Global Information Assurance Certification Paper

A Lightweight Personal Firewall

That Competes With the Heavyweights.

GIAC Security Essentials Certification

1.2f

James Manion

Introduction

As hacking becomes easier, more and more attempts will be made to access your office workstation as well as your personal home machine. Anti-virus software can detect those attacks that come in the form of an email attachment or a rogue file you downloaded off of the Internet that unknowingly contains malicious code. Unfortunately, anti-virus software will not protect your personal system from the direct attack when you have become specifically targeted. These attackers scan the Internet community looking for computers with unsuspecting ports and services running and attempt to exploit these communication lines. Your workplace computer is probably protected by a firewall; however, this protection is transparent to the user because it resides on the outskirts of the corporate network stopping these malicious probes and attacks. The casual home user does not have the luxury of a large network at home and must initiate a firewall protection system in order to safeguard their computers.

Background on Firewalls

Firewalls come in all shapes and sizes and offer a myriad of bells and whistles to protect your computer from reconnaissance and attacks. Essentially, if you look at the fundamentals of Internet communication you will see the basic element in protecting your system is to control the different types of inbound and outbound communications as well as the ports or open lines these communications are using. Using the telephone system as an analogy, each Internet communication consists of a source address and a destination address. Computers use IP addresses as an Internet “phone number” to communicate. But just like a radio station has one phone number but multiple phone lines for that number, your computer has over 65000 potential “phone lines” that it can communicate different kinds of data at the same time. This data can be in the form of streaming music, web pages, streaming stock quotes, or instant messages all concurrently communicating with your computer on the different phone lines or ports.

So to continue with the telephone analogy, there are phone numbers out there you may want to block. These numbers could be inbound phone numbers you do not wish to receive calls from or even

outbound phone numbers you want to block all calls made to those numbers. With a firewall you can block all communications to another system as well as inbound communication attempts from systems on the Internet. Now when you block an address it blocks all communication, but some specific communication between your system and the other system may be desired. You do not want to open up all the different phone lines to that system just specific lines (ports) carrying specific types of traffic. For instance maybe you want to allow inbound and outbound communication to a website but disallow other types of traffic between these systems. You want to permit communication between these addresses but limit it to a specific port. This port would probably be port 80, which is a well-known port for web servers.

When data comes through the Internet it is sent in small chunks or packets. These packets contain other information aside from the actual data. This is similar to the postal system. You send and receive mail all the time, however the mail comes in an envelope or package that contains the mailing information such as the sender and receiver information. Firewalls can act as the mail clerk who can filter out the junk and unsolicited mail from the “good” mail. There are several ways the firewall can filter this mail. It can look at the envelope and based on the sender/ receiver information it can choose to forward or reject the message. The firewall can also look at the type of message and make a forwarding decision based on the type of data being sent i.e email, web, and instant messaging traffic. The firewall can also look deeper into the package and make a forwarding decision based on the contents of data. As you can guess, this is very time consuming and can take special hardware and software to perform this type of inspection. However, today there are inexpensive and even free programs out there that allow the home user to utilize some of the above filtering techniques.

In comes AtGuard

(Figure 1.0)

AtGuard was born in the late 1990s by WRQ who make software for the corporate world.

AtGuard sticks to the basics which revolves around preventing applications from communicating with your computer or other remote devices without first alerting you that this communication is being attempted. Once alerted, you decide whether to permit or deny this type of traffic. These remote devices can be located on your network or even out in the vast world of the Internet. More and more people are getting connected to the Internet from their homes so this increases the list of potential targets an attacker can choose from. With the advent of high-speed Internet access technologies like cable and DSL modems, the attacker can scan the Internet for potential victims much faster.

AtGuard is a lightweight personal firewall geared towards the home user and does not require expensive hardware or software to run. AtGuard will run on a PC with Windows 95, 98, ME,XP, NT4.0 running Service Pack 3 and even Windows 2000. There is also a Macintosh version however it does not contain all the features that the Windows version contains. AtGuard was originally created by WRQ as a shareware product; however, Symantec bought the AtGuard technology and incorporated it into Norton Internet Security. You can still find the free WRQ version around on the Internet today and it even ships with some shareware CD-ROMs. Once installed and the machine is rebooted, AtGuard is working automatically. An icon in the system tray shows that AtGuard is enabled and configuration is easy by just right clicking on the system tray icon. Figure 1 shows the system tray icon, menu and dashboard

