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

James L. Mayne
GSEC v1.4b Option 2

Securing a University Environment; An Evolutionary Case Study

Abstract

Like many small private universities the one for which I work was for many years an open environment as far as network security was concerned. We quickly learned however, that the risks were too great. This case study outlines the steps that my university took to transition from an open network to one that balances the needs of faculty doing teaching and research, students needing to learn as well as be entertained and staff that require a secure and stable network environment to perform their business functions. Through focusing on the way one institution approached this problem I will provide some general methods that other similar institutions may use to aid in their transition. I will also discuss how our methodology drew on the principles that form the foundation of a good network security model. Lastly, I will look ahead and discuss some of the challenges that still face university network security.

Background

My university is a small, private historically liberal arts university. Founded in 1873, it currently has an enrollment of approximately 8,000 students of which 2,000 or so are graduate students. The university has eight colleges including liberal arts, sciences and engineering, communications, fine arts, education, nursing, business and theology. It covers over 260 acres.

The university network began with the wiring of our library and a new business school building in 1988. By 1989 we had acquired an IP address range and connected to a nearby university at 9600 baud. We were on the internet! The network increased enormously in 1996 when we wired all the dorms with 10/100 Mbps connections “to the pillow” (meaning that if the room has three occupants then it has three data drops). At this point the network contained about 7,000 drops. Like many educational institutions of the time the topology was that of one central router with connections to the various subnets and the internet.

During the growth of the network in the 1990’s we were excited to see the enormous increase in usage. The World Wide Web in particular brought many new users to the internet and to our network as well. The university set out to build its network as a tool for faculty, students, and staff and we were beginning to see that goal fulfilled. We also saw the network as a tangible representation of the principles of Academic Freedom. Our network was an extension of the ideals that helped create the internet; it was a way to share ideas like never before with

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

people around the globe. We felt very strongly that any constraints or barriers were not to be tolerated. This idea of academic freedom runs deep at universities. Institutes of learning are one place in particular where the free and open exchange of ideas is not only accepted it is required. The American Association of University Professors motto is “Academic Freedom for a Free Society”[1]. We were certain that there was no way a university could impede this principle. When vendors called on us to discuss a firewall we scoffed at them and told them that they just did not understand the environment in which we were operating.

In the late 1990’s we began to see however that the stability of the network was growing increasingly at risk by outside and inside forces. In addition, in 1998 we began the installation of the PeopleSoft software for use in Human Resources, Payroll and Student Administration. This software as never before would place the network at the center of our business model. It was imperative that the network be as secure and stable as possible for this new direction to succeed.

It was at this juncture that we sat down to review where our network stood and where it would go from here. At this time the Network Services group was comprised of myself, another staff member, and a manager. We began by looking at our population and found that the network had grown into a critical and integral part of TCU for faculty, students and staff. Faculty were using the network more and more as a tool for their research and as a way of communicating and collaborating with their colleagues. There were scores of different agencies, including governments, corporations, and foundations, sponsoring research and creative activities at the university. This meant dollars and prestige. In addition to programs in Computer Science and Engineering that revolved around computing and the network’s ability to deliver those resources to the students, other faculty were beginning to incorporate computing into their curriculum.

Students were using the network to do research on the internet as well as through our library. Of course students were also using the network for playing games, chatting and email. This too was important to us since it provided an incentive to stay in our dorms, which brought in revenue. It provided a selling point for us to help attract new students and it promoted an on-campus form of recreation to our students.

Of course as I mentioned, with the implementation of PeopleSoft most of our business processes including finance, human resources, and the administration of the students would require good solid network connectivity. In looking at our one router topology we saw that a worm, trojan, denial of service attack, or hacking could disrupt this critical link. We had nothing to detect a problem as it arose. We had only our router to stop an attack from continuing and we had no tools to look at our infrastructure after the attack to know what had been lost or

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

damaged. It was time to reassess our network security and take whatever steps were necessary to achieve that goal.

