



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
"Hacker Tools, Techniques, Exploits, and Incident Handling (Security 504)"
at <http://www.giac.org/registration/gcih>

Internal Espionage Incident: A perpetrator exploits the weakness in a Web Authentication System that uses NT Usernames and Passwords to authenticate staff before they are given Internet access. The perpetrator is able to obtain the NT Username and Password of another employee in order to access a secret document.

Daniel Mossman

GIAC Certified Incident Handler (GCIH) Practical Assignment
Version 2.1 (revised April 8, 2002)

Option 1 - Exploit in Action

Table of Contents

<u>Abstract</u>	3
<u>Part 1: The Exploit</u>	3
<u>Name</u>	3
<u>Operating System</u>	4
<u>Protocols/Services/Applications</u>	4
<u>Brief Description</u>	4
<u>Variants</u>	4
<u>References</u>	5
<u>Part 2: The Attack</u>	5
<u>Description and diagram of the network</u>	5
<u>Protocol description</u>	10
<u>How the exploit works</u>	13
<u>Description and diagram of the attack</u>	16
<u>Signature of the attack</u>	30
<u>How to protect against it</u>	34
<u>Part 3: The Incident Handling Process</u>	36
<u>Preparation</u>	36
<u>Identification</u>	37
<u>Containment</u>	44
<u>Eradication</u>	46
<u>Recovery</u>	47
<u>Lessons Learned</u>	49
<u>References</u>	52

© SANS Institute 2003. All rights reserved. Author retains full rights.

Abstract

The topic of this paper is an internal espionage incident in which a disgruntled employee, who we will call Ted Attacker, gains access to secret documents that contain information concerning valuable biological research.

Ted Attacker is a skilled member of the Widget Research's technical staff. In this incident, Ted Attacker uses his knowledge of the company's network to obtain the NT username and password of a fellow employee who we will call Joe Victim. Ted Attacker obtains Joe Victim's NT password by exploiting a weakness in the company's Web Authentication System. Widget Research uses a Web Authentication System that uses NT usernames and passwords to authenticate staff before they are given Internet access. This system has a security weakness in that at one point in the system the passwords are sent over the network in clear text. Once Ted Attacker gains the password and username of Joe Victim, he then logs into the company's Windows 2000 network with these credentials. Ted Attacker, thus gains access to the secret research paper that Joe Victim keeps on his network drive. Ted Attacker then copies the secret document and is in a position to seriously harm Widget Research as he intends to give this information to a competitor.

This paper will provide details on how the perpetrator carried out this internal espionage incident. It will also provide a detailed account of how the incident was handled by the IT security team at Widget Research.

This paper is based on a real organization and the associated information systems, corporate culture, security policy, and incident handling process of the organization. The incident that is described in this paper is meant as a demonstration of how the organizations systems could be exploited and does not describe an incident that actually happened. However, the organization has had real incidents that are very similar to the one being described in this paper.

Part 1 – The Exploit

Name

Internal Espionage Incident: A perpetrator exploits the weakness in a Web Authentication System that uses NT Usernames and Passwords to authenticate staff before they are given Internet access. The perpetrator is able to obtain the NT Username and Password of another employee in order to access a secret document.

Operating Systems

- Microsoft Windows 2000 Workstation with Service Pack 3
- Microsoft Windows 2000 Server with Service Pack 3
- Microsoft Windows 2000 Server, with Service Pack 3
- Red Hat Linux 7.0
- Red Hat Linux 7.3

Protocols/Services/Applications

- Cisco Secure Access Control Server v 3.0 using Radius Authentication running on Microsoft Windows 2000 Server with Service Pack 3
- Check Point VPN-1/FireWall-1 running on Red Hat Linux 7.0
- The Radius IEEE protocol Version 2
- Windows 2000 Kerberos Version 5 Authentication
- Misalliance command line utilities such as "net.exe" and "ping.exe".

Brief Description

An employee of the company, Widget Research, obtains access to a secret document that he plans to give to a competitor company. He accomplishes this task in a three-stage attack:

- 1) The attacker conducts surveillance in order to discover the victim's Internet Address (IP) and the victim's computer and Internet work habits.
- 2) The attacker exploits a weakness in the company's Web Authentication System in order to intercept the victim's NT username and password by deploying a "man in the middle attack". The organization uses a Network Authentication System that utilizes each employee's Windows 2000 username and password to authenticate the employee's ability to access the Internet.
- 3) The attacker then uses the victim's NT password and username to map a network drive to the victim's network drive in order to successfully copy a secret research document from the victim's drive to a 3.5 floppy diskette. The perpetrator then plans to give the document to a competitor of Widget Research.

Variations

This incident deals with an internal attacker taking advantage of a vulnerability in Widget Research's Web Authentication System. The incident is unique in that it involves an "in house" designed system. One key stage of the attack involves utilizing a well know "hacking" tool called Dsniff in order to "ARP spoof" the victim's workstation in order to intercept network traffic that was traveling from

the workstation to the Firewall. However, this paper is not focused on an ARP exploit by means of Dsniff tools, rather it focuses on a multiple stage attack against the company's Web Authentication System that requires the attacker to utilize multiple tools and his inside knowledge of the organization's Network.

References

The Following URLs have information related to this incident (Detailed References can be found at the end of this document.):

- <http://httpd.apache.org/docs/howto/auth.html> (Information on HTML Authentication)
- <http://www.microsoft.com/windows2000/docs/kerberos.doc> (Details on the Kerberos Authentication protocol.)
- <http://www.sans.org/rr/penetration/dsniff.php> (Details on dsniff signatures)
- <http://www.sans.org/rr/audit/dsniff.php> (Good overview of the dsniff tool)
- http://www.cisco.com/en/US/tech/tk583/tk59/technologies_tech_note09186a00800945cc.shtml (General information about the Radius Protocol)
- <http://www.untruth.org/~josh/security/radius/radius-auth.html> (Detailed information about the Radius Protocol).

Part 2 – The Attack

Description and Diagram of Network

There are several steps in this espionage incident but its success is based on a weakness in Widget Research's Web Authentication System. The company policy requires internal users to be authenticated before accessing resources on the Internet (e.g., web sites, FTP sites, Telnet hosts, etc.). To keep things simple for the end users Widget Research has designed a system that prompts users for their NT username and password when they attempt to access the Internet.

The authentication system is based on a Check Point Firewall that uses a combination of HTTP and Radius authentication. When a user attempts to access an Internet host they are first prompted for a user name and password via an HTTP authentication request that is sent by the Check Point Firewall. The Check Point Firewall then passes the user's credentials to a Cisco secure Access Control Server (ACS) v 3.0 for Radius Authentication. The ACS server then communicates with the Windows 2000 Active Directory domain controls to see if the username and password is valid. The ACS then communicates back to the firewall as to whether the credentials are valid or invalid and then the firewall permits or denies access to the Internet host.

There is major security issue in Widget Research's Web Authentication System in that the passwords are passed between the user's browser and the Check

Point Firewall in clear text. Although the communication is encrypted as part of the RADIUS protocol between the Check Point server and the Radius server, the HTTP authentication request to the end user's web browser is not encrypted. Hence, confidential NT usernames and passwords are sent from end-users workstations to the Firewall in clear text.

The attacker obtains the victim's NT password when the clear text password is passed between the victim's computer and the Check Point Firewall during an HTML authentication request. As the Network is Micro-switched (each computer and server is connected to a port on a Network Switch) the ARP poisoning tool "Dsniff" is required for the attacker to be able to sniff the password off the Network.

In order to demonstrate and provide further research on this exploitation, the attack was replicated in a lab that simulated Widget Research's Network. The lab network is a simplified but accurate model of Widget Research's Web Authentication System. The full attack was successfully conducted within this test lab. The following is the Network Diagram (Figure 0-1 Lab-Network) from the test lab and a detailed explanation of all the components in the lab is contained in Table 0-1 Details of Lab-Diagram.

-

© SANS Institute 2003, Author retains full rights.

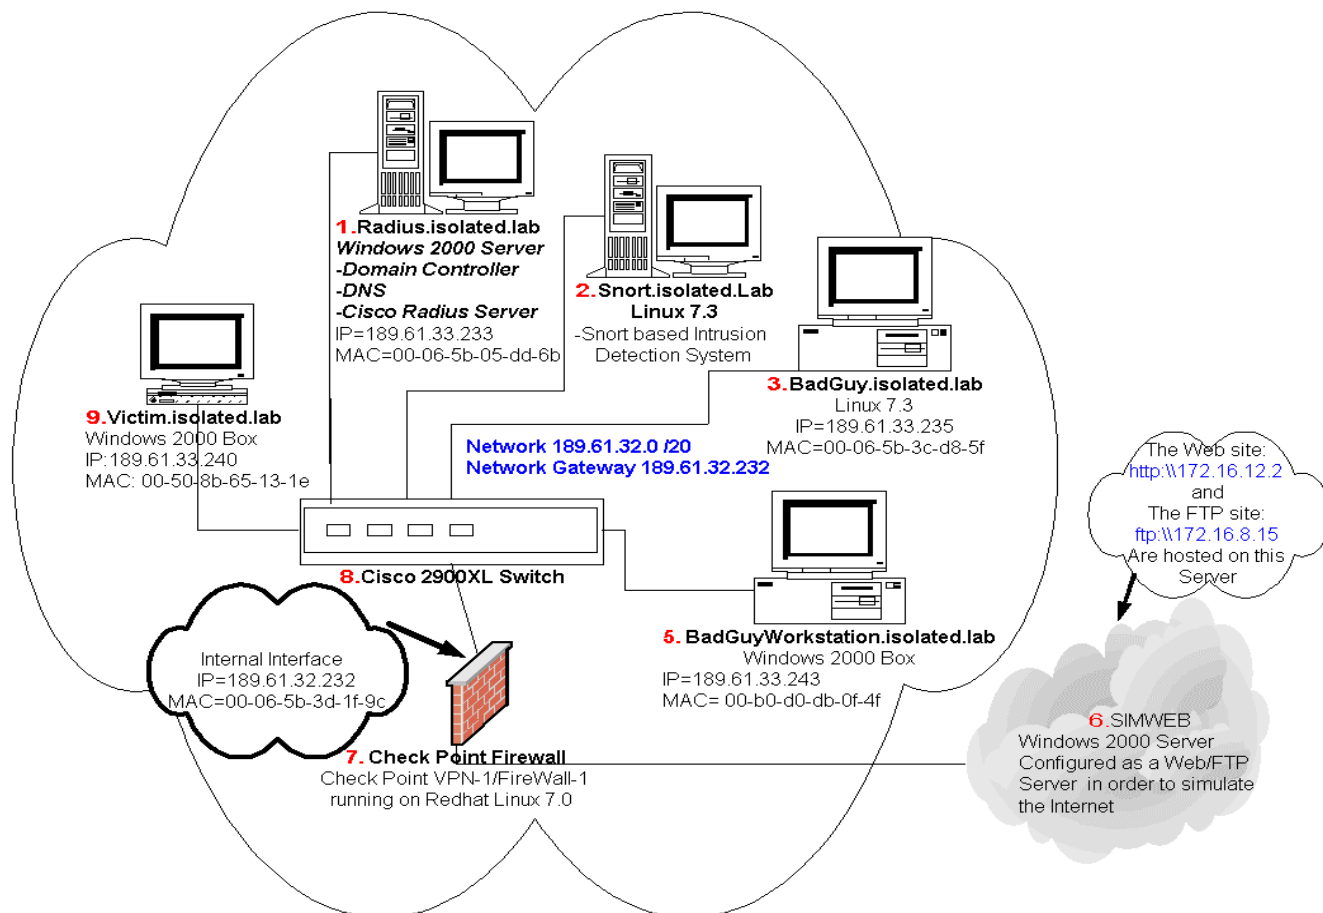


Figure 0-1 Lab-Network

DEVICE	DETAILS
1. Radius.isolated.la b	This Windows 2000 Server (189.61.33.233,MAC 00-06-5b-05-dd-6b) with Service Pack 2 has several functions: <ol style="list-style-type: none"> 1. It acts as the Windows 2000 Active Directory Domain Controller for the Domain “isolated.lab”. 2. It acts as the DNS server for lab network 3. It acts as the host for the Cisco Secure Access Control (ACS) Server v 3.0 for Radius Authentication.
2. Snort.isolated.lab	This Redhat Linux 7.3 server (No IP on interface—only sniffs traffic) serves as the Networks Intrusion Detection System (IDS) monitoring all traffic coming and leaving the local network. The server is running Snort v. 1.77.2.20 and is connected to the Port “FastEthernet 0/1” on the Cisco 2900X switch. The switch has being configured to mirror the port connected to the Check Point Firewall, which is “FastEthernet 0/15”, to port “FastEthernet 0/1”. Therefore, the IDS system can

	<p>monitor all traffic leaving and coming into the network 189.61.31.0/20</p>
3. BadGuy.isolated.lab	<p>This Redhat Linux 7.3 box (189.61.33.235,MAC 00-06-5b-3c-d8-5f) is located in the attacker's Work Cubical. It is the machine that will be used to attack the victim's Windows 2000 Workstation (victim. isolated. lab).</p> <p>The employee is suppose to use this box for testing and training but has decided to use it for malicious reasons.</p>
4. BadGuyWorkstation.isolated.lab	<p>This Windows 2000 Professional Workstation (189.61.33.243,MAC 00-b0-d0-db-of-4f) is used by the attacker as his day to day workstation.</p> <p>The workstation is a member of the Windows 2000 Active Directory domain "isolated. lab". The attacker access's the network via his Windows 2000 User account (attackert).</p>
6. SIMWEB	<p>This Windows 2000 Server with Service Pack 2 has been set up to simulate the Internet. A Crossover Ethernet Cable connects it to the Outside Interface of the Check Point Firewall.</p> <p>This server has being configured to run Internet Information Server (IIS) and several virtual web sites have been configured. During the Exploit, this server will be the Web Server that the victim accesses and unknowably has his NT password sniffed during the Web Authentication process.</p>
7. Check Point Firewall	<p>This Linux Redhat server 7.0 running "Check Point VPN-1/FireWall-1" acts as the firewall between the internal network 189.61.32.0/20 and the simulated internet.</p>
8. Cisco 2900XL Switch	<p>This Cisco Switch connects all of the internal Servers and Workstations together in the Lab. It is the same model that is most commonly used in the regular production network.</p>
9. Victom.isolated.lab	<p>This Windows 2000 Professional Workstation (189.61.33.240,MAC 00-50-8b-65-13-1e) is the Victim's workstation located in his work cubical.</p> <p>The workstation is a member of the Windows 2000 Active Directory domain "isolated. lab". The victim accesses the network via his User account that is part of the isolated. lab domain. The victims username is "victimJ" and his password "BigSky99"</p>

Table 0-1 Details of Lab-Diagram

As the Check Point Firewall configuration (item 7 on Figure 0-1) is what causes the vulnerability that leads to the attacker gaining access to the victim's NT username and password it is worth while to discuss this Firewall's configuration in detail.

Check Point VPN-1/FireWall-1 supports three different types of user authentication: Host Authentication, Session Authentication, and Client Authentication. The security policy of Widget Research specifies that Client Authentication should be used as the Check Point Firewall Authentication method; i.e., the employee is checked by Client Authentication before being allowed to connect to an Internet host.

