



Deploying and Tuning Network Intrusion Detection Systems

**A corporate white paper
from the product management and
sales engineering group
of Intrusion.com**



Intrusion.com SecureNet Pro Network IDS Sensor Deployment and Signature Tuning Recommendations

This document has been created to assist in answering the common questions encountered concerning how to best deploy Intrusion.com SecureNet Pro (SNP) network intrusion detection sensors in a network.

Every network is unique and requires thorough examination before any implementation can successfully take place. This document has been created to help directionally, as a recommendation for planning assumptions. When you begin to implement your network intelligence project using intrusion detection systems, take it step-by-step, work with a VAR, managed service provider or integrator who is familiar with IDSes and make sure your networking team is bought-in.

Table of Contents

Primer	3
Security = Visibility + Control.....	3
What are you looking for?	3
Deploying Network Intrusion Detection Systems	5
Deployment in a Simplified Network.....	6
First Priority: behind the perimeter firewalls	6
Second Priority: within the DMZ(s).....	6
Third Priority: between border-router(s) and perimeter firewalls	7
Fourth Priority: behind firewalled subnets of the primary LAN	7
Fifth Priority: behind remote/branch office firewalls	8
Special Note: e-business and partner connections	9
Tuning the Network Intrusion Detection System	9
Blinding the Sensor	9
Blinding the Operator	10
4-Step Tuning	10
Signature Tuning.....	11
How to Get Started.....	13
Start with what you know.	13
Start with one IDS or Sniffer.....	14
Practicum	14

Primer

Security = Visibility + Control

Security is not a deliverable. Security products provide two primary benefits: visibility and control. And, it is the combination of these two benefits that make it possible to create and enforce an enterprise security policy to make the private computer network secure.

- **Visibility:** the ability to see and understand the nature of the network and the traffic on the network.
- **Control:** the ability to affect network traffic including access to the network or parts thereof.

Visibility is paramount to decision making. Visibility makes it possible to create a security policy based on quantifiable, real-world data. This reality-based policy will enhance appropriate spending on security solutions and help guard the enterprise against unnecessary or premature expenditures. Additionally, visibility is the first element of predictive analysis, allowing the enterprise to invest in prevention technologies before a vulnerability is exploited. Visibility systems also enhance the value of control devices by providing quantitative validation of control system performance and the effectiveness of the security policies that drive them.

Control is paramount to enforcement. Control makes it possible to enforce compliance with security policy. Control systems allow security and IT professionals to shape and form the traffic in the network so that the enterprise may be reasonably assured that the information assets that live within the network are not vulnerable and have not been compromised.

Network Intrusion Detection Systems (NIDS) are the primary provider of security visibility and the resulting network intelligence required to make control devices effective and the enterprise network secure.


Deploying NIDS follows stereotypical network topology and provides different types of network intelligence depending on the position of the NIDS in the network. Each deployment provides visibility into different types of traffic and segregated network segments. Each sensor deployment can then be tuned to provide maximum performance with the minimal noise for the network intelligence required.

What are you looking for?

There are six types of events that a network IDS will report.

1. **Intrusion Attempt:** the attempted exploitation of a known vulnerability. Intrusion attempts are very rarely incorrectly alerted

upon as they are tied to specific behavior that is not common in the normal course of business.



Events are classified to help security professionals more accurately understand the nature and priority of network traffic. Each ring-level toward center raises the bar for threat and potential risk to the network.

- Intrusion attempts, DDoS and DoS attacks alert as **HIGH** priorities.
- Suspicious activity and protocol anomalies alert as **MEDIUM** priorities.
- Network events alert as **LOW** priority.

