Stephen Sims
GSEC – (GIAC Security Essentials Certification)
Version 1.4b – Option 2
11/07/2003

Case Study:

Meeting the Security Requirements of the Gramm-Leach-Bliley Act (GLBA)

Abstract

In November of 1999, the Gramm-Leach-Bliley Act (GLBA) was passed into law to protect consumers against the disclosure of financial data to unauthorized parties. Under this law, financial institutions are required to provide consumers with detailed information regarding their privacy policies, relative to the sharing of personal financial data. The mandate has many similarities to the Health Insurance Portability and Accountability Act (HIPPA), passed in 1996 to protect personal health information. While the GLBA hasn't experienced the attention or enforcement of HIPPA, I believe that it will soon be addressed with greater vigor.

The greatest effects of the GLBA's requirements are felt primarily in the information security and data networking fields. Financial institutions are faced with the rigorous task of encrypting all customer financial data while in transit and while in storage. While the first phase only concentrates on encrypting data in transit, financial institutions are faced with strenuous time constraints and serious penalties. Failure to meet the requirements on time could result in an audit by the Federal Trade Commission (FTC) or Office of the Comptroller of the Currency (OCC) and huge fines for each day that passes after the deadline. We will be focusing on the efforts and experiences I had while researching, designing and successfully fulfilling the GLBA first phase requirements. Phase two of the GLBA is currently in a negotiation phase with the FTC and OCC. This phase concentrates on the protection of customer data while stored on the server or other data repository.

Before

To most people in society, modern practices such as online banking feel safe. This is likely due to that little padlock that appears on the bottom of your browser while navigating your account. Security measures such as Secure Sockets Layer (SSL), have significantly improved client to server communications over the Internet by utilizing key exchanges and up to 168-Bit, Triple-DES encryption. There are still some risks to be aware of before relying completely on this technology. One risk is exposed during the initiation of an SSL connection. This example is extremely oversimplified. If Jason attempts to connect to Amy using SSL they must first perform a key exchange before they can communicate. If Rick, a third user, can intercept that key exchange, he can then pretend to be Amy and Jason. Rick will take Jason and Amy's public keys and respond back with his own. Since the data from Jason and Amy will be encrypted using Rick's public key, he will then be able to decrypt all of the data.[1] You must also account for the vulnerability level of the users computer. If the users computer has been broken into, the benefits of SSL will not help you. This is just one of a countless number of vulnerabilities.

The benefits of SSL usually end at the web server. Lets focus on Figure 1.1 for a moment. Looking at the communication between the customer and the web server, you'll see that SSL is being used to encrypt the data over the Internet. This encrypted connection protects the customer's data while in transit to and from the web server. Corporate Internet connectivity usually meets the GLBA requirements as long as the data is encrypted while customers are exchanging financial information to and from web servers.

Continuing with reference to figure 1.1, once the web server attempts to query the "customer-records" database, the sensitive data is no longer protected. In this scenario, the data flow between the web server and the DBMS (Database Management Server) must travel extensively to acquire the requested information. First, data leaves the web server unencrypted, such as account information and social security numbers. The data must travel through the banking institutions DMZ (Demilitarized Zone) and out to the California site's LAN (Local Area Network). In passing through this first area, we have already traveled through several switches, firewalls and routers, which all have their own vulnerabilities. The data is then routed out through the California Site's WAN (Wide Area Network) devices and sent through the carrier's public ATM (Asynchronous Transfer Mode) Network to the Texas Site's WAN devices. Now the data must travel through the Texas Site's LAN and to its final destination, the customer database servers. Once the DBMS receives the query from the web server, it will pull the information off of the unencrypted database. The data must then travel back along the path to the web server with customer information, such as account balances and credit card numbers. Anyone with a network sniffer could easily capture sensitive data.

The purpose of this diagram is to give you a logical perspective to what we will be addressing in the next few pages. Having a light understanding of the technical areas that will be addressed later may help to interpret the policy and procedural oriented practices of the GLBA. An understanding of the history and reasoning behind the GLBA must first be looked at in order for one to make the best decision on how to comply.

[1] Seifried, 17 December 2003.

Please review Figure 1.1 as we proceed into the specifications and mandates of the GLBA:

