



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 permitted 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>

The Tragedies of Misconfigured Networks and Post-Prevention.

James McMillan
Global Information Assurance Certification
GIAC Security Essentials Certification (GSEC)
Administrivia Version 2.3
Practical Paper Requirements - Version 1.4b
Option 2 - Case Study in Information Security
Friday, November 15, 2002

Introduction

This paper will introduce the dangers and possible tragedies of an enterprise-sized business, which was constructed in a hurry without proper pre-planning. Involved will be where to start, and the proper procedures to fix the issues once trouble has already started. Included in my presentation will be the removal of unneeded Static NAT, closing open relays on mail servers, applying proper Access control lists, initiating an Anti-virus server, and rolling out Anti-virus clients to provide consistent, enterprise-wide antiviral solutions.

The Beginning

As I worked for a company known to the readers of this article as “Company ‘X’”, over the course of time I noticed a dramatic decrease of work production due to IT problems. Users were being dropped from their connections to the server, major bandwidth problems occurred, at times users would never even establish a connection to the server, and it seemed like an endless battle against Nimda, SirCam, CodeRed and CodeRed II viruses on practically all client machines. At the time I was the company’s Sr. DBA and was not actually involved in the upkeep of the current network, I noticed there were major problems, but it seemed there was never any time to actually dig in and find out what the issues were due to the necessity to quickly remedy the client machines. To help offset the loss in production, Human Resources and the company’s Executive Board decided to hire more people and expand even further. What they failed to realize was that a decision like this would only make matters worse. If there were already bandwidth problems, adding more nodes to the current network would only stretch our available pipeline thinner, and compound the problem.

A brief history and topology of Company ‘X’

Refer to Diagram 1.1

- Currently 3 Separate sites connected via Point to Point T1s.
- Each site also had its own separate T1 connection to the internet.
- The main site’s (the middle leg of the 3) router was a Cisco 2600.
- The other two site’s routers were Intel ER9100 Express Router 9100.
- Site 1 contains both primary and external DNS servers, and the main Mail server.
- Site 2 contains the remote desktop server in which 75% of the company is based off; all web servers, 2 internal DNS machines, and the main fileserver.
- Site 3 is all client machines. No important servers here.

- The only firewall was a SonicWall located at Site 2.
- Company 'X' was all started with Site 1 and expanded very rapidly.
- The rapid expansion called for rapid roll-out of other services. T1s, routers, servers, etc...
- All roll-outs were never benchmarked nor further evaluated.
- Client machines had no common procedure of installation. Anyone that was available to install an OS would be the one to do so.




Figure 1.1

One could walk around the establishment and notice common problems like bad wiring methods, incorrect positioning of hubs (i.e.; tossed behind cubicles,) and client machines with half their casing removed. These were only the obvious, physical problems.

It seemed that so much time was wasted putting out fires, that there was no time to properly train the current IT staff. Because there was no security policy in place, all client machines were installed differently. Operating Systems varied from one client to the next. Most were Windows 95 clients, but there were some Windows 98, Windows NT Workstation, Windows 2000 Professional, and even a few -Windows NT Servers lurking about. The problem here was that it was impossible to roll-out mass patches and/or fixes to a network environment of this nature. The clients that actually had "Services" running were all running as default installs would intend. Some clients were actually older servers just handed down, like the NT 4 Server installs, which were still running IIS 4. If these machines were not properly patched then they would be DoS attackers as soon as CodeRed came around, running against its own network.

A denial of service vulnerability that could enable an attacker to cause the IIS 4.0 service to fail, if URL redirection has been enabled. The "Code Red" worm generates traffic that can in some cases exploit this vulnerability, with the result that an IIS 4.0 machine that wasn't susceptible to infection via the worm could nevertheless have its service disrupted by the worm.

[<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS01-044.asp>]

Many users never had the proper upgrade for their Web browsers, usually IE 5.0, and was subjective to Nimda attacks. The virus would do scheduled scans on the network causing many connections to have a large lag time or complete DoS. The current IT staff would only patch these systems as the problems occurred. It seemed much had to change here, and as soon as possible.

The Procedure

With a major holiday coming up and myself being caught up on all my database work, production halted for a week's vacation. This was an ideal time to run a few diagnostics. I took it upon myself to require all IT management to be in for proper training and maintenance of client machines. The network admin and I plotted out on how we were to resolve our current issues and plan the training. First we had to review to the IT staff on what patches needed to be installed on what machines. Any client machines that were running NT 4 server had to be wiped clean and a fresh, proper install of Windows 2000. We created a custom Windows 2000 bootable CD that was properly patched with the current Service Pack. It was not very practical to reinstall 70 client machines in 2 workdays.