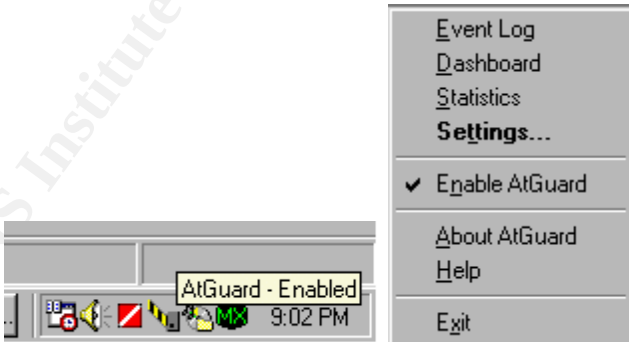

(Figure 1.1)

Here is a quick summary of the menu options:
Event Log - Displays a log of all activity the firewall has performed including the traffic blocked
Dashboard – Displays above window that can be docked at the top of the screen. You can enable and disable different components of AtGuard as well as view current stats
Statistics - Displays a window with various counters on it. Note that you can clear these by right-clicking in the window.
Settings – Displays the AtGuard settings where you can turn various blocking features on and off, add strings to the blocklist, add individual domains to allow cookies to, modify firewall rules, and password protect @Guard's settings. See Figure 1.2
Help – Will bring up the AtGuard online help.

Fighting the Bad Guys with AtGuard

There are lots of different types of traffic out there on the Internet and depending on what applications a user executes, different types of traffic can be generated. Using AtGuard the user can choose which rules to activate.

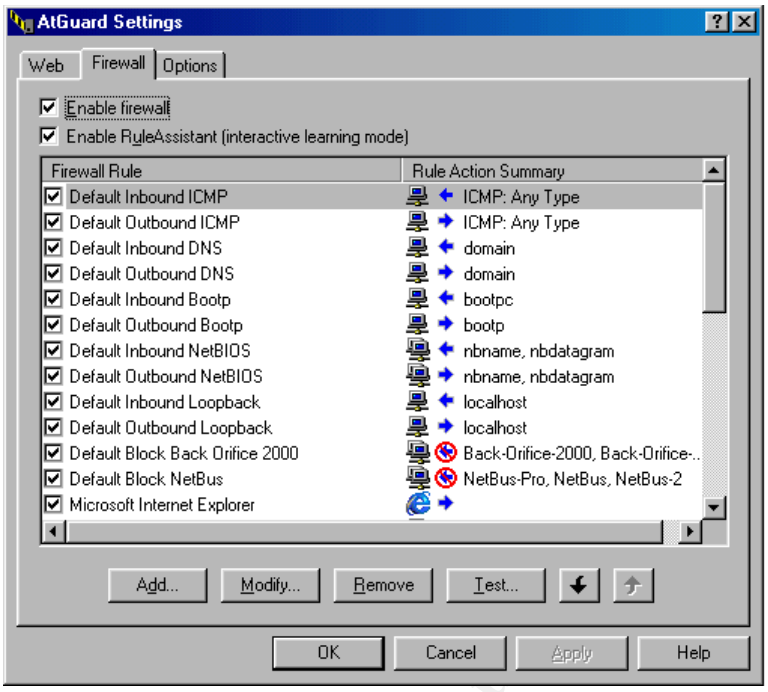

(Figure 1.2 AtGuard Settings)

AtGuard offers a RuleAssistant that guides the user through the configuration process. When AtGuard is running in Learning mode, each time a new type of traffic that doesn't have a corresponding rule attempts to start a connection, AtGuard will ask you to define a rule. Once a rule is defined, traffic matching that rule will either be permitted or denied into or out of the system. With this in mind, when you run AtGuard for the first time there are a lot of rules that need to be set to accommodate the most common types of

traffic and communication that is used with common applications. Web Surfing and email are the main types of traffic, so the first thing you need to do is to configure AtGuard to permit these types of traffic.

In order to configure AtGuard for web traffic, we need to define inbound and outbound communications. Since we will be using Internet Explorer to surf the web we will allow all outbound services from Internet Explorer. We will then configure AtGuard to allow inbound HTTP or web traffic.

