Global Information Assurance Certification Paper

Living with an Intrusion Protection System

Ronnie Wagers
GSEC Practical Version 1.4b

The goal of this paper is to outline how I went from being skeptical about Intrusion Detection to becoming an enthusiastic supporter of Intrusion Prevention. The road to my enthusiasm has taught me some lessons and given me some insight into possible future directions. The value of this paper is in sharing the knowledge I have gained on this journey. I became a security practitioner full-time in 2000. I have worked for a Midwestern University with a population of 10,000 students for 27 years in the IT field. This marks the ninth different position I have held there in my career.

According to Vic Wheatman of Gartner Group [1] who attended an RSA Conference in April of 2003, "Intrusion Detection is dead." Analyst Mike Rasmussen of Giga [1] agreed saying "Seventy five percent of IDS installations were failures." Mike blamed a failure to allocate enough resources to weed out false positives.

My own experiences can be summed up in what these two gentlemen observed. I knew that firewalls still allowed some undesirable traffic through on ports that you could not close. I could see the value in analyzing that traffic. I was not entirely without means as my bandwidth management solution was capable of seeing into packets. In fact, both code red and nimda could be detected and discarded. Yet this was not tight enough as the bandwidth manager is designed to let traffic flow and as such when stressed it allows traffic through that it does not have time to analyze.

I investigated what it would take in terms of equipment and time to set up a Snort IDS. It became clear that I could get the hardware without too much effort but finding the time to configure, deploy and tune the box proved too elusive. I made several attempts to line up student talent to help. They quickly dropped out as soon as the time requirements became clear to them.

A colleague in our library contacted me last summer saying he was going to deploy a Snort box in his area. I told him I would be very interested in how this turned out. A month or so passed before I heard from him again. He was trying to sort out what all the hits he was getting meant. After a few more exchanges we decided most of what we were seeing were false positives. In the end we decided it required too much effort to be of any practical value. He has since left the University and the Snort project died of neglect.

Simon Edwards of Top Layer Networks [2] wrote a white paper about network IDS. He lists three performance issues associated with nIDS. Whereas most vendors tell you how many attacks they can detect per second, Simon says the

real problem is how effectively the nIDS can pick out one attack against a mass of normal background traffic. The second consideration is the size of the packets. Most vendors assume a packet size of 1024 bytes. If your traffic mix consists of smaller packets this will have a negative effect on the performance as each packet has to be considered against the signature database. Smaller packets mean more packets in a given time frame. The final consideration is the number of signatures you have in the database. More signatures mean more criteria to match against.

Considering Simon's comments against my experience, I see his first point as the trip point for our failure. We simply had trouble picking out what to pay attention to against the background of normal traffic. The raw amount of data or the packet size that we had to look at did not tax the Snort box.

Simon also points out that where the nIDS is located dictates what traffic you can effectively monitor. My original hope with our Snort project was that once we learned to use it on the Library segment we could parlay that experience into a Snort box monitoring traffic passing thru the perimeter. Obviously since we failed on the smaller segment the idea of deploying on the perimeter was clearly a waste of our efforts.

Returning to the ZDNet article on the RSA Conference [1] that I quoted above, it goes on to say that most delegates on the show floor felt the term intrusion prevention system was merely an attempt by vendors to sell their old products with a fresher sounding buzzword. The article concluded by cautioning us to be skeptical of the new claims by old vendors.

Jon Oltsik writing for CNET News [3] had this to say about Intrusion Prevention: Those pithy scribes at Gartner decided that Intrusion Detection Systems was not conducive to intelligent conversation so they came up with "intrusion prevention".

Both the ZDNet and the CNET articles were written in April 2003. My mindset back in December 2002 was not much different when my boss asked me to investigate an intrusion prevention device recently brought to market in the fall of 2002. I was sure it was going to have too many false positives to be useful, only this time we were going to pay out big bucks to learn this lesson again. This device is meant to be installed inline making it a Network Intrusion Prevention System. It is this device I am speaking about when I use IPS in this paper.

My initial conversation with their sales representive piqued my interest far more than I was expecting. I found myself warming up to the idea even though my colleagues who joined in on the conversation were clearly not interested in another security initiative with the potential to drain dollars from their projects. What grabbed my interest early on was the ability to segment my campus with this device. As I listened I had grand visions of being able to separate the students from the faculty, the researchers from the administrators and all of the

above from everyone else on the planet. We would have lots of false positives but it would be intriguing knowing from which zone the false positive occurred.