As the staff went around to all client machines patching, upgrading and reinstalling, I cranked up Nmap on a local Linux machine running a broad TCP scan on all servers, turning off any services that were not being used. For example IIS installs come default with a SMTP service started which is vulnerable to Encapsulated SMTP Address Vulnerability

Portcullis has discovered that the Microsoft SMTP Service available with IIS 4.0 and IIS 5.0 is also vulnerable to the encapsulated SMTP address vulnerability even with anti-relaying features enabled. This vulnerability allows hosts that are not authorized to relay e-mail via the SMTP server to bypass the anti-relay features and send mail to foreign domains

[<http://www.securiteam.com/windowsntfocus/5QP0G0K7PE.html>]

I also decided to add a Netgear 100 base-T hub containing the internet T1, the Point-to-Point T1 and the LAN switch. This would provide me to watch all traffic coming from the outside and/or going outside by adding a Linux box with a few network analysis tools. I installed Ethereal, Nemesis, and Nmap for starters.

In case of a disaster during the day I could start up ethereal and watch the source of the problem. From there I am now *prepared* for an attack. If an attack does occur I can *identify* the problem, as well as the source, which leads to the *containment*. Once the machine has been *contained*, the company could get back to production and could start the cleaning process and *eradication*. Once cleaned, the machine is *recovered* and any data lose is replaced from backup and the machine is placed back in its area. From here we should do any additional research and discuss what we have *learned* and what we can do to prevent this from ever happening again.

I jumped in (at the time) over my head and started looking at the Intel routers. Logged on to the console of the Intel and dug through the menus for anything that didn't sound right. I found the NAT "tables" and noticed it was set for Static network translations translating the entire network one-to-one fashion. Refer to diagram 1.2. After I contemplated and checked the support Intel website, I read that,

When using Static Mapping, addresses are simply converted by translating the network part of the IP address between the internal and external address. The host part of the address remains the same, for example, an internal class B network address 10.10.4.8 (where 10.10 is the network part of the address) could be translated to the external class B network address 177.4.4.8.

[<http://support.intel.com/support/express/routers/9xxx/23797.htm>]

Realizing that this meant for every client machine's internal IP address there was a matching external IP. Continuing to read on through their site I found that Dynamic NAT can translate networks of different sizes. Searching online for a subnet calculator I found [<http://screamer.mobrien.com/net.shtml>] using it to find the subnet that would contain the least amount of hosts possible. Seems I could translate our internal subnet of 255.255.255.0 to the Class C range using a subnet of 255.255.255.252, allowing only two external host addresses to be accessible from the outside. Then deciding to create static NATs for the individual servers that were needed for public access, we ran thru all the FQDN we owned. We searched thru every zone gathering all external addresses and comparing them to the internal DNS names, creating a spreadsheet of the internal to external translations. We then added them to the router before removing the Static table, thus not creating DoS of any service we were running. After verifying that we did not miss any translations we removed the Static translation, wrote the configuration to flash RAM, and restarted it.

Dynamic Mapping can be used to translate between IP networks of different sizes, that is, a large internal network can be translated to smaller external network or vice versa (for example a class B internal network could be translated to class C external network addresses).

When using Dynamic Mapping, more addresses can be available to either the internal or external network. For example, if the internal network is class B and the external network address is class C, the internal network can have up to 65,536 network addresses while the external network address only offers up to 256 addresses. In this case, the entire internal address (network and host part)

must be translated to an assigned external address. External addresses are therefore assigned sequentially as they are required.

[<http://support.intel.com/support/express/routers/9xxx/23797.htm>]




Figure 1.2

Testing from the outside we noticed our translations went fine. The client machines that once have external translations were now inaccessible from the outside world. We ran a port scan against one of the web server's external translated addresses and noticed that there were ports open that should only be accessed from the internal LAN. The server was a Linux machine running Apache 1.3.20 and Sendmail 8.4.11; I successfully connected to the Sendmail daemon using telnet from an outside connection. This was a problem that needed immediate attention.

I connected to the web-based configuration page of the Sonicwall Firewall and looked through the menus; finally I found the Access and Rules menu. Looking at the current configuration, I could see that it was a wonder how the company was never truly attacked. There was only one rule, Allow Any, From Any, To Any, Protocol Any, Port Any. In other words, before the translation change there were approx 60-80 servers/clients standing out with all ports open to the "Big Bad Internet." With complete astonishment, we came up with the best plan for the time.

