



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Protecting Your Corporate Network from Your Employee's Home Systems

Todd Rosenberry

GIAC Security Essentials Certification (GSEC)

Version 1.4b

Option 1

21 December 2003

Table of Contents

<u>Abstract</u>	3
<u>The Threat</u>	3
<u>Policy & Education</u>	5
<u>Management Involvement in Security</u>	5
<u>Show Due Diligence</u>	6
<u>Consistency</u>	6
<u>Security Awareness</u>	6
<u>Home System Security</u>	7
<u>Physical Access</u>	7
<u>Passwords</u>	8
<u>Anti Virus Software</u>	8
<u>Patching</u>	8
<u>Personal Firewall</u>	9
<u>Operating System Hardening</u>	10
<u>Compliance Check</u>	10
<u>Network Configuration</u>	12
<u>Home User Network</u>	12
<u>Corporate VPN Network</u>	13
<u>Corporate Firewall Configuration</u>	15
<u>Web</u>	16
<u>Email</u>	17
<u>File Transfer</u>	17
<u>Name resolution</u>	17
<u>Custom Applications</u>	17
<u>Egress Filtering</u>	18
<u>Intrusion Detection and Prevention</u>	18
<u>Summary</u>	20
<u>List of References</u>	21

Abstract

In addition to the protection provided by a strong perimeter firewall, a security conscious corporation will often have strict control over the systems placed on employee desktops. This may include anti-virus software, patch management, configuration management, and removing the ability for employees to install unauthorized software. Maintaining this level of desktop control is not trivial for the Information Technology (IT) organization within a corporation, but for company owned systems located on a corporate campus it can be done. The challenge becomes much greater when employee home systems are allowed to access the corporate network via a Virtual Private Network (VPN). These systems spend most of their time connected to the wild wild Internet and the rest of their time directly connected to your corporate network. In addition to work related activities they may be used for many other purposes by any number of people.

While direct control of home systems is not always possible there are steps a corporation can take to ensure that a worm or hacker does not gain access to the corporate network by compromising a home system. These include creating a written policy for how home systems should be configured and maintained, designing a secure network configuration on the corporate side and at the users home, deciding which services need to be available to home users while disabling any others, and specifically monitoring these access points to your network using Intrusion Detection and Prevention systems. Finally a VPN solution can be chosen which allows the corporation to enforce its security policies onto systems over which they do not have direct control.

The Threat

Any system directly connected to the Internet will be routinely probed for vulnerabilities. Attackers can quickly scan millions of systems and then take advantage of any vulnerability they find. From April 2000 to February 2001 the Honeynet Project attached a small network to the Internet and monitored all traffic coming in and out. The machines on this network contained no special data and were not advertised in anyway to the Internet community. There was no reason for anyone to access them so any incoming traffic was considered suspicious. All of these machines were regularly probed and attacked. During this period the average number of unique monthly scans rose from 103 to 206 which is over 7 scans per day. This is a 100% increase in the 11 month period and there is no reason to expect the trend won't continue¹.

The probes and attacks may come from mindless worms that have been turned loose on the Internet or from a hacker who is looking for systems which can be

¹ Honeynet Project, "Know Your Enemy: Statistics. Analyzing the past ... predicting the future", 22 Jul 2001, <http://project.honeynet.org/papers/stats/>

"owned" possibly to be used later in a Distributed Denial of Service attack. Worse still the probe could be coming from a hacker who has specifically targeted your system. Perhaps they are looking for sensitive information about your company and know that the corporate network is too secure to be cracked directly but that an employee's unsecured home system may provide a much simpler way in. A recent article entitled "Home Workers Are Giving Hackers Open Access To Business Networks" quoted a study by the NCC group showing that one in six personal computers were *completely* without protection². It is likely that many of the others still had one or more easily exploitable vulnerabilities.

A hacker or worm has a wide range of vulnerabilities to choose from when trying to take over a system. The CERT Coordination Center reports that the number of available security vulnerabilities and the number of security incidents each rose by a factor of 4 from 2000 - 2002. The number of incidents in 2003 is already on track to be nearly double last year³. The Symantec Internet Security Threat Report for September 2003 shows many frightening trends on the rise. These include the number of new vulnerabilities, the number of new viruses and worms, and the frequency of attack. At the end of 2003 we are now averaging 60 new vulnerabilities each week. The report also shows that the majority of attacks are against easily exploitable services that are likely to be running on home user systems.⁴

In addition to unpatched vulnerabilities many home users are likely not to use other basic tools of defense such as anti virus software or personal firewalls. One of the most obvious places to attack a corporate network is at the point that it touches the Internet. This is obvious to the white hats as well as to the black hats, so most corporations have strong firewalls in place at this perimeter to tightly control and monitor the traffic that goes in and out. This can lead to a fortress-style defense where perimeter walls are fortified, but if someone does get inside, little is done to prevent them from doing damage. Allowing a home user to connect to your network is like constructing hundreds of backdoors into your fortress, and it would be foolish not to attempt to guard them with the same ferocity as the front door.