The process of Client Authentication will be explained in detail later in the paper but for now the following quote from the Check Point VPN-1/FireWall-1 manual summarizes the authentication method nicely. "Client Authentication grants access on a per host basis. Client Authentication allows connections from a specific IP address after successful authentication. It can be used for any number of connections, for any service and the authentication is valid for the length of time defined by the administrator. ...It is best used when the client is a single-user machine, such as a PC." (Check Point Software, page 490.)

The Check Point Firewall in the lab is configured to mimic the configuration of the production Networks firewall. The internal network card IP is 189.61.32.232 and this interface is the default gateway for the network 189.61.32.0/20. The Firewall's external Interface has the IP address 172.16.12.2 and is connected directly to the Windows 2000 server "SIMWEB" (Item 6 on Figure 0-1) by an Ethernet Crossover cable.

The server SIMWEB (Item 6 on Figure 0-1) hosts three virtual web sites using the Internet Information Service, and is used to simulate the Internet within the isolated Lab. The web sites hosted on SIMWEB have the IP addresses 172.16.12.2, 172.16.12.3, and 172.16.12.5. The SIMWEB server also hosts a FTP server with the IP address 172.16.8.15.

When a user inside the lab attempts to connect to either the web sites or the FTP site hosted on SIMWEB, the connection will be routed to the internal interface of the Check Point firewall as this is the default gateway for the internal network. The Firewall will then use Client Authentication before allowing the connection to the external web site or FTP site hosted on the external server, SIMWEB. The process of Client Authentication within the lab consists of the following steps:

1. The Firewall will use the built-in authentication protocols to get the security credentials from the end user. For example, when a user attempts to connect to a web site, that user will encounter a HTML authentication

dialog box that will be presented by the user's web browser as a result of an HTML authentication request from the firewall.

2. The Check Point Firewall would then pass the user's credentials to a Cisco Secure Access Control (ACS) Server v 3.0 for Radius Authentication (Item 1 on Figure 0-1).
3. The ACS server would then communicate with the Windows 2000 Active Directory domain controller (Item 1 on Figure 0-1) to see if the username and password is valid
4. The ACS would then communicate back to the firewall to signify whether the credentials are valid or invalid. Then the Firewall would permit or deny access to the end user to be allowed to connect to the Internet Web Server, FTP Server, and other sites. If the authentication succeeded, Check Point would allow connections to the Internet hosts using the protocol that the user has passed authentication from the IP address of the client's workstation until an idle time-out timer expired.

Protocol Description

The Web Authentication System that the attacker exploits uses several different protocols to authenticate users wishing to access a service on the Internet. The Authentication System has multiple components and utilizes multiple protocols. The main protocols the system uses are:

- 1) The Build in Authentication process of the Protocol that the end user attempts to use to connect to an Internet host, such as HTTP, FTP, etc. The Check Point Firewall will use the built-in authentication process of these protocols in order to gather the end users password and username. To keep things less complex, this paper will concentrate on the HTTP and FTP protocol authentication processes as these protocols are by far the most common protocol used by the Widget Research employees when connecting to Internet hosts.
- 2) The Radius IEEE protocol version 2, is used by the Check Point Firewall to validate the client's credentials as provided to it according to Item 1 on this list.
- 3) Microsoft Windows Kerberos Version 5 Authentication protocol is used by the Radius server to check if the user's credentials it has received from the Check Point Firewall, via the Radius protocol, is a valid username and password in the Windows 2000 Active Directory.

In this espionage incident, the attacker gains the victim's NT password by exploiting the security weakness in the protocol that the victim uses to attempt to connect to an internet host.

When a user opens a new HTTP connection, an opportunity is created for the attacker to intercept the victim's NT password and username as the password will be sent in plain text. On the Check Point firewall there is a rule that makes it mandatory for any connections from an inside machine to a outside machine to be first authenticated by Host Authentication. When a user first opens up a new HTTP connection, the Check Point Firewall will intercept the connection and pass it to the Check Point VPN-1/Firewall-1 HTTP security server. The HTTP security server is configured to contact a Radius server (in this case the Cisco Secure Access Control Server v 3.0) to authenticate the user, but it first needs to collect the username and password from the end user.

To accomplish the preceding, the HTTP security server will send an HTTP "401 response header" back to the end-user's browser. Assuming that the user is running a graphical browser, such as Netscape or Internet explorer, he or she will be presented with a graphical dialog box asking for a username and password. In an attempt to be user friendly, the company policy is to use NT usernames and passwords for Internet authentication, and therefore the end user will input his or her NT username and password in the dialog box. Once the user clicks on the submit button, the browser will send the username and password back to the Check Point VPN-1/Firewall-1 HTTP security server in clear text format. It is at this point in the web authentication process that the attacker has chosen to intercept the victim's NT username and password.

A very similar process is followed when the user attempts to connect to a FTP site hosted on the Internet, This process also provides an opportunity to intercept the NT passwords and usernames in clear text. For FTP sessions, the Check Point Firewall will intercept the initial request and send an FTP authentication request to the source host (the end user's computer that is trying to start the FTP session). For example, if the user starts a session to <ftp://widget.com>, the Check Point firewall will intercept the connection and pass it to the Check Point VPN-1/Firewall-1 FTP security server. The Check Point VPN-1/Firewall-1 FTP security server will send the FTP client a FTP authentication request which will prompt the user for their log in credentials. The user will type in his or her NT username and password which will traverse the network back to the Check Point firewall in clear text providing a chance for the attacker to sniff out the username and password. If the radius server validates the username and password, then the Check Point VPN-1/Firewall-1 FTP security server will allow the connection to continue to <ftp://widget.com> and the user will encounter <ftp://widget.com> authentication request for access to the ftp site. As a side note, in this situation, the attacker gets a bonus in that he can both intercept the end users NT username and password plus he can just as easily intercept the password and the username that the user uses to gain access to the FTP site.

The Radius and Kerberos protocols are the other two main protocols that are used in the Web Authentication System. Although the attacker has decided not to

try to exploit weaknesses in these protocols in this incident, it is worthwhile providing a brief description of them.

Check Point VPN-1/Firewall-1 has been configured to use the Radius IEEE version 2 protocol to validate the security credentials of the end users trying to access Internet web and FTP sites. Once the firewall has obtained the NT username and password via an HTTP or FTP authentication process as was described above, it can then pass validate those credentials by using the Radius protocol to the Cisco Secure Access Control (ACS) Server v 3.0 which has been configured to act as a Radius server.

The Remote Authentication Dial-In User Service (RADIUS) is an authentication and accounting protocol that was developed by Livingston Enterprises (Cisco Systems, Web Document: "How does Radius Work?"). RADIUS is a client/server protocol. In the Widget Research Web Authentication System, the Check Point Firewall is the RADIUS client and the Cisco Secure Access Control (ACS) Server v 3.0 has being configured to run as the RADIUS server. The RADIUS client will encrypt the end users NT password and username and then will pass the encrypted data to the Radius Server using the User Datagram Protocol. The radius server will then check to see if the NT passwords and usernames are valid in the Windows 2000 domain by contacting a Windows 2000 Active Directory Domain controller. The server will then reply to the client using the UDP protocol, to notify if the credentials are valid or not valid. The client will then allow or disallow the connection to the web hosts, he or she wishes to contact based on the response of the RADIUS server.

The Radius protocol encryption makes it inconvenient for the attacker to intercept and decode the usernames and passwords while they travel between the Firewall and the Radius server. The attacker would not likely choose this part of the Web Authentication System to exploit. However, it should be noted that this protocol does have some vulnerabilities that could be exploited in future attacks in order in intercept password and usernames. Joshua Hill discusses possible Radius issues in his paper "An Analyses of the RADIUS Authentication Protocol" found at <http://www.untruth.org/~josh/security/radius/radius-auth.html>. He outlines the following Radius vulnerabilities:

- Response Authenticator Based Shared Secret Attack
- User-Password Attribute Cipher Design Comments
- User-Password Attribute Based Shared Secret Attack
- User-Password Based Password Attack
- Request Authenticator Based Attacks
- Passive User-Password Compromise Through Repeated Request Authenticators
- Active User-Password Compromise through Repeated Request Authenticators
- Replay of Server Responses through Repeated Request Authenticators

- DOS Arising from the Prediction of the Request Authenticator
- Shared Secret Hygiene

Microsoft implementation of Kerberos Version 5 is also a key protocol used in the Web Authentication System. The Cisco Secure Server will use the Kerberos protocol to attempt to log into the Windows 2000 domain with the username and password that it has received from the Check Point Firewall via the Radius protocol. If the Kerberos authentication succeeds then the Cisco secure server advises the Check Point firewall that the user credentials are valid by means of the Radius protocol.

Kerberos is an authentication protocol that is highly secure. It is used as the default authentication protocol in Windows 2000 networks. The Microsoft White Paper document "Windows 2000 Kerberos Authentication" available at <http://www.microsoft.com/windows2000/docs/kerberos.doc>, describes the protocol in detail. The article describes the protocol as an authentication protocol that is based on "authentication technique involving shared secrets" (Microsoft Corporation, p.4). The document summarizes the protocol as follows "The Kerberos authentication protocol provides a mechanism for mutual authentication between a client and a server, or between one server and another, before a network connection is opened between them. The protocol assumes that initial transactions between clients and servers take place on an open network where most computers are not physically secure, and packets traveling along the wire can be monitored and modified at will." (Microsoft Corporation, p.4)

Exploiting the Kerberos protocol would be very difficult for the attacker to accomplish and therefore the attacker did not even consider exploiting this component of the Web Authentication System that uses the Kerberos protocol.

How the exploit works

This exploit works by a combination of three variables: the attacker utilizes his privileges associated with his Windows 2000 network account; his knowledge of networking protocols; and his internal knowledge of the company's Network infrastructures.

The attacker has a valid NT user account for the company's Windows 2000 based network. The organization has deployed a Windows 2000 network using one large domain in which to place all Active Directory objects. A "domain" in Microsoft networking is a grouping of user accounts, computer accounts and other objects in a logical group. By logging into the domain a user is granted privileges to access objects within the domain as configured by the Windows 2000 domain administrators. The attacker uses this privilege to gain surveillance information such as the victim's workstation name and IP address. Further the attacker uses his network privileges to figure out the time frame of when the victim would most likely be using his computer. Once the attacker obtained the

victim's username and password, he used them to log into the domain and in doing so has now the required privileges to access a file that contains the secret document that the attacker desires.

A key part of this incident occurs when the attacker sniffs the victim's username and password off the Network. "Sniffing" involves placing a network interface card into promiscuous mode and then running a "sniffer" program to gather network traffic off of the network. Promiscuous mode is a setting that tells the network card to capture all network traffic that it receives, not only traffic specifically addressed to it. In sniffing, usually the sniff program is configured with filters in order to filter out what traffic is being monitored.

Network sniffing is a relatively easy task in "hub" based networks. Historically, Ethernet networks were deployed using "hubs". A hub is a device that offers multiple ports in which network devices, such as servers and workstations are interconnected by network cabling. The hub acts as a multi-port repeater in that whatever traffic enters a port is repeated throughout all the other ports on the hub. In a hub-based network, all of traffic in the network traverses every section of the network cabling. Therefore, sniffing in a hub-based network is simple as one only has to plug the sniffer into any module of the hub in order to gain access to all traffic being produced on the network.

Network sniffing is much less useful in a Micro-switched network such as the one Widget Research has deployed. A switch is a more intelligent network than a hub in that when a network packet is sent to a hub the hub will send the packet out all of its ports, whereas a switch will only retransmit the packet out of the port that the destination network device is connected to. A switch will learn the MAC addresses of network devices attached to it and then will build a dynamic table that is used to keep track of what network devices are connected to what switch ports. In a Micro-switched network each network device is connected to a switch port by a point-to-point dedicated link (e.g., every workstation and server is directly connected to a switch port by means of network cable). In a Micro-switched network only the traffic addressed to network devices that are connected to a given port will be forwarded out the port, with the exception of network traffic being broadcasted or multicasted to multiple destinations. Therefore, a network sniffer is not able to capture network information being sent to a network device that is not connected to the same switch port that the sniffer is immediately connected to.

To get around the difficulties in sniffing traffic in a Micro-switched network, the attacker launched a "Man in the Middle" attack in order to intercept and sniff the network traffic being sent from the victim's workstation machine to the Check Point Firewall. The internal attacker knew that the traffic between the victim's machine and the organization's firewall would at some point contain the victim's NT username and password in plain text (not encrypted). The attacker knew this as he knows the organization uses a Web Authentication System that uses the

employee's NT username and password as the basis of Internet authentication, and that this system involves passing the username and password from the employee's workstation to the firewall in plain text. In this incident, the attacker used an ARP spoofing attack against the victim's workstation so that the workstation would send all traffic destined for the firewall directly to the attacker's Linux 7.3 box. This procedure allows the attacker to sniff all of the traffic that the victim's machine was sending to the firewall including the username and password. The details in how the attacker implemented this attack will now be discussed.

The attacker uses a program called "arp spoof" to send "Gratuitous ARP's to the victim's workstation in order to manipulate the ARP table of the victim's workstation. The airproof program is part of the Dsniff Active Sniffing suite written by Doug Song.

In order to explain how the airproof tool works it is first necessary to outline some of the workings of the IP protocol. In TCP/IP based networks each workstation is assigned a logical IP address either statically or through the DHCP protocol. The address assigned will fall within the Network range of IP addresses assigned to the network. For example, in the Widget Research network the IP address of the Check Point Firewall is 189.61.33.232/20 and it is part of the IP network 189.61.32.0/20. However when data is actually transmitted through the network, it is transmitted by using the sender's and destination's physical MAC addresses and not the logical IP address of the network devices. The physical MAC address is a 48 bit address that is coded into the Network Interface card of the network device. This address is unique in that all Network Interface cards manufactured though out the world are given a globally unique address.

When a TCP based network host wants to send information to a another network host, it most first figure out what the physical MAC address is that is associated with the destination's logical IP address. This task is accomplished by the ARP protocol. The ARP protocol works by broadcasting a IP address to all network hosts located on the local network and then having the workstation that is assigned that IP address respond with it's Mac address. For example, if the victim's workstation, that has being assigned the IP address 189.61.33.240, wishes to send information to the Check Point Firewall that has the IP address 189.61.33.232, it most first know the physical MAC address of the Check Point Firewall interface card. As a result, the victim's machine will send out an ARP request that says "Whoever has been assigned the IP address 189.61.33.232, please send me your MAC address". The Check Point Firewall will respond back to the victim's machine with an ARP response that says "I am the one with the IP 189.61.33.232 and my MAC address is '00-06-5b-3b-1f-9c". In order to save time, the victim's workstation will catch the IP to Mac association in its ARP catch table so that the ARP procedure does not have to be repeated in future communications with the firewall.

TCP/IP based network hosts will also accept unsolicited “Gratuitous APRs” as part of resolving the IP address to MAC address process. A Gratuitous ARP is an ARP response that was not asked for. Even though the ARP reply was not requested the network host will still make an entry in its ARP catch table.