Implementation

In designing a process to begin securing our network we looked at each of our three groups (faculty, students and staff) and decided that they were too different to be encompassed by one approach. The faculty was like a collection of self-employed users for which we were providing a service. They were doing teaching and research that was at the center of our mission and yet they had an autonomy that could not be ignored (their core concern for the preservation of academic freedom must be handled with great care).

The students were a group of consumers. We had to create and maintain a service they were willing to pay for. We had to create an environment that was secure on the one hand and open enough to be a selling point on the other. Unlike totally independent consumers however we did have some say in what students were allowed to do while enrolled. We also had a responsibility to their parents who, in most cases, were paying the bills.

Lastly, we had the staff. These were our corporation's employees. We had to give them the tools to do their job but they would accept any restriction that the administration deemed necessary.

With these three groups and their various requirements in mind it was time to look at what measures we could take to begin securing our network. We looked around at the technology and decided to focus on the following possible measures. These are similar to the measures discussed in the SANS Security Essentials[2] curriculum.

1. Physical security
2. Virus Protection
3. Router access lists
4. Firewall
5. Intrusion Detection Systems
6. VPN
7. Content filtering
8. Uniform system management

Physical Security

Physical security is always important and should not be overlooked. In our case however we could do very little at this point without large funding available to alter the physical security. We did eventually implement access card readers on doors into our central machine room and we will continue to look at what can be done here but since much of the computing resources are out in the hands of faculty and students this was not an area we concentrated on.

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

Virus Protection

Virus protection on the other hand was an area we could address. We had been able to gain financing for a site licensing to cover faculty and staff machines and we expanded that by turning to centrally administered anti-virus software. Users had seen more and more on the news about viruses and so it was visible. We also were able to argue for an email virus scanner for our exchange server as well. This was helped by the experiences with the “Melissa”[3] and “Love Letter”[4] worms in 1999 and 2000.

Router Access Lists

With the “Land”[5] and “Smurf”[6] denial of service attacks in 1997 and 1998 it became apparent that spoofed packets were a serious threat not only to us but also through us to others on the internet. Following the guidelines of RFC 2827[7] we added anti-spoofing access lists to our boundary router as well as disabling directed broadcasts. We could see no downside or negative impact to our population so we enacted this measure right away.

Firewall

Vendors had, of course, been pounding on us for years to install a firewall. We had always mentioned academic freedom and we again looked cautiously at firewalls In the end however we felt we had to know what was coming in and out of the university and we needed to protect ourselves from unwanted attacks and scans.

The faculty were our biggest concern in implementing a firewall. Since they had first begun connecting to the network they had been allowed to start up servers and communicate to/from the internet however and whenever they wanted to. How could we now block incoming traffic to servers they already had running and not disrupt their feeling of academic freedom? We also worried about instances such as the faculty that decided on Christmas day (when the university was closed for over a week and no one was around) to start up a server in collaboration with a distant colleague only to find it did not work because of our firewall.

We felt that there would also be strong resistance among the students who were no doubt running all manner of servers from their dorm rooms. We could of course just tell them that their servers were no longer permitted but we did not want to curtail activity just because we could.

As for the last of our groups, the staff, we knew that here we could take a corporate approach and just tell them that servers were not allowed.

Taking all this into account we decided that the need for protection was too great but to lessen the immediate impact we would use a phased approach to installing an internet firewall. We would insert it into the network in four steps:

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

1. Create a logging-only firewall
2. Block specific threats coming in and out
3. Explicitly allow all current services and deny everything else
4. Go back and certify all allowed incoming and outgoing services.

In addition to this internet firewall we would also protect our administrative servers that were running the PeopleSoft systems with a second internal firewall. This firewall we could install and setup from the start with specific rules allowing the necessary traffic and blocking all other traffic.

