



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 DETECTION IS DEAD.
LONG LIVE INTRUSION PREVENTION!

SANS GIAC Certification Practical

Version: GSEC 1.4b

Option 1 – Research on Topics in Information Security

Submitted by Timothy D. Wickham

April 21, 2003

© SANS Institute 2003, Author retains full rights.

ABSTRACT/SUMMARY

This practical will demonstrate the limitations and drawbacks of intrusion detection as well as the reasons why intrusion prevention is a vastly better method of securing a network. In summary, IDS (Intrusion Detection Systems) will soon be rendered obsolete by IPS (Intrusion Prevention Systems).

© SANS Institute 2003, Author retains full rights.

TABLE OF CONTENTS

I.	Introduction.....	1
II.	The Rise and Fall of Intrusion Detection.....	3
III.	The Rise of Intrusion Prevention.....	6
IV.	Conclusion.....	9

GLOSSARY

Blended threat. A worm or other exploit that can do damage in multiple ways (e.g., the Code Red worm, which defaces web pages on the exploited server and performs a distributed denial of service attack).

Distributed denial of service. A coordinated attack that utilizes multiple compromised systems to bombard a target system with various types of network traffic in an attempt to overwhelm its ability to handle that traffic.

False-Negative. Malicious network traffic (in the case of network-based systems) or system activity (in the case of host-based systems) that has been falsely labeled as legitimate. Also known as Type II error in the statistical and judicial fields.

False-Positive. Legitimate network traffic (in the case of network-based systems) or system activity (in the case of host-based systems) that has been falsely labeled as malicious. Also known as Type I error in the statistical and judicial fields.

IDS. Intrusion Detection System(s). System(s) that passively monitor a network or host for attacks launched against it.

Host-based IDS. IDS designed to detect attacks at a host level.

Network-based IDS. IDS designed to detect attacks at a network level.

IPS. Intrusion Prevention System(s). System(s) that actively monitor a network or host for attacks and block those attacks from occurring.

Host-based IPS. An IPS designed to prevent attacks at a host level.

Network-based IPS. An IPS designed to prevent attacks at a network level.

Multi-vector worm. A worm that uses multiple methods to propagate (e.g., Nimda, which spreads via open network shares, email attachments and other methods), thus making mitigation difficult.

SPAN. Switched Port Analyzer. A port on a switch configured to transmit a duplicate of the data from multiple ports. Otherwise known as a mirror port.

Chapter 1

INTRODUCTION

We've all heard the stories or experienced them ourselves. Long nights spent agonizing over logs, trying to pinpoint an exploit amid a sea of data. Weeks of fine-tuning and customization in a futile attempt to eliminate some of the noise. Untold hours analyzing exploited systems. Establishing the value of the compromised or altered data. Restoring altered files or completely rebuilding systems. Rehearsing how you're going to explain to management that the attack wasn't prevented. And, finally, trying to determine the best way to keep it from happening again. This is the lot of an IDS administrator.

Fortunately, the solution has arrived. Intrusion prevention promises to render intrusion detection irrelevant, allowing a fundamental change in the way that networks are secured. By blocking the attack rather than just detecting it, intrusion prevention allows an organization to shift from a reactive to a proactive security stance.

Relying on intrusion detection forces a security team into a reactive state. You are forced to simply wait for an exploit to occur before taking action to try to minimize the damage. Intrusion prevention, on the other hand, allows a security team to be proactive. By preventing the exploit from occurring in the first place, it frees you up to concentrate on the more productive aspects of the security field, such as security policies, business continuity planning, etc.

"Network 'signature-based' intrusion detection is a little like posting a guard outside the bank, and giving them pictures of all the known crooks

in the world. He scans the faces of the people walking past, and if he sees a known crook, he signals an alarm.

Host-based intrusion detection is like someone watching the gold bars in the vault to make sure they're still there."

- Leigh Purdie – Intersect Alliance Director and Principal Security Consultant

This quote describes an accurate analogy for host and network-based IDS. The problem here is, the IDS guards (both host-based and network-based) in this analogy are unarmed and can't do anything but write up a report when a burglary occurs. Taking this analogy a step further, implementing an IPS is like arming those guards, allowing them to actually thwart an attempted burglary.