In this attack, the attacker was able to use Gratuitous APRs to poison the victim’s ARP catch table. The attacker basically sent a Gratuitous APR to the victim’s workstation saying that “my IP address is 189.61.33.232 and my Mac address is “00-06-5b-3c-d8-5f”. As a result of this Gratuitous ARP the victim’s workstation made an entry in it’s ARP catch that the Mac address for the IP 189.61.33.232 is 00-06-5b-3c-d8-5f and as a result sent all traffic destined to 189.61.33.232 to the Network interface card 00-06-5b-3c-d8-5f. However, the Mac address 00-06-5b-3c-d8-5f is not the Mac address of the interface of the Check Point Firewall with the IP 189.61.33.232, rather it is the Mac address of the attacker’s Linux 7.3 box with an IP address of 189.61.33.235. The attacker uses the “arp spoof” command to send the Gratuitous APRs to the victim’s machine. (This tool will be explained in detail in the “Description and diagram of the attack” section). Hence, due to a poisoned ARP catch, the victim’s workstation sends all of the traffic destined to the firewall to the wrong network interface.

The attacker configuration his Linux box to do “IP Forwarding” so that all of the network traffic it was intercepting from the victim’s machine would be forwarded back to Check Point firewall, which was the victim’s workstation’s intended recipient of the data. This step is necessary as without IP forwarding configured all communication between the victim’s workstation and the firewall would halt as the two network hosts could not communicate with each other. As far as the victim is concerned, nothing unusual will be noticed but in fact all of the traffic is being routed through a third-party Linux box that is being configured to sniff all of the traffic.

The attacker used the Dsniff sniffer to filter out any passwords that it found in the network traffic that the victim’s machine was sending to the firewall. The Dsniff utility can be configured to just filter the traffic for passwords and usernames and then output this information when found. As will be detailed in “Description and diagram of the attack”, the attack was successful in sniffing out the victim’s NT username and password as it was passed to the firewall in a clear text HTML 401 response header authentication response.

Description and diagram of the attack.

This attack consisted of three stages, 1) surveillance, 2) obtaining the victims NT password, 3) using the password to gain access to the secret document held on the user’s Network drive. The three steps of this attack will first be summarized and then detailed:

Summary of the Attack

- 1) Stage one of the attack—surveillance. (In this stage of the attack, the attacker using Microsoft Windows protocols and utilities to discover the victim's workstation IP address. As well, he observed the victim and his computer use habits.)
- 2) Stage two of the attack, using Dsniff tools to obtain the user's NT password. (In this stage of the attack ,the attacker uses the information he gathered in Stage one of the attack in order to use Dsniff to obtain the NT username and password of the victim.)
- 3) Stage three of the attack, Steal the document by using the victim's username and password. (In this stage of the attack, the attacker uses the victim's username and password that was discovered in stage 2 of the attack to gain access to the secret document held on the victim's network drive.)

Details of the Attack

1. Stage one of the attack—surveillance

In order for the attack to be successful the attacker needs to know some preliminary information. The attacker must know a minimal amount of information about the victim's computer use such as the victim's working hours and when the victim most likely will be accessing the Internet. Further, the attacker needs to know the IP address of the victim's workstation. This information is easily obtainable for the attacker as the attacker is a technical employee of the organization and therefore can use a combination of social engineering and technical skills to obtain the perquisite information needed to conduct the attack.

The attacker obtains the victim's work habits by viewing the victim's Calendar. The organization uses GroupWise as it's Electronic Messaging system and utilizes the Calendar and Appointment features of GroupWise as the primary tool to schedule appointments and meetings amongst employees. When an employee wishes to schedule a meeting or an appointment, he or she will use GroupWise to send out the meeting invitation. When employees accept an invitation for a meeting or appointment, GroupWise automatically updates the employees' GroupWise Calendar. Therefore, each employee's GroupWise Calendar details a schedule of all of the employees' appointments and meetings. Further, the Calendar also displays the employees' working hours and vacations. The GroupWise system has being configured so that employees can view the Calendar of other employees in order to assist employees in scheduling appointment that do not have time conflicts. The attacker prefers to do the attack on a day that the victim will be spending a lot of time working in his office on his workstation, rather than a day that the victim will be spending most of the day at appointments and meetings. Therefore, the attacker will simply log into his GroupWise account and view the victim's Calendar and decide on the optimum

day to perform the attack. In this case, the attacker uses the busy search feature in GroupWise to determine that on January 13, 2003 the victim has only one appointment. As a result, the attacker decides that January 13, 2003 would be a good day to attempt the attack as the victim will most likely be spending a lot of time on his workstation that day. Figure 0-2 is a screen shot of the victim's Calendar that was obtained by the attacker by using the busy search feature that is available in the GroupWise Client program.

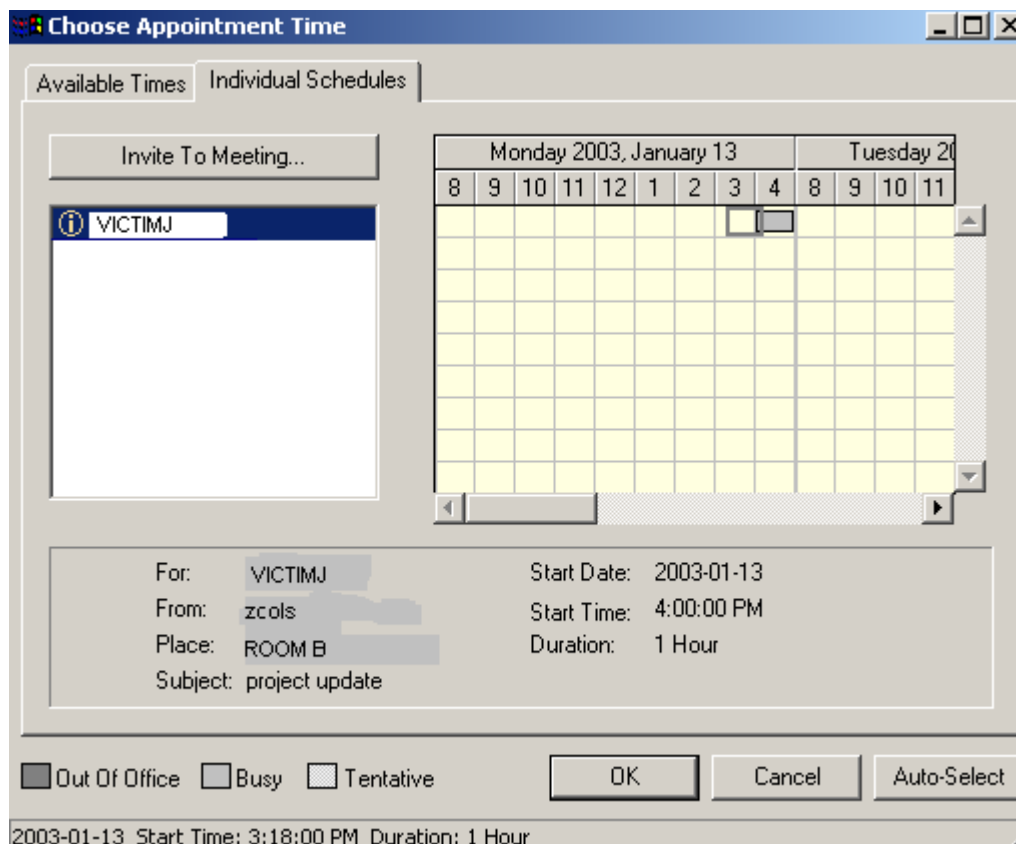


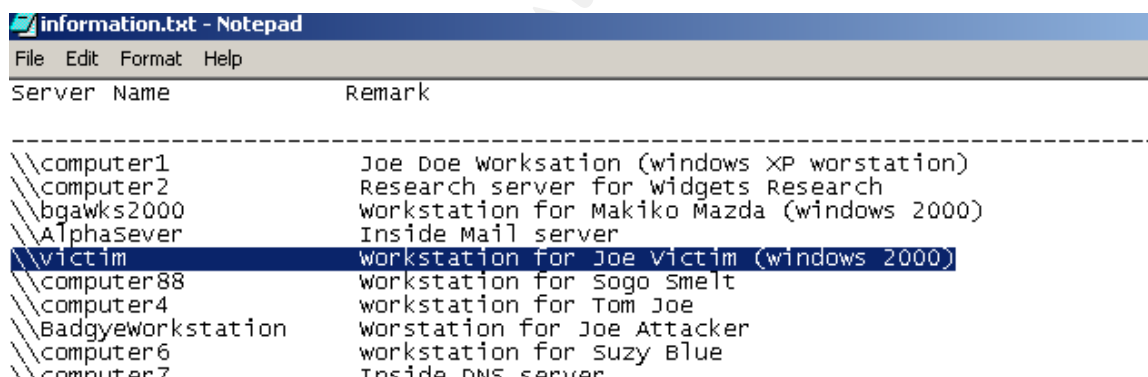
Figure 0-2 A screen shot of the Victim's work schedule, obtained by the attacker by using the "Busy search" feature of the attacker's GroupWise Client.

The attacker uses surveillance to gather additional information about the victim's work habits. Occasionally, the attacker will casually walk by the victim's work cubical in order to obtain information on the victim's work habits, such as; is there a favorite time for the victim to do research on the internet, what time does the victim take a break, does the victim always keep his computer powered on, etc.

In order to carry out the password sniffing part of the attack, the attacker needs the IP address of the victim's workstation. This information is obtained using the attacker's knowledge of the network and by utilizing some of the command line Utilities that are packaged with the Attacks Windows 2000 workstation. The company's usual procedure for names for Windows 2000 accounts is to use the employee's last name followed by the first letter of the employee's first name.

Further, the company’s policy dictates that all workstations within the Active Directory must have a description in order to simplify administration. In our example, the victim’s name is Joe Victim and his Windows 2000 workstation name is Victim.isolated.lab. The description for the workstation in the Windows 2000 Active Directory is “workstation for Joe victim (windows 2000)”. The attacker is aware of these policies, due to his position as a technical employee in the company. He uses this knowledge to figure out the IP address of the victim’s workstation.

The procedure the attacker used to get the IP address of the victim in this attack is as follows: The attacker uses the “net.exe” program that is bundled with Windows N/T and Windows 2000 on his Windows 2000 workstation to get a list of the workstations and their descriptions. The command “net.exe view > information.txt” will pipe a list of the workstation’s names and descriptions within the NT domain into the text file “information.txt”. The results are piped into information.txt to make it more convenient to locate the required information, as the list is rather long due to the large number of workstations. The attacker then simply scans the workstation descriptions in order to determine the name of the victim’s workstation. Below (Figure 0-3) is a sample of the output of the file information.txt that was produced by the command “net.exe view > information.txt”. The victim’s workstation is highlighted in the sample output.



```

information.txt - Notepad
File Edit Format Help
-----
Server Name          Remark
-----
\\computer1          Joe Doe workstation (windows XP workstation)
\\computer2          Research server for widgets Research
\\bgawks2000         workstation for Makiko Mazda (windows 2000)
\\AlphaServer        Inside Mail server
\\victim              workstation for Joe Victim (windows 2000)
\\computer88         workstation for Sogo Smelt
\\computer4          workstation for Tom Joe
\\Badgyeworkstation  workstation for Joe Attacker
\\computer6          workstation for Suzy Blue
\\computer7          Inside DNS server
  
```

Figure 0-3 Sample output from information.txt—a file generated by “net.exe view> information.txt” command on the attacker’s Windows 2000 workstation.

Now that the attacker knows the victim’s computer name, the attacker can now use this name to find the victim’s computer’s IP address. This can be done in several ways (e.g., nslookup, etc), but in this case the attacker issued the command “ping Victim” which returned the IP address of the victim’s machine. The ping command is used to test if a host is reachable by sending ICMP packets to the host computer. By typing, “ping Victim” the attacker’s Windows 2000 workstation will use name resolution to convert the computer name into an

IP address and then will send the ICMP packets. The output of the Ping command includes the IP address of the victim's workstation, which is 189.61.33.240. Even if the victim's workstation is powered down the attacker will still get the IP address. The ping command will contact the network Dynamic DNS server by automatically appending the "Connection-specific DNS Suffix", which is "isolated.lab", to the workstation name in order to make it possible to do a DNS query on the full qualified name ("victim.isolated.lab"). Therefore, whether or not the ICMP queries succeed or fail, the ping command output will display the IP address of the victim's computer, and this is demonstrated in Figure 0-4. The attacker now has all of the prerequisite information needed to perform the attack.

A) PING RESULTS WHEN VICTIM'S WORKSTATION IS POWERED ON.

Pinging victim.isolated.Lab [189.61.33.240] with 32 bytes of data:

```
Reply from 189.61.33.240: bytes=32 time<10ms TTL=128
Reply from 189.61.33.240: bytes=32 time<10ms TTL=128
Reply from 189.61.33.240: bytes=32 time<10ms TTL=128
Reply from 189.61.33.240: bytes=32 time<10ms TTL=128
```

Ping statistics for 189.61.33.240:

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

B) PING RESULTS WHEN VICTIM'S WORKSTATION IS POWERED OFF.

```
C:\>ping victim
```

Pinging victim.isolated.Lab [189.61.33.240] with 32 bytes of data:

```
Request timed out.
Request timed out.
Request timed out.
Request timed out.
```

