



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>

Abstract

With hacker attacks against well-known businesses and organizations on the rise, network security has made headlines. Of course, there are many attacks that do not make headlines and are not reported due to a loss of credibility or embarrassment. Then there are the attacks that are not even detected. The Defense Information Services Agency (DISA) states that up to 98% of attacks go unnoticed. These revelations have caused many businesses to rethink or to start thinking about the security of their own networks. For some organizations security has always been a concern, for these organizations they were ahead of the game and already had a basic security policy in place. Security of a network cannot be trusted to just one method of security, it must consist of many layers of security measures. These security measures may consist of, strong passwords, screening routers, firewalls, proxy servers, and intrusion detection systems. This paper will cover the last security measure mentioned, intrusion detection systems, also known as IDSs. An IDS is an important part of modern network security. Intrusion detection is the monitoring of a computer network with the goal of detecting an attack. IDSs do this well, but must be used in conjunction with other security measures. There are two major types of IDSs, Host and Network-based. Network-based IDSs monitor the whole network or a segment of the network, while host-based systems monitor a particular computer. If an IDS is to be installed it must be deployed with the network type and topology in mind. If not deployed correctly it may be ineffective, thus giving management a false sense of security.

Intrusion detection is a security technology that attempts to identify intrusions against a computer network. An intrusion is an unauthorized usage of or misuse of a computer system. In order to discover these intrusions a network administrator can employ an intrusion detection system (IDS). IDSs are available in two major types network-based and host-based. These types are available by many manufacturers and education facilities. Currently there are about 88 different IDS products; this paper will generally discuss generic functions and features. How these IDSs work and their deployment vary greatly with each major type.

The first type that will be covered is the network-based IDS. This type of IDS uses the raw network packets as the source for its' data. The IDS then analyzes the data looking for an attack signature. An attack signature is a known pattern in the packet or packets that match a model of a possible attack. This analyzing or recognition is done in real time, usually using the IDS's attack recognition module. There are four techniques that an IDS uses to recognize an attack (ISS, 1998). These four ways are:

- Pattern, expression or byte-code matching
- Frequency or threshold crossing
- Correlation of lesser events
- Statistical anomaly detection

The most used of the four techniques is the first one, pattern, expression or byte-code matching. In some writings this is also called signature analysis or misuse detection. Whatever it is called the IDS is programmed to interpret a certain series of packets or the payload of data in the packet as an attack. The IDS looks for a substring within a stream of data that is carried by network packets. When the IDS finds a substring that matches, it identifies it as an attack. This illustrates why it is very important to keep your attack signatures up to date.

Once a possible attack is discovered, the IDS will take action. The action that it takes is totally up to the capability of the IDS and the management's policy. Capabilities include; sending an alert to the console, logging the event, sending an E-mail, initiating a connection kill (TCP reset), reconfiguring a firewall or router, or using a SNMP trap. If the IDS were newly installed, it would be wise to only have it log events as they happen and analyze them for false positives.

A false positive is when an IDS finds a signature match that in reality is not an attack. False positives will occur at a surprisingly high rate when an IDS is first installed. A premium network-based IDS will have an extensive help library that will assist you in identifying these false positives.

In most network-based IDSs an agent, sensor, or engine is placed on a segment of a network that is to be monitored. This agent then sends back the semi-analyzed data to the console. This console is a centrally located computer that monitors all the agents. The data that is exchanged between agent and console should be encrypted to prevent an attacker from intercepting vulnerabilities of the network or disabling the agent by pretending to be the console. ISS's RealSecure uses RSA's public/private key for this encryption.

There are some system requirements for computers that are going to be used as an IDS. A network-based IDS console and agent uses a computer, a software package, and a

network interface card to monitor the network or segment of the network. The console computer needs to have a fast processor in order to keep up with the data being sent to it from the agents. This computer also needs a hard drive capable of not only holding the IDS software package, but also storing the logs and data from the agents. The agent's network interface card (NIC) needs to run in the promiscuous mode. NICs that run in promiscuous mode allow the agent computers to see all of the packet traffic on the network segment that the agent is on. The console and the agent can run on the same computer, I have set up IDS in this manner. Most vendors, however, do not recommend doing this. Running both on the same system can overwhelm the computer's resources, especially on a busy network.