The latest worm may continuously beat on your firewall but if it has been well configured the worm will probably not get through. A hacker is also likely to give up quickly on this method of attack, but home systems provide an attractive

² NCC Group, "Home Workers Are Giving Hackers Open Access To Business Networks", Jul 2003, <http://www.itsecurity.com/tecsnews/jul2003/jul158.htm>

³ CERT Coordination Center, "CERT/CC Statistics 1988-2003", 17 Oct 2003, http://www.cert.org/stats/cert_stats.html

⁴ Symantec Corporation, "Symantec Internet Security Threat Report Sees Increase in Blended Threats, Vulnerabilities and Internet Attacks", 01 Oct 2003, <http://www.symantec.com/press/2003/n031001.html>

target for both. Once an Internet connected system becomes infected by a worm it will attempt to infect other systems. As soon as this system connects to the corporate network it will begin trying to infect your other machines. This has become one of the most likely ways for an infection to enter a corporation. Similarly a hacker can take advantage of one of the many known vulnerabilities of an unpatched system to install something like a key stroke logger. Eventually this system will connect to corporate and all keystrokes will continue to be recorded. When the system is again available on the Internet the hacker can download the log and look for sensitive information including server names, user IDs, and passwords.

Policy & Education

The first step in getting your home users secure is to create a policy for how home systems should be configured and the second is to educate them about this policy. This section will deal with the importance of a security policy and not specifically what it should contain. The contents will be implied by the sections that follow. General security awareness training for your users is also an important part of keeping your networks safe.

In his article "Introduction to Security Policies, Part One: An Overview of Policies"⁵, Charl van der Walt lists many of the benefits to creating a security policy. The following section lists examples that can be easily related to home systems.

Management Involvement in Security

An IT Security department should have specific policies revolving around day to day operations, but the overall corporate security policy should focus at a higher level. It should not be concerned with the capabilities of existing tools or known problems within the organization. Protecting a corporation's assets is the responsibility of upper management and that is where the policy should begin. There may need to be some technological input but technology should not drive the overall policy. Getting upper management involved at this level will raise the awareness and importance of security for the whole company. Once a policy is established and supported by upper management it gives the IT organization an additional tool to prevent attacks.

If your network becomes infected with a worm and you are able track its entry point down to a home system you may want to remove that system from the network until you are sure it has been cleaned and is safe from similar future attacks. This user may then scream that the IT department is stopping them from getting their work done, but if a policy has been established it gives the IT Security department some leverage to take action against users who don't follow

⁵ Van der Walt, Charl, "Introduction to Security Policies, Part One: An Overview of Policies", 27 Aug 2001, <http://www.securityfocus.com/infocus/1193>

it.

You may want to require employees to physically sign off on the policy before issuing them a secure token or password which allows them to connect via VPN. By defining what is and is not acceptable your users will not be able to claim ignorance if they are found to be in violation of the policy. Even when users know what they should and shouldn't be doing sometimes they need to see it in writing before it sinks in.

Show Due Diligence

The existence of a policy can help protect your corporation in legal matters even when the policy has been breached. In "Introduction to Security Policies, Part One: An Overview of Policies", Charl van der Walt wrote, "Because policy reflects the philosophy and strategy of your company's management it is fair proof of the company's intention regarding information security." If some customer information was lost because an attacker was able to enter the corporate network via an unsecured home system the existence of the policy shows that the company had thought about this and had intentions to keep it from happening. Of course the real goal is to make sure your policies are followed and no information is lost.

Consistency

An IT organization should supply the tools and information that home users need to stay protected. Supporting systems that are not set up or generally controlled by the IT department is difficult but a security policy standardizes the tools and processes a home user should follow if they wish to connect to the corporate network. This might include things like which anti-virus software to use, how to configure a personal firewall, and when to apply security patches. A beneficial side effect of security policies is that home systems will become easier to support as they begin to take on similar characteristics.

Security Awareness

In addition to education about security policies, users should be given a more general education regarding security awareness. As an example a common way that viruses infect and spread is via email. If a stranger handed you a hand grenade you probably wouldn't be tempted to pull the pin to find out if it was real. However many people can't resist launching a program that a complete stranger sends them via email. This is somewhat understandable because such email messages are getting more and more clever. One might appear to come from support@microsoft.com and claim to contain the latest patches for your system. Another might appear to come from sales@somestore.com and tell you your credit card has just been charged five hundred dollars and if you believe a mistake has been made you should view the attached file to straighten it out. The attached file in this case could easily be a virus or backdoor program designed to give an attacker future access.

Many malicious attachments will also contain a double file extension. A file may appear to be called harmless.txt and a user may know that it is probably safe to open a text document. In reality many Windows systems are set up to hide file extensions. The real name of the attachment may be harmless.txt.exe and an exe (executable) file is much less likely to be harmless. Making sure your users understand the basic tricks attackers use and creating an environment where they are not afraid to ask questions can substantially raise the level of security on your network.

Nearly all computer users see the incredible level of information and services available to them on the Internet but many may not be aware what a dangerous place the Internet can be. It's often only the security team that understands how often a random system on the Internet is scanned and attacked. Tricking a user with an email message or taking advantage of the latest software vulnerability are just examples of the many techniques available to worms and hackers. An educated user is less likely to fall for the tricks and more likely to keep his system protected. The importance of security awareness is greatly amplified for users on home systems who most of the time don't have a strong corporate firewall and security team to protect them.

Home System Security

It can be hard for an IT organization to control the configuration of an employee's home system. Many VPN services were rolled out at a time when the threat from worms and hackers was not as pervasive as it is today. Automated patching and configuration is difficult because the system is not always connected to the corporate network. The home system is usually owned by the employee and therefore they control what gets installed. This also means that the corporation won't have administrator level access to the system which could be required to do automated patching, virus updates, or configuration changes. There may be other users such as roommates or children who play games or download unknown software.

Although it can be challenging a corporation should use policy and when possible technology to enforce a secure system. What should your policy say about home system configuration? There are many references that fully detail the configuration of a secure internet connected system so the following section will just hit the highlights.

