GIAC
CERTIFICATIONS

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 permited without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at http://www.giac.org/registration/gsec

The History and Evolution of Intrusion Detection

"*The information world is truly electronic-there's no turning back.*"
- Winn Schwartau

During the past five years, security of computer network has become mainstream in most of everyone's lives. Today, most discussions on computer security is centred on the tools or techniques used in protecting and defending networks. The aim of this paper is to examine the origins of detecting, analysing and reporting of malicious activity, where it is today and where it appears to be heading in the future. Some of the many techniques and tools presently used in Network defence will be explored as well.

There are a variety of tools providing a certain level of comfort with acceptable risks used in the defence and surveillance of computer networks. **Defence-in-Depth** is a term encompassing comprehensive analyst training, hardware deployed in strategic positions and a strong security policy necessary for achieving this objective. Everyday, we have tools at our disposal to reach this goal. The aggregation of data comes from routers, the host itself, firewalls, virus scanners and a tool strictly designed to catch known attacks; an Intrusion Detection System (IDS).

What is Intrusion Detection?

A simple definition: It is the unrelenting active attempts in discovering or detecting the presence of intrusive activities.

Intrusion Detection (ID) as it relates to computers and network infrastructure encompasses a far broader scope. It refers to all processes used in discovering unauthorized uses of network or computer devices. This is achieved through specifically designed software with a sole purpose of detecting unusual or abnormal activity.

The beginning

A USAF paper published in October 1972 written by James P. Anderson outlined the fact the USAF had "become increasingly aware of computer security problems. This problem was felt virtually in every aspect of USAF operations and administration".

During that period of time, the USAF had the daunting tasks of providing shared used of their computer systems, which contained various levels of classifications in a need-to-know environment with a user base holding various levels of security clearance. Thirty years ago, this created a grave problem that is still with us today. The problem remains: How to safely secure separate classification domains on the same network without compromising security?[1]

© SANS Institute 2000 - 2005 Author retains full rights.

In 1980, James P. Anderson published a study outlining ways to improve computer security auditing and surveillance at customer sites. The original idea behind automated ID is often credited to him for his paper on "How to use accounting audit files to detect unauthorized access". This ID study paved the way as a form of misuse detection for mainframe systems.[2]

The first task was to define what threats existed. Before designing an IDS, it was necessary to understand the types of threats and attacks that could be mounted against computers systems and how to recognized them in an audit data. In fact, he was probably referring to the need of a risk assessment plan to understand the threat (what the risks are or vulnerabilities, what the attacks might be or the means of penetrations) thus following with the creation of a security policy to protect the systems in place.

Between 1984 and 1986, Dorothy Denning and Peter Neumann researched and developed the first model of a real-time IDS. This prototype was named the Intrusion Detection Expert System (IDES). This IDES was initially a rule-based expert system trained to detect known malicious activity. This same system has been refined and enhanced to form what is known today as the Next-Generation Intrusion Detection Expert System (NIDES).[3]

The report published by James P. Anderson and the work on the IDES was the start of much of the research on IDS throughout the 1980s and 1990s. During this period, the U.S. government funded most of this research. Projects like Discovery, Haystack, Multics Intrusion Detection and Alerting System (MIDAS), Network Audit Director and Intrusion Reporter (NADIR) were all developed to detect intrusions.

Today

To better understand the terms used within the ID user and research community, some of the most commonly used terms are:

Host-Based: The data from a single host is used to detect signs of intrusion as the packets enters or exits the host.

Network-Based: The data from a network is scrutinized against a database and it flags those who look suspicious. Audit data from one or several hosts may be used as well to detect signs of intrusions.

Anomaly detection model: The IDS has knowledge of normal behavior so it searches for anomalous behavior or deviations from the established baseline. While anomaly detection's most apparent drawback is its high false positive, it does offer detections of unknown intrusions and new exploits.

© SANS Institute 2000 - 2005 Author retains full rights.

Misuse detection model: The IDS has knowledge of suspicious behavior and searches activity that violates stated policies. It also means looking for known malicious or unwanted behavior. In fact, its main features are its efficiency and comparably low false alarm rate.

