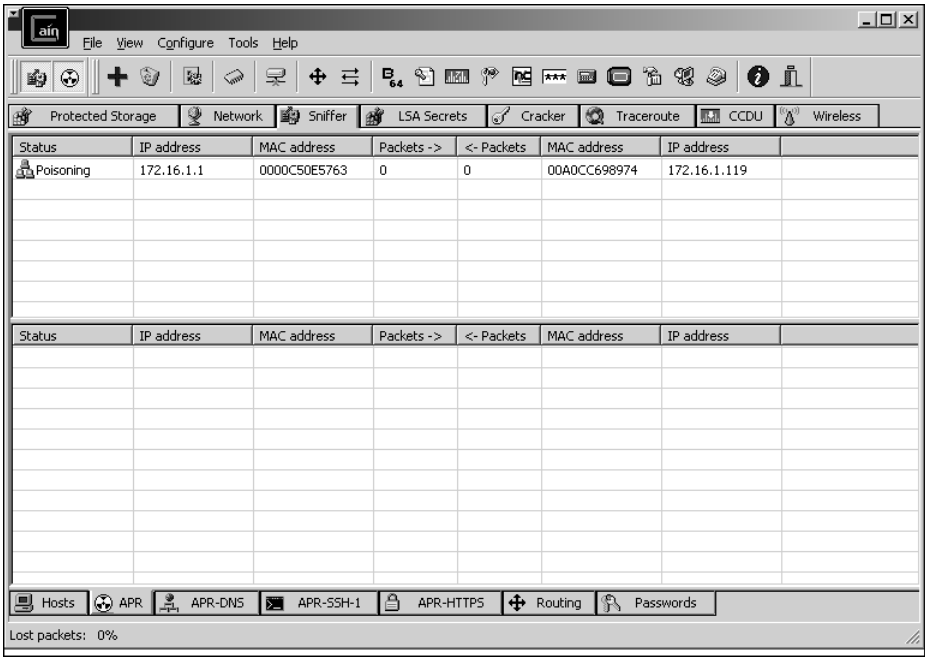

Status	IP address	MAC address	Packets ->	<- Packets	MAC address	IP address
Poisoning	172.16.1.1	0000C50E5763	0	0	00A0CC698974	172.16.1.119

Status	IP address	MAC address	Packets ->	<- Packets	MAC address	IP address

Lost packets: 0%

Figure 2.23 Man-in-the-Middle attack in process with ARP poisoning.

15. At this point, all traffic from server A to server B is going to server C first and then on its appropriate route. Server C can open up a network sniffer to view all the traffic. Additionally, Cain has a Passwords tab at the bottom that will capture the password of major protocols such as FTP, HTTP, IMAP, POP3, Telnet, and VNC. It will also capture password hashes such as Kerberos and IPsec using IKE (see Figure 2.24).

Layer 2 in the OSI model (Ethernet) is a key target for attackers. The attack becomes quite easy to execute with IP tools such as Windoze Interceptor, Dsniff, and Cain and Abel. Man-in-the-Middle attacks have also resurfaced in the application world using the same preceding techniques, but with cookies and certificates instead of ARP packets.

Similar to our discussion in Chapter 1 on how several attacks don't change but instead get modified, the idea of Man-in-the-Middle attacks can be applied to Fibre Channel frames also.

