



# 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>

## Intrusion Prevention: Are We Missing the Point?

Mark Orlando

GIAC Security Essentials (GSEC), Version 1.4b

February 27, 2003

In this paper, I intend to closely examine prevalent intrusion prevention technology and functionality. I will evaluate both as compared to other security devices and as a relevant solution in and of itself. I will illustrate that intrusion detection cannot be a replacement for intrusion detection or any other security tool commonly deployed in networked environments. The design and functionality of the intrusion prevention appliance does not meet the goals of the security-minded IT professional, as it is not a sufficient information-gathering tool nor does it contribute to network security in-depth. Instead, newly developed intrusion prevention-based appliances contribute to a false sense of security, and a false sense of strengthening security when in reality (in many cases) security is weakened. This paper is not meant as an attack on intrusion prevention technology and their place in the IT security world. Rather, it is meant as a question of whether or not this technology is ready to be a viable tool in the front lines of network defense. I submit that, due to the young and still-developing nature of the technology, it is not. There is still much room for other technologies (specifically, intrusion *detection*) to grow, and these technologies should not be abandoned for the newer, seemingly more effective, intrusion prevention. My goal is for the observations listed below to support this conclusion.

Anyone associated with IT security will tell you that new technologies and updated methods for countering network attacks and misuse are popping up on a daily basis. By the time the latest intrusion detection system or firewall has been deployed, chances are there is already something faster, more robust, and generally better on the market. It has become a race against the exponentially increasing threats facing information systems; a race that can never really be won, but one in which the main goal is to stay one step ahead of the mounting threats. It should be no surprise, then, that one of the hottest new technologies in the security market attempts to do the work of multiple appliances at once: the intrusion prevention system (IPS).

While this is a somewhat generic term, the intrusion prevention system essentially works as a combination intrusion detection and firewall appliance. However, instead of using a preset list of alert signatures in scanning for abnormal (and possibly malicious) traffic or allow/deny policies, these devices operate on a baseline of "normal" behavior within a network. Whenever traffic deviates from this baseline, a rule is created and enacted to block it. The idea is that the appliance is not bound by knowledge of known vulnerabilities or retroactive response to an attack in progress, but that it can adapt to new and developing attacks by monitoring the operations of individual applications.

Of course, no security solution is perfect. There is no “catch-all” appliance that can identify and prevent attacks and misuse while at the same time maintaining a transparent presence on a network. Why then is there such a push for one appliance that does the work of many with less management? Some may reply that it decreases cost, effort to deploy multiple security solutions, and time and manpower to manage the devices. Obviously this is a prevalent concern with protecting an environment, but is intrusion prevention really the answer?

First off, let us take a look at what it takes to support such a device (it should be fairly apparent even at this point). Approaching it from a personnel standpoint, is it really an improvement over other managed security devices such as firewalls or intrusion detection systems? Considering the young age of the IPS technology, it would at the very least require someone of considerable technical experience and knowledge to properly deploy and configure the appliance. Even then, it would really require no lesser amount of monitoring than an intrusion detection system or firewall. There are still logs to parse through (if the administrator is doing his/her job), and there is still the need for 24x7 personnel responsible for the device, unless the systems are powered off nightly (which is highly unlikely). The fact that the IPS has the ability to actively block traffic creates the need for immediate response if there is a denial or interruption of service caused by the device. Considering the support aspect, a decreased cost factor in implementing an IPS versus an IDS and firewall is rapidly diminished. Many IDS vendors sell products that merely require log analysis and alert monitoring. If the IDS appliance fails, at the very least, traffic will continue to be allowed into the network.

OKENA's *Stormwatch* and TopLayer's *Attack Mitigator* both promise protection from denial of service attacks, worms and other widely-known exploits, and offer product lines that can handle high-traffic and Gigabit Ethernet network segments. The appliances are designed to operate from pre-determined signatures and established application baselines with little management or full-time monitoring. Such are the stated advantages they hold over most firewalls and intrusion detection systems, both of which require parsing through (sometimes extensive) logs and constant tuning.

