Global Information Assurance Certification Paper

Discovering your Network: The First Step for Hacking into your Network

With the proliferation of the internet as a regular part of daily life, the effects of the level of information that can be accessed at moment's notice is staggering. Only a few minutes, a topic, and a search engine are needed to access an incredible amount of information. As a result, it is becoming imperative to know what information regarding your network is available on the internet. My hope is that you will use this paper to guide you as you discover what information is readily available concerning your network.

All that is really necessary to compromise most networks is time and information. There are several steps that a hacker must go through before they are able to gain access a network. The first step that they must take is to gain familiarity with the network in question. The next step for the hacker is to identify key systems, and applications running on those systems. Finally, once a hacker has identified key systems and applications, he may continue on his path and produce a general internet search on each system looking for key vulnerabilities. Each of these steps will be covered shortly.

Now to understand why gathering information on your target is one of the most important parts of attempting to illegally access a network. Most successful hackers will not attempt to break into a network on a whim; most successful hackers will attempt to learn as much about your network as is possible. This is similar to how successful bank robberies are conducted. First you familiarize yourself with the target, you learn its strengths, and weaknesses, you look for ways to exploit the weaknesses and then you go about committing your crime. The same steps are necessary when breaking into a network.

Below is a longer more in depth list of steps that a hacker goes through to gain information on his target.

1. Find Initial Information about the network
2. Find the IP address range of the network
3. Find Active machines
4. Find active ports
5. Discover what operating systems each active machine is using
6. Discover which services each port is running
7. Search for vulnerabilities for each service

Step 1: Find Initial Information about the network:

Before anyone can even begin to scan you network for vulnerabilities they have to familiarize themselves with your company. One way of familiarizing yourself with a network is to conduct an open-source search for information. An open-source search consists of information that is readily available on the internet. Open source data can be found on a company's website, a

periodical, or anywhere else. The data provided by this search is useful when conducting a social engineering attack*.

Open Source Information

Open source information can more easily be thought of as information you can get about a company by perusing their website, and other news articles. This method of information gathering is used to gather information that can be used mainly during a social engineering attack.

When an attacker starts their information search on a company, the first place they will look is the company's website. While at this site you should look for any information that may help you on your quest to enter the network. Key pieces of information you may try to look for are

- Personnel directory
- New hardware, operating systems, security products
- The type of information they contain
- What the company is trying to protect (what is valuable)
- How the company is organized
- Key personnel and phone numbers

The first piece of information you should look for is a personnel directory or the name of key personnel. Using this information and a search engine such as Google, you can start to plan a social engineering attack.

In addition to searching for the personnel directory you should also look at the news releases section. In this section you are looking for information regarding any new products or changes to the company infrastructure that has been introduced. Frequently cutting edge companies will try to leverage new software that they are working with to attract new customers (i.e. "All of our Servers now run Windows 2000 Advanced Server"). Using this information a hacker can begin to narrow their vulnerability search to Windows Advanced servers problems.

Using the data you learned from your searches you might have enough information to conduct a social engineering attack. For example, you may know that Bill Powers works in the technical support division by using information from a personnel directory on the website. Reading the news section, you may also know that many members of the technical staff recently attended a SANS conference in San Diego. Using that person's name, the conference name, location, and Google, you may be able to discover what workshops that person attended. While sometimes you may not learn anything important you may find out that Bill Powers attended a workshop on securing internet information servers. This information can lead you to believe that this company's internet information servers may be vulnerable. Or you may use the information you learned through your searches to gain someone's trust at the company, convince them that you

* Social Engineering Attack: Is an attack in which the purpose is to convince someone within a network to give you information or access that can be used to your advantage. This is similar to calling someone at the help desk and convincing them to give you a new account. This is usually accomplished by convincing the victim that you are an employee of the company that for some reason needs an account. This is a non-technical attack.

are a friend of Bill's and that you need information on a server within their network or to even create an account for you.

You will be amazed by what you can find that can be used for a social engineering attack while searching the internet. In order to safeguard your information you may want to review your website to make sure that not too much information is being released to the public[*].

Whois Information

"Whois" is a Unix based tool that is used to get information about a company from information servers on the internet. The data that these servers provide have to do with the information provided by the point of contact at the network when they registered the company's internet presence. Some of the information provided is listed below.