Sounds great in theory but there are reasons why intrusion detection has become such a popular component of an information security arsenal. There are also legitimate concerns about the viability of IPS. This paper will explain why IDS became popular, the drawbacks of IDS, and how the concerns about IPS are being overcome, paving the way for a changing of the guard in the security arena.

Chapter 2

THE RISE AND FALL OF INTRUSION DETECTION

The birth of IDS and its popularity are a direct result of concern over the security of data assets in an increasingly hostile world. Since the advent of the computer, people have been trying to find ways to exploit hardware and software bugs and misconfigurations for pride and profit as well as ideological, political, and theological reasons. Things have gotten considerably worse from the time the first boot-sector virus appeared (1981)¹ to the present threat environment of multi-vector worms, blended threats, flash worms², and distributed denial of service attacks. The quantity, variety, and potential disruptiveness of known attack techniques has been on the rise, particularly in recent years. The frequency with which they are being put to use has also increased dramatically. According to Carnegie Mellon's CERT Coordination Center, the number of reported vulnerabilities per year increased by almost 1500% from 1998 to 2002. Similarly, the number of reported computer security incidents per year increased by almost 2100% from 1998 to 2002³. And "reported" is the key word here, as there are various reasons why organizations are reluctant to report security incidents⁴. The actual numbers are probably even more dramatic.

¹ Slade, Robert M. "History of Computer Viruses." 1992. URL: http://www.claws-and-paws.com/virus/papers/slade_history.shtml

² Staniford, Grim, Jonkman. Silicon Defense. "Flash Worms: Thirty Seconds to Infect the Internet." 16 Aug 2001. URL: <http://www.silicondefense.com/flash/>

³ Carnegie Mellon Software Engineering Institute CERT Coordination Center. "CERT/CC Statistics 1988-2003." 16 Apr 2003. URL: <http://www.cert.org/stats/>

⁴ Sieberg, Daniel. Cable News Network. "FBI: Cybercrime Rising." 8 Apr 2002. URL: <http://www.cnn.com/2002/TECH/internet/04/07/cybercrime.survey/>

J.P. Anderson's 1980 paper [Computer Security Threat Monitoring and Surveillance](#) introduced the basic concept of IDS. Although it focused specifically on host-based statistical anomaly detection, the race was on to expand this thinking in a number of directions.

Paul Innella's 2001 paper [The Evolution of Intrusion Detection Systems](#) gives a thorough account of the progression of IDS from host-based to network-based. The current state of intrusion detection also includes a variety of techniques for detecting malicious traffic, including stateful pattern matching, protocol anomaly detection, and statistical anomaly detection.

Stateful pattern matching looks for specific signatures within a packet or across a stream of packets or IP fragments. For example, the string "vrfy root" can be used to gather reconnaissance information from an SMTP server. In order to properly detect this exploit, you must first reassemble any IP fragments. Then you must determine the state of the session, since this exploit is only effective in the control portion of an SMTP session (rather than in the text of an email, for example).

Protocol anomaly detection looks for traffic that bends the rules or guidelines of certain protocols. An example of this would be looking for IP packets longer than 65535 bytes, which generally indicates a "Ping of Death"⁵. Obviously you must be very careful here, since various custom and legacy applications as well as certain network equipment are known to bend these rules and most protocols have at least some "grey area" where interpretation is required.

Statistical anomaly detection looks for deviations from the normal traffic patterns on a network. Care must be taken here, as well, due to the inherently dynamic

⁵ Kenney, Malachai. Insecure.org. "Ping of Death." 21 Oct 1996. URL: <http://www.insecure.org/sploits/ping-o-death.html>

nature of computer networks. Detecting a large number of connections from a computer outside your network to a specific destination port on various systems inside your network would be indicative of a host sweep, for example.

There are a number of reasons why researchers and vendors initially decided to detect these types of traffic rather than prevent it. The main reason for this was hardware and software limitations. These limitations led to accuracy and performance problems.

