Top Layer

perfecting the art of network security

White Paper

Beyond IDS:
Essentials of Network Intrusion Prevention

Abstract

*Network **Intrusion Prevention Systems(IPS)** are able to take immediate action to stop attacks and intrusions missed by the firewall. IPS may be the best technology in which to make your next IT Security investment.*

November 2002

Table Of Contents

Top Layer Networks, Inc. reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of Top Layer™ to provide notification of such revision or change. Top Layer provides this documentation without warranty, term, or condition of any kind, either implied or expressed, including, but not limited to, the implied warranties, terms or conditions of merchantability, satisfactory quality, and fitness for a particular purpose. Top Layer may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

Introduction

A new breed of products, network Intrusion Prevention Systems (IPS), has emerged. For years, the philosophy behind Network Intrusion Detection could be summarized as "Detect as many attacks and intrusions as possible, and report them, so that others may take action." In contrast, Network Intrusion Prevention Systems have been designed with a new philosophy: "Take decisive action on those attacks or intrusions which can be accurately detected."

What makes them so different? First, they sit *inline on the network*, where they not only monitor traffic, but also *actively intervene* by limiting or dropping traffic that is deemed malicious, terminating suspicious sessions, or taking other action in response to an attack or intrusion.

Background - Network Security Infrastructure *Before* Intrusion Prevention

A generation ago, few organizations even had an IT security administrator; today it's one of the most critical jobs around -- and one of the toughest. In the past several years, as organizations have grown increasingly reliant on their data and their Internet connectivity, the IT team has faced intense demands on system and content availability. In turn, IT departments have rapidly expanded their system, storage, and network infrastructures.

The rate of Internet dependence for business-critical functions has simply outpaced the ability of IT security staff to adequately address the new challenges. Financial constraints have added to the problem. While most organizations understand that security is more than just products, the bottom line is that today's security infrastructure is not up to the task at hand.

When organizations first began experiencing the insecurity of internetworking, they placed barriers to entry on their networks – firewalls – akin to the locks on their doors and windows. Firewalls completely bar those entrances – TCP and UDP ports for specific Internet protocols – through which no traffic should be allowed to pass. In addition, they enforce access control over the ports they leave open, so that only traffic from desired IP addresses can get through.

Firewalls have proven effective against many types of intrusions. But of course, organizations can't use a firewall to simply block everything from passing through – the company might as well disconnect from the Internet-connected world. And we have learned that attackers will learn to exploit any entry left open. Hybrid attacks, Denial-of-Service (DoS) attacks, and protocol anomalies get through most firewall deployments.

Many companies have already installed (or are currently considering) a second set of security devices, analogous to posting closed-circuit video cameras. These are Network Intrusion Detection Systems (NIDS), which inspect the network traffic and report their findings to log files and databases.

Like the video camera, the NIDS stands to the side and watches all that transpires. It may send an alarm to the administrator. (Although, like the ubiquitous car alarms on city streets, they are so often false that people soon stop responding.)

According to the 2002 CSI/FBI Survey [1] ninety percent of respondents (primarily large corporations and government agencies) detected computer security breaches within the last twelve months. 80% acknowledged financial losses due to computer breaches. Yet 89% of these same respondents already have firewalls, and 60% already use NIDS. Clearly, those commonly installed products have not yet addressed today's network security challenges.

Costs of Cyber Crime – Inconvenience, Loss of Productivity, Outright Financial Loss

Financial losses due to cyber crimes occur when certain undesired activities take place, affecting network and computing resources.   The motivation for organizations to implement network security infrastructure should be to detect and stop these activities before they cause damage that will have negative financial implications on the organization. A single catastrophic attack, on a Web server or network, can cost a company immensely. An attack might prevent or delay business transactions, compromise proprietary information, idle workers, and force replacement of systems.

But smaller attacks hurt as well, often in ways that cannot be easily quantified. Workers' productivity decreases when the network experiences problems due to worms or other attacks. Even with a NIDS installed, an IT department must spend time and resources to determine what happened and its validity (or just another false positive), assess whether systems have actually been compromised, mitigate the problem, and return the network or server to its proper configuration. These can be difficult and time-consuming tasks.

In recent surveys, the vast majority of large organizations reported security breaches in the past twelve months. Attacks are inevitable. However, many attacks can be blocked before they intrude into network operations. And when intrusions do occur, the faster an organization mitigates the problem, the less the costs associated with downtime and remediation.

Mitigation does not simply mean stopping the ongoing exploitation of a vulnerability. An intrusion has multiple facets. For example, a worm entering the network will continue to spread and launch new attacks. All of the ongoing effects of the intrusion must cease. And that still leaves the task of cleaning up the damage done, and returning systems to their normal working function. This can take hours, days, or more – resulting in an ever growing cost to the organization.

For example, when the Nimda worm appeared, the resulting traffic caused most organizations to disconnect all their web servers from the network until systems could be patched. Since organizations have become highly dependent on web-based business processing applications, this had an enormous impact on the day-to-day running of those organizations. For example, *Computer Economics* [2] states the worldwide economic impact of Code Red was estimated at $2.62 billion, the worldwide economic impact estimate of SirCam was $1.15 billion, and the estimated worldwide economic impact of Nimda was $635 million, and growing.

Bad Traffic – How Network Activity Can Cause Financial Losses

Figure 1, below, illustrates a set of relationships between primary objectives of the Network Security Administrator and those network-based activities that can negatively affect them. For example, a Server-Resource-Exhaustion Targeted Denial of Service attack can impact the ability of an organization to complete Internet-based transactions and make its server systems unavailable. Likewise, unauthorized modification of content files on web servers can impact the same objectives, and is usually considered vandalism that must be repaired.