Figure 1.1

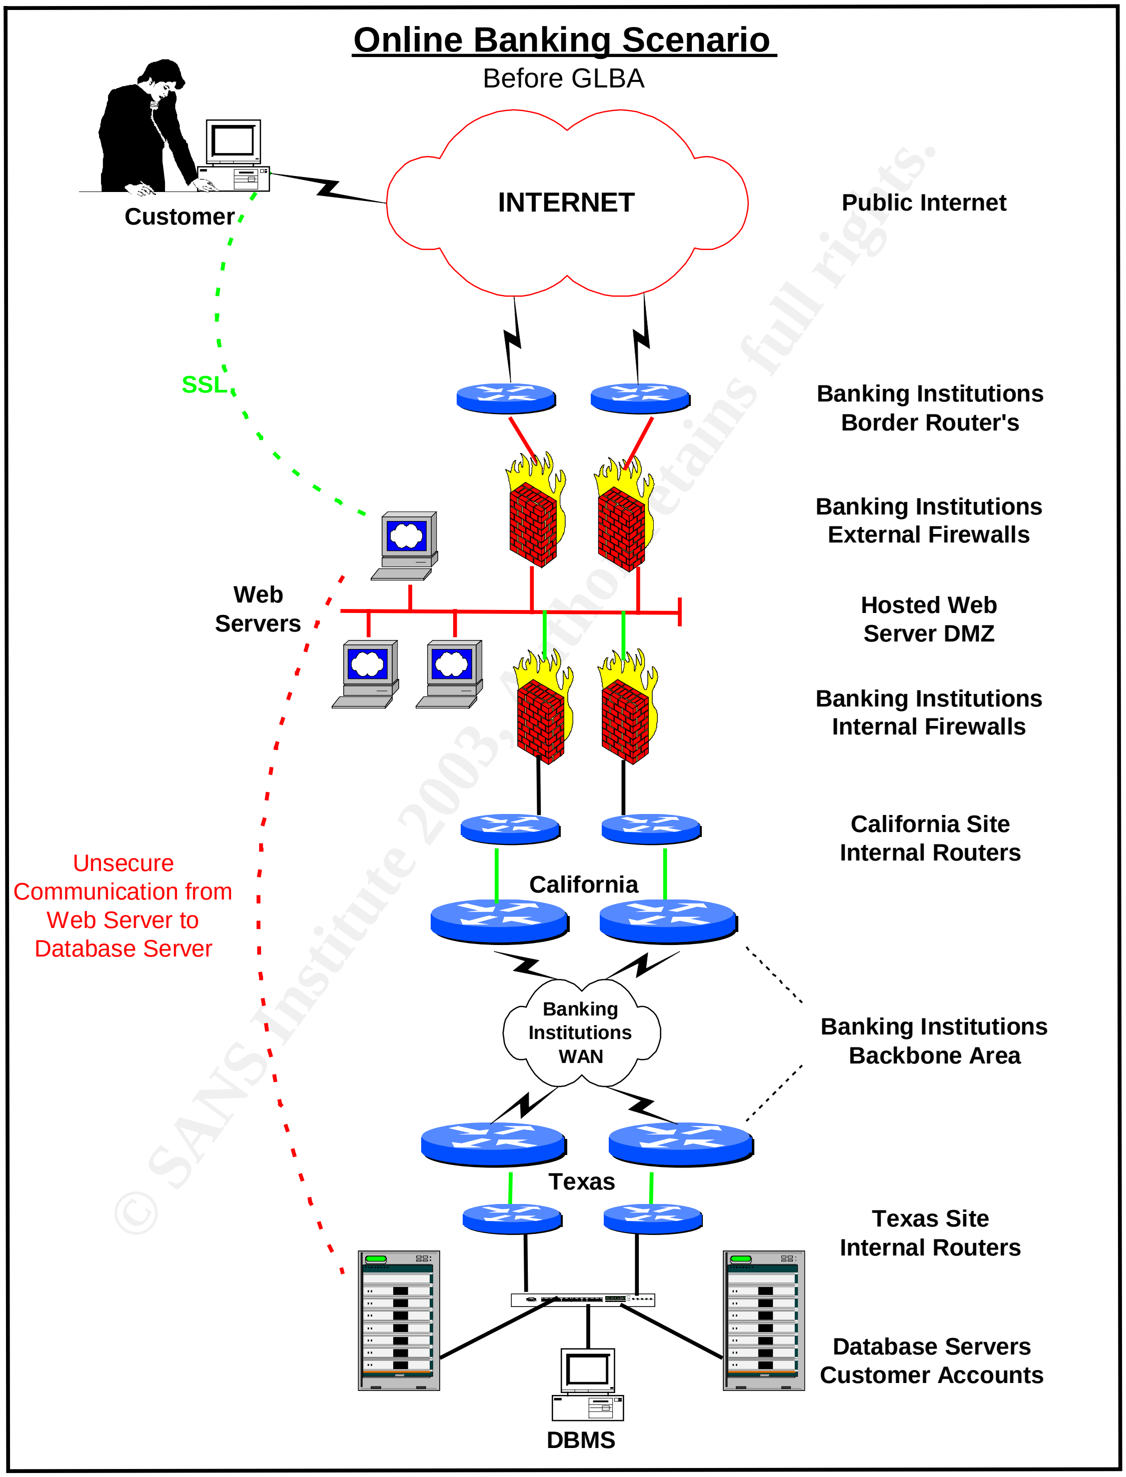
Online Banking Scenario
Before GLBA
Customer
INTERNET
Public Internet
SSL
Banking Institutions Border Router's
Banking Institutions External Firewalls
Web Servers
Hosted Web Server DMZ
Banking Institutions Internal Firewalls
California Site Internal Routers
California
Unsecure Communication from Web Server to Database Server
Banking Institutions WAN
Banking Institutions Backbone Area
Texas
Texas Site Internal Routers
Database Servers Customer Accounts
DBMS

While researching the requirements of the GLBA, I was struggling with an enormous amount of ambiguity from the FTC and OCC. The specifications were very high level and seemed to focus on addressing the requirement of financial institutions to inform consumers of their privacy policies. A large subsection focuses on the disclosure of financial data to unauthorized parties. The matrix below shows the requirements for customer notification, regarding privacy policies as stated by the FTC.

SUMMARY OF NOTICE REQUIREMENTS			
Type of Notice	To Whom	When	Contents
Initial	Customers	Not later than when you establish the customer relationship, unless it would substantially delay the transaction and the customer agrees	Description of information-collection and sharing practices, and opt-out notice (if you share NPI with nonaffiliated third parties outside of certain exceptions)
	Consumers who are not customers (including former customers)	Before you disclose their NPI to a nonaffiliated third party outside of certain exceptions	Full description of information-collection and sharing practices *or* "short-form" notice, along with opt-out notice
Annual	Customers	Delivery on a consistent basis at least once in any period of 12 consecutive months for the duration of the customer relationship	Description of information-collection and sharing practices, and opt-out notice (if you share NPI with nonaffiliated third parties outside of certain exceptions)

[2]

The only area on the FTC's website that addresses any technical requirements are in the Safeguards Rule portion of the Gramm-Leach-Bliley Act. The Commissions Financial Privacy Rule (Privacy Rule) and the rules regarding pretexting focus more on the consumer information areas of the GLBA. We will begin with addressing the Privacy Rule to allow for an understanding as to what was involved from a business perspective in the research portion of this project.

[2] FTC. July 2002.