Ping statistics for 189.61.33.240:

```
Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

Figure 0-4 Output of the "ping victim" command as executed on the attacker's Windows 2000 Workstation (the victim's IP address has being underlined).

Stage two of the attack, using Dsniff tools to obtain the victim's NT password.

As discussed in the section "How The Exploit works", the company's network is micro-switched and as a result it is necessary for the attacker to use the tool Dsniff to sniff out the victims NT password. In a micro-switched network each Network device, such as a workstation or a server, is attached directly to a

Network Switch port. Packets are switched from the source port to the destination port using a dynamically built switching table based on Mac addresses. Unlike a hub, a switch will forward a Mac address only out the port to where the destination host is attached, assuming that an entry for that destination's Mac-address exists in the switch's dynamically built switching table. Therefore, it is not possible to sniff all traffic in a switched network by simply plugging a sniffer into a switch port. As the sniffer will only see the traffic destined for specific ports. In order for the attacker to sniff the traffic between the victim and the Check Point Firewall, he must use a tool called "Dsniff". The Dsniff tool will allow the attacker to use Gratuitous APRs in order to intercept the traffic between the victim and the firewall in order to gain the password. Hence, the attacker has decided to use a "man in the middle" attack to intercept the NT username and password needed to access the secret documents on the victim's Network Drive.

The steps that the attacker used to intercept the password, via the Dsniff Tool, are as follows:

1. Install the Dsniff tool:

The attacker uses his home computer to download the Dsniff and burns it onto a CD. Further, the attacker also downloads the required Redhat Linux modules that are needed by the Dsniff tool. The attacker does the download on his own machine to avoid having the download logged by the company's firewall log files. The tool, Dsniff, is available for download at <http://www.monkey.org/~dugsong/dsniff>.

Like most technical workers at Widget Research, the attacker has a Windows 2000 Workstation (item 5 on Figure 0-1) and a Linux 7.3 (item 3 on Figure 0-1) box in his office. The Linux 7.3 box is supposed to be for research and administration but in this case the box will be used by the attacker to do the password sniffing. The attacker decides to install the Dsniff tool on his Linux 7.3 box in the evening time as there is less chance of a co-worker dropping by and taking an interest in to what the attacker is installing. The attacker has 24 hour access to his work cubical. The attacker uses the information in the Dsniff Read me files (available at "<http://www.monkey.org/~dugsong/dsniff/faq.html>") and the document "Introduction to dsniff " by Lora Danielle (available at "<http://www.sans.org/rr/audit/dsniff.php>") to install the tool.

2. Configure the 7.3 box as a router:

The attacker follows the direction in the document "IP Forwarding", by Steve Litt (available at http://www.troubleshooters.com/linux/ip_fwd.htm), to configure his Linux box as a router.

3. Use gratuitous APRs to poison the ARP table of the victim's machine so that the attacker can intercept the data traveling from the victim's machine to the Check Point Firewall:

On the day of the attack, the attacker uses the information he gained in the surveillance stage of the attack. The attacker launches the program "Arpspoof" on his Linux 7.3 box. Arpspoof, which is part of the Dsniff utility, uses Gratuitous APRs to poison the victim's ARP table. The attacker will run the Arpspoof as a background job on his Linux box and will run the job until the victim's password is successfully captured. The opportunity to capture the password will occur when the victim makes an initial connection to a web site or an FTP server because this is when the Check Point Firewall will use a HTTP response header to authenticate the victim as part of the Check Point Client Authentication process. As the Check Point Firewall will only ask for the password and username once per connection per protocol the attacker will have to continuously run the program in order for it to be running when the victim makes an initial connection. The victim will be completely unaware that all of his traffic to the Firewall is being directed through the attacker's computer as the tool does not cause any side effects on the victim's computer.

The actual command that the victim uses to start the Arpspoof job in the background of his Linux box is depicted and explained in Figure 0-5. As shown, the Arpspoof tool requires the IP address of the victim's workstation and this information is known by the attacker as a result of the attacker's surveillance activities.

© SANS Institute 2003,

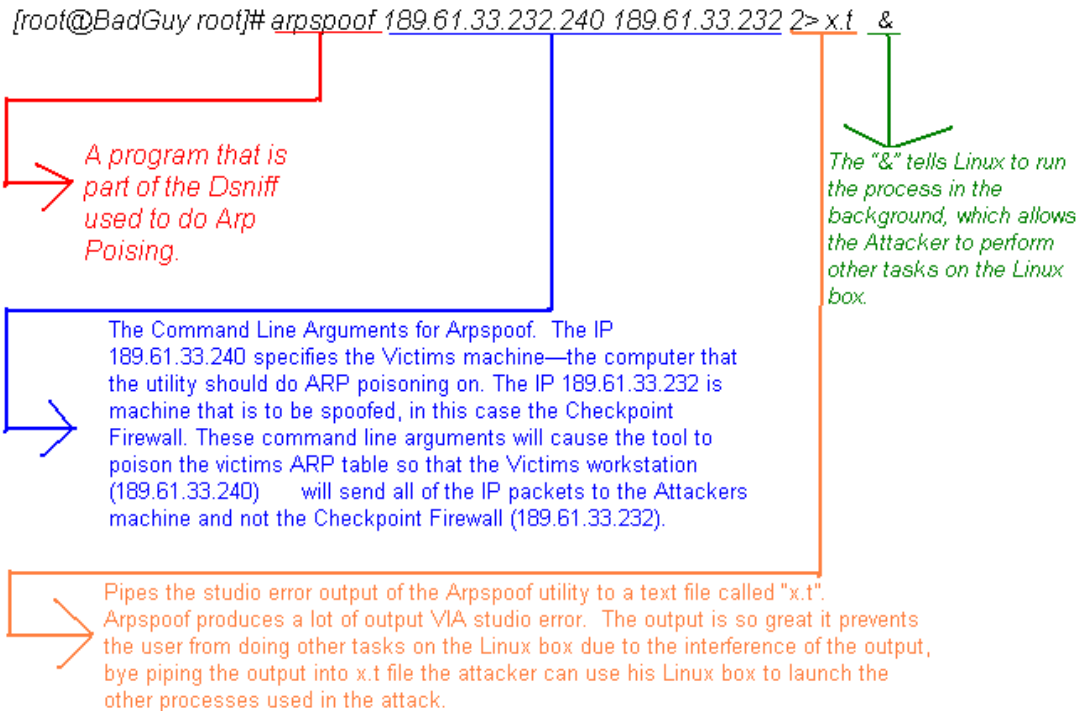


Figure 0-5 The command the attacker used to Launch the Arpspoof process on his Linux 7.3 box.

The attacker starts a "Dsniff -c" process in order to parse out the passwords from the traffic it is intercepting from the victim's machine. All traffic between the victim's machine and the Check Point Firewall is now being intercepted by the attacker's Linux box. The victim is most interested in the network traffic that contains passwords and usernames. To filter this information out the attacker executes the command "Dsniff -c > passwords.txt &" on his Linux box. This launches the program Dsniff, which will filter out passwords from the sniffed traffic when launched with the "-c" command line argument. As the attacker does not want to look at his screen for long periods of time waiting for the victim to do a web authentication, the attacker pipes the output of the Dsniff into the file "passwords.txt" and runs the process in the background by using the "&" at the end of the command. The attacker will then occasionally view the text file "Passwords.txt" to see if Dsniff has captured the victim's username and password.

After running the Dsniff process and the Arpspoof program for about one hour the attacker succeeds in discovering the victim's NT username and password. The attacker views the password.txt file by typing in the command "cat passwords.txt" and is pleased to see that Dsniff has parsed out the victims NT username and password. The attacker kills the two processes (Dsniff and Arpspoof) by using the "ps kill <process id>" command to immediately terminate the ARP processing of the victim's machine in order to reduce the chance of

being detected. The victim has connected to an external site and as a result has used his NT username and password to authenticate himself to the Check Point Firewall as part of the host authentication configured on the firewall. The Dsniff sniffer has parsed the HTTP authentication response from the victim's computer to the Check Point Firewall and has outputted the username and password to the "passwords.txt" file. The attacker is pleased to discover that he has received a bonus in that the victim has logged into a external FTP server and Dsniff has captured the victims username and password for that server. Hence, the attacker now has the user NT username and password and also has the username and password that the victim uses to log into the external FTP site.

When this attack was conducted in the lab, depicted in Figure 0-01, the victim initially opened up his internet explorer and attempted to connect to the web site "http://172.16.12.2" which is a web site hosted on the lab's simulated Internet (item 6 on Figure 0-1). As this is the initial connection, using the http protocol, to a Internet host the Firewall intercepted the connection and initiated the host authentication process. The Check Point Firewalls HTTP security server then sent a HTTP "401 response header" back to the victim's browser in order to collect the user's password and username for authentication before allowing the connection to the remote Internet web site. In response to the "401 response header" package the browser presented the victim with a Graphical box in order for username and password credentials to be inputted, this box is shown in Figure 0-6.

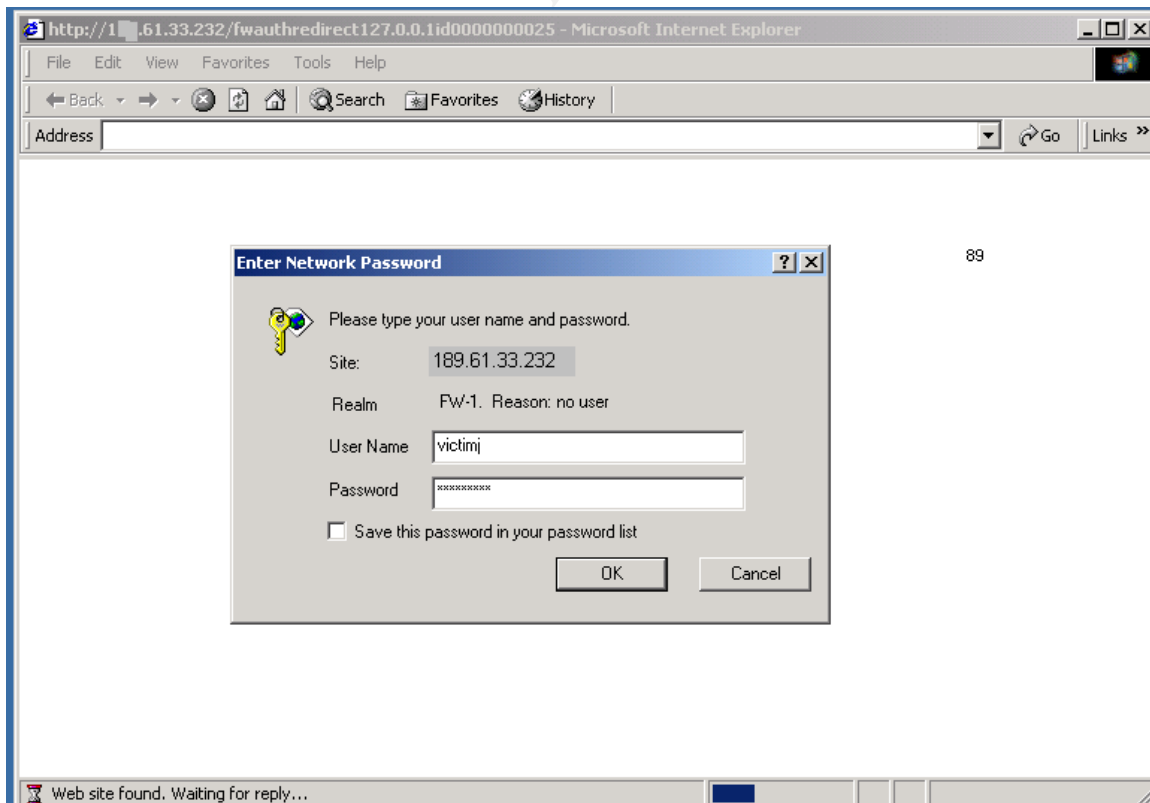


Figure 0-6 The victim's browser presenting him with a dialog box to enter his Log in Credential as a result of a "HTTP 401 response header" from the Check Point HTTP security server.

The victim inputted his NT user name (victimj) and password (FishBay66) into the "Enter Network Password" dialogue box and pressed "OK". The browser then sent the victim's credential back to the server requesting the "HTTP response header", which was the Check Point Firewall HTTP security server. The Check Point firewall then validated victimj password by means of the Cisco based Radius server (item 1 Figure 0-1) to validate the Username and password protocol. After the Cisco Radius server successfully logs into the Windows 2000 domain "isolated.lab" using victimj and the password "FishBay66" the Cisco server responded to the Check Point Firewall Radius authentication request that the log in Credentials of victimj are valid. The Check Point Firewall will then allow the connection from the victims workstation to the web-simulated Internet site "http:\\172.16.12.2".

