



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

Interested in learning more?

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

Case study in firewall and intrusion detection integration

Maged Mardini

GSEC 1.4 option 1

May 21, 2002

Introduction:

The purpose of this paper is to emphasize the benefit of using and integrating different security defense tools and procedures. The paper will focus on two main tools: The first one is a firewall and the second one is a Network intrusion detection tool.

We are going to present a sample configuration for both Symantec Enterprise Firewall Raptor 6.5 and Symantec Netproowler 3.5 to obtain an integrated defensive mechanism that is more efficient than using each tool separately.

Firewalls:

The firewall is a perimeter tool that isolates the internal LAN from the external Internet world. It acts like a configurable filter that can ensure control for both inbound and outbound traffic.

As an example we can consider a simple internal LAN consisting of a number of clients, a web server, a mail server and Internet access. We can use a firewall to divide this LAN into three isolated segments: The Internet segment, the servers segment and the clients segment. Then we can write rules to allow only the required services to be accessed from one segment to another one.

In this way we can make sure that clients and servers can access the Internet freely. The clients may have full access for the permitted services on the servers. But Internet users can only access the web service on the web server and the mail service on the mail server. They are not allowed to access other services or ports on the servers and they are not allowed to access the clients segment at all.

Intrusion Detection:

The firewall setup discussed in the previous section can limit the access between the different network segments according to the business needs, policies and procedures and provide us with a basic protection against abuses and attacks. However it does not ensure a 100% protection or in other words it is not the security magic bullet that we install, and we just sleep in peace.

I have seen so many firewall protected sites that had been easily compromised. We go and check the firewall expecting a bad configuration or a broken through firewall, we examine the configurations, the log files ... and everything seems to be in a good shape.

There are so many cases where a firewall fails to stop an attack: When it comes from inside the firewall, or from a back door like an Internet connected modem on one of the clients, or simply when the attack comes from outside through the firewall in the form of a malicious request to an allowed service. It might then fool its security policy or just cause this service to crash giving unpredicted responses and compromise the

security of the server and the whole LAN attached to it. A live example of such an attack will be discussed later in this paper.

Actually there is no silver bullet that achieves a full protection solution. But we can achieve higher protection levels by using additional procedures or tools. Intrusion detection systems are widely used these days to enhance the enterprise security. While firewalls are designed to prevent external or boundary attacks, Intrusion Detection systems have the advantage of detecting attacks that succeeded to pass through the firewall and discovering internal attacks originating from internal systems too.

From an architectural point of view, there are two main types of Intrusion Detection Systems: Host based and Network based Intrusion Detection Systems. Host based intrusion detection systems are installed on the host to be monitored and they are totally unaware of what is happening on the other hosts. They are operating system specific and can even be configured to monitor a specific application. They continuously monitor the operating system or application critical files, processes and log files for any suspicious activity. They can be programmed to take the appropriate actions in response to an attack, including an automatic repair or restore of a damaged file from a safe backup.

On the other side, network intrusion detection systems can monitor network traffic, looking for malicious packets. They can detect attacks targeting different hosts and applications on this network. In the following section we are going to have a deeper look on Network based Intrusion Detection Systems.

Network Intrusion Detection:

A Network intrusion detection system (NIDS) is a tool that acts as a network sensor or spy that listens to all the traffic on a given segment. The packets are analyzed and compared to known attack signatures. The identified attacks are logged and countermeasures can be taken starting from just paging the administrator up to ending the connection or hardening a firewall.

While a firewall performance can directly affect the overall network bandwidth, a NIDS does not load the network and has no effect on the network speed.

Most NIDS are signature based and therefore can detect only known signature attacks and hence need to be periodically updated to give a more efficient protection.

Another challenge is where to connect the NIDS sensor in switched network environments. Since switches do not broadcast all the network traffic to all ports they will blind off the NIDS sensor. We can overcome this shortage by using network taps that can mirror one network connection packets (only one connection is monitored). We can also use a switch that supports spanning ports. A spanning port can reflect one or more network ports traffic to the sensor port, but as the number of spanned ports increases, we risk to have bandwidth congestion on the sensor port, which means dropped packets and may be missed attacks.

Another compromise would be to connect the NIDS sensor to a port that mirrors a firewall leg that leads to the servers segment. This way the NIDS sensor will monitor all the traffic coming to the servers segment or going out from it.