The Privacy Rule concentrates primarily on the consumers right to know what a financial institutions information sharing policies are. This enables the consumer to make some decisions as to what data they wish to have shared and what they would like protected. Even if a company has the right to share specific types of consumer information, they must make it known in writing repeatedly based on defined recurring time increments. The following is the high level specification as to what must be shared, as written by the FTC.

> **The Privacy Notice**
>
> The privacy notice must be a clear, conspicuous, and accurate statement of the company's privacy practices; it should include what information the company collects about its consumers and customers, with whom it shares the information, and how it protects or safeguards the information. The notice applies to the "nonpublic personal information" the company gathers and discloses about its consumers and customers; in practice, that may be most - or all - of the information a company has about them. For example, nonpublic personal information could be information that a consumer or customer puts on an application; information about the individual from another source, such as a credit bureau; or information about transactions between the individual and the company, such as an account balance. Indeed, even the fact that an individual is a consumer or customer of a particular financial institution is nonpublic person information. But information that the company has reason to believe is lawfully public - such as mortgage loan information in a jurisdiction where that information is publicly recorded - is not restricted by the GLB Act.
>
> **Opt-Out Rights**
>
> Consumers and customers have the right to opt out of - or say no to - having their information shared with certain third parties. The privacy notice must explain how - and offer a reasonable way - they can do that. For example, providing a toll-free telephone number or a detachable form with a pre-printed address is a reasonable way for consumers or customers to opt out; requiring someone to write a letter as the only way to opt out is not.
>
> The privacy notice also must explain that consumers have a right to say no to the sharing of certain information - credit report or application information - with the financial institution's affiliates. An affiliate is an entity that controls another company, is controlled by the company, or is under common control with the company. Consumers have this right under a different law, the Fair Credit Reporting Act. The GLB Act does not give consumers the right to opt out when the financial institution shares other information with its affiliates.
>
> The GLB Act provides no opt-out right in several other situations: For example, an individual cannot opt out if:
>
> - a financial institution shares information with outside companies that provide essential services like data processing or servicing accounts;
> - the disclosure is legally required;
> - a financial institution shares customer data with outside service providers that market the financial company's products or services.

[3]

Pretexting is another condition under the GLBA. This is the act of impersonating a consumer to obtain private information about them. Pretexting also includes the sale of ones personal information to another person or business. A "pretexter" may attempt to acquire enough information about you to enable them the ability to impersonate you. This act could allow them to obtain personal information about you such as banking information, health information and other personal data. These requirements pertain primarily to the legal side of the business.[4]

The Safeguards Rule is the closest you will find to a definitive answer as to what the technical requirements are. The documentation regarding the Safeguards Rule clearly acknowledges that Information Security is an ever-changing field. With the perpetual growth in the knowledge and skills of hackers and the like, defining specifics would be counterproductive in the effort to protect consumer information. Rules and guidelines must be updated as the technology and know-how grows. There were proponents on both sides as to how detailed the document should be. There were also multiple sides to the many arguments in the effort to define a policy that could be applicable to both small and large businesses. To require a smaller company to achieve the same level of due diligence and compliance as those with a larger and broader knowledge base seemed unjust to some. One of many points addressed was that, "some commenters expressed concern about the ability of businesses – particularly smaller entities – to evaluate a service providers capabilities."[5] The primary defense for this was again, a smaller company may not have the technical expertise to ensure that a service provider is meeting the specifications of the GLBA. This defense was recognized by amending the document to state that a company must take reasonable steps and perform a high level of due diligence when selecting a service provider.

The final result in the attempt to develop the most comprehensive and fair Safeguards Rule is both definitive and general. It is definitive by listing specific, high-level security requirements. The general side resides with the ability to interpret those guidelines in many ways. When asked the question, "What is risk?", I was faced with a countless list of responses. When focusing from the higher level questions down to a specific area or system, the list got even larger. The Safeguards Rule states that you must perform a high level of risk assessment, require secure design policies and best practice, ensure employee awareness and education, and enable detailed monitoring and auditing.[6] (REFERENCE) That is what you are given to go on. The level of enforcement on these guidelines is based on the discretion of the FTC. The general idea is

[3] FTC. "In Brief: The Financial Requirements…"

[4] FTC. January 2001.

[5] FTC. 23 May 2003.

[6] FTC. September 2003.

that a large business should have a strong enough understanding and expertise as to how you must fulfill these requirements. Small businesses are given some educational materials, explaining how they can meet the specifications. As you could imagine, the development and push of the GLBA and HIPPA have resulted in a flurry of security consulting companies claiming to fully understand the requirements of both. In my experience, there was not a strong level of understanding by the consultants.

The OCC had given us a questionnaire to complete, as well as a Compliance Guideline Document. The questionnaire focused on the same areas as the FTC's Privacy Rule and is included at the end of this document as Appendix A. The Compliance Guideline Document was more of an auditing tool. It goes into great detail about investigating a companies security practices. Some examples of the questions are: "How does the institution assess risk?" "Does the Institution support its estimate of the potential damage posed by various threats?" "What is current the level of encryption of electronically transmitted and stored data?". [7] It is this document that should be used as the primary factor in determining your current level of compliance.

When meeting with the OCC to discuss the GLBA's security specifications and time constraints we must follow, the requirements became clearer. Consumer data must be encrypted while in transit and while at rest. The level of encryption was not specified. Only that the best practice is used and due diligence is performed when selecting the solution. All vendor connections, internet connections, intranet connections and WAN connections had to be identified and addressed. This was a time consuming process. The OCC broke the deadlines down into two phases. The first phase was to address the protection of customer data while in transit. This phase primarily focuses on data connections that leave a physical building or structure such as DSL (Digital Subscriber Line) connections, DS-3's, T1's and other various high bandwidth data circuits. We were given a strict eight-month deadline to meet this requirement by. The failure to comply on time would have resulted in a large fine for every day that passed from the due date. Phase two was to protect customer data while in storage. This primarily focuses on databases and workstations that contain and store financial data. There are a lot of gray area's with phase two. The most difficult being whether or not connections within a LAN, from a server or workstation to a database would need to be encrypted? Without application layer encryption to and from the database server, this makes for an extremely difficult task. Believe it or not, most companies do not use application layer encryption on a majority of their applications that are not web based. This case study focuses on how the first phase was identified and met. Phase two's terms and conditions are currently being researched and negotiated between the OCC, FTC and various other financial regulatory agencies. In the meantime, simple rules such as the requirement to lock workstations when unattended have been enforced. As well

[7] OCC, 2001.

as an attempt to ensure that all new applications used have security built into them.

