Global Information Assurance Certification Paper

Richard Ginski

SANS Security Essentials, GSEC Practical Assignment, Version 1.2e

Information Security Implementation for a Local Government

Table of Contents

Richard Ginski

SANS Security Essentials, GSEC Practical Assignment, Version 1.2e

Information Security Implementation for a Local Government

Introduction

This paper is a case study of a local government organization and its process of implementing information security. These are historically true accounts while maintaining anonymity. Our implementation of Information Security has been a learning process. Much of what has been implemented was based on a plan of action, but we are always learning more about Information Security. Therefore, our plan of action has been flexible.

Purpose

The purpose of this paper is to explain the uniqueness of local government as it relates to Information Security and explain what components of Information Security we implemented. Finally, the paper will describe the considerations made in the security components we chose to implement and why we chose the types of products we did.

Differences in Local Government

Some of the unique qualities for local governments are budget constraints, no single leader, equally autonomous agencies, public record, and what focus is placed on security. The combination of these factors can make implementation of information security a difficult challenge.

First, funding can be very difficult for any organization. However, in local government you can also add a factor of unpredictability. In other words, you never know what kind of budget you will have from one fiscal year to the next. This makes any type of long term planning, particularly in Information Security, very difficult.

Secondly, in some levels of local government such as county government, there is no clearly defined leader. In other words the entity has no president, no governor or mayor. This makes it particularly difficult to have direction and plan for the future. Another aspect of this issue is the fact that there can be many elected officials as part of single local government entity; all using the same enterprise network. Further, each elected official or group of elected officials are considered to have equal power. Therefore, consensus is absolutely necessary for anything to get accomplished. Of course, this impacts the development of information security. To further complicate matters, agencies can have their own IT departments and develop in many different technological directions. This autonomy can result in varying opinions on security, varying degrees of

security awareness, and varying degrees of security implementation. The important point to be made is if they all share the same enterprise network, it's imperative that all agencies be on the "same page" as far as Information Security.

Another difference I would like to point out is the issue of public record and the protection of privacy for its citizens. There is an extremely difficult balance that takes place here. Many states require information in possession by a local government to be public record. Yet, if all information was to be freely available, it could cause a lot of harm to the citizens and also the local government itself. So the issue is a balance between conforming to public records law without causing harm to citizens and the local governments they serve. For example, it is certainly important for our organization to make available an Internet web site that its citizens could use to find their flood zone by entering a street address, during the seasons of flooding. However, vehicle license plate registration information would not be good information to make available on an Internet site. What if someone were to be involved in a "road rage" incident? A potential perpetrator could go home, look up the license registration information, perhaps find where a person lives, and cause harm to the other person. Yet, license plate information is also deemed "public record."

Finally, there is the issue of threats. In private industry it is said that security can normally be emphasized in one of the "Three Bedrock Principles", *confidentiality*, *integrity*, and *availability*. As I learned at SANS, depending on their type of business, a company will probably place emphasis on one of these three principles while still addressing the other two. For example, a software development company, such as Microsoft, may focus on the principle of integrity. Microsoft's programming code for building applications and operating systems must have the trust of its customers. If Microsoft lost the trust of their customers because they felt the company's code was unstable or was dysfunctional; Microsoft would have a very difficult time making revenue. This is why it was such a big story when it was claimed that Microsoft was broken into. Not only was *integrity* threatened, when it appeared that someone had possibly accessed their source code, but *confidentiality* was violated since they saw their protected source code[1]. However, since local government organizations can be very diverse in nature, it is possible that all principles are important. Therefore, each component of the "Bedrock Principles" must be equally considered. In our local government, each elected official has certain primary responsibilities. For example, one official may be responsible for financials and court records. In this case the official may need to be concerned with all three bedrock principles, *integrity* for the accuracy of financial information, *availability* for the purpose of reliably producing court records as legal documents, and *confidentiality* for keeping "sealed records" expunged. Hopefully, you now can appreciate the challenges local governments face in developing Information Security. Next we will be discussing the evolution of security for our particular government organization.