- Registrant Name (Company Name)
- Registrant Address (Company Address)
- Domain Name
- Administrative/Technical Contact name, address, email, and phone number
- Billing name, address, email, and phone number
- Date the record information expires
- Date the record was created
- Date the last update was made
- Domain name server IP addresses

To conduct a Whois search you need a Unix based machine (Linux, etc) or a program that can conduct a Whois search. If you are running a Window's based machine you can use a program called Sam Spade to conduct a Whois search. Below you can see a printout from a Whois search run on Sam Spade.

[*] You should immediately remove any documents from your web server that you do not want the public to know. Remember that just because there is not a link to a document does not mean that someone cannot access that page.

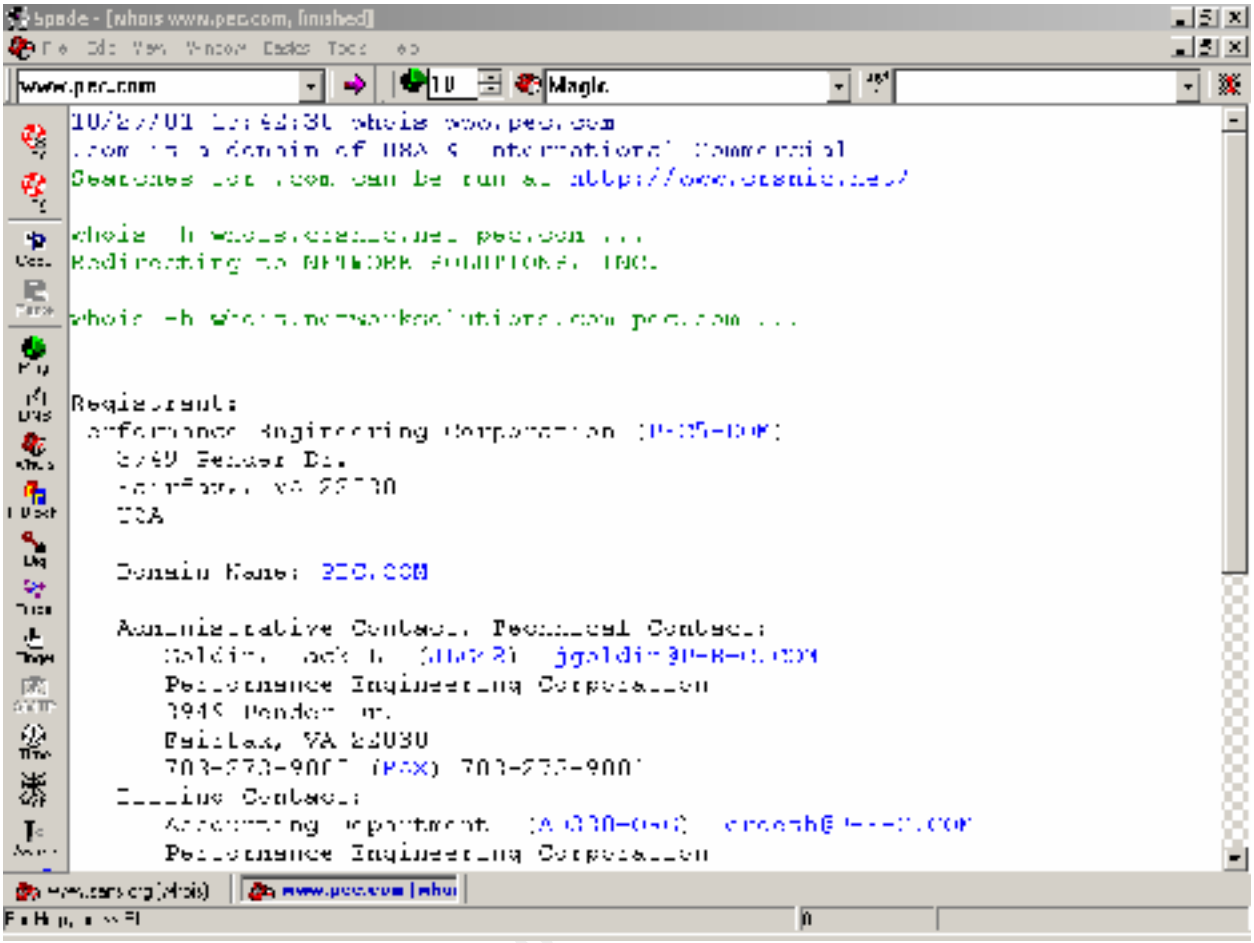

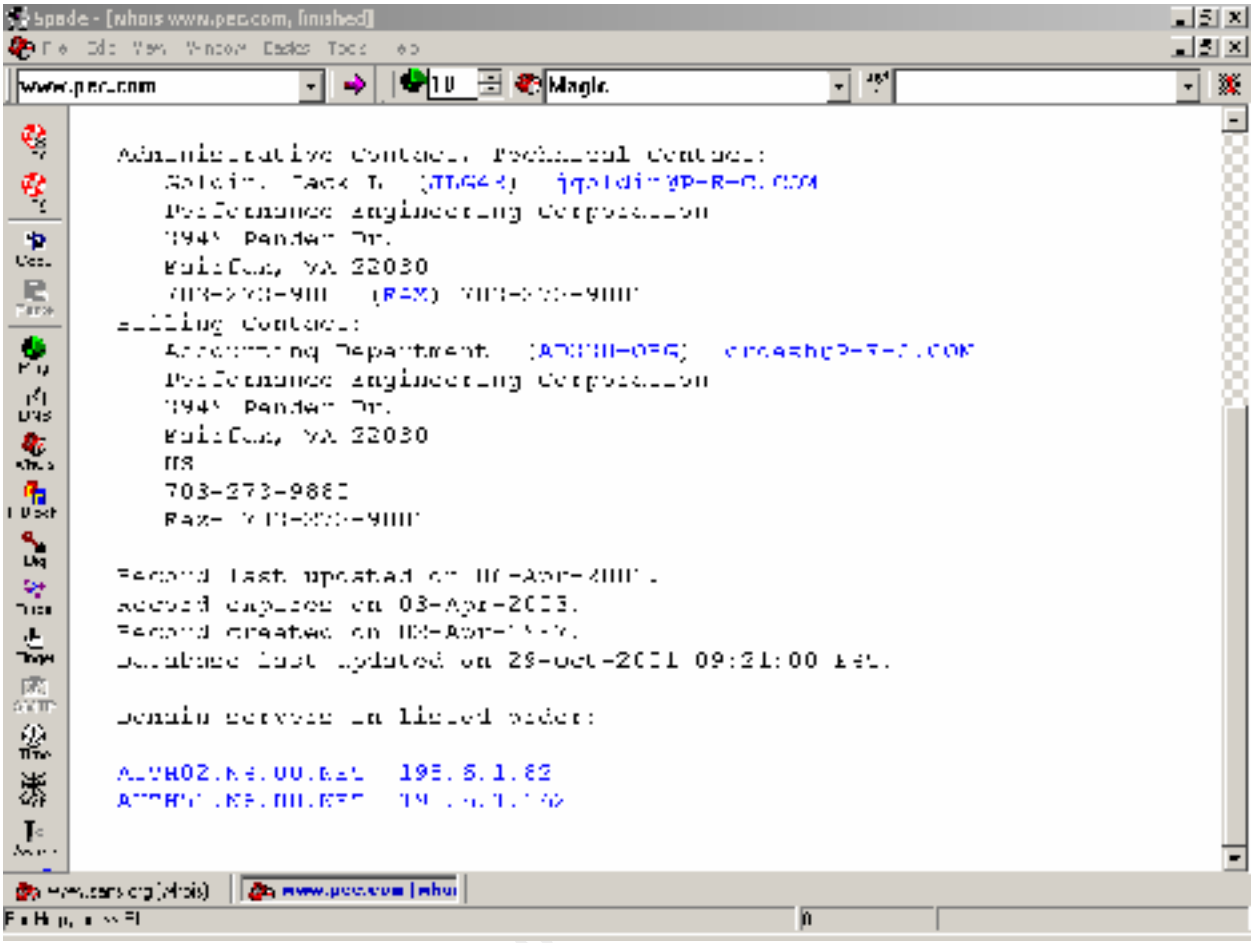

This information can be helpful on two fronts. First, the contact's name (under the administrative/technical heading) can be used as a start for a social engineering attack.

The other useful piece of information you gain from this is the address of the domain name servers for the website you are investigating. These addresses can be useful as a starting point for further reconnaissance of the target.

Step 2: Find the IP address range of your network

Now with the past two searches you have a fairly good starting point for a social engineering attack, but if you are looking to actually move on to a somewhat more technical reconnaissance you will have to visit a website. The site that you should visit is the American Registry for Internet Numbers (www.arin.net). This website provides the IP address ranges that an American company has been assigned. Using this information you can start to actually scan a network.