Physical Access

If an attacker has physical access to a home system it is unlikely that anything a home user can do will keep them from accessing the data it contains. The good news is that home systems are generally in a safe environment where it is unlikely that an attacker could get access without being noticed. Hopefully it goes without saying that employees should lock their homes when they are out. In the event that a physical break-in does occur your policy should require password protected screen savers for all systems. For extra security a BIOS, or

boot level password should be applied to stop an attacker from booting the system from their own media thus bypassing other security measures. The most popular web browsers will remember passwords as they are entered, so in the event that an attacker gets a hold of an employee's home system, these browsers should also be configured to store the passwords in an encrypted form and to require a master password to unlock the list.

Passwords

Information is often only as safe as the password used to protect it. Setting up a BIOS password can make your system tougher to crack but if the password is easy to guess it only gives a false sense of security. Passwords should not be written down and they should be changed with some regularity. Current best practices tell us that one way to create a secure password is to use the first letter of each word in a phrase, song lyric or any other easily remembered sentence. Punctuation and simple character substitutions can be thrown in to create a very secure password. Something like "4#a7ya,ofbfutcann" may seem impossible for a user to remember but when you see that it was derived from "Four score and seven years ago, our fathers brought forth upon this continent a new nation" it becomes much more realistic.

Anti Virus Software

You should maintain enough licenses for your corporate anti-virus software so that copies may be distributed to home users who will be connecting to the corporate office. Traditional Anti Virus Software works by maintaining a portfolio of known viruses and when it sees one it stops it from doing anything harmful. It is very important that the software is receiving regular virus updates. AV software is a necessary piece of the puzzle but it is still just a piece. Internet worms have appeared that can spread across the globe in a matter of minutes. It's unrealistic to think that an AV company could ever get a virus signature written and applied to all systems in that time. However once the signature has been applied it will help to control the spread of the virus and the damage caused.

Patching

Your security policy should require employee's to apply critical security patches to their systems within a reasonable amount of time after they are released. There are a variety of products on the market to make sure all your desktop systems have the most current patches but such systems can be problematic when dealing with remotely connected users. These systems are generally under the employee's full control and there will be a very wide range of hardware and software versions to support. The systems will only be available at unpredictable times when the user connects, and once the connection is established its quality may not be predictable and the connection could drop at any time. Finally there may be legal issues with patching systems the corporation doesn't own.

For these and other reasons patching will most likely be left up to the end users. An easy way for users to keep their Microsoft systems current is to properly configure Windows Update. Most Microsoft home user systems will be shipped so that they regularly check the Microsoft site for new security patches. When found the user will be prompted to download and install them. Your policy should call this out as a requirement. Microsoft's current strategy is to release patches monthly so this shouldn't end up being much of an inconvenience to your users.

Linux or other UNIX based users will need to work just a little harder. They may need to periodically check known web sites for updates or get on a mailing list to stay on top of the latest patches. The IT Security team should take it upon themselves to maintain a web site that contains the latest security patches and send notifications when a critical patch is released or when a new worm hits the Internet that exploits a known vulnerability. The Security team can also build packages or write wrapper scripts to make the installation of patches as easy as possible for the home users.

Personal Firewall

Home users should be required to install a host based personal firewall. This is an extra program that will inspect traffic as it enters or exits the PC and make decisions on what will be allowed. As with anti-virus software the corporate office should provide the software to home users as part of a "remote access package". Doing so makes the use of a firewall much more probable and provides some consistency to home user configurations.

One personal firewall product designed with the Microsoft home user in mind is ZoneAlarm from ZoneLabs. It will prompt the user to allow or deny traffic that it sees and learn a little more with each answer. Eventually it will be silently blocking unwanted traffic while requiring little to no user interaction. Linux products are also available including the built in and freely available combination of netfilter and iptables.

By default most personal firewalls will block incoming connections, and for the majority of systems connected to the Internet there will never be a reason for an unknown machine to initiate a conversation. One exception is the corporate VPN concentrator. A personal firewall will have to be set up to trust this device so that it can initiate communication such as key exchanges. Worms and scanners will constantly be checking out home systems looking for vulnerabilities to exploit but if this traffic never makes it past your firewall no answers will be given and the attacker will likely move on.

A personal firewall is a necessary part of a defense in depth strategy but it is not a magic bullet. A user who is not fully trained in security will likely be the one who sets up and maintains it. Virus writers are tricky and may get a user to allow malicious traffic by making it look legitimate. The home system is out of your control and it may for example be running a game server while connected to the

Internet. The firewall configuration could get in the user's way, and they may decide to turn it off altogether. This should be forbidden in your policy and you may want to go so far as to say that no system that can connect to the corporate network will be allowed to provide services to the Internet.

Operating System Hardening

In general end user systems are shipped so that the user will be able to quickly set up and use them and not so they will be secure. An IT department should make sure that they have tightened the security on all the desktop systems they control but this is pushing the limits of what can realistically be expected of home users. If mistakes are made during this process the user may not be able to recover and the IT staff may not have the resources to diagnose the wide range of potential problems.

You may want to provide highly technical information for your savvy users but in general this part of security should be kept simple. The first step in locking down an operating system is to shut down all services which are not being used. Many operating systems are shipped with a wide variety of services running just in case the user wants to take advantage of one of them. For example there is no reason to be running a web server if the system does not serve any web pages. Microsoft's Internet Information Server and to a slightly lesser degree the Apache web server that ships with Linux have historically been ripe targets for attacks that can lead to remote takeover of a system. If the web server is not needed, don't run it (and if it is needed make sure it's patched!). Another common point of attack is unprotected file shares. If a home wants to share files with other systems on their home network, they should be sure they are password protected and that their personal firewall only allows connections from other systems on their network.

