GIAC
CERTIFICATIONS

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permited without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at http://www.giac.org/registration/gsec

Case Study in Information Security

GIAC Security Essentials Certification Version 2.4 Option 2

Nathan Gehman
January 2003

© SANS Institute 2003, Author retains full rights.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

Nathan Gehman
GIAC Security Essentials Certification
Version 2.4 Option 2

Introduction

"Nathan, I have great news!" I turned to see my supervisor coming in the door with a big grin on his face. I returned a blank stare as he continued "We are moving to broadband!" I knew the change was coming, I had been working on configuring a new e-mail server as part of the changeover, but it still came as a surprise that it had all been approved and we could put the wheels in motion. My mind started spinning as all the security implications of broadband came flooding to my mind. What assets do I need to protect? Who am I protecting them from? What is the likelihood of an attack? What is the most cost effective way to go about this?

I knew the project was going to be involved and I was thankful for the resources available to those of us starting out in a brave new world. I was new to information security and was thankful for the opportunity before me. As illustrated so well over and over again in the SANS course, security needs to be a defense in depth, and needs to protect against a number of threats. The primary purpose of information security is to ensure confidentiality, integrity and availability. In this paper, I will deal with the steps I have taken in order to ensure the move to broadband proves to be a secure transition. I will also address the areas we have been lax in the past and propose plans to improve

Before

I began the process by looking at what we currently had in place. Our network consisted of about 50 workstations distributed among several switches. We had two NT servers acting as file and print servers for the network. Each workstation was assigned a private class IP address. Our external gateway was handled by a dial-on-demand modem using Network Address Translation (NAT). The modem was configured to hang-up after a set period of network inactivity. We ran a ccMail post office and ccMail clients for our office e-mail. Users in our office connected to our ccMail PO via the LAN. Our external Gateway PO handled the SPAM and virus removal.

Part of the transition to dialup involved the migration of our e-mail from a ccMail PO to an in-house pop3/smtp server. With ccMail we had been blessed and cursed by a proprietary program. We counted it a blessing that we were immune to the majority of worms, trojans and other uglies on the internet that targeted the Microsoft products. We were cursed because of the loop holes we had to jump through to get a older product to work with the newer Microsoft operating systems. Those loop holes was a large part of the reason we wanted to migrate to a non-proprietary system. I remember as the waves of viruses would wreck havoc on the corporate world thinking to myself, "I am glad we don't have to worry about it". I was well aware that all this was going to change as we migrated to Outlook and a pop3/smtp server.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

We were comfortable with the degree of external risk associated with our dialup connection. Our firewall consisted of network address translation (NAT) masquerading on the boarder router. NAT masquerading works by replacing the TCP/IP stack with its own IP stack. Since it maintains the state information of the packet, the router can effectively route packets to and from the internet while maintaining the privacy of the internal network. This protects the gateway and network by allowing the network side of the connection to use IP addresses reserved as private addresses. The obvious problems with NAT running by itself was the lack of protection of the gateway router itself and the inability to control what information left the network. I found some discussions on other vulnerabilities in NAT but most of the documents I found dealt with vulnerabilities with particular pieces of hardware or software. In all the cases the patches were all ready available or in the process of being fixed. It was a reminder however that information security is not a "set it and forget it" kind of operation. An integral part of any network security involves keeping hardware and/or software patched and up to date.

In the SANS course we learned that risk was equal to the threat times the vulnerability. Although our connection to the web introduced certain threats and vulnerabilities, we felt the slow connection, and NAT, created a level of risk we were willing to accept. We had no other firewall or access control list (ACL) controlling what protocols we allowed in and out of our network. We had no way of logging any web activity or filtering web sites for content. We controlled who had access to the internet by controlling the gateway IP information in the workstations TCP/IP settings. Our acceptable use policy was basically a code of conduct that relied on a relatively slow dialup connection to curtail improper use of network resources.

We ran a Windows NT 4.0 network with the majority of our workstations running Windows 98 and higher. We were in the process of updating the workstations to Windows 2000 Professional. Each user logged into and authenticated against the primary domain controller (PDC). We used login scripts to govern the mapping of network drives based on the authenticated user. Each workstation had antivirus software installed and was updated automatically by the PDC server which checked daily for updates.