	Malicious Network Activity																
Security Objectives	Employee Browsing to Non-Work-Related Sites	Network Eavesdropping by Employees or Outsiders	Session-Hijacking	Unauthorized System or File Access	Password Cracking from the Internet or Internally	Port Scanning from Internet	Computer Virus Infection on Desktops or Servers	Worm Propagation via email	Trojan Horse Program Infection via Network	Malicious Mobile Code execution via Web Server	Modification of Content Files on Web Server	Modification of System Files on Web Server	Server-Resource Exhaustion-Targeted DoS Attacks	Server-Vulnerability-Targeted DoS Attacks	Network Flood Denial of Service Attacks	DNS-based traffic redirection or site hijack attack	SPAM email from Internet
Enable Internet Communication															X	X	
Enable Internet transactions										X	X	X	X	X	X	X	
Provide reliable desktop and server computing resources							X	X	X	X	X	X	X	X	X	X	
Prevent Undesired or Inappropriate Use of Network or Computing Resources	X	X	X	X	X	X											X
Prevent Loss of Privacy		X	X	X													
Prevent Theft		X	X	X													
Prevent Fraud					X	X										X	
Prevent Vandalism				X	X	X	X	X	X	X	X	X	X	X	X	X	X

Figure 1.

It may never be possible to thwart *all* activities that can lead to cyber crime losses, but due diligence requires that organizations employ a variety of infrastructure elements, including firewall, VPN, Anti-Virus, Anti-Spam, network and host-based intrusion detection, and now, Intrusion Prevention Systems.

Unfortunately, there is often no single solution that will encompass all of these infrastructure elements, and in certain circumstances 'all in one' solutions can be detrimental. In fact, it is sometimes not clear which solution - network-based vs. host-based -- will provide the appropriate protection. Figure 2 illustrates the relative coverage against attacks and intrusions that can be obtained from host-based systems vs. network based systems. At Top Layer, we believe that customers will need to deploy both technologies in order to achieve comprehensive protection against malicious activities.

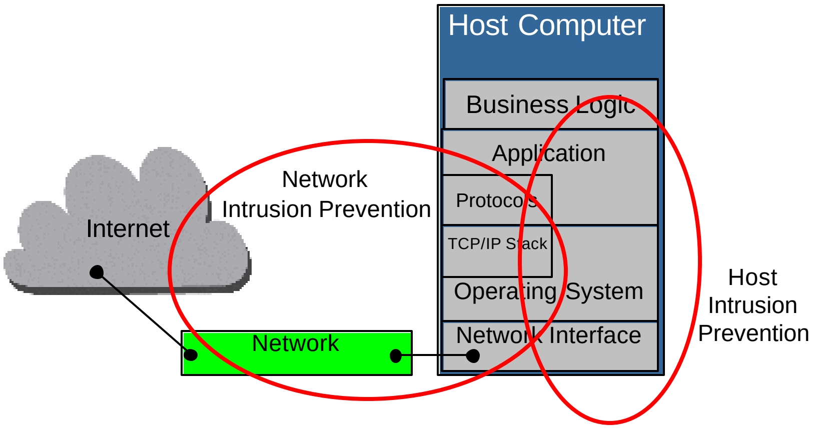

Figure 2.

There are thousands of attack tools created for performing these malicious network activities. In the next section, we'll take a closer look at some of the more common ones.

Malicious Network Activity - Common Attack Types

Attacks can affect a variety of critical network -- attached resources, including systems, files, programs, or the bandwidth that the organization relies upon for communication.

The following categories represent some common types of network security attacks. It is important to remember that while certain attack categories predominate today, that can change quickly. It took a very short time for HTTP worms like Nimda and Code Red to become the single most prevalent problem facing networks; another new attack could become number one just as quickly.

Protocol exploitations. The most common attacks on network resources come from a group of methods used to exploit the protocols an organization uses to communicate across the Internet. These attacks generally involve someone performing suspicious or malicious activities on the network. Some of these are simple probes – scanning ports to find vulnerabilities, for example. In themselves, such probes do not compromise the network, but often lead to attempts to exploit the openings, with login attempts or sneakier methods, such as IP spoofing. Further exploits using network protocols can occur once the network is first compromised.

Finally, users *inside* the organization may try some malicious behavior on the network; these violations can be the most damaging of all.

A NIDS often has no problem with detecting these sorts of attacks; however, they can be so frequent or generate so many alerts (for example, a simple port scan can generate 30-40 different alarms in some NIDS) that they can hide more serious attacks.

HTTP attacks. Worms are computer programs designed to replicate and propagate themselves automatically. Some merely spread, and spread, and spread, clogging up computer and network space. Others destroy their host computers as they move. HTTP worms like Nimda – which has been infecting machines by the thousands for more than a year – may spread first through an infected email, and then use that compromised host to scan for and infect other vulnerable systems. By using a host computer to look for other victim systems, Nimda actually works from *behind a company's firewall*, infecting trusted sites on the intranet, or even on an extranet shared with business partners. Antivirus software that scans emails will never see the resulting spread of the worm. A NIDS sensor often generates thousands of alerts from the activity of a single infected host, without stopping any of it.