I need to acknowledge that this company I was talking to was in fact a startup company. I don't believe they had made any sales at this point. However, they did have a number of their boxes out in the field on evaluations. The sales rep asked me if I was willing to sign up for their evaluation program. I decided I had nothing to lose because it was clear this project would not advance on mere talk alone. I entered into a 30-day evaluation period.

My vision of how we would hook up the IPS box was the first collision with reality. Our network topology did not allow for the neat partitions of users I had first pictured. What became of my four zones was a zone for traffic coming into campus, a zone for traffic leaving campus, a zone for our main server farm and a zone for our administrative support systems. This IPS box has the ability to support several sets of ports resulting in the ability to slice up your network into as many as 20 different zones. Each zone passes thru the same logic engine that is the heart of this beast. The claim is 2 gig of traffic a second can be processed.

The first thing I noticed once we were hooked up was the ability to locate code red users and nimda users. I quickly contacted the users on campus and was pleased to find out the IPS got it right. I continue to this day to get code red hits and nimda hits from outside campus. One future direction I see is to collect this data at one site for all users of this IPS and then have that central site try and contact each of these infected sites and try to get some of the noise on the internet in general to diminish. Perhaps it would make more sense to funnel this information to an existing site such as DShield but I like the idea of a separate site as I have come to believe it is highly accurate at finding these infections.

You are no doubt wondering what became of my fear of all the false positives. I did get one the first week we had the IPS. It incorrectly identified users updating web pages with Microsoft's Frontpage tool as being an exploit. The matter was quickly resolved with a dump of the traffic in question being sent to their tech support. It turned out to be related to how much of the packet they were examining. The distinction between what a legitimate packet from Frontpage and a packet that was part of this exploit occurred after the 200th byte in the packet. This is the only false positive that has occurred to my knowledge.

Think about what this means. After dealing with all the false positives from trying to use Snort to clue me in on what might be bad traffic, I have a box that not only gets this right but also then turns around and stops that traffic from continuing. I am out of the worrying business about what I am missing. Now you might consider turning over part of your protection to this IPS as very scary stuff. After all isn't it outside your control as to what is good traffic and what is bad traffic?

That is a very good point. First off I am not very talented at deciding what is good traffic and what is bad. That is one of the chief failings of Snort from my point of view. Unless you were good at understanding what it was telling you to begin with you had no chance of being able to use it effectively. I believe all successful Snort implementations have people skilled at understanding bad traffic signatures behind them.

My new IPS has a very talented signature team working for them. I was reading a SANS news bulletin one day and was surprised to see them bragging about a SANS members prowess at signature writing. The surprise was that my IPS has this person leading their signature team. The clear lesson to me is that I can never compete on that level against such talent. I am quite content to leave this piece of sleuthing to them.

Let's stop and define what makes up good traffic versus bad traffic. Good traffic is that traffic which enables us to meet our business goals. This does not mean what is left over is by default bad traffic. Some of our traffic also consists of activity that is personal in nature for our employees. Consider a parent who receives status messages from their children in the form of emails. That employee will be more relaxed with less worry about the children and therefore more productive. In general I feel the more access an employee has to the Internet at work, the happier that employee will be. It only becomes a problem if their productivity falls off.

The definition of what good traffic is versus bad traffic is even grayer at a University than it is in a typical business environment. It is very hard to make a judgment call on what traffic is promoting our business goals when our mission is to educate people. You can entertain the idea that pornography might have educational value if the goal of the exercise is to define what makes a given picture art.

Fortunately it is easier to define bad traffic. Bad traffic is traffic that puts at risk our ability to meet our business goals by denying us access to our network and the critical machines located there. A second white paper from Top Layer Networks [4] defines six types of bad traffic.

The first is traffic that seeks to exploit protocols. Often this type of traffic is in the form of probes against known ports looking for weakness or poor configurations in the machines they scan. Successfully locating a weakness, such as no password on a Win2k machine leads to that machine being compromised. Compromised machines are often used as a springboard for the other types of bad traffic. My University receives thousands of such scans an hour. A real life example of this would be someone who goes up and down your street trying all the doors and a windows looking for ones that are unlocked.