Raptor 6.5:

Symantec enterprise firewall Raptor 6.5 is an application -level proxy based perimeter security device. The secure proxies examine the entire data payload up to the application level. Though it is slower than simple or stateful packet filtering techniques but it is considered to be more secure as it can protect against application level attacks.

Raptor 6.5 can be installed on different operating Systems like Windows NT, Solaris, HP-UX and Tru64 Unix.

The Basic configuration of the firewall consists of two network interfaces. One network interface will be connected to the Internet or the outside world. The second interface will be connected to the internal LAN or the inside world. Each network interface is called a firewall leg. We can configure the Raptor firewall with more legs, usually we add a third leg on which we connect the servers segment that need to be accessed by the outside world like a web server for example. This segment can be called the service zone, previously called DMZ or demilitarized zone when used to remain outside the firewall. Cascaded firewalls can be connected to provide different levels of security depending on the network structure.

By default, after installing the Raptor software all legs are totally isolated from each other and no traffic is allowed to circulate between them through the firewall unless an explicit rule is written for this purpose.

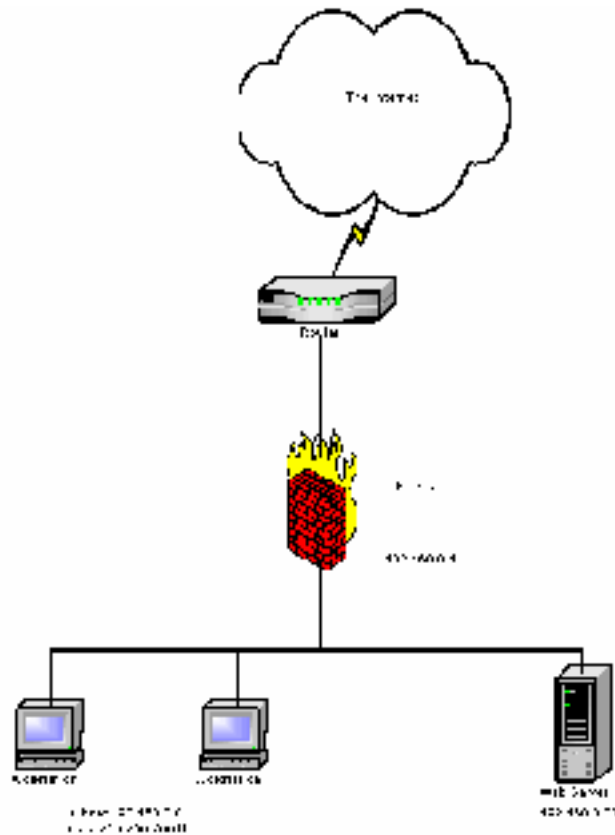
The raptor software comes with built in secure proxies for the following services [9]: Telnet, FTP, HTTP, NNTP, RealMedia, SMTP, DNS, NTP, Gopher, CIFS, SQL*net, H.323, Ping and NetBIOS datagram. Other protocols or services can be defined and a general -purpose proxy will handle them.

Raptor can authenticate network connections before giving them access to the network. It supports several authentication types either built in or using third party authentication servers [9]:

- Defender
- CRYPTOCARD
- Security Dynamics SecurID
- Bellcore S/Key
- Gateway password
- NT Domain
- TACACs+
- Radius
- LDAP

It also supports OOBAs, out of band authentication, which allows the firewall to authenticate a service that does not support authentication. The client accesses a special web page on the firewall using a simple web browser. This web page will allow the client to authenticate its IP for the required service.

Raptor also supports virtual private networks (VPNs) including: IPsec, IKE (ISAKMP/Oakley), and swIPE.



In the next section we will give a brief description for a sample Raptor configuration without going to the full depth of all the detailed capabilities.

Before writing a rule, we need to define the network entities between which there will be a flow of information. An entity can be mainly a host, a subnet, a domain or simply a group of other entities. A default entity is predefined in Raptor, that is 'universe' with address=0.0.0.0 which refers to any address inside or outside the firewall.

Assuming that we have a two legs Raptor with a number of PC Windows clients and a web server on leg1. The external leg connected to the Internet will be leg2. For the purpose of our exercise we are going to use two private IP ranges for both legs.