During phase 1 we moved our primary router out to the internet connection carrying the anti-spoofing access lists with it. We then installed the new firewall with three interfaces. The first was to the internet router, the second to the student subnet containing all the RESNET (residential students) users and the third was to the faculty/staff subnet. We turned on logging and spent some time becoming comfortable with the firewall, its software, and the traffic we now saw going into and out of the three subnets. This was unseen by our users and so had no negative impact.

During Phase 2 we stopped all the small services such as chargen, daytime, echo, etc. as well as services that we felt were risky and from the logging we knew were not in normal use. These included NFS, NBT, SNMP and other protocols. Several of these are included on the "Ten Most Often Exploited Internet Security Flaws"[8] list at SANS. This was also the point at which our student and faculty/staff policies diverged. We decided as consumers that we would continue to allow all incoming and outgoing services for students. We treated them as an ISP would treat their customers. This also meant that when we secured our faculty/staff subnet we would treat students as we would an internet user and we would always be mindful that their subnet was an unsecured area. This is to a large degree the policy in place today. This phase again had minimal impact on our faculty/staff and none on our students since we allowed all traffic by default.

With phase 3 we really made the paradigm shift from "deny some, allow all others" to "allow some, deny all others'". Of course we were not very exclusive in that we allowed anything we were seeing in the logs, no questions asked. We ran this way for a while as we began to pave the way for the real restrictions that would follow. During this time we went to the university computing committee representing faculty, staff and students and kept them informed as to what we were doing and what we had planned. We were again helped along by timing however as around this time we had begun to see security issues such as the "Code Red"[9] and "Nimda"[10] worms that exploited Microsoft IIS vulnerabilities and we could more easily argue that there were machines all over campus that were running software that was vulnerable. We argued that if we did not block all

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

traffic by default that any one of these vulnerable systems could become an inside threat to our entire network.

Phase 4 is really where the rubber hit the road so to speak. This was our first attempt to limit what a faculty member was doing. This is where the firewall rules begin to adhere to the Principle of Least Privilege. We went back and began reviewing the servers we did not centrally administer and yet had been allowing through the firewall under phase 3.

We had faculty that were running servers they did not even administer directly. They had their students doing the administrative work and very often did not know from one student to another what was happening on their systems. Before we went after the faculty we needed an unambiguous plan of attack. Our client/server group came up with a set of guidelines that described what a system should have on it and what a systems administrator should do to make and keep a machine somewhat secure. We required centrally administered anti-virus software be installed. We required that patch levels for the OS and the running software be at minimum levels and that these be checked through an automatic update or some other defined mechanism. The policy required there be current maintenance on the machines and that there was someone that had the knowledge necessary to responsibly administer the machine.

After we had the necessary policy in place we again went to the university computing committee and discussed our plans with them. We first asked all faculty running servers that ran services that we also provided in our central IS department to switch over to our servers. Some faculty agreed so they did not need to comply and keep complying with the new server policy. Some still resisted, feeling as though they would lose control of their research if they lost control of their servers. In these cases we assured the faculty that we would not pre-judge what they needed. If they had a server that would pass our requirements we would allow the traffic. They were educated to some extent about security and we had their assurance they would maintain their server's security.

We again allowed the students to continue as we had in phase 3, running their own servers and treating them as internet users. With the staff we were able to move most servers back into the central IT facility. We did run into some staff servers that for some reason or another we could not combine into our operations but there were only a couple of those and they were also required to comply with the client/server groups new server security policy.

Intrusion Detection Systems

So now we had central anti-virus software, routers with access lists and a firewall that blocked all traffic unless explicitly allowed. We next turned our eye to Intrusion Detection Systems (IDS). IDS systems can "Provides intelligent, automated integration of threat assessment, intrusion detection, active blocking

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

and data analysis within a self-contained, remotely managed application."[11]. Once again looking at our populations we saw minimal impact. We decided on two sensors. The first one we placed between the internet router and the firewall. This watches all traffic in and out of the internet. The second sensor we positioned watching the faculty/staff interface on the firewall. We did this to help watch traffic passing from the student subnet into the faculty/staff subnet as well as a sanity check on the internet sensor. If this sensor saw something we thought we were stopping at the firewall we knew we had a problem in our rule base.