2. **Distributed Denial of Service (DDoS)**: the coordinated flooding of packets from many computers to a single target. The goal of this type of attack is simply to block the targets ability to communicate with other systems, or to crash the target system. These signatures identify the use of specific known tools for DDoS attacks.
3. **Denial of Service (DoS)**: the exploitation of known vulnerabilities that allow a single attacker to stop another computer from communicating or to crash. These types of attacks can be incorrectly diagnosed if the alert has too low a threshold and alerts on normal business activity, categorized later as a network event.
4. **Suspicious Activity**: activities that have been shown to be precursors to attacks or that are not typical of normal business activities. Suspicious activity includes reconnaissance activities used for information gathering, network mapping and scanning activities as well as failed access attempts. The importance of this type of event will vary greatly by enterprise and sensor deployment and must be proactively addressed to provide accurate and valuable network intelligence. Suspicious activities occur before potentially damaging events and may be the security professional's best bet at proactively securing the network.
5. **Protocol Anomaly**: the detection of network traffic that doesn't conform to known standards. Anomaly detection differs from the prior four event types in that protocol anomaly detection alerts you to traffic that doesn't conform but cannot tell you what that means. The prior four event types are fairly accurate in describing the exploit and its potential resolution. Because protocol anomaly detection analyzes network traffic for deviation from standards rather than searching for known exploits, there is the potential for protocol anomaly events to serve as an early detector of undocumented exploits.
6. **Network Event**: if not specifically identified as one of the prior five types, everything is a network event. Sometimes incorrectly identified as "false positives," network events are events that are not perceived as direct attacks, but will play a part in forensic analysis and may hold importance to certain enterprises or organizations. Network events are most effectively reported by NIDS, like Intrusion.com SecureNet Pro, that provide a large number of decoded protocols – allowing the NIDS to correctly record and report on happenings within specific protocols.

Understanding the relative value and importance of each of these types of events in relation to the deployment of the sensor helps the security professional determine the appropriate tuning for each sensor deployment.



Tuning increases the “signal to noise” ratio, so that the security professional is able to easily find the events of all types that are of greatest importance within each unique network.

Deploying Network Intrusion Detection Systems

In the past, network IDSes were expensive, complicated and came in a one-size-for-all configuration. As a result of these pricing models and inflexible complexities, early NIDS deployments were limited to between the border-router and firewall where they attempted to monitor all enterprise traffic. This was acceptable to launch a new technology and fit the new networks and the single, 10Mb/s Ethernet architecture that dominated the market.

Today, there is Ethernet, fast-Ethernet and Gig-E network segments. Each speed of network plays a specific role in fulfilling financial and scalability needs – but they do not all warrant the same expenditure for security. Intrusion.com offers a family of network IDSes that fit the various speeds and deployment requirements of modern networks. Replacing the one-size-for-all model with a scalable, highly deployable range of IDS appliances allows the security professional to extend his/her network intelligence umbrella further into the network than ever before – thereby making the control devices more effective and the enterprise more secure.

Additionally, the expanse that networks cover today is far greater than ever before. These huge networks are being segregated by firewall technologies and connected via VPN technologies. Both firewalls and VPNs protect network segments and occlude visibility.

- When a subnet is segregated by a firewall, the security professional loses his/her visibility into the type and nature of the traffic in that segment.
- When network traffic is encrypted (VPN), the security professional loses his/her visibility into the nature of that traffic.

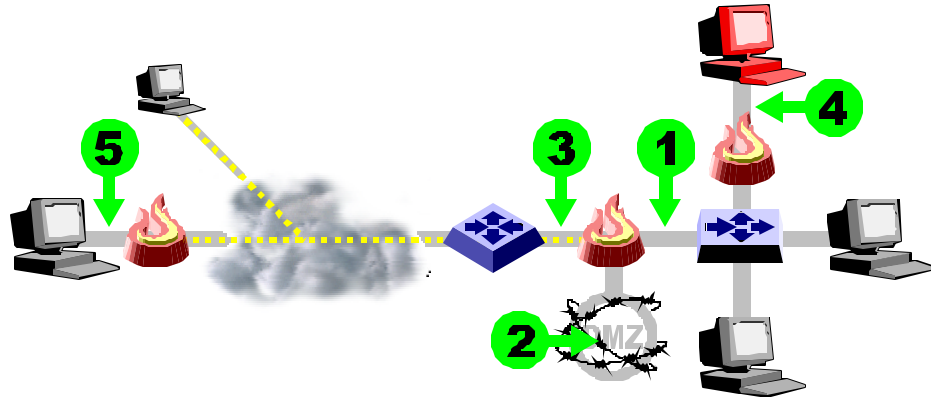
Placing a NIDS behind firewalls and VPNs, especially in widely dispersed and segregated networks, provides the security professional with the much needed visibility to ensure the security of the network.

Deployment in a Simplified Network


Viewing a simplified network, there are five typical deployments.