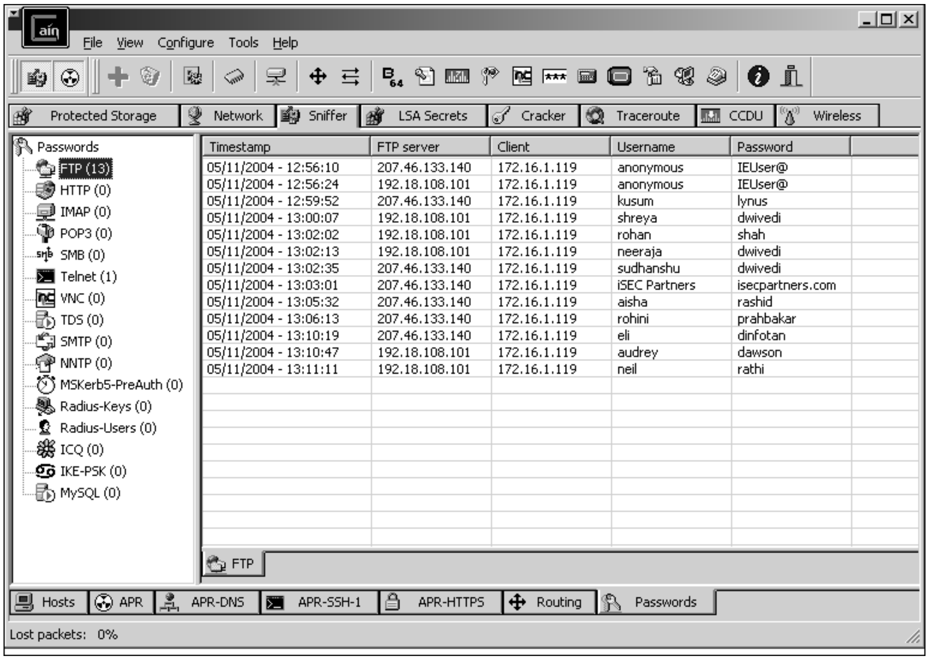

Timestamp	FTP server	Client	Username	Password
05/11/2004 - 12:56:10	207.46.133.140	172.16.1.119	anonymous	IEUser@
05/11/2004 - 12:56:24	192.18.108.101	172.16.1.119	anonymous	IEUser@
05/11/2004 - 12:59:52	207.46.133.140	172.16.1.119	kusum	lynus
05/11/2004 - 13:00:07	192.18.108.101	172.16.1.119	shreya	dwivedi
05/11/2004 - 13:02:02	192.18.108.101	172.16.1.119	rohan	shah
05/11/2004 - 13:02:13	192.18.108.101	172.16.1.119	neeraja	dwivedi
05/11/2004 - 13:02:35	207.46.133.140	172.16.1.119	sudhanshu	dwivedi
05/11/2004 - 13:03:01	207.46.133.140	172.16.1.119	iSEC Partners	isecpartners.com
05/11/2004 - 13:05:32	207.46.133.140	172.16.1.119	aisha	rashid
05/11/2004 - 13:06:13	207.46.133.140	172.16.1.119	rohini	prahbakar
05/11/2004 - 13:10:19	207.46.133.140	172.16.1.119	eli	dinfotan
05/11/2004 - 13:10:47	192.18.108.101	172.16.1.119	audrey	dawson
05/11/2004 - 13:11:11	192.18.108.101	172.16.1.119	neil	rathi

Figure 2.24 Capture password hashes due to the Man-in-the-Middle attack.

The MITM attack is possible due to the lack of authentication in ARP packets as well as the insecurities of IPv4. As demonstrated with Cain and Abel, the attack can be quite trivial, rendering the attack as a high security threat, but a high-risk item (see the SBR chart in Figure 2.25).

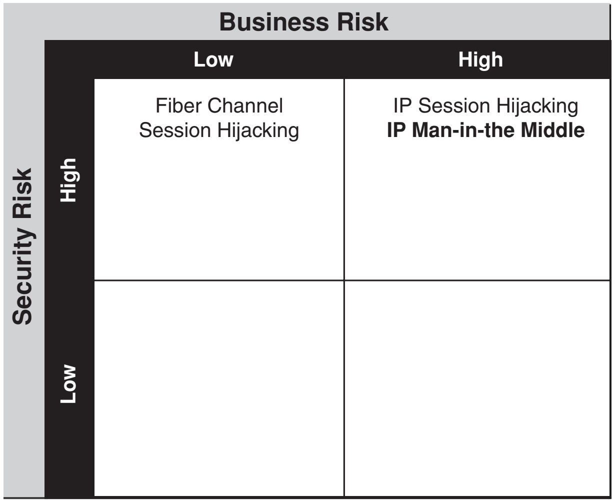

Figure 2.25 SBR chart—IP Man-in-the-Middle.

MAN-IN-THE-MIDDLE ATTACKS—FIBRE CHANNEL

In Fibre Channel fabrics, Man-in-the-Middle attacks are more difficult than IP and bear a smaller amount of risk; however, the weaknesses in the fabric are still very apparent.

NAME SERVER POLLUTION

In order to conduct a MITM attack on a Fibre Channel network, name server pollution is required. Described earlier in this chapter, there are significant weaknesses in the FLOGI and PLOGI processes that can be used to pollute the name server.