Leg1 subnet address is 192.168.0.0 with netmask 255.255.255.0

Leg2 subnet address is 192.168.10.0 with netmask 255.255.255.0

Now we can define the following entities:

Webserver : 192.168.0.33

InternalSubnet : 192.168.0.0 , mask=255.255.255.0

Until this point, there are no connections allowed between our LAN and the external world. We are going to create the following rules:

Rule1: from internalSubnet to universe allow : http, https, smtp, pop3, DNS, ping, telnet, ftp.

Rule2: from Universe to Webserver allow : http, https.

Address Transforms:

After adding the previous two rules our LAN can access the outside world for mail and web services in addition to telnet, ping and ftp.

However in spite of the permitting rule2, the outside world can not access our web server, since it has only an internal IP.

Raptor software allows us to use different types of address redirections and transforms. We are going to create an address redirection as follows:

Redirect requests to IP 192.168.10.5 - port 80 to IP 192.168.0.5 - port 80.

Now external users can address our internal web server using the new virtual IP 192.168.10.5.

With this sample configuration, our firewall protects our LAN and web server in a way that no external access is allowed to the PC's, and only http access is allowed to one host, which is the web server. More and above the real address of the web server is hidden from the external users.

When the firewall fails:

Our firewall is now configured to allow only http requests to the web server and no access at all from outside to our internal LAN. Does this mean that the web server is secure? Or the rest of the internal LAN is secure?

The answer is no. The firewall will filter a huge number of different attacks that are circulating all over the net, may be targeting our server or just looking for any vulnerable victim. But the servers might be running vulnerable applications that can be very easily exploited if the firewall permits access to these vulnerable applications. That would compromise not only the server's security but also the security of the whole LAN. This is what we are going to show in the next section.

Application vulnerability example:

IIS Unicode Vulnerabilities:

As an example of application vulnerabilities we will have a quick look on IIS Unicode vulnerabilities. We are going to install IIS version 4.0 or 5.0 on Windows NT or Windows 2000 and this is going to be our test web server. It has been installed using the default installation and it hasn't been hardened or patched. A huge numbers of web servers on the net share this status. Also a great number of non web servers have an unused default vulnerable IIS installation.

As stated by Bolotron in the bugtrack mailing list: [7]
<http://cert.uni-stuttgart.de/archive/bugtraq/2001/07/msg00537.html>

We are going to download the PHP interpreter source and compile it on a Linux or Solaris system. This system will act as the attacker and will be connected on the outside leg of the firewall.

Cut and paste the script on the mentioned mail and save it under the name 'iis-kabom'. iis-kabom contains 70 separate http requests that can be run sequentially to exploit different IIS Unicode vulnerabilities.

As agreed we are going to run this vulnerability test from the outside attacker machine and targeting our test web server:

```
# iis-kabom -t 192.168.10.5
```

The script will run and will stop at the first found positive vulnerability. The response will be something like this:

```
...  
Trying variant number 10 -----> No Vulnerable :(  
Trying variant number 11 -----> No Vulnerable :(  
Trying variant number 12 -----> No Vulnerable :(  
Trying variant number 13 -----> No Vulnerable :(  
Trying variant number 14 -----> No Vulnerable :(  
Trying variant number 15 -----> No Vulnerable :(  
Trying variant number 16 -----> No Vulnerable :(  
Trying variant number 17 -----> No Vulnerable :(  
Trying variant number 18 -----> vulnerable!!
```

Now we can try to run any command on this web server using a web browser or the same iis-kabom using the following format:

```
#iis-kabom 192.168.10.5 dir_c:\\ 68
```

where ‘_’ replaces a space and an extra ‘\’ precedes the normal ‘\’, 68 is the vulnerability number just discovered in the previous step.

Surprisingly the command will display a full list of the directory on drive c: \

You can try other commands like delete, or copy ... they all work.

You can also mount network drives or access already mounted network drives residing on the internal LAN. And all this happens normally through the firewall. The attacker can be any Internet user anywhere in the world.

Netprowler 3.5:

It is obvious that the firewall by itself is incapable of detecting attacks as the one described above and of course can not block them. Another tool is needed to help and provide more defensive power.

While firewalls offer perimeter and access controls, users can always attempt to exploit known vulnerabilities to circumvent security policies. Symantec Netprowler is a network intrusion detection software. It complements the existing security countermeasures and fortifies the system resistance against intruders.