During

Now that we have covered the basis behind the GLBA we can focus on the solution. From a manager's perspective, there were several deliverables that were deemed a priority. First, was to ensure that we met the deadline dictated by the OCC in order to not be penalized. Second, to find the best resolution at the lowest cost to the company. There were no quantitative drivers associated with this project to help acquire a large budget from upper management. There were only qualitative benefits and an obvious deadline requirement. The last priority was to make sure all potential areas of risk were identified and understood. One of the most difficult challenges with this project was to understand what would truly satisfy the OCC and the FTC's requirements. I was assigned with leading the Phase One initiative and seeing it to its completion. This involved the research, design, engineering, and implementation of the chosen solution.

The first objective was to identify all potential areas of risk. Starting with the WAN connections that were the most obvious, I began mapping out the entire data network and its twenty sites. There were a combination of Frame-Relay, ATM DS3's, ATM OC-3's, WAN Ethernet and point-to-point T1's. All areas were identified, logged and verified through other engineers. I then moved onto the Internet facing connections. There were multiple DS3's to the Internet. Some were for simple Internet access to allow customers to connect to our web servers and others were setup primarily for VPN connections to vendors and employee remote access. Finally, all vendor and extranet connections were identified. After all of the potential areas of risk were documented, I asked for representatives from various departments within the organization to verify my findings and to ensure there were no connections missing.

Being that the customer data needed to be encrypted while in transit, the first thought was to look into upper layer encryption such as SSL, which works at the Transport Layer. Encryption technologies such as SSL can be built into an application to provide end-to-end security. The primary driver for this type of solution was that it would take an enormous amount of strain and cost off of the network layer devices. Network layer encryption at a glance seemed to be the most expensive. A meeting was then arranged between the application development teams, server teams and data network teams. The first reaction from the application development team was that it would be impossible. There were multiple reasons for this decision. Most of the applications ran internally within our network did not support SSL or any other forms of encryption. They were not designed this way and cannot be easily altered to support it. The estimated timeframe given from the application developers was approximately two years, maybe longer. They did not have experience in developing this technology and estimated that the cost would be greater than that of a network layer solution. There were also concerns over the way in which we stored customer credit data off site to a vendor. The possibility of working with them on altering their methods of data storage to support only one client was not likely.

Management then removed the possibility of upper layer encryption due to the time constraints and concern over the success of such a solution. It was soon clear that the only method of security that could possibly meet the deadline was a lower layer solution. All attention was then focused on data-link and network layer based solutions. We had used data-link encryption before in other countries where it is a requirement by law on WAN connections. We had only used this technology on Frame-Relay connections. The first question was whether or not there were layer two link encryptors for ATM? A full line of products was found by a company called Cylink, currently known as SafeNet. This is the same company in which we had purchased the frame-relay encryptors from. They offered layer two encryption for frame-relay, ATM OC-3 & DS3 and T1 connections. Other vendors were found offering hardware encryption products, but were very limited. Most of them could not support high-speed connections with their products and did not have the customer track record to satisfy our requirements.

Some of our WAN connections were using an Ethernet-based solution offered by MCI and AT&T. Within certain distance limitations you have the option of using a long distance Ethernet connection for site-to-site connectivity. This provides the option of extending VLAN's (Virtual Local Area Network) across a WAN connection without the requirement of a router. It is a good benefit, but there are no products that offer a hardware based layer two encryption solution for Ethernet. There was never a demand for the product since there are alternative offerings such as VPN, along with the fact that Ethernet doesn't normally leave a building. This issue requires that there be a device on both sides of the WAN Ethernet connection to allow for a layer three encryption solution, such as VPN. Using a VPN device and forcing a layer three solution takes away from some of the reasoning and benefits of using a wide area Ethernet solution. Some of those benefits include having 100 Mbps of bandwidth site to site, extending a broadcast domain across a wide are connection, and paying a much lower cost than the use of more traditional wide are connections such as DS3.

With a layer three solution being the feasible option, the process began of looking into the available products to support this. There were many product options to support an IPSEC solution. We looked into Cisco, SafeNet, Nortel, NetScreen and a few others. All had comparable products that support thousands of simultaneous SA's (Security Associations) and a high throughput rate. At this point the decision was to request some pricing quotes from the various vendors and put together a proposal to management to obtain the funding.

At this time the business decided to bring in an outside contractor to assess our situation and requirements. They wanted to obtain a second opinion to satisfy some legal requirements. A vendor was awarded the short contract and began to research our obligations for the GLBA. After a couple weeks of assisting them with network diagrams and consultation, they came up with the same solution we had already proposed to management. At this point the decision was made to move forward with the proposed solution and begin the purchasing and acquisition process. Since there were not many opportunities for ROI (Return On Investment) associated with this project, the most affordable solution was the most desired.

During the investigation by the outside contractor, we began to bring in some of the various vendors to deliver presentations on their products. We had Cisco demonstrate their VPN solutions, SafeNet demonstrate their hardware encryptors and VPN solutions, as well as a couple of other vendors. We then asked SafeNet and Cisco to supply us with some demonstration models of their products so we can begin testing in a lab environment. Each device of interest was tested thoroughly with the support of a sales engineer. SafeNet's hardware encryptors all tested well when configured to model our network. The NetHawk devices did not have the same luck during that time. Even though they were designed for VPN and had a strong track record, we learned that our configuration was not going to be a standard use of the product. Most VPN devices are used to support either Business-to-Business (B2B) or client remote access connections. They are not designed to handle site-to-site Ethernet connections and the intricacies involved with them.