Many security vendors offer tools which will help users make the changes to tighten up their configuration and free tools are available at the Center for Internet Security's home page (<http://www.cisecurity.com>). Even though these tools can simplify the process it may still be overwhelming for a novice. An IT organization can't predict the configuration or use of a home system so this part of home system security is especially hard to automate.

There are many technical references that will describe further configurations that can be made to increase security but for the average home user anti-virus software, a personal firewall, up to date patching and common sense are all that can realistically be expected. Examples of guides for Microsoft operating systems can be found on the UK Security Online web site.⁶

Compliance Check

The previous sections talked about the importance of policy and education and

⁶ UK Security Online, "Home User Self-Defense Guide", <http://www.uksecurityonline.com/husdg/>

some of the elements your policy should contain with regard to home systems. This is very important and valuable but we know that the policies will not be followed 100% of the time. If an employee is having a problem while on the Internet and they find that by shutting down their personal firewall the problem goes away, they are likely to leave it shut down indefinitely without a second thought. Luckily a security engineer has the ability to choose a VPN solution that can enforce parts of their policy. This is a relatively new area but solutions already exist that can check whether certain programs are running on the home system before the VPN connection is established.

Security administrators may not have much control over the software running on a home system but they can usually enforce what VPN client is used to connect for the simple fact that usually only one will work with the installed VPN server. Many of today's most popular VPN clients have a personal firewall built in including those from Cisco and Symantec. A standard personal firewall policy can be preconfigured before the software is released to home users. This can be tricky because you can't predict all the uses for a home PC and you don't want your users to try and shut down the firewall to get around their issues.

Another limitation is that the firewall is only running while the VPN software is running and the most important time for a firewall to be running is when the system is connected to the Internet. However even if you have a user who shuts down their firewall while connected to the Internet you do get some protection by enforcing that it be running while connected via the VPN. Personal firewalls typically have rules about what can leave a system as well as enter it so if a home system becomes infected there is a good chance that the personal firewall will stop the infection from leaving the system and spreading into the corporate network.

VPN products are also emerging that can enforce a broader range of home system policies. There are products currently available that can check for the existence of a specific set of applications running on the home system before allowing a connection to be established. A VPN client and server from Cisco can be configured that will ensure a home system is running a personal firewall from Cisco, ZoneLabs, Network ICE or Sygate. Again these checks are only made while the system has an established VPN connection.

In the near future we can expect that VPN solutions will be able to check for a broader range of configurations before allowing a connection to be established. Symantec has announced a "client compliancy" initiative to achieve these goals.⁷ They will provide an Application Programmer Interface (API) that will allow them to partner with leaders in the security and networking industries to provide a wide

⁷ Symantec Corporation, "Symantec's New Client Compliancy Initiative Promotes Enforcement of Remote and Mobile Client Security Policies", 18 Nov 2003, <http://www.symantec.com/press/2003/n031118a.html>

range of possible compliancy checks before allowing a VPN connection. The entire security industry is likely to follow this trend towards interoperability. In Symantec's announcement, Chris Christiansen, the program vice president at IDC (a market intelligence firm specializing in IT) is quoted as saying "As customers increasingly look for client solutions, we expect the security industry to create industry standards for interoperability and ultimately, protection".

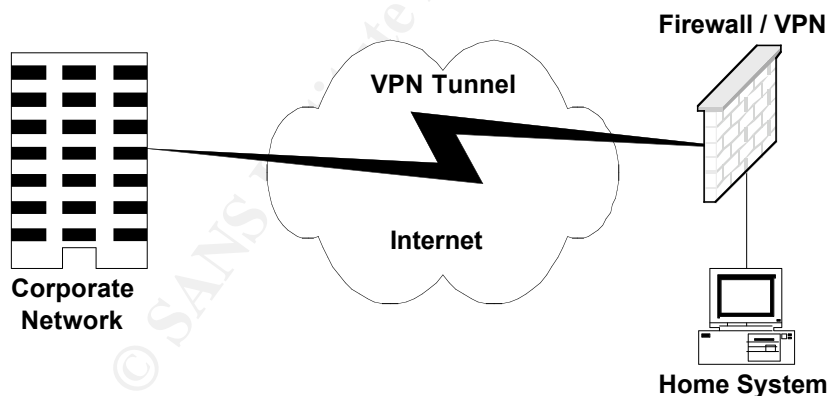
Any client software can potentially tap into a compliancy check engine. In addition to a check for a personal firewall and anti-virus software this could include checks for up to date patching, a securely configured operating system, or even a system whose critical systems files have not recently changed.

Network Configuration

The physical configuration of the VPN network at the end user and at the corporate sites will affect the options an administrator has for adding in security.

Home User Network

An advantage of a VPN is that it can make use of whatever Internet connectivity an employee has set up at home and allow them to use it to access the work place. When possible a secure configuration for the employee's home network should be enforced. One of the safest ways to allow an employee to access the data and services at the work place is to use a hardware VPN solution. A combination firewall router and VPN device can be configured by the IT staff and deployed to the employee's home network. If this device is configured to route *all* traffic back to the corporate network over the VPN then much of the risk presented by the home user has been removed.



In this scenario all traffic into or out of the home system will come to or from the corporation. It is as if the corporate network and the protection of the perimeter firewall have been extended to the employee's home. Even though the home system is connected to the Internet, firewall rules and network routing stop any other systems on the Internet from accessing it. The configuration of the firewall/VPN device can be locked down by the IT department such that only a very determined employee would be able to change it.