Netprowler is available for trial at Symantec site at:

<http://enterprisecurity.symantec.com/products/products.cfm?ProductID=50>

It installs on Windows NT platform and has three main components [11]:

- Agent
- Manager
- Console

The agent can be considered as the network sensor that gathers the information and communicates it to the manager. The manager can manage different agents distributed at different points of the network. Finally the console is the GUI interface that allows us to configure the manager and the agents parameters. Symantec recommends that the agents and the manager be installed on dedicated systems with a minimum of 128 MB RAM, while the console can be installed on a shared system.

Netprowler comes with built in attack signatures and a possibility to synchronize with new signatures updates through an HTTP connection to Symantec autoupdate site. An attack definition toolkit allows us to manually define new signatures. After installing the agent, it must be added to the manager and configured with the IP ranges it is going to monitor, the address of the firewall it can harden, the administrator email address or pager number to which it can send attack warnings.

The built in profiler can be scheduled to run at different time intervals. Its role is to probe the defined IP ranges to [11]:

- Discover live nodes on the network.
- Identify the type of operating system running on each host.
- Identify the type of applications running on each system.
- Automatically associates attack signatures based on the type of the operating system and the application running on each system.
- Adds the found host to the appropriate agent.

During consecutive profiler runs, new hosts and applications found are added to the configuration. Previously identified hosts and applications are kept, even if they were not found in the new run.

Attack association to each host can be modified manually for fine tuning the attack monitoring. We can cancel attacks, add attacks, change priority of an attack, add an authorized system and port or define an attack response.

Authorized system and port: For each attack, authorized systems and ports can be defined. They inform the agent not to consider the incident as an attack only if it comes from the specified source authorized address and port.

Attack responses: Once the agent identifies an attack, it can take one or more of seven possible responses [11]:

- Send email
- Page an administrator
- Send an SNMP trap
 - Harden a Firewall: Netprowler supports firewall notification for both Checkpoint Firewall-1 and Symantec Raptor.
- Terminate Session
- Capture the Session
- Spawn a command

Netprowler can terminate a session-based TCP attack by sending a TCP/IP reset command to the server on behalf of the client. However other types of attacks can not

be terminated. We should also note that the agent performs a near real time analysis for the captured packets. It may happen that during heavy network traffic periods, the agent becomes too much loaded, that it drops some frames or takes a relatively delayed response thus sending a TCP reset packet after the session has been normally terminated or in other words after the attack success. We'll see in the next section how to monitor and optimize the Netprowler performance.

NetProwler performance monitoring and optimizing:

From the Netprowler console we can monitor all the agents status. The agent statistics are displayed showing:

- Agent status: enabled/disabled
- Number of monitored sessions
- Number of TCP segments
- Number of UDP datagrams
- Number of ICMP datagrams
- Number of frames processed
- Number of frames dropped
- Number of invalid frames

As stated in the Netprowler user's guide [11]:

“The number of frames dropped refers to network traffic that NetProwler did not monitor because it did not have enough memory or available CPU cycles to analyze the traffic. During normal network conditions, NetProwler should not drop frames. If NetProwler is dropping frames, you should increase the system memory, monitoring fewer hosts, or add an additional installation of NetProwler. As a rule of thumb, you should upgrade your system's configuration if 3 percent or more of the total frames processed are being dropped.”

Raptor Firewall Hardening:

As stated in the Netprowler user's guide [11]:

“NetProwler uses an authenticated, relay -protected UDP protocol to communicate with the Raptor firewall. NetProwler uses a “shared secret” authentication string. “Shared secret” means that both NetProwler and Raptor use the same authentication string. This prevents a third party from inserting or modifying the communication data. For additional security, the UDP protocol's relay protection prevents a third - party from altering NetProwler or Raptor operations by capturing and retransmitting the conversation data.”

The act of Netprowler detecting an attack, notifying Raptor and Raptor blocking the attacker is called ‘firewall hardening’. The Raptor firewall processes the notification, adds the attacker address to the blacklist and starts immediately dropping all the packets coming from it for a configured period of time. This period is set up in Raptor for 24 hours as default, but can be changed by the administrator as needed. During this period, the administrator can have a chance to analyze any suspicious activities and decide if further securing steps are needed. The success of choosing the right attack responses specially connection termination, firewall hardening and the right blocking period can result out a quite strong defensive integrated tool that allows us to

block an attacking address in the early attacking phases while the attacks are still identifiable by our tools.

Netprowler Reporting:

The Netprowler manager uses a Mysql service to store its data bases: The Netprowler configuration database and the attacks database. All the attacks are stored in the database from the date of the installation or from the last database purge. At any time the administrator can query the alert database by clicking on the 'Find Alerts' button on the monitor menu of the console. Or he can generate and schedule preformatted reports. Customized reports can be designed using Crystal reports designer.