In the last few years, the ID field has grown considerably and therefore a large number of IDS have been developed to address specific needs[4]. The initial ID systems were once anomaly detection tools but today, misuse detection tools dominate the market. With an increasingly growing number of computer systems connected to networks, ID has become a necessity. In the mid 1990s, commercial products surfaced for the masses. Two of the most popular IDS in the mid 1990s were Wheelgroup's Netranger and Internet Security Systems' RealSecure. Both of these companies started out with network-base IDS.

Wheelgroup was formed in October 1995 to commercialize a security product initially prototyped by the U.S. Air Force then called Netranger. This product "scans traffic for "signature of misuse", providing real-time alarm and details of the furtive attacks that may plague a network".[5] In February 1998, Wheelgroup was acquired by Cisco to eventually become an integral part of Cisco's security architecture.

Internet Security Systems, Inc (ISS) was founded in April 1994 by Thomas Noonan and Christopher Klauss, after Mr. Klauss invented and released the first version of the Internet Scanner.[6] On 9 December 1996, ISS announced the release of a tool to augment network security with real-time attack recognition called RealSecure. On the 19 Aug 1997, they announced the first commercial released of their IDS called RealSecure 1.0 for Windows NT 4.0 a new commercial breakthrough.

Another point to consider is most commercially available systems are knowledge-based, which means matching signatures of known attacks against changes in systems or streams of packets on a network. However, their major weaknesses are, they are often helpless against new attacks, so they must be continually updated with new knowledge for new attacks signatures. Despite the fact these false positives are common with behaviour-based IDS, so is its ability to detect a previously unreported attack.

To help solve the knowledge-based problems, workshops have been held every year for the past four years to share information related to ID.[7]
The research topics are quite varied every year and they cover a wide range of subjects such as Lesson Learned, IDS and Law, Modeling Attacks, Anomaly Detection, etc. These workshops main objective are to find new solutions to new and challenging problems. The problems, the research community are now facing are high-speed networks and switching.

Today, more vendors are advertising they can process at gigabit speed. To name a few,

© SANS Institute 2000 - 2005 Author retains full rights.

Internet Security Systems (ISS), NetworkICE, and Intrusion.com advertise they can analyse and alert on gigabit traffic. As networks expand and get faster, network IDS may loose popularity.

To address this problem, vendors have turned to the host. How can the host be part of the equation and provide data when it is directly probed for information? The solution was to install host-based IDS. The advantages of this type of ID are: analysis of audit or data log, real-time and distributed processing. There are many forms such as host-based ID, TCP Wrappers, Tripwire, and a free tool such as Snort.

Snort is described as a lightweight ID system, with multi-platforms. This ID system can be used in two modes: host-based and network-based. However, when first released by Marty Roesch on 22 December 1998, it was available for UNIX systems only and had limited capabilities. Whereas' Snort as an ID system really took off during Y2K with the released of version 1.5 in December 1999, it was capable of performing real-time packet analysis and logging. During this period, usage included reporting abnormal activity to SANS' GIAC cell. This great success led to it being ported to Windows, by Michael Davis and released for the first time on 6 June 2000.

The rapid increase in network bandwidth from megabits to gigabits per second is making it progressively more difficult in carrying out analysis for detecting network attacks in a timely and accurate manner.

One major challenge network engineers face today is that most organisations are using switches and full duplex Ethernet network, complicating the task of deploying Network Intrusion Detection Systems (NIDS). Cisco solution was the invention and release of a blade, which fits into their Catalyst switch and reports to their Cisco Secure IDS Manager. This blade may not be the only solution for both switching and gigabit speed problems. The problem with data reduction and mining? How do we deal with such a challenge?

Another problem that emerged over the past two years is how to deal with denial-of-service (DoS) attack against perimeter defences? With the capabilities of IDS advances, attackers are finding new ways of detecting and bypassing or disabling ID Systems before attempting to penetrate more valuable targets (i.e. web or DNS server). A simple example would be a probe targeting the TCP DNS service against a class B block. The result would be the IDS alarms the console on every port probes, generating more than 65,000 alarms. You can see why it would overwhelm the console as well as the analyst. We will address this later in data consolidation.

The goal is to frustrate attackers by using an IDS architecture invisible to attackers' normal means of mapping a network. The most common way of accomplishing this "invisibility" is by restricting the communication allowed between different security components on a private network.

© SANS Institute 2000 - 2005 Author retains full rights.

What is in store for the future?

