



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>

Netprowler- A Look at Symantec's Network Based IDS.

Eric Biedermann

August 17, 2001

Introduction

System administrators today face many challenges presented by the increased use of data networks and the desire to connect private corporate networks to public networks. Certainly, one of the greatest challenges facing administrators is how to secure their networks against the ever-increasing threats facing computer networks today. The system administrator must balance the advantages of allowing access to seemingly endless quantities of information that the internet possesses, against the disadvantages presented by the fact that the same connection that allows employees and customers to access this information also allows an avenue of attack into the corporate network. The challenge facing enterprise network administrators today is to satisfy business needs by providing more access to more information, while maintaining the privacy and integrity of the corporate network.

Before we examine the features and capabilities of the Netprowler IDS, let us review the common types of attacks and look at an example of a typical intrusion scenario.

Types of Attacks

While a network is opened up, it can be exposed to various attacks. While there are many specific attacks used by the attacker today based on Operating system or specific information sought, they can usually be placed into one of three categories. The three categories of attacks are [2]:

- **Reconnaissance:** These include ping sweeps, DNS zone transfers, e-mail reconns, TCP or UDP port scans, and possibly indexing of public web servers to find cgi holes.
- **Exploits:** Intruders will take advantage of hidden features or bugs to gain access to the system.
- **Denial-of-Service (DoS) attacks:** Where the intruder attempts to crash a service (or the machine), overload network links, overloaded the CPU, or fill up the disk. The intruder is not trying to gain information, but to simply act as a vandal to prevent you from making use of your machine.

Now that we know what the three most common categories of attacks are and some specific attack types, we are faced with a question. How do the attackers use these attacks to gain access to networks to exploit information or cause malicious destruction?

The following is an example of a common intrusion scenario [2].

Step 1: **outside reconnaissance:** The intruder will find out as much as possible without actually giving themselves away. They will do this by finding public information or appearing as a normal user. In this stage, you really can't detect them. The intruder will do a 'whois' lookup to find as much information as possible about your network as registered along with your Domain Name (such as foobar.com. The intruder might walk through your DNS tables (using 'nslookup', 'dig', or other utilities to do domain transfers) to find the names of your machines. The intruder will browse other public information, such as your public web sites and anonymous FTP sites. The intruder might search news articles and press releases about your company.

Step 2: **inside reconnaissance:** The intruder uses more invasive techniques to scan for information, but still doesn't do anything harmful. They might walk through all your web pages and look for CGI scripts (CGI scripts are often easily hacked). They might do a 'ping' sweep in order to see which machines are alive. They might do a UDP/TCP scan/strobe on target machines in order to see what services are available. They'll run utilities like 'rcpinfo', 'showmount', 'snmpwalk', etc. in order to see what's available. At this point, the intruder has done 'normal' activity on the network and has not done anything that can be classified as an intrusion. At this point, a NIDS will be able to tell you that "somebody is checking door handles", but nobody has actually tried to open a door yet.

Step 3: **exploit:** The intruder crosses the line and starts exploiting possible holes in the target machines. The intruder may attempt to compromise a CGI script by sending shell commands in input fields. The intruder might attempt to exploit well-known buffer-overflow holes by sending large amounts of data. The intruder may start checking for login accounts with easily guessable (or empty) passwords. The hacker may go through several stages of exploits. For example, if the hacker was able to access a user account, they will now attempt further exploits in order to get root/admin access.

Step 4: **foothold:** At this stage, the hacker has successfully gained a foothold in your network by hacking into a machine. The intruder's main goal is to hide evidence of the attacks (doctoring the audit trail and log files) and make sure they can get back in again. They may install 'toolkits' that give them access, replace existing services with their own Trojan horses that have backdoor passwords, or create their own user accounts. System Integrity Verifiers (SIVs) can often detect an intruder at this point by noting the changed system files. The hacker will then use the system as a stepping-stone to other systems, since most networks have fewer defenses from inside attacks.

Step 5: **profit:** The intruder takes advantage of their status to steal confidential data, misuse system resources (i.e. stage attacks at other sites from your site), or deface web pages.