It has been written over and over, the key to success is simplicity. "The key to a secure firewall is a simple rulebase. The more rules you have, the more likely you or someone else will make a mistake. The fewer rules your rulebase has, the easier it is to understand and maintain. A good rule of thumb is to have no more than 30 rules. With 30 rules, its relatively easy to understand what is going on. Between 30 and 50 rules, things become confusing, the odds grow exponentially that something will be misconfigured. Anything over 50 rules and

you end up fighting a losing battle.” -Lance Spitzner The best way for us to go forward would be to sit down and construct exactly what would be needed to get the business flowing properly without future problems. To do this we picked up the scan for the entire subnet, the list of zones, and NAT tables we created earlier. From this we could construct an in-depth, yet basically simple list of what needed to be accessed and from where. Starting with hand written rules on a sheet of paper and noting the proper order in which rules subsided.

For a good base we created an opposite base rule than the one we had; DENY ALL FROM WAN TO LAN! From here we quickly constructed our way out, starting with allowing anyone on the LAN complete access to the WAN, followed by the highest priority services first. Looking at figure 1.3 you will see the “Add Network Access Rule.”

- Action -> Allow
- Service -> Web (port 80)
- Source -> WAN
- Address Range -> *
- Destination -> LAN
- Address Range -> x.x.x.x (internal IP of the web server)

The screenshot shows a window titled "Add Rule - Microsoft Internet Explorer". Inside is a form titled "Add Network Access Rule". The form contains the following elements:

- Action:** Two radio buttons, "Allow" (selected) and "Deny".
- Service:** A dropdown menu currently showing "Default".
- Source:** Labeled "Ethernet Addr Range Begin" and "Addr Range End". The "Begin" field contains an asterisk (*) and the "End" field is empty.
- Destination:** Labeled "Ethernet Addr Range Begin" and "Addr Range End". The "Begin" field contains an asterisk (*) and the "End" field is empty.
- Apply this rule:** A dropdown menu showing "always", followed by two empty fields, then "to", two more empty fields, and the text "(24-Hour Format)". Below these are two dropdown menus, both showing "Sun", with "to" between them.
- Inactivity Timeout in Minutes:** A text box containing the number "5".
- Allow Fragmented Packets:** A checkbox that is currently unchecked.
- Buttons:** "Update" and "Reset" buttons at the bottom.

Figure 1.3

The configuration was always updated on the fly by clicking Update after completing each new rule. We continued doing this until we were satisfied with our new ACL's.

Finally we took into consideration the current usage in Peer-2-Peer file sharing programs. Although the site was controlled with Microsoft Active Directory and Group policies, it could still be possible for a user to install and use one of the applications. So I decided to add a few extra specific services reflecting to rules in the ACLs directed to P2P Networks, including KaZaa, IRC, eDonkey, iMesh, etc See Figure 1.4. Also included were a lot of the most popular "Instant Messenger" programs. Of course there are privileges around these rules that the users could take in effect like changing the port numbers, but this is unlikely for the normal end-user. Also as Shawn Wood said in his GSEC practical "An external Gnutella client attempting a connection (download) first initiates the connection with a "pull" request to an internal client behind a firewall. Because the client is behind a firewall and the packets are not allowed to pass, the connection is refused. The problem lies in the fact that Gnutella clients are configured to initiate a "push request" if this initial "pull request" has been rejected." For issues like these we will have to depend on the Group Policies.