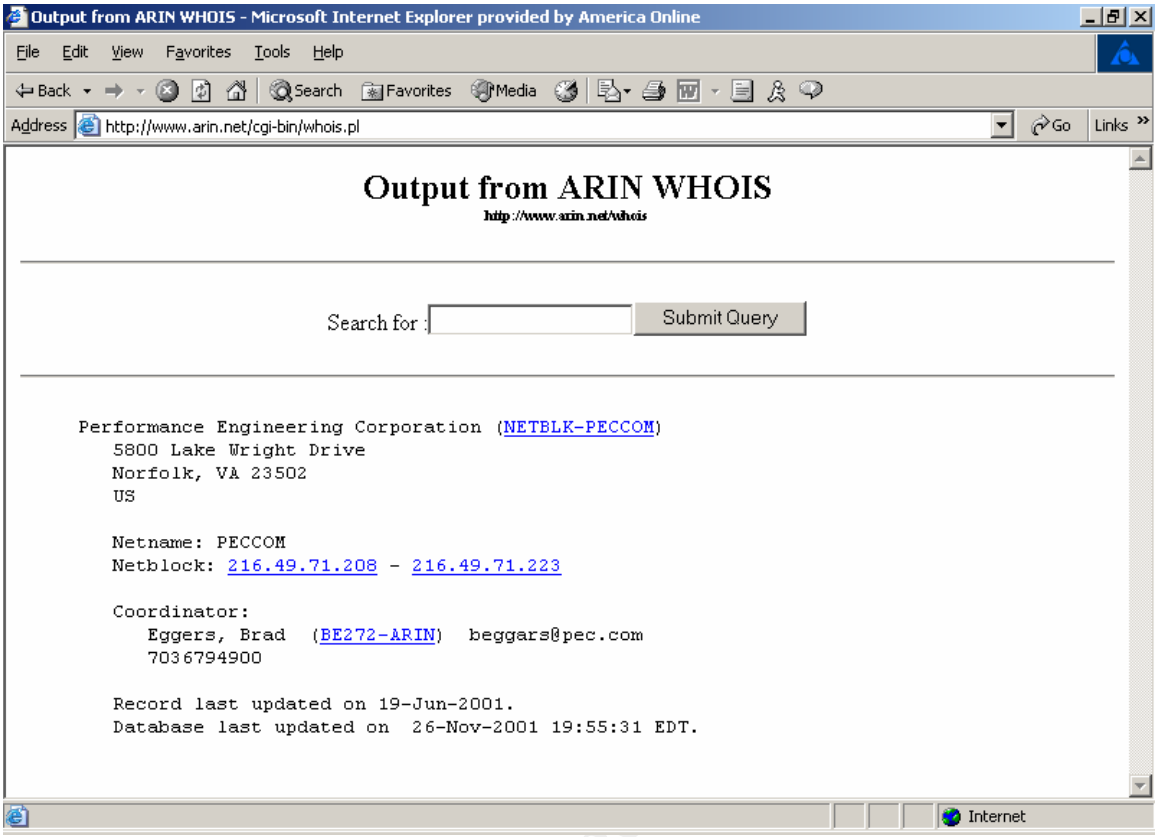

Above is a screen copy from the ARIN website concerning the IP addresses that a company owns. The important information concerning this screen has to do with the IP address range and the contact's name (social engineering attack). Remember to review all of the IP addresses that a company may own. Frequently, a company will have many different IP addresses assigned to it.

Finally, before you progress any further you should attempt to create at least a preliminary picture of the network you are trying to gather information on. Your goal at this point is to identify the external router, the firewall, and any other servers that are of interest (mail servers, web servers, DNS, etc). You can accomplish this by running a trace route to the web server or any of the servers. A trace route will show you the path that a packet takes from your computer to the computer you are trying to reach. Along with the IP addresses of the results you will see the name for layer 3 device (firewalls, routers) that the packet crosses. These addresses are important because these are the places where a company is most likely to try restrict you from getting any closer to actually reaching one of their internal hosts[1]. A key point here is to make sure that you compare the company's web server IP address with those assigned to it. Some companies have outsourced the hosting of their website to other company's who specialize in hosting web pages.

Step 3: Find Active Machines

While the previous steps have started to reveal a lot of information on your target the real information comes from running reconnaissance scans on a network. During the following scans you will be able to determine which IP addresses a company is actually using. The purpose of this section is not to establish a connection with a host, but just to verify that a host is actually using an IP address.