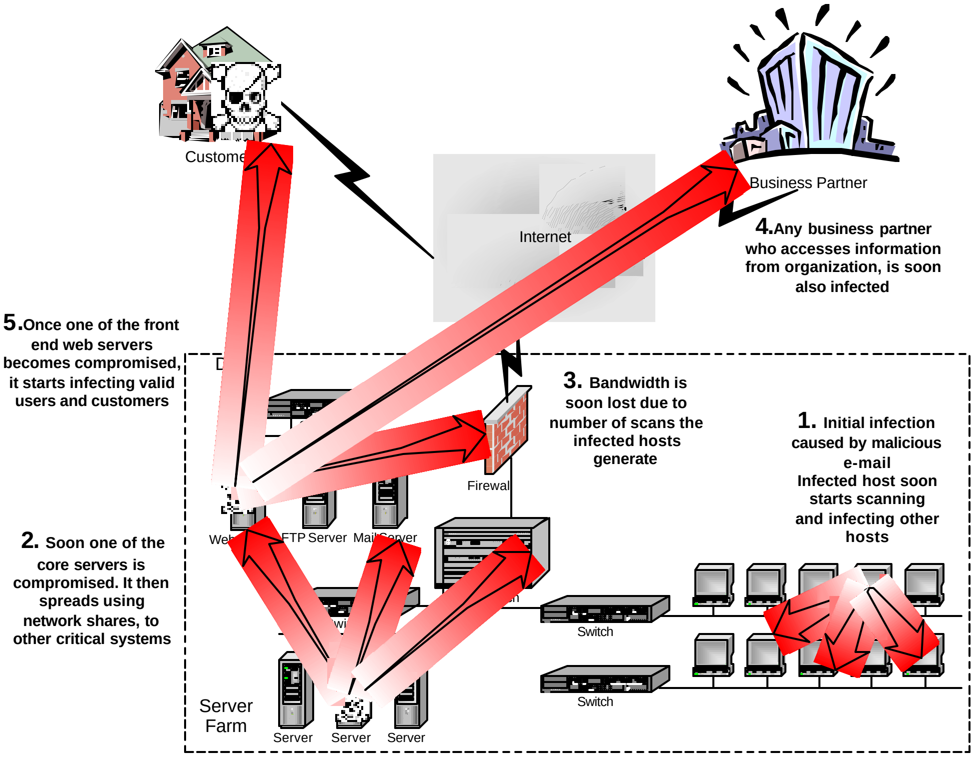
Custome
Business Partner
Internet
4.Any business partner who accesses information from organization, is soon also infected
5.Once one of the front end web servers becomes compromised, it starts infecting valid users and customers
3. Bandwidth is soon lost due to number of scans the infected hosts generate
1. Initial infection caused by malicious e-mail Infected host soon starts scanning and infecting other hosts
Firewal
Web
FTP Server
Mail Server
2. Soon one of the core servers is compromised. It then spreads using network shares, to other critical systems
Switch
Switch
Server Farm
Server
Server
Server

Figure 3.

SYN flood attacks. Denial-of-service attacks hit the headlines from time to time when they cripple or shut down high-profile web sites. But, they happen with great frequency and to organizations of all sizes. By flooding servers with repeated traffic, these attacks can slow the network, prevent legitimate use, or even completely stop the server from working.

SYN attacks demonstrate how simple it is to disrupt a network. SYN communications initiate contact between two computers that use TCP/IP protocols, like a handshake. The receiving computer responds to every SYN request, so when faced with an endless, repeating string of SYNs – perhaps with a phony source address – it keeps trying to shake hands with these elusive friends. It opens an ever-increasing number of incomplete connections, using up its resources, until finally the system crashes.

Standard firewall implementations are ineffective against stopping SYN flood attacks. Some vendors offer plugs-ins that monitor SYN activity, but these often come at a huge cost to the performance of the firewall.

Most NIDS sensors can detect SYN floods, but because they are often deployed on switch mirror ports passively (that is, outside the network looking in), they lack sufficient context to understand the changing nature of SYN activity. For example, a web cache normally creates many more open SYN connections than a file and print server. Lacking this context, NIDS sensors tend to generate numerous false positive alerts for SYN flood attacks.

FTP attacks. File Transfer Protocol (FTP) is a commonly used Internet protocol for transferring files. However, due to a design flaw, it is vulnerable to a "bounce" attack. All that the attacker needs is a target user's IP address. The attacker uses an FTP server to open a connection with the target user, and uses this seemingly innocent connection to send malicious code past the firewall. Often this code then establishes direct access back to the attacker, who then has full access to the target user's computer.

ICMP attacks. To perform diagnostics on the network, network administrators commonly use the ICMP protocol. Hackers, naturally, exploit that opening. They send a phony request (ping) through the port, prompting a reply from inside. This can be used to gain information, flood a system, or even plant bits of instruction that can result in greater vulnerability – for example, using ICMP responses to piggyback commands to compromised machines.

Application attacks. Applications running on Web servers are rarely bug-free, and when one hacker discovers how to exploit one, word quickly spreads through the hacker underground. Often the attacks focus on "buffer overflow" – a common vulnerability triggered by overloading system memory. Clever hackers have learned how to exploit these vulnerabilities to get malicious code, such as worms and Trojan horses, to execute on the target systems, under their control, and often unbeknownst to the system owner.

Defense Mechanisms – Detection Is Only The First Step

The limitations of firewalls were clear almost as soon as the products were first introduced. The Internet could clearly enable workers to do great things, but to do those things they needed to use all manner of protocols such as: SMTP, DNS, FTP, and HTTP. Network security administrators had to configure the firewalls to allow these protocols to some extent. This meant that the firewalls had no hope of stopping all attacks. Thus, network security managers needed a way to discover successful intrusions that had breached the firewall.

The category of products that arose to meet this need was Network Intrusion Detection Systems (NIDS). These products sit beside the network, offline, and watch the traffic as it goes by.

Initially, the strategy was to ask the NIDS to look for evidence of known threats taking place. Later, new methods were added to detect anomalies – network traffic that doesn't look like normal traffic.

The first step in stopping attacks and preventing intrusions is to accurately and reliably detect them. A variety of detection methods can be applied to detect certain kinds of attacks and intrusions:

Pattern-matching. Pattern-matching algorithms can be used to identify attacks and exploits. Traffic is scanned to determine the signatures of known attack patterns, and the NIDS sensor analyzes each packet that is sent to it constantly looking for evidence of those patterns. See Figure 4.