There are some drawbacks to this setup. All Internet traffic will be routed through the corporation so it is likely to be noticeably slower for the end user. Company resources will also be consumed when the employee or even a family member is browsing the Internet. A way to alleviate this is to allow a “split tunnel” configuration where all traffic from the home system that is not destined for the corporation will be sent directly to the Internet as opposed to traveling over the VPN tunnel. This again exposes the system to the dangers of the Internet and removes a lot of the benefit of this type of configuration. However extra care can be taken to configure the firewall to ensure that Internet based attacks will not be successful.

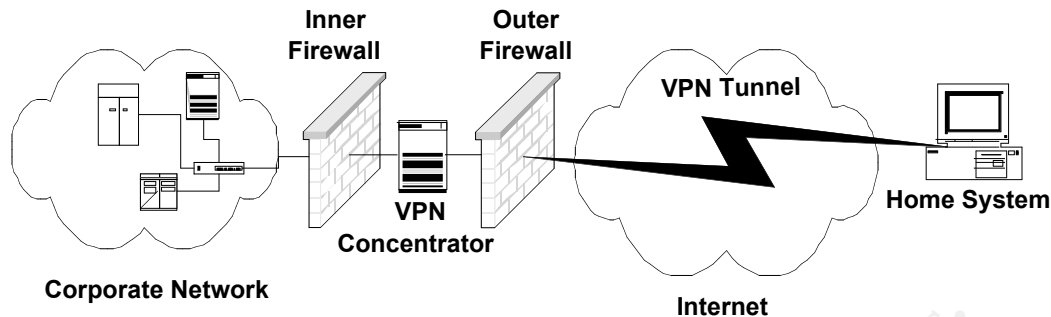
This configuration also opens you up to a different kind of attack. Worms and viruses will have a hard time entering your network through this route but if the connection is always live, a human attacker only needs access to the employee's home to gain access to the corporate network. This risk can be mitigated by configuring the VPN to request authentication via a password or hardware based token at regular intervals.

The employee will have access to the equipment involved and there is no guarantee that they will not physically bypass the hardware device when they want faster access to the Internet. They may also build a home network that has another connection point to the Internet which would effectively provide a full time backdoor into your corporation. Only policy and education can help an employee resist the urge to work around the approved setup.

Attempting to control the home user's network configuration is possible for telecommuters but for employees on the road it is not realistic to expect them to travel with all the necessary networking equipment. These road warriors will be expecting to be able to use their laptops and any available Internet connection to connect to the office and get some work done. The Internet connection could be provided by Starbucks, McDonalds or San Jose airport and the VPN connection will be made using software on the laptop as opposed to special hardware under a desk. In these situations it is even more important to take measures described in other sections of this paper.

Corporate VPN Network

How you set up your VPN infrastructure at the corporate office will affect its security. There are many all-in-one devices available but for this discussion we will look at each device separately. You may think that a firewall to protect your internal network from the VPN clients is unnecessary because you're talking about trusted employees on the other end of the line. Hopefully the information about worms and attackers has made you think otherwise.



The outer firewall between the Internet and the VPN concentrator protects the VPN device from attack. It would typically have a static configuration allowing only the protocols and traffic needed to maintain VPN sessions. The configuration necessary on this firewall will likely be explained in the documentation for your VPN concentrator. The inner firewall between the VPN concentrator and the corporate network will have a more complex and dynamic configuration that varies based on the services being provided to home users. In the real world the same firewall is capable of performing both functions but for simplicity two are shown.

The inner firewall gives you the ability to grant different privileges to different systems and users and will be the focus of the next section. If it is set up to do authentication it will be possible to define rules for specific users no matter what the IP address of their system is. If a VP desperately needs access to a sensitive application while on the road you can grant it without worrying about all your home systems having access. Also if you find that one of your home users has become infected with a worm the firewall will allow you to completely block that system from access until the infection has been cleaned. Adding an Intrusion Detection/Prevention System that has visibility to both sides of the firewall can also provide you with a wealth of useful information and protection as will be discussed later.

Before any traffic hits your firewall the session should be authenticated by your VPN concentrator. How this authentication is done will affect your overall VPN security. Like most applications adding additional security to your VPN will make it a little harder for your users to use. The least secure method of VPN authorization that is still reasonable would be to simply require a username and password. The biggest weakness here is that passwords can be guessed. If strong controls are not enforced, users will often pick a password such as a pet's name or anniversary date. If an attacker is specifically trying to crack your password they will gather whatever personal information they can. Even strong passwords are susceptible to a brute force attack where random passwords are tried one after another until one allows access. A password is also something that can be written down or easily shared with others if the owner doesn't take its security as seriously as you do.

The next level of authentication security is to use a secure token or certificate. This is called “two-factor” authentication because it requires you to physically have something (a token or digital certificate) and to know something before it can be used (a pin number or password to unlock the device). In the case of a SecureID token a piece of hardware is issued to each user that shows a random number which changes every minute. To authenticate, the user types in a pin number known only to them and whatever number appears on the token. Even if the token is stolen it can not be used without knowing the PIN number and the PIN number does no good without the token. The PIN can still be brute force guessed so the VPN server should be configured to lock out a token after its PIN number has been incorrectly guessed a set number of times. If a token or laptop is lost or stolen an administrator should permanently disable its access to the VPN server.