Some organizations, such as the University of California at Berkley, credit intrusion prevention systems with protecting them from potentially harmful attacks. Earlier this year, the Slammer Worm wreaked havoc on the Internet, causing denial of service conditions and disabling countless networks all over the world. It did not affect systems at UC Berkley, however, as they had recently deployed OKENA's *Stormwatch* intrusion prevention system on their network. In this situation, it seems that massive amounts of time, money, and man power were saved by averting an incident thanks to intrusion prevention. (Hulme)

Additionally, UC Berkley has admitted that many of their servers were in fact vulnerable to the Slammer Worm. This almost assuredly multiplies their relief that such a comprehensive solution was in place, but would a more layered approach to network security have had equally successful results? It is a well-known fact that the most effective security measures are those deployed in

layers. Firewalls with well-planned and managed policies, intrusion detection systems, anti-virus software, and comprehensive logging and auditing all contribute to a more secure network environment. While any good security guru knows that intrusion prevention will not take the place of all of these, there are dangers in trying to bundle too much into one appliance; moreover, a more automated solution is not always the most effective when it comes to network security. Had the servers at UC Berkley been sufficiently patched (Microsoft released a patch for the Slammer Worm in 2001), this would not have been a concern. Any denial of service attacks generated by the Worm could have just as easily been blocked by perimeter defenses.

Even so, many IT professionals argue that intrusion prevention appliances are the next generation of intrusion management and response technology; many more are moving towards implementing intrusion prevention in the place of intrusion detection. Examining each of these technologies more closely, we can better understand the differences, shortfalls, and advantages of both.

There are several types of intrusion detection solutions currently on the market. There are network intrusion detection systems (NIDS), host intrusion detection systems (HIDS), and hybrids of both technologies. Some operate as separate appliances and some as software. The basic technology of an IDS is to passively monitor network traffic and events, match each against a database of attack signatures and application baselines, and send an alert to a log file or management tool whenever there is a match.

For example, Network Flight Recorder's NFR 300 series sensors perform in-depth signature and stateful analysis for many different network protocols. They examine all packets and fragments and perform packet reassembly for fragmented data transmissions. As data passes through the monitoring piece of the sensor, it is compared to a list of attack signatures to determine suspicious or possible malicious activity. When alerts are triggered by traffic that matches these signatures, they are transmitted to a central management server where they are compiled and fed into the administrator's GUI. From the graphical user interface (GUI), the administrator can monitor all alerts generated by all network sensors currently deployed. Filters can be applied and signatures can be modified for each sensor.

The specific data that is provided and method of parsing through it varies among product and vendor, as does granularity of the information the administrator is able to view. Often, a default baseline is enabled, with alert filters and traffic recording pieces being subsequently enabled to tailor the IDS to a given network. Every major IDS product provides some method of tailoring the appliance not only to the environment, but to changing traffic patterns as well.

One of the major issues with IDS is that it is reactive. With new exploits being released daily, it is next to impossible to keep an IDS up to date with alert signatures for each one. However, new solutions are being deployed which attempt to alleviate this problem. For example, the ISS RealSecure network intrusion detection system has functionality that enables automatic signature updates as they are released from the vendor (unless the IDS vendor allows for in-house development of alert signatures by the IDS administrator, it is still up to it

to produce and disseminate those signatures). It is also difficult to filter out the considerable amount of false positives generated by an IDS. For this reason, many intrusion detection solutions are highly configurable to weed out legitimate activity that is causing the IDS to alert.

Intrusion prevention systems operate on a slightly different level, focusing more on the network applications. For example, OKENA's *Stormwatch* Intrusion Prevention Appliance monitors all system calls to file, network, COM, and registry sources. The appliance uses a specific set of rules for each application to determine what is and is not appropriate behavior for that application. When an operation is attempted, *Stormwatch* compares it against its own set of rules and subsequently makes a decision to allow or deny the traffic. The biggest difference between intrusion detection and intrusion prevention is the action being taken by each. An intrusion detection device is passive, simply alerting on certain traffic, whereas intrusion prevention involves acting on possibly malicious or suspicious activity.

