GIAC
CERTIFICATIONS

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

IDS, the Silver Bullet!?
A conversation with your CEO

SANS GIAC Security Essentials Certification
Practical
V.1.4b

Bobby J. Brown Jr.
CCNA, CCDA
February 22, 2004

Summary

The subject of this research paper has recently been a topic of discussion in the realm of network security the world over. With network security starting to come to the forethought of CEO's and CFO's, I wanted to put a paper together that would give them and idea of which security hardware they should include in their security infrastructure.

IDS's, IPS's and firewalls have been the latest "buzz" words thrown around in CIO's magazines. As we all know in these times of tight budgets, CFO's want to save money. We have all had these types of conversations with upper management. Have any of us heard this questions before, "Can I just buy this IDS thing and not have to worry about that virus and worm stuff anymore?" "I thought I bought a firewall last year?" This paper will give an example of a mock conversation with the big bosses and go into a detailed explanation of why an IDS is or is not the one thing that will save your network from intrusion. I will also compare and contrast IPSes. I will discuss what is needed to manage and maintain an IDS system. Hopefully when you finish reading this paper you too will be able to answer the question, "Is the IDS system the silver bullet we need to protect our network?"

IDS, what is it?

IDS stands for Intrusion Detection Systems. "What kind of intrusions can it protect us from?" Well, even the word intrusion can be substituted with the word misuse. "Who would try to break into my "little" network? People don't even know who we are." Hackers and crackers. Hackers could be a people or groups of people that want to "break in" to your network for any number of reasons. One hacker might want some type of financial gain. It could be a competitor of yours that wants to see the latest research that you are working on. Or it could be a cracker, someone who just wants prove to themselves that they could do it. It might just be someone who wants to be a little mischievous. A cracker could be someone that works for our institution. No matter the reason, these people are trying and will continue to try to get into your network. "Well what does this IDS thing do, how does it work?" The IDS is a piece of software or hardware that tries to detect any attempt to penetrate the virtual security perimeter of the network or any workstation (user's PC) on that network. IDSes are designed to distinguish between authorized entries and malicious intrusions.

"How is it able to do that?" It does this by looking through all network traffic (i.e.packets) and then determines if there are any signs of attacks or abnormalities. Let's discuss the different types of IDSes and where one would place one in the network. [1,8,10]

II

"What are the different types?"

There are two types of IDSes:

- NIDS
- HIDS

A Network Intrusion Detection System is just that. We would place this piece of hardware or software on the perimeter of our network in the DMZ at minimum, to monitor intrusions. We could also place them throughout the network in the switched Local Area Network backbones which are placed in various data closets across campus. "What is the DMZ? I know we are not talking about the one in Korea." The DMZ I'm referring to has the same principle as in Korea, the demilitarized zone, but I'm talking about the virtual portion of the network that is not trusted. It is between our private network and the public internet. I'm going to go into more detail a little later in our discussion. Let me tell you what HIDS stands for. Host Intrusion Detection System, this piece of software is actually placed on a particular user's workstation or server. "Which one is better?"

Let's go into the pros and cons of each. With an NIDS configuration, we would be able to monitor large amounts of network traffic and have the ability to monitor the broadest range of attacks, which may include the Denial of Service (DoS or DDos for Distributed) and the "Ping of Death" attacks with a single piece of equipment. They can even capture all traffic going to a targeted system. There is a downside, a higher amount of bandwidth can cause the IDS to miss or drop some packets. This is a processor intensive process which can overwhelm some IDSes. Current bandwidth limits of commercial NIDS are 80 MB/s to 120 MB/s. However the dropped packets will not result in blocked transmissions because NIDS are "passive" devices. This is different than a router or firewall. If these pieces of equipment drop some packets, connections could go "down".