Everyone now has no doubt that "Intrusion detection systems have become an essential component of computer security to detect attacks that can occur despite the best preventative measures."[8] Deploying the right tools to defend and protect a perimeter requires man-hours, patience and knowledge. Security is more complex than any one organization, business process, or any one person's view or agenda.

The IDS research community is developing better techniques for collecting and analyzing data in order to handle intrusions in large, distributed environments. In order to take advantage of this work, ID systems must be able to quickly adapt to new, improved components, and changes in the environment.

After many years in the security field, I believe no one product today or tomorrow will solve every security need. There are too many variable to take in considerations in knowing everything about security. That is why security teams exist such as CERT/CC® with many analysts, each with their own areas of expertise. Every member provides their own strengths and experiences to complement one another. With new intrusions appearing each day, it has become a race between upgrading the intrusion detection system and attackers finding new ways of getting into the various systems deployed on a network.

However, these security teams usually face obvious challenges. Organizations collect huge volumes of data in their daily operations. This wealth of information is often under-utilized because of economic reasons (weak or no database search capability) also, lack of trained personnel to correctly interpret the data. Therefore, in order to sift through large amount of data to discover hidden clues, data mining (also known as Knowledge Discovery in Databases) can be used to dissect the information.

Data mining helps revealing relationships or trends to answer specific questions too complex for traditional query and reporting tools. Recent years have seen a dramatic increase in the amount of information stored in electronic format. It has been estimated that the amount of information in the world doubles every 20 months and the size and number of databases are increasing even faster. The business world has provided some important research and testing by creating knowledge discovery database applications designed for managing the growth of on-line data volumes.

An IDS, a router, a firewall, or a server can generate mountains of data with very little means of merging the data to extract the centre and drill down on the attack. A security analyst's nightmare faced daily, is the amount of false positive data collected by IDS sensors. Being able to recognize low and slow reconnaissance probes or correlating information when amalgamated together. Therefore, yielding significant amount of intelligence is very important. Tools such as Intellitactics' Network Security Manager[9], can be used to drill down the correct information.

© SANS Institute 2000 - 2005 Author retains full rights.

The approach Intellitactics has taken regarding data mining and the manipulation of huge volumes of information is opening everything and letting the Network Security Manager (NSM) do all the work. NSM uses a six-step approach: collection and data consolidation (awareness process), normalize, classifies the assets, prioritize (understanding process) and analyze and response (appropriate response process).

Take a moment to assess the attacker's capabilities in collecting intelligence on the network being protected and defended by you. Is your IDS leaving a footprint making it vulnerable to reconnaissance through a port sweep? (i.e. the attacker is probing on a vendor-defined port, easily identifying the device).

This same ability in collecting the intelligence on anyone accessing, visiting or attacking your network is vital. The security of your network depends on it.

References:

1. Anderson, James P. "Computer Security Technology Planning Study Volume 2", October 1972 http://seclab.cs.ucdavis.edu/projects/history/papers/ande72.pdf

2. Anderson, James P. "Computer Security Threat Monitoring and Surveillance", 15 April 1980 http://seclab.cs.ucdavis.edu/projects/history/papers/ande80.pdf

3. Neumann, Peter. http://www.csl.sri.com/users/neumann/neumann.html

4. Sobirey, Micheal. http://www-rnks.informatik.tu-cottbus.de/~sobirey/ids.html

5. Red Herring Magazine. "Call in the air strike", 1 Dec 1996 http://www.redherring.com/index.asp?layout=story&channel=70000007&doc_id=1910016191

6. http://www.iss.net/company/profile/fact_sheet.php

7. http://www.raid-symposium.org/

8. Lippman, Richard et al. 'The 1999 DARPA off-line intrusion detection evaluation", Volume 34, Number 4, October 2000

9. http://www.intellitactics.com/html/nsm_feature.html

Bace, Rebecca. "Technology Series Intrusion Detection", Macmillan Technical Publishing, 2000

Zirkle, Laurie. "What is host-based Intrusion Detection", http://www.sans.org/newlook/resources/IDFAQ/host_based.htm

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

© SANS Institute 2000 - 2005 Author retains full rights.

Davis, Michael. “Port of Snort for Windows”
http://www.datanerds.net/~mike/snort.html

http://www.cerias.purdue.edu/coast/intrusion-detection/

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