The problem with OKENA's intrusion prevention system (and many others) is that it is designed essentially for the "textbook" attack. Its functionality is derived from the assumption that an attack will be preceded by suspicious activity, comprised of network and system reconnaissance. This could adversely affect network traffic, as what could be viewed by the appliance as reconnaissance could easily be valid traffic; traffic that, although unusual, is being generated by legitimate (or even essential) operations.

For example, the Internet Control Message Protocol (ICMP) is a tool used by countless applications for legitimate operations. It is also used in many ways by attackers with more malicious intent. As such, it is a protocol that should be blocked liberally but carefully - at a network perimeter to halt system reconnaissance on the internal network, yet to allow its useful implementations in troubleshooting and diagnostics. ICMP is known to cause many false alarms in intrusion detection systems, however due to the fact that it is used in many cases legitimately many of the alerts can be filtered out or ignored. It would stand to reason then, that there are many other protocols like ICMP that can be used for legitimate purposes or malicious ones. How wise is it to have an appliance making that determination when it affects an entire network, with possible widespread operations and monetary impact? It would seem much more reasonable to have someone who is familiar with the environment making that call, rather than have it blocked automatically. (Northcutt, 67)

Intrusion prevention systems also incorporate firewall functionality. In examining how a firewall is deployed, we can evaluate the intrusion detection system's relevance as a perimeter defense tool. The biggest way in which a firewall enforces an organization's security policy is by bringing the traffic going in and coming out of a network into compliance with that policy. It accomplishes this through a set of rules governed by the administrator; rules that dictate what kinds of traffic are allowed and which are denied. In this way, it may be easier to get an overall picture of how an environment falls in line with an established security policy. The firewall administrator knows exactly what is being allowed and what is being denied; rules are not created "on the fly." That being said, an

IPS requires much less attention in the face of malicious traffic being allowed into a network; obviously, this creates a trade off between protection and awareness. (Brenton)

Also, firewalls are not always used to prevent outsiders from coming in through the “front door.” They are used to protect networks from internal risks such as modem pools. Does an IPS present a better solution in this situation? Almost certainly, it does not. And, while IPS might not be marketed or sold as an internal protection device in this situation, the fact is that it does not yet provide the dynamic functionality of a well-placed firewall. An IPS cannot compete with the level of adaptability, or the preciseness in configuration, of a firewall in the hands of a knowledgeable administrator.

Firewall and IDS, like anything else in network security, are far from perfect. With regard to firewalls, they are only as effective as the policies that govern them. Intrusion detection (or any anomaly detection device) generally has about a 75% success rate. (Ranum, 35) Therein lies the problem with moving to intrusion prevention. While it may be a viable and useful tool in the future, proposals to migrate from intrusion detection to prevention would be much too hasty for many reasons.

First, intrusion detection is growing so quickly that each different product is moving in a slightly different direction. While this may be good from a development standpoint, it can be a source of major headaches to a security administrator attempting to use products from different vendors. Every major intrusion detection system has different methods of alerting on suspicious network activity. There needs to be some standardization among different IDS implementations with regard to how they monitor and alert, and specifically the alert signatures from which they do so. It is next to impossible to get a concrete benchmark for IDS capabilities, and very difficult to correlate/aggregate alerts from one vendor to the next considering they all alert differently.

Similarly, host-based intrusion detection is a very young technology. Many intrusion prevention systems boast more comprehensive event correlation functionality than host IDS; however, has the HIDS solution really been given its due? Whether or not the intrusion prevention technology can be a more effective and beneficial security solution than HIDS or a slyly-configured honeypot remains to be seen. It is a fact that both of these still-developing tools can be an extremely valuable source of information, effectively gathering detailed information about the anatomy of system-targeted attacks.

Second, intrusion detection still has considerable room to grow. Most products are moving from a solely signature-based alerting system to a hybrid anomaly/signature configuration. It is no longer enough to look for known exploits through character and string matching, and the evolution of IDS technology is reflecting that trend. There is still much to be done in moving toward a comprehensive logging and alerting tool for suspicious and possibly malicious activity.