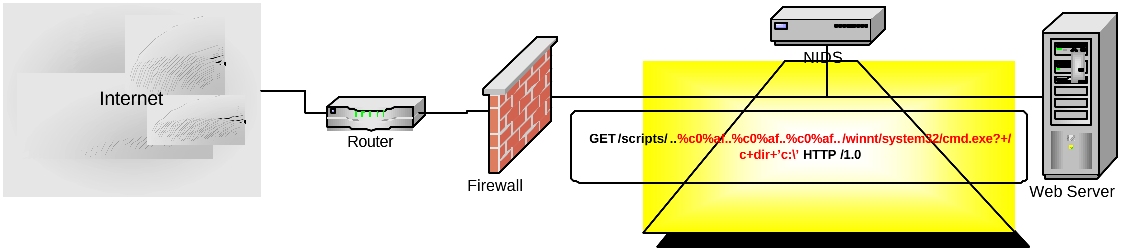

```
GET /scripts/ ..%c0%af..%c0%af..%c0%af.. /winnt/system32/cmd.exe?+/
c+dir+'c:\' HTTP /1.0
```

Figure 4.

NIDS sensors that search for attack signatures are like a cashier with a list of “bad check” names taped on the wall. The system catches only known perpetrators – and only if the list is constantly updated. Unfortunately, the people out there trying to exploit networks are neither lazy nor stupid; they constantly unleash new variations and new attacks. With each new attack, new signatures have to be “taught” to the NIDS. Over time, this has led to a need for NIDS products to hold literally thousands of attack signatures, and constantly scan for them all, while constantly updating the signature library. Attackers know exactly how the NIDS products work, so they continuously make slight alterations to elude detection. A more intelligent signature mechanism is needed to bring about more accurate results.

Stateful matching. Intelligent signature pattern matching greatly improves upon generic methods. It looks at context and placement of the signature to make smarter decisions about whether it represents an attack. The NIDS will keep track of the state of the connection with the outside entity, and consider the broader context of all the transactions initiated during the connection.

Protocol analysis. Newer systems attempt to save time by first identifying the protocol, and then looking specifically for anomalous activity or attack patterns relevant to that protocol. By doing so, it can do a much more targeted, and thus more effective search. This makes the NIDS better at finding attacks hidden within network traffic; however, it is still ultimately reliant on a list of known attack patterns.

Implementation Choices: NIDS vs. IPS

NIDS Deployments

NIDS products have undoubtedly helped organizations fight intrusions, but as attacks have increased, the limitations of the passive approach have become apparent. Deployment and operational issues plague the approach. Improvements to the products cannot change the inherent problems. It has become clear that *detection* is only part of the desired solution.

NIDS Advantages

NIDS have one major advantage: Since they are deployed offline, there is no way for them to cause network interruptions. However, taken to an extreme, doing nothing to protect your network has this same quality.

NIDS Disadvantages

Only detects, doesn't prevent. Because they sit off of the network, it is difficult for a NIDS to affect traffic. NIDS are generally powerless to stop, or even slow down, an attack in progress. And, given the preponderance of false alarms, nobody would want a NIDS to drop data packets without manual confirmation of an attack.

Lengthy time to mitigation. The need for manual intervention – and the overwhelming amount of data to comb through – results in a slow response to an attack or intrusion event. For effective protection, immediate action is required to minimize the window of vulnerability, and minimize the damage and costs associated with the attack.

False alarms. Anyone who has administered a NIDS knows the rush of emotion the first few times the system reports an intrusion in progress – and the frustration each time the event turns out to be benign. Eventually, the alarms spark less reaction, and ultimately, annoyance. NIDS sometimes simply mistake benign traffic for an attack – a "false positive" response. These mistakes often require considerable time to inspect the packets in question and determine that the NIDS has erred. Other times, the NIDS correctly identifies a particular traffic pattern matching a signature – but the pattern turns out to be normal for that particular organization. These "false alarm" alerts can cause a network security manager to temporarily block legitimate traffic being mistaken for an attack.

And even worse, false alarms can desensitize people to real intrusions. As a *Network World* report in July, 2002 [3] concluded: "When real attacks came along, some [NIDS] products...buried the reports so deep in false alarms that they were easy to miss."

Excessive log data. Because a NIDS does not actively stop most attacks, it must generate logs of reports of anomalous or questionable network activity. Network security managers must spend time reviewing these logs, as well as network traces and other diagnostic methods, to sort out what has happened and what must be done. A NIDS vendor's own Managed Service Provider has demonstrated the extent of the problem. The vendor reported over 21 *million* alarms in a three-month period – but in all that noise there were only 1,482 actual incidents requiring remediation. [4]

Here's another example: Top Layer frequently performs simple experiments by placing an unhardened host on the Internet (for example, a default build of Microsoft Windows 2000 with IIS, or Linux with Apache) with a NIDS to monitor what happens. Typically, the system will be compromised in less than four hours and sometimes in a matter of minutes. Figure 5 shows the data from such a test carried out with a single Windows 2000 web server with a leading NIDS to monitor what happened over a 12-hour period. Over 32,000 alerts were generated, 15,700 of which were High or Medium alerts (49%) and most were related to Nimda and Code Red. Note that all traffic recorded was malicious, as no legitimate services were actually available.

