



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 Systems:
Theory, Reality, and the Future**

By

Patrick J. Terry

**GIAC Security Essentials Certification
Practical Assignment
Version 1.4 (Amended April 8, 2002)**

© SANS Institute 2000 - 2002, Author retains full rights.

Overview: Can I afford not to have an Intrusion Detection System?

A CIO somewhere (ok, what the heck, lets say a beach in the Bahamas) is reading his magazine of interest, and sees - Intrusion Detection - I see that it's a good thing, but I hear that it costs a lot of money to do. Why do I really need it, nobody wants to attack us! Then reality sets in, as the CIO recalls all of the events that have been on the news and in startling rapid proliferation. Trojan Horses such as AOL4FREE, discovered April 16, 1997; Worms such as CodeRed, discovered July 16, 2001; Mass mailing Viruses such as Nimda discovered Sept 18, 2001; Distributed denial of service attacks, such as the one that put British Internet Service Provider CloudNine out of business¹, website defacements, identity theft at both a personal and corporate level!! It's as if there must be a group of people that get enjoyment out of other peoples pain and discomfort. Are those the Blackhats he has heard about?

Now the CIO is concerned, he wants an explanation of what IDS is, how it came about, and how to implement the processes and procedures that will be necessary to integrate this technology into his organization. Our CIO wants some real world experience of how this technology delivers, and he wants to understand if he should implement the current technology, or wait to see the new technology that he has read about can deliver. What can he do to ensure he has protected his assets from these Blackhats? Not the ones that are fairly benign, and find the different failures and vulnerability of network operating systems, applications, databases, etc for the excitement of the discovery, and the recognition of their talent (lets call them the Greyhats), but the others that are not so kind, and take advantage of their knowledge, often for personal gain. These are the ones that we need IDS for, and this paper will focus not only on the theory of IDS, but will also examine one Company's Security Engineers personal experiences with several leading vendors equipment in a complex outsourcing environment. The paper will conclude with some recommendations, and thoughts about the future.

Theory: Is this good stuff, or What !?!

As we begin this discussion of Intrusion detection, let us start with some simple definitions:

in·tru·sion [Pronunciation Key](#) (Īn-trŭ'zhən)
n.

1. The act of intruding or the condition of being intruded on.
2. An inappropriate or unwelcome addition.
3. Law. Illegal entry upon or appropriation of the property of another.
4. Geology.²

² Editors, The American Heritage® Dictionary of the English Language, Fourth Edition.

de·tec·tion [Pronunciation Key](#) (dĭ-tĕk'shən)
n.

1. The act or process of detecting; discovery: *detection of a crime*; *detection of radiation from a distant galaxy*.
2.²

² Ibid

Although the primary focus of this paper is Intrusion Detection systems as applied to the information age, it is important to note that humanity has had to deal with unwelcome entry into protected assets for a very long time. Early Intrusion Detection systems could be something as simple as a guard-dog or as complicated as a roving, random security person. As intruders became more wary and sophisticated about evading detection, systems had to improve, leading to the existence of physical security companies such as ADT, which has provided protection services against intrusion since the late 1800s.

Intrusion Detection systems in the information age have been around in one form or another since the early 1980s. With the advent of non-proprietary access to network resources, security issues immediately arose. One of the early papers that attempted to create a structure, that appears to have since been implemented by different vendors was “An Intrusion-Detection Model.”³ The paper defines in excellent detail how this model can be designed and implemented, but more importantly it defines why a real-time intrusion-detection system is needed:

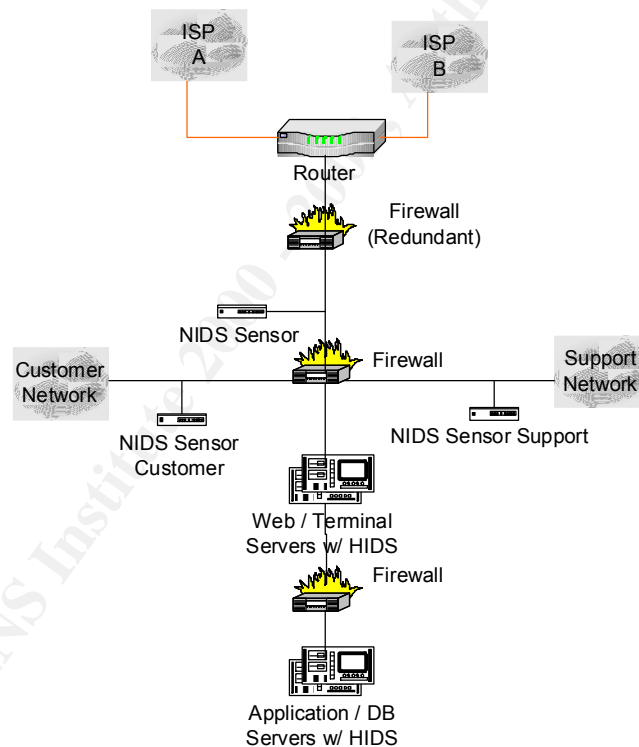
- Most existing systems have security flaws
- Existing systems with known flaws cannot easily be replaced
- Developing systems that are absolutely secure is very difficult
- Even the most secure systems are vulnerable by insiders

What is Intrusion Detection? Our reference that we refer to does a superb job of detailed explanation.⁴ Basically, Intrusion Detection attempts to monitor, alert and / or prevent unauthorized access to, and compromise of resources critical to your business or organization. These resources include not just the Corporate Inter- and Intra-net presences themselves, but the financial, health related or legal information that can be found through those presences. These resources need to be protected from both external and internal threats. Although significant amount of abuse of resources by trusted insiders does occur, often the external threat makes use of social engineering techniques⁵, such as the cycle described below:

- Information gathering phase
- Development of relationship
- Exploitation of relationship
- Execution to achieve objective

No piece of hardware, or any software yet created can prevent this type of attack from occurring, only good education of the companies' employees. Many people fall into these traps, meaning to be helpful and do well.

So, what can we do to protect our assets as well as we can? Create a Network Security Perimeter.⁶ While routers can be configured with Access Control Lists, and our Firewalls can have rule sets to restrict where incoming and outgoing traffic can travel, Intrusion Detection is what lets us know something has occurred that we should investigate in more detail. As most know, Intrusion Detection is broken into different categories: Network based and Host Based. A sample diagram explains the relationship:



This simplistic diagram created by the author shows the customers data being protected by network based IDS at the critical junctures, the Internet, the customers' own network, and the outsourcing company's support network. Networks based Intrusion Detection systems; such as those from various vendors and including open source products such as Snort operate by being located on the network near the system or systems being protected. They observe the network traffic by utilizing Network Interface Cards (NICs) operating in promiscuous mode, analyzing packets for signature matches, similar to what popular and well known virus protection software does. Unfortunately, the NIDS

signatures must be updated on a periodic basis, just as the virus protection software must be.

Security events can come from any or all of these locations shown above, and it is vital that we are aware of this, especially in the early alert phase of the security event. Network based IDS can be configured to log the information and create an alert, disconnect the offending connection via TCP reset, or modify access lists to prevent connections from the offender, either for a specific period, or permanently. This last method is very infrequently used, as it can create an unattended denial of service upon your environment without any human based intervention or knowledge.

Current host based IDS can in itself be broken down into two types: Network monitoring, which is similar to NIDS, but concentrates only on traffic entering the system, and host monitoring, which watches the file system structure, logs, and other 'static' information. These host monitoring systems require a type of cryptographic checksum be created for the information, and will alert if this checksum is modified or changed. Obviously this checksum itself needs to be well protected, or the offender can manipulate the checksum to make the modification invisible.

There is a new class of host based IDS coming into the picture, that utilizes proactive system protection by intercepting host system calls for system resources. This new technology will be discussed later in this paper.

To conclude the comparison of NIDS vs. HIDS, a simple analogy can be proposed: that the NIDS notices someone frequently walking by the candy store, and the HIDS notices that someone has now come into the store, and grabbed the candy. Utilizing both, and following proactive security approaches help ensure that your candy stays in the jar.

Reality: Living with two NIDS

We are a large globally located Outsourcing business, with over 1500 open systems and 60 large mainframe systems in just one data center, and as such have a very complex infrastructure. In this data center we support internal company systems, as well as our outsourced partners equipment. Our initial network security perimeter was similar to any normal center, with primary isolation of our customers systems from our internal systems with firewalls and DMZ partitioning. As we grew, and as threats became increasing active in the Internet world, our customers and we began incorporating Intrusion Detection capabilities into our organization.

