



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 (Security 401)"
at <http://www.giac.org/registration/gsec>

Implementation of a Secure Wireless Network on a University Campus

Greg Redder
Submitted for SANS GIAC Practical
Assignment Version: 1.2f

© SANS Institute 2000 - 2005, Author retains full rights.

Table of Contents

- Implementation of a Secure Wireless Network on a University Campus
- Figure 1: Diagram of wireless network layout at Colorado State University
- References
- Questions

© SANS Institute 2000 - 2005, Author retains full rights.

Several challenges confront the implementation of a wireless network on a University campus, but the challenge central to this topic is security. Henceforth, I have outlined in detail several possible solutions in maintaining a wireless network, the design of our network in order to encompass such solutions, the requirements within which our wireless infrastructure was created, and finally, various scenarios illustrating how specific security issues have been addressed.

Wireless networks are inherently insecure. One of the standards designed to address this inherent lack of security, the Wireless Encryption Protocol (WEP), is full of holes with new ones continually being found¹. A Virtual Private Network (VPN) connection over a wireless link is a viable alternative, however only for a limited type of devices. For example, numerous Personal Digital Assistants (PDAs) and other handhelds do not support VPN clients. Given these challenges, Colorado State University (CSU) worked to provide as secure a wireless network as possible. Colorado State University was not the first to tackle this challenge. Others, specifically SANS and the University of Utah have addressed many of CSU's initial concerns. Mr. Chris Hessing at the University of Utah enumerated these concerns and others in his paper entitled WAAC (Wireless Authenticated Access Control). Specifically, he stated²:

1. The solution must authenticate ANY and ALL wireless equipment wanting access to the University wireless infrastructure.
2. The solution cannot require any applications on the machine that desires authentication.
3. The solution cannot rely on a specific platform. It should work with all devices that can get a wireless network connection.
 - a. The solution must provide full line rate forwarding through the network.
 - b. The solution should not use any firewalls, or other packet inspection methods that interfere with the performance of the traffic.
4. The solution must allow all forms of network traffic.
5. The solution must allow administrators to identify the users/machine in the event of a problem report from someone on the Internet.

These points along with what was learned in the SANS GIAC material were taken into consideration in designing the wireless network at CSU. The design team thus developed four requirements: 1) to physically and virtually secure the network equipment; 2) to limit access to the network to only those individuals affiliated with the University; 3) to provide for VPN access; 4) to provide for access to the network for those devices without VPN clients but only after authentication. The last requirement included warning the user of the insecure nature of the wireless network and the need to link the wireless device to a specific individual.

The design of CSU's wireless network is depicted in figure 1. Cisco Wireless Access Points (APs) were installed throughout campus. These APs were all connected to the same Virtual Local Area Network (VLAN) and given IP addresses in the reserved address space of 10.10.0.0/255.255.0.0. Additionally, a Cisco VPN Concentrator 3000 was configured with an interface on this private network. The VPN concentrator had another interface connected to the campus' class B IP address space. The VPN concentrator provided a secure path between the wireless device and the VPN concentrator. It also acted as the intermediary between the wireless network and the Internet. The VPN concentrator unencrypted packets sent to it by the wireless device and converted them to packets with an IP address in the campus' routable class B address

space. It then passed them along toward their destination. A firewall comprising a Linux computer with IPChains was installed and connected in a similar manner as the VPN concentrator. The Linux computer also had an interface on the private network and another to the campus Local Area Network (LAN). The Linux firewall oversaw several services: routing and authentication for non-VPN clients, Domain Name System (DNS) service, Hyper Text Transport Protocol (HTTP) and Hyper Text Transport Protocol Secure Sockets Layer (HTTPS). Finally, this network was connected via a combination of HP4000 and 3Com3300 switches connected into a router interface on a Cisco 6509. Most of the 6509's routing functions for its wireless interface were disabled so that it only passed certain packets as will be described in detail later. This equipment completed the physical infrastructure necessary to provide for secure access to the campus wireless network. Next, we had to meet several requirements to ensure that the system was as secure as possible.