If you look at drawing Figure 2.1, you will see a simple scenario of three sites connected by wide area Ethernet through a carriers cloud. Each site has a single connection to the cloud for the MAN (Metropolitan Area Network) and a circuit to the corporate Internet. In order to reduce cost and still achieve a full mesh in the MAN, each site can peer with the other two while only having the single connection to the carrier. EIGRP (Enhanced Interior Gateway Routing Protocol) was the routing process used in this scenario. Prior to introducing the NetHawks into the design, each site had two adjacencies and we achieve the full mesh, thus removing the requirement for local intercommunications to traverse the corporate Internet to reach its destination. When introducing the NetHawks we immediately ran into some issues.

If you could first imagine having a corporate network with over 500 summarized networks being advertised across the backbone and between the various sites. Now focus in on the MAN portrayed in Figure 2.1. In the normal use of a VPN device for a B2B connection, you would create a rule allowing your public VPN address to establish an IPSEC tunnel with another businesses public VPN address. You would then define a crypto-map and access list with the networks permitted to enter the tunnel. Each time a user connects across the link a separate SA will be created. The problem in this scenario is exposed when trying to define and manage hundreds of separate VPN rules to account for the entire corporate Internet. Every site in the MAN has a connection to the corporate Internet, which means that their respective routers each know about all of the possible networks. Trying to configure all interconnected VPN devices to allow the same networks to pass on their protected sides simply does not work and will not be permitted in the configuration. If a properly secured VPN device sees a network that belongs on its trusted side coming in on an untrusted interface, it will drop the request due to the anti-spoofing rules. The solution developed was to create separate GRE Tunnels (Generic Routing Encapsulation) between each of the sites across the wide area Ethernet connections. This only required the need for two VPN rules in each device. Since the tunneling of network traffic occurs before it hits the physical interface, the VPN device will only see traffic from that physical interface. They must permit the physical interface of its local router to communicate with the physical interfaces of the other two sites routers. You must remove the address of the physical interface of the router from the routing process and add in the address of the tunnel interface. The adjacencies will establish using the tunnel interfaces and route all traffic though them between the sites in the MAN.

Note that this is only a requirement when connecting multiple sites through a single connection. When using a point-to-point wide area Ethernet connection you may define a rule stating “Any Behind” is permitted to communicate with “Any Destination”. This causes the device to ignore its anti-spoofing policy. The primary issue was now resolved and the configuration concerns removed. This would still make for a detailed troubleshooting process to be created, due to its complexity. A secondary area of concern was regarding the proper communication between the NetHawk devices and the NetHawk Manager. Routing must be set up to ensure that the path the manager takes to communicate with a NetHawk will be the same on the return back to the manager. Not ensuring this will cause potential stateful inspection issues depending on the return path. Stateful Inspection, created by Check Point Software, is a technology used by devices such as firewalls to verify the integrity of a connection. It works like packet filtering with added features. It analyzes the layers above the network and transport layers to evaluate factors other than only the source and destination stored in the header. It also keeps a table of all traffic passing through its interfaces to help prevent attackers from gaining unauthorized access. You must also ensure that the communication from the NetHawk manager to a remote NetHawk device does not get encrypted by another NetHawk on its path to the destination device. You must supply a rule stating not to encrypt the management traffic in any circumstance.