Live session monitoring:

From the Netprowler agent GUI screen, we can monitor live sessions like Telnet, FTP, echo, IRC, SMTP, POP3, rsh, and rlogin ... in real time. This process is CPU intensive and would affect the agent performance greatly. So, we can not configure it to monitor all the services on all the hosts, all the times. This feature is made to be used to closely monitor a session to get more information about it to create new attack signatures and enhance response capabilities.

Firewall hardening live test:

Back to our previous test environment, we are going to plug a Netprowler agent, manager and console inside the firewall on the internal LAN: 192.168.0.0.

In the agent properties we'll add the monitored IP range: 192.168.0.1 to 192.168.0.35 This range includes our web server (192.168.0.5) and the rest of our internal PC's. We right click on the agent and choose 'profile now' to run the profiler. The agent starts the network discovery and automatic attack association for each found host. We do not forget to configure the network switch to mirror the internal firewall leg to the Netprowler agent's port, to make sure that all traffic coming from the Internet can be monitored by the agent.

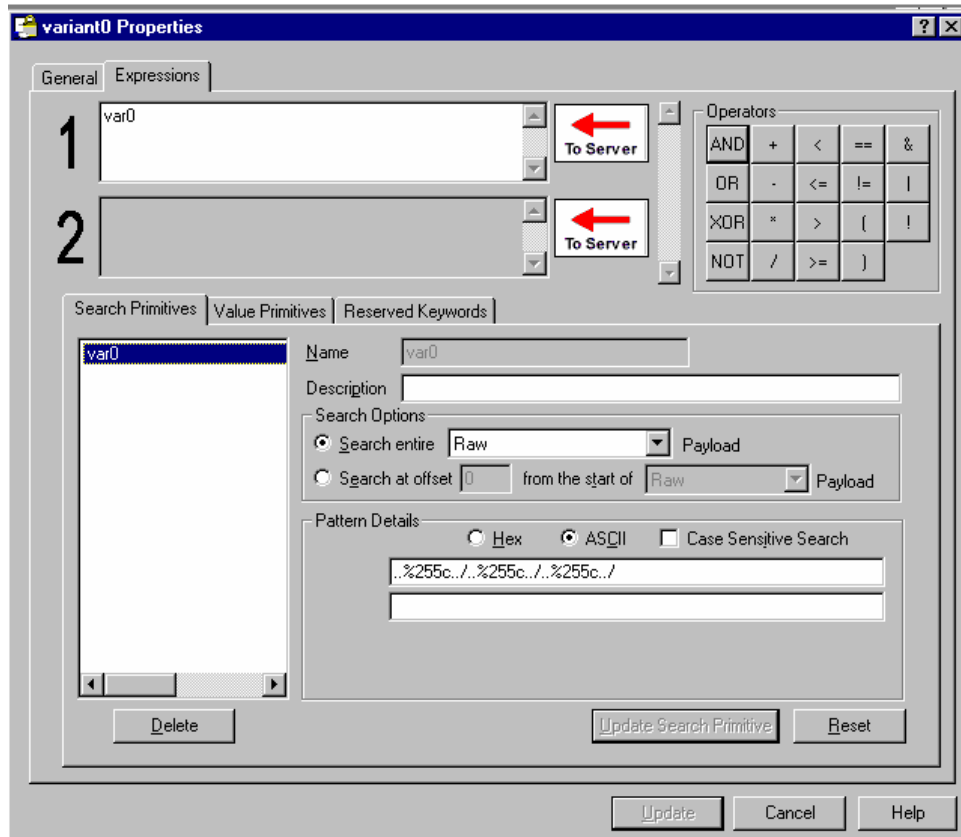
Now, from our outside attacker we rerun the previous iis -kabom test:

```
# iis-kabom -t 192.168.10.5
```