The first requirement was restricting physical and virtual access to the network equipment. Restricting the access points physically actually posed more of a challenge than anticipated. Devices had to be installed in classrooms, offices, and other public locations. For the most part, the devices were attached to the walls 12 to 15 feet above the floor. Other devices were placed in locked Intermediate Distribution Frames (IDFs) also known as "phone closets". Still others were installed in private offices. Unfortunately, it may be possible for someone with a ladder to steal one of the more public access points. This is an admitted area of weakness. However, all access points are monitored by network management software and alerts are generated as soon as a device becomes unresponsive due to power outage or theft. Meanwhile, all switches are located in locked IDFs or Main Distribution Frames (MDFs). Keys to these locations are only acquired after a background check is completed on the requesting individual and various agreements and contracts are signed regarding activity allowed in such locations. This provides for a reasonable level of physical security to the switches. The router is physically secured in the campus' main computer room. Access to this room is only provided to individuals who have equipment in the site. This area is staffed around the clock and monitored via closed circuit TV to the campus police department. Together, these locations and procedures provide for the physical security of the wireless equipment.

The next task was securing virtual access to the equipment. Access points and switches require login names and passwords for access. This access, however, is done in plain text. While the devices are on a switched network, a program such as `dsniff`³ could conceivably glean such usernames and passwords. The likelihood of this being done successfully is small, but it is possible. The Cisco 6509 router is on a secure VLAN with access to it restricted to the network staffs' equipment on that same VLAN. For additional security, the router is accessed via SSH. Connections from outside the VLAN are rejected via access lists. Spoofing is prevented via commands in Cisco's Internetwork Operating System (IOS). Finally, Simple Network Management Protocol (SNMP) access is permitted via a secret community string to the switches and access points. SNMP version 1 is used and its inherent security holes are recognized⁴. Fortunately, SNMP on the router is restricted to the secure VLAN. Thus, while virtual access control is not perfect, the network team has taken initial steps for making it so. The potential loopholes have been identified and methods of closing the holes discussed.

The third requirement centered on providing VPN access for the wireless network. As shown in figure 1, a VPN concentrator was connected to the wireless network and to the campus LAN. Instructions for connecting to the concentrator were documented at

<http://www.colostate.edu/acns/vpn> to aid individuals in configuring their VPN client. Once configured to connect with the VPN concentrator, the client requests a username and password pair that is securely transmitted between the VPN concentrator and its client. The username and password are then compared via Remote Authentication Dial-in User Service (RADIUS) to the central user database. Again, this is a point of some security laxness as usernames are passed in clear text via the RADIUS protocol (RFC 2138). Passwords, while MD5 encoded, can be easily decoded with tools found on the Internet. As with the access points and the switches, the ability to capture the proper packets, while technically possible, is practically infeasible. Nonetheless, as with the other security issues, it is a point of weakness and should be addressed. Once the authentication has taken place, the VPN concentrator completes the encryption process between the wireless device and itself. As mentioned, the VPN concentrator assigned IP addresses from the campus' routable class B address space to the VPN connection. This VPN connection then allowed for secure access at least over the wireless network.

The fourth requirement was to provide access for devices without VPN clients. This requirement included authentication, warning of the insecure nature of the wireless network, and an ability to link a wireless device to a specific individual. This was the most complicated, detailed component of the network. It was accomplished in part with a Linux computer running IPChains and some special configurations on the router interface. A Linux computer was installed and secured per the instructions in SANS' "Securing Linux step-by-step, version 1.0"⁵. That document and a SANS article on securing Linux⁶ describe the benefits of using IPChains to secure such a Linux computer. In CSU's case, IPChains would play a critical role in acting as a firewall and securing access to the campus wireless network. Additionally, both traditional HTTP and HTTPS web servers were installed on this firewall computer. These servers displayed pages to collect authentication information that was relayed to the central RADIUS server. Once authenticated, the device was allowed through the IPChains firewall while the web server would display pages warning of the insecure nature of wireless networks. This was done because the network team felt it was important to advise individuals of the security downfalls of wireless networking. Finally, access via the firewall and the VPN concentrator were logged and archived. This satisfied another part of the fourth goal, the ability to link wireless devices to a specific individual.

To ensure access to the network was only granted to those who authenticated with the firewall or VPN concentrator, the router interface was configured so that only Bootstrap Protocol (BOOTP) packets could transition in and out of its interface to the wireless Vlan. Its routing was disabled via access-lists. This was done to prevent an individual from manually entering an IP address on his wireless device and providing the router's IP as his gateway address. Since this method would circumvent the wireless LAN's security, it needed to be addressed and prevented. This forced all traffic from the wireless network to pass through either the firewall or the VPN concentrator.