Figure 2.1

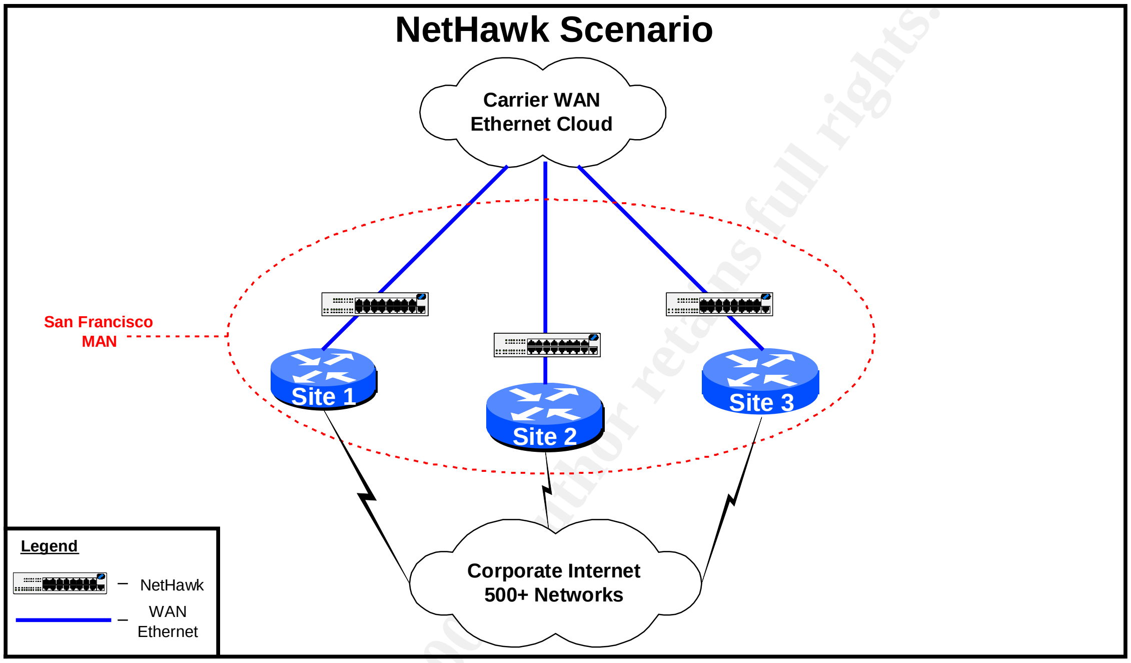

With all of these factors in mind we decided to move forward with the procurement of the hardware encryptors from SafeNet for our corporate Internet connections. We had a combination of ATM OC-3 and ATM DS3 circuits for a majority of the primary sites. There were approximately twenty-four separate ATM circuits that needed to be secured. Eight OC-3's and sixteen DS3's. Each ATM encryptor is capable of handling up to 4,096 virtual circuits with independent keys and up to 622 Mbps with the optional OC-12 module. They support Triple-DES encryption with Diffie-Hellman key exchange. These devices averaged about $40,000 each depending on the type of interface chosen. There were approximately ten point-to-point T1 connections that had to be addressed. We chose to use the T1 encryptors offered by SafeNet. These devices support AES or Triple-DES encryption with Diffie-Hellman key exchange. They support up to 52 Mbps on point-to-point connections with a series of optional interfaces. The average cost for these devices were $4,500 each. There were four frame-relay connections for which we purchased frame encryptors. These units cost approximately $8,000 each and also support Triple-DES encryption and Diffie-

Hellman key exchange. The purchase requisition was approved and the products were ordered.[8]

While the first requisition was making its way through our purchasing department, I began working on the other areas that needed a solution. These areas were the before mentioned wide area Ethernet connections, vendor connectivity and Internet connections. Lets first address the Internet connections, as they were the least time consuming. Customers that connect to our web servers over the internet must have SSL enabled on their browsers. In order to obtain any private information, each customer must supply their username and password and authenticate to our web servers. All of this is done using SSL encryption. This element already meets the GLBA requirements. Vendor and employee access over the Internet uses VPN. Our VPN IKE policy uses Diffie-Hellman type II key exchange and Triple-DES encryption. This also meets the GLBA requirements. Internet-facing connectivity, the most untrusted area of a network, seemed to require the least attention. This is highly based on the assumption that most attackers will attack from outside of a trusted network. This is a fairly arguable point.

Vendor connectivity was another time-consuming area to research. Some vendors, fortunately already used application layer encryption or other upper layer forms of security. These vendors were quickly identified and placed on the checklist as complete. At the time there were approximately thirty vendors that had to be investigated. Some had multiple connections for redundancy, and others with only a single connection. The vendors that were identified as still being a risk were contacted and put into multiple categories. A few of the vendors had the technology in place to allow us to switch their connections over to a B2B VPN solution using pre-existing Cisco or Nortel equipment. Some of the vendors offered the option to change their applications to a browser-based technology in order to allow SSL to assist in meeting the GLBA requirements. There were also the remaining few vendors that either required us to purchase the layer two encryption devices or did not have the technical skills to help us meet the requirements. The internal business units that held the contract with each vendor were required to negotiate with those who did not have the technical skills or did not wish to assist in purchasing the encryption solutions.

The wide area Ethernet connections were the last to be addressed. The decision was made to move forward with SafeNet on their NetHawk VPN solution. At this time most of the concerns had been rectified in the lab. We were able to work out a discount on the pricing due to our prior acquisition of the hardware encryptors. This was the primary driver behind choosing their product, combined with the NetHawks capabilities and customer references. They support ESP or Triple-DES encryption and Diffie-Hellman key exchange. Each device has two Ethernet ports on them, one for the inside network connection and the other out

[8] SafeNet, May 2003

to the carrier side. These devices were approximately $4,000 each and we purchased twelve of them.

It was at this point that we were at the five-month mark out of our eight-month deadline. We received word of some delivery dates for the ATM and T1 encryptors. About half of the encryptors ordered were to arrive in a month. Now began the process of submitting change requests to introduce the devices to the network. The first sections implemented were the ATM OC-3 connections, followed by the ATM DS3 connections. The SafeNet ATM Encryptor will allow you to map a large number of Permanent Virtual Circuits (PVC) off a single ATM circuit. With ATM, a single high bandwidth circuit can be divided with PVC's to connect to multiple locations. This is done by creating sub-interfaces off of the primary interface. By using a unique identifier provided by the carrier, you can map these sub-interfaces to many sites. We started with only two locations to minimize the impact in the event of a failure. The ATM encryptors have two possible modes when configuring PVC's. If you install an encryptor on a connection with multiple PVC's, you must define each PVC to be in either "secure-mode" or "clear-mode". If there is an encryptor on the other side of the PVC, the connection can be established in "secure-mode" and pass encrypted traffic. If there is not an encryptor on the other end of the PVC, it must be set to pass the traffic unencrypted in clear-mode. The ATM PVC's between the two sites where the encryptors were added were set to communicate in secure-mode. The remaining PVC's to all of the other sites were configured in clear-mode. This implementation was successful and followed with a series of implementations to complete the ATM encryptor installations. The only issue that arose in relation to the ATM encryptors was intermittent management connectivity failures. You can configure these devices via its console port or by the "Encryption Privacy Manager." Some of the devices would occasionally lock up the console port access and the remote management connectivity. The devices would still pass encrypted traffic, but their admin interface was inaccessible. This became an issue during implementation when we needed to modify the configuration to change a PVC from clear-mode to secure-mode. The issue was resolved by having our network operations group open up a ticket with SafeNet to replace the devices. This problem was apparently caused by a processor bug, which was later fixed by Motorola.