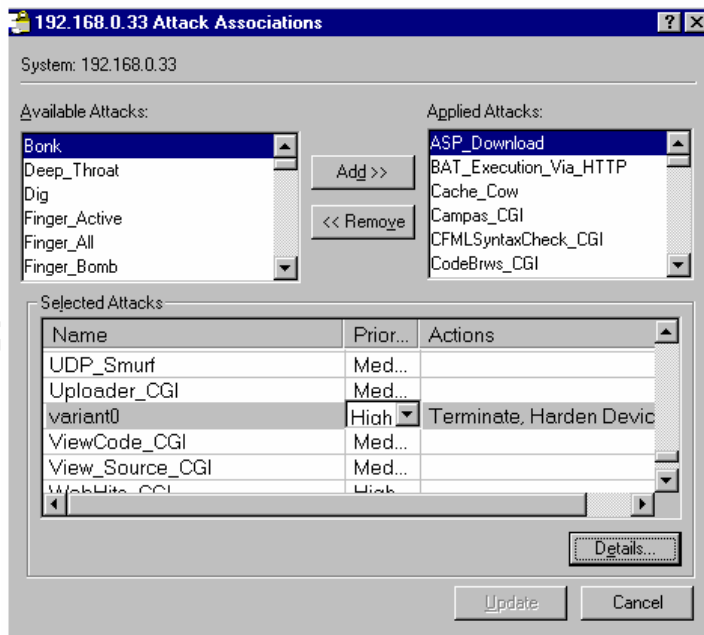
```
Trying variant number 0 -----> vulnerable!!
```

We still get that the web server is vulnerable starting at the first attack variant, and no alerts are triggered on the Netprowler. This means that the attack signature is not recognized. We can try to make a signature autoupdate from the Netprowler console menu, or for the purpose of our exercise, to manually create a user defined attack signature. The test attacks are all found in the source of iis -kabom.


We'll create a new signature called variant0 for which we are going to use the attack pattern of variant0 from the iis -kabom source. Right click on the web server IP and associate the new attack with it. Choose 'terminate session' as response and save the configuration.



Netprowler signature definition screen



Netprowler Attack association screen



Netproowler Attack response screen

Run the attack again: Now the result shows that the web server is not vulnerable anymore to the first attack but to the second one:

```
# iis-kabom -t 192.168.10.5
```

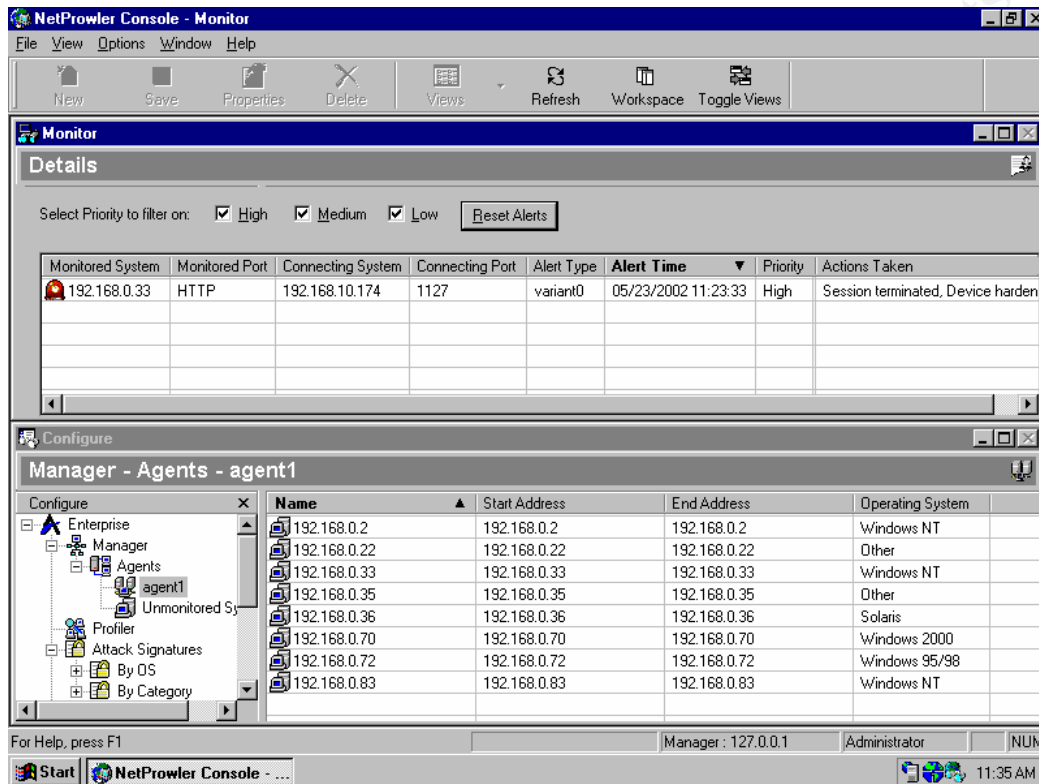
Trying variant number 0 -----> No Vulnerable :(
Trying variant number 1 -----> vulnerable!!