Figure 1.3 shows the rule set for outbound TCP traffic produced by Internet Explorer. After assigning a name to the rule, using the drop down menus we can configure the Action (Permit, Block, Ignore); Direction (Inbound, Outbound, Either); and Protocol (TCP, UDP, TCP or UDP, ICMP). AtGuard gives you the choice of having this rule apply to a specific application, service or address, or even let it apply to all applications, addresses or services. In the example shown below, Internet Explorer was the specified application, however, the rule will permit outbound TCP communication from Internet Explorer to "any address" and "any service". Since Internet Explorer is a trusted commercial application it is ok at allow such a "wide open" outbound communication rule. Now if this was a questionable application, further research should be done to determine the type of communication the application should permitted to perform.

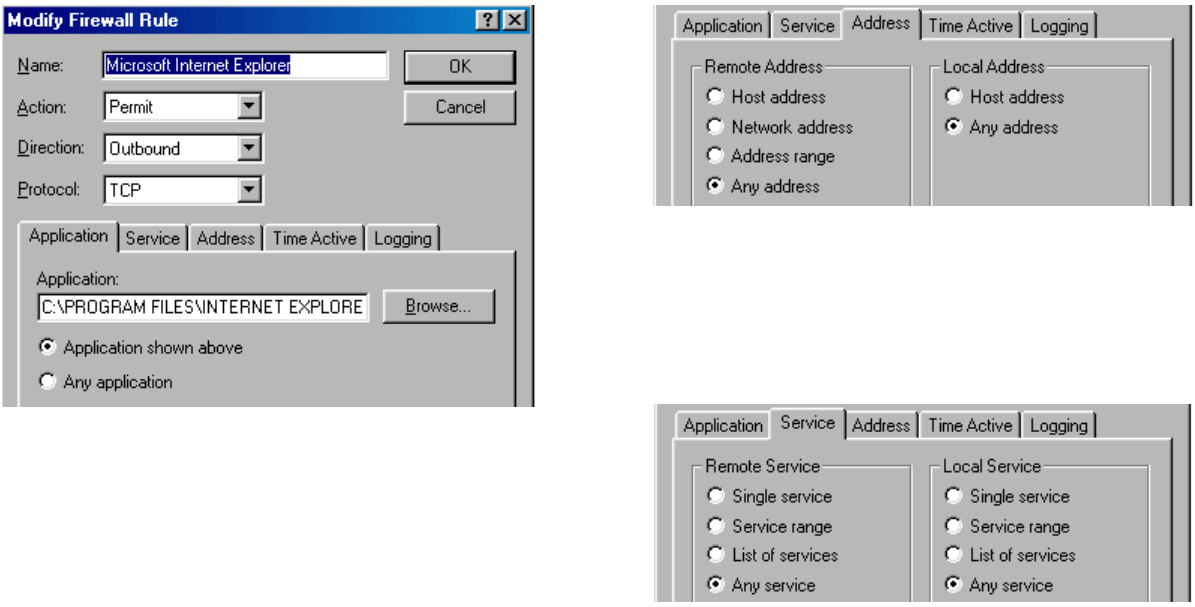

(Figure 1.3 Configuring IE for Outbound TCP Traffic)

The previous rule takes care of the outbound traffic when Internet Explorer to used to surf the web but Internet communication is a two-way street and the incoming packets must also have a rule. Since other applications use HTTP or web traffic this rule can be configured a little differently. Figure 1.4 shows the configuration settings for allowing inbound HTTP traffic. To date, traffic using the HTTP service is usually innocuous, however, a rule is still required. First, a name is given to the rule for logging purposes. Then we configure the rule to “Permit” “Inbound” “TCP” traffic for “Any Application” using the “http” service originating from “Any Address”. Table 1.0 below shows other common traffic that should be allowed by AtGuard. You can decide which application you commonly use in order to associate the program with the specific types of traffic. Separate rules will need to be created for each service.