The lack of specialized hardware made it difficult to deeply process packets quickly enough to accurately determine whether it was malicious or not. This can be attributed to the somewhat generic or imprecise signatures that IDS is restricted to as a result of these limitations. The bane of IDS has been the inability to weed out false positives and false negatives. This inaccuracy makes it extremely inefficient as a security solution. False positives lead to an enormous waste of time and skilled resources as the administrator searches through log files for legitimate exploits. False negatives expose the organization to undetected theft of or damage to internal assets.

Additionally, without specialized hardware, putting a device in-line would have introduced intolerable latency and throughput problems. To give you a feel for the processing power required, consider that to achieve zero latency on a link with 1 Gbps throughput with an average packet size of 512 bytes, a system has only 512 nanoseconds of processing time per packet.

As a result of these accuracy and performance problems, it was decided that the best place for IDS was off to the side (hanging off of a hub, SPAN, or network tap), rather than as an active part of the network. This put it in a position where it could passively monitor traffic but it wouldn't do any damage. However, it also put it in a

position where it couldn't *prevent* any damage from attacks. According to Gartner Research Director Richard Stiennon, "Legacy IDS technology was built on the belief that the number of security vulnerabilities and clever hackers targeting them is too daunting a task to prevent; thus enterprises have been relegated to monitoring activity, rather than attempting to block attacks."

In a June 2002 head-to-head evaluation of various IDS products, Network World magazine was so thoroughly unimpressed with the offerings that they decided not to declare a winner⁶. The result of these inefficiencies and inaccuracies was that many in the industry disavowed IDS. The fall of IDS was afoot.

⁶ Newman, Snyder, Thayer. Network World. "Crying wolf: False alarms hide attacks." 24 Jun 2002. URL: <http://www.nwfusion.com/techinsider/2002/0624security1.html>

Chapter 3

THE RISE OF INTRUSION PREVENTION

With specialized hardware (ASICs, FPGAs, network processors, etc.) now available, intrusion prevention is a reality. IPS can be implemented as part of the network fabric rather than passively off to the side. As each packet comes into the system, it is deeply analyzed and a “go/no-go” decision is made as to whether it should be allowed to continue on to its destination. If the packet is malicious, it is dropped and is never even seen by the victim.

Keep in mind that this has to be an in-line decision. According to Gartner Group Vice-President John Pescatore, there are a number of “snake oil” IPS vendors that have adopted the IPS mantle but don’t address the problems of IDS⁷. Some vendors claim to do intrusion prevention via TCP resets and firewall shunning. The problems with these methods are numerous. TCP resets are signals sent to the attacker and victim to tear down the connection before the exploit can occur. The problem is that, in a passive configuration, the IDS/IPS sees the attack at the same time that the victim does, so the damage is often already done by the time the reset is sent. Also, TCP resets are obviously ineffective against non-TCP exploits. The recent Slammer worm was a perfect example of a UDP-based attack which was completely impervious to TCP resets. Firewall shunning, on the other hand, is a signal sent from the IDS/IPS to a firewall (or router) to block traffic from the IP address that the attack appears to be originating from. This method suffers from the same timing problem as TCP resets. Additionally, by spoofing the attack to look like it is coming from a legitimate user, the attacker can force you to create

⁷ Fonseca, Brian. InfoWorld. “The IPS Question.” 4 Apr 2003. URL: http://www.infoworld.com/article/03/04/04/14ips_1.html?security

a firewall rule to block their access. Basically, the attacker has just forced you to do a denial of service on your own users.

When implemented properly, the high-performance nature of a hardware-based IPS solution allows you to have extremely precise and processor-intensive definitions of what is and what isn't malicious traffic. This helps to overcome the accuracy concerns of IDS. Better knowledge of the network that is being protected can help here as well. For example, TippingPoint Technologies' UnityOne line of products can do an inventory of the platforms and services that it is protecting to help weed out false positives. If it detects an Apache webserver exploit (e.g., the Slapper worm) being attempted on a Microsoft IIS webserver, it will block the traffic (since there is no reason for this traffic to be on your network, no matter how ill-targeted) but it will downgrade the alert to avoid waking up an administrator in the middle of the night.

