



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>

SANS GSEC Practical Assignment

Migdalia Salazar

San Antonio SANS - March 11, 2002 – March 16, 2002

GSEC

Version 1.1

Title: Intrusion Detection System: First line of Defense

Submitted 29 April 2002

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

TABLE OF CONTENTS

Abstract.....	2
Introduction.....	2
Enterprise Computing Environment.....	3
IDS -- A companion to the Firewall".....	4
Commonly Used Monitoring Systems.....	5
Network-based IDS.....	5
Host-based IDS.....	6
Some Known Intrusion Signatures.....	7
IDS Security Policy.....	8
Defense in Depth.....	9
Conclusion.....	9
Reference.....	11

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

Abstract

Intrusion Detection Systems (IDSs) have become increasingly popular security alarm systems over the past few years as more coordinated attacks on networks become a game of cat and mice. Any enterprise, small or large, should consider an IDS component as a first line of defense protecting the computer systems and networks from any attacks. These attacks on systems can come from insiders as well as outsiders, and these attacks can occur at anytime so enterprises are facing a lot of challenges today. This is compounded by the proliferation of Internet users and all of the networks they might be connected to. In order to protect the overall network infrastructure, an alarm should be placed in a strategic place. The two most common places are at the host and the network itself. There are two types of IDSs: network-based IDSs and host-based IDSs. These IDSs monitor network traffic and trigger alarms when your network is under attack. When an IDS is properly deployed it can provide warnings indicating that a system is under attack, even if the system is not vulnerable to the specific attack. It is an additional security component that best complements a firewall extending the security capabilities as well as protecting your assets.

Introduction

"An Intrusion Detection System (IDS) is a system that tries to detect and alert on attempted intrusions into a system or network, where an intrusion is considered to be any unauthorized or unwanted activity on that system or network" [1]. The IDS is like the new kid on the block regarding information security in both the government and commercial sectors. An IDS offers a process that can be used to monitor events occurring in a computer system or network that attempt to compromise the confidentiality, integrity, and availability of a computer system as well as the overall network. An IDS is an increasingly important segment in today's enterprise computing environment. Intrusion Detection Systems have gained acceptance as necessary additions to every organization's security infrastructure. The most common types of IDSs are network-based IDSs and host-based IDSs. These IDSs monitor network traffic and trigger alarms when your network is under attack. When an IDS is properly deployed it can provide warnings indicating that a system is under attack, even if the system is not vulnerable to the specific attack. These warnings can help users alter their installation's defensive posture to increase resistance to attack. It is an additional security system that complements a firewall.

The amount of work you might need to do will depend on the organization's perspective on information security. Installing an IDS is an important segment in the overall security architecture of organizations, especially in major organizations that are connected via the Internet. Installing and effectively using IDSs on networks and hosts requires a broad understanding of computer security. In today's computing environment, there will be a need for more servers, more connectivity, and more threats; therefore, IDSs can play a significant role in the overall security architecture of any organization

Enterprise Computing Environment

An IDS is needed in today's computing environment because it is impossible to keep pace with current and potential threats and vulnerabilities in our computing systems. The environment is constantly evolving and changing, fueled by new technology and the Internet. To make matters worse, threats and vulnerabilities in this environment are also constantly evolving and IDS products are viable tools to assist in managing threats and vulnerabilities in this changing environment.

Before an organization invests in an IDS, it is important to first understand which assets require protection and determine the real threats against those assets. If an organization chooses to deploy an IDS, a range of commercial and public domain products are available that offer varying deployment costs. Figure 1 illustrates the system structure of an IDS protected enterprise.




Figure 1. An IDS Protected Enterprise[2].

This architecture is an example of a small enterprise configured to isolate the Web server from the firewall. The network, host, and application sensors are configured in the computers to extract suspicious packets from the three main network segments known as the Network sensors. The packets are then forwarded to a network-specific analysis station which is the Analyzer. This particular layout shows some aspects of securing an enterprise network using a network sensor outside the protected network. The network sensors allow the administrator to sense the general threat level as indicated by probes

and the attempts that will be blocked by the outer firewall. The analyzer serves as the IDS's user interface and reports the activities to a management console. The management console alerts the system administrator/enterprise administration for further action.