Another scenario starts differently. Rather than attack a specific site, an intruder might simply scan random internet addresses looking for a specific hole. For example, an

intruder may attempt to scan the entire Internet for machines that have the SendMail DEBUG hole. They simply exploit such machines that they find. They don't target you directly, and they really won't even know who you are. (This is known as a 'birthday attack'; given a list of well-known security holes and a list of IP addresses, there is a good chance that there exists some machine somewhere that has one of those holes).

What is Netprowler?

Symantec describes Netprowler as follows [1]: “ Netprowler provides dynamic intrusion detection by transparently examining network traffic to detect, identify, log, and terminate unauthorized use or misuse of computer systems. Netprowler is the only network-based Intruder Detection System that combines “out of the box” use, attack signature extensibility, real-time signature deployment, and multi-platform host IDS integration.”

What does Intrusion Detection mean?

By definition, Intrusion Detection is the art and science of sensing when a system or network is being used inappropriately or without authorization. An intrusion-detection system (IDS) monitors system and network resources and activities and, using information gathered from these sources, notifies the authorities when it identifies a possible intrusion.

I have a firewall. Why do I need an Intrusion Detection System?

Computer security professional and author, Carolyn Meinel, points to the following example as a case highlighting a need for intrusion detection systems [4]:

You have the world's best firewall, your Windows computers update their antivirus software regularly and your Information Security staffers enforce your policies with an iron fist. Does this mean you're safe?

Maybe not. In 1998, a news story asserted that the firewall for the *New York Times* was one of the best. Yet at 7:08 a.m. on Sunday, Sept. 13, 1998, someone on the paper's network e-mailed reporters:

```
...COM3 V1S1T HTTP://WWW.NYTIMES.COM AND S33 0UR LAT3ST P13C3  
0F ART. 1F 1T D0ESN'T LOAD, JUST H1T 'RELOAD' A F3W T1MES. CL3V3R  
ADMINZ HAD S0M3 W3IRD CR0NTABZ OR S0METHING.
```

```
0H. W3 0WN YOU. YOU JUST HAV3NT N0T1C3D US 0N Y3R N3TWORK Y3T.  
UNT1L THE N3XT T1M3...
```

No one at the *Times* had noticed weeks worth of the Hacking for Girliez gang on their network. The intruders finally chose to go public by defacing the opening page of their Web site—on the day the *Times* expected millions of visitors to view the Monica Lewinsky transcripts. Instead, visitors encountered soft porn and an ad for Lewinsky-scented cigars.

This may have been avoided had the times been running a good intrusion detection system.

There are two general types of intrusion detection systems: Host based and Network based. A host based IDS is installed on the system or device that it is assigned to protect. Host based IDSs are designed to deal with attacks that are specific to the hosts operating system and applications, and any attacks that originate from the local console.

Like most IDSs, Netprowler is designed to monitor all traffic on a network segment and look for any attacks that are not identified or prevented through other security measures. IDSs work by monitoring networks or devices for evidence of attacks. An IDS identifies an attack by a pattern of behavior, known as a signature, which is associated with a specific attack or type of attack. The objective of an IDS is to identify an attack in progress and respond in a timely manner, before the systems are compromised.

Just as host based IDSs monitor activity on a specific host, a network based IDS monitors activity on a network. A network-based IDS's primary objective is to recognize attacks that exploit properties of networking protocols. Symantec highly recommends, that for maximum performance, network-based IDSs should reside on a dedicated system, one that is not used for server applications and provides network-based protection for some or all hosts on the network segment [1].

The combination of host-based and network-based IDSs provides greater protection than using either type by itself. A network-based IDS will often detect suspicious activity that represents an attacker's early probes as the attacker attempts to discover a network's resources prior to mounting an attack. This gives the security administrator some advanced warning to deal with the attack. On the other hand, a host based IDS is better designed to detect more operating system and application specific attacks, based on system information that is not available to a network-based IDS.

Netprowler Features

As a network-based IDS, Netprowler is designed to monitor TCP/IP traffic on a network segment, conduct near real time analysis of packet structures, identify suspicious activity, and take appropriate action.

Symantec designed many features for Netprowler with the goal being to improve network monitoring, these include [1]:

A Distributed Architecture that consists of a Console, Manager, and Agent. This provides centralized management and better network coverage.

The Netprowler Profile tool scans the network for live systems and adds them to the list of machines to be monitored, thus automatically configuring systems and ensuring coverage.