The IDS agent, also known as sensors or engines, need to be deployed in the proper place or places on the network that it is protecting. There is some conflicting information on just where to place them in reference to the firewall. Most information from ISS Corporation and from The US Army's Information Systems Security Officer (ISSO) level two training indicate that the agent should be put just inside the firewall. The rationale is to detect an attack coming through the firewall. The other school of thought is to put it on the outside of the firewall in the DMZ. The thought here is to detect an attack to your firewall. Of course this leaves your agent subject to pummeling! If you put your agent outside the firewall, you would not want your management console on the wrong side of the firewall exposed to attackers. Most firewalls will let you create a secure tunnel or a virtual private network (VPN) that will allow secure access to your agent. Usually all you will need to setup the tunnel is the port numbers that the IDS uses to communicate between console and agent.

If the organization can afford to deploy one agent inside the firewall and one outside the firewall, you gain a tremendous amount of detection capability. You are now able to see if an attack penetrated the firewall, you have eyes on both sides. With the additional agent on the inside you are able to detect inside jobs. Lastly, you will be able to see who is outside that needs to come in but can't due to a bad firewall configuration (Northcutt, S. 1999). Additional locations for sensors are high security areas of the network and subnets on the other side of a router.

Network-based intrusion detection systems have several advantages over host based IDS's.

The first advantage is the low cost of ownership; one console and one agent strategically placed can cover an entire network or the most vulnerable network segment. If using a totally host based detection system, the software would need to be loaded on every machine that needed protecting.

Network-based IDSs will detect several attacks that host-based IDSs won't. This is because the network based IDS will examine packet traffic looking for suspicious activity, such as denial of service and teardrop attacks. Host-based systems will not detect some of these attacks until they have done their damage to the attacked system.

With network-based systems it is harder for the attacker to cover their tracks. This is because the detection is real-time and notifications or actions have already happened. With host-based systems it is possible that an attacker could hack the logs to remove the record that they were there. This information is useful may be called into evidence to support prosecution against an attacker.

Another advantage over host-based systems is that unsuccessful attacks are detected, where as only successful attacks are normally detected by host-based IDS's.

The last advantage of network based over host based is that of operating system independence. Since the host-based IDS is running on the machine that it is protecting, it is dependent on its operating system (Bace, R. 1999).

Host-based IDSs do have a purpose and also have advantages to offer the network security manager. Host-based systems came about as an automated way to check computer based logs. Checking logs manually is probably a lost art. Famed hacker Kevin Mitnick was caught by a network engineer checking the system logs. Kevin Mitnick's attack would have been caught by either network-based or host-based IDSs (Northcutt, S. 1999).

Host-based IDSs monitor system, event and security logs on the Windows NT operating system. On UNIX based machines the syslog is log that is monitored. When a log entry is added to the log, the host-based IDS checks that entry to see if it matches an attack pattern. Some host-based systems can also check certain system files for changes using checksums every so often. If a change is detected, an alert is sent. Additionally, one host-based IDS can also monitor the ports on the computer that it is protecting, it will take action when certain ports are accessed. Once an attack is detected the host-based IDS can; log the event, alert the console, send an E-mail, initiate a SNMP trap, terminate the user login, disable user account.

The advantages of the host-based IDS are interesting. The first is that there are less false positives; this is because it logs events that actually occurred.

The host-based IDS also monitors only specific system activities. This allows for the detection of file accesses, attempts to install executables (such as Trojan Horses), and what a user is doing while logged on locally or remotely (ISS, 1998).