These implementations were followed by the addition of the T1 and Frame-relay encryptors. An interesting concern arose at this point with a pair of the Frame-relay connections. One of our sites uses Frame-relay on the site side and connects to ATM OC-3's on the remote side. This is a service provided by some of the telecommunications carriers known as FRF.8.1. It allows for the communication and translation between ATM and Frame-relay PVC's. This is basically done by mapping the Frame-relay frame into an AAL5 PDU and by stripping various attributes of the Frame-relay frame such as the flags and CRC-

16.[9] The reverse is done through a similar method. Basically, it enables a carrier to translate between the two communication types. This creates a problem when attempting to have an ATM encryptor communicate with a Frame-relay encryptor. There was an option with a Cisco supported FRF.8.1 module and IOS. When an ATM OC-3 encryptor receives traffic encrypted by a Frame-relay encryptor, it will still pass it to the next hop as it can still read the destination address. This hop being a Cisco device supporting FRF.8.1 would then forward the traffic out a separate interface that connects to a Frame-relay encryptor. This Frame-relay encryptor would decrypt the traffic and send it back to the Cisco router to be routed to its destination. The same would be done in the other direction. This option was not cost effective and the idea was discarded. We ended up going with a Cisco-based router-to-router VPN solution.

The final devices to be added were the NetHawks. There was a level of concern when rolling out these devices. They were carefully introduced into the network starting with the least impacting. Due to the high level of due diligence performed in the lab, the NetHawks were added with minimal problems. There were two issues that arose. The first issue was in relation to Maximum Transfer Unit (MTU) size and packet fragmentation. The default MTU on the Ethernet interfaces of most devices is 1500 bytes. When a 1500 byte packet arrives at a router using GRE it will attempt to add on 24 bytes for the GRE header during the encapsulation process. Since the MTU on the interface is set for 1500 bytes, the router will be forced to fragment the packet into two packets. In a lot of situations the originating device will set the Don't Fragment (DF) bit in the header of the packet. This causes the router to drop the packet and send an ICMP echo back to the sending device. As you could imagine this caused a lot of losses in connectivity and retransmission attempts. We were forced to set the MTU size on the Ethernet and tunnel interfaces down to 1400 bytes to allow for the GRE header size. The server teams were also required to set the MTU size on their servers down to 1400 to avoid fragmentation and the latency caused by it. The other issue was an anticipated one regarding asymmetrical routing to the manager. Manipulating the routing processes metrics easily rectified this problem. By changing the metrics using route maps you can force one path to be taken over another. At this time we were in compliance with phase one of the GLBA requirements approximately one week before the eight-month deadline.

After

Now that we were in compliance with phase one of the GLBA, we needed to ensure that the entire network was documented with all of the new additions and changes. The ability for an engineer, whom is unfamiliar with this type of solution, may have difficulty troubleshooting a potential problem. Data communications between all of our wide area connections are now fully

[9] FR Forum, February 2003.

encrypted and the mechanisms in place to ensure only intended parties and traffic are permitted in. The threat of an external party sniffing or tapping our communications was removed. All captured data would be encrypted with a 168 Bit cipher and the key would be virtually impossible to crack due to the Diffie-Hellman algorithm. The preshared key policy in place is using the strongest combination of characters, numbers and upper and lowercase letters with a high minimum length requirement.

With phase one complete, phase two would soon need to be addressed. Phase two has not been fully outlined by the OCC. Its primary focus is on the protection of consumer financial information while in storage. This mainly involves database servers and the communications with them. The OCC has not given us a deadline to be in compliance by, as there is some uncertainty to what the minimum requirements are. Would it require only that financial data stored on a database by encrypted, or would it also require client-to-database and server-to-database encryption? The later of the two would require an upper layer encryption solution. The idea of an IPSEC tunnel from every host or server that connects to the DBMS for information would be an almost impossible task. Phase One was enough to make me question my sanity. The overhead caused by and processor power needed to drive that solution would be great. It would also be difficult for the financial regulatory agencies to require every financial institution in the United States to change all applications to support upper layer encryption. Fortunately, a majority of sensitive applications are moving towards using this feature as a default, but it could be some time before we see enough to justify it as a requirement.

Database encryption will probably be the only requirement of phase two at this time. There are several concerns when understanding the topic of encrypting databases. It significantly increases the size of the stored data and delays the access time. This forces security vs. performance to be taken into serious consideration when deciding on a solution. Specific data must be analyzed to see what exactly should be encrypted and what can be left in the clear. There are a couple of reasons why you would not want to encrypt the data on the database itself. One reason is that the data would be in clear text until it reaches the database server, allowing someone to sniff the data as it traverses its way from the DBMS to the database. Another reason not to use “On-Database” encryption is that it would force all data to be encrypted and not allow the option to choose which data you want to encrypt. Encrypting all data will cause serious overhead drawbacks and an overall lack of performance. An alternative to encrypting on the database server itself is to encrypt at the database management system. There are still a couple of drawbacks, such as administrative access to the encryption keys, but overall it allows for much better performance. This allows us to choose which data to encrypt and which to leave in the clear.

There will always be potential risks to an organization as long as there are threats and vulnerabilities. We can only protect ourselves by using best practices and introducing various obstacles to reduce our level of threat. By performing recurring audits of your network, enforcing security policies and awareness, and staying current with the latest security vulnerabilities and patches, you can greatly mitigate your risk. Utilize all of the available tools and documentation on the GLBA to beat the push from the financial regulatory agencies to be in compliance. It is only my assumption that there will be an increase in activity for all applicable companies to move forward on meeting the requirement in the near future.