[1] Lynch, Ian and Craig, Andrew. "Hackers saw Microsoft source code." 30 Oct 2000. URL: http://www.vnunet.com/News/1113113 (21 June 2001).

The Beginning

I should tell you a little about myself; and how I started with Information Security. I used to work for a small application software developer that had only 25 employees. Its customer base was over 5000 and the firm attempted to be as diverse and flexible as it could for their customers. We supported many different operating systems. We also had a multiple-protocol network. This firm was one of the "unknown pioneers" in e-commerce back in 1994. They had developed a completely integrated e-commerce solution in regards to sales on the Net that included on-line credit card processing and sales order processing (which integrated through to their inventory control and accounts receivable). Also they had also developed a purchase order system that also had the same type of integration as the sales order components (integrated through their inventory and accounts payable). For this company, firewalls were access lists (packet filters) configured on a router. More emphasis were placed on creating bastion (hardened) host web servers, and using these bastion hosts as intermediaries for live client transaction processing to the protected servers behind the packet filtering routers. Needless to say, I gained a lot of "hands-on" experience in operating systems, networks, applications and secure frameworks to support their E-commerce solution.

I was network technician, when I first began working for a local government organization and its citizens. I consider the organization medium sized, as far as local government organizations go, with approximately 3000 employees. At that time, the local government had an Internet connection (56k frame relay) and a basic web site. The local organization had a router that did basic packet filtering. This was considered their firewall. They also had an E-mail system. Their RAS system used simple password authentication in which all users shared the same password (which never changed). The RAS system allowed employees to access internal hosts and the Internet at 28k. E-commerce, E-business, or E-government were nonexistent in this organization at this time. Finally, Information Security wasn't really a major consideration at that time.

The prevailing attitude in regards to Information Security was that much of the data they possessed was "public record". Therefore, if someone were to break in, they would be accessing information they had a right to access anyway. Further, it was just a local government (a grain of sand) why would we get attacked? This attitude has slowly changed with events I will discuss later and given the number of local organizations, even smaller than ours, which have fallen victim to attacks. Although my supervisor's boss had been "preaching" about the need for Information Security, it wasn't taken that seriously. However, the manager's agenda did get us a firewall and a RAS device that supported strong authentication, which I will also discuss a little later.

Through the support of IT management, I began a series of security presentations. Much of the content of the security presentations came from knowledge I acquired from SANS conferences. This included explaining potential *threats* and *vulnerabilities* that could lead to *compromise*. Little did I know, until this SANS Security Essentials course, I was talking about the "Threat Model". Besides breaking into systems, there were other threats

to consider that could affect availability, data integrity, and confidentiality. Further, I gave examples of these threats and related them to the business processes of the local government. This included items such as traffic management (availability, integrity), "911"(availability), manipulation of tax rolls (integrity), threats to the criminal justice systems (integrity, availability, confidentiality), and gaining access to information which was supposed protected as private (confidentiality). In giving my presentations, I attempted to individually address a potential threat for each of the elected officials in their primary area of responsibility. Further, I explained a generic process that someone could be taken in trying to break into our organization and displayed screen shots of the tools that are freely available to accomplish it. Also, the presentations included a security plan. The security plan was actually a strategic plan that entailed the creation of an Information Security Panel (security policy making body), the development of an organizational-wide Security Policy, an enhanced security infrastructure, and an IRT (incident response team).

These presentations were given to a board, which is the governing body for IT development across this local governmental organization. The board is comprised mainly of elected officials. The same presentations were also given to a technology advisory committees that supported "the board".

The months following these presentations were the much-nationally-publicized denial of service attacks[2]. The timing of lobbying for the development of Information Security could not have been better. Not only did these presentations get their attention, but the publicized attacks gave the presentations, and Information Security in general, credibility within our government organization.

The remainder of the paper will summarize the process we went through in selecting the various security components, the issues we considered during our reviews, and what technology we ended up with.