Chris Brenton and Cameron Hunt in their book state "Case studies have shown that a vast majority of attacks originate from within an organization. In fact some studies state that as much as 70 percent of all attacks come from someone within an organization" (Brenton, p.6) This was a risk we had not properly taken into account. Our servers were in the IT department but were vulnerable to anyone who was intent on physically accessing them. With only two people in the IT department there were often times that the department was empty and the servers unattended. The risk factor was hopefully curtailed by the relatively low threat factor. Being a relatively small non-profit organization, the quality and integrity of the individuals was assumed quite high. None the less it was only a matter of time before the trust factor became an exploited vulnerability.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

As I looked at the move to broadband I realized most of the things we had taken for granted were no longer true. It is my desire to address the security issues mentioned above and to provide a plan of action to shore up our defenses. Since we were moving to broadband and static IP's, I realized we are more prone to attacks and attention especially with an e-mail server sitting in the DMZ. However I not only wanted to address the external threats, but also the internal security risks as well. The lack of clear policies needed to be dealt with and the process maintaining and enforcing the polices once they were in place.

I divided the projects into two categories to help visualize the task at hand. The projects were as follows: The transition from dialup to broadband, and the transition from ccMail to an in house pop3/smtp server. I will touch on the revision of the policies governing the acceptable use of network resources in the AFTER part of this paper.

During

From Dial-up to Broadband

Before deciding on a firewall to sit between the Local Area Network (LAN) and the Wide Area Network (WAN), I needed to form the rules that would govern the flow of information in and out of our network,

> Before you can choose the type or brand of firewall to purchase, you have to ask yourself a very simple question: What are (or should be) the rules that deal with the flow of data traffic in and out of your network? The answers to this question will form your access control policy. (Brenton, p.144)

Up to this point we had no rules. We basically allowed everything in and out. Okay, maybe not exactly, we did have NAT that provided a certain degree of protection. However we were certainly wide open when it came to controlling what traffic went out. I broke the picture into sections or zones. I wanted to be able to control the traffic independently from LAN to WAN, LAN to e-mail, e-mail to WAN, WAN to e-mail.

I decided the most practical way to solve this was to use a DMZ or demilitarized zone. In its simplest form a DMZ uses a subnet to isolate a "public" part of the network from the "Private" or internal part of the netwrok. By putting public servers (such as FTP or Mail) on the DMZ, it allows the admin. to allow traffic needed for particular services into the DMZ without compromising the security of the internal network.

There are two approaches to creating a firewall. Deny everything and explicitly allow the traffic you need through or allow everything and deny selected traffic. I took Robert Ziegler's approach and set everything to "deny".

> The deny-everything policy is the recommended approach. This approach makes it easier to setup a secure firewall, but each service and related protocol transaction you want must be enabled explicitly. (Ziegler p.23)

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

To begin my rule list I created the following defaults:

Default - Deny – LAN:* - Deny all traffic from the LAN to anywhere
Default - Deny – WAN:* - Deny all traffic from the WAN to anywhere
Default - Deny – DMZ:* - Deny all traffic from the DMZ to anywhere

This sets the default policy of the rule list to deny. Often people make the assumption that this is the default configuration of their "firewall". THIS IS NOT ALWAYS THE CASE. By setting everything to deny, only traffic explicitly approved will be allowed in and out of the network. This created problems when I first went live since I was not aware of all the ports I needed open, but by watching the logs I was able to add the rules necessary for our small network to operate "securely" and efficiently.

The next step in creating my rule list involved allowing the LAN access to the internet. Certain things need to be able to happen for a user to "surf" the web. First of all, the machine on the LAN needs to know the IP address of the machine on the web hosting the website. For instance, when a user types www.google.com in their browser, the computer performs a DNS lookup to find the IP address associated with that domain name. This service uses port 53 and the UDP protocol.

Along with port 53(UDP), we need to open ports 80(TCP) and 443(TCP) on the firewall from the LAN to the WAN. Normal web traffic utilizes port 80 (TCP). Https protocol, used on secure web sites, makes use of port 443(TCP). Without opening these ports from the LAN to the WAN, users will be unable to surf the web. If you desire more control over your users, and have manually assigned them IP addresses you can allow access based on individual IP's or a range of IP addresses.