Also the NIDS are unable to look into encrypted traffic (i.e. SSH or VPN) therefore, if an attack, virus or worm, has been sent during one of these sessions, the NIDS will not be able to see it. If there are any back-doors set-up on a network for example modems and or remote management software, for good or evil purposes, a NIDS would not be able to stop or catch this type of "insider" attacks. It can be very hard to do a thorough job of detecting and protection with a NIDS alone.

HIDS on the other hand, as discussed earlier, are placed on the workstation or server, so it is able to monitor all traffic sent to that host, and it is able to identify unauthorized attempts to access the host. It does this by analyzing specific files, logs and registry settings of the workstation. HIDS will also monitor "legit" use, which can help

III

determine if the user is going to websites that may not be appropriate or related to work. HIDS can also see the end results of an encrypted session that is sent to its host. One of the biggest downsides to HIDS is what the name implies. We would have to deploy the HIDS to every host in our organization, if we were to implement this technology as a single line of defense. "Sounds costly and it still doesn't answer my question, which one is better?" **(7,8,10,12)**

Which one is better?

Well, there are still more things to consider. The NIDS and HIDS can be broken into two other categories:

- Signature based
- Abnormality based.

Their titles explain how they work. The signature based IDS's use a pattern-matching algorithm which compares the traffic analyzed with known sets of attacks. If one of these patterns were matched, it would indicate that an attack has occurred or is occurring. That is one of the major problems with these types of systems, signature based IDS's have a lot of false positives. When it sees a signature it recognizes it fires an alarm. No matter where the traffic is coming from or going to. One has to monitor and "tune" these alarms. I'll talk about this a little more in depth shortly.

Anomaly-based IDSes perform a baseline and then "fires" an alarm when it sees abnormal exceptions to the baseline. The baseline information includes important statistics such as CPU utilization, disk activity, user logins, file activity, and so forth. It perceives these anomalies as possible attacks. "Ok??."

There are more pros and cons. NIDS are much easier to deploy and maintain. They can provide far greater detail into the nature of network traffic whereas a HIDS does not. NIDS are a more mature technology while HIDS are a newer technology. Some of the commercial NIDS on the market have the ability to not only detect an attack but can also stop it in real-time. It can then "block" the attackers' IP address for a specified period of time. HIDS however, do have the ability to replace a file or return a workstation/server back to their original configuration after the attack has occurred.

Alright sir, maybe I should talk about something that is close to your heart; **cost**. HIDS can range from $50 to $1000 per host whereas NIDS can range from $10,000 to $30,000. "Wow." That's not all sir, an IDS is like any other new server "out of the box". Let's use a Web server for an example. Initially when we bought our Web server it was just a server. We had to populate it with content specific to our needs. Keeping the content accurate and up to date requires time and talent. The IDS is the same way. I have

IV

read that it can take up to six months to "tune" a system and its going to take anywhere from one-half of an FTE (Full-time employee) to a whole FTE to continue to manage the system or systems. "So it's actually going to cost us more? I don't know about this." Yes and no sir. I have also read studies that state the Return on Investment (ROI) can be up to 145%, with an operational cost savings of 41% and business benefits of 33%. These dramatic turn-arounds could occur in 14 months. Those kinds of numbers could at least deserve a second look right sir? "Hmm, possibly." "How did they come up with those numbers?" This particular study used the Total Economic Impact (TEI) methodology to examine the possible returns. **(6, 8, 9)**

Since there are some legislation regulatory requirements that are coming down the pike, such as the Gramm-Leach-Bliley Act, which affects the financial services industry, the Health Insurance Portability and Accountability Act (HIPAA), which affects the healthcare industry, and the Sarbanes-Oxley Act, which will require CEO's to sign off on financial statements. "It might behoove us to get one of these products." Some states like California are passing legislation that requires a company to report all security violations. If a company doesn't report it, they would set themselves up for a class action suit. Congress is using the California legislation as a national model. "Ok, you got me thinking seriously about acquiring one of these things." **(9)**

As far as to which one is better, both the NIDS and HIDS have their advantages and disadvantages as you can see. It really depends on what we feel the primary threats to the organization are:

- Do you want to know who is attacking our organization?
- How are we being attacked?
- Do we have a threat on the inside?
- Do we want to look at specific threats on a particular VLAN (i.e. Finance)?
- Do we want to have the ability to collect evidence for eventual prosecution? **(7)**

HIDS may be more fitting if the organization is more concerned with internal "crackers". NIDS might be more fitting if the organization is more concerned with external "hackers". I strongly suggest that we try and deploy a combination of host and network intrusion detection systems because their strengths are complementary. Possibly deploying a NIDS at our perimeter and HIDS on important, maybe all, servers, and special workstations throughout our facility. This would go with what the network security community calls a "layered defense or defense in depth." I will also talk more about this a little later.

Before for we discuss which one we want to go with, I wanted talk to you about another technology you might have been hearing about, Intrusion Prevention Systems (IPS). "IPS, IDS what is the difference?" **(6,7,8,9)**

v

IDS vs. IPS

"You're right, I did get a chance to read an article. It said something like IDS is dead. What's that all about?" Well sir that is a good point. I thought you might have read that article. You have just touched on a topic that is still up for great debate throughout the industry. But what I think is most important thing to remember is, in order to prevent an attack, one must be able to detect it first. Before I talk about that I wanted to take a step back and discuss the differences between the two. An IPS is an inline tool that monitors network traffic, and much like an anomaly HIDS, it determines the "normal" traffic patterns of applications and operating systems. Once the learning period is over, it then incorporates the behavior patterns into a policy. The policy then determines what traffic will get through and which will not. This is a very important point. While IDSes, for the most part, passively monitor traffic "out of band", the IPS will take action on the traffic and stop it. "Well that's a good thing, right?"

As I stated earlier, the main complaint about IDS systems in the industry is all of the false positives that are being reported by these types of hardware. This is just a part of life at this time when dealing with detection and prevention systems. Since false positives are inherent in the architecture of this equipment it also could be bad "thing" because IPSes can stop legitimate traffic which would in turn effectively shut down the network. The false positives are created secondary to the whole underlying foundation of pattern matching. If the hardware sees a pattern, of course it has no idea if that traffic is for "good" or for "evil"; it just knows that the pattern matches an intrusion signature that resides in its database. The IDS or IPS then responds by either sending an alarm or, in the IPSes case, denying the suspect traffic. One of the key benefits of being "out of band", like an IDS, is that one has the ability to flag traffic that looks even the slightest bit suspicious. On the other hand IPSes don't have that ability. If it is set too sensitively it will have negative effects on the network. There is no room for any wrong decisions, especially when it will "shut down" the network. Hence, the IPS can become the single point of failure for your network. Unlike IPSes, an IDS is more discriminating, minimizing the risk of blocking authorized users who might have made a small mistake.

One more thing about this that I wanted to point out. Security administrators can use an IDS to look for patterns of malicious activity, for example, if there is someone on staff that is trying to break in and crack passwords on a variety of servers on the network via a stealth attack. Security administrators can build a "case" by having the ability to use audit logs and then performing forensics, while making correlations to specific activities. Once these activities have been identified, the administrator can then make the decision to continue to monitor the behavior or to "tune" a signature to give him or her more "clues", eventually finding a suspect and stopping the activities. Human operators should be responsible for determining what traffic to stop and what traffic they should let

traverse their network. Some IDSes can be configured to shun particular threats to the network. The IDS would tell a perimeter gateway or router to stop traffic, after being approved by an experienced operator during the tuning process. A large percentage of malicious activity cannot even be detected and prevented on the fly. If a new virus or worm comes out today, it will take at least a day before vendors can release a signature update. Therefore the attack will not be prevented. That is why it is imperative to have more than one of these types of security technologies on campus. Some pundits even say that it is unfair or not realistic to compare IPSes with an IDS. The reality is that an IPS is like an extension to a firewall and not an IDS. The last point I wanted to make about this is that all IPSes have and IDS at their core. [2,4,5,13]