Typical Deployments in Priority Order

1. Behind perimeter firewall
2. Inside the DMZ
3. Between border-router and perimeter firewall
4. Behind internal, departmental firewalls
5. Behind remote/branch office firewalls



First Priority: behind the perimeter firewalls



SecureNet Gig for Gig-E



SecureNet PDS 5000 Series for 100Mb/s segments with local management capabilities

Should you only be able to afford one NIDS sensor, this is the place to put it. This is a central chokepoint of aggregate traffic that passes into and out of the enterprise private network. This position allows the NIDS sensor to provide a gross level alarm that something is wrong or has made it through the firewall and into the private network. It will not provide visibility into suspicious activity that remains within subnets nor monitor outside traffic to the DMZ. The level of exposure at this deployment warrants activation of signatures from suspicious activity through intrusion attempt.



It should be noted throughout that logging should be turned on for all attacks, especially those of a buffer overflow type of attack.

Second Priority: within the DMZ(s)




SecureNet Gig for Gig-E



SecureNet PDS 5000 Series for 100Mb/s segments with local management capabilities

This is where the enterprise's service for and to access the outside world are kept including web servers, ftp servers and email servers. This position allows the NIDS sensor to provide information about the network traffic and activity that affects these outward-facing servers where the majority of denial of service, web exploit and email attacks will be targeted. The level of exposure at this deployment warrants activation of signatures from protocol anomaly through intrusion attempt.



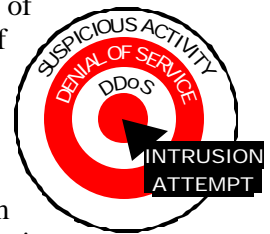
A specific area of signature tuning that can be particularly useful within the DMZ is taking advantage of Intrusion.com SecureNet Pro's module decodes which allow the security professional to track

all network events and log them. This provides a large amount of data about all activities within the DMZ, not just security events. This log data will be especially valuable for forensic analysis of any breach that affected or molested DMZ physical and/or information assets.

Third Priority: between border-router(s) and perimeter firewalls



This deployment provides visibility into reconnaissance and exploits attempted directly on the firewall. Some administrators desire this added visibility into the external attacker's pattern as part of national defense, corporate intelligence and overall heightened security. Much like a honey pot, this deployment provides information about what people are trying to do while attacking the enterprise network. The level of exposure at this deployment warrants activation of signatures from suspicious activity through intrusion attempt.




Deployment outside of the firewall presents a number of significant risks that may deter some administrators from electing it. The risks involve exposing intelligence information or perhaps another hole through the perimeter firewall. If the risk is worth the value of the visibility, there are six final recommendations.

- The monitoring interface should have no IP address.
- Change the communications port for SecureNet Pro (default 975) to mask the identity of the NIDS.
- Utilize a non-routable address (RFC 1918 like 10.x.x.x) for the management interface on this sensor.
- Apply an aliased interface on the outside of your firewall, or
- Add an interface to your firewall specifically for the IDS, or
- Create a separate V-LAN for intrusion detection management (ie: that routes around the firewall).

This will allow traffic to pass back to the console, but not make the sensor as readily apparent on the outside of your network.

Fourth Priority: behind firewalled subnets of the primary LAN

Typically, this deployment is for the protection of mission critical servers like ERP, CRM, PDM and accounting systems. In addition to placing mission critical servers behind firewalls, it is becoming more common that critical departments in the enterprise are having their network segments



SecureNet Pro Sensors for this deployment



SecureNet PDS 5000 Series for 100Mb/s segments with rack mount wiring closets.



SecureNet PDS 2345 for 100Mb/s segments that require a desktop appliance.

segregated from common traffic. Firewalls are used internally to segregate departments like the following.

- Executive
- Finance
- Human Resources
- Research and Development
- Engineering
- Patents
- Legal

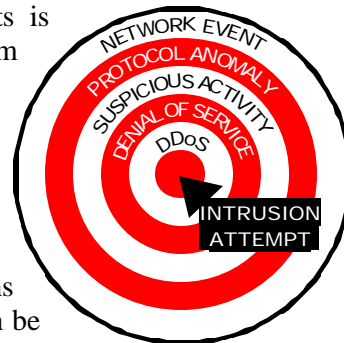
This deployment may also be used to cover ebusiness connections to partners, which is covered in the following section.

These firewalls not only protect these network segments, but as discussed in the primer they also stop other IDSes from monitoring the traffic behind these firewalls. No department is immune from disgruntled employees and back door attacks creating holes in the network perimeter and putting information assets at risk.