Netprowler's Automatic Application of Attack Signatures will automatically assign each new system a set of attack signatures that will allow them to recognize an attack. This ensures that the new machines are protected as soon as they come on line.

Agent Status Monitoring will keep track of the Agent's status and notify an administrator of any interruption in service. This ensures that no segment of the network will be left unprotected.

Understanding Netprowler's Architecture

Netprowler consists of the following components [1]:

- Console
- Manager
- Agent

The Netprowler Console is the main graphical user interface (GUI) that allows the administrator to manage the individual agents dispersed throughout the network. The Console allows the management of all the Agents assigned to a specific Manager, monitor attacks detected by the Agents, generate reports, create new attack signatures, set access restrictions, and monitor the status of each Agent.

The Netprowler Manager stores all configuration and attack information for the Agents assigned to it. Based on the information it receives from the administrator via the Console, it directs the configuration of the Agents, stores attack signature definitions, stores alert data from the Agents, and sets database configuration options.

The Netprowler Agent is where the bulk of the work is done in Netprowler. The Agent is an application that monitors the network segment assigned to it for suspicious behavior and responds to intrusion attacks in near real time. The Agent monitors network traffic between servers, clients, and other network devices, detects attacks such as TCP/IP spoofing and SYN flooding, responds to identified attacks with preconfigured actions, such as terminating a session or hardening a firewall, and monitors network sessions in real time.

How the Agent Detects Network Attacks

Symantec recommends installing the Netprowler Agent on a dedicated system connected to an Ethernet 10 Mbps or 100 Mbps network using the TCP/IP protocol suite. Since the Agent is placed on a dedicated system, it can detect network attacks without degrading the performance or accessibility to other network systems or devices, and is completely transparent to network users. The Agent functions by having the systems network interface card (NIC) placed in "promiscuous" mode. This means that the NIC has no IP address or protocol suite bound to it, thus, rendering the NIC invisible to other systems. This allows the Agent to monitor all network traffic and detects attacks by comparing the packet data to its database of attack signatures.

As the Agent scans the network traffic, it analyzes the TCP packets looking at the destination address to see if it belongs to the range of addresses it is responsible for. If the address does not fall within this range then the Agent ignores the packet. If the address in the packet belongs to the Agent's assigned range it will compare the packet to a library of attack signatures looking for a match. If there is a match, the Agent sends an alert and takes appropriate action. If there is no match, the packet continues on to its destination.

Netproowler comes equipped with an impressive set of tools for attack recognition, including:

Stateful Dynamic Signature Inspection (SDSI): a detection method in which Netproowler is able to remember the context, or state, of a network session. This capability allows Netproowler to detect, identify, and prevent sophisticated attacks that send multiple series of packets, each of which seem harmless when looked at individually.

Predefined Attack Signatures: allows Netproowler to be used pretty much right out of the box to detect, identify, and prevent many attacks. This allows for immediate protection without having to configure additional signatures.

Signature Sync feature: allows the user to download new attack signatures from the Symantec web site and load them into the existing signature library.

A Custom Attack Signature Definition tool: allows the user to design custom attack signatures.

Notification and Response Tools

Once an attack has been detected, Netproowler provides several response options, including:

Notifying an administrator via email or pager.

Capturing the packets for the rest of the session. This is important for post analysis.

Terminating the attacker's session by sending a TCP/IP reset command to the server.

Reconfiguring or Hardening a firewall. Netproowler will alert the firewall to the suspicious activity taking place and reconfigure the firewall if necessary. This feature is currently supported on only Checkpoint and Symantec Enterprise firewalls.

Generating SNMP traps. These can be sent to two SNMP managers notifying them of the attack.

Spawn a command or execute a batch file.

Installation

Netproowler, consisting of the Console, Manager, and Agent, can be installed on systems that meet the following minimum requirements:

Windows NT 4.0* server/workstation (with service packs 3,4,5,and 6 installed)

300 MHz Pentium II processor (233 MHZ Pent. II for Console)

128 MB RAM

70 MB free disk space for storage
150 MB virtual memory
10/100 MB Ethernet
Static IP address, gateway, and DNS server.
Internet Explorer 5.x or greater (Console only)
Modem and phone line for agent (optional)

* The Netproowler Manager and Agent are designed only for Windows NT 4.0. The Netproowler Console will operate on Windows 2000 workstations as well.