A look on the Netproowler alerts screen, shows that attack variant0 has been identified and that the action 'session terminated' has been taken. The Netproowler succeeded to stop the newly defined attack variant0 as configured.


Now, comes the last step of our test. Add 'harden firewall' to the responses for variant0 attack. In the Netproowler agent's properties add the firewall IP, and an authentication string. On the raptor firewall, run the 'rempass' utility to add the IP of the Netproowler agent and the same authentication string. Note that a new address transform rule need to be created on raptor to allow external IP's to pass transparently to the web server instead of being hidden and replaced by the firewall interface IP. This will enable the Netproowler to identify the attacker IP and communicate the correct notification to the firewall.

Now if we run again the variant0 attack, the Netproowler agent will notify the firewall, which will add the attacker IP to the blacklist. All traffic coming from the attacker IP will be blocked for the next 24 hours or whatever period the administrator would configure raptor to use. On the raptor active connections screen we can see all the black listed IP's, the time they started at and the elapsed time. At any time the administrator can kill this connection and the corresponding IP is allowed to pass through again. The raptor log file shows the failed http request, the blacklist process authentication between the Netproowler agent and the firewall, the blacklisted IP and

the timeout duration. The blacklist info is lost if the firewall is rebooted but is kept if the firewall is shutdown and restarted from the console. The administrator has the chance to further investigate the attacks and decide whether to block this IP permanently or for any amount of time.



Netproler Console showing the Alerts, the agents and the monitored IP's of agent1



A clip from Netproler Attacks - summary report

Now, no more attacks will be allowed from this attacking IP, thus preventing the attacker from compromising our system if the Netprowler fails to stop the attack in time or if the attacker uses other non identifiable attacks.

Conclusion:

A firewall protects the internal network from unauthorized access, but can miss malicious network traffic. Intrusion detection systems are designed to detect that kind of attacks and warn about them. Though network intrusion detection systems can end TCP sessions, they can not block all suspicious traffic but they can strengthen the perimeter by notifying the firewall to update its policy to block a threatening address at the very early identified suspicious activities.

There is no security silver bullet that can give us all the protection we need all the time, one tool can miss an attack that can be detected by another one. Using different tools enhances our defenses. The presence of security skilled staff can help managing, integrating and tuning the existing tools.

Having good tools does not mean to relax. Continuous efforts are required to provide defense in depth. We should always keep updated regarding the latest security news, tools, bug tracks and new attacks. We must make sure that we have security and backup policies, they are strictly followed, that our systems have the latest patches applied and our applications and services are securely configured.

References:

- [1] Dennis Carter " Raptor Firewall 6.5 Running on NT 4.0 "
<http://rr.sans.org/firewall/raptor.php>
- [2] Eric Biedermann "Netprowler- A Look at Symantec's Network Based IDS "
<http://rr.sans.org/tools/netprowler.php>
- [3] Symantec Netprowler trial download page:
http://enterprisesecurity.symantec.com/products/products_.cfm?ProductID=50
- [4] Ryon Packer "A Basic Guide to Intrusion Detection"
http://www.ciscoworldmagazine.com/webpapers/2001/08_intrusion.shtml
- [5] Char Sample, Ian Poynter and Mike Nickle
"Why Security Products Fail"
<http://www.jerboa.com/whitepapers/whysecurity.pdf>
- [6] Mark Cooper "An Overview of Intrusion Detection Systems "
http://www.xinetica.com/tech_explained/general/ids/wp_ids.pdf
- [7] Bolotron , Bugtraq Mailing List.
"Yet another UNICODE exploit code and vulnerability test for IIS 4.0/5.0."
<http://cert.uni-stuttgart.de/archive/bugtraq/2001/07/msg00537.html>

[8] Tom Rodriguez “Understanding IIS Unicode Vulnerabilities”
http://www.infosecalliance.com/resources/whitepapers/iis_unicode-vuln.pdf

[9] Raptor Firewall and PowerVPN V6.5 Configuration Guide for NT

[10] Raptor Firewall and PowerVPN V6.5 Reference Guide

[11] Netprowler 3.5 user's guide

© SANS Institute 2000 - 2002, Author retains full rights.