Intrusion detection systems are and always will be an invaluable source of information. If an attack or network reconnaissance never reaches perimeter devices, how can it be used to an administrator's advantage. The point of having

comprehensive logs of malicious activity is so that it can be used as a learning tool. True, successful compromises cost millions of dollars each year, and incident response is infinitely more costly than prevention, however if an attack is never seen then how do we know what to look for?

Also, it is easy for complacency to cause more problems on a network than a determined hacker. Placing a device inline in front of a network or external firewall that has the capability of blocking and allowing traffic is always dangerous. A poorly-configured firewall can stop legitimate services from entering or leaving a network, causing more harm than an attack. The only saving grace here is that the firewall is not making its own decisions; it is working from a policy of explicit rules set forth by someone who is familiar with the environment. This is not always true in the case of intrusion prevention. This appliance can only make a determination based on an internal baseline of "normal" activity. If it was to block legitimate traffic that happened to clash with this baseline, it could potentially do more harm than good – especially if an attacker were to figure out that the IPS is in-line and use it to their advantage.

Any security professional has seen what can happen when an attacker gains control of a critical server or network resource. The attacker is then free to cause whatever destruction he/she wants with impunity. The problem with giving so much control to one network appliance is that it can be used against the organization it was installed to serve. No matter how secure an appliance is, there is always the chance that someone will gain access or, at the very least, cause a denial of service that will prevent valid systems from utilizing it. This can cause major problems for the rest of the environment, not the least of which being a denial of service for the entire network. The only answer is deployment of multiple (or at least redundant) appliances, again decreasing the cost-effectiveness of the solution as a whole.

As with non-redundant perimeter firewalls, it is also never a good idea to introduce what could become a single point of failure. Intrusion detection does not come with such risks, as it does not have the capability to deny traffic. And, although increased network traffic may cause an overloaded IDS to drop packets and become less effective, it cannot create a denial of service on the monitored network.

Such a situation can and has occurred with the advent of intrusion prevention. Last November, a major educational institution deployed TopLayer's Attack Mitigator (AM) to provide defense against DOS, DDOS, and URL/URI exploits. Soon after the first device was up and running, however, the complaints began rolling in from all over campus. Network performance was degraded, packets were being lost, and connections were being dropped. The AM's CPU utilization did not indicate a problem, which made it difficult to troubleshoot. In the end, further investigation revealed that the denial of service had been caused by the overloaded AM. Although the appliance was equipped to monitor Gigabit Ethernet segments, it was just not designed to handle the heavy saturation often present on a campus network. Deploying load balancers for a cluster of AMs rendered the solution cost-ineffective, and even so it was evident the appliance was unable to prevent many types of malicious attacks (such as fragmented

Nimda traffic). This does not mean that TopLayer's product is badly designed or even that it is ineffective; however in this situation the technology was unable to successfully scale to a large, heavy-traffic environment. Intrusion detection systems can have similar shortcomings in such situations, though none would likely cause the denial of service seen here. (Paquette)

Many organizations have an extremely tight budget for security. Considering this, it comes as no surprise that many top-level managers might opt for a more comprehensive (read:inclusive) security solution. It is also somewhat naïve to think that every network administrator (and his/her management) makes network security a priority. There is definitely something to be said for increased peace of mind through automation. The problem occurs when a false sense of security is gained through a solution that, in effect, is designed to prevent attacks.

Is there really any way to prevent an attack? According to Martin Roesch (President and CEO of Sourcefire), "The time delta between when a vulnerability is discovered and when an auto-attack tool is developed and put into wide distribution will get shorter and shorter...there's going to be less response time from the vendor and admin community to get patches out." Some might argue that this supports the deployment of more automated response solutions like intrusion prevention, however it stands to reason that this diminishing timeline with regard to new exploits would affect that technology as well. It would be presumptuous to think that simply because a new exploit would be "caught" by an intrusion prevention appliance, there would be no way for it to enter an environment. Only by performing event analysis and correlation can an organization really know whether or not an attacker is circumventing its security measures. Only through real-time analysis and response can that organization be certain it is mitigating future risks.