When performing a FLOGI, a Fibre Channel node will use the source address of 0x000000 because it does not have a valid S_ID yet. The node will send its frame to the destination address (D_ID) of 0xFFFFFE, which is similar to a broadcast address for Fibre Channel fabrics. After the switches receive the frame at the address of 0xFFFFFE, it will return an Accept frame, known as an ACC, to the node with its new 24-bit address, giving the node a valid fabric address. After the node has received the ACC frame and its new 24-bit address, it will then perform a PLOGI. The PLOGI will send its new 24-bit address to the address of 0xFFFFFC, registering its new address to the switch's name servers.

The security weakness is that a malicious node can craft a spoofed PLOGI frame and send it to the address of 0xFFFFFC. The malicious node could complete the FLOGI process, but instead of responding with its real 24-bit address, it could use a spoofed 24-bit address of a target during the PLOGI. Since the malicious node knows the address to send PLOGI responses to (0xFFFFFC), the act of inserting the 24-bit address is not a challenge. The switch name server would receive the spoofed PLOGI frame at the address 0xFFFFFC and will update its name server with the incorrect information. For a persistent attack, the malicious node would continue to send PLOGI frames at the address of 0xFFFFFC, continuously updating the name server with incorrect information and leaving the target with the actual 24-bit address completely out of the process. A detailed description of the contents of each frame is depicted in Figure 2.26.

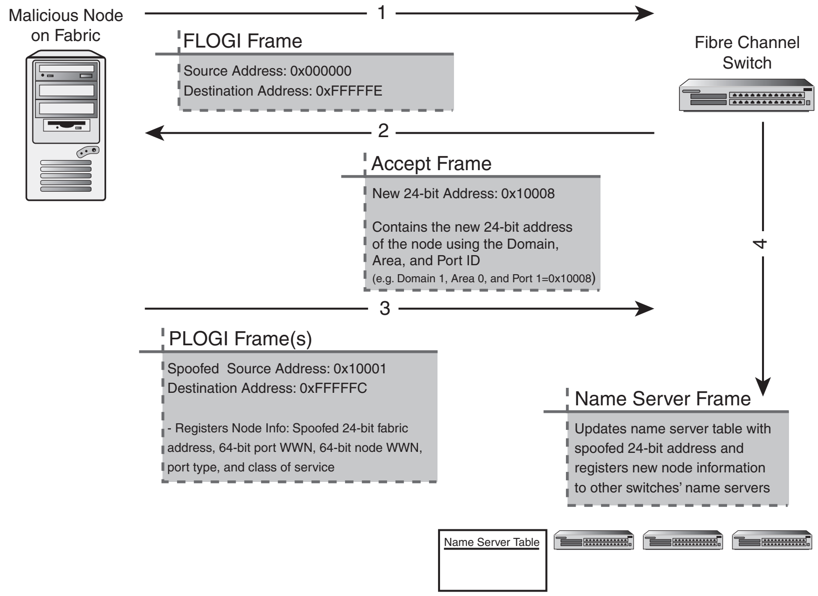
Malicious Node on Fabric
1
FLOGI Frame
Source Address: 0x000000
Destination Address: 0xFFFFFE
Fibre Channel Switch
2
Accept Frame
New 24-bit Address: 0x10008
Contains the new 24-bit address of the node using the Domain, Area, and Port ID
(e.g. Domain 1, Area 0, and Port 1=0x10008)
4
3
PLOGI Frame(s)
Spoofed Source Address: 0x10001
Destination Address: 0xFFFFFC
- Registers Node Info: Spoofed 24-bit fabric address, 64-bit port WWN, 64-bit node WWN, port type, and class of service
Name Server Frame
Updates name server table with spoofed 24-bit address and registers new node information to other switches' name servers
Name Server Table

Figure 2.26 Name server pollution process.

MITM ATTACK

In order for two Fibre Channel nodes to communicate with each other, they must know the others' 24-bit address and port ID on the fabric. The fabric address is given and updated during the FLOGI and PLOGI processes, which also has no authentication process (similar to ARP). The port ID is the physical port number that the node is connected to on the switch. If one node wanted to communicate to another node, it would then send frames to the other node's 24-bit address, which would be the Destination address (D_ID) in a frame header. The switch would receive the frame, match the 24-bit address to the correct port ID, which is completed via the switches' name server table, in order to find the correct physical port of the destination node, and then pass the frame on to the correct port. See Figure 2.27 for normal communication in a fabric.