These deployments frequently go into networks that are switched, which present the classic network monitoring problems that are inherent in all switched environments. Deploying SecureNet Pro within the switched environment is accomplished via the spanning port of the switch. Another option is deploying a small hub or switch on the backbone between switches. This allows the full duplex traffic to pass over the monitored segment. Deployment in switched environments is outlined in technical documentation available from www.intrusion.com.

The level of exposure at this deployment warrants activation of all signatures from network event through intrusion attempt.

Additionally, because of the additional restrictions of the interdepartmental firewalls, these areas can be more prone to self-configured PCs or user efforts to gain access to other departments or the Internet. Protocol anomaly detection and network event signatures can come in particularly handy to detect suspicious events that comply with security policies or are behavior based, rather than signature based.



Fifth Priority: behind remote/branch office firewalls

The enterprise is now a virtual organization that is spread across many remote and branch offices – all needing to be connected to headquarters with access to many of the enterprise information assets.

This deployment may also be used to cover e-business connections to partners, which is covered in the following section.



SecureNet Pro Sensors for this deployment

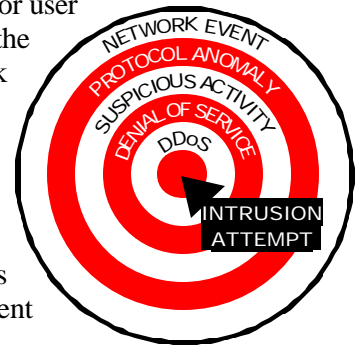


SecureNet PDS 5000 Series for 100Mb/s segments with rack mount wiring closets.



SecureNet PDS 2345 for 100Mb/s segments that require a desktop appliance.

As in the deployment scenario above, the firewalls as a perimeter for remote and branch offices not only protect them, but also stop the enterprise IDSes from monitoring the traffic at these remote sites. No department is immune from disgruntled employees and back door attacks creating holes in the network perimeter and putting information assets at risk. Additionally, because of the distance from headquarters, remote offices can be more prone to self-configured PCs or user efforts to gain access to other departments or the Internet. Protocol anomaly detection and network event signatures can come in particularly handy to detect suspicious events that comply with security policies or are behavior, rather than signature based.



The level of exposure at this deployment warrants activation of all signatures from network event through intrusion attempt.

Special Note: e-business and partner connections

Many businesses today are dependent on sharing enterprise proprietary information with partners in the industry via Internet or WAN connectivity. Depending on the network layout, these partner connections may be similar to the fourth priority deployment where it is within the company’s perimeter or it may be treated as an extranet connection that would be similar to the fifth priority deployment. Regardless of the deployment, partner connections pose a different type of threat for the enterprise. The military adage of “trust but verify” is the framework for these e-business connections and they require a higher priority for network intrusion detection class intelligence.

When a #4 or #5 deployment is an e-business or partner connection it may warrant higher priority than even DMZ deployments. Connections to other networks add all of the vulnerabilities of that network to your own.

Tuning the Network Intrusion Detection System

There are two primary weaknesses within a network intrusion detection system that can be exploited to get an attack into the network without detection: (1) blinding the sensor and (2) blinding the operator.

Blinding the Sensor

Hackers attempt to blind the sensor to make it difficult or impossible for the NIDS to detect a real attack by flooding the network with dummy

traffic to mask the real attack or using malformed traffic to evade the IDS. Certain types of network traffic, mainly small, fragmented and/or malformed packets can bring a NIDS to its knees or take it completely out of commission. The recently publicized “stick” attack was able to completely blind the industry leading NIDS – making it possible for another attack to be sent through completely undetected. Additionally, some NIDS sensors don’t reassemble fragmented packets or use simplified means, gaining them marketing without actually delivering a solution. Some NIDS claims to reassemble packets but actually only do partial re-assembly, leaving the user still vulnerable to a knowledgeable attacker. By implementing advanced intrusion detection technologies like stateful inspection, multi-path packet reassembly and high-speed packet analysis, some of the more advanced, second generation NIDS, like Intrusion.com SecureNet Pro, are more resilient and effective at avoiding blinding or evasion attacks and are capable of detecting attacks on normal network traffic at near-wire speed.