So to my original rule list I added the following:

Allow – UDP:53 – LAN:WAN - Allow DNS lookups. (Should specify your DNS servers here rather then allowing outbound UDP to any internet server)
Allow – TCP:80 – LAN:WAN - Allow web traffic from the LAN to the Internet
Allow – TCP:443 – LAN:WAN - Allow https traffic from the LAN to the internet
Allow – ICMP:3 – WAN:LAN - Allow "destination unreachable" from the internet to the LAN

In my rule list you will notice ICMP and the number 3. ICMP protocol is handled differently then UDP and TCP traffic. Unlike UDP and TCP which can be bound to a port, ICMP is referred to by type. For instance type ICMP type 8 is an echo request or "ping" and type 0 the reply. ICMP type 3 refers to destination unreachable. When an individual on the LAN tries to access a website that is unavailable, it is preferred to allow the message back into the LAN telling the machine that the website is unavailable. If you do not allow ICMP back in, the user will have to wait for his/her browser to timeout. There is some controversy and misunderstanding surrounding ICMP traffic.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

The controversy surrounding ICMP and firewalls stems from a lack of understanding of the ICMP protocol and a lack of understanding of TCP/IP in general. ICMP is vital for IP networks including the internet to work properly. ICMP however has also been used to scan networks, launch denial of service attacks and even tunnel into remote networks. Because of this Ofir Arkin in his paper "ICMP Usage in Scanning the complete know-how " suggests only permitting one of the 52 types through the firewall to the LAN and no ICMP traffic out of the LAN. "If you want to maintain strong ICMP filtering rules with your Firewall/Filtering device I suggest you block all incoming ICMP traffic except Type 3 Code 4" (Arkin, p.187)

One of the most common types of ICMP is type 8 otherwise known as Ping. Ping is a tool that uses the IP stack to send a request to a remote IP stack. If the remote IP stack responds to the request, the remote stack is up and running and the connection between you and the remote stack is working. Ping is sometimes used in this way to scan networks and based on the reply, create a list of active systems in that netwrok. Ping can also be used in a denial of service (DOS) attack. A DOS attack works on the assumption that if you send enough information to a system it will run out of processor and will be unable to respond to legitimate work load. It can also be used to tie up a connection to the web. In order for this to be effective the attacker would have to have more bandwidth then the victim or via a distributed denial of service (DDOS). In a DDOS the attacker makes use of several location to launch the attack. In this way the attacker gains more bandwidth then the individual under attack. More and more people are choosing to block incoming type 8. Others have chosen a single machine on their network that responds to all ping traffic to enable them to perform automated checks on their network availability.

CMP packets are commonly used in other forms of DOS. The common attacks involve using ICMP redirects to adjust the machines router table. Another example of a DOS involves the source quench (type 4) packet. In this DOS attack the attacker continuously commands the server to slow down until it almost stops. DOS attacks are an example of an attack on information availability.

ICMP packets can also be used to identify the operating system of a remote machine. ICMP packets not only have a type but, as mentioned in Ofir's paper, they also have subtypes called Codes. Microsoft and Unix operating systems vary in how they handle these codes. For example an echo request (type 8) with a code value set, sent to a Microsoft box will clear the code in the response packet where as a Unix box leaves the code intact. By examining the returned packet the hacker has an idea of the operating system and can send further probes to further classify the OS.

I decided to go with Ofir Arkin's advice in order to allow web browsers to function correctly on our LAN and chose to allow destination unreachable ICMP (type 3) packets into the network. This should speed up the web surfing since users will not have to wait for the browser to time out before they receive a failed connection. I also followed Ofir's advice and did NOT ALLOW destination unreachables out of the network since this can be used by an attacker to reverse map your network.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

Now that I had my basic rule set I went back to the choice of firewalls. I noticed I had 2 basic choices of firewall products. These choices were server-based, or appliance-based. The advantages of a server-based firewall were the advanced configuration options and the price, often a cheaper solution than the appliance-based solution. The drawbacks being the strength of the firewall depended on the proper configuration of the OS and its current security level. Appliance-based firewalls tend to be more expensive and usually more robust. The disadvantage is you are relying on the vender to keep the firmware up to date. It is strongly suggested to go with a reliable vender that is constantly working on improving and patching their product. If there is a known vulnerability, but the vender isn't fixing it you are up the creek. Both solutions shared the common strengths and weaknesses of proper or improper configuration of the rules controlling the network traffic.

I took a trip to my local bookstore and spent some time looking for books on network security and/or firewalls. I settled on 2 books. One book focused on Unix firewalls and the other on information security in general. After installing a flavor of Unix and playing with it for a while I decided I would be fooling no one if I pretended to have a good enough grasp of the OS to configure it securely. I was also uncomfortable with my understanding of Microsoft's OS. Therefore I abandoned the "software" firewall route for the appliance based firewall from a well known company. I had a fairly limited budget of $2000. I did a couple of searches on the web and was impressed with the reports and products from SonicWall. I then did searches on the web to find any complaints or reviews of their products and found they were tested and approved by ICSA labs.