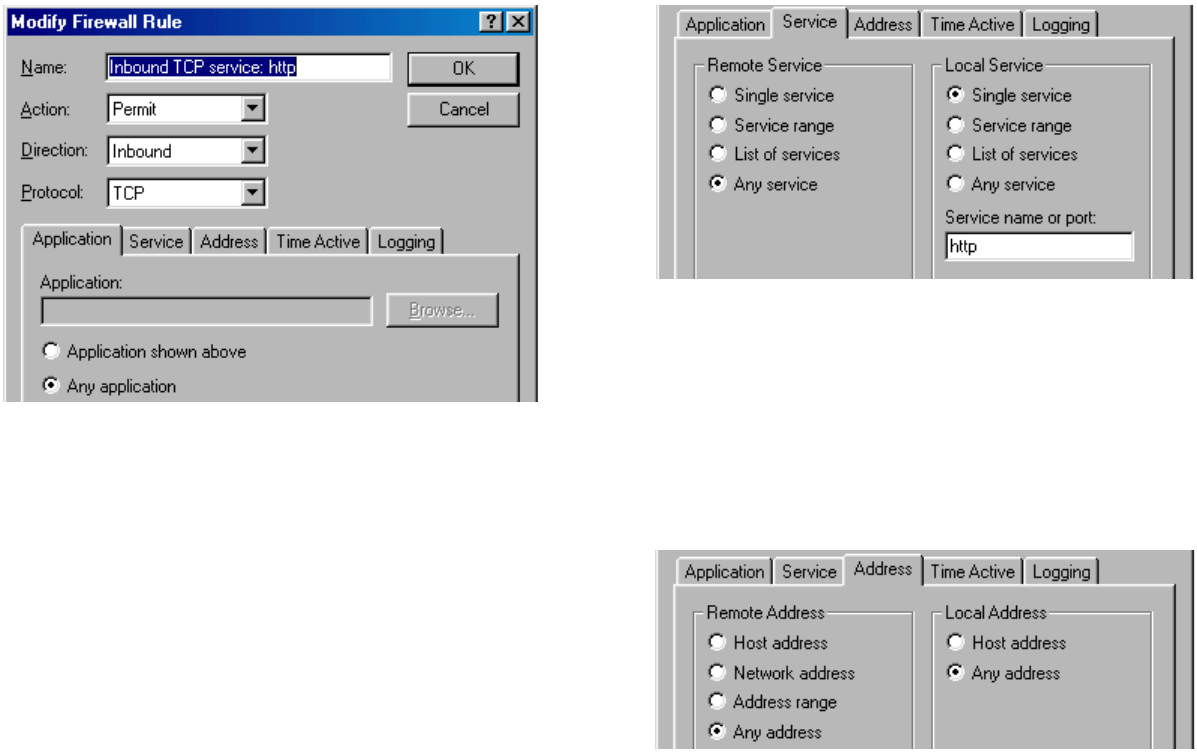

(Figure 1.4 Configuring HTTP for Inbound TCP Traffic)

HTTP Web server	80/tcp
DNS Server	53/udp
FTP server	21/tcp
IMAP3 Email server	220/tcp
IMAP4 Email server	143/tcp
SMTP Email server	25/tcp
POP3 Email server	110/tcp
Remote desktop (Terminal Services)	3389/tcp
SSL web server	443/tcp
Telnet server	23/tcp

(Table 1.0)

After the initial configuration of AtGuard, you will be prompted from time to time to create a rule. Whenever unknown traffic tries to enter or leave your system, AtGuard will prompt you to create a rule. AtGuard will implicitly block all traffic unless there is a rule association to that address, application or service. Figure 1.5 shows the RuleAssistant prompting the user to create a rule for an Inbound TCP DNS

packet. The IP address of the inbound packet has been masked in order to not display any public address. All the black boxes on the images represent the host IP addresses of the connecting inbound traffic. Now back to the RuleAssistant. After the first prompt by the RuleAssistant, the user can choose to block or permit this connection for just this one attempt and will be prompted for all future attempts or a rule can be created to block or permit this traffic for all future attempts. If a rule is desired, the RuleAssistant walks the user through the decision process by asking the user to specify settings for the service and address. Finally, a summary screen is shown (Figure 1.6) before the rule is created. Once the rule is created it is instantly activated and the traffic is filtered accordingly.