IDS – A companion to the “Firewall”

An IDS is considered by many to be the logical complement to network firewalls, thus extending the security management capabilities of system administrators to include security audit, monitoring, attack recognition and response. Firewalls examine and control the flow of information and services between a protected subnetwork and/or hosts and the outside world. They protect one network from another by blocking specific traffic while allowing other traffic. Most people think of the firewall as their first line of defense, however there are intruders/hackers that can figure out how to bypass one.

Firewalls are designed to lock out as many intruders as possible, but they are still not enough. The decision as to which traffic to allow is based upon the content of the traffic itself. Typical decision criteria include traffic direction, a network address, a port, protocol type, and service type. One study found that the goal of the firewall is to provide efficient and authorized access for users "inside" the firewall to the outside world while controlling the access of "outside" users to protected resources by exporting limited and precisely controlled services[3]. Firewalls are best implemented on separate hardware for performance and security reasons, and thus there is the expense of acquisition and maintenance.

As every firewall administrator knows, the system is under almost constant attack from the Internet. Hackers from all over the world are constantly probing the system for weaknesses. There are a variety of different probes hackers will attempt to use, and port scanning is very popular among hackers. Port scanning occurs when an individual attempts to connect to a variety of different ports, usually from ports 1 to 1024, but there are over 65000 ports to choose from. The most popular ports/protocols that are susceptible to attacks from the Internet are: 21 (ftp), 23 (telnet), 25 (SMTP), 53 (DNS), 79(finger), 80 (HTTP), and 143 (IMAP). These are some of the protocols that should be identified as part of the basic security policy in your IDS configuration. Scans can be used on a specific target, or used to scan entire IP ranges, often chosen at random. Moreover, every few months a new security vulnerability is found in popular products, at which point the hackers simply scan the entire Internet looking for that port of entry.

While some log file analysis programs do scan firewall logs for signs of intrusions, most intrusion detection systems get their information elsewhere. Remember that firewalls are simple rule-based systems that allow/deny traffic going through them. There are some robust firewalls which are designed to go beyond that, but even those do not have the capability to clearly say whether the traffic constitutes an attack; they only determine whether it matches their rules or not. Having a firewall is not the dynamic defensive system that users imagine it to be whereas an IDS is much more of a dynamic system.

An IDS recognizes attacks against the network that firewalls are unable to see. For example, if a number of sites all had firewalls that restricted access only to the web server at port 80, the web server could still be hacked. Thus, the firewall would provide no defense. On the other hand, an IDS system would have discovered the attack, because it matched the signature configured in the system. Although firewalls protect against external access attempts, they can leave the network unprotected as well as vulnerable from internal intrusions if not configured for inside protection.

Commonly Used Monitoring Systems

There are several types of IDSs available today, characterized by different monitoring and analyses approaches. Some IDSs analyze network packets captured from network backbones or LAN segments to detect attackers. Other IDSs analyze information sources generated by the operating system or application software for signs of intrusion. The two most commonly used in most network enterprise today are Host-based IDS and Network-Based IDS.

Like any networking or security products, they have both positive attributes and legitimate shortcomings.

Network-based IDS: Network-based systems examine the individual packets flowing through a network. Unlike firewalls, which typically only look at IP addresses, ports and ICMP (Internet Control Message Protocol) types, network based intrusion detection systems ("NIDS") are able to understand all the different flags and options that can exist within a network packet [4]. A NIDS can therefore detect maliciously crafted packets that are sometimes overlooked by a firewall's relatively simplistic filtering rules. A network-based IDS sits on the LAN (or a LAN segment) and monitors network traffic packet by packet in real time, to see if that traffic conforms to predetermined attack signatures. An advantage of a network-based monitoring system is that it has the benefit of seeing and coordinating attacks that are occurring across your entire network, and it only examines the traffic from your network. There are advantages and disadvantages of a network-based IDS. Here are some of the advantages [5]:

- Network based IDS can monitor large networks
- Network-based IDSs has little impact on the performance of an existing network.
- Network-based IDSs are typically passive devices that listen on a network wire without interfering with normal network operation.