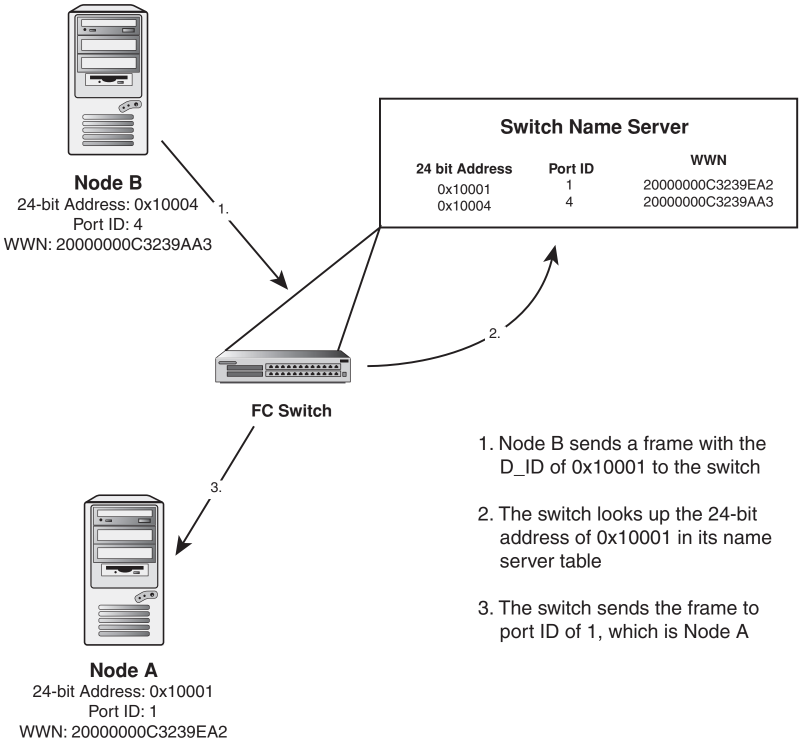

24 bit Address	Port ID	WWN
0x10001	1	20000000C3239EA2
0x10004	4	20000000C3239AA3

1. Node B sends a frame with the D_ID of 0x10001 to the switch
2. The switch looks up the 24-bit address of 0x10001 in its name server table
3. The switch sends the frame to port ID of 1, which is Node A

Figure 2.27 Normal fabric communication.

In order to perform a Fibre Channel MITM attack, a malicious node would spoof its 24-bit address to match the address of its target node (Node A). Because name server information can be automatically updated during the PLOGI process (remember that the FLOGI and PLOGI processes update name server information without authentication), the malicious user would then perform a PLOGI, sending their port ID, WWN, and spoofed 24-bit address to the fabric address of 0xFFFFFC for all the switches in the fabric to accept. The switches, with the incorrect information for the 24-bit address, would update their name servers with the port ID, WWN, and the spoofed 24-bit address. When another node wants to communicate to the real node, the switch's routing table will map the 24-bit address, which was spoofed, to a different port ID—hence, routing the frame to a different node. See Figure 2.28 for details.