The Firewall

When we went through the selection process of purchasing a firewall, we weighed many factors such as types of firewalls, throughput; and what other organizations, similar to ourselves, were using. Simply put, the different classifications of firewalls are based on how each type of firewall scrutinizes the different layers of the OSI model (the foundation for data communications).

We reviewed packet filtering firewalls, stateful inspection firewalls, and application proxy firewalls. First, a packet filtering firewall is known to inspect the IP and TCP (or UDP) header information. Essentially packet filtering firewall look at IP addresses and port numbers to determine whether certain traffic is permitted through it. It typically only looks at each individual packet without consideration of the preceding or succeeding

[2] Hopper, Ian D. "FBI investigation swamped with tips, continue to seek Midwest 'Coolio'." 16 Feb 2000. URL: http://europe.cnn.com/2000/TECH/computing/02/16/dos.attacks.coolio/index.html (21 June 2001).

packets and their "state". In regards to how much traffic can pass through it, since it scrutinizes traffic the least (packet headers only), it can handle the most volume of the three. Unfortunately, since a packet firewall is least critical in analyzing packet, it is considered the least secure. The next type of firewall we reviewed was "stateful inspection" firewalls. Stateful inspection firewalls go one-step further in the OSI model by analyzing not only header information but also scrutinizes the state of the packet streams. For example, with TCP packets, a stateful inspection firewall will scrutinize sequence numbers as well as TCP flags such as ACK, SYN and FIN. This type of firewall has a little more overhead than a packet filtering firewall because it goes a little further in ensuring the correct traffic is passing through it. However, it is considered more secure that a packet filtering firewall. The last type of firewall we reviewed was called an application proxy firewall. An application proxy firewall not only scrutinizes further than a stateful inspection firewall, but also is capable of hiding protected ip addresses behind it. The application proxy firewall communicates to a user behind the firewall as though it's a server. It takes the request of the user then acts as a client to the remote host the user was originally trying to connect to. This type of firewall operates at the application layer of the OSI model while the other two firewalls operate at the network and transport layer. An application proxy firewall "speaks" the various protocols such as http, telnet, ftp, etc. Therefore, not only can it scrutinize the packets, but also limits how protocols can communicate to it. For example, if someone trying to ftp, and is permitted by the application proxy to do so, it will only accept ftp commands because that is all the "ftp proxy" understands. Therefore, someone cannot go outside the boundaries of what the application proxies are designed for. This makes this type of firewall more secure as far as scrutinizing traffic, however it consumes has the most overhead and takes the most resources of the three firewalls we reviewed. Finally, it is also the most limited in the amount of volume that it can handle because no traffic passes through it. It is all processed by the proxies which act as intermediaries.[3]

At the time we had upgraded our Internet connection from a 56k frame relay circuit to 256k frame relay circuit. Regarding selecting a firewall, we also considered the future needs of bandwidth for our Internet connection. We anticipated that the bandwidth would not significantly increase anytime soon. Therefore, throughput was not a major factor in our consideration of a firewall. We also thought about the flexibility of hardware-based and software-based firewalls. Further, it was found that many similar organizations were using an application proxy type of firewall. Finally, we decided to purchase a software-based application proxy firewall. It gave us the flexibility and robustness we were looking for; and our bandwidth requirements were not going to negatively impact the firewall or ourselves.

RAS

We realized that the current RAS (Remote Access Server) we had, with the shared single password was inadequate. When considering another RAS solution, we felt that authentication was the key. Our goal was to implement strong, two factor authentication.

[3] Steinke, Steve. "Firewalls." Network Magazine. 12 Jun 2000. URL: http://www.networkmagazine.com/article/NMG20000613S0010 (21 June 2001)

Two factor authentication being "something you know" and "something you possess". This type of authentication is similar to ATM machines[4]. Although hard tokens are nice and could be considered a stronger form of authentication, they were considered quite a burden and more tedious to the user community. The soft tokens, that came with the RAS, used an automated "challenge response" type of authentication and were installed on the user's PC's. The user would have to enter a password, and the soft token would authenticate using challenge response. Finally, the positioning of the RAS device is such that all traffic authenticated to the RAS has to pass through our firewall.