Here are some of the disadvantages:

- Network-based IDSs may have difficulty processing all packets in a large or busy network.
- As the networks become increasingly larger, it becomes more difficult to place a network-based IDS at a single location. It will then require the utilization of more sensors throughout your network, more configuration management and it will definitely increase the costs of maintaining the IDS on your network.

A NIDS tends to be more distributed than a host-based IDS. Instead of analyzing information that originates and resides on a computer, a network-based IDS uses techniques like "packet-sniffing" to pull data from TCP/IP or other protocol packets traveling along the network [6]. This surveillance of the connections between computers makes a network-based IDS great at detecting access attempts from outside the trusted network. In general, network-based systems are best at detecting the following activities [7]:

- Unauthorized outsider access: When an unauthorized user logs in successfully, or attempts to log in, they are best tracked with a host-based IDS. However, detecting the unauthorized user before their log on attempt is best accomplished with a network-based IDS.
- Bandwidth theft/denial of service: These attacks from outside the network single out network resources for abuse or overload. The packets that initiate/carry these attacks can best be noticed with the use of a network-based IDS.

Host-based IDS: The primary benefit of a host-based monitoring system is that it can detect both external and internal misuse, something that network monitors and firewalls can't do. It is the first type of IDS to be developed and implemented. This real time monitor operates on a system and receives application or operating system audit logs for evidence of malicious or suspicious application activity.

The appeal of such a tool is obvious, as security breaches are more likely to come from an internal user than from a hacker outside the network. Host agents are powerful tools for addressing the authorization and access issues that make internal security so complex. They are able to detect such things as repeated failed access attempts or changes to critical systems files. In some cases, an agent can halt an "attack" on a system, though a host agent's primary function is to log events and send alerts [8].

The basic concept of a host-based IDS is to analyze activities on your operating system but there are advantages as well as disadvantages of a host-based IDS. Here are some of the advantages [9]:

- Host-based IDSs can detect attacks that are not detectable by a network-based IDS because this type has a view of events that are local to a host.
- Host-based IDSs can operate in a network that is using encryption when the encrypted information is decrypted on (or before) reaching the host that is being monitored.
- Host-based IDSs can operate in switched networks.

Here are some of the disadvantages [10]:

- The collection mechanisms must usually be installed and maintained on every host that is to be monitored.
- Host-based IDS are not well suited for detecting network scans of all the hosts in a network.
- Host-based IDSs frequently have difficulty detecting and operating in the face of denial of service attacks.

Not all IDSs are created equal. Depending on the level of expertise within your organization, the ease of use and strong reporting as well as the ease of creating new attack signatures or price should be considered. It is very important to consider the performance of an IDS so that the IDS can consistently look at every IP packet and respond accordingly.

While IDS vendors feverishly try to brand their wares as intrusion prevention systems, users snicker that today's solutions are more like intrusion reaction systems—useful as a log of previous hacks, but not as a means of preventing attacks in progress [11]. Before investing in an IDS perform a study evaluating the products. There are so many different types of attacks and not all IDSs are configured to detect them all. It is also important to know the incidence of false-positives which occurs when the IDS generates an alarm from normal user activity. If your IDS generates too many false positives, then you will lose confidence in the capability of the IDS. These are just a few examples that should be investigated when evaluating IDS products.

Over 300 Infosecurity Professionals participated in a poll to find out about their experience using an IDS. Based on their experience, about two-thirds of these professionals said their IDS is “very important” to their security infrastructure. About 35 percent said it is “somewhat important” and about one-fifth of respondents said they were either “not at all confident” or “not very confident” that their IDS is protecting their mission-critical systems from cyberattack [12].

Running a network-based IDS and a host-based IDS simultaneously is strongly recommended. They both complement each other on core hosts such as the domain name server (DNS), Mail, Web, and other high value servers. As the size and value of the organization increases, the importance of detecting and protecting your network security infrastructure increases as well.

That is why it is so important to ask the vendors some tough questions before deciding on an IDS product that will fit your security infrastructure. Some questions to ask the vendors is whether the rules have to be updated, who provides them, and how much do these updates cost. Keeping up with the latest attacks can be so overwhelming that it's hard to manage, especially when there are so many false positives. Most information security professionals think that an IDS should be 100% reliable, but it's difficult to get to the point where you're savvy enough to determine what's a significant event and what isn't.