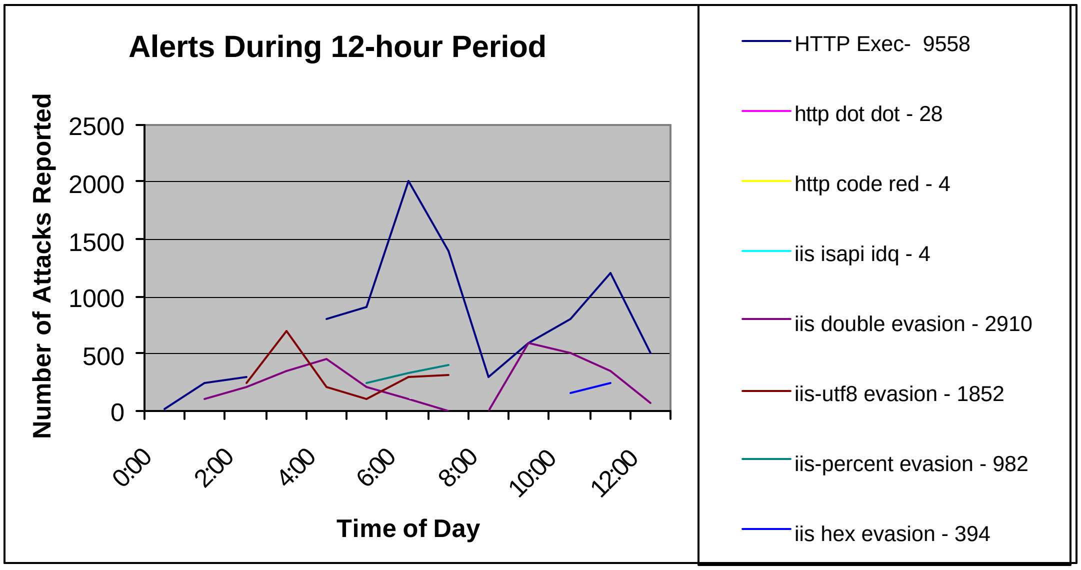

Figure 5.

Deployment issues. NIDS solutions must be placed at proper points at all entryways to the organization's systems and servers, must be properly configured, and must be constantly updated. In practice, this has proven frustratingly cumbersome and expensive. This results in additional work time to administer the sensors, as well as lower performance from a poor configuration.

Poor results. Even as vendors have improved NIDS through several generations of products, and even as they have compiled vast libraries of attack signatures, few organizations feel more secure now than they did eight years ago when the first NIDS products began to appear. Attacks – even known, common attacks – continue to occur and succeed. Subtle variations of attacks have confounded the products. Perhaps most importantly, new attacks, like Nimda, have moved so quickly that tremendous damage has been done before NIDS users can respond.

Top Layer believes that the goal of the organization should be to shrink the window of exposure caused by an attack. Intrusions into the network or server should be reined in quickly – ideally, the window will be zero as the attack is blocked. It is imperative that the intrusion event immediately triggers a sequence of actions, often called an Intrusion Response, which results in rapid mitigation and correction of the situation.

IPS Deployments

An IPS sits in-line, ideally inspecting all packets going inbound or outbound. It performs a range of detection analyses, not only on each individual packet, but also on network conversations and patterns, viewing each transaction in the context of others that have come before or will go after.

If the IPS deems the packet harmless, it forwards it as a traditional Layer 2 or Layer 3 network element. End users are unaware of any effect. However, when the IPS detects suspicious traffic, it can then initiate one of many response mechanisms. It may limit the traffic, by forwarding it normally up to a certain bandwidth or a certain number of TCP connections. Or, the IPS can discard the packet completely.

Of course, an IPS must also have an extensive reporting mechanism – but this must be more than a simple log of activity. The IPS can create an alarm and transmit it to appropriate destinations. It can send copies of the actual traffic out through a forensic port for immediate analysis and diagnosis by IT security personnel. It can even create an entire, ongoing Flow Mirror ™ copy of the session traffic to send to a mirror port.

IPS Advantages

Speedy end to intrusions. As discussed earlier, an intrusion event begins a process of harm to an organization's computing resources – not to mention potential legal liabilities. By stepping in at the moment of detection, an IPS rapidly ends the intrusion and minimizes the overall time before the network is back to normal.

Accurate and reliable detection. By using multiple detection methods, and utilizing its position in the line of network traffic, the IPS can detect attacks and intrusions more accurately and reliably. By relying less on signatures and more on intelligent methods of detection, the IPS generates far fewer false alarms. This focuses the organization's time and effort on only the true threats.

Active prevention. Whereas a NIDS simply announces the presence of suspicious or anomalous traffic, an IPS can instigate a variety of response mechanisms as described earlier. This reduces the costs of administering network security, and reduces the risk of the organization suffering damage or loss due to cyber attacks.

IPS Requirements – What To Look For

Unfortunately, the term "Intrusion Prevention System" is being used indiscriminately to describe a variety of security technologies and solutions. This paper focuses on Network Intrusion Prevention Systems that are able to automatically take action to block attacks and intrusions without manual intervention. Top Layer recommends that organizations look for Network Intrusion Prevention Systems that have the following characteristics:

- An Inline device capable of accurately and reliably detecting and precisely blocking attacks – ***Accuracy and Precision***

- Operates at line speed with no negative impact to network performance or availability – ***Good Network Citizenship***

- Integrates effectively into security management environment – ***Effective Security-Focused Management***

- Needs to easily accommodate prevention for future attacks – ***Anticipates unknown attacks and easily accepts signatures for newly discovered attacks***

Figure 6 illustrates that Network Intrusion Prevention is a new layer of protection in the network security infrastructure, blocking the attacks and intrusions that pass through the firewall.