The system worked out well except for the fact that the company has since been bought out and the "new owner" dropped this device from its product line. A word to the wise, ensure the company you are going to buy from (and its products) are going to be around for a while. A lot of time, energy, and expense can be wasted. We since have chosen a more mainstream system with similar qualities as the original RAS had, but with more flexibility and more authentication schemes.

Intrusion Detection

One of the biggest "sells" when we were trying to implement intrusion detection was the fact that we didn't know whether or not we were getting broken into. We expressed this concern in the presentations we gave. Sure, there were firewall logs, but intrusion detection had the capability of knowing who was knocking at the doors, identifying what types of attacks we being attempted, and provided reporting to management (pretty graphs, etc). With reporting, we could actually strengthen our efforts in developing Information Security.

When we considered intrusion detection, we wanted something that offered both host-based intrusion detection and network-based intrusion detection. Further, we wanted both IDS's to report to a single console. That way, ideally, we would track where the intruder was going. Also, we wanted a system that offered the ability to create custom signatures, be able to tune out false positives, and one that contained a lot of already-made signatures.

Many of the network-based intrusion detection systems use "sensors". They are similar to packet sniffers except that they compare the network traffic to a database of signatures. Then, if the traffic matches a particular signature it sets off an alarm to the console. Host based intrusion detection is just that. It monitors the host it resides on as to whether it senses break-in attempts. Some systems also include the ability to automate responses. This could include sending reset packets to the "assumed attacker", or automatic configuration of a firewall. We decided that it would be a good idea to have more than one "sensor". One would be placed outside of the firewall, in the DMZ, and one behind the firewall, on our protected network. That way, we could detect the attempts outside of the firewall and whether a particular attack managed to get through the firewall.

[4] Duksta, John. "We know who you are." Network World. 8 Aug 1998. URL: http://www.nwfusion.com/reprints/0824review.html (21 June 2001).

E-mail Virus Gateways

One of the threats we felt we needed to address was e-mail attachments. Even though we had virus protection on desktops and servers, we knew the users' tendencies of just opening an attachment; instead of saving the attachment to disk and having the already-installed virus scanner scan the saved files. Therefore, we searched for a virus scanner that could scan virus attachments in E-mail. One of the features we were also hoping to include was the ability to scan for harmful Java and Active X code. Also, we wanted to ensure that the product we selected included the capability of automatically downloading pattern files updates. Then, we would always have the current pattern files provided by the vender. Also, we wanted a product that would allow us to specify what attachment file names we wanted to block. That way, if the vendor did not provide a pattern file in a timely manner, and if we became aware of a virus threat, we would still be able to do something about it by blocking E-mail with that attachment name, subject line, or some other known pattern of the virus that was "in the wild". Another feature we were looking for was the ability for the virus scanner to E-mail an alert to the recipient and the sender when a virus was detected. This not only would warn the sender that they something that contained a virus, but also would make the users feel they had been protected. We managed to locate such a product, however, due to budget constraints and concerns over latency, we could not opt for the component that would scan for harmful Java and Active X code. We hope, in the very near future, to be able to add this component.

A product was selected just before the "Love Bug" virus hit. Unfortunately, we chose to accept the defaults of how frequently the pattern files were updated (weekly). We had to contain and clean our several mail systems. Since that time, we automatically update our pattern files on a daily basis, whether or not the vendor supplies one. We also keep an eye out for virus alerts. The E-mail virus gateway has been a very good investment and has kept us out of harms way on many occasions, especially with all of the variants that came after the "Love Bug" virus.

Security Policy

As part of the security presentations, I emphasized the need for an organizational-wide security policy. The organization as whole did not have one. We felt it was imperative that the organization had a baseline for information security. The security policy was expected to create that baseline. It was suggested to have each "agency head" appoint a representative to sit on a security panel. I became the chair of that panel.