The second type of traffic concerns exploiting HTTP. Code red and nimda are two examples of this type of bad traffic. Each consisted of a special crafted URL sent to a web server, in this case Microsoft's IIS. Both in turn generate large amounts of traffic looking for other machines to infect. This extra traffic is definitely bad news for your network as it has the potential to choke out the good traffic much in the same way as weeds can choke out the good plants in a garden.

The third type of bad traffic seeks to overwhelm the hosts at which they are targeted. The classic three way hand shake of setting up a TCP/IP connection can be exploited by continuously sending the SYN packet without ever acknowledging the SYN/ACK from the host. The host will keep resources tied up waiting for the ACK that never comes. Imagine if you would, answering your phone only to find no one on the line. How much useful work can you accomplish with such a continuous interruption? Add to that the complication of having to hold your phone line open for 30 seconds just in case someone did decide to speak.

The fourth type of bad traffic seeks to exploit FTP. Bad configuration might result in an FTP server that accepts anonymous logins. Such an FTP server can be used to store and send files that benefit others outside of your business. This is such a common activity that it has a name, Warez. How happy would you be if someone were moving stolen property in and out of your garage? If your garage got popular enough, one day you would find you could not park your car there anymore.

The fifth type of bad traffic uses ICMP packets. ICMP packets are useful when performing diagnostics on your network by your network administrators. They become bad traffic when they are used to map your network by hackers who study the different responses different types of host operating machines produce. It is also possible to hide information inside ICMP packets. Enough ICMP packets flowing on your net will result in denial of service attack against your net. The Nachi/Welchia worm has that effect on us now. If you think freeway traffic jam here, you get close to the idea. Proper realism would require everyone to drive the same car and all moving as fast as we could towards that toll gate with only one correct change booth open!

The last type of bad traffic is traffic that exploits flaws in application software. A common program error is to not check the input data. Hackers look for this situation and find ways to send carefully crafted data that results in the program having its original instructions over written with new instructions of the hackers choosing. This buffer overrun is one of the more popular ways of compromising host machines. Suppose you were following a recipe to bake cookies. You start the recipe and someone manages to alter the recipe as you are following it. When the oven is opened at the end of the cook cycle you find lima bean casserole on the rack instead of those yummy cookies.

Top Layer Network's white paper [4] goes on to say that detecting this is not good enough. Indeed that would have been one of the problems with our Snort project had we managed to successfully implement it. We would know this activity was occurring but we would be overwhelmed trying to do something about it. In most of the examples of bad traffic listed above, my IPS does that 'something' for us. The area where it too fails is in stopping probing behavior described in bad traffic type one.

The reason this can be a problem is if a hacker were sending a continuous flood of probes that examined a series of ports and then moved on to the next machine we could see that pattern in a small enough time frame to react. But suppose a hacker varied the order or even which ports he looked at from machine to machine. Further suppose he varied the sequence of which machine he looked at next. He might even vary the time between packets. Each of these tactics increases the difficulty of being able to see the pattern. Our IPS cannot be expected to remember each traffic flow long enough to be able detect the pattern of abuse. My IPS has an anomaly detector but it finds good traffic as much as bad traffic. Because this is such a difficult problem it only reports the anomaly. This is one area where perhaps some better algorithms and better storage structures might result in an improvement to the point where this activity could be blocked with out disrupting normal traffic. Indeed my IPS vendor has promised this improvement is on its product development path.

Cisco has a white paper where they discuss Heuristic-Based Analysis entitled The Science of Intrusion Detection System Attack Identification [5]. They say that an algorithm of this type would use statistical evaluations of the type of traffic being presented. A good example of this type of signature would be to detect the port sweeps mentioned in the preceding paragraph. They go on to say such a signature looks for the presence of a threshold number of unique ports being touched on a particular machine. Their "pros" statement says this is the only way to detect some types of suspicious/malicious activity. Their "cons" statement says these algorithms may require tuning or modification in order to conform to network traffic and limit false positives.

My IPS has these types of threshold settings for its anomaly detector. Perhaps some tweaking is in order to see if I can eliminate the reports against machines on my network that should be performing good traffic. I have not invested a lot of time in this activity because it is a given that as a University we know we are being scanned. Quantifying this activity with no real way to stop it seems like an exercise in futility at worst. At best we are wasting time that can be better put to use with other security concerns where a difference could indeed be made.