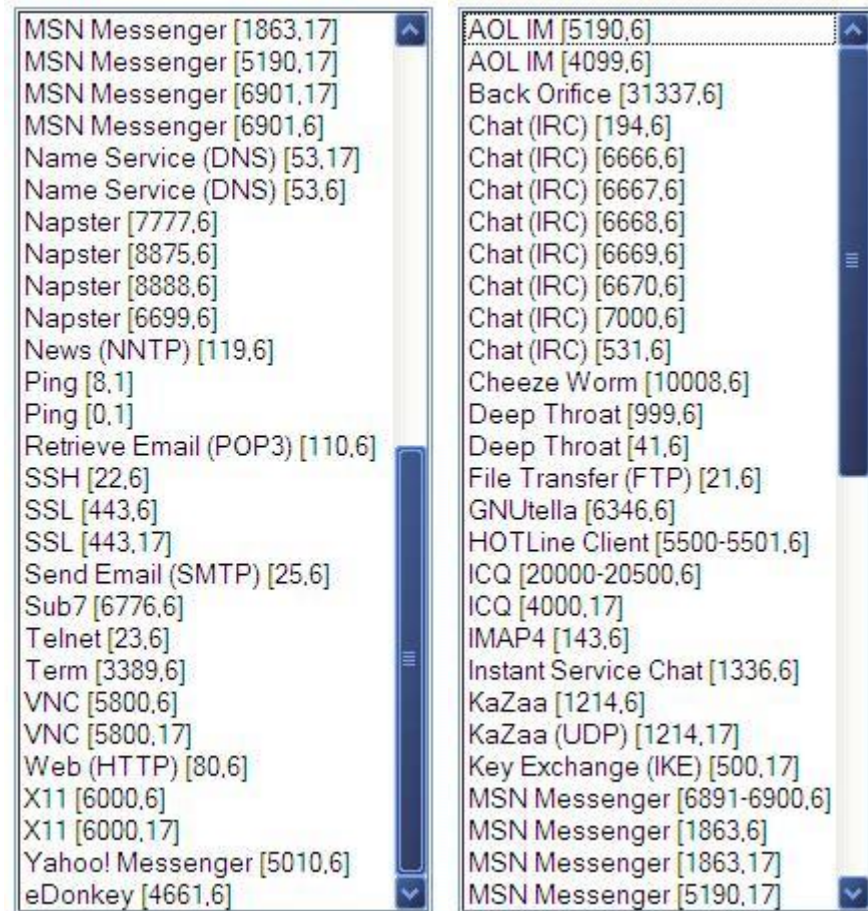


Figure 1.4

So now the network has been tied down properly and much more securely than that of the previous setup, but there is still a major issue that we have not considered, internal security. The users could maliciously or accidentally bring into the LAN a virus or another threat from an outside source, like an email or a floppy disk. A few of these clients have a home version of NAV, usually out-dated by a year or so, but the majority of them are protection free. I spent a few weeks researching the options of Enterprise level Anti-virus software and decided on choosing Symantec Norton Anti-Virus Corporate Edition. This package provided remote administration, roll outs, central quarantine stations, real-time virus scanning, and the possibility of having a central "Live Update" server located on the LAN for client machines to pull new definitions from. This seems ideal and once it was purchased the roll out began.

The major issue here was the Windows 95 and Window 98 clients would not accept the roll-out because of improper "shares" located on the machines and the absence of "Services". The only logical thing to do would be to run fresh clean installs of Windows 2000 Pro on each 95/98 client machine. This would be a timely project and taken into consideration, we decided to set a deadline of 2

weeks on all machines. This did not however prolong our efforts in getting the anti-virus server up and running on all clients possible. We dedicated a machine to be the Primary AV server for each location. These machines were Pentium III 500s with ½ gig of RAM, as the application did not need a lot of process power to run.

Once the base server and the “Snap-in” management console were installed, next the roll-out install package was installed. After the installs were in order and AV servers were rebooted as needed, everything could be managed from a central machine and client installs were to take place. It was as simple as selecting “Roll-out Anti Virus Install” and clicking the clients and servers that needed the install via SMB names. Selecting approximately 30 clients at one location, 10 in another, and finally 60 at the main customer service location, the push installs took only 20 minutes per building at max, reporting that each client had to reboot, advising every user to shut their machines down at the end of the day. Once rebooted, scheduling a full system scans would be initiated, setting critical system directories every night, and complete system scans every Saturday night.

Instantly the next day, we checked the AV server and noticed a significant amount of various virus activities, all which had already been properly quarantined/contained. The only job left here is to properly dispose of the quarantined viruses and constantly check client installs incase something goes wrong. This provided us with the utmost relief in tedious tasks, and provided extra time for the IT staff to research other solutions that could benefit the company, than the wasted time patching every system as needed.

The Aftereffect

Within the past few months, the additions and improvements have come to keep “Company X” doing better than ever. All users as well as the servers are having relatively no problems and the workload from the company has thinned out because of the efficiency of the network and stations. The IT staff are being trained the proper procedures of installs, basic security policies, and have more time to research upcoming security issues.

Daily duties now, replacing the frantic running around patching machine by machine, include

- Scanning for needed updates on client machines as well as servers using the Microsoft HFNetChk (Network Security Hot Fix Checker Tool) described in Larry Nicholl GSEC paper entitled Remote Scanning Utilities for Microsoft Hot Fixes and Service Packs.
- Creating custom Microsoft Management Consoles (MMC) for bulk actions such as user groups, group policies, DNS etc. This helps contain specific ways of doing security procedures, since all client and servers can be managed from a single machine.