The make-up of this panel ranged from those who were technically savvy to those who weren't. However, the non-technical individuals made their own contributions with their experience in the political climate. Further, they gave us insight as to what we took for granted, how a typical user would understand the technical content of a security policy.

One of the valuable references we used in preparing a security policy was a book called "Information Security Policies Made Easy" by Charles Cresson Wood[5]. The book not only outlines the procedures in creating information security polices, but also includes hundreds of such policies. Further, the book (and accompanying CD) it is organized by technology type. Which was very helpful to us in addressing the different aspects of technology. We then used the book to build information security policies.

Our Security Panel began listing our security concerns to ensure we included them in the policy. One of the predicaments was that there were a lot of monies spent on individual agency's IT. Each agency guarded what they had implemented and didn't want to "give up" something because they were out of compliance. Therefore, in order to make the security policy more palatable, it is initially designed to be a little lax. I realize, according to SANS this is not the correct approach, and I would tend to agree. However, a security policy is a "living document", and during this process we have been creating a list of future considerations for the policy. The idea is to get the policy ratified with minimal friction, and then begin tightening the policy. This will give agencies the chance to conform to the new additions to the policy, instead of everyone considered being out of compliance. With everyone "out of compliance" we would endure more friction in trying to get the initial document passed. We also asked agencies to present their own security policies, for those had taken the time and expense to develop their own policies. That way, "good components" were selected from those policies and included in the organizational policy. This also helped in the area of "buy in".

As this paper is written, we are almost through our first edits of the rough draft. The plan is to edit the document a second time. Once the policy has been passed by a vote within the panel, the next phase is to have each of the appointees "sell" the policy to their "agency head". The appointees will have already prepped their respective board member prior to ratification of the policy by "the Board". It was felt that the policy needed unified support in order to "give it legs". Addressing each member's concerns, aids in having that member support the documents as a whole. The process of creating security policy in local government organizations can be very painstaking, but in the long run, we know it will be very worthwhile.

Best Practices

These documents are currently "work in process". They are designed to supplement the security policy. The best practices documents are the more technical "how to do's" of information security. They are geared toward administrators and are step-by-step procedures to harden systems, networks, and applications. The Security Panel will review the best practices documents to ensure compatibility with the security policy.

Security Awareness Presentations

[5] Wood, Cresson Charles. Information Security Policies Made Easy, Version 7. San Diego: Trade Service Publications, 1999.

Along with the IT security components and policies, comes user education. We have just begun to give security awareness programs to our user community. The goals of these security awareness programs are to combat "social engineering", secure the desktop, improve the password strength of user accounts, reduce virus infection, and emphasize the need to have limited access. Our future plans for these security awareness programs are to include them as part of new employee training and to have an interactive presentation on the Intranet.

Conclusions

Naturally, the security policy should have been the first component implemented. A security policy is the foundation of information security. In our case, we had to play with the cards we were dealt.

The security components that were implemented were in the order of perceived threats. We continued to analyze what security components we had implemented versus what threats we still felt we had. We attempt to prioritize from greater threats to minor threats. Then we addressed these threats by order of priority.

Finally, an important issue in the development of Information Security is credibility. We tried not to overreact when it came to discussing information security with the organization. Further, we ensured that adequate technical justification was made for every security component implemented. We didn't want to be perceived as "someone crying wolf" and have the organization turn a deaf ear. Even though funds are always tight, somehow we manage to be able to obtain what we need, as long as we have properly justified it. We continue to maintain this type of demeanor and it has been paying off.

Citation of Sources:

[1] Lynch, Ian and Craig, Andrew. "Hackers saw Microsoft source code." 30 Oct 2000. URL:
http://www.vnunet.com/News/1113113 (21 June 2001).

[2] Hopper, Ian D. "FBI investigation swamped with tips, continue to seek Midwest 'Coolio'." 16 Feb 2000. URL:
http://europe.cnn.com/2000/TECH/computing/02/16/dos.attacks.coolio/index.html (21 June 2001).