The following scenarios show how the security issues detailed above were addressed. The first scenario assumes that the connection originates from a wireless device with VPN capability. This wireless device is powered on and sends a BOOTP request that is passed on by the router interface to the Dynamic Host Configuration Protocol (DHCP) server on another subnet. The request is returned, via the router interface, back to the wireless network and accepted by the wireless device. The device is assigned an address on the 10.10.20.0/255.255.252.0, "private", network. The DHCP server also returns the IP of the firewall as the configured gateway. Thus,

further restricting the device to only access other devices on the private network until a VPN connection is established or is authenticated with the firewall, as explained later.

Once a device has a valid address on the private network, the device's operator establishes a VPN connection. This VPN connection is established with the VPN concentrator that has an interface on the private network. The VPN server authenticates usernames and passwords with the central RADIUS server that contains all individuals associated with the University. These usernames and access times are recorded and archived. Once authenticated, the VPN connection is assigned an IP address from the campus' routable Class B network that gives the device full access to the Internet.

The second scenario revolves around a device that does not support VPN connections. Initially, the University considered only allowing devices that supported VPN connections access to the wireless network. However ideal that would be, this was impractical given the proliferation of PDAs, handhelds and other wireless devices for which VPN clients are either unavailable or in their infancy. Practically and politically, their access could not be denied, yet they could not just be left to roam freely. It was imperative that access to the wireless network be restricted to those affiliated with the University and to make them aware of the insecure nature of wireless networks.

This access control was accomplished by configuring the Linux server/firewall to provide customized DNS replies. In conjunction with the firewall, all DNS packets coming from computers that have not yet authenticated with the firewall would be redirected to the DNS server configured on that same firewall. This DNS server would be authoritative for any domain queried. In other words, it would return its own IP address for any address it was asked about. The firewall, as mentioned, has both http and https servers. Thus, when a web request is made from an unauthenticated computer, first a DNS query goes out and is captured by the firewall that redirects it to its own DNS server. The DNS server returns its IP address in the place of whichever address the wireless owner typed in. Thus, the ensuing web request is sent unknowingly from the client to the IP of the firewall. There, an http server awaits any incoming request. Since http requests are not secure, a web page is displayed explaining the insecure nature of wireless networks along with a link to the secure server for authentication to take place. The individuals' web browser is then redirected to the https server on that same firewall for which a valid certificate has been purchased and installed. This server presents a web page asking the individual to provide a username and password. Since this is a secure page, such information can be entered confidentially, even over the wireless network. As with the VPN concentrator, the login information is passed onto the central RADIUS server. Upon successful authentication, the IP assigned to the wireless device is placed in the firewall such that the firewall will permit packets to pass through untouched. Also, the firewall does Network Address Translation (NAT) and masquerades as an IP on the campus' class B network. Usernames, dates, times, MAC addresses and corresponding IP addresses are logged and archived. Finally, after being authenticated, the individual is presented with a web page that reemphasizes the lack of security on the wireless network and explains that if the device supports VPNs such a connection should be used instead. The page also explains that all information sent to non-secure sites is vulnerable to interception. After which, the wireless device's information is logged and it is allowed access to the campus LAN and the Internet.

CSU recognizes that the human component, while not a part of any network diagram, is still a critical link. It is imperative that information about VPNs and wireless networking be made available to users. To complete the project, web pages and documentation were written to help

campus members understand how a VPN could help them, how to recognize secure web pages, and to raise awareness about security issues surrounding the wireless part of the campus LAN. VPN documentation is available at <http://www.colostate.edu/acns/vpn>⁷. Wireless documentation is still underway as the project nears completion, but preliminary documentation is available at <http://www.colostate.edu/acns/wireless>⁸.

After completing the design and installation of any network, it is critical that security issues be identified. As mentioned earlier, various security holes existed in the network. The flaws in SNMP version 1, physical access to wireless access points, susceptibility of switched traffic to dsniff, are three already identified examples. The following issues have further been identified:

A computer on the wireless network, without authenticating, has access to all other computers on the wireless network. This would be a problem, if, for example, a wireless laptop with the W.32.Nimda.A@mm virus was connected into the wireless network and began searching for and possibly found another laptop to infect on that same network. The proliferation of worms internal to the wireless network is conceivable. While this same problem is present on virtually any LAN, it is a problem and even more so in a wireless network where the device is mobile.