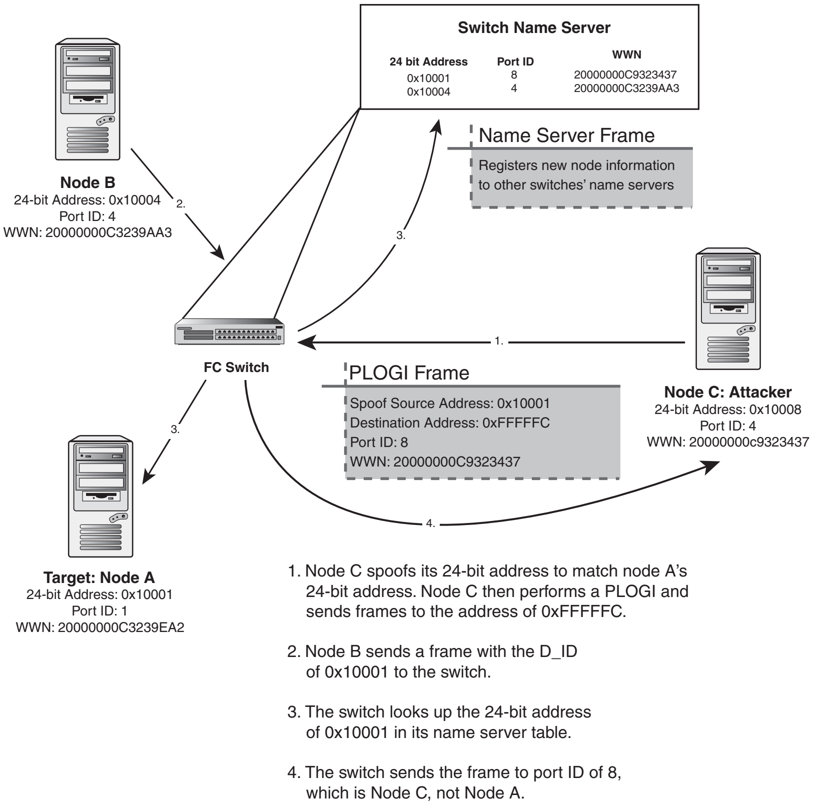

24 bit Address	Port ID	WWN
0x10001	8	20000000C9323437
0x10004	4	20000000C3239AA3

1. Node C spoofs its 24-bit address to match node A's 24-bit address. Node C then performs a PLOGI and sends frames to the address of 0xFFFFFC.

2. Node B sends a frame with the D_ID of 0x10001 to the switch.

3. The switch looks up the 24-bit address of 0x10001 in its name server table.

4. The switch sends the frame to port ID of 8, which is Node C, not Node A.

Figure 2.28 Man-in-the-Middle attack on a fabric.

The primary security weakness is the lack of authentication when sending FLOGI or PLOGI frames that consequently update name server information on switches. In Figure 2.28, node A has a 24-bit fabric address of 0x10001 and node B has a 24-bit fabric address of 0x10002. Fabric routing tables and rules would allow the two entities to communicate with each other quite easily using port ID 1 and port ID 2. When the malicious node, node C, performed a Man-in-the-Middle attack to intercept the traffic between node A and node B, the following steps were performed:

1. Node C did not perform a FLOGI, because it does not care to have a real 24-bit fabric address, but will be using the 24-bit address of its target, which is node A.
2. Using a traffic analyzer, node C crafts a frame mimicking a PLOGI frame, as if it were registering its own 24-bit address to the fabric and adjoining switches, but actually updating its spoofed 24-bit address to the authorized name server.
3. Node C performs a PLOGI using the 24-bit fabric address of 0x10001, allowing name servers to think that the 24-bit address of 0x10001 now correlates to node C, port ID 8, and WWN of 20000000c9323437.
4. Once the switches update their name servers, correlating the 24-bit address of 0x10001 with node C, any traffic destined to the 24-bit address of 0x10001, which should be node A but now is node C, will be redirected to the malicious node for interception, enumeration, and compromise.
5. When the address of 0x10004 (node B) tries to communicate to the 24-bit address of 0x10001 (node A), the traffic will actually go to node C, since the name server table in the switch thinks that port ID 8 has the 24-bit address of 0x10001.
6. In order for the Man-in-the-Middle attack to be fully complete, once node C receives the traffic from node B, it must then actually route the frames to the real destination (node A) in order for both parties to continue communication without any suspicion and for node C to continue to receive traffic from node B. If node C fails to transmit the traffic to node A, node B will realize the communication it is trying to perform is not working and stop sending frames, thus leaving node C without any frames to compromise. (Note: The last routing portion of the attack is extremely difficult due to the speeds of 2gb/sec.)

The Fibre Channel MITM attack is possible due to the lack of authentication in PLOGI frames, as well as the security weaknesses during the name server update process. As demonstrated with the preceding examples, the attack is quite possible in SAN fabrics; however, it is significantly difficult due to the speeds that an attacker would have to emulate in order to switch frames in the SAN at 2gb/sec. The throughput/performance part of the attacks makes its risk value lower, rendering the attack as a high security threat, but a low-risk item (see Figure 2.29).