After the victim finishes with the web site "http:\\172.16.12.2", he uses internet explorer to go to the FTP site "172.16.8.15 " which is a web site hosted on the simulated internet. The Check Point Firewall intercepts this connection attempt and passes it to the Check Point FTP server which sends back a FTP authentication request to the FTP client (in this case it is Internet Explorer's web browser built in FTP client). The victim then uses his NT user name (victimj) and password ("FishBay66") to pass the Check Point Host Authentication procedure. Once the victim passes the Authentication procedure the Check Point Firewall allows the connection to the remote FTP site "ftp::\\172.16.12.2". The remote FTP site then does a FTP authentication to validate the victim. The victim uses the username "victimFTP" and password "tree", which are the log in credentials provided to the victim by the FTP server administrator of the remote FTP site.

During the authentication procedure described above all traffic was being routed to the firewall via the attacker's machine (item 3 on Figure 0-1), due to ARP poisoning. As a result, the Dsniff tool running on the Attacker's computer succeeded in parsing the username and password of the victim from the traffic being sent to the firewall. The IP address of the inside interface of the firewall is 189.61.33.232 and has a Mac-Address of 00-06-5b-3d-1f-9c. However, due to the ARP poisoning the victim's ARP table entry for the IP 189.61.33.232 is now 00-06-5b-3c-d8-5f which is the Mac address of the attacker's Linux machine 189.61.33.235. The victim's workstation sends all of the traffic to the attackers Linux box's Network card giving the opportunity for all of the traffic to be sniffed by the attacker before being rerouted, via IP Forwarding, to the intended destination of the data which is the Check Point Firewall server. During the time that the attacker poisoned the victim's workstation ARP table, the attacker was able to sniff the victim's NT username and password twice, once when the user

authenticated himself to the firewall when attempting to connect to the web site "http://172.16.12.2", and a second time when the victim attempted to connect to the FTP server "172.16.8.15". Further, the attacker gained the victim's FTP username and password for the FTP server "172.16.8.15" when the victim logged into this server. A diagram of the ARP spoofing stage of the attack can be seen in Diagram 0-1 and Diagram 0-2. Diagram 0-1 shows the network traffic prior to the launch of the attack and Diagram 0-2 shows the network traffic during the attack.

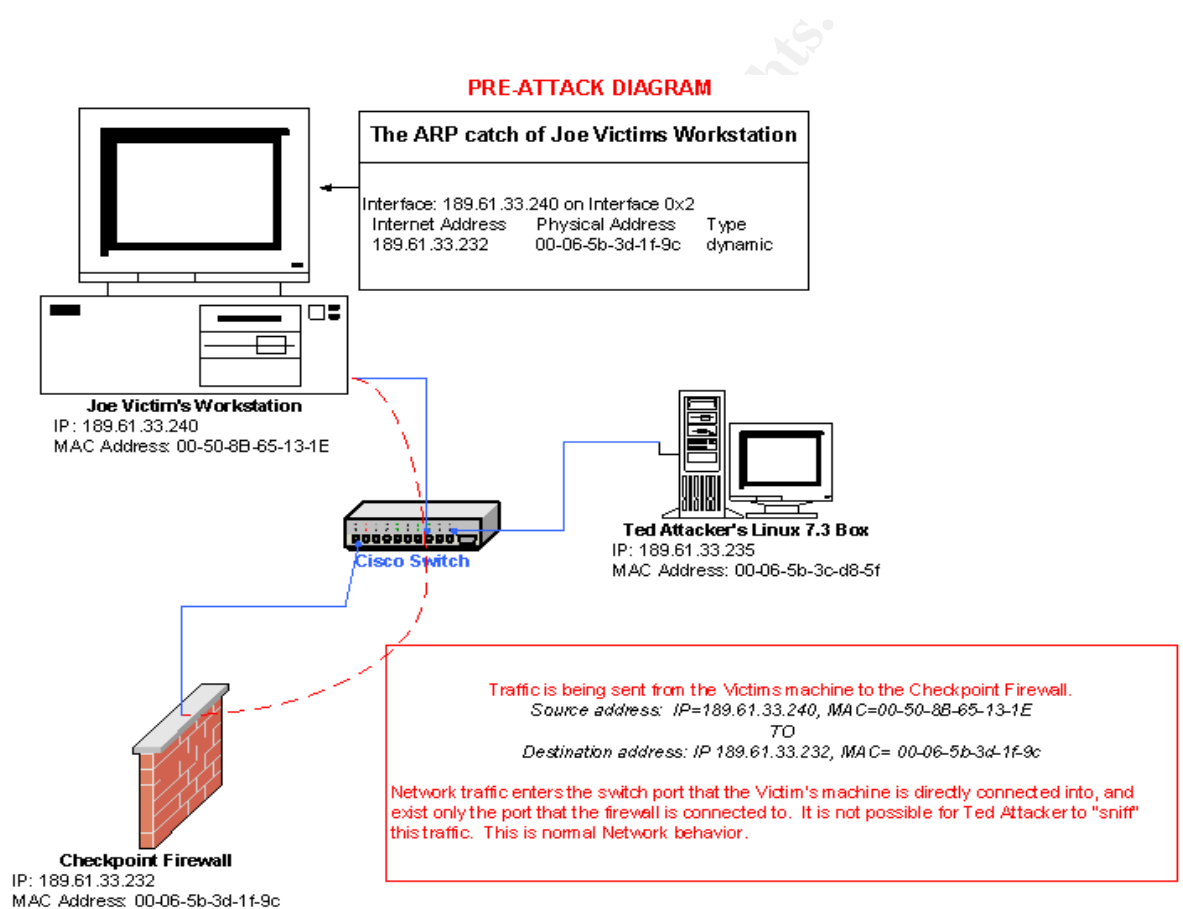


Diagram 0-1 A network Diagram showing normal network traffic—this is a “before the ARP spoofing attack” network diagram.

© SANS




Diagram 0-2 A network Diagram showing network traffic during the time the ARP Spoofing attack is occurring.

4. The attacker uses the victim's NT username and password to obtain a copy of the secret document held on the victim's network drive:

In this stage of the attack, the attacker uses his Windows 2000 Workstation (item 5 on figure 0-1) to access the victim's network drive, which resides on the domains PDC (item 1 on figure 0-1). The attacker is able to do this as he is now in possession of the victim's NT networking username and password.

Widget Research has configured their Windows 2000 network so that users are each assigned a network drive that is hosted on a Network file server. All employees are to keep all of their data on this network drive to simplify data backup procedures. The network drive is mapped when users log into their Windows 2000 machines as this is how the user accounts have being configured.

The attacker could simply log into his Windows 2000 machine using the victim's username and password and then copy the secret document from the victim's

network drive. However, the attacker prefers to be more subtle about stealing the information to minimize the chances of being detected. For example, logging in to the victim's account would create log entries in the server logs and create a local profile for the account on the attacker's machine.

The attacker has decided that the best way of stealing the information, in terms of minimizing the chances of being caught, is to log into the network using his own Windows 2000 account and then map drive to the victim's network drive using the victim's credentials. The steps that the attacker takes to accomplish this task were reproduced in the laboratory: The attacker logged into his Windows 2000 workstation using his own username (attacker) and password (victimj). The attacker right clicks on the "my computer" icon on his desktop and chooses "Map Network Drive" from the pop-up menu (see Figure 0-7). This results in the Graphical Interface for "Map Network Drive" appearing.




Figure 0-7 The attacker launches the "Map Network Drive" on his Windows 2000 Workstation.

The attacker knows that all the employees' network drives are contained in a file share called "home" that is hosted on the server "radius" (item 1 on Figure 0-1). The attacker also knows that the naming standard for network drive shares are the same as the users account name. Therefore the attacker maps a drive to the folder [\\radius\home\victimj\](#). As the attacker has no access rights to access this folder he must use the "connect using a different user name" option and then input the victim's username ("victimj") and password ("FishBay66") into the

“connect AS” Dialog box. Figure 0-8 depicts the attacker mapping a drive to the victim’s network drive.

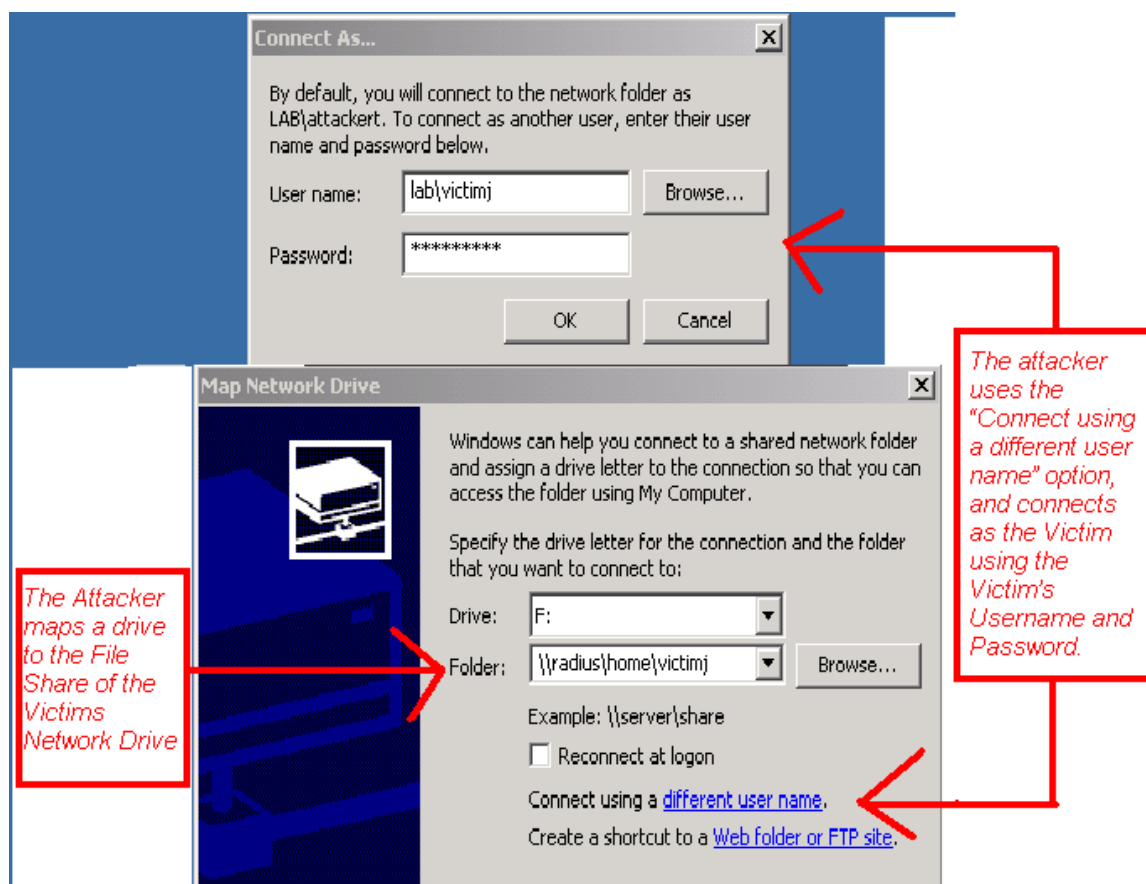


Figure 0-8 The attacker mapping a drive to the victim’s Home Share using the victim’s username and password to gain access.

Now that the attacker has mapped a drive to the victim’s Network Drive, he can look through the contents of the drive until he finds the secret document. In the reproduction of the attack in the lab environment the attacker is looking for the document “A top Secret Document.txt” and he finds this document under the folder “project DOCS” (See Figure 0-9). The attacker then makes a copy of this document onto a floppy disk as this is the easiest way to transport the document out of the building without being detected. In this case, the attacker right clicked the file and used the “send to” feature to send a copy of the file to his floppy disk. Further, the attacker knows that by doing a copy of the file, the file accessed and modified flags, will not be changed and therefore his chances of detection is minimized.

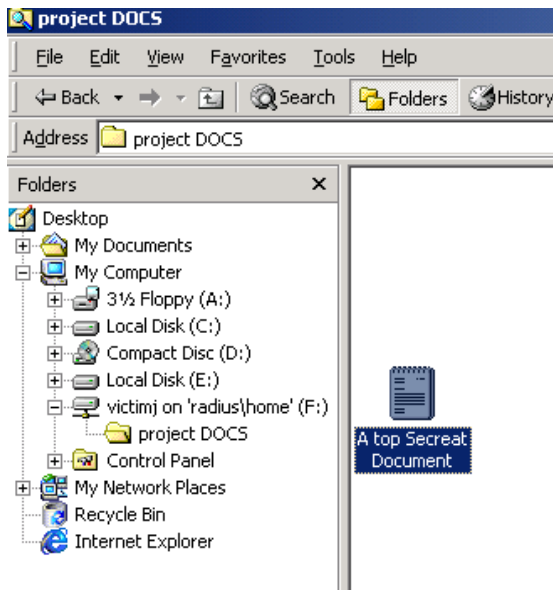


Figure 0-9 The attacker finds the secret document “A top Secret Document.txt” under the folder “project Docs” on the victim’s network Drive.

The attacker gives the secret document to a competitor of Widget Research and as a result seriously harms Widget Research. The attacker has now gained the revenge he was looking for against the company and in the mean time has being rewarded by the competitor for providing the document to them. The attack is now completed.

Signature of the Attack

Although the attacker has tried to be as careful as possible in not being caught, there are a few options that the security staff could use to detect the attack while it is happening or after the attack is finished. These options include:

Noticing abnormalities in the victim’s Windows 2000 ARP table

While an ARP poisoning attack is being conducted the most blatant evidence of the attack is found by viewing the Victims ARP table. The ARP table serves as a catch of what Mac address are associated with what IP addresses. The ARP table is efficient in that the workstation does not have to resolve an IP address to a MAC address every time it wishes to send a IP packet to the host. If it is noticed that an ARP table has the wrong Mac address for a given IP then someone is most likely doing an ARP spoofing attack. Further, it is often the case in ARP poisoning attacks that two IP addresses are associated with

identical MAC-address-tables and this also acts as evidence of an ARP poisoning attack in progress.

To clarify the above it is worthwhile to analyze the before and after ARP tables of the victim's workstation from the attack that was conducted in the test lab (item 9 on Figure 0-1). The output of the ARP tables was produced by the command "ARP -a" on the workstation command prompt.

THE "BEFORE" ARP TABLE OF THE WINDOWS 2000 WORKSTATION
"Victim.isolated.lab" PRIOR TO LAUNCHING THE ARP SPOOFING ATTACK:

```
Interface: 189.61.33.240 on Interface 0x2
Internet Address  Physical Address  Type
189.61.33.232    00-06-5b-3d-1f-9c  dynamic
189.61.33.233    00-06-5b-05-dd-6b  dynamic
```

It should be noted that in the "BEFORE" ARP table, the physical interface for the IP 189.61.33.232 switch is the Check Point Firewall (item 7 Figure 0-1). It is set to "00-06-5b-3d-1f-9c" which is the correct Mac address of the firewall's network interface card. The entry 189.61.33.233 is the DNS server and the Domain Controller for the network and, as the workstation is in communication with this server, there is an entry in the ARP table.

THE "After" ARP TABLE OF THE WINDOWS 2000 WORKSTATION
"Victim.isolated.lab" WHILE THE ARP SPOOFING ATTACK WAS BEING
CONDUCTED:

```
Interface: 189.61.33.240 on Interface 0x2
Internet Address  Physical Address  Type
189.61.33.232    00-06-5b-3c-d8-5f  dynamic
189.61.33.233    00-06-5b-05-dd-6b  dynamic
189.61.33.235    00-06-5b-3c-d8-5f  dynamic
```

The "after" ARP table of the victim's workstation indicates that the incorrect MAC address is associated with the Check Point Firewall network interface. In fact the network interface of the firewall is associated with the MAC address of the attacker's Linux 7.3 box (item 3-Figure 0-1), which is "00-06-5b-3c-d8-5f". Further, it can be seen that identical MAC addresses exist for both the IP 189.61.33.232 (the Firewall) and for 189.61.33.235 (the attacker's Linux 7.3 box). As it is impossible to have the same MAC address on two separate network cards this is a decisive clue that someone is running a ARP poisoning tool on the network. The attacker's IP address 189.61.33.235 is in the victim's ARP table as the Dsniff utility "Arpspoof" spoofing system involves IP communication between

the attacker's computer and the victim's workstation, which causes the ARP table entry.

Using abnormalities in the victim's ARP table to detect the attack would only be useful if the security staff had prior intelligence that some one was going to do the attack. It would be very unusual for a technical or a non-technical worker to be monitoring their ARP tables while they are being victimized by a "man in the middle" ARP spoofing attack. However, if security expects that the victim will be attacked at a particular time, ARP table monitoring would be a way of identifying the attack and source of the attack.

Unusual network traffic such as incorrect HOP counts in network packets and packet delays.

If data is available from a packet Sniffer application during the time of the ARP spoofing attack, the network traffic could be analyzed for evidence of the attack. The data could either be collected by the investigator while the attack is occurring or from a Sniffer that is running on the network on a constant basis, such as a Sniffer dedicated to monitoring the traffic going to the Internet. The investigator could look for packets with an improper Hop count. In the TCP/IP protocol, the Hop count is increased by a router as part of the routing process. If the attack had captured the packets and then rerouted them to the firewall then the Hop count would be one higher than it was supposed to be. Further, the source Mac address entry in the TCP headers would be the Mac address of the attacker's machine and not the victim's machine. Finally, the time the packet takes to reach the firewall will be longer than usual due to the slight delay of the packet traveling to the attacker's machine and then routed to the firewall. Richard Duffy (November 2001) has written a paper on the topic of using the Dsniff tool on the network, which can be found at <http://www.sans.org/rr/penetration/dsniff.php>.

Analyzing the Windows 2000 Log files

As part of the Widget Research security policy, the organization has deployed the Auditing feature available in the Windows 2000 Active Directory. The network team has implemented a Domain Wide Group policy that all log in failures and success are logged as well as failed and successful file access on employee's network drive. The auditing results can be viewed by means of the "Event Viewer" which is used to view Windows 2000 logs. Further, the auditing policy has been deployed secretly and only a few employees involved in administering the Windows 2000 network and the security department is aware of how extensive the logging is. These logs are a valuable tool in detecting the attacker as it will have logged all network access attempts by the victim and the attacker. Included in these logs will be the IP address of the computer that the access was attempted from. The attacker had accessed the network with the victim's credentials when he used the "connect as a different user" option when mapping a drive to the victim's network drive. Therefore, investigators could

parse the logs looking for successful or unsuccessful log-in attempts from the user "victimj" that originated from an IP address that is not used by the victim's workstation. Further, the investigators will look for the file access log entries for the victim's files that do not match the times when the victim was working on these files. In this case, they probably will not find anything in term of unusual file access times as copying a file does not seem to change the file's modified or accessed stamp, however it is worth a try in case the attacker accidentally accessed a file on the victim's network drive. The results from the log analysis in this case will be discussed fully in the Incident Response section of the paper.

Finding evidence of the attacker's activities by checking the attacker's Windows 2000 box and his Linux 7.3 box.

The attacker would most likely try to cover his tracks by removing any evidence of his questionable activities that he had conducted on both his Window 2000 box and the Linux 7.3 box. However, if the investigator had enough evidence that it was the attacker who did the attack both machines could be seized and "combed" through to find evidence that the machines had being used in the attacks.

On the Linux machine the investigators could look for files that belong to the Dsniff utility and the supporting Linux files that are needed to run Dsniff on Linux. The investigator could gain the root level password by using the password recovery procedure for Linux. The investigator could also check to see if the Box had being configured for IP forwarding which is needed to make an ARP spoofing "man in the middle" attack work. Further, in Linux, recent commands are kept in a buffer that can be accessed by pushing the up-arrow key. The investigator could use this command buffer to see what commands have recently been run on the Linux box. The system logs could also be reviewed to see if they offer any traces of the Dsniff tool being run or installed.

On the Windows 2000 box, the investigator could check the local systems log to see if it has information on connecting to the victim's network share. If the attacker was not willing to give his password to the investigator this password could be reset on the Windows 2000 domain controller to give the investigator access. Further, on Windows 2000, a history of mapped drives is kept on the Map Network utility by means of a drop down menu with all of the recent drives mapped. Finally, all the drives on the computer could be searched for files belonging to the victim just in case the attacker accidentally left one on his hard drive.

Using the companies IDS system to detect the attack.

An IDS system was set up in the laboratory to see if it would detect the ARP spoofing activity. The IDS system was Snort v. 1.77.2.20 running on a Linux 7.3 Box (item 2 of Figure 0-1). The Snort configuration of the lab is identical to the

one used in the production network of Widget Research. Snort was configured to listen to the interface that was connected to FastEthernet 0/1 on the lab's Cisco 2900XL switch. The 2900XL switch was configured to mirror all traffic traveling through the port that the Check Point Firewall was connected to (FastEthernet 0/15). Therefore, the Snort process saw all of the traffic coming and going from the simulated Internet. Unfortunately, no alerts were found in the Snort Alert Log concerning the victim's machine, the attacker's machine, or the Check Point Firewall. The attack was conducted without even triggering one alert by the Snort Application. It is possible that there are Dsniff signatures for Snort but they did not come with the Standard Snort Signatures as downloaded from www.snort.org.

How to protect against the exploit

There are several ways to deal with the vulnerability: to make modifications to the organization's Web Authentication System, to consider using a different authentication database for NT authentication, to deploy personal firewalls on workstations, and to make policy changes on the storing of sensitive electronic documents,

The biggest problem with the Widget Research Web Authentication System is that passwords are sent between the client and the firewall in clear text as part of Check Point Firewalls Host Authentication. Widget Research should change the Firewall Authentication method to "Session Authentication" as this authentication schema can be configured to use Secure Sockets layers (SSL) in order to encrypt the data. In Session Authentication a "Session Authentication agent" is installed on peoples' workstations. The Session Authentication agent handles the task of sending the user's username and password to the Check Point Firewall. When a client attempts to access an Internet host, such as a web site or an FTP site, the Check Point Firewall will attempt to contact the Session Authentication agent running on the user's computer. The Session Authentication agent will then pop up with a graphical interface for the user to enter his or her username and password. The agent will then send the information back to the Firewall using Secure Socket layers. By using the Session Authentication agent, the Check Point Firewall no longer has to use the weak authentication methods of the connection protocol but rather can utilize the authentication agent application to collect the password and username and transport the data in a secure encrypted manner. Therefore, Widget Research should immediately begin a project to install the Check Point Authentication agent on all employees' desktops and reconfigure the Check Point Firewall to use "Session Authentication". By this means the employees' sensitive NT passwords and usernames would be encrypted at all times by the Internet Authentication system.

The organization should also address potential vulnerabilities in the part of the Web Authentication system that utilizes the Radius protocol. The Check Point

Firewall uses Radius to validate the username and password of the user, and although Radius encrypts data, there are vulnerabilities that could be exploited to intercept user's NT usernames and passwords. The server running the Cisco Secure Server, which is acting as the Radius server, should be directly connected to the Check Point Firewall either by a crossover cable or a switch that is only used to connect the Firewall and Radius server. The Cisco Secure server would then need to be directly connected to a Windows 2000 domain controller in order to validate the usernames and passwords on the Windows 2000 domain. Under this setup, all of the communications between the Firewall Radius Server and Windows 2000 domain controller would traverse by isolated point-to-point rather than travel through the production network that is accessible to all internal users. By deploying this setup there is no opportunity for an attacker to exploit weaknesses in the Radius protocol or the Kerberos protocol as the attacker would have no way to access the network traffic or to directly connect to the Radius servers since the traffic would no longer be traversing the production traffic.

The organization should also examine its policy on storing sensitive data. Such highly sensitive documents should not have been stored on the victim's network drive in clear text. The document should have been kept on a floppy disk stored in a safe or a reinforced locked drawer. Alternately, the company should deploy a system for heavily encrypting sensitive documents if they wish to continue to store the information on network drives. One system may be to deploy Public Key Infrastructure technology such as that offered by Entrust Research .

The organization also needs strict policies in transferring sensitive information by means of unsecured protocols such as SMTP (email), HTTP, and FTP. This policy should be that sensitive information should not be electronically transported using such protocols. Rather, a system using encryption technology should be used to transport the information such as PKI, or SSH file transfer.

The company should deploy personal firewalls on the workstations of employees who are working with sensitive data. A personal Firewall would have alerted the victim when the attacker pinged the workstation during the surveillance stages of this attack. Further, as described in Richard Duffy's paper "Finding Dsniff on your Network" the personal firewall "ZoneAlarm" will alert the victim when the Default gateway is being spoofed (as was the case in this incident) as the Dsniff application running on the attacker's computer attempts to communicate with the victim's workstation, which causes ZoneAlarm to alert the user.

Further, for those computers that are operated by people working on sensitive data it may be worth providing such employees with a second computer that is connected to a high secure secondary network. Although it would be expensive, a new network that is running on fiber optic media and not connected to any external networks could be used for highly sensitive data. Fiber optic media is considered more secure than copper media as it is harder to tap into the network.

The user would have one computer to connect to the company's internal network and the Internet and a second one that is only connected to a highly secure network designed to protect sensitive data.

Finally, the organization may decide not to use the Windows 2000 user database as the basis for NT authentication or to remove Internet Authentication all together. The Windows 2000 usernames and passwords of employees are very sensitive security credentials. It may be worthwhile to deploy a second user database that could be used for web authentication such as a LDAP database. The disadvantage is that users would have to use a different username and password to access the Internet. However, the positive aspect is that if someone intercepts the Internet authentication password all that he or she would gain is access to the Internet. Even more drastic, the company may decide that the security risk of uncontrolled access to the Internet is less of a risk than creating a potential chance of user accounts and passwords being intercepted.

Part 3: The Incident Handling Process

Preparation

Widgets Research's security system relies on several independent groups working together. The Security Services group is responsible for investigating incidents and developing security policy. The Internet Services group is in charge of maintaining all systems related to the Internet including the firewall, DNS servers, SMTP servers, IDS systems, switches and routers within the Internet system architecture and all servers and components involved in the Web Authentication system. The Internet Group provides information to the Security Group when requested to do so as well as report any unusual or suspect activities observed within the Internet systems. The Security Services activities are restricted to making policy and carrying out investigations whereas, the Internet Group are restricted to technical tasks and the collecting of technical material such as log file analyses. A third group, the Windows 2000/LAN team maintains and administers anything to do with the Windows 2000 network such as the Windows 2000 Active Directory, the Windows 2000 DNS structure, Windows 2000 domain servers (Domain Controllers, Print Servers, File Servers etc), They also configure and support the Windows 2000 professional workstations and support all of the applications used by the employees on their Windows 2000 workstations. The Windows 2000 group also provides information to security when needed. All three groups are managed independently and members from one group do not have any access to administrative accounts, or log files that fall under the jurisdiction of any other group. Hence, when the security group needs information or system access, they have to request it from the respective group's manager who may or may not provide it to them.

Emergency response and incident handling are coordinated by the Security team. After an incident is reported to them, they will get any information and technical assistance needed from the Internet group and the Windows 2000 group. Security will also contact the organization's Legal Services Department and outside authorities if needed. Security will advise other groups on how to prevent similar incidents in the future after the investigation is complete. The Security team is understaffed and has a low budget, which limits their investigation and policy development capabilities. Further, they do not get full cooperation from some of the groups involved in the company's network and information systems. There is no formal documentation on how computer related incidents are to be handled.

After business hours, emergency response capabilities are limited to one of the Internet team members being on 24-hour on-call status. The web masters and LAN managers have a phone number to page this person in an emergency. The on-call staff member is primarily equipped to fix failures in the Internet system, such as a server crash, but will deal with an exploited system in minimal manner such as unplugging the machine or blocking an IP address. The incident will not be investigated until the start of the next business day as no members from the Security staff or Windows 2000 group are available to be called in after business hours.

The organization has developed very lengthy policies on acceptable use of the company's computer information system, and Internet access usage. A log-in banner has been configured for all workstations reminding the users that the company's computer systems must be used in a way that does not contravene the company's "computer and network" policy. However, new employees are not provided with the policy and it takes some effort to get a copy of this lengthy document. Furthermore, the employees are represented by powerful union and disciplining an employee is a very difficult process. The organization requires all employees to sign a contract on remote network access if the employee requests VPN access to the network from their home computer. The organization's outlook on security is focused on preventing and investigating any external attacks to the network and has spent a lot of resources on preventing these attacks.

Identification

Unfortunately this incident was identified after a highly secret research document was stolen and given to a competitor of the company. It was found out that the information was provided to the competitor company after the competitor took control of the market using the information in the stolen document. The incident was not discovered by technical means but rather by the painful realization that the competitor must have this document as the competitor was giving the valuable information contained in the document to its clients free of charge as a customer bonus. The competitor could afford to provide this information free of

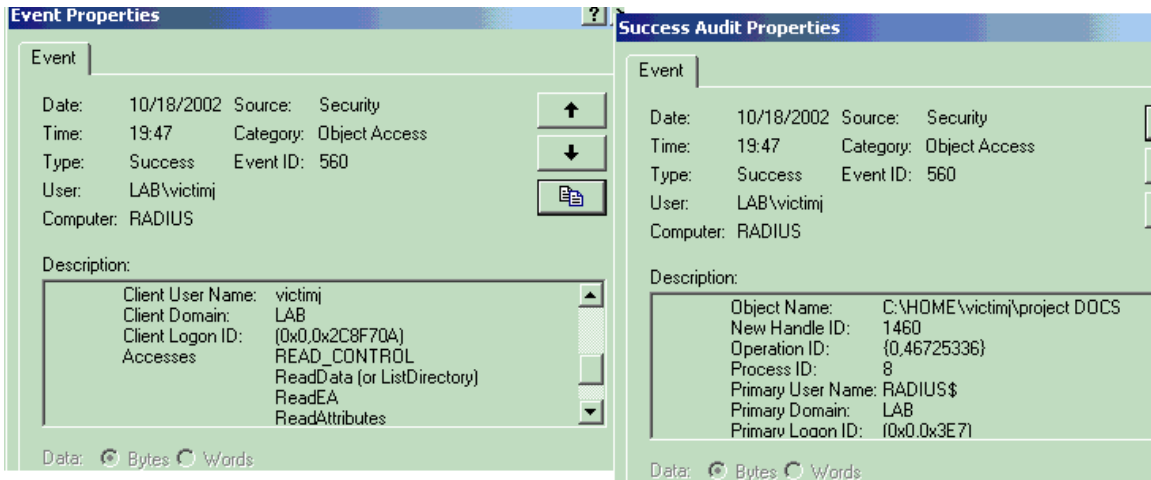
charge as it did not have to pay for the research. Thus the competitor took potential business away from the Widget Research Company. Further, Widget Research was forced to give the information away for free to their clients just to stay competitive.

It was three weeks before it was noticed that someone had gained access to the document. At this point, the investigation started into how the document was accessed. This investigation involved interviewing staff, reviewing log files, analyzing suspect's computers, and by reviewing non technical resources such as unusual employee work habits. Eventually the offender, who we will call Ted Attacker, was identified as a result of the investigation.

The first step was to interview the victim in the incident, who we will call Joe Victim. The interview convinced the investigator, who was a member of the Security Service team, that it was not Joe Victim or any of Joe Victim's colleagues that leaked the document to the competitor company. The investigator concluded that a third party had gained access to the document.

As part of the interview with Joe Victim, the investigator got as much information as possible as to the times and dates that Joe Victim accessed the secret document. The investigator contacted the Windows 2000 group manager and requested the Event Viewer logs. As the Windows 2000 team had enabled "File Auditing" on all files and folders contained in the employees' network drives, an entry was made in the log every time someone accessed or modified a file or folder. The investigator got one of the junior members of the security team to study all log entries related to accessing files or folders on Joe Victim's network drive. The object access time, file access times and modify times found in the logs were compared to times that Joe Victim recalls working on the files. This lengthy procedure did produce some useful information as the logs indicated that the directory had been accessed at a time that Joe Victim was positive he had not been accessing his network drive.

In Figure 0-10 a sample of the finding from the log search is displayed. These log samples (taken from the reproduction of the attack in the lab tests) indicate that the victim's network drive was accessed by the user account victimj about 19:47. The log entries indicated that victimj's home directory had been accessed by victimj by a type "ReadData or ListAccess" successfully. Further, the access included the specific directory where the secret document was held was found. The log refers to "c:\Home\Victimj\project DOCS" as these logs are taken from the local event viewer's security logs for the server that contained the employees network drive share. During the time frame that the logs indicate that the directory was accessed, Joe Victim was at home and not working on any files on his network drive. As a result of these log findings the investigator became very suspicious that some one had discovered Joe Victim's password and used it to access his network drive.



Log Name	Event ID	Date and Time	Source	Category	Event ID	User	Computer
Security Log	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS
System Log	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	562	SYSTEM	RADIUS
Directory Services	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS
DNS Server	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	562	SYSTEM	RADIUS
File Replication	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	562	SYSTEM	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	562	SYSTEM	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	562	SYSTEM	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	562	SYSTEM	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	562	SYSTEM	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	562	SYSTEM	RADIUS
	Success Audit	10/18/2002 7:47:42 PM	Security	Object Access	560	victimj	RADIUS

Figure 0-10 A sample of the Event Viewer Security Log entries that show the user “victimj” accessing his network directory at a time that the victim is positive he was not doing any work on the network drive. The sample is taken from the file server hosting the employees network drives.

The next step in the investigation was to compare the times that Joe Victim stated he logged into the network to times that the Windows 2000 logs indicated he logged into the Windows 2000 domain and look for inconsistencies. The Event Viewer Security log files on one of the networks Domain Controller, were parsed for any entries that involved account “victimj” logging into the network.

The results of this log searching led to the discovery of what computer was used to access Joe Victim's network using Joe Victim's user account at the time Joe Victim was at home not doing any work on his network drive.

A sample of these log entries can be seen in Figure 0-11. These log samples, taken from the Event viewer on one of the domain's Windows 2000 Server Domain Controllers, indicate that victimj logged into the network at about 19:20. Further the logs indicate that victimj logged in from the workstation 189.61.33.235 which is the Windows 2000 workstation for the employee Ted Attacker. It is certain that Joe Victim did not use Ted Attacker's computer to access his network drive and therefore Ted Attacker became the lead suspect as the person who stole the secret document. Later in the investigation it would be discovered that the "victimj" log in entries were caused by Ted Attacker using the "connect as a different user" feature when mapping the drive to Joe Victim's user account using Joe Victim's username and password.

Type	Date	Time	Source	Category	Event	User	Computer
Success Audit	10/18/2002	7:20:33 PM	Security	Account Logon	673	SYSTEM	RADIUS
Success Audit	10/18/2002	7:20:33 PM	Security	Account Logon	673	SYSTEM	RADIUS
Success Audit	10/18/2002	7:20:33 PM	Security	Account Logon	672	SYSTEM	RADIUS
Success Audit	10/18/2002	7:20:33 PM	Security	Object Access	560	victimj	RADIUS
Success Audit	10/18/2002	7:20:33 PM	Security	Logon/Logoff	540	victimj	RADIUS
Success Audit	10/18/2002	7:20:33 PM	Security	Privilege Use	576	victimj	RADIUS

Figure 0-11 Event viewer Logs indicating that victimj had logged into the network from the workstation 189.61.33.243 (the Windows 2000

professional workstation used by Ted Attacker) as a result of the Attacker mapping a drive to victimj network drive using the victimj username and password.

Due to the evidence found in the log files Ted Attacker's Windows 2000 Workstation and Linux 7.3 box were seized by security and Ted Attacker was told to stay home while the investigation was conducted. Ted Attacker's Windows 2000 accounts were disabled to prevent him from accessing the network system. Further, all accounts for servers and systems that Ted Attacker knew the password for, such as the systems he administered as part of his job duties, were immediately changed. The investigator was confident that Ted Attacker was the one who accessed Joe Victim's network drive and gained access to the secret document.

The next job of the investigator was to find out how Ted Attacker was able to obtain Joe Victim's password. The investigators quickly ruled out Ted Attacker guessing Joe Victim's password as Joe Victim had used a fairly complex password, which was "FishBay66".

The investigator attended a meeting with technical members of the Internet team. The Internet team advised the Investigator that they had a "hunch" that the password may have been gained by Ted Attacker exploiting a weakness in the Company's Web Authentication system. This weakness being that the NT passwords are transferred in clear text from the employees' workstations to the organization's Check Point Firewall as part of the Web Authentication system. The team advised that the simplest way to intercept the passwords was to do a "man in the middle" attack using the Dsniff tool. The team had worried about this type of attack but senior management did not believe the threat of an internal employee exploiting the weakness to be worth the resources needed to address the system weaknesses.

The investigator had the Windows 2000 team manager reset the username "attackerT" password so that he could log into Ted Attacker's machine to look for traces of the attack. The investigator hoped something in Ted Attacker's local profile on his Windows 2000 box would lead to more evidence. After some searching, the Investigator discovered that Ted Attacker had tried to map a drive to Joe Victim's network drive as the drive mapping was present in the drop down menu of the "map network drive" menu. The "map network drive" utility in Windows 2000 keeps a history of recently mapped drives to make it convenient to re-map the drives in the future. In this case it added further evidence that Ted Attacker was the attacker. Figure 0-12 is a screen shot of this finding.




Figure 0-12 The “Map Network Drive” utility history feature indicates that the user, Ted Attacker, had mapped a drive to victimj ‘s network drive. The investigator logged into the Attacker’s workstation using Ted Attacker’s account to gain this evidence.

Next, the investigator logged into Ted Attacker’s Linux 7.3 box. The investigator gained the root password by having a member of the Internet Team do a passwords recovery procedure on the box. The investigator searched for any files associated with the Dsniff utility as this is one of the most popular utilities for doing internal “man in the middle attacks”. The investigator executed the command “find . –name “dnss*” at the root of the Linux file system. The command revealed that the Dsniff utility had being installed on the system. Next the investigator executed the command “find . -name “arpspoof” to verify that the utility that the Dsniff utility “arpspoof” that is commonly used for this type of attack was installed on the system. Figure 0-13 shows the results of the file searches conducted by the Investigator. Further, the investigator found that the Box had being configured with IP forwarding to route intercepted data to the Check Point firewall. This configuration is consistent with a machine that is being used for an Arpspoof “man-In-the middle” attack.