Buying the right NIDS starts with the sensor. There is nothing that can be done at the customer premises or with creative management tools and GUIs that can improve the performance of a sensor or increase its ability to detect attacks under network saturation or attacks that are deliberately crafted to evade IDSes. Above all other metrics, the core performance and suitability of the sensor will set the foundation for your intrusion detection architecture.

Blinding the Operator

In contrast to blinding the sensor, blinding the operator is substantially affected by the deployment and tuning of the NIDS, which is completely under the control of the security administrator.

Blinding the operator is simply done by having too much information being sent to the network intrusion detection console. Too much data makes it difficult or impossible for the security administrator to recognize immediate threats within the data being presented.

To control the possibility of blinding the operator, the sensors must be “tuned.” Sensor tuning is the process of determining what signature, under what parameters – called a policy, should be deployed and how the network IDS should be configured to increase the number of pertinent events reported and a percentage of network events.

4-Step Tuning

To minimize the likelihood of either the sensor or the console being blinded, Intrusion.com SecureNet Pro provides multiple levels of tuning that provide the security professional unsurpassed capability to find the

events that are relevant to his/her unique network. Tuning comes in four phases:

1. **Limit the number of signatures you are looking for.** Especially with a product like SecureNet Pro that does extensive module decoding, you need to decide what elements of the protocol you view as a threat to your network or network segment.
2. **Use SecureNet Pro's global filtering.** Global filters allow you to limit the amount of data brought into the system from the sensor. Global filters allow security professionals to use a single signature policy throughout the organization and then customize the way policies are executed at the sensor level by filtering packets by Ethernet, IP and protocol before signature processing occurs.
3. **Filter events at the console.** To limit the amount of data presented to the security professional – console filtering leaves the most amount of data in the database for forensic analysis but limits the effect of “blinding the operator” by presenting only the events deemed relevant.
4. **Tune the actual signature.** Signatures may be tuned so that appropriate network events are not incorrectly reported as other, more threatening event types (sometimes called false positives) or alerted as less threatening event types or not alerted at (sometime called false negatives).

Signature Tuning

Using Intrusion.com SecureNet Pro, the parameters of the signatures may be changed to make them more applicable to each unique network. This includes editing the description and other text fields to be able to add site-specific information as well as other items.

The first level of signature tuning is to modify the **priority** of the event to suit the level of importance for the network. As discussed previously, in some networks, a suspicious activity event may be of a “medium” priority while in another network this same event could be a “high” priority. To increase the likelihood that the security professional will see the information they are looking for, the signature priority should be analyzed to match the priorities of the security professional.

Additionally, the security professional may choose to analyze the **classifications** of the signatures. Each signature is classified as either an: intrusion attempt, DDoS, DoS, suspicious activity, protocol anomaly or network event. In general, these event classifications are typical of all organizations. Each enterprise's security professional should study signature classifications in context of his/her network. For example, an



event classified by Intrusion.com as a network event may be reclassified as suspicious activity to fit the security policy of a specific enterprise.

When using Intrusion.com open source or user created **string-matching signatures**, the search strings and parameters for these signatures may be customized to make the signatures more accurate. Open source string-matching signatures may also be duplicated and renamed to create additional identification capabilities. Open source signatures may be created in the SecureNet Pro Linux console or using text editors. Users may create string-matching signatures to address ad hoc and network specific intrusion detection and network intelligence needs.

When a string matching signature needs to be more specific than the native filtering capabilities will allow, Intrusion.com provides a utility that makes it fast and easy to add **complex IP filtering** to any signature. Using this utility, IP filtering can be done on an inclusive, exclusive, range, list and ordered basis.

Context analysis signatures may be parameterized like other signatures, to vary the priority, classification, description, IP, MAC and other variables. In the rare instance where an Intrusion.com context analysis signature is not accurately identifying events for a specific network, these signatures may be made inactive and replaced with user created signature that use different triggers to identify events.

The following section walks through the process of deploying and tuning NIDS. It is intended for people who have a high level of understanding of network IDSes and signatures.

“Deploying and Tuning Network Intrusion Detection Sensors” was written in concert by the product management and sales engineering groups of Intrusion.com to provide a blend of product technology with real-world experience. To contact the authors, please email them directly.

Ryon Packer is the vice president of product management and can be reached at rpacker@intrusion.com.

Michael J. Staggs is a sales engineer for the western states and can be reached at mstaggs@intrusion.com.