Similarly a digital certificate can be issued to a user which gets installed onto their home computer. It must be unlocked with a user defined password before it can be used. The mechanics of how the certificate works are well beyond the scope of this paper but the mathematical formulas involved ensure that a certificate can not be brute force guessed with today’s technology. If the certificate (i.e. a user’s laptop) is stolen it must still be unlocked with a password so strong passwords for certificates should be enforced. Two-factor authentication adds overhead for administrators that need to issue and track tokens or certificates and to the users who must now remember their password and make sure they have the hardware they need to make a VPN connection. The added security and peace of mind outweigh these inconveniences.

Corporate Firewall Configuration

Hopefully your employee’s home systems have been configured to resist attack but it is nearly impossible to be assured of this. In some environments VPN solutions were rolled out without any thought to enforcing a configuration on the home systems. One of the easiest ways to take control and enforce security across all home systems is to control what they are allowed to send in and out of the corporate network. This can be achieved using a firewall positioned between the corporate side of the VPN connection and the rest of the corporate network.

It is important that remote employees have enough access to do their jobs but it may not be necessary to give them complete access to the network as if they were sitting at a corporate facility. In general home systems are less secure than those on the corporate network and they should be treated as such. A security engineer should spend time figuring out what services are required for remote users to be productive and block all others. If we assume that these systems are usually connected directly to the Internet this becomes easier. When they do connect to the corporate networks all they need access to are things that will help them do their job. Everything else is available to them once they disconnect.

For example in August 2003 the Blaster worm infected over 1 million systems

and caused a large amount of lost productivity across the globe. This worm spread itself using TFTP (Trivial File Transfer Protocol). Many security engineers had configured their perimeter firewalls to block this protocol because there is almost no reason to allow a system on the Internet to access an internal system using TFTP. However many home systems became infected while connected to the Internet and then easily spread the infection to the corporate network once they joined it over the VPN. There are a handful of uses for TFTP within a corporate network but probably none that apply to remote users. By blocking this service the Blaster worm and any other threat that uses TFTP to spread itself would not have been able to sneak into the corporation using the VPN.

The example shows the need to block unused services but the right policy is to figure out what services are necessary and drop all others by filtering traffic through a firewall before it hits your network. In his article “Implementing and Managing a VPN Security Policy”, Stuart Broderick proposes a number of requirements when setting up VPN connectivity. Three that illustrate this point follow:

Policy Requirement: Business owners need to explicitly define their need for VPN connectivity.

Policy Requirement: Protocols and their supported configurations should meet but not exceed business requirements.

Policy Requirement: Organizations should mandate network security architectures and access controls that meet their business protection requirements.⁸

Once you decide which services and ports should be allowed it is equally important to determine the servers on which they run. You could allow all your VPN clients to talk HTTP to the internal network, but if an internal employee brings up an unpatched IIS server on their machine it can easily be used by an attacker or worm to invade your corporate network. You should only allow your VPN clients to talk HTTP to your known and trusted web servers. The same is true for all the services to be described below.

Another point is that any standard service can be made to run on a non-standard port. This is probably most likely to happen with web traffic. If you’re going to maintain tight control over your home users with firewall rules you should make sure you know all the servers and ports where important web services are available.

Web

There is almost no question that you will need to allow your home users to retrieve corporate information using a web client. There are an unlimited number of services including email, file download, employee self service applications and

⁸ Broderick, Stuart, PhD, “Implementing and Managing a VPN Security Policy”, 23 Apr 2002, <http://enterprisesecurity.symantec.com/article.cfm?articleid=1298&EID=0>

of course static web pages that can be served to users via a web browser. All of this is done using one protocol called HTTP, Hypertext Transfer Protocol (TCP port 80 or port 443 if encryption is enabled). Web browsers and the HTTP protocol have become so common and powerful that it is possible that this is the only access you will need to grant to your home users. Nearly all other services can be made run over HTTP. This can make your firewall configuration simple and seemingly restrictive but as a security engineer it should also scare you. By allowing this one port you are actually allowing an unlimited number of services across your firewall. A deep inspection firewall may be required if you are concerned with what is contained in the HTTP traffic you are passing.

Email

This is another service that your users will demand. Many companies rely on email for the majority of business communication. It's possible to enable this service in a web browser but if your users use a thick client such as Outlook or Mozilla to read email you will need to allow the popular email protocols to cross the boundary into your network. These are IMAP4, Internet Mail Access Protocol (TCP port 143), and POP3, Post Office Protocol (TCP port 110), SMTP, Simple Mail Transfer Protocol (TCP port 25) may be required in UNIX environments although it is unlikely to be a true requirement for home users. If you are in an environment where only one of these protocols is supported, only one should be enabled.

File Transfer

File transfer is well supported using HTTP but FTP, File Transfer Protocol (TCP port 20 and 21) is still popular as well especially for uploads. This protocol can end up using any port to send and receive web traffic but modern stateful firewalls are built to understand and control this. You can allow FTP traffic without worrying about having to allow a large number of random ports.

Name resolution

In order for other services to work your home systems will need to be able to find the corporate servers that provide them. The most popular way for one system to find another is to use DNS, Domain Name Service (TCP and UDP port 53) which maps names to IP addresses. To fully extend the corporate environment to the home you would also need to allow NIS, Network Information Service, in a UNIX environment and WINS, Windows Internet Naming Service, in a Windows environment. Each of these can complicate and weaken the policy on your firewall because they are not as straightforward in their port usage.

Custom Applications

There will almost certainly be other services that your VPN users will need access to but the thing to remember is that the services should be identified and analyzed before the access is given. Don't give your VPN users access to a wide range of services that they may not even need because it makes firewall

maintenance easier. Also don't grant every request without taking a close look at it. Your users may have a valid reason to access a specific internal web server but before it is granted you should make sure that internal web server has the latest patches.