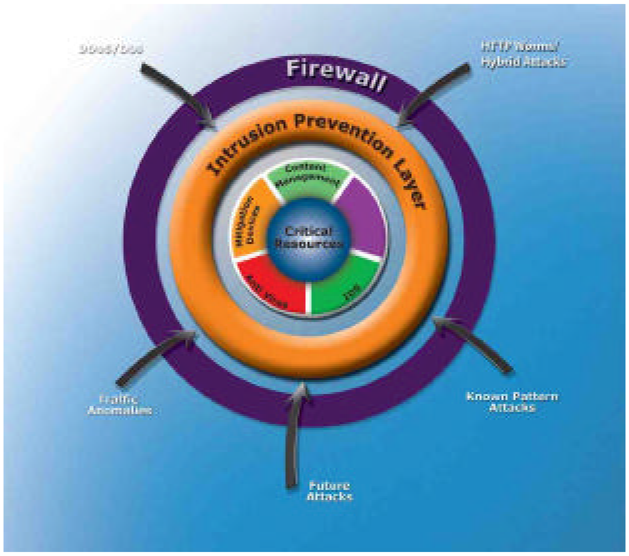

Figure 6.

Accuracy and Precision. As mentioned earlier, a significant problem with NIDS products to date has been the numerous false results generated by the detection methods. While this is extremely problematic in a NIDS, it is absolutely unacceptable in an IPS. Inaccurate detection can result in response mechanisms affecting legitimate traffic, frustrating users. Top Layer believes that its combination of detection methods – methods beyond those used by NIDS sensors -- achieves the necessary high level of accuracy and reliability.

Good Network Citizenship. The IPS is not a bystander; it is an integral part of the network. As such, it must stand up to any strain the organization may place on it. It must be a good network citizen, as judged by performance, reliability, and availability. Performance describes the ability of the IPS to keep the traffic flowing on the network. Poor performance in a heavy -traffic environment will result in slowed network performance, or even lost packets. Reliability refers to the ability of the IPS to perform its functions properly, without interfering with other systems on the network. Availability refers to the amount of downtime of the product, due to shutdown, crashes, or maintenance.

Effective Security-Focused Management. An IPS gives the network security administrator a great many options, since it is capable of not only detecting attacks and intrusions, but also directly affecting network traffic through limiting or blocking. It must give the administrator an easy interface for setting and changing configurations on the devices.

In addition, a true IPS solution should not simply stand alone, but operate as an integral part of a Security Integrated Management suite, ultimately cooperating with firewall, NIDS, anti-virus, and vulnerability-assessment products and functions.

Anticipates unknown attacks and easily accepts signatures for newly discovered attacks. An IPS must have flexible and seamless methods to update not only new attack signatures, but also capabilities to respond to entirely new classes of attacks using firmware or software upgrades. In addition, IPS systems should have methods that are able to respond to new attacks without the need for signature updates. Such methods may include inverse exclusion, where all requests, except those that are legal for a given destination, are dropped; protocol validation, where illegal request methods are dropped; or attack-independent-blocking, where hostile attackers are identified, and all traffic from the attacker is dropped, regardless of whether the attacks are known or not.

Top Layer's Intrusion Prevention System Solution

The best and most reliable IPS solutions will come from companies that have proven experience handling inline network traffic. Top Layer has been deploying successful inline networking solutions for customers worldwide since 1999. We have leveraged that experience, along with our intrusion and attack expertise, gained from years of deploying IDS balancing solutions, to bring to market a high-performance and reliable network intrusion prevention system, called the Top Layer Attack Mitigator™ IPS.

Top Layer believes that its extensive and thorough approach to accurate and reliable detection, flexible and powerful response, and detailed reporting and management distinguishes its IPS solution from others.

How the Attack Mitigator–IPS Works

Top Layer's Attack Mitigator IPS is a highly reliable, high performance, ASIC-based, forwarding element that is able to perform accurate and reliable detection of a large and growing number of network attacks and intrusions, and is able to precisely control traffic by forwarding, limiting, or discarding.

Top Layer's IPS is based on its TopFire™ Network Intrusion Prevention and Response Engine architecture, which is optimized to operate on both inbound and outbound traffic, using accurate detection and precision control to allow good traffic through as fast as possible, stop known bad traffic as fast as possible, and analyze suspicious traffic until it can be classified as one or the other. Figure 7 illustrates the TopFire architecture.

Top Layer Stateful Network Intrusion Prevention and Response Engine

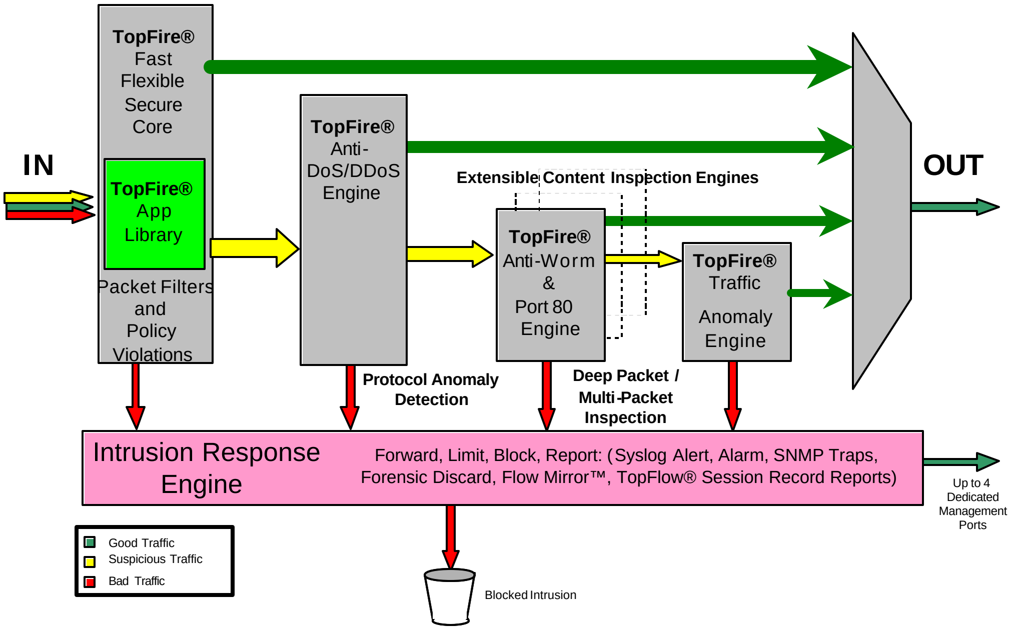
TopFire® Fast Flexible Secure Core
TopFire® App Library
Packet Filters and Policy Violations
IN
TopFire® Anti-DoS/DDoS Engine
Extensible Content Inspection Engines
TopFire® Anti-Worm & Port 80 Engine
TopFire® Traffic Anomaly Engine
OUT
Protocol Anomaly Detection
Deep Packet / Multi-Packet Inspection
Intrusion Response Engine
Forward, Limit, Block, Report: (Syslog Alert, Alarm, SNMP Traps, Forensic Discard, Flow Mirror™, TopFlow® Session Record Reports)
Up to 4 Dedicated Management Ports
Good Traffic
Suspicious Traffic
Bad Traffic
Blocked Intrusion