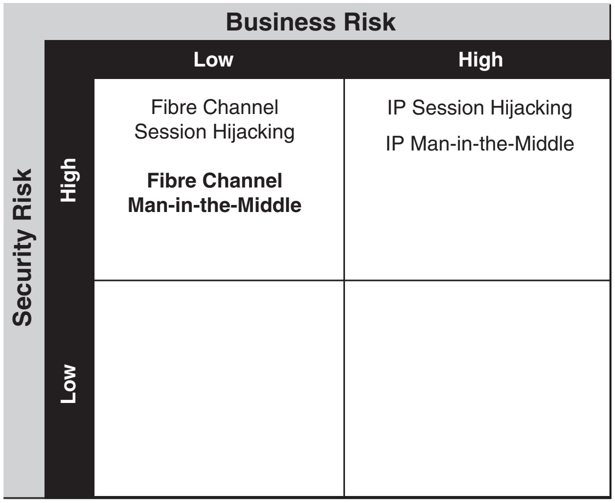

Security Risk \ Business Risk	Low	High
High	Fibre Channel Session Hijacking **Fibre Channel Man-in-the-Middle**	IP Session Hijacking IP Man-in-the-Middle
Low		

Figure 2.29 SBR chart—Fibre Channel Man-in-the-Middle.

ATTACK SUMMARY: MAN-IN-THE-MIDDLE

Attack description—Sending a fake PLOGI frame to the switch in order to register a target's 24-bit address to the attacker's WWN and port ID; hence, pollute the name server to route traffic incorrectly to the malicious node.

Risk level—Low. An unauthorized entity could gain access to unauthorized frames.

Difficulty—High. This is a sophisticated attack that requires deep knowledge of Fibre Channel frames and the use of a hardware and software traffic analyzer.

Best practice—None to date; however, the use of authenticated FLOGI and PLOGI frames would mitigate this issue in the future. Ask your storage vendor about frame authentication or integrity options.

ATTACK SUMMARY: NAME SERVER POLLUTION

Attack description—Corrupting the name server information on Fibre Channel switches where an attacker registers its 24-bit address to a target's WWN. If any legitimate node attempts to communicate to the target, the traffic is redirected to the attacker's machine by the incorrect name server information (similar to a Man-in-the-Middle attack in the IP architecture).

Risk level—High. An unauthorized entity could gain access to sensitive data with trivial attacks.

Difficulty—High. This is a sophisticated attack that requires deep knowledge of Fibre Channel frames and the use of a hardware and software traffic analyzer.

Best practice—Ensure malicious PLOGI frames, which are used to update switch name servers, cannot corrupt name server tables. Ask your storage vendor about frame authentication or integrity options.

SUMMARY

In this chapter, we discussed the risks of Fibre Channel communication in SANs. This chapter is the first of three chapters that will describe the risks of SANs and describe how to actually expose each risk identified.

Three different aspects of security were addressed that are important for any entity, including the overall risks of the entity, the method of communication that is used within the entity, and the objects that are used in the entity. Fibre Channel security risks were addressed overall, including risks of Fibre Channel as a medium of networking and risks of devices that are used in Fibre Channel storage networks.

The chapter was able to identify some of the key overall issues of Fibre Channel as they pertain to the six areas of security that can be applied to any entity, including authentication, authorization, encryption, auditing, integrity, and availability. The chapter also identified the security strengths and weaknesses of each category in order to determine the level of risks that can be exposed. Unfortunately, most SANs are missing some of the major security entities that are required for proper security, including authentication, encryption, and integrity. Furthermore, many security entities do not exist, such as authorization, are not ideal, and do not hold up to many SAN attacks, such as spoofing.

This chapter was also able to discuss some of the risks associated with Fibre Channel as a medium of communication and networking. The chapter demonstrated how clear-text communication can be a big issue in terms of SAN protection. Furthermore, the weaknesses in Fibre Channel frames can hurt the overall security of a SAN architecture. The chapter also discussed Fibre Channel layer 2 as a target for various attacks on Fibre Channel frames, including spoofing, man-in-the-middle, session hijacking, PLOGI/FLOGI attacks, and name server corruption.

In the next chapter, we will describe the details of the risks identified with HBA and LUN masking. The next chapter will also describe the details of each risk and what factors need to exist in order to perform any attacks that lead to data compromise.