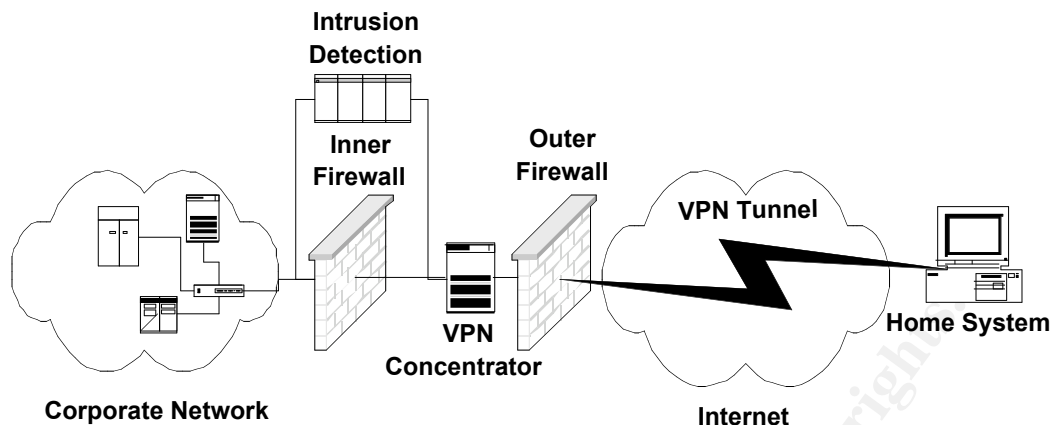
Egress Filtering

This is an often overlooked part of firewall configuration. A security engineer should pay attention to what can get out of the network in addition to what can get in. In a typical home user environment all connections will be initiated by the home user system so the firewall should be configured to block connections from being initiated in the other direction. If a worm does find its way in you should keep it from getting out while you work to squash it. This applies to the Internet as a whole and to your home users. Many of the corporate services supplied to home users may be unavailable until the worm is eradicated but there is no reason to directly expose home systems to the attack. Setting up this type of filter also cuts down on IP address spoofing. Your firewall should know which networks are plugged into which subnets and drop any traffic that shows up in the wrong place under the assumption that it is spoofed.

Intrusion Detection and Prevention

The goal of an Intrusion Detection system is to alert you to attacks and scans, possibly in real time. The logs generated by an IDS system will be indispensable when you are trying to figure out how an attack made it onto your network. IDS systems generally have two methods of detecting an attack. The first is to watch all network traffic for known signatures which is similar to how anti-virus solutions work. When an attack is detected you will know exactly what is happening and how to combat it because someone has already analyzed this attack so that a signature could be created for it. If the attack is brand new and no signature exists it will not be detected.

The second way an IDS system can detect an attack is by checking for protocol anomalies. For example in the HTTP protocol a URL is allowed to be 1024 characters long. If a client is continually requesting URLs that are a million characters long it may be trying to take advantage of a buffer overflow error in the web server to gain control of the system. At a minimum this traffic is not valid so there is no good reason to let it continue on to the web server. In the near future most IDS systems will likely use both methods of attack detection. The below logical network diagram shows where an Intrusion Detection device can be placed in your system.



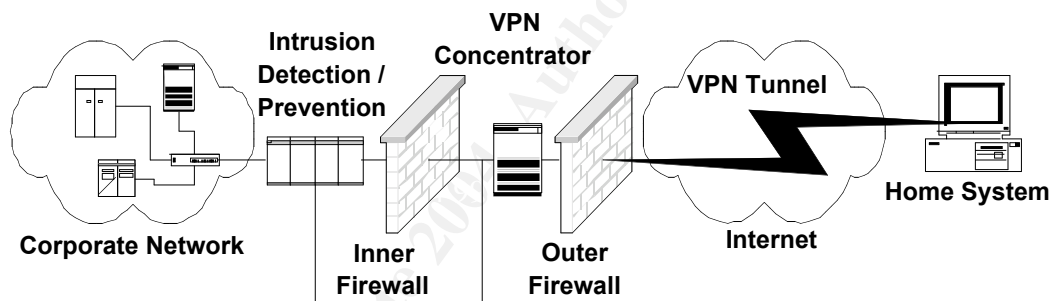
If you do use one device to monitor both sides of the VPN firewall you need to be careful that you haven't created an alternate path for packets to travel that completely bypasses your firewall. Some simple steps to take include not assigning IP addresses to the public side of the device. This will allow it to view the traffic as it passes but will not allow any device to make a connection directly to it. The private side will likely need an IP address so that it can be managed unless this is done via a third management interface. The system where the IDS system is running should be configured such that it never routes packets from one interface to another. If the IDS device is an appliance designed for intrusion detection it most likely will be designed such that it does not route packets. This is a simple configuration change for Windows and UNIX based systems. Great care should also be taken to ensure that your IDS system is one of the most secure on your network. If an attacker gets control of this system they may be free to launch whatever other attacks they want without being detected.

Placing an Intrusion detection system on the far side of your Internet firewalls may not provide you with much useful information. You already know that your network and systems are going to be constantly scanned and attacked. This has become so common on the Internet that it is just noise. Realistically there is not much you can do about it. However placing an Intrusion Detection system on the far side of your inner firewall will provide you with a wealth of useful information. Unlike the Internet example the only people that should be on the other side of the VPN firewall are fellow employees. If you see attacks or probes coming from them you should take action.