Our initial product was Internet Security Systems RealSecure Network IDS. It has one of the largest installation bases, and was well known in the industry, and that was the primary selection criteria at the time. We have upgraded several times

throughout its evolution. RealSecure consists of the Management console, and one or more remote sensors, configured with promiscuous mode NICs. In our environment we also utilize taps and 10/100 hubs, which allow the bi-directional analysis of network data. We connect the tap into a the hub, and then connect the NIC into the hub, theoretically making the IDS system invisible, since the sensor itself does not appear to have a valid TCP/IP address. Of course, the traffic that is analyzed does have to be reported back to the management station, and so there is a secondary NIC in the sensor that is on a security segment that is isolated to the console and sensors. Both the management stations, and the sensors are based on Windows NT 4.0 Intel platforms, with our more recent systems utilizing Windows 2000. The latest versions utilize a MSDE database, and are more robust in the logging and reporting of data. The install of the software is very straight forward, and initial configuration can be completed quickly. The most time consuming part is the elimination of false positives. This is an essential part of the process, and is a tuning step that is continually occurring. RealSecure, like other products base the signatures that they examine as low, medium, or high. Signatures are classified in these categories as shipped from Internet Security System Inc, and can be modified individually. This begins the process of false positive elimination and network entity awareness.

What is a false positive? Simply put, an event occurs that matches the signature set of the IDS, which may not actually be the real attack. This triggers the system to issue an alert inappropriately. But false positives are not the only issue in becoming network entity aware. Not all signatures are necessarily as important in one network in comparison to another. For example, if the network has no systems running Microsoft Internet Information Systems product, do you care if you detect attacks based on vulnerabilities based on the product? These signatures can be modified to be low, with no notification other than possibly logging.

In our organization, we have determined that we require notification of any signature classified as "High" by immediate paging of our network IDS security group, the immediate on call network engineer, as well as automated creation of trouble tickets in our Trouble-ticket system. Our helpdesk staff, insuring that the network security group is aware of the issue, and stays involved throughout the process that we follow, owns these tickets.

These pages and Trouble-ticket system cases have the following format:

"Customer-abbreviation" IDS Sensor: "Sensor-id" SIGTYPE: "Intrusion event, either type by name for RealSecure, or signature numeric ID for CSPM" SRCIP:" Source IP address" DSTIP: "Destination IP address", SRCPORT: "Source port number", DSTPORT: "Destination port"

So a typical page could be:

ABC IDS Sensor: ABC001 SIGTYPE: HTTP_POST_FILENAME_PASSWD
SRCIP: 111.111.111.111 DSTIP: 222.222.222.222 SRCPORT: 2191 DSTPORT:
80

By readily having this information available, it decreases reaction time to security events immensely. During a security audit that was recently performed, this notification method was considered unique and valuable.

Obviously, if we did not eliminate false positives as much as possible, this would quickly cause gridlock of people, pagers, and would be unsupportable. Any signatures classified as "Medium" are issued as Trouble-ticket system tickets, and are investigated the new business day. "Low" signatures are logged for forensic purposes and investigated as needed. The ongoing analysis and investigation are a critical part of the false positive elimination cycle, as well as an interesting way of discovering what activities are actually occurring in a network. In one customer network, it was discovered that known, but unauthorized tools were being utilized, and when the RealSecure product detected a signature indicating the use, and the specific source and destination IP address of this tool, the customer had to address the issue (and we had to create a new closure code for our Trouble-ticket system!).

Reporting of events that have been observed, investigated, classified, and resolved in the Intrusion Detection systems is an important part of the process. In our environment, we typically provide customers with a monthly report of all events, and resolution of each. Some of the normal codes that we classify closure of events are:

- IDS – Misconfiguration
- IDS – False Positive
- IDS – Penetration test
- IDS – User generated event
- IDS – Potential Intrusion
- IDS – Actual Intrusion

The detailed report that customer receives includes graphing of the mixture of events, and closures, as well as a indication of date, time and other specific to customer requirements for their use and potentially to customers that they provide support for.