[3] Steinke, Steve. "Firewalls." Network Magazine. 12 Jun 2000. URL:
http://www.networkmagazine.com/article/NMG20000613S0010 (21 June 2001)

[4] Duksta, John. "We know who you are." Network World. 8 Aug 1998. URL:
http://www.nwfusion.com/reprints/0824review.html (21 June 2001).

[5] Wood, Cresson Charles. Information Security Policies Made Easy, Version 7. PentaSafe Security Technologies, 2000.

Last Updated: July 23rd, 2017

Upcoming Training

SANS Prague 2017	Prague, Czech Republic	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS Boston 2017	Boston, MA	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS Salt Lake City 2017	Salt Lake City, UT	Aug 14, 2017 - Aug 19, 2017	Live Event
Community SANS Omaha SEC401*	Omaha, NE	Aug 14, 2017 - Aug 19, 2017	Community SANS
SANS New York City 2017	New York City, NY	Aug 14, 2017 - Aug 19, 2017	Live Event
SANS Chicago 2017	Chicago, IL	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS Virginia Beach 2017	Virginia Beach, VA	Aug 21, 2017 - Sep 01, 2017	Live Event
SANS Adelaide 2017	Adelaide, Australia	Aug 21, 2017 - Aug 26, 2017	Live Event
Community SANS Trenton SEC401	Trenton, NJ	Aug 21, 2017 - Aug 26, 2017	Community SANS
Virginia Beach 2017 - SEC401: Security Essentials Bootcamp Style	Virginia Beach, VA	Aug 21, 2017 - Aug 26, 2017	vLive
Community SANS Pasadena SEC401 @ NASA	Pasadena, CA	Aug 23, 2017 - Aug 30, 2017	Community SANS
Mentor Session - SEC401	Minneapolis, MN	Aug 29, 2017 - Oct 10, 2017	Mentor
SANS San Francisco Fall 2017	San Francisco, CA	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Tampa - Clearwater 2017	Clearwater, FL	Sep 05, 2017 - Sep 10, 2017	Live Event
Mentor Session - SEC401	Edmonton, AB	Sep 06, 2017 - Oct 18, 2017	Mentor
SANS Network Security 2017	Las Vegas, NV	Sep 10, 2017 - Sep 17, 2017	Live Event
Mentor Session - SEC401	Ventura, CA	Sep 11, 2017 - Oct 12, 2017	Mentor
Community SANS Albany SEC401	Albany, NY	Sep 11, 2017 - Sep 16, 2017	Community SANS
Community SANS Dallas SEC401	Dallas, TX	Sep 18, 2017 - Sep 23, 2017	Community SANS
Community SANS Columbia SEC401	Columbia, MD	Sep 18, 2017 - Sep 23, 2017	Community SANS
SANS Copenhagen 2017	Copenhagen, Denmark	Sep 25, 2017 - Sep 30, 2017	Live Event
Community SANS Boise SEC401	Boise, ID	Sep 25, 2017 - Sep 30, 2017	Community SANS
Baltimore Fall 2017 - SEC401: Security Essentials Bootcamp Style	Baltimore, MD	Sep 25, 2017 - Sep 30, 2017	vLive
Community SANS New York SEC401	New York, NY	Sep 25, 2017 - Sep 30, 2017	Community SANS
Rocky Mountain Fall 2017	Denver, CO	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS London September 2017	London, United Kingdom	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Baltimore Fall 2017	Baltimore, MD	Sep 25, 2017 - Sep 30, 2017	Live Event
Community SANS Sacramento SEC401	Sacramento, CA	Oct 02, 2017 - Oct 07, 2017	Community SANS
SANS DFIR Prague 2017	Prague, Czech Republic	Oct 02, 2017 - Oct 08, 2017	Live Event
Community SANS Charleston SEC401	Charleston, SC	Oct 02, 2017 - Oct 07, 2017	Community SANS
Mentor Session - SEC401	Arlington, VA	Oct 04, 2017 - Nov 15, 2017	Mentor