Despite this improved accuracy, it should still be easy to tune the IPS to avoid blocking legitimate traffic in the event that a false positive should occur. For example, your security administrator might regularly run vulnerability scans from his/her workstation to determine points of weakness on the network. You should be able to easily configure the system to allow this traffic from this particular system to pass freely rather than being blocked.

Regarding performance, the specialized network processors can process traffic at Gbps rates with switch-like performance. This performance should be achievable under any condition. The protocols present, the mix of those protocols and the size of the packets on the network shouldn't matter. In most products, you can enable or disable protection against each type of attack. The fewer attacks you are trying to protect against, the lower the processing required, making it easier for the IPS to keep up. However, the system should be able to maintain performance

even if everything is turned on. Again, taking TippingPoint as an example, a recent Tolly Group evaluation showed that TippingPoint achieved 100% of their claimed throughput under a variety of traffic conditions with all protection enabled while introducing a maximum latency of 215 microseconds⁸.

And if this must be an in-line device, the first question any conscientious network administrator is going to ask is “What if this thing crashes? The last thing I need is another single point of failure to cost me my uptime bonus.” The high availability concern can be addressed in a variety of ways. First, you can rely on tried and true routing protocols like Cisco’s proprietary HSRP or the more vendor-neutral protocols like VRRP, IGP and others to route around a failed link. Second, you can have a stateful failover setup by using a dedicated Ethernet connection between redundant systems that allows them to keep all session and configuration information synchronized. If the IPS on the primary link fails, traffic is re-routed down the secondary link and through the back-up IPS⁹. Finally, you can have the IPS “fail-open” (allow all traffic to pass unchecked) in the event of a software or hardware problem or if the system becomes oversubscribed.

But if the performance, accuracy, and availability concerns are properly addressed, the potential for increased network security and the efficiency with which this is achieved are enormous. Anyone who is familiar with patching on almost any size network will tell you that it is a monumental task. On a network with a variety of platforms and applications, lacking massive resources dedicated to the task, it can be impossible. First, you must wait for the hardware or software vendor to which

⁸ The Tolly Group. “TippingPoint UnityOne Intrusion Prevention Appliances Performance Evaluation.” 1 Feb 2003. URL: <http://www.tolly.com/DocDetail.aspx?DocNumber=203101>

⁹ Cisco Systems. “IPSec VPN High Availability Enhancements.” 6 Feb 2003. URL: <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122limit/122y/122ye/1229ye/12yipse.htm>

the patch applies to make it available for download. This generally occurs within a few days of the vulnerability being exposed but can sometimes take much longer. Second, the patch must be tested on all of the possible combinations of hardware and software. Third, the patch must be implemented, with proper documentation, on each system. The expense in terms of man-hours and the hardware and software licenses required to set up the test environment, perform the proper testing, and perform the rollout is enough to make many organizations give up. Testament to this is the fact that systems are still being exploited by Nimda and Code Red even though the vulnerabilities that these worms exploit are over a year old. In the case of the SQL Slammer/Sapphire worm, Microsoft (and just about everyone else) was hit by a worm that exploited a 6-month old vulnerability in their own software for which they had readily made a patch available. The problem was that they failed to properly implement it on their own servers. According to Bruce Schneier, CTO for Counterpane Internet Security, "This shows that the notion of patching doesn't work."¹⁰ And the patching problem is not isolated to Microsoft products. They are actually making great strides towards making their software more secure but the whole concept of patching is fundamentally flawed.

When implemented properly, an IPS can provide protection against a newly discovered vulnerability within hours of its discovery. This prevents such attacks from occurring in the first place, giving an organization time to plan and test properly for patch rollout rather than calling in their administrators in the middle of the night or on a weekend to haphazardly apply an untested patch. Recent experience has shown that this is a dangerous strategy¹¹.