Network Associates has a white paper entitled Intrusion Prevention: Myths, Challenges, and Requirements [6]. The myths they list are first that Intrusion Detection and Intrusion Prevention Are Two Separate Solutions. I certainly do not

hold to this view. The only places detection alone is useful is in a forensic situation and trend analysis. In all other cases it is information that has arrived too late.

The second myth is Intrusion Prevention is ALL or NOTHING. This is certainly not true in the case on my ISP. The signatures are divided up into two categories. I have an absolute category in which the signatures are strongly believed to be always dangerous. They are all customizable to the point I can remove them from the absolute category completely or I can change their setting. The setting choices are alert only, alert and block, block only, or the final choice of ignore. There is also a policy category where we can pick and choose signatures from the perspective of what will support our University policy statements such as our Acceptable Use Policy. One such policy signature is the ability to block psexec.exe from being used to remote control a system. Psexec.exe is a favorite of hackers for allowing commands to be executed on remote systems. It also may have a legitimate use for your system administrators to remotely administer machines. Even here it is not an all or nothing choice. I can configure the running of psexec.exe to be denied except for these exceptions. The exceptions are of course machines we want the administrator to administrate right down to source destination pairs if desired.

The third myth is Intrusion Prevention is TCP Kills/Resets or Modify Firewalls by IDS. There is no need to let your hacker have a clue as to what has happened to his malicious packets. Dropping them is the preferred course of action in that regard and one that my IPS follows. There is one area where this proves to be a potential problem. Email based worms such as Sobig and Bugbear. Both these worms are detected and blocked by my IPS. The problem is the sender gets no feed back that the email was accepted or rejected ergo the assumption is the email was lost in transmission. Our good and faithful email servers truly have the postmen's creed in their programming. Neither wind, nor rain, nor dark of night nor drop packets shall stay their course to deliver the mail. The email server sends the message again and my IPS drops it again.

Now our email servers are only slightly more tireless than flesh and blood postmen. They will reach the point where they will give up. The problem from the IPS's perspective is that I now have at least hundreds if not thousands of extra log entries from what should have been one incident. A kind soul at Purdue University helped me sort this out one evening. My apologies to him as I don't believe he ever truly appreciated his role (or responsibility) in this incident. However, he did stay patient with me long enough for at least one of us to uncover the truth. Perhaps donations to Purdue in my name might help make amends.

One last comment about the third myth, it needs to stay a myth as far as modifying the firewall rules go. They are separate but complementary security devices. I would be more than uncomfortable to allow such an event to occur.

Happily my IPS has never considered this to be something it wanted to try. If it ever did would Cisco extended access control lists be far behind?

The fourth and final myth that Network Associates puts forth is Intrusion Prevention is Losing Control Over Intrusion Detection and Response. In light of my previous examples I think it is clear that this is not the case. The one exception is the email response. I still remember hearing some nameless SANS instructor (their initials are EC) saying not to give them a clue by responding if you are ultimately going to deny the traffic. Of course the subject was firewalls but I believe it applies equally to any security device.

Jamil Farshchi wrote a paper for SANS called Statistical Based Approach to Intrusion Detection [7]. He makes his case for the Statistical Approach by pointing out the shortcomings of signature-based systems. Then he discusses why stat-based systems would be a benefit over signature-based systems. My experience with my IPS and my work background at a University make me uniquely qualified to comment on two of his observations.

The lynch pin of a successful stat-based system is the ability to learn normal behavior and then treat deviations to this normal behavior as anomalies that need action. I have already stated that Universities are scanned on a regular basis. This scanning is from outsiders looking for some weakness to exploit as well as insiders who may be practicing some thing they learned in computer security class. How can a stat-based system ever learn good behavior in an environment that is almost by definition has scanning as a normal behavior? In fairness to Jamil, he does point out this possible weakness but labels it as a drastic scenario. Sorry Jamil but it is normal life at a University. (Please look me up at SANS Jamil… I no doubt owe you a beer.)

Jamil correctly points out that a serious shortcoming of signature-based systems is that it is only as strong as its signature set. I whole heartily see his point when I consider my Snort experience. However, when I consider my experience with my IPS then I have to point out that this “weakness” is in fact one its strengths. I had a slammer worm defense within 8 hours of the outbreak. Now to be fair I also had help from the firewall on this one but the bottom line is that no machine at my University got this worm. I knew this because once the signature had been loaded I did not see any of my machines show up in the hit lists. I was in fact quite puzzled by the lack of inside infections at first. It was only later I made the observation that UDP port 1434 used by slammer was not an open port through the firewall. I had time to contact the community and insist that the patches be installed before someone introduced the worm by some other vector besides our perimeter.