Another potential security hole exists in the possible “hijacking” of an authenticated device. For instance, when a non-VPN enabled wireless device is assigned an IP address and consequently authenticated, that IP address has full access to resources. If that person was to shut their device off and another individual was to manually configure his device with that same IP address, he would have essentially “hijacked” the IP address and have access to the network without him or his device actually being authenticated.


Once such security flaws were identified, steps were taken to address them. Unfortunately, solutions did not exist for each flaw. In these cases, reasonable effort needed to be taken to make the cost of exercising the flaw far outweigh the benefit. First, the flaws in SNMP version 1 are simply inherent to the protocol. Later versions of SNMP address these concerns, but major vendor implementation and backwards compatibility is lacking. Thus as mentioned, CSU used secret community strings, designed a switched network to prohibit sniffing, and where possible restricted the computers that could even use SNMP to query devices. Second, physical access to the wireless access points has been restricted as much as practical and the devices monitored as closely as possible in a University setting. Third, dsniff’s ability to sniff packets on a switched network is disturbing. This potential security threat has been identified, but not resolved. Fourth, virus and worm spreading on the wireless network is as much a battle there as it is on the rest of the campus. Ongoing user education and vigilant patching of systems appears to be the only practical solution to that effort. Fifth, the potential hijacking of an IP was limited. The firewall continually watches the Media Access Control (MAC) addresses on the wireless network. When it sees a device “disappear” for a specified amount of time, that device’s access through the firewall is removed. This reasonably limits the amount of time such an attacker has to actually detect a free IP address and hijack the connection. While not eliminating this vulnerability, it does restrict it severely. Other issues may still exist that were not enumerated here. Nonetheless, the first and important steps in the identification and elimination of security threats on the wireless network have begun.

The challenges surrounding the securing of a wireless network on a university campus are many. CSU recognized these challenges and presented the networking team with four concrete

requirements. These requirements were met through a combination of securing physical access, securing virtual access, providing for VPN connections and authenticating individuals. Since not all devices can maintain secure VPN connections, they must at least have a secure method for authenticating and some ability to inform the device's owner that further information may be compromised. Many of the parts of the wireless network were engineered based on wireless networks deployed at other universities and information gleaned from the Internet, including SANS. Currently, the University has 73 access points deployed across one square mile of campus serving 23,000 students and 5,000 staff/faculty. Not all buildings are covered, but many classrooms, dormitory dining halls and the Student Center are part of the wireless network. Through thoughtful engineering, research, planning, education and documentation, the security model detailed herein has provided a solid foundation that can be further built upon.

© SANS Institute 2000 - 2005, Author retains full rights.

Figure 1



References

- ¹ Fisher, Dennis and Nobel, Carmen, “Wireless LANS Dealt New Blow”, eWeek, August 10, 2001, URL: <http://www.zdnet.com/eweek/stories/general/0,11011,2803615,00.html>
- ² Hessing, Chris, “WAAC (Wireless Authenticated Access Control)”, URL: http://www.research.utah.edu/networking/anl/current_projects/wireless/WAACwhitepaper.html
- ³ Song, Dug, “dsniff”, URL: <http://www.monkey.org/~dugsong/dsniff/>
- ⁴ Camacho, Jose Luis, “SNMP Security Enhancement”, March 28, 2001,, URL: http://www.sans.org/infosecFAQ/netdevices/SNMP_sec.htm
- ⁵ “Securing Linux Step-by-Step, Version 1.0”, the SANS Institute.
- ⁶ Retallack, Roger, “Securing Linux Installations”, June 15, 2001, URL: http://www.sans.org/infosecFAQ/linux/sec_install.htm
- ⁷ Wilson, Nancy, “VPN at Colorado State University”, September 12, 2001, URL: <http://www.colostate.edu/acns/vpn>
- ⁸ Wilson, Nancy, “Wireless Networking At Colorado State University”, September 17, 2001, URL: <http://www.colostate.edu/acns/wireless>

© SANS Institute 2000 - 2005, Author retains full rights.