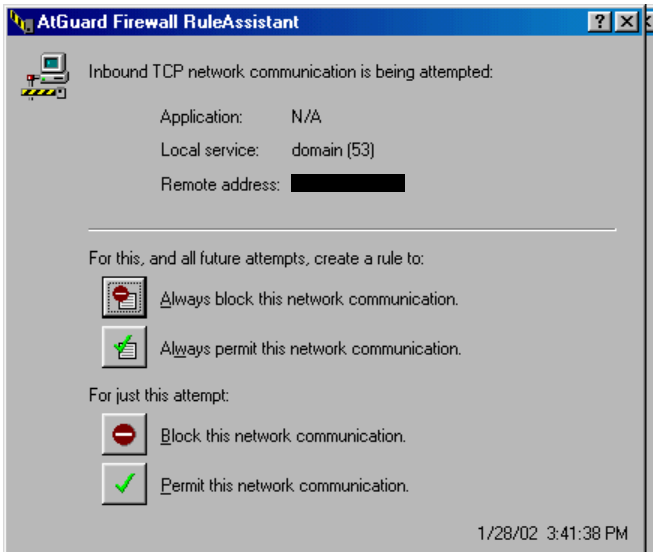

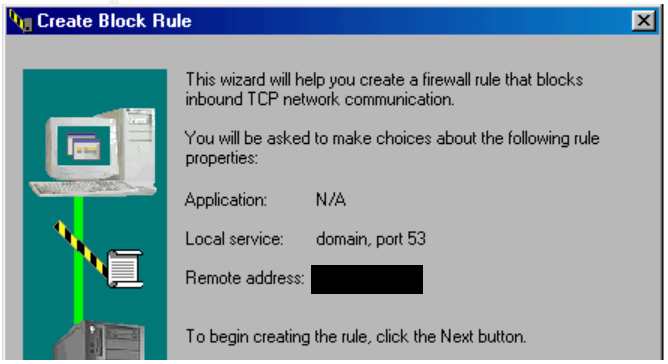

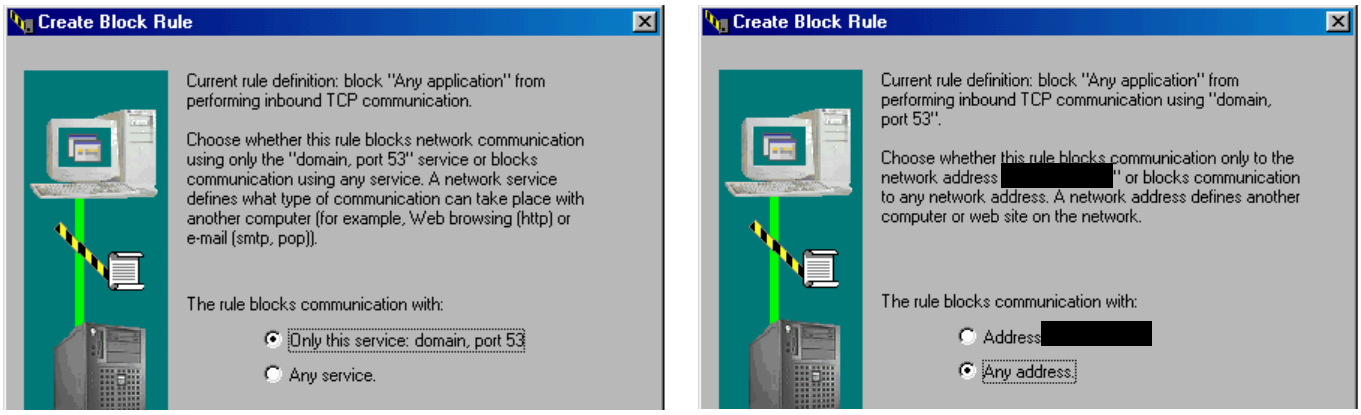

(Figure 1.5 Configuring DNS Traffic)

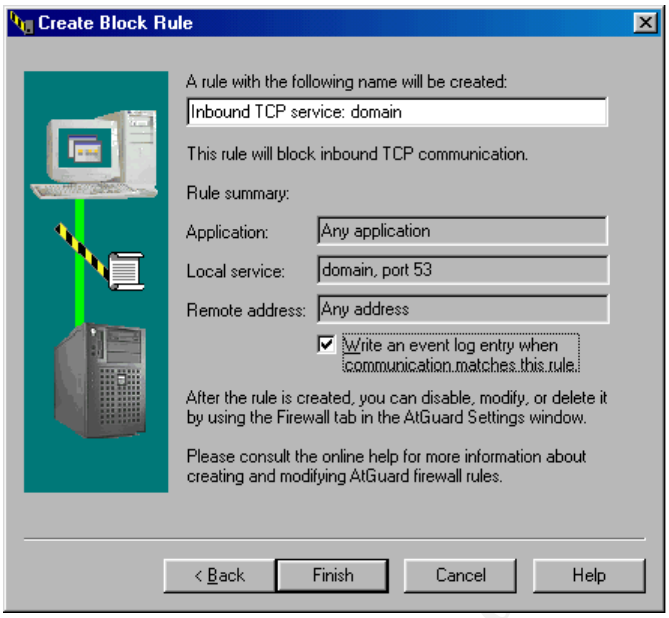

(Figure 1.6 Summary Screen for DNS rule)

Example of Blocking Potentially Malicious Traffic.