Current events have provided me both with a major distraction to finishing this paper as well as providing me an excellent example for this part of the paper – our friend the blaster worm. My IPS’s signature team is made up of world-class

signature writers. I have already mentioned SANS bragging about the team's leader. They had a signature to protect against blaster prior to the outbreak. I have come to expect them to pay better attention to news than I do and act accordingly. Unfortunately for all of us, blaster exploited a weakness present in our newer windows desktop machines. Slammer was only effective against machines running MS-SQL resolution. Slammer had a smaller potential audience. Blaster hit my University hard despite me having a perimeter defense against it. How did this happen?

Thanks to my signature-based IPS I know exactly what occurred. It dutifully recorded internal machines trying to infect the outside world. All of these early hits are concentrated on one subnet, a subnet that is used by our helpdesk hardware engineers. Someone brought a machine in to investigate why it was rebooting constantly. That machine had been at the owner's home plugged into a cable modem prior to bringing it to campus. Before anyone on our staff had knowledge that blaster was occurring on the net they plugged the misbehaving machine in to our net to download the latest patches. Now my IPS was being used as a forensic device to help us locate the infected machines on campus before anyone had released a scanning tool to locate them. The positive side to this adventure is that I now have a concrete example I can point to saying no matter how good my perimeter defenses are, there is absolutely no substitute for keeping our machine patches up to date. Somehow I don't think I would have as strong a case if I had to sit down and prove my stat-based system learned about this anomaly quick enough to prevent outside infections. The signature-based system had without a doubt a defense that was in place prior to the blaster outbreak.

To help me close, TippingPoint Technologies has a white paper entitled The Profound Benefits of Network-Based Intrusion Prevention [8]. They state in their conclusion that having an IPS will improve your corporate productivity and profitability. My experience supports their first statement. By having my IPS in place I have saved significant amounts of time in locating infections that by their nature are flooding our net with unwanted traffic. It is also possible to configure a machine before it gets infected in our test labs thanks to the help from my IPS. They next state it will protect sensitive information from being stolen. I don't have any examples where I can say this has occurred, but I can say we have not had any loss of sensitive information since my IPS has been in place. They also say it will protect key-infrastructure from imminent global cyber-attacks thus preserving standards of living and ways of life. My observations are that if everyone had an effective IPS we could go along way towards preventing and/or confining some of the outbreaks we have all suffered through. You can't get warm and fuzzy feelings like that from an IDS. I have converted from an IDS pyrrhonist to an IPS disciple.

References

[1] Judge, Peter. “Intruders: Is Detection or Protection the Answer?” 15 April 2003. URL: https://secure.zdnet.com.au/itmanager/strategy/story/0,2000029582,20273747,00.htm (10 July 2003)

[2] Edwards, Simon. “Must Know Information on Realizing and Overcoming the Risks.” Network Intrusion Protection Systems: Important IDS Network Security Vulnerabilities. September 2002. URL: http://www.toplayer.com/pdf/WhitePapers/wp_network_intrusion_system.pdf (7 August 2003)

[3] Oltsik, Jon. ”Off the Hype Meter.” 29 April 2003. URL: http://news.com.com/2102-1071_3-998529.html (10 July 2003)

[4] “Beyond IDS: Essentials of Network Intrusion Prevention.” November 2002. URL: http://www.toplayer.com/pdf/WhitePapers/IPS_Whitepaper_112602.pdf (7 August 2003)

[5] ”The Science of Intrusion Detection System Attack Identification.” 21 May 2002. URL: http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/prodlit/idssa_wp.htm (8 August 2003)

[6] “Intrusion Prevention: Myths, Challenges, and Requirements.” April 2003. URL: http://www.intruvert.com/technology/pdf/IntrusionPrevention_WhitePaper.pdf (12 August 2003)

[7] Farshchi, Jamil. “Statistical Based Approach to Intrusion Detection.” Intrusion Detection FAQ. URL: http://www.sans.org/resources/idfaq/statistic_ids.php (16 August 2003)

[8] “The Profound Benefits of Network Based Intrusion Prevention.” URL: http://www.tippingpoint.com/resource_library/pdfs/PNBIPS80321.pdf (19 August 2003)