Script configuration and implementation in RealSecure is a time-consuming process. The initial policy configuration of what scripts are available is to identify to the management console what the script is, where it is located - in the RealSecure product, the script must be on every sensor managed, and finally what signatures will activate the script or scripts involved. This final step involves examining every signature, and based on priority or need to know checking the

check box in the GUI interface for that signature to activate the script. When the hundreds of signatures are considered the effort to do this as a part of not only an implementation, but an ongoing support issue, as new signatures are released quickly becomes an issue. Network engineers become swamped with the effort, especially as the technology becomes more prevalent in customers networks. We currently have multiple RealSecure instances, from a simple two sensor, one management console up to a 9 sensor (deployed worldwide), one management console environment.

Our other product came to us by a customer requirement. This customer is a established healthcare provider, and is subject to HIPAA requirements, so security of their information is critical. They were utilizing the Cisco CSPM (Cisco Secure Policy Manager) with NetRanger sensor appliances. The CSPM console operates on a Windows NT 4.0 based Intel platform, and the NetRanger sensor operates on a Solaris x86 Intel platform. We currently have the 4210 version, which is 45 Mbps capable, and the 4230 version which is 100 Mbps capable. Both models have recently gone end of life support, and we have begun the upgrade process. As RealSecure, this is a signature-based product, and has faults similar in nature. The configuration and updating of signatures is still a time consuming effort. Our script notification however, is much easier to implement, since the console directs the script activation, rather than the sensors, as is the case of RealSecure. Activating scripts is by signature classification (High, Medium, Low) rather than by each signature, regardless of classification, so only two steps are required to activate all High and Medium signatures. When updating signatures is required, no further work is necessary to ensure that our notification of events is necessary, unlike RealSecure.

Unfortunately, we experienced significant amounts of trouble keeping the system up. From the initial classroom experience of the 2.3.3i version into actual production, the product crashed repeatedly, especially when the logging database would fail. This database appeared to be a collection of flat files, without a true database engine, such as that used in the RealSecure product. Only backing up and cleanup of the database allowed the system to not crash, but it was a painful experience in the beginning, as signing onto the application would no longer work, and our only option was to re-install the application, reload topology information of the affected network, and accept the data loss. Luckily, most of our pains were eliminated in our setup phase, and loss of production data has been minimum. We currently have multiple instances of the Cisco product, running combinations of the 4210 and 4230 sensors. We have had good success in support from Cisco, but their new products need to deliver.

Are we alone in this pain? Evidently not, as discussed in one recent article:

Because no product distinguished itself, we are not naming a winner (See "No cigar"). The eight products we tested - from Cisco, Intrusion, Lancope, Network Flight Recorder (NFR), Nokia (running on OEM version of

Internet Security Systems RealSecure 6.5), OneSecure, Recourse Technologies and the open-source Snort package - all ask too much of their users in terms of time and expertise to be described as security must-haves.⁷

⁷ (Newman, p. 2)

Does this mean that current IDS is worthless and 'dead'? Some sources⁸ think that is, and point out the following flaws:

- Dependency on shared network segments, current switched environments cause complications in deployment of IDS sensors.
- The need for higher network speeds increases the volume of traffic that the IDS must scan and analyze, both legitimate, and various Denial of Service techniques.
- Known fragmentation and reassembly techniques.
- The dependency on the signature base itself, being only as good as the creator of the signatures. How does an operator of one products ensure that their vendor is better than another?
- Last, but not least is the growth of encryption of data. Once encrypted, the IDS no longer can match signatures to the data packet content, so is allowed to pass.

Will these flaws cause organizations to walk away from these products? No, for two basic reasons: legally, the organization is expected to practice it's best effort in protecting it's own and its customers resources, and ethically is it reasonable to ignore a technology because it is flawed? As mentioned earlier, Intrusion Detection has been around a long time, have we quit posting security guards, dogs, and using other physical techniques? Those techniques have been bypassed on numerous occasions, but we continue to see them, both in the private and public sector, especially after recent terrorist actions.

An interesting sideline in the development of Intrusion Detection seems to show that the open-source products, such as Snort often appear to be quicker to react to events and have signatures released faster than the proprietary vendors. The acceptance of these open-source products would seem to be logical, but most organizations tend to stay away from open source, feeling that having a proprietary product is safer, more supportable, and better 'imaged'. After all, our CIO at the beginning of this discussion might feel better telling his customers that Cisco (or any of the recognized brands) is protecting their data, rather than something called Snort!