Remember that just because a company has access to a lot of IP addresses does not mean that they actually have to use them. In addition a lot of companies are using network address translation[ii*] to increase the number of IP addresses that they have access to and to keep hackers from being able to scan their network. Using the information you learn here, you will be able to differentiate which computers are servers, workstations and other pieces of equipment along with that you will shorten the list of computers that you have to port scan.

In this section you will learn how to scan a network. The purpose of this section is to gain a list of possible targets for your hack attack. To create your list you will have scan the allotted IP address range of a network. If you remember you already have this list of IP addresses from the your research from The American Registry of Names database.

The simplest method of deciding whether a machine is actually at an IP address is to use the ping scan. Ping works by sending an ICMP packet whose primary purpose is to discover whether or not a computer can be reached physically over a network. There are a lot programs that you can use to run ping scans on a network, but be forewarned that scans that rely on ping will rarely work because pings are frequently blocked at the router or at the firewall (as are most ICMP commands). In addition, scans that use ping are easily picked up intrusion detection systems. For this reason, you will most likely want to use a different set of scans.

There are three main types of packets that you can use to scan a network. You can use ICMP to scan a network, but since almost all routers and firewalls block these packets, you will most likely not want to use these packets. Since you are not able to use ICMP you're left with using transmission control packets (TCP), or user datagram packets (UDP). For our scans we will use TCP since it is a connection-orientated protocol. This means that before any data is sent between two hosts the two hosts must decide how they plan to communicate[**]. It is during these preliminary stages of the communication process that we are able to discover whether a host is using a specific IP address.

To understand why TCP is used for stealth scans (scans that are more difficult to detect than ICMP based scans) is because TCP is based on the principle of a three-way handshake. The three way hand-shake refers to a process by which TCP based communication will not begin

[*] Network Address Translation is the process by which a layer 3 device (router, server) changes the IP address of an internal machine (using a private non-routable IP address) to that of a public routable IP address. The purpose is to prevent anyone from the outside direct access to an internal machine. This is done by using IP addresses for internal machines that cannot be reached from the internet. This prevents network scanning of a network.

[**] For more information regarding TCP, UDP, ICMP refer to the book TCP Illustrated: Volume 1 by W. Richard Stevens (ISBN 0-201-63346-9). This is one of the most in thorough books regarding protocols.

until a three step process of setting the rules for communication between two hosts is completed. The process begins with the initiating host (client) attempting to establish communication with another host by sending a packet that contains a flag that is set to SYN; this is similar to the client computer saying, "hello I would like to speak to you." The receiving host (usually a server) then responds with another packet containing a SYN/ACK, similar to saying "I will listen to you, and I am interested in speaking to you." The original host (client) then responds with another ACK, this is like saying "I will also listen to you." A lot of other information is also communicated during this process but that information is not necessarily needed to understand how a TCP based scan operates. In any case if the server was not interested in speaking with the client he would respond with a RST message. This effectively halts all communication immediately.

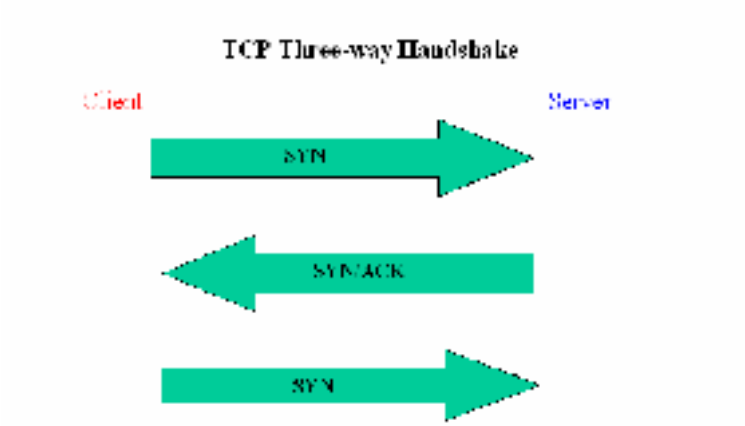

Since most communication involves TCP there are a lot of scanners available that will attempt to exploit this setup to gain a response from a host. Remember that the purpose of a scan is to elicit any type of response from a host. TCP being a complicated communication protocol has a lot of methods which you can use to get a response from a system. Most scanners will use a lot of different methods to try to get a response from a host. Some of the better scanners will even attempt to mix up the methods that they use to get a response from a host in order to not set off any alarms with the firewall or the intrusion detection system.