Defense in Depth

Now this is the most important thing I want you to take away from our conversation today. No matter what particular IDS platform we decide to go with, it should be just one part of our overall security architecture. Network security is not just a destination but an ongoing process. It should also be set-up in a layered format. "You sound like we are talking about more money." Not necessarily sir. Let's go into more detail. When I speak of a layered defense I usually categorize it into five main areas:

- Perimeter
- Network
- Host
- Applications
- Data

Of course the perimeter acts as the first and last point of contact between our network and the "world". It includes the DMZ that we spoke of earlier. The DMZ usually has a Web server, a DNS server and what seems to be the most important server here on campus, the email server. "Was that supposed to be funny?" Sorry sir. Now as far as security hardware, the perimeter should include the firewall we bought last year, an IDS and some type of corporate anti-virus product. I've talked about it a briefly but I wanted to go into a little more detail about the firewall. A firewall performs three basic jobs: traffic control, IP address translation and it can also be a VPN endpoint. "Wait a minute, we talked about the access control and the Internet Protocol address but, what exactly is a VPN?" A Virtual Private Network. It's an encrypted tunnel-like connection between a network and some kind of remote device. Either a business partner or even a telecommuting employee can use this type of technology. "I'm glad you mentioned that because "tele-working" is high on my agenda for the next fiscal year."

The next area we should concentrate on is security for our private network itself. This includes our Local Area Network (LAN) and our multiple Wide Area Network (WAN) connections. Because we are one big happy family, it is pretty "open" inside. We share connections to the corporate offices that we consider safe. But why should we. Some researchers state that almost 90% of attacks occur from the inside. That's why I suggested that we place not only a few NIDS throughout the network but HIDS as well. We should also think about having some kind of vulnerability assessment tool. "Ok, I'll bite. What is that?" Well the name kind of says it all. It scans our network for flaws. For example if we have unknown services (ports) that are open to exploitation (i.e. ftp, epmap) we would be notified by this tool's audit.

The third level is the hosts itself. As we discussed earlier, host are the individual workstations and servers on the network. Security here should include HIDS, host vulnerability scanners and anti-virus software. These types of technologies allow system administrators to quickly identify which device settings require updating, "Patches or Hot fixes"; this would harden and fortify our defenses.

Security at the Application level is the fourth area. Poorly protected applications can easily provide an intruder with "top secret" information. The reality is most of the "off the shelf" applications are programmed with security as an afterthought, if at all. Some of our departments even place these applications on our website. We should think about including things like input validation tools and application shields. These tools can provide more accountability and log information on what or who is getting access to your confidential documents. The word used in information security today is non-repudiation. This means we would have the ability to provide proof to a third party that any data traveling across our network is assured to have been checked for its integrity and origin.

Last but not least, level five. The actual data. Data security includes a combination of policy and encryption. Encryption is one of the most important aspects of security because if all of the other measures we spoke of fails, at least the data will be protected from general access. The policies have an effect on which users are authorized to see specific data and what they can do with it. (i.e. read, write, delete etc.) The policies should clearly state who the owner of the data is and who is ultimately responsible for it.

Types of data encryption include: PKI (Public Key Infrastructure), which involves a sharing of Keys (passwords) with users throughout the network to encrypt mail or other data, PGP (Pretty Good Privacy) which enables you to share messages, secure files and disk volumes with strong authentication and SSL (Secure Sockets Layer) handshaking protocol which allows protocols like HTTP, FTP and Telnet to transmit across the network encrypted. "Ok, Ok, Ok, I see where you going with this." Sir I just wanted to point out a few of the other things we might want to take a look at to possibly implement. Not all of the technologies I just spoke of cost money. There are some shareware and freeware products out there. "Great!". They will of course cost in time to install and configure. [11,12]