```

[root@BadGuy]#cd /
[root@BadGuy]#find . -name "dnss*"
./usr/sbin/dnssec-makekeyset
./usr/sbin/dnssec-keygen
./usr/sbin/dnssec-signkey
./usr/sbin/dnssec-signzone0./usr/share/doc/bind-9.2.0/misc/dnssec.rc1stuff
./usr/share/doc/bind-9.2.0/misc/dnssec
./usr/share/man/man8/dnssec-keygen.8.gz0
./usr/share/man/man8/dnssec-signzone.8.gz
./usr/share/man/man8/dnssec-signkey.8.gz
./usr/include/dns/dnssec.h
./usr/local/lib/dnsspoof.hosts
./usr/local/sbin/dnsspoof
./usr/local/man/man8/dnsspoof.80

[root@BadGuy]#cd /
[root@BadGuy]#find . -name "arpspoof"
./usr/local/sbin/arpspoof

```

Figure 0-13 The results of the command “find . -name “dnssf*” and “find . -name “arpspoof” as executed on Ted Attacker’s Linux 7.3 Box. This command will search the Linux Box’s file system for any file starting with “dnss”. The results prove that the hacking tool “Dsniff” is installed on the Linux Box.

The investigator asked the Internet Team to search the Snort based IDS log files for any alerts concerning the IP address 189.61.33.240 (Joe Victims machine), 189.61.33.235 (Ted Attacker’s Linux 7.3 machine), 189.61.33.243 (Ted Attacker’s Windows 2000 workstation), 189.61.33.233 (the file server hosting the network drives), and 189.61.33.232 (the internal interface for the firewall). Unfortunately, new alerts that relate to this incident were found by the Internet team.

Ted Attacker and his union representatives were called in for an interview. After the evidence was presented to Ted Attacker it was agreed that Ted Attacker would be given a chance to resign rather than face firing. In exchange for allowing Ted Attacker to resign, Ted Attacker agrees to provide a detailed account of all the steps he took to steal the NT username and password of Joe Victim and how he used Joe Victim’s account to steal the secret document. Ted Attacker signed an agreement not to release any information about the system and is escorted from the building. The deal was good for both the company and Ted Attacker as neither has to go through a long drawn-out embarrassing employee termination procedure. Also sensitive information would have to be released in such a process. Ted Attacker did feel bad about being caught and having to resign his position. However, he feels a little better knowing he has a good chance at getting a new job, especially with the company to which he gave the secret document. Ted Attacker vows that in future attacks he will do better surveillance as he made a huge mistake in not establishing that the Windows 2000 team had configured such very detailed logging procedures.

Containment

To contain the damage the company decided to turn off the Web Authentication System and to force all employees to change their NT passwords. The logic of these steps was that other employees may also have figured out how to exploit the organization's Web Authentication System. Further, it was possible that the passwords of other employees may have been discovered either by Ted Attacker or by other people who knew about the exploit. The Internet Group reconfigured the rules on the Check Point Firewall to allow access from the workstations to Internet hosts without having to first be authenticated. To accommodate for the turning off of the Web Authentication System, extra staff resources were dedicated to analyzing the Check Point Firewall log files to determine what web sites and FTP sites employees were accessing.

Containment was also achieved during the Identification stage of the Incident handling process when management suspended Ted Attacker from his work and removed all network privileges by changing system passwords and locking his Windows 2000 account. Further, the computer was seized in the Identification stage preventing this machine from being used for future attacks.

The password change was conducted by adjusting the Windows 2000 Domain security policy. The policy was configured so that all passwords would expire in two days. This policy change forced all 10,000 + employees to change their passwords in the next two business days. When an employee logged into his or her Windows 2000 machine, they encountered a dialog box indicating that they must now change their password. After the two days had passed the policy was changed back to 48 days to avoid employees having to change passwords every two days indefinitely. Further, employees with highly sensitive documents were contacted and told to immediately change their passwords. Figure 0-14 is a screen shot of the Windows 2000 domain administrator adjusting the policy.

© SANS Institute

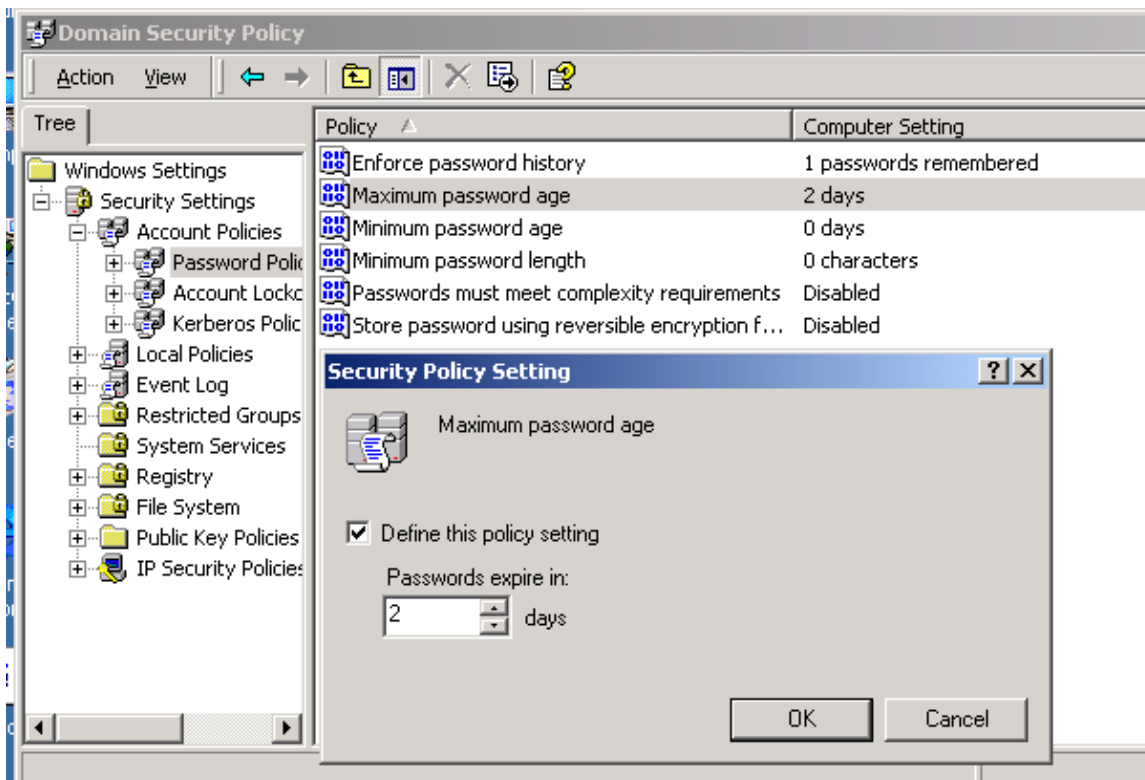


Figure 0-14 The Domain Wide security policy attribute “Maximum password age” was set to 2 days to force all employees to change their passwords.

The investigator worried that Ted Attacker had used his password interception knowledge to intercept the passwords of other employees. He also worried that other employees had exploited the web authentication system. The investigator also considered the possibility of people placing hostile code on employees’ Network drives, although the investigation had not provided evidence of this as yet.

It was decided to audit every system that Ted Attacker had been allowed to use as part of his job. Further, even systems that Ted Attacker had no access to were audited, such as the systems involved in the organization’s Internet system (e.g., firewall, DNS server, etc), as well as all of the Windows 2000 domain controllers. The audit involved checking for signs that the systems had been “hacked” into or compromised and that all recent service packs and patches had been installed. It was also ordered that a database be compiled of which employees had secret documents saved on their network drives. For the most sensitive of these documents, the network access and network drive accesses were reviewed via the Windows 2000 Log Files. Finally, an outside consultant with years of experience in packet and network analyzing was hired to check for signs that any more employees were running hacking tools, such as Dsniff, from within the internal network.

Also as precaution, it was ordered that all backups of the Network Drives for the period three months before the incident and for three months after be kept indefinitely. The Windows 2000 team does a full backup of the Network drives weekly and a differential backup daily. These backup are kept for anywhere from one month to six months. However, all backup tape media containing data from three months before to three months after the exploit were labeled and kept indefinitely in storage.

Eradication

Several of the steps discussed in the Identification and Containment section also serve to eradicate the problem:

- a) In the Identification stage of the incident handling process, the employee responsible for the attack was identified. As Ted Attacker was immediately stripped of his Network privileges and suspended he was no longer able to do further damage to the organization. Hence, the source of the attack was eradicated by suspending and eventually pressuring Ted Attacker to resign.
- b) In the containment stage of the incident handling process the Web Authentication System was shut down to prevent any more employees from discovering other employees' NT usernames and passwords. Although this step opened the opportunity for anonymous Internet access to people inside the organization, it did serve to eradicate the vulnerability that was exploited to gain the NT password that lead to the document being stolen.
- c) Also, in the containment stage, all employees where forced to change their passwords. This eradicated the problem that other employees' passwords may have been discovered by Ted Attacker or a different person who had decided to exploit the Web Authentication system.

As this attack did not involve infecting workstations or servers with malicious code there was not a lot of clean up to undertake. The auditing described in the containment section did not uncover any evidence that the attacker had placed malicious code or files on Joe Victims' Network Drivers or any of the systems he had access to as part of his job duties. It seems that Ted Attacker was only interested in stealing the secret document and nothing more. The attack consisted of directing network traffic using Apr spoofing, which does not involve infecting or corrupting servers and workstations within the network. Rather the attacker used the properties of the TCP/IP protocol to fool workstations to send data to the attacker's machine rather than the correct destination. Therefore, the workstations and servers that where targeted in the attack do not have to re-installed.

Ted Attackers Linux 7.3 box is the only system with potentially damaging code and utilities installed on it as the Dsniff tool is installed on the box as well the

machine is configured with IP forwarding. This computer will be stored in a secure location along with Ted Attacker's Windows 2000 workstation box in case they are needed for evidence in the future. As the Linux box will be held in secure storage for an indefinite period of time, the source computer of the attack has being eradicated from Widgets Research's network.

The root cause of this incident was an unhappy employee, combined with a weakness in the companies Web Authentication System that gave the chance for the employee to get revenge on the company by stealing information. The secondary cause of this incident was that secret documents were being kept on employees' network drives without any encryption. The company thought it was doing a positive thing in setting up an authentication system for internet access that was convenient for employees to use. Employees could use the same username and password they use to log into the Windows 2000 network. However, by deploying a system that addressed the security problem of anonymous Internet access, they opened up an even more severe security vulnerability. In the next section it will be discussed how the company re-deploys Web authentication in a much more secure manner.

Recovery

Shutting down the system addresses the security vulnerability of NT usernames and passwords traversing the network. However, now that the Web Authentication System is shut down any one connected to the internal network can access the Internet anonymously. This situation is a problem for the company as it is important to be able to precisely pin point exactly who accessed a web site or FTP site at given time. The ability to access the Internet without first going through an authentication procedure is just not acceptable to the company's management and security staff. Further, it was a requirement that employees be able to use their NT usernames and password to authenticate to the Internet to avoid the costs of maintaining a separate database of user accounts and passwords used for only Internet access. In order for the company to recover from this incident the Web Authentication must be re-deployed but in a way that sensitive NT usernames and passwords do not traverse the network in clear text.

A decision is made to re-deploy Web Authentication but this time to use Check Point "Session Authentication" rather than "host authentication". As explained in detail in the section "Part 1-How to protect against it", Session Authentication provides a means of encrypting the data between employees' workstations and the Check Point Firewall using Secure Sockets Layers (SSL). By deploying Session authentication the passwords and usernames will be encrypted at all times when traversing the Network. A project team is formed in order to deploy the automated installation of the Check Point Session Authentication agent on all employees' workstations, as this agent is needed when using Session Authentication. The Session Authentication agent is a program that runs on the

employees' workstations that will communicate with the Check Point Firewall when the employee tries to connect to an Internet site. The agent handles retrieving the username and password from the user and then sending it to the Firewall using SSL.