Below you will see AA tools, a Windows based network scanner. This scanner not only scans a list of IP addresses but it also scans each host to see which ports are open. At this point we are more interested in which machines are active than which ports are in a listening state on each computer. For Unix inclined users, you should download a program called Network Mapper (NMAP). Either program will allow you to conduct a TCP based scan on a network.

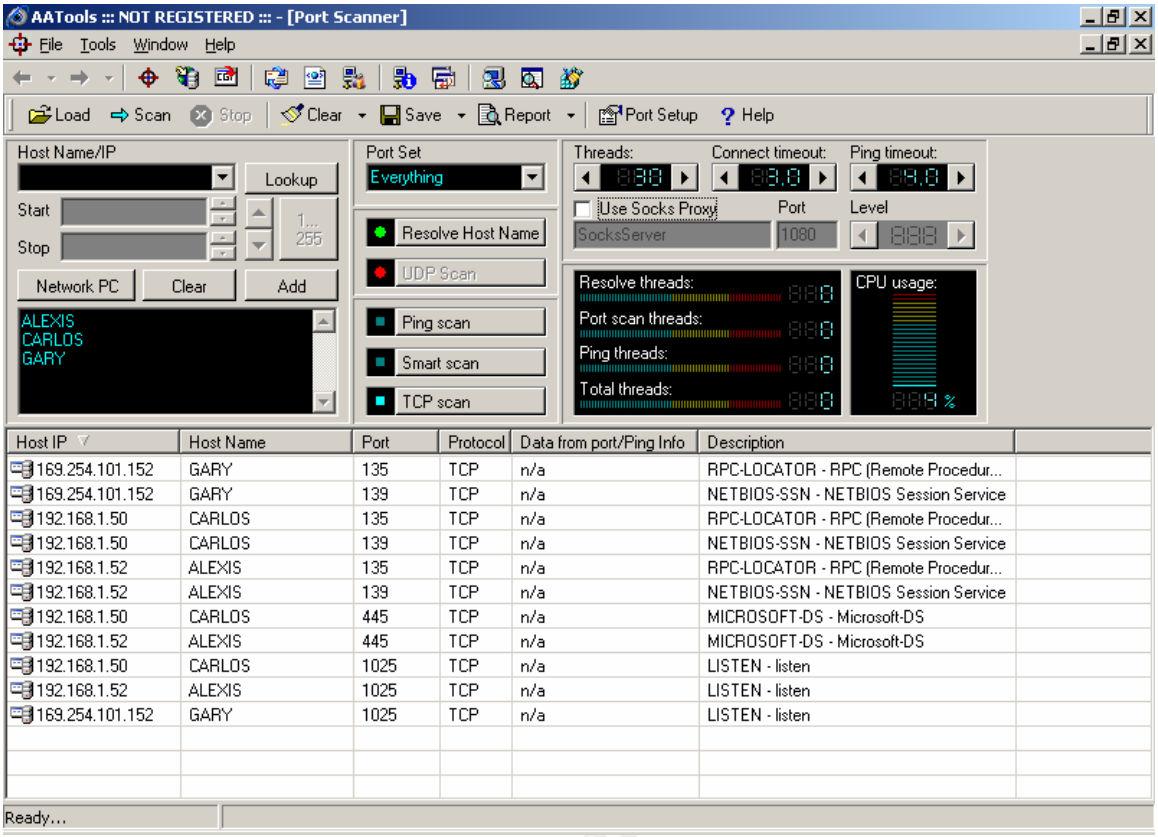

Whatever scanner you use you may be want to run the scan twice, once during normal business hours and once afterwards. The reason for this is because most users shut down their workstations at night, while most servers are on all the time. While this is not a steadfast rule, it may help you in shortening the list of hosts you will have to run a port scan on later.

At this point I would like to warn most users that running frequent scans against an IP address range will trigger alarms on an intrusion detection system or a firewall. For this reason you may want to vary the order in which you scan certain IP addresses, ports, and methods you use.

Step 4: Finding Active Ports

Once you have identified a list of targets within your network (IP addresses that are actually being used), you can begin the process of discovering the weaknesses on each host. Before you can begin exploiting each weakness you have to discover each weakness. The weaknesses on computers are ports. Ports are very important when dealing with computers, without the use of ports; computers would not be able to communicate. Ports are similar to a person's ears, each time a host starts communicating with another computer it makes available another set of ears (port). Each port has a number, so each packet bound for a computer is marked with the port number on the receiving host that is assigned to listen to their conversation. Through the use of ports, computers are able to maintain several conversations at one time; there are about 65,000 available ports.