As mentioned earlier, from time to time attackers will inevitably scan you. AtGuard will detect these attempts and ask you to create a rule. And just like the DNS example above, you can chose to permit or block this traffic. You can also specifically block the individual from other types of attacks. Many hackers will scan a system for multiple open ports to exploit so blocking traffic from this address may be warranted. As shown below, Figure 1.7 depicts an inbound attempt of a TCP packet destined for port 1243 a port commonly used for the SubSeven Trojan. AtGuard was able to detect this incoming packet and prompt the

user for the desired action. Since this is not a commonly used port for any other traffic except backdoor Trojans, it is best to block this port from all addresses.

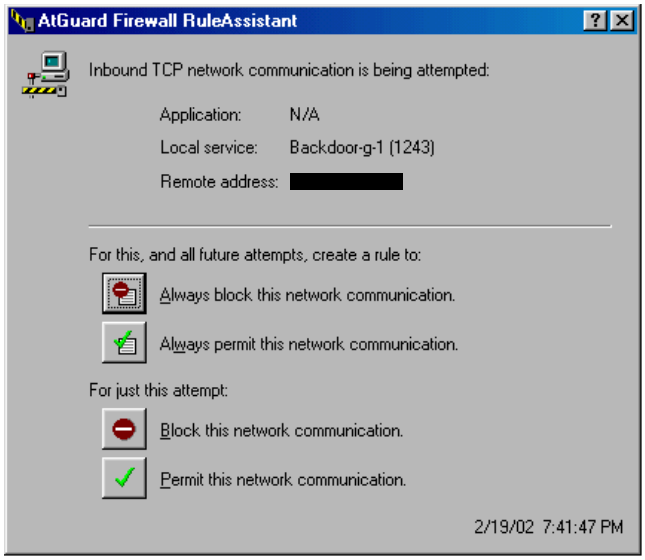

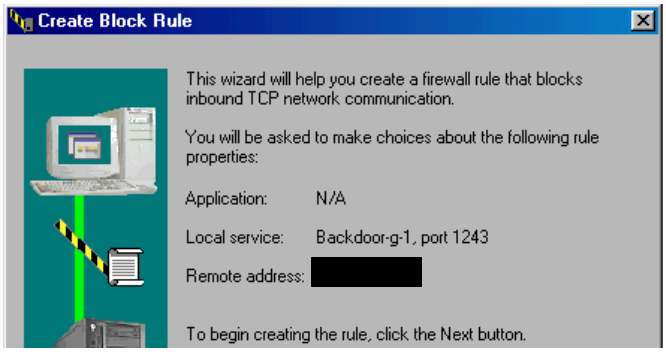

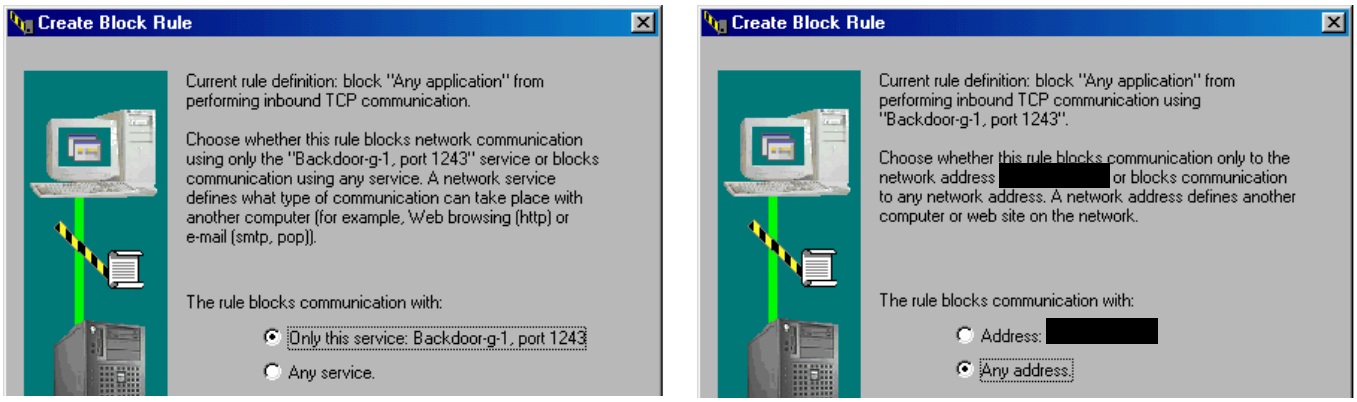

(Figure 1.7 Blocking SubSeven Trojan)

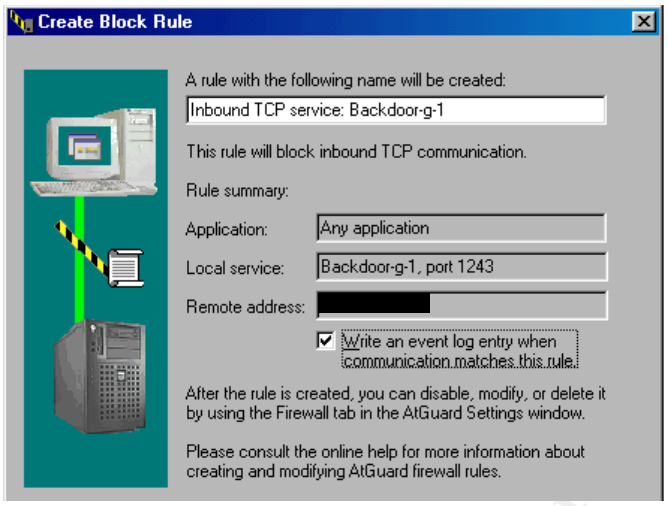

(Figure 1.8 Blocking SubSeven Trojan summary.)

What else can AtGuard do

Now lets take a look at some of the other features of AtGuard. Through the use of a simple GUI you can configure the firewall to block banner ads and cookies. The blocking of this excess traffic can speed up your web page viewing by blocking superfluous traffic. Cookies are bits of information that web servers store on your computer for their later use. Web servers can use cookies to keep track of how many times you've visited and when, what sort of info you have been surfing for on their site. Web servers can even use cookies to pass your information on to other web servers, such as advertisement servers. On the positive side, cookies can be used to store your own web site configuration, to remember items placed in your "shopping cart" at an on-line shopping site, or to store account and password information for subscription sites. You may not want to block ALL cookies, so AtGuard uses a cookie "allow" list. Cookies

are blocked on the way OUT of your computer, NOT on the way in. Incoming cookies are accepted, but the information that they contain is not allowed to be sent back to a web server unless you explicitly put the domain name of the server into the cookie allow list. A number of cookie-blockers work this way, perhaps because it is easier to implement. There are several ways for web servers to set cookies on your computer, but there is only one way that browsers give cookies back to web servers. If they're blocked on the way out, the blocker catches all of them. To permit cookies for a specific site, just add the site's "domain name" to the cookie allow list.

AtGuard also offers an easy to view statistical manager that lets the user view network connections as well as statistics about inbound and outbound communications. You can also view a history of the rules that have been used and the number times a packet has been blocked or allowed to pass through the firewall.

Summary

AtGuard gives the everyday common user, the ability to configure a firewall to intercept both inbound and outbound connection attempts and decide whether to allow or deny these attempts based on a list of rules that you define. A RuleAssistant acts as an interactive wizard to guide you through making rules for specific types of traffic, services or addresses. Having a firewall on your home system can ensure that no data will be sent across the Internet without your knowledge and consent. A firewall can also protect you from malicious attacks by catching the attempts during the attackers reconnaissance. With the configuration of a personal firewall being this easy, there is no reason why everyone should not be using a firewall on his or her home computers.