Acronym Guide

ATM	–	Asynchronous Transfer Mode
B2B	–	Business-to-Business
DBMS	–	Database Management Server
DES	–	Data Encryption Standard
DMZ	–	Demilitarized Zone
ESP	–	Encapsulating Security Payload
FTC	–	Federal Trading Commission
GLBA	–	Gramm-Leach-Bliley Act
IPSEC	–	IP Security Protocol
LAN	–	Local Area Network
OCC	–	Office of the Comptroller of the Currency
PVC	–	Permanent Virtual Circuit
ROI	–	Return On Investment
SA	–	Security Association
VLAN	–	Virtual Local Area Network

VPN	–	Virtual Private Network
WAN	–	Wide Area Network

References

-1-
Seifried, Kurt. The End of SSL and SSH. 17 December 2000.
URL: http://www.seifried.org/security/cryptography/20011108-end-of-ssl-ssh.html

-2-
Federal Trade Commission. How To Comply with the Privacy of Consumer Financial Information Rule of the Gramm-Leach-Bliley Act July 2002.
URL: http://www.ftc.gov/bcp/conline/pubs/buspubs/glblong.htm

-3-
FTC. In Brief: The Financial Privacy Requirements of the Gramm-Leach-Bliley Act.
URL: http://www.ftc.gov/bcp/conline/pubs/buspubs/glbshort.htm

-4-
FTC. Pretexting: Your Personal Information Revealed. January 2001
http://www.ftc.gov/bcp/conline/pubs/credit/pretext.htm

-5-
Federal Trade Commission. Standards for Safeguarding Customer Information; Final Rule 23 May 2002.
URL: http://www.ftc.gov/os/2002/05/67fr36585.pdf

-6-
FTC. Financial Institutions and Customer Data: Complying with the Safeguards Rule. September 2003
URL: http://www.ftc.gov/bcp/conline/pubs/buspubs/safeguards.htm

-7-

OCC. Examination Procedures to Evaluate Compliance with the Guidelines to Safeguard Customer Information. 2001
URL: http://www.occ.treas.gov/ftp/bulletin/2001-35a.pdf

-8-
SafeNet. Safe Enterprise ATM Encryptor. May 2003
URL: http://safenet-inc.com/products/documents/safe_enterprise/SAE_ATMEncryptor_ProductBrief_may2003.pdf

-9-
FR Forum. Frame Relay/ATM PVC Service Internetworking Implementation Agreement. February 2000
URL: http://www.frforum.com/5000/Approved/FRF.8/FRF.8.1.pdf

APPENDIX A

Privacy Preparedness Questionnaire

Assessing Existing Information Practices

1. What are your information-sharing practices?

* What information is shared with affiliates and nonaffiliates (including sharing within and outside of the regulatory exceptions contained in 12 CFR 40.13, 40.14, 40.15), what is the purpose of the sharing, and is information shared on former customers?
* Are account numbers or access numbers/codes disclosed to nonaffiliated third parties?
* What information do you share on consumers who are not customers?
* Do you route requests for nonpublic personal information to a central point or use other control measures?
* Will any of your current information-sharing practices be prohibited by the regulation?

2. What kinds of information do you collect from consumers and customers for the various financial products and services offered by the bank?

3. Do you obtain information about consumers and customers from other financial institutions? If so, do you use or share the information for other purposes?

4. Are your safeguards for protecting customer information consistent with Section 501(b) of the Gramm-Leach-Bliley Act?

* Has the board approved the written information security program?
* Are your safeguards adequate to: a) ensure security and confidentiality of customer records and information, b) protect against any anticipated threats or hazards to the security or integrity of customer records and information, and c) protect against unauthorized access to, or use of, such records or information that could result in substantial harm or inconvenience to any customer?
* Has your information security program been tested in accordance with the regulatory guidelines?

Evaluating Agreements with Nonaffiliated Third Parties that Involve Disclosure of Consumer Information

5. What arrangements, agreements, or contracts exist with nonaffiliated third parties that involve disclosing consumer information? Do contracts or agreements detail responsibilities regarding the use, disclosure, and protection of consumer information?

6. What changes need to be made to conform the arrangements, agreements, or contracts to the regulation?

Establishing Mechanisms to Handle Opt-Out Elections

7. If applicable, how will you administer the opt-out provisions of the regulation?

* Is the opt-out mechanism reasonably convenient for the consumer to use?
* How will you document those consumers who opt out or later change their opt-out status, and how will you segregate their information?
* How much time will you allow for consumers to opt out and how quickly will you process opt-outs?
* What are your opt-out arrangements for consumers who jointly hold a financial product or service?
* Will you allow partial opt-outs? If so, under what circumstances, and are your record- keeping systems capable of handling that level of complexity?

Developing a Privacy Policy

8. Have you developed a privacy policy? If so, what is it?

* Does the policy contain all relevant disclosures required by the privacy regulation?
* Is the information in the privacy policy stated clearly and in a way that consumers are likely to understand? Is it presented in a way that is likely to call the consumer's attention to the nature and significance of the information in the notice?
* Has the policy been reviewed by the board and senior management, the compliance officer, and legal counsel?
* Does it reflect your actual practices?
* Do you think your customers will accept your privacy policy?
* Does the institution have a process to ensure that privacy policies are kept current?

Delivering Privacy Notices

9. How will you deliver initial, annual, and revised privacy notices, and opt-out notices to customers, consumers, and customers who jointly hold a financial product or service?

* Will you hand deliver notices to individuals conducting transactions in person?
* Will you mail the notices, and if so, will you mail them with other information, such as account statements, or separately?
* Do you intend to deliver any notices electronically? If so, how will you obtain the consumer's/customer's agreement to receive electronic delivery?

Establishing a Training Program

10. Describe your plan to train employees on privacy.

* Who will be trained, when, and what information will be covered?
* Will there be different levels of training depending upon job responsibilities?

Establishing a Compliance Program

11. Describe audit's/compliance review's role in developing and implementing the bank's privacy program.

12. Have internal controls, policies, procedures, and audit programs been established to ensure a satisfactory level of compliance?

Developing an Implementation Plan

13. Describe your implementation plan.

* Has the plan been approved by senior management and the board?
* Does it contain target dates, responsibilities, responsible parties, testing procedures, and progress reports?
* Is the plan on schedule?
* Does the plan ensure delivery of the privacy policy prior to July 1, 2001, and afford customers a reasonable time to exercise any opt-out rights before that date?

Source:
OCC. Privacy Preparedness Questionnaire.
URL: http://www.occ.treas.gov/ftp/advisory/2001-2a.txt