While ports are great for communication, they are also unfortunately a weakness that computers have. They are a weakness because it is here that a computer can be broken into.

You goal during this stage is to scan hosts for ports that are in a listening state. The best tool for scanning a host is NMAP (Unix based), NMAP. Using this tool, you may want to scan the ports numbered 1024 and below. The reason for this is that most servers use ports number 1024 and below for connections. From there you may want to reduce the ports you have to scan to some of the more important ones (80, 23, 21, 20, etc). The ports you may want to scan are the ones commonly used by servers[iii]. Shortening the list of ports to a few ports will also lessen the likely hood of being caught by an intrusion detection system. On the Windows side, you may want to use StOrM or PortPro to conduct the same search.

Step 5: Discover what operating system each host is using

Now that you have a list of ports you are almost at the end of the line. Now you have to identify a few more pieces of information. For each machine you should be able to identify the operating system, and the application running on each port. The reason that you need this information is to help narrow your search for vulnerabilities.

To discover which operating system a host is running you can use one of two programs. You can either use NMAP, or another program called Queso (Spanish for cheese). Either one of these two programs has the ability to discover which operating system is running on each host. These programs work by sending each host a set of malformed packets. Since there is no set rule on how a computer responds to these packets each host will respond in its own way. By comparing each host's responses against a database you can narrow what type of operating system a host is running to a narrow list.

Step 6: Discover which services each port is running

Although knowing which operating system is running on each host is a good start we really need to know more information. From here you need to learn more about the programs using each port. The best way to learn this information is by telnetting to each individual port. While this is slow, it will provide the most thorough response.

Telnet is a terminal emulation program that comes by default on most systems. Using this program you will be allowed to attempt to connect to each port. Telnet will display on its screen the banner that appears when you connect to each port. The banner usually by default contains the version number of the program along with a brief description of what type of traffic that port uses (i.e. SNMP, SMTP, FTP, etc).

If you do not have the time to spend telnetting to each port you can search for vulnerability scanners. Vulnerability scanners will scan a network for vulnerabilities that are listed in its database. On the positive side using a scanner is much quicker than using telnet, on the downside the number of false vulnerabilities you receive is much higher.

Step 7: Search for Vulnerabilities

Now that we have a list of open ports and the operating system running on them, we can move onto the task of finding vulnerabilities for each program. There are several good places to start but I would have to recommend Bugtraq. BugTraq is a mailing list where people can post bugs or vulnerabilities found in different programs. You can also visit the Security Focus website and narrow your search as you look for vulnerabilities focused on specific operating systems. You can scan their database for previously reported vulnerabilities corresponding to software.