This type of IDS will also detect walk up to the keyboard attacks. Keyboard attacks should really not be a problem if the machine is in a secured area, is logged off, and has a strong password that is changed often.

In a host-based IDS deployment there is no hardware to buy, the system is a software package that runs on the host machine.

Lastly, but more importantly is that host-based IDSs work well on encrypted networks. They also work well on high-speed networks; network-based IDS cannot handle these two types of networks. Since the host-based IDS resides on the computer it cares nothing about the topology or the network protocol in use.

Looking at the advantages of network-based and host-based intrusion detection systems, it is easy to figure out that a combination of both will give an analyst all the information needed to stop attacks and prevent future attacks. This can be costly, but Stephen Northcutt states that by using "minimum reasonable intrusion detection capability", the cost of intrusion detection can be kept to a minimum. The next generation of intrusion detection tools will incorporate both host-based and network based systems in their design.

When looking to purchase an IDS, the buyer needs to do some research and be sure that they are buying what they needs. I say this because IDS software is not cheap. Currently there are over eighty IDS products available. Listed below are just some of the prices that I could find.

- Centrax by CyberSafe (Network and Host based) - Starting at \$2,500
- RealSecure by Internet Security Systems (Network and Host) about \$8,500
- NFR by Network Flight Recorder (Network Based) about \$3,400
- CyberCop by Network Associates (Network Based) about \$9,400

When shopping for an IDS there are two essential features that an IDS should have. The first one is that the system should educate the user; you should not have to be a security expert to protect your network. The second must have is the ability to be updated. This updating is necessary because new attacks are being discovered constantly, therefore attack signature files need to be updated (Bragg, R, 2000). These two points are very valid, in addition to those two points an IDS should also be easy to install and use. The IDS will not be used if it is too difficult to do so. Also, the IDS should allow you to define your own security policy. After you run the IDS for a while you will learn some of the anomalies that occur in your network, you then should be able to lower that event priority or turn off reporting on that event type. The last feature that I would look for is the ability of the IDS to generate good graphical and textual reports. A graphical report is not just eyewash, it is a valuable tool. Everyone has a boss and management needs to see attack attempts or network problems, you can't expect management to decipher raw data.

Many problems on a network can and will be found once an IDS is installed. I installed an IDS on a network in a particular building and one of the first problems found was duplicate Internet Protocol (IP) addresses residing on the network. Normally this is a sign of an attack as this could be someone pretending to be someone that they are not. Since I had just installed the IDS, I could deduce that this duplicate IP had been like this for sometime. I contacted the LAN manager who was more than happy to improve the health of his network.

In Summary, Intrusion Detection Systems can be used to detect attacks against a network. IDSs are available in two major types; network-based and host-based. The network-based looks for security problems on the network, specifically analyzing the packets and looking for attack signatures. The host-based IDS is installed on the computer or server that it is protecting, scanning the system logs for security warnings and other anomalies. The type of network that an organization has is important regarding the selection of IDSs. If you operate an encrypted or fast network, the host-based IDS may be the right choice. If using a conventional network with ties to the Internet the network-based IDS may be the right choice. Placement of the IDS within the network is critical to the effectiveness that it will have in detecting problems. An IDS will take action once an attack is identified, the action that it takes is generally user defined. No matter which type of IDS is chosen, the IDS is not a silver bullet. The IDS must be maintained and used in conjunction with other security measures.

References

Bace, Rebecca. (1999). An Introduction to Intrusion Detection & Assessment. [On-line], Available: <http://www.icsa.com>

Bragg, Roberta. (2000, Feb). Who's Lookin' at You. Microsoft Certified Professional Magazine, 6, 23-27.

ISS article (1998). Network- vs. Host-based Intrusion Detection: A Guide to Intrusion Detection Technology. [On-line], Available: <http://www.iss.net>

Northcutt, Stephen (1999). Network Intrusion Detection: An Analyst's Handbook. Indianapolis: New Riders.

© SANS Institute 2000 - 2005, Author retains full rights.