The most likely cause will be that a system has become infected while connected to the Internet and is now trying to infect the corporate network. The employee should be notified to disinfect and patch their system as soon as possible. This is a service to the employee who may have been wondering why their system was "acting funny". Using a signature based IDS system an administrator will most likely be able to tell the employee exactly what virus or worm has infected them and how to take care of it. This is also a great service to the corporate network. You were lucky enough to catch this attack but what about the next

one? The IDS has probably identified a home system that does not have adequate security. Unless it is taken care of it may eventually succeed in sneaking in an undetected attack.

Hopefully the configuration on your firewall has been locked down to the point that most of the attacks detected by the IDS do not get into your corporate network. If the latest worm uses TFTP to transfer itself to a new host and you are not allowing this traffic, you are safe, but for the reasons mentioned above you should still take action. Any service you are allowing your home systems to connect to can potentially be attacked. Your IDS system may detect an attack but if the firewall lets it in because it is attacking an allowed service then it has a chance at being successful. This is where an Intrusion Prevention System comes in. An IPS not only detects attacks but will stop them in their tracks. It will know the difference between a valid HTTP request and one that is attempting to exploit a known web server vulnerability. The invalid request will be dropped and depending on the configuration the host that made the request may be blocked from all future communication until an administrator intervenes. The network in front of the firewall should still be monitored to help find home systems that have been exploited.



One drawback of IPS systems is that they can become a network bottleneck. They must examine all network traffic and make decisions on what to let through. They must get much deeper into the data than a typical firewall would. For example one common way to get around an IPS system is to break an attack up into many small network packets. This way no one packet looks like an attack but when the final destination system receives and assembles all the data it may be in trouble. For this reason an IPS may need to hold a series of packets until it has a chance to look at them as a whole and then pass them on. Adding this overhead may not be practical at many points in your network but unless you have a very large remote workforce a reasonably sized IPS system should be able to keep up with the load and home users shouldn't notice much of a slowdown.

Summary

No system connected to the Internet is safe from attack. It doesn't matter how insignificant the system may appear. It will be hit with automated scans and

attacks on a daily basis. When an unprotected home system is also used to access a corporate network via a Virtual Private Network connection all the electronic functions of the corporation are put at risk.

There is no one solution that can fully remove this risk so the best way to approach it is with many layers of security forming a cohesive defense in depth strategy. The first step is policy and education. This lets home users know what is acceptable and gives the security department some leverage when policies aren't followed. In addition to basic security common sense the policy should layout an acceptable configuration for home systems that includes anti-virus software, a personal firewall and up to date patching. A VPN solution can also be chosen that will enforce parts of the policy by refusing connections when minimum requirements aren't met.

The VPN server is on the perimeter of your network so it should be designed such that a firewall is utilized to screen incoming traffic. Instead of blocking dangerous traffic a security engineer should take time to determine what should be allowed in and out and block everything else. An Intrusion Detection System should be deployed on the far side of the firewall to identify home systems that have become infected with a worm and to detect any suspicious activity. A firewall typically makes decisions based on what services are allowed and which side initiated the conversation but can not detect attacks to services that have been permitted. Therefore an Intrusion Prevention System should be deployed on the inside of the firewall to drop any dangerous traffic that has gotten through.

If each of these steps is taken the risk that an attack will jump from a home system to your corporate network will be greatly reduced.

List of References

Broderick, Stuart, PhD, "Implementing and Managing a VPN Security Policy ", 23 Apr 2002,

<http://enterprisesecurity.symantec.com/article.cfm?articleid=1298&EID=0>

The Center for Internet Security, <http://www.cisecurity.com/>

CERT Coordination Center, "CERT/CC Statistics 1988-2003", 17 Oct 2003,

http://www.cert.org/stats/cert_stats.html

CERT Coordination Center, "Home Network Security", 22 Jun 2001,

http://www.cert.org/tech_tips/home_networks.html

Cisco Systems, "CISCO VPN 3000 SERIES CONCENTRATORS (User Management)",

http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/products_configuration_guide_chapter09186a008015ce2d.html

Honeynet Project, "Know Your Enemy: Statistics. Analyzing the past ... predicting the future", 22 Jul 2001, <http://project.honeynet.org/papers/stats/>

NCC Group, "Home Workers Are Giving Hackers Open Access To Business Networks", Jul 2003, <http://www.itsecurity.com/tecsnews/jul2003/jul158.htm>

Netfilter, "Firewalling, NAT and packet mangling for Linux 2.4", <http://www.netfilter.org/>

Red Hat, Inc., "Red Hat Linux 9. Red Hat Linux Security Guide", 2002, <http://www.redhat.com/docs/manuals/linux/RHL-9-Manual/security-guide/>

Symantec Corporation, "Symantec's New Client Compliancy Initiative Promotes Enforcement of Remote and Mobile Client Security Policies", 18 Nov 2003, <http://www.symantec.com/press/2003/n031118a.html>

Symantec Corporation, "Symantec Internet Security Threat Report Sees Increase in Blended Threats, Vulnerabilities and Internet Attacks", 01 Oct 2003, <http://www.symantec.com/press/2003/n031001.html>

Tanase, Matt, "The Great IDS Debate: Signature Analysis Versus Protocol Analysis", 05 Feb 2003, <http://www.securityfocus.com/infocus/1663>

UK Security Online, "Home User Self-Defense Guide", <http://www.uksecurityonline.com/husdg/>

Van der Walt, Charl, "Introduction to Security Policies, Part One: An Overview of Policies", 27 Aug 2001, <http://www.securityfocus.com/infocus/1193>

© SANS Institute