Symantec strongly recommends that the agent and the manager be installed on dedicated systems. If they are loaded on a system with other applications running, their performance can be degraded. This is important because the Agent does the majority of the work and will consume up to 100% of the CPU's resources processing the large amount of data necessary to do its job. The necessity to inspect each packet requires the Agent to monopolize the CPU and not allow resident applications access to the CPU.

Symantec recommends that for maximum coverage and performance, that installation should consist of (1) Agent on each network that is being monitored. These should be networks where sensitive and critical information reside. It is the Agent's job to examine the packets of all network traffic on the segment. First, the Agent looks at the destination address to see if it belongs to one of the workstations that the Agent is configured to monitor. If the address does not belong to a workstation the Agent is responsible for the Agent will ignore the packet. If the address is a match, the Agent will compare the packet to a library of attack signatures and look for a match. If there is a match, then the Agent sends an alert and takes appropriate action. If there is no match, the packet continues on to the destination address.

One manager is installed for every 20 Agents. Actually, Managers can support up to 20 Agents, but depending on the network configuration and the traffic load, increasing the number of Managers may be necessary to reduce the workload for each Manager. Symantec strongly recommends that the Manager be local to the Agents that it manages, instead of forcing Agents to report to Managers at a remote sight. It is the Manager's job to monitor the status of each Agent, push updates of attack signatures to them, and implement any configuration changes received from the console.

The Console provides the graphical user interface for administrators to monitor the network, configure Agents, and to define and apply signatures. The Console can be placed on the same machine as the Manager or placed at a location convenient for the administrator.

Netproowler Deployment

The deploying of Netproowler is more complicated than simply placing one Agent on each segment that you want traffic monitored on, or placing one Manager for every 20 Agents. The deployment is more complex due to the fact that there are additional considerations in placing Netproowler Agents, and often-additional steps must be taken to be able to view all the traffic that

you desire. In the following section we will look at some issues involved in deploying Netprowler on several types of network infrastructures.

Symantec recommends the following deployment strategies [1]:

Deploying the Manager and Console:

As a rule of thumb, both the Manager and the Console should be located on protected networks to prevent corruption or manipulation of the manager's database. Besides that, the only other requirement is that the Console be able to communicate with the Manager and that the Manager be able to communicate with the Console and all Agents that it manages.

Symantec's security engineers have studied the deployment of Netprowler and where to best place Netprowler to take advantage of the many tools that it provides. Several of these locations are discussed next.

Placing an Agent behind the Firewall:

Security consultant David "Del" Elson states: "Placed between the firewall and the system being secured, a network based intrusion detection system can provide an extra layer of protection to that system. For example, monitoring access from the Internet to the sensitive data ports of the secured system can determine whether the firewall has perhaps been compromised, or whether an unknown mechanism has been used to bypass the security mechanisms of the firewall to access the network being protected [5]."

By placing the Agent behind the firewall it will serve to backup the firewall where you can use Netprowler's stateful inspection abilities and advanced attack signature recognition to supplement the firewall's security. Symantec claims that this is a distinct advantage due to the fact that many firewalls cannot detect certain types of attacks, such as port scans or complex attacks performed with multiple packets. Placing the Agent behind the firewall allows the Agent to harden the firewall by detecting that an attack has penetrated the firewall. The Agent can instruct the firewall to take the necessary steps to prevent further penetrations from the attacker.

Placing an Agent outside the Firewall:

In this situation, the Agent is placed between the firewall and the device that provides connectivity to the Internet, usually a router. In this location, the Agent can monitor all incoming and outgoing traffic.

Symantec feels that this placement has a few distinct advantages. You are able to detect when your network is under attack and to monitor the effectiveness of your firewall configuration. This provides verification of exactly which attacks are successful in penetrating the firewall and which are not. However, due to the large amount of unfiltered traffic that the Agent must monitor from this location it is fairly easy for the Agent to become overwhelmed. Due to this limitation, this location is best suited for verifying the firewall configuration, not providing network protection.

Placing an Agent on the Services Network:

Typically, a services network is located on a separate firewall port. A services network usually consists of web servers, mail servers, or FTP servers, that are readily accessible to the public. These systems are frequent targets of attack due to the relative ease of access to them.