Some Known Intrusion Signatures

Attack signatures are activities that match known attack patterns. For example: TearDrop Denial of Service (DoS) attack sends packets that are fragmented in such a way as to crash the target system whereas the network monitor will recognize packets that conform to the TearDrop signature and take action [13].

The IDS vendor provides a database of attack signatures, and administrators can also add customized signatures. If the IDS recognizes an attack, it alerts an administrator. In some cases, the IDS can also respond, for example, by terminating a connection. In addition to its monitoring and alarm functions, the IDS also records attack sessions for later analysis. Network IDSs can also be linked to other security features, such as firewalls, to make sure those systems haven't been breached.

A network monitor has two main benefits. The first is the real-time nature of the alarm, which can give administrators an opportunity to halt or contain an attack before it does significant harm. This is especially valuable for DoS attacks, which must be dealt with immediately to mitigate damages.

The second benefit is evidence collection. Not only can administrators analyze the attack to determine what damage might have been done, the attack session itself can point out security flaws that need addressing [14]. Because many hackers first scan a target network for known vulnerabilities, a hacker's choice of attack may indicate that such vulnerabilities exist on your network. A simple example is an operating system that has yet to be secured with the latest vendor patch.

To date, commercial IDSs have concentrated mostly on string matching and other forms of signature identification to detect classes of outsider attacks. These techniques all rely on previously encountered attack signatures for their defense. While the easy part of the problem of signature identification seems to have been addressed by the commercial community, research advances in the community at large appear to have slowed, resulting in an increased emphasis on detecting known types of outsider attacks.

Unfortunately, as soon as one method of attack is taken care of, a new one starts to circumvent systems' defenses. Detecting, identifying, and responding to these unknown attacks continue to be a challenging problem. Providing global rather than local analysis is still a very important research area that is relatively uncharted.

IDS Security Policy

A security policy should be in place to help define and implement the actions an organization must take to protect its information. In formulating a security policy, the essential steps consist of implementing user authentication and access controls, eliminating unnecessary services, applying patches to eliminate known vulnerabilities, deploying firewalls, and using file integrity checking tools. Because most real-time commercial IDSs base their detection approach on known attempts to exploit known vulnerabilities, an administrator's time is often better spent minimizing vulnerabilities through the application of patches or other security measures.

An effective security policy should consider the following: operating systems, services (web servers, e-mail servers, and databases), network IDSs, firewalls, and the network management platform (such as Open View) [15]. IDSs should be included as part of the overall security policy of an organization because they enforce the security policy by

detecting prohibited traffic and/or activities. It plays an active role in the identification of incidents for which the security policy outlines specific responses.

Defense in Depth

When establishing an IDS, a "defense in depth" process concentrating on software, networks, and hardware becomes the key to a successful business [16]. It begins with the establishment of appropriate and effective security policies. An IDS is only one aspect of a layered defensive posture. Having effective policies in place helps ensure that the threats are minimized so that the critical assets are treated and protected, and that the managers and users are adequately trained in order to take an action when an intrusion is detected and identified. A good security policy puts an IDS in its proper perspective and context and, whenever possible, the policy should reflect the mission of the organization. Therefore, it should codify the rules governing enterprise operations as they are reflected in its information infrastructure.

"Defense in depth, and overkill paranoia, are your friends" [17]. Hackers are much more capable than you think; the more defenses you have, the better. And they still won't protect you from the determined hacker. They will, however, raise the bar on the determination needed by the hackers. An organization committed to defense in depth should be employing good information security practices. This will demonstrate that the benefit of an IDS is just good business sense because the investment is well worth the return you get.

"Military history teaches us to never rely on a single defensive line or technique. The firewall serves as an effective noise filter, stopping many attacks before they can enter our network. Within our internal net, the router or switch can be configured to watch for signs of intrusion or fraud. When a detect occurs, the switch can either block the session and then seal off the host, or can just send a silent alarm. We can improve our model further by adding the host-based layer of defense. Here we can detect the insider with a legitimate login (whether or not it is really theirs) accessing files he shouldn't. Toss in a couple more network-based-detection systems, including a few stealthy ones, and you have an architecture sufficient to counter the increasing threat" [18].