"The Candidate Firewall Product Passed the Firewall Product Certification Criteria version 3.0a"

http://www.icsalabs.com/html/communities/firewalls/certification/rxvendors/sonicwallpro/labreport_cid289.shtml

I then called one of our suppliers who included me in on a call to his contact at SonicWall. I appreciated the time they took to talk me through my requirements and to explain the product(s) that would best fit my needs and situation. I also read through the product summary report at:

http://www.icsalabs.com/html/communities/firewalls/certification/rxvendors/sonicwallpro/pfd.pdf

I was impressed by the features and the scalability of the product. I decided that this was the kind of product we were interested in and have been pleased with the results since setting it up.

From ccMail to Outlook

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

The second phase of the project involved configuring an e-mail server to sit on the DMZ and act as our e-mail server. There were several things that needed to be done in-order to secure the server. Part of the decision process involved the decision of which OS to run on the server and the software I chose to handle the mail. That whole process is a paper in itself. In an attempt to stick with my topic and to refrain from starting any flame wars I will remain neutral. For informational purposes however, I started with a minimal install of RedHat Linux 7.3 I performed a "custom" install and selected I wanted to chose what services I wanted installed. I wanted to make sure I only had the services I needed running.

There are plenty of helpful documents on the web that can aid in the setting up and securing a Linux box. I found a helpful guide by Werner Puschitz at http://www.puschitz.com/Security.shtml. I suggest reading through a couple of these guides as individuals vary on some of their approaches to security and by reading a few it may help you to broaden or narrow your approach to security.

I further tightened the server by customizing the built in firewall to only allow traffic in and out of the specific ports I need. I also denied access from other machines on the DMZ. I then stopped all the services that I did not need running. We are going to need access to the following services on our server and require rules to address each one.

TCP - 25 – SMTP
TCP - 465 – SMTPS
TCP - 636 – LDAPS
TCP - 995 – POP3S

All of our users are required to make a encrypted connection using SSL We have required our users to use port 465 for SMTP instead of allowing them to use SSL on port 25. I also need to open port 25 so we can receive mail from other servers not taking advantage of SSL. Opening the ports on the server firewall was not enough since I had to tell the boarder firewall as well that I wanted to allow this traffic in from the LAN and from the WAN to the DMZ.

As part of our migration we were losing our external gateway that handled all the spam and virus removal for our e-mail users. I needed to find a product that worked well with our server and our system so I decided on RAV antivirus. I found the software relatively user friendly on the administrative end but more importantly the software has proved reliable and effective. The software checks for updates every half hour. I have also been able to configure RAV using reg expressions to act as a basic spam filter. It was important to me to use a product from a different manufacturer then we use on the LAN. This goes back to the defense in depth. If the e-mail server missed the virus hopefully the desktop software will catch it.

To finish off the process I installed and configured tripwire. This process was a bit more complicated then I expected as by default, tripwire assumed I had performed a

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

full or “everything” install instead of a custom “minimal” install. This meant quite a bit of customization of the configuration files. Rather then try to figure out manually which ones I needed I had tripwire run a test and let me know which ones it thought I should have that I didn’t have. I then went through the configuration files and commented out the ones it could not find.

On the client side we had a couple of requirements as well. Our e-mail server was set to require authentication. This was to discourage the use of our server as an open relay and to attach activity in the log to a specific account. We also required our users to make connections over SSL. In this way we could provide confidentiality in the transfer of e-mail via the web to and from people on the system. Even the users that accessed the LDAP were required to access the information over SSL with the intent to protect the confidentiality of the information being exchanged.

Along the same confidentiality lines, our users in the office were required to password protect their .pst files. This provides an additional layer of protection. This is one area we still need to improve our policies. Simple steps like requiring them to password protect and lock the screen when they are not at their desk, will go a long way in securing our internal network.