Darren Harlow is also a sales engineer for the western states and can be reached at dharlow@intrusion.com.

How to Get Started

How do you know which signature modules are appropriate for each SNP sensor? Since each network is unique in its requirements, you must treat each network differently. We will use the two-step approach: Start with what you know and then add to it.

TABLE 1

SNP Sensor 1

Active Modules

- All UDP
- All telnet
- All tftp
- All TCP Scans
- All finger
- Set email alerts on

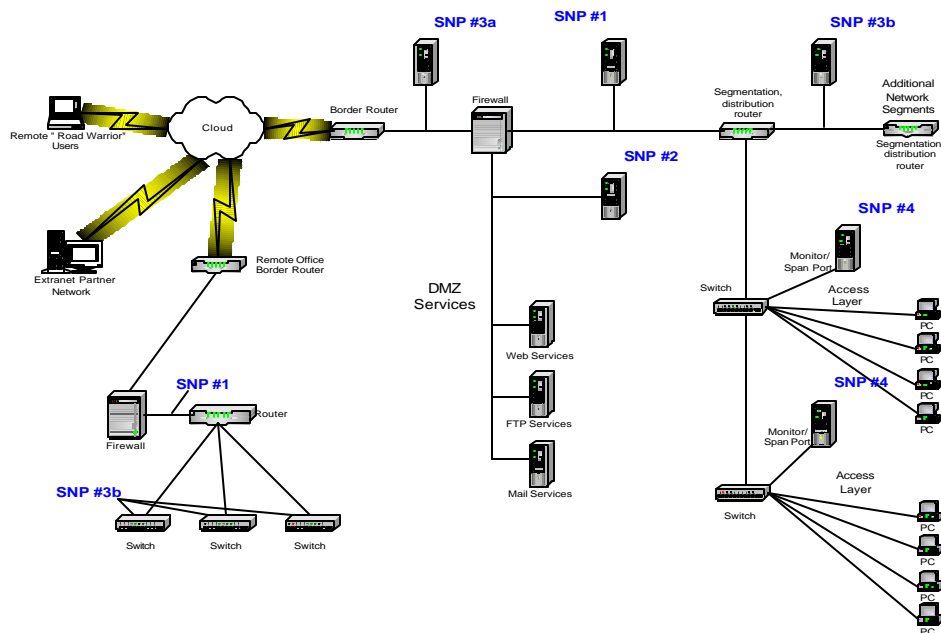


TABLE 2

SNP Sensor 2 Active Modules

- | | |
|--|---|
| All SMTP modules | All Telnet |
| All FTP modules with the exception of: | All UDP |
| 4dgifts | All HTTP except: |
| FTP Abort | HTTP cookie sent |
| Anonymous Login | HTTP bash_history attack |
| Bisonware | HTTP client .rhosts seen |
| Broker Probe | HTTP client AnyForm attack |
| CSM Proxy Overflow | HTTP Client Cold fusion attacks |
| FTP client password overflow | HTTP Client get cmd |
| FTP client retrieve data | HTTP Client Lotus Notes delete document cmd |
| All finger modules | HTTP Client Lotus Notes edit services cmd |
| ICMP decoder | HTTP Client Netscape Page services probe |
| ICMP ping flood | HTTP Client Novell convert attack |
| ICMP ping Sweep | HTTP Client Post |
| All MSTREAM modules | Both HTTP Client Sambar attacks |
| All Portmapper modules | HTTP Client Spaceview attack |
| All rlogin modules | HTTP Client Web+ attack |
| TCP decoder | HTTP client aglimse attack |
| TCP FIN port scan | HTTP client bizdb1-search attack |
| TCP NULL Port Scan | HTTP Client campas attack |
| TCP SYN flood | HTTP client campus attack |
| TCP SYN scan | HTTP client cat%20seen |
| TCP XMAS scan | HTTP client cgimail.exe attack |
| All TFTP | Set email alerts on all buffer overflow attacks Win DoS attacks |

Start with what you know.

Place a diagram of your network in front of you and make a list of all the protocols that pass each aggregate point in your network. For example, let's assume that you are deploying SNP Sensor #4. This is the sensor that monitors the access layer of your network.

Begin by listing all the protocols that you know you use. Do you use windows browsing? Do you allow Internet web access? FTP? Mail? Keep brainstorming until you run out of ideas. Write them down as you do so. Now list the Operating Systems (OS's) that you know are on the network segment. Windows 95? NT? 2K? What flavors of Unix? Netware? Mac?