While the Session agent deployment project was underway the Internet team made some architectural changes to the network that added even more security to the Web Authentication system. Before the attack, the Check Point Firewall communicated with the Cisco Secure server, that was configured as the Radius server, using the Widgets Research internal network. A decision was made to move the Cisco Secure Server off the regular network and connect it to an isolated network that consisted of only the Cisco Secure Servers, a Windows 2000 domain controller, and a Interface to the Check Point Firewall. This new design eliminated the opportunity for any one to interfere in the traffic between the Firewall and the Cisco Secure server and therefore reduced the chances of future attacks based on exploiting the Radius protocol. Further, it prevents any network access to the Cisco server network intercept from potential internal hackers. The diagram for the redesigned Network Authentication system is in Figure 0-15.

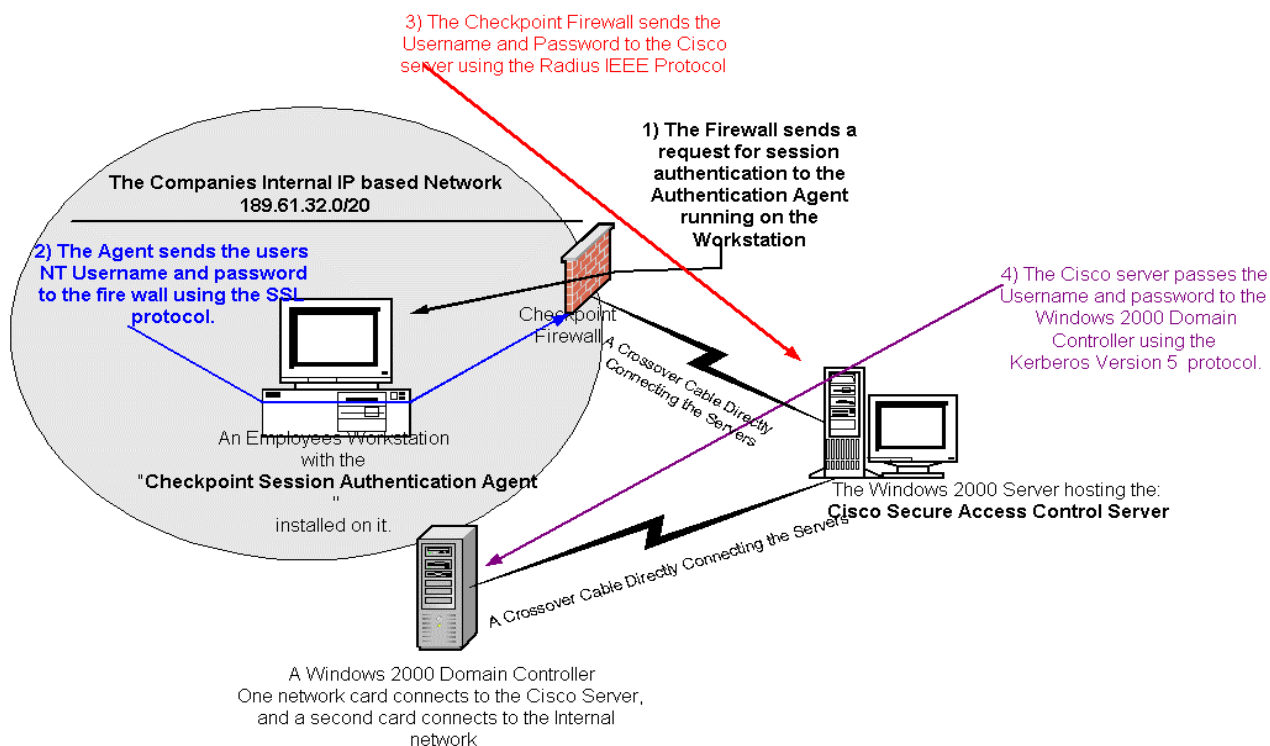


Figure 0-15 A diagram of the re-designed Web Authentication system.

The Web Authentication Agent deployment and the architectural changes to the system are completed in about four weeks. After the new design is tested with a

small group of employees, Network Authentication is reconfigured on the Check Point Firewall this time using the Session Authentication Schema. Prior to deploying the new system, the company had hired an outside consultant to assess the security of the new system. The company has now recovered from the exploit incident as they now have a Web Authentication System that allows employees to authenticate to the Internet using their NT usernames and passwords. The company is more secure as the passwords are encrypted throughout the authentication system. However, it will take some time to recover from the reputation and business loss that resulted when the document was stolen and given to the competitor company.

Lessons Learned

There were several lessons learned from this incident. The main ones were that the company's security policies and incident detection processes need improvement.

The incident was not detected until three weeks had passed after the attacker accessed the document. Further, the Incident Identification process was only started after it was noticed that the information in the document had been given to a competitor company. It would have been preferable to have an infrastructure in place that would have detected the incident much sooner.

In future incidents, the company's management wants to have systems able to detect a similar attack while the attack is actually occurring rather than discover the attack after it finished. The investigator was able to find the attacker by doing log parsing but there was no system to alert security of the unusual network traffic that was occurring during the time period that Ted Attacker was redirecting traffic destined for the firewall through his Linux Box.

The decision is made to invest in a project that is dedicated to detecting the "man-in-the-middle" style attacks. This project is of high importance given that there are "man in the middle" attacks for intercepting data sent using SSL, which is the protocol that the new Web Authentication system uses to transport NT passwords and usernames to the firewall. The company also wants to detect employees who use "man in the middle" attacks to gain access to SMTP (email) traffic going being sent out to the Internet. The project involves running a network Sniffer constantly in order to have a log of all network packets traversing critical parts of the network such as the connection to Firewall and other key servers.

Steps to be taken involve: installing personal firewalls on each employee's workstations; the purchasing of a product called "AntiSniff" that is useful in detecting where Sniffers are being ran on the network; and the purchase of commercial heavy weight Intrusion Detection systems to complement the Snort Ids systems currently running on the network.

Although it was Log files that eventually lead to finding the attacker, the process of finding the information was very time consuming. The security investigator had to wait for the Windows 2000 team to provide the log and then had to devote a lot of his staff's time in manually parsing through the log to find the information needed. In order to stream-line future investigations, it is decided that an automated log transfer process will be deployed so that Security will always have a up-to-date copy of the log files. Further, it is decided to purchase or design a utility that will parse the huge Log files into a clear summary report, preferably in HTML format. In addition, Security will try to find a product that does log analyzers of the Windows 2000 event viewer logs in order to find suspicious network and file access activities.

This incident has also prompted the company to change the way that secret documents are electronically stored. The company has decided to encrypt all secret documents using Public Key Infrastructure (PKI) technology. Therefore, if a file is obtained by an attacker then this file has some protection in that it will be encrypted.

The company needs a policy to protect itself from internal attacks. The current security policy is heavily concerned with avoiding and detecting attacks from people outside of the organization. However, this incident reinforced the concept that critical components of an internal network must be secured from internal attackers. Policy must be written and distributed so as to secure the company's network from it's own employees and not just concentrate on outside attackers.

The incident handling process needs more "fine tuning" in order to detect, respond and investigate incidents sooner. The investigator from the Security Group was constantly delayed by waiting for Log files and other information from other groups. Further, it took longer than it should for Securities' request to have the password reset for the attacker's password in terms of getting senior management approval and contacting the Windows 2000 administrator to have the work done. Also, the Security department does not have a staff member who is an expert in the technical aspects of deploying and preventing computer systems related incidents and as a result security has to constantly contact people outside the group for advice.

It was decided that a "Incident handling" team had to formed that was made of people from all parts of the company and an incident handling procedure document would be written for this group to follow when dealing with an incident. The group would be made up of people from Security, the Windows 2000 Group, the Internet Group, the GroupWise group, and Legal Services. The diversity of the group would insure that information could be accessed quickly and technical and non-technical procedures are done in a smooth and fast manner. A diversity of expertise would be readily available to the investigator due to the make up of the group. As part of the deployment of this group a "jump kit" would be put

together that consisted of any tools that would be needed in the incident response process. Further, an after hour procedure of calling in this team would be implemented. The actions that the team will take in an incident will be dictated by the detailed incident handling process document.

Finally, it was decided that Human Resources would look into why employees get so upset that they decide to do damage to their own company. This study will examine if Ted Attacker's anger towards the company is isolated to him or if this is a wide scale problem. It may be the case that the company needs to implement programs to boost employee moral and help them deal with the huge amount of changes being made within the company.

© SANS Institute 2003, Author retains full rights.

References

- Song, Dug. "dsniff Frequently Asked Questions" 7 December 2001
URL: <http://monkey.org/~dugsong/dsniff/faq.html> (19 January, 2003).
- "Authentication, Authorization, and Access Control" The Apache Software Foundation URL: <http://httpd.apache.org/docs/howto/auth.html> (19 January, 2003)
- Lammle, Todd., Porter, Donald., & Chellis, James. Cisco Certified Network Associate Study Guide SYBEX, Inc, 1999. 77-104.
- Chacon, Michael, Chellis, James, Donald, Lisa, Desai, Anil, & Robichaux. MCSE Accelerated Windows 2000 Study Guide SYBEX, Inc, 2000.
- Lammle, Todd., & Hales, Kevin. CCNP Switching Study Guide SYBEX, Inc, 2001.
- Skoudis, Ed., & Cole, Eric, Computer and Network Hacker Exploits, Parts 1, 2 and 3 Ed Skoudis & SANS, 2002. 175-195.
- SANS. Incident Handling Step-by-Step and Computer Crime Investigation SANS, 2001-2002.
- Microsoft Corporation. Windows 2000 Kerberos Authentication White Paper, Microsoft Corporation, 1999.
URL: <http://www.microsoft.com/windows2000/docs/kerberos.doc> (19 January, 2003)
- Cisco Systems. How Does RADIUS Work?, 12 December 2002 URL: http://www.cisco.com/en/US/tech/tk583/tk59/technologies_tech_note09186a00800945cc.shtml (19 January, 2003)
- Hill, Joshua. "An Analysis of the RADIUS Authentication Protocol" 24 November 2001
URL: <http://www.untruth.org/~josh/security/radius/radius-auth.html> (19 January, 2003)
- Litt, Steve. "IP Forwarding" 1999
URL: http://www.troubleshooters.com/linux/ip_fwd.htm (19 January, 2003)
- Duffy, Richard. "Finding dsniff on Your Network" 28 November 2001
URL: <http://www.sans.org/rr/penetration/dsniff.php> (19 January, 2003)

Dunston, Duane. "Network Monitoring with Dsniff" 29 April 2001
URL: http://www.linuxsecurity.com/feature_stories/feature_story-89.html
(19 January, 2003)

Check Point Software Check Point VPN-1/Firewall-1 Administration Guide.
Check Point Software Technologies, Ltd., 1999-2000. 489-562

Danielle, Lora. "Introduction to dsniff" 1 June 2001
URL: <http://www.sans.org/rr/audit/dsniff.php> (19 January 2003)

© SANS Institute 2003, Author retains full rights.

Upcoming Training

Click Here to
{Get CERTIFIED!}



SANS Madrid 2017	Madrid, Spain	May 29, 2017 - Jun 03, 2017	Live Event
SANS Atlanta 2017	Atlanta, GA	May 30, 2017 - Jun 04, 2017	Live Event
SANS Houston 2017	Houston, TX	Jun 05, 2017 - Jun 10, 2017	Live Event
SANS San Francisco Summer 2017	San Francisco, CA	Jun 05, 2017 - Jun 10, 2017	Live Event
Community SANS Virginia Beach SEC504*	Virginia Beach, VA	Jun 05, 2017 - Jun 10, 2017	Community SANS
SANS Charlotte 2017	Charlotte, NC	Jun 12, 2017 - Jun 17, 2017	Live Event
SANS Rocky Mountain 2017 - SEC504: Hacker Tools, Techniques, Exploits and Incident Handling	Denver, CO	Jun 12, 2017 - Jun 17, 2017	vLive
SANS Thailand 2017	Bangkok, Thailand	Jun 12, 2017 - Jun 30, 2017	Live Event
SANS Rocky Mountain 2017	Denver, CO	Jun 12, 2017 - Jun 17, 2017	Live Event
Mentor Session - SEC504	Reston, VA	Jun 13, 2017 - Aug 01, 2017	Mentor
SANS Minneapolis 2017	Minneapolis, MN	Jun 19, 2017 - Jun 24, 2017	Live Event
SANS Paris 2017	Paris, France	Jun 26, 2017 - Jul 01, 2017	Live Event
SANS Cyber Defence Canberra 2017	Canberra, Australia	Jun 26, 2017 - Jul 08, 2017	Live Event
SANS Columbia, MD 2017	Columbia, MD	Jun 26, 2017 - Jul 01, 2017	Live Event
SANS London July 2017	London, United Kingdom	Jul 03, 2017 - Jul 08, 2017	Live Event
Cyber Defence Japan 2017	Tokyo, Japan	Jul 05, 2017 - Jul 15, 2017	Live Event
SANS ICS & Energy-Houston 2017	Houston, TX	Jul 10, 2017 - Jul 15, 2017	Live Event
SANS Los Angeles - Long Beach 2017	Long Beach, CA	Jul 10, 2017 - Jul 15, 2017	Live Event
SANS Cyber Defence Singapore 2017	Singapore, Singapore	Jul 10, 2017 - Jul 15, 2017	Live Event
Community SANS Seattle SEC504	Seattle, WA	Jul 10, 2017 - Jul 15, 2017	Community SANS
Community SANS Sacramento SEC504	Sacramento, CA	Jul 17, 2017 - Jul 22, 2017	Community SANS
Community SANS Ottawa SEC504	Ottawa, ON	Jul 17, 2017 - Jul 22, 2017	Community SANS
SANSFIRE 2017	Washington, DC	Jul 22, 2017 - Jul 29, 2017	Live Event
Community SANS Des Moines SEC504	Des Moines, IA	Jul 24, 2017 - Jul 29, 2017	Community SANS
Community SANS Annapolis SEC504	Annapolis, MD	Jul 24, 2017 - Jul 29, 2017	Community SANS
Community SANS Phoenix SEC504	Phoenix, AZ	Jul 24, 2017 - Jul 29, 2017	Community SANS
Security Awareness Summit & Training 2017	Nashville, TN	Jul 31, 2017 - Aug 09, 2017	Live Event
SANS San Antonio 2017	San Antonio, TX	Aug 06, 2017 - Aug 11, 2017	Live Event
Community SANS Raleigh SEC504	Raleigh, NC	Aug 07, 2017 - Aug 12, 2017	Community SANS
SANS Boston 2017	Boston, MA	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS New York City 2017	New York City, NY	Aug 14, 2017 - Aug 19, 2017	Live Event