The other part of the process involved the configuration of the firewall to allow the necessary traffic in and out of my DMZ. I added the following to my rule list:

Allow – TCP:443 – WAN:DMZ – Allows HTTPS from internet to server (web mail)
Allow – TCP:636 – WAN:DMZ – Allows LDAPS from internet to server
Allow – TCP:465 – WAN:DMZ - Allows SMTPS from internet to server
Allow – ICMP:3 – WAN:DMZ – Allows destination unreachable messages from internet to the server.
Allow – UDP:53 – LAN:WAN - Allow DNS lookups. (Should specify your DNS servers here rather then allowing outbound UDP to any internet server)
Allow – TCP:995 – LAN:DMZ – Allows users on the LAN to POP mail with SSL
Allow – TCP:465 – LAN:DMZ – Allows users on the LAN to send mail with SSL
Allow – TCP:636 – LAN:DMZ – Allows users on the LAN to search the directory with SSL

After

As mentioned earlier in the paper keeping things up to date is a vital part of any information security program. There are a number of steps I have taken to keep informed on what is going on and to keep our systems up to date. The first step is subscribing to security and bug lists. Three lists that come to mind

1) SANS Newsletters and Digests – available at:
   http://www.sans.org/newsletters/
2) Bugtraq – A number of lists customizable to your needs at:
   http://www.securityfocus.com/

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

3) CERT - the first computer security incident response team at:
http://www.cert.org/contact_cert/certmaillist.html

I also subscribe to the mailing lists provided by my firewall manufacturer and software venders in order to keep up to date on firmware, exploits, and patches as they become available.

There are some automated tools that can be used to aid in the process of keeping products up to date. RedHat offers a tool called up2date. I have turned off the service which checks and downloads and updates packages automatically, but I do use up2date to simplify the download and verification of packages while installing them manually.

In order to update packages and firmware you may need to open outbound ports on the firewall. You can make the decision based on your risk tolerance verses ease of use as to if the ports are opened on "as needed" basis or if you want to leave them open.

We have migrated from basic NAT masquerading to an appliance based Stateful Packet Inspection firewall. We now combine ACL, NAT, and other proprietary algorithms in our firewall to protect against a number of common attacks including Syn flood, Ping of death, IP Spoofing, Land attack, Smurf amplification, sequence number prediction. We have set the default policy of the firewall to deny and log all our network traffic infringements. We also maintain a database of network activity from which we can better monitor and supervise the information coming in and going out of our network.

The weakness in our system is the inability to decide the legitimacy of traffic. We assume that because the traffic is using expected protocols on expected ports it is okay. The problem is Trojans and other rouge software could exploit our trust factor and funnel valuable information outside the network. We are only vulnerable to this if the connection is initiated by the internal machine.

We are in the process of moving our servers and switches to a secure and climate controlled room to strengthen the internal defenses. We are in the process of planning and implementing a disaster/recovery plan including alternative office space, internet access, servers, and backups at a remote location. This should aid in our pursuit of availability and integrity.

I am in the process of rewriting some policies and procedures to give management and the end user a clearer understanding of what our acceptable use policy is. This will give management the foothold to take action when needed to enforce the policy. The security policies themselves were the biggest challenge. That came as a surprise to me. I expected the policy part to be as simple as putting the unspoken polices down on paper however the more I worked on them the more frustrated I got. One of the biggest frustrations was the apparent lack of standards. The more polices I downloaded and

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

read through the more apparent it became that everyone does it their own way. Some organizations had one policy that covered the various policies. Others had separate policies for each activity

When dealing with information security we are dealing with three issues. They are: Confidentiality, Integrity, and Availability. I believe the steps I have taken and the products I have recommended have made a significant improvement to the overall security of our network. I have created the groundwork for a "Defense in Depth" and trust that the things I am learning from SANS and other resources will enable me to gain a deeper understanding of the rules of the road and enable me to effectively manage the security of our network.

© SANS Institute 2003, Author retains full rights.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

Refrences

1. Arkin, Ofir. “ICMP Usage in Scanning The Complete Know-How” Version 3.0. June 2001. http://www.syssecurity.com/archive/papers/ICMP_Scanning_v3.0.pdf (January 2003)

2. Brenton, Chris & Hunt, Cameron. Active Defense A Comprehensive Guide to Network Security. San Francisco: Sybex, 2001. p6, p144

3. Puschitz, Werner. “Securing Red Hat Linux 7.2, 7.3, and RedHat 2.1 Advanced Server” Dcember 18, 2002. http://www.puschitz.com/Security.shtml (December 2002)

4. Ziegler, Robert L., Linux Firewalls. Indianapolis: New Riders 2000. p23

5. Firewall Lab Report – SonicWALL Inc, SonicWALL Pro 200 http://www.icsalabs.com/html/communities/firewalls/certification/rxvendors/sonicwallpro/labreport_cid289.shtml (December 2002)

6. Firewall Product Functional Summary - ICSA Labs http://www.icsalabs.com/html/communities/firewalls/certification/rxvendors/sonicwallpro/pfd.pdf (January 2003)

© SANS Institute 2003, Author retains full rights.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.