Start with one IDS or Sniffer

Start up some sort of sniffer on this network segment. Capture a trace file over about an hour of normal work time and look at the traffic. What protocols are you seeing? Add these to your list of protocols.

Then, you need to look at what you want to know that may not be covered by what is already in your network. You may not be running IIS servers, but still want to know that someone is trying IIS exploits in your DMZ – as this could be a precursor to other web server exploits.

Add the results of your existing network with what other items you want to know and this composite list becomes the list of potential vulnerabilities or hacking alerts. Looking at the total database of signatures by both classification and priority, then determine if signatures are required for the deployments you have planned.

Practicum

We have a network that provides Internet web access (TCP destination port 80 and 443 for HTTPS), FTP (20 and 21 TCP), DNS queries (UDP 53) from the inside out. Our inside machines are a mix of Windows 95/98/NT/2K machines. These machines mostly use IE5X as a browser and Outlook as the mail client. We provide mail services via a Solaris 2.6 Sendmail box in our DMZ, Anonymous FTP read only via wu-ftp on Redhat Linux 6.2 and we serve up a modest website via Win2k/ IIS5. DNS is provided for us via our ISP. You have sniffed your net and find that some internal tftp is happening between your routers; and telnet is being used as well to manage machines remotely. Because of your windows machines, you also see a ton of network traffic from the Server Message Block (SMB) suite on the inside.

Step 1:

Write down the protocol services that you access from inside to out. We have a list containing HTTP, HTTPS, FTP and DNS. Now add to that the protocols that you serve up to the net- SMTP, FTP and HTTP. Add the internal SMB's. Finally, list the ones you found by sniffing: telnet and tftp.

Step 2:

Now match the protocols served with the OS's and Applications that you have active in you network. In this case, the OS's are Solaris 2.6 , Redhat 6.2 and the Windows cousins. The apps are IIS5, wu-ftp, sendmail, MSOutlook, IE5x.

Step 3:

Now go through the active signature database, expanding the list to see the gross level granularity of Electronic Mail, File Transfer, Misc and Web.

Step 4:

Factor Budget. Can you afford one, two, three, etc Sensors? If only one, this step is easy. As you navigate down the active module list, right click each one and read the description. Does it mention any of the protocols or applications that you listed above? If so, enable it on your single sensor. You must decide if this single sensor should be deployed in position #1 or position #3a. Keep in mind when you decide this that placing it in position #1 allows for no visibility into your DMZ. Don't forget to save the module configuration after you have done all this work on the sensor configuration.

It is assumed that the sole sensor is located in position #3a. It is also assumes that no specific protocols have been expressly forbidden by company policy and that said policy also does not specify all protocols and services not specifically authorized are unauthorized. Should you encounter such a policy, enable all of the signature modules. In fact, it is strongly suggested that you run the SecureNet Pro with all signatures enabled for at least a few days. Take the time to weed through it all, and then tweak it down to what you want to see.

There is also an example of what you can do with 2 SNP sensors in table 2. It splits the load between the sensors- mostly looking for windows exploitation within the local net and concentrating on watching services in the DMZ. Should budget allow 3 to 5 sensors, you can extrapolate the configurations in tables 1 and 2.



Step 5:

To filter IP addresses, specify the IP ranges of the systems being monitored in the sensor's global filters found in the filter.cfg file.

To define the packet reassembly order, adjust IP ranges or addresses of the systems being monitored by specifying "old" or "new" in the hosts.cfg file.

To define ranges of systems being monitored with stateful inspection, use the ranges.cfg file.

Details for making these changes are in the SecureNet Pro Users Guide. Again, it is wise to start out observing everything, focusing on detail only after a holistic view is understood.

Step 6:

Wean out the false positives. No matter what you do, you will run up against some false positives. This is most common when an application harmlessly violates some part of an RFC that can be exploited or one of your own machines is relaying mail or doing ping polling of a network.

Other white papers available from www.intrusion.com include:
Maximizing the Value of Network Intrusion Detection
Why Firewalls Are Not Enough
Applying Network Intrusion Detection Under HIPAA
And more...



1101 East Arapaho Road
Richardson, TX 75081

(888) 637-7770
(972) 301-3607

www.intrusion.com