Supplement - Weird Traffic

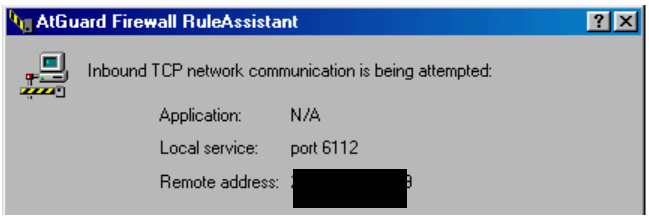

Scan for Battlenet Game Ports

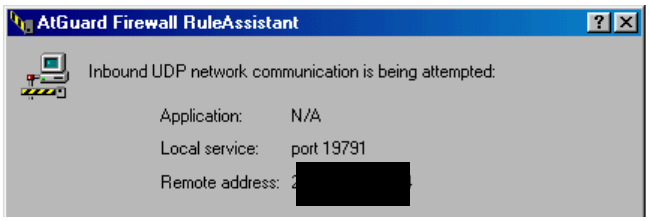

Port 19791

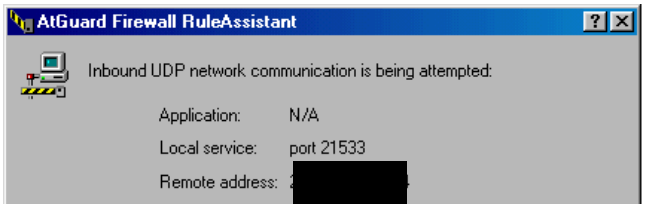

Free BSD Port

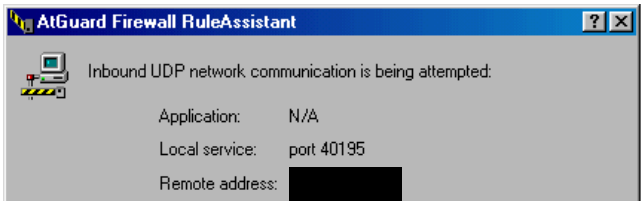

Port 40195

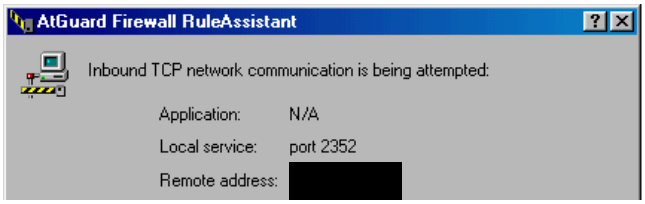

Port 2352

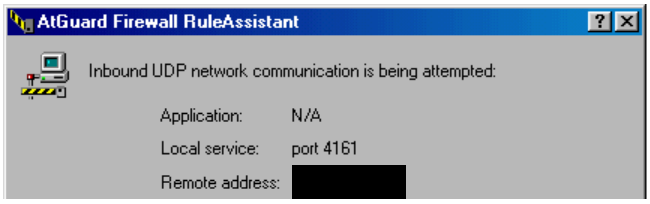

Port4161

Wingate Proxy Server Scan

References

PCWorld – How it Works: Personal Firewalls June 5, 2000 by Robert L. Hummel
http://www.pcworld.com/hereshow/article.asp?aid=17012

Microsoft Technet – Protecting yourself online October 30, 2001 by Steve Riley
http://www.microsoft.com/technet/columns/security/aus1001.asp?frame=true#a

ZdNet Reviews AtGuard August 9, 1999 by PC Magazine
http://www.zdnet.com/products/stories/reviews/0%2C4161%2C410455%2C00.html

ZdNet Firewall Comparison Table August 9, 1999 from PC Magazine
http://www.zdnet.com/products/stories/reviews/0,4161,2311256,00.html

Download AtGuard 3.22
http://www.peregate.com/anvil/atguard.html

WRQ – Creators of AtGuard
http://www.wrq.com

Firewall Scoreboard by Steve Gibson, Gibson Research Corporation
http://grc.com/lt/scoreboard.htm

Security Magazine Review of AtGuard
http://www.scmagazine.com/scmagazine/2000_03/survey/products_01.html

vbeGone Firewall Review by TruchiSoft
http://www.vbegone.com/virushelp/firewalls/

Whois database
http://www.arin.net

IANA Port Numbers
http://www.iana.org/assignments/port-numbers

SANS Institute Trojan horse Ports to Block
http://www.sans.org/newlook/resources/IDFAQ/oddports.htm

© SANS Institute 2000 - 2005 Author retains full rights.

DOShelp.com Trojan Horse Ports
http://www.doshelp.com/trojanports.htm

FreeBSD Ports
http://www.freebsd.org/cgi/query-pr.cgi?pr=21533

SubSeven
http://subseven.slak.org/