Conclusion

Network attacks have increased in numbers and severity over the past few years and the need for an Intrusion Detection System has increased. They have become necessary as additions to the security infrastructure of most organizations. Attacks can come from the outside as well as from the inside, which means that IDS need to be strategically located behind the firewall as well as inside the firewall. Known vulnerabilities and security gaps have gone unfixed and are due to poor security practices. Implementing an IDS should be a standard protection method on most enterprise networks today. An IDS should be definitely be included among your considerations when developing a security

policy for your organization. Protecting your network with an IDS is definitely your first line of an effective “Defense in Depth”.

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

REFERENCES

- [1] Cooper, Mark. "An Overview of Intrusion Detection Systems". Autumn 2000 Newsletter. URL: www.xinetica.co.uk/tech_explained/general/ids/wp_ids.htm (19 April 2002).
- [2] McHugh, John, Christie, Alan & Allen Julia. "Defending Yourself: The Role of Intrusion Detection Systems". 2000 October/November. URL: www.computer.org/software/so2000/pdf/s5042.pdf (16 April 2002).
- [3] Halme, Lawrence R. & Bauer, Kenneth R. "AINT Misbehaving: A Taxonomy of Anti-Intrusion Techniques". No date. URL: www.sans.org/network/resources/IDFAQ/aint.htm (21 April 2002).
- [4] Cooper, Mark. "An Overview of Intrusion Detection Systems". Autumn 2000 Newsletter. URL: www.xinetica.co.uk/tech_explained/general/ids/wp_ids.htm (19 April 2002).
- [5] Mell, Peter. "Acquiring and Deploying Intrusion Detection Systems". 5 November 1999. URL: <http://csrc.nist.gov/staff/mell/idsBulletin.pdf> (21 April 2002).
- [6] Bace, Rebecca and Mell, Peter. "Intrusion Detection Systems. NIST Special Publication on Intrusion Detection System". Special Publication 800-31. URL: <http://online.securityfocus.com/data/library/idsdraft.pdf> (21 April 2002).
- [7] Bace, Rebecca and Mell, Peter. "Intrusion Detection Systems. NIST Special Publication on Intrusion Detection System". Special Publication 800-31. URL: <http://online.securityfocus.com/data/library/idsdraft.pdf> (21 April 2002).
- [8] Conry-Murray, Andrew. "Intrusion Detection". 5 December 2000. URL: <http://www.networkmagazine.com/article/NMG20001130S0007> (15 April 2002).
- [9] Mell, Peter. "Acquiring and Deploying Intrusion Detection Systems". 5 November 1999. URL: <http://csrc.nist.gov/staff/mell/idsBulletin.pdf> (21 April 2002).
- [10] Mell, Peter. "Acquiring and Deploying Intrusion Detection Systems". 5 November 1999. URL: <http://csrc.nist.gov/staff/mell/idsBulletin.pdf> (21 April 2002).
- [11] Briney, Andy. "New Directions In Intrusion Detection." Information Security Magazine, August 2001: 49-50.

[12] Briney, Andy. "New Directions In Intrusion Detection." Information Security Magazine, August 2001: 49-50.

[13] Conry-Murray, Andrew. "Intrusion Detection". 5 December 2000.
URL: <http://www.networkmagazine.com/article/NMG20001130S0007> (15 April 2002).

[14] Conry-Murray, Andrew. "Intrusion Detection". 5 December 2000.
URL: <http://www.networkmagazine.com/article/NMG20001130S0007> (15 April 2002).

[15] Wiens, Richard. "Realistic Expectations for Intrusion Detection Systems". 19 March 2001. www.rivalpath.com/kb/greyarts/docs/985039995:3051.html (16 April 2002).

[16] Manderscheid, Scott. "An Intrusion Detection System Process: Defense in Depth". 2 February 2001. URL: www.rr.sans.org/intrusion/process.php (16 April 2002).

[17] Graham, Robert. "FAQ: Network Intrusion Detection Systems". 21 March 2000.
URL: www.robertgrahm.com/pubs/network-intrusion-detection.html (16 April 2002).

[18] Northcutt, Stephen and Novak, Judy. Network Intrusion Detection – An Analyst's Handbook. New Riders Publishing, 2001. p. 214.

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