Perhaps this accounts for the recent rise in intrusion prevention technologies. This is further reinforced by the selling point that such technologies will prevent attacks as they occur, whereas intrusion detection devices and firewalls come with crucial response delays. The fact is that there is no substitute for multi-layered security architecture. Likewise, there is no substitute for being in control of your own network – having network and system administrators that monitor according to their own baselines and react to deviations and incidents as necessary. At the most basic level, network devices act only as they are configured and programmed to act; it is ridiculous to assume that any network security solution can analyze and respond to every exploit an attacker could have at his/her disposal. Security is only as good as the person in charge of the network. That is to say, security devices are tools that are used by a person to combat other people trying to attack an environment (be it through scripts, worms, reconnaissance, etc.). These tools should not be relied upon to counteract every malicious operation launched against a system or network.

The bottom line is that an intrusion prevention solution, while beneficial in some situations, is not yet a more comprehensive solution to overall network security. As advanced as the technology has become in such a short time, it is still far from replacing a human being who is actively monitoring network activity and filtering, alerting, or ignoring network traffic as necessary. Anyone who says

that intrusion prevention is the next big thing- that it will replace any aspect of security in-depth (including IDS) - is missing the point. Security is not a solution or an end-state, but a process. While an intrusion prevention appliance may help once it becomes more versatile, effective, and reliable, it cannot, by itself, secure an environment. Not every organization can afford a comprehensive managed security service, however an increased level of automation is not the solution (as of yet). Even as intrusion prevention becomes more prevalent and easily-integrated into any environment, it should be used in conjunction with (not in place of) other diagnostic systems and perimeter defenses. (Fratto)

Hopefully, in the future, IPS will present itself as a valid and invaluable part of in-depth network security. It certainly has its place among other security devices, however still needs time for further development before becoming a major player in relation to firewall and intrusion detection technology. It would be hard to argue against a comprehensive network intrusion prevention tool, however the most prevalent and advances IPS still has far to go before the benefits outweigh the potential hazards to an environment and its users. Intrusion detection systems and firewalls have both proven themselves as being an integral part of network security; a gauntlet from which IPS has yet to emerge victorious.

© SANS Institute 2003, Author retains full rights.

## Sources

Benton, Chris et. al. Active Defense: A Comprehensive Guide to Network Security. San Francisco: Sybex, 2001. 144-172.

Cummings, Joanne. "From Intrusion Detection to Intrusion Prevention." Network World. URL: <http://www.nwfusion.com/buzz/2002/intruder.html> (Sept. 23, 2002).

Fratto, Mike. "Keep Out." Network Computing. URL: [http://www.nwc.com/1322/1322f1.html?ls=TW\\_103002\\_fea](http://www.nwc.com/1322/1322f1.html?ls=TW_103002_fea) (Oct. 21, 2002)

Hulme, George. "Attacks Averted." InformationWeek. URL: <http://www.informationweek.com/story/IWK20030202S0002> (Feb. 3, 2003)

Mun, Jeffery et. al. "Intrusion Prevention vs Intrusion Detection." ISS Forum. URL: <http://cert.uni-stuttgart.de/archive/issforum/2002/12/msg00014.html> (Dec 3, 2002)

Network Flight Recorder Corporate Web Site  
URL: <http://www.nfr.com/>

Northcutt, Stephen and Novak, Judy. Network Intrusion Detection: An Analyst's Handbook. Indianapolis: New Riders, 2001.

OKENA Stormwatch Corporate Web Site  
URL: <http://www.okena.com/>

Paquette, Mike. "Intrusion Prevention Bolsters Network Security." URL: <http://www.toplayer.com/pdf/faq.pdf>

Ranum, Marcus et. Al. "IDS at the Crossroads." Information Security Magazine June 2002 (2002): 32-41.

© SANS Institute 2003, Author retains full rights.