We looked around and decided on a sensor that could communicate back to our firewall adding temporary rules to block traffic judged to be hostile. Of course IDS's are good and bad. They are notorious for false alarms and they need to be tended but the information is invaluable in terms of seeing a pattern occurring over a period of time. We began the operation of the sensors with a period of just watching and reporting. Once we were comfortable with their operation we did decide to allow some temporary rules to be added to the firewall. We do this only for obvious scanning of ports or specific threats like Nimda and Code Red that are not as susceptible to the false signals. The combination of the firewall and the IDS sensors and their logs are helping us implement the Principle of "Prevention is Ideal but Detection Is a Must".

VPN

We now had centrally administered anti-virus software, routers with access lists, a firewall and two IDS sensors. We were definitely building up the defenses in depth. Now we looked at a Virtual Private Network (VPN). When we clamped down on the allowed incoming traffic using the firewall it meant that users, especially administrative programmers, that did some occasional work from home had no way to mount drives and run software from home as they could from the university. We turned to a VPN concentrator as the answer. With the VPN in place we could handle requests from staff to work outside the university using broadband connections.

Of course with the VPN in place you now have a tunnel into your network so placement is critical. We originally placed the VPN concentrator inside the faculty/staff subnet so it was easy for users to work as they did from their offices. The concentrator was exactly where their office was in the overall topology. This opened up the risk however that they or someone with control of their home machine also had a tunneled connection through our firewall and IDS sensors into our network. For this reason we will soon be moving the VPN concentrator outside the firewall. We will add the necessary rules to the firewall to allow the traffic we feel is necessary. We will then also be logging that traffic and it will no longer be encrypted either.

Content Filtering

Content filtering is a method that some companies use to help reduce security risks. Again going back to the idea of academic freedom this was simply not an

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

option for us. We are currently having discussions in fact over spam filtering and even in this arena of content filtering we are treading softly. We are again talking with the university committees and getting buy in from the administration and faculty before we proceed.

Uniform System Management

Lastly we looked at the way in which servers were administered. Even with all we had done if one system on the faculty/staff subnet was compromised it left all others vulnerable to attack. This is a difficult area when it comes to faculty. As mentioned before faculty often have servers, some with multi-user interactive operating systems such as Linux that are not actually administered by the faculty but by their students. Sometimes these machines are not even the property of the university but came as part of a grant or award or are a necessary component of some large instrument. How do you protect one faculty member from another? We are dealing with this issue through education. We will never be able to prevent local administration of machines but hopefully we can give the faculty and staff some easy to follow guidelines on how to best secure their machines. In addition we are looking at ways to remotely monitor these machines periodically so we can find non-secure servers before a hacker does.

In terms of our own central servers we are looking for ways to create server images that are secure from the beginning. These images are used to setup new servers. We are also looking at auditing and fingerprinting techniques so we can tell when a system has been altered and what may have been damaged. This will help us with the principle of "Know Your System". The client/server group is also going back over privileges and looking at ways to implement the "Principle of Least Privilege". This is difficult because it is taking away something that someone already has. It is certainly better if this principle can be kept in mind from the beginning.

Current Status

Today we have a network that in a few years has transitioned from a single router open network to one that has several lines of defense. We have centrally administered anti-virus software that we load on end user machines.

We have an internet router with ingress and egress access lists to help stop spoofed address packets from entering and leaving our network.

We have installed two firewalls: the first is an internet firewall that through a phased in approach explicitly allows traffic we want to pass and denies by default and an internal firewall which protects our business servers.