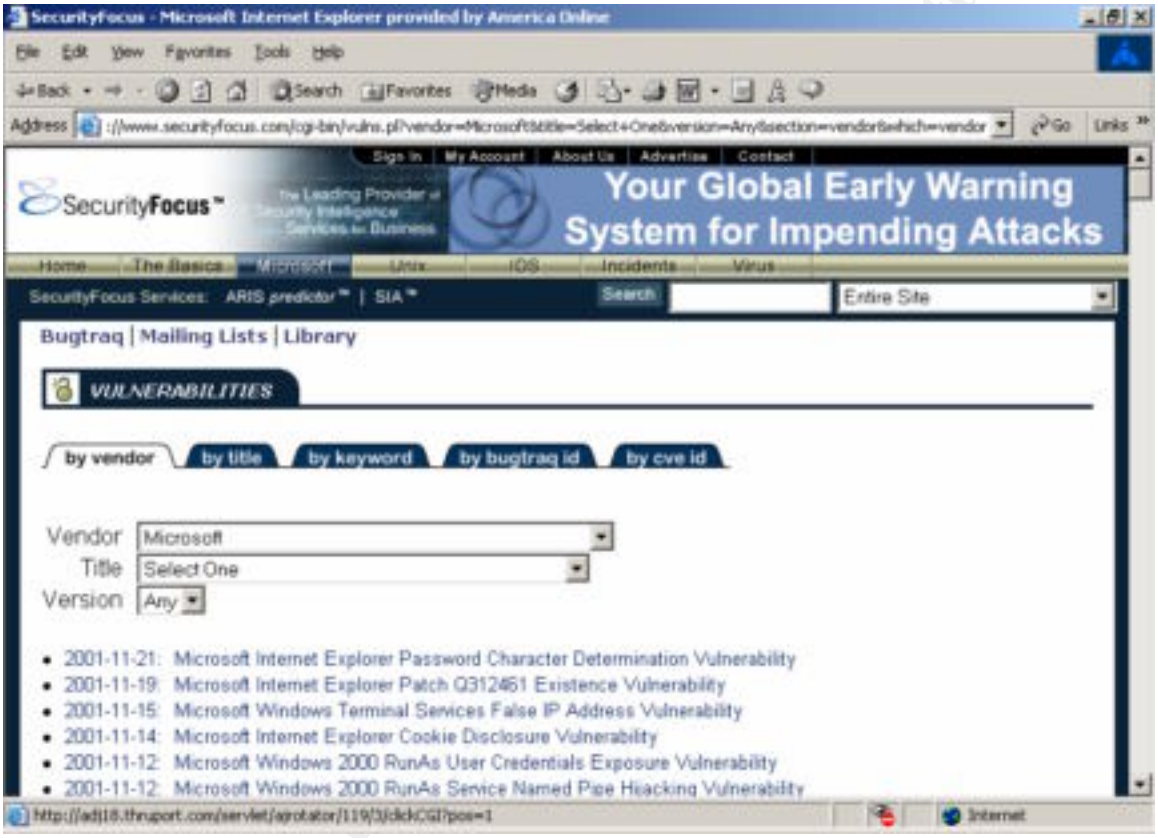

Another place you can look for vulnerabilities is the Google search engine. Google will direct you to a lot of sites that detail certain exploits for different systems.

Once you find your vulnerability you can go about testing your system to see how well it responds to the attempted exploit. Remember some systems will not respond to the vulnerability because they have already been patched. Other systems will respond to the vulnerability and grant you access. Another issue to consider is that most systems when exploited will only grants you access to the highest security level that they have. What this means is that if you break into a computer whose highest rights only deal with backing up data, then the highest power you will have are backup rights not system administrator rights.

Conclusion

While this is not meant to be a completely thorough look into information gathering it is meant to be a good introduction to the subject. It is my hope that you will take what you have learned here and apply it to your network. Use this knowledge to test and learn more about the vulnerabilities that can be accessed from within your own network, and from the internet. If you do not run these tests then you will never be able to really protect your network because you will not know what hackers know about your network.

Even after conducting these tests you really don't want to rest on that, you really should move further, try to understand more advanced concepts. Look into running UDP based scans, look at sending malformed packets and seeing how certain hosts react. Take unneeded services off of servers (don't just disable them, uninstall them). Remember; if you don't understand how your network operates then how will you ever be able to keep it secure.

Information is Power

Downloads/Websites

Sam Spade
URL: http://www.tucows.com/preview/195088.html

American Registry of Internet Names (ARIN)
URL: http://www.arin.net

U.S. Government Registry of Internet Names
URL: http://whois.nic.gov

U.S. Military Registry of Internet Names
URL: http://whois.nic.mil

European Registry of Internet Names
http://whois.ripe.net

AA Tools
URL: http://www.webattack.com/get/advadmin.shtml

Network Mapper (NMAP)
URL: http://www.insecure.org/nmap

Google Search Engine
URL: http://www.google.com

SANS Homepage
URL: http://www.sans.org

BugTraq Site
http://www.securityfocus.com/archive/1

References

[i] Zwicky, Elizabeth. Simon Cooper, Brent Chapman. Building Internet Firewalls. Beijing: O'Reilly, 2000.
[ii] Cisco: How NAT Works: URL: http://www.cisco.com/warp/public/556/nat-cisco.shtml
[iii] Network Ice: Port Knowledgebase: URL http://www.networkice.com/advice/Exploits/Ports/

Bibliography

Bhamidipati. "The Art of Reconnaissance – Simple Techniques."
URL: http://www.sans.org/infosecFAQ/audit/recon.htm

Cole, Eric. Hackers Beware: Defending your Network against the Wiley Hacker. Boston: New Riders, 2001.

Fydor. "The Art of Port Scanning."
URL: http://www.insecure.org/nmap/nmap_doc.html

Harris, Danny. Security Essential: TCP Concepts. SANS Institute, 2001.

McClure, Stuart, and Joel Scambray, and George Kurtz. Hacking Exposed: Network Security Secrets and Solutions. Berkeley: Osbourne/McGraw Hill, 1999.

Stevens, W. Richard. TCP/IP Illustrated, Volume 1. Reading: Addison Wesley Longman, Inc, 1994.