- Reading server auditing logs and all error logs. Any malicious successful or failed access logs, we take note of the IP address and host name, applying any needed missing rule into the firewall or denying access from the attacked machine itself. Any unrecognizable errors we locate and seek the reasoning behind it. If there is a way to prevent it, we will apply the patch and/or setting change.
- Reading firewall logs to notice any malicious activity coming from a specific place, and the types of attacks running. Again if there are known attacks coming at a dedicated rate we will take the proper procedure and deny everything from the remote IP. If the attack seems to be coming from an entire subnet we will block the entire subnet for certain amount of time while we research another way of stopping the attack.
- Briefing the anti-virus quarantine server, checking "Live Update", and checking for any corrupted anti-virus client installs. If a virus is noticed that does not seem familiar one will research the virus, collecting any information regarding itself including, the transport procedure, the growth rate, destruction rate, and what else can be done to stop it and create a brief report to pass out or discuss with the remainder of the IT staff.

Since the improvements, "Company X," has nearly increased its revenue rate by two-hundred percent, and promoting me to Sr. Systems Admin. This provides for the everlasting job of upgrading hardware as well as providing the proper pay rate for more experienced employees. Regardless of how well a company is tied down in the security aspect, it would all be irrelevant without the means to maintain the hardware and hire/pay fluent employees. As they say "A chain is as strong as its weakest link," it is our job, as IT professionals, to seek that weakest link, and remedy the situation causing it to be weak, hence the work "hardening."

In Conclusion

Continuing to work everyday on ways to improve security and efficiency of all network traffic, I have come to the conclusion that one will always have a concern on security and privacy. "Company X," like most business networks, is not perfect and will probably never be. The best one can do is to stay trained and influenced in the world of IT security, have the proper time to research the need in which to keep this ongoing training and last but not least, have a fully trained staff that is ready to do what has to be done to complete the job in the most thorough of ways.

For anyone reading this paper, the best way I have found to learn and correct any unknown problems is to jump right in and find the weak link. Patch it, clean it, reinstall it, or whatever needs to be done, and then go to the next problem. You will most likely find that the problems are not the users that we are all fond of blaming but instead the tragedies of a misconfigured network. It is your job to jump and learn to reverse engineer the ongoing problem one step at a time. It may take a few hours or possibly a few months depending on the rate of the problems and the size of the network. Regardless of the pain and struggles

that take place during that time, then ending reward of thinking back at the before and after picture and noticing what you have created is more than worth your effort.

For myself in the IT profession, I feel that every problem has an infinite amount of lessons to be learned. It all depends on how deep you want to dive into the world of Information Security. There has not yet been a computer that is completely invulnerable, and attached to a network. Therefore it's an open world to go into and learn from. The world is dependent upon us, to plan properly, learn, and solve. Do the job to the best of your ability and learn to enjoy what it is you are working for.

© SANS Institute 2003, Author retains full rights

References

Moore, David. "The Spread of the Code-Red Worm (CRv2)." 14 June 2002
http://www.caida.org/analysis/security/code-red/coderedv2_analysis.xml (27 Oct 2002).

Carnegie Mellon University. "'Code Red II: Another Worm Exploiting Buffer Overflow In IIS Indexing Service DLL" 6 Aug 2001
http://www.cert.org/incident_notes/IN-2001-09.html (27 Oct 2002).

Spitzner, Lance. "Building Your Firewall Rulebase." 9 Dec 1999
<http://www.ussrback.com/docs/papers/firewall/rules.html> (29 Oct 2002)

Wood, Shawn. "The Dangers of Peer to Peer Applications within the Enterprise." 31 July 2002
http://www.giac.org/practical/Shawn_Wood_GSEC.doc (2 Nov. 2002).

"Intel® Express 9xxx Series Routers." How to Configure Network Address Translation (NAT)
<http://www.intel.com/support/express/routers/9xxx/23797.htm> (29 Oct 2002).

Mobrien.com. "Subnet Mask Calculator"
<http://screamer.mobrien.com/net.shtml> (1 Nov 2002).

Nicholl, Larry. "Remote Scanning Utilities for Microsoft Hot Fixes and Service Packs." 25 Sept 2002
http://rr.sans.org/win2000/remote_scan.php (4 Nov 2002)