One last comment that must be made to any security conscious group – If you don't have an internal vulnerability-scanning group – Build one! – For a minimal cost, a dual-boot W2K/ Linux system can be deployed with all of the free / shareware scanning tools. These systems are critical in determining where your

organization is weak, and where best to prevent future conflicts. Don't forget to have the management chain above buy in to the testing, and become familiar with the tools slowly. For an excellent reference to building and using these tools Chapter 6 of the GSEC Security essentials⁹ is highly recommended.

We have discussed in some respect and detail physical Intrusion Detection, as well as network and host based electronic Intrusion Detection. Briefly, since this is a topic and paper of its own, lets consider a new technology that has been seeing lots of attention: Wireless access to network resources, especially the 802.11b version. Few policies exist, and fewer resources are made available to identify and resolve issues of this type. Since the wireless access is 'invisible', how many organizations actually know what is occurring in their environments? We were quite surprised when an informal 'war-walk' of our building identified 4 different 802.11b wireless access points, two of which allowed DHCP addressing and limited internal access, but full Internet capabilities. Not even the minimum protection of WEP was utilized. Our security team located the access points, and educated the users as to the ramifications, and has since found no other WAP in use, but how many organizations out there even check? Even if the segments that the access points had been scanned by internally placed IDS sensors, the traffic would have appeared to have been generated by legitimate employees.

The Future: Blackhats vs Whitehats, what's next??

As with most technologies, as flaws are uncovered, solutions for these flaws are created. Intrusion Detection is becoming re-created as Intrusion Prevention.¹¹ These new technologies work by various means, such as intercepting application interface calls to operating systems and classifying the calling activity. If the Intrusion Prevention system thinks that the caller is inappropriate, the access can be denied, allowed, logged, or a combination of such. These systems are being crafted for both HIDS and NIDS type environments, and show that vendors are listening to the needs of the security community.

Nokia, Cisco and other groups are incorporating their technologies into blade type devices / appliances as they attempt to cut down on the total cost of ownership issues that all organizations have. Some of the problems that will cause the acceptance of these new technologies to be slower than preferred are the uncomfortable experiences with the current products that organizations have endured, the cost of these current products that organizations have already absorbed, and of course the cost of the new technology itself.

True, the Blackhats are out there, creating havoc for their own pleasure, but the Whitehat community is becoming more aware of the fact that they are not out there on their own. The efforts of SANS and other key security concerns have opened eyes, and have let the Bad guys know that we are watching them, and are working towards technologies that will not only let us be aware of what they are doing, prevent what they are doing, and ultimately catch them and make

them accountable for their actions. They may have a head start, but as they become less anonymous, the fun and excitement that they live for will disappear, replaced by dread, apprehension and fear.

© SANS Institute 2000 - 2002, Author retains full rights.

References:

1. Warner, Bernard "Internet firm hacked out of business" Reuters Feb. 1, 2002 URL: <http://techupdate.zdnet.com/techupdate/stories/main/0,14179,2844881,00.html>
2. Online Dictionary, information from The American Heritage Dictionary of the English Language, Fourth Edition. URL's: <http://www.dictionary.com/search?q=intrusion>
<http://www.dictionary.com/search?q=detection>
3. Denning, Dorothy E. "IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, VOL. SE-13, NO. 2, FEBRUARY 1987, 222-232 URL: <http://www.cs.georgetown.edu/~denning/infosec/ids-model.rtf>
4. Elson, David "Intrusion Detection, Theory and Practice" Security Focus Mar. 27, 2000 URL: <http://online.securityfocus.com/infocus/1203>
5. Mogull, Rich "Employees Are Unwitting Victims of Social Engineering" Gartner URL: <http://security1.gartner.com/story.php.id.39.s.1.jsp>
6. Ellis, Chris "NSP Special Report" Network World Oct. 10, 2002
7. Newman, David; Synder, Joel; Thayer, Rodney "Crying wolf: False alarms hide attacks" Network World Fusion June 24, 2002 URL: <http://www.nwfusion.com/techinsider/2002/0624security1.html>
8. McClure, Stuart; Scambray, Joel "Once-promising intrusion detection systems stumble over a myriad of problems" Infoworld URL: <http://www.infoworld.com/articles/op/xml/00/12/11/001211opswatch.xml>
9. Cole, Eric; Newfield, Mathew; Millican, John M. GSEC Security Essentials Toolkit Indianapolis: Que Publishing, March 2002 125 -149