We added two IDS sensors, one between the firewall and the internet router and one between the firewall and the faculty/staff subnet. We would also like to add a third sensor between the student subnet and the firewall at some later time. We

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

have a VPN concentrator in place and will be re-positioning it between the internet router and the firewall as soon as possible.

We rejected the idea of explicit content filtering however we are looking at spam email filtering appliance. The client/server group is attempting to create more uniform secure servers and looking for ways to audit those servers on a regular basis and when a compromise is suspected.

At the time this process began, I did not know the four basic security principles taught in the SANS GSEC course namely “Defense in Depth,” principle of “Least Privilege,” “Know Your System” and “Prevention is Ideal but Detection is a Must” however you can see how those principles are woven into the approach we chose to take.

Today there are still challenges we must overcome. We need to continue pruning down our firewall rules, allowing fewer servers to receive internet traffic. We need to continue to improve the IDS signatures and their sensitivities. Of course they also have to be continually updated and reviewed. We will need to look more at host-based security solutions such as host-based IDS software and host-based firewalls.

We are also being pressured to deploy wireless technologies across the campus. Administrators feel compelled to do so partly to compete with other universities that can say they have it. The growing number of devices with wireless capabilities built-n is pressuring us. We will begin to see applications that will grow from this critical mass and we need to be ready to secure a network that is no longer wired and geographically secured.

Conclusion

Like all universities we were faced with the reality that we could no longer go without defenses and yet we had to deal with a population that contained three distinct groups of users with three distinct requirements. We were also always mindful of the unique environment of a university and the importance of not only allowing but also promoting free speech and the exchange of ideas.

I believe that through an evolutionary process and with the help of university committees that included all parties we were able to balance these needs of the faculty researcher, the teacher, the student and the staff member and create an environment that provides all these users with a more stable and usable network.

References

[1] American Association of University Professors, “AAUP Home Page.” URL: http://www.aaup.org/ (15 Dec. 2002)

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

[2] SANS Institute. SANS Security Essentials + CISSP CBK. SANS Institute. 2002

[3] CERT. “CERT® Advisory CA-1999-04 Melissa Macro Virus”. March 1999. URL: http://www.cert.org/advisories/CA-1999-04.html (12 Dec. 2002)

[4] CERT. “CERT® Advisory CA-2000-04 Love Letter Worm.” May 2000. URL: http://www.cert.org/advisories/CA-2000-04.html (12 Dec. 2002)

[5] CERT. “CERT® Advisory CA-1997-28 IP Denial-of-Service Attacks.” May 1998. URL: http://www.cert.org/advisories/CA-1997-28.html (12 Dec. 2002)

[6] CERT. “CERT® Advisory CA-1998-01 Smurf IP Denial-of-Service Attacks.” March 2000. URL: http://www.cert.org/advisories/CA-1998-01.html (Dec 12 2002)

[7] Ferguson, Paul and Senie, Daniel. “RFC 2827, Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing.” May 2000. URL: http://www.rfc-editor.org/rfc/rfc2827.txt (10 Dec. 2002)

[8] SANS Institute. “How To Eliminate The Ten Most Critical Internet Security Threats The Experts’ Consensus.” June 25, 2001. URL: http://www.sans.org/topten.htm (15 Dec. 2002)

[9] CERT. “CERT® Advisory CA-2001-19 "Code Red" Worm Exploiting Buffer Overflow In IIS Indexing Service DLL.” July 2001. URL: http://www.cert.org/advisories/CA-2001-19.html (10 Dec. 2002)

[10] CERT. “CERT® Advisory CA-2001-26 Nimda Worm.” Sep. 2001. URL: http://www.cert.org/advisories/CA-2001-26.html (10 Dec 2002)

[11] Internet Security Systems. “RealSecure® Network Protection.” URL: http://www.iss.net/products_services/enterprise_protection/rsnetwork/index.php (19 Dec. 2002)

[12] Cisco Systems. “Design Guide Cisco IOS Firewall”. URL: http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/firew_dg.htm (15 Dec.2002)

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