Figure 7.

In the above illustration, all packet flow occurs from left to right regardless of whether the traffic is in the inbound or outbound direction. The Attack Mitigator IPS performs its accurate detection and precision control on traffic in both directions.

Fast, Flexible, Secure Core is Top Layer's patented Layer 4-7 stateful inspection engine. This engine identifies traffic on a conversation-by-conversation basis, creating a detailed internal session record for each conversation. Once a conversation (such as a TCP connection) has been established, it is characterized by this engine, and subsequent packets in this conversation are immediately identified by the TopFire ASICS as belonging to the same session, and processed at extremely high speeds.

This engine also implements most of the single-packet attack and IP fragmentation attack filters, as well as blocking traffic when identified by policies using IP addresses and/or TCP/UDP Port numbers. Traffic belonging to known good sessions is forwarded at high rates to the output. Traffic identified as malicious is sent to the intrusion response engine. Traffic that is suspicious is sent along to the next engine.

Anti-DoS/DDoS engine. The Attack Mitigator IPS uses its Anti-DoS/DDoS engine to identify Denial of Service attacks, including distributed attacks that can confound other devices by appearing as a series of apparently harmless individual packets from disparate sources. This engine uses patent-pending TCP connection behavior analysis to detect DoS/DDoS attacks such as SYN flood attacks, even in cases where the attack is initiated by multiple, random source addresses.

This engine automatically assesses the threat level of up to 200,000 Source-IP addresses based on their individual and collective patterns of TCP connection behavior. Anomalous TCP connection behavior is an excellent indicator of a Denial of Service attack. Traffic from trusted IP sources is forwarded normally (assuming there is no other reason to block it); Traffic from suspicious IP sources is proxied by this engine, waiting for a complete TCP handshake before passing the session along to the destination server; Traffic from malicious IP sources is sent to the Intrusion Response Engine, usually for discarding. Figure 8 shows how the Attack Mitigator IPS management interface provides a snapshot of current threat-level assessment to the security administrator.

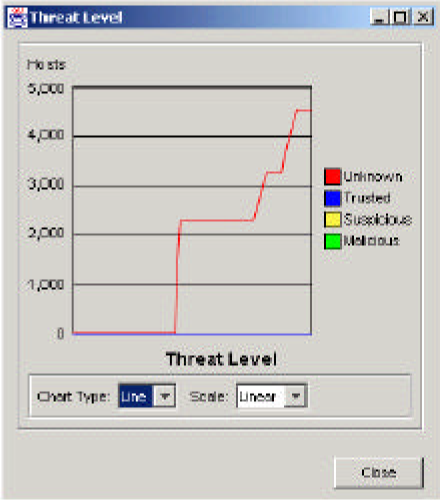

Chart Type: Line
Scale: Linear

Figure 8.

Anti-worm and Port 80 engine. This is where the Top Layer Attack Mitigator IPS puts its intelligent application signatures to work. This engine uses stateful normalized matching algorithms that allow it to identify exploit variations that generic pattern-matching can only spot with new signatures. Furthermore, Top Layer has developed the industry's best methods for accurately detecting HTTP attacks, using HTTP protocol validation, normalized string matching and URI length-checks.

Top Layer's use of HTTP URI signatures to detect HTTP exploits provides a good example of its ability. Most NIDS require a different signature for each variant of each exploit. To reliably detect a simple (but very common) directory traversal attack, the NIDS needs as many as twenty different signatures for each Unicode variant. In contrast, the Attack Mitigator IPS actually decodes the URI (as a web server would) and looks at the result to see whether the '../..' exploit is included. This requires just one signature.

Traffic Anomaly engine. Using an intelligent traffic-analysis system, the Attack Mitigator IPS performs not only protocol-anomaly detection and signature detection, but also traffic-anomaly detection. This means that it goes beyond identifying suspicious packets, to catching suspicious traffic behavior. When specified applications exceed normal traffic levels, as previously calibrated by the Attack Mitigator IPS and the security administrator, traffic can be limited or blocked. In addition, the number of simultaneously active sessions (such as TCP connections) can be limited as well. This powerful attack mitigation feature is especially adept at mitigating application layer attacks, where all protocols are adhered to, and no obvious attack is present, but the goal of the attack is to overwhelm system resources.

Intrusion Response / Management engine. The key to successful deployments of Intrusion Prevention Systems is the Intrusion Response Engine. The Top Layer Attack Mitigator IPS lets network security managers configure detection, response, and reporting mechanisms according to the needs and the policies of the organization. Settings can be heightened from "disable" (no detection, no response, no reporting) to "monitor" (detection and selected reporting mechanisms, but no response), and finally to "mitigate" (detection and selected response and reporting mechanisms) as administrators become more confident that the IPS can do all this without adverse effects.