Since the services network is a high-risk area for intrusion, placing an Agent on the services network allows you to monitor all traffic to systems on the services network regardless of the origin (internal or external). By being placed on the services network the Agent has the ability to terminate sessions to these servers.

Placing an Agent on a switched network:

By using a switch or switching network hub, an administrator can provide an additional amount of security for the network because packets for a given destination are only transmitted out that system's port, and not out other ports. Because of this, other devices attached to the same switch do not see the traffic.

This placement, however, creates a problem for the Agent. By attaching the Agent to one of the switch's ports, it will be unable to monitor the traffic for the whole network. In this placement the IDS will only see that one port's traffic. For the Agent to be able to monitor traffic on a switched network, you must attach the agent to the switch's monitored port. A monitored port allows you to duplicate traffic from selected ports onto the monitored port. This way the Agent can monitor the traffic on the network segment because it will be able to see the traffic.

This placement does have some performance considerations. By binding or mirroring several ports to the port where Netrowler is listening, creates the potential to surpass the bandwidth available on that port. Careful monitoring of the specific port and possible deployment of additional Agents will avoid any decrease in performance.

Connecting an Agent using dual network interface cards (NIC):

Earlier we briefly discussed using dual NICS on the Agent to monitor traffic. In this configuration, the Agent uses one NIC for monitoring network traffic and the other NIC for communicating. Symantec strongly recommends that this configuration be utilized in unsecured areas of the network, such as a DMZ.

The first NIC is not configured with a TCP/IP address and is connected to the network segment being monitored. It can only perform profiling, system resets and terminations, and passive monitoring. The transmissions from this NIC use a special Netrowler driver that creates special packets that do not reveal any protocol addresses.

The second NIC is connected to a secure network, and is configured with a TCP/IP address. This NIC is used to send attack notifications, SNMP traps, system resets, to the Netprowler Manager.

By configuring the Agent in this manner you can provide a greater degree of security for the Agent system. The absence of a TCP/IP address for the first NIC prevents it from being subjected to attacks or from being exploited. Since it does not make any transmissions and does not respond to any queries, it is invisible to other devices on the network. This NIC is in what is termed “promiscuous” mode.

Once the deployment of the Manager, Console, and Agents has been completed, the administrator can now use the Console to configure the Agents and Managers. This may include selecting Agent responses to detected attacks, adding additional Agents, assignment of attack signatures, managing alerts, scheduling reports, importing new Agent signatures, and many other day to day tasks necessary to update and maintain Netprowler’s maximum performance. As stated earlier, Netprowler is configured to perform intrusion detection right out of the box, with a fairly comprehensive list of attack signatures already loaded. This fact allows the security administrator to provide a certain level of security for the network immediately, thus, providing time to customize Netprowler to compliment existing network defenses.

In Conclusion

It is the intent of this document to serve as an introduction to Symantec’s Netprowler Intrusion Detection System and to present Netprowler as a powerful tool in the security administrator’s arsenal to use to provide a secure network environment. However, no single security tool can provide the amount of protection needed in today’s dynamic threat environment. Installing an intrusion detection system (IDS), an increasingly common component of layered security, will “raise the bar” even higher on the level of determination needed by an attacker to effectively penetrate your network [3]. Intrusion detection systems are only one tool in the network security administrator’s toolbox, but it can be a powerful tool when configured and implemented properly. However, IDSs cannot do it alone, and like all security tools, should be used along side other security measures to provide defense in depth.

References

- [1] Symantec, “ Getting Started- Netprowler 3.5” support manual. URL: ftp://ftp.symantec.com/public/English_us_Canada/products/netprowler/3.5/manuals/getting_started.pdf
- [2] Robert Graham homepage. FAQ: Network Intrusion Detection Systems URL: <http://www.robertgraham.com/pubs/network-intrusion-detection.html>
- [3] Symantec homepage. Wells, Mark and Thrower, Woody. Intrusion Detection. “The

Importance of Layered Security.” URL:

<http://enterprisesecurity.symantec.com/content/featurearticles.cfm?PID=7758820>

[4] Meinel, Carolyn “The ABCs of IDSs” URL:

http://www.messageq.com/security/meinel_2.html

[5] Elson, Del. “ Intrusion Detection, Theory and Practice” URL:

<http://www.securityfocus.com/focus/ids/articles/davidelson.html>

© SANS Institute 2000 - 2005, Author retains full rights.