VIII

Conclusion

Security is an ongoing dynamic process. IDS products deliver the ability to provide two of the most important aspects of network monitoring: visibility and control. Being able to "dig" into data packets and understand the nature of network traffic is necessary to make correct decisions. The information gathered from an IDS removes much of the guesswork for a security professional (the visibility), and it can give you most of data required to build a reality based policy. Once you can see the traffic, you can then prevent certain kinds by shunning unwelcome attacks (the control). You don't have to hypothesize about what should be in a security policy or guess at the possible misuse of the network. IDSes provide this information in real time and the information they collect is shared with routers and firewalls. **(11)**

Last but not least. IDSes are not just for security. "What do you mean?" We can also use it to provide information about the way employees use the network. Which makes it possible to mitigate another serious problem in business today, a decrease in productivity because of web surfing not related to work. It also can protect your human resources. "How?" With the visibility we talked about. The administrator has the ability to see outbound traffic that could indicate who is looking for a job, or watch inbound traffic to determine what recruiters are trying to locate candidates. The company suffers from the void of an employee leaving. The cost of that includes candidate searches, placing advertisements, and placement fees and in some cases relocation fees can cut into the overall budget of our organization. Not to mention the training of a new employee can cause a further decrease productivity of the department. **(7)**

In today's environment, in which hackers and cracker launch network attacks with increasing frequency and sophistication, there is not one silver bullet to solve all of our network security concerns. But I believe that an IDS is a "must have" along with selectively installing some of the other security measures I spoke of to help in the threat that gathers against us every day.

References:

1. Barker, Meg. "IDS choices you make the call" August 2003
    http://www.sheshunoff.com/email/archive/0803/oper_new1.html
    (February 22, 2004)

2. Golomb, Gary. "IDS vs. IDS commentary" June 16, 2003
    http://www.linuxsecurity.com/articles/forums_article-7476.html
    (February 22, 2004)

3. Venezia, Paul. "Intrusion Detection or Protection" Info World. December 12, 2003.
    http://www.infoworld.com/article/03/12/12/49FEgovforneeds_1.html
    (February 22, 2004)

4. Messmer, Ellen. "Don't Dis My IDS" Network World Fusion. June 16, 2003.
    http://napps.nwfusion.com/weblogs/security/002959.html (February 22, 2004)

5. Messmer, Ellen. "Security Debate Rages" Network World Fusion. October
    06, 2003 http://www.nwfusion.com/news/2003/1006ids.html
    (February 22, 2004)

6. Network Ice Cooperation. "Protocol Analysis vs. Pattern Matching". 2000.
    http://www.anitian.com/corp/papers/protocol%20analysis.pdf (February 22, 2004)

7. Parker, Ryan. "Maximizing the Value of Network Intrusion Detection" 2001.
    http://www.intrusion.com/products/download/MaximizingValueIDS.pdf
    (February 22, 2004)

8. Internet Security Systems. "Network Vs Host-based Intrusion Detection"
    October 2, 1998
    http://www.isskk.co.jp/customer_care/resource_center/whitepapers/nvh_ids.pdf
    (February 22, 2004)

9. Intruvert. "Building an IDS Business Case, Q & A with IDS Experts: David
    Piscitello and Michael Rasmussen"
    http://www.networkassociates.com/us/_tier2/products/_media/sniffer/nww_insert
    _____issue3.pdf (February 22, 2004)

10. Intrusion Detection System Group. "An Introduction to Intrusion Detection
    Systems" 2001
    http://www.intrusion-detection-system-group.co.uk/ (February 22, 2004)

11. Ashley, Mitchell. "Layered Network Security: A best-practices approach"
January 2003. Latis Networks, Inc. Whitepaper

12. Maiwald, Eric. Network Security A Beginner's Guide, New York, New York
Osborne/McGraw-Hill 2001

13 "IDS vs. IPS"
http://www.nfr.net/resource/IDSvsIPS.php (February 22, 2004)