Mitigation settings in the Intrusion Response Engine include limiting traffic by bandwidth, limiting the number of simultaneously active TCP sessions, or blocking packets or sessions outright.

Top Layer's suite of products provides several important security management features. A dedicated management port is separate from three maintenance ports: one configurable Flow Mirror port where the organization might place a network analyzer; a forensic discard port where filtered packets can be copied for safe, offline observation; and a configurable port for session reports. Summary Security Reports provide an at-a-glance overview without the need to pour through cumbersome logs. And device management is provided through secure protocols such as HTTPS, SSH, password-protected local serial console, and syslog messages. SNMP, HTTP, and Telnet are disabled by default for added security.

Attack Mitigator IPS Deployment Scenarios

Figure 9 illustrates the many possible deployment scenarios for the Top Layer Attack Mitigator IPS, including behind the perimeter gateway firewall, in front of the DMZ, in front of internal server farms, and on extranet links, behind the VPN endpoint. Since the Attack Mitigator IPS provides both inbound and outbound protection, it is useful against both internal and externally originated attacks and intrusions.

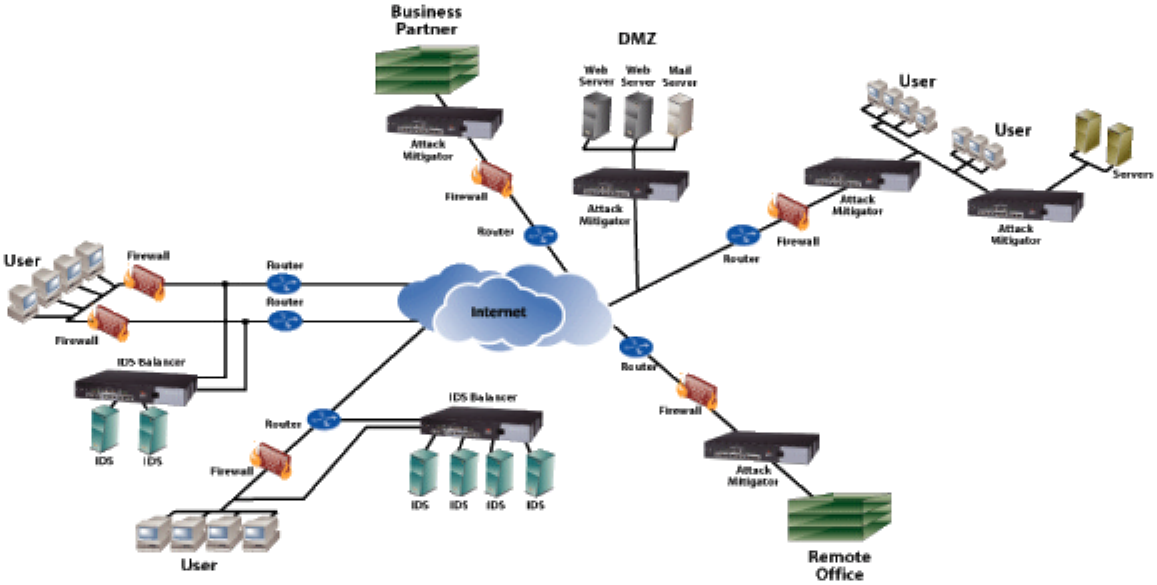

Figure 9.

Conclusion

Intrusion Prevention Systems represent a new and promising technology in network security. Network intrusion prevention devices can automatically take action to stop attacks and intrusions. IPS offers protection that a NIDS cannot.

An IPS should be an inline device capable of accurately and reliably detecting and precisely blocking attacks; it must operate at line speed with no negative impact to network performance or availability; it must integrate effectively into security management environments; and finally, it needs to easily accommodate prevention for future attacks, both known and unknown.

Organizations looking for an IPS should pay special attention to the network capabilities of the products. The network citizenship aspects of the product will determine the ultimate success of the deployment.

Top Layer believes that its Attack Mitigator IPS meets all of these important criteria, providing the accurate and reliable detection, precision blocking, excellent network performance, effective management, and support for future attacks. Top Layer is strongly positioned to provide this solution, with its extensive network experience, its inline intrusion prevention experience, and its extensive attack-identification methods.

With the Attack Mitigator IPS, an organization can finally realize the protection benefits that have been promised, but not delivered, by the existing security infrastructure.

References

[1] CSI/FBI Computer Crime and Security Survey SI Computer Security Issues & Trends 2002.
[2] Computer Economics (www.computereconomics.com).
[3] Crying wolf: False alarms hide attacks Newman, Snyder ,& Thayer *Network World, 06/24/02*
http://www.nwfusion.com/techinsider/2002/0624security1.html

[4] ISS Internet Risk Impact Summary – June 2002.

About Top Layer Networks

Founded in 1997, Top Layer Networks delivers proven network security solutions worldwide, enabling enterprises to protect against cyber threats and scale their infrastructure to meet new, ever increasing security demands. The Company's intrusion prevention products are built on a patented, ASIC-based architecture. The products are engineered to block high-volume DoS and DDoS attacks, HTTP worms, traffic anomalies and unknown attacks; improve the effectiveness of intrusion detection systems through intelligent balancing and distribution of traffic; and enhance the availability and performance of firewalls through firewall/VPN balancing technology. Top Layer Networks is headquartered in Westboro, Massachusetts with sales and support presence in Australia, France, Germany, Japan, Korea, Malaysia, Singapore and the United Kingdom.

Top Layer Networks, Inc. **www.TopLayer.com**
2400 Computer Drive • Westboro, MA 01581 USA • 508.870.1300 • Fax 508.870.9797