¹⁰ Lemos, Robert. CNETNews. "Microsoft fails Slammer's security test." 27 Jan 2003. URL: <http://news.com.com/2100-1001-982305.html>

¹¹ Middleton, James. Vnunet.com. "Experts warn not to apply Microsoft patch" 17 Apr 2003. URL: <http://www.vnunet.com/News/1140296>

Chapter 4

CONCLUSION

Imagine your weekly security war-room meetings. With IDS, they are highly stressful. Everyone spends a great deal of time preparing, sorting through logs, and preparing answers or excuses for why an exploit occurred or a system wasn't patched properly. With IPS, they are more of a show-and-tell session to calmly discuss the problems that were avoided and how to intelligently apply the necessary patches in an organized and cost-effective manner.

The shortcomings of IDS are apparent and IPS is clearly the superior security solution. According to Pescatore, "We think IDS is dead. It's failed to provide enterprise value."

However, when looking at potential IPS vendors, skepticism should be the rule. Will this solution accurately block attacks and nothing else? Will it do this without hampering network throughput or introducing latency? Is it going to perform dependably? What happens if it doesn't? Will it be kept current with the ever-changing threat environment? Any solution that effectively answers these questions will be leading the charge.

Tom Danford, CIO of the University of Dayton, Ohio, says his organization realized that IPS was the only way to go. According to Danford, "We were hit by all those [viruses], and it brought the university to its knees on a couple of occasions... We had classes that were affected and [a large] expense in paying people to clean up the machines and damage. There's also all that lost time and productivity. We

decided that prevention was going to keep our security where we wanted it to be."¹² The changing of the guard is clearly underway.

© SANS Institute 2003, Author retains full rights.

¹² Fonseca, Brian. InfoWorld. "The IPS Question." 4 Apr 2003. URL:
http://www.infoworld.com/article/03/04/04/14ips_1.html?security

LIST OF REFERENCES

- Slade, Robert M. "History of Computer Viruses." 1992. URL: http://www.claws-and-paws.com/virus/papers/slade_history.shtml
- Staniford, Grim, Jonkman. Silicon Defense. "Flash Worms: Thirty Seconds to Infect the Internet." 16 Aug 2001. URL: <http://www.silicondefense.com/flash/>
- Carnegie Mellon Software Engineering Institute CERT Coordination Center. "CERT/CC Statistics 1988-2003." 16 Apr 2003. URL: <http://www.cert.org/stats/>
- Anderson, J. P. "Computer Security Threat Monitoring and Surveillance." 15 Apr 1980. URL: <http://seclab.cs.ucdavis.edu/projects/history/papers/ande80.pdf>
- Innella, Paul. "The Evolution of Intrusion Detection Systems." 16 Nov 2001. URL: <http://www.securityfocus.com/infocus/1514>
- Sieberg, Daniel. Cable News Network. "FBI: Cybercrime Rising." 8 Apr 2002. URL: <http://www.cnn.com/2002/TECH/internet/04/07/cybercrime.survey/>
- Kenney, Malachai. Insecure.org. "Ping of Death." 21 Oct 1996. URL: <http://www.insecure.org/splotts/ping-of-death.html>
- Newman, Snyder, Thayer. Network World. "Crying wolf: False alarms hide attacks." 24 Jun 2002. URL: <http://www.nwfusion.com/techinsider/2002/0624security1.html>
- Fonseca, Brian. InfoWorld. "The IPS Question." 4 Apr 2003. URL: http://www.infoworld.com/article/03/04/04/14ips_1.html?security
- The Tolly Group. "TippingPoint UnityOne Intrusion Prevention Appliances Performance Evaluation." 1 Feb 2003. URL: <http://www.tolly.com/DocDetail.aspx?DocNumber=203101>
- Cisco Systems. "IPSec VPN High Availability Enhancements." 6 Feb 2003. URL: [Cisco Systems. "IPSec VPN High Availability Enhancements." 6 Feb 2003. URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122limit/122y/122ye/1229ye/12yipsec.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122limit/122y/122ye/1229ye/12yipsec.htm)
- Lemos, Robert. CNETNews. "Microsoft fails Slammer's security test." 27 Jan 2003. URL: <http://news.com.com/2100-1001-982305.html>
- Middleton, James. Vnunet.com. "Experts warn not to apply Microsoft patch" 17 Apr 2003. URL: <http://www.vnunet.com/News/1140296>

© SANS Institute 2003, Author retains full rights.